Leveraging User Perception for 6G Edge-Cloud Orchestration of Networked Extended Reality

José Santos, Sam Van Damme, Javad Sameri, Susanna Schwarzmann, Qing Wei, Riccardo Trivisonno, Filip De Turck, and Maria Torres Vega

Abstract

As extended reality (XR) technologies evolve, optimizing quality of experience (QoE) for end-users becomes increasingly challenging due to the stringent network requirements of these applications. Traditional quality of service (QoS) metrics often fall short, making user perception central to effective network management. This challenge becomes even more pronounced with the advent of 6G networks, where XR applications will be deployed on distributed computing resources across the compute continuum (CC). This article presents our vision for edge-cloud orchestration in 6G, featuring a hybrid QoE model that integrates network and physiological metrics with subjective user feedback. In addition, we leverage reinforcement learning (RL) to efficiently manage QoE across edge-cloud infrastructures within the CC. By addressing the dynamic and complex nature of XR applications from both technical and human perspectives, our approach aims to meet the stringent low-latency requirements of these immersive applications. The article also highlights recent advancements in QoE modeling through a collaborative virtual reality (VR) use case, demonstrating the importance of integrating objective and subjective metrics for effective QoE management. This article aims to provide a foundational framework for QoE management in 6G networks, paving the way for future research and innovation in this rapidly evolving domain.

Introduction

Recent advances in extended reality (XR) systems have generated significant interest across various verticals, including Education, Entertainment, Manufacturing, and Healthcare [1]. These sectors are leveraging immersive XR experiences to enhance productivity through virtual training systems and remote collaboration among employees. In the consumer market, immersive entertainment experiences, such as gaming and live sporting events, are also gaining substantial traction. Financial institutions such as Goldman Sachs project immense growth potential for the XR ecosystem, predicting it will become an eighty-billion-dollar market by 2025 [2]. With expectations of widespread adoption in the coming years, the rapid advancement of XR technologies is transforming the landscape of immersive experiences.

Nonetheless, quality of experience (QoE) assessment and optimization for end-users remain a significant challenge [3, 4]. The complexity arises from the subjective nature of QoE, which is influenced by a variety of factors, including user expectations, network conditions, and the specific context in which services are consumed. This is even further exacerbated in emerging 6G compute continuum (CC) scenarios [5, 6], such as cloud XR gaming and holographic telepresence [7], where the dynamic and heterogeneous nature of the network introduces additional variability in user experience. 6G CC environments will involve diverse network conditions and a wide range of computing resources, from edge devices to cloud data centers (DCs). The variability in computing power and geographical location adds complexity, requiring advanced orchestration methods to ensure consistent performance and low latency for real-time XR applications. Existing methods often struggle to capture and optimize these dynamic user experiences, emphasizing the need for more sophisticated models that better align technical metrics with actual user perceptions and satisfaction, especially in 6G CC contexts.

This article presents our vision for QoE management in 6G networks, integrating subjective user feedback with network and physiological metrics to create a more holistic approach to edge-cloud orchestration. In contrast to 5G solutions, by combining these different elements, we seek to bridge the gap between technical performance indicators and user satisfaction in highly dynamic and heterogeneous environments. This hybrid QoE model leverages advanced machine learning (ML) techniques and real-time data analytics to continuously assess and adapt to the dynamic conditions of modern networks, partic-

José Santos, (corresponding author) Sam Van Damme, Javad Sameri, and Filip De Turck are with Ghent University — imec, Belgium; Susanna Schwarzmann, Qing Wei, and Riccardo Trivisonno are with Huawei Technologies, Germany; Maria Torres Vega is with eMedia Research Lab, Belgium.

Digital Object Identifier: 10.1109/MCOM.003.2400594

ularly in 6G CC environments. While the dynamic and heterogeneous nature of the CC introduces complexity in resource management, it also provides an opportunity for more sophisticated orchestration strategies. Our approach not only enhances the accuracy of QoE predictions but also enables proactive management strategies that can anticipate and mitigate potential issues before these impact the end-user experience.

We also propose reinforcement learning (RL) [8] as the main orchestrator for QoE optimization in 6G, enabling real-time decisions to adapt network configurations and resources according to network performance and user perception. The envisioned RL-based orchestrator will handle key service scheduling tasks using the hybrid QoE model as input. The article presents results from a collaborative virtual reality (VR) use case [9], which involves a virtual kitchen scenario where participants collaboratively bake a pizza (detailed further below). By fostering seamless interaction between network conditions and user expectations, our hybrid QoE vision aims to establish a new standard for QoE management that is both robust and user-centric.

The remainder of the article is organized as follows: The next section discusses the literature and explores the motivation behind our hybrid QoE approach for XR applications. Then, we present the proposed 6G cloud-native (CN) infrastructure for QoE management within the CC while describing novel components for the proper orchestration of XR applications. Following that, we introduce the collaborative VR use case centered around a virtual kitchen and then discuss the obtained results. Last, we focus on open challenges and future directions. The final section concludes this article.

