

RITSA: Toward a Retrieval-Augmented Generation System for Intelligent Transportation Systems Architecture

Afef Awadid¹, André Meyer-Vitali², Dominik Vereno³ and Maxence Gagnant¹

¹*Technological Research Institute SystemX, Palaiseau, France*

²*German Research Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany*

³*Josef Ressel Centre for Dependable System-of-Systems Engineering, Salzburg, Austria*

{afef.awadid, maxence.gagnant}@irt-systemx.fr, andre.meyer-vitali@dfki.de, dominik.vereno@fh-salzburg.ac.at

Keywords: Intelligent Transportation Systems (ITS), ITS Architecture Design and Modeling, Large Language Model (LLM), Retrieval-Augmented Generation (RAG), Modeling Assistant, ITS Reference Architectures.

Abstract: Intelligent Transportation Systems (ITS) have significantly transformed the transportation domain by addressing critical challenges such as traffic safety, cost, and energy efficiency. However, the increasing complexity of ITS—arising from the extensive range of applications and technologies they encompass—has made their architectural design modeling time-consuming and challenging, particularly for modelers lacking specialized expertise. Recent advancements in the literature suggest that large language model (LLM)-based modeling assistants offer a promising solution to mitigate these challenges. In this context, this paper introduces the RAG for Intelligent Transportation Systems Architecture (RITSA) project, which seeks to develop a retrieval-augmented generation (RAG) system to support ITS designers/ modelers throughout the architecture design process.

1 INTRODUCTION

Intelligent Transportation Systems (ITS), such as autonomous vehicles, are defined as “those systems utilizing synergistic technologies and systems engineering concepts to develop and improve transportation systems of all kinds” (Giesecke et al., 2016). The emergence of ITS has significantly transformed transportation, addressing critical challenges such as cost efficiency, traffic safety, speed, and user comfort (Waqar et al., 2023), (Elassy et al., 2024). However, the growing complexity of ITS stemming from the extensive range of applications and services they encompass (Damaj et al., 2022) has rendered their architectural design increasingly intricate.

To manage this complexity effectively, Model-Based Systems Engineering (MBSE) has been widely recognized as a pivotal approach (Friedenthal et al., 2014). MBSE leverages models to define the composition and interfaces of a system’s architectural layers, including functional, logical, and physical layers (Haomin et al., 2021). These system models serve as central artifacts in the systems engineering process, providing comprehensive representations of

the system and its environment. They incorporate multiple views to support activities such as planning, requirements specification, architecture definition, design, analysis, verification, and validation (INCOSE, 2023).

Nevertheless, building these models is time-consuming and challenging, particularly for modelers who lack expertise in the field of intelligent transportation. In this context, intelligent modeling assistants present a promising solution to address this issue. Against this backdrop, recent studies have investigated the potential of large language models (LLMs), such as OpenAI’s GPT-4, in providing support for modeling tasks (Combemale et al., 2023). LLMs can assist in model development, particularly in an ‘acceleration’ mode, where an initial model already exists (Barke et al., 2023). In such cases, LLMs contribute by extending the existing model through the integration of new entities or functionalities, or by enriching specific elements with additional characteristics. Thus, the primary role of LLMs in this context is to enhance and complete pre-existing models.

However, it is well established that LLMs can ‘hallucinate’ outputs—that is, generate false or

misleading content—particularly when applied outside their domain of expertise or when addressing complex or ambiguous topics (Aslam and Nisar, 2024). In safety-critical domains such as intelligent transportation systems, the consequences of inaccurate outputs can be significant. This highlights the importance of fine-tuning LLMs and grounding their responses using information derived from reference architectures specifically designed for intelligent transportation, such as ARC-IT (Architecture Reference for Cooperative and Intelligent Transportation) (ARC-IT Version 9.3, 2025). This underscores the importance of developing a Retrieval-Augmented Generation (RAG) system to specialize large language models (LLMs) by incorporating information retrieved from domain-specific sources in the field of intelligent transportation. RAG is an AI framework designed to retrieve facts from an external knowledge base, enabling LLMs to be grounded in accurate, up-to-date information while providing users with insight into their generative processes (Martineau, 2023).

