

Charting the Trajectory of Artificial Intelligence Agents in Air Traffic Control

Seda ÇEKEN^{1*}, Arif TUNCAL²

^{1*} Istanbul University, Institute for Aviation Psychology Research, Department of Aviation Psychology, İstanbul, Türkiye. sedaceken@istanbul.edu.tr
ORCID: 0000-0002-5870-2246

² International Science and Technology University, Department of Aviation Systems and Technologies, Warsaw, Poland. arif.tuncal@istu.edu.pl,
atuncal@gmail.com ORCID: 0000-0003-4343-6261

Article Info

Received: 13.06.2026

Revised: 18.09.2026

Accepted: 20.09.2026

Published: 15.10.2026

Keywords:

Artificial intelligence, AI agent,
Bibliometric analysis, Multi-agent,
Reinforcement learning, Human-AI.

Jel Code:

L93, O32, O33, R41

Corresponding Author:

Seda ÇEKEN

Research Article

<https://doi.org/10.5281/zenodo.23372598>

Abstract

The aim of the study was to chart the evolution and the current research landscape of artificial intelligence (AI) agents within the domain of air traffic control. Using bibliometric analysis, the study identified major research trends and gaps and provided a consolidated reference for researchers. The analysis included 58 publications retrieved from the Web of Science Core Collection. The findings showed that the earliest publication in the dataset appeared in 1998 and that research output increased noticeably after 2022. Publication output peaked in 2025, with 16 articles, indicating growing scholarly interest in AI agents for air traffic control. Citation analysis identified Agogino and Tumer's (2012) article as the most influential publication, with 68 citations. The article introduced a multi-agent approach to air traffic flow management. Brittain and Wei's (2022) article ranked second, with 54 citations, highlighting the increasing role of reinforcement learning. At the country level, the United States had the highest total citation count, followed by China, France, England, and Spain. By examining AI agents as a distinct research topic, the study provides a focused overview of an underexplored area relevant to researchers, practitioners, and regulators.

Introduction

Global aviation relies on air traffic controllers whose cognitive performance is essential to operational efficiency and safety. As traffic demand increases, maintaining this performance becomes more challenging. According to data released by the International Air Transport Association (IATA) (2026), the year 2025 witnessed record-high global aviation demand with a 5.3% increase in total revenue passenger kilometers (RPKs) and a 5.2% rise in capacity. International traffic increased by 7.1%, while the global passenger load factor reached a record 83.6% (IATA, 2026). Rising traffic demand may increase operational complexity and controller workload across routes and sectors. To support human performance and maintain operational safety, the aviation sector is increasingly exploring artificial intelligence (AI)-based solutions.

This growing interest is also reflected in industry investment. The integration of AI into aviation has attracted substantial investment and accelerated the shift toward advanced automation. The global market for AI in aviation was valued at approximately USD 7.45 billion in 2025. It is projected to reach USD 36.68 billion by 2034, corresponding to a compound annual growth rate of 19.48% (Fortune Business Insights, 2026). This projected growth reflects increasing efforts to improve operational efficiency, including the use of algorithms in air traffic management (ATM). AI-based architectures may support ATM capacity through real-time data processing, predictive traffic modeling, and proactive conflict resolution.

AI has become an important part of the air traffic control modernization agenda. This development is supported by the Single European Sky ATM Research (SESAR) program, a major European initiative for ATM modernization. Recent studies show that SESAR projects increasingly examine the integration of AI and machine learning into key ATM functions (Çeken, 2025; Tuncal, 2025). Autonomous AI agents are increasingly being investigated as tools for reducing controller workload, improving safety, and optimizing traffic flow across different airspace systems (Ali et al., 2025; Brittain & Wei, 2019; Brittain & Wei, 2022; Mohamed et al., 2026; Wang et al., 2019). Furthermore, the operational integration of these autonomous agents relies heavily on several related socio-technical concepts. Human-AI teaming refers to collaboration between human operators and intelligent systems in the management of safety-critical tasks (Kirwan, 2025). Explainable AI (XAI) aims to make the reasoning of complex learning algorithms more transparent and

verifiable for human controllers. This transparency may support appropriate trust in AI-based systems (Degas et al., 2022; Mohamed et al., 2026).

The growth of this research agenda has attracted increasing attention from the bibliometric community. In recent years, many mapping studies have explored the landscape of AI in aviation and transportation. Chen et al. (2023) mapped the knowledge diffusion and network characteristics of AI applications specifically in traffic flow prediction. Demir (2026) analyzed a decade of literature to highlight the integration of autonomous systems and the critical need for sustainable, ethical AI. Demir et al. (2024) reviewed AI's role in aviation safety, identifying key trends in accident analysis and pilot behavior. Similarly, Ertekin et al. (2024) mapped scientific collaborations to underscore the central role of machine learning in improving both safety and operational efficiency. Karakavuz (2025) categorized the intellectual structure of the air transportation sector, distinguishing between core foundations like deep learning and driving themes such as air traffic control and unmanned aerial vehicles. Lopes et al. (2025) provided a comprehensive mapping that culminated in a novel conceptual framework spanning human factors, environmental initiatives, and flight operations. Sadou and Njoya (2023) identified five dominant research clusters in air transport, revealing that prediction and optimization account for the vast majority of current studies. Şahin (2024) evaluated recent publication metrics, demonstrating that the performance-enhancing impact of AI has become the primary focus of aviation literature. Despite this growing literature, a structural gap remains. Most of these studies treat AI in aviation or transportation as a broad, aggregate category. With the exception of Sadou and Njoya (2023), who briefly mention the "agent" concept without incorporating it into their analytical framework, existing mapping studies largely overlook this specific domain. The AI agent paradigm relies on unique theoretical features, such as autonomous perception, goal-directed reasoning, reinforcement learning, and multi-agent coordination. Agent-based approaches are fundamentally different from conventional machine learning (Abou Ali et al., 2025). They have distinct operational logic, safety/security implications, and regulatory needs (Abou Ali et al., 2025; Adabara et al., 2025; Biswas & Sarkar, 2026; Raza et al., 2025). Today, these systems are moving from academic simulations to real-world applications. A clear understanding of the research community developing these systems is therefore needed. Mapping its collaborative architecture, thematic clusters, and emerging frontiers is essential for coherent scientific planning and governance. The existing literature lacks a systematic bibliometric evaluation dedicated to the AI agent paradigm.

The study addresses the gap through a focused bibliometric analysis of research at the intersection of air traffic control and AI agents. The findings provide a structured overview of an emerging and underexplored field for researchers, practitioners, and regulators. To establish a theoretical foundation for the analysis, the study initially reviews the extant literature to conceptualize the core elements.

