

Large-Scale Sensor Systems for Sentient Environments: Emerging Research Directions

Boris Sedlak¹, Víctor Casamayor Pujol², Naser Hossein Motlagh³, *Senior Member, IEEE*, Roberto Morabito⁴, *Member, IEEE*, Francesc Wilhelmi⁵, Ildefons Magrans de Abril⁶, Ilir Murturi⁷, *Senior Member, IEEE*, Boris Bellalta⁸, *Senior Member, IEEE*, Praveen Kumar Donta⁹, *Senior Member, IEEE*, and Schahram Dustdar¹⁰, *Fellow, IEEE*

Abstract—The last decade of research in sensor systems was dominated by two parallel directions: first, improved data collection architectures, revolving around pervasive computing devices, and second, efficient methods to process data and extract values from them, ultimately through artificial intelligence (AI) technologies. The maturation and convergence of these two paths heralds a new generation of sentient environments for supporting humans in their daily interactions. Following a vision of pervasive, ambient intelligence, this article identifies nine research directions that will enable these environments to scale sustainably and reach a ubiquity, that will allow them to weave themselves into the fabric of everyday life.

Index Terms—Computing continuum (CC), edge computing, intelligence, sensor networks, sentient environments.

Received 12 June 2026; revised 13 July 2026; accepted 17 July 2026. Date of publication 23 July 2026; date of current version 11 September 2026. The work of Boris Sedlak, Víctor Casamayor Pujol, and Ildefons Magrans de Abril was supported by CNS2023-144359 and the European Union NextGenerationEU/PRTR under Grant MICIU/AEI/10.13039/501100011033. The work of Naser Hossein Motlagh was supported by the Research Council of Finland under Grant 362594. The work of Francesc Wilhelmi and Boris Bellalta was supported by projects TRUE Wi-Fi PID2024-155470NB-I00 (MICIU/AEI/10.13039/501100011033/FEDER,UE) and ICREA Academia 2024 (00077 AGAUR). Recommended by Lead Guest Editor Hossein Fotouhi. (Corresponding author: Boris Sedlak.)

Boris Sedlak, Víctor Casamayor Pujol, and Ildefons Magrans de Abril are with the Distributed Intelligence and Systems-Engineering Lab, Universitat Pompeu Fabra, 08002 Barcelona, Spain (e-mail: boris.sedlak@upf.edu; victor.casamayor@upf.edu; ildefons.magrans@upf.edu).

Naser Hossein Motlagh is with the Department of Computer Science, University of Helsinki, 00100 Helsinki, Finland (e-mail: naser.motlagh@helsinki.fi).

Roberto Morabito is with the Networked Systems Group, EURECOM, 06410 Biot, France (e-mail: roberto.morabito@eurecom.fr).

Francesc Wilhelmi and Boris Bellalta are with the Wireless Networking Research Group, UPF, 08002 Barcelona, Spain (e-mail: francisco.wilhelmi@upf.edu; boris.bellalta@upf.edu).

Ilir Murturi is with the University of Prishtina, 10000 Prishtinë, Kosovo (e-mail: ilir.murturi@uni-pr.edu).

Praveen Kumar Donta is with the Department of Computer and Systems Sciences, Stockholm University, 11419 Stockholm, Sweden (e-mail: praveen@dsv.su.se).

Schahram Dustdar is with the Distributed Systems Group, Vienna University of Technology, 1040 Wien, Austria, and also with the ICREA, 08010 Barcelona, Spain (e-mail: dustdar@dsg.tuwien.ac.at).

Digital Object Identifier 10.1109/JSAS.2026.3716613

I. INTRODUCTION

OVER the last decade, sensor networks continued to grow into increasingly dense architectures that support a myriad of applications using distributed sensing and processing. While this massive scaling—and the consequent challenges—were foreseen [1], [2], they remain directions of active research. For example, *security* and *privacy* gained even more importance under the prevalence of sentient environments [3]; similarly, *architectural design* and its implications on *scalability* are discussed alongside every innovation. Looking back at the research progress of the last decade, we identified two directions that, on the one hand, addressed many contemporary problems, but on the other, opened up research directions with countless new challenges: the emergence of *Edge Computing* as a new architectural paradigm, and the long-awaited maturation [4] of everyday artificial intelligence (AI) systems.

The push toward Edge computing was fueled by two main factors [5], [6]: performance and compliance. From a functional perspective, performing computations close to the data source supports real-time applications, avoids network congestion, and unlocks scalability; from a nonfunctional perspective, it supports privacy because data can be (pre)processed nearby, without exposure to the Internet. While Edge computing benefits latency-sensitive applications, it faces limitations (e.g., resource constraints or volatility) that Cloud systems are well-suited to fill. To combine the strengths of these different computational tiers, processing infrastructures are increasingly evolving toward composite platforms, such as the computing continuum (CC) [7]. The CC thus emerges as the architectural backbone for collecting and processing high-density data across large scale sensing networks.

At the same time, driven by increasing requirements of fully interactive, sentient spaces (e.g., smart cities), the density and volume of sensing data is skyrocketing [8]. Research and industry have agreed that this increasing data volume cannot be tackled with compute power alone, making AI technologies indispensable to extract value from sensing data.

Altogether, the precision and variety of sensor data have enabled applications that 20 years ago were futuristic [9], including smart cities with autonomous mobility. In recent years, societies have devoted a significant effort to develop such sentient environments [10], [11], which aim to improve the quality of life




Fig. 1. Sentient environments are becoming prevalent due to a confluence of high-quality sensing data and ubiquitous, low-latency computing. To support autonomous decision-making without considerable human intervention, AI technologies evolve within these boundaries.

(health, living) or the efficiency of resource use (energy, water), and production. Just as the Internet of Things (IoT) unified these disciplines [1], sensing and processing networks will support emerging applications like digital twins and autonomous cyber-physical systems [12], [13]. As shown in Fig. 1, the intelligence for operating and optimizing such environments is achieved through AI systems: autonomous components train AI models on continuous data streams and improve decision-making across distributed hardware and software components [14], [15].

In brief, the evolution of sensing and processing networks, together with the big push for AI systems, enables a new generation of interactive and responsive spaces. Still, integrating these systems into our everyday lives goes along with significant technical and societal frictions that span across various disciplines—including cost-efficient sensing, democratization of data and compute power, and network resilience. This article examines nine research directions that emerge from our vision of fully functional, sentient environments. Thus, it provides a comprehensive overview of technical hurdles and open research problems that will shape the next decade of research, rather than an exhaustive list of action points.

II. VISION AND SCOPE

Looking ten years into the future, we envision a society where smart environments have become seamlessly integrated into everyday life. They guide and support human activity through pervasive sensing and intelligent decision-making processes that often remain in the background. Such deep-rooted technologies—as envisioned by Mark Weiser [16]—will “*weave themselves into the fabric of everyday life until they are indistinguishable from it*”. While the communication and networking of sensor systems already happen largely unnoticed, we believe that this will also extend to their applications. For example, consider the friction from dealing with smartwatches and personal apps for tracking vital parameters, in combination with the manual programming of ambient ventilation, illumination, and acoustics to our well-being.

We envision a radical shift in current human–computer interaction: away from manual, user-based application setups, toward agent-based setups that act according to human intents, interpreted through an empathetic mixture of natural language, gestures, and observed habits. To monitor environments and human intents, sensory networks form a dense, global nervous system that integrates multimodal inputs, such as hyperspectral imaging and ambient wireless sensing; in resource-restricted environments, virtual sensing offers a cost-efficient alternative to resolve missing features. Sentient environments constantly monitor human comfort and evolve accordingly. For instance, after a user’s sudden drop in blood–oxygen saturation, and a consequently predicted drop in cognitive focus, an agent is tasked with countering this through the environmental control and life support systems. This conversion to agent-based application management removes the friction of using physical interfaces, like handhelds and keyboards, for programming computers and applications. Agents can rely on pervasive computing infrastructure, hence, providing personal compute power (e.g., through a smart watch) will become merely a user choice. Under the hood, this builds upon a zero-touch service orchestration fabric that manages a highly fragmented and heterogeneous computing and sensing infrastructure. This sentient reality is powered by specialized hardware, like neuromorphic chips or AI-accelerators, that enable complex inference within strict budgets (energy, latency). Crucially, the environment is not a static installation but a self-aware lifecycle of continuous AI model updates and verified interactions, ensuring data sovereignty and privacy through decentralized, verifiable protocols.

This vision should not be interpreted as an inevitable or frictionless trajectory. First, maintaining and coordinating among fragments sensing and processing infrastructure adds substantial costs and complexity for large-scale deployments. Second, interoperability remains a major barrier because sentient environments will depend on heterogeneous devices, communication technologies, data models, and governance regimes that are often controlled by different vendors or administrative domains. Finally, user acceptance may stall if autonomous decisions or




Fig. 2. Three layers that enable pervasive sentient environments: applications, management, and infrastructure. This article identifies and discusses core research challenges in these three domains.

data-sharing mechanisms are perceived as opaque, or misaligned with human preferences.

As part of this vision, we identified three core domains that must be addressed—as shown in Fig. 2. Although presented separately for clarity, these directions are tightly coupled: infrastructure capabilities enable management mechanisms, management mechanisms coordinate and safeguard application-level intelligence, and application requirements feed back into infrastructure and orchestration decisions.

Across these three layers, we identify *nine research directions*: specialized hardware and wireless networking, as part of the infrastructure layer; service/infrastructure management, intelligence lifecycles, and security/privacy, as part of the management layer; and agentic ecosystems, data economics, virtual sensing, and multimodal sensing, as part of the application layer. These directions provide a structured perspective on how advances in sensing, networking, distributed computing, and AI can jointly enable practical, scalable, and trustworthy sentient environments over the next decade. The remainder of this article will discuss these research directions in more detail. Finally, Section VI will synthesize our findings, extract a ten-year research outlook, and discuss cross-cutting challenges around sustainability and interoperability.