Background and Motivation

QoE modeling [10] (Fig. 1) for interactive multi-user, multi-device XR applications presents a complex challenge due to the multifaceted nature of human perception. Based on our previous work [9], it is evident that a single parameter cannot adequately capture the nuances of user experience. Instead, this article proposes that QoE is influenced by several factors, encompassing not just technological parameters but also contextual and human physiological responses. Technological aspects include network performance, device capabilities, and software efficiency, which can be estimated through network key performance indicators (KPIs). These objective metrics provide a foundation for understanding the technical aspects of user experience in real-time collaborative environments. On the other hand, human-related factors, such as physiological responses, are crucial for capturing the subjective aspects of QoE. In our experiments [9], heart activity seems a reliable indicator of user engagement and stress levels, which are critical for assessing QoE in immersive VR environments. Thus, the QoE in these networked XR environments can be represented by combining technological and human factors, offering a more comprehensive and accurate reflection of user experience. This approach allows for a practical and scalable way to estimate QoE by integrating both network and physiological data with subjective human responses, thereby advancing the precision of QoE models for XR applications.

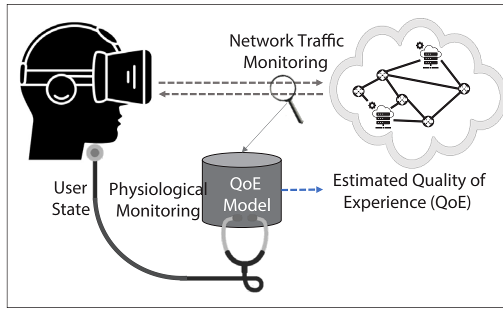

FIGURE 1. Conceptual overview of a QoE modeling approach.

Cloud-Native Extended Reality [11] represents the convergence of XR technology with cloud computing infrastructure, enabling the creation, distribution, and consumption of immersive experiences over the internet. In contrast to traditional XR setups that rely on locally hosted infrastructure, CN XR will leverage the scalability and flexibility offered by cloud providers to deliver immersive experiences to end users across various devices and locations. By leveraging cloud resources, developers can offload processing-intensive tasks, such as sampling and encoding, to remote cloud locations, thus enabling end users to access XR content without requiring high-end hardware [11]. *Kubernetes (K8s)* [12] has emerged as a fundamental building block in the development and deployment of CN applications, including XR experiences. As the most popular container orchestration platform, K8s simplifies the management of containerized workloads, providing automated scaling and deployment capabilities while ensuring high availability, fault tolerance, and efficient resource utilization. By leveraging K8s, developers can deploy XR workloads across multiple computing nodes, dynamically scale resources based on user demand, and seamlessly manage updates and rollbacks, thereby optimizing the delivery of immersive XR experiences to end users [11].

The Compute Continuum [5] represents a paradigm shift in computing, where processing power and computational resources are seamlessly distributed across a spectrum of devices, from edge to cloud. This continuum of virtual resources enables a more flexible and responsive infrastructure between centralized cloud DCs and decentralized edge devices [6]. In the context of 6G and next-generation networks, the CC is crucial for supporting latency-sensitive applications, such as XR experiences and real-time collaborative environments. XR Developers and industries agree that application roundtrip latency needs to remain below 20 ms, and the bandwidth should be able to scale up to 30 Gb/s [11]. By dynamically allocating resources across the CC, applications can achieve the necessary computational efficiency and responsiveness, adapting to varying workloads and user demands in near real-time. The CC will play a pivotal role in the future of networked systems, enabling new levels of performance, scalability, and user experience across diverse applications and services.

Reinforcement Learning [8] is a powerful ML-based technique that enables agents to learn optimal behaviors through trial and error, guided by rewards and penalties from their environment.