Literature Review

AI agents and conceptual foundations

The transition from specialized rule-based programs to dynamic autonomous systems has brought the concept of the AI agent to the forefront of contemporary technological research (Qu et al., 2025). Broadly defined, an AI agent is a computational system capable of autonomously perceiving its environment, reasoning over contextual information, and taking goal-directed actions without continuous human intervention (Krishnan, 2025; Sapkota et al., 2025). The concept is based on the perception-action loop, in which an agent receives information from its environment and responds through appropriate actions. Modern AI agents may also integrate large language models (LLMs) and large image models (LIMs). These models can support complex reasoning, decision-making, and planning (Sapkota et al., 2025). The concept of "agency" in this context refers to the degree to which an AI system can directly impact the world and achieve long-horizon goals with minimal external specification on how to do so (Chan et al., 2024).

Modern AI agents rely on several fundamental capabilities, including perception, reasoning, memory, and tool use. Perception mechanisms form the interface between an agent and its environment, allowing the system to process unstructured inputs such as text and visual data. After receiving input, the reasoning and decision-making modules evaluate available alternatives and select an appropriate action. This process may draw on the inference capabilities of foundation models (Krishnan, 2025). Furthermore, because foundational models inherently suffer from constraints like static knowledge and limited computational precision, agents are augmented with "tools" or external application programming interfaces (APIs). These tools enable agents to execute operations such as web searches, code interpretation, or calendar management, fundamentally transforming passive language models into interactive problem-solvers (Qu et al., 2025; Sapkota et al., 2025). To perform these operations, agents rely on supporting structures known as scaffolding. This concept involves

structural methodologies like strategic prompting, systematic planning routines, and memory frameworks that help them achieve complex goals (Chan et al., 2024).

Memory enables AI agents to preserve context and learn across extended interactions (Sapkota et al., 2025). Agentic memory systems are generally categorized into short-term (or working) memory, which maintains task-relevant information during active execution, and long-term (or parametric and episodic) memory, which stores interaction histories, factual knowledge, and procedural skills for future retrieval (Krishnan, 2025; Qu et al., 2025). Without these memory architectures, agents would be relegated to stateless, single-session interactions; with them, agents can reflect on past actions, recognize patterns, and iteratively refine their behaviors through self-evaluation and feedback loops (Nisa et al., 2025). Combined with dynamic memory access, this reflective capability can help agents recover from errors and maintain coherence during complex multi-step tasks (Sapkota et al., 2025).

As the complexity of tasks escalates, the paradigm shifts from isolated, single-agent systems to the broader concept of "Agentic AI". While an AI agent typically refers to an individual entity executing a narrowly defined task with external tools, Agentic AI refers to broader collaborative and distributed systems (Jiang et al., 2025; Sapkota et al., 2025). In an Agentic AI framework, a complex goal is dynamically decomposed into smaller sub-tasks, which are then assigned to an ensemble of specialized sub-agents acting cooperatively (Sapkota et al., 2025; Zhang et al., 2026). This orchestration is managed by meta-agents or central controllers that regulate inter-agent communication, resolve conflicts, and ensure that the collective actions align with the overarching objective (Sapkota et al., 2025). Agentic AI encapsulates not just individual autonomy, but system-level social ability, negotiation, and emergent intelligence, moving beyond basic reactive automation toward coordinated, proactive problem-solving (Nisa et al., 2025; Sapkota et al., 2025). The classification and evaluation of AI agents and Agentic AI systems are fundamentally tied to their levels of autonomy and goal complexity (Kasirzadeh & Gabriel, 2025). Autonomy, originating conceptually from philosophical notions of negative and positive freedom (Berlin, 2017), measures the capacity of the AI to perform actions without human direction or control. This spectrum ranges from restricted autonomy, where humans closely supervise automated tasks, to full autonomy, where systems operate independently in dynamic environments. Similarly, the sophistication of an agent is defined by its goal complexity, representing the shift from simple stimulus-response behaviors to the unbounded ability to generate, interpret, and decompose underspecified objectives (Kasirzadeh & Gabriel, 2025). Together, these concepts describe the transition from automated software tools to reasoning-capable, adaptable, and collaborative digital systems (Krishnan, 2025).

AI agents in air traffic control

Modernizing air traffic control systems is an important challenge in contemporary aviation, particularly because of growing traffic demand and the integration of new aerial vehicles. To manage this growing complexity, the aviation sector is increasingly examining AI-based methods that may reduce controller workload and support the safe and efficient management of aircraft. In this domain, automated decision-support systems and intelligent algorithms are generally classified into four primary categories based on their underlying operational logic: 1) Rule-based agents, 2) Optimization agents, 3) Reinforcement learning agents, and 4) Hybrid agents (Mohamed et al., 2026). Rule-based agents use predefined expert logic, decision trees, and heuristics to generate clearance recommendations. Although these agents are generally interpretable and consistent with standard operating procedures, they may be too rigid for unexpected or highly complex traffic situations. The Hawk agent is an example of this approach. It uses a decision tree to apply safety and conflict-resolution rules derived from expert interviews. Optimization agents, conversely, plan trajectories dynamically by treating the airspace as a mathematical problem subject to hard constraints, such as sector containment and safe separation limits. The Falcon agent illustrates the optimization approach. It uses algorithmic search to recalculate efficient trajectories while maintaining safety constraints (Carvell, Thomas, et al., 2026). Reinforcement learning agents use a different approach based on repeated interaction with an environment and feedback from a reward function. Hybrid agents combine different architectures. Their purpose is to integrate the safety constraints of explicit rules with the adaptability of machine-learning methods (Mohamed et al., 2026).

Within the realm of learning algorithms, reinforcement learning and multi-agent reinforcement learning (MARL) have received increasing research attention for resolving both local conflict detection and broader demand-capacity balancing problems. In highly dynamic en-route sectors, individual aircraft or traffic management units are modeled as autonomous agents that must navigate stochastic environments. For instance, the K-Control Actor-Critic algorithm processes two-dimensional continuous action spaces to generate precise heading changes, effectively resolving potential collisions while minimizing the required

number of tactical interventions (Wang et al., 2019). When scaled to high-density airspaces with multiple intersecting routes, deep distributed MARL (DD-MARL) frameworks, utilizing simulators like BlueSky, have successfully demonstrated the ability to decentralize policy execution and resolve nearly all conflicts autonomously (Brittain & Wei, 2019). The success of these agents mirrors the breakthroughs seen in other complex AI domains, such as DeepMind's AlphaGo (Silver et al., 2016). To comprehensively prevent MARL policies from overfitting to specific, rigid airspace geometries, researchers have introduced polar state representations and transformer-based encoder architectures; by encoding ownship and intruder data as computational tokens, these models adapt robustly to unstructured airspaces and varying intersection angles, successfully outputting speed advisories that prevent near mid-air collisions (Aziz & Wei, 2026).

