

Integrating Neurosymbolic AI in Advanced Air Mobility: A Comprehensive Survey

Kamal Acharya¹, Iman Sharifi², Mehul Lad¹, Liang Sun³ and Houbing Song¹

¹Department of Information Systems, University of Maryland, Baltimore County

²Department of Mechanical and Aerospace Engineering, The George Washington University

³Department of Mechanical Engineering, Baylor University

{kamala2, du72811, songh}@umbc.edu, i.sharifi@gwu.edu, liang_sun@baylor.edu

Abstract

Neurosymbolic AI combines neural network adaptability with symbolic reasoning, promising an approach to address the complex regulatory, operational, and safety challenges in Advanced Air Mobility (AAM). This survey reviews its applications across key AAM domains such as demand forecasting, aircraft design, and real-time air traffic management. Our analysis reveals a fragmented research landscape where methodologies, including Neurosymbolic Reinforcement Learning, have shown potential for dynamic optimization but still face hurdles in scalability, robustness, and compliance with aviation standards. We classify current advancements, present relevant case studies, and outline future research directions aimed at integrating these approaches into reliable, transparent AAM systems. By linking advanced AI techniques with AAM's operational demands, this work provides a concise roadmap for researchers and practitioners developing next-generation air mobility solutions.

1 Introduction

Artificial Intelligence (AI) has undergone significant development, marked by a "Cold War" between symbolic and neural approaches. The symbolic approach dominated the early stages of AI research, but in recent years, neural models have gained prominence. Both approaches face distinct challenges, and this has paved the way for the emergence of the third wave of AI, known as Neurosymbolic AI [Garcez and Lamb, 2023]. Neurosymbolic AI combines the learning capabilities of neural models with the reasoning abilities of symbolic models. Simultaneously, the aviation and transportation industries are experiencing a transformative shift with the advent of Advanced Air Mobility (AAM). AAM envisions a future where air transportation is seamlessly integrated into everyday life, enabling efficient, safe, and sustainable movement of people and goods. This vision is being realized through technological advancements in electric propulsion, automation, and materials science, fostering the development of innovative aircraft such as electric vertical take-off and landing (eVTOL) vehicles [Johnson and Silva, 2022].

AAM faces several obstacles like regulatory complexity, the need for safe and explainable automated decision-making, and the integration of heterogeneous data from demand forecasting, infrastructure constraints, weather, and traffic conditions [Cohen and Shaheen, 2024]. Neurosymbolic AI offers a promising solution by combining the pattern-recognition power of neural networks with the rule-based logic of symbolic systems, allowing for a more interpretable, high-level reasoning framework that can encode safety protocols, regulatory constraints, and domain-specific knowledge. As a result, decisions made by Neurosymbolic AI are not only data-informed but also grounded in explicit safety requirements and operational guidelines critical for fostering growth in AAM by improving trust, compliance, and the overall reliability of autonomous aerial operations.

Neurosymbolic AI offers significant potential for AAM by combining neural learning from large, heterogeneous datasets (e.g., sensor streams, flight trajectories, air traffic patterns) with symbolic reasoning that enforces regulatory and safety constraints. This dual approach supports both Urban Air Mobility (UAM) and Regional Air Mobility (RAM) by ensuring models not only adapt to complex, real-time data but also adhere to airspace rules, vehicle performance limits, and scheduling constraints. In particular, Neurosymbolic Reinforcement Learning [Acharya *et al.*, 2023] can optimize flight paths, manage dynamic air traffic, and coordinate ground infrastructure in real time, balancing efficiency with safety and noise reduction. Embedding domain knowledge into the symbolic layer alongside adaptive neural policies allows continuous strategy refinement in AAM systems. However, developing robust Neurosymbolic architectures for challenges like multi-vehicle coordination, real-time constraint satisfaction, and large-scale deployment remains an open research area. This article categorizes key research domains essential for advancing AAM from engineering design and electrification to demand modeling, autonomy, cybersecurity, and public acceptance and outlines how Neurosymbolic AI can be applied across these areas, along with potential case studies and associated risks.

While surveying existing literature, we observed no comprehensive review on the integration of Neurosymbolic AI within AAM. Prior studies have focused on isolated topics, demand modeling [Long *et al.*, 2023], aircraft design [Kiesewetter *et al.*, 2023], electrification [Dorn-Gomba *et al.*,

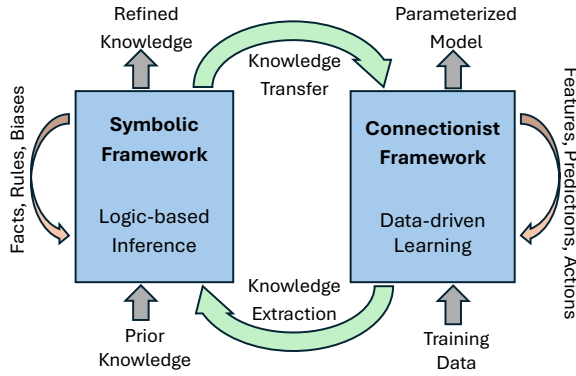


Figure 1: General interactions between symbolic and connectionist frameworks in Neurosymbolic AI

2020], and safety [Yoo *et al.*, 2022], but none have explored how Neurosymbolic AI can accelerate and enhance research across all AAM domains. This survey paper is the first to provide a detailed analysis of ongoing research efforts in AAM and to propose ways Neurosymbolic AI can drive advancements in this field.

2 Background

2.1 Neurosymbolic AI

Neurosymbolic AI represents a hybrid approach that integrates neural networks’ learning capabilities with symbolic AI’s structured reasoning, overcoming the limitations of both paradigms. Symbolic AI, dominant from the 1950s to the 1980s, offers explainable reasoning and structured knowledge representation, but it requires extensive manual rule definitions [Mao *et al.*, 2019]. In contrast, connectionist AI, or deep learning, enables data-driven pattern recognition but suffers from black-box decision-making and interpretability challenges [Acharya *et al.*, 2023]. Neurosymbolic AI leverages the strengths of both by enabling efficient knowledge extraction and conceptual reasoning, requiring less training data while maintaining transparency and adaptability in uncertain environments [Garcez and Lamb, 2023]. Due to these advantages, Neurosymbolic AI is often regarded as the third wave of AI, capable of achieving error correction, enhanced explainability, and logical inference beyond deep learning alone. Figure 1 depicts the general interaction between the symbolic and connectionist frameworks in Neurosymbolic AI.

