

# JOURNAL OF DEFENCE AND SECURITY INDUSTRIES: STRATEGY AND TECHNOLOGY



VOLUME: 1   ISSUE:2   YEAR: 2026

### **Chairman**

Prof. Dr. Haluk GÖRGÜN

### **Editors-in-Chief**

Editor-in-Chief (Strategy & Policy Section)

Prof. Dr. Hakan KARATAŞ

Editor-in-Chief (Technology & Science Section)

Prof. Dr. İhsan KAYA

### **Section Editors**

Prof. Dr. Cenk AKTAŞ

Assoc. Prof. Dr. Sinan KESKİN

Dr. Veda DUMAN KANTARCIOĞLU

### **Journal of Defence and Security Industries: Strategy and Technology (JDSI)**

Journal of Defence and Security Industries: Strategy and Technology (JDSI) is an open-access, peer-reviewed academic journal dedicated to publishing scientific research in the fields of engineering, technology, management, and strategy as they relate to the defence and security sectors. Established by Defence Industries Academy Publications, JDSI aims to foster academic and industrial collaboration by providing a platform where scholars and professionals can exchange knowledge, insights, and innovative approaches to complex challenges in defence and security.

### **Publication Frequency**

The JDSI accepts the articles in English and is being published three times a year.

### **Correspondence Address**

Journal Contact: [editorial@jdsi.org](mailto:editorial@jdsi.org)

### **Advisory Board**

Prof. Dr. Ahmet Turan Özdemir (Aspilsan Energy Industry and Trade Inc.)  
Prof. Dr. Cemal Zehir (Yıldız Technical University)  
Prof. Dr. Ebru Tümer (Gebze Technical University)  
Prof. Dr. Ebru Menşur (Gebze Technical University)  
Prof. Dr. Ercüment Arvas (ASELSAN Inc.)  
Prof. Dr. Erhan Afyoncu (National Defence University)  
Prof. Dr. F. Duygu Özel Demiralp (Türkiye Biotechnology Institute)  
Prof. Dr. Fahrettin Öztürk (TUSAS Engine Industries Inc.)  
Prof. Dr. İbrahim Kılıçaslan (İşbir Electric Industry Inc.)  
Prof. Dr. Lokman Kuzu (Turkish Aerospace Industries Inc.)  
Prof. Dr. Mahmut F. Akşit (TUSAS Engine Industries Inc.)  
Prof. Dr. Mehmet Çelik (Konya Karatay University)  
Prof. Dr. Mehmet Ali Tekiner (Ordu University)  
Prof. Dr. Nurettin Acır (National Defence University)  
Prof. Dr. Serdar Salman (National Defence University)  
Dr. Ertaç Koca (Presidency of The Republic of Türkiye - Secretariat of Defence Industries)  
Dr. Fatih Say (EHSIM Electronic Warfare System Inc.)  
Dr. Mehmet Akif Nacar (HAVELSAN Inc.)  
Dr. Mehmet Demiroğlu (Turkish Aerospace Industries Inc.)  
Dr. Murat Ceran (Presidency of The Republic of Türkiye - Secretariat of Defence Industries)  
Dr. Şaduman Aziz (STM Inc.)  
Reha Ufuk ER (Presidency of The Republic of Türkiye - Secretariat of Defence Industries)  
Taha Yücel (ASELSAN Inc.)

### **Editorial Board**

Assoc. Prof. Dr. Ahmet Fazıl Yağlı (ASELSAN Inc.)  
Assoc. Prof. Dr. Arif Furkan Mendi (HAVELSAN Inc.)  
Assoc. Prof. Dr. Emrah Sarıca (Baskent University)  
Assoc. Prof. Dr. Seval Aksoy Kürü (Marmara University)  
Dr. Ali Emre Güleç (ASELSAN Inc.)  
Dr. Bahadır Küçükçök (ASELSAN Inc.)  
Dr. Begüm Şahin (Medipol University)  
Dr. Kaan Arda (Presidency of The Republic of Türkiye - Secretariat of Defence Industries)  
Dr. Mevlut Bulut (Presidency of The Republic of Türkiye - Secretariat of Defence Industries)  
Dr. Şafak Bilgi Akdemir (ASELSAN Inc.)  
Dr. Oğuz Menekşe (ROKETSAN Inc.)  
Dr. Sefa Özkaya (National Defence University)  
Selin Tamer (ROKETSAN Inc.)  
Levent Erikan (HAVELSAN Inc.)  
Zeynep Cenk Kuru (Presidency of The Republic of Türkiye - Secretariat of Defence Industries)

## Articles

### *Strategy & Policy Section*

Strategic Importance of Industrial Espionage in the Turkish Defence Industry and the Development of Counterintelligence Awareness Against Espionage Methodologies ..... 1-26

Energy Espionage 5.0: Corporate Intelligence, Critical Metals, and Artificial Intelligence ..... 27-58

### *Technology & Science Section*

Volume Over Lethality: A Matrix-Exponential Framework for Defending Drone Swarms..... 59-87

Comparative Analysis of Artificial Intelligence-Based Learning Systems Used in Air Combat Simulation Environments..... 88-109

ML-Based RF Calibration of Military T/R Modules.....110-128

Domain-Specific Summarization to Enhance LLM Effectiveness .....129-156

Design and Experimental Validation of an EMI Filter for Military DC-DC Converters Achieving MIL-STD-461G CE102 Compliance.....157-176

Mechanically Alloyed Nd-Fe-B Nanoparticles with Graphitic Carbon Shell: Synthesis, Characterization and Magnetic Properties .....177-200

## **Editorial**

### **Foundations of Intelligent and Strategic Defence Technologies**

We are delighted to present the second issue of the Journal of Defence Science and Industry. This issue represents another significant milestone in our mission to establish a multidisciplinary platform for the dissemination of high-quality research in defence science, engineering, and emerging strategic technologies. As a newly established journal, our ambition extends beyond publication alone; we seek to foster a collaborative ecosystem that unites academia, industry, governmental organisations, and research institutions in addressing the complex technological challenges of modern defence.

The nature of defence is undergoing a profound transformation. Technological superiority is no longer determined by the performance of individual platforms or weapon systems, but by the effective integration of artificial intelligence, advanced electronics, resilient materials, autonomous systems, secure digital infrastructures, and strategic awareness into a cohesive defence ecosystem. This holistic perspective is closely aligned with the strategic vision of the Presidency of the Republic of Türkiye Secretariat of Defence Industries, which continues to promote artificial intelligence, cyber resilience, and quantum technologies as key drivers of future national defence capability.

The papers presented in this issue collectively embody this integrated vision. Although they span a broad range of scientific disciplines, they are united by a common objective: advancing the technological foundations of next-generation defence systems through complementary and interdisciplinary research.

We are particularly encouraged by the strong participation of researchers from leading universities, Turkish defence companies, governmental organisations, and national research institutions. Their contributions exemplify the collaborative research environment required to address rapidly evolving technological and geopolitical challenges and demonstrate the growing synergy between scientific excellence, industrial innovation, and national strategic priorities.

As we publish this second issue of the Journal of Defence Science and Industry, we reaffirm our commitment to serving as an international forum for high-impact research in defense science, engineering, and emerging technologies. We extend our sincere gratitude to our authors, reviewers, editorial board members, and readers for their continued trust and support. We hope that this issue will inspire new collaborations, encourage interdisciplinary research, and contribute to the development of intelligent, resilient, and strategically significant defence technologies for the benefit of both the scientific community and society.

Sincerely,

*Prof. Dr. Hakan Karataş & Prof. Dr. İhsan Kaya*

Editors-in-Chief

# Strategic Importance of Industrial Espionage in the Turkish Defence Industry and the Development of Counterintelligence Awareness Against Espionage Methodologies

Kadir Murat Altıntaş<sup>1,\*</sup> 

## Abstract

In the 21st century, the qualitative and quantitative diversification of information technology products has significantly increased the vital dependence of all institutions, both public and private, on information and technology. Although this transformation has provided companies with substantial resource savings and operational efficiency, it has also exposed them to a cyber environment characterized by limited oversight and insufficient control. In other words, the information technology systems and infrastructures established by companies that could not remain indifferent to this technological transformation have become vulnerable to threats such as highly sophisticated data theft, information exfiltration, and the misappropriation of technological know-how. Consequently, large-scale corporations have been compelled to protect, on the one hand, all digital activities transferred into virtual environments, and, on the other hand, the technologies produced through Research and Development (R&D) activities against cyber-originated threats. The aim of this study is to conduct a strategic analysis of the adoption of Counter-Intelligence (CI) principles and the construction of CI awareness—both at the corporate and employee levels—in order to protect the industrial property rights and commercial/technological secrets legally possessed by Turkish defence industry companies, whose economic and political significance is considerable, against illegal industrial espionage attacks and their methodologies. Within this framework, the study tests the following hypothesis: the establishment of counter-intelligence awareness among managers and employees in Turkish defence industry companies regarding industrial espionage methodologies constitutes a strategically significant precaution for preventing potential industrial espionage attacks and minimizing possible damage. All data and information obtained in this study were acquired from open sources, and thematic analysis and document review approaches, which are widely accepted qualitative analysis methods in the literature, were utilized. The fact that sufficient analysis or evaluation on the subject has not been reflected in either the literature or open sources is among the most important limitations of the study, but it also adds to its originality. In this context, the content of the limited publications and documents on the subject that are reflected in open sources has been analysed according to certain rules, and an attempt has been made to reach final conclusions.

**Keywords:** Defence industry, industrial espionage, counter-intelligence, counter-espionage, espionage methodologies.

Received: 14.02.2026

Accepted: 07.05.2026

Research Article

<sup>1</sup> Faculty of Applied Sciences, Abant İzzet Baysal University, Bolu, Türkiye

\*Corresponding author.

✉ kadiralintas@ibu.edu.tr

**Cite this article as:**

<https://doi.org/10.65834/jdsi.12.62>

Altıntaş, K.M., (2026). Strategic Importance of Industrial Espionage in the Turkish Defence Industry and the Development of Counterintelligence Awareness Against Espionage Methodologies. *Journal of Defence and Security Industries: Strategy and Technology* 1(2), 1-26.

## 1. Introduction

It would not be an exaggeration to assert that at no other period in human history have technology and innovation played such a decisive and far-reaching role in shaping our world and human life. In this process, technology not only provides strategic advantage to those who produce and effectively utilize knowledge within the information society, but also emerges as a determining variable influencing economic, cultural, and political spheres—in short, every dimension of life—by shaping levels of welfare, competitiveness, growth, and development. The knowledge, technology, and innovation that define our future, however, appear to be generated predominantly by a limited number of developed countries and corporations (Adıgüzel, 2012). Nevertheless, one of the domains most profoundly affected by technological globalization is information security (TASAM, 2016).

In today's world, the qualitative and quantitative diversification of information technology products has significantly increased the vital dependence of all institutions—without exception, whether public or private—on information and technology. While this development has provided organizations, particularly corporations, with substantial resource savings and enhanced operational efficiency, it has simultaneously exposed them to a cyber environment characterized by limited regulation and constrained oversight. In other words, companies that cannot remain indifferent to this technological transformation have developed extensive information technology systems and infrastructures that are increasingly vulnerable to highly sophisticated threats, including data and information theft as well as the misappropriation of technological know-how. Consequently, large-scale corporations have been compelled to protect, on the one hand, all digital activities relocated to virtual environments, and on the other, the technologies generated through Research and Development (R&D) processes against cyber-originated threats targeting data and information systems (Altıntaş & Asal, 2020). Information security has thus evolved beyond a concept limited to commercial and economic interests; it now intersects with political and administrative interests within the framework of the state's public relations and is further conceptualized in conjunction with national security, particularly in terms of intelligence. When this expansive sphere of influence is considered alongside the rapid transformation accompanying the materialization of information, digital networks emerge as the primary environment in which potential damage may occur (Törenli, 2007).

In corporate structures, both organizational and environmental risks exist in terms of information security; therefore, the protection of organizational and technical data is also required. Particularly, organizational culture and awareness regarding the security and privacy of big data play a crucial role in preventing human-induced data breaches (Salleh & Janczewski, 2016). Although safeguarding the confidentiality of data and information is not an easy task in today's context, the continuous flow of information has stimulated intelligence production and revealed the strategic importance of information security. Consequently, in Türkiye, intelligence activities emerge as a primary constraint that must be taken into

consideration in the process of ensuring data and information security (Şeşen & Kuzucuoğlu, 2021).

Türkiye's location at the intersection of three continents within a geostrategic geography has, for centuries, rendered the ancient lands of Anatolia a focal point of intelligence rivalries. In the twenty-first century, commercial competition among corporate entities has, in many respects, evolved into forms of technological warfare. Beyond the technological capacities of states themselves, the protection of qualified data and information belonging to private enterprises—classified as commercial or technological secrets—against hostile foreign-origin attacks has become an indispensable component of contemporary economic security. In this context, safeguarding proprietary knowledge and strategic industrial capabilities is no longer solely a matter of corporate resilience but a broader issue intertwined with national interests. Indeed, due to the technological advancements and commercial achievements observed in the Turkish defence industry in recent years, it is assessed that these developments have attracted scrutiny not only from the intelligence services of advanced Western states but also from large-scale global defence industry corporations. Such dynamics further reinforce the necessity of integrating economic security, technological sovereignty, and information security within a comprehensive national security framework.

The oligopolistic characteristics that have long prevailed among global defence industry actors—often accompanied by monopolistic market responses—have entrenched all participants within a rigid and consolidated distribution of market shares. Consequently, the entry of new global actors with the capacity not only to influence but potentially to alter the status quo in international markets is perceived as one of the most serious threats to the strategic interests of incumbent firms. Particularly in a period when political and societal expectations regarding the advancement of the defence industry have reached unprecedented levels, the presence of the Republic of Türkiye and its companies—prepared to activate substantial engineering capabilities in the development of “new products”—is regarded as a significant risk factor by established global defence manufacturers. From a market dynamics perspective, a new entrant inevitably signifies, for traditional producers, both a “loss of customers,” as former client states or institutions diversify their procurement sources, and a “loss of market share,” as heightened competition disrupts previously stabilized competitive balances. In this context, the expansion of Türkiye's defence-industrial and technological capabilities introduces structural pressures into a global market historically characterized by concentrated power, limited entry, and high barriers to competition.

On the other hand, in the contemporary context, the adoption of physical and digital safeguards in line with general Counter-Intelligence (CI) principles for the protection of registered assets (including trade secrets and industrial property rights) owned by Turkish defence industry companies is of critical importance; however, such measures alone are insufficient. The identification, detection, examination, investigation, as well as the verification or refutation of potential acts of industrial espionage originating from rival corporations must fall squarely within the mandate and responsibility of in-house CI units. In clearer terms, these functions

should be explicitly defined within the formal job descriptions and operational frameworks of corporate CI personnel. Accordingly, systematically informing and raising awareness among all company employees in accordance with general CI principles constitutes the essential “cement” that reinforces the protective walls surrounding registered corporate assets. By cultivating organizational awareness of industrial espionage threats, employees can internalize CI principles not merely as procedural requirements but as an integral component of corporate culture and professional conduct.

The primary objective of this study is to conduct a strategic analysis of the processes through which CI principles are adopted and a CI awareness culture is constructed among companies and employees within the Turkish defence industry. Particular emphasis is placed on the protection—within a legal framework—of industrial property rights and commercial/technological secrets against unlawful industrial espionage attacks and their associated methodologies. Given the significant economic and political importance attributed to Turkish defence industry companies, safeguarding these proprietary assets constitutes not merely a corporate priority but a matter of broader national interest. In other words, the study seeks to address the following questions: ‘What is the strategic significance of industrial espionage attacks for Turkish defence industry companies, and why is the establishment of CI awareness among managers and employee imperative within the framework of espionage methodologies?’ Within this scope, the study tests the hypothesis that ‘the construction of CI awareness among managers and employees in Turkish defence industry companies—particularly with regard to the methodologies of industrial espionage—constitutes a strategically significant measure for preventing potential espionage attacks and minimizing possible damage’.

Statistical data and qualitative information obtained within the scope of this study were sourced exclusively from open sources and widely accepted qualitative analysis methods in the literature specifically thematic analysis and document review—were employed. One of the main limitations of the study is that sufficient analyses or evaluations on the subject have not been extensively reflected in either the existing literature or available open sources. At the same time, however, this limitation contributes to the originality of the study. Within this framework, the study sought to reach final conclusions by analysing the content of the limited publications and documents available in open sources according to specific methodological rules and criteria.


### **1.1. Recent Commercial/Technological Developments in the Turkish Defence Industry and Future Projections**

The defence industry is a strategic sector that plays a critical role in ensuring national security and directly affects a country’s level of independence from both economic and technological perspective as well as strategic point of view. Nations make substantial investments in the defence industry to strengthen their defence capabilities, enhance deterrence power, and reduce external dependency (Kalkan, 2025). A country’s capacity in the defence industry serves as an

important indicator in determining its global power status. Within this framework, two fundamental characteristics of the defence industry stand out. First, it relies on high-technology systems that require confidentiality to ensure national security. Second, by developing these costly systems with domestic resources, the defence industry significantly contributes to the country's industrial, scientific, and technological capacity (Yılmaztürk, 2023).

This sector, through high-technology-based production processes and intensive R&D investments and activities, significantly contributes to the development of the national technological infrastructure while simultaneously supporting economic growth and generating substantial employment. The export of defence industry products provides states with greater diplomatic leverage and manoeuvring space in the international arena, thereby transforming the defence industry into a strategic instrument that supports not only security concerns but also economic development and foreign policy objectives (Kurç et al., 2017). In this context, national defence projects contribute to enhancing local engineering capabilities, increasing technological innovation, strengthening the competitiveness of the defence industry, preventing brain drain, and boosting employment. These projects also serve as a guarantee of national security, as the country's capacity to produce its own defence equipment ensures greater independence against external threats (SASAD, 2023).

According to the various data comparisons (based on the years 2023-2025) made, the Turkish Defence and Aerospace Industry achieved sales revenue exceeding USD 15 billion. Of this total, international sales accounted for USD 5.5 billion, while R&D expenditure reached USD 2.6 billion in 2023 (Defence and Aerospace Industry Manufacturers Association, 2023). In the 2024 edition of the Defence News Top 100, which ranks the world's largest defence companies, five Turkish firms are listed within the top percentile. ASELSAN, ranked 42nd, reported USD 3 billion in revenue; TUSAŞ, ranked 50th, reported USD 2.2 billion; ROKETSAN, at 71st, generated USD 1.2 billion; MKE, at 84th, had USD 905 million; and ASFAT A.Ş., at 94th, reported approximately USD 650 million in revenue. Similarly, in the 2023 list of the largest 100 companies compiled by the Stockholm International Peace Research Institute (SIPRI), three Turkish companies are featured: ASELSAN at 54<sup>th</sup>, BAYKAR at 69<sup>th</sup>, and TUSAŞ at 78<sup>th</sup> (Defence News, 2024). In 2023, the annual sales revenue of the defence and aerospace sector reached USD 15.072 billion, with the domestic production rate in the defence industry achieving 81% (Presidency of the Republic of Türkiye, Secretariat of Defence Industries, 2025).



**Figure 1.** Sectoral Overview, 2024-2028 Defence Industry Sectoral Strategy Document, 2023, p. 35 and Presidency of the Republic of Türkiye, Secretariat of Defence Industries, Annual Report, 2025, p. 3.

As of 2024, there are approximately 1,300 R&D laboratories and 11 design centres in Türkiye, 66 of which operate specifically within the Defence and Aerospace sector (Ministry of Industry and Technology, 2024). Through collaborative R&D projects conducted by universities and industrial organizations under the coordination of the Secretariat of Defence Industries, the sector has seen a continuous increase in product diversity, export capacity, and technology development processes each year. As a result of R&D and defence industries ovation investments in the Turkish defence industry, domestically produced, technology-oriented defence products have been developed, while at the same time enabling the implementation of geostrategic policies that prioritize security in a region often described as a “fire circle.” Furthermore, the numerous registered assets generated through these R&D and innovation efforts have contributed to the emergence of an ecosystem that also supports broader economic development policies.

The foundation of this progress lies in R&D investments and innovative production processes. In 2023, R&D expenditure increased by 27%, reaching USD 2.6 billion, while employment in the sector rose to 90,969 personnel (up from 75,660 in 2021 and 81,132 in 2022). R&D and innovation investments play a critical role, particularly in ensuring the sustainability of innovations in the defence industry and securing a competitive advantage in the global market. In 2023, the defence industry accounted for 2.5% of Türkiye’s total exports (Figure 1). In terms of the regional distribution of foreign sales revenue, exports to Europe and the Americas constituted 46% of total sales abroad. Specifically, exports to the United States (U.S.) amounted to USD 1.331 billion, exports to Europe reached USD 1.220 billion, and exports to

other regions—including the Middle East, Asia, Africa, and South America—totalled USD 2.994 billion. Meanwhile, imports in 2023 rose by 22.35%, reaching USD 3.303 billion. Regionally, imports from Europe increased by 23.11%, from the America's by 22.95%, and from other countries by 18.42% (SASAD, 2023).

According to the April 2023 report by the SIPRI, Türkiye allocated 1.31% of its GDP to defence industry expenditures in 2023, with approximately 25.4% of the defence budget dedicated to equipment procurement and R&D activities (SIPRI, 2024). In 2025, defence industry expenditures are projected to increase by 12% (SIPRI, 2025).

In 2024, Türkiye's total defence industry expenditure reached USD 25 billion, while the export value of the defence and aerospace sector rose to USD 7.2 billion. As a result, the Turkish defence industry evolved into the world's 11<sup>th</sup> largest arms supplier. By 2024, more than 3,500 companies were actively engaged in a wide range of critical areas within the defence and aerospace sector, from original design and domestic production to modernization and modification, as well as R&D activities and international projects. By 2025, total exports of the Turkish defence industry had further increased to USD 9.87 billion (Presidency of the Republic of Türkiye, Investment and Finance Office, 2024).

Through R&D and Technology Management activities, Türkiye aims to achieve full technological independence by localizing foreign-dependent components and critical technologies in defence systems used by its security forces, as well as by acquiring advanced technologies for future systems that will provide strategic superiority (Presidency of the Republic of Türkiye, Secretariat of Defence Industries, 2023). In line with this objective, the statements of Prof. Dr. Haluk Görgün, Secretary of the Defence Industry, in the 2024–2028 Defence Industry Sectoral Strategy Document are particularly noteworthy:

“In the new century of our nation, within the framework of our goals to ‘increase the technological depth and domestic content of the defence industry, minimize foreign dependency as much as possible, ensure the sustainability of the defence industry ecosystem, and focus on breakthrough technologies,’ we will focus, through the policies and strategies we implement, on enhancing the technological depth of our defence industry—which has reached a certain level of maturity at the system level—on disseminating the acquired knowledge and capabilities throughout the sector to ensure continuity, and on increasing the leverage effect of the defence industry in the economy through our visionary export models” (Presidency of the Republic of Türkiye, Secretariat of Defence Industries, 2023).

In the process of transforming from a regional power into a global actor, the Republic of Türkiye derives unique contributions and opportunities from its defence industry. In this context, both the implementation of Türkiye's regional policies and interests, as well as the execution of strategic decisions on a global scale, are supported by the combined deterrence capabilities and capacities of the Turkish Armed Forces and the Turkish defence industry.

These capabilities provide strategic support to both Türkiye's regional security and its economic development. Accordingly, it is anticipated that, particularly through high value-added products and technology-oriented production supported by R&D and innovation, the Turkish defence industry will continue to develop comprehensively in the medium and long term.

On the other hand, legally protected registered assets—classified as intangible fixed assets owned by commercial entities—have become a fundamental factor enhancing the competitive capacity and economic sustainability of globally operating companies. The primary purpose of intellectual property rights is to ensure the legal protection of high value-added products, innovations, or technological infrastructure. However, merely achieving advancements in ground-breaking technological fields through R&D and innovation investments, or creating “new products” that generate added value in the commercialization process, is no longer sufficient. Some global companies, without conducting adequate engineering work or significant R&D and innovation activities, acquire new technologies through unlawful means—often without any payment—and import them to the market, thereby creating “unfair competition.” By resorting to unethical and illegal methods, these companies gain sectoral advantages. While this situation provides significant competitive superiority to the offending companies, it simultaneously causes the victimized firms to lose their commercial and technological capabilities and accumulated expertise in the short to medium term.

The consequences of the "industrial espionage" survey conducted by a non-governmental organization-BITKOM, representing more than 2200 companies in Germany's information, communication and new media sectors, have been announced that seven out of every ten industrial companies in the Germany have been the victims of sabotage, data theft or industrial espionage attacks in recent years (BITKOM, 2021). Elizabeth et.al. finds out that ‘unfair competition’ in the domain of ‘Intellectual Property Rights’ involves actions by an operator that infringe the legislation's rules, harm operators' legitimate rights and interest on the market, thereby threatening socio-economic order. The study concluded that unfair competition acts in the field of intellectual property rights (like industrial espionage, infringement of trade secrets, trademark infringement) impede innovation, stifle, and affect competition in the real market. And the occurrence of these unfair competition practices is common, even though there are existing legal frameworks that regulate this sector (Elizabeth et al., 2021).

Today, illicit industrial espionage attempts target not only the defence industry infrastructures and large-scale companies of developed western countries but also smaller and medium-sized enterprises in developing nations that, despite having limited capacity, are more dynamic and hold significant future potential. In some cases, even local commercial organizations that support macroeconomic development and possess growth potential can become victims of industrial espionage attacks. In other words, beyond the Turkish defence industry, technology-oriented firms operating in sectors such as wireless communications, consumer electronics, acrylic fiber, image processing and artificial intelligence technologies, and global gaming

software—even those with relatively limited production and export capacities—are also exposed to these risks.

## **1.2. The Strategic Importance and Methodology of Industrial Espionage within the Context of Illegal/Unethical Technology Transfer**

The trade secrets<sup>2</sup> and intellectual property rights<sup>3</sup>—arguably the most important assets of companies—essentially protect the economic value of firms’ technologically developing resources. For the qualified data and information possessed by companies to be recognized as “registered assets,” they are generally expected to meet three fundamental criteria: they must be confidential, they must hold commercial value due to their confidentiality, and they must be subject to specific preventive measures to ensure their secrecy (Lippoldt & Schultz, 2014).

In today’s world, where the commercialization of technology remains a prevailing trend, companies that undertake large-scale R&D and innovation investments, thereby dedicating substantial time, effort, and financial resources to enhancing their technological capabilities—have gained a competitive edge in international markets and achieved notable financial success. However, industrial espionage attacks targeting legally protected registered assets—such as trade and technological secrets or intellectual property rights—enable malicious competitors to attain the desired level of technological advancement without incurring comparable investments of time, labour, or capital. In this sense, industrial espionage initiatives, frequently observed in the twenty-first century, have evolved into a global phenomenon that effectively challenges the existence of companies operating within legal frameworks, notwithstanding prevailing international norms and domestic legal restrictions. By contrast, companies may lawfully obtain access to a close equivalent of an original product by supporting limited technical data acquired through legitimate means with methods such as reverse engineering and competitive intelligence. When conducted within established legal boundaries, such practices do not constitute a criminal offense.

The historical origins of industrial espionage activities go back to ancient times. The expansion of international trade as well as technological advancements have increased the theft of data

---

<sup>2</sup> According to the definition of “trade secret” set forth in the third section of the October 2006 Draft Law on Trade Secrets, the term encompasses “data, information, and documents related to the field of activity of a commercial enterprise or company that are known or accessible only to a limited number of its members and other authorized personnel; that are likely to cause harm if learned by competitors; that must not be disclosed to third parties or to the public; and that are of substantial importance for the company’s success and efficiency in economic life.”

<sup>3</sup> Commercial and/or technological data and information relating to registered assets—such as trademarks, patents, industrial designs, integrated circuit topographies, geographical indications, or utility models—are legally protected against unauthorized reproduction for the duration of their registration under intellectual property regimes. However, when hostile actors originating from rival companies reinforce the limited technical data disclosed to the public within the framework of patent law with confidential (technological) information unlawfully obtained through industrial espionage, they do not hesitate to introduce ostensibly “new” products to the market. This practice fosters an environment of “unfair competition” in international markets while simultaneously posing a direct threat to the commercial sustainability of the victimized companies.

and information between societies. In 1474, Venice, one of the leading states of that time, granted significant commercial privileges to merchants who brought technological innovations from various parts of the world to the country (Altıntaş & Asal, 2020). In 1712, Reverend François Xavier d'Entrecolles, who visited Jingdezhen in China, secretly learned the methods of producing Chinese porcelain and brought this knowledge back to Europe. Subsequently, manufacturers used this technology-based information to begin and develop porcelain production in Europe (Rowe and Brook, 2009).

In the 21st century—an era marked by an unprecedented rise in the global perception of industrial espionage as a serious threat—intense competition and hostile conduct have emerged among international corporations, often supported directly or indirectly by rival states. In this context, industrial espionage activities frequently target companies operating in sectors such as defence and aerospace, medical technologies, software, automotive, and electronics. For technology-oriented corporate entities, potential threats may originate not only from competitors but also from suppliers, stakeholders, and even dissatisfied current or former employees. These insider risks, combined with external hostile actors, significantly increase the vulnerability of high-technology firms. A case reported by the news agency Reuters illustrates this risk: the European aerospace giant Airbus dismissed 16 employees and launched an internal investigation in 2019 following an industrial espionage incident. The suspected employees were accused of unlawfully attempting to obtain communication systems and classified documents belonging to the German armed forces (Reuters, 2019). In another case, a military contractor named Hohn Murray Rowe was held responsible in September 2025 for sharing sensitive U.S. defence industry information—particularly classified data related to U.S. Air Force electronic warfare technologies—with foreign nationals. A court in South Dakota found the individual guilty and sentenced him to 126 months' imprisonment along with a USD 25,000 fine (U.S. Department of Justice, 2025). These cases demonstrate that industrial espionage constitutes not merely a theoretical or abstract risk but a tangible and recurring threat capable of generating severe legal, financial, and strategic consequences.

Although industrial espionage resembles economic espionage (Wagner, 2012; Soilen, 2016), it is generally conducted covertly and deceptively by private companies acting on their own behalf. In other words, industrial espionage refers to the acquisition of economic intelligence by the private sector using clandestine and unlawful means (Porteous, 1994). According to Baltacı (2015), industrial espionage encompasses activities carried out by a private company—or by individuals acting on its behalf—aimed at obtaining information of an economic intelligence nature that cannot be acquired from open sources, through covert and illegal methods directed against competitors. Ultimately, industrial espionage can be defined as the covert and deceptive acquisition—through ethically and legally questionable methods—of competitors' trade and technological secrets or confidential data and information protected under intellectual property rights (and not accessible through open sources), by commercial entities seeking to gain advantage in international competition (Altıntaş & Asal, 2020). The primary goal of industrial espionage is to obtain confidential and strategic information about rival companies (Toriola, 2011). According to another definition, industrial espionage is the

theft of a business's confidential or protected information (such as trade secrets or customer lists that are not publicly available) for use by a competitor or a foreign country (Benny, 2014).

On the other hand, the possible expectations and ultimate objectives of companies engaging in industrial espionage activities may be summarized as follows (Altıntaş & Asal, 2020):

- a. To gain competitive superiority and prestige within the sector,
- b. To save resources allocated particularly to R&D activities,
- c. To cause reputational damage to rival firms,
- d. To enhance effectiveness in internal decision-making processes.

A recent case illustrates these motivations in practice. A Chinese national, Ji Wang, who worked on projects funded by U.S. Defence Advanced Research Projects Agency (DARPA) and the private company Corning Incorporated, was involved in the development of optical fibers for high-powered lasers with both military and commercial applications. He unlawfully used non-public data and information—including secret trade production technologies required for manufacturing advanced optical fibers—for his personal commercial future and for establishing his own company. Following his arrest by law enforcement authorities, Ji Wang was found guilty, and at a hearing scheduled for April 15, 2026, he is expected to face a potential sentence of up to ten years' imprisonment (U.S. Department of Justice, 2025). This case demonstrates how the illicit acquisition and exploitation of confidential technological knowledge can directly serve competitive, financial, and strategic ambitions, while simultaneously generating severe criminal consequences.

Parties engaged in economic espionage do not refrain from spying on one another, even when they are friendly or allied states (Moyer, 1994). Countries allied with the U.S.—such as France, Israel, Germany, South Korea, and Japan—have reportedly targeted U.S. industry over the past several decades and have not hesitated, when deemed necessary for national interests, to appropriate commercial and technological secrets or to infringe intellectual property rights (Schweizer, 1993). Industrial espionage disputes between corporations likewise impose significant legal, financial, and temporal costs on both parties. For example, in 2006 the Polish roof-window manufacturer Fakro initiated legal proceedings against Velux on grounds of unfair competition. However, after more than a decade of litigation, the European Commission ultimately rejected the complaint in 2018. Similarly, in 2018, Uber reached a settlement with Waymo by transferring equity valued at approximately USD 244.8 million—equivalent to 0.34% of Uber's shares based on a USD 72 billion valuation (Toren, 2018). These cases illustrate those industrial and technological disputes, whether framed as unfair competition or trade secret litigation, often result in protracted legal battles and substantial economic costs, even when resolved through settlement mechanisms.

Another illustrative case concerns the Swiss-based private company Crypto AG, which produced encrypted communication devices intended to enable secure military and diplomatic communications among NATO member states and other countries. Toward the end of the twentieth century, the company reportedly sold encrypted devices to more than 120 countries

worldwide. However, what its customers were unaware of was that these cryptographic systems allegedly contained a “backdoor” mechanism accessible to the Central Intelligence Agency (CIA). In 1992, a Crypto AG engineer was detained and interrogated by Iranian intelligence authorities, bringing increased scrutiny to the company’s activities. Despite various disclosures and allegations over time, the company continued its operations for years and ultimately ceased activities in 2018 (Miller, 2020). This case underscores the complex intersection between intelligence operations, private sector technology providers, and international trust in secure communication systems, illustrating how commercial cryptographic technologies may become instruments of state-level intelligence competition.

The number of disclosed industrial espionage cases represents only the “tip of the iceberg” (Androulidakis & Kioupakis, 2016). Cybersecurity attacks generating victims across 139 countries were analysed in the 2025 Data Breach Investigations Report. According to the report, a total of 22,052 security incidents formally reported and subjected to legal proceedings by companies of varying sizes and sectors were examined, and 12,195 of these cases were confirmed to involve verified data breaches (Verizon, 2025). However, victims of industrial espionage are often highly reluctant to report incidents to judicial authorities due to concerns over reputational damage in the eyes of customers and shareholders. Consequently, companies frequently tend to pursue out-of-court settlements with attackers or rival firms (Waziri & Yerima, 2011). Moreover, because victimized companies are generally unable to immediately detect data exfiltration or information misappropriation attempts, it is often difficult to determine the probable scope and extent of the damage incurred.

However, according to a 2018 report prepared by the Centre for Strategic and International Studies (CSIS), the estimated financial losses incurred by companies because of cybercrime and cyber espionage amount to approximately USD 600 billion (Lewis, 2018). Beyond the direct losses suffered by companies due to industrial espionage attacks, it is also necessary to consider indirect damages—although difficult to quantify—such as long-term harm to business prospects, erosion of competitive positioning, and the loss of customers (McKown, 2017). For this reason, relying solely on legal remedies to compensate for damages arising from industrial espionage is not a realistic strategy. Instead, awareness training and the cultivation of CI consciousness among managers and employees constitute far more functional and proactive measures in preventing industrial espionage activities (Kahn, 2019).

On the other hand, the absolute advantages that defence industry companies in Türkiye enjoy in accessing a highly skilled workforce—primarily stemming from the long-established engineering education infrastructure and project experience provided by Turkish universities—constitute the primary reason for their success in international competition today. However, it must be remembered that this highly qualified workforce is simultaneously the main target and potential victim of industrial espionage attacks. Therefore, a foundational awareness should be instilled among these individuals through education and training aligned with general CI principles. There have been public allegations that the Dutch government has incentivized engineers working in Türkiye’s defence industry to relocate to the Netherlands and, moreover,

“deliberately” persuaded engineers involved in strategic defence projects to leave the country (Acar, 2018). Reports also indicate that the number of Turkish engineers who have moved to ASML—a company specializing in the production of machinery and facilities for semiconductors used in computer manufacturing—has reached approximately 1,300 (Nebil, 2024). This situation underscores the strategic importance of protecting intellectual and human capital, as the exfiltration of skilled personnel can directly facilitate technology transfer to foreign competitors and compromise national defence capabilities.

Within this framework of Altintas Pyramid (2021a) industrial espionage methodologies initially commence with Open Source<sup>4</sup> and Competitive Intelligence<sup>5</sup> activities. After acquiring preliminary information—serving as the foundation for potential illicit industrial espionage targeting specific companies—through open source and competitive intelligence efforts, a subsequent phase emerges that requires significantly greater expertise and experience.

At this stage, which requires a specific knowledge base and skill set in espionage, the core activity revolves around “Social Engineering,” often succinctly described as the “art of human deception/hacking.” Within this context, personnel considered the “weak link” inside a target company are identified and “profiled,” after which the attacker exploits the individual’s personal vulnerabilities to gain unauthorized access to the company’s systems. In more precise terms, social engineering involves the deliberate manipulation of human emotions and deficiencies to coerce a target employee into disclosing confidential data and information about their employer, using various persuasion and deception techniques while bypassing ethical standards (Gehl & Lawson, 2022). Operations that exploit emotions such as fear, excitement, patriotism, or concerns over reputation and trust with clients or the public compel employees to make mistakes. These social engineering campaigns not only result in significant data and

---

<sup>4</sup> According to the official website of the Milli İstihbarat Teşkilatı (Turkish National Intelligence Service), open sources are defined as any information, documents, materials, or individuals, institutions, and organizations that can provide such information, whose disclosure, acquisition, and use pose no security concern (i.e., non-classified). Open-Source Intelligence (OSINT) is explained as the process of carefully analysing these open sources, collecting information from effective and reliable channels, and delivering it to the analyst in a timely manner and proper format to produce actionable intelligence ([www.mit.gov.tr/sozluk.html#A](http://www.mit.gov.tr/sozluk.html#A), Accessed: 30.01.2026). Sources utilized to generate OSINT include the cyber domain, professional reports, trade fairs, land registry records, publicly available patent information, social media, print and broadcast media, and similar publicly accessible platforms.

<sup>5</sup> In a highly competitive business environment, decision-makers need to protect their company’s strategies from competitors while simultaneously acquiring as much information as possible about competitors’ objectives and plans. Consequently, in today’s dynamic markets, it has become essential for companies to anticipate and understand competitors’ actions to effectively position their products and services. Understanding competitors is also a critical component in developing a viable corporate strategy. Effective strategic planning requires the integration of high-quality competitive intelligence. To achieve this, companies must develop alternative intelligence sources, scan various aspects of their environment, and sensitize employees to the necessity of gathering intelligence. As observed, company executives need to collect and process information that is valuable for strategy formulation, including insights about competitors, customers, suppliers, governments, technological trends, and broader environmental developments. In summary, the generation and processing of information regarding a company’s external environment for strategic purposes is referred to as competitive intelligence (Özdemir, 2010).

information loss but can, in some cases, render the systems of competitor companies entirely inoperative.

In October 2010, Glenn Duffie Shriver was convicted of conspiring to communicate classified U.S. national defence information to foreign actors and was sentenced to four years' imprisonment. Born in Virginia in November 1981, Shriver studied at a university in Michigan and participated in a 45-day summer program in Shanghai at East China Normal University. He later had the opportunity to complete an additional academic year there and developed advanced Chinese language proficiency. Shriver was reportedly approached with an offer to write a political article on U.S.–China relations for a Chinese newspaper. Subsequently, the newspaper's representative introduced him to other Chinese officials, and a cooperative relationship gradually developed. Rather than recruiting him directly as an intelligence asset, they allegedly encouraged him to apply for positions within U.S. law enforcement or intelligence institutions such as the CIA. Shriver received approximately USD 70,000 in several payments from Chinese officials. In June 2010, while preparing to travel to South Korea, he was detained at an airport. According to reports, a covert personality assessment conducted by Prof. Geling Shang—one of the coordinators of the summer program in China—identified Shriver's pronounced financial motivations. This vulnerability was strategically exploited during the recruitment process, and his primary weakness was manipulated to facilitate his engagement. The case demonstrates how social engineering and long-term cultivation strategies—particularly those leveraging financial incentives and psychological profiling—can be employed to infiltrate sensitive national security institutions (Reuters, 2010).

Another prevalent method within industrial espionage methodology is “Parallel Employment” a highly effective attack vector with significant damage potential. This technique involves identifying personnel within the target company who occupy operationally critical positions and persuading them to engage in dual employment—effectively, transforming them into de facto agents acting on behalf of a rival entity. Accordingly, employees at all organizational levels may be assessed as potential targets. Particular attention is often directed toward individuals who exhibit financial vulnerabilities or personal weaknesses, including susceptibility to inappropriate relationships. Once the primary motivational driver of the selected individual is identified and exploited, the recruitment (or engagement) process becomes considerably easier. Through this mechanism, access to the desired confidential data and proprietary information can be systematically secured. Parallel employment thus represents a hybrid insider threat model, combining elements of recruitment, coercion, and exploitation, and underscores the critical importance of robust internal CI awareness and personnel security practices within strategically sensitive sectors.

In February 2024, Chenguang Gong, a naturalized U.S. citizen since 2011, was arrested for allegedly stealing trade-secret technologies developed by a private R&D company on behalf of the U.S. government. Gong illegally transferred over 3,600 files onto personal storage devices, inflicting substantial financial damage on the victim company, which had invested more than seven years and tens of millions of dollars into developing these technologies. Among the

stolen files were designs for infrared sensors intended for satellite-based systems used to track ballistic and hypersonic missiles, as well as sensor schematics designed to implement countermeasures to interfere with missiles' infrared tracking capabilities (U.S. Department of Justice, 2024). This case highlights the critical vulnerabilities of proprietary technologies and the potentially severe financial and strategic consequences of insider-enabled industrial espionage.

Another method within industrial espionage methodology may be summarized as the strategic "Placement of a Spy" within the target organization's critical technical positions. In such operations, the individual's true motivation is deliberately concealed and often masked by carefully constructed cover narratives. These infiltration operations, which require considerable patience and a high degree of expertise, are typically associated with individuals possessing professional intelligence backgrounds or the discipline derived from formal military training. Within this framework, unlawful technology transfers are executed directly by an individual formally employed on the organization's payroll, thereby embedding the espionage activity within the legitimate operational structure of the company. This method represents a sophisticated insider threat model, in which the organization's own human capital is systematically leveraged to facilitate covert acquisition and exfiltration of sensitive technological assets.

Chinese electronics engineer Wei Sun was convicted for illegally transferring sensitive defence technology data from Raytheon Missiles and Defence, where he had worked for over a decade, to China. Due to his senior engineering role, Sun had access to most of the company's classified information related to missile guidance systems and other defence technologies. Between December 2018 and January 2019, during business trips to China, he removed sensitive data from the company laptop without implementing any protective measures, despite being aware of established security protocols, and delivered the information to Chinese counterparts (South China Morning Post, 2020). As a result, Sun was sentenced in the U.S. to 38 months of imprisonment. This case exemplifies the insider threat dimension of industrial espionage, highlighting how privileged access and technical expertise can be exploited to exfiltrate critical defence-related intellectual property across international borders.

The increasing use of the Internet for data transmission and electronic communication in the final decades of the 20<sup>th</sup> century has simultaneously facilitated unauthorized interference by malicious actors, rendering traditional notions of crime and punishment largely ineffective in the uncontrolled and insecure cyberspace. Consequently, ill-intentioned parties capable of gaining illegal access to sensitive data and information through 'Cyberattacks' inevitably began targeting technological assets produced through large-scale R&D investments and activities. In recent years, cyberattacks have become a significant concern—indeed, a near-nightmare—for defence industry companies seeking to protect their intellectual property. These attacks can be organized remotely from geographically distant locations at almost zero cost, making the legal prosecution of perpetrators extremely difficult and, in many cases, preventing the effective administration of justice. The content and methodology of such cyberattacks

predominantly involve the interception of competitors' email traffic and the unauthorized control of electronic communication systems, providing attackers with high-impact results relative to the minimal resources expended (Altıntaş, 2021).

In March 2025, the U.S. Department of Justice filed multiple charges against 12 Chinese cyber actors. Referred to as 'Cyber Mercenaries' this group was accused of serious offenses, including the theft of sensitive information belonging to the U.S. Department of Defence Intelligence and unauthorized access to the computer networks of defence industry suppliers. The indictment, focused on cyber espionage, specifically alleges that Yin Kecheng and Zhou Shuai, between 2011 and 2024, identified vulnerabilities in the target companies' computer networks, deployed malicious software, and established dedicated virtual servers to exfiltrate sensitive information (Cyber Security News, 2025).

In recent decades, Türkiye's defence and aerospace industry efforts to develop 'new products' have attracted significant attention in international markets and are being closely monitored. Consequently, in line with CI principles, the security of technological assets and trade secrets within the Turkish Defence and Aerospace Industry has become more critical than ever, making it imperative to implement the necessary measures to protect registered intellectual property and proprietary assets.

## **2. Discussion**

Industrial espionage targeting the trade secrets or intellectual property rights of defence industry companies—often indirectly facilitated by adversary states or deliberately aimed at by competing firms—constitutes one of the most serious threats to both national and economic security. Approaching the issue solely from the perspective of criminal law, or enforcing stricter judicial measures, does not guarantee the prevention of potential industrial espionage incidents in the future. Similarly, physical security measures implemented by defence industry companies, or law enforcement efforts to safeguard sensitive company assets, do not ensure the complete protection of registered intellectual property and proprietary information. In summary, for defence industry companies, both internalizing counter-intelligence awareness at the individual level and establishing counter-intelligence principles at the corporate level are of strategic importance.

CI should not be the exclusive domain of intelligence agencies. Highly sensitive institutions, such as diplomatic missions, economic bodies, and energy authorities, must also establish dedicated CI units (Kavıracı & Demirbaş, 2020). Similarly, for private sector entities, it is imperative to cultivate CI awareness both at the organizational level and among individual employees. The adoption of a CI mindset by managers and staff, and its integration as a habitual practice, constitutes a more effective protective measure against industrial espionage attacks than even the most stringent legal sanctions.

Intelligence can be succinctly defined as the collection of necessary information by a state to ensure its security or gain strategic advantages. Since other states naturally seek similar

information for their own interests, it is expected that they will also engage in information-gathering activities. At this point, the protection of strategically important state information and the prevention of its acquisition by foreign intelligence organizations—i.e., counterintelligence—becomes critically important (Tezsever, 1999). CI can be described as the process of identifying potential threats posed by foreign intelligence services, neutralizing these threats, and manipulating organizations attempting to collect intelligence on the state in ways that serve one's own national interests (Godson, 2004). Just as intelligence activities are employed to gain superiority over or neutralize enemy elements, CI functions as a “radar” activity aimed at detecting and neutralizing the adversary's tools or capabilities (Özdağ, 2018). In essence, CI activities encompass all protective security measures for safeguarding data, information, documents, personnel, materials, and facilities, alongside active counterespionage operations.

CI activities, which must inevitably be considered by defence industry companies, encompass a variety of operations with distinct characteristics. Activities that primarily involve taking traditional measures to physically and digitally protect trade secrets or intellectual property rights—where the “defensive reflex” predominates—fall under the scope of CI operations. Such traditional CI measures include classifying documents according to confidentiality levels, implementing physical security and protective measures, selecting personnel for critical positions and subsequently monitoring their financial behaviours including those of family members, and enforcing encrypted access protocols for both building/room entry and intranet/internet usage. For instance, in August 2021, the British Intelligence Service discovered that David Smith, assigned to the United Kingdom's Berlin Embassy, had not used bank or credit cards for an extended period. Following an investigation by British CI authorities, it was revealed that Smith had leaked information to the Russian Federation Intelligence Service and was subsequently arrested (Conner, 2021).

Counter-Espionage (CE), a narrower subset of CI, specifically focuses on detecting, neutralizing, and ensuring the legal or operational resolution of espionage activities. This includes not only identifying infiltration or spy attempts conducted by adversarial actors but also, where applicable, manipulating the opposing side through deliberately leaking or misleading information and attempting to convert an operational asset into a double agent. For example, in 1994, Aldrich Ames, who served as a CI officer at the U.S. CIA, was convicted of long-term espionage on behalf of the Soviet Union and sentenced to life imprisonment. During this period, Ames received approximately \$4.6 million in unofficial payments from Russian intelligence, which allowed him to maintain a lifestyle far beyond that of his colleagues. Factors such as financial difficulties arising from his second wife's penchant for luxury and his longstanding struggle with alcohol addiction were reportedly exploited by Russian operatives in recruiting him (Encyclopaedia of Intelligence and Counterintelligence, 2015).

The philosophy of CI, encompassing both defensive and offensive strategies, is structured around four core functions (Van Cleave, 2013, 60). Interpreting these stages in the context of

commercial enterprises allows for a better understanding of industrial espionage threats and enables more effective responses. These processes can be briefly summarized as follows:

**a. Identification and Assessment of Risks and Threats:** The ability to discern the true intentions behind competitors' activities and to analyse the potential (both short and long-term) risks and threats these activities may pose,

**b. Identification of Vulnerabilities and Assessment of Hostile Activities:** Following the detection of existing weaknesses, determining the areas in which competitors' actions—disguised under the guise of commercial activities—are concentrated in a hostile or competition-impeding manner,

**c. Neutralization and Legal Prosecution of Responsible Parties:** Following the detection of industrial espionage attempts by competitor companies through defence measures, taking action to neutralize the perpetrators and subsequently initiating legal proceedings.

**d. Turning Attacks Against Competitors (Counter-Exploitation):** As an alternative to legal referral, engaging operational personnel to compel a change of allegiance, thereby misleading the competitor company with erroneous data and information—in short, “baiting” the adversary.

Establishing CI awareness among employees and standardizing protective measures against both physical and cyber threats are indispensable for safeguarding a company's trade secrets and intellectual property rights. This is because the human element lies at the core of CI activities. It is particularly noteworthy that nearly all potential interactions in any malicious attack originate from human actions. In this context, the human factor simultaneously assumes both the roles of victim and perpetrator. If employees are not fully informed about the security measures in place or fail to fully embrace them, the functionality of a company's CI policies cannot be ensured. Therefore, it is essential to equip personnel with comprehensive CI training programs and to periodically emphasize the strategic importance of these measures. However, training alone is insufficient. The ultimate objective is to cultivate genuine CI awareness among employees, fostering a mindset in which the philosophy of CI becomes an integral part of their professional conduct and daily practice.

CI activities constitute a comprehensive set of protective measures employed by companies to safeguard against the “unfair competition” initiatives of rival firms. Simultaneously, they establish a defensive barrier against the competitive intelligence efforts of competitors. However, from a proactive perspective, effectively preventing industrial espionage requires the early detection of hostile attempts. Successfully achieving this challenging objective demands substantial experience, accumulated knowledge, and intuitive capacity, alongside the strategic utilization of various tools and methodologies.

Piecing together small fragments of information to form the complete picture—or recognizing the significance of seemingly trivial behaviours and materials—is almost a way of life for CI specialists who question everything and probe behind the scenes. Consequently, it is useful to

briefly outline certain techniques that hold strategic importance in CI operations and are utilized in the detection of industrial espionage attempts:

**a. Dumpster Diving:** Companies may consider the discarded materials of competitors—such as paper scraps, flash drives, non-functional prototypes, old notebooks, and post-it notes that have not been fully destroyed—as valuable sources of data and information. In particular, employees of high-tech firms with substantial R&D investments may inadvertently create opportunities for industrial espionage if they fail to exercise sufficient care with materials disposed of in office trash or waste areas. In highly competitive global markets, the covert collection and analysis of physically discarded items whose structural integrity is partially intact can provide critical insights into a company's current operations and future strategies (Altıntaş, 2021). For example, a paper collector repeatedly retrieving waste from a company's surrounding trash areas is highly suspicious. Therefore, defence industry personnel should always destroy obsolete materials using shredders rather than simply discarding them (Prunckun, 2019).

**b. Determining Vulnerability:** This activity aims to identify potential weaknesses of a target individual and subsequently exploit those vulnerabilities or manipulate them for personal gain (Prunckun, 2019). Human vulnerabilities may arise from financial incentives, physical limitations, past traumas, or even sexual preferences and other personal shortcomings (Davies, 2005, 23). Consequently, employing individuals in defence industry companies—particularly in managerial positions—whose motivations are primarily self-serving and whose vulnerabilities are pronounced is extremely risky. At a minimum, however, senior managers who are aware of these vulnerabilities should closely monitor and supervise such personnel.

**c. Reconnaissance and Observation:** CI personnel within a company are responsible for monitoring an internal target employee, when necessary, but without engaging in any physical intervention. In situations where data or information leakage is suspected, it is critical to track and analyse the locations visited and individuals contacted by the employee under surveillance. In addition to utilizing in-house cameras for observation, the detection of any dead-drop<sup>6</sup> points can only be achieved through close tracking of the suspect employee. During monitoring operations, CI specialists may employ legally permissible electronic surveillance devices but must also ensure the strict protection of the target individual's legally protected personal data.

Moreover, it is the responsibility of company CI officers to train critical technical personnel or managers in defensive reactions before events such as overseas business meetings or trips, particularly when there is a risk of physical or digital surveillance by competitor company representatives or foreign intelligence services. Employees in key positions should also be

---

<sup>6</sup> **Dead-Drop:** A dead-drop is a secure communication method used between parties (e.g., a case officer and an agent) in intelligence operations. It involves prearranging a secret location—such as behind a rock, beside a statue, or under a bench—where classified or sensitive information obtained through espionage can be deposited by one party and later retrieved by another. This technique is employed when direct contact between parties could arouse suspicion or when one party is believed to be under surveillance.

educated about “black bag operations”<sup>7</sup> during foreign accommodations. Competitor personnel may sometimes persuade hotel staff to gain covert access to an important technical employee’s room or luggage, potentially leading to copying or even disappearance of sensitive documents.

**d. Denial and Deception:** A common practice in competitive business environments, denial and deception involves manipulating data and information to gain an advantage over rival companies, obscure future plans, and conceal key capabilities. Considering the high speed of social media dissemination and the significant time required to detect falsified content, the digital realm has effectively become a lawless platform, making legal tracking extremely difficult. This environment also conceals the behind-the-scenes support of public entities (e.g., intelligence agencies) in ostensibly private corporate industrial espionage operations, providing the so-called “Plausible Deniability.” Deliberate dissemination of false information via social media as a deception tool has become an effective competitive instrument, causing potentially severe consequences for those unable to distinguish reality from fabrication. In denial and deception operations, which must be coordinated and complementary, the primary requirement is strict adherence to confidentiality principles while compelling the target company to make strategic errors. For example, deception campaigns suggesting the existence of lucrative business opportunities in a specific region are inherently deniable in both content and outcome.

**e. Agent/Source Exploitation:** Individuals within a competitor organization—or within a target company—may be persuaded to provide confidential information in exchange for strong personal incentives. Acting as auxiliary intelligence assets, these insiders systematically relay information regarding planned, ongoing, or potential corporate actions to the controlling intelligence officer (or the operational manager of the organization’s espionage initiative) (Urhal, 2008, 558). Similarly, in a defence industry company, a CI officer may obtain sensitive information from personnel in strategic departments in exchange for specific motivators. While rigorous internal inspections, strict security measures, or employee monitoring may erode trust and loyalty among staff, continuous management guidance emphasizing that protective measures are not personalized can yield short- to medium-term benefits. Otherwise, employees may begin perceiving one another as potential informants, inevitably decreasing productivity in the medium to long term. Directly sourced intelligence is generally more secure, faster, and cost-effective than other technical surveillance methods, explaining the method’s prevalence. Despite its advantages in crime prevention and evidence collection, source exploitation remains susceptible to manipulation and potential ethical or legal concerns and therefore should only be designed and implemented by experienced CI personnel.