With this in mind, this paper provides an overview of our ongoing research project, RITSA (Retrieval-Augmented Generation for Intelligent Transportation Systems).

The remainder of the paper is organized as follows: Section 2 discusses the theoretical background and reviews related work. Section 3 outlines the RITSA project, detailing its motivations, objectives, and solution architecture. Finally, Section 4 concludes the paper and outlines future work.

2 BACKGROUND

2.1 ITS and Reference Architectures

As previously noted, Intelligent Transportation Systems (ITS) are defined as 'systems utilizing synergistic technologies and systems engineering concepts to develop and improve transportation systems of all kinds' (Giesecke et al., 2016). ITS integrate advanced technologies, data analytics, and communication systems to enhance the efficiency, safety, and environmental sustainability of transportation networks (Khalid et al., 2019). By leveraging real-time data, sensor networks, and intelligent algorithms, ITS aim to alleviate traffic congestion, reduce travel times, enhance safety, and minimize environmental impacts (Barba et al., 2012), (Khalid et al., 2019).

ITS encompass a wide range of applications, from traffic management and control to autonomous

vehicles, all aimed at enhancing mobility experiences and addressing challenges associated with urbanization. Among the key components of ITS that contribute to improving transportation efficiency, safety, and sustainability are Vehicular Adhoc Networks, Intelligent Traffic Lights, Virtual Traffic Lights, and Mobility Prediction (Elassy et al., 2024).

Given the complexity of ITS, designing their architectures poses significant challenges, particularly for designers unfamiliar with such systems. A reliable approach to address this issue is to leverage existing domain-specific reference architectures. A reference architecture represents a shared and agreed system description, used by a community of interest as a guide for the development and evolution of systems (ISO/IEC/IEEE 42020, 2019). It is a key outcome of the domain design process (ISO/IEC 26552, 2019), encapsulating knowledge on designing system architectures within a specific application domain (Nakagawa et al., 2011). By incorporating proven architectural elements, reference architectures help mitigate risks and enhance reliability (Martínez-Fernández et al., 2015).

A prominent reference architecture in the field of intelligent transportation systems (ITS) is ARC-IT (Architecture Reference for Cooperative and Intelligent Transportation). ARC-IT provides a common framework for planning, defining, and integrating (ARC-IT Version 9.3, 2025), offering a common language for describing the architecture of these systems. It is structured around four key viewpoints: enterprise, functional, physical, and communication. The scope of ARC-IT is delineated by a set of ITS services, which refer to transportation services enabled by intelligent transportation systems. ARC-IT employs the concept of Service Packages to define the portions of the architecture necessary to implement a specific service. These Service Packages encompass elements from each of the four architectural views required to describe the service comprehensively. The areas of Services include Commercial Vehicle Operations, Public Transportation, Vehicle Safety, Weather, Public Safety, Traffic Management, Traveler Information, Data Management, Sustainable Travel, Parking Management, Maintenance and Construction, and Support (Kotsi et al., 2020).

Building on this foundation, an effective approach to assist in the architectural design of ITS is to leverage a RAG framework. This effectiveness arises from the framework's ability to combine the generative capabilities of large language models (LLMs) for producing comprehensive responses to

needs with the accuracy and reliability of reference architectures, which offer proven solutions tailored to the domain.

2.2 LLMs and RAG

Unlike traditional AI systems, large language models (LLMs) are designed to process and interpret unstructured data—such as natural language, images, and complex sensor data—in a manner that closely resembles human cognition (Katz et al., 2024), (Liu et al., 2024), (Liu, 2024). LLMs are extensively utilized across various subfields of natural language processing (NLP), addressing diverse tasks by conditioning the models on a few examples (few-shot learning) or on instructions describing the task (zero-shot learning). This method of guiding the language model is known as 'prompting,' and the development of effective prompts, whether manually or automatically, has emerged as a prominent area of research in NLP (Kojima et al., 2022).

These capabilities enable LLMs to process the vast and diverse data inputs characteristic of modern transportation systems, enhancing their adaptability and responsiveness to the dynamic challenges inherent in transportation. Furthermore, the integration of LLMs into ITS represents a significant advancement, facilitating the development of smarter and more responsive transportation systems that are better equipped to address future demands (Wandelt et al., 2024).