As these multi-agent systems evolve, they must accommodate the operational realities of the aviation industry, which involves competing airlines and heterogeneous objectives. Advanced frameworks, such as the Multi-Agent Autonomous Separation Assurance (MAASA) model, specifically address heterogeneous environments where agents optimize entirely different, often proprietary, reward functions. By employing intention-learning mechanisms, these agents can infer the planned maneuvers of intruder aircraft without requiring explicit knowledge of their underlying algorithms, facilitating cooperation across diverse commercial entities (Brittain & Wei, 2022). This cooperative capability reflects earlier theoretical models of principled negotiation, wherein intelligent agents utilize overlapping declarative functions to optimize flight operations through mutually beneficial trajectory proposals (Wangermann & Stengel, 1998). In the context of the ground holding problem and air holding problem, the careful design of reward functions is equally critical; RL models that balance safety separation imperatives with financial fairness among airlines have proven highly effective in mimicking human traffic manager decisions (Cruciol et al., 2013). The use of difference rewards may reduce localized congestion more effectively than simple global reward signals (Agogino & Tumer, 2012). To manage the immense computational complexity of these tasks, hierarchical MARL schemes dynamically partition state and action spaces into multiple levels of abstraction, allowing agents to efficiently adhere to airspace capacity constraints while reducing overarching operational delays (Spatharis et al., 2023).

Despite their potential, unconstrained RL and MARL policies may generate outputs that are inconsistent with operational requirements. Standard RL agents tend to issue high-frequency, micro-adjustments that induce severe cognitive load on pilots and congest radio communication frequencies. To address this problem, online action-stacking can be used as an inference-time wrapper. It combines multiple primitive actions into a single clearance that is more consistent with operational practice (Carvell, De Ath, et al., 2026). Furthermore, to ensure that the strategic preferences of AI agents align with expert behavior, intelligent interactive conflict solvers have been deployed to digitally capture human demonstrations during separation tasks; by integrating these human-provided trajectory change points into the reward signal, RL agents learn to issue resolutions that closely mirror human cognitive strategies (Tran et al., 2020). Understanding the behavioral dynamics of these interactions is equally critical. Recent large-scale experimental evidence highlights how AI fundamentally reshapes teamwork. Compared to traditional human-human pairs, collaboration with AI agents fosters highly task-oriented communication and significantly increases task delegation. This dynamic enhances overall productivity and specific quality metrics, although it can introduce a risk of output homogeneity (Ju & Aral, 2025). Within complex operational environments, these underlying mechanisms of delegation and task-orientation become vital for managing human cognitive capacity.

These operational requirements have increased interest in human-AI teaming and human-AI hybrid systems. A human-AI hybrid combines human judgment, ethical oversight, and situational adaptability with the computational capabilities of AI (Ali et al., 2025). Regulatory frameworks, such as those established by the European Union Aviation Safety Agency (EASA), classify these interactions into distinct levels of AI autonomy, ranging from simple machine learning support to highly collaborative agents that can dynamically negotiate tasks with human operators. The Horizon Europe HAIKU project has examined these concepts through several applied case studies. The FOCUS (Flight Operational Companion for Unexpected Situations) system acts as an intelligent assistant in the cockpit, utilizing psychophysiological sensors and eye-tracking to detect a pilot's startle response during sudden emergencies, dynamically highlighting critical flight displays to accelerate the recovery of situational awareness. The COMBI agent serves as a collaborative teammate for flight crews, instantly computing optimal diversion airports during severe weather events by balancing fuel constraints and passenger connections, reducing a 15-minute manual task to less than a minute. Similarly, the Intelligent Sequence Assistant (ISA) operates in air traffic control towers, using neural networks to optimize runway utilization by dynamically resequencing arriving and departing aircraft based on live radar data (Kirwan, 2025).

The operational viability of human-AI teaming is intrinsically dependent on the systemic transparency of the AI, elevating the necessity of XAI. For controllers to rely appropriately on AI-supported tactical recommendations, the system must communicate the basis of its decisions clearly. Comprehensive XAI models in aviation utilize a Descriptive, Predictive, and Prescriptive (DPP) framework. Descriptive XAI elucidates the algorithmic logic behind a specific decision; predictive XAI anticipates how changes to inputs will affect future trajectories; and prescriptive XAI identifies systemic flaws and suggests actionable remedies to mitigate errors (Degas et al., 2022). While traditional rule-based systems inherently possess high interpretability through auditable decision trees, complex deep reinforcement learning agents require sophisticated post-hoc analytical techniques. The mathematical computation of Shapley values allows developers to isolate and visualize the specific state features that drive a neural network's clearance recommendation. For instance, they can identify how the conflicting altitude of an intruder affects the outcome. This process establishes a verifiable logic for decisions that would otherwise remain opaque (Mohamed et al., 2026). Additionally, sophisticated visual analytics tools are heavily integrated into MARL pipelines to render the dense decision-making outputs interpretable, allowing human supervisors to visually inspect how large swarms of agents distribute delays and resolve localized demand-capacity hotspots (Kravaris et al., 2023).

Before any intelligent agent can be safely deployed, it must undergo exhaustive human-in-the-loop testing within simulated environments that faithfully replicate real-world uncertainties. Within the UK's Project Bluebird, researchers have engineered BluebirdDT, a probabilistic digital twin of the London Area Control Centre's en-route airspace. By fusing live surveillance feeds, historic flight plans, meteorological data, and physics-informed machine learning models, this digital twin accurately models both epistemic and aleatoric uncertainties, such as wind speed variability and pilot response delays (Pepper et al., 2026). Within this simulated environment, Agent Mallard operates as a forward-planning, rule-based agent designed for systemized airspace. Agent Mallard follows a 'simulate before you commit' approach. Using a depth-limited backtracking search, the agent simulates each proposed tactical resolution in the digital twin before implementation to assess whether it could create secondary conflicts (Kent et al., 2026). Furthermore, to rigorously validate these agents, they are exposed to the exact regulator-certified formative and summative assessment scenarios utilized in actual human controller training curricula, ensuring that machine competence meets strict operational safety standards (Carvell, Thomas, et al., 2026).

The successful integration of autonomous agents into air traffic control extends beyond pure technical validation; it is profoundly influenced by public perception, algorithm aversion, and the complex psychological dynamics of trust and blame attribution. Empirical analyses comparing autonomous vehicles and air traffic control environments reveal a pronounced cognitive bias among the general public, who tend to attribute significantly greater foundational trust to human operators while disproportionately assigning blame to AI agents for equivalent operational errors. Notably, within an AI trainee model configuration, the human supervisor often absorbs substantially more blame for failing to intervene than the AI itself, highlighting the complex socio-legal nuances of shared responsibility in human-AI interactions (Mei et al., 2025). The modernization of air traffic control is a multidisciplinary process. It requires reliable multi-agent learning, transparent XAI interfaces, human-in-the-loop validation, and careful consideration of human-machine interaction.