---

<sup>7</sup> **Black Bag Jobs:** In human intelligence operations, black bag jobs refer to illicit infiltration attempts into physical locations or cyber environments that are otherwise restricted or secured. The term “black bag” originates from the black bag carried by burglars, which contains tools intended for use during their covert activities. These operations aim to gain unauthorized access to sensitive areas or systems to obtain information or assets.

### 3. Conclusion

In the second quarter of the 21st century, within a global climate where international economic, technological, and political competition is intensifying day by day, private enterprise managers who have internalized intelligence and espionage concepts—and view intelligence as a force multiplier—are indispensable for the sustainability of companies. Meanwhile, industrial espionage, which has increased both quantitatively and qualitatively in recent years, poses a serious threat to the commercial and technological assets of all medium- and large-scale companies, particularly those operating in technology-focused defence sectors, and can cause disproportionately high damage relative to its cost. In this context, the technological know-how and trade secrets currently possessed by the Turkish defence industry are under significant threat due to widespread and destructive industrial espionage attacks on an international scale.

Developing specific standards for the implementation of physical and cyber measures to protect the proprietary assets of Turkish defence industry companies, or in other words reducing the CI philosophy merely to defence against physical or cyber-attacks and designing an all-encompassing protective system accordingly—is useful but insufficient. In addition, it is necessary to establish internal corporate training programs aimed at informing and raising awareness among employees at all levels of the organization in accordance with general CI principles. In other words, organizing educational programs for managers and employees of Turkish defence industry companies on general CI principles and current industrial espionage methodologies constitutes a highly valuable corporate initiative. However, more importantly, at a minimum, all employees must comprehend the strategic significance of the CI doctrine in the face of industrial espionage threats and develop corresponding awareness, thereby internalizing and normalizing CI practices as a way of life; this represents the most effective measure to prevent data and information leakage. In the contemporary environment, it cannot be claimed that severe legal sanctions alone effectively prevent unfair competitive practices among companies in the context of industrial espionage.

Furthermore, within a proactive CI framework, it is essential to plan protective measures against all forms of industrial espionage in advance and to anticipate potential subversive activities, thereby fostering both individual and organizational defensive culture. Moreover, in the 21st century—where perception management on an international trade scale has become more critical than ever—commercial competitive strategies, which can be conceptualized as an “information war” by rival companies, naturally include psychological operations. Consequently, it is mandatory to establish a defensive infrastructure and culture against potential propaganda campaigns and disinformation efforts, particularly via “new media” platforms, as part of CI activities. In this regard, the establishment of strategically important CI units within Turkish defence industry organizations constitutes a highly valuable initiative; in their absence, similar responsibilities may need to be delegated to human resources departments and the organizational structure reconfigured.

Finally, based on strategic analyses of CI activities and document reviews within the relevant literature, the hypothesis that “the development of CI awareness among managers and

employees regarding industrial espionage methodology in Turkish defence industry companies is a strategically important measure for preventing potential industrial espionage attacks and minimizing probable losses” has been qualitatively confirmed.

In this context, Chinese-origin companies with a historically extensive record of involvement in industrial espionage increasingly target technologically advanced Western companies and innovation hubs. The most pressing questions in this regard concern the degree to which legally registered commercial entities, which outwardly appear as private companies, act independently from the state, and the extent to which Chinese intelligence services are involved in industrial espionage operations on behalf of these companies. Analysis of numerous judicial cases, processed by the U.S. Department of Justice and reflected in open-source documentation, clearly indicates that Chinese intelligence services are positioned at the very centre of these industrial espionage activities.

## Conflict of Interest

There is no conflict of interest among authors.

## Funding

This research received no funding.

## Ethical Statements

Not applicable.

## Data and Code Availability

Not applicable.

## Supplementary Materials

Not applicable.

## References

- Acar, İ. (2018, January 16). *Savunma sanayinden Hollanda'ya beyin göçü. Sabah Gazetesi.* <https://www.sabah.com.tr/ekonomi/2018/01/16/savunma-sanayiinden-hollandaya-beyin-gocu>.
- Adıgüzel, M. (2016). *Teknolojinin küreselleşmesi ve Türkiye açısından bir değerlendirme.* Nobel Akademik Yayıncılık.

- Altıntaş, K. M., & Asal, S. (2020). *Ekonomik istihbarat kapsamında endüstriyel espionaj: Ticari tüzel kişilikler açısından endüstriyel casusluğun stratejik önemi* (5. baskı). Nobel Akademik Yayıncılık. ISBN: 978-625-439-223-8
- Altıntaş, K. M. (2021a). Comparative analysis of strategic relationship between industrial versus corporate espionage within the framework of implementation methods. *Global Security and Intelligence Studies*, 6(1), 105–124. <https://doi.org/10.18278/gsis.6.1.5>
- Altıntaş, K. M. (2021b). İstihbarata karşı koyma prensipleri bakımından atık istihbaratının stratejik önemi ve karşılaştırmalı analizi. *Uluslararası Kriz ve Siyaset Araştırmaları Dergisi*, 5(2), 879–914.
- Baltacı, H. (2015). Ekonomik güvenlik bağlamında ekonomik istihbaratın rolü ve önemi. In S. Aydın (Ed.), *Ekonomik istihbarat* (pp. 81–132). Adalet Yayınevi.
- Benny, D. J. (2014). *Industrial espionage: Developing a counterespionage program*. Taylor & Francis Group.
- BITKOM. (2024). *German businesses under attack: Losses of more than 220 billion euros per year*. URL: <https://www.bitkom.org/EN/List-and-detailpages/Press/German-business-losses-more-than-220-billion-euros-per-year>.
- Conner, N. (2021, January 27). *MI5 began probe into British embassy security guard accused of being 'Putin spy' in Berlin after he stopped using his credit and debit cards*. Daily Mail. <https://www.dailymail.co.uk/news/article-9893287/MI5-began-probe-accused-Putin-spy-stopped-using-credit-cards.html>
- Cyber Security News. (2025, March 13). *US charges 12 Chinese hackers for hacking national security infrastructure*. <https://cybersecuritynews.com/us-charges-12-chinese-hackers>.
- Defence News. (2024). *Top 100 for 2024*. <https://people.defencenews.com/top-100>.
- Carlisle, R. (2015). *Encyclopedia of intelligence and counterintelligence*. Routledge.
- Elizabeth, Y. N., Yu, L. M., Palmer, D. P., & Anyomi, B. (2021). Unfair competition in the field of intellectual property rights: Analyzing concepts, acts of unfair competition and laws. *Journal of Politics and Law*, 14(3).
- Gehl, R. W., & Lawson, S. T. (2022). *Social engineering: How crowdmasters, phreaks, hackers, and trolls created a new form of manipulative communication*. MIT Press.
- Godson, R. (2004). *Dirty tricks or trump cards: U.S. covert action & counter intelligence*. Transaction Publishing.
- Kahn, R. A. (2017). Economic espionage in 2017 and beyond: 10 shocking ways they are stealing your intellectual property and corporate mojo. *Business Law Today*, 1–5.
- Kalkan, D. (2025). Savunma sanayinde milli teknoloji hamlesi: Ankara ili örneği. *Dumlupınar Üniversitesi Sosyal Bilimler Dergisi*, 86, 582–596.
- Kavıracı, O., & Demirbaş, M. (2020). İstihbarat faaliyetlerinin devlet güvenliği açısından incelenmesi. *Anadolu Strateji Dergisi*, 2(1).
- Kurç, Ç., & Neuman, S. G. (2017). Defence industries in the 21st century: A comparative analysis. *Defence Studies*, 17(3), 219–227. <https://doi.org/10.1080/14702436.2017.1353537>

- Lewis, J. (2018). *Economic impact of cybercrime, no slowing down*. McAfee.
- Lippoldt, D. C., & Schultz, M. F. (2014). Trade Secrets, Innovation and the WTO. *E-15 Think Piece*. Geneva: ICTSD and WEF, with IMD.
- McKown, C. (2021, January 27). *The American Greed report: Corporate spying costs billions, can it be stopped?* CNBC. <https://www.cnbc.com/2017/05/13/the-american-greed-report-corporate-spying-costs-billions-can-it-be-stopped.html>
- Miller, G. (2020). The intelligence coup of the century. *The Washington Post*. <https://www.washingtonpost.com/graphics/2020/world/national-security/cia-crypto-encryption-machines-espionage/>
- Nebil, F. S. (2024, April 1). *Savunma sanayii'nden Hollanda'ya giden Türk mühendis sayısı 1.300 oldu*. Türk-İnternet.com. <https://turk-internet.com/savunma-sanayiinden-hollandaya-giden-turk-muhendis-sayisi-1-300-oldu/>
- Özdağ, Ü. (2018). *İstihbarat teorisi*. Ankara: Kripto Basım Yayım Dağıtım.
- Özdemir, E. (2010). Rekabet istihbaratı toplama ve etik: Bir alan araştırması. *İstanbul Üniversitesi Siyasal Bilgiler Fakültesi Dergisi*, 43, 67–95.
- Porteous, S. D. (1994). Economic espionage: Issues rising from increased governmental involvement with private sector. *Intelligence and National Security*, 9(4), 735–752.
- Reuters. (2010, October 23). *American pleads guilty to trying to spy for China*. <https://www.reuters.com/article/us-usa-china-spying-idUSTRE69L3CK20101022/>
- Reuters. (2019, December 1). *Airbus dismisses 16 employees in German compliance investigation*. <https://www.reuters.com/article/business/airbus-dismisses-16-employees-in-german-compliance-investigation-idUSKBN1Y5149/>
- Rowe, W. & Brook, T. (2009). *China's Last Empire: The Great Qing*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press. p. 368. ISBN 978-0-674-03612-3.
- Salleh A. K. & Janczewski L. (2016) Technological, Organizational and Environmental Security and Privacy Issues of Big Data: A Literature Review. *Procedia Computer Science*. December. 100: 19-28. DOI:10.1016/j.procs.2016.09.119.
- Sanayi ve Teknoloji Bakanlığı. (2021–2024). *Girişimci bilgi sistemi*. <https://gbs.sanayi.gov.tr/>
- Sevinç, S. (2022). *İstihbarata karşı koyma faaliyetleri kapsamında bir örnek olay incelemesi: Rubicon operasyonu*. (Yüksek Lisans Tezi, Jandarma ve Sahil Güvenlik Akademisi Güvenlik Bilimleri Enstitüsü, Ankara).
- South China Morning Post. (2020, November 18). *Former Raytheon engineer jailed for exporting military-related tech to China*. <https://www.scmp.com/news/china/diplomacy/article/3110339/former-raytheon-engineer-jailed-exporting-military-related>
- Savunma ve Havacılık Sanayi İmalatçılar Derneği (SASAD). (2023). *Performans raporu*. <https://www.sasad.org.tr/sasad-sektor-performans-raporu-2023-1751544308>

- SIPRI. (2024). *Military expenditure database: Armaments, disarmament and international security*. [https://www.sipri.org/sites/default/files/2024-06/yb24\\_summary\\_en\\_2\\_1.pdf](https://www.sipri.org/sites/default/files/2024-06/yb24_summary_en_2_1.pdf)
- SIPRI. (2025). *Military expenditure database: Armaments, disarmament and international security*. [https://www.sipri.org/sites/default/files/2025-06/yb25\\_summary\\_en.pdf](https://www.sipri.org/sites/default/files/2025-06/yb25_summary_en.pdf)
- Solberg Söilen, K.S., (2016). Economic and industrial espionage at the start of the 21st century –Status quaestionis. *Journal of Intelligence Studies in Business*, 6(3), 51–64.
- Schweizer, P. (1993). *Friendly spies: How America's allies are using economic espionage to steal our secrets*. New York: Atlantic Monthly Press.
- Söilen, K.S. (2016) Economic and industrial espionage at the start of the 21st century–status quaestionis. *Journal of Intelligence Studies in Business*, 6 (3) pp. 51-64. DOI: <https://doi.org/10.37380/jisib.v6i3.196>
- Şeşen, Y., Kuzucuoğlu, A.H. (2021). Veri ve Bilgi Güvenliği Bağlamında İstihbarat Faaliyetleri. *Lamre Journal*, 2(2), 93-110.
- TASAM-Türk Asya Stratejik Araştırmalar Merkezi. (2016). *Küreselleşmenin boyutları ve etkileri*. [https://tasam.org/tr-TR/Icerik/211/kuresellesmenin\\_boyutlari\\_ve\\_etkileri](https://tasam.org/tr-TR/Icerik/211/kuresellesmenin_boyutlari_ve_etkileri)
- T.C. Cumhurbaşkanlığı Yatırım ve Finans Ofisi. (2024). *Savunma ve havacılık*. [https://www.invest.gov.tr/tr/sectors/sayfalar/defensandaerospace.aspx#:~:text=Savunma a%20sanayiinin%20i%C5%9F%20hacmi%20son,den%20fazla%20%C5%9Firket](https://www.invest.gov.tr/tr/sectors/sayfalar/defensandaerospace.aspx#:~:text=Savunma%20sanayiinin%20i%C5%9F%20hacmi%20son,den%20fazla%20%C5%9Firket)
- T.C. Cumhurbaşkanlığı Savunma Sanayi Başkanlığı. (2023). *2024–2028 savunma sanayi sektörel strateji dokümanı*. [https://www.ssb.gov.tr/Images/Uploads/MyContents/F\\_20240917164305314800.pdf](https://www.ssb.gov.tr/Images/Uploads/MyContents/F_20240917164305314800.pdf)
- T.C. Cumhurbaşkanlığı Savunma Sanayi Başkanlığı. (2025). *2024 yılı faaliyet raporu*. [http://www.sp.gov.tr/upload/xSPRapor/files/ut8N6+SSB\\_2024\\_Yili\\_Faal\\_Rap.pdf](http://www.sp.gov.tr/upload/xSPRapor/files/ut8N6+SSB_2024_Yili_Faal_Rap.pdf)
- Tezsever, S. (1999). Milli Güvenliğimiz İçinde İstihbarat. *İstanbul: İstanbul Üniversitesi Basımevi*.
- Toren, P. (2018, February 14). Some lessons from the Waymo (Alphabet) versus Uber theft of trade secret litigation. *Ipwatchdog*. <http://www.ipwatchdog.com/2018/02/14/waymouber-theft-trade-secret-litigation/id=93528/>
- Toriola, T. (2011). *Industrial Espionage or Competitive Intelligence*. Purdue University. [https://www.cerias.purdue.edu/assets/pdf/bibtex\\_archive/2011-10report.pdf](https://www.cerias.purdue.edu/assets/pdf/bibtex_archive/2011-10report.pdf).
- Törenli, N. (2007). Küreselleşen dünyada bilgi güvenliği ve Türkiye. III. İletişim Teknolojileri Ulusal Sempozyumu (ITUSEM) Bildiriler Kitabı, Adana.
- Urhal, Ö. (2008). *Kamu güvenliği açısından istihbarat ve örgütlü suçlar*. Adalet Yayınevi.
- U.S. Department of Justice. (2025, September 15). Former defence contractor sentenced to over 10 years in prison for attempted espionage. Office of Public Affairs. <https://www.justice.gov/opa/pr/former-defence-contractor-sentenced-over-10-years-prisonattempted-espionage>
- U.S. Department of Justice. (2025, November 5). Fiber laser expert convicted by federal jury of economic espionage and theft of trade secrets.

<https://www.justice.gov/opa/pr/fiber-laser-expert-convicted-federal-jury-economic-espionage-and-theft-trade-secrets>

Verizon Business. (2025). *Data breach investigations report-2025: Executive summary*.

<https://www.verizon.com/business/resources/reports/2025-dbir-executive-summary.pdf>

Yılmaztürk, A. (2023). Türkiye’de savunma sanayi sektörü ve ekonomi üzerindeki etkisinin değerlendirilmesi. *Enderun Dergisi*, 7(2), 139–165.

Wagner, R. E. (2012). Bailouts and the potential for distortion of federal criminal law: Industrial espionage and beyond. *Tulane Law Review*, 86(5), 1017–1055.

Waziri, K. M., & Yerima, T. F. (2011). Industrial espionage and intellectual property rights protection: How legal is it legal. *International Journal of Sustainable Development*, 4(3), 1-22.

## Energy Espionage 5.0: Corporate Intelligence, Critical Metals, and Artificial Intelligence

Tuğba Çakıroğlu<sup>1,\*</sup> 

### Abstract

The strategic significance of critical minerals in the global energy transition has intensified, transforming competition among states and corporations into a geopolitical security challenge. This study examines intelligence practices and industrial espionage strategies targeting critical mineral resources. It analyses how lithium, cobalt, nickel, and rare earth elements influence national security agendas and supply-chain protection policies, focusing on cyber-espionage, insider information leakage, and covert data acquisition methods. The paper evaluates the use of satellites, high-resolution drones, AI-driven mineral exploration systems, and remote-sensing technologies for resource detection and production monitoring. It further explores the capacity of multinational mining firms to shape energy diplomacy and public policy, positioning them as strategic non-state actors. Additionally, the research investigates open-source intelligence systems, AI-based commercial data platforms, and competitive intelligence tools as mechanisms for gaining informational advantage in the energy sector. Findings indicate that the convergence of digital surveillance technologies and geoeconomic rivalry has created a new security architecture, where artificial intelligence emerges as a decisive force in mineral geopolitics.

**Keywords:** Energy espionage, critical minerals, artificial intelligence, satellite surveillance, drone systems, corporate intelligence, OSINT.

---

Received: 16.12.2025

Accepted: 04.06.2026

Research Article

<sup>1</sup> Turkish Energy, Nuclear and Mineral Research Agency, Ankara, Türkiye

\* Corresponding author.

✉ safrantugba78@gmail.com

**Cite this article as:**

<https://doi.org/10.65834/jdsi.12.49>

Çakıroğlu, T. (2026). Energy Espionage 5.0: Corporate Intelligence, Critical Metals, and Artificial Intelligence. *Journal of Defence and Security Industries: Strategy and Technology* 1(2), 27-58.

## 1. Introduction

The global energy transition has accelerated the shift toward sustainable and carbon-neutral energy systems, significantly increasing the strategic demand for critical minerals. Minerals such as lithium, cobalt, nickel, and Rare Earth Elements (REEs) have become essential components in the production of advanced battery technologies, defence systems, electric vehicles, and artificial intelligence-based hardware (IEA, 2021). This trend has not only intensified economic competition but also exacerbated geopolitical power struggles, driving both states and multinational corporations to design comprehensive overt strategies and covert mechanisms aimed at securing critical mineral supply chains. As the security of these supply networks becomes central to national economic and strategic policies, there has been a notable increase in state-led interventions, regulatory measures, and competitive actions aimed at diversifying and protecting access to critical mineral resources (Nakano, 2021).

To contextualize the concept of Energy Espionage 5.0, it is useful to briefly outline the historical evolution of energy-related intelligence practices. Energy espionage initially emerged in the twentieth century as a state-centred activity focused on securing access to oil reserves, pipeline routes, and strategic chokepoints, often relying on human intelligence and diplomatic reporting. In subsequent phases, particularly during the late Cold War and post-Cold War periods, energy intelligence expanded to include industrial espionage, corporate infiltration, and the protection of proprietary extraction technologies. With the digitalization of energy systems and the globalization of supply chains in the early twenty-first century, cyber espionage and open-source intelligence became integral to monitoring energy markets and infrastructure. The current phase conceptualized in this study as Energy Espionage 5.0 represents a qualitative shift characterized by the convergence of artificial intelligence, commercial satellite data, corporate intelligence, and large-scale data analytics, thereby redefining both the actors and the instruments of energy-related intelligence.

The intelligence dimension of competition in the energy and resource sectors encompasses a broad range of techniques, from classical espionage practices to advanced digital surveillance and analytical tools. Modern remote sensing technologies—including satellite imagery, Unmanned Aerial Vehicles (UAVs), and other airborne platforms—combined with artificial intelligence-based data processing and pattern recognition software, play an increasingly critical role in detecting resource deposits, monitoring production activities, and generating high-resolution geospatial intelligence that enhances situational awareness for both corporate and state actors. These developments reflect the growing integration of geospatial analytics and machine learning in strategic resource governance and decision-making processes (Yang et al., 2024). In this context, energy espionage has emerged as a new strategic domain with implications for both national security and economic sustainability. This study examines the intelligence and technological aspects of the competition over critical minerals and analyses the interplay between corporate, geopolitical, and digital instruments.

This research employs descriptive analysis and document review methods within the scope of a qualitative research design. Academic articles, reports published by international

organizations, defence strategy documents, and the corporate intelligence literature were first surveyed. Sources were selected through Google Scholar, Scopus, and Web of Science indexes. During the review process, key concepts were identified within the thematic framework of energy security, the geopolitics of critical minerals, cyber espionage, OSINT, and artificial intelligence–supported monitoring technologies.

The selected sources were classified through content analysis, and three main research questions were formulated:

1. How does competition over critical minerals shape international security dynamics?
2. How are power relations between mining corporations and state actors evolving?
3. How do artificial intelligence and surveillance technologies transform energy intelligence practices?

A systematic literature approach was employed in compiling the data, with a particular emphasis on studies published in the last decade to evaluate recent technological and strategic developments. Several analyses were supported by up-to-date industry reports and policy documents. The results suggest that competition over critical minerals has increasingly become a central issue in global geopolitics and economic security debates. The concentration of rare earth element production and processing capacity in a limited number of countries particularly China has prompted the United States and the European Union to treat mineral supply chains as a strategic priority and to develop policies aimed at diversifying and securing access to these resources (IEA, 2021). Consequently, alternative supply routes have been developed, strategic stockpiling programs have been strengthened, and mining investments have increasingly been utilized as geopolitical instruments.

Multinational mining corporations are emerging not merely as economic entities but increasingly as significant geopolitical actors with the capacity to shape international relations and strategic policy outcomes. Through their involvement in critical mineral extraction and processing, these firms influence diplomatic engagements, contribute to infrastructure financing, and intersect with national energy security strategies. Their strategic investments and operations are embedded within broader geopolitical frameworks, where access to resources such as lithium, cobalt, and rare earths has implications for technological competition, supply-chain resilience, and state foreign policy priorities (Müller, 2023). As a result, corporate diplomacy and the influence of non-state actors in security and policy-making processes are expanding.

From a technological standpoint, strategic competition in energy and minerals has been increasingly shaped by digitization, as satellite remote sensing, drone-based field surveys, and artificial intelligence–enhanced analytical tools transform how reserves are estimated and production performance is monitored. These advanced technologies enable more transparent mapping of geological features and real-time data acquisition, while machine learning models accelerate the processing and interpretation of large geospatial datasets, thereby increasing both the efficiency and competitive intensity of exploration and monitoring activities in the mining

sector (Yang, et al., 2024). OSINT platforms and commercial data analytics firms enable the acquisition of strategic insights through market movements, logistics networks, and corporate datasets. This dynamic deepens debates surrounding data privacy, energy security, and national sovereignty.

The findings of this research indicate that artificial intelligence–based intelligence methods will assume an increasingly decisive role in the geopolitics of energy and minerals in the coming period. The interplay between digital surveillance and strategic data analytics will redefine geopolitical competition and bring new security risks to the forefront.

This study adopts a qualitative and analytical research design grounded in geopolitics, security studies, and political economy. Rather than relying on empirical field data, the analysis is based on secondary data sources in order to conceptualize and examine the emerging phenomenon of Energy Espionage 5.0. The methodological approach aims to synthesize fragmented discussions across energy security, critical mineral geopolitics, corporate intelligence, and artificial intelligence into an integrated analytical framework.

The data utilized in this study consist of peer-reviewed academic literature, policy documents, official government publications, reports by international organizations, and publicly available cybersecurity and intelligence assessments. Sources were selected based on their relevance to critical minerals, energy security, industrial espionage, satellite surveillance, and artificial intelligence–driven intelligence practices. Particular emphasis was placed on materials published by institutions such as the International Energy Agency, national security agencies, and leading cybersecurity firms to ensure analytical validity and contemporary relevance.

The analysis was conducted through qualitative content analysis and thematic synthesis. Selected sources were systematically examined to identify recurring patterns related to resource competition, surveillance technologies, corporate intelligence practices, and data sovereignty. These themes were subsequently interpreted through a geopolitical and security-oriented analytical lens to demonstrate how artificial intelligence and digital surveillance technologies reshape traditional forms of energy-related industrial espionage. This structured analytical process enables the transformation of descriptive material into a coherent theoretical interpretation.

### **1.1. Mineral Geopolitics and Industrial Espionage**

The global economy has historically been shaped by access to resources. From antiquity to the present, power dynamics have revolved around gold, silver, spices, oil, and, in today's context, strategic metals and rare earth elements (Rare Earth Elements, REE). In the twenty-first century, developments such as the transformation of energy systems, the expansion of the knowledge-based economy, and the rapid advancement of artificial intelligence technologies have carried traditional energy competition into a new strategic dimension. Competition is no longer limited to oil and natural gas resources; critical minerals including lithium, cobalt,

nickel, graphite, titanium, palladium, platinum, and neodymium have increasingly emerged as essential components of global power projection (Gielen, 2021).

In this regard, mineral geopolitics has become an integral component of national security doctrines and economic strategies, while industrial espionage activities aimed at controlling the supply chains of critical minerals have intensified. The concept of “Energy Espionage 5.0” represents a new paradigm of competition shaped by artificial intelligence–supported data analytics, digital surveillance, corporate penetration techniques, and hybrid operations targeting global supply chains.

The geopolitical positioning of critical minerals carries strategic significance due to their use in high value–added sectors such as energy storage, the defence industry, telecommunications, semiconductor manufacturing, artificial intelligence hardware, and space technologies. The International Energy Agency (IEA) projects that the demand for lithium, cobalt, and nickel will increase by a factor of 40, 20, and 25, respectively, by 2040, driven by the expansion of electric vehicles and battery storage systems (Kim, et al, 2021).

A review of global production structures reveals that the geographical distribution of critical minerals is highly asymmetric. For example (U.S. Geological Survey, 2024):

- 56% of the world’s lithium reserves are located in Chile,
- 70% of global cobalt production takes place in the Democratic Republic of Congo,
- More than 60% of rare earth elements are processed in China.

This asymmetric distribution brings supply chain dependencies and geoeconomic vulnerabilities to the forefront. States simultaneously employ diplomatic, financial, and intelligence instruments to secure access to strategic minerals. This approach reflects a shift from Joseph Nye’s concepts of “soft power” and “hard power” toward a “smart power” model (Nye, 2011).

Competition over strategic minerals involves not only access to physical resources but also processing technologies, advanced metallurgy, R&D infrastructure, and data-driven exploration capabilities. Accordingly, the geopolitics of minerals is situated at the intersection of natural resource diplomacy and technology geopolitics.

The development of industrial espionage and competition over strategic resources encompasses actions taken by state and private actors to acquire competitors’ critical information and technologies for the purpose of gaining commercial or strategic advantage (Holmström, 2020). Throughout the 20th century, espionage in the energy sector primarily targeted oil reserves, pipeline infrastructures, and refinery technologies. Today, however, the focus has shifted toward a digitalized information ecosystem involving mineral exploration data, material science innovations, and battery technologies.

Table 1 summarizes the principal espionage methods currently employed in strategic resource competition.

**Table 1.** Contemporary industrial espionage tactics in strategic resource competition.

| <i>Espionage Method</i>                                       | Description                                                                                                                                                                                                                                                                     |
|---------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Digital penetration and cyber intelligence</i>             | The monitoring of unauthorized access attempts targeting corporate networks and industrial control systems is conducted to gather information and enhance adversarial influence; in this context, cyber intelligence focuses on identifying threat vectors and vulnerabilities. |
| <i>Corporate infiltration and insider agents</i>              | The covert acquisition of strategic information through individuals positioned or placed within a company to gain competitive advantage targets organizational security vulnerabilities; corporate infiltration operates on exploiting these internal weaknesses.               |
| <i>Supply chain manipulation</i>                              | It refers to the disruption of supply security by manipulating material flows, supplier relations, or logistics processes through informational or physical interventions, thereby placing competitors at a disadvantage.                                                       |
| <i>Information leakage through academic partnerships</i>      | It refers to the transfer of scientific data obtained through university–industry collaborations and research projects to third parties outside ethical and legal frameworks, or its indirect use for profit-driven purposes.                                                   |
| <i>Intelligence operations disguised as legal investments</i> | It refers to the intertwining of legitimate economic activities—such as direct investment, joint ventures, or acquisitions—with intelligence operations aimed at obtaining strategic information and securing competitive advantage.                                            |
| <i>The theft of geological modelling data</i>                 | The seizure of databases containing reserve estimates, sample analyses, and model configurations weakens companies’ exploration investments and market positioning.                                                                                                             |

Source: Nasheri (2005), Hannas et al. (2013), Noonan and Archuleta (2008), Hou and Wang (2020).

A publicly documented example that illustrates the evolution of industrial espionage in the energy domain is the cyber operation widely known as *Stuxnet*. Although originally targeting Iran’s nuclear infrastructure, this operation demonstrated how cyber capabilities could be used to infiltrate industrial control systems, manipulate operational data, and disrupt strategic energy-related assets without direct physical intervention (Langner, 2011). The Stuxnet incident is widely regarded as a turning point in revealing the feasibility of digitally driven sabotage and intelligence operations against critical energy infrastructures.

More recent publicly reported cases indicate that similar cyber-espionage patterns have targeted mining companies involved in critical minerals such as lithium, rare earth elements, and battery technologies. Reports by cybersecurity firms and governmental institutions document attempts to access geological survey data, mineral processing technologies, and

supply chain intelligence, underscoring the relevance of Energy Espionage 5.0 as a data-driven and hybrid form of strategic competition (U.S. Department of Justice, 2023).

The competition between China and the United States over technology and critical raw materials is extensively examined in the academic literature as a key arena where strategic rivalry and industrial espionage have intensified. Scholars note that China's dominant position in the mining, processing, and supply of rare earth elements and other strategic minerals which are essential inputs for advanced technologies and defence systems has prompted concern in Washington about risks to intellectual property, supply chain vulnerabilities, and the potential use of covert methods to secure technological advantage. As a result, economic security debates increasingly frame resource and technology competition as intertwined domains of the broader Sino-American strategic rivalry (Stango, 2024). Economies such as the European Union and Japan are also developing intelligence-driven policies to ensure the security of critical mineral supplies.

Artificial intelligence-driven analytical systems are increasingly transforming mineral exploration and extraction processes. The integration of satellite imagery, extensive geophysical datasets, machine-learning-based geological modelling, and automated or remotely operated mining technologies is progressively transforming the sector into a data-driven and digitally governed industry. Although such technological advancements improve efficiency and support more informed decision-making in resource exploration and production processes, the increasing reliance on digital infrastructures simultaneously heightens the exposure of mining systems to cyber espionage and data-related security vulnerabilities (Marshall et al., 2015). Table 2 presents the key domains in which artificial intelligence intersects with intelligence activities in the mining sector.

**Table 2.** Artificial intelligence applications and intelligence purposes in the context of energy espionage in the mining sector.

| <i>Domain</i>                         | <i>Purpose</i>                                                        |
|---------------------------------------|-----------------------------------------------------------------------|
| <i>Digital geological data mining</i> | New reserve estimation and the theft of competitors' exploration data |
| <i>Machine learning</i>               | Supply chain risk analysis and speculation                            |
| <i>Cyber intelligence</i>             | Infiltration into industrial control systems                          |
| <i>Autonomous operations</i>          | The risk of remote sabotage or manipulation                           |
| <i>Deepfake technologies</i>          | Social engineering attacks in executive communications                |

Source: Gaber et al. (2021), Frankovits and Mirsky (2023), and Pedersen et al. (2025).

This situation demonstrates that information security in the mining industry is not confined to the physical domain; rather, it requires a broad security ecosystem ranging from data centres to academic collaborations.

State–corporate cooperation and the new geoeconomic competition over access to critical minerals have shifted from classical state-cantered policies toward hybrid models supported by corporations. State-backed mining initiatives conduct global competition through financial

instruments, geostrategic partnerships, and international legal mechanisms (Bridg, 2008). For example:

- The United States provides federal support for battery metals under the Defence Production Act.
- China controls the entire rare earth elements supply chain under its “Made in China 2025” program.
- The European Union has developed strategic reserve policies through the Critical Raw Materials Act.

These policies demonstrate the rise of state capitalism in the energy sector and the incorporation of corporate intelligence units into national security strategies. From ethical, legal, and security perspectives, the growing competition over energy resources increasingly brings debates in international law to the forefront. Cyberattacks, corporate infiltration, and data manipulation create risks that have the potential to undermine both the stability of the international trade system and the functioning of democratic oversight mechanisms (Clarke & Knake, 2021). Ethical questions are particularly concentrated along the following axes:

1. How does the use of artificial intelligence in resource competition transform state sovereignty?
2. How will human rights be protected in regions where critical minerals are extracted?
3. Where is the boundary between industrial espionage and the protection of strategic information?

These questions make it necessary to evaluate mineral geopolitics not only as an economic and strategic domain but also within the context of human security, sustainability, and ethics. The era of Energy Espionage 5.0 represents a period in which competition over natural resources has become digitalized, mineral geopolitics has converged with technology geopolitics, and artificial intelligence has emerged as both a source of power and vulnerability. As demonstrated in this section, critical minerals carry not only economic value but also lie at the centre of national security strategies, global political balances, and corporate power relations.

As competition is expected to intensify in the coming years, the development of ethical, legal, and international cooperation mechanisms will be critical for achieving sustainable and equitable resource governance.

## **1.2. Satellites, Drones, and Resource Monitoring Technologies**

Across the history of energy geopolitics and competition over mineral resources, states have consistently relied on advanced surveillance and reconnaissance capabilities to secure valuable extraction sites and protect strategic frontiers. From early aerial reconnaissance missions in the mid-20th century such as the U.S. CORONA satellite program, which pioneered space-based photographic intelligence to the widespread use of satellites, unmanned aerial vehicles (UAVs/drones), and data-driven monitoring systems in the twenty-first century, surveillance technologies have been integral to shaping strategic economic and military dynamics.

Contemporary remote sensing and aerial surveillance platforms now serve as essential tools for gathering real-time intelligence, enhancing situational awareness, and safeguarding critical resource infrastructure (Raghuvanshi et al., 2025).

Within the framework of contemporary strategic resource competition, the integration of satellite observation systems, geospatial analytics, unmanned platforms, and artificial intelligence-enhanced data fusion has transformed how critical mineral deposits are detected, monitored, and protected. For strategic resources such as lithium, cobalt, nickel, graphite, and rare earth elements, space-borne remote sensing and AI-assisted processing provide real-time, high-resolution geological insights that reduce dependence on traditional field campaigns and offer continuous, verifiable data independent of conventional diplomatic or ground intelligence methods. As such, these advanced monitoring technologies have emerged as pivotal instruments in securing resource supply chains and reinforcing geopolitical influence in the twenty-first century (Garcia-del-Real & Alcaráz, 2024).

Although satellite technologies were initially developed for military purposes during the 1960s, they are now widely utilized in areas such as natural resource management, environmental monitoring, mine planning, and national security (Prost, 2025). High-resolution optical satellites, hyperspectral imaging systems, synthetic aperture radar (SAR), and thermal sensors provide powerful tools for detecting subsurface resources and observing operational activities.

Hyperspectral satellites, in particular, provide critical data for analysing rock composition, mineral geochemistry, and detecting land-use changes (Yan, et al., 2022). This allows mineral exploration processes to be conducted much faster, more cost-effectively, and more comprehensively than classical geological field studies.

Table 3 outlines the principal applications of satellite and remote sensing data in mineral exploration and resource management.

**Table 3.** Applications of satellite and remote sensing data in mineral exploration and management.

| <i>Application Area</i>                       | <i>Operational Description</i>                                                                                                                                                                                                      |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Identification of metallogenetic zones</i> | The analysis of geological, geochemical, and remote sensing data to determine mineral formation processes, ore-forming zones, and potential reserve areas through scientific modelling.                                             |
| <i>Expansion analysis of open-pit mines</i>   | It refers to the assessment of the growth rate of open-pit areas, production expansion trends, and land-use changes using satellite imagery, drone data, and time-series analyses.                                                  |
| <i>Detection of illegal mining activities</i> | It refers to the early detection of unauthorized excavation sites, concealed mining equipment, and illicit production lines through satellite, radar, and thermal imaging technologies.                                             |
| <i>Monitoring of environmental impacts</i>    | It refers to the continuous assessment of the impacts of mining operations on ecosystems—such as vegetation loss, water resource degradation, soil alteration, and carbon emissions—using remote sensing and environmental sensors. |

**Table 3.** (Continued)

|                                                                      |                                                                                                                                                                                                                                                           |
|----------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Identification of geological fault lines and alteration zones</i> | It refers to the identification of subsurface fault lines, hydrothermal alteration zones, and structurally significant features with high ore-forming potential using spectral analysis, geophysical data, and geographic information systems.            |
| <i>Analysis of infrastructure development</i>                        | It refers to the monitoring of roads, power lines, storage facilities, processing plants, and other operational infrastructure in mining regions through remote sensing methods to evaluate capacity expansions and indicators of strategic preparedness. |

Source: Bedini (2017), Peyghambari and Zhang (2021), and Balaniuk et al. (2020).

Hyperspectral imaging systems are capable of resolving and differentiating mineralogical components across hundreds of narrow spectral bands, enabling the detailed identification of surface composition, alteration zones, and mineral distributions. This advanced remote sensing technique is particularly valuable in the exploration of minerals such as lithium and copper, whose surface signatures can be detected and mapped with high spectral fidelity, thereby enhancing the accuracy and efficiency of targeting prospective ore bodies (Van der Meer et al., 2021).

In satellite intelligence and energy espionage, satellite systems provide not only reconnaissance but also strategic advantages in espionage activities. In the context of geopolitical competition, satellite data can:

Within the context of geopolitical competition, satellite data has become an important tool for monitoring strategic activities related to natural resources and industrial production. Remote sensing technologies enable the observation of mining investments carried out by rival states, the tracking of supply chain infrastructures such as railways and pipelines, and the surveillance of port and logistics operations that facilitate the global movement of critical materials. In addition, satellite imagery allows the early identification of newly planned mining sites and provides indirect indicators of industrial production levels through the analysis of infrastructure expansion, transportation intensity, and energy-related facilities. Such capabilities make satellite data a valuable component of contemporary geoeconomic intelligence and resource governance (Finer, Jenkins, Pimm, Keane & Ross, 2018).

For instance, China and the United States utilize space-based observation systems to monitor critical mineral regions in Africa and South America in order to detect potential investment opportunities at an early stage. Such practices point to the emergence of a new geoeconomic framework in which cooperation between states and corporations increasingly overlaps with intelligence activities. In parallel, drones have quickly become indispensable operational instruments within the mining sector. Supported by cloud-based data processing, lidar sensors, and high-resolution imaging systems, these platforms have significantly transformed both exploration processes and security operations (Shahmoradi et al., 2020). Table 4 summarizes the principal operational applications of drone technology in the mining sector.

**Table 4.** Operational applications of drone technology in the mining sector.

| Drone Application Area                                             | Operational Description                                                                                                                                                                                      |
|--------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Topographic mapping</i>                                         | It refers to the three-dimensional modelling of elevation data, land morphology, and surface changes using satellite-, lidar-, and drone-based measurement techniques, and their use in mining planning.     |
| <i>Open-pit volume calculations</i>                                | It refers to calculating the volumetric changes of excavated areas through photogrammetry, lidar scanning, and geotechnical analyses, thereby optimizing production efficiency and reserve utilization.      |
| <i>Blast planning</i>                                              | It refers to the analysis of geological structure, explosive types, rock resistance coefficients, and site safety data in order to design safe, controlled, and efficient blasts in ore-breaking operations. |
| <i>Environmental monitoring</i>                                    | It refers to the monitoring of water, air, soil, and noise parameters through sensor networks, remote sensing, and field measurements to evaluate the environmental impacts of mining activities.            |
| <i>Environmental risk control (erosion, dust, water pollution)</i> | It refers to the implementation and monitoring of environmental protection strategies such as erosion barriers, dust suppression practices, water filtration systems, and biotechnical measures.             |
| <i>Security patrols</i>                                            | It refers to the regular surveillance and threat detection activities conducted by physical security personnel and unmanned patrol systems around critical equipment, boundary lines, and operational areas. |

Source: Shahmoradi et al. (2020) and Said et al. (2020).

In such environments, Table 5 illustrates how drones function as real-time surveillance instruments across various operational dimensions.

**Table 5.** Drone-Based surveillance applications in mining security and operations.

| Surveillance Application                      | Operational Description                                                                                                                                                                                                                                    |
|-----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Strengthening border and site security</i> | It refers to the reduction of unauthorized entry and sabotage risks along mining sites and boundary lines through the integration of physical barriers, access control systems, patrol arrangements, and electronic surveillance.                          |
| <i>Detecting unauthorized intrusions</i>      | It refers to the early detection of unauthorized intrusions around site boundaries and perimeters using motion sensors, thermal cameras, radar systems, and data analysis algorithms, triggering appropriate response mechanisms.                          |
| <i>Detecting infrastructure damage</i>        | It refers to the early detection of damage and deterioration in critical infrastructure components—such as pipelines, bridges, and storage facilities—through structural health monitoring (SHM), IoT sensors, satellite imagery, and regular inspections. |

**Table 5.** (Continued)

|                                        |                                                                                                                                                                                                                                      |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Monitoring extraction equipment</i> | It refers to ensuring production continuity by monitoring the operational status, maintenance needs, and potential failures of machines through equipment telemetry data, operational performance indicators, and remote monitoring. |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Source: Ajayi & Ikemefuna, 2025

Artificial intelligence algorithms integrated into drone systems enhance autonomous security capacity by performing anomaly detection and threat classification. However, alongside their security benefits, drones also carry a significant potential for use in espionage. Table 6 illustrates how drones and artificial intelligence can be leveraged for strategic intelligence and espionage purposes in the mining sector.

**Table 6.** Drone and ai applications in strategic intelligence and mining espionage.

| <i>Espionage Application</i>               | <i>Operational Description</i>                                                                                                                                                                                                                                                                |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Monitoring extraction rates</i>         | It refers to analytical activities aimed at gaining strategic competitive advantage by monitoring production volumes, ore extraction rates, and reserve depletion dynamics in real time through satellite, sensor, and data analysis tools.                                                   |
| <i>Documenting facility designs</i>        | It refers to the examination and archiving of engineering layouts, production lines, and technical infrastructure related to mining facilities through imaging, infiltration, or open-source data tracking.                                                                                   |
| <i>Analysing production infrastructure</i> | It refers to the assessment of the operational capacity and vulnerabilities of critical components—such as power lines, processing equipment, transportation networks, and storage systems—using technical, logistical, and digital tools.                                                    |
| <i>Monitoring personnel movements</i>      | It refers to the analysis of operational priorities and security vulnerabilities by tracking workforce distribution in the field, the activity intensity of specialist teams, and the movement patterns of critical personnel.                                                                |
| <i>Collecting geological data</i>          | It refers to the collection of samples, sensor outputs, radar measurements, artificial intelligence-derived data, and geophysical information from the field to evaluate reserve quality, ore body structure, and geological formation characteristics for scientific and strategic purposes. |

Source: Shahmoradi et al. (2020), Sihag et al. (2023), and Alsadie (2025).

Artificial intelligence algorithms embedded in unmanned aerial systems enhance autonomous security capacities by enabling advanced anomaly detection, behavioural classification, and predictive threat analysis. Despite these operational benefits, drones also pose significant risks for misuse in covert intelligence gathering. Both competing firms and state-linked actors can exploit such platforms to carry out a range of surveillance and strategic data collection tasks, including real-time monitoring of extraction and production metrics; documentation of plant layouts and critical infrastructure; assessment of operational vulnerabilities in processing and logistics networks; analysis of personnel movement patterns for security profiling; and

collection of geological information to infer reserve characteristics (Ocelik, et al., 2024). This situation has made drone countermeasures a necessity in corporate security strategies. Advanced sensing technologies such as radar-based detection, radio-frequency monitoring, and laser-based scanning systems are increasingly applied at critical mineral sites to enhance situational awareness and security. The vast volume of data generated by satellite and unmanned aerial systems necessitates sophisticated computational frameworks for effective processing, classification, and interpretation. By integrating machine learning, deep learning, and Geographic Information Systems (GIS), contemporary mineral exploration and monitoring workflows have become significantly more automated and data-driven, improving both target identification and operational efficiency (Yang et al., 2024). Data fusion encompasses the following components:

1. Hyperspectral analysis + topographic mapping
2. SAR radar + optical imagery analysis
3. Satellite data + drone imagery + seismic data integration
4. AI-based reserve estimation

These systems shorten mineral exploration timelines and minimize investment risks. The same technologies can also be used for espionage-oriented data collection. For instance, the production volume, stockpile areas, and logistical traffic of a mining company can be analysed to conduct an economic assessment. This blurs the boundaries between energy security and intelligence activities.

Although satellite observations are often categorized as a form of Open-Source Intelligence (OSINT), the unauthorized use of such imagery to reveal proprietary or trade-secret information raises significant ethical and legal concerns. The expanding availability of high-resolution commercial satellite data has outpaced existing legal frameworks governing privacy, intellectual property, and competitive information rights, prompting debate over how OSINT practices should be regulated to balance transparency with the protection of confidential data (Pedram, Chandler & Georgiades, 2025). Space law generally regulates data collection activities at the state level; however, the rise of commercial satellite companies has created grey areas.

Regulatory frameworks governing drone operations remain insufficient in numerous jurisdictions, particularly within many developing nations where critical mineral sites are concentrated. In these contexts, weak or poorly enforced aviation laws create legal gaps that can be exploited for unauthorized drone flights, including the illicit collection of aerial data and surveillance over sensitive areas. The absence of comprehensive legal structures and enforcement mechanisms underscores the challenges that developing countries face in regulating unmanned aerial systems, which in turn facilitates operational vulnerabilities and unauthorized uses (Ayamga, Tekinerdogan & Kassahun, 2021).

The governance and ownership of data generated through resource monitoring systems has become a matter of strategic importance for states. Control over geological and mineral

resource data—particularly when managed by foreign companies or external technological infrastructures—can create vulnerabilities in terms of national security and economic sovereignty. For this reason, many countries increasingly treat critical geological information as a strategic asset and seek to ensure that its collection, storage, and analysis remain under national oversight (Bradshaw, 2010). Satellites and drones are invisible yet decisive actors in energy and mineral competition. While these technologies offer revolutionary opportunities for mineral exploration and security, they also introduce new dimensions to industrial espionage and intelligence operations.

The increasing reliance on commercial satellite imagery in critical mineral exploration and energy security has intensified debates surrounding data sovereignty. Unlike traditional state-operated reconnaissance systems, commercial satellite platforms operate across jurisdictions, enabling foreign actors to collect, process, and monetize geospatial data related to strategic resources without the explicit consent of the territorial state. This raises fundamental legal questions regarding ownership, control, and jurisdiction over resource-related data generated beyond national airspace but directly affecting national economic security (Jakhu, Pelton, & Nyampong, 2017).

From an ethical perspective, the cross-border use of commercial satellite data challenges existing principles of state sovereignty and informed consent, particularly in developing countries where critical mineral sites are located. The asymmetry between technologically advanced data-collecting states and data-subject states may result in extractive data practices that mirror historical patterns of resource exploitation, albeit in digital form. Consequently, data sovereignty has emerged as a core concept linking satellite governance, energy security, and ethical resource management in the era of Energy Espionage 5.0 (Fischer, 2023).

In the context of Energy Espionage 5.0, the satellite–drone–artificial intelligence triad shifts strategic resource oversight toward a cyber-physical security axis. In the future, these technologies are expected to become more autonomous, more intelligent, and more covert. This process makes the urgent development of ethical, legal, and diplomatic regulations essential; otherwise, the deepening of global vulnerabilities becomes inevitable.

### **1.3. The Geostrategic Role of Mining Companies**

One of the fundamental pillars of the global order lies in the capacity of states to project economic and military power through control and utilization of natural resources. Since the onset of the industrial era, minerals have functioned as concrete indicators of national power: coal and iron enabled the expansion of the locomotive industry, hydrocarbons shaped the structure of modern geopolitics, and in the contemporary period, strategic metals and rare earth elements have gained critical importance for artificial intelligence, defence systems, and advanced technological production (Haglund, 2025).

In this context, mining corporations have transcended their traditional roles as mere economic actors to become influential participants in shaping national security policies, geoeconomic

competition strategies, and the global balance of power. Their extensive financial resources, international investment networks, control over strategically vital mineral assets, and integration into geopolitical supply chains grant these firms a geostrategic significance traditionally associated with states. Consequently, major mining companies are increasingly embedded within national and transnational security frameworks, reflecting their role in securing supply chains critical to defence, energy transition, and high-technology sectors (Critical Minerals and Great Power Competition, SIPRI, 2024).

Although the geopolitical significance of minerals has transformed over time, their strategic importance has remained constant. During the nineteenth century, coal strengthened the United Kingdom's global dominance by supporting steam-powered industry and steel production, while in the early twentieth century petroleum resources played a decisive role in the rise of the United States as a major global power (Dannreuther, 2013). In the 21st century, lithium, cobalt, nickel, and rare earth elements have become essential inputs for the digital economy, electric vehicles, defence industries, and artificial intelligence infrastructures.

Multinational mining corporations have increasingly taken on roles that extend beyond traditional economic functions to encompass strategic political and security influence in the regions where they operate. Large extractive firms such as Anglo American, Rio Tinto, BHP, Vale, and Glencore are deeply embedded in global value chains and geopolitical competition, often shaping public policy, local governance, and state–corporate relations through their investment decisions and interactions with host governments. In some contexts, these firms assume quasi-governmental functions by influencing policy outcomes, negotiating regulatory frameworks, and embedding themselves in political-economic structures in resource-rich territories (Bechtum, 2024).

Since the latter half of the twentieth century, resource nationalism has gained prominence, leading numerous states in Latin America, Africa, and Asia to assert greater control over their mineral sectors through nationalization and restrictive ownership policies. Nonetheless, the forces of globalization and the liberalization of capital markets have subsequently facilitated the reintegration of foreign and transnational mining firms into these economies, enabling them to regain influence and operational strength within global extractive networks. This interplay between nationalist policies and global capital flows highlights the evolving nature of state-firm relations in the extractive industries (Radetzki, 2024).

Today, geopolitical competition has transformed into a multi-layered struggle involving states, transnational corporations, and financial consortia. For this reason, the functions of mining companies are not limited to the economic sphere:

- Geographical exploration capabilities = geopolitical projection
- Supply chain control = strategic dependency
- Technology and R&D investments = competitive advantage
- Intelligence activities = an element of information warfare
- Diplomatic networks = economic influence

In the current global arena, competition for critical mineral supply chains has evolved into a complex geopolitical contest that engages not only states but also powerful non-state and corporate actors. Control over these resources has shifted from mere economic competition to a broader strategic framework where multinational companies and financial consortia function as integral parts of international supply networks. Their capacities in exploration, supply-chain governance, technological investment, intelligence provision, and diplomatic engagement embed them in the geopolitical strategies of states, effectively positioning dominant mining firms as actors with influence comparable to traditional sovereignty mechanisms in international relations (Vivoda et al., 2024). The strategic competition over critical minerals is increasingly reflected in national and multilateral strategy documents, where key actors such as the United States, the European Union, China, Japan, and South Korea have elevated rare earth elements and battery metals to matters of national security and economic resilience. These governments have introduced dedicated policies, strategic partnerships, and cooperative frameworks to diversify supply chains, strengthen domestic processing capacities, and pursue joint mining and investment projects with allied partners. Such initiatives form part of broader efforts to reduce dependence on dominant suppliers and embed critical mineral industries within national security priorities (European Parliamentary Research Service, 2023; U.S.–Japan Framework for Securing Critical Minerals, 2025). Table 7 provides an overview of selected national strategies and state-supported mining instruments adopted by major economies to secure critical mineral supply chains.

**Table 7.** National strategies and state-supported mining instruments for critical minerals.

| <i>Country</i>             | <i>Strategy</i>                       | <i>Mining Instrument</i>                      |
|----------------------------|---------------------------------------|-----------------------------------------------|
| <i>United States (USA)</i> | Defence Production Act                | State-backed credit for mining investments    |
| <i>European Union (EU)</i> | Critical Raw Materials Act            | Joint strategic fund and diplomatic support   |
| <i>China</i>               | Made in China 2025                    | State-owned enterprises + global acquisitions |
| <i>Australia</i>           | Critical Minerals Strategy            | Supply diplomacy independent of China         |
| <i>Canada</i>              | Critical Minerals Infrastructure Fund | Northern mineral corridor investments         |

Source: European Commission (2024), IEA (2023)

In this context, states deploy a variety of instruments such as financial diplomacy, infrastructure development support, international legal agreements, military security assurances, and intelligence cooperation to secure access to critical mineral resources that are vital for economic and strategic objectives. This multifaceted approach is reflected in initiatives like the Minerals Security Partnership and bilateral frameworks aimed at strengthening resilient and trusted supply chains, alongside diplomatic and regulatory efforts that align geopolitical interests with sustainable resource access (Minerals Security Partnership, 2022).

These units do not merely collect economic information; they gather and process a wide range of data, including analyses of regional instability, changes in government policy, investment security, social movements, and environmental risks. For this reason, mining companies have become the corporate component of energy espionage. China's Belt and Road Initiative is not only an infrastructure investment program but also a systematic strategy for controlling global mineral supply chains (Zhao & Zhao, 2025). Investments in cobalt in Africa, lithium in Latin America, and rare earth sites in Asia are projects carried out by Chinese state-owned enterprises with diplomatic and military support.

The United States and the European Union have increasingly aligned their strategies to secure critical mineral supply chains through cooperative frameworks that emphasize trusted partnerships, diversification of sources, and supply-chain resilience. As part of this effort, transatlantic coordination under initiatives such as the Minerals Security Partnership and bilateral memoranda of understanding aims to strengthen joint mining and processing projects with democratic partners while reducing strategic dependencies on dominant suppliers. These "secure supply chain" and allied sourcing approaches reflect broader economic security priorities and friend-shoring principles in critical minerals policy (U.S. Department of Commerce & European Commission Joint Statement, 2024). This competitive strategic landscape has fostered a broad framework in which coercive diplomacy, financial sanctions, technology transfer agreements, and intelligence activities are employed to influence modern mining operations. Contemporary mining practices increasingly rely on advanced technological tools such as satellite remote sensing for geological and operational surveillance, unmanned aerial reconnaissance systems for detailed site monitoring, machine learning-based geological modelling for data-driven exploration, IoT-enabled field systems for real-time equipment and infrastructure monitoring, and blockchain enhanced supply chain solutions to ensure traceability and transparency throughout mineral value chains (Onifade, Adebisi & Zvarivadza, 2023). Table 8 presents the principal technological tools and their strategic functions in modern mining operations.

**Table 8.** Technological tools and strategic functions in modern mining operations.

| <i>Technology / Tool</i>                           | <i>Operational and Strategic Purpose</i>                                                                                                                                                                        |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Satellite imagery</i>                           | It refers to the detection of mineral reserves, the monitoring of land-use changes, and the analysis of illegal mining activities through non-destructive and large-scale remote sensing techniques.            |
| <i>Drone reconnaissance systems</i>                | These are autonomous or semi-autonomous aerial platforms that collect high-resolution imagery and sensor data in hard-to-access areas, enabling geological analysis, site security, and operational monitoring. |
| <i>Machine learning-based geological modelling</i> | This involves the use of artificial intelligence algorithms to analyse geological datasets for mineral reserve estimation, ore quality classification, and the reduction of exploration risks.                  |

**Table 8.** (Continued)

|                                               |                                                                                                                                                                                                       |
|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>IoT-based field systems</i>                | These are connected device networks that enable real-time monitoring of equipment, sensors, and infrastructure in mining sites, collecting data to enhance operational security and efficiency.       |
| <i>Blockchain-based supply chain security</i> | These are distributed ledger technology mechanisms that provide traceability, transparency, and anti-counterfeiting capabilities from the extraction of critical minerals to the final product stage. |

Source: Onifade et al. (2023), Zvarivadza et al. (2024), and Onifade, Adebisi and Zvarivadza (2024).