Several taxonomies have been proposed to classify Neurosymbolic AI systems. Henry Kautz proposed a classification system based on six distinctive categories [Kautz, 2022]. More recently, a survey categorized Neurosymbolic AI based on efficiency, generalization, and interpretability, leading to three primary approaches: learning for reasoning, reasoning for learning, and learning-reasoning [Yu *et al.*, 2023]. Learning for reasoning uses neural networks to enhance symbolic reasoning, reasoning for learning incorporates symbolic logic to refine neural learning processes, and learning-reasoning represents a fully integrated system where both paradigms influence each other dynamically. These frameworks provide a

comprehensive approach to structuring Neurosymbolic AI for robust, explainable decision-making in various domains.

Neurosymbolic AI has broad applications, particularly in intelligent transportation systems (ITS) and AAM. In ITS, Neurosymbolic AI enables real-time sensor integration with pre-defined traffic rules, ensuring adaptive yet regulatory-compliant AI-driven transportation solutions [Garcez and Lamb, 2023]. This improves traffic condition prediction, autonomous vehicle safety, and public transit optimization. Similarly, in AAM and UAM, Neurosymbolic AI enhances decision-making, safety, and airspace management. It aids Unmanned Aircraft Systems (UAS) in navigating complex environments by incorporating expert-defined flight safety rules and predictive modeling [Gilpin and Ilievski, 2021]. By fusing symbolic reasoning with deep learning, Neurosymbolic AI improves reliability, demand forecasting, and autonomous decision-making in emerging aviation technologies, ensuring intelligent, safe, and interpretable air mobility operations [Kohaut *et al.*, 2024].

2.2 Advanced Air Mobility

AAM represents a paradigm shift in transportation, promising to revolutionize the movement of people and goods, particularly in urban and regional environments. This transformative vision encompasses a diverse array of applications, extending beyond the initial concept of air taxis to include UAM, cargo delivery, emergency medical services, and various other aerial missions [Goyal and Cohen, 2022]. Figure 2 provides the overview of technologies in advancement of AAM along with the applications areas.

Technological Advancements Driving AAM

Electric Vertical Take-Off and Landing Aircraft eVTOL aircraft are at the heart of the AAM revolution. Their electric propulsion systems offer several compelling advantages over traditional helicopters and airplanes. These include reduced noise pollution, significantly lower emissions contributing to a more sustainable transportation system, and the potential for autonomous operation [Long *et al.*, 2023]. The reduced noise footprint is particularly important for urban environments, where noise pollution is a major concern [Rizzi and Scata Jr, 2022]. However, substantial challenges remain in developing batteries with sufficient energy density and range to support commercially viable flight times and distances. Ensuring the safety and reliability of battery technology, including addressing potential failure modes and mitigating risks, is also crucial. Extensive research is underway to address these technological hurdles, focusing on improvements in battery technology, electric motor design, and overall aircraft design optimization [Barrera *et al.*, 2022].

Autonomous Flight and Traffic Management The safe and reliable integration of numerous autonomous aircraft into existing airspace demands the development of sophisticated Unmanned Traffic Management (UTM) systems [Namuduri, 2023]. These systems are crucial for managing air traffic flow, preventing mid-air collisions, and ensuring compliance with safety regulations. Research is actively exploring various approaches to UTM, including centralized and decentralized architectures. Centralized systems offer a more controlled ap-

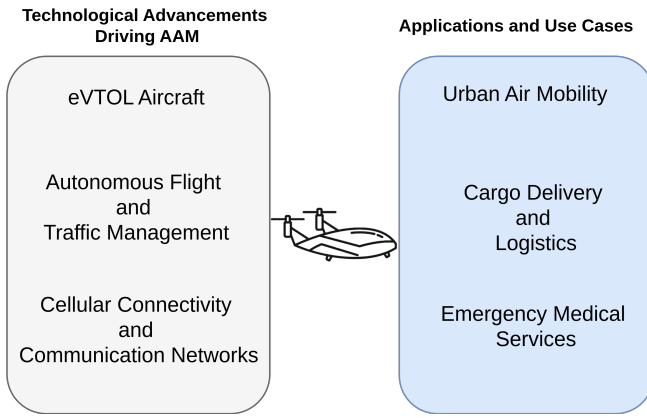


Figure 2: AAM Overview

proach to managing air traffic, while decentralized systems offer greater resilience and scalability. The choice between these architectures involves trade-offs between control and resilience, and the optimal approach may vary depending on the specific operational context [Murça *et al.*, 2024]. The integration of machine learning techniques is being explored to improve the efficiency and adaptability of UTM systems [Deniz *et al.*, 2024]. Furthermore, the development of advanced sensor technologies, such as lidar and radar, is crucial for enhancing the situational awareness of autonomous aircraft [Karampinis *et al.*, 2024].

Cellular Connectivity and Communication Networks

Seamless data exchange between aircraft, ground control, and other stakeholders is essential for real-time monitoring, traffic management, and overall system safety [Do *et al.*, 2024]. The integration of cellular networks and other communication technologies plays a pivotal role in achieving this. The evolution of cellular technology, with 5G and the anticipated arrival of 6G networks, is expected to significantly enhance the capabilities of AAM communication systems. The use of advanced beam forming techniques can improve the efficiency and reliability of cellular connectivity in AAM scenarios. Researchers are also exploring the use of other communication technologies, such as satellite communication, to augment cellular networks and provide broader coverage [Do *et al.*, 2024]. Vehicle-to-vehicle (V2V) communication is essential for enabling safe and efficient coordination between aircraft [Gür *et al.*, 2024].