However, it is widely acknowledged that LLMs can generate hallucinated content—outputs that contradict established factual knowledge—when they lack sufficient information (Mondal et al., 2025). Such inaccuracies can have severe consequences, especially in safety-critical systems like Intelligent Transportation Systems (ITS). Safety-critical systems are defined as those whose failure can result in significant harm, including the loss of human life (Awadid et al., 2024).

To address this limitation of LLMs, one solution is to ground them in existing proven references, such as domain-specific reference architectures. This can be achieved through Retrieval-Augmented Generation (RAG), a methodology that integrates the generative capabilities of LLMs with information retrieval (Lewis et al., 2020). RAG allows the model to dynamically access and incorporate relevant external information during the generation process (Xia et al., 2024). This approach has been recognized as a promising strategy, not only for mitigating hallucination but also for enhancing the domain-specific expertise and temporal relevance of LLMs

(Chen et al., 2024), (Zhao et al., 2024). By doing so, RAG improves the controllability and interpretability of model outputs, making them more reliable and aligned with the application domain.

2.3 Related Work

The immense potential of LLMs in the field of transportation has sparked significant interest, not only among the general public but also within the research community, as evidenced by a growing body of literature. One prominent area of research focuses on the applications of LLMs in Intelligent Transportation Systems (ITS), with a particular emphasis on Autonomous Driving (AD). Autonomous Driving refers to 'the technology and systems that enable a vehicle to perceive its environment and make decisions independently, using artificial intelligence, computer vision, and sensor technologies to ensure safe driving' (Li et al., 2024).

The integration of autonomous driving (AD) and large language models (LLMs) holds significant promise, with the potential to enhance both user experience (e.g., through in-vehicle voice assistants and driving-related decision-making) and vehicle performance (e.g., by enabling complex environmental perception, intelligent anomaly detection, and optimized battery management) (Lei et al., 2023), (Singh, 2023), (Yang et al., 2023).

From the user perspective, large language models (LLMs) offer significant potential to enhance interactions with autonomous vehicles (AVs), including providing feedback on driving quality, incorporating personal preferences to influence driving behavior, and retrieving contextual route information such as traffic conditions or weather updates (Chen et al., 2023), (Du et al., 2023), (Xu et al., 2024). These capabilities are particularly valuable in scenarios where an autonomous driving (AD) system encounters uncertainties or requires user input to support decision-making. In such cases, LLMs could facilitate comprehensive dialogues with users, addressing safety, comfort, or route modification concerns in real time. Currently, communication between users and vehicles remains rudimentary, typically limited to keyword-based commands or a restricted set of predefined use cases. LLMs promise to enable more natural and intuitive interactions with broader functionality, potentially exceeding the original design intent of the vehicle manufacturer or operator. However, this expanded interaction scope raises potential security and safety concerns, as users

may inadvertently interfere with the AD system in unintended ways.

From the vehicle's perspective, integrating large language models (LLMs) holds significant promise for enhancing the performance and safety of autonomous driving (AD) systems (Wang et al., 2023). As outlined in (Fu et al., 2024), AD systems aiming to emulate human driving require three key abilities: reasoning, interpretation, and memorization. To explore the feasibility of employing LLMs in AD scenarios, a closed-loop system was developed, demonstrating the exceptional reasoning capabilities of LLMs in addressing long-tail scenarios. In a similar vein, (Cui et al., 2023) introduced a decision-making framework called DriveLLM, which integrates AD stacks with LLMs to enable commonsense reasoning in decision-making processes. DriveLLM's cyber-physical feedback system allowed it to iteratively learn from mistakes, thereby improving performance. Additionally, (Wen et al., 2023) proposed the DiLu framework, leveraging the emergent abilities of LLMs and real-world datasets to enhance AD tasks. Complementing this, (Sha et al., 2023) developed cognitive pathways and algorithms designed to bridge LLM-based reasoning with actionable driving commands. Their approach outperformed baseline methods in single-vehicle tasks and proved effective for managing complex behaviors, including multi-vehicle coordination.