III. APPLICATION AND INTELLIGENCE CHALLENGES

The realization of sentient environments is fundamentally driven by the diverse and increasingly autonomous requirements of the application layer. This section examines the challenges involved in translating our 10-year vision into operational capabilities, where sensing and intelligence must manifest as a cohesive experience. We structure this discussion around the evolution of agent-driven interaction, the democratization of data across heterogeneous devices, and the advancement of multimodal and virtual sensing. These topics advance the primary interface through which the system perceives and interacts with the physical world.

A. Agentic Ecosystems

The transition from manually configured applications to agent-mediated applications enables greater autonomy and resilience by allowing systems to adapt to users and their context [17]. This vision requires mechanisms that translate high-level human intents into runtime decisions over distributed sensing and processing resources [18]. For example, an agent maintaining indoor comfort may combine occupancy sensing, air-quality estimation, local inference, and ventilation control, requesting additional validation when observations remain ambiguous [19], [20]. Such applications cannot be built only from fixed endpoints under a single provider's control; instead, they may combine physical or virtual sensors, Edge and Cloud processing, actuators, communication channels, other agents, or human-backed operations. As in the broader CC, these components are heterogeneous, dynamic, distributed across providers, and only partially observable [21]. Agentic ecosystems must therefore enable dependable composition and accountable interaction under decentralized ownership. However, decentralized ownership also introduces adoption barriers: providers must agree on interoperable descriptions and accountability mechanisms before agents can safely compose capabilities across organizational boundaries.

1) *Composable and Dependable Service Offerings*: Existing sensing and actuation capabilities may expose interfaces through which applications retrieve data streams or issue commands. However, invocation alone is insufficient for agent-mediated application formation: even in IoT environments, discovering and selecting suitable services requires descriptions of functionality, constraints, and quality attributes [22]. An autonomous agent must additionally determine whether a capability satisfies an intent, composes with other capabilities, can be instantiated under applicable constraints, and is likely to remain suitable during operation.

Future sensor ecosystems should therefore expose capabilities as *machine-actionable service offerings* [23]. An offering should make a capability discoverable, evaluable, instantiable, composable, and monitorable by agents. It should distinguish immediate adequacy from expected dependability over the application's operating horizon. For example, a battery-powered sensor may currently provide measurements with the required accuracy and latency, yet its estimated remaining energy may be insufficient to sustain operation for the next 6 h required by a safety-relevant application.

Offerings must support heterogeneous and recursive composition. An agent may combine an occupancy sensor, a virtual air-quality estimator, nearby inference capacity, connectivity, and a ventilation actuator into an operational capability that can itself be exposed as an offering for reuse in a broader application. Such compositions may also incorporate user-governed resources: private sensing streams or wearable observations may be made available under user-controlled access conditions rather than transferred to an application provider [24]. Users may likewise contribute sensing or local processing resources under explicit usage, governance, and economic conditions, as discussed in Section III-B.

Dependable composition also requires adaptation. If an air-quality estimate becomes unreliable, an agent may activate a physical sensor, select another inference resource, reduce functionality to preserve energy, or request inspection. Supporting such adaptation requires continuous monitoring and fallback-based recomposition over machine-readable information about capability quality, health, energy, connectivity, capacity, and access constraints.

2) Accountable Cross-Provider Execution: Consider the indoor-comfort agent introduced above. Confirming an air-quality anomaly may require temporary access to a privately owned sensor, inference on a commercial edge accelerator, and activation of a ventilation system controlled by the building operator. Although this composition may be technically valid, its execution crosses boundaries of ownership, trust, policy, cost, and responsibility. This creates an additional practical bottleneck: even when the technical orchestration is feasible, the deployment may be delayed by incompatible service-level agreements or reluctance among stakeholders to expose operational capabilities to autonomous agents. Agentic ecosystems must therefore support not only composition, but also the accountable acquisition and use of independently governed capabilities.

Keeping oversight in agentic ecosystems will mean that autonomous decisions must be inspectable, interruptible, and auditable, rather than requiring a human to approve every step. For each interaction, the ecosystem must establish 1) who requested the capability; 2) whether it was authorized; 3) which privacy and safety conditions applied; 4) what resources were consumed; 5) what was delivered; and 6) with what measured outcome. In that context, human oversight should be risk-adaptive rather than imposed as a permanent bottleneck in every control loop. Routine actions can be executed automatically within predefined authority boundaries, such as limits on cost, privacy exposure, physical actuation, or service degradation. Agents should escalate decisions when confidence is insufficient, policies conflict, preferences are unclear, or actions may have safety-, privacy-, or operation-critical consequences. Existing mechanisms—including zero-trust (ZT) authorization, decentralized identity, auditable agreements, provenance records, metering, and attestable execution—provide building blocks for this purpose [25], [26], [27], [28]. Their detailed realization belongs to the data-governance, management, and security challenges discussed in Sections III-B, IV-A, and IV-C.

Creating this emerging, agentic ecosystem will require application providers to expose programmable and verifiable offerings rather than only human-facing dashboards or application-specific integrations. The management layer can then compose, monitor, and adapt these offerings, while the infrastructure layer supports their secure activation, isolation, and runtime observability. In this way, heterogeneous sensing, computing, communication, and actuation resources become first-class components of agent-formed applications.

B. Democratizing Data and Devices

Contemporary smart environments are often driven by business models that provide “free” services in exchange for personal data; consequentially, companies create value through

large-scale data aggregation rather than innovating on the actual data-processing services [29], [30]. This paradigm is problematic for three reasons: first, it creates strong vendor lock-in effects because users cannot fully dictate the access—and potential commercialization—of their personal data anymore; second, this blocks open innovation because competitors lack access to these accumulated data silos; and third, it introduces a strict resource lock-in at the hardware layer. While consumers purchase and maintain increasingly capable edge hardware, its processing power remains entirely monopolized by proprietary application silos, reducing powerful local nodes to passive data-generation pipes for centralized clouds.

As part of our 10-year vision, we will transition from data-centric to service-centric models, where applications are decoupled from the data they operate on [24]. To foster maximum sovereignty and privacy, data remains within user-governed wallets; to process data and derive value from it, e.g., in a crowdsensing use case, service providers can negotiate access to users’ wallets or computing devices.

1) Decoupling Data From Services: Clients should not merely be informed—or kept in the loop—about how their personal data is handled [31]; they should have utmost power to orchestrate data sharing and processing. This demands an architectural shift toward a *hub-and-spoke* model: the central *hub* is not a Cloud instance, but a localized, physical cluster of hardware representing the user’s digital *persona*—a bounded network encompassing their personal devices, wearables, and the ambient sensors capturing their daily actions. The external services and applications act as the *spokes* that must query or interact with this personal cluster. On the positive side, there is progress to realize this through personal data wallets [32] or data sources authorized through decentralized identifiers (DIDs) [27]. On the flip side, there are numerous open-ended questions: *where should data wallets be physically hosted, how do we avoid that this becomes a single point of failure, and how to grant and verify access to wallets?*

Considering the development of pervasive computing environments, a logical choice is hosting wallets directly across users’ personal device cluster, and embedding localized peer-to-peer redundancy mechanisms to mitigate hardware failures. This physically anchors the hub-and-spoke architecture. Instead of extracting raw data streams to a remote cloud, the *data gravity* [33] of a personal cluster pulls the processing services inward. The application are brought to the devices where the data is created, executing locally within the safe boundary of the persona. Another question is: *how to avoid that data is copied from a wallet?* To this, there can be architectural solutions—only allowing local processing within the wallet and never share data—or cryptographic solutions, like trusted execution environments (TEEs) and differential privacy.

2) Crowdsensing and Crowdprocessing: Crowdsensing and crowdprocessing unlock unprecedented analytical and economic opportunities by transforming pervasive devices (e.g., drones, or city bikes) into a distributed network of sensing and computing nodes. On the one hand, these platforms can span across areas for solving complex computational problems at higher accuracy and lower cost, compared to traditional, centralized infrastructure [34]. On the other hand, it allows using

idle processing power—also called *dark compute* [35]—and leverage real-time data from a highly distributed sensing substrate. However, integrating millions of heterogeneous, uncalibrated devices introduces extreme technical fragmentation compared to uniform infrastructure. Consequentially, crowdsensing can yield low signal-to-noise ratio (SNR) [36], or demand considerable computational overhead for filtering noise, calibrating erratic readouts, and deduplicating data.

To ensure a reasonable number of devices, sensing and computing platforms must first overcome participation barriers (e.g., downloading and installing applications), and then keep users engaged. These barriers are not only technical, but also economic and behavioral, since users may reject participation if rewards are too small, battery or bandwidth costs are unclear, or the platform's privacy and trust guarantees are difficult to understand. Consequentially, platform operators must incentivize and compensate users for their contributed resources [34]. Micro-payments, e.g., through decentralized blockchain technologies, are a promising solution [37] for streaming fractions of a cent to user wallets based on contributed compute cycles or the utility of sensor data [26].

Finally, even with perfect technical integration and incentives, a platform will collapse without securing privacy and trust. This creates a structural tension because platforms must on one side protect user data [38], e.g., through differential privacy. On the other side, platforms must ensure their integrity, e.g., by monitoring user trust to protect against data poisoning and Sybil attacks [39]. Hence, the challenge is distinguishing between tampered data and such that was manipulated for privacy protection. The data- and resource-sharing models are inseparable from the security/privacy mechanisms in Section IV-C: user participation depends on verifiable access control, provenance, revocation, and privacy-preserving execution.

C. Virtual Sensors

The continued growth of sentient environments over the next 10 years requires sensor networks to grow into increasingly dense architectures. However, as density increases, scaling them further becomes inherently limited by physical infrastructure constraints—discussed in Section V-B—and skyrocketing operational costs. To overcome the limitations and management bottlenecks of hyper-dense physical deployments, virtual sensing (or proxy sensing) methodologies can provide a cost-efficient alternative. Using AI technologies, it is possible to infer unmeasured environmental variables from existing, heterogeneous data sources [40]. By doing so, we can provide a holistic and accurate digital twin of the environment, offering the high-fidelity insights of a dense network without the practical and financial burdens of high-density hardware setups [41]. In addition, besides improving deployment flexibility, reducing the reliance on additional physical sensing infrastructure may also contribute to lowering long-term energy consumption and hardware footprint.