Applications and Use Cases

Urban Air Mobility UAM aims to alleviate traffic congestion, reduce commute times, and improve accessibility, particularly for areas with limited ground transportation infrastructure [Goyal *et al.*, 2021]. However, challenges remain in integrating UAM into existing urban environments, managing air traffic density, and addressing public concerns about noise and safety. UAM systems require the development of sophisticated airspace management strategies to handle the high density of aircraft operations. This includes defining flight corridors, establishing separation standards, and managing potential conflicts with other air traffic [Murça *et al.*,

2024]. The development of robust communication and navigation systems is also crucial for ensuring the safety and efficiency of UAM operations [Do *et al.*, 2024]. UAM operations need to be integrated with ground transportation systems to provide seamless multimodal travel options. This involves the development of vertiports and other infrastructure to facilitate the transfer between air and ground transportation. The integration of UAM into existing urban environments requires careful consideration of land use planning, community engagement, and public acceptance.

Cargo Delivery and Logistics AAM offers significant potential for improving cargo delivery and logistics, particularly in remote or underserved areas. Drones and other AAM vehicles can transport goods more efficiently and cost-effectively than traditional ground transportation, especially in areas with challenging terrain or limited road infrastructure. Applications include medical supplies delivery, package delivery, and the transportation of time-sensitive goods [Dulia *et al.*, 2021]. The integration of AAM into existing logistics networks requires coordination with ground transportation and warehousing systems [Raghunatha *et al.*, 2023]. The use of autonomous drones for cargo delivery can reduce labor costs and improve efficiency. However, ensuring the safety and security of autonomous drone deliveries is a critical concern [Gür *et al.*, 2024].

Emergency Medical Services (EMS) eVTOL aircraft and drones can provide rapid transport of patients to hospitals, reducing response times and potentially improving patient outcomes. The use of AAM in EMS also has the potential to improve access to healthcare in underserved communities [Bridgelall and Tolliver, 2024]. AAM based EMS systems require the integration of AAM vehicles with existing EMS infrastructure and protocols. The economic feasibility of AAM based EMS depends on factors such as operating costs, patient transport times, and the potential for improved patient outcomes. The strategic placement of vertiports near hospitals and other healthcare facilities can optimize response times and improve access to care. The development of efficient routing algorithms and real-time traffic management systems is crucial for optimizing AAM-based EMS operations [Bridgelall and Tolliver, 2024].

3 Application Areas of Neurosymbolic AI in Advanced Air Mobility

Neurosymbolic AI has been gaining popularity worldwide due to its efficient learning and reasoning capability in almost every domain. In context of AAM, main applications of Neurosymbolic AI is shown in the Figure 3.

3.1 Electrification

Electrifying AAM is pivotal for achieving sustainable aviation by reducing emissions, lowering operational costs, and enabling quieter operations. However, the success of eVTOL aircraft is constrained by significant technological and operational challenges, notably the limitations of current lithium-ion battery (LIB) technology. Despite LIBs' high energy density (300–400 Wh/kg for short-range applications and over

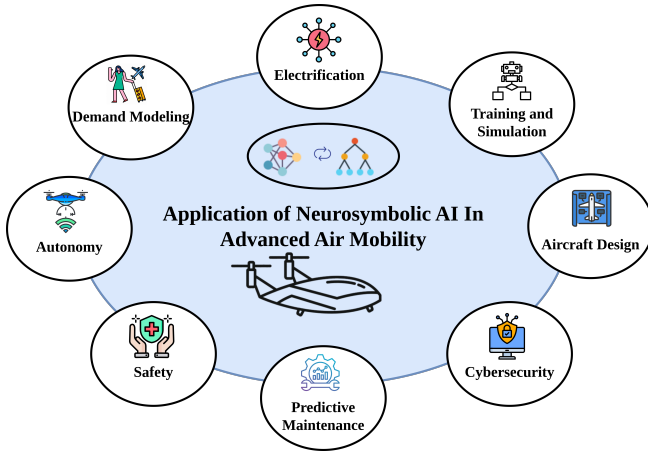


Figure 3: Applications of Neurosymbolic AI in AAM

800 Wh/kg for narrow-body aircraft), issues such as weight constraints, energy inefficiency, thermal instability, and inadequate charging infrastructure persist [Barrera *et al.*, 2022]. These challenges are exacerbated during high-demand phases like takeoff and landing, necessitating advanced energy management strategies. To address these issues, innovations such as wide bandgap semiconductors, improved thermal management systems, and advanced control strategies are being pursued to enhance the reliability and efficiency of electric propulsion systems [Dorn-Gomba *et al.*, 2020]. Furthermore, Neurosymbolic AI is emerging as a promising tool for optimizing energy usage, extending flight range, and managing battery lifecycles through techniques like predictive energy optimization, smart grid integration, and fault detection.

3.2 Aircraft Design

The design of eVTOL aircraft for air taxis and metro systems hinges on advancing battery storage, propulsion efficiency, and noise reduction to meet both safety and urban sustainability criteria. While manufacturers like Airbus are developing specialized eVTOL platforms, challenges persist in certification, infrastructure integration, and public acceptance [Silva *et al.*, 2018]. Regulatory frameworks need to evolve alongside these technologies, as highlighted by NASA/Deloitte Concept of Operations (CONOPS) [Hill *et al.*, 2020], and user-centric design elements, such as cabin noise and ride comfort, are critical for adoption. Design trade-offs between payload capacity, flight range, and vehicle weight continue to constrain the scalability of eVTOL systems [Kiesewetter *et al.*, 2023]. Moreover, distinct requirements for intra versus inter-city passenger transport, along with the differing design challenges of UAM aircraft and last-mile delivery drones, call for a multifaceted engineering strategy that incorporates aerodynamic efficiency, AI-driven control systems, and optimized energy management. Despite promising advances in control mechanisms, propulsion, and VTOL transition efficiency, bridging the gap between technological potential and real-world deployment remains a significant challenge, underscoring the need for an interdisciplinary approach integrating engineering, policy, and urban planning.