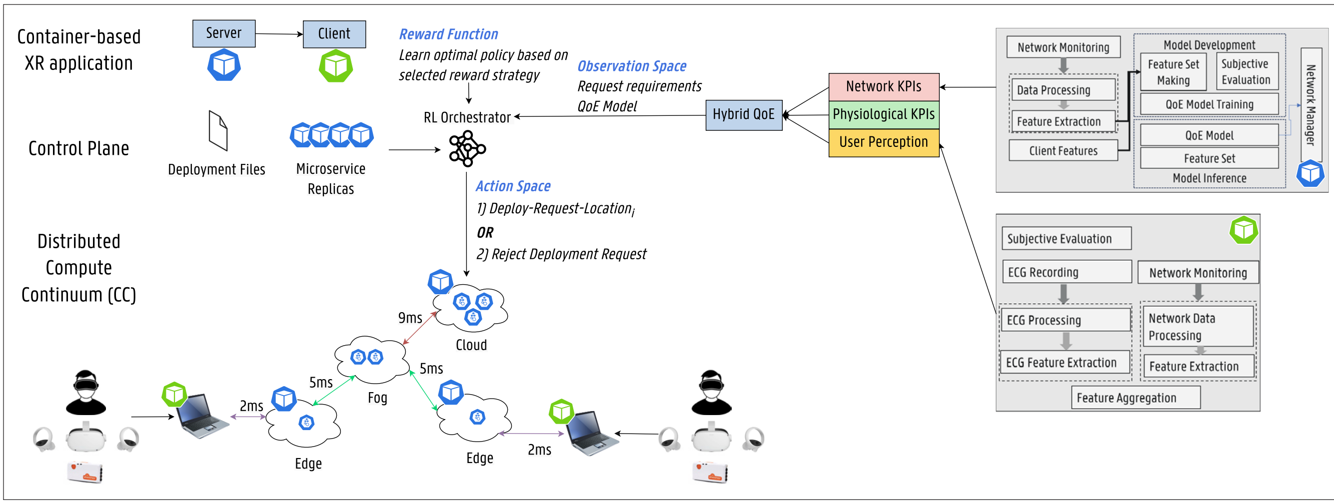

FIGURE 1. High-level view of the deployment of a container-based XR application within the 6G CC.

Unlike traditional supervised learning, where a model is trained on a fixed dataset, RL involves an agent interacting with its environment to discover the best strategies for achieving specific goals. Thus, RL is well-suited for dynamic and complex decision-making tasks where actions must be determined in near real-time. In the context of 6G, RL can be applied to optimize QoE for end users by dynamically adjusting network configurations and allocating computational resources based on real-time information, ensuring that performance metrics such as latency, bandwidth, and user satisfaction are continuously optimized. In addition, the versatility of RL methods to handle multi-objective optimization problems makes it ideal for scenarios where multiple, often conflicting, goals must be balanced, such as maximizing user experience while minimizing energy consumption of the infrastructure. As technologies continue to evolve, RL is expected to play a critical role in enabling intelligent, autonomous systems capable of adapting to the ever-changing demands of modern networks and applications.

QoE optimization in 6G networks presents both opportunities and challenges. The convergence of advanced technologies such as K8s, RL, and the CC is essential in addressing the increasing complexity and variability of user experiences. As 6G networks aim to provide ultra-reliable, low-latency communications for a wide range of applications, including immersive real-time collaborative XR environments, optimizing QoE becomes increasingly demanding. RL can facilitate the creation of intelligent, adaptive systems that dynamically optimize network resources and configurations in near real-time, ensuring precise alignment with user demands. The CC can also offer a flexible infrastructure that efficiently distributes computational tasks between the edge and the cloud, minimizing latency and enhancing overall performance. This article discusses further these open challenges while proposing solutions to address existing limitations in the coming years. The aim is to overcome today's bottlenecks to efficiently deploy XR applications in distributed cloud environments, not only meeting but exceeding user expectations and paving the way for unparalleled XR experiences across a broad spectrum of applications.

Toward Efficient QoE Management in 6G Networks for Networked Extended Reality

This section outlines the envisioned 6G CN distributed infrastructure, in which containerized XR components are deployed across multiple K8s clusters to deliver immersive experiences. A hybrid QoE model that combines network and physiological data with subjective user responses is also proposed for a comprehensive view of user experience. Lastly, we explore the potential of RL for efficient QoE optimization in 6G networks within the CC.

System Overview

The advent of novel architectural paradigms has paved the way for the deployment of containerized applications on distributed computing resources from the edge to the cloud, forming the known CC. Figure 2 illustrates the envisioned 6G CN infrastructure, which integrates key architectural concepts based on the K8s orchestration platform. This vision aims to deploy and manage XR applications within a CN architecture, leveraging K8s to efficiently optimize its components across the CC, including three main locations: edge, fog, and cloud. In contrast to traditional clouds, the fog and edge layers serve as intermediate computing zones, positioned closer to end users to reduce latency and improve performance. These layers bring processing power, storage, and memory capacity closer to the end devices, allowing more responsive and immersive applications, especially in XR. Each location consists of a K8s cluster managed via multi-cluster orchestration solutions [11]. In this setup, an XR application is decomposed into loosely coupled microservices (e.g., server and client), each independently developed, deployed, and maintained. Within K8s, these microservices run in containerized groups known as *Pods*, the smallest operational unit in K8s running in a shared execution environment. Compared to monolithic architectures, microservices offer several advantages: they enable independent scaling of components to meet demand, such as scaling the XR server instances without affecting clients, improve resilience by isolating failures to specific services, and allow modular updates and feature additions without disrupting the entire XR application. The ability to deploy microservices across heterogeneous layers ensures latency-sensitive components are placed closer to end users while utilizing cloud locations for computationally intensive tasks.

Hybrid QoE Modeling

In contrast to conventional QoE assessments that often focus on video quality or 2D imagery, XR environments require a more nuanced approach due to their immersive nature, which involves a range of sensory inputs and interaction dynamics. Traditional methods for modeling user perception in XR often overlook the human-centric nature of perception in these environments. Motivated by the potential of physiological signals to reflect human perception, this work introduces a novel hybrid QoE model that integrates network metrics with physiological responses, providing a comprehensive analysis of user perception in collaborative XR environments.