Method

The study employs a bibliometric analysis to map the intellectual structure, thematic evolution, and collaborative architecture of research on AI agents in air traffic control. Bibliometric analysis is a well-established quantitative methodology that enables systematic examination of scientific literature through the measurement of publication outputs, citation patterns, co-authorship networks, and keyword co-occurrence structures (Donthu et al., 2021). Its application in the present context is particularly appropriate given the interdisciplinary and rapidly evolving nature of the field, where large volumes of dispersed publications cannot be adequately synthesized through conventional narrative review.

Research purpose

The study aims to provide a comprehensive bibliometric account of scholarly work at the intersection of AI agents and air traffic control, with the goal of identifying key publication trends, influential contributors, leading countries, and emerging thematic clusters. The findings are intended to serve as a structured reference for researchers entering this domain and to inform research prioritization efforts by highlighting areas of intellectual concentration as well as gaps that remain underexplored. Specifically, the study seeks to answer the following research questions:

1. How has the annual volume of publications on AI agents in air traffic control evolved over time, and what temporal patterns characterize the field's development?
2. What are the collaboration structures among authors, and which scholars occupy central positions within the co-authorship network?
3. Which publications and countries have accumulated the highest citation counts, and what do these patterns reveal about the field's intellectual foundations and geographic distribution?
4. Which authors share the strongest bibliographic similarities, and what research clusters emerge from bibliometric coupling analysis?
5. Which keywords appear most frequently across the corpus, and what thematic clusters and conceptual relationships do co-occurrence analyses reveal?

Data collection

Data were retrieved from the Web of Science (WoS) Core Collection, which is widely regarded as one of the most comprehensive and rigorously curated multidisciplinary scientific databases and is extensively employed in bibliometric research (Mongeon & Paul-Hus, 2016). The search was conducted on June 7, 2026, using the following query string applied to topic fields (title, abstract, author keywords, and Keywords Plus): TS=("air traffic control*" OR "air traffic management" OR "air traffic controller*") AND TS=("artificial intelligence" OR "machine learning" OR "reinforcement learning") AND TS=("AI agent*" OR "MARL" OR "multi-agent*" OR "multiagent*" OR "intelligent agent*" OR agent). These specific keywords and search criteria were systematically derived from a preliminary literature review and expert consultations to align directly with the aim of the study. No restriction was initially imposed on WoS categories, allowing the search to capture the interdisciplinary nature of research on AI agents in air traffic control. The retrieved records were distributed across several WoS categories, predominantly including Computer Science, Artificial Intelligence; Engineering, Aerospace; Engineering, Electrical & Electronic; Transportation Science & Technology; Automation & Control Systems; Computer Science, Interdisciplinary Applications; Computer Science, Theory & Methods; and Operations Research & Management Science. The records were covered by the Science Citation Index Expanded (SCI-EXPANDED), Social Sciences Citation Index (SSCI), Emerging Sources Citation Index (ESCI), and Conference Proceedings Citation Index. The results were subsequently restricted to the document type "Article", and all publications retained for analysis were in English. The initial search of the WoS Core Collection identified 101 records. The results were first restricted by document type to articles, reducing the dataset to 65 records. The titles and abstracts of these records were then screened for relevance to the scope of the study. Following this relevance assessment, 7 records were excluded, resulting in a final dataset of 58 publications included in the bibliometric analysis.

Data analysis

The exported dataset was analyzed using VOSviewer (version 1.6.20), a widely adopted software tool for constructing and visualizing bibliometric networks (van Eck & Waltman, 2010). Prior to analysis, the dataset was cleaned to resolve author name ambiguities and standardize keyword variants. Five analytical procedures were applied in alignment with the research questions.


Annual publication output was examined through descriptive statistics to characterize temporal trends in research activity. Co-authorship network analysis was conducted at the author level to map collaborative structures and identify central actors within the field, with a minimum threshold of one publication and one citation per author. Citation analysis was performed at both the article and country levels to identify the most influential contributions and assess the geographic distribution of research impact; the minimum threshold was set to one publication and one citation per unit of analysis. Bibliographic coupling analysis was conducted at the author level to identify intellectual similarities among researchers based on shared reference lists, applying the same minimum threshold criteria. Keyword co-occurrence analysis was performed on author-assigned keywords, with a minimum occurrence threshold of two, to identify dominant thematic clusters and trace conceptual linkages across the corpus. All network maps were generated using the VOS layout algorithm, with cluster assignments determined by the modularity-based community detection method implemented within VOSviewer.

Findings

Annual distribution of articles

A total of 58 publications related to AI agents in air traffic control, published between 1998 and June 2026, were analyzed. Figure 1 presents the annual distribution of these publications. Although only sporadic publications were observed during the early years of the field, a noticeable increase in research output emerged after 2022. The publication gap between the 1998 and 2012 studies could be linked to the technological limitations of early AI models and the conservative nature of the aviation industry. The number of publications reached its highest level in 2025 ($n = 16$), indicating substantial growth in scholarly attention toward the application of AI agents in air traffic control environments. The upward trajectory observed in recent years suggests that the integration of advanced AI technologies, including intelligent agents, multi-agent systems, and reinforcement learning approaches, has become an increasingly prominent research topic within the ATM domain.

Figure 1. Distribution of articles per year




Co-authorship analysis of authors

Co-authorship analysis is one of the most widely used bibliometric techniques for examining the intellectual and social structure of a research field by identifying collaboration patterns among scholars (Melin & Persson, 1996; van Eck & Waltman, 2010). In the study, author collaboration networks were analyzed using VOSviewer. The inclusion thresholds were set at one publication and one citation per author. Based on these threshold values, six authors met the inclusion criteria and were incorporated into the co-authorship network.

As illustrated in Figure 2, the network included six authors and 11 collaboration links, with a total link strength of 22. The six authors formed two clusters. The relatively small size of the network indicates that research on AI agents in air traffic control remains an emerging and highly specialized domain, with a limited number of recurring collaborations among scholars. Similar patterns are commonly observed in developing research areas where the knowledge base is still forming and the number of active contributors remains relatively low (Donthu et al., 2021). The largest collaboration cluster is composed of Pham Duc-Thinh, Sameer Alam, Yanjun Wang, and Hasnain Ali, who demonstrate the highest degree of interconnectedness within the network. Among these authors, Pham Duc-Thinh and Sameer Alam occupy central positions, maintaining collaborative ties with authors from both clusters and thereby functioning as important connectors within the network structure. According to social network theory, such central actors facilitate knowledge diffusion and contribute to the integration of otherwise fragmented research communities (Borgatti & Halgin, 2011). The second cluster consists of Vu Duong and Phu N. Tran, whose collaboration appears relatively isolated from the broader network. Although these authors maintain strong intra-cluster collaboration, their connections with the remaining authors are comparatively limited. This finding suggests the existence of geographically or institutionally concentrated research groups that operate with limited interaction across the broader scientific community.

