Leveraging NVIDIA AI technologies in the development of REQAPP: a machine learning platform for gathering and defining requirements for smart cities applications

Andra Paula AVASILOAIE,
POLITEHNICA Bucharest, Splaiul Independentei Street, No. 313, 6th District, 060042 Bucharest, Romania
andra.avasiloaie@stud.etti.upb.ro

Augustin SEMENESCU,
POLITEHNICA Bucharest, Splaiul Independentei Street, No. 313, 6th District, 060042 Bucharest, Romania
Academy of Romanian Scientists, Ilfov 4, 5th District, 050044 Bucharest, Romania
augustin.semenescu@upb.ro

Eduard Cristian POPOVICI*,
POLITEHNICA Bucharest, Splaiul Independentei Street, No. 313, 6th District, 060042 Bucharest, Romania
eduard.popovici@upb.ro *(*Corresponding author)*

Razvan CRACIUNESCU,
POLITEHNICA Bucharest, Splaiul Independentei Street, No. 313, 6th District, 060042 Bucharest, Romania
razvan.craciunescu@upb.ro

Ionut COSMIN CHIVA,
POLITEHNICA Bucharest, Splaiul Independentei Street, No. 313, 6th District, 060042 Bucharest, Romania
ionut_cosmin.chiva@upb.ro

Abstract
This paper introduces REQAPP, an innovative machine learning (ML)-powered platform utilizing NVIDIA AI technologies, such as Retrieval-Augmented Generation (RAG) chatbots, to optimize the elicitation and definition of software requirements. The platform is designed for smart city applications, including e-government, social innovation, urban planning, and urban development. REQAPP addresses the challenge of accurately gathering and refining user needs, especially for complex public sector projects, by learning and adapting from iterative interactions. The study builds on advancements in requirement engineering and the integration of ML in software development. By leveraging NVIDIA AI frameworks and tools, it extends existing research on adaptive learning systems and interactive AI solutions to meet the unique needs of smart city stakeholders. The development and evaluation of REQAPP are demonstrated through four case studies, each focusing on a distinct application domain, such as online stores, ticket booking systems, resource-sharing platforms, and urban planning tools. Using NVIDIA AI technologies, including RAG chatbots, the platform offers a conversational interface that guides users through the requirements elicitation process while dynamically refining its models based on feedback and context. The results showcase REQAPP's ability to reduce ambiguity in requirement definitions and improve user engagement. The incorporation of RAG chatbots enhances the system's capacity to provide accurate and context-aware suggestions, accelerating the requirements gathering process and ensuring alignment with stakeholder expectations. REQAPP presents a significant advancement for academics, practitioners, and policymakers involved in smart city projects. This study contributes to the field by introducing an original AI-driven framework that combines state-of-the-art NVIDIA technologies with a machine learning-centric approach. REQAPP's adaptability and focus on real-world smart city applications make it a valuable tool for future AI-enhanced development processes.

Keywords: smart city software solutions, public sector digitalization, urban innovation systems, context-aware application development, AI-driven requirement gathering.

1. Introduction

Artificial intelligence (AI) integration into software development has become increasingly important, especially in complex domains like applications for smart cities. As urban environments seek to enhance efficiency, sustainability, and overall quality of life, there has been a notable increase in the demand for advanced software systems. However, getting people to talk about and define software requirements in such complicated settings is very hard, and this often leads to confusion and wasted time during the development process.

This paper presents REQAPP, a novel platform powered by machine learning that utilises NVIDIA AI technologies, such as Retrieval-Augmented Generation (RAG) chatbots, to improve the requirements engineering process for smart city applications [1]. In contrast to earlier methodologies that emphasise the use of AI agents for the generation and refinement of user stories during the requirements analysis phase [2], REQAPP prioritises the essential initial step of thorough information gathering. This emphasis is especially crucial in intricate fields where critical information may be missed in preliminary evaluations.

REQAPP operates as an AI-powered business analyst, engaging with clients to guarantee that all essential enquiries are resolved prior to the development of user stories. By highlighting the thoroughness and precision of essential information, REQAPP seeks to reduce ambiguities and possible misinterpretations in the subsequent stages of development, thereby contributing to more resilient and successful project results. This proactive, AI-enhanced methodology sets REQAPP apart from approaches that mainly focus on automating the generation of user stories.

While authors such as those of paper [3] recognises the capacity of AI and intelligent technologies to enhance cost-effectiveness in infrastructure development—emphasizing wider economic consequences like housing affordability and construction expenses—REQAPP distinctly concentrates on the requirements engineering process within the realm of smart city software development. By simplifying the intricate and frequently unclear process of collecting and articulating software requirements, REQAPP provides a valuable resource for improving efficiency and precision in the initial phases of smart city initiatives.

Similarly, although authors [4] advocates for integrating formally stated requirements into machine learning models during design or as part of the loss function, REQAPP takes a different approach. The application makes advantage of RAG chatbots from NVIDIA to help humans and AI work together during the requirements elicitation and specification phase. This interactive approach, which is not specifically covered in [4], enables REQAPP to dynamically adjust to user input and context prior to formal requirements being generated. REQAPP is especially well-suited for complicated projects where defining initial needs and expectations is crucial because of its emphasis on the pre-formalization stage.