Network Metrics provide insights into the operational health of a network and the user experience, especially in complex applications such as XR. Furthermore, *physiological responses* assess how network conditions impact users physically and emotionally. This study focuses on the following metrics:

- *Throughput* is the amount of data successfully transmitted over the network in a given time, reflecting the overall network capacity.
- *Packet size* of the packets sent through the network can affect the latency of communication.
- *Inter-Arrival Time* refers to the time between the arrival of consecutive packets, which can also impact the latency.
- *Heart Rate Activity* represents the physiological state of the user.

Subjective user feedback is based on users' personal perceptions and experiences, offering qualitative insights, in contrast to network metrics, which often rely on quantitative data. Based on our previous work [9], four key aspects have been identified to assess the user experience in XR settings:

- *Latency* refers to the delay between a user's action and the system's response.
- *Jerkiness*, or the visual stuttering caused by burst impairments, disrupts smooth movement within the XR environment.
- *Synchronization* involves the alignment of sensory inputs, ensuring that the visual and haptic feedback in an XR environment is cohesive.
- *Network perception* refers to how users perceive the overall performance and reliability of the XR system.

In summary, our hybrid approach aims to offer a more comprehensive estimate of the end user's QoE by combining both objective and subjective data. By integrating diverse data streams, the hybrid model captures a broader spectrum of user responses and provides a more personalized and adaptive assessment.

QoE Optimizations via Reinforcement Learning

QoE optimizations are crucial in ensuring end-users receive the best possible service, particularly in immersive XR applications. In the context of the 6G CC, where XR applications are deployed across distributed computing resources, RL offers a powerful approach to optimize QoE. The envisioned RL-based orchestrator will leverage an observation space that combines network-level KPIs, physiological data, and subjective feedback. Using this rich observation space, RL methods can learn to select appropriate actions that maximize a complex reward function that includes numerous objectives, such as minimizing latency, ensuring high throughput, and maximizing the user's QoE. Thus, the hybrid QoE model is a critical component of the RL agent's observation space, providing the agent with a comprehensive understanding of user QoE for efficient decision-making.

The RL-based orchestrator focuses on key decision-making processes to ensure appropriate service provisioning:

- Determining the optimal placement of XR microservices across edge, fog, and cloud locations when receiving new user requests
- Predicting when a user will need to be migrated to another location to mitigate performance issues (For example, if a user's QoE deteriorates due to latency, the agent can relocate the microservices to an edge node closer to the user, improving responsiveness and overall experience.)
- Adapting the network configuration to meet individual user needs according to subjective feedback received as a reward.

User satisfaction, measured through subjective feedback, is possibly the most direct indicator of QoE. For example, if a user consistently reports low satisfaction due to video quality, the RL system can deploy microservices at a particular location, ensuring higher-resolution streaming for that user, even if it requires trade-offs with other KPIs.

RL-based methodologies offer several advantages over traditional approaches. RL can continuously learn and adapt to the dynamic network environment, ensuring optimal resource utilization and user satisfaction. By focusing on multiple KPIs, the proposed RL-driven QoE optimization can provide a more responsive and personalized user experience in 6G networks. As networks become more complex and demand more from the underlying infrastructure, RL offers a flexible and powerful mechanism to meet open challenges, ensuring that the promise of next-generation immersive XR applications is fully realized.

Use Case — Collaborative VR Application

Figure 3 illustrates the developed VR application [9], where participants collaboratively bake a pizza in a virtual kitchen. The kitchen layout features tables arranged in an "H" shape to separate the two users. All necessary ingredients are placed on these tables, requiring users to pass items to each other, fostering close collaboration. The virtual kitchen also features an oven for baking the pizza, but it is only accessible to one user. A blackboard on the wall guides participants through the step-by-step process. The virtual environment is intentionally designed so that neither user can access all the necessary utensils alone, making object handover crucial for task completion. The main steps are described below (illustrated in Fig. 3b):

User satisfaction, measured through subjective feedback, is possibly the most direct indicator of QoE. For example, if a user consistently reports low satisfaction due to video quality, the RL system can deploy microservices at a particular location, ensuring higher-resolution streaming for that user, even if it requires trade-offs with other KPIs.