1) Multimodal Deployment Gaps and Resource Constraints:

In many real-world smart environments, it is not always feasible to deploy a diverse set of sensing devices due to cost, infrastructure, or operational constraints. For example, while cameras

can effectively support tasks such as space monitoring and occupancy detection, they cannot directly measure environmental variables, such as light intensity, noise levels, temperature, or air quality indicators including CO₂ and PM_{2.5}. Conversely, relying solely on specialized ambient sensors limits the ability to comprehensively characterize the environment, as no single sensor provides total coverage of all relevant parameters. As a result, deploying multiple specialized sensors becomes impractical and costly when attempting to obtain a complete picture of environmental conditions.

2) Cross-Variable Inference and Heterogeneous Data Fusion:

Addressing the limitations of physical sensing requires advancing virtual sensing approaches that utilize AI to estimate environmental variables from disparate data sources. Research in this direction focuses on exploiting relationships among available measurements; for instance, surveillance camera frames can serve as inputs for hybrid deep learning models that leverage temporal continuity to estimate outdoor air quality [42]. Furthermore, predictive models can be developed to estimate complex indicators using measurements from simpler environmental variables, such as using temperature, humidity, and noise levels to infer CO₂ and black carbon concentrations [40]. Such cross-variable estimations are essential for sustaining high-fidelity sensing in environments where specific modalities are missing or restricted.

3) Reliability, Calibration, and Robustness of Proxy Models:

Sensor measurements used as inputs for virtual sensors are often susceptible to noise, drift, and measurement errors, which can significantly degrade the accuracy of proxy models. This necessitates the development of robust calibration methodologies to ensure that the underlying data remains reliable. Future research must advance beyond current Bayesian framework-based calibration approaches—which employ hybrid white-box and black-box calibrators to improve measurements [43]—to develop even more sophisticated methods. Ensuring that sensor readings are accurate enough to support virtual sensing is a critical hurdle for the long-term deployment of dependable and autonomous sentient systems.

Virtual sensing does not eliminate deployment costs entirely, instead, it shifts part of the burden toward model training, calibration, and validation. Thus, virtual sensing is closely linked to multimodal sensing because proxy estimates often depend on combining heterogeneous physical measurements. Advances in multimodal alignment, fusion, and missing-data handling directly improve the accuracy and robustness of virtual sensors, while virtual sensing can compensate for unavailable, costly, or failed modalities in multimodal systems.

D. Multimodal Sensing

Building on virtual sensing, multimodal sensing integrates heterogeneous physical and inferred signals to construct a more complete representation of the environment. To achieve the described vision over the next 10 years, key challenges in multimodal sensing will be the seamless integration of a diverse set of sensing technologies, ranging from high-bandwidth modalities such as hyperspectral cameras to ambient wireless signals

transmitted by technologies like Wi-Fi and 5G-and-beyond networks. By synthesizing heterogeneous inputs from video, LiDAR, audio, and time-series signals, multimodal AI systems can achieve the robust, adaptive, and intelligent decision-making necessary to support human activity recognition and environmental optimization. This fusion of diverse data sources is critical for creating a holistic and effective sensing infrastructure capable of perceiving and responding to the complex dynamics of future smart environments. However, the practical integration of these modalities may be slowed by incompatible data formats, vendor-specific application programming interfaces (APIs), calibration procedures, privacy constraints, and the cost of deploying and maintaining high-bandwidth sensing hardware at scale.

1) Structural Heterogeneity and Data Representation: Multimodal sensing systems generate data with fundamentally different structures, ranging from high-resolution spatial video streams and geometric 3-D point clouds to temporal time-series measurements and symbolic data. This diversity in representation creates significant hurdles for data preprocessing and normalization due to the conflicting spatial, temporal, and geometric properties of each modality [44]. Future research must address how to harmonize these disparate structures without losing the unique contextual information inherent to each individual sensing source.

2) Temporal Synchronization and High-Precision Alignment: Sensors often operate at vastly different sampling frequencies—such as video at 60 Hz versus inertial measurement units at 1000 Hz—and may not be naturally synchronized. Achieving accurate temporal alignment across modalities is essential, particularly in applications requiring high precision where millisecond-level offsets can significantly degrade performance, such as in smart industrial environments [45]. Hence, sensing architectures that ensure rigorous synchronization are a prerequisite for dependable multimodal fusion.

3) Annotation Complexity and Resilience to Data Loss: The varying volume and velocity of multimodal datasets make data labeling and annotation complex and time-consuming, where minor mislabeling can propagate inconsistencies across the entire system. Furthermore, real-world systems must contend with sensor dropouts caused by hardware failures, communication issues, or battery depletion. These conditions lead to incomplete data, necessitating the development of robust algorithms and effective data imputation strategies to maintain continuity in the environmental state estimation [46], [47].

4) Strategic Context-Preserving Fusion Methodologies: The diversity of data structures makes feature extraction and fusion particularly difficult, as effective methods must preserve the contextual information of each modality while combining them into a unified representation [48]. Selecting the appropriate strategy—whether early fusion of raw data, intermediate fusion of learned features, or late fusion of model outputs—remains a nontrivial research direction. Each approach presents significant trade-offs in terms of accuracy, computational complexity, and system robustness.

5) Mitigation of Dimensionality and Signal Imbalance: Multimodal systems frequently combine high-dimensional data,

such as 4 K video, with low-dimensional signals like temperature or humidity measurements. This imbalance can cause machine learning models to overemphasize high-dimensional modalities while underutilizing critical low-dimensional signals [49]. Addressing these dimensionality disparities requires novel AI methodologies for balancing disparate signal strengths and ensuring that environmental variables contribute effectively to the final intelligent output [50].

6) Computational Scalability for Edge Intelligence: High-resolution sensors and multiple high-bandwidth data streams generate massive volumes of data, often reaching terabytes per day. Processing and modeling these datasets require substantial computational resources, such as GPUs or TPUs, which creates a significant barrier for deployment on resource-constrained edge devices. Future research must focus on managing the trade-offs between computational efficiency, real-time performance, and model accuracy to enable multimodal intelligence at the edge of the network. The computational and communication demands of multimodal fusion also create direct requirements for the management layer—as discussed in Section IV-A—which must dynamically place processing tasks, select sensing modalities, and adapt data flows according to latency, energy, bandwidth, and reliability constraints.

IV. MANAGEMENT CHALLENGES

Bridging the gap between high-level application intent and raw physical resources requires a sophisticated management layer capable of autonomous orchestration. As sensor systems evolve toward a CC, the complexity of coordinating distributed resources exceeds manual oversight, necessitating a shift toward self-governing platforms. This section details the governance challenges inherent in this transition, focusing on intent-based orchestration, real-time discovery of dynamic resources, and the continuous lifecycle management of embedded AI. This hierarchical structure reflects the need for a distributed management fabric that can maintain consistency, trust, and performance across a fragmented and rapidly evolving technological landscape.

A. Infrastructure and Service Management

The massive scaling and inherent heterogeneity of modern computing environments demands a transition from manual application deployment to autonomous, intent-based management platforms. To weave sentient environments into our daily lives—as part of our 10-year vision—systems must evolve to orchestrate resources in milliseconds [51]. In particular, this means autonomously synthesizing data sources, computing nodes, and actuators according to high-level intents [6]. This shift is critical because ubiquitous infrastructure alone does not guarantee seamless service delivery across the CC. Instead, the management layer must mediate conflicting requirements, navigate environmental deterioration, and dynamically route workloads to the optimal processing nodes. To achieve this, we pose four research challenges alongside Fig. 3; most critically, the deployment of such management platforms may be delayed without standardized control interfaces, limited observability




Fig. 3. Ecosystem for managing infrastructure and services: an engine converts human intents to hierarchical goals, which affects services adaptation in real-time; for operation, services use pervasive sensing and computing devices that can be owned or provided by the crowd.

and unclear responsibilities across providers, and the cost of migrating from manually managed legacy systems.

1) Real-Time Discovery and Integration of Dynamic Infrastructure: Large segments of pervasive computing and sensing infrastructure are highly dynamic, with devices such as autonomous vehicles, satellites, and crowdsensors appearing and leaving the network erratically [52]. To create a unified platform, both transient infrastructure and static infrastructure—such as Road Side Units—must be discovered and integrated in real-time. This process is hindered by extreme device heterogeneity, which prevents standard load-balancing strategies as workload performance on unknown device types cannot be predicted a-priori [15]. Furthermore, as the extension and density of IoT-Edge networks increase, the system must overcome the complexities of platform-wide trust management and the network congestion caused by the continuous exchange of onboarding messages.

2) Intent Parsing and Hierarchical Goal Transformation: A core requirement for keeping operators in the loop is the continuous alignment of infrastructure behavior with high-level requirements through automated intent parsing. While the emergence of large language models (LLMs) offers new tools for extracting requirements and converting them into dynamic actions or hierarchical goals [53], [54], significant research is required to refine these systems for integration into fast-paced control loops. Moving away from cumbersome, manual requirements elicitation is essential for allowing the management platform to respond effectively to stakeholder needs in rapidly changing environments. This requires the management platform to autonomously synthesize resources according to high-level intents [18].

3) Platform Optimization Under Fragmentation and Uncertainty: Optimizing performance across the CC is challenging because no single entity possesses the full system state due to high distribution and dynamism [21]. While current research often aims to optimize this through reinforcement learning and partially observable markov decision processes (POMDPs) [55], [56], [57], these approaches fail to account for the actual fragmentation and uncertainty across providers, at every level of the infrastructure [58]. Future research must develop

optimization methods that are robust to this inherent uncertainty while maintaining platform-wide efficiency. Energy-aware orchestration may become an additional optimization objective alongside latency and reliability.

4) Rapid Orchestration and Fine-Grained Scaling: Dynamically changing environments demand near-instantaneous responses across increasingly smaller software components. Emerging approaches such as serverless computing and lightweight runtime engines like WebAssembly support fast, function-level scaling [59], [60], which is promising for low-latency applications. However, the orchestration layer must move beyond simple workload placement to also parameterize the dynamic switching between sensors and actors. This requires sophisticated placement strategies for microservices and serverless functions that are context-aware and topology-optimized to meet dynamic demands of sentient systems.