Data has increasingly become a strategic resource that generates both economic value and geopolitical leverage. As a result, companies operating in resource-intensive sectors such as mining are becoming active actors within the broader competition over data and digital capabilities. Artificial intelligence technologies are applied across multiple areas, including large-scale data analytics, OSINT processing, automated monitoring systems such as drones, and production optimization in industrial operations. These developments demonstrate how digital infrastructures and data analytics are reshaping strategic decision-making in resource industries (Bughin et al., 2018). At the same time, it opens the door to hybrid threats such as data manipulation and competitor sabotage.

Mining companies carry the responsibility of securing not only economic legitimacy but also social and environmental legitimacy in the communities where they operate. Engagement with local populations, investment in sustainability initiatives, effective management of environmental impacts, and adherence to human rights and corporate social responsibility protocols significantly influence the operational environment and long-term viability of mining firms. The literature emphasizes that constructive relations with surrounding communities and responsive environmental practices are necessary for acquiring a “social license to operate,” shaping both corporate reputation and regulatory legitimacy in the extractive sector (Rodrigues, et al, 2022).

However, in some regions, the discourse of social investment often masks broader geostrategic influence-building strategies. Corporations may deploy corporate social responsibility initiatives, cultivate relationships with local elites, and leverage communication networks to shape local political dynamics and secure operational legitimacy in resource-rich areas (Csevár, 2021). As competition over critical minerals intensifies, risks are also increasing. Table 9 categorizes the principal geostrategic risks and operational threats confronting critical mineral operations.

**Table 9.** Geostrategic risks and operational threats in critical mineral operations.

| <i>Risk / Threat</i>            | <i>Description / Operational Impact</i>                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Cyberattacks</i>             | This refers to the targeting of critical mineral databases, production systems, and supply chain infrastructures to steal information, disrupt operations, or conduct strategic manipulation.      |
| <i>Logistics sabotage</i>       | This involves disrupting supply security through physical or digital interventions targeting mining and energy transport lines, ports, railways, and storage facilities.                           |
| <i>Drone espionage</i>          | This refers to the use of unmanned aerial vehicles to monitor mining sites, access confidential reserve information, or record field operations.                                                   |
| <i>Local militia influence</i>  | This describes the activities of non-state armed groups operating in mining regions to exert control over security, access, or economic gains, thereby placing pressure on companies.              |
| <i>Disinformation campaigns</i> | This refers to the deliberate spread of false information through social media and media channels to discredit companies, destabilize the investment environment, or manipulate public perception. |
| <i>Political intervention</i>   | This involves the strategic direction of mining policies and corporate activities by governments or international actors through laws, permits, taxation, or diplomatic processes.                 |

Source: IRENA (2023)

This trend has intensified not only economic competition but also geopolitical power struggles, compelling both states and multinational corporations to formulate an array of explicit strategic policies and covert operational methods to secure critical mineral supply chains. As the security and resilience of these supply networks have become central elements of national strategic agendas, governments are increasingly integrating political intervention into mining policies, regulatory frameworks, and foreign investment rules. Such dynamics have, in turn, driven corporate actors to adopt more robust security postures sometimes resembling quasi-military protection arrangements to defend their assets, operations, and geopolitical interests in contested resource environments (Yeh & Ni, 2025). Today, mining corporations have moved beyond their traditional identity as economic actors and have become geopolitical players, components of intelligence architectures, and agents of security strategy. In the era of Energy Espionage 5.0, mining corporations assume multifaceted strategic roles that extend well beyond conventional commercial operations. Table 10 delineates these evolving corporate functions within the framework of energy espionage.

**Table 10.** Strategic roles of mining corporations in energy espionage 5.0.

| <i>Corporate Role</i>                                 | <i>Operational and Strategic Description</i>                                                                                                                                                  |
|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>An instrument of nation-state power projection</i> | This refers to the use of commercial intelligence by states to gain strategic influence in the international arena and to expand their economic and geopolitical spheres of effect.           |
| <i>A data-driven economic intelligence actor</i>      | This denotes the capacity to generate strategic data by monitoring economic trends through artificial intelligence, big data, and financial analytics in order to gain competitive advantage. |
| <i>A stakeholder in technological competition</i>     | This describes the strategic role played in advancing global technological leadership through R&D processes, critical minerals, and advanced production technologies.                         |
| <i>An implementer of hybrid security systems</i>      | This involves conducting multi-layered security and intelligence operations by integrating cybersecurity, corporate security, satellite surveillance, and human intelligence.                 |
| <i>A diplomatic negotiation actor</i>                 | This refers to the role played in economic agreements, resource access arrangements, supply chain accords, and strategic partnerships in line with national interests.                        |

Source: IRENA (2023), Holmström (2020), and Vivoda et al. (2025).

As illustrated in Table 10, mining corporations have risen to the status of systemic power actors. As global competition over critical minerals intensifies, the strategic significance of these corporate roles will continue to expand, making robust legal frameworks, ethical governance structures, and integrated geopolitical risk management increasingly decisive for both corporate sustainability and national security.

#### 1.4. Commercial Data Collection Intelligence

With the digitalization of the global economy, commercial information has become one of the most critical power factors of the modern era. Commanding data plays a decisive role not only in economic decision-making processes but also in achieving geopolitical, industrial, and technological superiority. For this reason, international corporations, state institutions, and strategic sectors are increasingly investing in commercial data collection intelligence, employing big data analytics, artificial intelligence, satellite systems, Human Intelligence (HUMINT), and cyber-monitoring systems in an integrated manner.

The foundation of commercial intelligence lies in the systematic collection, classification, and transformation of diverse data sources into strategic insight. In this context, the data ecosystem is multidimensional and includes the following primary categories:

1. Digital monitoring systems: Competitive analyses are generated through corporate networks, digital platforms, and online user behaviour. Large-scale data mining is used to understand companies' production, distribution, and investment strategies.
2. Supply chain networks: Logistics flows, port activities, supplier linkages, and global trade data are examined. The directional movement of critical raw materials identifies competitive threats and opportunity areas.

3. Corporate communication infrastructures: Press releases, investor reports, industry meetings, and executive discourse are analysed to assess strategic orientations.
4. Market behaviours: Consumer trends, pricing strategies, and demand fluctuations are monitored to interpret competitive dynamics.
5. Artificial intelligence monitoring systems: AI-based algorithms detect patterns within data sets and generate predictive models.
6. Satellite and drone data: Mining sites, energy infrastructures, logistics corridors, and production facilities are observed remotely.
7. Financial transaction networks: Foreign direct investment flows, commercial bank transfers, and capital market movements are analysed.
8. Cybersecurity sensors: Data leaks, spyware activity, and indicators of corporate network penetration are identified.

The coordinated use of these data types enables companies not only to observe the current state of competition but also to shape their future strategic orientations. As Clarke and Knake (2019) emphasize, modern competition is no longer physical; it is structured around cognitive and digital dominance. Commercial data collection intelligence differs fundamentally from traditional market research. Table 11 provides a comparative analysis that highlights the fundamental distinctions between conventional market research and commercial intelligence in terms of purpose, scope, methodology, and risk.

**Table 11.** Comparative analysis of market research and commercial intelligence.

| Criterion              | Market Research                      | Commercial Intelligence                                |
|------------------------|--------------------------------------|--------------------------------------------------------|
| <i>Primary purpose</i> | Analysing customer behaviour         | Ensuring competitive and strategic security            |
| <i>Data type</i>       | Sales, advertising, customer surveys | Digital networks, supply chain data, financial signals |
| <i>Scope</i>           | Consumer-focused                     | Includes national security and industrial sovereignty  |
| <i>Method</i>          | Legal open sources                   | Authentic open sources + grey data channels            |
| <i>Focus</i>           | Micro-market analysis                | Macro geo-economic analysis                            |
| <i>risk factor</i>     | Low                                  | High (cyber, diplomatic, industrial)                   |

Source: Clarke and Knake (2019), Naseri (2005), and Holmström (2020).

According to Alnoukari and Hanano (2017), commercial intelligence occupies a strategic space at the intersection of contemporary economic espionage and legitimate data analysis. Consequently, while states and corporations strive to maintain legal boundaries in their information-gathering activities, they must simultaneously navigate the reality that competing actors are constantly testing and extending those limits. In the current era marked by energy transition, the rising significance of critical minerals, and rapid technological transformation, the strategic value of information continues to grow. The critical mineral strategies of China, the United States, and the European Union demonstrate that state power and corporate

intelligence are becoming deeply intertwined (Stango, 2024). In this context, the core motivations of commercial intelligence are as follows:

- Strategic resource security: Information dominance is achieved to secure the supply of critical minerals such as rare earth elements, lithium, cobalt, and graphite.
- Technological competition: Battery technologies, artificial intelligence models, refining solutions, and production innovations are monitored.
- Financial superiority: Market manipulation risks and capital flow directions are anticipated.
- Geopolitical positioning: State policies in energy and mineral regions are closely tracked.
- Defence and national security: Threats targeting industrial infrastructure are identified in advance.

Within this framework, commercial data collection is not merely an economic activity; it is also an instrument of strategic power projection. OSINT constitutes the cornerstone of the commercial information ecosystem. The primary open-source tools used include:

- Official reports: Government and international organization reports that contain strategic projections (Kim, et al, 2022).
- Commercial registries: Corporate ownership structures, capital distribution, and managerial history are analysed.
- Academic studies: R&D orientations and scientific capacity are monitored.
- Satellite maps: Production and logistics facilities are examined remotely.
- Patent databases: Emerging technological trends are identified.
- Industrial forums: Professional knowledge-sharing platforms serve as sources of strategic signals.
- Port and logistics data: Supply chain flows are interpreted through vessel and container movements.
- Financial reports: Investment strategies and profitability projections are assessed.
- Blockchain analytics: Digital financial movements and ledger chains are tracked.

The data obtained from these sources is transformed into strategic foresight at both corporate and state levels. The following methods, while legally and ethically contentious, are widely used in global competition:

- Fake buyer profiles: Used as identity-masking methods to obtain pricing and production details from rival companies.
- Supply chain infiltration: Gaining access to competitor supplier networks to acquire information on materials and cost structures.
- Industrial espionage: Attempts to access technologies, production formulas, and strategic plans.

- Financial monitoring: Analysing financial movements to identify capital flow directions.

Such methods are particularly widespread in the critical minerals sector, where competition over resources has increasingly assumed a strategic dimension (Maus et al., 2021). Increasing remote sensing capabilities in geo-economic domains have revolutionized corporate analysis:

- Reserve monitoring: Mineral capacity and extraction rates are estimated through satellite data.
- Site activity control: Production rhythm is measured by tracking personnel and machinery activity.
- Logistics monitoring: Ports, road networks, and storage facilities are observed.
- Infrastructure development: new mining wells, energy facilities, and transportation projects are monitored.

Recent studies indicate that digital technologies and artificial intelligence have significantly improved the capacity to collect and analyse geological and operational data in the mining sector. Nevertheless, intelligence gathering in strategic industries still relies heavily on human-centred channels in addition to technological monitoring systems (Davenport & Ronanki, 2018). Key human-based information channels include:

1. Talent mobility and expert recruitment: Firms may gain strategic knowledge through the hiring of specialists who previously worked for competitor companies.
2. Professional conferences and industry events: Informal exchanges and presentations often reveal technological trends and operational strategies.
3. Supplier and subcontractor networks: Secondary actors in the supply chain can provide insights into production processes, procurement strategies, and logistical operations.
4. On-site observers and field representatives: Direct observation of facilities, project sites, and infrastructure developments may generate valuable operational information.

With the expansion of data processing capacity, AI has become the primary component of intelligence:

- Demand forecasting systems: The future direction of market demand is predicted.
- Mineral price modelling: Commodity price fluctuations are analysed.
- Risk analytics: Geopolitical and logistical risks are modelled through scenario analysis.
- Cyber threat detection: Anomalies within corporate networks are identified.
- Deepfake and manipulative media analysis: Protection is provided against external information threats.

AI-driven intelligence provides states with rapid decision-making superiority and enhanced early-warning capabilities. At the same time, critical minerals constitute the primary raw material base of the emerging energy economy (Gielen, 2021). For this reason, data collection

domains in the mining and energy sectors have expanded considerably. Table 12 outlines the strategic value of key data types in mining and energy intelligence.

**Table 12.** Strategic value of data types in mining and energy intelligence.

| <i>Data Type</i>           | <i>Strategic Value</i>                                        |
|----------------------------|---------------------------------------------------------------|
| <i>Reserve information</i> | Provides power projection for states and corporations         |
| <i>Supply chain routes</i> | Logistic disruptions can evolve into national security risks  |
| <i>R&amp;D data</i>        | Emerging technologies form the basis of competitive advantage |
| <i>Price movements</i>     | Enable financial manipulation and market dominance            |
| <i>Diplomatic data</i>     | Resource agreements generate geopolitical consequences        |

Source: (2021), IRENA (2023), and Holmström (2020).

IEA 2022 reports indicate that competition over critical minerals will stand at the centre of security doctrines after 2030 (Kim et al, 2021). The most sensitive aspect of commercial intelligence is the preservation of legal boundaries. The fundamental framework includes:

1. Competition law: Protection against unfair competition and corporate trade secrets.
2. Personal data protection: Limits of data processing under GDPR and national data regulations.
3. Intellectual property law: Protection of patents, know-how, and trade secrets.
4. Cyber law: Penalization of digital breaches and data theft.
5. Ethical governance: Transparency and human-rights principles in data processing.

Commercial intelligence requires a delicate balance between power dynamics and ethical norms.

Commercial data collection intelligence is a strategic domain of knowledge that bridges the global economy and national security. Critical minerals, energy security, artificial intelligence, and supply chain dynamics are shaping the competitive landscape of the future. Data superiority signifies not only economic growth but also geopolitical sovereignty. For this reason, states and corporations will continue to enhance their data-collection capacities and maintain their competitive advantage while upholding legal and ethical frameworks.

## 2. Discussion

The strategic competition of the 21st century is being reshaped within the framework of safeguarding energy supply and securing technological superiority. The global climate and sustainability goals driving the transition from fossil fuels to renewable energy systems have created a new economic and geopolitical paradigm. This transformation surpasses traditional energy security doctrines and places critical mineral security, data-driven competition, corporate intelligence strategies, and AI-supported surveillance systems at the centre. In this context, a new strategic environment—one that can be described as the Energy Espionage 5.0 paradigm—is emerging. This environment represents a multidimensional power domain in which not only states, but also multinational corporations, financial institutions, private military security companies, and technology giants have become influential actors.

Within the scope of this study, the four principal components—mineral geopolitics and industrial espionage; satellites, drones, and resource-monitoring technologies; the geostrategic role of mining companies; and commercial data-collection intelligence—were systematically analysed to understand the nature and future trends of contemporary energy competition. The findings demonstrate that the global energy order is not merely an economic or technical matter but is deeply intertwined with political, military, intelligence, and technological processes.

Lithium, cobalt, nickel, graphite, and rare earth elements have become the most critical inputs of the green energy transition. Used in electric vehicle batteries, wind turbine magnets, precision sensors, defence technologies, and artificial intelligence chips, these minerals form the strategic arteries of the global economic system. While this situation evokes the oil geopolitics of the Cold War era, it has produced a far more complex and technology-centred form of competition.

Operations carried out by multinational corporations, states, and private intelligence networks targeting critical mineral markets extend across a broad spectrum, ranging from the theft of trade secrets and cyber intrusions to geological data espionage and the manipulation of logistics systems. The acquisition of mining-related information is no longer based solely on field research conducted by geologists; instead, it increasingly relies on satellite imagery, electromagnetic scanning technologies, AI-assisted modelling, and large-scale data mining. Consequently, the convergence of geopolitical interests and advanced information technologies forms a central component of contemporary energy strategy (Binnemans et al., 2021; Gielen, 2021).

Satellite systems, unmanned aerial vehicles, and intelligent sensor networks have set the new standard for surveillance in the energy and mining sectors. Today, states and corporations can monitor their competitors' reserve areas, extraction capacities, infrastructural activity, shipment patterns, and environmental impacts in real time.

Not only physical operations but also logistical routes, financial flows, port movements, and even employee behaviour leave digital traces. These traces hold strategic value in energy espionage (Yang et al., 2024). At this point, the rise of drone operations specific to the energy sector occupies a critical position at the intersection of military and commercial applications. China's dominance in drone technologies, U.S. export restrictions, and debates over aerial data security stand as visible examples of this growing competition.

Alongside these technological and strategic developments, a growing body of literature highlights the normative tension between data privacy, commercial confidentiality, and national security imperatives. Scholars argue that the increasing reliance on commercial satellite data, open-source intelligence, and AI-driven surveillance blurs traditional boundaries between civilian and military intelligence domains, raising critical concerns regarding data governance and individual as well as corporate privacy (Deibert, 2019; Lyon, 2022). From a national security perspective, unrestricted access to high-resolution commercial data may

expose critical infrastructure and strategic assets, while from a privacy-oriented viewpoint, expansive surveillance practices risk enabling forms of digital overreach and extraterritorial monitoring (Kello, 2017; Zuboff, 2024). This debate underscores that contemporary energy espionage is not solely a technical or geopolitical issue, but also a legal and ethical challenge situated at the intersection of security governance, corporate intelligence, and data sovereignty.

Traditional security literature has positioned states as the primary actors in the international system. However, in the transformational era of energy, mining corporations, technology giants, and financial institutions have emerged as strategic power centres capable of influencing national interests. Energy and mining giants such as Chevron, Glencore, Rio Tinto, and BHP are not merely economic actors; they are geostrategic players involved in diplomatic, military, and informational operations (Müller, 2023).

Power relations in the energy sector are no longer shaped solely by intergovernmental agreements but increasingly through corporate mergers, investment strategies, patent ownership, data dominance, and supply chain control. While this does not eliminate the role of states, it indicates the emergence of a non-hierarchical, multi-layered power architecture in the international system.

One of the most competitive areas in the energy sector is the pursuit of informational superiority. Open-source intelligence (OSINT), commercial data analytics platforms, AI-based market prediction algorithms, and information leaks constitute the invisible infrastructure of energy espionage.

In particular, big data analytics, blockchain-based supply chain monitoring systems, and deep learning models shape the decision-making processes of both state intelligence agencies and corporations. This creates a new geoeconomic order in which the actor who possesses information gains power (Zhao et al., 2022). However, this competition also increases the risks of data manipulation, disinformation operations, and economic sabotage.

This analysis demonstrates that energy security is no longer simply about ensuring the continuity of energy supply. In the new era, energy security is:

- Securing critical minerals
- Strategic technological investments
- Information superiority and cybersecurity
- Supply chain resilience
- Geopolitical alliance engineering
- Artificial intelligence and data dominance
- Corporate diplomacy and green regulation warfare

are the core components that define this new paradigm.

Taken together, these policy directions provide a practical framework through which developing economies can better protect strategic mineral assets in the era of Energy Espionage 5.0.

### **3. Conclusion**

Accordingly, *Energy Espionage 5.0* represents not merely the modernization of espionage activities but the integration of economic, technological, diplomatic, and military strategies into a unified framework. The contribution of this study lies in redefining energy security not as a simple supply–demand equilibrium, but as a multi-actor, multi-layered geo-intelligence ecosystem.

Ultimately, strategic success for nations and corporations will no longer be measured solely by possessing resources, but by their capacity to collect, process, protect, and operationalize information about those resources. The future extends far beyond competition over physical minerals: whoever controls the data mines will also control the energy system. For this reason, the literature on energy espionage—located at the intersection of international politics, technology strategies, and industrial security—must deepen both theoretically and practically. This study makes an original contribution to the field of energy geopolitics by placing the informational dimension at the centre and integrating corporate intelligence with energy security. Future research should expand to include the legal, ethical, and security dimensions of AI-supported intelligence systems.

#### **Conflict of Interest**

The authors declare no conflict of interest.

#### **Funding**

This research received no external funding.

#### **Acknowledgments**

The authors thank all colleagues who contributed to discussions during this study.

#### **Ethical Statements**

This study did not involve human participants or animals and therefore did not require ethical approval.

#### **Data and Code Availability**

All data and code used in this study are publicly available at

#### **Supplementary Materials**

Not applicable.

## References

- Ajayi, R., & Ikemefuna, C. D. (2025). Drone-based detection systems for resource exploration. *Int. J. Adv. Res. Publ. Rev*, 2, 590-615.
- Akporiaye, A., & Webster, D. G. (2022). Social license and CSR in extractive industries: A failed approach to governance. *Global Studies Quarterly*, 2(3), ksac041.
- Alnoukari, M., & Hanano, A. (2017). Integration of business intelligence with corporate strategic management. *Journal of Intelligence Studies in Business*, 7(2).
- Alsadie, D. (2025). Cybersecurity and Artificial Intelligence in Unmanned Aerial Vehicles: Emerging Challenges and Advanced Countermeasures. *IET Information Security*, 2025(1), 2046868. <https://doi.org/10.1049/ise2/2046868>
- Ayamga, M., Tekinerdogan, B., & Kassahun, A. (2021). Exploring the challenges posed by regulations for the use of drones in agriculture in the African context. *Land*, 10(2), 164. <https://doi.org/10.3390/land10020164>
- Balaniuk, R., Isupova, O., & Reece, S. (2020). Mining and tailings dam detection in satellite imagery using deep learning. *Sensors*, 20(23), 6936. <https://doi.org/10.3390/s20236936T>
- Balaniuk, R., Isupova, O., & Reece, S. (2020). Mining and tailings dam detection in satellite imagery using deep learning. *Sensors*, 20(23), 6936. <https://doi.org/10.3390/s20236936>
- Bechtum, A. (2024). Transnational corporate influence and local politics: How mining companies intervene in neighbouring communities. *Zeitschrift für Friedens-und Konfliktforschung*, 13(1), 193-219. <https://doi.org/10.1007/s42597-024-00121-6>
- Bedini, E. (2017). The use of hyperspectral remote sensing for mineral exploration: A review. *Journal of Hyperspectral Remote Sensing*, 7(4), 189-211. <https://doi.org/10.29150/jhrs.v7.4.p189-211>
- Bradshaw, M. J. (2010). Global energy dilemmas: a geographical perspective. *Geographical Journal*, 176(4), 275-290. <https://doi.org/10.1111/j.1475-4959.2010.00375.x>
- Bughin, J., Seong, J., Manyika, J., Chui, M., & Joshi, R. (2018). Notes from the AI frontier: Modeling the impact of AI on the world economy. *McKinsey Global Institute*, 4(1), 2-61.
- Clarke, R. A., & Knake, R. K. (2019). *The fifth domain: Defending our country, our companies, and ourselves in the age of cyber threats*. Penguin.
- Dannreuther, R. (2013). Geopolitics and international relations of resources. In *Global resources: Conflict and cooperation* (pp. 79-97). London: Palgrave Macmillan UK.
- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard business review*, 96(1), 108-116.
- Deibert, R. J. (2020). *Reset: Reclaiming the internet for civil society*. House of Anansi.

- Ericsson, M., Löf, O., & Löf, A. (2024). Locus of control over global mine production—developments between 1985 and 2018 against a historical background. *Mineral Economics*, 37(3), 633-643. <https://doi.org/10.1007/s13563-024-00454-x>
- European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1252 of the European Parliament and of the Council establishing a framework for ensuring a secure and sustainable supply of critical raw materials (Critical Raw Materials Act)*. Official Journal of the European Union. <https://eur-lex.europa.eu/>
- European Parliamentary Research Service. (2023). *EU-US critical minerals agreement: Mineral Security Partnership and policy approaches*. European Parliament. Retrieved from [https://www.europarl.europa.eu/RegData/etudes/BRIE/2023/754617/EPRS\\_BRI%282023%29754617\\_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2023/754617/EPRS_BRI%282023%29754617_EN.pdf)
- Finer, M., Jenkins, C. N., Pimm, S. L., Keane, B., & Ross, C. (2008). Oil and gas projects in the western Amazon: threats to wilderness, biodiversity, and indigenous peoples. *PloS one*, 3(8), e2932. <https://doi.org/10.1371/journal.pone.0002932>
- Fischer, A. (2023). Data sovereignty and e-governance: The legal implications of national laws on digital government systems. *Legal Studies in Digital Age*, 2(4), 1-12.
- Frankovits, G., & Mirsky, Y. (2023). Discussion paper: The threat of real time deepfakes. In *Proceedings of the 2nd Workshop on Security Implications of Deepfakes and Cheapfakes* (pp. 20-23). <https://doi.org/10.1145/3595353.3595881>
- Gaber, T., El Jazouli, Y., Eldesouky, E., & Ali, A. (2021). Autonomous haulage systems in the mining industry: cybersecurity, communication and safety issues and challenges. *Electronics*, 10(11), 1357. <https://doi.org/10.3390/electronics10111357>
- Garcia-del-Real, J., & Alcaráz, M. (2024). Unlocking the future of space resource management through satellite remote sensing and AI integration. *Resources Policy*, 91, 104947. <https://doi.org/10.1016/j.resourpol.2024.104947>
- Gielen, D. (2021). Critical minerals for the energy transition. *International Renewable Energy Agency, Abu Dhabi*, 42.
- Haglund, D. G. (1986). The new geopolitics of minerals: An inquiry into the changing international significance of strategic minerals. *Political Geography Quarterly*, 5(3), 221-240. [https://doi.org/10.1016/0260-9827\(86\)90035-2](https://doi.org/10.1016/0260-9827(86)90035-2)
- Hannas, W. C., Mulvenon, J., & Puglisi, A. B. (2013). *Chinese industrial espionage: Technology acquisition and military modernisation*. Routledge.
- Holmström, L. (2010). Industrial espionage and corporate security: the Ericsson case.
- Hou, T., & Wang, V. (2020). Industrial espionage—A systematic literature review (SLR). *computers & security*, 98, 102019. <https://doi.org/10.1016/j.cose.2020.102019>
- International Energy Agency. (2021). *The role of critical minerals in clean energy transitions*. <https://www.iea.org/reports/the-role-of-critical-minerals-in-clean-energy-transitions>

- International Energy Agency. (2023). *Critical minerals market review 2023*. OECD Publishing. <https://www.iea.org/reports/critical-minerals-market-review-2023>
- International Renewable Energy Agency (IRENA). (2023). *Geopolitics of the Energy Transition: Critical Materials*. IRENA, Abu Dhabi. <https://www.irena.org/publications/2023/Jul/Geopolitics-of-the-Energy-Transition-Critical-Materials>
- Jakhu, R. S., Pelton, J. N., & Nyampong, Y. O. M. (2017). *Space mining and its regulation* (Vol. 106). Cham, Switzerland: Springer International Publishing.
- Kello, L. (2017). *The virtual weapon and international order*. Yale University Press.
- Kim, T. Y., Gould, T., Bennet, S., Briens, F., Dasgupta, A., Gonzales, P., ... & Lagelee, J. (2021). The role of critical minerals in clean energy transitions. *International Energy Agency: Washington, DC, USA*, 70-71.
- Langner, R. (2011). Stuxnet: Dissecting a cyberwarfare weapon. *IEEE security & privacy*, 9(3), 49-51.
- Lyon, D. (2021). *Pandemic surveillance*. John Wiley & Sons.
- Marshall, J. A., Bonchis, A., Nebot, E., & Scheduling, S. (2016). Robotics in mining. In *Springer handbook of robotics* (pp. 1549-1576). Cham: Springer International Publishing.
- Maus, V., Giljum, S., Gutschlhofer, J., da Silva, D. M., Probst, M., Gass, S. L. & McCallum, I. (2020). A global-scale data set of mining areas. *Scientific data*, 7(1), 289.
- Müller, M. (2023). The ‘new geopolitics’ of mineral supply chains: A window of opportunity for African countries. *South African journal of international affairs*, 30(2), 177-203. <https://doi.org/10.1080/10220461.2023.2226108>
- Nakano, J. (2021). *The geopolitics of critical minerals supply chains* (pp. 19-22). Washington, DC, USA: Center for Strategic and International Studies (CSIS).
- Nasheri, H. (2005). *Economic espionage and industrial spying* (Vol. 1). Cambridge University Press.
- Noonan, T., & Archuleta, E. (2008). The Insider Threat to Critical Infrastructures. The National Infrastructure Advisory Council.
- Nye, J. S. (2011). *The future of power*. PublicAffairs.
- Ocelík, V., Kolk, A., & Ciulli, F. (2023). Multinational enterprises, Industry 4.0 and sustainability: A multidisciplinary review and research agenda. *Journal of Cleaner Production*, 413, 137434. <https://doi.org/10.1016/j.jclepro.2023.137434>
- Onifade, M., Adebisi, J. A., & Zvarivadza, T. (2024). Recent advances in blockchain technology: Prospects, applications and constraints in the minerals industry. *International Journal of Mining, Reclamation and Environment*, 38(7), 497-533. <https://doi.org/10.1080/17480930.2024.2319453>

- Onifade, M., Adebisi, J. A., Shivute, A. P., & Genc, B. (2023). Challenges and applications of digital technology in the mineral industry. *Resources Policy*, 85, 103978. <https://doi.org/10.1016/j.resourpol.2023.103978>
- Organisation for Economic Co-operation and Development. (2022). *Global critical mineral supply chains and economic security*. OECD Publishing.
- Pedersen, K. T., Pepke, L., Stærmose, T., Papaioannou, M., Choudhary, G., & Dragoni, N. (2025). Deepfake-driven social engineering: Threats, detection techniques, and defensive strategies in corporate environments. *Journal of Cybersecurity and Privacy*, 5(2), 18. <https://doi.org/10.3390/jcp5020018>
- Pedram, M., Chandler, S., & Georgiades, E. (2025). When open-source information backfires: satellite imagery and privacy breaches. *Vand. J. Transnat'l L.*, 58, 119.
- Peyghambari, S., & Zhang, Y. (2021). Hyperspectral remote sensing in lithological mapping, mineral exploration, and environmental geology: an updated review. *Journal of Applied Remote Sensing*, 15(3), 031501-031501. <https://doi.org/10.1117/1.JRS.15.031501>
- Prost, G. L. (2025). *Remote sensing for geoscientists: image analysis and integration*. CRC Press.
- Raghuvanshi, N., Johri, A., Saxena, M., & Rout, P. (2025). *Space Based Surveillance and Reconnaissance Technologies*.
- Rodrigues, M., Alves, M. C., Silva, R., & Oliveira, C. (2022). Mapping the literature on social responsibility and stakeholders' pressures in the mining industry. *Journal of Risk and Financial Management*, 15(10), 425. <https://doi.org/10.3390/jrfm15100425>
- Shahmoradi, J., Talebi, E., Roghanchi, P., & Hassanalain, M. (2020). A comprehensive review of applications of drone technology in the mining industry. *Drones*, 4(3), 34. <https://doi.org/10.3390/drones4030034>
- Sihag, V., Choudhary, G., Choudhary, P., & Dragoni, N. (2023). Cyber4Drone: A systematic review of cyber security and forensics in next-generation drones. *Drones*, 7(7), 430. <https://doi.org/10.3390/drones7070430>
- Stango, A. (2024). The geopolitical competition between China and the US in new technologies. *Luiss School of Government, Working Paper Series, SOG-WP12/2024.(mayo de 2024)(en línea)* <https://sog.luiss.it/sites/sog.luiss.it/files/stango%20AS%20REV%20The%20geopolitical%20competition%20between%20China%20and%20the%20U.S.%20in%20new%20technologies.pdf>
- Stockholm International Peace Research Institute (SIPRI). (2024). *Critical Minerals and Great Power Competition*. SIPRI.
- U.S. Department of Commerce & European Commission. (2024). *U.S.–EU Joint Statement of the Trade and Technology Council: Collaboration on Critical Minerals Supply Chains and Economic Security*. Retrieved from <https://www.commerce.gov/news/press-releases/2024/04/us-eu-joint-statement-trade-and-technology-council>

- U.S. Department of Justice. (2023). *Chinese cyber espionage campaigns targeting critical infrastructure and natural resource sectors*. <https://www.justice.gov/>
- U.S. Department of the White House. (2025). *United States-Japan Framework for Securing the Supply of Critical Minerals and Rare Earths through Mining and Processing*. Retrieved from <https://www.whitehouse.gov/briefings-statements/2025/10/united-states-japan-framework-for-securing-the-supply-of-critical-minerals-and-rare-earths-through-mining-and-processing/>
- U.S. Geological Survey. (2024). *Mineral Commodity Summaries 2024*. Reston, VA: U.S. Geological Survey
- Van der Meer, F. D., Van der Werff, H. M., Van Ruitenbeek, F. J., Hecker, C. A., Bakker, W. H., Noomen, M. F., ... & Woldai, T. (2012). Multi-and hyperspectral geologic remote sensing: A review. *International journal of applied Earth observation and geoinformation*, 14(1), 112-128. <https://doi.org/10.1016/j.jag.2011.08.002>
- Vivoda, V., Matthews, R., & Andresen, J. (2025). Securing defense critical minerals: Challenges and US strategic responses in an evolving geopolitical landscape. *Comparative Strategy*, 44(2), 281-315. <https://doi.org/10.1080/01495933.2025.2456427>
- Vivoda, V., Matthews, R., & McGregor, N. (2024). A critical minerals perspective on the emergence of geopolitical trade blocs. *Resources Policy*, 89, 104587. <https://doi.org/10.1016/j.resourpol.2023.104587>
- Yan, Z. B., Liu, S. W., & Wu, J. (2022). Hyperspectral remote sensing of plant functional traits: monitoring techniques and future advances. *Chinese journal of plant ecology*, 46(10), 1151-1166. <https://doi.org/10.17521/cjpe.2022.0223>
- Yang, F., Zuo, R., & Kreuzer, O. P. (2024). Artificial intelligence for mineral exploration: A review and perspectives on future directions from data science. *Earth-Science Reviews*, 258, 104941. <https://doi.org/10.1016/j.earscirev.2024.104941>
- Yeh, K. H., & Ni, G. (2025). Critical mineral supply chain security and multi-layered regional governance: insights from Central Asia. *Trames: A Journal of the Humanities and Social Sciences*, 29(3), 271-294.
- Zhao, P., & Zhao, T. (2025). The relationships between geopolitics and global critical minerals shipping: A literature review. *Ocean & Coastal Management*, 262, 107559. <https://doi.org/10.1016/j.ocecoaman.2025.107559>
- Zuboff, S. (2024). The age of surveillance capitalism: The fight for a human future at the new frontier of power. *Journal of Information Ethics*, 33(1), 84-85.
- Zvarivadza, T., Onifade, M., Dayo-Olupona, O., Said, K. O., Githiria, J. M., Genc, B., & Celik, T. (2024). On the impact of Industrial Internet of Things (IIoT)-mining sector perspectives. *International Journal of Mining, Reclamation and Environment*, 38(10), 771-809. <https://doi.org/10.1080/17480930.2024.2347131>

# Volume Over Lethality: A Matrix-Exponential Framework for Defending Drone Swarms

Mustafa Tekinay <sup>1</sup> 

## Abstract

We study large, heterogeneous drone swarms composed of both threat and decoy drones. The defender's engagement is modelled as a multi-stage stochastic process that captures uncertainty and variability in response times. A non-discriminatory defender is considered in which threat and decoy drones are treated identically. We adopt a matrix-exponential framework that yields tractable, closed-form expressions for key performance measures. The results indicate that swarm size has a substantially stronger impact on defence success than threat-decoy composition. For a typical defender, swarms exceedingly roughly ten drones may exceed single-defender swarm-size capacity under the baseline modelling assumptions, even with high decoy fractions. While a typical defender can reliably handle seven threat drones, augmenting the swarm with seven decoys reduces the probability of successful defence from 0.93 to 0.09. From a cost perspective, heterogeneous drone swarms can reduce overall cost by up to 67% when decoy drones are priced at approximately one tenth of the cost of threat drones. Finally, performance depends not only on aggregate rates but also on individual stage bottlenecks. Under idealised conditions where each engagement stage completes in approximately one second on average, single-defender capacity increases to approximately 35-40 drones.

**Keywords:** Drone swarms, counter-UAV systems, threat drones, decoy drones, matrix-exponential distributions, stochastic performance analysis, swarm defence capacity.

---

Received: 24.01.2026

Accepted: 25.05.2026

Research Article

<sup>1</sup> ASELSAN Inc., Ankara, Türkiye

\*Corresponding author.

✉ m.tekinay@aselsan.com

**Cite this article as:**

<https://doi.org/10.65834/jdsi.12.58>

Tekinay, M. (2026). Volume Over Lethality: A Matrix-Exponential Framework for Defending Drone Swarms. *Journal of Defence and Security Industries: Science & Technology*1(2), 59-87.

## 1. Introduction

Unmanned Aerial Systems (UAS) have rapidly become a central element of modern warfare. Compared to manned platforms and long-range precision weapons, drones are cost-effective, low-risk, and scalable, enabling persistent operations with limited human exposure (Special Competitive Studies Project, 2023). Their relatively low unit cost allows attackers to trade individual platform survivability for mass and persistence.

The ongoing Russia-Ukraine war provides a clear and contemporary example of this shift. Both sides have repeatedly employed large-scale drone attacks, often involving hundreds of Unmanned Aerial Vehicles (UAVs) in a single engagement, aimed at overwhelming air defence systems through saturation rather than precision (Kyiv Post, 2026; Reuters, 2025a, 2025b, 2025c, 2026). These engagements demonstrate that swarm size itself has become an operational variable.

A key feature of contemporary drone warfare is the deliberate use of heterogeneous swarms, in which a fraction of the platforms is designed to act as decoys or expendable assets. These drones are significantly cheaper and intended to trigger sensors, consume interceptors, and degrade command-and-control effectiveness, thereby increasing the likelihood that a smaller number of high-value threat drones penetrate defensive layers (Center for European Policy Analysis, 2024; Espresso Global, 2024). Such tactics exploit structural limitations of air defence systems, which must sequentially perform detection, identification, tracking, engagement preparation and neutralisation for each incoming target (Joint Air Power Competence Centre, 2021; Joint Research Centre (European Commission), 2025). This sequential processing structure implies that defender performance is fundamentally governed by the cumulative time required to process incoming targets. Each engagement consists of multiple stochastic stages, and under swarm conditions, these stage durations accumulate across targets, creating a compounded timing constraint that directly limits system capacity. When confronted with sufficiently large swarms, even advanced systems experience performance degradation due to finite processing rates and resource constraints. This mechanism is analogous to Distributed Denial-of-Service (DDoS) attacks in computer networks.

Replacing threat drones with decoy drones further enables higher production rates and sustained operational tempo, contributing to prolonged pressure and operational fatigue. Decoy drones are typically reported to cost on the order of a few thousand US dollars, whereas long-range strike drones may cost on the order of two hundred thousand US dollars (Center for European Policy Analysis, 2024). Russia has reportedly introduced several decoy or imitation UAV platforms since 2023, commonly referred to as *Geran/Gerbera* and *Parodiya*, which are intended to emulate the general shape and flight behaviour of the *Shahed-136* drone (Anokhin & Faragasso, 2025; Espresso Global, 2024).

From a modelling perspective, this leads to a structured stochastic system in which individual engagements are composed of multiple sequential phases, and overall defence performance depends on the aggregation of these processes across multiple targets. Capturing such

behaviour requires a representation that can model multi-stage dynamics while remaining tractable under repeated convolution and mixture. In particular, the ability to characterise cumulative engagement time in closed form is essential for assessing system effectiveness under varying swarm sizes and threat compositions.

Motivated by these observations, this work develops a heterogeneous drone swarm model to study the impact of swarm size and threat composition on defender performance. Drone defence is modelled as a multi-stage process, and the impact of the number of drones and the proportion of threat drones on system effectiveness is examined. To capture this structure, we adopt a Matrix-Exponential (ME) representation, which enables modelling of multi-stage engagement dynamics while preserving analytical tractability. While simulation-based approaches can reproduce engagement scenarios, they provide limited analytical insight into scaling behaviour and often require extensive computation across parameter spaces. In contrast, an analytical formulation enables direct characterisation of system performance and exposes the structural drivers of defence effectiveness. In this formulation, the sequential engagement stages correspond to phases of the ME representation, repeated neutralisation of multiple targets corresponds to convolution of these processes, and heterogeneous swarm composition induces a mixture structure over the number of threat drones.

Existing work on UAV swarm defence can be broadly categorised based on modelling approach, including agent-based simulation, optimisation-driven deployment strategies, and system-level dynamic models. These approaches differ in how they represent engagement dynamics and in the level of analytical tractability they provide. Overall, while the literature on UAV swarm defence is growing, analytically tractable models that capture engagement-time dynamics remain relatively limited.

A first class of work focuses on modelling interactions between attacking and defending agents at the swarm level. Brust et al. (2021) propose a swarm-based counter-UAV defence system in which coordinated defensive agents intercept intruding UAVs. Similarly, Gaertner (2013) study UAV swarm tactics using an agent-based simulation framework combined with Markov process modelling. These approaches capture detailed interaction dynamics and coordination effects among agents but rely primarily on simulation-based evaluation and do not provide closed-form characterisation of cumulative engagement time or defender capacity under large-scale swarms.

A second class of work addresses optimal deployment and strategic decision-making for countering UAV swarms. Li et al. (2023) and Wang et al. (2023) investigate optimal placement and positioning of defensive resources to improve coverage and interception effectiveness. Mutzari et al. (2025) model urban multi-drone defence as a sequential Stackelberg security game, focusing on strategic allocation and adversarial decision-making. While these studies provide valuable insights into spatial configuration and resource allocation, they do not explicitly capture the temporal accumulation of engagement processes that governs defender throughput under saturation conditions.

A third line of research emphasises high-fidelity simulation and system-level effectiveness evaluation. Hotters et al. (2025) and Yan et al. (2025) employ detailed simulation frameworks for red-blue confrontation scenarios, while NATO methodologies and related studies provide structured approaches for modelling UAV swarm engagements (Sosa, 2025). Jia et al. (2019) use a system dynamics model to evaluate operational effectiveness of swarming UAV combat systems, capturing aggregate behaviour and feedback effects over time. Although these approaches offer realistic representations of engagement scenarios, they predominantly rely on computational simulation and do not yield analytically tractable expressions for key performance measures such as engagement time distributions or success probabilities under heterogeneous swarm composition.

In contrast to the above approaches, the present work focuses explicitly on the temporal accumulation of engagement processes in drone defence. Rather than modelling detailed interactions or optimising spatial deployment, we represent each neutralisation as a multi-stage stochastic process and analyse how these processes compose under sequential servicing of multiple targets. By adopting an ME framework, the proposed model enables closed-form characterisation of cumulative neutralisation time, including its distribution, moments, and quantiles, while explicitly incorporating uncertainty in swarm composition. This analytical perspective provides complementary insight to simulation-based studies by revealing how defence performance scales with swarm size and identifying the structural bottlenecks that limit system effectiveness.

ME and phase-type distributions are well established in queueing theory and stochastic service systems. In this work, we build upon these existing tools by utilising them in a non-conventional setting. Rather than modelling a system with homogeneous service requirements, we construct a layered formulation that combines a multi-stage ME representation of sequential neutralisation with a Bernoulli-driven mixture over the unknown number of true threats embedded within a finite swarm containing decoys. This structure captures both engagement dynamics and threat uncertainty within a unified analytical framework. Furthermore, the ME representation provides structural flexibility, enabling straightforward modification of engagement stages and parameters without requiring a complete reformulation of the model.

The main contributions of this paper are as follows:

- We introduce a novel analytical framework for modelling drone swarm defence systems using ME representations, a modelling approach not previously applied in this area.
- The proposed framework yields tractable, closed-form expressions for key performance measures while explicitly distinguishing between threat and decoy drones.
- The model captures the probabilistic nature of warfare, including uncertainty in resource availability, sensing and lock-on success, and environmental effects.
- The results provide concrete operational insights, showing that swarm size has a substantially stronger impact on defence success than threat-decoy composition. Under the baseline modelling assumptions, swarms exceedingly roughly ten drones begin to

surpass single-defender swarm-size capacity. This occurs even when the swarm consists of up to 90% decoy drones.

- Despite identical average neutralisation time, higher decoy ratios increase dispersion in neutralisation durations, leading to reduced predictability in defender performance. Defender effectiveness is governed by the slowest stages, which in the adopted model correspond to tracking and neutralisation.

## Nomenclature

|                     |                                                               |
|---------------------|---------------------------------------------------------------|
| $B$                 | subgenerator matrix of the ME representation                  |
| $e'$                | summation column vector of ones                               |
| $E[\cdot]$          | expectation operator                                          |
| $f(t)$              | probability density function                                  |
| $F(t)$              | cumulative distribution function                              |
| $J$                 | strict super diagonal shift matrix                            |
| $K$                 | index of the last threat drone in the neutralisation sequence |
| $N$                 | total number of drones in the initial swarm                   |
| $p$                 | probability that an incoming drone is a threat drone          |
| $p$                 | initial state row vector of the ME representation             |
| $\mathbb{P}(\cdot)$ | probability measure                                           |
| $Q$                 | infinitesimal generator matrix                                |
| $S_i$               | stage $i$ associated with the single-drone neutralisation     |
| $t_q$               | $q$ -quantile of the neutralisation time                      |
| $(\cdot)'$          | transpose operator                                            |
| $\delta_1$          | first canonical row basis vector                              |
| $\lambda_i$         | transition rate associated with stage $i$                     |
| $\sigma$            | standard deviation                                            |
| $\tau$              | prescribed time threshold for successful drone neutralisation |

## 2. Methodology

This section describes the modelling approach used to represent the neutralisation of an incoming drone swarm. A single-drone neutralisation process is treated as a sequence of consecutive operational stages, namely *detection and identification*, *tracking*, *engagement preparation*, and *neutralisation*, whose combined duration is modelled using ME distributions. This representation allows complex, multi-stage processing behaviour to be captured within a unified analytical framework.

The modelling procedure is developed progressively, beginning with the neutralisation of a single drone and then extending to scenarios involving multiple drones. Within the swarm, drones may correspond either to genuine threats or to decoys that are intended to consume defensive resources without posing direct harm, and both types are incorporated within the same framework. Building on the single-drone formulation, the methodology is extended to account for arbitrary swarm sizes and mixed compositions of threat and decoy drones.

In this formulation, a non-discriminatory defence assumption is adopted, whereby threat and decoy drones are processed identically by the defender. This reflects operational settings in which decoy UAV platforms are intentionally designed to emulate the observable characteristics of threat systems, including shape, flight behaviour, and signatures, thereby limiting reliable discrimination prior to engagement (Anokhin & Faragasso, 2025; Espresso Global, 2024). Under such conditions, both threat and decoy drones consume defensive resources in the same manner and are therefore modelled within a unified framework.

### 2.1. ME Distributions

ME distributions are a class of probability distributions characterised by rational Laplace transforms (Bladt & Neuts, 2003; van de Liefvoort & Heindl, 2005). They generalise the family of phase-type distributions. An ME distribution that can be interpreted as the absorption time in a continuous-time Markov chain with  $m$  transient states is of phase type.

Consider a continuous-time Markov chain whose state space is partitioned into  $m$  transient states and a single absorbing state. Without loss of generality, index the transient states as  $1, \dots, m$  and the absorbing state as  $m + 1$ . The infinitesimal generator can then be written in the standard block form

$$Q = \begin{bmatrix} B & -Be' \\ 0 & 0 \end{bmatrix} \quad (1)$$

where  $B$  is an  $m \times m$  matrix,  $-Be'$  is an  $m$ -dimensional column vector, and  $0$  is an  $m$ -dimensional row vector of zeros. Matrix  $B$  describes the transitions between transient states and is called the subgenerator matrix. Vector  $-Be'$  describes the transitions from the transient states to the absorbing state. Vector  $e'$  is a column vector of ones and is called the summation vector.

Throughout this paper, boldface lowercase letters denote row vectors, while boldface uppercase letters denote matrices.

Let  $\mathbf{p}$  be an  $m$ -dimensional row vector that describes the initial distribution of the transient states. This vector is called the initial state vector. Starting from state  $i$ , the probability of being in state  $j$  at time  $t$  is  $\exp(\mathbf{B}t)_{ij}$  (Bladt, 2005). The probability that the process has not yet jumped to the absorbing state by time  $t$  is expressed as

$$\begin{aligned} 1 - F(t) &= \sum_{i=1}^m \sum_{j=1}^m p_i \exp(\mathbf{B}t)_{ij} \\ &= \mathbf{p} \exp(\mathbf{B}t) \mathbf{e}' \end{aligned} \quad (2)$$

Thus, the Cumulative Distribution Function (CDF),  $F(t)$ , and the Probability Density Function (PDF),  $f(t)$ , are given as

$$F(t) = 1 - \mathbf{p} \exp(\mathbf{B}t) \mathbf{e}', \quad t \geq 0 \quad (3)$$

$$f(t) = -\mathbf{p} \exp(\mathbf{B}t) \mathbf{B} \mathbf{e}', \quad t \geq 0 \quad (4)$$

The results shown in (3) and (4) hold for any  $\langle \mathbf{p}, \mathbf{B}, \mathbf{e}' \rangle$  as long as they satisfy the properties of distribution and density functions.

ME distributions have been widely adopted in applied probability and engineering, including risk theory, reliability analysis, and survivability modelling of communication and wireless networks. They enable tractable representation of non-exponential timing behaviour while preserving Markovian structure (Fackrell, 2009; He & Zhang, 2007; Horvath et al., 2016; Kaja et al., 2021; Tekinay et al., 2020). They are also closely related to phase-type distributions and have been extensively used in queueing-based performance modelling (Asmussen, 2003; Latouche & Ramaswami, 1999; Neuts, 1981).


## 2.2. Single-drone Neutralisation Process

Drone neutralisation systems span a wide range of sensing and mitigation methods. Comprehensive surveys of drone neutralisation technologies are available in the literature (Joint Research Centre (European Commission), 2025; Khawaja et al., 2025; Liang et al., 2025; Petsioti et al., 2024; Tu et al., 2024; Zhu et al., 2025). Information on commercial counter-drone systems can be found in existing review studies (González-Jorge et al., 2024). Regardless of the specific technologies employed, such systems inherently operate as a sequential process, since a drone must first be detected, subsequently identified and tracked, and only then subjected to an engagement decision and neutralisation. This staged engagement logic is explicitly reflected in doctrinal counter-UAS frameworks (Joint Air Power Competence Centre, 2021).

The durations of these stages are inherently stochastic, as they depend on a range of uncertain operational factors. Detection and identification latencies are influenced by environmental

conditions, clutter, propagation effects, and sensor performance, while tracking duration depends on target dynamics and sensing quality. Engagement preparation and engagement times are further affected by resource availability, system load, and operational constraints. These sources of variability motivate a probabilistic, stage-based modelling approach.

We model the time required to neutralise a single drone as a four-stage hypoexponential distribution, corresponding to the stages of *detection and identification*, *tracking*, *engagement preparation*, and *neutralisation* as illustrated in Figure 1.



**Figure 1.** Markov representation of a multi-stage drone neutralisation process.

The initial state vector, the subgenerator matrix, and the summation vector corresponding to the single-drone neutralisation time are given by

$$\mathbf{p} = [1 \quad 0 \quad 0 \quad 0], \quad \mathbf{B} = \begin{bmatrix} -\lambda_1 & \lambda_1 & 0 & 0 \\ 0 & -\lambda_2 & \lambda_2 & 0 \\ 0 & 0 & -\lambda_3 & \lambda_3 \\ 0 & 0 & 0 & -\lambda_4 \end{bmatrix}, \quad \mathbf{e}' = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \quad (5)$$

From an operational perspective, the vector  $\mathbf{p}$  indicates that each engagement starts in the detection and identification stage. The subgenerator matrix  $\mathbf{B}$  encodes the progression of the defence system through successive stages (detection, tracking, engagement preparation, and neutralisation), where each diagonal entry represents the rate of leaving a stage and off-diagonal elements represent transitions to the next stage. The vector  $\mathbf{e}'$  aggregates all transient states and enables computation of the probability that the engagement process has completed, i.e., that the drone has been neutralised.

The probability that a drone is neutralised by time  $t$  is given by  $\mathbb{P}(T \leq t)$  and is obtained by substituting the parameters in (5) into the ME distribution function in (3). Let  $\tau$  denote the time threshold within which a threat-drone must be neutralised to prevent harm. The probability of successful neutralisation is then given by  $F(\tau)$ .

### 2.3. Multi-drone Neutralisation Process

An extension of the single-drone neutralisation model is developed for the case where the system is initially presented with an arbitrary number of threat drones. Neutralisation is assumed to proceed sequentially, so that the total neutralisation time for  $k$  drones is given by the sum of  $k$  independent and identically distributed single-drone neutralisation times. Leveraging the closure property of the ME class under convolution, we construct an equivalent ME representation described by a single parameter triplet  $\langle \mathbf{p}_k, \mathbf{B}_k, \mathbf{e}'_k \rangle$  that characterises the neutralisation of  $k$  consecutive drones. This representation provides a compact state-space

formulation and enables exact analysis of the resulting service-time distribution for arbitrary swarm sizes.

Let the neutralisation time of a single drone follow a ME distribution with representation  $\langle \mathbf{p}, \mathbf{B}, \mathbf{e}' \rangle$ . Then, the total time required to neutralise  $k$  consecutive drones admit a ME representation  $\langle \mathbf{p}_k, \mathbf{B}_k, \mathbf{e}'_k \rangle$  where  $\mathbf{B}_k$  is constructed to capture the sequential progression through  $k$  identical neutralisation phases,  $\mathbf{p}_k$  denotes the corresponding initial-state probability vector, and  $\mathbf{e}'_k$  is a column vector of ones of appropriate dimension. Operationally, this construction represents the defender processing drones one at a time, where each phase corresponds to a full detection-to-neutralisation cycle for a single drone.

Let  $\mathbf{J}_k$  denote a  $k \times k$  strict superdiagonal shift matrix, defined by  $(\mathbf{J}_k)_{i,i+1} = 1$  for  $i = 1, \dots, k-1$ , and zeros elsewhere. Matrix  $\mathbf{J}_k$  has ones on the first superdiagonal and zeros on the diagonal and all other entries. We express the  $k$ -fold subgenerator matrix as

$$\mathbf{B}_k = \mathbf{I}_k \otimes \mathbf{B} + \mathbf{J}_k \otimes (-\mathbf{B}\mathbf{e}'\mathbf{p}) \quad (6)$$

The first term  $\mathbf{I}_k \otimes \mathbf{B}$  represents  $k$  independent copies of the single-drone engagement process, while the second term  $\mathbf{J}_k \otimes (-\mathbf{B}\mathbf{e}'\mathbf{p})$  links these copies sequentially, ensuring that the system transitions to the next drone only after completing the current neutralisation.

The resulting  $\mathbf{B}_k$  matrix has a block upper-bidiagonal structure, with each diagonal block equal to the single-drone subgenerator matrix  $\mathbf{B}$  and each super diagonal block representing the transition between consecutive drone-neutralisation stages. The matrix  $\mathbf{B}_k$  has dimension  $km \times km$  where  $m$  denotes the dimension of the single-drone subgenerator matrix  $\mathbf{B}$ . For completeness, the corresponding block matrix form of  $\mathbf{B}_k$  is shown below.

$$\mathbf{B}_k = \begin{bmatrix} \mathbf{B} & -\mathbf{B}\mathbf{e}'\mathbf{p} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{B} & -\mathbf{B}\mathbf{e}'\mathbf{p} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{B} & -\mathbf{B}\mathbf{e}'\mathbf{p} \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{B} \end{bmatrix} \quad (7)$$


In this representation, each diagonal block corresponds to the internal stages of neutralising a single drone, while each superdiagonal block represents the transition from completing one drone engagement to initiating the next. This directly reflects the sequential nature of the defence system.

