38th International Electric Vehicle Symposium and Exhibition (EVS38)
Göteborg, Sweden, June 15-18, 2025

Ensuring Trustworthy and Compliant AI for Electric Fleet Predictive Maintenance: An Argument-Based Assurance Framework Aligned with the EU AI Act

Ahmet Bilal Aytekin[1], Ahu Ece Hartavi [2]
[1] *University of Genoa, Tarello Institute for Legal Philosophy*
[2] *Ahu Ece Hartavi (corresponding author), University of Surrey, Automotive Engineering Centre,*
a.hartavikarci@surrey.ac.uk

Executive Summary

The EU AI Act aims to regulate high-risk AI systems, with a particular focus on sectors like automotive, where electric fleets demand advanced predictive maintenance (PdM) to ensure reliability and operational efficiency. In electric vehicle (EV) fleets, where PdM is essential, compliance must extend beyond mere technical specifications to encompass ethical principles that prioritize safety, transparency, and trust. Traditional compliance approaches often risk superficial adherence to regulations, potentially overlooking critical ethical considerations in the context of high-stakes AI applications. An ethical assurance framework is thus proposed to embed ethical reflection throughout the development of AI systems used in EV fleets, particularly for predictive maintenance. This framework leverages a model based on ethical principles, fostering multidisciplinary dialogue and promoting explainability as a core requirement. Explainability is vital for user empowerment and regulatory compliance, offering transparency in AI decisions that impact safety protocols within EV fleets.

AI – Artificial intelligence for EVs, Electric Vehicles, Electric autonomous vehicle safety and security, Standardization, Health and Safety Considerations

1 Introduction

The automotive industry is undergoing a significant transformation driven by electrification and the increasing integration of Artificial Intelligence (AI) systems. Within this context, Predictive Maintenance (PdM) utilizing AI has emerged as a critical technology for managing Electric Vehicle (EV) fleets, promising enhanced operational efficiency, component longevity, and overall reliability [1]. However, the deployment of such AI systems, particularly those influencing safety-critical functions or maintenance schedules, introduces complex regulatory and ethical considerations. Notably, under the European Union's AI Act, components related to vehicle safety and systems governed by specific Union harmonization legislation may be classified as 'high-risk' pursuant to Article 6(1) and Annex I. Satisfying the stringent requirements for these systems demands more than demonstrating technical functionality; it necessitates robust evidence of trustworthiness, encompassing principles like safety, transparency, fairness, and accountability [2]. Conventional approaches to compliance, often relying on standardized checklists or post-hoc audits, frequently prove inadequate for capturing the nuanced ethical risks associated with adaptive AI systems and may lead to superficial adherence rather than genuine ethical integration [3].

This paper addresses the challenge of operationalizing ethical principles within the development lifecycle of high-risk AI for EV PdM. We present a structured, argument-based ethical assurance framework designed to systematically integrate ethical reflection and provide verifiable evidence of trustworthiness, using our project as a case study. Our primary contribution is the articulation and preliminary application of this framework, which, unlike general AI ethics guidelines or purely technical validation procedures, leverages goal-based argumentation to explicitly link ethical requirements (aligned with the AI Act's underlying principles) to concrete system design features, operational parameters, and generated evidence. This approach aims to enhance transparency and build justifiable confidence in AI-driven safety and maintenance protocols for EV fleets.

2 Ethical Assurance: Foundational Concepts and Principles

Ethical assurance represents a systematic and evidence-based process for establishing justifiable confidence in the ethical properties of AI systems throughout their lifecycle. It builds upon the established practices of traditional safety assurance, which focuses primarily on preventing physical harm, but extends the scope significantly to encompass a broader range of ethical considerations, including fairness, accountability, transparency, and the preservation of human autonomy. [4] This approach fundamentally rejects the notion of ethics as a peripheral concern or a compliance layer applied retrospectively. Instead, ethical assurance advocates for embedding ethical analysis and mitigation strategies proactively within the entire development process. It champions, as articulated by Burr et al., an "end-to-end habit of critical reflection and deliberation throughout all stages of a research or innovation project's lifecycle" [5] demanding consideration of ethical implications from initial conception and requirements gathering, through data management, model design and training, to deployment, operation, and decommissioning.