B. Embedded Intelligence Lifecycle Management

The evolution of AI in sensing systems from static, offline-trained components to dynamic, agentic entities necessitates a fundamental shift in lifecycle management strategies. To realize sentient environments over the next 10 years, we must move beyond the traditional model of infrequent, centralized updates toward a paradigm where AI components—leveraging foundation models and generative capabilities—evolve continuously through retraining and prompt adaptation. Achieving this objective is critical for maintaining the reliability and safety of large-scale sensor systems, ensuring that distributed intelligence can adapt to environmental changes while adhering to the strict energy, compute, and connectivity constraints inherent to the CC. Altogether, this recurrent retraining, rollback, and certification adds substantial operational costs for long-lived deployments, unless lifecycle responsibilities and update policies are clearly defined.

1) Structural Tension and Compatibility Drift: A significant hurdle arises from the mismatch between the rapid evolution of AI components and the long-term stability of physical sensing infrastructure. While models and tool interfaces may change on timescales of days, the underlying embedded devices are often maintained over several years. This mismatch leads to progressive compatibility drift and fragile deployments that accumulate operational debt [61]. In large-scale environments, this drift is not merely a software concern but a threat to system reliability and safety guarantees, requiring research into architectures that can bridge the temporal gap between fluid software intelligence and rigid hardware lifecycles. Beyond reliability, this also extends the lifetime of deployed infrastructures, thus reducing the cost of hardware replacement.

2) Resource-Constrained Optimization and Hardware Coupling: In the IoT and edge domain, updating intelligence is not a simple matter of redeployment; it involves complex trade-offs between accuracy, latency, and resource consumption [62]. Lifecycle management is inherently fragmented by heterogeneous hardware platforms that require distinct optimization pipelines. Consequently, the AI lifecycle becomes tightly coupled with specific hardware capabilities and deployment conditions rather




Fig. 4. Lifecycle stratification in pre-agentic and agentic embedded intelligence systems. Traditional embedded intelligence follows a stable, layered lifecycle with clear separation across system components. In contrast, agentic AI introduces fast-evolving artifacts (models, prompts, and tools) and cross-layer dependencies, where changes propagate across runtime, software, and hardware layers.

than being a purely software-driven process. Future research must address this fragmentation through methods that automate the adaptation of artifacts across heterogeneous targets without violating strict memory or power budgets.

3) *Ecosystem-Wide Propagation in Agentic AI*: The emergence of agentic AI transforms models from passive components into active entities that coordinate with other agents and external APIs [63]. This shift expands lifecycle management from individual models to entire ecosystems of interacting reasoning modules and orchestration logic. Within such systems, a change in a single part—such as a modified prompt structure or tool interface—can propagate in complex ways, affecting the behavior and safety of the entire system. Managing these interactions requires a shift from layered versioning to a system-wide stack approach as illustrated in Fig. 4, where cross-layer dependencies are explicitly modeled and managed.

4) *Fragility of Automated Deployment and Continuous Validation*: While leveraging LLMs to automate deployment stages promises “write once, run anywhere” flexibility, these abstractions often break down under real-world constraints [64]. Small variations in hardware targets or compiler behaviors can lead to inconsistent or nonexecutable outputs, particularly in resource-constrained environments [65]. Addressing this requires lifecycle-aware architectures that incorporate continuous validation and repair as core functions. Research must focus on observability frameworks that capture model decision quality, adaptive policies that balance reliability with performance, and coordination mechanisms to ensure consistent evolution across the CC.

C. Security and Privacy

Security and privacy have evolved from supporting requirements into the fundamental conditions that determine the deployability of large-scale sentient environments [66]. In the coming 10 years, the objective is to establish a resilient framework that can safeguard dense sensing infrastructures, edge-to-cloud pipelines, and AI-driven decision-making processes across multiple administrative domains. This transformation is critical because the attack surface now encompasses not only physical devices and networks, but also models, data pipelines, and orchestration logic—artifacts from Section IV-A. Achieving this goal requires moving beyond static trust assumptions to

a continuous, end-to-end verification architecture that protects against emerging risks (e.g., prompt injection in agentic AI and privacy leakage through ambient sensing) which ensures systemic integrity and user trust in increasingly autonomous smart environments. At the same time, overly complex security mechanisms can discourage the very users and operators they are meant to protect, which makes transparency and explainable policy enforcement core deployment concerns, not just nice-to-have usability features.

1) *End-to-End Attack Surface and ZT Architecture*: As sensing and computing infrastructures become more integrated, compromises in one segment can propagate throughout the entire system. Addressing this requires a shift toward the ZT paradigm, which treats every interaction as potentially exposed and requires explicit, continuous verification of identities, devices, and workloads [25], [67], [68], [69]. The technical hurdle lies in maintaining coherent policy enforcement across the CC under low-latency requirements, intermittent connectivity, and decentralized ownership. This difficulty is amplified when autonomous agents are introduced, shifting the problem from simple access control to defining the specific conditions, scope, and oversight under which a model may act upon a protected resource [70], [71].

2) *Model Layer Integrity and Agentic Risk Mitigation*: The model layer has become a primary operational attack surface as LLMs and agentic systems take on roles in retrieving information and triggering actions [72]. Protecting this layer requires more than conventional input filtering; it necessitates policy-aware mediation of tool access and explicit control over privilege to prevent indirect prompt injection attacks, where malicious instructions in external content manipulate an agent’s execution [71], [73]. Furthermore, ensuring the long-term safety of these systems requires verifiable lifecycle trust, necessitating stronger support for signed and attestable software updates, supply chain transparency, and provenance that spans data, code, and model artifacts over time.

3) *Privacy-Preserving Sensing and Multilayer Leakage*: While moving inference to the edge reduces central exposure, privacy risks remain through shared model updates, memorization, and membership inference attacks [74], [75]. Privacy protection must therefore be treated as an architectural principle extending from the sensing layer to the interaction layer. Implementing mechanisms such as Federated Learning (FL) with differential privacy, secure aggregation, and clipping is essential, yet these introduce difficult privacy-utility trade-offs, particularly when dealing with non-IID data and resource-constrained devices. Research must focus on balancing these protections with system performance while navigating the hardware fragmentation and interoperability issues inherent in integrating hardware-backed security features like TEEs [76].

V. INFRASTRUCTURE CHALLENGES

Supporting the layers of intelligence and management necessitates a robust and highly optimized physical foundation. The infrastructure layer represents the critical bottleneck where theoretical capabilities meet the limits of silicon and spectrum.

This section addresses the fundamental challenges of the underlying hardware and connectivity that underpin the entire sensor system architecture. We explore the transition toward domain-specific hardware architectures—essential for overcoming the stagnation of general-purpose computing—and the evolution of wireless networking to support the massive device densities and deterministic reliability required for sentient environments. By grounding the paper in these physical realities, we complete the cross-layer analysis necessary for a viable long-term vision.

A. Specialized Hardware

The realization of sentient environments over the next 10 years is fundamentally contingent upon an infrastructure layer optimized for the extreme responsiveness and efficiency demands of large-scale applications. General-purpose architectures are increasingly hitting physical and economic ceilings due to the fading of Moore’s Law and the breakdown of Dennard Scaling, creating a level of friction that threatens the user experience [77], [78]. To overcome these limitations, the industry must pivot toward domain-specific architectures (DSAs) and functional specialization, meticulously tailoring hardware to the requirements of the applications they serve [79]. This shift is critical for achieving the peak efficiency and energy autonomy required to support dense, ubiquitous sensor networks, and complex AI models operating in resource-constrained niches. However, specialization can also slow adoption when hardware platforms become expensive or obsolete before reaching the end of their operational lifetime.

1) Domain-Specific Acceleration and Reconfigurable Logic:

To sustain the computational throughput required for modern AI and high-velocity data streams, hardware must move beyond general-purpose processing. While GPUs and TPUs have revolutionized training and inference [80], the migration of these tasks to the edge necessitates NPUs and specialized micro-controllers capable of low-quantization neural network execution [81]. However, such specialization risks hardware obsolescence as models evolve, requiring research into “software-defined hardware,” where FPGAs or coarse-grained reconfigurable architectures (CGRAs) can be dynamically reconfigured to adapt to shifting application needs [82].

2) In-Situ Processing and Energy Autonomy:

The densification of sensor networks creates a significant strain on traditional energy budgets, particularly for deployments relying on battery reserves or energy harvesting. To maintain efficient data flows, architectures must eliminate the Von Neumann bottleneck, the physical separation of memory and processor, through emerging paradigms such as processing-in-memory, which executes logic directly within memory arrays specialized for stream-heavy workloads [83]. Furthermore, neuromorphic architectures offer a compelling path toward energy autonomy through event-driven, quiescent operation, though significant research is required to bridge the gap from laboratory prototypes to manufacturing scales necessary for ubiquitous deployment [84].

3) Hardware-Level Security and Cryptographic Resilience:

The distributed nature of sentient networks necessitates robust encryption at the hardware level, especially when data must




Fig. 5. Integrated communication and sensing technologies to support future smart/sentient environments.

transit between proximal compute nodes. Because traditional cryptographic standards are often too power-intensive for limited sensors, there is a growing need for specialized modules implementing lightweight encryption primitives [85]. Future research must explore chiplet-based approaches to co-locate sensors and dedicated encryption engines in a single package [86], while simultaneously integrating post-quantum cryptographic primitives into the silicon to ensure long-term resilience against emerging computational threats [87].

4) *Software Ecosystem Fragmentation and Codesign:* As hardware becomes increasingly diverse—spanning NPUs, FPGAs, and neuromorphic processors—the lack of unified programming abstractions creates a “silo” effect that locks software to specific vendors [88]. These new architectures demand a fundamental shift toward hardware-software codesign, where algorithmic design is fundamentally reimaged for the underlying hardware, such as transitioning from frame-based deep learning to asynchronous spiking neural networks for neuromorphic chips [89]. Overcoming this fragmentation and the high fabrication complexity of integrated chiplet packages remains a formidable hurdle for achieving the economic viability required for large-scale deployment.