The co-authorship structure reveals a fragmented but gradually developing collaboration landscape. The presence of only two small clusters and a limited number of authors indicates that scholarly collaboration in the field of AI agents for air traffic control is still at an early stage of development.

Figure 2. Co-authorship network of authors



Citation-based findings

Citation analysis is a widely used bibliometric technique for identifying the most influential studies within a research field and assessing their intellectual impact on subsequent research (Bornmann & Daniel, 2008). In the study, citation performance was evaluated for 58 publications, and the ten most cited articles are presented in Table 1.

The most cited publication was Agogino and Tumer (2012) with 68 citations, introducing a multi-agent approach for air traffic flow management. The study is considered one of the foundational contributions demonstrating how autonomous agents can collaboratively manage traffic flow and reduce congestion. The second most cited article, Brittain and Wei (2022) with 54 citations, focused on scalable autonomous separation assurance using heterogeneous MARL, highlighting the growing role of reinforcement learning in air traffic control applications. Similarly, Cruciol et al. (2013) and Tran et al. (2020) contributed to the literature by addressing learning-based traffic flow control and conflict resolution problems, receiving 37 and 34 citations, respectively.

Several highly cited studies, including Wang et al. (2019), Pham et al. (2022), and Spatharis et al. (2023), focused on reinforcement learning-based conflict detection, conflict resolution, and ATM optimization. In contrast, the study by Wangermann and Stengel (1998) represents one of the earliest publications in the dataset and examines negotiation mechanisms between intelligent agents for ATM. More recent publications, such as Chen et al. (2024) and Kravaris et al. (2023), have already attracted considerable citation attention despite their recent publication dates, reflecting the increasing scholarly interest in XAI, urban air mobility, and MARL applications.


The most influential publications in the field predominantly focus on multi-agent systems, reinforcement learning, conflict management, and autonomous decision-making, indicating that these themes constitute the intellectual foundation of research on AI agents in air traffic control.

Table 1. Most cited articles on AI agents for air traffic control

Ranking	Article	Citations
1	Agogino & Tumer (2012)	68
2	Brittain & Wei (2022)	54
3	Cruciol et al. (2013)	37
4	Tran et al. (2020)	34
5	Wang et al. (2019)	33
6	Pham et al. (2022)	32
7	Chen et al. (2024)	30
8	Spatharis et al. (2023)	30
9	Wangermann & Stengel (1998)	24
10	Kravaris et al. (2023)	23

A citation network analysis was conducted to identify the intellectual structure of the literature on AI agents in air traffic control. The minimum threshold was set to one publication and one citation. Based on these criteria, 31 interconnected publications were included in the network visualization. The resulting citation network consisted of 8 clusters, 49 links, and several highly connected publications that form the intellectual core of the field as shown in Figure 3. Among the analyzed publications, Agogino and Tumer (2012) emerged as the most influential study, occupying the most central position within the citation network. Other prominent publications include Brittain and Wei (2022), Cruciol et al. (2013), Tran et al. (2020), Pham et al. (2022), Wang et al. (2019), and Spatharis et al. (2023). These studies exhibit strong citation connections and constitute the primary knowledge base upon which subsequent research has been developed.

Figure 3. Citation network of articles




Country-level citation analysis

The citation network of publications was analyzed at the country level using VOSviewer. The minimum threshold was set at one publication and one citation per country. Based on these criteria, 17 countries met the inclusion requirements and were incorporated into the citation network. The resulting network consisted of 7 clusters, 49 links, and a total link strength of 114, indicating a moderate level of citation connectivity among contributing countries as shown in Figure 4.

Among the countries represented in the network, the United States emerged as the most influential contributor with 281 citations, followed by the People's Republic of China (157 citations), France (90 citations), England (89 citations), and Spain (79 citations). These countries occupy the most central positions within the citation map and exhibit strong citation connections with multiple countries, reflecting their significant influence on the development of research concerning AI agents in air traffic control.


Figure 4. Country-level citation network



Bibliographic coupling of authors

Bibliographic coupling analysis is a widely used bibliometric technique that identifies intellectual similarities among authors based on the extent to which they cite the same references (Zupic & Čater, 2015). Authors sharing a substantial proportion of references are considered to have similar research interests and knowledge foundations. Within the scope of this analysis, the minimum threshold was set to one publication and one citation. As a result, 40 authors met the inclusion criteria and were incorporated into the bibliographic coupling network. The resulting network consisted of 6 clusters, 593 links, and a total link strength of 10,372, indicating a highly interconnected intellectual structure within the field as shown in Figure 5. The relatively high number of links and total link strength suggest that researchers frequently rely on a common body of literature, reflecting the emergence of a shared knowledge base in research on AI agents in air traffic control. Among the authors, Sameer Alam, Minghua Hu, George A. Vouros, Hasnain Ali, and Hong-Cheol Choi occupy prominent positions within the network. Their central locations and strong bibliographic connections indicate that their publications share substantial similarities in referenced literature and contribute significantly to the intellectual development of the field. The cluster structure further reveals several interconnected research streams focusing on topics such as multi-agent systems, reinforcement learning, conflict resolution, autonomous decision-making, and intelligent ATM.

Figure 5. Bibliographic coupling of authors



Keyword co-occurrence analysis of author keywords

Keyword co-occurrence analysis is a widely used text mining technique that identifies conceptual relationships among terms by examining their simultaneous occurrence within scientific publications. The method enables the visualization of thematic structures, research hotspots, and intellectual trends within a research domain (Cobo et al., 2011). In VOSviewer maps, the proximity between keywords indicates the strength of their conceptual association, while node size reflects the frequency of occurrence.

In the study, a keyword co-occurrence analysis was performed on publications related to AI agents in air traffic control. The resulting network consisted of 44 keywords, 6 clusters, 214 links, and a total link strength of 301, revealing several interconnected thematic areas within the literature as shown in Figure 6. Among all keywords, “air traffic control” emerged as the most frequently occurring term (19 occurrences), followed by “reinforcement learning” (16 occurrences) and “air traffic management” (14 occurrences). The prominence of these terms indicates that the literature is primarily focused on the application of intelligent learning algorithms to operational air traffic control and management challenges.