The initial state vector  $\mathbf{p}_k$  is given in (8). The triplet  $\langle \mathbf{p}_k, \mathbf{B}_k, \mathbf{e}'_k \rangle$ , where  $\mathbf{e}'_k$  is a column vector of ones of length  $km$ , thus defines a valid ME representation for the neutralisation time of  $k$  consecutive drones.

$$\mathbf{p}_k = [\mathbf{p} \quad \mathbf{0} \quad \dots \quad \mathbf{0}] \quad (8)$$

This structure indicates that the process always begins with the first drone entering the detection stage, after which the system progresses through subsequent drones sequentially.

We assume that the process begins when a swarm of  $N$  drones appear at time 0. The sequential neutralisation process by a single defender is illustrated in Figure 2 where each state represents the remaining number of drones.



**Figure 2.** ME state-transition representation of the drone neutralisation process, starting from an initial swarm size of  $N$ .

In contrast to the single-drone case, the defender must neutralise all  $N$  drones in a sequential manner by time  $\tau$  to be successful. The aggregate neutralisation time distribution function is obtained by substituting (7) and (8) into (3). The probability of successful neutralisation is then given by  $F_k(\tau) = 1 - p_k \exp(B_k \tau) e'_k$ .

#### 2.4. Threat and Decoy Drone Modelling

Heterogeneous drone populations are incorporated into the proposed ME framework by distinguishing between threat and decoy drones. Specifically, each incoming drone is assumed to be a threat with probability

$$p \triangleq \mathbb{P}(\text{drone is a threat}) \quad (9)$$

and a decoy with complementary probability  $1 - p$ .

The defender does not have the ability to distinguish between decoy and threat drones. Therefore, the defender adopts the same drone neutralisation process to both kinds of drones. At each engagement, the defender selects a drone to neutralise at random. We are interested in the probability of neutralising all the threat drones by some time  $\tau$  in the presence of decoy drones. As a result, the relative ordering of threat and decoy drones within the neutralisation sequence influences the total time required to neutralise all threats. We define the random variable

$$K = \max\{i \leq N: X_i = 1\}, \quad X_i \sim \text{Bernoulli}(p), \quad \text{i.i.d.} \quad (10)$$

where  $K$  represents the index of the last threat drone in a sequence of  $N$  drones and follows the distribution of the last success index in a finite sequence of Bernoulli trials. Its Probability Mass Function (PMF) is given by

$$\mathbb{P}(K = k) = \begin{cases} (1 - p)^N, & k = 0, \\ p(1 - p)^{N-k}, & k = 1, 2, \dots, N \end{cases} \quad (11)$$

It follows that all threat drones are neutralised by time  $\tau$  with probability

$$\mathbb{P}(\text{success by } \tau) = \mathbb{P}(K = 0) + \sum_{k=1}^N \mathbb{P}(K = k) F_k(\tau) \quad (12)$$

As shown in (12), the overall success probability is a weighted mixture of the neutralisation time distributions corresponding to different numbers of required threat neutralisations. The success criterion in (12) can be naturally incorporated within the ME framework by allowing for a defective initial condition. In particular, with probability  $(1 - p)^N$ , the swarm contains no threat drones, corresponding to a point mass at zero neutralisation time. With the complementary probability, at least one threat drone is present, and the time to success follows a ME distribution obtained as a mixture over the possible numbers of required threat neutralisations. This formulation preserves the analytical tractability of the ME framework while explicitly accounting for uncertainty in swarm composition.

Let  $\delta_1^{(k)}$  denote the first canonical row basis vector of dimension  $k$ . It follows that

$$\mathbf{p}_k = \delta_1^{(k)} \otimes \mathbf{p} \quad (13)$$

Let  $\mathbf{B}_{\text{mix}}$  denote the block-diagonal subgenerator matrix mixing  $\{\mathbf{B}_k\}_{k=1}^N$ . It is given by

$$\mathbf{B}_{\text{mix}} = \mathbf{B}_N \oplus \mathbf{B}_{N-1} \oplus \dots \oplus \mathbf{B}_1 \quad (14)$$

Operationally,  $\mathbf{B}_{\text{mix}}$  captures all possible engagement scenarios corresponding to different numbers of threat drones within the swarm, with each block representing a specific case where a given number of drones must be neutralised.

Each  $\mathbf{B}_k$  has a dimension  $km \times km$ . Consequently,  $\mathbf{B}_{\text{mix}}$  is a square matrix of dimension  $mN(N + 1)/2$ . Let  $w_k \triangleq \mathbb{P}(K = k)$ . The initial state vector of the ME representation is then given by

$$\mathbf{p}_{\text{mix}} = [w_N \mathbf{p}_N \quad w_{N-1} \mathbf{p}_{N-1} \quad \dots \quad w_1 \mathbf{p}_1] \quad (15)$$

The weighting coefficients  $w_k$  reflect the likelihood of each scenario, so that the overall model accounts for uncertainty in swarm composition by combining all possible threat realisations.

By stacking the  $k$  consecutive ME representations into the block-diagonal subgenerator matrix  $\mathbf{B}_{\text{mix}}$  and the weighted initial state vector  $\mathbf{p}_{\text{mix}}$ , the mixture admits a closed-form ME representation. The success-time CDF (time to neutralise all threats) can be written as

$$F_{\text{succ}}(t) = 1 - \mathbf{p}_{\text{mix}} \exp(\mathbf{B}_{\text{mix}} t) \mathbf{e}'_{\text{mix}}, \quad t \geq 0 \quad (16)$$

The resulting success-time distribution is defective, with an atom of size  $w_0$  at  $t = 0$ ; the remaining continuous component is represented by the ME triplet  $\langle \mathbf{p}_{\text{mix}}, \mathbf{B}_{\text{mix}}, \mathbf{e}'_{\text{mix}} \rangle$ , where  $\mathbf{e}'_{\text{mix}}$  is a column vector of ones of length  $mN(N + 1)/2$ .

### 3. Results

This section applies the analytical expressions derived above to evaluate a drone swarm defence system. The stage-wise rate parameters are initialised as  $\lambda_1 = 5$ ,  $\lambda_2 = 0.5$ ,  $\lambda_3 = 2$ , and  $\lambda_4 = 0.5$ , corresponding to mean durations of 200 ms for *detection and identification*, 2 s for *tracking*, 500 ms for *engagement preparation*, and 2 s for *neutralisation*, respectively.

Modern counter-UAS systems exhibit rapid operational timelines consistent with the parameter values adopted. Implementations indicate that detection and identification are performed within sub-second timescales, with recent perception systems demonstrating single-frame processing latencies on the order of 50-100 ms and conventional approaches often exceeding 200 ms (DroneShield Ltd., 2024; Joint Air Power Competence Centre, 2021; Lockheed Martin, 2025; Zhu et al., 2025). Tracking is typically realised through successive updates of target state information, and practical systems describe track formation over short temporal windows involving multiple observations (Argustec Information Technology Co., Ltd., 2024; Blighter Surveillance Systems and Chess Dynamics and Enterprise Control Systems, 2017; Joint Research Centre (European Commission), 2025).


Engagement preparation is generally performed rapidly following track establishment, with system descriptions indicating short handoff and targeting delays prior to countermeasure deployment (Blighter Surveillance Systems and Chess Dynamics and Enterprise Control Systems, 2017; Lockheed Martin, 2025).

Countermeasure mechanisms further indicate that neutralisation timelines depend on the engagement modality: electronic attack and high-power RF techniques are often characterised by near-instantaneous or sub-second effects once applied, whereas directed energy approaches and kinetic responses may require longer but still short-duration engagements on the order of seconds (DroneShield Ltd., 2024; Joint Air Power Competence Centre, 2021; Joint Research Centre (European Commission), 2025).

#### 3.1. Impact of $N$ and $p$

This subsection examines the impact of swarm size  $N$  and per-drone threat probability  $p$  on the distribution of neutralisation time. Let  $F_{succ}(t)$  in (16) denote the CDF of the swarm neutralisation time (time to neutralise all threat drones) for a swarm size of  $N$  with per-drone threat probability  $p$ . By the ME representation, the density function is readily obtained as  $f_{succ}(t)$ .


Figure 3 illustrates the effect of increasing swarm size under the worst-case assumption that all incoming drones are threat drones  $p = 1$ . As the swarm size grows, the neutralisation time distribution shifts progressively to the right, indicating that achieving high-confidence neutralisation becomes increasingly time-consuming.



**Figure 3.** Swarm neutralisation time distribution for  $N \in \{1, 3, 5, 7\}$  for all-threat case  $p = 1$ . **(a)** Probability density function  $f_{succ}(t)$  and **(b)** cumulative distribution function  $F_{succ}(t)$ .

Figure 4 presents a joint view of the discrete and continuous effects induced by mixed swarms with  $N = 7$ . Figure 4(a) shows the probability mass function of the last threat index  $K$  for  $p \in \{0.15, 0.35, 0.55, 0.75\}$ , while Figure 4(b) depicts the corresponding cumulative distribution functions of the neutralisation time.

A first observation is the presence of a jump in the CDF at time zero, whose magnitude increases as  $p$  decreases. This jump corresponds to the probability mass at  $K = 0$ , i.e., the event that no threat drones are present in the swarm. In practical terms, this represents scenarios where all drones are decoys, so the defender achieves immediate success without performing any neutralisation. The size of this atom is directly reflected in Figure 4(a) at  $k = 0$  by the probability mass  $(1 - p)^N$ . This correspondence provides a consistent interpretation of the defective behaviour observed in the time-domain CDFs.



**Figure 4.** Swarm size  $N = 7$  with threat-drone probability  $p \in \{0.15, 0.35, 0.55, 0.75\}$ . **(a)** Probability mass function of the last threat index  $K$  and **(b)** cumulative distribution function  $F_{succ}(t)$ .

In addition, the tail behaviour of the neutralisation time distributions exhibits a notable convergence pattern. Although the CDFs begin at different initial levels, their upper tails approach that of the  $p = 0.75$  case, indicating similar high-confidence neutralisation times. This behaviour contrasts with the earlier scenario of fixed  $p$  and increasing  $N$  where the upper tail progressively shifted rightward. The convergence of the tails can again be understood through Figure 4(a): the PMF of  $K$  is right-skewed (excluding  $k = 0$  mass), with non-negligible probability mass at large indices  $k = 6$  and  $k = 7$ . As a result, even swarms with relatively low per-drone threat probability can, with appreciable likelihood, induce neutralisation dynamics comparable to those of more heavily armed swarms.

### 3.2. Operational Significance

We are interested in answering the following question: will a defender be successful for a given swarm size  $N$  and per-drone threat probability  $p$ ? To address this, we define a success measure as

$$\mathbb{P}(\text{success}) = F_{succ}(\tau) \quad (17)$$


where success corresponds to neutralising the swarm within a prescribed engagement time  $\tau$ . Assuming detection at standoff ranges on the order of several hundred metres and drone speeds of tens of metres per second, this yields an engagement window on the order of tens of seconds. Accordingly,  $\tau = 45$  s is adopted as a representative operational threshold.

Reported counter-UAS capabilities indicate that detection and engagement ranges in practice span from a few hundred metres to several kilometres depending on the sensing modality and system configuration. For example, compact radar systems report drone detection ranges on the order of 250-400 m, whereas integrated counter-UAS platforms provide detection ranges up to several kilometres and engagement or recognition ranges around 1 km (Aaronia AG, 2024; A-Bots, 2026; Argustec Information Technology Co., Ltd., 2024; Joint Air Power Competence Centre, 2021; Joint Research Centre (European Commission), 2025; SpotterRF / Spotter Global, 2024). At the same time, representative UAS platforms, including loitering munitions such as the Shahed-136, operate at speeds on the order of tens of metres per second (Army Recognition, 2024; Army Technology, 2024). These values correspond to engagement timelines ranging from on the order of 10-20 s for short-range sensing to over a minute for longer-range detection, placing a 45 s threshold within a realistic and representative operational regime.

Figure 5 presents the behaviour of the success probability  $F_{succ}(45 \text{ s})$  with respect to  $N$  and  $p$ . Note that the plots do not depict the cumulative distribution function with respect to  $\tau$ ; rather, they show the cumulative distribution function evaluated at  $\tau = 45$  s.

When all drones are threat drones ( $p = 1$ ), the defender can reliably neutralise swarms up to seven drones, achieving a success probability of approximately 0.93. Reducing  $p$  to 0.25 yields a success probability of 0.98. Even when  $p$  is as low as 0.1, the defender can handle a swarm

size of approximately ten drones with high confidence. For swarm sizes exceedingly roughly ten drones, the deployment of additional defenders becomes necessary to maintain a high probability of successful neutralisation.



**Figure 5.** Probability of successful neutralisation  $F_{succ}(45 \text{ s})$  as a function of  $N$  and  $p$ . (a) Success-probability surface over the  $(N, p)$  parameter space and (b) corresponding heatmap with 90%, 95%, and 99% success boundaries.

Instead of a swarm of seven all-threat drones (which represents the defender's capacity for purely threatening swarms), introducing seven decoy drones to form a fourteen-drone swarm increases the probability of defence evasion sharply from 7% to 91%.

As evident from the heatmap in Figure 5(b), swarm size  $N$  has a substantially stronger impact on defence success than the threat probability  $p$  (provided  $p$  is not vanishingly small). The success boundaries are primarily driven by variations in  $N$  rather than  $p$ , indicating a higher sensitivity to swarm size than to threat-drone probability.

Table 1 presents several swarm configurations that achieve the same defence evasion probability. Two interesting observations can be made. First, the presence of a single threat drone in a swarm of size fifty can be as effective as a swarm composed entirely of ten threat drones.

**Table 1.** Different swarms achieving the same evasion probability.

| Swarm | $N$ | $p$  | $\mathbb{E}[\text{threats}]$ | $\mathbb{E}[\text{decoys}]$ | $1 - F_{succ}(45 \text{ s})$ |
|-------|-----|------|------------------------------|-----------------------------|------------------------------|
| 1     | 10  | 1    | 10                           | 0                           | 0.56                         |
| 2     | 11  | 0.48 | 5                            | 6                           | 0.56                         |
| 3     | 15  | 0.14 | 2                            | 13                          | 0.56                         |
| 4     | 50  | 0.02 | 1                            | 49                          | 0.56                         |

Second, the effectiveness of decoys is strongly dependent on  $p$ . Taking Swarm 1 as the baseline configuration, Swarm 2 replaces nearly half of the threat drones with decoys. In this case, the number of introduced decoys is approximately equal to the number of removed threat drones.

Swarm 3 removes eight threat drones and introduces thirteen decoys. In the extreme case of Swarm 4, forty-nine decoys are introduced to substitute nine threat drones.

The trend indicates that adding decoys in place of threats becomes progressively less effective as  $p$  decreases. Therefore, maintaining a fixed level of evasion probability requires a disproportionately large increase in swarm size  $N$  at very small values of  $p$ . This behaviour is reflected in the right-tail structure of the boundaries shown in Figure 5(b), where increases in  $N$  are progressively less likely to cross the success boundary at low values of  $p$ .

Consequently, decoy-based strategies are most effective in an intermediate operating regime where  $p$  remains sufficiently large to avoid this tail behaviour. Beyond this regime, further reductions in  $p$  yield diminishing returns, as increasingly large swarms are required to sustain the same evasion performance.

From a cost standpoint, assuming decoys cost 10% of a threat drone, Swarm 3 is the most economical one: the total attacker cost is reduced by approximately 67% (corresponding to the equivalent of 3.3 threat drones). Even at a higher decoy cost of 20%, the cost reduction remains substantial at roughly 54%. This comparison reflects only unit platform costs and does not account for interceptor usage, sensing resources, operational workload, or command-and-control processes. A more comprehensive study would be required to capture the full cost implications.

### 3.3. Dispersion Effects in Swarm Neutralisation Time

This section examines the temporal behaviour of the neutralisation-time distribution, rather than evaluating performance at a fixed prescribed time. In particular, the dispersion characteristics of the neutralisation-time distribution are analysed. In this context, high quantiles and Coefficient of Variation (CV) serve as informative indicators of dispersion.

Let  $t_q$  denote the  $q$ -quantile of the neutralisation-time distribution, defined as

$$t_q \triangleq \inf\{t \geq 0: F_{\text{succ}}(t) \geq q\} \quad (18)$$


The quantity  $t_{95}$  therefore specifies the minimum time required to successfully neutralise a swarm in 95% of engagements.

Figure 6 presents four mean-equivalent swarm configurations, all exhibiting nearly identical expected neutralisation times of  $18 \pm 0.05$  s, yet displaying markedly different levels of distributional dispersion. The expected neutralisation time is obtained directly from the ME representation as

$$\mathbb{E}[T] = -\mathbf{p}_{\text{mix}} \mathbf{B}_{\text{mix}}^{-1} \mathbf{e}_{\text{mix}} \quad (18)$$

where  $T$  denotes the neutralisation time.

Table 2 provides detailed information for the swarms shown in Figure 6. All four swarm configurations correspond to scenarios in which the defender is expected to succeed with high confidence. On average, the defender requires 18 s to finalise neutralisation in all cases. Nevertheless, the configurations exhibit substantially different  $t_{95}$  values. In particular, neutralisation is completed within  $\approx 50.5$  s for Swarm 1 and within  $\approx 28.7$  s for Swarm 4 in 95% of engagements.



**Figure 6.** CDFs of four mean-equivalent configurations with varying  $t_{95}$  quantiles.

This indicates that Swarm 4 exhibits a more concentrated neutralisation-time distribution than Swarm 1. The reduced variability observed for Swarm 4 is attributable to its large  $p$ , under which nearly all drones are threat drones and must be neutralised to achieve success. In contrast, for Swarm 1, the sequence of neutralisation events plays a more significant role: threat drones may be neutralised either before or after a large number of decoy drones. This uncertainty introduces a mixture of distinct neutralisation-time scales, resulting in increased dispersion.

This effect is also reflected in Table 2 in terms of CV, defined as the ratio of the standard deviation  $\sigma$  to the mean  $\mathbb{E}[T]$ . The second moment of the neutralisation time is obtained using  $\mathbb{E}[T^2] = 2 \mathbf{p}_{\text{mix}} \mathbf{B}_{\text{mix}}^{-2} \mathbf{e}_{\text{mix}}$ .


**Table 2.** Equivalent-mean swarm configurations with varying dispersion.

| Swarm | $N$ | $p$  | $\mathbb{E}[T]$ | $t_{95}$ | CV   |
|-------|-----|------|-----------------|----------|------|
| 1     | 10  | 0.09 | 17.98 s         | 50.5 s   | 1.03 |
| 2     | 7   | 0.2  | 18.04 s         | 39.5 s   | 0.74 |
| 3     | 5   | 0.45 | 18.04 s         | 32.1 s   | 0.48 |
| 4     | 4   | 0.85 | 17.97 s         | 28.7 s   | 0.33 |

### 3.4. Defence Stage Sensitivity Analysis

In this section, the sensitivity of neutralisation performance to the internal defence stages is examined. Specifically, we study how variations in the stage rates  $\{\lambda_i\}$  affect both expected performance and dispersion, with the objective of identifying dominant bottlenecks and stages with limited operational leverage.

Figure 1 shows the defence stages. The baseline rate values are  $\lambda_1 = 5$ ,  $\lambda_2 = 0.5$ ,  $\lambda_3 = 2$ ,  $\lambda_4 = 0.5$ . We are interested in evaluating the relative importance of stages. To quantify their contribution to the overall neutralisation time, we sweep individual rates whilst keeping the remaining rates at their baseline values. Specifically, each baseline rate is multiplied by factors of 0.2, 0.5, 1.25 and 2. Figure 7 presents the resulting sensitivity for a swarm of size  $N = 7$  with a threat probability of 0.2.




**Figure 7.** Defence stage sensitivity analysis for a swarm of size  $N = 7$  with threat drone probability  $p = 0.2$ . Each subfigure shows the impact of a single defence stage. The baseline rate value is multiplied by 0.2, 0.5, 1.25, and 2, while the other parameters are fixed. **(a)** Sensitivity to  $\lambda_1$ , **(b)** sensitivity to  $\lambda_2$ , **(c)** sensitivity to  $\lambda_3$  and **(d)** sensitivity to  $\lambda_4$ .

Sensitivity is highest for stages with the lowest baseline rates, namely stages 2 and 4. These stages behave identically because their baseline rates are the same. These two stages represent the *tracking* and *neutralisation* processes, respectively. This identical behaviour demonstrates that the ordering of stages does not affect overall neutralisation performance. Improvements to these stages yield substantially larger gains than improvements to faster stages, as they act as system bottlenecks. Once a bottleneck stage is sufficiently sped up, performance becomes dominated by the next slowest stage, and the marginal benefit diminishes.

Figure 8 presents the sensitivity analysis in terms of  $t_{95}$  and  $F_{succ}(45\text{ s})$  for the same swarm configuration. The maximum observed changes in  $t_{95}$  and  $F_{succ}(45\text{ s})$  over the rate sweeps are shown, further highlighting the dominant influence of bottleneck stages on overall defence performance.

With the baseline rates, are  $\lambda_1 = 5$ ,  $\lambda_2 = 0.5$ ,  $\lambda_3 = 2$ ,  $\lambda_4 = 0.5$ , the defence success probability is 0.66 for a swarm of size  $N = 10$  with threat drone probability  $p = 0.4$ . Increasing the tracking-stage rate  $\lambda_2$  by 100% to a rate value of 1 raises the success probability to 0.93 while jointly increasing both  $\lambda_2$  and  $\lambda_4$  by 50% to rate values of 0.75 improves it to 0.97. These results demonstrate that performance is not determined solely by aggregate rate sums; rather, individual stage rates, particularly bottlenecks, play a decisive role.



**Figure 8.** Defence stage sensitivity analysis for a swarm of size  $N = 7$  with threat drone probability  $p = 0.2$ . **(a)** Change in  $t_{95}$  when stage rate values are multiplied by 0.2 and 2 and **(b)** change in defence success probability  $F_{succ}(45\text{ s})$  when stage rate values are multiplied by 0.2 and 2.

### 3.5. System Performance Across Parameter Regimes

In this section, we examine the system behaviour at a holistic level. The objective is to study how the system performs as a whole and to draw practical conclusions by jointly considering swarm size  $N$ , threat probability  $p$ , and stage rates  $\lambda_i$ . In addition, this section examines the robustness of the observed results across different defender and attacker characteristics, as well as across varying operational time scales and scenarios.

Capacity is defined as the maximum swarm size  $N$  for which the defence success probability satisfies  $F_{\text{succ}}(45) \geq 0.93$ . This threshold represents a high-confidence operational criterion and provides a consistent basis for comparing performance across different configurations.

Table 3 presents the defender capacity under varying stage rates and threat probabilities. The results indicate that capacity is primarily governed by the stage rates, with substantial gains achieved as processing speeds increase. Faster stage configurations lead to a significant expansion in the maximum manageable swarm size. In contrast, variations in the threat probability  $p$  result in only moderate changes in capacity. Even for low values of  $p$ , the increase in capacity remains limited compared to the improvements obtained through faster stage rates, indicating that processing bottlenecks dominate system performance.

**Table 3.** Defender swarm-size capacity  $N$  for  $F_{\text{succ}}(45) \geq 0.93$  under varying stage rates and threat probabilities  $p$ .

| Stage Rates |             |             |             | Capacity $N$ |           |            |           |
|-------------|-------------|-------------|-------------|--------------|-----------|------------|-----------|
| $\lambda_1$ | $\lambda_2$ | $\lambda_3$ | $\lambda_4$ | $p = 1$      | $p = 0.7$ | $p = 0.45$ | $p = 0.1$ |
| 5           | 0.5         | 2           | 0.5         | 7            | 7         | 7          | 9         |
| 3           | 1           | 2           | 1           | 12           | 13        | 13         | 15        |
| 8           | 4           | 8           | 1           | 24           | 24        | 25         | 28        |
| 10          | 5           | 4           | 2           | 37           | 37        | 38         | 40        |

For example, when all stages—detection and identification, tracking, engagement preparation, and neutralisation—are performed on the order of one second on average, the defender's swarm-size capacity increases substantially, reaching up to 37 drones. This highlights the direct operational benefit of reducing stage durations, as even moderate improvements across stages can translate into significant gains in achievable swarm capacity.

Table 4 provides a complementary perspective by evaluating the success probability for fixed swarm sizes.

**Table 4.** Success probability  $F_{\text{succ}}(45)$  under varying swarm size  $N$ , threat probability  $p$ , and stage rates.

| Stage Rates |             |             |             | $p$ | $F_{\text{succ}}(45)$ |          |          |
|-------------|-------------|-------------|-------------|-----|-----------------------|----------|----------|
| $\lambda_1$ | $\lambda_2$ | $\lambda_3$ | $\lambda_4$ |     | $N = 5$               | $N = 20$ | $N = 35$ |
| 5           | 0.5         | 2           | 0.5         | 0.7 | 0.9975                | <0.001   | <0.001   |
|             |             |             |             | 0.1 | 0.9996                | 0.3293   | 0.0678   |
| 3           | 1           | 2           | 1           | 0.7 | >0.999                | 0.0653   | <0.001   |
|             |             |             |             | 0.1 | >0.999                | 0.6390   | 0.1320   |
| 8           | 4           | 8           | 1           | 0.7 | >0.999                | 0.9977   | 0.1331   |
|             |             |             |             | 0.1 | >0.999                | 0.9996   | 0.6043   |
| 10          | 5           | 4           | 2           | 0.7 | >0.999                | >0.999   | 0.9879   |
|             |             |             |             | 0.1 | >0.999                | >0.999   | 0.9975   |

For small swarms (e.g.,  $N = 5$ ), the system consistently achieves near-certain success across all parameter settings, indicating operation well below capacity. However, as swarm size increases, performance degrades rapidly. For intermediate and large swarm sizes, the success probability drops sharply, often reaching values effectively indistinguishable from zero under slower stage configurations. While reducing the threat probability improves performance, this effect remains secondary and does not offset the impact of increasing  $N$ .

Overall, these results show that swarm size appears to be the dominant factor governing system performance within the considered modelling assumptions. While threat composition influences outcomes, its effect is comparatively limited. The consistency of this behaviour across all tested configurations indicates that the dominance of swarm size is a robust property of the system.

### **3.6. Simulation-Based Validation of the Analytical Framework**

To assess the accuracy of the analytical derivations and underlying stochastic model, Monte Carlo simulations are performed and compared against the analytical results. In each trial, the number of threat drones is generated according to the Bernoulli model, and the total neutralisation time is obtained by summing independent realisations of the four-stage sequential process.


Figure 9 presents the analytical CDF together with empirical CDFs obtained from Monte Carlo simulations for increasing numbers of trials. For lower sample sizes (e.g., 100 and 200 trials), the empirical curves exhibit visible fluctuations due to sampling variability. As the number of trials increases, the Monte Carlo estimates converge toward the analytical distribution, with near-perfect agreement observed for 10,000 trials.

This agreement is observed across the full support of the distribution, indicating that the analytical model accurately captures both the central behaviour and the tail of the neutralisation time. The maximum deviation between analytical and empirical CDFs decreases with increasing sample size, demonstrating convergence of the simulation to the theoretical result.

A similar trend is observed in Figure 10, where the expected neutralisation time is evaluated as a function of the threat probability  $p$ . Even with a relatively small number of trials, the Monte Carlo estimates follow the analytical curve closely, while higher sample sizes yield nearly indistinguishable agreement across the entire range of  $p$ .

## **4. Discussion**


This work studies heterogeneous drone swarm engagements against a single defender and characterises defender swarm-size capacity under a range of operational scenarios. Using an ME modelling framework, closed-form expressions are derived for the probability of defence success and the expected neutralisation time. The threat-drone probability, swarm size, and per-stage defence rates are treated as tunable parameters, enabling systematic exploration of their impact on system performance.



**Figure 9.** Analytical and Monte Carlo CDFs of the neutralisation time for a swarm of size  $N = 7$  with threat probability  $p = 0.55$ . The Monte Carlo estimate converges to the analytical distribution as the number of trials increases. **(a)** 100 trials, **(b)** 200 trials, **(c)** 1,000 trials, and **(d)** 10,000 trials.

The results presented in the previous section should be interpreted in the context of the modelling assumptions adopted, including the sequential engagement structure and single-defender abstraction.

Our results show that swarm size is a more critical factor than the threat-drone ratio. Increasing the decoy ratio primarily increases uncertainty in defender performance rather than proportionally extending capacity. When cost considerations are taken into account, decoy compositions in the approximate range of 50%-90% appear to constitute a practically relevant regime. From a practical perspective, this range represents a reasonable operating region for the UAV types and pricing assumptions considered in this simplified model. For extremely large swarm sizes and vanishing threat probabilities, economic and logistical constraints are expected to dominate, and a more comprehensive cost analysis would be required to obtain a clearer system-level picture.



**Figure 10.** Expected neutralisation time as a function of the threat probability  $p$  for a swarm of size  $N = 7$ , with  $p \in [0.1, 0.9]$ . Analytical results are compared with Monte Carlo estimates. **(a)** 200 trials and **(b)** 10,000 trials, showing improved agreement with increasing sample size.

The defender's swarm-size capacity is largely governed by the bottleneck stage within the sequential defence process. Improving the slowest stage yields the most significant performance gains, whereas accelerating non-limiting stages provides diminishing returns. Given the growing number of reported drone attacks, this result suggests that each defence stage must operate, on average, on the order of one second or faster. If this requirement cannot be met, deploying additional defenders with overlapping coverage becomes necessary to maintain a high probability of successful defence.

Several modelling assumptions should be noted. The defender is assumed to be unable to distinguish between decoy and threat drones. This assumption is intentional, as the objective of this work is to quantify the impact of high-volume swarms, and it has been reported that decoy drones may exhibit signatures and footprints similar to those of threat drones in operational environments. The problem of reliable decoy-threat discrimination is therefore considered orthogonal to the present analysis. In addition, cooperative drone behaviours and advanced swarm tactics are not modelled, and the framework focuses exclusively on a single defender.

In practice, however, defence systems may exhibit partial discrimination capability. This can be incorporated within the proposed framework by augmenting the single-target service model with an initial classification stage, represented as an additional phase in the ME structure, after which non-threatening targets exit the process early while predicted threats proceed to the remaining engagement stages. In this interpretation, the modification corresponds to an adjusted initial distribution and transition structure with an early absorbing branch, while preserving the overall probabilistic mixture formulation. Such a formulation reduces the effective load on the defender and shifts the system from a volume-driven regime toward a threat-driven regime as discrimination improves. Accordingly, the present formulation can be interpreted as a conservative upper-bound stress case.

It should also be noted that, even for a single target, the neutralisation process may in practice be more complex than the four-stage sequential representation adopted here. Real-world defence chains may include feedback loops, repeated identification or tracking updates, and partial overlap between functions enabled by sensor fusion and pipelined processing. Such effects could be represented through a more general state-transition structure, for example by introducing backward transitions between phases or additional coupled states within the underlying ME model. While this would provide a richer description of single-target engagement dynamics, it would also increase model complexity substantially and move the analysis away from the parsimonious staged structure adopted in this work. For this reason, we retain the present four-stage sequential abstraction as a first-order representation of the single-target defence timeline and leave explicitly feedback-driven or partially parallel stage dynamics for future work.

Furthermore, the defence process is modelled as strictly sequential, corresponding to a single effective engagement capability. This assumption provides a conservative baseline and is consistent with a single-server interpretation in queueing terms. In practice, modern systems may support partial parallelism, including concurrent tracking or multi-target engagement capabilities, which would effectively increase service capacity and shift the system bottleneck. From a queueing perspective, transitioning to a multi-server regime reduces the effective service time per target and increases the maximum swarm size that can be handled within a given time window. Asymptotically, this leads to an approximately linear scaling in expected defender capacity with the number of parallel engagement resources. In the limiting case of high parallelism, system behaviour becomes increasingly governed by the slowest stage rather than the sum of all stages. Importantly, the qualitative insights of this work, including the identification of bottleneck stages and the dominant role of swarm size, remain unchanged, and the reported results can be interpreted as conservative lower-bound estimates of defender performance.

In addition, multi-layered or networked defence architectures can be interpreted within this framework as interconnected service stages or parallel service nodes operating across different layers. Such configurations effectively distribute the engagement process spatially and functionally, further increasing system capacity and resilience. A detailed analytical treatment of layered or networked defence structures is left for future work.

Despite these limitations, the proposed framework remains flexible and readily admits extensions. Multiple defenders, successive swarm arrivals, target differentiation, feedback mechanisms between stages, and probabilistic engagement failures can all be incorporated within the same ME modelling structure without fundamentally altering the analytical form.

In most cases, these extensions can be represented through appropriate modifications of the initial distribution vector and subgenerator matrix, which define the phase structure and transition dynamics of the process. Since all parameters are tuneable, the framework is also applicable to a wide range of operational settings, including urban warfare scenarios where reduced engagement times and tighter performance requirements are critical.

From this perspective, the present model can be viewed as one instantiation within a broader modelling class. The key advantage of this approach is that increasingly complex system behaviours can be captured through parameter and structural adjustments, rather than requiring a complete reformulation of the analysis. This provides a structured way to incorporate additional system features while preserving analytical tractability and allows the framework to be adapted to a range of defence configurations under uncertainty.

## **5. Conclusion**

This paper introduced a ME framework for modelling a single defender engaging drone swarms composed of both decoy and threat drones, as observed in recent conflict environments. The results indicate that high-volume swarms with sizes exceeding approximately ten drones may overwhelm a typical defender under the baseline modelling assumptions, even when the decoy composition is as high as 90%. Achieving a swarm-handling capacity on the order of 40 drones requires, under the same assumptions, that the defender complete detection, identification, tracking, engagement preparation, and neutralisation stages on average within one second each.

The proposed framework provides a baseline for extending the analysis to more complex scenarios, including coordinated drone swarms, multiple defenders, and successive swarm arrivals over time. A natural extension is the incorporation of parallel or multi-target engagement capabilities, which would require extending the ME framework to multi-channel service structures and a modified mixture construction. Such an extension would enable explicit modelling of systems capable of simultaneous tracking and engagement, providing a more complete characterisation of defender capacity under modern operational conditions.

## **Conflict of Interest**

The author declares no conflict of interest related to this work.

## **Funding**

This research received no external funding.

## **Acknowledgments**

The author acknowledges ASELSAN Inc. for providing research, computational, and software resources supporting this work. The author also appreciates general technical discussions with colleagues during the course of the study.

## References

- Aaronia AG. (2024). *Aartos anti-drone jammer*. Retrieved March 25, 2026, from <https://drone-detection-system.com/sensor-types-overview/anti-drone-jammer/>
- A-Bots. (2026). *Drone detection in 2026*. Retrieved March 25, 2026, from <https://abots.com/blog/drone-detection-2026/>
- Anokhin, I., & Faragasso, S. (2025). *Updated analysis of Russian Shahed-136 deployment against Ukraine*. Institute for Science and International Security. <https://isis-online.org/isis-reports/may-2025-updated-analysis-of-russian-shahed-136-deployment-against-ukraine>
- Argustec Information Technology Co., Ltd. (2024). *Argustec wukong anti-drone system*. Argustec Information Technology Co., Ltd. Retrieved March 25, 2026, from <https://www.argustecn.com/antidronesystem.html>
- Army Recognition. (2024). *Shahed-136 loitering munition technical data*. <https://armyrecognition.com/military-products/army/unmanned-systems/unmanned-aerial-vehicles/shahed-136-loitering-munition-kamikaze-suicide-drone-technical-data>
- Army Technology. (2024). *Shahed-136 kamikaze UAV (Iran)*. <https://www.army-technology.com/projects/shahed-136-kamikaze-uav-iran/>
- Asmussen, S. (2003). *Applied probability and queues*. Springer.
- Bladt, M. (2005). A review on phase-type distributions and their use in risk theory. *ASTIN Bulletin*, 35(1), 145–161. <https://doi.org/10.2143/AST.35.1.583170>
- Bladt, M., & Neuts, M. F. (2003). Matrix-exponential distributions: Calculus and interpretations via flows. *Stochastic Models*, 19(1), 113–124. <https://doi.org/10.1081/STM-120018141>
- Blighter Surveillance Systems, Chess Dynamics, & Enterprise Control Systems. (2017). *AUDS anti-UAV defence system brochure*. <https://www.atri.it/wp-content/uploads/2022/05/auds-brochure.pdf>
- Brust, M. R., Danoy, G., Stolfi, D. H., & Bouvry, P. (2021). Swarm-based counter UAV defense system. *Discover Internet of Things*, 1(1), 2. <https://doi.org/10.1007/s43926-021-00002-x>
- Center for European Policy Analysis. (2024). *The phony war: Ukraine and Russia's decoy drones*. Retrieved January 7, 2026, from <https://cepa.org/article/the-phony-war-ukraine-and-russias-decoy-drones/>
- DroneShield Ltd. (2024). *RfAI: Radio frequency artificial intelligence*. Retrieved March 25, 2026, from <https://www.dronesshield.com/capabilities#rfai>
- Espresso Global. (2024). *Weapons of mass deception: What are decoy drones and how Russia and its allies are trying to bypass air defense*. Retrieved January 7, 2026, from

- <https://global.espreso.tv/russia-ukraine-war-weapons-of-mass-deception-what-are-decoy-drones-and-how-russia-and-its-allies-are-trying-to-bypass-air-defense>
- Fackrell, M. (2009). An alternative characterization for matrix exponential distributions. *Advances in Applied Probability*, 41(4), 1005–1022. <https://doi.org/10.1239/aap/1261669583>
- Gaertner, U. (2013). *UAV swarm tactics: An agent-based simulation and markov process analysis* (Master's thesis). Naval Postgraduate School. Monterey, CA, USA.
- González-Jorge, H., Aldao, E., Veiga-López, F., Fontenla-Carrera, G., Balvís, E., & Ríos-Otero, E. (2024). Counter drone technology: A review. *Preprints*, 2024. <https://doi.org/10.20944/preprints202402.0551.v1>
- He, Q. M., & Zhang, H. (2007). On matrix exponential distributions. *Advances in Applied Probability*, 39(1), 271–292. <https://doi.org/10.1239/aap/1175266478>
- Horvath, A., Scarpa, M., & Telek, M. (2016). Phase type and matrix exponential distributions in stochastic modeling. In L. Fiondella & A. Puliafito (Eds.), *Principles of performance and reliability modeling and evaluation: Essays in honor of kishor trivedi on his 70<sup>th</sup> birthday* (pp. 3–25). Springer International Publishing. [https://doi.org/10.1007/978-3-319-30599-8\\_1](https://doi.org/10.1007/978-3-319-30599-8_1)
- Hotters, D., Kalliatanakis, N., Kuhnert, F., Prakasha, P. S., van Heur, R., & Kaya, B. (2025). A novel approach to quantify counter-drone system effectiveness against UAS swarms. *Proceedings of CEAS Aerospace Europe Conference (CEAS 2025)*.
- Jia, N., Yang, Z., & Yang, K. (2019). Operational effectiveness evaluation of the swarming UAVs combat system based on a system dynamics model. *IEEE Access*, 7, 25209–25224. <https://doi.org/10.1109/ACCESS.2019.2898728>
- Joint Air Power Competence Centre. (2021). *A comprehensive approach to countering unmanned aircraft systems*. Retrieved January 9, 2026, from <https://www.japcc.org/wp-content/uploads/A-Comprehensive-Approach-to-Countering-Unmanned-Aircraft-Systems.pdf>
- Joint Research Centre. (2025). *Technical developments in counter-drone technology: C-UAS detection, tracking and identification technology* (EUR Scientific and Technical Research Report). Publications Office of the European Union. [https://publications.jrc.ec.europa.eu/repository/bitstream/JRC140692/JRC140692\\_01.pdf](https://publications.jrc.ec.europa.eu/repository/bitstream/JRC140692/JRC140692_01.pdf)
- Kaja, H., Paropkari, R. A., Beard, C., & Van De Liefvoort, A. (2021). Survivability and disaster recovery modeling of cellular networks using matrix exponential distributions. *IEEE Transactions on Network and Service Management*, 18(3), 2812–2824. <https://doi.org/10.1109/TNSM.2021.3090355>
- Khawaja, W., Ezuma, M., Semkin, V., Erden, F., Ozdemir, O., & Guvenc, I. (2025). A survey on detection, classification, and tracking of UAVs using radar and communications systems. *arXiv*. <https://arxiv.org/abs/2402.05909>

- Kyiv Post. (2026). *Russia fires iskander missile, swarms Ukraine with 95 drones overnight*. Reuters. Retrieved January 7, 2026, from <https://www.kyivpost.com/post/67565>
- Latouche, G., & Ramaswami, V. (1999). *Introduction to matrix analytic methods in stochastic modeling*. SIAM.
- Li, N., Su, Z., Ling, H., Karatas, M., & Zheng, Y. (2023). Optimization of air defense system deployment against reconnaissance drone swarms. *Complex System Modeling and Simulation*, 3(2), 102–117. <https://doi.org/10.23919/CSMS.2023.0003>
- Liang, X., Qi, Y., Chen, G., & Meng, G. (2025). Research on a distributed strategy for UAV detection and tracking of specific targets. *Artificial Intelligence Review*, 58(12), 388. <https://doi.org/10.1007/s10462-025-11364-x>
- Lockheed Martin. (2025). *Counter unmanned aerial systems (C-UAS)*. Lockheed Martin Corporation. Retrieved March 25, 2026, from <https://www.lockheedmartin.com/en-us/capabilities/counter-unmanned-aerial-systems.html>
- Mutzari, D., Deb, T., Molinaro, C., Pugliese, A., Subrahmanian, V. S., & Kraus, S. (2025). Defending a city from multi-drone attacks: A sequential stackelberg security games approach. *Artificial Intelligence*, 349, 104425. <https://doi.org/10.1016/j.artint.2025.104425>
- Neuts, M. F. (1981). *Matrix-geometric solutions in stochastic models: An algorithmic approach*. Dover Publications.
- Petsioti, P., Życzkowski, M., Brewczyski, K., Cichulski, K., Kaminski, K., Razvan, R., Mohamoud, A., Church, C., Koniaris, A., De Cubber, G., & Doroftei, D. (2024). Methodological approach for the development of standard C-UAS scenarios. *Open Research Europe*, 4(240). <https://doi.org/10.12688/openreseurope.18339.1>
- Reuters. (2025a). *Russia launches war's largest drone attack on Ukraine, Kyiv says*. Retrieved January 7, 2026, from <https://www.reuters.com/world/europe/russia-launches-warlargest-drone-attack-ukraine-kyiv-says-2025-02-23/>
- Reuters. (2025b). *Russia pounds Kyiv, other regions in mass drone and missile attack*. Retrieved January 7, 2026, from <https://www.reuters.com/business/aerospace-defense/russianlaunches-major-drone-missile-attack-ukraine-still-ongoing-2025-09-28/>
- Reuters. (2025c). *Three killed in biggest Ukrainian drone attack on Moscow region*. Retrieved January 7, 2026, from <https://www.reuters.com/world/europe/ukraine-launches-droneattacks-targeting-moscow-russia-says-2025-03-11/>
- Reuters. (2026). *Ukraine targets Moscow daily with drones this year, Russia says, in apparent escalation*. Retrieved January 7, 2026, from <https://www.reuters.com/business/aerospace-defense/ukraine-targets-moscow-daily-with-drones-this-year-russia-saysapparent-2026-01-04/>
- Sosa, C. (2025). *Swarm of UAS modeling and simulation methodology* (tech. rep.) [STO-MPSAS- 192]. NATO Science & Technology Organization, Systems Analysis & Studies Panel.

<https://publications.sto.nato.int/publications/STO%20Meeting%20Proceedings/STO-MP-SAS-192/MP-SAS-192-31.pdf>

Special Competitive Studies Project. (2023). *The drone advantage in modern warfare*. Retrieved January 7, 2026, from <https://thesvi.org/the-drone-advantage-in-modern-warfare/>

Spotter Global. (2024). *360° counter-UAS radar datasheet*.  
<https://www.milestonesys.com/globalassets/marketplace/uploaded-assets/00120000013cnzqaai/360-counter-uas-radar-datasheet-rev0042.pdf>

Tekinay, M., Beard, C., & van de Liefvoort, A. (2020). Partial packet duplication: Control of fade and non-fade duration outages using matrix exponential distributions. *IEEE Transactions on Vehicular Technology*, 69(5), 5652–5656.  
<https://doi.org/10.1109/TVT.2020.2978848>

Tu, X., Zhang, C., Zhuang, H., Liu, S., & Li, R. (2024). Fast drone detection with optimized feature capture and modeling algorithms. *IEEE Access*, 12, 108374–108388.  
<https://doi.org/10.1109/ACCESS.2024.3438991>

van de Liefvoort, A., & Heindl, A. (2005). Approximating matrix-exponential distributions by global randomization. *Stochastic Models*, 21(2-3), 669–693.  
<https://doi.org/10.1081/STM-200056231>

Wang, C., Wu, A., Hou, Y., Liang, X., Xu, L., & Wang, X. (2023). Optimal deployment of swarm positions in cooperative interception of multiple UAV swarms. *Digital Communications and Networks*, 9(2), 567–579.  
<https://doi.org/10.1016/j.dcan.2022.04.002>

Yan, J., LI, J., Li, J., Liu, C., & Guo, Y. (2025). Simulation and deduction method for UAVs swarm in air-to-ground attack scenarios. *Proceedings of the 2025 17th International Conference on Computer Modeling and Simulation*, 135–144.  
<https://doi.org/10.1145/3761668.3761690>

Zhu, Y., Yin, L., & Zhao, K. (2025). Millisecond-response tracking and gazing system for UAVs: A domestic solution based on "phytium + cambricon". *arXiv*  
<https://arxiv.org/abs/2509.04043>

# Comparative Review and Analysis of Artificial Intelligence Based Learning Systems Used in Air Combat Simulation Environments

Emre Şafak<sup>1,\*</sup> , Tolga Erol<sup>1</sup> , Hüseyin Furkan Ceran<sup>1</sup> , Emre Ayvaz<sup>2</sup> 

## Abstract

The advancement of Artificial Intelligence (AI) technology is also manifesting itself in the defence sector. This study comprehensively reviews and comparatively analyses AI-based learning systems used in air combat simulation environments. The studies were categorised according to within-visual-range and beyond-visual-range scenarios and compared based on the methods used and the results obtained. These studies demonstrate that AI agents can develop realistic tactics, share situational awareness, and compete effectively with human pilots. In this study, methods such as Proximal Policy Optimisation (PPO) and Hierarchical Multi-Agent Reinforcement Learning (HMARL) were observed to be particularly prominent in terms of performance. Significant problems such as long training durations, high computational costs, partial observability, and sensitivity to sensor noise are the major challenges encountered. As a result of the study, it is concluded that intelligent agents developed in air combat simulation environments can enhance training efficiency by serving as teammates and opponents in pilot training.

**Keywords:** Air combat simulation, reinforcement learning, military simulation training, deep reinforcement learning, multi-agent systems, AI-driven tactical decision making, human-in the-loop.

---

Received: 10.11.2025

Accepted: 24.04.2026

Review Article

<sup>1</sup> HAVELSAN Inc., Ankara, Türkiye

<sup>2</sup> Turkish Air Force Command, Ankara, Türkiye

\* Corresponding author.

✉ esafak@havelsan.com.tr

**Cite this article as:**

<https://doi.org/10.65834/jdsi.12.30>

Şafak, E., Erol, T., Ceran, H.F. & Ayvaz, E. (2026). Comparative Review and Analysis of Artificial Intelligence Based Learning Systems Used in Air Combat Simulation Environments. *Journal of Defence and Security Industries: Strategy and Technology* 1(2), 88-109.

## 1. Introduction

Air combat environments are evolving towards sixth-generation platforms, collaborative Manned-Unmanned Teaming (MUM-T) concepts, and increasingly complex rules of engagement. Therefore, the simulation infrastructures used for pilot training and tactical development must also become "learning" and adaptable at the same pace. This necessity represents the evolution from rule-based Computer Generated Forces (CGF) approaches to today's Deep Reinforcement Learning (DRL) and Multi-Agent Reinforcement Learning (MARL)-based autonomous agents. Current studies integrating DRL into air combat aim to generate more realistic, adaptable, and less labour-intensive tactics. The ability to generate more engagements in a virtual environment than a human pilot could experience in real life will usher in a new era in pilot training. Unlike High Level Architecture and Distributed Interactive Simulation (HLA and DIS) simulators that run predefined scenarios, the Live-Virtual-Constructive (LVC) paradigm demands intelligent simulators. These systems must generate tactics, share situational awareness, and interact with humans online (Neville et al., 2015).

AI agents in high-fidelity air combat simulators are no longer static entities executing predefined scenarios. They have evolved into learning systems capable of reacting to opponents' manoeuvres, explaining discovered tactics, and sharing situational awareness across platforms. Mei et al. (2024) automated basic air-to-air manoeuvre decisions by framing single-aircraft manoeuvre generation with DRL. Kong et al. (2023) focused on coordinated behaviour by designing a Hierarchical Multi-Agent Reinforcement Learning (HMARL) framework for short-range, multi-aircraft scenarios. Sun et al. (2021) developed a self-play-based multi-agent policy gradient that enabled the emergence of tactics. Dantas et al. (2025) contributed to the DRL pipeline regarding engagement timing by using supervised learning to determine the optimal conditions and timing for missile launches in the Beyond Visual Range (BVR) environment. Toubman and Roessingh (2015) strengthened the practical aspects of guiding agent behaviour through novel reward-method development in training simulations, bringing it closer to application. Thus, these diverse yet complementary studies indicate that DRL/MARL-based agents for both Within Visual Range (WVR) and BVR air combat are approaching the maturity required to meet operational needs (Mei et al., 2024; Kong et al., 2023; Sun et al., 2021). Recent work has also extended this picture by addressing explainability, partial observability, and human-in-the-loop integration in increasingly realistic training settings. Despite this rapid progress, a significant gap remains in the literature: because these AI-based learning approaches are reported with different assumptions, fidelity levels, and human-in-the-loop integrations, it is still unclear which type of learning system a force structure should prefer. The objective of this article is to systematically explain AI-based learning systems used in air combat simulation environments according to the criteria of learning paradigm (single-agent DRL, MARL, hierarchical/curriculum learning), human-machine interaction (human-in-the-loop, explainability, super-agent integration), and operational realism (partial observability, obscured vision, multi-aircraft coordination, missile-launch timing). This study will enable command, control, and training authorities to distinguish

the differences between high-fidelity but costly-to-train systems and simpler yet rapidly adaptable systems. Thus, it will allow them to select the most suitable agent type for their specific network-centric warfare architecture.

This article reviews and provides a comparative analysis of studies utilizing artificial intelligence (AI) technology in air combat simulation environments. The second section covers the AI algorithms used in air combat simulation environments and the examined studies. The third section presents a detailed comparative analysis of the reviewed studies. The fourth section contains the conclusions reached from the examined studies and suggests potential future research directions.

## **2. Methodology**


In the study, AI algorithms used in air combat simulation environments were first explained. Studies were categorized into WVR, BVR, and hybrid scenarios. Literature review was conducted using the IEEE Xplore, ScienceDirect, SpringerLink, Scopus, Web of Science, and ACM Digital Library databases. The search included English and Turkish journal and conference publications from the period 2012–2024, with selected early-access 2025 records included only where publication details could be verified. Specific Inclusion Criteria (IC) and Exclusion Criteria (EC) were applied to the selected studies to enhance the effectiveness of the research. These criteria are presented as:

- **IC1.** Studies must have been published in a conference, journal, press, or similar venue.
- **IC2.** Sufficient information must be provided in the abstract and the full text of the study.
- **IC3.** Studies using AI in air combat simulations are included in the research.
- **EC1.** Books, letters, notes and patents are not included in the review.
- **EC2.** Studies published in languages other than English and Turkish between 2012 and 2025 were excluded.
- **EC3.** Studies must be accessible; otherwise, they are not included in the review.
- **EC4.** Purely, theoretical articles lacking simulations are not included in the research.

### **2.1. AI Algorithms Used in Air Combat Simulation Environments**

#### **2.1.1. Reinforcement Learning**

Reinforcement Learning (RL) is a branch of machine learning that focuses on how agents can learn to make decisions through trial and error to maximize the rewards received. RL enables machines to learn by interacting with an environment and receiving feedback based on their actions. This feedback occurs in the form of rewards or penalties (Sivamayil et al., 2023). In reinforcement learning, the algorithm works with an unlabelled dataset focused on a specific outcome. Each action the algorithm takes to explore the dataset generates positive, negative, or neutral feedback. This feedback constitutes the reinforcement part of the learning process. In the final stage, the model will be able to determine the best way to achieve the outcome. The general outline of how RL algorithm works is shown in Figure 1.



**Figure 1.** RL architecture (Sivamayil et al., 2023).

As shown in Figure 1, the RL process consists of four fundamental components: the agent, environment, action, and reward. The agent serves as the learner and decision-maker within the system. The external world with which the agent interacts is referred to as the environment. At each step, the agent makes choices through an action. The feedback the agent receives after each action, which indicates whether the outcome meets the expected level, is called a reward (Sivamayil et al., 2023). RL is divided into two types: model-based and model-free. In model-based reinforcement learning, the agent attempts to build a model of the environment. This model enables the agent to predict the consequences of its actions before executing them, allowing for a more planned and strategic approach. In model-free reinforcement learning, an explicit model of the environment is not constructed. Instead, the agent directly learns the optimal policy by associating actions with rewards based on the feedback received. RL algorithms are commonly used in dynamic environments due to their ability to adapt quickly to changing conditions and develop new strategies (Shakya et al., 2023).

#### 2.1.1.1. DRL

DRL is a method that combines deep learning and RL algorithms. Thanks to deep learning, agents can learn rules directly from unstructured data. DRL is fundamentally based on Q-Learning, policy gradients, and actor-critic systems (Terven, 2025). Q-learning is a model-free RL algorithm that aims to enable an agent to make optimal decisions by interacting with the environment. The agent does not require a model of the environment; instead, it learns from its experiences by exploring different actions through trial and error. In Q-Learning, the goal is for the agent to build a Q-table that stores Q-values. Each Q-value estimates how correct it is to perform a specific action in a given state in terms of future rewards. The agent updates this table based on the feedback it receives (Jang et al., 2019). Policy gradient is a method that


directly models and optimizes the policy for the agent to obtain optimal models. Since the environment is generally unknown, it aims to address the challenge of the effect a policy update has on the state distribution (Francois-Lavet et al., 2018). The actor-critic algorithm provides a parallel training method. Each agent interacts with its own environment, using multiple actors that learn in parallel. The agents add updates to a shared global model, enabling faster and more stable training (Grondman et al., 2012). Curriculum learning is a method where training data is ordered from simple to complex examples and then used in the training process. By sequencing the data, the model used in Deep RL can learn better and achieve improved generalisation (Soviany et al., 2022).

#### **2.1.1.2. MARL**

While single-agent RL has made significant progress, real-world problems often involve multiple agents. This has created a need for multi-agent learning systems. MARL is an approach in which multiple agents can communicate and simultaneously influence the environment. The MARL method enables a multi-agent structure by allowing agents to learn, cooperate, and compete. A key challenge faced by MARL is the non-stationarity of the environment, caused by agents learning and adapting to one another (Ning & Xie, 2024). In MARL, agents can exhibit cooperative, competitive, or mixed behaviours depending on their interactions and goals. In cooperative MARL, agents work in coordination to complete a task, share a reward, and find the optimal policy for the overall system. Success depends on effective coordination and communication. Competitive MARL is a method in which agents have conflicting goals: they try to maximise their own rewards while minimising those of other agents. Mixed MARL combines cooperation and competition, where agents may cooperate but also face competition (Huh & Mohapatra, 2023). In multi-agent learning, the iterative process in which an agent improves its own performance, thereby leading to environmental variability that affects both itself and other agents, is called self-curriculum. The self-curriculum cycle results in different learning stages, where each stage is built upon the previous one. Self-curriculum is particularly evident in competitive environments where groups of opposing agents continuously change their strategies in response to their opponents' actions (Baker et al., 2019).

#### **2.1.2. Long Short-Term Memory**

Long Short-Term Memory (LSTM) is a type of Recurrent Neural Network (RNN) designed to prevent the decay of feedback loop signals between network outputs based on the input. Thanks to these feedback loops, recurrent networks demonstrate better performance in pattern recognition. LSTM mitigates the vanishing gradient problem, a situation where the activation function becomes ineffective, thereby achieving better performance compared to other RNN architectures (Hochreiter & Schmidhuber, 1997). The LSTM architecture, which consists of three gates input, forget, and output is shown in Figure 2.