B. Wireless Networking

Future smart and sentient ambient environments are envisioned as complex cyber-physical systems where vast numbers of heterogeneous sensors and actuators must collaborate to perceive and react to their surroundings, as illustrated in Fig. 5. To realize this vision over the next 10 years, wireless infrastructures must evolve from passive coexistence to AI-driven, multi-RAT convergence, seamlessly integrating terrestrial solutions—such as Wi-Fi, cellular, and LoRa—with satellite constellations to provide ubiquitous coverage. This evolution is critical for supporting the anticipated expansion of immersive extended reality, autonomous mobility, and intelligent industrial automation, all of which demand a fundamental shift in how connectivity is coordinated to ensure scalability, determinism, and resilience in increasingly dense and unpredictable operating environments.

1) Massive Network Scalability and Interference Coordination: As smart environments scale to support thousands of devices, networks must accommodate dense deployments and highly dynamic traffic patterns, further complicated by emerging paradigms such as joint sensing and communication. A primary hurdle remains the management of neighboring interference in shared spectrum, where small payloads often incur disproportionately high medium-access overhead [90]. While multi-AP cooperation offers a path toward improved spectrum utilization through coordinated resource allocation, achieving this without excessive signaling latency or overhead remains unresolved. Research must focus on managing device density across multiple coexisting networks while ensuring that frequent updates from distributed AI models and digital twin telemetry do not lead to systemic bottlenecks.

2) Deterministic Quality of Service in Unlicensed Bands: Emerging mission-critical applications, such as collaborative robotics and remote healthcare, require reliable and predictable communication with strict latency guarantees [91]. However, unlicensed wireless technologies have traditionally prioritized peak data rates over the deterministic performance required by cyber-physical systems. While multilink strategies and scheduled access mechanisms—such as R-TWT and L4S—aim to enhance predictability, the stochastic nature of contention-based access and unpredictable interference from noncooperative devices continue to impede strict reliability bounds. The challenge lies in translating application-level timing requirements into enforceable guarantees at the MAC and physical layers while maintaining high spectral efficiency across distributed, synchronized nodes.

3) Resilient Operation and Fault-Tolerant Architectures: Sensor nodes frequently operate under harsh environmental conditions that induce measurement drift, bias, and hardware failures, significantly degrading network performance [92], [93]. In networked control systems, achieving fault tolerance is increasingly complex as failures may arise from both individual node malfunctions and communication disruptions. To sustain reliable operation, the system must move beyond simple redundancy toward active fault-tolerant control and predictive maintenance strategies that can detect, isolate, and accommodate faults in real-time. Future research directions include the development of neural-network-based observers and adaptive reconfiguration mechanisms capable of maintaining system stability and data integrity despite the inevitable failures of sensors and network links in unpredictable environments.

VI. DISCUSSION AND SYNTHESIS

The research directions identified in this article reveal a fundamental shift toward a deeply integrated CC. While the challenges have been categorized into layers for clarity, their resolution ultimately requires a cross-layer codesign philosophy spanning applications, management, and infrastructure. Across these layers, realization will depend not only on technical progress, but also on reducing deployment complexity, ensuring interoperability, proving economic viability, and building user and operator trust.

TABLE I
RESEARCH LANDSCAPE FOR SENTIENT ENVIRONMENTS

Layer	Research Topic	Core Technical Focus
App.	Agentic Ecosystems	Composable Offerings, Accountable Execution
	Data Democratization	Discovery, Governance
	Virtual Sensing	Cross-modal Inference
	Multi-modal Fusion	Synchronization, Efficiency
Mgmt.	Service Orchestration	Intent-based Automation
	AI Lifecycle	Continuous Validation, Repair
	Security & Privacy	Zero-Trust, Model Integrity
Infra.	Specialized HW	Domain-Specific Efficiency
	Wireless Networks	Determinism, Scalability

It is important to note that the directions discussed throughout this paper should not be interpreted as an exhaustive roadmap for every aspect of future smart environments. Rather, we have deliberately focused on a set of research areas that we believe will play a particularly important role in shaping the next decade of large-scale sensor systems. Also, such directions reflect both emerging technological trends and the collective experience of the authors across sensing systems, distributed computing, AI, networking, and cyber-physical systems. The broader landscape is undoubtedly much wider, and many additional challenges and opportunities will certainly emerge as these technologies continue to mature.

This section synthesizes these findings and discusses how advances across these domains may collectively drive the transition from today's smart environments toward increasingly sentient ones. Table I provides a compact index of the research landscape explored in this work.

A. Synthesizing the Application Layer

The primary objective for the application layer is the movement toward autonomous, agentic entities that can perceive and react to the world with high fidelity. The technical hurdles centered on *Agentic Ecosystems* and *Multimodal Sensing* necessitate research into machine-actionable composition, accountable cross-provider execution, and sophisticated temporal alignment. To make these systems viable, we must overcome the *Virtual Sensing* challenge—bridging modal gaps through AI-driven inference—while ensuring that the entire stack remains *Sustainable* through zero-power architectures and survival-aware orchestration policies. Furthermore, the *Democratization of Data* remains a key societal hurdle, requiring standardized metadata and open governance models to prevent infrastructure silos.

Ten-Year Outlook — Infrastructure Layer

- Agent ecosystems coordinate sensing and actuation autonomously.
- Virtual sensors routinely replace expensive or unavailable physical measurements.
- Data sharing becomes incentive-driven and interoperable by design.

B. Evolving Management and Lifecycle Strategies

As intelligence becomes embedded within the environment, the management layer must evolve from static configuration to dynamic, intent-based governance. The most significant hurdles identified in *Infrastructure Management* involve the transformation of high-level human intents into hierarchical control loops capable of millisecond-scale execution. Closely coupled with this is the *Intelligence Lifecycle*, where research must address the “compatibility drift” between fluid AI models and long-lived physical hardware. Finally, as the attack surface expands, *Security and Privacy* research must transition toward Zero-Trust Architectures (ZTA), ensuring the integrity of not only the data, but also the model logic and agentic interactions.

Ten-Year Outlook — Infrastructure Layer

- High-level intents are automatically translated into coordinated actions spanning devices, networks, and AI services.
- AI lifecycle management becomes autonomous, continuously validating, adapting, and repairing models and generated artifacts across heterogeneous deployments.
- Data sharing becomes incentive-driven and interoperable by design.

C. Grounding the Infrastructure Foundation

The infrastructure layer represents the physical constraints of the 10-year vision. Overcoming the stagnation of general-purpose computing requires a focus on *Specialized Hardware*, specifically DSA and Neuromorphic computing, to eliminate the Von Neumann bottleneck and achieve energy autonomy. Simultaneously, *Wireless Networking* research must pivot from best-effort paradigms to deterministic quality of service (QoS) and massive scalability. This requires coordination across unlicensed bands to provide reliable connectivity for safe, real-time cyber-physical interaction.

Ten-Year Outlook — Infrastructure Layer

- Domain-specific and neuromorphic architectures enable pervasive intelligence under strict energy and latency constraints.
- Wireless networks provide deterministic, scalable, and AI-assisted connectivity capable of supporting real-time cyber-physical interactions.
- Sensing, communication, and AI inference operate as a unified infrastructure substrate, seamlessly supporting distributed intelligence across the computing continuum.

D. Convergence Across Layers

Ultimately, the transition toward sentient environments will not be driven by advances in any single layer, but by the convergence of applications, management, and infrastructure into a cohesive CC. The presented research directions are

not independent, but form a hierarchy of mutually reinforcing capabilities. Specialized hardware and wireless networking provide the foundational compute, energy, and connectivity substrate. Infrastructure and service management, embedded intelligence lifecycle management, and security and privacy operate above this substrate to discover, coordinate, validate, and protect distributed resources and AI components. Built on these foundations, the intelligence layer supports adaptive and autonomous applications through agentic ecosystems, data and device democratization, virtual sensing, and multimodal sensing. The dependencies among research directions can also be bidirectional: the application-level requirements for latency, reliability, privacy, and accuracy do influence orchestration, networking, and hardware design; conversely, infrastructure constraints determine which sensing and reasoning capabilities can be deployed in practice.

Evaluating this multilayered convergence requires a robust validation infrastructure; however, to date, few institutions have access to such an infrastructure, forcing the majority of researchers and developers to seek remedy in lower-fidelity simulations [21]. However, the community needs large-scale testbeds, real-world deployments, interoperable simulators, and representative benchmarks designed specifically for the large network sensors and the CC. While developing this foundational testing infrastructure is often underappreciated by the research community compared to purely theoretical contributions, prioritizing this effort is essential. Without these practical validation tools, translating these complex, conceptual architectures into verifiable and scalable systems will remain out of reach.

If the challenges outlined in this article are successfully addressed, future sensor systems will move beyond their traditional role as data collection platforms. Instead, they will become adaptive and trustworthy digital companions, reasoning, and acting in support of human activities. In this sense, the long-standing vision of ubiquitous computing may gradually materialize, with sensing and embedded intelligence becoming seamlessly nested into the environment, supporting society while largely disappearing into the background.

E. Sustainability as a Cross-Cutting Concern

One topic that intentionally cuts across all research directions discussed in this article is sustainability. While sustainability is often treated as an independent research area, we argue that, in the context of future sentient environments, it should instead be viewed as a system-wide property emerging from decisions made across applications, management, and infrastructure.

Consequently, we do not attempt to provide an exhaustive roadmap for sustainable sensing systems. Such a goal would require a dedicated treatment extending far beyond the scope of this article. Instead, we try to highlight several areas where the challenges discussed throughout this manuscript are likely to have a direct impact on the environmental footprint and long-term viability or large-scale sensor infrastructures.