3.3 Training and Simulation

Training and simulation are essential for advancing AAM by creating adaptive environments for skill development and operational planning. Neurosymbolic AI, which synergizes the robust pattern recognition capabilities of deep learning with the interpretability and reasoning strengths of symbolic AI, can significantly enhance simulation platforms [Siyaev *et al.*, 2023]. By integrating Neurosymbolic AI, simulation systems can not only process vast amounts of sensor and operational data through deep learning but also incorporate domain-specific rules and logical reasoning to better replicate the intricate dynamics of AAM environments. This approach enables the development of training scenarios that are both realistic and adaptable, addressing complex operational variables such as weather variations, urban infrastructure constraints, and emergency protocols. Moreover, Neurosymbolic AI facilitates the creation of digital twins that can reason about potential system failures, optimize route planning, and simulate multi-agent interactions with high fidelity [Siyaev *et al.*, 2023].

3.4 Predictive Maintenance

Predictive maintenance is crucial for enhancing safety and operational efficiency in AAM. By leveraging vast arrays of sensor data and advanced analytics, predictive maintenance systems can identify early indicators of potential component degradation or system malfunctions before they escalate into critical failures [Nguyen *et al.*, 2022]. Neurosymbolic AI is used in fault diagnosis by integrating a case-based reasoning (CBR) method with a Bayesian network (BN) approach, where fault cases are initially formalized for similarity-based diagnosis and later transformed into a probabilistic model for dynamic multi-fault detection under uncertainty [Yang *et al.*, 2018]. Neurosymbolic reasoning [Siyaev *et al.*, 2023] enhances digital twins by enabling natural language interaction, allowing users to manipulate 3D components, read maintenance manuals, and perform installation and removal procedures autonomously. The result is not only a reduction in unexpected downtime, but also a significant improvement in the lifecycle management of critical systems, ultimately leading to safer and more efficient air mobility operations.

3.5 Safety

Safety is a critical aspect of AAM due to the inherent risks and complexities associated with integrating new aerial systems into urban and regional airspaces [Yoo *et al.*, 2022]. According to NASA's In-Time System-wide Safety Assurance (ISSA) framework, hazards in AAM can arise from multiple sources, including vehicle malfunctions, environmental factors, operational contexts, and broader aviation system challenges [Ellis *et al.*, 2020]. The National Transportation Safety Board (NTSB) emphasizes that AAM safety standards must exceed those of general aviation, given the high volume of short-distance flights expected in AAM operations [National Academies of Sciences *et al.*, 2018]. To ensure safety in AAM, one approach involves using symbolic logic constraints driven by common rules, often derived from Federal Aviation Administration (FAA) regulations, to filter or "shield" the outputs of neural networks, preventing unsafe

actions during the training phase [Jansen *et al.*, 2020]. To this end, a Neurosymbolic Deep Reinforcement Learning approach has been proposed to enhance safety by embedding symbolic logic constraints directly into the learning process [Sharifi *et al.*, 2023]. Another strategy involves guiding neural networks using logical safety rules, providing feedback to the AI agent regarding the safety of its actions [Kimura *et al.*, 2021].

3.6 Autonomy

AAM envisions high-density, low-altitude operations in urban and regional environments, requiring aerial vehicles to navigate dynamic obstacles, weather conditions, and air traffic without direct human intervention. Autonomous systems reduce pilot workload, lower operational costs, and enhance safety by making real-time decisions based on sensor data and predefined regulations. Autonomy ensures seamless integration with existing Air Traffic Management (ATM) systems, facilitating coordination between manned and unmanned aerial vehicles (UAVs). By embedding symbolic representations of air traffic rules, safety protocols, and operational guidelines, Neurosymbolic systems can ensure compliance with aviation standards, avoiding collisions and potential conflicts, while dynamically adapting to changing conditions. Moreover, Neurosymbolic AI tackles the overlying scalability issue in traditional connectionist frameworks by integrating symbolic knowledge (e.g., air traffic rules, safety protocols) with neural network-based perception and control systems, allowing autonomous AAM vehicles to scale across diverse environments with minimal retraining [Wang *et al.*, 2024].

3.7 Cybersecurity

AAM cybersecurity challenges stem from their reliance on interconnected systems, autonomous decision-making, and real-time data exchange. Critical threats include GPS spoofing, wireless attacks (e.g., jamming and man-in-the-middle), AI adversarial manipulations, and cyber-physical intrusions, all of which can compromise navigational accuracy and flight control [Petit and Shladover, 2014]. Unlike conventional deep learning models, which are prone to adversarial vulnerabilities, Neurosymbolic AI enhances security by integrating logical rules, knowledge graphs, and domain constraints. For instance, Logic Tensor Networks can merge formal security policies with real-time anomaly detection to filter malicious network activity [Grov *et al.*, 2024], while Neurosymbolic methods also enable cross-validation of multiple navigational inputs to detect and correct GPS spoofing. Neurosymbolic AI driven Security Operations Centers have proven effective in cyber threat intelligence and adaptive incident response, offering a robust defense framework for AAM’s digital infrastructure.

3.8 Demand Modeling

AAM demand modeling is challenged by complex, multi-faceted data and dynamic operational environments that traditional models often fail to capture [Long *et al.*, 2023]. Neurosymbolic AI, which fuses deep neural networks’ pattern

recognition with the structured reasoning of symbolic systems, has emerged as a promising yet still exploratory approach in this domain [Garcez and Lamb, 2023]. Key hurdles include aligning unstructured sensor and user behavior data with the symbolic representations of regulatory and operational constraints, thereby affecting model interpretability, scalability, and robustness. Recent advancements, such as the integration of decision trees with neural networks [Acharya *et al.*, 2025] and the development of Probabilistic Mission Design architectures [Kohaut *et al.*, 2024], have demonstrated improved forecasting accuracy and novel methods for embedding legal frameworks into AI models.

4 Potential Case Studies

4.1 FAA Roadmap for Artificial Intelligence Safety Assurance

The FAA Roadmap for AI Safety Assurance¹ establishes a structured, phased approach for integrating artificial intelligence into aviation, a domain traditionally governed by deterministic safety protocols. The roadmap advocates for an incremental, risk-based strategy, initially focusing on low-criticality systems to gather operational data and refine safety assurance processes before transitioning to high-stakes, safety-critical applications. This iterative engineering approach ensures that early-stage insights inform the deployment of more complex AI-driven systems. A key aspect of the roadmap is its distinction between learned AI (static, pre-trained models) and learning AI (adaptive, real-time evolving systems). While learning AI enhances adaptability and responsiveness, it poses significant certification challenges due to its inherent unpredictability, limited explainability, and potential for bias issues that conventional safety methodologies struggle to address.