(Deng et al., 2023) proposed an end-to-end test generation framework designed to automatically construct test scenarios within an autonomous driving simulator. The framework leverages ChatGPT-4 to extract key information from traffic rules, enabling the generation of targeted test cases. Experimental evaluations across diverse autonomous driving systems, traffic regulations, and road maps demonstrated the framework's effectiveness in identifying rule violations and uncovering known issues. To address information barriers arising from system and module heterogeneity, (Tian et al., 2023) introduced a Transformer-based unified framework called VistaGPT. This framework integrates Modular Federations of Vehicular (M-FoV) Transformers with Automated Autonomous (AutoAuto) Driving Systems to enable seamless interoperability.

Building on the preceding analysis of related work, existing studies have predominantly focused on the potential applications of LLMs and augmented LLMs in Intelligent Transportation Systems (ITS) from two primary perspectives: enhancing user experience and improving vehicle performance. While these perspectives are invaluable, they leave a

critical gap in understanding and leveraging LLMs to support ITS architecture designers and modelers. Specifically, the role of LLMs in facilitating the design, evaluation, and evolution of ITS architectures has been largely unexplored. Addressing this gap forms a core motivation for our research project, RITSA (Retrieval Augmented Generation for ITS), which aims to pioneer the integration of augmented-LLMs (RAG) into the ITS architecture design process, thereby extending their potential beyond user and vehicle-focused applications.

3 RITSA PROJECT: AN OVERVIEW

3.1 Motivations for RITSA Project

The motivations behind our research project RITSA (RAG for Intelligent Transportation Systems Architecture) are twofold.

3.1.1 Expanding Beyond the ITS end-User Perspective

The existing body of research on LLMs and augmented LLMs in Intelligent Transportation Systems (ITS) demonstrates a growing interest in leveraging these technologies to enhance the user experience and optimize vehicle performance. These studies underscore the transformative potential of LLM-driven applications in addressing end-user needs, such as improving in-vehicle communication, personalizing driving behaviors, and supporting real-time navigation. However, a critical gap remains: current research predominantly focuses on the ITS end-user perspective, overlooking other essential roles, particularly that of ITS architecture designers and modelers.

From a systems engineering standpoint, the early stages of the ITS lifecycle—encompassing operational analysis, system specification, and architectural design—are fundamental to ensuring the system's efficiency, reliability, and adaptability. Despite their importance, these phases remain largely unexplored in the context of LLMs and Retrieval-Augmented Generation (RAG) frameworks. Addressing this gap constitutes a primary motivation for the RITSA project. By integrating RAG into the ITS architecture design process, RITSA seeks to pioneer a new frontier where LLM-based tools empower designers and modelers, broadening the applications of these technologies to encompass not

only end-user and vehicle-centric solutions but also system-level design and modeling.

3.1.2 Bridging the Gap in Utilizing Reference Architectures

The concept of "reference architecture" is well-established and highly valued within the standardization community, where it is defined as “a shared and agreed reference system description used by a community of interest to achieve its business purposes. It is typically generic and instantiated as system architectures specific to individual business purposes” (ISO/IEC/IEEE 42010, 2022). Despite its recognized significance, feedback from joint research and industry projects conducted by our research institute —such as TAM (Trusted Autonomous Mobility)¹, SCA (Secure Cooperative Autonomous Systems)², and RTI (Resilience of Intelligent Transport)³ reveals that ITS reference architectures are often underutilized.

A key factor behind this underutilization is the lack of intuitive and interactive approaches for instantiating and tailoring these architectures in real-world projects. This challenge highlights the need for tools that make reference architectures more accessible and user-friendly. The RITSA project addresses this gap by leveraging RAG

methodologies, which promise to facilitate more natural and intuitive user interactions. By doing so, RITSA aims to empower ITS stakeholders to effectively exploit the full potential of reference architectures, meeting the practical needs of their users and enhancing their application in the ITS domain.

3.2 A High-Level Overview of the RAG Framework for RITSA

This section presents the high-level architecture of the RAG framework for the RITSA project, focusing on its role in assisting ITS designers during the early phases of the ITS development lifecycle. The RAG framework is designed to enable ITS designers to express their needs in terms of the desired ITS service, such as cooperative autonomous driving, traffic management, or vehicle-to-infrastructure communication. Based on this input, the framework retrieves and synthesizes relevant architecture elements to generate tailored architecture designs for the specified ITS service. The high-level overview of the RAG framework for the RITSA project is given in Figure 1.