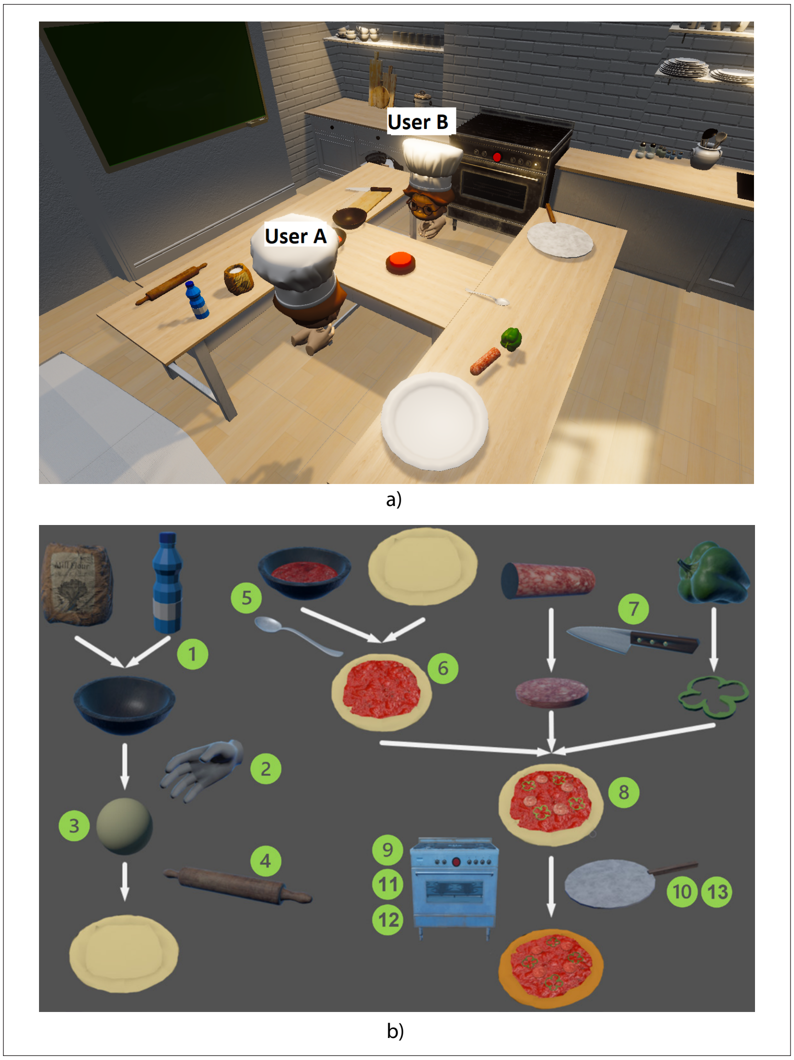

FIGURE 3. Overview of the collaborative pizza-baking application, illustrating the VR environment and the sequence of steps required to complete the task [9]: a) The collaborative VR environment; b) Schematic representation of the various steps.

- User A hands water and flour to User B, who mixes them to form a dough.
- User B kneads the dough and places it in the designated area.
- User A passes the rolling pin for User B to roll the dough.
- User B hands a spoon to User A, who spreads tomato sauce on the dough, refilling the spoon as needed.
- User A passes sausage and bell pepper to User B, who chops and places the toppings on the pizza.
- User B opens the oven and bakes the pizza using a pizza shovel passed by User A.

The use case has been chosen based on its gamified and pleasant nature, inspired by popular cooking games such as "Overcooked!." The design helps prevent user boredom during repeated playthroughs. Moreover, the use case integrates key collaborative VR elements essential for completing the task and ensuring a meaningful collaborative experience as described by Perez *et al.* [13]: *Deliberation* (communication between users), *Exploration* (locating and identifying objects), and *Manipulation* (interacting with and passing objects). Table 1 shows the timing and content of subjective evaluations during a single experimental session across three phases: *Pre-session*, *In-session*, and *Post-session*. The pre-session focuses on gathering participant demographics, assessing prior experience with VR, self-assessing technological proficiency, and conducting baseline tests, including the virtual reality sickness questionnaire (VRSQ) and *Ishihara* color vision tests. *In-session* evaluations used a 5-point Likert scale to measure various aspects of the experience of participants, such as playthrough time, VRSQ scores, and perceived latency and jerkiness on object interaction and collaboration. However, *In-session* questionnaires may impact task performance, particularly in scenarios where user engagement is critical. A few strategies could mitigate these effects by dynamically scheduling these questionnaires during natural breaks in the XR session to minimize disruption. The aim is to create concise and visually intuitive questionnaires to reduce cognitive load and response time. Also, passive monitoring techniques (e.g., physiological data) can complement or even replace in-session questionnaires, reducing interruptions. *Post-session* ranked perceived network performance and cybersickness.

The study involved 10 sessions with 20 participants (30% female, 70% male) between 22 and 47 years of age. Each participant used a Meta Quest 2 VR headset in the virtual environment. Wireshark has been used for deep traffic analysis, while two Shimmer GSR+ devices recorded physiological signals. Four different scenarios (*A*–*D*) have been assessed:

- *A* (No impairments).
- *B* (500 ms delay).
- *C* (500 ms throttling at a 50% rate).
- *D* (500 ms delay and 500 ms throttling at 50%).

All participants experienced each scenario for comprehensive analysis. For more details, refer to related studies [9, 14] on the impact of VR on user experience.