The authors of [5] specifically developed REQAPP for practical implementation in real-world smart city projects, in contrast to their examination of the student perspective on the use of generative AI in requirements engineering education. The authors of [5] investigate

the potential applications of AI tools such as ChatGPT in educational settings. Meanwhile, REQAPP helps people who work on building smart cities get clear information about what software they need. For example, REQAPP uses NVIDIA's AI technologies to create an adaptive and interactive system that better understands what users want. This addresses issues not addressed in the educational framework of [5].

Furthermore, while the authors of [6] examine the impact of AI on customer data collection in digital marketing, REQAPP solely focuses on leveraging AI to determine the necessary features for smart city software. Through the utilisation of NVIDIA's AI technologies to promote a conversational methodology, REQAPP tackles the distinct challenges associated with articulating intricate software requirements within the public sector, highlighting the importance of collaboration and clarity from the initial stages of project development.

This study builds on earlier work on adaptive learning systems and interactive AI solutions by using NVIDIA AI frameworks and tools. It does this by taking advantage of progress in requirements engineering and the use of machine learning in software development. In the context of smart cities, REQAPP offers a novel approach that bridges theoretical frameworks with real-world applications.

The rest of the paper follows this structure. The methodology used in the creation and evaluation of REQAPP is described in Section 2, with particular attention given to the integration of NVIDIA AI technologies, such as Retrieval-Augmented Generation (RAG) chatbots. Section 3 presents the findings of four case studies, each focusing on a different application domain: online retailers, ticketing platforms, resource-sharing websites, and urban planning tools. These studies show that REQAPP makes the process of gathering requirements better by clearing up any confusion and making sure that everyone's expectations are the same across a number of smart city scenarios. Section 4's analysis of the results examines the platform's contributions to requirements engineering techniques and the challenges encountered during implementation. It addresses REQAPP's present shortcomings and areas in need of improvement, as well as how well it promotes teamwork and enhances user involvement. Section 5 concludes the work by offering suggestions for future research topics and a summary of the main findings. These include extending REQAPP's applicability to other areas, improving its machine learning-driven methodology to support smart city developments, and maybe integrating advanced AI.

2. Methodology

This section outlines the fundamental methodologies and technologies utilised in the development and assessment of REQAPP, emphasising its significance in tackling the distinct challenges associated with requirements engineering in the realm of smart city applications. Through the utilisation of NVIDIA AI technologies, specifically Retrieval-Augmented Generation (RAG) chatbots, REQAPP improves the precision, flexibility, and effectiveness of the processes involved in gathering and defining requirements. The methodology incorporates a conceptual framework for AI-driven requirements gathering, a comprehensive system architecture, and iterative machine learning techniques designed for contextual interaction. The assessment is carried out via qualitative feedback and

quantitative performance metrics, thereby ensuring the platform's alignment with the varied needs of stakeholders.

2.1. Challenges in requirements engineering for smart cities

The process of requirements engineering for smart cities necessitates the resolution of distinct complexities that arise from the dynamic nature of urban systems and the varied needs of stakeholders. Navigating these intricacies demands innovative approaches and tools. Below, we explore key obstacles that underscore the necessity for advanced solutions like REQAPP.

Managing stakeholders' varied and perhaps competing priorities is one of the main challenges. Smart city initiatives frequently bring together individuals, business organisations, and public sector bodies, each with different technological expertise and goals. Authors in [7] highlight how this diversity might lead to misalignments that slow down progress. The inherent ambiguity in early-stage needs presents another major obstacle. Iterative cycles of clarification caused by stakeholders' recurrent difficulties to express specific needs delay development and raise expenses. These changing needs are frequently not well satisfied by traditional approaches, especially in dynamic metropolitan settings [2].

Integrating cutting-edge technologies like AI and IoT further complicates the requirements engineering process. It is essential to adhere to established standards throughout the lifecycle of AI-driven systems, yet existing methods usually lack the clear protocols needed to manage reliability and safety in these high-stakes situations [4].

Scalability and adaptability are also critical considerations. As cities grow and change, the systems supporting them must evolve in tandem. Generative AI offers potential solutions by dynamically adjusting requirements to shifting contexts. However, effectively utilising such technologies requires platforms capable of synthesising diverse inputs and prioritising needs based on real-time data [5].

Finally, fostering effective collaboration among stakeholders remains a core issue. Transparent and inclusive interactions are essential to the success of smart city initiatives. Tools such as NVIDIA's RAG chatbots are particularly well-suited to facilitating these interactions, offering context-aware recommendations that bridge gaps between technical and non-technical participants [6].

2.2. Overview of REQAPP and NVIDIA AI technologies

The advanced machine learning platform REQAPP enhances the requirements engineering process for smart city applications. By integrating NVIDIA's cutting-edge AI technologies, REQAPP offers a dynamic and adaptive approach to gathering and defining software requirements, tailored to the complex needs of urban development projects.

REQAPP's functionality fundamentally relies on the deployment of NVIDIA's Retrieval-Augmented Generation (RAG) chatbots. These chatbots enable interactive and context-aware conversations with stakeholders by combining retrieval-based and generative

models. This integration makes sure that answers are correct and useful, which greatly enhances the quality of the data gathered during requirements elicitation.