This integrated and anticipatory stance distinguishes ethical assurance from potentially superficial methods, basic ethics checklists such as The Assessment List for Trustworthy Artificial Intelligence (ALTAI) [6], which may fail to capture system-specific nuances or the dynamic nature of AI behavior [7]. It also moves beyond purely technical verification and validation, acknowledging that ethical risks often arise from the interplay between the technology, its users, and the socio-organizational context [8]. By encouraging developers to proactively identify, analyze, and address potential ethical challenges before they manifest as harms, ethical assurance aims to foster the development of AI systems that are not only functional and compliant but also demonstrably trustworthy.

Central to operationalizing ethical assurance is the adoption of a shared ethical framework that allows diverse stakeholders to reason consistently about complex ethical trade-offs. For this purpose, we leverage the well-established "4+1 Principles" model, drawing from bioethics and adapted for AI ethics, which includes: Beneficence (promoting well-being), Non-maleficence (avoiding harm), Respect for Human Autonomy (ensuring human control and dignity), Justice (fairness in distribution and process), and the crucial addition of Transparency/Explainability[4]. This set of principles provides a robust and widely recognized taxonomy for identifying and evaluating the potential ethical impacts of AI systems, such as the PdMAI application discussed herein. Crucially, it offers a common vocabulary, facilitating essential dialogue and collaboration among multidisciplinary teams composed of engineers, data scientists, legal experts, ethicists, and domain specialists (e.g., fleet managers). While not explicitly mandated in this form by the EU AI Act, these principles resonate strongly with the fundamental rights, safety considerations, and ethical values underpinning the regulation, thereby providing a practical lens for navigating its requirements.

3 Aligning Ethical Assurance with the EU AI Act for Automotive AI

Grounded in the principles outlined earlier, the ethical assurance framework provides a practical and structured methodology for addressing the complex requirements of the EU AI Act, particularly those arising between Articles 8 and 28, with a specific focus on the demanding context of high-risk automotive AI systems such as PdMAI. Its primary contribution lies in translating high-level regulatory obligations and ethical principles into verifiable claims supported by concrete evidence.

3.1 Enhancing Transparency and Explainability

A core tenet of the AI Act is the requirement for transparency and the provision of information, particularly for high-risk AI systems, as mandated under Article 13, enabling users and regulators to understand the systems' functioning and intended use. Ethical assurance directly supports this by prioritizing explainability (XAI) not merely as a technical feature, but as a fundamental component of trustworthy AI. As highlighted by Sovrano et al., explainability serves crucial "user-empowering" functions (allowing operators to understand, trust, and appropriately interact with PdMAI predictions) and "compliance-oriented" functions (facilitating audits and regulatory assessments)[9]. Our case-study (detailed in Section 4 and Figure 1) provides a systematic means to evaluate the quality and appropriateness of the explanations generated by the AI system. It allows developers to structure arguments demonstrating how the provided explanations meet necessary criteria – such as being "risk-focused, model-agnostic, goal-aware, and intelligible and accessible" [9] – ensuring they are meaningful to diverse stakeholders, from fleet managers and maintenance technicians to auditors and regulators, and tailored to the potential risks involved. This moves beyond simply stating a system is explainable, towards demonstrating why and how it achieves adequate transparency for its context.

3.2 Addressing High-Risk Automotive Systems

The significance of this approach is amplified for AI systems integral to motor vehicles. Safety-critical components and systems subject to existing Union harmonization legislation, such as Regulation (EU) 2019/2144 concerning the general safety of motor vehicles, trigger the 'high-risk' classification under Article 6(1) and Annex I of the AI Act when AI is incorporated. While the Act outlines specific requirements for high-risk AI (e.g., concerning data governance, technical documentation, human oversight, accuracy, robustness, cybersecurity), navigating these in the context of complex, adaptive AI systems presents challenges. Standard technical compliance checks may not suffice to build comprehensive trust or fully address underlying ethical considerations. Ethical assurance emerges as a necessary complementary methodology. It provides a structured framework to argue explicitly about how the AI system meets not only technical benchmarks but also broader trustworthiness criteria – including safety, fairness in predictions, accountability, and respect for operator autonomy – using reasoned arguments backed by diverse evidence (test results, simulation data, design documents, user feedback). This offers a more robust pathway to demonstrating genuine conformity and trustworthiness than relying solely on minimum regulatory adherence.