Results

Effective RL-based orchestration depends on accurate QoE estimation, as unreliable estimations can hinder policy learning and compromise service placement. This work evaluates the hybrid QoE model, laying the foundation for its future integration into RL. Three ML algorithms, namely Decision Tree (DT), Support Vector Machine (SVM), and Extreme Gradient Boosting (XGB) have been evaluated based on their ability to predict QoE criteria using different feature sets: network, physiological, and hybrid. These models produce binary predictions of user perception, distinguishing between the presence and absence of perceived issues as reported by participants. Table 2 summarizes the performance of each model measured using balanced accuracy (BAC), which provides a balanced evaluation by accounting for class imbalances, ensuring an accurate assessment of both positive and negative predictions.

Jerkiness criterion DT achieved the highest BAC (0.65) using network features, indicating that network metrics are more predictive of perceived *Jerkiness* than physiological data. This trend was consistent across SVM and XGB, where network features generally outperformed other models. *Latency* criterion Models trained on combined features consistently outperformed those trained on either network or physiological features alone. SVM achieved the highest BAC (0.79) with hybrid data, highlighting the importance of combining network metrics and physiological responses for accurate latency prediction. *Sync* criterion Hybrid models performed best for synchronization perception. Both SVM and DT achieved a BAC of 0.79, indicating that synchronization is influenced by both network and physiological data. *Network Preference* Criterion Network-only models performed better than physiological or hybrid models. DT and SVM achieved BAC values of 0.60 and 0.55, respectively, emphasizing the dominant role of network metrics in shaping user perceptions of network quality.

The results emphasize the value of hybrid models in improving the accuracy of QoE predictions, particularly for latency and synchronization-related criteria. Network-only features have been more effective for *Jerkiness* and *Network Preference*, suggesting that the relevance of physiological data depends on the specific QoE criterion being modeled. Overall, SVM and XGB emerged as the most effective classifiers, highlighting the potential of ML in refining QoE modeling for collaborative XR environments. By leveraging this hybrid QoE approach, RL-based methodologies can efficiently manage QoE in 6G networks. This work establishes the foundation for integrating QoE models into the observation space of RL to ensure adequate actions for numerous performance factors, with practical implementation and testing of RL strategies left for future work.

Timing	Content
Pre-session	Demographics (e.g., age, gender). Prior experience with VR (Never, Once, Quite, Very). Technological proficiency (Low, Medium, High). Baseline VRSQ and Ishihara tests.
In-session (5-point Likert scale)	Estimation of playthrough time (in seconds). VRSQ. Perceived Latency, Jerkiness on object interaction. Perceived task difficulty.
Post-session (Ranking)	Perceived network optimality. Perceived cybersickness inducement.

TABLE 1. Overview of the timing and content of the subjective evaluations during a single experimental session.

QoE Criterion	ML Classifier	Network (BAC)	Physiological (BAC)	Hybrid (BAC)
Jerkiness	DT	0.65	0.47	0.43
	SVM	0.36	0.62	0.34
	XGB	0.62	0.56	0.40
Latency	DT	0.63	0.58	0.78
	SVM	0.57	0.55	0.79
	XGB	0.57	0.54	0.68
Sync	DT	0.44	0.51	0.77
	SVM	0.55	0.45	0.79
	XGB	0.43	0.56	0.60
Network Preference	DT	0.60	0.51	0.40
	SVM	0.55	0.43	0.53
	XGB	0.54	0.45	0.45

TABLE 2. Results for the hybrid QoE modeling approach based on different features. Higher BAC is better.

Open Challenges and Future Directions

The use of RL for optimizing QoE in 6G networks presents not only challenges but also opportunities. As 6G networks aim to support ultra-low latency, massive device connectivity, and high reliability, leveraging RL for edge-cloud orchestration becomes critical but also complex. Scalability is a major challenge for RL algorithms in dynamic 6G environments. As edge devices and network nodes increase, the observation and action spaces for RL models grow exponentially, complicating real-time decision-making. Future research should explore hierarchical and distributed RL to address these scalability issues. Real-time adaptation is also crucial, as RL algorithms need to adjust to fluctuating user demands, network loads, and mobility patterns, helping to mitigate the impact of cybersickness. Investigating techniques such as offline and online learning strategies and transfer learning methods could improve the adaptability of RL models in these scenarios.

Security in RL-based orchestration for 6G networks is vital due to the risk of adversarial attacks and inherent uncertainties in network conditions. Future studies should focus on secure RL methods, including adversarial training, to improve model robustness by simulating and defending against potential attacks. Resilient frameworks are also needed to provide RL systems with the capability to detect and recover from malicious activities or unexpected disruptions. In addition, Multi-Agent Coordination will be key as edge-cloud orchestration in 6G involves multiple agents (e.g., base stations, edge servers) making simultaneous decisions. Future research should explore multi-agent RL and game theory approaches to address the challenges of coordination and cooperation among distributed agents.