The roadmap underscores the importance of regulatory alignment with existing aviation safety frameworks, advocating for industry collaboration among regulatory bodies, government agencies, and stakeholders to establish standardized AI assurance protocols. However, it highlights critical gaps, including the lack of robust certification methods for evolving AI, inadequate bias mitigation strategies, and the absence of comprehensive regulatory guidelines for autonomous AI-driven aviation systems. Ethical considerations, essential for public trust and system accountability, receive minimal focus, indicating an area requiring further exploration. Addressing these challenges is essential for ensuring safe, reliable AI integration into aviation while maintaining compliance with rigorous industry standards. By integrating Neurosymbolic AI, the FAA can establish a robust, certifiable AI framework that enhances aviation safety while accelerating the adoption of AI in autonomous flight operations.

4.2 EASA AI Roadmap

EASA AI Roadmap 1.0² is the initial roadmap established a foundational strategy for integrating AI—primarily ma-

¹https://www.faa.gov/aircraft/air_cert/step/roadmap_for_AI_safety_assurance

²<https://www.easa.europa.eu/en/downloads/109668/en>

chine learning and deep learning—into aviation, addressing its transformative potential across aircraft design, operations, maintenance, air traffic management, drone applications, and safety risk management. It underscored the critical need for robust trustworthiness frameworks through enhanced explainability, human-AI collaboration, and phased certification processes to ensure safety and reliability. EASA AI Roadmap 2.0³, expands the scope by incorporating hybrid AI approaches such as Neurosymbolic methods, while further emphasizing cybersecurity and AI safety assurance within a human-centric framework. This version reaffirms AI’s role as an augmentation tool, integrating advanced digital twin technologies and predictive maintenance to enhance operational efficiency and safety. Aligned with the regulatory requirements of the EU AI Act, it introduces risk-based categorization of AI systems, mandating transparency, explainability, and human oversight.

5 Risk and Challenges

Neurosymbolic AI has the potential to enhance the efficiency, safety, and scalability of AAM, but it also presents several challenges and risks as depicted in the Figure 4.

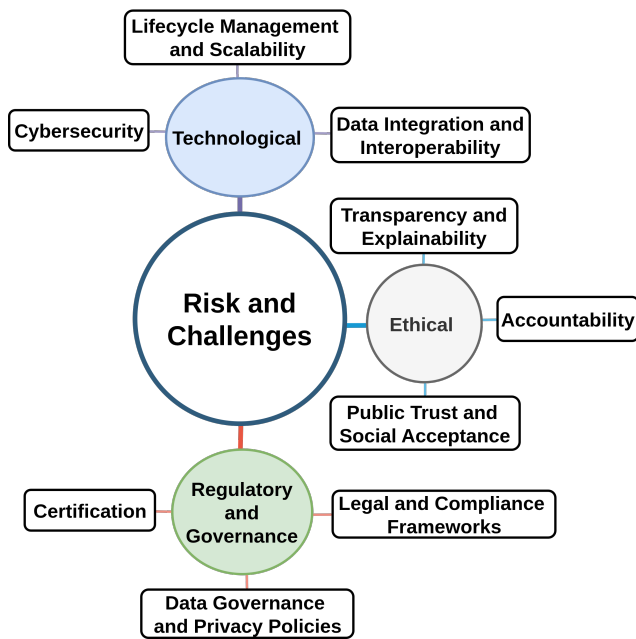


Figure 4: Risks and Challenges

5.1 Technological Challenges

Data Integration and Interoperability

Data integration in aviation face challenges due to heterogeneous sensor sources. Multi-sensor fusion can be homogeneous or heterogeneous, with the latter requiring precise synchronization. AI enhanced integration with air traffic management demands legacy system re-engineering while ensuring

safety compliance[Khanmohamadi and Guerrieri, 2024]. Interoperability is critical, necessitating standardized data formats, exchange protocols, and interface standards. The FAA’s Next Gen program focuses on digital system interoperability to support AI-driven airspace operations⁴. Advanced fusion methods leverage variance and covariance estimations for optimal integration, while AI driven standardization enhances cross-platform data accessibility, as seen in the U.S. military’s Combined Joint All Domain Command and Control (CJADC2)⁵ system.

Cybersecurity

The integration of neural networks and symbolic reasoning, although potentially powerful, introduces new complexities in system design and validation[Jalaian and Bastian, 2023]. Ensuring the robustness of these hybrid systems against adversarial attacks targeting both the neural and symbolic components requires extensive testing under diverse, real-world conditions [Hagos and Rawat, 2024]. Moreover, the interpretability of Neurosymbolic models, crucial for safety-critical applications like AAM, remains an ongoing research challenge. The dynamic nature of cyber threats also necessitates continuous adaptation of these systems, raising questions about long-term reliability and maintenance [Hagos and Rawat, 2024]. The increased computational requirements of Neurosymbolic AI may pose challenges for real-time decision-making in resource-constrained AAM environments[Jalaian and Bastian, 2023].

Lifecycle Management and Scalability

As AAM systems continually evolve, the necessity for robust version control, formal verification, and comprehensive testing becomes paramount to maintain both operational integrity and certification. Moreover, the scalability of these frameworks is tested by the heterogeneous nature of AAM fleets, which encompass vehicles with diverse sensor modalities and performance profiles, potentially leading to compatibility issues and performance bottlenecks if not managed through modular and adaptive architectures [Kohaut *et al.*, 2024]. While the modularity of Neurosymbolic systems offers a pathway to seamless integration and system evolution, their practical implementation in AAM must address the inherent trade-offs between adaptability and the rigorous demands of safety and regulatory compliance.