**Figure 2.** LSTM architecture (Mohsen et al., 2021).

As shown in Figure 2, the input gate determines which information will be added to the memory cell. The forget gate determines which information will be removed from the memory cells. The output gate identifies which information will be provided as output (Mohsen et al., 2021). Bidirectional LSTM (BiLSTM) is a type of LSTM that enables the bidirectional learning of time series and sequential data. BiLSTM networks pass the input data through LSTM layers twice, which allows for additional training and enhances performance. BiLSTM consists of two components: a forward LSTM and a backward LSTM. The forward LSTM facilitates learning from the first-time step to the last time step. The backward LSTM facilitates learning from the last time step to the first-time step. After the data passes through the two LSTM components, the outputs are merged (Chang et al., 2020).

## 2.2. AI Algorithms Used in Air Combat Simulation Environments

### 2.2.1. WVR

WVR is an air combat scenario that takes place within the range where the pilot can visually see the enemy aircraft. Due to the close distance, short-range weapons are used. The pilot's physical skill, situational awareness, and the aircraft's manoeuvrability are the most decisive factors. AI-based studies conducted for the WVR scenario are explained in this section.

Teng et al. (2012) compared the effectiveness of an Adaptive Computer-Generated Force (CGF), capable of learning and adapting through real-time interactions, is compared against traditional, doctrine-based CGFs used in aviation simulations for pilot training. The proposed method developed an Adaptive CGF model based on RL, capable of learning counter-tactics from human pilots in real-time. The foundation of the model is a self-learning neural network called FALCON, based on Adaptive Resonance Theory (ART). The value functions (Q-values)

of state-action pairs are learned using the FALCON algorithm integrated with the Temporal Difference method (TD-Falcon). An Adaptive  $\epsilon$ -Greedy Policy was used for action selection. The experiments utilized the STRIVE CGF Studio simulation platform. The TD-FALCON model was connected to the simulator using a client-server architecture. In the initial trials, the Adaptive CGF generally achieved losses and draws against novice pilots. Over time, it learned effective strategies and achieved more victories by the 3<sup>rd</sup> trial. However, it could not maintain this advantage as the novice pilots also adapted quickly. Against veteran pilots, the Adaptive CGF could not establish a clear superiority and mostly suffered defeats and draws. Yet, over time, it learned defensive strategies and managed to increase the number of draws. Compared to traditional, non-adaptive CGFs, the Adaptive CGF demonstrated surprising manoeuvres and unexpected behaviours, thereby providing a more realistic and challenging training environment for pilots. However, its learning speed was found to be insufficient, especially against experienced veteran pilots (Teng et al., 2012).

Bae et al. (2023) developed an AI model for WVR air combat manoeuvres. The model operates in a realistic setting with incomplete enemy data, sensor limits, and measurement errors. The Soft Actor-Critic (SAC) (Haarnoja et al., 2018) algorithm was used for the training process due to its effectiveness in continuous action spaces and robustness against noise. The JSBSim aerodynamic model was used as the simulation environment, and OpenAI Gym with the Python programming language was utilized for developing the reinforcement model. SAC-LSTM demonstrated better performance in all experiments compared to the model using only fully connected layers. When the visual sensor range was reduced from 10 km to 1 km, SAC-LSTM models completed the training and achieved a win rate of over 75% against the rule-based model. In contrast, SAC-FC, which does not use LSTM, had a low success rate of 27.3% at the 1 km range. When noise was added to sensor measurements, the SAC-LSTM model achieved a win rate of over 80%, while the SAC-FC model remained at 49.2%. The model trained without the proposed curriculum learning learned more slowly and achieved lower final performance compared to the model trained with the curriculum. The study focused solely on one on-one air combat scenarios and was not extended to multi-agent or swarm scenarios (Bae et al., 2023).

Selmonaj et al. (2024) focused on the explainability of MARL models used in air combat scenarios. The HMARL model developed consists of a high-level commander policy and low-level control policies such as attack, engagement, and defence. The agents were trained using the Proximal Policy Optimisation (Schulman et al., 2017) algorithm. The study utilised a custom 2D simulation platform developed in Python, based on Dassault Rafale aircraft dynamics. Explainability was pursued through methods such as policy simplification, reward decomposition, feature-contribution analysis, hierarchical modelling, and causal modelling. Experiments showed that integrating the hierarchical commander policy improved air combat performance. In explainability experiments, it was observed that as the commander's detection range increased, it selected more cautious (defensive) manoeuvres against numerically superior enemies. At lower detection ranges, situational awareness decreased, leading to persistence in aggressive tactics. In an analysis conducted for a specific aircraft, the bearing angle was

identified as the most critical factor in the commander's decision-making. The commander selected attack mode when the agent was in an advantageous position and defence mode when the enemy agent was directly confronting it. However, the realism of the study is limited by the fact that the simulations were conducted only in a 2D environment (Selmonaj et al., 2024).

Selmonaj et al. (2023) proposed a HMARL-based air combat manoeuvre framework for teams consisting of multiple and heterogeneous agents. The study addresses challenges such as high-dimensional state/action space, partial observability, and nonlinear flight dynamics by dividing the decision-making process into two abstraction layers: commander and unit level. In the proposed HMARL method, lower-level policies control individual aircraft, while a higher-level commander policy issues macro commands for the entire mission. The model was trained using PPO on two basic policies: engagement and disengagement. The policies utilised a sophisticated Actor-Critic neural network incorporating self-attention and Gated Recurrent Unit modules. Parameter sharing exists between agent types. A custom, lightweight 2D Python simulation platform was developed for fast simulation, based on Dassault Rafale dynamics, for use in training and experiments. Despite being trained in 2v2 scenarios through curriculum learning and the self-attention mechanism, the lower-level policy agents were able to operate in 5v5 scenarios as well. The commander policy improved combat performance in small team sizes (3v3). However, in larger team sizes, performance decreased due to partial observability, often resulting in draws or near-equal win/loss ratios. The simulations are limited to a 2D environment, which restricts realism (Selmonaj et al., 2023).

Kong et al. (2023) developed a hierarchical manoeuvre control architecture that was developed for multi-aircraft close air combat. In contrast to traditional one-on one air combat studies, the research focused on variable-sized formations, which are common in modern air combat. The study employed a HMARL structure. Three subtasks were defined: attack, defence, and firing. The Recurrent Soft Actor-Critic (RSAC) algorithm and self-play (SP) method were used to learn the sub-strategies. The open-source JSBSim flight dynamics model was used for aircraft dynamics, with the F-16 Fighting Falcon selected as the aircraft model. The simulation frequency was set to 50 Hz, and the command input frequency was set to 10 Hz. Experimental results achieved an average success rate of 75% in defence and a 50%-win rate in attack. In aggressive firing, the average damage per episode was 0.8. Since the action and state spaces were simplified, complex scenarios are not yet supported (Kong et al., 2023).

Zhu et al. (2024) developed an agent that was developed for close air combat using DRL. To address the problem of neural network plasticity loss in traditional curriculum learning and to develop an effective air combat strategy in a realistic digital twin environment, a Motivational Curriculum Learning based Distributed Proximal Policy Optimization (MCLDPPO) algorithm was proposed. This algorithm observes the agent's performance deficiencies, adds appropriate rewards, and ensures the agent's continuous improvement. A mechanism was designed that allows the agent to interrupt its current action and make a new decision when its threat perception changes (such as an enemy missile warning or entering/exiting the enemy's visual field). A Deep LSTM-based Actor Critic network was chosen to process local and global

observations. Using Unity3D, a high-fidelity digital twin air combat simulation environment was created, incorporating aerodynamics, radar, infrared, and missile systems. The simulation was modelled with the F-22 fighter jet and AIM-120 missiles. In tests against an expert rule-based Finite State Machine (FSM), the MCLDPPO agent achieved a 66%-win rate. The addition of the interruption mechanism improved the agent's performance by approximately 10%. The study focused solely on scenarios involving a single friendly aircraft against a single enemy aircraft (Zhu et al., 2024).

Saldiran et al. (2024) achieved global explainability was achieved for RL-trained AI agents in close air combat. In the safety-critical domain of air combat, understanding AI decision processes is of great importance alongside performance improvement. A Deep Q-Network-based RL approach, combined with a reward decomposition method that extends local explainability to a global level, was employed. The agent's decisions were explained by decomposing them into different reward types. The state space was divided into tactical regions such as attack, defence, neutral, and head-on, allowing for regional interpretation of the agent's behaviour. The relative contributions of different reward types were visualized to enhance comprehensibility. A 3DOF model was adopted for aircraft dynamics, and Python with the CleanRL library was used for training. The reward decomposition method was found effective in understanding the agent's tactical preferences. However, the method was not directly applied to policy-based RL methods, and the visualization approach was limited to only three reward types (Saldiran et al., 2024).

#### **2.2.1.1. Dogfight**

Israelsen et al. (2018) addressed the optimisation of behavioural parameters for AI-based sparring partners in aviation combat training. The primary focus is on optimising and adapting AI pilots, particularly in air-to-air dogfight scenarios, within a variable and noisy environment. The study employed the Gaussian Process Surrogate-based Bayesian Optimisation (GPBO) method. To address GPBO's limitations in handling highly variable and noisy objective functions, a new sampling technique called Hybrid Repeat/Multi-point Sampling (HRMS) was developed. HRMS combines Repeat Sampling (RS) and Multi-point Sampling (MS) strategies to both enhance the reliability of optimisation and enable the construction of a global model of the objective function. Acquisition functions used included Expected Improvement (EI), Upper Confidence Bound (UCB), and Thompson Sampling (TS). The GPBO method supported by HRMS produced more consistent and repeatable optimisation results compared to traditional Single-point Sampling (SS) methods. However, in high-dimensional spaces ( $d \geq 6$ ), the performance of methods like Thompson Sampling decreases (Israelsen et al., 2018).

Pope et al. (2023) examined AlphaDogfight, developed as part of DARPA's Air Combat Evolution (ACE) program, is examined. A competitive experimental framework was established to measure the capability of AI agents to conduct air combat against human pilots in a simulation environment. Within the scope of the AlphaDogfight Trials, autonomous agents based on DRL were developed. The study proposes a novel Hierarchical Deep Reinforcement Learning (HDRL) approach called PHANG-MAN (Policy Hierarchy for Adaptive Novel

Generation of MANeuvers). In the proposed method, agents are trained separately at lower and upper levels. In lower-level policies, agents are trained separately for different combat situations (e.g., attack, defence, control), such as Control Zone, Aggressive Shooter, and Conservative Shooter. The high-level policy selector chooses the most suitable sub-policy based on the current context. The JSBSim open-source software was used for basic flight dynamics. The environment was extended based on OpenAI Gym, and multi-aircraft scenarios were supported. A realistic damage model was implemented by defining the Weapon Engagement Zone (WEZ). The adversarial self-play method, which involves agents competing against themselves and their past versions, was employed. In a match between the trained agent and an F-16 pilot, the AI agent defeated the human pilot with a score of 5-0. The most successful agents demonstrated high-precision weapon usage, rapid decision-making, and low response times. While the training only covered 1vs1 (single combat) scenarios, multi-agent coordination was not tested (Pope et al., 2023).

### **2.2.2. BVR**

BVR is an air combat scenario where the target is not visually seen but is detected and engaged solely through radar, data links, or sensors. The decision-making process is based on sensors, radar warning systems, electronic warfare, and tactical data sharing. The missiles used are radar-guided medium/long-range air-to-air missiles. The pilot typically engages the target without ever seeing it. AI-based studies conducted for the BVR scenario are explained in this section.

Toubman et al. (2015) developed RL-based behaviour generation was developed for CGF used in air combat training simulations. Firstly, to address the problem of the chance factor in missile hits leading to incorrect rewards during the learning process, a new reward function named Probability-of-kill Rewards is proposed. The Probability-of-kill Rewards function rewards agents based on the expected outcome (the missile's probability of kill) rather than the actual outcome (whether the missile hits or not). The Dynamic Scripting RL agent method was used for training. The study utilized a customized air combat simulation environment modelling F-16 fighter jets. The scenario is based on a 2v1 air combat engagement. The proposed Pk reward function demonstrated statistically significant higher performance compared to traditional binary (win/loss) reward functions and expert knowledge-based reward functions. The proposed method consistently showed performance improvement of approximately 10% to 19.4% across all scenarios. The highest performance increase, 19.4%, was achieved against an enemy using close-range tactics when compared to the expert knowledge-based reward function. However, the performance is limited by the existing rules in the rule base, and the absolute optimal behaviour cannot be discovered (Toubman et al., 2015).

Piao et al. (2020) successfully taught BVR air combat tactics were successfully taught to AI agents using end-to-end RL and self-play, without relying on prior human knowledge. The Proximal Policy Optimization algorithm was used in training to improve performance. To address the cold start problem and enable rapid acquisition of basic air combat skills, a Key Air Combat Event Reward Shaping (KAERS) mechanism was developed. Instead of relying

solely on sparse end-of episode win/loss signals, KAERS provides rewards based on objective and sporadically occurring key air combat events such as radar lock, missile launch, and evasion. The simulation environment utilized an RL-focused air combat simulation framework developed by Northwestern Polytechnical University (NPU). By the end of the training, agents trained with KAERS demonstrated higher win rates compared to agents trained with traditional dense reward functions. However, the scope of the study remains limited to 1v1 scenarios only (Piao et al., 2020).

Sun et al. (2021) developed a MARL method was developed for AI-assisted air combat. The proposed Multi-Agent Hierarchical Policy Gradient (MAHPG) algorithm utilizes a PPO based training process. MAHPG features a hierarchical neural network architecture where the high-level makes discrete manoeuvre and firing decisions, while the low-level determines continuous parameters (normal load factor, velocity). RL agents enable learning against each other through self-play. The Key Event based Reward Shaping (KAERS) mechanism helps mitigate the sparse reward problem. The air combat simulation environment named WUKONG, developed by Northwestern Polytechnical University (NPU), was used. WUKONG can simulate both BVR and WVR scenarios. The MAHPG algorithm achieved a higher win rate compared to contemporary algorithms such as P-DQN, MADDPG, PPO, and A2C. MAHPG increased its win rate from 0% to 68% against an expert-based bot as training progressed. However, the framework is only built upon fully observable states (MDP) and has not been adapted for partially observable (POMDP) scenarios (Sun et al., 2021).

Dantas et al. (2025) trained autonomous fighter aircraft agents were trained using DRL for BVR air combat, specifically in two-versus-two (2v2) scenarios. A customized training environment named AsaGym was developed for BVR air combat simulations and DRL-based agent training. AsaGym includes an F-16 aircraft model, radar, missiles, and behaviour trees. For training, PPO, Soft Actor-Critic (SAC), Twin Delayed Deep Deterministic Policy Gradient (Fujimoto et al., 2018) (TD3), and Advantage Actor Critic (A2C) algorithms were used and compared. In experiments conducted with only the leader using DRL and setups with both leader and wingman using DRL, PPO demonstrated the best performance in terms of average episode reward and episode length metrics. SAC learned stably but with lower performance than PPO, while TD3 and A2C exhibited inconsistent or limited performance. PPO achieved the highest average reward (0.75-0.80) and the shortest engagement times (125-175 steps), showing the best performance. However, PPO also had the longest training time (6.3-6.8 hours). SAC learned with lower rewards (0.5) than PPO but in a stable manner and had a shorter training time (5.0-5.7 hours). TD3 and A2C were the fastest algorithms to train (3.1-4.7 hours), although they achieved significantly lower rewards and had longer/less effective engagement times. The simulations were conducted at a fixed altitude (25,000 feet), and the complexity of three-dimensional manoeuvres was not addressed (Dantas et al., 2025).

Piao et al. (2024) proposed a novel algorithm named deep Excitation Inhibition Factored Manoeuvre (ENACTIVE), which integrates DRL with prior knowledge, to enhance air combat decision-making processes. Inspired by the Excitatory-Inhibitory (E/I) balance in

neuroscience, the algorithm enables the agent to adapt its decision-making timing based on situational awareness, learning when to act. The RL agents were trained on Offensive/Defensive Semantic Manoeuvre (ODSM) and Energy Semantic Manoeuvre (ESM). The experiments in the study were conducted using the WUKONG air combat simulation environment. The proposed ENACTIVE algorithm achieved a 62%-win rate against state-of-the-art algorithms such as MAHPG, PPO, A2C, and TempoRL. The results demonstrated that the ENACTIVE algorithm enables the agent to maintain E/I balance by making timely decisions at critical moments, rather than exhibiting overly reactive or overly passive behaviour. However, the proposed method has so far only been tested in one-on-one and two-on-two scenarios. Its scalability to more crowded scenarios such as 3v3 and 4v4 has not yet been verified (Piao et al., 2024).

### **2.2.3. BVR & WVR**

Toubman et al. (2015) used transfer learning was used to transfer air combat behaviours from a simpler scenario (2v1) to a more complex one (2v2). The Reinforcement Learning-based Dynamic Scripting (DS) technique was used for the training process. Rule weights learned in the source task (2v1) were directly copied to the rule bases of the agents in the target task (2v2). To measure the success of the transfer, metrics such as Jumpstart, Asymptotic Performance, Total Trials Won, Transfer Ratio, and Time to Threshold were used. A specially developed air combat simulation environment was used for training, which modelled F-16-based fighter aircraft. Behaviours were controlled using scripts consisting of if-then rules. In the default, evading, and close-range tactics, the agents that underwent transfer showed higher performance across all five metrics compared to those without transfer. However, in the new and challenging lead-trail tactic scenario, although the transferred agents showed better initial and final performance, their learning process progressed more slowly. Furthermore, they demonstrated lower performance in terms of the total number of trials won and the time to reach the threshold metrics compared to the non-transferred agents (Toubman et al., 2015).

Selmonaj et al. (2025) developed a system that enables real-time interaction between pilots and AI-supported combat pilots. The system aims to facilitate human-AI teamwork, military training scenarios, and the discovery of innovative tactics. The study utilised MARL for training the agents. The model incorporated Multi-Agent Proximal Policy Optimization (MA-PPO), an Actor-Critic network architecture, the Centralized Training with Decentralized Execution (CTDE) paradigm, curriculum learning, and self-play methods to enhance performance. The agents were trained on attack, engage, and defend policies. Additionally, a commander policy was implemented to decide which control policy to activate. A custom 3D flight dynamics simulation environment based on JSBSim was used for the training process. The agents achieved an 80%-win rate in 10-versus-10 air combat scenarios. However, training the AI agents directly within VR-Forces has not yet been fully realised, and a separate training environment was used instead.

### 3. Discussion

WVR, BVR, and combined WVR&BVR studies have been compared according to defined criteria. Table 1 presents studies focused on WVR air combat scenarios.

**Table 1.** Studies on air combat scenarios within visual range.

| Study                  | Primary Focus                                                                        | Method                                                                                                 | Simulation Environment                              | Performance Results                                                                                                                                                        |
|------------------------|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Teng et al., 2012      | The development of Adaptive CGF for pilot training.                                  | RL was used together with TD FALCON (ART-based neural network).                                        | CAE STRIVE CGF Studio                               | Effective strategies were developed against novice pilots over time.                                                                                                       |
| Bae et al., 2023       | Manoeuvre generation was enabled in a realistic environment with sensor limitations. | SAC-LSTM (Soft Actor Critic with LSTM networks) and Curriculum Learning algorithms were used together. | JSBSim and OpenAI Gym                               | It demonstrated significantly better performance than non-LSTM models, with over 75%-win rate at 1 km sensor range and 80%-win rate under sensor noise.                    |
| Selmona j et al., 2024 | Explainability of MARL decisions in air combat tactics was achieved.                 | PPO and HMARL algorithms were used together.                                                           | Custom 2D Python simulation environment             | The hierarchical commander policy improved performance. Agent behaviour was explained based on detection range and aspect angle. Agents trained in 2v2 were scaled to 5v5. |
| Selmona j et al., 2023 | HMARL has been used to manoeuvre heterogeneous agents.                               | Hierarchical MARL, PPO, personal attention and GRU networks were used together.                        | Custom 2D Python Simulator to model Dassault Rafale | While the commander policy improved performance in small teams, it was challenged by partial observability in larger teams.                                                |

**Table 1.** (Continued)

|                        |                                                                                              |                                                                                                             |                                                                 |                                                                                                                                                |
|------------------------|----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Kong et al., 2023      | Multi plane within range combat studied in variable conditions.                              | Hierarchical MARL, Repetitive Soft Actor-Critic (RSAC) and Self-Play methods were used.                     | JSBsim with F-16 Model                                          | 75% success rate in defence and a 50%-win rate in offense was achieved.                                                                        |
| Zhu et al., 2024       | Deep RL agents are trained in a high-quality digital twin environment.                       | The MCLDPPO (Motivational Curriculum Learning DPPO) method suitable for action cuts is presented.           | High Quality F-22 and AIM 120 digital twins created in Unity3D. | A 66%-win rate was achieved against a rule based expert agent. Performance increased by 10% with interrupt rates.                              |
| Saldiran et al., 2024  | An attempt has been made to ensure the explainability of AI agent tactics.                   | Deep Q Network with reward decomposition.                                                                   | 3DOF Aircraft Model                                             | The reward decomposition method has been effective in understanding tactical choices.                                                          |
| Israelsen et al., 2018 | Optimized AI sparring partner parameters for dogfight training.                              | Bayesian Optimization with Gaussian Process Surrogate (GPBO) and Hybrid Repeat/Multi point Sampling (HRMS). | Not specified                                                   | HRMS has shown more consistent and reproducible results compared to single-point methods.                                                      |
| Israelsen et al., 2018 | AI decision makers are adaptively trained in dogfight simulations.                           | Gaussian Process Bayesian Optimization (GPBO) with Hybrid Repeat/Multi point Sampling (HRMS).               | Mission Simulation System                                       | HRMS showed 0.6 times lower variances and required 30-40% less simulation training than classical GPBO.                                        |
| Pope et al., 2023      | AI agents have been developed for simulated dog fighting aerial combat against human pilots. | Hierarchical Deep RL (HRL) and Adversarial Self-Play.                                                       | JSBSim                                                          | The AI agent defeated an experienced F-16 human pilot 5-0. Demonstrated high accuracy of weapon selection and ability to make quick decisions. |

Table 1 shows a clear shift from basic RL (Teng et al., 2012) towards more complex methods such as HMARL (Selmonaj et al., 2024) and advanced actor-critic algorithms combined with LSTM to handle partial observability (Bae et al., 2023). Recent studies address real-world challenges such as sensor noise, limited range, and partial information (Bae et al., 2023). Although dogfight studies (Israelsen et al., 2018) and some WVR studies focus on 1v1 combat, there is a significant research trend towards multi-agent scenarios (Selmonaj et al., 2023). New research areas are emerging that focus not only on performance but also on the explainability and trustworthiness of AI decisions (Saldiran et al., 2024; Selmonaj et al., 2024). The environments used have evolved from proprietary platforms (STRIVE, MSS) to open-source platforms (JSBSim, OpenAI Gym) and high-fidelity digital twins (Unity3D). This evolution has increased both accessibility and realism. From a performance perspective, RL-based methods have demonstrated high adaptability, while PPO and MARL-based models have provided stable learning curves. However, training durations remain long, and computational costs are high. Although computer vision-based models have improved perceptual accuracy, they have shown sensitivity to sensor noise. Overall, WVR research has made significant progress in tactical flexibility and autonomous decision-making; however, there is a future need for hybrid systems that are effectively integrated with real flight data. Studies focusing on BVR air combat scenarios are presented in Table 2.

**Table 2.** Studies on BVR air combat scenario.

| Study                | Primary Focus                                                                    | Method                                                                                                              | Simulation Environment                                              | Performance Results                                                                   |
|----------------------|----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Toubman et al., 2015 | An adaptive reinforcement algorithm was used for Computer Generated Forces.      | A reinforcement learning-based dynamic script generation technique with a Probability of Kill (Pk) reward was used. | A custom air combat simulation environment built with an F-16 model | The Pk reward improved performance by between 10% and 19.4% compared to the baseline. |
| Piao et al., 2020    | BVR tactics were taught through end-to-end RL without any prior human knowledge. | Key Air Combat Event Reward Shaping (KAERS) mechanism and PPO.                                                      | The NPU air combat simulation environment                           | KAERS trained agents achieved a higher win rate than dense reward agents.             |
| Sun et al., 2021     | RL was used to develop multi agent forces for large-scale air combat tactics.    | The Multi-Agent Hierarchical Policy Gradient (MAHPG), PPO, and KAERS algorithms were used together.                 | WUKONG simulation framework                                         | 68% win rate was achieved against the expert bot after self-play training.            |

**Table 2.** (Continued)

|                     |                                                                                                                                                                                                                       |                                                                                   |                                          |                                                                                                             |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| Dantas et al., 2025 | An autonomous Deep RL agent with collective situational awareness was trained for 2v2 BVR air combat. The ENACTIVE algorithm, which combines inhibition balance inspired by neuroscience with Deep RL, was developed. | PPO, SAC, TD3, and A2C methods were trained and their performances were compared. | AsaGym simulation environment            | PPO achieved the highest win rate, ranging between 75% and 80%, and the fastest convergence times. ENACTIVE |
| Piao et al., 2024   |                                                                                                                                                                                                                       | ENACTIVE (Deep Inhibitory Factorized RL) algorithm was used.                      | WUKONG air combat simulation environment | achieved a better win rate, at 62%, compared to the PPO, A2C, and TempoRL algorithms.                       |

As can be seen in Table 2, recent studies on BVR air combat AI demonstrate diversity ranging from rule-based behaviour generation (Toubman et al., 2015) to DRL with multi-agent coordination and neuroscience-inspired decision models (Piao et al., 2024). While early studies largely focused on hand-crafted rewards or scenario constraints, subsequent studies have been adapted to enable the exploration of tactical manoeuvres. Studies using PPO are observed to achieve better results. Comparative studies (Dantas et al., 2025) demonstrate that PPO performs better than SAC, TD3, and A2C in both convergence speed and tactical consistency. The integration of self-play strategies into the training environment (Sun et al., 2021) enhanced generalisation capabilities in dynamic environments. While better results were obtained in single-agent (1v1) scenarios, scalability issues emerged in multi-agent (2v2, 3v3, etc.) environments. Dantas et al. (2025) directly address the coordination problem in 2v2 combat. Electronic warfare, radar jamming, and communication constraints have been included to a limited extent in BVR studies. Most simulation environments simplify some aspects of three-dimensional engagement dynamics. Table 3 presents studies that encompass both BVR and WVR air combat scenarios.

**Table 3.** Studies on Beyond Visual Range & Within Visual Range air combat scenarios.

| Study                | Primary Focus                                                                                                                  | Method                                                | Simulation Environment                                        | Performance Results                                                                   |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|---------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Toubman et al., 2015 | It was ensured that the learned air combat behaviours were transferred from simple (2v1) scenarios to complex (2v2) scenarios. | Transfer Learning applied to Dynamic Scripting agents | Custom developed F-16 model air combat simulation environment | Transfer-trained agents outperformed in all but the most complex lead trail scenario. |

**Table 3.** (Continued)

|                       |                                                                                                                    |                                                                      |                                                     |                                                                                                                 |
|-----------------------|--------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Sun et al., 2021      | A HMARL framework was developed for tactics that span both within visual-range and beyond visual-range air combat. | MAHPG, PPO, self-play, and KAERS were used together.                 | WUKONG simulation framework                         | The framework achieved a 68%-win rate against an expert bot and showed adaptability across both combat regimes. |
| Selmonaj et al., 2025 | A system that enables real-time human-AI air combat interaction and teamwork has been developed.                   | A human-AI environment has been prepared through the MARL framework. | Custom 3D flight dynamics simulator based on JSBSim | Agents achieved 80%-win rate in 10v10 air combat scenarios.                                                     |

As seen in Table 3, there is a clear trend towards multi-agent reinforcement learning, transferability, and human-AI collaboration in air combat modelling. The studies demonstrate a progression from single-task tactical learning (Toubman et al., 2015), to hierarchical multi-agent tactical coordination across mixed combat regimes (Sun et al., 2021), and ultimately to human-AI integrated combat systems (Selmonaj et al., 2025). This pattern represents the current frontier in the use of AI technology in air combat.

#### 4. Conclusion

The study presented a comprehensive comparative analysis of AI-based learning systems used in air combat simulation environments. Research in WVR, BVR, and hybrid scenarios were examined. It was demonstrated that there is a transition from basic RL and rule-based systems to advanced DRL and MARL architectures. Hierarchical structures, curriculum learning, self-play, and LSTM memory mechanisms have greatly improved agents' decision-making, adaptability, and realism in dynamic air combat. DRL and MARL demonstrated superior performance in generating complex air combat manoeuvres and ensuring coordination among multiple agents. The importance of explainability has emerged for the operational use of the developed systems, to ensure their decisions are understandable and reliable. In the studies conducted, high computational costs and long training durations are required, especially in complex multi-agent scenarios. The intelligent agents developed for air combat simulation environments have the potential to revolutionize pilot training as adaptive, challenging, and realistic AI opponents and teammates.

The AI-based learning systems examined in this study have been comparatively analysed for their strengths and limitations across different scenarios. In WVR scenarios, LSTM-based algorithms show high robustness against sensor noise and achieve successful results even in limited visual ranges. In contrast, MARL approaches excel particularly in multi-aircraft coordination but suffer from performance degradation in large-scale team scenarios due to

partial observability. In BVR studies, the design of reward functions plays a critical role, directly affecting the realism of missile launch timing and tactical success. While some research in the literature focuses only on WVR or only on BVR scenarios, fewer studies address these two environments together, investigating the algorithms' adaptability to different tactical conditions. Studies that jointly consider both environments demonstrate the scenario-independent generalizability and multidimensional decision-making capacity of AI-based learning systems. Analytically, scenarios combining BVR and WVR offer greater operational realism compared to single-focus studies. However, this approach requires more complex reward functions and more powerful simulation infrastructure. While SAC-LSTM models, which are resistant to sensor noise, are preferred for generating realistic scenarios in pilot training, PPO-based HMARL models have been found to be more efficient in tasks requiring small-team coordination.

Among the reviewed studies, SAC-LSTM models, which exhibit high resilience to sensor noise, stand out for their use in realistic scenarios for pilot training. These models demonstrate stable performance even under sensor errors and limited range conditions, providing reliable results in individual manoeuvre decisions. In contrast, PPO-based HMARL models have shown more efficient performance in tasks requiring team coordination. Thanks to their commander-unit level decision-making structure, these models enhance coordination in multi-aircraft scenarios and strengthen tactical harmony.

Furthermore, it paves the way for the development of advanced autonomous systems for mission planning and execution by enhancing overall combat effectiveness. There is a need to develop more sample-efficient and computationally cost-effective training algorithms in the future. The development of standardized, high-reliability, open-source simulation benchmarks will significantly accelerate progress and enable fair comparisons. Finally, to enhance the reliability of agents, the explainability of AI agents' decisions must be ensured.

## **Conflict of Interest**

The authors declare that they have no conflict of interest related to this study.

## **Funding**

This research received no funding.

## **Ethical Statements**

Not applicable.

## **Data and Code Availability**

Data sharing is not applicable to this article as no new datasets were generated or analysed. No custom code was developed or used in this study.

## Supplementary Materials

Not applicable.

## References

- Bae, J. H., Jung, H., Kim, S., Kim, S., & Kim, Y.-D. (2023). Deep reinforcement learning-based air-to-air combat maneuver generation in a realistic environment. *IEEE Access*, 11, 29249–29265. <https://doi.org/10.1109/ACCESS.2023.3257849>
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., & Mordatch, I. (2019). Emergent tool use from multi-agent autocurricula. *arXiv*. <https://doi.org/10.48550/arXiv.1909.07528>
- Chang, K.-C., Chu, K.-C., Wang, H.-C., Lin, Y.-C., & Pan, J.-S. (2020). Agent-based middleware framework using distributed CPS for improving resource utilization in smart city. *Future Generation Computer Systems*, 108, 445–453. <https://doi.org/10.1016/j.future.2020.03.006>
- Dantas, J. P. A., Medeiros, F. L. L., Samersla, A. R., Botelho, P. L. R., Gomes, V. C. F., Silva, S. R., Ferreira, Y. D., Arantes, A. O., Aquino, M. R. C., & Maximo, M. R. O. A. (2025). Deep reinforcement learning agents with collective situational awareness for beyond visual range air combat. *IEEE Access*, 13, 143052–143069. <https://doi.org/10.1109/ACCESS.2025.3597199>
- Francois-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G., & Pineau, J. (2018). An introduction to deep reinforcement learning. *arXiv*. <https://doi.org/10.48550/arXiv.1811.12560>
- Fujimoto, S., van Hoof, H., & Meger, D. (2018). Addressing function approximation error in actor-critic methods. *arXiv*. <https://doi.org/10.48550/arXiv.1802.09477>
- Grondman, I., Busoniu, L., Lopes, G. A. D., & Babuska, R. (2012). A survey of actor critic reinforcement learning: Standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6), 1291–1307. <https://doi.org/10.1109/TSMCC.2012.2218595>
- Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv*. <https://doi.org/10.48550/arXiv.1801.01290>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>

- Huh, D., & Mohapatra, P. (2023). Multi-agent reinforcement learning: A comprehensive survey. *arXiv*. <https://doi.org/10.48550/arXiv.2312.10256>
- Israelsen, B., Ahmed, N., Center, K., Green, R., & Bennett, W. Jr. (2018). Adaptive simulation-based training of artificial-intelligence decision makers using Bayesian optimization. *Journal of Aerospace Information Systems*, 15(2), 38–56. <https://doi.org/10.2514/1.I010553>
- Jang, B., Kim, M., Harerimana, G., & Kim, J. W. (2019). Q-learning algorithms: A comprehensive classification and applications. *IEEE Access*, 7, 133653-133667. <https://doi.org/10.1109/ACCESS.2019.2941229>
- Kong, W., Zhou, D., Du, Y., Zhou, Y., & Zhao, Y. (2023). Hierarchical multi-agent reinforcement learning for multi-aircraft close-range air combat. *IET Control Theory & Applications*, 17(13), 1840–1862. <https://doi.org/10.1049/cth2.12413>
- Mei, J., Li, G., & Huang, H. (2024). Deep reinforcement-learning-based air-combat maneuver generation framework. *Mathematics*, 12(19), 3020. <https://doi.org/10.3390/math12193020>
- Mohsen, S., Elkaseer, A., & Scholz, S. G. (2021). Industry 4.0-oriented deep learning models for human activity recognition. *IEEE Access*, 9, 150508–150521. <https://doi.org/10.1109/ACCESS.2021.3125733>
- Neville, K. J., Thom McLean III, A. L., Sherwood, S., Kaste, K., Walwanis, M., & Bolton, A. (2015). Live-Virtual-Constructive (LVC) Training in Air Combat: Emergent Training Opportunities and Fidelity Ripple Effects. In *International Conference on Augmented Cognition* (pp. 82-90). Springer. [https://doi.org/10.1007/978-3-319-20816-9\\_9](https://doi.org/10.1007/978-3-319-20816-9_9)
- Ning, Z., & Xie, L. (2024). A survey on multi-agent reinforcement learning and its application. *Journal of Automation and Intelligence*, 3(2), 73–91. <https://doi.org/10.1016/j.jai.2024.02.003>
- Piao, H., Sun, Z., Meng, G., Chen, H., Qu, B., Lang, K., Sun, Y., Yang, S., & Peng, X. (2020). Beyond-visual-range air combat tactics auto-generation by reinforcement learning. In *2020 International Joint Conference on Neural Networks (IJCNN)* (pp. 1–8). IEEE. <https://doi.org/10.1109/IJCNN48605.2020.9207088>
- Piao, H. Y., Yang, S., Chen, H., Li, J., Yu, J., Peng, X., Yang, X., Yang, Z., Sun, Z., & Chang, Y. (2024). Discovering expert-level air combat knowledge via deep excitatory-inhibitory factorized reinforcement learning. *ACM Transactions on Intelligent Systems and Technology*, 15(4), Article 65, 1–28. <https://doi.org/10.1145/3653979>
- Pope, A. P., Ide, J. S., Mićović, D., Diaz, H., Twedt, J. C., Alcedo, K., Walker, T. T., Rosenbluth, D., Ritholtz, L., & Javorsek, D. II. (2023). Hierarchical reinforcement

- learning for air combat at DARPA's AlphaDogfight Trials. *IEEE Transactions on Artificial Intelligence*, 4(6), 1371–1385. <https://doi.org/10.1109/TAI.2022.3222143>
- Saldiran, E., Hasanzade, M., Inalhan, G., & Tsourdos, A. (2024). Towards global explainability of artificial intelligence agent tactics in close air combat. *Aerospace*, 11(6), 415. <https://doi.org/10.3390/aerospace11060415>
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv*. <https://doi.org/10.48550/arXiv.1707.06347>
- Selmonaj, A., Antonucci, A., Schneider, A., Rügsegger, M., & Sommer, M. (2025). Explaining strategic decisions in multi-agent reinforcement learning for aerial combat tactics. *arXiv*. <https://doi.org/10.48550/arXiv.2505.11311>
- Selmonaj, A., Del Rio, G., Schneider, A., & Antonucci, A. (2025). Towards human engagement with realistic AI combat pilots. In *Proceedings of the 13th International Conference on Human-Agent Interaction (HAI '25) (pp. 1–3)*. ACM. <https://doi.org/10.1145/3765766.3765827>
- Selmonaj, A., Szeher, O., Del Rio, G., Antonucci, A., Schneider, A., & Rügsegger, M. (2023). Hierarchical multi-agent reinforcement learning for air combat maneuvering. *arXiv*. <https://doi.org/10.48550/arXiv.2309.11247>
- Shakya, A. K., Pillai, G., & Chakrabarty, S. (2023). Reinforcement learning algorithms: A brief survey. *Expert Systems with Applications*, 231, 120495. <https://doi.org/10.1016/j.eswa.2023.120495>
- Soviany, P., Ionescu, R. T., Rota, P., & Sebe, N. (2022). Curriculum learning: A survey. *International Journal of Computer Vision*, 130(6), 1526–1565. <https://doi.org/10.1007/s11263-022-01611-x>
- Sivamayil, K., Rajasekar, E., Aljafari, B., Nikolovski, S., Vairavasundaram, S., & Vairavasundaram, I. (2023). A systematic study on reinforcement learning based applications. *Energies*, 16(3), 1512. <https://doi.org/10.3390/en16031512>
- Sun, Z., Piao, H., Yang, Z., Zhao, Y., Zhan, G., Zhou, D., Meng, G., Chen, H., Chen, X., Qu, B., & Lu, Y. (2021). Multi-agent hierarchical policy gradient for Air Combat Tactics emergence via self-play. *Engineering Applications of Artificial Intelligence*, 98, 104112. <https://doi.org/10.1016/j.engappai.2020.104112>
- Teng, T.-H., Tan, A.-H., Ong, W.-S., & Lee, K.-L. (2012). Adaptive CGF for pilots training in air combat simulation. In *Proceedings of the 2012 15th International Conference on Information Fusion (FUSION) (pp. 2263–2270)*. IEEE.
- Terven, J. (2025). Deep reinforcement learning: A chronological overview and methods. *AI*, 6(3), 46. <https://doi.org/10.3390/ai6030046>

- Toubman, A., Roessingh, J. J. M., Spronck, P., Plaat, A., & van den Herik, H. J. (2015). Rewarding air combat behavior in training simulations. *In 2015 IEEE International Conference on Systems, Man, and Cybernetics* (pp. 1397–1402). IEEE. <https://doi.org/10.1109/SMC.2015.248>
- Toubman, A., Roessingh, J. J. M., Spronck, P., Plaat, A., & van den Herik, H. J. (2015). Transfer learning of air combat behavior. *In 2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)* (pp. 226–231). IEEE. <https://doi.org/10.1109/ICMLA.2015.61>
- Zhu, J., Kuang, M., Zhou, W., Shi, H., & Han, X. (2024). Mastering air combat game with deep reinforcement learning. *Defence Technology*, 34, 295–312. <https://doi.org/10.1016/j.dt.2023.08.019>

## ML-Based RF Calibration of Military T/R Modules

Alperen M. Ünal<sup>1</sup> , Yusuf Yiğitbaşı<sup>1</sup> , Akil Kutlu<sup>1</sup> , Fatma Hatipoğlu<sup>1,\*</sup> 

### Abstract

In this work, Machine Learning (ML) techniques are incorporated into the Radio Frequency (RF) calibration workflow with the objective of reducing associated costs and enhancing operational efficiency for the military grade Transmitter/Receiver (T/R) modules. RF calibration entails a series of measurements that verify the module's performance at certain temperatures, power levels and frequencies; when performed with conventional procedures, it is both time-consuming and expensive. In this study for a single unit, the traditional RF calibration process requires approximately 17 h and involves ten dedicated test and calibration stations as well as seven operators. Ensuring RF calibration, seven distinct Received-Signal-Strength-Indicator (RSSI) values are recorded; in this study focuses on one of these parameters, denoted RSSI F, which is essential for the correct functioning of the target module. In validation, the ML predicted RSSI F values are directly benchmarked against measured references under a  $\pm 12$  mV error envelope; the comparison yielded a 66.80% success rate.

**Keywords:** Machine learning, artificial intelligence, RF calibration, regression, XGBoost, LightGBM, and TabNet.

Received: 12.03.2026

Accepted: 21.05.2026

Research Article

<sup>1</sup> ASELSAN Inc., Ankara, Türkiye

\* Corresponding author.

✉ fatmakart1407@gmail.com

**Cite this article as:**

<https://doi.org/10.65834/jdsi.12.52>

Ünal, A.M., Yiğitbaşı, Y., Kutlu, A. & Hatipoğlu, F. (2026). ML-Based RF Calibration of Military T/R Modules. *Journal of Defence and Security Industries: Strategy and Technology 1(2)*, 110-128.

## 1. Introduction

Communication is a vital nerve network for the army's command-control, intelligence sharing and unit coordination; without timely and secure information flow, decisions are delayed, synchronization between units deteriorates and operational effectiveness is significantly reduced. Communication between every unit has an impact on the strength of operability in the army. Considering that many units sending and receiving voice and digital data, linearity is critical. The army is obliged to fulfil its duties in different places, under different climatic conditions for long periods of service time. Therefore, military technology Radio Frequency (RF) devices must be capable of delivering high performance across the operating temperature range from  $-40^{\circ}\text{C}$  to  $+55^{\circ}\text{C}$ . The system shall operate without performance degradation across specified temperature extremes and therefore must successfully satisfy the high-temperature and low-temperature qualification tests defined in MIL-STD-810H, 501.7 and 502.7 (United States Department of Defence [DOD], 2019). In view of RF Transceiver device as a system, expecting multiple subsystem, such as power supply, Power Amplifier (PA) and filters have to perform proportionally.

The basis of the transferring an information in the environment, carrying the frequency and amplifying of signals have high level of linearity. Amplified the signals according to varied required signal levels, bandwidths, waveform have to meet the performance metrics. These metrics are 1 dB compression point  $P_{1\text{dB}}$ , power gain, Power-Added Efficiency (PAE), Adjacent Channel Power Ratio (ACPR), Third-Order Intercept Point ( $IP_3$ ), Peak-to-Average Power Ratio (PAPR). Managing broad bandwidths and performance metrics while preserving the signal integrity and efficiency is a significant issue.

Modern communications systems cover High Frequency (HF), Very High Frequency (VHF) and Ultra High Frequency (UHF) communication transmitters. The communication units generate signals in baseband and amplify these signals to the medium by carrying the frequencies. According to the technical standards and specs, PAs amplify the generated signals with varying amplitude and phase components. This amplifying process must preserve signal integrity, minimize interference, be accurate and linear. Therefore, the PAs are of critical importance by influencing the power efficiency of radio transmitters and receivers. (Ortega-González, Pño-Gómez, & Madueño-Pulido, 2025).

RF calibration is the process of setting specific parameters into digital units so that a device or module operates stably and efficiently. The RF calibration process can be performed in three different approaches. The first approach relies on iterative testing and exhaustive search techniques, which are characterized by high cost and significant time consumption. To overcome these limitations, recent research emphasizes a shift toward Machine Learning (ML) and fully integrated self-calibration strategies. Unlike manual calibration that necessitates extensive data collection, ML-based techniques can attain near-optimal performance settings using only a fraction of the measurements required by exhaustive searches. Bayesian optimization with both global and local search algorithms is used to predict control voltages of circuits (Chang, Wen, Hong, Padmanabhan, Franzon, & Floyd, 2025).

In the work of Dikmese et al., modern ML methods are proposed to predict the dynamic nonlinear behaviour of wideband RF PAs. Neural networks, k-nearest neighbours and several tree-based ML algorithms are adapted to handle complex-valued signals and applied to PA modelling. Evaluation with measured data from two commercial base-station PAs shows that gradient boosting performs best, yielding comparable or better performance than memory-polynomial reference models in terms of Normalized Mean-Squared Error (NMSE) and Adjacent-Channel Error-Power Ratio (ACEPR). The authors suggest that applying techniques for Digital Pre-Distortion (DPD) linearization could significantly reduce predictor computational complexity by developing a single predictor capable of modelling a PA across different output-power levels (Dikmese et al., 2019).

Sriram and Tuffy present an ML-based accelerated optimization framework for RF PA design that utilizes Max-Min Latin Hypercube Sampling with CatBoost gradient boosting to intelligently explore multidimensional parameter spaces. This approach reduces simulation requirements by 65% while maintaining  $\pm 0.4$  dBm accuracy for most modes. The trained CatBoost model predicts P2dB performance across the entire design space, transforming PA design from exhaustive search to intelligent exploration and achieving a 58.24%–77.78% reduction in total simulation time with an average  $R^2$  of 0.901 across 15 PA operating modes, enabling design cycles twice as fast without compromising accuracy (Sriram & Tuffy, 2025).

Tong et al., propose a measurement-driven automated PA design approach that leverages load-pull measurement data and ML methods for the design of harmonically controlled PAs. By introducing a two-stage optimization strategy using the Particle-Swarm Optimization (PSO) algorithm, accurate matching of harmonic impedances is achieved. The automated flow demonstrates high performance in a 2.4 GHz PA, attaining a maximum Power-Added Efficiency (PAE) of 73.4% (or 73.6%) and a maximum output power of 40 dBm, confirming that the process is able to avoid multiple simulation-tool iterations (Tong et al., 2024).

Cai et al., present a dynamic behavioural-modelling technique for multi-device RF PAs that often exhibit strongly nonlinear behaviour with long-term memory effects. The method combines a real-valued decomposed-piecewise approach with Support-Vector Regression (SVR) to accurately model distinct characteristics of the multi-device system across different power levels. Experimental results show that the decomposed piecewise SVR model provides significant performance improvements over standard SVR models, achieving approximately 10 dB–14 dB reduction in Normalized Mean Square Error (NMSE) for a single-transistor PA and an NMSE of 50 dB for a multi-device Doherty PA (Cai et al., 2020).

Uğur used a dataset consisting of measurements from 200 power-converter modules obtained during prototype and pilot production at three power levels, three temperatures and 128 frequency points. The features included driver-bias voltage, output-bias voltage, target output power, temperature, gain and the current drawn from the main power source. Hyper-parameter optimization is performed automatically with the third-party AX platform, which explored the hyper-parameter space using Sobol sampling to identify the configuration yielding the best performance. Among the evaluated models, a Multilayer Perceptron (MLP) achieved the

lowest prediction error 1.07% Mean Absolute Error (MAE) on the second dataset, outperforming both a Convolutional Neural Network (CNN) and a linear-regression baseline. The inclusion of the primary-current feature in the second dataset further improved the MLP's gain-prediction accuracy, confirming the relevance of current-related variables for Printed Circuit Board (PCB) characterization. Conversely, the CNN's performance deteriorated when the larger, current-augmented dataset is used and its training cost is markedly higher than that of the MLP, indicating that CNNs are less suited to high-dimensional tabular data in this context (Uğur, 2023).

Artificial Intelligence (AI) can optimize non-linearity. Military tech communication devices and systems need optimization to prevent any performance decreases. Climatic conditions, supply current and voltages, waveforms, aging, design limitations cause non-linearity. Although the hardware is the same design, units identified by different component which has different serial number may differ slightly in their test results. To remove non-linearity, RF calibration is essential and highly recommended.

An ML-based prediction of the RF calibration via digital control operations is presented in this paper. While ML techniques have been previously explored for RF PA modelling and calibration-related tasks, their specific application to predicting digital control parameter values for T/R-modules using large-scale production data remains limited (Cai et al., 2020; Chang et al., 2025; Dikmese et al., 2019; Sriram et al., 2025; Uğur et al., 2023). In this study, we address this application-specific gap by generating a ML-based model that predicts digital control parameter values of the RF circuit of the settings with a negligible error margin, replacing conventional and self-test calibration methods. To achieve this, thousands of data rows are processed, interpolated and combined to create benchmark comparisons across multiple models, resulting in a prediction success rate of 66.80% within the defined error bounds. Although ML techniques have been previously used for RF PA behavioural modelling and circuit design optimization, their application to predicting digital control knob parameters for military-grade T/R modules directly from large-scale production test data without requiring post-fabrication calibration runs remains an unexplored domain. Among the seven Received-Signal-Strength-Indicator (RSSI) parameters recorded during RF calibration, RSSI F is identified by the RF designers as the most critical parameter governing stable device operation and is therefore selected as the sole prediction target. The RSSI F value of fabricated T/R PA modules is predicted using the LightGBM algorithm, which significantly reduces the cost and labour force associated with iterative calibration approaches. The methodology consists of large-scale data acquisition, interpolation to fill missing entries, training and benchmarking of several regression algorithms (including XGBoost, LightGBM, and TabNet) with hyper-parameters tuned via Bayesian Optimization and selection of the best performing LightGBM based on error metrics such as Root Mean Squared Error (RMSE), Root Mean Square Percentage Error (RMSPE) and Mean Absolute Percentage Error (MAPE). The proposed framework offers high repeatability and robustness to process variations; however, it depends on the availability of a comprehensive measurement database and its accuracy may degrade with limited data, while the current implementation addresses only a single

performance metric, leaving multi-objective calibration as a future challenge. In conclusion, the study confirms that ML-driven calibration specifically gradient-boosting regression can reliably predict digital control parameter values of T/R module with negligible error, streamlining the calibration workflow.

The key contributions of this work are:

- An ML-based RF calibration framework that reduces calibration time from 17 hours to minutes.
- Early application of gradient-boosting-based ML methods for RF PA calibration in production-scale military T/R modules.
- Systematic comparison of XGBoost, LightGBM, and TabNet on harmonized datasets.
- 66.80% prediction accuracy within  $\pm 12$  mV error bounds using LightGBM regression.

The remainder of this paper is organized as follows: Section II describes the methodology and RF calibration workflow. Section III presents the dataset characteristics and pre-processing steps. Section IV details the ML models and performance metrics. Section V discusses the experimental results and model comparisons. Section VI concludes the paper with summary remarks and future research directions.

## 2. Methodology

RF calibration is the process of setting specific parameters so that a device operates stably and efficiently. With the current method, calibrating the output power of a single device takes about 17 h and incurs additional costs such as labour and test-setup expenses. A regression-based ML approach is developed to shorten the RF calibration time and make it more cost-effective. The model is built in Python, using NumPy and Pandas for data pre-processing and manipulation and Scikit-Learn for applying the ML algorithms. The algorithms employed are XGBoost (using the XGBoost library), LightGBM (using the LightGBM library), and TabNet (implemented with the PyTorch/Torch library).

During model training, the data are partitioned into training, validation and test sets with proportions of 70%, 15% and 15% respectively and 5-fold cross-validation is employed to mitigate overfitting. The ML models are trained on an Intel Xeon Gold 6254 CPU and an NVIDIA Grid T4-4Q GPU.

### 2.1.Dataset & Data Pre-processing

The calibrated measurements including frequency, power level, temperature and waveform labels are retrieved from a Structured Query Language (SQL) database. The database includes 187 PA modules records. For 167 of these modules, the data are partitioned into training, validation and test sets, while the remaining 20 modules are never exposed to the model and used for only testing the model to obtain the performance metrics of the model.

The experiment employed two distinct datasets. The first is referred to as the calibration-dataset, includes four fundamental measurement attributes: waveform, temperature, power

level and frequency. The second is designated as the module-test dataset and contains the same four attributes and additionally, five RSSI values obtained during module-level testing to characterize module-specific behaviour.

The module-test dataset, an independent source of T/R-module measurements, comprises a single temperature level, one waveform, two distinct power levels and 40 frequency points. In contrast, the calibration-dataset comprises three temperature levels, two waveforms, five power levels and 64 frequencies. To harmonize the two datasets, the 40 frequencies present in the module-test data are linearly interpolated to match the 64 frequencies of the RF calibration values and the two power levels are interpolated to five power levels. To capture the intrinsic characteristics of each module, a comprehensive attribute set consisting of 41 distinctive features is derived using an interpolated dataset.

## **2.2. Feature Engineering**

To effectively capture the complex RF dependencies and module-specific signatures embedded in the T/R-module measurements, a systematic feature engineering strategy is employed on the harmonized dataset. The raw module-test records provide five distinct Received-Signal-Strength-Indicator (RSSI) channels denoted RSSI A, RSSI B, RSSI D, RSSI F, and RSSI G each representing an independent receiver path. Rather than relying solely on these absolute values, the relative interactions between paths are encoded by constructing pairwise features across four physically meaningful antenna pairings: A–B, A–D, B–D, and F–D. For each pair, eight mathematical transformations are computed, yielding 32 derived descriptors.

These transformations are deliberately chosen to expose complementary aspects of the RF behaviour. The arithmetic mean and sum quantify common-mode signal strength and aggregate energy content, whereas the difference highlights relative imbalances that may indicate differential phase shifts or asymmetric attenuation between paths. The product captures multiplicative channel interactions, while the ratio and inverse ratio furnish scale-invariant metrics that suppress absolute power fluctuations and emphasize proportional disparities. The geometric mean characterizes multiplicative central tendency, and the harmonic mean provides an appropriate averaging measure for rate-like quantities such as signal-to-noise contributions. To ensure numerical stability particularly under weak-signal or near-null conditions a small regularization constant ( $\epsilon = 1 \times 10^{-8}$ ) is incorporated into all ratio and harmonic mean computations, thereby preventing division-by-zero artefacts.

This set of 32 engineered interaction features is augmented with the original five RSSI measurements and the four environmental/operational covariates: ambient temperature, waveform type, transmit power level, and operating frequency. The resulting 41-dimensional feature vector constitutes a rich, informative representation that enables the learning algorithm to discern intricate nonlinear mappings between the module's operating context and the target RSSI F calibration response. By explicitly modelling both intra-channel coupling phenomena and extrinsic operational conditions, this feature space enhances discriminative capacity and supports robust generalization, particularly for unseen modules during inference.

### **2.3. ML Models**

This study addresses a regression problem based on tabular dataset. In such datasets, variables typically exhibit non-linear and highly complex interactions, which often renders classical linear models inadequate. Consequently, tree-based gradient boosting algorithms tend to demonstrate superior prediction performance under these conditions. In this study, XGBoost and LightGBM are used as base gradient boosting models. Additionally, the TabNet algorithm, based on deep learning is evaluated to provide a comparative analysis between traditional boosting methods and a neural network-based approach.

#### **2.3.1. XGBoost**

XGBoost, proposed by Chen and Guestrin, is an advanced gradient boosting algorithm known for its high scalability. This method uses second-order, i.e., Hessian gradient information and a regularized objective function containing both  $L_1$  and  $L_2$  penalties. These features provide fast and stable performance on large datasets, reducing overfitting (Chen and Guestrin, 2016).