The platform is enhanced by utilising the NVIDIA AI Enterprise suite, which offers a complete array of AI tools and frameworks optimised for enterprise applications. This suite supports the efficient operation of REQAPP, allowing it to scale seamlessly to meet the demands of diverse and intricate smart city projects.

The architecture diagram available in [8] provides a detailed representation of how NVIDIA's AI technologies integrate into solutions like REQAPP, demonstrating how to configure RAG chatbots and other AI tools to maximise adaptability and contextual relevance in real-world applications.

RAG consists of two processes: 1. Ingestion of documents from document repositories, databases, or APIs that are all outside of the foundational model's knowledge. 2. Retrieval of relevant document data and generation of responses during inference.

The following graphic describes these processes, the ingestion of documents, and the generation of responses further.

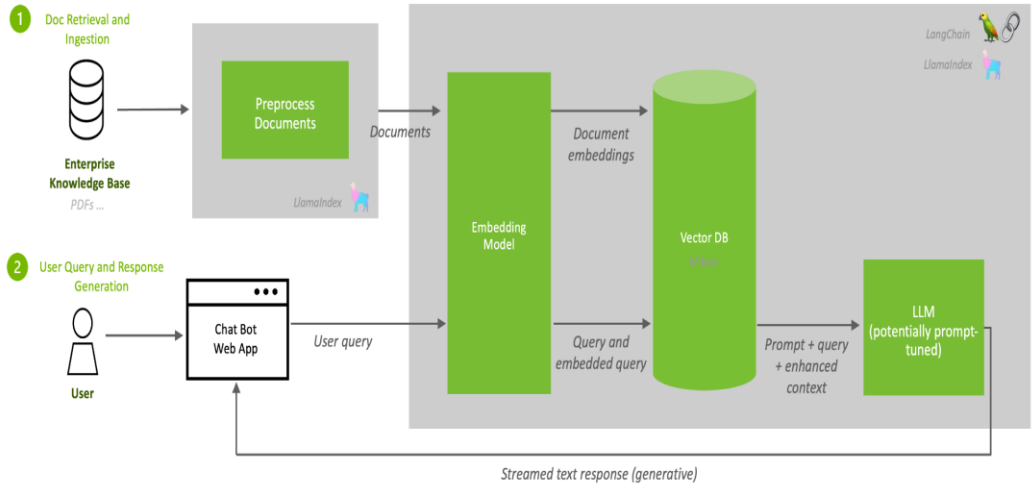

Fig. 1. NVIDIA RAG Architecture.
Source: NVIDIA Generative AI Chatbot Technical Brief [8].

2.3 Conceptual framework for REQAPP

Traditional requirements engineering methods often struggle with ambiguities and inefficiencies, particularly in dynamic and complex domains such as smart cities. REQAPP introduces a transformative approach by embedding AI capabilities at the core of its requirements gathering process.

The platform employs machine learning models that iteratively adapt based on stakeholder inputs, refining the understanding of user needs over time. This iterative approach ensures

a continuous improvement cycle, reducing the risk of misinterpretation and enhancing the accuracy of requirements. By analysing patterns in stakeholder interactions and previously gathered data, REQAPP identifies potential gaps in requirements, prompting users to provide additional details.

Furthermore, the AI-driven approach prioritises the contextual relevance of gathered requirements. Instead of relying solely on predefined templates or static processes, REQAPP dynamically adjusts its strategies based on the specific project domain, stakeholder preferences, and real-time feedback. This adaptability is particularly valuable in the smart city context, where requirements often span diverse domains such as e-government, social innovation, and urban planning.

A cornerstone of REQAPP's framework is the integration of NVIDIA's Retrieval-Augmented Generation (RAG) chatbots. These chatbots facilitate contextual and interactive communication between stakeholders and the platform, enabling a richer and more effective requirements elicitation process.

RAG chatbots use both retrieval-based techniques, which use huge knowledge bases, and generative models, which can come up with complex responses that are based on the situation. This hybrid approach ensures that the platform not only retrieves the most relevant information but also provides explanations and suggestions that align with the specific needs of the project.

For instance, in an urban planning scenario, the chatbot might retrieve zoning laws or infrastructure requirements from its knowledge base while simultaneously generating tailored questions to clarify stakeholder preferences for a proposed development. This dual capability boosts stakeholder engagement and guarantees the early capture of critical details in the requirements engineering process.

Additionally, the contextual intelligence of RAG chatbots allows them to adapt dynamically to ongoing conversations. The chatbot adjusts its line of questioning to thoroughly explore all aspects of a requirement by analysing the dialogue flow and identifying recurring themes or ambiguities. This iterative refinement process supports a collaborative and inclusive approach to requirements gathering, bridging gaps between technical and non-technical stakeholders.

With these cutting-edge AI tools built into REQAPP's basic idea, the platform is now a top-of-the-line way to handle the tricky parts of smart city requirements engineering. REQAPP is a strong and scalable way to define software requirements because it combines adaptive learning with smart, context-aware interaction.