3.3 Facilitating Conformity Assessment and the Role of Intermediaries

A key output of applying this ethical assurance framework is the development of a structured 'ethical assurance case'. This documented artifact presents the claims about the system's ethical properties, the arguments supporting these claims, and the corresponding evidence. Such assurance cases can serve as invaluable documentation for conformity assessments, particularly when conducted by the independent, impartial third-party intermediaries (notified bodies) envisaged by the AI Act for certain high-risk systems [10]. By providing a transparent and logically structured justification for the system's trustworthiness, the assurance case can facilitate objective evaluation, mitigate the risks of conflicts of interest associated with purely internal assessments, and provide regulators with greater confidence. The credibility and competence of these notified bodies are, therefore, critical pillars for the effective functioning of the AI Act's regulatory ecosystem [10]. Ensuring their independence and expertise is paramount for leveraging ethical assurance to genuinely foster public trust.

3.4 Acknowledging Challenges and Future Directions

Despite its significant potential, ethical assurance is an evolving practice and faces notable challenges that warrant ongoing research and development. A primary difficulty lies in the operationalization and measurement of inherently qualitative ethical parameters. Developing robust, repeatable methodologies for quantifying aspects like non-financial benefits (e.g., driver well-being), assessing risks of non-physical harm (e.g., discriminatory outcomes, erosion of skills), or defining acceptable boundaries for human autonomy in AI-assisted decision-making remains a complex task [5]. This complexity makes establishing universally accepted thresholds for ethical acceptability challenging.

Furthermore, the empirical validation of ethical assurance's effectiveness is crucial. Rigorous studies are needed to determine how effectively these processes lead to the development of more ethical AI systems in practice and prevent the deployment of harmful ones [5]. Investigating the causal links between applying assurance methodologies and achieving desired ethical outcomes is a key area for future work.

Finally, practical implementation presents hurdles, especially concerning resource allocation and expertise. Integrating comprehensive ethical assurance requires time, skilled personnel (including ethicists, legal experts, and domain specialists alongside engineers), and organizational commitment, which may be challenging for smaller enterprises or public institutions[5]. Consequently, the development of supporting resources and tools is vital. This includes creating standardized argument patterns for common ethical challenges in AI, developing user-friendly software platforms for collaborative assurance case construction and management, and building repositories of relevant evidence types and assessment techniques. Such advancements would significantly enhance the feasibility, scalability, efficiency, and widespread adoption of ethical assurance practices across the industry.

4 Case Study: Predictive Maintenance in ESCALATE

To illustrate the practical application and demonstrate the utility of the ethical assurance framework, this section presents a case study centered on an AI-driven Predictive Maintenance system (hereafter referred to as PdMAI) developed within the context of the ESCALATE project. The primary objective within this case study was to construct a robust ethical assurance case focusing specifically on the explainability of the PdMAI system's outputs to its intended end-users: EV fleet operators and maintenance personnel. This focus aligns directly with the ethical principle of transparency, identified in Section 2 as crucial for trustworthy AI and resonant with the requirements of the EU AI Act discussed in Section 3.

The methodology employed follows the argument-based structure central to ethical assurance framework. Figure 1 visually represents this structure, adapted for the specific goal of demonstrating PdMAI explainability. The structure follows a hierarchical decomposition:

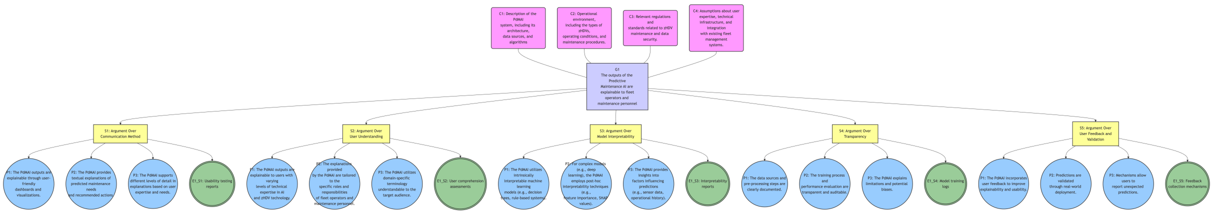

Figure 1: Argument-based assurance for Explainablity in PdM AI

4.1 Goal and Context: Establishing the Foundations for Trust

The overarching goal of the ethical assurance process for the PdMAI in the ESCALATE project, as defined in the provided document, is: The outputs of the PdMAI are explainable to fleet operators and maintenance personnel. This goal claim is deliberately focused on the practical needs of the users, who are key stakeholders in this process, and is not merely a statement about the technical capabilities of the system. This goal claim is also clear, specific and therefore, it provides a solid starting point for ethical analysis. As stated in the abstract, the context for this pattern is the ESCALATE project, specifically the development and deployment of PdMAI for zero-emission heavy-duty vehicles (zHDVs). It is explicitly stated that the PdMAI will be used by fleet operators and maintenance personnel to optimize maintenance schedules, reduce downtime, and improve the overall efficiency of zHDV fleets. However, this context statement needs further elaboration to more fully capture the intricacies of the ESCALATE project. Such a development requires considering the context claims:

- **C1:** It is crucial to provide a description of the PdMAI system, including its architecture, data sources, and algorithms. This documentation serves as the foundation for understanding the inner workings of the AI system, thus, it allows for proper evaluation of the system's properties[9].