#### **2.3.2. LightGBM**

LightGBM, proposed by Ke et al., is a gradient-boosting framework that achieves high training speed and predictive accuracy on large-scale datasets. It reduces computational burden substantially through two innovative techniques: Gradient-based One-Side Sampling, which retains instances with large gradients while randomly discarding a large portion of instances with small gradients; and Exclusive Feature Bundling, which merges mutually exclusive sparse features into a single composite feature, thereby decreasing dimensionality without loss of information. Consequently, LightGBM delivers fast, memory-efficient learning while maintaining state-of-the-art performance (Ke et al., 2017).

#### **2.3.3. TabNet**

TabNet, proposed by Arik and Pfister, is a deep learning architecture for tabular data that performs feature selection step-by-step through an attention-based mechanism. Thus, it provides interpretability by nature. By applying a sparse mask to determine the most important features at each decision step, the model achieves high prediction accuracy while additionally offering strong explainability (Arik and Pfister, 2019).

### **2.4. Performance Metrics**

The study employed MAPE, RMSE and RMSPE as performance metrics to evaluate the models from different perspectives. RMSE measures the absolute error in the unit of the target variable. It is sensitive to large residuals, which can cause scale effects when the dataset is extensive and non-uniform. MAPE expresses the error as a percentage, offering a simpler and more interpretable measure, yet it becomes unstable when the true values approach zero due to its sensitivity to small magnitudes. RMSPE combines percentage-based error measurement with sensitivity to relative errors, thereby providing a more stable and comparable performance assessment for datasets with a wide range of values.

### 2.4.1. RMSE

As emphasized by Chai and Draxler and Willmott and Matsuura, RMSE is a regression performance metric defined as the square root of the mean of the squared errors; it penalizes large deviations more heavily and, because it shares the same unit as the target variable, is straightforward to interpret (Chai and Draxler, 2014; Matsuura, 2005).

This formulation is presented in Equation (1).

$$RMSE = \sqrt{\frac{\sum (y_i - \hat{y}_i)^2}{n}} \quad (1)$$

$y_i$ : real value

$\hat{y}_i$ : predicted value

$n$ : number of values

### 2.4.2. MAPE

As emphasized by Hynman&Koehler, and Armstrong&Collopy, MAPE expresses forecast errors in percentage terms, making it easy to interpret, but it tends to inflate errors when the true value approaches zero (Hynman and Koehler, 2006; Armstrong and Collopy, 1992).

The MAPE formula is presented in Equation (2).

$$MAPE = \frac{\%100}{n} \sum_{i=0}^n \left| \frac{y_i - \hat{y}_i}{y_i} \right| \quad (2)$$

$y_i$ : real value

$\hat{y}_i$ : predicted value

$n$ : number of values

### 2.4.3. RMSPE

As highlighted by De Myttenaere et al., RMSPE is a regression performance metric that computes the square root of the mean of squared percentage errors, thereby imposing heavier penalties on large relative deviations. It is especially advantageous in tabular-data problems where scale disparities are evident (De Myttenaere et al., 2016).

The RMSPE formula is presented in Equation (3).

$$RMSPE = \sqrt{\frac{1}{n} \sum_{i=0}^n \left( \frac{y_i - \hat{y}_i}{y_i} \right)^2} \quad (3)$$

$y_i$ : real value

$\hat{y}_i$ : predicted value

$n$ : number of values

## 2.5. Statistical Validation via Cluster Bootstrap

To address the inherent uncertainty associated with reporting performance metrics from a fixed test set of 20 modules, a cluster bootstrap procedure is employed at the module level. This approach preserves the within-module correlation of the 320 measurement points per module, providing a statistically valid estimate of metric variability. The 20 unseen test modules are resampled with replacement to generate  $B = 5000$  bootstrap samples. For each resample, the success rate (within  $\pm 12$  mV tolerance), outlier count, RMSE, MAPE, and RMSPE are recomputed for all four models. The 95% confidence intervals are derived using the percentile method, taking the 2.5th and 97.5th quantiles of the bootstrap distributions. This procedure quantifies the sensitivity of the reported metrics to the specific composition of the test set.

## 2.6. Hyperparameter Optimization

Hyperparameter tuning is carried out using a Bayesian Optimization framework. This probabilistic search strategy builds a surrogate model of the objective function and explores the hyperparameter space efficiently, thereby locating near-optimal configurations with a reduced number of evaluations. During the optimization, RMSE served as the primary performance metric for all models. For XGBoost, however, both RMSE and the RMSPE are employed, allowing the model's error to be assessed in absolute as well as relative (percentage) terms and enabling a direct comparison of the two evaluation criteria. This dual-metric scheme highlights how the choice of performance metric influences model selection while demonstrating the capability of Bayesian Optimization to discover high performance hyperparameter settings efficiently. Consequently, optimal hyperparameter sets are identified for each algorithm, and the overall predictive accuracy is maximized.

Table 1 lists the hyperparameter settings that achieved the best performance for XGBoost in the context in this study. The table shows the configurations obtained when the model is separately with respect to RMSE and RMSPE, indicating that these are the optimal XGBoost hyperparameters for this work.

**Table 1.** The optimal hyperparameters for XGBoost.

| Model   | Hyperparameter   | Value   | Hyperparameter   | Value   |
|---------|------------------|---------|------------------|---------|
|         | RMSE             |         | RMSPE            |         |
| XGBoost | n Estimators     | 1414    | n Estimators     | 1998    |
|         | Max Depth        | 12      | Max Depth        | 12      |
|         | Learning Rate    | 0.0580  | Learning Rate    | 0.03527 |
|         | Min Child Weight | 3       | Min Child Weight | 9       |
|         | Subsample        | 0.7659  | Subsample        | 0.7847  |
|         | Colsample Bytree | 0.8824  | Colsample Bytree | 0.8670  |
|         | Reg Alpha        | 0.4147  | Reg Alpha        | 3.1873  |
|         | Reg Lambda       | 38.0753 | Reg Lambda       | 36.0197 |
|         | Gamma            | 0.9294  | Gamma            | 0.8727  |
|         | Max Bin          | 249     | Max Bin          | -       |

Table 2 lists the hyperparameters of LightGBM that are obtained through the optimization process and represent the best-performing configuration for this study.

**Table 2.** The optimal hyperparameters for LightGBM.

| Model    | Hyperparameter    | Value   |
|----------|-------------------|---------|
|          | <i>RMSE</i>       |         |
| LightGBM | n Estimators      | 2465    |
|          | Max Depth         | 16      |
|          | Learning Rate     | 0.0589  |
|          | Num Leaves        | 305     |
|          | Min Child Samples | 14      |
|          | Subsample         | 0.7956  |
|          | Colsample Bytree  | 0.9755  |
|          | Reg Alpha         | 0.7870  |
|          | Reg Lambda        | 44.2262 |
|          | Max Bin           | 265     |

Table 3 lists the hyper-parameters of TabNet that are obtained through the optimization process and represent the best-performing configuration for this study.

**Table 3.** The optimal hyperparameters for TabNet.

| Model  | Hyperparameter     | Value  |
|--------|--------------------|--------|
|        | <i>RMSE</i>        |        |
| TabNet | n d                | 36     |
|        | n a                | 26     |
|        | n Steps            | 4      |
|        | Gamma              | 1.4718 |
|        | Lambda Sparse      | 0.0020 |
|        | Batch Size         | 256    |
|        | Virtual Batch Size | 128    |

### 3. Results

The predictive performance is assessed on the forecasts of 20 modules, which together comprise 6400 individual data points. For the outlier detection analysis, the 12mV tolerance threshold is defined based on domain expertise provided by the RF system designers, who identified this value as the critical limit for ensuring stable module operation; any prediction deviating from the reference value by more than this limit is classified as an outlier.

Four regression models are trained for the forecasting task, two XGBoost models are built, each with its hyperparameters tuned separately for the RMSE and RMSPE loss functions. The LightGBM and TabNet models are tuned with respect to RMSE. All hyperparameter searches

are carried out using Bayesian Optimization to ensure systematic and reproducible selection of the optimal configurations.

As illustrated in Table 4, the four models are compared across three error metrics MAPE, RMSE and RMSPE. In this evaluation the LightGBM model achieved the best overall performance, followed by XGBoost model whose hyperparameters are optimized for RMSPE. The TabNet model exhibited the poorest performance among the four approaches.

**Table 4.** Model performance evaluation.

|                      | RMSE (mV) | MAPE (%) | RMSPE (%) |
|----------------------|-----------|----------|-----------|
| <b>XGBoost-RMSE</b>  | 13.152    | 0.605    | 0.761     |
| <b>XGBoost-RMSPE</b> | 13.151    | 0.598    | 0.754     |
| <b>LightGBM</b>      | 12.292    | 0.570    | 0.714     |
| <b>TabNet</b>        | 15.742    | 0.712    | 0.925     |

Figure 1 and Table 5, the XGBoost model optimized for RMSPE identified 2,215 outliers and had a maximum deviation of 48.52 mV, whereas the model optimized for RMSE detected 2,254 outliers with a maximum deviation of 43.47 mV; thus, the RMSE-optimized model exhibits a lower maximum deviation. The RMSE-optimized XGBoost model is less effective at limiting the total number of outliers and detects more outliers with larger deviations. In contrast, the RMSPE-optimized model finds fewer outliers, and the deviations of those outliers are smaller.


**Table 5.** Outlier count of models.

|                      | Outlier Count |
|----------------------|---------------|
| <b>XGBoost-RMSE</b>  | 2254          |
| <b>XGBoost-RMSPE</b> | 2215          |
| <b>LightGBM</b>      | 2124          |
| <b>TabNet</b>        | 2704          |

LightGBM and XGBoost exhibit divergent performance depending on the evaluation metric. As illustrated in Table 4 the LightGBM attains lower overall error metrics (e.g., MAPE, RMSE, RMSPE) than XGBoost model whose hyperparameters are tuned for RMSE. However, when the outlier detection results shown in Figure 1 and Table 5 are considered, LightGBM identifies 2,124 outlier fewer than the 2,254 detected by the RMSE-optimized XGBoost and its mean absolute error is 18.32 mV versus 19.45 mV for XGBoost. However, LightGBM maximum deviation is 50.25 mV, which exceeds the 43.47 mV maximum deviation of the RMSE-optimized XGBoost. Consequently, LightGBM is superior in outlier count and mean deviation, whereas XGBoost-RMSE performs better in terms of maximum deviation; the preferred model thus depends on the primary evaluation criterion.

In this study, in addition to the gradient boosting algorithms, the deep-learning-based TabNet model is also evaluated and its results are obtained. As shown in Figure 1 and Table 5, TabNet

produced the highest number of outliers among the tested models and are not able to achieve a performance comparable to that of the tree-based gradient boosting algorithms.




**Figure 1.** Results of models according to 12 mV limit.

**Table 6.** Bootstrap-based performance metrics with %95 confidence intervals (B=5000).

| Model                | Success Rate (%)       | RMSE (mV)              | MAPE (%)             | RMSPE (%)           |
|----------------------|------------------------|------------------------|----------------------|---------------------|
| <b>XGBoost-RMSE</b>  | 64.77<br>(63.63-65.90) | 13.13<br>(12.93-13.34) | 0.60<br>(0.60 -0.62) | 0.76<br>(0.74-0.78) |
| <b>XGBoost-RMSPE</b> | 65.38<br>(64.25-66.51) | 13.15<br>(12.92-13.38) | 0.60<br>(0.59-0.61)  | 0.75<br>(0.74-0.77) |
| <b>LightGBM</b>      | 66.80<br>(65.66-67.95) | 12.29<br>(12.08-12.50) | 0.57<br>(0.59-0.58)  | 0.71<br>(0.70-0.72) |
| <b>TabNet</b>        | 57.74<br>(56.52-58.95) | 15.74<br>(15.40-16.08) | 0.71<br>(0.56-0.58)  | 0.92<br>(0.90-0.94) |

Table 6 shows that the three gradient-boosting models (XGBoost-RMSE, XGBoost-RMSPE, and LightGBM) consistently outperform TabNet. Their success rates are higher (64.8%–66.8% versus 57.7%), their RMSE values are lower (12.3–13.2 mV versus 15.7 mV), and their percentage errors are smaller (MAPE  $\leq$  0.60%, RMSPE  $\leq$  0.76%). The near-identical RMSE values between XGBoost-RMSE and XGBoost-RMSPE (13.15 mV for both) stem from the fact that the target variable F lies in the range of approximately 1200–1500 mV, while the prediction errors are around 13 mV. At this scale, percentage errors ( $\sim$ 0.6–0.8%) are nearly proportional to absolute errors, so optimizing for RMSE and optimizing for RMSPE yield very similar models and performance. Nevertheless, XGBoost-RMSPE achieves a slightly higher

success rate (65.38% versus 64.77%) and a marginally lower RMSPE (0.75% versus 0.76%), confirming that the relative-error objective provides a small but measurable advantage. Moreover, the 95% confidence intervals for these metrics are narrow (e.g., success-rate width  $\approx 2$  pp, RMSE width  $\approx 0.4$  mV), indicating that performance is stable across different test splits. In contrast, TabNet has wider intervals and larger errors, confirming that tree-based boosting methods provide both greater accuracy and higher reliability for this regression task.




**Figure 2.** Per-Module RMSE Comparison Across Four Models for the 20 Test Modules.


The per-module RMSE values in Figure 2 demonstrate that the LightGBM ensemble provides the most consistently low error across the 20 test modules; it achieves the smallest RMSE for 7 of the 20 units and its overall mean RMSE ( $\approx 9.6$  mV) is lower than those of the two XGBoost variants (mean  $\approx 12.0$  mV) and TabNet (mean  $\approx 14.8$  mV). XGBoost-RMSE and XGBoost-RMSPE show comparable average performance but exhibit larger variability, with several modules (e.g. 1, 8, 9, 15) producing RMSEs above 20 mV, whereas LightGBM never exceeds 22 mV. TabNet attains the highest RMSE on several modules (14) and yields the greatest spread of errors, confirming its inferior predictive accuracy for this dataset. Consequently, the table indicates that gradient-boosting models particularly LightGBM are more accurate and more robust to module-to-module variations than the TabNet architecture.

The per module MAPE values in Figure 3 show that the LightGBM model consistently delivers the lowest errors for the majority of the 20 test modules (13 out of 20), attaining a mean MAPE of  $\approx 0.49\%$ , which is lower than the averages of XGBoost-RMSE ( $\approx 0.62\%$ ), XGBoost-RMSPE ( $\approx 0.66\%$ ) and TabNet ( $\approx 0.73\%$ ). XGBoost-RMSE and XGBoost-RMSPE produce comparable performance, but their errors are more dispersed, with several modules (9, 15) exceeding 1% MAPE. TabNet exhibits the poorest consistency, recording the highest MAPE on several modules (5 and 14) while only marginally outperforming the boosting models on a

few easy cases (7179 and 7185). These observations indicate that tree-based gradient-boosting ensembles particularly LightGBM provide more accurate and robust percentage-error predictions across heterogeneous modules than the TabNet architecture.



**Figure 3.** Per-Module MAPE Comparison Across Four Models for the 20 Test Modules.



**Figure 4.** Per-Module RMSPE Comparison Across Four Models for the 20 Test Modules.

The per-module RMSPE values in Figure 4 demonstrate that the LightGBM ensemble yields the most accurate and stable percentage-error performance across the 20 test units. LightGBM attains the lowest RMSPE on 13 of the 20 modules, with a mean RMSPE of  $\approx 0.65\%$ , whereas the two XGBoost variants achieve comparable average errors ( $\approx 0.72\%$  for XGB-RMSE and  $\approx 0.73\%$  for XGB-RMSPE) but display a broader spread, reaching values above 1 % on several modules (6 and 15). TabNet exhibits the poorest overall consistency, presenting the highest RMSPE on multiple hard cases (14 and 5) and a mean RMSPE of  $\approx 0\%$ . These observations confirm that the tree-based gradient-boosting approach, particularly LightGBM, provides superior and more reliable RMSPE predictions compared with both XGBoost configurations and the TabNet architecture.

The predictive performance is evaluated on the forecasts of twenty modules, corresponding to a total of 6400 measurement points. Four regression models XGBoost optimized for RMSE, XGBoost optimized for RMSPE, LightGBM and TabNet are assessed on outlier metrics (outlier count, outlier percentage, and mean absolute error). The LightGBM model achieved the lowest outlier occurrence (outlier count = 2124, outlier percentage = 33.19%, mean absolute error = 18.32 mV), followed by the XGBoost model optimized for RMSPE (outlier count = 2215, outlier percentage = 34.61%, mean absolute error = 19.48 mV). The XGBoost model tuned for RMSE yielded outlier count = 2254, outlier percentage = 35.22%, mean absolute error = 19.45 mV, whereas TabNet performed the worst (outlier count = 2704, outlier percentage = 42.25%, mean absolute error = 20.90 mV).

Outlier detection is carried out over the 6400 measurement points. The number of outliers identified by each model is 2124 for LightGBM (33.19%), 2215 for XGBoost–RMSPE (34.61%), 2254 for XGBoost–RMSE (35.22%) and 2704 for TabNet (42.25%). Accordingly, the proportion of predictions flagged as outliers is lowest for LightGBM, followed by the XGBoost configurations, and highest for TabNet. Among the outliers, LightGBM exhibits the lowest mean absolute error (18.32 mV) and the tightest dispersion (standard deviation = 5.53 mV), whereas XGBoost–RMSE and XGBoost–RMSPE record comparable mean values (19.45 mV and 19.48 mV, respectively) with standard deviations of 5.80 mV and 6.19 mV; TabNet, meanwhile, shows the highest mean absolute error (20.90 mV), the largest maximum deviation (68.94 mV) and the widest spread (standard deviation = 9.15 mV). These results indicate that LightGBM provides the most reliable performance overall, achieving the fewest outliers and the smallest error magnitude among them. The XGBoost models also deliver strong control, with the RMSPE-optimized variant incurring a slightly lower outlier count than the RMSE-optimized one, while the RMSE-optimized model records a marginally lower mean absolute error and maximum deviation on the outlier subset. In contrast, TabNet exhibits a markedly larger outlier count and higher error statistics, suggesting the deep-learning-based tabular approach is less suitable for this dataset.

The study therefore demonstrates that, for the examined tabular forecasting problem, gradient-boosting methods (particularly LightGBM and the XGBoost variants) outperform the deep-learning-based TabNet both in terms of outlier frequency and outlier magnitude reduction. Moreover, among the XGBoost models, hyperparameter tuning with RMSPE as the objective yields a slightly lower outlier percentage than tuning with RMSE, whereas the

RMSE-optimized model achieves a marginally better mean absolute error on the identified outliers. These findings support the preference for gradient boosting algorithms when addressing tabular prediction tasks in this domain, where interpretability, training efficiency and outlier control are jointly critical.

## 4. Discussion

The results show that the proposed ML-based approach can predict RF calibration parameters with a notable degree of accuracy. Within a  $\pm 12$  mV tolerance envelope, 66.80% of the predicted values agree with the reference measurements. This performance is attributable to the selected features, which directly capture the operating dynamics of the module. Importantly, the ML predictor achieves this agreement using calibration and module-test data that are substantially less extensive than the measurement sequences required by conventional calibration procedures. The 17 h manual calibration routine is reduced to a matter of minutes, yielding measurable savings in labour and turnaround time.

Prior work in RF PA design has similarly highlighted the time-intensive nature of conventional measurement-based calibration. Chang et al. (2025) report that traditional measurement-weighted methods for PA efficiency and linearity characterization, while accurate, require lengthy measurement sequences. Sriram and Tuffy (2025) address a related challenge in PA design, employing a CatBoost model with Max-Min Latin Hypercube Sampling to reduce simulation requirements by 65% and achieve  $\pm 0.4$  dBm prediction accuracy. While the tasks and metrics differ from those in the cited work, the present study shares the same motivation: replacing time-consuming measurement campaigns with data-driven prediction.

The results reveal a clear performance–complexity trade-off among the evaluated models. TabNet, despite its greater architectural complexity, yielded the weakest predictive performance and the highest outlier count, indicating that deep learning complexity does not confer an advantage for this particular tabular dataset. LightGBM and the RMSE-optimized XGBoost, by contrast, delivered superior accuracy with lower computational demands, consistent with the broader literature on gradient boosting versus deep learning for tabular data. A notable limitation is that model performance may degrade when applied to production batches with module characteristics substantially different from those in the training set. This generalizability concern underscores the need for future work to evaluate the models on larger and more diverse production datasets.

## 5. Conclusion

In this work a data-driven calibration framework is presented in which the raw calibration measurements and module-test results are directly employed as inputs to a ML predictor. This approach achieved a prediction success rate of 66.80% with an error bounded to  $\pm 12$  mV, demonstrating that the extracted features capture the intrinsic behaviour of the RF PA without the need for extensive additional measurement. By replacing the traditional 17 h manual calibration routine with an automated process that completes within a few minutes, the method

substantially reduces module related expenses, labour effort and overall turnaround time. The outcome proves that high-accuracy, rapid calibration is feasible for rugged military communication equipment operating under demanding environmental constraints.

In this study, four regression approaches are compared for tabular RF calibration forecasting and found that gradient boosting methods clearly outperform the deep-learning-based TabNet. Among the evaluated models, LightGBM achieved the most reliable outlier performance, while the XGBoost variant tuned for RMSE provided a robust overall error profile. Consequently, gradient boosting algorithms are the preferred solution for this type of task and future research should investigate ensemble techniques that combine the algorithms complementary strengths.

While the current results are promising, several approaches are able to be explored to further improve the solution. Expanding the training database with more data and diverse ML algorithms will increase the model's success. More characteristic feature extraction from the module-test data is able to provide richer representations and improve prediction accuracy, especially in tight tolerance windows. To address the slight performance degradation observed at ultra-low tolerances, alternative learning paradigms such as deep neural networks, ensemble methods and hybrid models are able to be investigated. Finally, integrating an online adaptation mechanism that continuously updates the predictor with field-collected calibration results will enable long term performance monitoring and further reduce maintenance costs of deployed systems.

## **Conflict of Interest**

The authors declare no conflicts of interest.

## **Funding**

No funding has been received for this research.

## **Acknowledgment**

The authors gratefully acknowledge ASELSAN A.Ş. for providing the computational resources and data used in this study, as well as for its support during the experimental phase. Also, the authors would like to thank to all individuals who contributed to the execution of this study and provided technical assistance during the measurement and testing procedures particularly the ASELSAN Communication and Information Technologies Business Sector, Department of Hardware Test and Verification Engineering.

## **Ethical Statements**

The study is performed solely on electronic hardware and software; no ethical committee approval is required.

## Data and Code Availability

The calibration data, feature definitions and software used in this work are not publicly available. The all data is provided by ASELSAN Communication and Information Technologies Business Sector.

## Supplementary Materials

Not applicable.

## References

- Arik, S. O., & Pfister, T. (2019). TabNet: Attentive interpretable tabular learning. arXiv preprint arXiv:1908.07442.
- Armstrong, J. S., & Collopy, F. (1992). Error measures for generalizing about forecasting methods: Empirical comparisons. *International Journal of Forecasting*, 8(1), 69–80.
- Cai, J., King, J. B., Yu, C., Pan, B., Sun, L., & Liu, J. (2020). Dynamic behavioural modelling of RF power amplifiers based on decomposed piecewise machine learning technique. *International Journal of Microwave and Wireless Technologies*, 1-7. <https://doi.org/10.1017/S1759078720001208>
- Chang, Y., Wen, Y., Hong, Z., Padmanabhan, B., Franzon, P. D., & Floyd, B. A. (2025). A 27–30 GHz T/R module with reflection-type phase shifting and machine-learned calibration. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 72(9), 4649-4660. <https://doi.org/10.1109/TCSI.2024.3524280>
- Chai, T., & Draxler, R. R. (2014). Root Mean Square Error (RMSE) or mean absolute error (MAE)? – Arguments against avoiding RMSE in the literature. *Geoscientific Model Development*, 7(3), 1247–1250.
- Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining* (pp. 785-794).
- De Myttenaere, A., Golden, B., Le Grand, B., & Rossi, F. (2016). Mean absolute percentage error for regression models. *Neurocomputing*, 192, 38-48.
- Dikmese, S., Anttila, L., Campo, P. P., Valkama, M., & Renfors, M. (2019). Behavioral modeling of power amplifiers with modern machine learning techniques. In *2019 IEEE MTT-S International Microwave Conference on Hardware and Systems for 5G and*

- Beyond (IMC-5G)* (pp. 1-3). IEEE. <https://doi.org/10.1109/IMC-5G47857.2019.9160381>
- Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22(4), 679–688.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., ... & Liu, T. Y. (2017). Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.
- Ortega-González, F. J., García, J. A., Patiño-Gómez, M., & Madueño-Pulido, D. (2025). High-Efficiency Transmitters Operating at HF, VHF, and UHF Frequencies. *IEEE Microwave Magazine*. <https://doi.org/10.1109/MMM.2025.3592159>
- Sriram, A., & Tuffy, N. (2025). Accelerating RF power amplifier design via intelligent sampling and ML-based parameter tuning. arXiv. <https://arxiv.org/abs/2507.11928>
- Tong, Y., Wang, J., Fu, T., Su, J., & Liu, J. (2024). Measurement-Driven Automated Design of High-Efficiency Power Amplifiers. In *2024 9th International Conference on Integrated Circuits and Microsystems (ICICM)* (pp. 129-133). IEEE. <https://doi.org/10.1109/ICICM63644.2024.10814353>
- Uğur, A. B. (2023). *Derin öğrenme tabanlı güç yükseltici modülü kalibrasyonu* (Master's thesis, Gazi University).
- U.S. Department of Defense. (2019). *MIL-STD-810H: Environmental engineering considerations and laboratory tests*. [http://everyspec.com/MIL-STD/MIL-STD-0800-0899/MIL-STD-810H\\_55998/](http://everyspec.com/MIL-STD/MIL-STD-0800-0899/MIL-STD-810H_55998/)
- Willmott, C. J., & Matsuura, K. (2005). Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance. *Climate research*, 30(1), 79-82.

## Domain-Specific Summarization to Enhance LLM Effectiveness

Hamza Emin Hacıoğlu<sup>1,\*</sup> , İsmail Karakaya<sup>1</sup> , Ensar Erdoğan<sup>1</sup> ,  
Serdar Kalaycı<sup>1</sup> , Müge Ak<sup>1</sup> , Özgür Umut Vurgun<sup>1</sup> , Arif Furkan Mendi<sup>1</sup> ,  
Mehmet Akif Nacar<sup>1</sup> 

### Abstract

Large Language Models (LLMs) have become increasingly prevalent in the field of natural language processing, particularly for text summarization tasks. Their ability to capture semantic relationships and generate coherent summaries has made them a powerful alternative to traditional extractive and abstractive summarization methods. However, the effective summarization of long and domain-specific texts, such as petitions, remains a challenging problem. This study explores the problem of petition text summarization by applying state-of-the-art LLMs and evaluating various training strategies on a custom-built dataset. The paper investigates different model training approaches, including Low-Rank Adaptation (LoRA) and prompt engineering techniques, to optimize summarization quality. Human-in-the-loop summarization methods are also incorporated to refine the models, ensuring high-quality outputs in real-world scenarios. Furthermore, the model's performance is assessed through expert evaluation, which highlights areas for potential improvement in terms of both accuracy and coherence. The research further delves into methods for effectively summarizing long-form petition texts, exploring ways to preserve critical information while maintaining brevity and readability. The results demonstrate the feasibility of leveraging LLM-based summarization methods for complex, lengthy texts while providing insights into the challenges and opportunities within this evolving field.

**Keywords:** LoRA, fine-tuning, SFT, context length, petition text, prompt engineering.

Received: 12.11.2025

Accepted: 18.05.2026

Research Article

<sup>1</sup> HAVELSAN Inc., Ankara, Türkiye

\*Corresponding author.

✉ hamzaeminh@havel-san.com.tr

**Cite this article as:**

<https://doi.org/10.65834/jdsi.12.32>

Hacıoğlu, H. E., Karakaya, I., Erdoğan, E., Kalaycı, S., Ak, M., Vurgun, O. U., Mendi, A. F., & Nacar, M. A. (2026). Domain-Specific Summarization to Enhance LLM Effectiveness. *Journal of Defence and Security Industries: Strategy and Technology* 2, 129-156.

## 1. Introduction

Text summarization refers to the task of condensing long texts while preserving key information, and significant progress has been made with the emergence of deep learning and transformer-based models. Traditionally, summarization techniques have been divided into two main approaches: extractive and abstractive summarization. Extractive methods select and extract important sentences directly from the source text, while abstractive methods aim to generate new sentences that capture the essence of the input. Although both approaches have proven effective in various domains, challenges remain when dealing with long, complex, and domain-specific texts that require specialized understanding. Recent advances in artificial neural network architectures, particularly transformer-based models, have led to substantial improvements in summarization tasks. Models such as BERT, GPT, and their derivatives excel at capturing long-range dependencies and contextual nuances in language. However, despite these advancements, challenges persist in summarizing long texts, especially legal documents such as petitions. These challenges stem from the need to effectively integrate both local (sentence-level) and global (document-level) context. To address these issues, researchers have developed hybrid approaches that combine extractive and abstractive methods, providing more robust solutions for handling longer and more complex documents. Despite progress in the field of text summarization, domain-specific characteristics of the documents to be summarized still present significant challenges. This is particularly true for petition texts, which are often written in diverse styles and legal language by different individuals. Traditional summarization approaches often fail to effectively summarize such texts. Petition documents tend to be lengthy, lack a standardized structure, and may contain unclear or vague claims from the author. Consequently, automatically generating summaries through artificial intelligence remains a difficult task. Traditional summarization methods, whether extractive or abstractive, struggle to capture the subtle nuances of such documents. Challenges are especially prominent in maintaining the integrity of the text, balancing realism with detail, and ensuring brevity without losing essential information. This study explores innovative techniques for summarizing petition texts that combine local and global contexts, incorporate discourse-aware structures, and implement fine-tuning strategies for domain-specific performance. The goal of this research is to bridge the gap between current summarization techniques and the specific requirements for summarizing petition texts, ensuring that the generated summaries are not only concise but also coherent, relevant, and contextually rich. By focusing on the efficient processing of long documents, this work demonstrates how the combination of extractive and abstractive methods, along with domain adaptation, can significantly enhance the quality and usability of petition text summarization.

Research on text summarization has undergone a significant evolution, progressing from traditional extractive and statistical methods to advanced neural and transformer-based approaches. Early models focused on sentence salience and topic distribution to identify the most representative information within a document. With the rise of deep learning, sequence-to-sequence architectures and attention mechanisms enabled more coherent and contextually aware abstractive summaries. Subsequent developments in transformer-based models

introduced improved contextual understanding, enabling hybrid frameworks that combine extractive precision with abstractive fluency. More recent efforts have concentrated on the challenges of long document summarization, exploring efficient attention mechanisms, hierarchical encoding strategies, and multi-stage frameworks that balance local and global contextual information. These advances have substantially improved the ability of models to handle extended inputs such as reports, dialogues, and multi-section documents. Parallel to advancements in summarization architecture, fine-tuning strategies for LLMs have emerged as a key enabler of task specialization and efficiency. Parameter-efficient adaptation methods, such as low-rank approximation and quantization, have been shown to significantly reduce computational overhead while maintaining or improving summarization performance. The integration of discourse-aware structures and task-specific adaptation techniques has enhanced coherence and factual consistency, particularly in long-form and domain-specific summarization tasks. Furthermore, the combination of supervised, unsupervised, and in-context learning approaches has allowed models to better generalize across diverse datasets while maintaining adaptability to specialized domains. Together, these developments reflect a broader shift toward efficient, scalable, and contextually rich summarization systems grounded in large language model fine-tuning.

Text summarization methods for long documents aim to preserve both local sentence-level information and global document-level coherence. A study shows that this challenge can be addressed by integrating local context, derived from sentence proximity, with global context representing document structure and themes (Xiao & Carenini, 2019). By jointly modelling these two levels of information, their extractive approach improves sentence selection and produces summaries that better reflect the central ideas of long documents while maintaining coherence across sections. Another line of work focuses on improving the efficiency of attention mechanisms in Transformer-based models for long inputs. Another study shows that standard full attention becomes impractical for long sequences due to quadratic computational and memory costs (Huang et al., 2021). To address this limitation, they propose HEPOS, a head-wise stride-based attention mechanism that distributes token coverage across attention heads. This design reduces attention complexity while preserving both local and global dependencies, enabling efficient processing of long documents with competitive summarization performance. Multi-stage summarization frameworks further mitigate context-length limitations by combining extractive and abstractive strategies. For example, SummN, which segments long inputs into manageable chunks, applies extractive filtering to identify salient content, and progressively refines summaries through staged compression and abstractive rewriting (Zhang et al., 2022). This approach allows coherent summarization of long, information-dense documents without exceeding model input limits. Earlier extractive approaches rely on structural importance rather than generation. For instance, LexRank represents sentences as nodes in a similarity graph and selects salient sentences based on graph centrality (Erkan & Radev, 2004). While effective for identifying important content, such graph-based methods may produce disjointed summaries and can be computationally expensive for very long documents.

Topic modelling methods, such as latent Dirichlet allocation, support summarization by identifying dominant themes within documents (Blei et al., 2003). Although these methods provide strong thematic coverage, they often overlook fine-grained details and local discourse structure, leading to overly coarse summaries. Neural sequence-to-sequence models with attention have significantly advanced abstractive summarization. Also, there are pointer-generator networks that combine text generation with copying mechanisms, improving coherence and factual consistency (See et al., 2017). However, these models are computationally intensive and require substantial training data. Sentence embedding-based extractive methods select salient sentences by comparing contextual sentence representations in a continuous semantic space. This approach fine-tunes BERT for extractive summarization, enabling improved semantic matching beyond surface-level similarity, albeit at increased computational cost (Liu, 2019). Finally, hybrid approaches seek to balance factual accuracy and linguistic fluency by combining extractive and abstractive techniques. Furthermore, a study proposes a multi-task framework that jointly optimizes both objectives, improving readability while preserving content coverage, at the cost of increased system complexity (Chen et al., 2019).

Recent work on fine-tuning LLMs has increasingly emphasized achieving strong task performance with reduced computational cost, especially for summarization. A dual-stage framework for news summarization combines in-context learning and parameter-efficient fine-tuning (Guan et al., 2024). In this approach, the model is first adapted to the document-specific context using in-context learning (ELearn), which can be particularly useful when annotated examples are limited (Guan et al., 2024). A second stage then applies an efficient fine-tuning step (EFit) to update only a small subset of parameters, aiming to retain general capabilities while improving task-specific summarization quality under constrained resources (Guan et al., 2024). Overall, this design targets a practical balance: fast adaptation to the current input distribution through prompting, followed by lightweight learning that stabilizes and improves output quality. Large-scale evidence for the effectiveness of parameter-efficient adaptation is provided by an extensive evaluation of 310 LoRA-fine-tuned LLMs across diverse benchmarks (Zhao et al., 2024). The results highlight how LoRA can achieve strong performance without updating full model weights, significantly reducing training compute and time compared with full fine-tuning (Zhao et al., 2024). This work suggests that LoRA supports scalable specialization, allowing models to be adapted to tasks or domains while preserving broadly useful language capabilities, which is attractive for organizations that need multiple specialized models but cannot afford repeated full fine-tuning. Beyond efficiency, recent research has explored incorporating higher-level structure into fine-tuning for long-document summarization. A discourse-aware approach integrates Rhetorical Structure Theory (RST) into LoRA-based adaptation to better capture hierarchical structure and rhetorical dependencies in long texts, improving coherence and reducing omissions or inconsistencies in abstractive summaries (Liu & Demberg, 2024). This line of work reflects a broader trend toward adapting LLMs not only to domain vocabulary but also to document organization and discourse-level constraints that matter for summary quality. Another practical constraint is deploying fine-tuned models under quantization. A method called IR-QLoRA mitigates the accuracy


degradation commonly observed when LoRA-based fine-tuning is performed under ultra-low-bit quantization (e.g., 2–3 bits), by combining Information Calibration Quantization (ICQ) and Information Elastic Connection (IEC) to preserve and recover information more effectively (Qin et al., 2024). Reported results on LLaMA-family models indicate measurable gains compared with prior approaches, while adding minimal overhead, which supports the feasibility of deploying quantified, fine-tuned LLMs in resource-constrained settings (Qin et al., 2024). Fine-tuning strategies have also been examined more systematically in the context of report summarization. A distinction is drawn between Knowledge Fine-Tuning (KFT), which adapts models to domain-specific language using specialized corpora (e.g., government or legal text), and Format Fine-Tuning (FFT), which optimizes task behaviour using paired inputs and reference summaries to reinforce output structure and content selection patterns (Rallapalli et al., 2025). Further discussion covers unsupervised and semi-supervised methods, such as contrastive learning, self-training, and hybrid training using labelled and unlabelled data, as practical alternatives when high-quality labelled summaries are limited (Rallapalli et al., 2025). Together, these studies motivate fine-tuning pipelines that combine efficient adaptation, structural awareness, and training-data pragmatism to make LLM summarization both effective and deployable in real-world scenarios.

## 2. Methodology


The methodology outlined in this study establishes a structured approach for training a summarization model tailored to Turkish public procurement petitions. The dataset used in this work is confidential and accessible only within institutional authorization limits. The workflow integrates dataset preparation, supervised fine-tuning with controlled prompt variation, efficient parameter adaptation using LoRA, and systematic strategies for handling documents that exceed the model's effective context window. To ensure that the model produces summaries aligned with official Public Procurement Authority writing standards, a two-stage inference procedure is used to separate argumentative content extraction from linguistic normalization. Throughout the process, human expert evaluation was incorporated to assess factual completeness, structural fidelity, and stylistic conformity, enabling iterative refinement of both the model and the prompting strategy.

### 2.1. Data Source

The dataset comprises 5215 petition text-summary pairs written in Turkish. The average length of the petition texts is 8449 characters, while the average length of the summaries is 3251 characters. The distribution of text lengths is illustrated in Figure 1. In addition, the most frequent words, excluding conjunctions, are presented in Figure 2.



**Figure 1.** Distribution of petition vs. summary text lengths.



**Figure 2.** Most frequent words in the dataset, excluding conjunctions ('bağlaçlar' in Turkish language).

### 2.1.1. Characteristics of Petition Texts and Summaries

Petitions are primarily submitted by companies whose tender applications have been rejected, often written from a first-person perspective. These texts are marked by linguistic and structural inconsistencies, including misspellings, syntactic errors, inverted sentence constructions, and unrelated narrative digressions. Many petitions lack clear claim numbering, with some repeating the same argument or combining multiple issues into one claim. These


inconsistencies introduce variability in text quality, making automatic summarization challenging. In contrast, the summaries are formally consistent, well-structured, and written in clear, standardized Turkish language. Claims are numbered and separated into distinct statements, even when the original petition merges them. Repeated claims in the petition are consolidated in the summary. Written in the third person, the summaries use passive constructions and often end with phrases such as “... yapıldığı,” “... edildiği,” or “... iddialarına yer verilmiştir”. Company names are mentioned once and referred to generically (e.g., “istekli,” “birinci istekli,” “ikinci istekli”), and the Public Procurement Authority (Kamu İhale Kurumu) is consistently referred to as “idare”.

### 2.1.2. Dataset Selection and Filtering

To prepare the dataset for fine-tuning, several filtering and selection steps were applied. An initial examination of petition lengths showed that approximately 4000 petitions were shorter than 10 pages, 500 petitions ranged between 10 and 20 pages, and around 300 petitions exceeded 20 pages.

- **Claim Numbering:** Some petitions lacked explicit numbering of claims and instead presented arguments in long, unstructured paragraphs. Since claim numbering provides structural cues that facilitate alignment with summaries, petitions without numbering were deprioritized.
- **Length Harmony:** To ensure proportionality between source and target texts, the summary-to-petition length ratio was examined. Pairs with ratios between 33% and 47% were retained, reflecting realistic summarization behaviour and avoiding extreme length imbalances.

After filtering, the final fine-tuning dataset consisted of 2314 numbered petitions and 500 unnumbered petitions. A limited set of unnumbered samples was intentionally included to expose the model to both structured and unstructured examples, improving generalization. The remaining petitions were held out for evaluation on unseen data.



**Figure 3.** Distribution of petition vs. summary text lengths in the dataset after filtering.

## **2.2. Fine-Tuning Process**

### **2.2.1. Supervised Fine-Tuning (SFT) Setup**

The model was fine-tuned using a supervised, instruction-based approach inspired by instruction following training practices in prior work (Ouyang et al., 2022). Each training instance paired a petition text with its corresponding human-written summary, and the objective was to train the model to produce structured, numbered summaries in passive voice, using formal Turkish legal language while preserving factual content and numerical accuracy.

Early experiments used a single fixed instruction prompt for all samples, which increased prompt overfitting: the model tended to reproduce prompt phrasing rather than internalize the intended summarization behaviour. To reduce this effect, multiple semantically equivalent instruction variants were introduced and randomly assigned across the training set, encouraging the model to generalize the task specification instead of relying on surface-level prompt patterns.

Several prompt-design configurations were evaluated before finalizing the training setup. Early trials using long, inference-style prompts resulted in unstable behaviour, including verbatim text reproduction and overly short or structurally weak summaries. The final configuration employed ten syntactically distinct, but semantically aligned instruction prompts, each conveying the same objective: generating numbered, passive-voice Turkish summaries while preserving dates, percentages, and other factual details. These prompts were rotated across samples to reduce prompt dependency and improve robustness.

This prompt diversification strategy improved training stability and stylistic consistency, allowing the fine-tuned model to learn task-level summarization behaviour, while fine-grained stylistic control was deferred to inference-time system prompts.

### **2.2.2. Parameter Settings and LoRA Configuration**

Fine-tuning was performed using LoRA applied selectively to the feed-forward layers of the LLaMA-3.3-70B-Instruct model (Hu et al., 2021). LoRA enables parameter-efficient adaptation by inserting trainable low-rank matrices into the model while keeping many pretrained parameters frozen, which substantially reduces memory usage and training compute compared with full fine-tuning (Hu et al., 2021).

The experiments were conducted using the LLaMA-3.3-70B-Instruct model as the base architecture, trained with mixed bfloat16 precision on two NVIDIA H200 GPUs providing a total of 288 GB memory. The optimizer was 8-bit Paged AdamW with hyperparameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$  and  $\epsilon = 1 \times 10^{-8}$ , and the learning rate was set to  $1 \times 10^{-6}$  with a cosine scheduler and a warm-up ratio of 0.3%. The batch size per device was 1 with a gradient accumulation step of 4, while gradient checkpointing was enabled to reduce memory usage. Weight decay was fixed at 0.0001, and LoRA was applied using a LoRA rank  $r = 16$ , LoRA  $\alpha$  of 32, and a dropout rate of 0.05. Training was performed for a single epoch, and no evaluation strategy

was integrated into the training loop; instead, evaluation was conducted separately after fine-tuning.

Regarding numerical precision, the model was trained using bfloat16 (bf16) rather than standard float16 (fp16) precision. Although both formats halve memory consumption compared to full precision, bf16 offers a substantially wider dynamic range while retaining the same 16-bit storage footprint, which helps prevent gradient underflow and instability in large-model training, particularly under very low learning rates and long sequence lengths. Since NVIDIA H200 GPUs natively support bf16 operations, this choice enabled numerically stable optimization and removed the need for loss-scaling techniques typically required by fp16. Overall, this configuration achieved a favourable balance between stability and efficiency: mixed precision, gradient accumulation, and gradient checkpointing together allowed reliable fine-tuning of long petition texts without exceeding the available GPU memory.

### **2.2.3. Layer Selection for LoRA Adaptation**

Initial experiments applied LoRA to all transformer layers, including attention and feed-forward components. While this increased adaptation capacity, it led to overfitting: the model began echoing long input segments instead of producing compact summaries. This indicated that modifying attention projections caused it to over-focus on surface lexical details rather than abstract structure. LoRA was therefore restricted to the feed-forward (FFN) layers (gate, up, and down projections). The attention mechanisms already modelled contextual relations effectively; the summarization task instead required refinement in semantic abstraction and reformulation. Adapting the FFN layers strengthened the model's ability to compress, paraphrase, and structurally organize content while preserving factual accuracy. This selective tuning reduced overfitting and improved generalization: summaries became more concise, consistently numbered, and aligned with legal-administrative style. In short, FFN-only LoRA enhanced semantic reasoning and controlled linguistic realization, achieving a balanced trade-off between precision and fluent summarization.

### **2.2.4. Model Capacity and Token Limit Testing**

Comprehensive testing of model input capacity indicated a pronounced decline in performance once petition texts exceeded approximately 16,000 tokens. While longer sequences were technically accepted by the tokenizer and inference pipeline, the model's effective attention span deteriorated, leading to partial omission of later content or a breakdown in logical coherence across sections. On the output side, generation stability was maintained up to about 4,000 tokens. Beyond this point, completions exhibited structural drift, repetition, or premature truncation, suggesting a saturation of the model's active context window. To mitigate these limitations, the inference process was adapted into a hierarchical structure. Long petitions were first segmented into semantically self-contained sections, each treated as an independent summarization unit. The resulting partial summaries were then recursively merged through a secondary synthesis stage, ensuring that both local detail and global continuity were preserved.

This multi-pass summarization strategy enabled complete document coverage without exceeding the model’s operational token and attention constraints.

### 2.2.5. Base Model Comparison

A comparative evaluation was conducted prior to fine-tuning to identify the most suitable base architecture. The evaluated candidates included GPT-OSS-20B (OpenAI, 2025), LLaMA-3.1-70B-Instruct and LLaMA-3.3-70B-Instruct (Meta AI, 2024), Aya-32B-Expanse (Dang et al., 2024), and Qwen3 variants (Qwen3-235B-A22B, Qwen3-30B-A3B, and Qwen3-32B) (Yang et al., 2025). All models were tested under identical conditions using the same petition sample and the same summarization prompt. Native Turkish-speaking legal experts evaluated the outputs in terms of factual completeness, clarity of argumentation, and structural coherence. During the evaluation, it was observed that model size strongly influenced generation length. Smaller models such as GPT-OSS-20B, Aya-32B-Expanse, Qwen3-30B, and Qwen3-32B consistently produced summaries that were disproportionately short relative to the source texts. In contrast, LLaMA-3.1-70B-Instruct, LLaMA-3.3-70B-Instruct, and Qwen3-235B-A22B generated summaries with sufficiently proportional length and preserved the overall argumentative structure of the petitions. Among these candidates, Qwen3-235B-A22B demonstrated the highest factual accuracy and generation stability; however, its computational requirements exceeded the available infrastructure. The LLaMA models provided comparably strong outputs, though both LLaMA-3.1-70B-Instruct and LLaMA-3.3-70B-Instruct exhibited a recurring issue: the insertion of foreign-language loanwords—particularly toward the end of the summaries—despite explicit Turkish-only prompting. Considering overall output quality, linguistic controllability, and resource feasibility, LLaMA-3.3-70B-Instruct was selected as the base model for fine-tuning. A quantitative comparison between the selected base model and its domain-adapted fine-tuned variant is presented in Table 1 below.

**Table 1.** Quantitative comparison between the base pretrained model (LLaMA-3.3-70B-Instruct) and the domain-adapted fine-tuned variant under identical inference and decoding settings on a held-out evaluation set of administrative petitions.

| Metric         | Base Model Mean | Base Std | Fine-Tuned Model Mean | Fine-Tuned Std | $\Delta$ (Fine-Tuned Minus Base) | Relative Improvement (%) |
|----------------|-----------------|----------|-----------------------|----------------|----------------------------------|--------------------------|
| BERT Precision | 0.6672          | 0.3758   | 0.7010                | 0.0903         | +0.0338                          | +5.07%                   |
| BERT Recall    | 0.6719          | 0.0899   | 0.6974                | 0.0955         | +0.0255                          | +3.79%                   |
| BERT F1        | 0.6688          | 0.0834   | 0.6976                | 0.0898         | +0.0288                          | +4.31%                   |
| ROUGE-1 F1     | 0.5394          | 0.1380   | 0.5779                | 0.1386         | +0.0385                          | +7.14%                   |
| ROUGE-2 F1     | 0.3629          | 0.1390   | 0.3991                | 0.1404         | +0.0362                          | +9.97%                   |
| ROUGE-L F1     | 0.3758          | 0.1422   | 0.4090                | 0.1363         | +0.0332                          | +8.84%                   |

Note: M = mean; SD = standard deviation. Relative improvement is computed as  $(M_{\text{fine-tuned}} - M_{\text{base}}) / M_{\text{base}} \times 100$ . Positive  $\Delta$  values indicate improvement.

## 2.3. Handling Long Petition Texts

One of the primary challenges encountered during this study was the summarization of long petitions, particularly those exceeding eight to ten pages. Beyond this length, the models were unable to generate outputs longer than approximately 4000 tokens, regardless of the input size or prompt configuration. This inherent limitation prevented the models from capturing all claims and contextual details within a single generation, often leading to incomplete or overly compressed summaries. To overcome this constraint, a degree of human intervention was initially required to manually divide the petitions into smaller, processable segments. Building upon this necessity, several automated strategies were subsequently developed and tested to minimize manual effort while maintaining summary completeness and coherence. The following subsections present these methods in detail, outlining their underlying logic, implementation steps, and empirical outcomes.

### 2.3.1. Fixed Length Segmentation and Merging Attempts

The initial approach for handling long petitions used fixed-length segmentation, where each document was divided into consecutive parts based on a character or word threshold. Each segment was summarized independently to ensure full coverage within the model's context window. In a subsequent step, the model was prompted to merge the partial summaries into a single coherent output, as shown in Algorithm 1. However, instead of aggregating the content, the model consistently treated the merge instruction as a new summarization task, producing an output comparable in length to a single partial summary and discarding information from earlier segments. For instance, a multi-page petition split into three segments yielded three detailed partial summaries, but the merged output retained only a fraction of the total claim content. While the segmentation stage itself was reliable, the merging process failed to preserve proportional detail and claim coverage. Consequently, this strategy was deemed unsuitable for long-document summarization.

---

#### **Algorithm 1:** Fixed-Length Segmentation and Merging Procedure

---

Require : petition\_text, segment\_size

---

Step-1 : segments  $\leftarrow$  SPLITBYLENGTH(petition\_text, segment\_size)

Step-2 : partial\_summaries  $\leftarrow$  []

Step-3 : for seg  $\in$  segments do

Step-4 : s  $\leftarrow$  LLM\_SUMMARIZE(seg)

Step-5 : APPEND(partial\_summaries, s)

Step 6 : end for

Step 7 : merged\_summary  $\leftarrow$  LLM\_MERGESUMMARIES(partial\_summaries)

Step 8 : return merged\_summary

### 2.3.2. Context-Preserving Sequential Summarization

To improve coherence across long documents, a sequential summarization strategy was applied at Algorithm 2 in which each segment was summarized together with contextual information from the previous segment. Specifically, the model received both the final portion of the preceding text and its generated summary, allowing it to maintain continuity of arguments and narrative flow. This process was repeated iteratively for all segments, resulting in summaries that were more cohesive and less fragmented than those produced by isolated segmentation. However, as contextual information accumulated across iterations, GPU memory usage increased substantially, frequently leading to Out-of-Memory (OOM) errors for long petitions. As a result, despite improved coherence, this approach was not operationally stable or scalable for large-scale summarization.

---

#### **Algorithm 2:** Context-Preserving Sequential Summarization

---

Require : petition\_text, segment\_size, context\_overlap = 500

---

Step-1 : segments  $\leftarrow$  SPLITBYLENGTH(petition\_text, segment\_size)

Step-2 : summaries  $\leftarrow$  []

Step-3 : prev\_summary  $\leftarrow$  ""

Step-4 : prev\_tail  $\leftarrow$  ""

Step-5 : for  $i \leftarrow 1$  to |segments| do

Step 6 : current  $\leftarrow$  segments[i]

Step 7 : if  $i > 1$  then

Step 8 : context  $\leftarrow$  CONCAT(prev\_tail, prev\_summary, current)

Step 9 : else

Step 10 : context  $\leftarrow$  current

Step 11 : end if

Step 12 : summary  $\leftarrow$  LLM\_SUMMARIZE(context)

---

```

Step 13 : APPEND(summaries, summary)

Step 14 : prev_summary ← summary

Step 15 : prev_tail ← GETTAIL(current, context_overlap)

Step 16 : end for

Step 17 : final_summary ← LLM_MERGE(summaries)

Step 18 : return final_summary

```

---

### 2.3.3. Character-Position Claim Segmentation

This experiment attempted to segment petitions based on the approximate character positions of individual claims rather than using fixed-length boundaries. The primary objective was to identify where one claim ended and the next began within the raw text, thereby allowing the model to generate precise and semantically meaningful split indices. To achieve this, a specialized prompt was designed instructing the model to locate the approximate start and end positions of each distinct claim within the petition. The model was expected to return outputs in the following structured form:

- **Claim 1** → characters 136–1121;
- **Claim 2** → characters 1121–2530;
- **Claim 3** → characters 2530–6390.

Using these character ranges, each claim section was extracted from the original petition text and summarized independently. The individual summaries were then combined sequentially to form a complete, claim-by-claim summary of the petition. For instance, in one petition concerning a contract award, the model produced the following segmentation:

- **Claim 1:** Objection to the rejection of the bidder's offer (characters 220–1410).
- **Claim 2:** Argument that the awarded bidder's proposal was abnormally low (characters 1410–2890).
- **Claim 3:** Assertion that evaluation criteria were not applied consistently (characters 2890–4870).

Although this approach was conceptually sound, the practical results were highly inconsistent. Different models, and even multiple runs of the same model, returned divergent character boundaries for identical petitions. In many cases, the identified indices did not align with actual claim boundaries or fell within the middle of a sentence, leading to fragmented and incoherent sections. This instability made it impossible to reproduce the segmentation reliably. Consequently, despite its structural elegance, as illustrated in Algorithm 3, the character-

position method was deemed unsuitable for systematic summarization, as it failed to provide consistent or interpretable split locations across models and runs.

---

**Algorithm 3:** Character-Position Based Claim Segmentation and Summarization

---

Require : petition\_text

---

Step-1 : boundaries  $\leftarrow$  LLM\_DETECTCLAIMBOUNDARIES(petition\_text)

Step-2 : summaries  $\leftarrow$  []

Step-3 : for all (start, end) in boundaries do

Step-4 : claim\_text  $\leftarrow$  petition\_text[start : end]

Step-5 : claim\_summary  $\leftarrow$  LLM\_SUMMARIZECLAIM(claim\_text)

Step 6 : APPEND(summaries, claim\_summary)

Step 7 : end for

Step 8 : final\_summary  $\leftarrow$  LLM\_MERGE(summaries)

Step 9 : return final\_summary

---

### 2.3.4. Main Subclaim Decomposition

This experiment investigated a structured approach in which the model was instructed to identify the internal argumentative hierarchy of each petition by extracting main claims and their associated subclaims in a single pass. Each main claim, together with its supporting subclaims, was treated as a coherent semantic unit for subsequent summarization, and unit-level summaries were later merged in the original document order. For instance, in one petition the model produced the following hierarchical outline:

- **Main Claim 1:** Objection to the tender decision due to the rejection of the bidder's offer.
  - **Subclaim 1.1:** The bidder's proposal should have been considered the most economically advantageous offer.
  - **Subclaim 1.2:** The awarded company submitted an abnormally low bid that should have been disqualified.
  - **Subclaim 1.3:** The tender commission applied the evaluation criteria inconsistently with the relevant procurement regulation.
- **Main Claim 2:** The tender specifications were not applied uniformly to all participants.
  - **Subclaim 2.1:** Certain bidders were permitted to correct documentation errors, while others were not.

While this approach produced summaries with a clear logical structure, it exhibited significant variability in the number and granularity of extracted claims across petitions and runs. In some cases, redundant subclaims inflated summary length, while in others overly coarse claims omitted procedural details. This instability limited controllability, and the approach was therefore deemed unsuitable for large-scale or production use despite its conceptual advantages, exhibited in Algorithm 4.