2.4 System architecture

The REQAPP frontend is designed to provide an intuitive and user-friendly experience for stakeholders. This component includes the User Interface (UI), which is responsive and accessible, enabling seamless interaction across multiple devices and platforms. The User Interaction Analytics module tracks and analyses user behaviours, providing insights to

optimise and refine the interface based on real-time feedback. The Stakeholder Feedback Collection system also makes sure that all interactions record important user inputs and preferences. This encourages a collaborative and open process for gathering requirements. Together, these elements create an interface that empowers stakeholders to communicate their needs effectively, reducing misunderstandings early in the project lifecycle.

The REQAPP backend, which supports its core functionalities and ensures the seamless integration of various components, is the core of its processing capabilities, ensuring the platform can handle complex and diverse workloads. The Data Processing Pipeline manages the transformation of raw stakeholder inputs into actionable insights, preparing them for analysis and refinement. The Machine Learning Models are designed to adapt and improve iteratively, learning from interactions to provide more precise recommendations and suggestions over time. The NVIDIA RAG Chatbots component effectively bridges the gap between user interactions and technical data processing by facilitating intelligent and context-aware conversations. Meanwhile, the Database Management system securely stores all data related to requirements, ensuring its accessibility and integrity. Finally, Scalability Services guarantee that the platform can handle increased loads as the complexity and volume of smart city projects grow.

The Data Processing Pipeline is a crucial component within the back-end architecture, responsible for converting raw data into structured, actionable formats. This module ingests data from multiple sources, including stakeholder inputs, chatbot interactions, and external databases. Through a series of transformation processes, such as cleaning, normalisation, and enrichment, the pipeline prepares data for machine learning models to analyze. Its efficiency ensures that the platform delivers timely and accurate insights, even as the volume and diversity of inputs scale up.

The Machine Learning models embedded within REQAPP are designed for adaptive learning and decision-making. These models analyse historical and real-time data, extracting patterns and refining the requirements engineering process iteratively. By continuously learning from stakeholder feedback and interaction data, the models enhance their accuracy and relevance. This ability to adapt makes sure that the platform stays in line with changing project needs. This makes it a key part of dynamic and responsive requirements engineering.

At the core of REQAPP's interactive capabilities are the NVIDIA RAG Chatbots, which combine retrieval-based and generative AI approaches. These chatbots leverage an extensive knowledge base to retrieve relevant information while dynamically generating tailored responses based on user queries and context. This hybrid design allows the chatbots to engage in meaningful, context-aware conversations, ensuring stakeholders provide comprehensive and accurate input. Their adaptability during interactions effectively bridges the gap between technical data processing and user requirements, thereby streamlining the elicitation process.

The Scalability Services in REQAPP ensure the platform's ability to handle the growing complexity and volume of data associated with smart city projects. By dynamically

allocating computing resources, these services keep up smooth performance even as urban systems grow and stakeholder interactions rise. Built on cloud-based infrastructure, the scalability services adapt to changing workloads. This makes sure that the platform can handle large datasets, multiple users at the same time, and a lot of machine learning operations without slowing down. This feature is important for meeting the changing and wide-ranging needs of smart city applications.

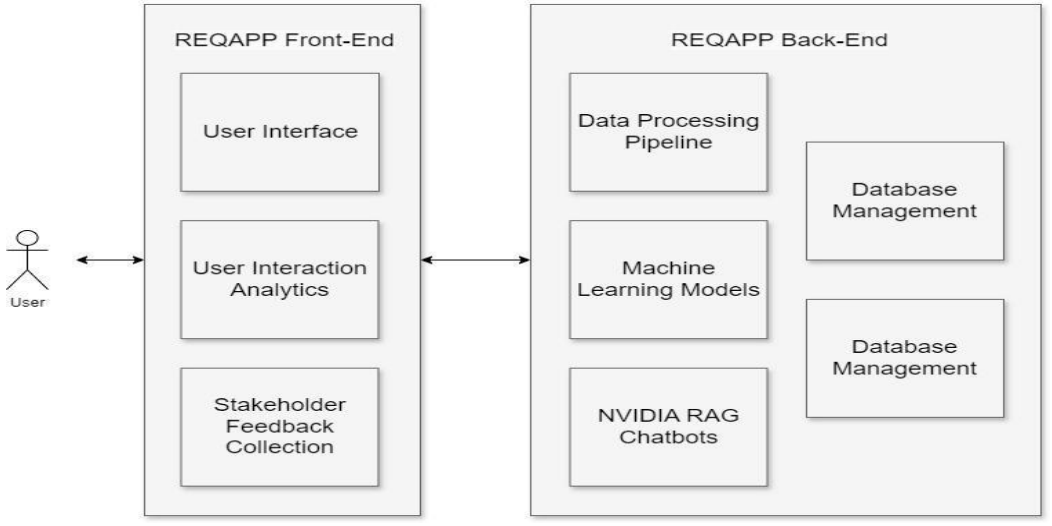

Fig. 2. REQAPP System Architecture.
Source: Author's representation of the architecture

2.5 Retrieval-Augmented Generation (RAG) in requirements engineering

Refer to the architectural diagram in [9] to see how NVIDIA's Retrieval-Augmented Generation (RAG) technology integrates into REQAPP.

This diagram provides a comprehensive view of how RAG pipelines function, detailing the flow from document ingestion to query processing and response generation. It shows how different parts, like the document ingestion pipeline, embedding models, and the LLM inference pipeline, work together to give correct information that is relevant to the situation.