5.2 Ethical Challenges

Accountability

Accountability issues remain particularly complex in scenarios where autonomous decisions lead to system failures or incidents, determining responsibility among manufacturers, programmers, and operators is nontrivial, highlighting the urgent need for clear regulatory frameworks and transparent decision-making processes that can apportion accountability effectively [Novelli *et al.*, 2024]. The integration of neural and symbolic methods, while promising enhanced interpretability, does not inherently mitigate risks related to bias

⁴https://www.faa.gov/aircraft/air_cert/step/roadmap_for_AI_safety_assurance

⁵<https://www.ai.mil/Initiatives/CJADC2/>

³<https://www.easa.europa.eu/en/downloads/137919/en>

and fairness [Ferrer *et al.*, 2021]. Inadequate or skewed training data may propagate discriminatory outcomes, thereby compromising equitable service delivery across diverse operational contexts. Addressing these concerns requires embedding explicit fairness constraints and ethical auditing protocols within the Neurosymbolic framework to ensure decisions are both transparent and impartial.

Transparency and Explainability for Trust

Transparency and explainability foster trust and enable scrutiny of automated decision-making processes. By integrating symbolic reasoning with neural networks, these systems can provide more interpretable decision pathways, allowing stakeholders to understand the rationale behind actions taken. Furthermore, implementing robust auditability mechanisms through comprehensive logging and traceability is essential for post-incident reviews, system improvement, and regulatory compliance. These features are vital for maintaining public trust and developing a resilient framework that can adapt to the evolving demands of AAM environments. However, it is important to note that achieving a balance between transparency and system performance remains a challenge, as increasing explainability may impact the efficiency of complex models [Acharya *et al.*, 2023].

Public Trust and Social Acceptance

Public trust and social acceptance are crucial for addressing concerns related to data privacy, ethics, and system accountability in the rapidly evolving field of AAM. Despite hybrid approaches like Probabilistic Mission Design (ProMis) [Kohaut *et al.*, 2024] aiming to enhance transparency by merging symbolic reasoning with neural networks, significant risks remain. Current systems still fall short of fully explaining complex decision-making processes, undermining accountability and eroding public trust. Moreover, without sustained, effective collaboration among regulators, manufacturers, and the public, the ethical and responsible deployment of AAM systems is jeopardized, potentially stalling their broader acceptance and market integration. The transparency is particularly important when dealing with sensitive operational data and ensuring compliance with data protection regulations⁶.

5.3 Regulatory and Governance Challenges

Certification

The FAA has made limited progress in determining appropriate certification paths for AAM aircraft, partly due to the lack of established airworthiness standards and operational regulations for novel features⁷. This regulatory uncertainty is compounded by the need for new approaches to verify and validate Neurosymbolic systems, which combine probabilistic neural components with deterministic symbolic rules. The development of robust testing protocols that can rigorously verify system behavior under rare or extreme conditions remains an unresolved challenge, requiring advanced simulation environments, stress-testing methodologies, and formal verification techniques.

⁶<https://rm.coe.int/prems-107320-gbr-2018-compli-cahai-couv-texte-a4-bat-web/1680a0c17a>

⁷https://www.oig.dot.gov/sites/default/files/FAA_AAM_Certification_Final_Audit_Report_w_NGO_response.pdf

Legal and Compliance Frameworks

The dual nature of Neurosymbolic systems, merging the adaptive, data-driven insights of neural networks with the clear, rule-based logic of symbolic reasoning, complicates the clear attribution of accountability in incident scenarios. This complexity necessitates the establishment of legal guidelines that delineate responsibility among manufacturers, programmers, and operators, ensuring that each stakeholder is adequately accountable for AI-driven decisions. Moreover, given the global nature of AAM operations, harmonizing international standards becomes imperative to ensure that AI systems adhere to uniform safety and legal norms across diverse regulatory landscapes [Bengio *et al.*, 2025]. Achieving this will require coordinated efforts between legal experts, regulatory bodies, and technical developers to create transnational frameworks that not only standardize certification protocols but also provide mechanisms for continuous oversight and post-incident analysis.

Data Governance and Privacy Policies

The implementation of dynamic, adaptive policy frameworks is essential to ensure compliance with evolving data protection laws and international privacy standards. This approach requires continuous updates to data management practices, aligning with regulations like General Data Protection Regulation (GDPR) and California Consumer Privacy Act (CCPA) [Van Dalsem *et al.*, 2021]. Transparency in data use is achieved through standardized protocols for data sharing among stakeholders, while sensitive information is protected using advanced encryption, anonymization, and access control measures. This focus on compliance and transparency not only builds trust among stakeholders but also creates a resilient ecosystem with meticulous documentation of data provenance and maintenance of audit trails for rigorous oversight⁸.

6 Conclusion

Neurosymbolic AI holds significant promise for advancing AAM by delivering transparent, adaptable, and high-performance solutions tailored to meet the stringent demands of modern air transportation. However, empirical insights from FAA and EASA initiatives expose pressing issues like scalability, data integration, cybersecurity, and transparent decision-making that must be rigorously addressed through standardized protocols and robust validation. Continued interdisciplinary research, along with sustained collaboration among academia, industry, and regulatory agencies, is essential to fully harness the potential of these hybrid approaches and to drive the next generation of safe, efficient, and sustainable air mobility solutions.

Acknowledgments

This material is based upon work supported by the NASA Aeronautics Research Mission Directorate (ARMD) University Leadership Initiative (ULI) under cooperative agreement

⁸https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/06/ai-data-governance-and-privacy_2ac13a42/2476b1a4-en.pdf

number 80NSSC23M0059. This research was also partially supported by the U.S. National Science Foundation through Grant No. 2317117 and Grant No. 2309760.