---

**Algorithm 4:** Main- Subclaim Decomposition and Per-Unit Summarization

---

Require : petition\_text

---

Step-1 : structure  $\leftarrow$  LLM\_EXTRACTHIERARCHY(petition\_text)

Step-2 : summaries  $\leftarrow$  []

Step-3 : for all (mc, sc\_list) in structure do

Step-4 : claim\_unit  $\leftarrow$  BUILDUNIT(petition\_text, mc, sc\_list)

Step-5 : unit\_summary  $\leftarrow$  LLM\_SUMMARIZEUNIT(claim\_unit)

Step 6 : APPEND(summaries, unit\_summary)

Step 7 : end for

Step 8 : final\_summary  $\leftarrow$  LLM\_MERGE(summaries)

Step 9 : return final\_summary

---

### 2.3.5. Semantic Chunking – Final and Stable Approach

The final and most successful method was semantic chunking, which divided petitions into meaning-based sections rather than fixed lengths or indexed positions. The goal was to preserve the logical flow and internal consistency of each claim group while keeping the model's input size within its processing limit. In this method, the text was first analysed as a whole, and then divided into semantically coherent sections (e.g., Claim Group-1, Claim Group-2, . . . ). Each section was summarized independently, and the resulting partial summaries were merged sequentially to form the complete output. This approach yielded the most consistent and stable results for several reasons:

- Claim integrity and logical order were preserved,
- Information loss was minimized across parts,
- The process remained within the model's effective limits (approximately 16,000 tokens input, and 4000 tokens output),
- No OOM issues were observed.

Extensive internal testing confirmed that semantic chunking produced the most balanced and reliable long summaries. Consequently, this approach was adopted in Algorithm 5 as the standard procedure for handling petitions beyond the model's natural context window.

---

**Algorithm 5:** Semantic Chunking and LLM-Based Summarization

---

Require: petition\_text, max\_words = 5000, flex\_window = 1000

---

```

Step-1 : words ← TOKENIZETOWORDS(petition_text)

Step-2 : i ← 0

Step-3 : chunks ← []

Step-4 : while i < |words| do

Step-5 : base_end ← min(i + max_words, |words|)

Step 6 : win_start ← max(0, base_end - flex_window)

Step 7 : win_end ← min(|words|, base_end + flex_window)

Step 8 : candidate ← JOIN(words[win_start : win_end])

Step 9 : split_info ← LLM_SUGGESTSPLIT(candidate)

Step 10 : local_split ← MAPSENTENCESTOWORDINDEX(candidate, split_info)

Step 11 : split_idx ← win_start + local_split

Step 12 : chunk ← JOIN(words[i : split_idx])

Step 13 : APPEND(chunks, chunk)

Step 14 : i ← split_idx

Step 15 : end while

Step 16 : summaries ← []

Step 17 : for all c in chunks do

Step 18 : s ← LLM_SUMMARIZE(c)

Step 19 : APPEND(summaries, s)

Step 20 : end for

```

Step 21 : final\_summary  $\leftarrow$  LLM\_MERGE(summaries)

Step 22 : return final\_summary

### 2.3.6. Summary of Findings

Overall, the experiments demonstrated that the model's summarization capability does not scale proportionally with input length. Beyond a certain threshold, longer inputs consistently produced compressed or incomplete summaries, regardless of model size or configuration. Fixed-length segmentation generated usable partial summaries but failed during the merging stage, where the model re-summarized rather than concatenated content. The context-preserving sequential method improved coherence but suffered from rapid memory accumulation and OOM errors. Structural decomposition techniques, such as character-position segmentation and main subclaim extraction, introduced interpretability but proved unstable due to inconsistent boundary detection and variable claim granularity. In contrast, the semantic chunking approach addressed these limitations collectively by aligning split points with semantic boundaries instead of fixed lengths or character indices. It produced coherent, detailed, and reproducible summaries across petitions while remaining computationally stable. As a result, semantic chunking was adopted as the standard strategy for handling long petition texts in subsequent stages of the project. A comparison of chunking and summarization strategies is shown in Table 2 below.

**Table 2.** Comparison of chunking and summarization strategies evaluated for petition summarization.

| Method                                      | Core Idea and Implementation                                                                                                                                               | Observed Issues                                                                                      | Outcome / Stability                               |
|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| Fixed-Length Segmentation and Merging       | Petitions were split into equal-length parts processed independently. Each part was summarized separately, and the partial summaries were merged into one complete output. | The model summarized instead of concatenating, producing shorter merged summaries and losing detail. | Conceptually simple but inconsistent; discarded.  |
| Context-Preserving Sequential Summarization | Each segment was summarized using the previous part's final 500 characters and its summary to maintain continuity.                                                         | Improved coherence, but cumulative context caused high GPU memory usage and frequent OOM errors.     | Conceptually strong but computationally unstable. |
| Character-Position Claim Segmentation       | The model was asked to identify the start and end character indices of each claim, and these ranges were summarized separately.                                            | Character boundaries varied across models and runs, producing misaligned and fragmented segments.    | Structurally elegant but non-reproducible.        |

**Table 2.** (Continued)

|                                             |                                                                                                                                                                               |                                                                                                                              |                                                               |
|---------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------|
| Main–<br>Subclaim<br>Decomposition          | The entire petition was analysed in one pass to extract hierarchical pairs of main claims and subclaims. Each main subclaim unit was summarized independently and merged.     | The number and depth of extracted claims varied; over- and under segmentation caused inconsistent lengths and detail levels. | Provided structural insight but unstable for large-scale use. |
| Semantic<br>Chunking<br>(Final<br>Approach) | Petitions were divided into meaning-based chunks within model limits (approx. 16k input / 4k output tokens). Each chunk was summarized independently and merged sequentially. | No significant computational issues; minimal information loss.                                                               | Most balanced and reliable; adopted as standard procedure.    |

Note: OOM = out-of-memory error.

## 2.4. Inference-Time Prompt Design

This section describes the development and refinement of the inference-time prompt used for summarizing long petitions. While the fine-tuning stage focused on learning the general summarization behaviour, inference required a standardized instruction that balanced factual completeness, legal tone, and stylistic consistency across multi-part summaries. The process involved approximately 30 prompt candidates, iteratively refined through evaluation with large-scale models such as ChatGPT, Claude, and internal deployments of LLaMA-3.3-70B-Instruct. The final prompt served as the foundation for all experiments presented in Section 4.

### 2.4.1. Inference Objectives

At inference time, the model was required to generate outputs that conformed to formal Turkish legal administrative style while preserving the full factual and legal content of the petitions. Despite task-aligned fine-tuning, output quality remained sensitive to instruction phrasing for long, multi-claim inputs. Accordingly, the inference objectives were defined along five axes:

- **Factual and Legal Completeness:** Every factual claim, legislative citation, specification article, date, percentage, monetary value, annex reference (e.g., Ek-1, Ek-2), and procedural justification had to be retained. No evidentiary or contextual element was to be omitted or paraphrased to the point of losing meaning.
- **Structural Uniformity:** The summary had to begin with the exact phrase “Başvuru sahibinin dilekçesinde özetle;” and end with a single formal closing line. Claims were to be presented as a single continuous numbered list (1), 2), 3), ... without restarting, with each claim occupying exactly one paragraph.

- **Passive-Voice Enforcement:** The entire text was required to use passive constructions consistent with official administrative style. Each numbered paragraph ended with a neutral passive clause (e.g., “... uygun olmadığı,” “... geçersiz olduğu,” “... tespit edildiği,” “... değerlendirme dışı bırakılması gerektiği”). Endings such as “iddia edilmektedir” or “verilmektedir” were disallowed within the claim paragraphs.
- **Linguistic and Stylistic Fidelity:** Outputs were expected to reflect the tone and rhythm of official Public Procurement Authority summaries: formal, objective, and free of colloquial or advocative phrasing. Long, multi-clause passive sentences were acceptable so long as supporting facts were integrated clearly and consistently.
- **Context Preservation Within Claims:** When a claim included justifications, evidence, examples, or consequences, these components were to be included within the same numbered paragraph to maintain causal and evidentiary coherence without fragmenting the argument.

These objectives guided the final inference instructions toward a balance of completeness, structural reliability, and stylistic consistency, enabling reproducible summaries across petitions of varying length and complexity. Across approximately thirty iterations, prompt design experiments explored different stylistic and structural formulations. The main design dimensions included:

- **Instruction Framing:** Whether the model was instructed to “summarize” vs. “rewrite intro claims”.
- **Output Structuring:** Numbered vs. unnumbered, single sentence vs. multi-sentence per claim.
- **Stylistic Control:** Passive voice enforcement, tone formality, or explicit mention of Turkish language.
- **Length Regulation:** Inclusion or omission of token, paragraph, or claim count limits.
- **Factual Anchoring:** Explicit emphasis on preserving percentages, dates, and references to legislation.
- **Each Prompt Variant Was Tested on a Representative Subset of Long Petitions.** Native Turkish legal experts qualitatively assessed the outputs for factual accuracy, linguistic formality, and consistency with the official summary style.

#### 2.4.2. Final Prompt and Rationale

The final inference process employed a two-stage prompting strategy designed to preserve factual completeness while enforcing grammatical and stylistic consistency in Turkish legal-administrative language. This approach was introduced after observing that large models often struggled to maintain fully passive and nominalized sentence endings (-dığı/-diği/-duğu/-düğü) when summarizing long legal petitions in a single pass. To address this, inference was separated into a content-generation stage followed by a linguistic normalization stage.

##### 2.4.2.1. Stage 1 – Argumentative Summarization

In the first step, the model was instructed to act as a legal summarization specialist, producing a detailed, argumentative summary of the petition. The goal was to extract all factual claims, legal citations, evidence, and quantitative details, with each numbered paragraph representing a coherent argumentative unit. This stage reliably preserved evidential reasoning and claim structure but required subsequent grammatical standardization to align with passive-voice norms.

#### 2.4.2.2. Stage 2 – Linguistic Normalization

The second step reformatted the Stage 1 output into a legally consistent linguistic form. Here, the model acted as a legal text formatter, transforming the preliminary summary without altering its content. Each claim was rewritten as a single Turkish sentence ending with an appropriate nominalized passive form (-dığı/-diği/-duğu/-düğü), while varied endings (e.g., belirtildiği, tespit edildiği, ifade edildiği, öne sürüldüğü) were encouraged for naturalness and stylistic authenticity. This stage also enforced fixed structural conventions: the standardized opening “Başvuru sahibinin dilekçesinde özetle;”, numbered one-sentence claims, and the closing line “İddia edilmekte ve itirazın şikâyet başvurusunun Kurumumuzca incelenerek gereğinin yapılması istenmektedir.” In operational terms, the pipeline treats Stage 1 as a high-recall extraction step and Stage 2 as a constrained rewriting step with strict invariants. Specifically, Stage 2 is instructed to preserve the original ordering, claim count, and all factual entities (dates, quantities, legal references, and identifiers) while only regularizing syntax and surface form. To minimize error propagation, Stage 2 explicitly prohibits adding, removing, merging, or splitting claims, and it enforces one sentence-per-claim formatting with a fixed opening and closing line. Normalization makes the method robust to long-document effects: even when the input is processed in multiple semantic chunks (Algorithm 5), the normalization step yields a uniform legal style across chunks without sacrificing evidentiary content.

#### 2.4.2.3. Rationale

The two-step procedure was essential for achieving both semantic fidelity and linguistic precision. Stage 1 ensured complete preservation of claims and evidentiary relationships, while Stage 2 guaranteed syntactic regularity and conformity to institutional writing norms. Separating reasoning from formatting mitigated the model’s tendency toward mixed verb forms and yielded summaries that were consistent, grammatically controlled, and legally admissible. This two-stage prompt was applied uniformly to all semantically segmented chunks (Algorithm 5), minimizing stylistic drift across the final output.

#### 2.4.2.4. Illustrative Example

The effectiveness of the approach is demonstrated by comparing the raw argumentative output of Stage 1 with the linguistically normalized version from Stage 2. Firm and personal identifiers have been anonymized for publication.

- **Stage 1 Output — Content-Focused Argumentative Summarization (Anonymized):** *Başvuru sahibinin dilekçesinde özetle;*

- İdari Şartnamenin 7.1(h) maddesine göre isteklilerin teklif kapsamında Gıda Üretim İzni Belgesi veya İşletme Kayıt Belgesi sunmaları gerektiği, ancak Firma A ve Firma B tarafından sunulan belgelerin faaliyet konularının “Hazır Yemek ve Tabldot Yemek Üretimi” kapsamını karşılamadığı ve belgelerde yer alan adres beyanlarının gerçeği yansıtmadığı, bu nedenle söz konusu isteklilerin teklif dosyalarının değerlendirme dışı bırakılması gerektiği ileri sürülmektedir.
  - Firma B’nin işletme kayıt belgesinin Tarım ve Orman Bakanlığının resmi kayıtları ile uyumlu olmadığı, adres bilgilerinin gerçeği yansıtmadığı ve belgenin geçersiz olduğu, bu nedenle ilgili isteklinin teklifinin değerlendirme dışı bırakılması gerektiği belirtilmektedir.
  - İdari Şartnamenin 7.4.1 maddesi uyarınca sunulması gereken banka referans mektuplarının, Firma A ve Firma B tarafından ilan tarihinden önce alındığı ve kredi tutarlarının şartnameye uygun olmadığı değerlendirilmektedir.
  - İdari Şartnamenin 7.5.2 maddesi uyarınca sunulması gereken 1.000 kişilik yemek üretim kapasitesine ilişkin kapasite raporlarının, Firma A ve Firma B tarafından geçerli şekilde sunulmadığı ve adres beyanlarının diğer ihale belgeleri ile örtüşmediği ifade edilmektedir.
  - Sonuç olarak, ihale sürecinde belirtilen aykırılıklar doğrultusunda düzeltici işlem yapılması gerektiği iddia edilmektedir.
- **Stage 2 Output — Passive-Voice Legal Normalization (Anonymized):** *Başvuru sahibinin dilekçesinde özetle;*
- İdari Şartnamenin 7.1(h) maddesi gereği sunulması gereken Gıda Üretim İzni Belgesi veya İşletme Kayıt Belgesi kapsamında, Firma A ve Firma B tarafından sunulan belgelerde “Hazır Yemek ve Tabldot Yemek Üretimi” faaliyet konusunun yer almadığı ve adres beyanlarının gerçeği yansıtmadığı belirtilerek, ilgili isteklilerin teklif dosyalarının değerlendirme dışı bırakılması gerektiği belirtildiği,
  - Firma B tarafından sunulan işletme kayıt belgesinin resmi kayıtlarla uyumlu olmadığı ve belgenin geçersiz olduğu tespit edildiği,
  - İdari Şartnamenin 7.4.1 maddesi uyarınca sunulması gereken banka referans mektuplarının Firma A ve Firma B tarafından ilan tarihinden önce alındığı ve kredi tutarlarının şartname koşullarını karşılamadığı belirtildiği,
  - İdari Şartnamenin 7.5.2 maddesi uyarınca sunulması gereken 1000 kişilik yemek üretim kapasitesine ilişkin kapasite raporlarının Firma A ve Firma B tarafından geçersiz sunulduğu ve adres bilgilerinin ihale belgeleri ile örtüşmediği vurgulandığı,
  - İddia edilmekte ve itirazen şikâyet başvurusunun Kurumumuzca incelenerek gereğinin yapılması istenmektedir.

### 2.4.3. Empirical Evaluation of the Two-Stage Inference Design

To justify the proposed two-stage inference strategy, a comparative evaluation was conducted against a single-pass summarization baseline on 25 administrative petitions selected to capture structural and linguistic variability. Both approaches were applied under identical decoding settings and evaluated in collaboration with domain experts. Quantitative results are reported in Table 3, comparing single-pass (Stage 1) and two-stage (Stage 2) outputs. The two-stage approach consistently outperforms the baseline across all automatic metrics, including BERT Score and ROUGE-1/2/L (F1), indicating improved semantic alignment with reference summaries while preserving lexical and structural consistency. Error-type analysis revealed that single-pass summaries frequently exhibited non-compliant sentence endings and occasional non-Turkish lexical forms. These issues were consistently corrected in the second stage through linguistic normalization enforcing standardized passive constructions and Turkish lexical usage. Finally, blind pairwise human evaluations by domain experts showed a clear preference for the two-stage summaries in all 25 cases, citing higher legal accuracy, greater compactness, and closer alignment with expert-written administrative summaries. Overall, the results in Table 3 provide empirical justification for the proposed two-stage inference design.

**Table 3.** Quantitative comparison between single-pass (Stage 1) and two-stage (Stage 2) summarization.

| Metric         | Stage 1 Mean | Stage 1 Std | Stage 2 Mean | Stage 2 Std |
|----------------|--------------|-------------|--------------|-------------|
| BERT Precision | 0.6672       | 0.0822      | 0.7010       | 0.0903      |
| BERT Recall    | 0.6719       | 0.0899      | 0.6974       | 0.0955      |
| BERT F1        | 0.6688       | 0.0834      | 0.6976       | 0.0898      |
| ROUGE-1 F1     | 0.5349       | 0.1380      | 0.5779       | 0.1386      |
| ROUGE-2 F1     | 0.3629       | 0.1390      | 0.3991       | 0.1404      |
| ROUGE-L F1     | 0.3758       | 0.1422      | 0.4090       | 0.1363      |

Note: Stage 1 corresponds to single-pass summarization, while Stage 2 corresponds to the proposed two-stage summarization pipeline. Values are reported as mean and standard deviation across the evaluation set.

### 2.4.4. Common Failure Patterns and Lessons Learned

During prompt evaluation, several recurring failure types were observed:

- **Under-Summarization:** Some prompts caused the model to compress multi-claim petitions into overly brief outputs.
- **Over-Generation:** Others led to redundant restatements or the introduction of fabricated subclaims.

- **Inconsistent Formatting:** Minor template variations occasionally resulted in missing numbering or mixed sentence structures.
- **Style Drift:** Without strict phrasing, some models reverted to active-voice constructions or informal tone.

Through iterative refinement and domain-expert review, these issues were systematically reduced. The final inference prompt reflects an empirically optimized balance between content fidelity, structural consistency, and linguistic control. While earlier configurations produced acceptable results for some petitions, achieving stable consistency across all document types required multiple cycles of prompt adjustment and evaluation.

## **2.5. Human-in-the-Loop Evaluation**

In addition to automated consistency checks, a human-in-the-loop evaluation process was integrated to ensure legal and linguistic adequacy of the generated summaries. A panel of procurement law specialists and linguistic reviewers examined approximately ten petition outputs produced under different inference configurations. These sessions provided continuous expert feedback on the model's handling of sentence endings, Turkish grammatical correctness, claim separation, and the overall conformity of tone to formal administrative language. The experts' feedback was systematically incorporated into prompt and model refinements across several iterations. Attention was given to achieving natural variation in passive endings, accurate reflection of legal claims, and balanced sentence structure without loss of factual completeness. Through this cycle, both the model's inference behaviour and the prompt design evolved toward higher linguistic precision and legal reliability. Ultimately, this collaborative process yielded a configuration that consistently met expert expectations for clarity, structural consistency, and adherence to institutional writing standards. It also established a feedback-driven evaluation loop that can be reused for future fine-tuning and validation cycles.

### **2.5.1. Data and Code Availability**

The dataset utilized in this study consists of confidential petition texts that are not publicly available due to institutional data protection requirements. These documents contain legally sensitive information and are subject to internal confidentiality policies, which prohibit public release or unrestricted distribution. For this reason, the fine-tuned model, training corpus, preprocessing scripts, and intermediate training artifacts derived from the dataset cannot be shared externally.

### **2.5.2. Ethical, Privacy and Data Governance Considerations**

Given the sensitive nature of administrative and legal petitions, careful ethical and privacy-aware data handling practices were applied throughout this study. All documents were obtained through authorized institutional channels and were used exclusively for research purposes. Full anonymization was not applied prior to modelling, as preserving the identities of the parties involved (e.g., contracting authority and bidder sides) was necessary to correctly interpret roles, claims, and contextual relationships required for accurate legal summarization. However, the

data was not modified, redistributed, or exposed beyond the authorized research environment. Access to the dataset was strictly limited to the research team, and all processing was conducted in secured, access-restricted systems. The study relies on retrospective administrative documents and does not involve human subjects or intervention; therefore, formal ethics committee approval was not required under applicable institutional regulations.

### **2.5.3. Main Contributions**

The main contributions of this work can be summarized as follows:

- A domain-specific summarization framework for Turkish public procurement petitions. A complete, end-to-end summarization pipeline is presented, tailored to the structural, linguistic, and legal characteristics of administrative petition texts, addressing challenges such as inconsistent claim structure, long document length, and strict institutional writing requirements.
- A two-stage inference strategy separating legal reasoning from linguistic normalization. A two-stage inference design is introduced and empirically justified in which argumentative content extraction is decoupled from stylistic and grammatical normalization, enabling both factual completeness and strict adherence to Turkish legal-administrative passive-voice conventions.
- A systematic evaluation of long-document handling strategies for petition summarization. A detailed empirical analysis of multiple long-text processing approaches is conducted, including fixed segmentation, sequential context propagation, claim-based decomposition, and semantic chunking, and semantic chunking combined with staged inference is identified as the most stable and scalable solution.
- Domain-informed fine-tuning with controlled prompt diversification. It is demonstrated that supervised fine-tuning with LoRA, combined with prompt diversification rather than fixed instruction templates, improves generalization and reduces prompt overfitting in domain-specific summarization tasks.
- Expert-driven evaluation and error analysis beyond automatic metrics. A human-centred evaluation framework is provided involving domain experts, including quantified comparison, error-type analysis, and blind preference studies, showing that the proposed approach consistently outperforms single-stage summarization in legal admissibility, linguistic accuracy, and stylistic conformity.

## **3. Conclusion**

This study presented a complete, domain-focused summarization framework for legal petition texts, combining dataset preparation, model training, algorithmic design, and expert-driven evaluation. A detailed analysis of petition and summary structures was first conducted, examining length distributions, variability in linguistic clarity, and the role of numbering and legal reference patterns. Based on these observations, a dataset was curated and filtered that

preserved realistic claim organization while reducing noise and structural ambiguity. On the model development side, supervised fine-tuning (SFT) was employed on LLaMA-3.3-70B-Instruct using parameter-efficient LoRA adaptation, supported by carefully diversified system prompts to prevent memorized prompting behaviour and ensure stylistic generalization. In parallel, multiple strategies were developed and tested for handling long petition texts, including fixed segmentation, sequential summarization with continuity signals, hierarchy-based claim decomposition, and character-position boundary heuristics. Through iterative evaluation, semantic chunking combined with a two-stage inference prompt emerged as the most stable approach, providing coherent claim separation, consistent passive-voice legal style, and improved factual retention without exceeding the model's operational context limits. Finally, legal experts reviewed model outputs to validate correctness, tone, and structural alignment with institutional summarization standards. The resulting system demonstrates that domain adaptation, prompt engineering, and structured long document processing can significantly improve the reliability and legal validity of summarization models in real-world procurement contexts. Future work will focus on extending the system beyond supervised summarization toward a preference-aligned, context-stable model. Future work will aim to improve long-context generation within a single inference window through attention-extension and continuity-aware training, reducing reliance on chunking. In addition, Reinforcement Learning from Human Feedback using Proximal Policy Optimization will be applied with a summarization-specific reward model targeting claim completeness, factual accuracy, and legal-administrative passive voice style. Finally, user-configurable summarization profiles will be explored to control compression level, evidential detail, and linguistic formality across institutional contexts.

## **Conflict of Interest**

The authors declare that they have no conflict of interest.

## **Funding**

This research was supported by HAVELSAN / Artificial Intelligence Group Leadership under internal R&D support. No external funding was received.

## **Acknowledgements**

The authors would like to express their sincere gratitude to Dilek BİLİR for her expert review, critical evaluation, and invaluable guidance throughout the study. They also wish to acknowledge Burcu ŞIKLAR and İrem BURGAÇ PEKSAYIN for their detailed feedback and domain-specific assessments, which greatly contributed to enhancing the quality of the summarization output. Finally, the authors extend their appreciation to Necmettin YÜKSEL, Head of the Electronic Procurement Department, and Serdar HALICI, Deputy President of the Institution, for their continued support and constructive contributions during the research and implementation process.

## Ethical Statements

The study involved analysis of petition-style documents and expert evaluation of anonymized model outputs. Any personal or institutional identifiers were removed or masked prior to experimentation and reporting, and results are presented only in aggregated or anonymized form. No human subject experimentation was conducted.

## Data and Code Availability

Due to confidentiality constraints associated with domain-specific petition data and institutional usage, the dataset and trained model artifacts cannot be publicly released. An anonymized sample and/or evaluation templates may be made available from the corresponding author upon reasonable request, subject to approval and data-sharing policies.

## Supplementary Materials

Supplementary materials include additional prompt variants, example anonymized outputs for Stage 1 and Stage 2, and extended qualitative evaluation notes. These materials can be provided upon request.

## References

- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.
- Chen, Y., Ma, Y., Mao, X., & Li, Q. (2019). Multi-task learning for abstractive and extractive summarization. *Data Science and Engineering*, 4(1), 14–23. <https://doi.org/10.1007/s41019-019-0087-7>
- Dang, J., Singh, S., D'souza, D., Ahmadian, A., Salamanca, A., Smith, M., Peppin, A., Hong, S., Govindassamy, M., Zhao, T., Kublik, S., Amer, M., Aryabumi, V., Campos, J. A., Tan, Y.-C., Kocmi, T., Strub, F., Grinsztajn, N., Flet-Berliac, Y., ... Hooker, S. (2024). *Aya Expanse: Combining research breakthroughs for a new multilingual frontier*. arXiv. <https://doi.org/10.48550/arXiv.2412.04261>
- Erkan, G., & Radev, D. R. (2004). LexRank: Graph-based lexical centrality as salience in text summarization. *Journal of Artificial Intelligence Research*, 22, 457–479. <https://doi.org/10.1613/jair.1523>
- Guan, C., Chin, A., & Vahabi, P. (2024). *Enhancing news summarization with eLearnFit through efficient in-context learning and efficient fine-tuning*. arXiv. <https://doi.org/10.48550/arXiv.2405.02710>
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2021). LoRA of LLMs. arXiv. <https://doi.org/10.48550/arXiv.2106.09685>

- Huang, L., Cao, S., Parulian, N., Ji, H., & Wang, L. (2021). Efficient attentions for long document summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 1419–1436). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.112>
- Liu, D., & Demberg, V. (2024). RST-LoRA: A discourse-aware LoRA for long document abstractive summarization. arXiv. <https://doi.org/10.48550/arXiv.2405.00657>
- Liu, Y. (2019). Fine-tune BERT for extractive summarization. arXiv. <https://doi.org/10.48550/arXiv.1903.10318>
- Meta AI. (2024). The LLaMA 3 herd of models. arXiv. <https://doi.org/10.48550/arXiv.2407.21783>
- OpenAI. (2025). gpt-oss-120b & gpt-oss-20b model card. arXiv. <https://doi.org/10.48550/arXiv.2508.10925>
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. arXiv. <https://doi.org/10.48550/arXiv.2203.02155>
- Qin, H., Ma, X., Zheng, X., Li, X., Zhang, Y., Liu, S., Luo, J., Liu, X., & Magno, M. (2024). Accurate LoRA-finetuning quantization of LLMs via information retention. arXiv. <https://doi.org/10.48550/arXiv.2402.05445>
- Rallapalli, S., Gallagher, S., Mellinger, A. O., Ratchford, J., Sinha, A., Brooks, T., & Brown, B. (2025). Fine-tuning LLMs for report summarization: Analysis on supervised and unsupervised data. arXiv. <https://doi.org/10.48550/arXiv.2503.10676>
- See, A., Liu, P. J., & Manning, C. D. (2017). Get to the point: Summarization with pointer-generator networks. arXiv. <https://doi.org/10.48550/arXiv.1704.04368>
- Xiao, W., & Carenini, G. (2019). Extractive summarization of long documents by combining global and local context. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)* (pp. 3011–3021). <https://doi.org/10.18653/v1/D19-1298>
- Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., Zheng, C., Liu, D., Zhou, F., Huang, F., Hu, F., Ge, H., Wei, H., Lin, H., Tang, J., ... Qiu, Z. (2025). Qwen3 technical report. arXiv. <https://doi.org/10.48550/arXiv.2505.09388>
- Zhang, Y., Ni, A., Mao, Z., Wu, C.-H., Zhu, C., Deb, B., & Zhang, R. (2022). SummN: A multi-stage summarization framework for long input dialogues and documents. In

*Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 1592–1604). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.112>

Zhao, J., Wang, T., Abid, W., Angus, G., Garg, A., Kinnison, J., Sherstinsky, A., Molino, P., Addair, T., & Rishi, D. (2024). *LoRA Land: 310 fine-tuned LLMs that rival GPT-4, a technical report*. arXiv. <https://doi.org/10.48550/arXiv.2405.00732>

# Design and Experimental Validation of an EMI Filter for Military DC-DC Converters Achieving MIL-STD-461G CE102 Compliance

İsmail Ovalı<sup>1,\*</sup> , İres İskender<sup>2</sup> 

## Abstract

This paper presents the design, implementation, and experimental validation of a practical Electromagnetic Interference (EMI) filter developed for military DC-DC converter applications in order to satisfy the conducted emission limits defined by MIL-STD-461G CE102. High-frequency switching behaviour, fast voltage and current transitions, and unavoidable parasitic elements in DC-DC converters make compliance with these limits extremely challenging without dedicated filtering, while practical military platforms impose severe constraints on size, stability, and realizability. In this paper, a systematic and iterative design methodology based on established EMI filter design principles is proposed, and the resulting attenuation performance is experimentally validated through CE102 measurements. The procedure includes determination of the design resistance considering the negative incremental impedance of the converter, design of Differential-Mode (DM) and Common-Mode (CM) stages, and suppression of resonance through properly selected damping networks. The resulting filter is implemented using commercially available components on a compact Printed Circuit Board (PCB) compatible with practical military installations. The performance of the proposed filter is verified experimentally through CE102 measurements performed under identical operating conditions for both filtered and unfiltered cases. The unfiltered converter is shown to exceed the CE102 limits, whereas the filtered configuration achieves compliance with sufficient margin across the entire frequency range. In addition, frequency-domain attenuation characteristics are measured using a Bode analyser to reveal the insertion loss behaviour of the filter. The consistency between analytical expectations, attenuation measurements, and compliance tests demonstrates that the proposed approach provides a reliable, repeatable, and application-oriented framework for EMI filter design in military power electronic systems.

**Keywords:** EMI, EMI filter, DC-DC converter, MIL-STD-461G CE102, insertion loss.

---

Received: 19.02.2026

Accepted: 22.05.2026

Research Article

<sup>1</sup> OB Sensor Technologies, Ankara, Türkiye

<sup>2</sup> Çankaya University, Ankara, Türkiye

\* Corresponding author.

✉ ovaliismail@gmail.com

## Cite this article as:

<https://doi.org/10.65834/jdsi.12.64>

Ovalı, İ., & İskender, İ. (2026). Design and Experimental Validation of an EMI Filter for Military DC-DC Converters Achieving MIL-STD-461G CE102 Compliance. *Journal of Defence and Security Industries: Strategy and Technology 1(2)*, 157-176.

## 1. Introduction

Military electronic systems are required to operate reliably under stringent Electromagnetic Compatibility (EMC) constraints due to dense integration, harsh operating environments, and mission-critical functionality. Power electronic converters used in such platforms are among the primary sources of conducted Electromagnetic Interference (EMI), which may propagate through power lines and adversely affect neighbouring subsystems (Williams, 2001). To ensure compatibility, military standards impose strict emission limits; among them, MIL-STD-461G defines the CE102 requirement for power line conducted emissions over a wide frequency range (U.S. Department of Defence, 2015).

Modern DC-DC converters inherently generate significant conducted EMI due to high switching frequencies, fast voltage and current transitions, and unavoidable parasitic elements (Hartal, 1995). In practical military applications, compliance with CE102 limits is rarely achievable without dedicated EMI filtering. However, EMI filter design remains challenging because attenuation requirements must be satisfied while preserving system stability, minimizing volume, and maintaining practical realizability.

In the recent power electronics literature, there is a significant shift toward Active EMI Filters (AEFs) and Hybrid EMI Filters (HEFs) to address the volumetric constraints of passive components. Fishta and Fiori (2026) utilized active capacitors to achieve a 50% reduction in magnetic volume for traction inverters, while Tian et al. (2026) proposed cascaded active structures to enhance CM attenuation for wide-bandgap drives. Although these methodologies offer substantial miniaturization, they introduce significant control complexity, require auxiliary power supplies, and are subject to bandwidth limitations imposed by the gain-bandwidth product of operational amplifiers, often leading to performance degradation above 3 MHz. For mission-critical military platforms, where extreme reliability, structural simplicity, and robust performance under harsh environments are paramount, passive EMI filters remain the preferred and most reliable solution.

However, existing passive EMI filter design procedures in the literature often exhibit specific limitations in practical military applications, particularly regarding size, high-frequency performance, and system stability. Traditional approaches, such as those detailed by Cadirci et al. (2005) and Kichouliya et al. (2023), primarily rely on baseline noise measurements to determine the required insertion loss. While effective for their specific setups, this strictly attenuation-driven approach often leads to sub-optimal designs. For example, to meet a 60 dB low-frequency attenuation target, Kichouliya et al. (2023) employed disproportionately large magnetic components, requiring a massive 21.1 mH CM choke. In contrast, the proposed methodology in this study achieves a superior CM attenuation of 75 dB at 500 kHz using a significantly smaller 390  $\mu$ H choke. Furthermore, conventional filters often suffer from high-frequency degradation; Cadirci et al. (2005) reported significant performance loss and limit violations above 5 MHz due to the unsuppressed self-resonances of the filter components. By deliberately incorporating RC damping networks, the proposed design effectively suppresses

these high-Q resonances, maintaining robust attenuation across the entire CE102 frequency spectrum.

Beyond physical size and resonance issues, many existing studies focus on achieving theoretical insertion loss but completely neglect the dynamic stability of the overall system. Specifically, the interaction between the EMI filter's output impedance and the DC-DC converter's negative incremental input resistance ( $R_n$ ) is frequently overlooked in classical designs. Ignoring this impedance interaction can lead to unforeseen oscillations and severe system instability under varying load conditions. The proposed methodology addresses this critical gap by integrating the converter's  $R_n$  into the core of the design process. By establishing a design resistance ( $R_d$ ) that is strictly smaller than the magnitude of the negative input resistance ( $R_d \leq |R_n|$ ), the proposed framework guarantees robust stability while simultaneously optimizing the component values. This stability-oriented approach ensures that the filter not only meets stringent military emission limits but also preserves the reliable operation of the converter.

To address these gaps, this study proposes a systematic and iterative passive design methodology that prioritizes both compliance and system stability. Unlike previous works that rely on oversized inductors, the proposed framework integrates the determination of a design resistance ( $R_d$ ) based on the converter's negative incremental input resistance ( $R_n$ ) to guarantee robust dynamic stability. Furthermore, it utilizes RC damping networks to suppress high-Q resonances, optimizing the CM and DM stages to achieve superior attenuation. This structured design flow allows for the realization of the filter using compact, Commercially Available Off-The-Shelf (COTS) components within a size envelope suitable for military platforms. In order to define realistic design targets, publicly available datasheet information of representative military-grade DC-DC converters and their associated EMI filters was reviewed. Observations indicating that dominant emissions typically concentrate around switching frequencies near 250 kHz and their harmonics guided the identification of the most critical emission frequencies.

In this study, by bridging the gap between theoretical filter principles and the practical constraints of military DC-DC converters, the effectiveness of the proposed filter is evaluated through both simulation and rigorous experimental investigations. Conducted emission measurements are performed according to the MIL-STD-461G CE102 setup under realistic operating conditions to assess compliance. In addition, frequency-domain insertion loss characteristics are obtained using a Bode analyser, enabling a detailed physical observation of the filter's attenuation behaviour.

The primary contributions of this study are summarized as follows:

- A systematic and practical EMI filter design methodology tailored to MIL-STD-461G CE102 requirements, explicitly incorporating dynamic stability ( $R_n$ ) and resonance damping.

- Comprehensive experimental characterization of the filter's attenuation capability through both standard CE102 compliance measurements and Bode analyser-based frequency response analysis.
- Demonstration of direct applicability to real military systems through a size-constrained hardware implementation and successful validation on a high-frequency DC-DC converter.

## 2. Design Methodology

This section presents the proposed EMI filter design methodology for a military DC-DC converter targeting compliance with the MIL-STD-461G CE102 conducted emission limits. The procedure follows a systematic and iterative framework derived from established EMI filter design principles (Nave, 1991; Ozenbaugh & Pullen, 2012) while accounting for practical implementation constraints.

To establish realistic design expectations, publicly available datasheet information from representative military-grade DC-DC converters and EMI filters was reviewed. This reference provides practical guidance regarding dominant frequency regions and achievable attenuation levels.

### 2.1. Background on EMC and MIL-STD-461G CE102

EMC is defined as the ability of an electronic system to function correctly in its electromagnetic environment without introducing intolerable EMI to other equipment. In Switch-Mode Power Supplies (SMPS) like military DC-DC converters, EMI is inherently generated by the high  $dv/dt$  and  $di/dt$  characteristics of semiconductor switching actions. This noise propagates through conduction along power lines (conducted emissions) or through radiation (radiated emissions).

To ensure reliable operation in dense military platforms, the MIL-STD-461G standard establishes strict testing procedures and emission limits. Specifically, the CE102 requirement evaluates conducted emissions on power leads within the frequency range of 10 kHz to 10 MHz. The measurement methodology requires the use of Line Impedance Stabilization Networks (LISNs), typically standardized at 50  $\mu$ H, placed between the power source and the Equipment Under Test (EUT). The LISNs serve a dual purpose: they provide a defined, stable Radio-Frequency (RF) impedance to the EUT across the measurement band, and they isolate the test setup from ambient noise present on the facility's power grid. Compliance with CE102 guarantees that the high-frequency switching noise of the DC-DC converter will not degrade the performance of adjacent sensitive military electronics sharing the same power distribution bus.

### 2.2. Design Methodology Overview

The proposed approach consists of the following sequential steps:

- Determination of the design resistance from converter input characteristics,
- Design of the DM filter,
- Suppression of DM resonances through damping networks,
- Design of the CM filter,
- Suppression of CM resonances, and
- Integration and simulation-based verification of the EMI filter design.

This structured process enables predictable attenuation performance while maintaining system stability and realizability.

### 2.3. Determination of Design Resistance

Switch-mode DC-DC converters exhibit a negative incremental input impedance due to constant power behaviour. It can be approximated as

$$R_n = - \left| \frac{V_{in}}{I_{in}} \right| \quad (1)$$

where  $V_{in}$  and  $I_{in}$  denote the input voltage and current.

For stable interaction, the output impedance of the EMI filter must remain significantly smaller than the magnitude of the converter input impedance (Middlebrook, 1976):

$$|Z_o| \leq |R_n| \quad (2)$$

Since the filter design resistance  $R_d$  equals the filter's output impedance  $Z_o$  at a quality factor  $Q = 1$ ,  $R_d$  must be chosen to be smaller than the magnitude of the negative incremental input resistance,  $|R_n|$ .

$$R_d \leq |R_n| \quad (3)$$

To obtain a realistic reference, a representative military-grade DC-DC converter was considered. Based on publicly available datasheet specifications of SynQor MCOTS-C-28VE-05-QT (SynQor, 2016), the converter operates over an input voltage range of 9-70 V with a rated output power of 85 W. The worst-case condition in terms of minimum input impedance occurs at the lowest input voltage, where the input current reaches its maximum value. Using this operating point and accounting for the converter's rated efficiency of 86%, the maximum input power is approximately 98.8 W. This yields a minimum negative input resistance magnitude ( $|R_n|$ ) of approximately 0.82  $\Omega$ .

For the proposed filter,  $R_d = 0.3 \Omega$  is adopted, which corresponds to approximately 36% of the minimum calculated  $|R_n|$ . This ratio ensures that the filter's output impedance remains well below the converter's input impedance, providing sufficient margin against instability while also helping to limit excessive resonance within the filter network.

### 2.4. DM Filter Design

When selecting the filter topology, L, T, and  $\pi$ -type configurations are the primary candidates. An L-type filter is a second-order network providing an attenuation slope of 40 dB/decade,

which is insufficient to meet the stringent high-frequency attenuation targets of military applications. Both T and  $\pi$ -type filters are third-order networks offering a steeper 60 dB/decade roll-off. However, a T-type filter presents a high inductive impedance to the converter, which can cause excessive voltage spikes when interacting with the pulsed input current of switch-mode DC-DC converters. Therefore, a  $\pi$ -type filter topology is selected for the DM stage; its capacitive output provides a low-impedance path for high-frequency switching currents while achieving the required 60 dB/decade attenuation slope (Tarateeraseth, 2011, 2012a, 2012b). The attenuation requirement is defined with respect to the dominant emission band identified for the targeted application.

The cut-off frequency of the DM filter is derived from the attenuation objective according to

$$f_{CDM} = \frac{f_{DM}}{\sqrt[6]{64 \times \left(10^{\frac{Att_{DM}}{10}} - 1\right)}} \quad (4)$$

where  $f_{DM}$  denotes the frequency of interest and  $Att_{DM}$  represents the desired attenuation level.

As observed in the baseline measurements (detailed in Section 3), the unfiltered system initially exceeds the CE102 limit at the converter's switching frequency of 83 kHz by approximately 25 dB. Although the minimum required attenuation could be derived directly from these unfiltered noise levels, merely targeting such relatively low attenuation values may not provide sufficient margin against the rigorous demands of military applications. To establish a more realistic and robust target, the performance levels of military-grade EMI filters were considered. According to the publicly available datasheet of the SynQor MCOTS-F-28-P-DM, DM attenuation values on the order of 80 dB are specified around 500 kHz (SynQor, 2024). Consequently, this information was adopted as a practical benchmark, and the DM attenuation target for the proposed design was selected as 80 dB at 500 kHz.

Using  $f_{DM} = 500 \text{ kHz}$  and  $Att_{DM} = 80 \text{ dB}$ , the theoretical cut-off frequency becomes

$$f_{CDM} = \frac{500k}{\sqrt[6]{64 \times \left(10^{\frac{80}{10}} - 1\right)}} \cong 11.6 \text{ kHz} \quad (5)$$

To accommodate parasitic elements, component tolerances, and modelling uncertainties, the practical cut-off frequency is conservatively selected below the theoretical estimate.

Once the cut-off frequency is defined, the total inductance and capacitance values are calculated from the design resistance as

$$L_t = \frac{R_d}{2 \times \pi \times f_{CDM}} \quad (6)$$

$$C_t = \frac{1}{2 \times \pi \times R_d \times f_{CDM}} \quad (7)$$

Using  $R_d = 0.3 \Omega$  and a practical cut-off frequency of 10 kHz, the resulting component values are

$$L_t = 4.77 \mu H \quad (8)$$

$$C_t = 53.05 \mu F \quad (9)$$

Finally, the nearest standard component ratings are selected to enable practical hardware implementation.

## 2.5. DM Damping Design

High quality-factor resonances may deteriorate attenuation performance and destabilize the interaction between the converter and the EMI filter. To mitigate this effect, RC shunt damping branches are connected across the DM capacitors.

The damping resistance is selected equal to the design resistance,

$$R_{DD} = R_d \quad (10)$$

which ensured adequate energy dissipation once the damping network becomes active.

The damping capacitance is chosen such that the resistive action is effectively introduced in the vicinity of the resonance frequency while minimizing additional losses outside this band. In practical EMI filter implementations, empirical design heuristics are frequently employed to guarantee robust behaviour. A widely adopted approach is to select the damping capacitance several times larger than the main DM capacitor.

In the present design, a factor of approximately four is used. This choice enables strong resonance suppression without significantly affecting low-frequency operation or increasing unnecessary power dissipation.

## 2.6. CM Filter Design

Similar to the DM stage, the topology of the CM filter must be carefully selected to provide sufficient attenuation slope. To achieve a steep attenuation roll-off of approximately 60 dB/decade at high frequencies, a third-order  $\pi$ -type structure composed of a CM choke and two sets of Y-capacitors is implemented. This configuration provides superior high-frequency suppression compared to a standard 40 dB/decade L-type CM filter. In practical military applications, however, safety and leakage current regulations impose strict limits on the total Y-capacitance connected to chassis ground. These constraints directly influence achievable cut-off frequencies and CM choke values. The attenuation objective is defined with respect to the dominant emission region identified for the application.

Similar to the DM design, to establish a realistic target, performance levels of military-grade EMI filters are considered. According to publicly available datasheet information of SynQor MCOTS-F-28-P-DM, CM attenuation values on the order of 60 dB are specified around 500 kHz (SynQor, 2024). This information is therefore adopted as a practical benchmark for the proposed design.

The cut-off frequency is calculated from the attenuation requirement as

$$f_{CCM} = \frac{f_{CM}}{\sqrt[6]{64 \times \left(10^{\frac{Att_{CM}}{10}} - 1\right)}} \quad (11)$$

Using  $f_{CM} = 500 \text{ kHz}$  and  $Att_{CM} = 60 \text{ dB}$ , the theoretical cut-off frequency becomes

$$f_{CCM} \cong 25 \text{ kHz} \quad (12)$$

Due to safety and leakage current limitations, the total allowable Y-capacitance is restricted. In the present design, the aggregate value is selected as

$$C_y = 0.60 \mu\text{F} \quad (13)$$

Based on the selected cut-off frequency of 15 kHz and total Y-capacitance, the required CM inductance is obtained as

$$L_{CM} = \frac{1}{2 \times \pi^2 \times f_{CCM}^2 \times C_y} \quad (14)$$

which yields approximately

$$L_{CM} \cong 375 \mu\text{H} \quad (15)$$

For practical implementation, the nearest standard value of 390  $\mu\text{H}$  is selected.

It should be noted that the physical implementation of the CM choke introduces a leakage inductance, which appears in series with the DM path. As discussed later in Section 4, this leakage inductance inherently contributes to the total DM inductance, lowering the effective cut-off frequency and enhancing low-frequency DM attenuation. However, to maintain a conservative design, this added additional attenuation is not factored into the initial ideal calculations.

## 2.7. CM Damping Design

To suppress high-Q resonance behaviour in the CM network, an RC damping branch is introduced. A practical and widely used approach is to select the damping resistance close to the characteristic impedance of the CM filter.

Accordingly, the damping resistance is approximated as

$$R_{CD} = \sqrt{\frac{2 \times L_{CM}}{C_y}} \quad (16)$$

Using  $L_{CM} = 390 \mu\text{H}$  and  $C_y = 0.60 \mu\text{F}$ , the resulting value becomes

$$R_{CD} \cong 36 \Omega \quad (17)$$

The damping capacitance is then chosen such that the resistive damping action becomes effective around the resonance frequency while minimizing additional losses outside the targeted band.

## **2.8. Integration and Simulation-Based Design Verification**


The DM and CM sections are combined into a unified EMI filter structure in order to evaluate overall response prior to hardware realization. The complete circuit is implemented in LTspice using the component values obtained in the previous subsections.

Figure 1 illustrates the DM, CM, and integrated filters modelled in LTspice for the simulation. The simulated DM attenuation is presented in Figure 2. The response indicates strong suppression in the dominant emission region while maintaining a smooth transition toward lower frequencies. No excessive resonance amplification is observed, confirming the effectiveness of the damping strategy.


Similarly, the CM attenuation characteristic is given in Figure 3. The results demonstrate that the CM stage provides the intended reduction across the targeted frequency band and that high-Q behaviour is effectively controlled.

At 500 kHz, the simulated attenuation reaches approximately 156.09 dB for the DM case and 94.80 dB for the CM case. These values correspond to idealized simulation conditions in which parasitic couplings, layout effects, and non-ideal component characteristics are not fully represented. Consequently, lower attenuation levels are expected in practical hardware realization (Chiu et al., 2024; Wang et al., 2004). Nevertheless, the simulation results clearly indicate that sufficient design margin exists with respect to the targeted attenuation requirements.


Overall, the simulation study confirms that the selected parameters are appropriate for meeting the intended attenuation objectives before proceeding to experimental validation. Hardware-based conducted emissions measurements are provided in the following section.



(a)




(b)




(c)

**Figure 1.** Filters modelled in LTspice for simulation: (a) DM filter, (b) CM filter, and (c) integrated filter.



**Figure 2.** Simulated DM attenuation of the proposed EMI filter obtained from the integrated LTspice model.




**Figure 3.** Simulated CM attenuation of the proposed EMI filter obtained from the integrated LTspice model.

### 3. Results

This section presents the experimental evaluation of the proposed EMI filter and demonstrates its effectiveness under realistic operating conditions. The validation is performed through conducted emissions measurements according to the CE102 requirement and complementary frequency-domain insertion loss analysis. Together, these tests provide verification of both standard compliance and the underlying attenuation mechanisms.

The EMI filter is realized using commercially available components and assembled on a dedicated printed circuit board. The prototype is designed with clear separation between CM and DM sections, while damping networks are integrated to suppress resonance phenomena. The compact implementation confirms that the proposed methodology leads to a solution compatible with practical military integration constraints.

The fabricated prototype is shown in Figure 4(a) and Figure 4(b), illustrating the top and bottom views of the board, respectively.




**Figure 4.** EMI filter prototype: (a) top view and (b) bottom view.

Following prototype realization, conducted emissions tests and frequency-domain measurements are carried out. The corresponding results are detailed in the subsequent subsections.

### 3.1. CE102 Conducted Emission Results


Conducted emission measurements are performed to experimentally verify compliance with the CE102 requirement. The test configuration is illustrated in Figure 5. Figure 5(a) shows the baseline setup without the EMI filter, while Figure 5(b) presents the configuration with the proposed filter inserted at the converter input. The DC–DC converter is powered through LISNs), and the conducted noise is measured separately on the positive and negative input lines. All tests are conducted under identical input voltage and load conditions to enable a fair comparison between filtered and unfiltered operation.




**Figure 5.** CE102 test configuration: (a) without EMI filter and (b) with the proposed EMI filter.

The baseline results obtained without the EMI filter are presented in Figure 6(a) and Figure 6(b) for the positive and negative lines, respectively. The spectra clearly exceed the CE102 limits; specifically, the first significant violation occurs at the converter's switching frequency of 83 kHz, where the emissions exceed the limit by approximately 25 dB. These observations confirm that additional filtering is mandatory for standard compliance.

Following integration of the proposed EMI filter, the measurements are repeated. The corresponding results are shown in Figure 7(a) and Figure 7(b). A substantial reduction in conducted noise is observed across the entire frequency range. The emission levels remain below the CE102 limits with visible margin for both lines, demonstrating successful mitigation of interference.



**Figure 6.** Baseline CE102 conducted emission results without the EMI filter: (a) positive line and (b) negative line.



**Figure 7.** CE102 conducted emission results with the proposed EMI filter: **(a)** positive line and **(b)** negative line.

Moreover, the close similarity between the positive and negative line responses indicates symmetrical and consistent filter behaviour, further validating the robustness of the design.

### 3.2. Frequency-Domain Insertion Loss Results

While CE102 testing provides direct confirmation of standard compliance, it does not fully reveal how attenuation is distributed across frequency or how the CM and DM sections contribute to the overall performance. For this reason, insertion loss measurements are performed using a Bode analyser.


The measured CM insertion loss is presented in Figure 8 and Figure 9 for the low and high-frequency ranges, respectively. In the lower frequency region, the experimental results generally follow the trend predicted by simulations, and the filter exhibits the expected attenuation behaviour. Around 500 kHz, approximately 75 dB attenuation is obtained, exceeding the initial 60 dB objective.

At higher frequencies, however, the attenuation deviates from the ideal response and decreases to nearly 55 dB in certain regions. These variations are primarily attributed to non-ideal effects such as parasitic capacitances, leakage paths, component self-resonances, and layout-related couplings. Such mechanisms become dominant in the MHz range and limit the achievable performance compared to ideal simulations.


The DM insertion loss is shown in Figure 10 and Figure 11. In the low-frequency region, the measured attenuation is notably higher than predicted by simulations, indicating that the effective filtering action begins earlier than anticipated during the design stage. This behaviour is mainly associated with additional parasitic inductive contributions that are not fully captured in the idealized model.

At 500 kHz, the measured attenuation is approximately 55 dB, which is lower than the 80 dB benchmark adopted during design. Furthermore, the attenuation profile is not strictly


monotonic across frequency; depending on resonance interactions, local improvements or degradations are observed. Such deviations from ideal predictions are expected in practical implementations due to parasitic couplings, leakage inductances, component self-resonance, PCB layout effects, and measurement fixture influences.




**Figure 8.** Measured CM insertion loss in the 0-500 kHz frequency range.



**Figure 9.** Measured CM insertion loss in the 500 kHz-10 MHz frequency range.



**Figure 10.** Measured DM insertion loss in the 0-500 kHz frequency range.



**Figure 11.** Measured DM insertion loss in the 500 kHz-10 MHz frequency range.

Despite these deviations, the measurements confirm that substantial suppression is achieved within the frequency region critical for CE102 compliance. The presence of frequency-dependent variations further highlights the importance of parasitic and cross-coupling effects, which are investigated in more detail in the Discussion section.

Overall, the Bode analyser results provide physical interpretation of the successful CE102 performance and demonstrate that the realized filter maintains effective attenuation even under non-ideal practical conditions.

## 4. Discussion

The experimental and frequency-domain results provide valuable insight into the practical behaviour of the proposed EMI filter, highlighting both the differences between idealized predictions and real hardware performance, as well as its advantages over existing methodologies in the literature.

One of the most notable observations is the variation between simulation and measurement outcomes. In the low-frequency region, the measured attenuation, particularly in DM, is in several cases higher than predicted. This indicates that additional inductive contributions, such as leakage inductance from the CM choke and unintended parasitic elements, effectively shift the filter action toward lower frequencies.

In contrast, as frequency increases, the measured attenuation becomes limited compared to simulation expectations. The reduction in performance is primarily associated with parasitic capacitances, component self-resonances, and coupling mechanisms introduced by the physical implementation. These effects create alternative propagation paths and resonance conditions that cannot be fully captured in ideal circuit representations.

A dominant contributor to the discrepancy between simulation and measurement is the leakage inductance of the CM choke. Experimental measurements show that this parasitic inductance ( $\sim 15 \mu\text{H}$ ) is several times larger than the nominal DM inductance assumed during the design stage. Consequently, the effective inductance in the DM path increases, shifting the cut-off frequency toward lower values. This mechanism explains the improved attenuation observed at low frequencies while simultaneously moving self-resonance and interaction effects into earlier regions of the spectrum, thereby constraining high-frequency performance. However, unlike the models proposed by Cadirci et al. (2005), which showed significant performance degradation above 5 MHz due to self-resonance effects, the filter in this study maintained effective attenuation across the entire CE102 spectrum due to optimized layout considerations and resonance suppression.

Beyond these physical insights, the results demonstrate a significant advantage of the proposed framework in achieving superior attenuation and volumetric efficiency through a stability-oriented design flow. For instance, while high-performance active designs such as the cascaded FF-FB AEF reported by Tian et al. (2026) achieved a CM attenuation of 42 dB, the passive filter in this study reached 75 dB at 500 kHz. Furthermore, conventional passive approaches, such as the one detailed by Kichouliya et al. (2023), often require very large magnetic components—such as 21.1 mH CM chokes—to meet low-frequency limits in high-power systems. By integrating the converter's negative incremental resistance into the determination of the design resistance and utilizing RC damping networks, this study achieves compliance with a significantly smaller 390  $\mu\text{H}$  inductor. This demonstrates that prioritizing system stability allows for the use of larger capacitances and smaller inductances without compromising the compact physical envelope required for military platforms.

Despite the non-ideal limitations, both CE102 and Bode analyser measurements verify that the realized filter maintains sufficient attenuation in the frequency range critical for compliance. The existence of margin under realistic parasitic conditions confirms the robustness of the proposed design methodology. From an engineering standpoint, the findings emphasize that accurate prediction of high-frequency performance requires detailed modelling of parasitic elements, component construction, and PCB layout. Nevertheless, despite the added complexity of active alternatives, this study confirms that a systematic passive approach provides a more robust, reliable, and simpler framework that leads to successful MIL-STD-461G compliance without excessive redesign effort.