In the context of REQAPP, this architecture can be adapted to facilitate effective requirements engineering by enabling the system to retrieve pertinent information from extensive datasets and generate precise, context-aware responses during stakeholder Platformteractions. This approach ensures that the requirements elicitation process is both comprehensive and aligned with the specific needs of smart city applications.

This example demonstrates the orchestration of NVIDIA microservices within a scalable framework, utilising containerised environments to enhance workflows. Key parts of the architecture, like data ingestion, model processing, and response generation, work together smoothly to support high-performance applications. This method is a strong example for putting Retrieval-Augmented Generation (RAG) pipelines into action, showing how to use advanced AI-based solutions in real life.

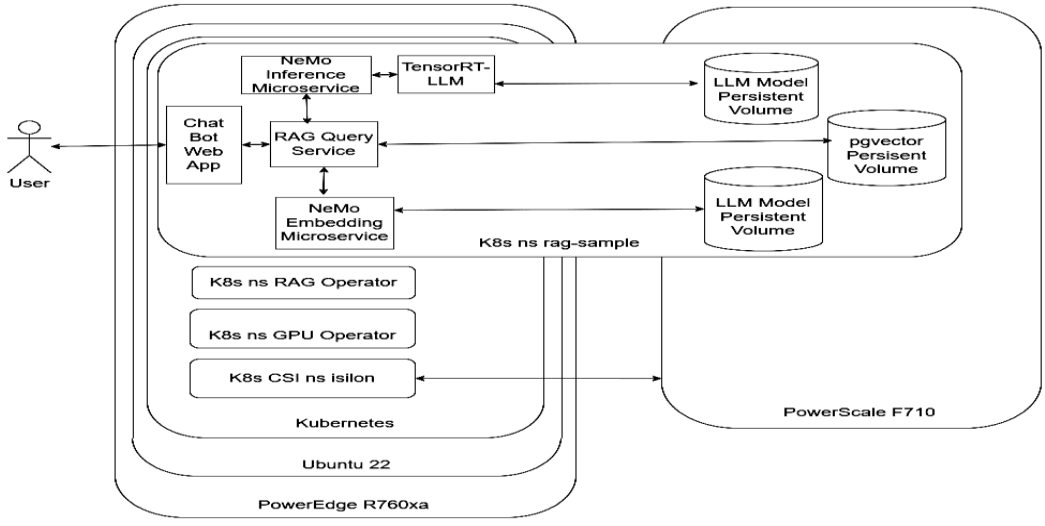

Fig. 3. Advanced AI Deployments with NVIDIA RAG and Dell Technologies Diagram
Source: Dell Scalable Architecture for Retrieval-Augmented Generation (RAG) with NVIDIA Microservices

3. Case studies

To evaluate the effectiveness and adaptability of REQAPP in diverse smart city contexts, we conducted four case studies, each focusing on a distinct application domain. These studies show how the platform uses NVIDIA AI technologies to solve the problems of requirements engineering, especially in situations that change quickly and are driven by stakeholders. By examining the Online Store Application, Ticket Booking System, Resource-Sharing Platform, and Urban Planning Tool, we explore the platform's ability to enhance requirements elicitation, improve engagement, and align stakeholder expectations across varying domains.

3.1 Online store application

The online store application represents a critical domain where stakeholder requirements are multifaceted, involving end-users, vendors, and administrative teams. REQAPP used its Retrieval-Augmented Generation (RAG) chatbots to have dynamic, context-aware conversations with these stakeholders during the initial elicitation process. The platform identified essential requirements such as product categorisation, payment gateway integration, and delivery logistics by querying its knowledge base and tailoring questions to stakeholder inputs.

Even as stakeholders struggled to articulate specific functionalities, this iterative interaction ensured that critical needs remained unnoticed. For example, when vendors expressed general concerns about inventory management, the system generated targeted queries about preferred tracking methods and integration with existing software. This proactive approach reduced ambiguity and laid a solid foundation for subsequent development stages.

After gathering the initial requirements, REQAPP used machine learning models to analyse user interaction patterns and improve the platform's design. The system synthesized insights from vendor feedback and end-user preferences to prioritize features that would

enhance engagement, such as personalized product recommendations and streamlined checkout processes.

The RAG chatbots played a pivotal role in this refinement phase by continuously interacting with stakeholders to validate assumptions and incorporate real-time feedback. For instance, the chatbot highlighted discrepancies in how users navigated product categories, prompting adjustments to the UI for better accessibility. Furthermore, the platform's adaptability allowed it to respond to evolving requirements, such as the integration of regional payment methods, ensuring the application remained relevant and user-centric.

When REQAPP dealt with these problems using AI-driven insights and flexible feedback systems, it showed that it could help build strong and scalable online store applications as part of smart city development as a whole.

3.2 Ticket booking system

The ticket booking system presents a domain where requirements evolve rapidly due to fluctuating user demands and external factors, such as event scheduling changes or transportation disruptions. During the requirements elicitation process, REQAPP utilised its Retrieval-Augmented Generation (RAG) chatbots to engage stakeholders, including event organisers, transportation providers, and end users. The system dynamically adjusted its queries based on stakeholder inputs, enabling the identification of key requirements like seat allocation preferences, dynamic pricing mechanisms, and real-time availability updates.