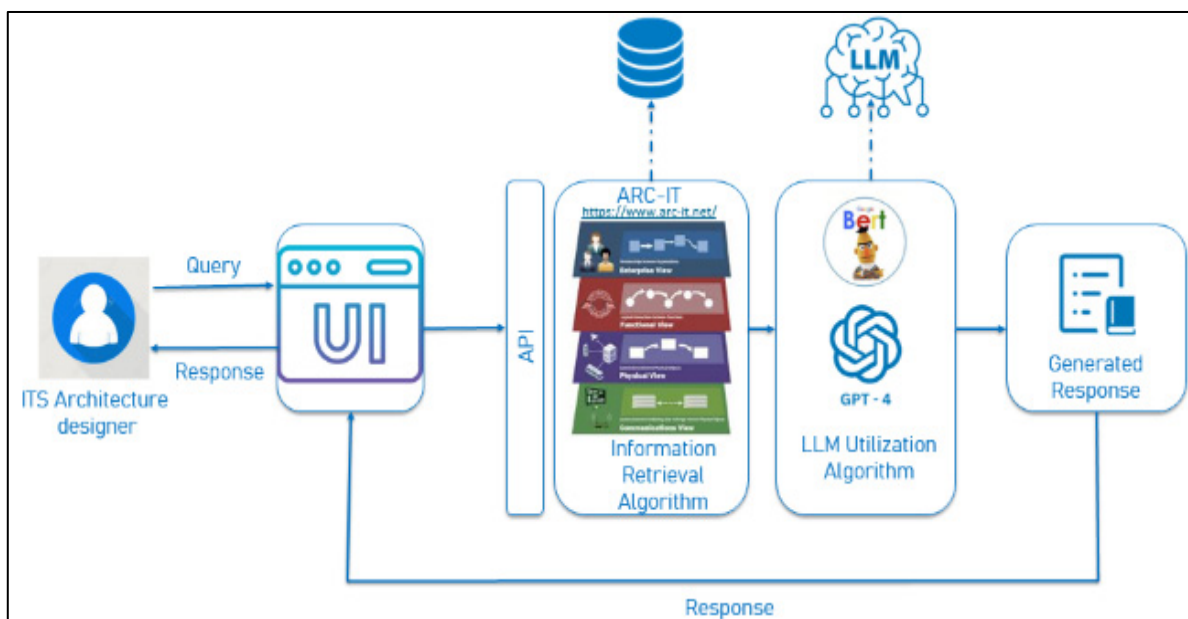


Figure 1: High-level overview of the RAG framework for the RITSA project.

¹<https://www.irt-systemx.fr/projets/tam/>

²<https://www.irt-systemx.fr/en/projets/sca/>

³<https://www.irt-systemx.fr/en/projets/rTI/>

The framework presented in Figure 1 is described as follows:

1. **ITS designer interface and prompt input:** The framework begins with an intuitive user interface (UI), where ITS designers can express their needs as natural language prompts. Specifically, the designer articulates their query in terms of the desired ITS service that requires architectural design. For instance, the designer might request the architecture for a cooperative traffic signal system or a secure vehicle-to-infrastructure communication module. In this context, it is noteworthy that designers are not required to possess deep technical expertise in ITS architecture modeling, as prompts serve as the primary mechanism for initiating the retrieval and generation of architecture-specific elements.
2. **Query processing via the API:** The designer's prompt is transferred to the backend through the application programming interface (API), which acts as a communication bridge between the UI and the processing components. The API translates the natural language prompt into formats understandable by the retrieval and generation algorithms. It ensures compatibility between the user-facing interface and the underlying algorithms.
3. **Information retrieval algorithm:** The Information retrieval algorithm is the first processing step. It identifies and extracts relevant information from reference architecture such as the ARC-IT (Architecture Reference for Cooperative and Intelligent Transportation) framework. It serves to match the designer's query to the appropriate ITS service-related content to ensure that retrieved elements are relevant to the specific ITS service described in the designer's prompt.
4. **LLM utilization algorithm:** Once the relevant information is retrieved, it is processed by the LLM Utilization Algorithm, which leverages advanced Large Language Models (LLMs) like GPT-4 and BERT. It synthesizes the retrieved information with the context of the original designer's query, and then generates a comprehensive response that outlines the architecture design elements for the specified ITS service.
5. **Generated response:** The result of the LLM processing is a Generated Response, which is tailored to the designer's query. This response is passed back through the API to the UI,