- **C2:** The framework requires an understanding of the operational environment, including the types of zHDVs, operating conditions, and maintenance procedures. The types of zHDV may range from buses to cargo trucks, all of which may operate in highly divergent environments. Moreover, the range of operating conditions, such as variations in weather and climate as well as maintenance procedures, may also impact the efficacy and trustworthiness of the system.

- **C3:** There must be careful consideration of relevant regulations and standards related to zHDV maintenance and data security. By including this context claim, the framework acknowledges that technical and ethical standards do not operate in isolation. This helps bring transparency to any relevant regulatory expectations on the system itself.

- **C4:** Finally, assumptions about user expertise, technical infrastructure, and integration with existing fleet management systems" must be explicit. It is important to be clear about the technical skills and the infrastructure that are available to the user as the "usability, intelligibility, and accessibility"[9] of the system must vary according to these assumptions.

Considering these claims, the project team should be able to produce a more detailed context statement that guides them throughout the various stages of development and deployment. This context must be revisited and updated periodically to ensure that it remains relevant to the ongoing nature of the project.

4.2 Strategies and Property Claims: Breaking Down the Path to Explainability

Having established the high-level goal and context, the framework is further specified through strategies and property claims. Each strategy represents a specific line of reasoning aimed at demonstrating that the PdMAI meets its explainability goals.

- **Strategy 1: Argument Over Communication Method.** This strategy explores how the PdMAI communicates its findings to users. Its property claims include that "PdMAI outputs are explainable through user-friendly dashboards and visualizations", that "the PdMAI provides textual explanations of predicted maintenance needs and recommended actions", and that "the PdMAI supports different levels of detail in explanations based on user expertise and needs. These property claims demonstrate that different needs are considered in the design of the communication methods and that all stakeholders in the system have access to relevant information in varying degrees.

- **Strategy 2: Argument Over User Understanding.** This focuses on tailoring the explanations to the specific needs and abilities of the users. The property claims here include: The PdMAI outputs are explainable to users with varying levels of technical expertise in AI and zHDV technology, "The explanations provided by the PdMAI are tailored to the specific roles and responsibilities of fleet operators and maintenance personnel, and that The PdMAI utilizes domain-specific terminology that is understandable to the target audience. These property claims intend to ensure that the different stakeholders from the project are able to utilize and make sense of the output, regardless of their background knowledge.

- **Strategy 3: Argument Over Model Interpretability.** This focuses on the internal workings of the AI and seeks to make them more readily understandable. Its property claims include: The PdMAI utilizes intrinsically interpretable machine learning models (e.g., decision trees, rule-based systems) whenever possible, that For complex models (e.g., deep learning), the PdMAI employs post-hoc interpretability techniques (e.g., feature importance, SHAP values) to explain predictions", and finally that "The PdMAI provides insights into the factors influencing predictions, such as vehicle sensor data, operational history, and environmental conditions. As with the other claim, this set of properties seeks to provide as full a picture as possible of the various ways in which the model reaches conclusions.

- **Strategy 4: Argument Over Transparency.** Here, the focus is on the transparency of the data and model development processes. The associated property claims include that "The data sources and pre-processing steps used by the PdMAI are clearly documented, The training process and performance evaluation of the PdMAI models are transparent and auditable", and The PdMAI provides mechanisms for users to understand the limitations and potential biases of the system. As noted by Sovrano et al.a system's conformity to the regulations must be assessable [4], and this section of the ethical assurance framework is aimed to meet that assessment.

- **Strategy 5: Argument Over User Feedback and Validation.** This strategy demonstrates a commitment to ongoing improvement and refinement of the PdMAI system. The associated property claims include: "The PdMAI incorporates feedback from fleet operators and maintenance personnel to improve the explainability and usability of the system", "The accuracy and effectiveness of the PdMAI predictions are validated through real-world deployment and continuous monitoring", and finally, The PdMAI provides mechanisms for users to report and investigate unexpected or incorrect predictions. This ensures that ethical compliance is an iterative process and that is not achieved at a certain point only to be ignored going forward.