For example, when stakeholders were worried about how to handle busy booking times, REQAPP came up with specific ideas for putting in place load-balancing systems and using predictive analytics to plan for events with a lot of demand. The platform effectively addressed even nuanced stakeholder needs by iteratively refining its understanding of these requirements.

In addition to gathering requirements, REQAPP enhanced the decision-making process by providing real-time insights derived from its machine learning models. The platform analysed historical booking patterns and external data, such as weather forecasts and event popularity trends, to offer actionable recommendations for optimising the system’s functionality.

For example, REQAPP suggested dynamic adjustments to seat availability based on real-time booking trends, ensuring maximum occupancy during events. Additionally, the RAG chatbots facilitated continuous interaction with stakeholders, validating these recommendations and incorporating new insights as they emerged.

By leveraging NVIDIA’s AI technologies, the platform also supported the integration of advanced features like personalised ticket suggestions and adaptive pricing models, ensuring a user-centric experience. These enhancements not only improved the efficiency

of the ticket booking system but also demonstrated REQAPP's ability to adapt to dynamic and time-sensitive contexts, making it a valuable tool for smart city applications.

3.3 Resource-sharing platform

The resource-sharing platform targets scenarios such as shared transportation, workspace allocation, or equipment leasing, where collaboration between stakeholders and efficient resource distribution are critical. Using its Retrieval-Augmented Generation (RAG) chatbots, REQAPP made it easier to find needs like scheduling conflicts, fair resource distribution, and connecting to outside platforms.

For example, during initial interactions, stakeholders highlighted challenges in managing overlapping requests for high-demand resources. REQAPP dynamically tailored its queries to extract specific preferences, such as priority rules for resource allocation or the ability to set blackout periods. The platform's machine learning models analysed historical usage patterns and projected demand to propose optimised schedules and allocation strategies.

By aligning its recommendations with stakeholder inputs, REQAPP ensured that the resource-sharing system minimised conflicts and maximised utilisation, fostering collaboration among users with potentially competing interests.

A key strength of REQAPP in this domain is its ability to adapt to real-time feedback and evolving needs. The platform employed machine learning models to continuously refine its understanding of stakeholder priorities and usage behaviors. Through iterative interactions, the RAG chatbots validated initial assumptions and identified new requirements as they emerged.

For instance, when stakeholders reported delays in booking confirmations, REQAPP suggested enhancements to streamline the approval workflow and provide real-time status updates. The platform also adapted its algorithms to accommodate unexpected changes, such as surges in demand due to special events or resource outages.

These adaptive feedback mechanisms ensured that the resource-sharing platform remained flexible and responsive, accommodating the diverse and dynamic needs of stakeholders. By integrating NVIDIA's AI technologies, REQAPP enabled a collaborative and efficient solution that enhanced resource allocation while fostering trust and satisfaction among users.

3.4 Urban planning tool

Urban planning projects often involve intricate requirements due to the diverse priorities of public institutions, private developers, and citizen groups. The Urban Planning Tool powered by REQAPP tries to deal with these problems by using NVIDIA's AI technologies, especially its Retrieval-Augmented Generation (RAG) chatbots, to make it easier to find out what people need.

Public sector needs frequently include compliance with zoning laws, environmental regulations, and infrastructure integration. During the elicitation phase, REQAPP

dynamically generated targeted questions to uncover specific constraints and objectives. For instance, it highlighted regulatory overlaps when municipalities required alignment between local and regional development plans. Furthermore, the platform's machine learning models processed historical project data and ongoing stakeholder feedback to identify potential bottlenecks, such as budgetary constraints or logistical challenges.

In analyzing urban planning strategies for smart cities, authors in [10] highlights the importance of integrating advanced technologies, sustainable practices, and citizen participation. Using thematic maps generated with QGIS, the study evaluates the policies and characteristics that differentiate the most successful smart cities globally.

These insights underscore the critical role of data-driven decision-making and participatory approaches in developing urban planning tools that address complex public sector needs and enhance stakeholder collaboration. In the context of urban planning, artificial intelligence has become a cornerstone for enhancing decision-making and improving the efficiency of public organizations.

Authors in [11] emphasizes that AI technologies enable public sector entities to process large volumes of data, predict future trends, and engage stakeholders more effectively. These capabilities are particularly relevant for smart city initiatives, where the integration of AI-driven tools like REQAPP can facilitate the identification of complex urban challenges and provide adaptive solutions tailored to diverse stakeholder needs.

The urban planning tool's adaptability to the complex requirements of public sector projects ensured the identification and resolution of all critical aspects prior to the design phase. Stakeholder collaboration is a cornerstone of successful urban planning. REQAPP enhanced this process by providing a centralised platform where stakeholders, including urban planners, policymakers, and community representatives, could interact in real-time. Through its RAG chatbots, the platform facilitated transparent and inclusive conversations, enabling stakeholders to express their concerns and preferences clearly.

For example, when citizens raised questions about proposed green spaces or transportation hubs, REQAPP retrieved relevant urban planning guidelines and generated context-aware suggestions to address these concerns. It also allowed policymakers to prioritise citizen needs by analysing recurring themes in feedback and presenting actionable insights.