where the ITS designer can access it. The output is delivered as a generated architecture design, comprising a detailed description of the architecture elements for the specified ITS service, an explanation of how the elements fit together to support the desired service, and references to the ITS reference architecture for traceability and verification.

4 CONCLUSIONS AND FUTURE WORK

Due to the inherent complexity and evolving nature of Intelligent Transportation Systems (ITS), their architectural design and modeling remain challenging and time-consuming, particularly for modelers with limited expertise in these systems. This highlights the need for intelligent assistants to support the ITS architecture design and modeling process. In response to this need, and to address the research gap in this area, this paper introduces the RAG for Intelligent Transportation Systems Architecture (RITSA) project. RITSA aims to assist ITS designers and modelers during the architectural design process by leveraging a Retrieval-Augmented Generation (RAG) framework. By integrating the generative capabilities of large language models (LLMs) with information retrieval from proven ITS reference architectures such as ARC-IT, RITSA seeks to pioneer a new frontier in ITS design. This approach broadens the applications of LLM-based tools beyond end-user and vehicle-centric solutions to encompass system-level design and modeling. Through this innovative integration, RITSA empowers ITS stakeholders to effectively exploit the full potential of reference architectures, addressing practical needs and enhancing their applicability within the ITS domain.

This paper provided a comprehensive overview of the RITSA project, covering its background—including Intelligent Transportation Systems (ITS), reference architectures, large language models (LLMs), Retrieval-Augmented Generation (RAG), and related existing work—along with its motivations and the high-level architecture of the underlying RAG framework. The project is driven by two key motivations: (1) extending the application of LLM-based technologies beyond the ITS end-user perspective to encompass the ITS designer and modeler perspective; and (2) addressing the underutilization of ITS reference architectures by providing intuitive and interactive means for their

instantiation and tailoring. Building on the high-level RAG framework architecture presented in this paper, the immediate next steps for the RITSA project will focus on two main objectives: (1) implementing the RAG system and (2) evaluating its performance based on well-defined criteria.

ACKNOWLEDGEMENTS

The RITSA project is a collaborative research initiative between the French Technological Research Institute SystemX, the German Research Center for Artificial Intelligence (DFKI), and the Austrian Josef Ressel Centre for Dependable System-of-Systems Engineering.