As shown in Table II, we expect that many of the proposed research directions will simultaneously create opportunities for

TABLE II

EXAMPLES OF SUSTAINABILITY OPPORTUNITIES AND TRADE-OFFS ASSOCIATED WITH THE RESEARCH DIRECTIONS DISCUSSED IN THIS ARTICLE

Research Direction	Sustainability Opportunities	Sustainability Trade-offs
Virtual Sensing	Reduced physical sensing operations and communication overhead through AI-driven inference	Additional computational demand, model maintenance, and validation overhead
Multi-modal Sensing	Adaptive selection of sensing modalities based on quality and energy requirements	Increased sensing, synchronization, and data processing complexity
Agentic Ecosystems	Autonomous workload placement and resource optimization across the continuum	Additional reasoning, coordination, and communication overhead
Intelligence Lifecycle	Extended useful lifetime of deployed systems through adaptation and continuous maintenance	Continuous monitoring, validation, and update costs throughout the system lifecycle
Infrastructure Management	Energy-aware and carbon-aware orchestration across heterogeneous resources	Additional control, telemetry, and decision-making overhead
Specialized Hardware	Improved performance-per-watt and energy-efficient AI execution	Manufacturing footprint, hardware heterogeneity, and replacement costs
Wireless Networking	Efficient coordination and data transport across distributed sensing systems	Communication overhead may dominate energy consumption at large scale

improving sustainability while introducing new sources of overhead. We therefore advocate that future research will need to balance performance, autonomy, reliability, and environmental impact simultaneously, recognizing that improvements in one dimension may create trade-offs in another.

F. Interoperability and Standardization

Future sentient environments will integrate diverse sensing, computing, networking, and actuation technologies originating from heterogeneous devices, platforms, and communication ecosystems. Consequently, interoperability and standardization are essential enablers, extending beyond mere device connectivity to include common data models, semantic context representations, AI services, and orchestration frameworks. Standardized interfaces, communication protocols, and service descriptions are vital for facilitating seamless interactions across distributed infrastructures without requiring vendor-specific adaptations [94].

Despite advancements, challenges persist due to the coexistence of diverse communication technologies, proprietary data models, and fragmented standards, all of which complicate large-scale deployments. In addition, integrating with legacy systems designed for conventional operations necessitates backward compatibility and incremental migration strategies. Addressing these issues will require interoperable architectures that support diverse devices throughout their operational lifecycles while minimizing costs.

To enhance interoperability, future sentient environments should adopt open, modular, and extensible architectures [95]. Standardized protocols and service interfaces, combined with semantic interoperability, will empower heterogeneous systems to share data and a common understanding of context and capabilities. These features are crucial for achieving zero-touch service orchestration and autonomous management in sentient environments.

Moreover, emerging standards should support decentralized trust management, privacy-preserving data exchange, and the secure lifecycle management of evolving AI models [96]. Establishing open and interoperable standards across these layers

will be essential for avoiding vendor lock-in, reducing deployment complexity, and enabling scalable, economically sustainable sentient environments. Ultimately, interoperability and standardization will be fundamental to realizing human-centric sentient environments that seamlessly integrate heterogeneous technologies and adapt to evolving user needs.

VII. CONCLUSION

In this article, we outlined a vision ultimately inspired by a long-standing ambition of ubiquitous computing: to make technology so deeply integrated into everyday life that it becomes indistinguishable from it. As envisioned more than three decades ago, future sensing and computing systems may eventually disappear into the background, while seamlessly supporting human activities through pervasive intelligence and autonomous decision-making. Although such a vision once appeared distant, recent technological advances suggest that many of its foundational building blocks are materializing. Our goal has not been to accurately predict the future research landscape, but to provide a structured perspective on how recent advances may converge toward increasingly sentient environments. We have described nine research directions that must be pursued to realize this vision in a trustworthy, interoperable, and scalable fashion. We believe that progress toward large-scale sentient environments will depend not only on advances within individual disciplines, but also on stronger interactions across traditionally separate research communities. We hope that the perspectives presented in this article help researchers connect these different pieces and identify both the opportunities and the challenges that lie ahead.

REFERENCES

- [1] J. A. Stankovic, "Research directions for the Internet of Things," *IEEE Internet Things J.*, vol. 1, no. 1, pp. 3–9, Feb. 2014.
- [2] E. Borgia, "The Internet of Things vision: Key features, applications and open issues," *Comput. Commun.*, vol. 54, pp. 1–31, 2014.
- [3] A. Rejeb, K. Rejeb, S. Simske, H. Treiblmaier, and S. Zailani, "The big picture on the Internet of Things and the smart city: A review of what we know and what we need to know," *Internet Things*, vol. 19, 2022, Art. no. 100565.

- [4] N. Maslej et al., "The AI index 2025 annual report," AI Index Steering Committee, Inst. Hum.-Centered AI, Stanford Univ., Tech. Rep., Apr. 2025.
- [5] W. Shi and S. Dustdar, "The promise of edge computing," *Computer*, vol. 49, no. 5, pp. 78–81, May 2016.
- [6] M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 50, no. 1, pp. 30–39, Jan. 2017.
- [7] P. K. Donta, I. Murturi, V. Casamayor Pujol, B. Sedlak, and S. Dustdar, "Exploring the potential of distributed computing continuum systems," *Computers*, vol. 12, no. 10, Oct. 2023, Art. no. 198.
- [8] Huawei, "Intelligent world 2035," Huawei Technologies Co., Ltd., Tech. Rep., Sep. 16, 2025. [Online]. Available: <https://www.huawei.com/en/giv>
- [9] C. Urmson et al., "Autonomous driving in urban environments: Boss and the urban challenge," in *The DARPA Urban Challenge: Autonomous Vehicles in City Traffic*, M. Buehler, K. Iagnemma, and S. Singh, Eds., Berlin, Heidelberg: Springer, 2009, pp. 1–59.
- [10] D. J. Cook and S. K. Das, "How smart are our environments? An updated look at the state of the art," *Pervasive Mobile Comput.*, vol. 3, no. 2, pp. 53–73, Mar. 2007.
- [11] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash, "Internet of things: A survey on enabling technologies, protocols, and applications," *IEEE Commun. Surv. Tut.*, vol. 17, no. 4, pp. 2347–2376, Fourth Quarter 2015.
- [12] E. Pretel, E. Navarro, V. C. Pujol, and S. Dustdar, "Digital twins and artificial collective intelligence: Synergies for the future," *IEEE Internet Comput.*, vol. 29, no. 1, pp. 75–85, Jan./Feb. 2025.
- [13] C. Qian, X. Liu, C. Ripley, M. Qian, F. Liang, and W. Yu, "Digital twin—Cyber replica of physical things: Architecture, applications and future research directions," *Future Internet*, vol. 14, no. 2, 2022, Art. no. 64.
- [14] J. Kephart and D. Chess, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, Jan. 2003.
- [15] B. Sedlak, V. C. Pujol, P. K. Donta, and S. Dustdar, "Equilibrium in the computing continuum through active inference," *Future Gener. Comput. Syst.*, vol. 160, pp. 92–108, 2024.
- [16] M. Weiser, "The computer for the 21st century," *SIGMOBILE Mobile Comput. Commun. Rev.*, vol. 3, no. 3, pp. 3–11, Jul. 1999.
- [17] J. C. Augusto, V. Callaghan, D. Cook, A. Kameas, and I. Satoh, "Intelligent environments: A manifesto," *Hum.-Centric Comput. Inf. Sci.*, vol. 3, no. 1, 2013, Art. no. 12.
- [18] A. Morichetta, N. Spring, P. Raith, and S. Dustdar, "Intent-based management for the distributed computing continuum," in *Proc. IEEE Int. Conf. Serv.-Oriented Syst. Eng.*, 2023, pp. 239–249.
- [19] D. J. Cook, M. Youngblood, and S. K. Das, "A multi-agent approach to controlling a smart environment," in *Designing Smart Homes: The Role of Artificial Intelligence*, Berlin, Germany: Springer, 2006, pp. 165–182.
- [20] B. Dong, V. Prakash, F. Feng, and Z. O'Neill, "A review of smart building sensing system for better indoor environment control," *Energy Buildings*, vol. 199, pp. 29–46, 2019.
- [21] B. Sedlak, V. C. Pujol, I. M. de Abril, P. K. Donta, A. N. Toosi, and S. Dustdar, "Service orchestration in the computing continuum: Structural challenges and vision," *IEEE Internet Comput.*, vol. 30, no. 2, pp. 88–98, Mar. 2026.
- [22] M. Achir, A. Abdelli, L. Mokdad, and J. Benothman, "Service discovery and selection in IoT: A survey and a taxonomy," *J. Netw. Comput. Appl.*, vol. 200, 2022, Art. no. 103331.
- [23] I. Magrans de Abril, B. Sedlak, V. C. Pujol, and S. Dustdar, "Composable autonomic offerings for decentralized service provisioning in the computing continuum," in *Proc. Workshop Adv. Tools, Program. Lang., Platforms Implementing Evaluating Algorithms Distrib. Syst.*, 2026, pp. 1–9.
- [24] R. Verborgh, "Re-decentralizing the web, for good this time," *Linking World's Inf.: Essays Tim Berners-Lee's Invention the World Wide Web*, vol. 52, pp. 215–230, 2023.
- [25] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," *NIST Special Pub.*, vol. 800-207, pp. 1–52, 2020.
- [26] A. Dawod, D. Georgakopoulos, P. P. Jayaraman, A. Nirmalathas, and U. Parampalli, "IoT device integration and payment via an autonomic blockchain-based service for IoT device sharing," *Sensors*, vol. 22, no. 4, 2022, Art. no. 1344.
- [27] B. Zhang, R. Shi, X. Li, and M. Zhang, "Decentralized identifiers based IoT data trusted collection," *Sci. Rep.*, vol. 15, no. 1, Feb. 2025, Art. no. 4796.
- [28] M. Calvo and M. Beltrán, "Remote attestation as a service for edge-enabled IoT," in *Proc. IEEE Int. Conf. Serv. Comput.*, 2021, pp. 329–339.
- [29] S. Geisler et al., "Knowledge-driven data ecosystems toward data transparency," *J. Data Inf. Qual.*, vol. 14, no. 1, pp. 3:1–3:12, Dec. 2021.
- [30] R. Verborgh, "A data ecosystem fosters sustainable innovation," Dec. 7, 2020. [Online]. Available: <https://ruben.verborgh.org/blog/2020/12/07/a-data-ecosystem-fosters-sustainable-innovation/>
- [31] R. Mortier, H. Haddadi, T. Henderson, D. McAuley, and J. Crowcroft, "Human-data interaction: The human face of the data-driven society," *SSRN Electron. J.*, Oct. 2014, doi: [10.2139/ssrn.2508051](https://doi.org/10.2139/ssrn.2508051).
- [32] S. A. de Chaves and F. B. V. Benitti, "Improving personal data protection interaction in digital identity wallets: A user-centric analysis," in *Proc. Hum.-Comput. Interact.-INTERACT*, 2026, pp. 213–233.
- [33] M. Campbell, "Smart edge: The effects of shifting the center of data gravity out of the cloud," *Computer*, vol. 52, no. 12, pp. 99–102, 2019.
- [34] A. Capponi, C. Fiandrino, B. Kantarci, L. Foschini, D. Kliazovich, and P. Bouvry, "A survey on mobile crowdsensing systems: Challenges, solutions, and opportunities," *IEEE Commun. Surv. Tut.*, vol. 21, no. 3, pp. 2419–2465, Third Quarter, 2019.
- [35] M. B. Taylor, "A landscape of the new dark silicon design regime," *IEEE Micro*, vol. 33, no. 5, pp. 8–19, Sep. 2013.
- [36] F. Restuccia, N. Ghosh, S. Bhattacharjee, S. K. Das, and T. Melodia, "Quality of information in mobile crowdsensing: Survey and research challenges," *ACM Trans. Sensor Netw.*, vol. 13, no. 4, pp. 34:1–34:43, Nov. 2017.
- [37] M. Sober, G. Scaffino, S. Schulte, and S. S. Kanhere, "A blockchain-based IoT data marketplace," *Cluster Comput.*, vol. 26, no. 6, pp. 3523–3545, Dec. 2023.
- [38] T. M. Nguyen and X.-S. Vu, "Privacy and trust in IoT ecosystems with Big Data: A survey of perspectives and challenges," in *Proc. IEEE 9th Int. Conf. Big Data Comput. Serv. Appl.*, 2023, pp. 215–222.
- [39] D. Kalui, D. Zhang, G. Muketha, and J. Onsomu, "Simulation of trust-based mechanism for enhancing user confidence in mobile crowdsensing systems," *IEEE Access*, vol. 8, pp. 20870–20883, 2020.
- [40] M. A. Zaidan et al., "Intelligent calibration and virtual sensing for integrated low-cost air quality sensors," *IEEE Sensors J.*, vol. 20, no. 22, pp. 13638–13652, Nov. 2020.
- [41] H.-O. Choe and M.-H. Lee, "Digital twin-based virtual sensor data prediction and visualization techniques for smart swine barns," *Sensors*, vol. 25, no. 24, Jan. 2025, Art. no. 7690.
- [42] X. Wang, M. Wang, X. Liu, Y. Mao, Y. Chen, and S. Dai, "Surveillance-image-based outdoor air quality monitoring," *Environ. Sci. Ecotechnol.*, vol. 18, 2024, Art. no. 100319.
- [43] M. A. Zaidan et al., "Intelligent air pollution sensors calibration for extreme events and drifts monitoring," *IEEE Trans. Ind. Inform.*, vol. 19, no. 2, pp. 1366–1379, Feb. 2023.
- [44] P. Nanjundan, S. Singh, J. George, and S. Eslamian, "Challenges in preprocessing and normalization of heterogeneous data," in *Multimodal Learning Using Heterogeneous Data*. New York, NY, USA: Elsevier, 2026, pp. 75–88.
- [45] T. R. Bennett, N. Gans, and R. Jafari, "Data-driven synchronization for Internet-of-Things systems," *ACM Trans. Embedded Comput. Syst.*, vol. 16, no. 3, pp. 1–24, 2017.
- [46] R. Wu, S. D. Hamshaw, L. Yang, D. W. Kincaid, R. Etheridge, and A. Ghasemkhani, "Data imputation for multivariate time series sensor data with large gaps of missing data," *IEEE Sensors J.*, vol. 22, no. 11, pp. 10671–10683, Jun. 2022.
- [47] Y. Liu, T. Dillon, W. Yu, W. Rahayu, and F. Mostafa, "Missing value imputation for industrial IoT sensor data with large gaps," *IEEE Internet Things J.*, vol. 7, no. 8, pp. 6855–6867, Aug. 2020.
- [48] D. Lahat, T. Adali, and C. Jutten, "Multimodal data fusion: An overview of methods, challenges, and prospects," *Proc. IEEE*, vol. 103, no. 9, pp. 1449–1477, Sep. 2015.
- [49] T. Hasanin, T. M. Khoshgoftaar, J. L. Levey, and R. A. Bauder, "Severely imbalanced Big Data challenges: Investigating data sampling approaches," *J. Big Data*, vol. 6, no. 1, pp. 1–25, 2019.
- [50] M. Altalhan, A. Algarni, and M. T.-H. Alouane, "Imbalanced data problem in machine learning: A review," *IEEE Access*, vol. 13, pp. 13686–13699, 2025.
- [51] M. M. Cummings, "Man versus machine or man machine," *IEEE Intell. Syst.*, vol. 29, no. 5, pp. 62–69, Sep./Oct. 2014.
- [52] H. Kokkonen et al., "Autonomy and intelligence in the computing continuum: Challenges, enablers, and future directions for orchestration," Feb. 2023.
- [53] C. Capova, A. Morichetta, A. Lackinger, and S. Dustdar, "Intent-to-learning translation for computing continuum management," in *Proc. IEEE 33rd Int. Conf. Netw. Protoc.*, 2025, pp. 1–6.
- [54] D. Wu et al., "NetLLM: Adapting large language models for networking," in *Proc. ACM SIGCOMM 2024 Conf.*, New York, NY, USA, 2024, pp. 661–678.