In the realm of urban planning, the integration of AI-driven tools has profound implications for enhancing local governance and citizen participation. Authors in [12] introduce the concept of E-Government 3.0, an AI-enabled model designed to revolutionize public service delivery and strengthen local democracies.

By leveraging AI to analyze citizen petitions and prioritize community needs, this model facilitates a more inclusive and responsive urban planning process. Such approaches align with the objectives of tools like REQAPP, which aim to foster collaboration among stakeholders while addressing complex public sector challenges in smart city development.

Real-time monitoring and predictive analysis of urban systems, as highlighted by [13], enable cities to anticipate disruptions, optimize resource allocation, and implement proactive solutions. By integrating IoT sensor data with advanced analytics, these systems enhance operational efficiency and safety, proving essential for effective infrastructure management and rapid emergency response in modern urban environments.

Digitalization plays a crucial role in enhancing urban planning processes and fostering innovation in smart cities. Authors in [14] highlight Romania's progress in public administration digitalization, comparing it with European trends. The study emphasizes the significance of e-government solutions in improving service delivery and citizen engagement. These findings underscore the potential of AI-driven tools like REQAPP to align with digital governance initiatives, enabling more transparent and efficient collaboration among stakeholders while addressing complex urban challenges.

The platform's iterative learning mechanisms further strengthened collaboration by dynamically refining its models based on ongoing interactions. This adaptability ensured that the Urban Planning Tool could align diverse stakeholder perspectives, balancing technical, social, and environmental considerations effectively. By integrating NVIDIA's AI technologies and focusing on collaboration, REQAPP empowered stakeholders to co-create urban planning solutions that were both innovative and responsive to the unique needs of smart city development.

4. Results and discussion

The results from the case studies demonstrate the adaptability and effectiveness of REQAPP across diverse smart city applications. By leveraging NVIDIA AI technologies, including Retrieval-Augmented Generation (RAG) chatbots, the platform showcased its ability to streamline requirements engineering processes, reduce ambiguities, and improve collaboration among stakeholders.

These outcomes highlight REQAPP's potential to address the unique challenges of smart city projects while delivering scalable and contextually aware solutions. Furthermore, the platform's dynamic feedback mechanisms ensured continuous alignment with evolving project demands, proving its utility as a comprehensive tool for requirements gathering and refinement.

4.1 Comparison across case studies

The case studies revealed distinct patterns and challenges across different application domains, illustrating the versatility of REQAPP. For the online store application, the platform excelled in capturing granular user preferences and adapting to domain-specific requirements, such as payment integration and inventory management. In contrast, the ticket booking system highlighted REQAPP's strength in managing dynamic requirements, such as real-time availability updates and demand prediction.

The resource-sharing platform emphasised collaboration and equitable resource allocation, demonstrating the platform's ability to integrate multiple stakeholder priorities effectively. Finally, the urban planning tool showcased REQAPP's capacity to address complex

regulatory requirements and foster inclusive stakeholder engagement, critical for public sector projects. These findings underscore the platform's adaptability to diverse smart city scenarios, providing robust solutions that meet technical and social requirements across various domains.

One of REQAPP's most significant contributions was its ability to adapt to context-specific challenges through its AI-driven mechanisms. The RAG chatbots played a pivotal role in tailoring interactions, retrieving domain-relevant data, and generating responses that aligned with stakeholder expectations. This adaptability was particularly evident in domains such as urban planning, where regulatory complexity and diverse stakeholder needs required a nuanced approach.

The iterative learning framework enabled REQAPP to refine its understanding of requirements over time, responding dynamically to feedback and evolving project contexts. For example, in the Ticket Booking System, the platform adapted to high-demand scenarios by integrating real-time analytics to adjust availability and pricing. Similarly, in the resource-sharing platform, REQAPP prioritised equitable resource distribution by continuously analyzing usage patterns and stakeholder inputs.

These insights demonstrate that context-aware AI not only enhances the accuracy of requirements gathering but also fosters stakeholder trust and collaboration by aligning outputs with their unique needs and expectations. This capability positions REQAPP as a transformative tool for advancing smart city projects through intelligent, adaptive, and scalable solutions.

4.2 Challenges and limitations

While REQAPP demonstrated significant strengths across the case studies, several algorithmic and operational challenges were observed. A primary constraint lies in the reliance on high-quality training data for its machine learning models. The quality and relevance of the data used to train its Retrieval-Augmented Generation (RAG) chatbots directly impact the platform's performance. In scenarios where domain-specific data was sparse or inconsistent, the chatbot struggled to provide contextually accurate responses, highlighting the need for continuous curation and enrichment of the knowledge base.

Operationally, the platform's computational requirements can be demanding, particularly in projects with large-scale datasets or high-frequency stakeholder interactions. Although NVIDIA AI technologies provide robust infrastructure, ensuring scalability in resource-constrained environments remains a challenge. Additionally, time-sensitive projects or environments with limited stakeholder engagement may not always be feasible for the iterative learning approach, which requires consistent feedback loops.

Another critical limitation is related to stakeholder adoption of AI-driven tools like REQAPP. While the platform is designed to simplify requirements engineering, some stakeholders, particularly those with limited technical expertise, found it challenging to engage effectively with the RAG chatbots. Misinterpretations of chatbot outputs or reluctance to trust AI-driven recommendations created barriers to effective collaboration.