REFERENCES

- ARC-IT Version 9.3 (2025). Available at <https://www.arc-it.net/>.
- Aslam, M. S., & Nisar, S. (2024). Ethical Considerations for Artificial Intelligence Tools in Academic Research and Manuscript Preparation: A Web Content Analysis. In *Digital Transformation in Higher Education, Part B: Cases, Examples and Good Practices* (pp. 155-196). Emerald Publishing Limited.
- Awadid, A., Robert, B., & Langlois, B. (2024, February). MBSE to Support Engineering of Trustworthy AI-Based Critical Systems. In *12th International Conference on Model-Based Software and Systems Engineering*.
- Barba, C. T., Mateos, M. A., Soto, P. R., Mezher, A. M., & Igartua, M. A. (2012, June). Smart city for VANETs using warning messages, traffic statistics and intelligent traffic lights. In *2012 IEEE intelligent vehicles symposium* (pp. 902-907). IEEE.
- Barke, S., James, M. B., & Polikarpova, N. (2023). Grounded copilot: How programmers interact with code-generating models. *Proceedings of the ACM on Programming Languages*, 7(OOPSLA1), 85-111.
- Chen, H., Yuan, K., Huang, Y., Guo, L., Wang, Y., & Chen, J. (2023). Feedback is all you need: from ChatGPT to autonomous driving. *Science China Information Sciences*, 66(6), 1-3.
- Chen, J., Lin, H., Han, X., & Sun, L. (2024, March). Benchmarking large language models in retrieval-augmented generation. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 38, No. 16, pp. 17754-17762).
- Combemale, B., Gray, J., & Rumpe, B. (2023). ChatGPT in software modeling. *Software and Systems Modeling*, 22(3), 777-779.
- Cui, Y., Huang, S., Zhong, J., Liu, Z., Wang, Y., Sun, C., ... & Khajepour, A. (2023). Drivellm: Charting the path toward full autonomous driving with large language models. *IEEE Transactions on Intelligent Vehicles*.
- Damaj, I., Al Khatib, S. K., Naoos, T., Lawand, W., Abdelrazzak, Z. Z., & Mouftah, H. T. (2022). Intelligent transportation systems: A survey on modern hardware devices for the era of machine learning. *Journal of King Saud University-Computer and Information Sciences*, 34(8), 5921-5942.
- Deng, Y., Yao, J., Tu, Z., Zheng, X., Zhang, M., & Zhang, T. (2023). Target: Traffic rule-based test generation for autonomous driving systems. *arXiv preprint arXiv:2305.06018*.
- Du, H., Teng, S., Chen, H., Ma, J., Wang, X., Gou, C., ... & Wang, F. Y. (2023). Chat with chatgpt on intelligent vehicles: An ieev tiv perspective. *IEEE Transactions on Intelligent Vehicles*, 8(3), 2020-2026.
- Elassy, M., Al-Hattab, M., Takturi, M., & Badawi, S. (2024). Intelligent transportation systems for sustainable smart cities. *Transportation Engineering*, 100252.
- Elassy, M., Al-Hattab, M., Takturi, M., & Badawi, S. (2024). Intelligent transportation systems for sustainable smart cities. *Transportation Engineering*, 100252.
- Friedenthal, S., Moore, A., & Steiner, R. (2014). *A practical guide to SysML: the systems modeling language*. Morgan Kaufmann.
- Fu, D., Li, X., Wen, L., Dou, M., Cai, P., Shi, B., & Qiao, Y. (2024). Drive like a human: Rethinking autonomous driving with large language models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (pp. 910-919).
- Giesecke, R., Surakka, T., & Hakonen, M. (2016, April). Conceptualizing mobility as a service. In *2016 eleventh international conference on Ecological Vehicles and Renewable Energies (EVER)* (pp. 1-11). IEEE.
- Haomin, W., Haomin, L., Xinai, Z., & Miao, Y. (2021, April). High-Frequency Data Link System Architecture Design based on SysML Modeling Language. In *Journal of Physics: Conference Series* (Vol. 1884, No. 1, p. 012014). IOP Publishing.
- INCOSE (Ed.). (2023). *INCOSE systems engineering handbook*. John Wiley & Sons.
- ISO/IEC 26552:2019: *Software and systems engineering — Tools and methods for product line architecture design*. Int. Organ. Stand. Geneva, Switzerland (2019).
- ISO/IEC/IEEE 42010:2022: *Software, Systems and Enterprise Architecture Description*. ISO. <https://www.iso.org/standard/74393.html>
- ISO/IEC/IEEE 42020:2019:2019-07.: *Software, systems and enterprise — Architecture processes*. Int. Organ. Stand. Geneva, Switzerland (2019).
- Katz, D. M., Bommarito, M. J., Gao, S., & Arredondo, P. (2024). Gpt-4 passes the bar exam. *Philosophical Transactions of the Royal Society A*, 382(2270), 20230254.
- Khalid, T., Khan, A. N., Ali, M., Adeel, A., ur Rehman Khan, A., & Shuja, J. (2019). A fog-based security framework for intelligent traffic light control system. *Multimedia Tools and Applications*, 78, 24595-24615.

- Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2022). Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35, 22199-22213.
- Kotsi, A., Lusco, T., Mitsakis, E., & Sill, S. (2020). Harmonization and interoperability of C-ITS Architectures in Europe and USA. *arXiv preprint arXiv:2010.08405*.
- Lei, L., Zhang, H., & Yang, S. X. (2023). ChatGPT in connected and autonomous vehicles: benefits and challenges. *Intell. Robot*, 3(2), 145-148.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., ... & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33, 9459-9474.
- Li, M., Liu, X., Chen, S., Yang, L., Du, Q., Han, Z., & Wang, J. (2024). MST-YOLO: Small Object Detection Model for Autonomous Driving. *Sensors*, 24(22), 7347.
- Liu, X., Zheng, Y., Du, Z., Ding, M., Qian, Y., Yang, Z., & Tang, J. (2024). GPT understands, too. *AI Open*, 5, 208-215.
- Liu, Y. (2024). Large language models for air transportation: A critical review. *Journal of the Air Transport Research Society*, 100024.
- Martineau, K., Explainable, A. I., & Generative, A. I. (2023). What is retrieval-augmented generation?. *IBM Research Blog*, 22.
- Martínez-Fernández, S., Ayala, C. P., Franch, X., Nakagawa, E. Y.: A Survey on the Benefits and Drawbacks of AUTOSAR. In: *Proceedings of the First International Workshop on Automotive Software Architecture*, pp. 19-26 (2015).
- Mondal, H., Mondal, S., & Behera, J. K. (2025). Artificial intelligence in academic writing: Insights from journal publishers' guidelines. *Perspectives in Clinical Research*, 16(1), 56-57.
- Nakagawa, E. Y., Antonino, P. O., Becker, M.: Reference architecture and product line architecture: A subtle but critical difference. In: *European Conference on Software Architecture*. Springer, Berlin, Heidelberg, pp. 207-211 (2011).
- Sha, H., Mu, Y., Jiang, Y., Chen, L., Xu, C., Luo, P., ... & Ding, M. (2023). Languagempc: Large language models as decision makers for autonomous driving. *arXiv preprint arXiv:2310.03026*.
- Singh, G. (2023). Leveraging chatgpt for real-time decision-making in autonomous systems. *Eduzone: International Peer Reviewed/Refereed Multidisciplinary Journal*, 12(2), 101-106.
- Tian, Y., Li, X., Zhang, H., Zhao, C., Li, B., Wang, X., & Wang, F. Y. (2023). VistaGPT: Generative parallel transformers for vehicles with intelligent systems for transport automation. *IEEE Transactions on Intelligent Vehicles*.
- Wandelt, S., Zheng, C., Wang, S., Liu, Y., & Sun, X. (2024). Large Language Models for Intelligent Transportation: A Review of the State of the Art and Challenges. *Applied Sciences*, 14(17), 7455.
- Wang, Y., Jiao, R., Zhan, S. S., Lang, C., Huang, C., Wang, Z., ... & Zhu, Q. (2023). Empowering autonomous driving with large language models: A safety perspective. *arXiv preprint arXiv:2312.00812*.
- Waqar, A., Alshehri, A. H., Alanazi, F., Alotaibi, S., & Almujiabah, H. R. (2023). Evaluation of challenges to the adoption of intelligent transportation system for urban smart mobility. *Research in Transportation Business & Management*, 51, 101060.
- Wen, L., Fu, D., Li, X., Cai, X., Ma, T., Cai, P., ... & Qiao, Y. (2023). Dilu: A knowledge-driven approach to autonomous driving with large language models. *arXiv preprint arXiv:2309.16292*.
- Xia, Y., Jazdi, N., & Weyrich, M. (2024, September). Enhance FMEA with Large Language Models for Assisted Risk Management in Technical Processes and Products. In *2024 IEEE 29th International Conference on Emerging Technologies and Factory Automation (ETFA)* (pp. 1-4). IEEE.
- Xu, Z., Zhang, Y., Xie, E., Zhao, Z., Guo, Y., Wong, K. Y. K., ... & Zhao, H. (2024). Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*.
- Yang, Z., Jia, X., Li, H., & Yan, J. (2023). A survey of large language models for autonomous driving. *arXiv preprint arXiv:2311.01043*.
- Zhao, S., Yang, Y., Wang, Z., He, Z., Qiu, L. K., & Qiu, L. (2024). Retrieval Augmented Generation (RAG) and Beyond: A Comprehensive Survey on How to Make your LLMs use External Data More Wisely. *arXiv preprint arXiv:2409.14924*.