The keyword cluster map further illustrates the conceptual structure of the field. The central green cluster is organized around “air traffic control”, “deep reinforcement learning”, “conflict resolution”, and “urban air mobility”, highlighting research related to autonomous decision-making and conflict management in increasingly complex airspace environments. The blue cluster is dominated by “reinforcement learning”, “decision support systems”, and “conflict detection and resolution”, reflecting studies that employ learning-based approaches to enhance operational decision support. The yellow cluster centers on “air traffic management” and “multi-agent reinforcement learning”, emphasizing optimization, coordination, and intelligent traffic flow management. The red cluster includes terms such as “decision making”, “real-time

systems”, “reinforcement learning”, and “graph neural networks”, indicating growing interest in advanced AI techniques for dynamic and real-time air traffic applications. Finally, the turquoise cluster, represented by “machine learning” and “ETA prediction”, reflects predictive analytics and forecasting applications within air traffic operations. The density visualization further supports these findings by identifying “air traffic control”, “reinforcement learning”, and “air traffic management” as the most conceptually dense and interconnected areas within the literature, confirming their central role in shaping research on AI agents in air traffic control as shown in Figure 6 and Figure 7.

Figure 6. Author keyword co-occurrence network





Figure 7. Author keyword co-occurrence density visualization



Discussion and Conclusion

The study set out to chart the intellectual trajectory of AI agent research within the domain of air traffic control by conducting a systematic bibliometric analysis of the available scholarly corpus. The research questions examined the development of the field, its main thematic and methodological foundations, and its leading contributors. The findings provide a structured quantitative overview of these dimensions.

The temporal distribution of publications reflects a broader transition in the field from exploratory inquiry to applied research. The foundational work of Wangermann and Stengel (1998) established the conceptual

premise of autonomous negotiation among intelligent agents in ATM, yet sustained scholarly engagement did not materialize for over two decades. The acceleration observed after 2022 aligns with concurrent advances in deep reinforcement learning architectures and the growing recognition that rule-based automation is insufficient for managing the complexity of increasingly dense airspace environments. The dominance of reinforcement learning as a methodological paradigm is evident in both citation patterns and keyword co-occurrence structures. This trend is consistent with findings reported in adjacent bibliometric surveys of AI in aviation (Degas et al., 2022; Karakavuz, 2025). Nevertheless, the present analysis shows that narrowing the unit of inquiry specifically to AI agents yields distinct results. Under this focused scope, the thematic concentration around conflict resolution, separation assurance, and multi-agent coordination becomes considerably more pronounced. The concentration of citation impact in the United States and China reflects not only the scale of national research investment in autonomous systems but also the alignment of these countries' aviation authorities with modernization programs that create institutional demand for agent-based solutions. The increasing focus on explainability and human-AI collaboration in the relevant literature, as evidenced in Kravaris et al. (2023) and Kirwan (2025), indicates that the field is beginning to address the operational and regulatory aspects of deployment. The co-authorship and bibliographic coupling networks further reinforce this picture of an early-stage research community. The presence of only two small co-authorship clusters and a limited number of recurring collaborators indicates that cross-institutional and cross-national collaboration remains underdeveloped. Nevertheless, the considerable total link strength observed in the bibliographic coupling network suggests that, despite fragmented collaboration structures, researchers increasingly draw from a shared intellectual foundation. This pattern is consistent with that observed in fields in which a common methodological core precedes the formation of broader collaborative ties (Donthu et al., 2021).

The primary contribution of the study is its specific analytical focus. Previous bibliometric studies on AI in aviation generally examined machine learning, optimization, and agent-based approaches within broad AI categories. In contrast, the present study examines AI agents as a distinct research topic. AI agents are characterized by autonomy, perception–action reasoning, and goal-directed behavior rather than by a specific technical method alone. This targeted approach reveals thematic clusters and citation structures often hidden in broader analyses. Key areas uncovered include the importance of human-AI teaming, the role of XAI in operational acceptance, and the use of probabilistic digital twins for validating autonomous agents. Consequently, these findings provide researchers with a clear map of an underexplored subfield. Furthermore, the identified themes highlight critical transition points from simulation to real-world deployment for practitioners and regulators working on system certification.

The study has several limitations that warrant consideration. The analysis is restricted to publications indexed in the WoS Core Collection. Databases such as Scopus and PubMed, national repositories including TÜBİTAK ULAKBİM and the Council of Higher Education Thesis Archive, and non-indexed grey literature fall outside the scope of the present corpus. This reflects a deliberate choice to prioritize source quality and longitudinal comparability, but it may underrepresent certain research communities. Additionally, despite systematic data cleaning, the possibility that a small number of relevant records were inadvertently excluded during duplicate removal or keyword standardization cannot be fully ruled out. The relatively small final corpus of 58 publications, a function of the query's specificity rather than a methodological weakness, does impose constraints on the robustness of network-level analyses.

Future studies could replicate this analysis using Scopus or a combined database approach to broaden coverage and enable cross-database comparison. Systematic reviews or meta-analyses built around the thematic clusters identified here would allow for deeper synthesis of methodological choices and empirical findings within specific sub-areas, such as MARL for separation assurance. Periodic updates to this bibliometric map would also be valuable, given the pace at which the field is developing.

Author Contributions

Author 1: Conceptualization, Writing – Original Draft, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – Review & Editing, Visualization

Author 2: Conceptualization, Writing – Original Draft, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – Review & Editing, Visualization.

Declaration 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.

Ethical Approval

Not applicable.