Together, these strategies and property claims create a detailed framework for ensuring that the PdMAI is not only technically proficient but also ethically acceptable and trustworthy. This is achieved by carefully considering the various stakeholders and the types of information they need to engage with AI systems as meaningfully as possible.

4.3 Evidence for the ESCALATE PdMAI: Substantiating the Claims

The above argument pattern is only as good as the evidence that supports it. Therefore, several categories of evidence must be gathered throughout the design, development, and deployment of the system. The document, "Evidences for Explainability Argument Pattern of PdMAI in ESCALATE Project" provides a detailed list of evidence types for each strategy:

- **For Strategy 1:** To support this, "usability testing reports" will be necessary. Moreover, the project team will need user interface design documentation" to outline their design choices as well as "examples of PdMAI outputs. These items will showcase the usability of the communication methods.

- **For Strategy 2:** This strategy needs to be supported by "user comprehension assessments" to showcase user understanding, as well as user interviews and focus groups to gather information regarding the user perception of the AI. Finally, these need to be accompanied with "training materials and documentation" that explain the functionality in an accessible way.

- **For Strategy 3:** The system's architecture needs "model documentation" as well as "interpretability reports" that visualize the factors of the decisions Finally, a "sensitivity analysis" is needed to showcase the robustness of the system's explanations.

- **For Strategy 4:** "Data provenance documentation" is required along with "model training logs" and "bias assessment reports". This will show that the data, the training process, and the testing process are thoroughly documented to ensure the transparency of the process.

- **For Strategy 5:** types include: "feedback collection mechanisms" used to gather data for system improvements, "performance monitoring reports" which are intended to showcase the system's efficacy, and finally "case studies" which help to demonstrate the value of the system in concrete examples . This will illustrate both the efficacy and continuous improvement of the AI as an ongoing process.

By gathering this diverse range of evidence, the ESCALATE project can demonstrate that the PdMAI not only possesses the technical capability but also adheres to core ethical principles outlined in this paper. In doing so, it can build confidence among stakeholders and provide the necessary justification for the PdMAI's trustworthiness and compliance. This will be particularly beneficial in the long run by enabling a feedback loop for continuous improvement that can further help navigate the complexities of this field.

In summary, this section has illustrated the practical application of the ethical assurance framework to a real-world project, i.e., the ESCALATE's PdMAI. By adopting the argument pattern and by collecting evidence to support the property claims, ESCALATE project team will demonstrate the trustworthiness and compliance of its product, while also navigating ethical challenges and building stakeholder confidence [9] [11]. This will demonstrate a practical approach for bridging the gap between the technical and ethical considerations when designing responsible AI systems, paving the way for the wider adoption of trustworthy AI solutions within the EV sector.

5 Conclusion

This paper has addressed the critical and timely challenge of ensuring that AI systems deployed for PdM in EV fleets are not only effective but also trustworthy, ethical, and compliant with burgeoning regulations like the EU AI Act. We argued that traditional compliance methods, often reliant on technical checklists, risk falling short of addressing the nuanced ethical considerations inherent in high-risk AI applications, potentially leading to systems that meet regulatory requirements superficially but lack genuine trustworthiness.

In response, we detailed an argument-based ethical assurance framework. Distinct from generic ethical guidelines or purely technical verification procedures, this framework provides a practical and systematic methodology for embedding ethical reflection throughout the AI lifecycle. Its core strength lies in its structured approach, adapting principles from safety assurance cases to explicitly link high-level ethical principles—such as transparency, fairness, accountability, and safety, which align with the values underpinning the AI Act—to concrete design choices, operational procedures, and verifiable evidence.

The primary contribution of this work is the articulation and practical demonstration of this methodology. As illustrated through the PdMAI explainability case study within the ESCALATE project (Section 4, Figure 1), the framework operationalizes abstract ethical requirements into specific, assessable claims supported by evidence. This structured argumentation facilitates a more rigorous and transparent approach to demonstrating AI trustworthiness compared to less formal methods. It provides developers with a clear process for building confidence in their systems and offers stakeholders, including auditors and regulators, a traceable justification for why an AI system should be considered ethically sound and compliant in its specific context of use. This directly addresses the need for detailed methodologies and demonstrations of effectiveness.