References

- [Acharya *et al.*, 2023] Kamal Acharya, Waleed Raza, Carlos Dourado, Alvaro Velasquez, and Houbing Herbert Song. Neurosymbolic reinforcement learning and planning: A survey. *IEEE Transactions on Artificial Intelligence*, 2023.
- [Acharya *et al.*, 2025] Kamal Acharya, Mehul Lad, Liang Sun, and Houbing Song. Neurosymbolic ai for travel demand prediction: Integrating decision tree rules into neural networks. *arXiv preprint arXiv:2502.01680*, 2025.
- [Barrera *et al.*, 2022] Thomas P Barrera, James R Bond, Marty Bradley, Rob Gitzendanner, Eric C Darcy, Michael Armstrong, and Chao-Yang Wang. Next-generation aviation li-ion battery technologies—enabling electrified aircraft. *The Electrochemical Society Interface*, 31(3):69, 2022.
- [Bengio *et al.*, 2025] Yoshua Bengio, Sören Mindermann, Daniel Privitera, Tamay Besiroglu, Rishi Bommasani, Stephen Casper, Yejin Choi, Philip Fox, Ben Garfinkel, Danielle Goldfarb, et al. International ai safety report. *arXiv preprint arXiv:2501.17805*, 2025.
- [Bridgelall and Tolliver, 2024] Raj Bridgelall and Denver Tolliver. Transforming healthcare delivery with advanced air mobility: A rural study with gis-based optimization. *Sustainability*, 16(13):5709, 2024.
- [Cohen and Shaheen, 2024] Adam Cohen and Susan Shaheen. Advanced air mobility: Opportunities, challenges, and research needs for the state of california (2023-2030). 2024.
- [Deniz *et al.*, 2024] Sabrullah Deniz, Yufei Wu, Yang Shi, and Zhenbo Wang. A reinforcement learning approach to vehicle coordination for structured advanced air mobility. *Green Energy and Intelligent Transportation*, 3(2):100157, 2024.
- [Do *et al.*, 2024] Hieu Do, Ramon Delgado Pulgar, Gábor Fodor, and Zhiqiang Qi. Cellular connectivity for advanced air mobility: Use cases and beamforming approaches. *IEEE Communications Standards Magazine*, 8(1):65–71, 2024.
- [Dorn-Gomba *et al.*, 2020] Lea Dorn-Gomba, John Ramoul, John Reimers, and Ali Emadi. Power electronic converters in electric aircraft: Current status, challenges, and emerging technologies. *IEEE Transactions on Transportation Electrification*, 6(4):1648–1664, 2020.
- [Dulia *et al.*, 2021] Esrat F Dulia, Mir S Sabuj, and Syed AM Shihab. Benefits of advanced air mobility for society and environment: A case study of ohio. *Applied Sciences*, 12(1):207, 2021.
- [Ellis *et al.*, 2020] Kyle Ellis, John Koelling, Misty Davies, and Paul Krois. In-time system-wide safety assurance (issa) concept of operations and design considerations for urban air mobility (uam). 2020.
- [Ferrer *et al.*, 2021] Xavier Ferrer, Tom Van Nuenen, Jose M Such, Mark Coté, and Natalia Criado. Bias and discrimination in ai: a cross-disciplinary perspective. *IEEE Technology and Society Magazine*, 40(2):72–80, 2021.
- [Garcez and Lamb, 2023] Artur d’Avila Garcez and Luis C Lamb. Neurosymbolic ai: The 3 rd wave. *Artificial Intelligence Review*, 56(11):12387–12406, 2023.
- [Gilpin and Ilievski, 2021] Leilani H Gilpin and Filip Ilievski. Neuro-symbolic reasoning in the traffic domain. *J AI Res*, 15(3):123–145, 2021.
- [Goyal and Cohen, 2022] Rohit Goyal and Adam Cohen. Advanced air mobility: Opportunities and challenges deploying evtols for air ambulance service. *Applied Sciences*, 12(3):1183, 2022.
- [Goyal *et al.*, 2021] Rohit Goyal, Colleen Reiche, Chris Fernando, and Adam Cohen. Advanced air mobility: Demand analysis and market potential of the airport shuttle and air taxi markets. *Sustainability*, 13(13):7421, 2021.
- [Grov *et al.*, 2024] Gudmund Grov, Jonas Halvorsen, Magnus Wiik Eckhoff, Bjørn Jervell Hansen, Martin Eian, and Vasileios Mavroeidis. On the use of neurosymbolic ai for defending against cyber attacks. In *International Conference on Neural-Symbolic Learning and Reasoning*, pages 119–140. Springer, 2024.
- [Gür *et al.*, 2024] Gürkan Gür, Kamesh Namuduri, Stefano Savazzi, Saba Al-Rubaye, Ashwin Ashok, Sven G Bilén, Ivan Petrunin, Marco Hernandez, Shane Nicoll, and Batool Dalloul. V2v uas communications and use cases for advanced air mobility. In *2024 IEEE Conference on Standards for Communications and Networking (CSCN)*, pages 103–107. IEEE, 2024.
- [Hagos and Rawat, 2024] Desta Haileselassie Hagos and Danda B Rawat. Neuro-symbolic ai for military applications. *IEEE Transactions on Artificial Intelligence*, 2024.
- [Hill *et al.*, 2020] Brian P Hill, Dwight DeCarme, Matt Metcalfe, Christine Griffin, Sterling Wiggins, Chris Metts, Bill Bastedo, Michael D Patterson, and Nancy L Mendonca. Uam vision concept of operations (conops) uam maturity level (uml) 4. 2020.
- [Jalaian and Bastian, 2023] Brian Jalaian and Nathaniel D. Bastian. Neurosymbolic ai in cybersecurity: Bridging pattern recognition and symbolic reasoning. In *IEEE MILCOM 2023*, pages 268–273, October 2023. ISSN: 2155-7586.
- [Jansen *et al.*, 2020] Nils Jansen, Bettina Könighofer, Sebastian Junges, Alex Serban, and Roderick Bloem. Safe reinforcement learning using probabilistic shields. In *31st International Conference on Concurrency Theory (CONCUR 2020)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2020.
- [Johnson and Silva, 2022] W Johnson and C Silva. Nasa concept vehicles and the engineering of advanced air mobility aircraft. *The Aeronautical Journal*, 126(1295):59–91, 2022.