References

- Abou Ali, M., Dornaika, F., & Charafeddine, J. (2025). Agentic AI: a comprehensive survey of architectures, applications, and future directions. *Artificial Intelligence Review*, 59(11), 1-37. <https://doi.org/10.1007/s10462-025-11422-4>
- Adabara, I., Olaniyi Sadiq, B., Nuhu Shuaibu, A., Ibrahim Danjuma, Y., & Maninti, V. (2025). Trustworthy agentic AI systems: a cross-layer review of architectures, threat models, and governance strategies for real-world deployment. *FI000Research*, 14, 905. <https://doi.org/10.12688/fi000research.169927.1>
- Agogino, A. K., & Tumer, K. (2012). A multiagent approach to managing air traffic flow. *Autonomous Agents and Multi-Agent Systems*, 24(1), 1-25. <https://doi.org/10.1007/s10458-010-9142-5>
- Ali, H., Pham, D. T., Alam, S., Schultz, M., Li, M. Z., Wang, Y., Itoh, E., & Duong, V. N. (2025). Human-AI hybrids in safety-critical systems: concept, definition and perspectives from air traffic management. *Advanced Engineering Informatics*, 65, 103256. <https://doi.org/10.1016/j.aei.2025.103256>
- Aziz, A., & Wei, P. (2026). Transformer-based Multi-agent Reinforcement Learning for Separation Assurance in Structured and Unstructured Airspaces. *arXiv preprint arXiv:2601.04401*. <https://doi.org/10.48550/arXiv.2601.04401>
- Berlin, I. (2017). Two concepts of liberty. In *Liberty reader* (pp. 33-57). Routledge.
- Biswas, B., & Sarkar, S. (2026). Responsible agentic artificial intelligence governance: Risk, safety, and ethical challenges in autonomous systems. *International Journal of Applied Resilience and Sustainability*, 2(2), 142-167. <https://doi.org/10.70593/deepsci.0202005>
- Borgatti, S. P., & Halgin, D. S. (2011). On network theory. *Organization science*, 22(5), 1168-1181. <https://doi.org/10.1287/orsc.1100.0641>
- Bornmann, L., & Daniel, H. D. (2008). What do citation counts measure? A review of studies on citing behavior. *Journal of documentation*, 64(1), 45-80. <https://doi.org/10.1108/00220410810844150>
- Brittain, M., & Wei, P. (2019). Autonomous air traffic controller: A deep multi-agent reinforcement learning approach. *arXiv preprint arXiv:1905.01303*. <https://doi.org/10.48550/arXiv.1905.01303>
- Brittain, M., & Wei, P. (2022). Scalable autonomous separation assurance with heterogeneous multi-agent reinforcement learning. *IEEE Transactions on automation science and engineering*, 19(4), 2837-2848. <https://doi.org/10.1109/TASE.2022.3151607>
- Carvell, B. J., De Ath, G., Benjamin, E., & Everson, R. (2026). Online Action-Stacking Improves Reinforcement Learning Performance for Air Traffic Control. In *AIAA SCITECH 2026 Forum* (p. 2746). <https://doi.org/10.2514/6.2026-2746>
- Carvell, B. J., Thomas, M., Pace, A., Dorney, C., De Ath, G., Everson, R., Pepper, N., Keane, A., & Cannon, R. (2026). Human-In-The-Loop Testing of AI Agents for Air Traffic Control with a Regulated Assessment Framework. In *AIAA SCITECH 2026 Forum* (p. 2558). <https://doi.org/10.2514/6.2026-2558>
- Çeken, S. (2025). Hava Trafik Yönetiminde Yapay Zekâ Uygulama Süreci. In Önen, V. (ed.) 2025. Havacılık Yönetiminde Yapay Zeka Uygulamaları, 113-146. <https://doi.org/10.58830/ozgur.pub1142.c4710>
- Chan, A., Ezell, C., Kaufmann, M., Wei, K., Hammond, L., Bradley, H., ... & Anderljung, M. (2024, June). Visibility into AI agents. In *Proceedings of the 2024 ACM conference on fairness, accountability, and transparency* (pp. 958-973). <https://doi.org/10.1145/3630106.3658948>
- Chen, S., Evans, A. D., Brittain, M., & Wei, P. (2024). Integrated conflict management for UAM with strategic demand capacity balancing and learning-based tactical deconfliction. *IEEE Transactions on Intelligent Transportation Systems*, 25(8), 10049-10061. <https://doi.org/10.1109/TITS.2024.3351049>
- Chen, Y., Wang, W., & Chen, X. M. (2023). Bibliometric methods in traffic flow prediction based on artificial intelligence. *Expert Systems with Applications*, 228, 120421. <https://doi.org/10.1016/j.eswa.2023.120421>

- Cobo, M. J., López-Herrera, A. G., Herrera-Viedma, E., & Herrera, F. (2011). Science mapping software tools: Review, analysis, and cooperative study among tools. *Journal of the American Society for information Science and Technology*, 62(7), 1382-1402. <https://doi.org/10.1002/asi.21525>
- Cruciol, L. L., de Arruda Jr, A. C., Weigang, L., Li, L., & Crespo, A. M. (2013). Reward functions for learning to control in air traffic flow management. *Transportation Research Part C: Emerging Technologies*, 35, 141-155. <https://doi.org/10.1016/j.trc.2013.06.010>
- Degas, A., Islam, M. R., Hurter, C., Barua, S., Rahman, H., Poudel, M., Ruscio, D., Ahmed, M. U., Begum, S., Rahman, M. A., Bonelli, S., Cartocci, G., Di Flumeri, G., Borghini, G., Babiloni, F., & Arico, P. (2022). A survey on artificial intelligence (AI) and explainable AI in air traffic management: Current trends and development with future research trajectory. *Applied Sciences*, 12(3), 1295. <https://doi.org/10.3390/app12031295>
- Demir, G. (2026). Bibliometric analysis in artificial intelligence and transportation: Evaluation of research and future perspectives. In *Reliable Decision-Making for Sustainable Transportation* (pp. 255-273). *Academic Press*. <https://doi.org/10.1016/B978-0-443-33740-6.00005-0>
- Demir, G., Moslem, S., & Duleba, S. (2024). Artificial intelligence in aviation safety: Systematic review and biometric analysis. *International Journal of Computational Intelligence Systems*, 17(1), 279. <https://doi.org/10.1007/s44196-024-00671-w>
- Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. *Journal of Business Research*, 133, 285–296. <https://doi.org/10.1016/j.jbusres.2021.04.070>
- Ertekin, R., Rodoplu, H., & Gürsel, S. (2024). The use of artificial intelligence in aviation: A bibliometric analysis. *International Journal of Computational and Experimental Science and Engineering*, 10(4). <https://doi.org/10.22399/ijcesen.747>
- Fortune Business Insights. (2026). Artificial intelligence (AI) in aviation market size, share & industry analysis. Retrieved from <https://www.fortunebusinessinsights.com/ai-in-aviation-market-113289>
- International Air Transport Association (IATA). (2026). Strong 2025 Passenger Demand Masks Ongoing Capacity Constraints. <https://www.iata.org/en/pressroom/2026-releases/2026-01-29-02/>
- Jiang, P., Lin, J., Shi, Z., Wang, Z., He, L., Wu, Y., ... & Han, J. (2025). Adaptation of agentic ai. A Survey of Post-Training, Memory, and Skills. *arXiv preprint arXiv:2512.16301*. <https://doi.org/10.48550/arXiv.2512.16301>
- Ju, H., & Aral, S. (2025). Collaborating with ai agents: Field experiments on teamwork, productivity, and performance. *arXiv preprint arXiv:2503.18238*. <https://doi.org/10.48550/arXiv.2503.18238>
- Karakavuz, H. (2025). Bibliometric Analysis of Studies on Artificial Intelligence in the Air Transportation Sector. *Journal of Aviation*, 9(1), 118-136. <https://doi.org/10.30518/jav.1583144>
- Kasirzadeh, A., & Gabriel, I. (2025). Characterizing AI agents for alignment and governance. *arXiv preprint arXiv:2504.21848*. <https://doi.org/10.48550/arXiv.2504.21848>
- Kent, P., De Ath, G., Carvell, B. J., Everson, R., Layton, M., & Hart, A. (2026). A future capabilities agent for tactical Air Traffic Control. In *AIAA SCITECH 2026 Forum* (p. 1203). <https://doi.org/10.2514/6.2026-1203>
- Kirwan, B. (2025). Human factors requirements for human-AI teaming in aviation. *Future Transportation*, 5(2), 42. <https://doi.org/10.3390/futuretransp5020042>
- Kravaris, T., Lentzos, K., Santipantakis, G., Vouros, G. A., Andrienko, G., Andrienko, N., ... & Martinez, E. I. (2023). Explaining deep reinforcement learning decisions in complex multiagent settings: towards enabling automation in air traffic flow management. *Applied Intelligence*, 53(4), 4063-4098. <https://doi.org/10.1007/s10489-022-03605-1>
- Krishnan, N. (2025). Ai agents: Evolution, architecture, and real-world applications. *arXiv preprint arXiv:2503.12687*. <https://doi.org/10.48550/arXiv.2503.12687>
- Lopes, N. M., Aparicio, M., & Neves, F. T. (2025). Challenges and prospects of artificial intelligence in aviation: a bibliometric study. *Data Science and Management*, 8(2), 207-223. <https://doi.org/10.1016/j.dsm.2024.11.001>