By systematically integrating ethical considerations and prioritizing properties like explainability, this framework empowers users, enhances accountability, and supports meaningful alignment with the stringent demands of the EU AI Act for high-risk systems. The insights shared offer tangible benefits for developers navigating the complex intersection of AI innovation and ethical responsibility in the automotive domain and beyond.

However, we acknowledge that ethical assurance is an evolving discipline. Significant challenges remain, particularly in developing standardized metrics for qualitative ethical parameters, empirically validating the long-term effectiveness of assurance processes in mitigating harm, and overcoming practical implementation hurdles related to resources and expertise within organizations. Future research should focus on refining measurement techniques, conducting empirical studies on the impact of ethical assurance, and crucially, developing accessible tools and standardized argument patterns to support the broader adoption of these vital practices.

Ultimately, adopting rigorous, evidence-based ethical assurance practices, as advocated herein, represents a crucial advancement. It moves beyond mere compliance towards the development and deployment of genuinely trustworthy AI systems – systems that not only enhance the efficiency and reliability of critical infrastructures like EV fleets but do so in a manner that respects human values and fosters societal trust in this transformative technology.

Acknowledgments

This research was funded by the ESCALATE (Grant Agreements No. 101096598 and 10063997) and FASTEST (Grant Agreements No. 101103755 and 10078013) projects, co-funded by the European Union's Horizon Research and Innovation Programs and Innovate UK.

References

[1] P. Arévalo, D. Ochoa-Correa, and E. Villa-Ávila, "A systematic review on the integration of artificial intelligence into energy management systems for electric vehicles: Recent advances and future perspectives," *World Electric Vehicle Journal*, vol. 15, no. 8, p. 364, 2024.

[2] N. Díaz-Rodríguez, J. D. Ser, M. Coeckelbergh, M. L. de Prado, E. Herrera-Viedma, and F. Herrera, "Connecting the dots in trustworthy artificial intelligence: From ai principles, ethics, and key requirements to responsible ai systems and regulation," 2023. [Online]. Available: https://arxiv.org/abs/2305.02231

[3] K. Lyle, S. Weller, G. Samuel, and A. M. Lucassen, "Beyond regulatory approaches to ethics: making space for ethical preparedness in healthcare research," *Journal of Medical Ethics*, vol. 49, no. 5, pp. 352–356, 2023.

[4] Z. Porter, I. Habli, J. McDermid, and M. Kaas, "A principles-based ethics assurance argument pattern for ai and autonomous systems," *AI and Ethics*, vol. 4, no. 2, pp. 593–616, 2024.

[5] C. Burr and D. Leslie, "Ethical assurance: a practical approach to the responsible design, development, and deployment of data-driven technologies," *AI and Ethics*, vol. 3, no. 1, pp. 73–98, 2023.

[6] P. Ala-Pietilä, Y. Bonnet, U. Bergmann, M. Bielikova, C. Bonefeld-Dahl, W. Bauer, L. Bouarfa, R. Chatila, M. Coeckelbergh, V. Dignum *et al.*, *The assessment list for trustworthy artificial intelligence (ALTAI).* European Commission, 2020.

[7] P. Brey and B. Dainow, "Ethics by design for artificial intelligence," *AI and Ethics*, vol. 4, no. 4, pp. 1265–1277, 2024.

[8] G. Bakirtzis, A. A. Tubella, A. Theodorou, D. Danks, and U. Topcu, "Navigating the sociotechnical labyrinth: Dynamic certification for responsible embodied ai," *arXiv preprint arXiv:2409.00015*, 2024.

[9] F. Sovrano, S. Sapienza, M. Palmirani, and F. Vitali, "Metrics, explainability and the european ai act proposal," *J*, vol. 5, no. 1, pp. 126–138, 2022.

[10] H. Fraser and J.-M. B. y Villarino, "Acceptable risks in europe's proposed ai act: reasonableness and other principles for deciding how much risk management is enough," *European Journal of Risk Regulation*, vol. 15, no. 2, pp. 431–446, 2024.

[11] A. EU, "European commission white paper on artificial intelligence–a european approach to excellence and trust," *URL https://ec. europa. eu/info/sites/info/files/commissionwhitepaper-artificial-intelligence-feb2020 en. pdf*, 2021.

Presenter Biography

Ahmet Bilal Aytekin had both its Bachelor from Istanbul Sehir University and Master of Laws degrees from University of Genoa. He is currently a Doctoral Researcher at University of Genoa and Visiting Researcher at University of Surrey working on AI applications in Regulatory Impact Assessment.