## 5. Conclusion

This paper presented a systematic and practical EMI filter design methodology for military DC–DC converter applications targeting compliance with the MIL-STD-461G CE102 conducted emission limits. The approach links attenuation objectives, stability considerations, and realizable component values within a structured design flow.

The proposed filter was implemented using commercially available components and validated through both CE102 measurements and frequency-domain insertion loss analysis. The experimental results demonstrated that, while deviations from ideal simulations arise due to parasitic effects and component non-idealities, sufficient attenuation margin is maintained within the frequency region critical for compliance. The agreement between analytical expectations, simulations, and hardware measurements confirms that the methodology provides a reliable and repeatable framework for practical EMI mitigation.

The study also highlighted the strong influence of parasitic elements, particularly the leakage inductance of the CM choke, on the overall filter behaviour. These findings underline the necessity of considering cross-coupling between CM and DM paths in real implementations.

Future work may focus on incorporating detailed parasitic models into early-stage simulations to improve prediction accuracy, especially in the high-frequency region. In addition, optimization of magnetic component design, including custom winding strategies for CM chokes, represents a promising direction for enhancing performance. Further research may also explore mechanical integration concepts that facilitate direct deployment of compact EMI filtering solutions in military and aerospace platforms, as well as the inclusion of extended protection features against abnormal operating conditions.

## Funding

This research received no external funding.

## Acknowledgments

The authors would like to thank Aselsan, Anova, and OB Sensör Teknolojileri for their support during the measurement and testing activities.

## Ethical Statements

Not applicable.

## Data and Code Availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

## Supplementary Materials

Not applicable.

## References

- Cadirci, I., Saka, B., & Y. Eristiren. (2005). Practical EMI-filter-design procedure for high-power high-frequency SMPS according to MIL-STD 461. *IEE Proceedings*, 152(4), 775–775. <https://doi.org/10.1049/ip-epa:20045079>
- Chiu, P.-T., Pong, M. H., Yao, C.-J., & Chiu, H.-J. (2024). The effect of parasitic elements on EMI common mode filter insertion loss. In *Proceedings of the International Conference on Industrial Electronics and Applications (ICIEA)* (pp. 1–6). <https://doi.org/10.1109/ICIEA61579.2024.10664692>
- Fishta, M., & Fiori, F. (2026). A volume-optimized hybrid EMI filter for automotive traction inverters. *IEEE Transactions on Electromagnetic Compatibility*, 68(2), 321–329. <https://doi.org/10.1109/TEMPC.2026.3652665>
- Hartal, O. (1995). *EMC by design* (3rd ed.). R & B Enterprises.
- Kichouliya, R., Sundar, S., & Reddy, P. (2023). Power line EMI filter design of power converters to meet the conducted emission specifications of MIL-STD-461G. In *Proceedings of the National Power Electronics Conference (NPEC)* (pp. 1–6). <https://doi.org/10.1109/NPEC57805.2023.10385027>
- Middlebrook, R. D. (1976). Input filter considerations in design and applications of switching regulators. *Proc. IEEE IAS'76*, 91-107. <https://ci.nii.ac.jp/naid/80014902952/>
- Nave, M. J. (1991). *Power line filter design for switched-mode power supplies*. Van Nostrand Reinhold.
- Ozenbaugh, R. L., & Pullen, T. M. (2012). *EMI filter design* (3rd ed.). CRC Press.
- SynQor. (2016). MCOTS-C-28VE-05-QT: Military COTS DC-DC Converter Datasheet. <https://www.synqor.com/products/mil-cots/mcots-c-28ve-05-qt>

- SynQor. (2024). MCOTS-F-28E-P-DM: Military COTS EMI filter Datasheet. <https://www.synqor.com/products/mil-cots/mcots-f-28e-p-dm>
- Tarateeraseth, V. (2011). EMI Filter Design Part I: Conducted EMI Generation Mechanism. In *IEEE Electromagn. Conf* (pp. 44-50).
- Tarateeraseth, V. (2012a). EMI filter design part II: Measurement of noise source impedances. *IEEE EMC Magazine*, 1(1), 42–49. <https://doi.org/10.1109/MEMC.2012.6244944>
- Tarateeraseth, V. (2012b). EMI filter design part III: Selection of filter topology for optimal performance. *IEEE EMC Magazine*, 1(2), 60–73. <https://doi.org/10.1109/MEMC.2012.6244975>
- Tian, Y., Jiang, Y., Liao, Y., & Tan, Y. K. (2026). Highly-Efficient Cascaded Active EMI Filter for WBG-based integrated Motor Drive. *IEEE Journal of Emerging and Selected Topics in Power Electronics*, 1–1. <https://doi.org/10.1109/jestpe.2026.3677211>
- U.S. Department of Defense. (2015). MIL-STD-461G: *Requirements for the control of EMI characteristics of subsystems and equipment*. U.S. Government Printing Office.
- Wang, S., Lee, F. C., Chen, D. Y., & Odendaal, W. G. (2004). Effects of parasitic parameters on EMI filter performance. *IEEE Transactions on Power Electronics*, 19(3), 869–877. <https://doi.org/10.1109/TPEL.2004.826527>
- Williams, T. (2001). *EMC for product designers* (3rd ed.). Newnes.

# Mechanically Alloyed Nd-Fe-B Nanoparticles with Graphitic Carbon Shell: Synthesis, Characterization and Magnetic Properties

Mustafa Bellek<sup>1,\*</sup> , Sina Kavak<sup>1</sup> , Duygu Ağaoğulları<sup>2</sup> 

## Abstract

In this study, Nd-Fe-B nanoparticles with permanent magnet properties were synthesized from elemental raw materials using the mechanical alloying (MA) method and subsequently encapsulated with graphitic carbon shell to enhance their thermal and chemical stability for potential electronic applications. Alloy composition Nd  $\approx$  23.3 wt.%, Fe  $\approx$  75.7 wt.%, B  $\approx$  1.0 wt.%, Nd<sub>10</sub>Fe<sub>84</sub>B<sub>6</sub> atomic composition, were milled for varying durations 8-10 and 12 h to evaluate the effects of milling duration on magnetic performance. Due to high-energy milling, powder yield remained between 41% and 53%, influenced by particle adhesion and partial melting. X-ray Diffraction (XRD) confirmed the formation of the tetragonal Nd<sub>2</sub>Fe<sub>14</sub>B phase in all MA-synthesized powders, while Dynamic Light Scattering (DLS) and Scanning Electron Microscopy (SEM) analyses verified nanoparticle size and agglomerated morphologies characteristic of welding-fracture-rewelding cycles in MA. Annealing at 900°C for 1 h improved crystallinity and promoted the formation of additional phases, as supported by XRD, SEM and Energy Dispersive Spectroscopy (EDS) results. A Graphitic Carbon Shell (GCS) was successfully deposited on the nanoparticle surfaces via Chemical Vapor Deposition (CVD) under a CH<sub>4</sub>/H<sub>2</sub> atmosphere at 950°C; the presence of the carbonaceous coating was confirmed by XRD and elemental mapping. Transmission Electron Microscopy (TEM) analysis revealed the core-shell structure, the number of graphitic carbon shell layers, and the uniformity of the encapsulation surrounding the nanoparticles. Magnetic characterization using Vibrating Sample Magnetometer (VSM) showed variations in coercivity, saturation magnetization and remanence before and after encapsulation. Among all samples, the 10 h milled, annealed and graphitic carbon shell encapsulated sample (NFB2TG) exhibited the most favorable magnetic properties, reaching a coercivity of 185.71 Oe and saturation magnetization of 46.3 emu/g. These results demonstrate that MA followed by annealing and encapsulation GCS is a promising route for producing stable Nd-Fe-B nanoparticles.

**Keywords:** Nd-Fe-B nanoparticles, mechanical alloying, graphitic carbon shell encapsulation, chemical vapor deposition, magnetic properties.

Received: 25.11.2025

Accepted: 22.05.2026

Research Article

<sup>1</sup>ASELSAN Inc., Ankara, Türkiye

<sup>2</sup> Istanbul Technical University, Istanbul, Türkiye

\* Corresponding author.

✉ mubellek@aselsan.com

Cite this article as:

<https://doi.org/10.65834/jdsi.12.46>

Bellek, M., Kavak, S. & Ağaoğulları, D. (2026). Mechanically Alloyed Nd-Fe-B Nanoparticles with Graphitic Carbon Shell: Synthesis, Characterization and Magnetic Properties. *Journal of Defence and Security Industries: Science & Technology Section* 2(1), 177-200.

## 1. Introduction

The main goal in studies on magnetism has been to develop more powerful permanent magnets. Rare Earth (RE)-Fe-B permanent magnets are known for their superior magnetic properties, in which Neodymium (Nd) is generally used as the RE element. They are frequently preferred in industry due to their lower cost compared to Sm-Co magnets (Wu et al., 1997). Various production methods exist for Nd-Fe-B permanent magnets: Powder Metallurgy (PM), melt spinning, hot deformation and Hydrogen Decrepitation Desorption Recombination (HDDR) (Rahimi et al., 2017; Zhou et al., 2019). The primary objective of these methods is to produce anisotropic Nd-Fe-B magnets with improved magnetic properties. Chemical synthesis of Nd<sub>2</sub>Fe<sub>14</sub>B nanoparticles with tetragonal lattice structure is challenging due to the high reduction potential of Nd<sup>3+</sup> and the oxidation tendency of Nd<sub>2</sub>Fe<sub>14</sub>B (Rahimi et al., 2017).

It is anticipated that Nd-Fe-B nanoparticles produced by solid-state synthesis methods will exhibit favorable saturation magnetization and coercivity values (De Moraes et al., 2024; Schultz et al., 1987). However, their high surface area makes them chemically reactive, necessitating surface coating (Jung et al., 2024). Several methods have been investigated to increase the oxidation resistance of Nd-Fe-B magnets (Lee et al., 2005; Liu et al., 2012; Pich et al., 2005). While alloying can also enhance oxidation resistance, it frequently degrades magnetic properties, making protective coatings the preferred approach for practical applications (Chen et al., 2007; Tung et al., 2015).

In the literature, Nd-Fe-B magnetic materials have been synthesized in a variety of forms through distinct production routes. Nanoparticles synthesized from a Nd-Fe-B alloy using hydrogen decrepitation and high-energy ball milling have been investigated for their efficacy as Magnetic Hyperthermia (MH) agents. Périgo et al. (2012) found that the resulting materials, characterized by a mixed-phase structure, demonstrated superior heating capabilities. Under an applied Alternating Current (AC) magnetic field (3.7 kA m<sup>-1</sup>, 228 kHz), a maximum estimated Specific Absorption Rate (SAR) of approximately 2500 W kg<sup>-1</sup> was recorded. The study also confirmed acceptable in vitro biocompatibility attributed to a protective coating applied during milling. The development of high-performance magnetic nanoparticles necessitates continuous optimization of synthesis techniques, particularly ball milling methods, to control both size and shape. Chakka et al. (2006) utilized surfactant-assisted ball milling in an organic carrier liquid to synthesize Nd-Fe-B nanoparticles with a narrow size distribution below 30 nm. Their investigation highlighted a key morphological distinction: while Fe and Fe-Co particles were largely spherical, Nd-Fe-B, Co and Sm-Co nanoparticles exhibited elongated rod shapes, with most systems displaying superparamagnetic behavior at room temperature. Zhang et al. (2018) successfully employed a sol-gel process, using citric acid and glycol, to prepare Nd-Fe-B nanoparticles consistently below 100 nm. Their method relied on a two-step thermal treatment: an initial annealing to form oxide precursors and a subsequent reductive annealing to yield the final metallic Nd-Fe-B nanoparticles. As confirmed by Vibrating Sample Magnetometer (VSM) analysis, the saturation magnetization increased substantially from 1.08 emu/g to 14.17

emu/g, validating successful hard magnetic phase formation. Mehrifar et al. (2025) demonstrated the synthesis of Nd<sub>2</sub>Fe<sub>14</sub>B nanoparticles using ethanol-assisted wet ball milling, efficiently reducing the starting powder size through prolonged milling.

In addition to synthesis investigations, coating studies have been conducted to enhance oxidation resistance, chemical and thermal stability. Hosseinabadi et al. (2019) reported that Nd-Fe-B powders developed a thin amorphous carbon layer derived from oleic acid surfactant during high-energy wet ball milling. Increasing the surfactant concentration produced a thicker carbonaceous coating, enhancing the imaginary part of dielectric permittivity by approximately 95%, increasing absorption shielding efficiency by about 40% and total shielding efficiency by nearly 30% at certain frequencies, attributed to improved impedance matching and strengthened microwave absorption. Ouyang et al. (2019) demonstrated that a polydopamine-based magnetic fluid coating provided exceptional corrosion inhibition, with corrosion current density of  $1.74 \times 10^{-3}$  mA/cm<sup>2</sup> in 3.5 wt.% NaCl solution over 30 days, roughly three orders of magnitude lower than bare Nd-Fe-B (1.01 mA/cm<sup>2</sup>), and excellent self-healing behavior after mechanical damage. Shen et al. (2018) improved corrosion resistance of sintered Nd-Fe-B magnets using Ni-Al<sub>2</sub>O<sub>3</sub> nanocomposite coatings via pulsed current-jet electrodeposition, finding that nanoparticle incorporation enhanced both adhesion strength and corrosion resistance over pure nickel coatings. Yang et al. (2022) modified multiwall Carbon Nanotubes (CNTs) via noncovalent functionalization with tannic acid to form Tannic Acid Functionalized Carbon Nanotubes (TCNTs) hybrids and incorporated them into epoxy coatings on sintered Nd-Fe-B magnets. Electrochemical tests showed specimens with 2 g/L TCNTs maintained a high impedance ( $|Z|_{0.01\text{Hz}} \approx 10^8 \Omega \text{ cm}^2$ ) after 36 days in 3.5 wt.% NaCl solution, demonstrating excellent corrosion resistance. Bystrzejewski et al. (2007) produced carbon nanoencapsulates of 20-50 nm using an arc discharge method with graphite anodes doped with Nd<sub>2</sub>Fe<sub>14</sub>B, achieving a maximum coercive force of approximately 413 Oe.

Nd-Fe-B nanoparticles have a wide range of potential applications. In biomedicine, they are used for targeted drug delivery, MH and Magnetic Resonance Imaging (MRI) enhancement (Iacovacci et al., 2015; Omelyanchik et al., 2020; Saritas et al., 2012). They have also been investigated for Micro-Electro-Mechanical Systems (MEMS) devices (Dempsey, 2009; Lecerf et al., 2025) and magnetically recoverable catalysts (Yang et al., 2019). Their main limitation is rapid oxidation, necessitating protective encapsulation with carbon, silica or polymer shells.

In this study, Nd-Fe-B nanoparticles were synthesized via Mechanical Alloying (MA), a process driven by cyclical cold welding, fracturing and re-welding of elemental powders to achieve a refined, homogeneous microstructure (Ağaoğulları et al., 2011; Ağaoğulları et al., 2017; Suryanarayana, 2010). The Nd<sub>10</sub>Fe<sub>84</sub>B<sub>6</sub> composition was selected because it is slightly Fe-rich relative to the stoichiometric Nd<sub>2</sub>Fe<sub>14</sub>B phase (Nd<sub>11.76</sub>Fe<sub>82.35</sub>B<sub>5.88</sub> at. %). This Fe-rich, Nd-lean composition promotes the simultaneous formation of hard magnetic Nd<sub>2</sub>Fe<sub>14</sub>B and soft magnetic  $\alpha$ -Fe phases during MA and annealing. The exchange coupling between these two nanoscale phases is known to enhance remanence beyond the single-phase limit, and this composition has been widely investigated in nanocrystalline Nd-Fe-B systems (Manaf et al.,

1993; Schultz et al., 1987). Nanoparticles were milled for 8, 10 and 12 h, followed by annealing at 900°C. Comprehensive characterization using X-Ray Diffraction (XRD), Dynamic Light Scattering (DLS), Scanning Electron Microscopy (SEM) and Energy Dispersive Spectroscopy (EDS) confirmed the successful synthesis of nanoscale particles containing the Nd<sub>2</sub>Fe<sub>14</sub>B tetragonal phase. Following annealing, the powders were subjected to Chemical Vapor Deposition (CVD) at 950°C for Graphitic Carbon Shell (GCS) encapsulation. Transmission Electron Microscopy (TEM) analysis revealed a distinct core-shell architecture with a continuous GCS surrounding the nanoparticles. The optimized sample (NFB2TG) exhibited a coercivity of 185.71 Oe and a saturation magnetization of 46.3 emu/g, demonstrating that GCS-encapsulated, MA-synthesized Nd-Fe-B nanoparticles represent a promising route for producing chemically stable magnetic nanoparticles.

## 2. Methodology

### 2.1. Powder Preparation

Neodymium powders (Nanografi A.Ş., 99.5% purity, ~325 µm) were combined with Fe powders (Alfa Aesar, ~44 µm) and B powders (Pavezyum Kimya, <1 µm). Powder mixtures of 3 g each were prepared in the Nd<sub>10</sub>Fe<sub>84</sub>B<sub>6</sub> composition, corresponding to Nd ≈ 23.3 wt.%, Fe ≈ 75.7 wt.% and B ≈ 1.0 wt.%, giving 0.70 g Nd, 2.27 g Fe and 0.03 g B per batch. All powder handling was performed in a glove box (MBraun™ LABstar) under a high-purity Ar atmosphere (Linde™, 99.999%), with O<sub>2</sub> and H<sub>2</sub>O levels below 5 ppm. The powders were transferred into stainless steel vials sealed under Ar and prepared for milling at a powder-to-ball ratio of 1:10 using 6 mm diameter stainless steel balls.

### 2.2. Mechanical Alloying

Powder mixtures were milled for 8, 10 and 12 h in a high-energy ball mill (Spex™ 8000D Mixer/Mill). After milling, the vials were cooled at room temperature for approximately 12 h in air, then opened inside the glove box under Ar to prevent oxidation. Samples milled for 8, 10 and 12 h were designated NFB1, NFB2 and NFB3, respectively. After annealing, the letter 'T' was appended as NFB1T, NFB2T, NFB3T; after graphitic carbon shell encapsulation, 'G' was further appended, yielding NFB1TG, NFB2TG and NFB3TG. Sample designations are summarized in Table 1.

**Table 1.** Sample designations as a function of milling duration and applied processing steps.

| Sample Name | Milling Time (h) | MA  | Annealed | Graphitic Carbon Shell Encapsulated |
|-------------|------------------|-----|----------|-------------------------------------|
| NFB1        | 8                | Yes | No       | No                                  |
| NFB2        | 10               | Yes | No       | No                                  |
| NFB3        | 12               | Yes | No       | No                                  |
| NFB1T       | 8                | Yes | Yes      | No                                  |
| NFB2T       | 10               | Yes | Yes      | No                                  |
| NFB3T       | 12               | Yes | Yes      | No                                  |
| NFB1TG      | 8                | Yes | Yes      | Yes                                 |
| NFB2TG      | 10               | Yes | Yes      | Yes                                 |
| NFB3TG      | 12               | Yes | Yes      | Yes                                 |


### 2.3. Annealing Process

Following mechanical alloying, the  $\text{Nd}_{10}\text{Fe}_{84}\text{B}_6$  powders containing the tetragonal  $\text{Nd}_2\text{Fe}_{14}\text{B}$  phase were annealed to increase the proportion of this phase and improve crystallinity. Annealing was performed in a temperature- and pressure-controlled tube furnace (PROTHERM™ Furnaces), as shown in Figure 1, under flowing Ar (100 ml/min) at 50 mbar.

The annealing procedure (Figure 2) consisted of three sequential stages: (1) heating from room temperature to 900°C over 1 h, (2) isothermal holding at 900°C for 1 h, and (3) controlled cooling from 900°C to 300°C over 1 h. Ar flow and vacuum were maintained until room temperature was reached, after which the system was switched off and allowed to cool naturally.



**Figure 1.** Annealing experimental setup (PROTHERM™ Furnaces).




**Figure 2.** Annealing process regime.

## 2.4. Graphitic Carbon Shell Encapsulation

Following annealing, all samples were subjected to CVD encapsulation in a temperature-controlled tube furnace (PROTHERM™ Furnaces) under 50 mbar vacuum, using H<sub>2</sub> (Linde™, 99.99%) and CH<sub>4</sub> (Linde™, 99.99%) as carrier and carbon precursor gases, respectively.


The CVD procedure (Figure 3) comprised three stages: (1) Heating from room temperature to 950°C over 1 h, (2) Isothermal holding at 950°C for 1 h to promote carbon deposition, and (3) Controlled cooling from 950°C to 300°C over 1 h. H<sub>2</sub> was introduced at room temperature at 100 ml/min; CH<sub>4</sub> was introduced at 550°C at the same flow rate. Both gases were shut off at 550°C during cooling while the vacuum was maintained until room temperature was reached. The resulting coating is referred to throughout as a graphitic carbon shell, as confirmed by XRD and TEM.



**Figure 3.** Chemical Vapor Deposition (CVD) process parameter regime.

### 3. Results

The Nd-Fe-B powder mixtures were prepared in the  $\text{Nd}_{10}\text{Fe}_{84}\text{B}_6$  atomic composition, corresponding to 0.70 g Nd, 2.27 g Fe and 0.03 g B per 3 g batch. Figure 4 presents the XRD patterns following mechanical alloying for 8, 10 and 12 h. The  $\text{Nd}_2\text{Fe}_{14}\text{B}$  phase with tetragonal crystal structure was confirmed in all three samples. Alongside the primary  $\text{Nd}_2\text{Fe}_{14}\text{B}$  phase,  $\alpha$ -Fe and  $\text{Fe}_2\text{O}_3$  phases were identified. The presence of  $\text{Fe}_2\text{O}_3$  indicates partial surface oxidation during the air-cooling period following milling. Given the high surface area and chemical reactivity of  $\text{Nd}_2\text{Fe}_{14}\text{B}$  nanoparticles, even brief air exposure is sufficient to promote oxidation, which is expected to influence the magnetic properties and phase distribution.




**Figure 4.** X-Ray Diffraction (XRD) pattern of the samples; (a) NFB1, (b) NFB2, (c) NFB3.

DLS analysis of the NFB1, NFB2 and NFB3 samples yielded average hydrodynamic sizes of 15.48 nm, 20.83 nm and 146.3 nm for the 8, 10 and 12 h milled samples, respectively. It should be noted that DLS measures the hydrodynamic diameter of agglomerates in suspension rather than true primary particle size. As reported by Lim et al. (2013), DLS is highly sensitive to agglomeration in magnetic nanoparticle systems and can significantly overestimate primary particle dimensions due to agglomerate contributions to the light scattering signal. Furthermore, Jia et al. (2023) noted that scattering intensity scales with the sixth power of particle diameter, meaning agglomerates dominate the signal and mask the true primary particle size. The apparent increase in measured size with longer milling durations is therefore attributed to progressive agglomeration rather than grain growth, consistent with findings by

Mehrfar et al. (2025) for dry ball-milled  $\text{Nd}_2\text{Fe}_{14}\text{B}$  systems. A TEM-based primary particle size analysis would be required to accurately determine individual particle dimensions, and this is acknowledged as a limitation of the current characterization approach.


Figure 5 presents SEM images of NFB1, NFB2 and NFB3 at 500 $\times$ , 1000 $\times$  and 3000 $\times$  magnifications. Pronounced agglomeration is evident in all samples, forming irregularly shaped clusters. This is characteristic of the MA process, in which repeated cold welding, fracturing and re-welding promote strong interparticle bonding (Ağaoğulları et al., 2011; Suryanarayana, 2010). Strong magnetic dipole-dipole interactions inherent to Nd-Fe-B systems further drive agglomeration upon air exposure. SEM reveals agglomerate morphology rather than individual primary particles; agglomerate sizes are consistent with the DLS measurements. The degree of agglomeration increases with milling duration, with NFB3 displaying larger and more densely packed structures, consistent with progressive oxidation-enhanced surface reactivity at longer milling times.



**Figure 5.** Scanning Electron Microscopy (SEM) images at 500X, 1000X and 3000X magnifications: (a) NFB1, (b) NFB2 and (c) NFB3.

Figure 6 shows the EDS analysis points on SEM images of NFB1, NFB2 and NFB3, with elemental results in Table 2. NFB1 and NFB3 show close agreement between two analysis points, indicating relatively homogeneous elemental distributions. NFB2 shows lower Fe content and greater variability between points. EDS is inherently unreliable for quantitative boron detection due to the low atomic number of boron ( $Z = 5$ ), resulting in characteristic X-

rays of insufficient energy for accurate quantification (Goldstein et al., 2018). The boron values reported in Table 2, including elevated B concentrations in NFB2, should be treated as semi-quantitative estimates only. The target composition was Nd  $\approx$  23.3 wt.%, Fe  $\approx$  75.7 wt.%, B  $\approx$  1.0 wt.%; among the three samples, NFB1 shows the closest agreement with this intended distribution.




**Figure 6.** Energy Dispersive Spectroscopy (EDS) analysis points marked on the SEM images of the mechanically alloyed powders: (a) NFB1, (b) NFB2 and (c) NFB3.

**Table 2.** Energy Dispersive Spectroscopy (EDS) analysis results correspond to the points shown in Figure 6.

| Sample      |          | Elemental Ratio |            |           |
|-------------|----------|-----------------|------------|-----------|
| <b>NFB1</b> |          | <b>Nd%</b>      | <b>Fe%</b> | <b>B%</b> |
|             | <b>1</b> | 25.83           | 74.17      | 0         |
|             | <b>2</b> | 26.64           | 74.36      | 0         |
| <b>NFB2</b> |          | <b>Nd%</b>      | <b>Fe%</b> | <b>B%</b> |
|             | <b>1</b> | 25.10           | 69.44      | 5.46      |
|             | <b>2</b> | 21.71           | 57.55      | 20.74     |
| <b>NFB3</b> |          | <b>Nd%</b>      | <b>Fe%</b> | <b>B%</b> |
|             | <b>1</b> | 25.92           | 74.08      | 0         |
|             | <b>2</b> | 27.08           | 72.92      | 0         |


Figure 7 presents the VSM hysteresis curve of NFB2. The sample exhibits a coercivity of 54.81 Oe, a saturation magnetization of 133.43 emu/g and a remanent magnetization of 2.83 emu/g under approximately 2 T. While the saturation magnetization falls within the expected range for Nd-Fe-B systems, the coercivity of 54.81 Oe is notably below values reported for bulk sintered Nd-Fe-B magnets, typically exceeding 10,000 Oe (Coey, 2010). The as-milled sample behaves as a magnetically soft material, reflecting: (1) the heavily strained, partially amorphous microstructure limiting Nd<sub>2</sub>Fe<sub>14</sub>B crystallization and magnetocrystalline anisotropy development; (2) soft magnetic  $\alpha$ -Fe and non-magnetic Fe<sub>2</sub>O<sub>3</sub> secondary phases reducing effective anisotropy; and (3) disordered surface spins at the high-surface-area nanoparticle

boundaries suppressing coercivity. The low remanence ratio ( $M_r/M_s \approx 0.021$ ) confirms an isotropic, exchange-decoupled system.



**Figure 7.** Vibrating sample magnetometer (VSM) analysis result shown on the hysteresis curve of the NFB2 sample.


Figure 8 presents XRD patterns of NFB1T, NFB2T and NFB3T following annealing at 900°C for 1 h. The  $\text{Nd}_2\text{Fe}_{14}\text{B}$  tetragonal phase is confirmed in all three samples, with notably increased peak intensities compared to as-milled counterparts (Figure 4), indicating significant improvement in crystallinity. The sharpening of  $\text{Nd}_2\text{Fe}_{14}\text{B}$  diffraction peaks is consistent with grain growth and the transformation of the partially amorphous as-milled structure into a more ordered crystalline state. The  $\alpha\text{-Fe}$  and  $\text{Fe}_2\text{O}_3$  phases remain present. Persistence of  $\text{Fe}_2\text{O}_3$  confirms that oxidation is not eliminated by thermal treatment under Ar, suggesting surface oxidation had progressed significantly prior to annealing. The continued  $\alpha\text{-Fe}$  is expected for the  $\text{Fe-rich Nd}_{10}\text{Fe}_{84}\text{B}_6$  composition, as excess iron segregates as a soft magnetic phase. While the resulting  $\text{Nd}_2\text{Fe}_{14}\text{B}/\alpha\text{-Fe}$  nanocomposite can theoretically enhance remanence through exchange coupling, the simultaneous presence of non-magnetic  $\text{Fe}_2\text{O}_3$  dilutes the effective magnetic moment and limits coercivity development.




**Figure 8.** XRD patterns of the annealed samples produced in the Nd-Fe-B composition: (a) NFB1T, (b) NFB2T, (c) NFB3T.

Figure 9 presents SEM images of NFB1T, NFB2T and NFB3T. Agglomeration persists after annealing, consistent with as-milled morphology (Figure 5). However, the annealed samples display a more defined spherical particle shape, attributed to surface energy minimization during annealing, whereby atomic diffusion and rearrangement drive the formation of more energetically favorable spherical geometries. The degree of agglomeration is comparable across the three conditions, suggesting that 900°C annealing did not promote significant interparticle sintering. The pre-existing oxidized surface layer is expected to act as a diffusion barrier limiting grain boundary migration.

Figure 10 shows EDS analysis points for NFB1T, NFB2T and NFB3T, with results in Table 3. In NFB1T and NFB2T, an increase in Nd content accompanied by a decrease in Fe content is observed relative to the as-milled state. This compositional shift is consistent with Nd redistribution during annealing, whereby Nd-rich phases segregate toward grain boundaries and particle surfaces at elevated temperatures, a phenomenon well documented in Nd-Fe-B systems (Sepehri-Amin et al., 2014). NFB3T shows the opposite trend, suggesting a different local phase distribution resulting from more extensive structural changes during 12 h milling. Among the annealed samples, NFB2T shows the closest agreement with the  $\text{Nd}_{10}\text{Fe}_{84}\text{B}_6$  target composition. As previously established, EDS boron values should be treated as semi-quantitative only (Goldstein et al., 2018).



**Figure 9.** SEM images of the annealed Nd-Fe-B samples at 500X, 1000X and 3000X magnifications: (a) NFB1T, (b) NFB2T, (c) NFB3T.




**Figure 10.** EDS analysis points marked on the SEM images of the mechanically alloyed and annealed Nd-Fe-B powders: (a) NFB1T, (b) NFB2T, (c) NFB3T.

**Table 3.** EDS results correspond to the points shown in Figure 10.


| Sample       |  | Elemental Ratio |            |           |
|--------------|--|-----------------|------------|-----------|
| <b>NFB1T</b> |  | <b>Nd%</b>      | <b>Fe%</b> | <b>B%</b> |
| <b>1</b>     |  | 30.08           | 69.92      | 0         |
| <b>2</b>     |  | 27.93           | 72.07      | 0         |
| <b>NFB2T</b> |  | <b>Nd%</b>      | <b>Fe%</b> | <b>B%</b> |
| <b>1</b>     |  | 26.86           | 73.14      | 0         |
| <b>2</b>     |  | 27.67           | 72.33      | 0         |
| <b>NFB3T</b> |  | <b>Nd%</b>      | <b>Fe%</b> | <b>B%</b> |
| <b>1</b>     |  | 24.95           | 75.05      | 0         |
| <b>2</b>     |  | 18.94           | 73.23      | 7.83      |

Figure 11 presents VSM hysteresis curves of NFB1T, NFB2T and NFB3T under a 2 T field. Annealing produced a clear coercivity improvement across all samples. For NFB2T, coercivity increased from 54.81 Oe to 107.77 Oe, an approximate doubling, attributed to crystallinity restoration and recovery of magnetocrystalline anisotropy in the Nd<sub>2</sub>Fe<sub>14</sub>B phase (Popa et al., 2013). Conversely, saturation magnetization decreased slightly from 133.43 to 128.90 emu/g, likely due to continued growth of non-magnetic Fe<sub>2</sub>O<sub>3</sub> during annealing, which dilutes the net magnetic moment. This coercivity-M<sub>s</sub> trade-off is a recognized behavior in mechanically alloyed Nd-Fe-B systems (Savchenko et al., 2018). NFB2T exhibits the highest coercivity among the annealed samples, while NFB3T records the highest M<sub>s</sub>, consistent with a greater  $\alpha$ -Fe proportion from 12 h milling. All annealed samples remain in the semi-hard magnetic regime, well below bulk sintered Nd-Fe-B values (>10,000 Oe), due to residual oxidation, secondary phase presence and isotropic grain orientation.

Figure 12 presents XRD patterns of NFB1TG, NFB2TG and NFB3TG following CVD encapsulation. A distinct carbon peak is observed at  $2\theta \approx 27^\circ$  in all three samples, corresponding to the (002) reflection of graphitic carbon, confirming successful carbonaceous layer deposition. The interlayer spacing associated with this peak is consistent with ordered graphitic stacking, suggesting a graphitic carbon shell rather than amorphous carbon. However, XRD alone is insufficient to conclusively determine graphene layer number or coating quality; Raman spectroscopy would be required for definitive assessment, an acknowledged limitation. The Nd<sub>2</sub>Fe<sub>14</sub>B tetragonal phase and  $\alpha$ -Fe peaks are retained, confirming that CVD at 950°C did not degrade the hard magnetic phase. Fe<sub>2</sub>O<sub>3</sub> peaks persist, indicating the pre-existing oxide layer was not reduced under the CH<sub>4</sub>/H<sub>2</sub> atmosphere.




**Figure 11.** Hysteresis curves and VSM analysis results of the samples: **(a)** NFB1T, **(b)** NFB2T and **(c)** NFB3T.




**Figure 12.** X-ray diffractometry results of the encapsulated samples produced from elemental raw materials: (a) NFB1TG, (b) NFB2TG and (c) NFB3TG.

Figure 13 presents SEM images of NFB1TG, NFB2TG and NFB3TG at 500 $\times$ , 1000 $\times$  and 1500 $\times$  magnifications. The agglomerated spherical morphology observed in the annealed samples (Figure 9) is retained after CVD, indicating that 950 $^{\circ}$ C treatment did not significantly alter particle arrangement or promote further sintering. The graphitic carbon shell is not resolvable at these magnifications; confirmation of the core-shell architecture relies on TEM analysis (Figure 16). The absence of observable morphological changes is consistent with a thin, conformal carbon coating that does not substantially modify particle dimensions, as expected for a thin, conformal graphitic carbon shell.

Figure 14 presents SEM elemental mapping of NFB1TG, NFB2TG and NFB3TG. In NFB1TG and NFB3TG, Nd, Fe, B and C show overlapping spatial distributions, indicating relatively uniform elemental co-distribution across agglomerate surfaces. In NFB2TG, Nd and Fe co-distribute as expected for the  $\text{Nd}_2\text{Fe}_{14}\text{B}$  phase, while C and B signals are concentrated in coincident regions, suggesting locally non-uniform carbon distribution at the mapping scale. The spatial correlation of the carbon signal with the Nd-Fe-B nanoparticle regions supports successful CVD deposition, consistent with XRD findings. Oxygen is present in all samples, confirming persistence of the  $\text{Fe}_2\text{O}_3$  surface oxide after encapsulation. Boron mapping should be considered semi-quantitative only (Goldstein et al., 2018).




**Figure 13.** SEM images of the graphitic carbon shell encapsulated samples at 500X, 1000X and 1500X magnifications: (a) NFB1TG, (b) NFB2TG and (c) NFB3TG.



**Figure 14.** SEM images and elemental mapping analysis results of the encapsulated samples: (a) NFB1TG, (b) NFB2TG and (c) NFB3TG.


Figure 15 presents the VSM hysteresis curve of NFB2TG. Coercivity increased to 185.71 Oe from 107.77 Oe (NFB2T), attributed to the secondary annealing effect of CVD at 950°C, which promotes additional  $\text{Nd}_2\text{Fe}_{14}\text{B}$  crystallization and enhances magnetocrystalline anisotropy (Popa et al., 2013). In contrast, saturation magnetization decreased substantially to 46.3 emu/g from 128.90 emu/g, a reduction of approximately 64%. Three contributing factors are proposed: the non-magnetic graphitic carbon shell adding sample mass without contributing to magnetization; further surface oxidation during CVD increasing  $\text{Fe}_2\text{O}_3$  content; and possible partial surface disordering of the  $\text{Nd}_2\text{Fe}_{14}\text{B}$  phase at elevated temperatures. The low  $M_r/M_s$  ratio of  $\sim 0.079$  confirms an isotropic, exchange-decoupled nanoparticle system with significant non-magnetic phase content. Although 185.71 Oe represents the highest coercivity in this study, it remains well below values reported for bulk sintered Nd-Fe-B magnets ( $>10,000$  Oe) and chemically synthesized  $\text{Nd}_2\text{Fe}_{14}\text{B}$  nanoparticles (500-3000 Oe) (Rahimi et al., 2017). NFB2TG therefore exhibits semi-hard magnetic behavior, making it suitable for applications requiring moderate coercivity with enhanced chemical stability.



**Figure 15.** Hysteresis curve and VSM analysis results of the NFB2TG.

Figure 16 presents the TEM image of NFB2TG, confirming the core-shell architecture in which the Nd-Fe-B nanoparticle core is surrounded by a thin graphitic carbon shell. The carbon layers appear as continuous lattice fringes at the particle periphery, indicating a well-defined interfacial structure and successful carbon deposition. The regularity of these fringes suggests ordered graphitic stacking rather than amorphous carbon, though Raman spectroscopy or High-

resolution transmission electron microscopy (HRTEM) analysis would be required to conclusively determine the layer number and quality, an acknowledged limitation. The graphitic carbon shell is expected to act as a physical barrier against further oxidation of the Nd-Fe-B core and to improve thermal and chemical stability, supporting the suitability of this encapsulation approach for electronic applications.



**Figure 16.** Transmission Electron Microscopy (TEM) image of the NFB2TG sample showing the core-shell nanostructure, with the Nd-Fe-B nanoparticle as the core surrounded by a thin graphitic carbon shell deposited via CVD.

#### 4. Discussion

Nd<sub>10</sub>Fe<sub>84</sub>B<sub>6</sub> powders were synthesized by MA and subsequently annealed and graphitic carbon shell encapsulated. The apparent increase in DLS-measured particle size with longer milling durations reflects progressive agglomeration rather than true grain growth. Finer particles produced at extended milling times have higher surface areas and greater reactivity, making

them more susceptible to oxidation upon air exposure, which enhances interparticle attractive forces and promotes stronger agglomeration (Mehrfar et al., 2025).

Despite progressive coercivity improvements across the as-milled (54.81 Oe), annealed (107.77 Oe) and encapsulated (185.71 Oe) states, values remain well below those of bulk sintered Nd-Fe-B magnets (>10,000 Oe). This is attributed to three concurrent factors: the isotropic grain orientation inherent to MA-produced powders, which prevents texture-driven anisotropy development (Coey, 2010); the presence of non-magnetic Fe<sub>2</sub>O<sub>3</sub> and soft magnetic  $\alpha$ -Fe secondary phases, which create domain reversal nucleation sites and reduce the effective anisotropy field; and nanoparticle surface effects including disordered surface spins that suppress bulk magnetocrystalline anisotropy. The superior performance of the 10 h milled sample is attributed to an optimal balance between sufficient mechanical energy input for homogeneous phase distribution and the avoidance of excessive structural damage, agglomeration and oxidation associated with 12 h milling (Popa et al., 2013).

The substantial  $M_s$  decrease from 128.90 emu/g (NFB2T) to 46.3 emu/g (NFB2TG) following CVD encapsulation is attributed to the non-magnetic carbon shell contributing to sample mass without contributing to magnetization, further surface oxidation during processing at 950°C and possible partial surface disordering of the Nd<sub>2</sub>Fe<sub>14</sub>B phase. Despite this reduction, the graphitic carbon shell provides a physical barrier against further oxidation and its secondary annealing effect improved crystallinity, reflected in the coercivity enhancement. This coercivity- $M_s$  trade-off is a recognized challenge in carbon-encapsulated magnetic nanoparticle systems (Hosseinabadi et al., 2019).

The main limitations of this work are: (1) the absence of Raman spectroscopy and XPS and coating quality, so the carbon shell is referred to as a graphitic carbon shell throughout; (2) DLS reflects agglomerate hydrodynamic size rather than primary particle dimensions; (3) surface oxidation was not directly quantified by thermogravimetric or XPS analysis; and (4) only a single TEM image was presented, limiting statistical representativeness of the core-shell structure assessment.

## 5. Conclusion

Nd-Fe-B nanoparticles with the Nd<sub>10</sub>Fe<sub>84</sub>B<sub>6</sub> composition were successfully synthesized from elemental raw materials via mechanical alloying and subsequently annealed and encapsulated with a graphitic carbon shell via SCVD. Comprehensive characterization by XRD, DLS, SEM, EDS, VSM and TEM confirmed the formation of the Nd<sub>2</sub>Fe<sub>14</sub>B tetragonal phase, with annealing at 900°C improving crystallinity and magnetic hardness. Among all samples, NFB2TG, produced by 10 h milling, annealing and encapsulation, exhibited the most favorable magnetic properties, with a coercivity of 185.71 Oe and a saturation magnetization of 46.3 emu/g. TEM confirmed a core-shell nanostructure with a continuous graphitic carbon layer surrounding the magnetic core. While coercivity values remain below those of bulk permanent magnets due to isotropic grain orientation, secondary phase content and nanoparticle surface effects inherent to the MA process, the combined approach of MA, annealing and CVD encapsulation

demonstrates a viable route for producing chemically stable Nd-Fe-B nanoparticles with semi-hard magnetic behavior, suitable for MEMS devices and magnetically recoverable systems.

Future work should focus on: (1) Raman spectroscopy and XPS to verify graphitic carbon shell quality and layer number; (2) TEM-based primary particle size analysis to distinguish primary particle dimensions from agglomerate sizes; (3) direct oxidation quantification by thermogravimetric analysis; (4) optimization of CVD parameters to minimize Ms loss during encapsulation; and (5) magnetic field-assisted annealing to introduce crystallographic texture and improve coercivity toward practical permanent magnet performance levels.

## **Conflict of Interest**

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## **Funding**

This research was supported by TENMAK National Boron Research Institute (BOREN) under the 2018 BK-1 Bor Scholarship Project and TÜBİTAK through the 1001 Project (Project Code: 118F430).

## **Acknowledgments**

The authors would like to thank TENMAK National Boron Research Institute (BOREN) for supporting this thesis within the scope of the 2018 BK-1 Bor Scholarship Project. The authors also thank ASELSAN Electronics Industry Inc. for their support of this work within the framework of the TOHUM Project. Additionally, the budget of the TÜBİTAK 1001 project (Project Code: 118F430) was used to cover consumables and service costs for this thesis. The authors appreciate TÜBİTAK and their advisors for their contributions and guidance in deriving the thesis topic from this project. The authors also thank their colleagues Mustafa İlgen Kızılöz and Ayşegül Afal Geniş for their support during this work.

## **Ethical Statements**

This study did not involve human or animal subjects. All experiments were conducted following institutional safety and laboratory regulations.

## **Data and Code Availability**

All experimental data generated and analyzed during this study are available from the corresponding author upon reasonable request.

## **Supplementary Materials**

Not applicable.

## References

- Ağaoğulları, D., Gökçe, H., Duman, İ., & Öveçoğlu, M. L. (2011). Characterization investigations of ZrB<sub>2</sub>/ZrC ceramic powders synthesized by mechanical alloying of elemental Zr, B and C blends. *Journal of the European Ceramic Society*, 32(7), 1447-1455. <https://doi.org/10.1016/j.jeurceramsoc.2011.04.026>
- Ağaoğulları, D., Madsen, S. J., Ögüt, B., Koh, A. L., & Sinclair, R. (2017). Synthesis and characterization of graphite-encapsulated iron nanoparticles from ball milling-assisted low-pressure chemical vapor deposition. *Carbon*, 124, 170-179. <https://doi.org/10.1016/j.carbon.2017.08.043>
- Bystrzejewski, M., Lange, H., & Huczko, A. (2007). Carbon encapsulation of magnetic nanoparticles. *Fullerenes, Nanotubes and Carbon Nanostructures*, 15(3), 167-180. <https://doi.org/10.1080/15363830701236357>
- Chakka, V. M., Altuncevahir, B., Jin, Z. Q., Li, Y., & Liu, J. P. (2006). Magnetic nanoparticles produced by surfactant-assisted ball milling. *Journal of Applied Physics*, 99(8), Article 08E912. <https://doi.org/10.1063/1.2170593>
- Chen, W., Kim, J., Sun, S., & Chen, S. (2007). Composition effects of FePt alloy nanoparticles on the electro-oxidation of formic acid. *Langmuir*, 23(22), 11303-11310. <https://doi.org/10.1021/la7016648>
- Coey, J. M. D. (2010). *Magnetism and magnetic materials*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511845000>
- De Moraes, I. G., Fischbacher, J., Hong, Y., Naud, C., Okuno, H., Masseboeuf, A., Devillers, T., Schrefl, T., & Dempsey, N. M. (2024). Nanofabrication, characterisation and modelling of soft-in-hard FeCo-FePt magnetic nanocomposites. *Acta Materialia*, 274, Article 119970. <https://doi.org/10.1016/j.actamat.2024.119970>
- Dempsey, N. M. (2009). Hard magnetic materials for MEMS applications. In J. P. Liu, E. Fullerton, O. Gutfleisch, & D. J. Sellmyer (Eds.), *Nanoscale magnetic materials and applications* (pp. 661-683). Springer. [https://doi.org/10.1007/978-0-387-85600-1\\_22](https://doi.org/10.1007/978-0-387-85600-1_22)
- Goldstein, J. I., Newbury, D. E., Michael, J. R., Ritchie, N. W. M., Scott, J. H. J., & Joy, D. C. (2018). *Scanning electron microscopy and X-ray microanalysis* (4th ed.). Springer. <https://doi.org/10.1007/978-1-4939-6676-9>
- Gubin, S. P., Spichkin, Y. I., Yurkov, G. Y., & Tishin, A. M. (2002). Nanomaterial for high-density magnetic data storage. *Russian Journal of Inorganic Chemistry*, 47. <https://amtc.ru/publications/articles/5rus.pdf>
- Hosseinabadi, S., Jafari, M. J., Kokabi, M., & Mohseni, M. (2019). Improving the electromagnetic shielding of fabricated Nd-Fe-B particles by a coating thin carbonaceous

- p layer.
- Chemical Physics Letters*
- , 739, Article 137015.
- 
- <https://doi.org/10.1016/j.cplett.2019.137015>
- Iacovacci, V., Lucarini, G., Innocenti, C., Comisso, N., Dario, P., Ricotti, L., & Mencias, A. (2015). Polydimethylsiloxane films doped with Nd-Fe-B powder: Magnetic characterization and potential applications in biomedical engineering and microrobotics. *Biomedical Microdevices*, 17(6), Article 112. <https://doi.org/10.1007/s10544-015-0024-0>
- Jia, Z., Li, J., Gao, L., Yang, D., & Kanaev, A. (2023). Dynamic light scattering: A powerful tool for in situ nanoparticle sizing. *Colloids and Interfaces*, 7(1), Article 15. <https://doi.org/10.3390/colloids7010015>
- Jung, Y., Lee, Y., Yoon, S., & Choi, J. (2024). Synergistic effect of core/shell-structured composite fibers: Efficient recovery of rare-earth elements from spent Nd-Fe-B permanent magnets. *Advanced Fiber Materials*, 6(6), 1729-1745. <https://doi.org/10.1007/s42765-024-00442-4>
- Lee, S., Jeong, J., Shin, S., Kim, J., Chang, Y., Lee, K., & Kim, J. (2005). Magnetic enhancement of iron oxide nanoparticles encapsulated with poly(d,l-lactide-co-glycolide). *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 255(1-3), 19-25. <https://doi.org/10.1016/j.colsurfa.2004.12.019>
- Lecerf, I., Angulo-Cervera, J. E., Orlandini-Keller, F., Moritz, P., Mathieu, F., Bourrier, D., Charlot, S., Nicu, L., Leïchl  , T., Devillers, T., Haettel, R., Dempsey, N. M., Blon, T., & Lacroix, L. (2025). A MEMS electromagnetic vibration energy harvester with monolithically integrated Nd-Fe-B micromagnets. *Advanced Materials Technologies*, 10(10), Article 2401817. <https://doi.org/10.1002/admt.202401817>
- Lim, J., Yeap, S. P., Che, H. X., & Low, S. C. (2013). Characterization of magnetic nanoparticle by dynamic light scattering. *Nanoscale Research Letters*, 8, Article 381. <https://doi.org/10.1186/1556-276X-8-381>
- Liu, Y. D., Choi, H. J., & Choi, S. (2012). Controllable fabrication of silica encapsulated soft magnetic microspheres with enhanced oxidation-resistance and their rheology under magnetic field. *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, 403, 133-138. <https://doi.org/10.1016/j.colsurfa.2012.04.002>
- Manaf, A., Al-Khafaji, M., Zhang, P. Z., Davies, H. A., Buckley, R. A., & Rainforth, W. M. (1993). Microstructure analysis of nanocrystalline Fe-Nd-B ribbons with enhanced hard magnetic properties. *Journal of Magnetism and Magnetic Materials*, 128(3), 307-312. [https://doi.org/10.1016/0304-8853\(93\)90476-1](https://doi.org/10.1016/0304-8853(93)90476-1)
- Mehrifar, Y., Moqtaderi, H., Hamidi, S. M., Golbabaie, F., Hasanzadeh, M., & Dehghan, S. F. (2025). Magnetic nanoparticles of Nd<sub>2</sub>Fe<sub>14</sub>B prepared by ethanol-assisted wet ball milling technique. *Scientific Reports*, 15(1), Article 3257. <https://doi.org/10.1038/s41598-025-87301-3>

- Omelyanchik, A., Lamura, G., Peddis, D., & Canepa, F. (2020). Optimization of a Nd-Fe-B permanent magnet configuration for in-vivo drug delivery experiments. *Journal of Magnetism and Magnetic Materials*, 522, Article 167491. <https://doi.org/10.1016/j.jmmm.2020.167491>
- Ouyang, Y., Qiu, R., Xiao, Y., Shi, Z., Hu, S., Zhang, Y., Chen, M., & Wang, P. (2019). Magnetic fluid based on mussel inspired chemistry as corrosion-resistant coating of Nd-Fe-B magnetic material. *Chemical Engineering Journal*, 368, 331-339. <https://doi.org/10.1016/j.cej.2019.02.126>
- Périgo, E. A., Silva, S. C., De Sousa, E. M. B., Freitas, A. A., Cohen, R., Nagamine, L. C. C. M., Takiishi, H., & Landgraf, F. J. G. (2012). Properties of nanoparticles prepared from Nd-Fe-B-based compound for magnetic hyperthermia application. *Nanotechnology*, 23(17), Article 175704. <https://doi.org/10.1088/0957-4484/23/17/175704>
- Pich, A., Bhattacharya, S., Ghosh, A., & Adler, H. (2005). Composite magnetic particles: 2. Encapsulation of iron oxide by surfactant-free emulsion polymerization. *Polymer*, 46(13), 4596-4603. <https://doi.org/10.1016/j.polymer.2005.01.017>
- Popa, O., Dorolti, E., Borodi, G., Marcu, G., Coldea, M., & Turcu, R. (2013). The influence of milling and annealing on the structural and magnetic behavior of Nd<sub>2</sub>Fe<sub>14</sub>B/ $\alpha$ -Fe magnetic nanocomposite. *Journal of Alloys and Compounds*, 574, 496-501.
- Rahimi, H., Ghasemi, A., Mozaffarinia, R., & Tavoosi, M. (2017). Magnetic properties and magnetization reversal mechanism of Nd-Fe-B nanoparticles synthesized by a sol-gel method. *Journal of Magnetism and Magnetic Materials*, 444, 111-118. <https://doi.org/10.1016/j.jmmm.2017.08.011>
- Saritas, E. U., Goodwill, P. W., Croft, L. R., Konkole, J. J., Lu, K., Zheng, B., & Conolly, S. M. (2012). Magnetic particle imaging (MPI) for NMR and MRI researchers. *Journal of Magnetic Resonance*, 229, 116-126. <https://doi.org/10.1016/j.jmr.2012.11.029>
- Savchenko, A. G., Menushenkov, V. P., Plastinin, A. Y., Zhukov, A. A., & Kuznetsov, P. A. (2018). Phase composition and magnetic properties of Nd<sub>2</sub>Fe<sub>14</sub>B/ $\alpha$ -Fe nanocomposites prepared by mechanical alloying. *Russian Metallurgy (Metally)*, 2018(4), 354-358. <https://doi.org/10.1134/S0036029518040134>
- Schultz, L., Wecker, J., & Hellstern, E. (1987). Formation and properties of Nd-Fe-B prepared by mechanical alloying and solid-state reaction. *Journal of Applied Physics*, 61(8), 3583-3585. <https://doi.org/10.1063/1.338708>
- Sepehri-Amin, H., Ohkubo, T., Gruber, M., Schrefl, T., & Hono, K. (2014). Micromagnetic simulations on the grain size dependence of coercivity in anisotropic Nd-Fe-B sintered magnets. *Scripta Materialia*, 89, 29-32. <https://doi.org/10.1016/j.scriptamat.2014.06.020>
- Shen, L., Fan, M., Zhao, K., Qiu, M., & Tian, Z. (2018). Preparation and properties of nanocomposite coatings on sintered Nd-Fe-B magnets. *Materials Research Express*, 5(8), Article 086401. <https://doi.org/10.1088/2053-1591/aad121>

- Suryanarayana, C. (2010). Synthesis of nanocomposites by mechanical alloying. *Journal of Alloys and Compounds*, 509(Suppl. 1), S229-S234. <https://doi.org/10.1016/j.jallcom.2010.09.063>
- Tung, D. K., Manh, D. H., Phong, P. T., Phong, L. T. H., Dai, N. V., Nam, D. N. H., & Phuc, N. X. (2015). Structural and magnetic properties of mechanically alloyed Fe<sub>50</sub>Co<sub>50</sub> nanoparticles. *Journal of Alloys and Compounds*, 640, 34-38. <https://doi.org/10.1016/j.jallcom.2015.04.022>
- Wu, K., Yao, Y., & Klik, I. (1997). Electrical and magnetic properties of Nd-Fe-B films. *Applied Surface Science*, 113-114, 174-177. [https://doi.org/10.1016/s0169-4332\(96\)00843-4](https://doi.org/10.1016/s0169-4332(96)00843-4)
- Yang, C., Ren, B., Wang, D., & Tang, Q. (2019). Synthesis of Nano-Fe@Nd-Fe-B/AC magnetic catalytic particle electrodes and application in the degradation of 2,4,6-trichlorophenol by electro-assisted peroxydisulfate process. *Environmental Technology*, 41(19), 2464-2477. <https://doi.org/10.1080/09593330.2019.1567826>
- Yang, H., Duan, L., Zhang, P., Xu, G., Cui, J., Lv, J., Sun, W., Li, B., Wang, D., & Wu, Y. (2022). Corrosion resistance of functionalized carbon nanotubes enhanced epoxy coatings on sintered Nd-Fe-B magnets. *Journal of Coatings Technology and Research*, 19(5), 1317-1329. <https://doi.org/10.1007/s11998-022-00641-x>
- Zhang, J., Wu, W., Meng, F., Ding, H., & Dong, J. (2018). Sol-gel-based chemical synthesis of Nd-Fe-B hard magnetic nanoparticles. *Modern Physics Letters B*, 32(34n36), Article 1840070. <https://doi.org/10.1142/s0217984918400705>
- Zhou, X., Tian, Y., Yu, H., Zhang, H., Zhong, X., & Liu, Z. (2019). Synthesis of hard magnetic Nd-Fe-B composite particles by recycling the waste using microwave assisted auto-combustion and reduction method. *Waste Management*, 87, 645-651. <https://doi.org/10.1016/j.wasman.2019.02.050>