Standardization and interoperability will also play a major role. Standardizing protocols, interfaces, and performance metrics is essential for QoE optimization in 6G networks, requiring collaboration between industry and academia to ensure interoperability across vendors and infrastructures. Energy Efficiency is also a critical concern in 6G networks due to the deployment of dense networks and edge devices. Optimizing the energy efficiency of RL-driven edge-cloud orchestration without compromising QoE is another essential area for future research. To address this challenge, novel methodologies focused on energy-aware RL could be explored further. In conclusion, overcoming these challenges is key to fully realizing the potential of RL for optimizing QoE in 6G networks.

Conclusions

As VR and XR technologies continue to push the boundaries of immersive experiences, optimizing QoE for end-users in the era of 6G networks presents both challenges and opportunities. This article presents a vision for QoE optimization in 6G, focusing on a hybrid model that integrates network metrics, physiological data, and subjective user feedback. Results from a collaborative VR use case demonstrate the potential of combining network and physiological data with subjective user feedback to predict the user experience more accurately. By leveraging RL for edge-cloud orchestration, our approach addresses the dynamic and complex nature of VR applications, ensuring that the stringent low-latency requirements of these applications are met. Despite promising advancements in RL-driven QoE optimization, several challenges remain, such as the scalability of RL algorithms, real-time adaptation to network conditions, and resilience against security threats, which are crucial for ensuring secure RL. Overcoming these obstacles is critical to fully realizing the potential of RL in managing QoE within 6G networks. This article offers a comprehensive overview and a forward-looking perspective on efficient QoE management in the CC.

By leveraging RL for edge-cloud orchestration, our approach addresses the dynamic and complex nature of VR applications, ensuring that the stringent low-latency requirements of these applications are met. Despite promising advancements in RL-driven QoE optimization, several challenges remain, such as the scalability of RL algorithms, real-time adaptation to network conditions, and resilience against security threats, which are crucial for ensuring secure RL.

Acknowledgment

José Santos and Sam Van Damme are funded by the Research Foundation Flanders (FWO), with respective grant numbers 1299323N and 1SB1822N.

References

[1] M. Eswaran and M. R. Bahubalendruni, "Challenges and Opportunities on Ar/Vr Technologies for Manufacturing Systems in the Context of Industry 4.0: A State of the Art Review," *J. Manufacturing Systems*, vol. 65, 2022, pp. 260–78.
[2] H. Bellini, "The Real Deal With Virtual and Augmented Reality," accessed 10 July 2024; available: https://www.goldmansachs. com/intelligence/pages/virtual-and-augmented-reality.html.
[3] A. A. Barakabitze *et al.*, "QoE Management of Multimedia Streaming Services in Future Networks: A Tutorial and Survey," *IEEE Commun. Surveys & Tutorials*, vol. 22, no. 1, 2019, pp. 526–65.
[4] M. S. Anwar *et al.*, "Subjective QoE of 360-Degree Virtual Reality Videos and Machine Learning Predictions," *IEEE Access*, vol. 8, 2020, pp. 148,084–99.
[5] D. Balouek-Thomert *et al.*, "Toward a Computing Continuum: Enabling Edge-to-Cloud Integration for Data-Driven Workflows," *Intl. J. High Performance Computing Applications*, vol. 33, no. 6, 2019, pp. 1159–74.
[6] J. Santos *et al.*, "Toward Low-Latency Service Delivery in a Continuum of Virtual Resources: State-of-Theart and Research Directions," *IEEE Commun. Surveys & Tutorials*, vol. 23, no. 4, 2021, pp. 2557–89.
[7] A. M. Aslam *et al.*, "Metaverse for 6G and Beyond: The Next Revolution and Deployment Challenges," *IEEE Internet of Things Mag.*, vol. 6, no. 1, 2023, pp. 32–39.
[8] X. Wang *et al.*, "Deep Reinforcement Learning: A Survey," *IEEE Trans. Neural Networks and Learning Systems*, vol. 35, no. 4, 2022, pp. 5064–78.
[9] S. Van Damme *et al.*, "Impact of Latency on QoE, Performance, and Collaboration in Interactive Multi-User Virtual Reality," *Applied Sciences*, vol. 14, no. 6, 2024, p. 2290.
[10] G. Kougioumtzidis *et al.*, "A Survey on Multimedia Services Qoe Assessment and Machine Learning-Based Prediction," *IEEE Access*, vol. 10, 2022, pp. 19,507–38.
[11] J. Santos, "Toward Cloud-Native Virtual Reality Applications: State-of-the-Art and Open Challenges," *Proc. NGMSE2024, the Next-Generation Multimedia Services at the Edge: Leveraging 5G and Beyond workshop, co-located with ISCC2024*, 2024.
[12] M. Luksa, *Kubernetes in Action*, Simon and Schuster, 2017.
[13] P. Perez *et al.*, "Emerging Immersive Communication Systems: Overview, Taxonomy, and Good Practices for Qoe Assessment," *Frontiers in Signal Processing*, vol. 2, 2022, p. 917684.
[14] J. Sameri *et al.*, "Collaborative Cooking in VR: Effects of Network Distortion in Multi-User Virtual Environments," *Proc. 15th ACM Multimedia Systems Conf.*, 2024, pp. 509–15.