- [55] N. Filinis et al., "Intent-driven orchestration of serverless applications in the computing continuum," *Future Gener. Comput. Syst.*, vol. 154, pp. 72–86, May 2024.
- [56] Z. Pan et al., "MagicScaler: Uncertainty-aware, predictive autoscaling," *Proc. VLDB Endow.*, vol. 16, no. 12, pp. 3808–3821, Aug. 2023.
- [57] G. Quattrocchi, E. Incerto, R. Pincioli, C. Trubiani, and L. Baresi, "Autoscaling solutions for cloud applications under dynamic workloads," *IEEE Trans. Serv. Comput.*, vol. 17, no. 3, pp. 804–820, May/Jun. 2024.
- [58] N. Toumi, M. Bagaa, and A. Ksentini, "Hierarchical multi-agent deep reinforcement learning for SFC placement on multiple domains," in *Proc. IEEE 46th Conf. Local Comput. Netw.*, 2021, pp. 299–304.
- [59] T. Pusztai and S. Nastic, "ChunkFunc: Dynamic SLO-aware configuration of serverless functions," *IEEE Trans. Parallel Distrib. Syst.*, vol. 36, no. 6, pp. 1237–1252, Jun. 2025.
- [60] C. Marcelino, T. Pusztai, and S. Nastic, "Roadrunner: Accelerating data delivery to WebAssembly-based serverless functions," in *Proc. 26th Int. Middleware Conf.*, New York, NY, USA, 2025, pp. 354–368.
- [61] R. Morabito and M. Tatipamula, "The internet of physical AI agents: Interoperability, longevity, and the cost of getting it wrong," 2026, *arXiv:2603.15900*.
- [62] V. Janapa Reddi et al., "Edge impulse: An MLOPS platform for tiny machine learning," *Proc. Mach. Learn. Syst.*, vol. 5, pp. 254–268, 2023.
- [63] F. Jiang, C. Pan, K. Wang, P. Michiardi, O. A. Dobre, and M. Debbah, "From large AI models to agentic AI: A tutorial on future intelligent communications," *IEEE J. Sel. Areas Commun.*, vol. 44, pp. 3507–3540, 2026.
- [64] G. Wu, S. Tarkoma, and R. Morabito, "Consolidating tinyML lifecycle with large language models: Reality, illusion, or opportunity," *IEEE Internet Things Mag.*, vol. 8, no. 5, pp. 88–96, Sep. 2025.
- [65] R. Morabito and G. Wu, "When the code autopilot breaks: Why large language models falter in embedded machine learning," *Computer*, vol. 58, no. 11, pp. 32–41, 2025.
- [66] T. Gebremichael et al., "Security and privacy in the industrial Internet of Things: Current standards and future challenges," *IEEE Access*, vol. 8, pp. 152351–152366, 2020.
- [67] C. Bicer, I. Murturi, P. K. Donta, and S. Dustdar, "Blockchain-based zero trust on the edge," in *Proc. Int. Conf. Comput. Sci. Comput. Intell.*, 2023, pp. 1006–1013.
- [68] J. Weber, I. Murturi, X. Bajrami, R. Farahani, P. K. Donta, and S. Dustdar, "Lightweight webassembly-based intrusion detection for zero trust edge networks," *IEEE Access*, vol. 14, pp. 35337–35355, 2026.
- [69] D. Kaltenböck, I. Murturi, and S. Dustdar, "A zero trust single sign-on framework with attribute-based access control," in *Proc. 26th Int. Conf. Bus. Informat.*, 2024, pp. 149–157.
- [70] I. Murturi, P. K. Donta, V. C. Pujol, A. Morichetta, and S. Dustdar, "Learning-driven zero trust in distributed computing continuum systems," in *Proc. IEEE Int. Conf. Dependable, Autonomic Secure Comput.*, 2023, pp. 44–49.
- [71] A. Yehudai et al., "Survey on evaluation of LLM-based agents," 2025, *arXiv:2503.16416*.
- [72] W. Xu, C. Huang, S. Gao, and S. Shang, "LLM-based agents for tool learning: A survey," *Data Sci. Eng.*, Berlin, Germany: Springer, 2025, pp. 1–31.
- [73] Z. Zhang et al., "Agent-safetybench: Evaluating the safety of LLM agents," 2024, *arXiv:2412.14470*.
- [74] H. Li, L. Ge, and L. Tian, "Survey: Federated learning data security and privacy-preserving in edge-Internet of Things," *Artif. Intell. Rev.*, vol. 57, no. 5, 2024, Art. no. 130.
- [75] M. Miranda, E. S. Ruzzetti, A. Santilli, F. M. Zanzotto, S. Bratières, and E. Rodolà, "Preserving privacy in large language models: A survey on current threats and solutions," *Trans. Mach. Learn. Res.*, 2025. [Online]. Available: <https://openreview.net/forum?id=Ss9MTTN7OL>
- [76] A. Muñoz, R. Riós, R. Román, and J. López, "A survey on the (in) security of trusted execution environments," *Comput. Secur.*, vol. 129, 2023, Art. no. 103180.
- [77] M. M. Waldrop, "The chips are down for Moore's law," *Nature News*, vol. 530, no. 7589, 2016, Art. no. 144.
- [78] R. Dennard, F. Gaensslen, H.-N. Yu, V. Rideout, E. Bassous, and A. LeBlanc, "Design of ion-implanted MOSFET's with very small physical dimensions," *IEEE J. Solid-State Circuits*, vol. 9, no. 5, pp. 256–268, Oct. 1974.
- [79] J. L. Hennessy and D. A. Patterson, "A new golden age for computer architecture," in *Proc. Turing Lecture Del. 45th ACM/IEEE Annu. Int. Symp. Comput. Architecture*, Los Angeles, CA, 2018, pp. 48–60, http://iscaconf.org/isca2018/turing_lecture.html, Video: <https://www.youtube.com/watch?v=3LVEjsn8Ts>
- [80] N. P. Jouppi et al., "In-datacenter performance analysis of a tensor processing unit," in *Proc. 44th Annu. Int. Symp. Comput. Archit.*, 2017, pp. 1–12.
- [81] A. Gholami, S. Kim, Z. Dong, Z. Yao, M. W. Mahoney, and K. Keutzer, "A survey of quantization methods for efficient neural network inference," in *Proc. Low-Power Comput. Vis.*, 2022, pp. 291–326.
- [82] R. Prabhakar et al., "Plasticine: A reconfigurable architecture for parallel patterns," in *Proc. 44th Annu. Int. Symp. Comput. Archit.*, 2017, pp. 389–402.
- [83] O. Mutlu, S. Ghose, J. Gómez-Luna, and R. Ausavarungnirun, *A Modern Primer on Processing in Memory*. Berlin, Germany: Springer, 2023, pp. 171–243.
- [84] M. Davies et al., "Loihi: A neuromorphic manycore processor with on-chip learning," *IEEE Micro*, vol. 38, no. 1, pp. 82–99, Jan./Feb. 2018.
- [85] V. A. Thakor, M. A. Razzaque, and M. R. A. Khandaker, "Lightweight cryptography algorithms for resource-constrained IoT devices: A review, comparison and research opportunities," *IEEE Access*, vol. 9, pp. 28177–28193, 2021.
- [86] G. H. Loh and R. Swaminathan, "The next era for chiplet innovation," in *Proc. Des., Automat. & Test Eur. Conf. Exhib.*, 2023, pp. 1–6.
- [87] H. Nejatollahi, N. Dutt, S. Ray, F. Regazzoni, I. Banerjee, and R. Camarota, "Post-quantum lattice-based cryptography implementations: A survey," *ACM Comput. Surv.*, vol. 51, no. 6, pp. 1–41, Jan. 2019.
- [88] J. Auerbach et al., "A compiler and runtime for heterogeneous computing," in *Proc. 49th Annu. Des. Automat. Conf.*, 2012, pp. 271–276.
- [89] J. Shalf, "The future of computing beyond Moore's law," *Philos. Trans. Roy. Soc. A: Math., Phys. Eng. Sci.*, vol. 378, no. 2166, Jan. 2020, Art. no. 20190061.
- [90] G. Durisi, T. Koch, and P. Popovski, "Toward massive, ultrareliable, and low-latency wireless communication with short packets," *Proc. IEEE*, vol. 104, no. 9, pp. 1711–1726, Sep. 2016.
- [91] M. K. Atiq, R. Muzaffar, Ó. Seijo, I. Val, and H.-P. Bernhard, "When IEEE 802.11 and 5G meet time-sensitive networking," *IEEE Open J. Ind. Electron. Soc.*, vol. 3, pp. 14–36, 2021.
- [92] P. K. Donta, A. Lapkovskis, E. Mingozzi, and S. Dustdar, "Resilient by design—active inference for distributed continuum intelligence," in *Proc. IEEE Int. Conf. Pervasive Comput. Commun. Workshops Affiliated Events (PerCom Workshops)*, 2026, pp. 477–480.
- [93] K. Shaheen, A. Chawla, and P. S. Rossi, "Trust-enhanced distributed Kalman filtering for sensor fault diagnosis in sensor networks," *IEEE Trans. Signal Inf. Process. Netw.*, vol. 11, pp. 1178–1187, 2025.
- [94] S. S. Albouq, A. A. Abi Sen, N. Almasf, M. Yamin, A. Alshanqiti, and N. M. Babhouh, "A survey of interoperability challenges and solutions for dealing with them in IoT environment," *IEEE Access*, vol. 10, pp. 36416–36 428, 2022.
- [95] M. Sharma, A. Tomar, and A. Hazra, "Edge computing for Industry 5.0: Fundamental, applications, and research challenges," *IEEE Internet Things J.*, vol. 11, no. 11, pp. 19070–19093, Jun. 2024.
- [96] D. Kreuzberger, N. Kühl, and S. Hirschl, "Machine learning operations (MLOPS): Overview, definition, and architecture," *IEEE Access*, vol. 11, pp. 31866–31879, 2023.