- [Karampinis *et al.*, 2024] Vasileios Karampinis, Anastasios Arsenos, Orfeas Filippopoulos, Evangelos Petrongonas, Christos Skliros, Dimitrios Kollias, Stefanos Kollias, and Athanasios Voulodimos. Ensuring uav safety: A vision-only and real-time framework for collision avoidance through object detection, tracking, and distance estimation. *arXiv preprint arXiv:2405.06749*, 2024.
- [Kautz, 2022] Henry Kautz. The third ai summer: Aai robert s. engelmore memorial lecture. *Ai magazine*, 43(1):105–125, 2022.
- [Khanmohamadi and Guerrieri, 2024] Masoud Khanmohamadi and Marco Guerrieri. Advanced sensor technologies in cavs for traditional and smart road condition monitoring: A review. *Sustainability*, 16(19), 2024.
- [Kiesewetter *et al.*, 2023] Lukas Kiesewetter, Kazi Hassan Shakib, Paramvir Singh, Mizanur Rahman, Bhupendra Khandelwal, Sudarshan Kumar, and Krishna Shah. A holistic review of the current state of research on aircraft design concepts and consideration for advanced air mobility applications. *Progress in Aerospace Sciences*, 142:100949, 2023.
- [Kimura *et al.*, 2021] Daiki Kimura, Subhajit Chaudhury, Akifumi Wachi, Ryosuke Kohita, Asim Munawar, Michiaki Tatsubori, and Alexander Gray. Reinforcement learning with external knowledge by using logical neural networks. *arXiv preprint arXiv:2103.02363*, 2021.
- [Kohaut *et al.*, 2024] Simon Kohaut, Benedict Flade, Daniel Ochs, Devendra Singh Dhami, Julian Eggert, and Kristian Kersting. Probabilistic mission design in neuro-symbolic systems. *arXiv preprint arXiv:2501.01439*, 2024.
- [Long *et al.*, 2023] Qi Long, Jun Ma, Feifeng Jiang, and Christopher John Webster. Demand analysis in urban air mobility: A literature review. *Journal of Air Transport Management*, 112:102436, 2023.
- [Mao *et al.*, 2019] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B Tenenbaum, and Jiajun Wu. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. *arXiv preprint arXiv:1904.12584*, 2019.
- [Murça *et al.*, 2024] Mayara Condé Rocha Murça, Wallace Silva Sant’Anna Souza, and João Vitor Turchetti Ribeiro. Characterizing performance tradeoffs from air traffic management services for advanced air mobility. *Journal of Aerospace Information Systems*, pages 1–14, 2024.
- [Namuduri, 2023] Kamesh Namuduri. Digital twin approach for integrated airspace management with applications to advanced air mobility services. *IEEE Open Journal of Vehicular Technology*, 2023.
- [National Academies of Sciences *et al.*, 2018] National Academies of Sciences *et al.* *In-time Aviation Safety Management: Challenges and Research for an Evolving Aviation System*. National Academies Press, 2018.
- [Nguyen *et al.*, 2022] Van-Thai Nguyen, Phuc Do, Alexandre Vosin, and Benoit Iung. Artificial-intelligence-based maintenance decision-making and optimization for multi-state component systems. *Reliability Engineering & System Safety*, 228:108757, 2022.
- [Novelli *et al.*, 2024] Claudio Novelli, Mariarosaria Taddeo, and Luciano Floridi. Accountability in artificial intelligence: what it is and how it works. *Ai & Society*, 39(4):1871–1882, 2024.
- [Petit and Shladover, 2014] Jonathan Petit and Steven E Shladover. Potential cyberattacks on automated vehicles. *IEEE Transactions on Intelligent transportation systems*, 16(2):546–556, 2014.
- [Raghunatha *et al.*, 2023] Aishwarya Raghunatha, Patrik Thollander, and Stephan Barthel. Addressing the emergence of drones—a policy development framework for regional drone transportation systems. *Transportation Research Interdisciplinary Perspectives*, 18:100795, 2023.
- [Rizzi and Scata Jr, 2022] Stephen A Rizzi and Donald S Scata Jr. Urban air mobility community noise test planning. In *183rd Meeting of the Acoustical Society of America*, 2022.
- [Sharifi *et al.*, 2023] Iman Sharifi, Mustafa Yildirim, and Saber Fallah. Towards safe autonomous driving policies using a neuro-symbolic deep reinforcement learning approach. *arXiv preprint arXiv:2307.01316*, 2023.
- [Silva *et al.*, 2018] Christopher Silva, Wayne R Johnson, Eduardo Solis, Michael D Patterson, and Kevin R Antcliff. Vtol urban air mobility concept vehicles for technology development. In *2018 Aviation Technology, Integration, and Operations Conference*, page 3847, 2018.
- [Siyaev *et al.*, 2023] Aziz Siyaev, Dilmurod Valiev, and Geun-Sik Jo. Interaction with industrial digital twin using neuro-symbolic reasoning. *Sensors*, 23(3):1729, 2023.
- [Van Dalsem *et al.*, 2021] William Van Dalsem, Sandeep Shetye, Aditya N Das, Kalmanje S Krishnakumar, Sandy Lozito, Kenneth Freeman, Aaron Swank, Peter Shannon, and Luka Tomljenovic. A data & reasoning fabric to enable advanced air mobility. In *AIAA Scitech 2021 Forum*, page 2033, 2021.
- [Wang *et al.*, 2024] Wenguan Wang, Yi Yang, and Fei Wu. Towards data-and knowledge-driven ai: a survey on neuro-symbolic computing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [Yang *et al.*, 2018] Shunkun Yang, Chong Bian, Xing Li, Lin Tan, and Dongxiao Tang. Optimized fault diagnosis based on fmea-style cbr and bn for embedded software system. *The International Journal of Advanced Manufacturing Technology*, 94:3441–3453, 2018.
- [Yoo *et al.*, 2022] Jaeho Yoo, Yunseon Choe, and Soo-i Rim. Risk perceptions using urban and advanced air mobility (uam/aam) by applying a mixed method approach. *Sustainability*, 14(24):16338, 2022.
- [Yu *et al.*, 2023] Dongran Yu, Bo Yang, Dayou Liu, Hui Wang, and Shirui Pan. A survey on neural-symbolic learning systems. *Neural Networks*, 2023.