- Mei, P., Cannon, R., Everett, J., Liu, P., & Awad, E. (2025). Public trust and blame attribution in human-AI interactions: a comparison between air traffic control and vehicle driving. *Transportation Research Interdisciplinary Perspectives*, 32, 101545. <https://doi.org/10.1016/j.trip.2025.101545>
- Melin, G., & Persson, O. (1996). Studying research collaboration using co-authorships. *Scientometrics*, 36(3), 363-377. <https://doi.org/10.1007/BF02129600>
- Mohamed, E., Carvell, B. J., Procter, R., Benjamin, E., De Ath, G., & Everson, R. (2026). Towards Transparent AI Agents for Air Traffic Control. In *AIAA SCITECH 2026 Forum* (p. 2869). <https://doi.org/10.2139/ssrn.6042354>
- Mongeon, P., & Paul-Hus, A. (2016). The journal coverage of Web of Science and Scopus: A comparative analysis. *Scientometrics*, 106(1), 213–228. <https://doi.org/10.1007/s11192-015-1765-5>
- Nisa, U., Shirazi, M., Saip, M. A., & Pozi, M. S. M. (2025). Agentic AI: The age of reasoning—A review. *Journal of Automation and Intelligence*, 5, 69-89. <https://doi.org/10.1016/j.jai.2025.08.003>
- Pepper, N., Keane, A., Hodgkin, A., Gould, D., Henderson, E., Lauritsen, L., Vlahos, C., De Ath, G., Everson, R., Cannon, R., Sierra Castro, A., Korna, J., Carvell, B. J., & Thomas, M. (2026). A Probabilistic Digital Twin of UK en Route Airspace for Training and Evaluating AI Agents for Air Traffic Control. In *AIAA SCITECH 2026 Forum* (p. 1794). <https://doi.org/10.2514/6.2026-1794>
- Pham, D. T., Tran, P. N., Alam, S., Duong, V., & Delahaye, D. (2022). Deep reinforcement learning based path stretch vector resolution in dense traffic with uncertainties. *Transportation research part C: emerging technologies*, 135, 103463. <https://doi.org/10.1016/j.trc.2021.103463>
- Qu, X., Damoah, A., Sherwood, J., Liu, P., Jin, C. S., Chen, L., ... & De Souza, C. (2025). A comprehensive review of AI agents: Transforming possibilities in technology and beyond. *arXiv preprint arXiv:2508.11957*. <https://doi.org/10.48550/arXiv.2508.11957>
- Raza, S., Sapkota, R., Karkee, M., & Emmanouilidis, C. (2025). Trism for agentic ai: A review of trust, risk, and security management in llm-based agentic multi-agent systems (2025). *arXiv preprint arXiv:2506.04133*. <https://doi.org/10.48550/arXiv.2506.04133>
- Sadou, A. M., & Njaya, E. T. (2023). Applications of artificial intelligence in the air transport industry: a bibliometric and systematic literature review. *Journal of Aerospace Technology and Management*, 15, e2223. <https://doi.org/10.1590/jatm.v15.1312>
- Şahin, Y. (2024). Havacılık ve yapay zeka alanındaki çalışmaların bibliyometrik analizi. *International Journal of Social and Humanities Sciences Research (JSHSR)*, 11(110), 1637-1648. <https://doi.org/10.5281/zenodo.13732160>
- Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). Ai agents vs. agentic ai: A conceptual taxonomy, applications and challenges. *Information Fusion*, 126, 103599. <https://doi.org/10.1016/j.inffus.2025.103599>
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., ... & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587), 484-489. <https://doi.org/10.1038/nature16961>
- Spatharis, C., Bastas, A., Kravaris, T., Blekas, K., Vouros, G. A., & Cordero, J. M. (2023). Hierarchical multiagent reinforcement learning schemes for air traffic management. *Neural Computing and Applications*, 35(1), 147-159. <https://doi.org/10.1007/s00521-021-05748-7>
- Tran, P. N., Pham, D. T., Goh, S. K., Alam, S., & Duong, V. (2020). An interactive conflict solver for learning air traffic conflict resolutions. *Journal of Aerospace Information Systems*, 17(6), 271-277. <https://doi.org/10.2514/1.1010807>
- Tuncal, A. (2025). Artificial Intelligence Integration in Air Traffic Management: A Qualitative Content Analysis of the SESAR Research. *AIPA's International Journal on Artificial Intelligence: Bridging Technology, Society and Policy*, 1(2), 61-72. <https://doi.org/10.5281/zenodo.17391128>
- van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics*, 84(2), 523–538. <https://doi.org/10.1007/s11192-009-0146-3>
- Wang, Z., Li, H., Wang, J., & Shen, F. (2019). Deep reinforcement learning based conflict detection and resolution in air traffic control. *IET Intelligent Transport Systems*, 13(6), 1041-1047. <https://doi.org/10.1049/iet-its.2018.5357>

-
- Wangermann, J. P., & Stengel, R. F. (1998). Principled negotiation between intelligent agents: a model for air traffic management. *Artificial Intelligence in Engineering*, 12(3), 177-187. [https://doi.org/10.1016/S0954-1810\(98\)80001-0](https://doi.org/10.1016/S0954-1810(98)80001-0)
- Zhang, R., Qin, P., Cao, Q., Zhang, L., & Xie, P. (2026). AIBuildAI: An AI Agent for Automatically Building AI Models. *arXiv preprint arXiv:2604.14455*. <https://doi.org/10.48550/arXiv.2604.14455>
- Zupic, I., & Čater, T. (2015). Bibliometric methods in management and organization. *Organizational research methods*, 18(3), 429-472. <https://doi.org/10.1177/1094428114562629>