Biographies

José Santos obtained his M.Sc. degree in Electrical and Computer Engineering in July 2015 from the University of Porto, Portugal. He completed his doctoral studies at Ghent University in April 2022, and he is currently a Postdoctoral Researcher in the Internet Technology and Data Science Lab (IDLab) Research Group at Ghent University — imec, Belgium. His research interests include Cloud and Fog Computing, IoT, Service Function Chaining, and Reinforcement Learning. His work has been published in more than 35 scientific publications. He received the imec Ph.D. Excellence Award in 2022, the Best Dissertation Award at NOMS 2023, and the IBM Innovation Award in 2023 based on the research conducted during his Ph.D.

Sam Van Damme obtained his B.Sc. (2016) and M.Sc. (2019) in Computer Science Engineering at Ghent University, Belgium. In 2024, he received joint Ph.D. degrees in Computer Science Engineering and Engineering Technology from Ghent University and KU Leuven, Belgium, respectively. He is currently affiliated as a Postdoctoral Researcher with the Internet Technology and Data Science Lab (IDLab) Research Group at INTEC, Ghent University — imec, Belgium. He has a strong interest and expertise in Quality-of-Experience (QoE) modeling and assessment, Human-Computer Interaction (HCI), Haptic Feedback, and Multi-User Interaction with a main focus on immersive multimedia technologies such as Virtual and Augmented Reality (VR/AR).

Javad Sameri obtained his M.Sc. degree in Artificial Intelligence and Robotics in November 2019 from the Iran University of Science and Technology. He is currently a doctoral researcher at the Internet Technology and Data Science Lab (IDLab), Ghent University — imec, Belgium. His research interests include physiological signal processing, human-machine interfaces, Virtual Reality (VR), and Quality of Experience (QoE) modeling.

Susanna Schwarzmann received her bachelor's and master's degrees in computer science from the University of Wurzburg. In 2022, she received her Ph.D. degree from TU Berlin, where she worked as a Research Assistant in the Network Architectures group. Since 2021, she has worked for Huawei Technologies, Munich Research Center, where she holds the position of a Research Engineer in the Advanced Wireless Technologies Lab.

Qing Wei received her M.Sc. in communication engineering with honors from TU Munich, Germany, in 2001 and her B.Sc. in Electrical Engineering from Shanghai JiaoTong Univ., China, in 1997. In 2002, she started working at DOCOMO Euro Labs, leading the mobile optical network team. In 2015, she joined Huawei Technologies as a principal researcher and worked in the area of 5G mobile network architecture. She has 50+ publications and 100+ IPRs in the area of mobility management and network slicing, and she also contributes directly to standardization activities such as NGMN, 3GPP, and ETSI, with 100+ accepted TDoC submissions. She is a task leader in numerous EU projects, such as SASER-SaveNet, 5G-Xhaul, and 5G-MoNArch. Her research interests are 6G network architectures with embedded AI and enriched network services.

Riccardo Trivisonno received M.Sc. and Ph.D. degrees in telecommunications engineering from the University of Bologna in 2000 and 2005, respectively. He joined Huawei Technologies in 2011 and is currently serving as Head of Network Architecture — Research and Standardization — for the Advanced Wireless Technologies Laboratory, Munich Research Center. Over the past ten years, he has been a leading contributor to the definition and the standardization of 5G network architecture and technologies, delivering to 3GPP Releases 15–18 in the areas of architecture modularization, network slicing, network analytics and QoS filing 100+ IPRs. He has been the Chairman of 6G-IA Pre-standardization WG since 2020, a Board Member of one6G Association since its foundation in 2021, and was elected one6G Board Vice-Chair in 2023.

Filip De Turck [F] leads the network and service management research group at Ghent University, Belgium, and imec. He (co-)authored over 700 peer-reviewed papers, and his research interests include the design of efficient softwarized networks and cloud systems. He is involved in several research projects with industry and academia, served as chair of the IEEE Technical Committee on Network Operations and Management (CNOM), and serves as a steering committee member of the IM, NOMS, CNSM, and NetSoft conferences. He served as Editor-in-Chief of IEEE Trans. Network and Service Management (TNSM).

Maria Torres Vega received the M.Sc. degree in telecommunication engineering from the Polytechnic University of Madrid, Spain, in 2009, and the Ph.D. degree from the Eindhoven University of Technology, The Netherlands, in 2017. She is currently a Tenure Track Assistant Professor with KU Leuven, Belgium, where her research focuses on devising human-driven control and management mechanisms for enhancing the perception of immersive systems. Her research interests include quality of service and QoE in immersive multimedia systems and autonomous network management.