Moreover, resistance to change was evident among stakeholders accustomed to traditional methods of requirements gathering. This resistance was particularly pronounced in public sector projects, where procedural inertia and regulatory constraints often hinder the adoption of innovative tools. To address these barriers, extensive onboarding, training, and demonstration of REQAPP's capabilities are necessary to build stakeholder confidence and encourage active participation.

By identifying these challenges, future iterations of REQAPP can focus on improving algorithmic robustness, optimising operational efficiency, and implementing user-centric strategies to enhance stakeholder engagement and trust in AI-driven requirements engineering.

5. Conclusions and future directions

This study has demonstrated the potential of REQAPP as an innovative platform for enhancing requirements engineering in smart city applications. By integrating NVIDIA AI technologies, including Retrieval-Augmented Generation (RAG) chatbots, the platform addressed key challenges such as stakeholder collaboration, dynamic requirements, and contextual relevance.

Through its case studies, REQAPP showcased its adaptability across diverse domains, including online retail, ticket booking, resource sharing, and urban planning, delivering accurate, scalable, and stakeholder-centric solutions. The iterative learning mechanism and AI-driven insights further solidified REQAPP's position as a transformative tool for smart city development.

While REQAPP has proven its effectiveness, several opportunities for improvement exist. Enhancing the platform's ability to handle low-quality or sparse datasets could increase its applicability in underresourced domains. Moreover, optimising the computational requirements for deployment in resource-constrained environments would make REQAPP accessible to a broader range of stakeholders, including smaller municipalities or organisations with limited infrastructure. In future versions of REQAPP, neuro-symbolic AI techniques could be added to improve its ability to reason and understand.

By combining symbolic reasoning with deep learning, the platform could better address complex, high-stakes requirements that involve logical constraints and formal verification processes. This integration would be particularly beneficial for domains such as regulatory compliance or infrastructure safety, where precise adherence to formal requirements is critical.

Expanding REQAPP's application scope to emerging domains in smart cities presents a significant opportunity for future research and development. Areas such as climate resilience planning, energy optimisation, and public health infrastructure could greatly benefit from the platform's ability to manage complex, interconnected requirements. Additionally, adapting REQAPP to support cross-domain projects that involve multiple stakeholders and overlapping objectives could further demonstrate its versatility and scalability.

References

[1] A. Ullah, G. Qi, S. Hussain, I. Ullah and Z. Ali, "The Role of LLMs in Sustainable Smart Cities: Applications, Challenges, and Future Directions," 2024.

[2] M. Sami, M. Waseem, Z. Zhang, Z. Rasheed, K. Systä and P. Abrahamsson, "AI based Multiagent Approach for Requirements Elicitation and Analysis," 2024.

[3] P. Akinola, "Leveraging Cost-Effective AI and Smart Technologies for Rapid Infrastructural Development in USA," 2024.

[4] E. Giunchiglia, F. Imrie, M. van der Schaar and T. Lukasiewicz, "Machine learning with requirements: A manifesto," 2024.

[5] N. Mellqvist and P. Mozelius, "Implementing Generative AI in Requirements Engineering Education: The Student Perspective," 2024.

[6] J. Eniola, O. Enoch, B. William, G. Olaoye and M. Idowu, "Elevating Digital Marketing Strategies: The Impact of Advanced AI on Customer Insight Gathering," 2024.

[7] M. A. Nadeem, S. Lee and M. Younus , "A Comparison of Recent Requirements Gathering and Management Tools in Requirements Engineering for IoT-Enabled Sustainable Cities,," in *Smart Grid Analytics for Sustainability and Urbanization in Big Data*, 2024.

[8] NVIDIA, " “Generative AI Chatbot Using RAG Technical Brief”," NVIDIA Generative AI Workflow, [Online]. Available: [https://resources.nvidia.com/en-us-generative-ai-chatbot-workflow/knowledge-base-chatbot-technical-brief]. [Accessed November 2024].

[9] Dell Technologies Info Hub, "Dell Scalable Architecture for Retrieval-Augmented Generation (RAG) with NVIDIA Microservices," [Online]. Available: https://infohub.delltechnologies.com/en-us/t/dell-scalable-architecture-for-retrieval-augmented-generation-rag-with-nvidia-microservices-1. [Accessed November 2024].

[10] S. Carboni, "Smart Cities in Comparison: An Analysis of the Best Smart Cities"," *Smart Cities and Regional Development (SCRD) Journal,,* vol. 8, no. 3, pp. 65-78, 2024.

[11] C. Vrabie, ""Artificial Intelligence Promises to Public Organizations and Smart Cities.," in *Digital Transformation. Lecture Notes in Business Information Processing*, 2022.

[12] C. Vrabie, " "E-Government 3.0: An AI Model to Use for Enhanced Local Democracies,"," *" Sustainability, ,* 2023.

[13] G. Suciu and C. Stalidi, "Digital Twins, the Software Solution for Safer Cities," *Scientific Bulletin of Communication and Networking Systems,* vol. 1, no. 1, 2023.

[14] R. Damaschin and M. G. Mihaila, "Digitalizarea administratiei publice din Romania in raport cu tendintele europene," in *Smart Cities International Conference (SCIC) Proceedings*, 2023.