Boris Sedlak received the Ph.D. degree in computer science from TU Wien, Vienna, Austria, in 2025.

He is a Postdoctoral Researcher with Universitat Pompeu Fabra (UPF), Barcelona, Spain, at the Distributed Intelligence and Systems-Engineering Lab (DISL). His research interests include intelligent and adaptive mechanisms for resource-constrained Edge devices, causal methods for the computing continuum, and service-oriented computing.



Víctor Casamayor Pujol received the Ph.D. degree in robotics from Universitat Pompeu Fabra (UPF), Barcelona, Spain, in 2025.

He is a tenure-track Professor with Universitat Pompeu Fabra (UPF), Barcelona, Spain. Prior to that, he worked as a Postdoctoral Researcher with the Distributed Systems Group at TU Vienna. His research interests revolve around self-adaptive methodologies for computing continuum systems, including SLO-based definitions, causal and machine learning inference, and robotics.



Naser Hossein Motlagh (Senior Member, IEEE) received the D.Sc. degree in networking technology from the School of Electrical Engineering, Aalto University, Espoo, Finland, in 2018.

He is an Academy Research Fellow with the Department of Computer Science, the University of Helsinki within the Nokia Center for Advanced Research (NCAR). Previously, he was a Postdoctoral fellow within Helsinki Institute for Information Technology (HIIT). His research interests include the Internet of Things, wireless sensor networks, environmental sensing, smart buildings, and aerial and underwater vehicles.



Ilir Murturi (Senior Member, IEEE) received the Ph.D. degree in computer science from TU Wien, Vienna, Austria, in 2022.

He is currently an Assistant Professor with the University of Prishtina and an external Senior Researcher with the Distributed Systems Group, TU Wien. Previously, he worked as a Postdoctoral Researcher and University Assistant with TU Wien. His current research interests include the Internet of Things, Edge Computing, EdgeAI, and privacy in distributed, self-adaptive, and cyber-physical systems.

Dr. Murturi has authored more than 40 peer-reviewed publications in international journals and conferences and has served as a TPC member for several international events.



Roberto Morabito (Member, IEEE) received the Ph.D. degree in networking technology from Aalto University, Espoo, Finland, in 2019.

He is currently an Assistant Professor with the Communication Systems Department, EURECOM, France. Before joining EURECOM, he was a Senior Researcher with the University of Helsinki and a Postdoctoral Researcher with Princeton University, USA. He also spent eight years with Ericsson Research, Finland, where he worked on cloud systems, IoT, and cyber-physical systems. His research focuses on the design and orchestration of AI-driven distributed systems under computing and networking constraints.



Boris Bellalta (Senior Member, IEEE) received the Ph.D. degree in information and communication technologies from Universitat Pompeu Fabra (UPF), Barcelona, Spain, in 2007.

He is currently a Professor with Universitat Pompeu Fabra (UPF), where he leads the Wireless Networking group. His research interests lie in the area of wireless networks and performance evaluation, with a particular emphasis on Wi-Fi technologies and machine learning-based adaptive systems.



Francesc Wilhelm received the Ph.D. degree in information and communication technologies from the Universitat Pompeu Fabra (UPF), Barcelona, Spain, in 2020.

He is currently a Tenure Track Professor with Universitat Pompeu Fabra (UPF). Before, he was with Nokia Bell Labs and the Catalan Technology Center for Telecommunications (CTTC).

Dr Wilhelm was a recipient of the IEEE Com-Soc Best Young Professional Award, in 2024.



Praveen Kumar Donta (Senior Member, IEEE) received the Ph.D. degree in computer science and engineering from IIT (ISM) Dhanbad, Dhanbad, India, in 2021.

He is a Professional Member of the ACM. He is currently an Associate Professor (Docent) with the Department of Computer and Systems Sciences, Stockholm University, Sweden. He serves on the editorial boards of the IEEE Internet of Things Journal & Network Magazine, Computing (Springer), Scientific Reports (Nature), Wiley's ETT, and Engineering Reports.



Ildefons Magrans de Abril received the Ph.D. degree in electronic engineering from the Universitat Autònoma de Barcelona (UAB), Barcelona, Spain, in 2008.

He is a Postdoctoral Researcher with Universitat Pompeu Fabra, Barcelona, Spain, at the Distributed Intelligence and Systems-Engineering Lab. His research interests include intelligent and adaptive mechanisms for the computing continuum, sequential optimization methods and computational neuroscience.



Schahram Dustdar (Fellow, IEEE) received the Ph.D. degree in business informatics from the University of Linz, Linz, Austria, in 1992.

He is a Full Professor of computer science, focusing on internet technologies, and heads the Distributed Systems Group at TU Wien, Vienna, Austria.

Prof. Dustdar is Chairperson of the Informatics Section of the Academia Europaea. He has been an Associate Editor for the IEEE TRANSACTIONS ON SERVICES COMPUTING, *ACM Transactions on the Web*, and *ACM Transactions on Internet Technology*. He has received several accolades, including the ACM Distinguished Scientist and IBM Faculty awards.