Anticipatory Edge Intelligence: A Foundational Enabler for Resilient Critical Infrastructure Systems

Shruti Lall and Nelishia Pillay

University of Pretoria, Pretoria, 0002, South Africa

ABSTRACT

Critical infrastructure systems (CISs), such as power grids, water networks, and transportation systems, operate under stringent requirements for timeliness, resilience, and coordinated response. As these systems become increasingly data-driven and automated, decisions must often be made under uncertainty and with limited tolerance for delay. This position article advocates for anticipatory edge intelligence (AEI) as a conceptual framing for designing edge-enabled intelligence in CISs, with resilience and containment as primary objectives. AEI emphasizes the generation, exchange, and operationalization of short-horizon anticipatory information at the edge to enable coordinated, preemptive action before degradation propagates. The article examines key challenges faced by CISs, identifies opportunities where anticipatory coordination can enhance system-level resilience, and uses an illustrative scenario to motivate this perspective. By articulating AEI as a research and design lens, this work aims to guide future investigation into resilient, edge-enabled infrastructure systems.

Keywords: Edge AI, Resilience, Degradation, Delays, Critical Infrastructure, Cloud Computing, Real Time Systems, Decision Making, Computer Architecture, Edge Computing, Infrastructure Systems, Critical Infrastructure, Critical Infrastructure Systems, Modernity, Transport System, Internet Of Things, Power Grid, Traffic Congestion, Tight Coupling, Water Network, Pre-Emptive Action, Control System, Disaster, Forecasting, Physical System, Online Learning, Noisy Data, Learning Mechanisms, Smart Grid, Graph Neural Networks, Cascading Failures, Edge Nodes, Deep Uncertainty, Digital Twin, Static System, Multi Agent Reinforcement Learning, Cyber Physical Systems, Interdependent Components, Disaster Scenarios, Flight Path

Critical infrastructure systems (CISs) quietly sustain modern life. Their operation is often taken for granted until it falters. A brief power fluctuation cascades into communication outages, a delayed traffic signal compounds congestion during an evacuation, or a missed early warning escalates into a system-wide emergency. In such settings, timing is not merely a performance metric; it is the difference between containment and cascade.

As CISs are modernized through cyber-physical technologies and dense Internet of Things sensing, they generate unprecedented volumes of real-time data and increasingly rely on intelligent, automated decision making. Yet these systems must operate under profound uncertainty: Sensing is imperfect, connectivity is fragile, and disturbances evolve faster than centralized coordination can often respond. The challenge is not only to react quickly, but to act early enough to prevent escalation across tightly coupled subsystems.

Edge artificial intelligence (Edge AI) has emerged as a natural response to these pressures, bringing learning-based perception, detection, and control closer to physical processes by deploying models on embedded controllers, gateways, and sensor nodes. By reducing reliance on cloud connectivity, Edge AI enables faster local responses and greater autonomy. However, most existing deployments remain oriented around reacting to observed conditions or using short-term prediction to optimize local objectives, such as latency, energy, or resource utilization. These approaches work well in many domains, but their underlying assumptions are strained in critical infrastructure settings.

In CISs, actionable signs of degradation may be weak, ambiguous, or spatially distributed until escalation is already underway. By the time thresholds are crossed or faults are detected, opportunities for effective intervention may have passed. What is needed in such systems is not simply faster reaction or more accurate prediction, but a shift in perspective: from intelligence that responds to the present to intelligence that anticipates near-future risk and coordinates action before degradation becomes visible. Although Edge AI has already been applied to many individual CIS components, these efforts largely treat infrastructure systems as application domains for existing techniques, rather than rethinking how edge intelligence itself should be designed when cascading risk, coordination timing, and resilience are dominant concerns.

This article advocates for anticipatory edge intelligence (AEI) as a lens for rethinking how edge intelligence is designed and evaluated in CISs. We use the term anticipatory deliberately to emphasize a focus on foresight, coordination timing, and early containment, rather than isolated task-level optimization. AEI does not propose a new algorithm or architecture; instead, it reframes edge intelligence around a simple but consequential question: How can anticipatory information be generated, shared, and acted upon early enough to preserve system-level resilience?

The remainder of this article develops this perspective. We examine why conventional Edge AI framings fall short in CIS contexts, articulate the principles of AEI for CISs, and outline open research challenges that must be addressed to realize anticipatory coordination at the edge. Our aim is not to replace existing approaches, but to encourage a shift in how researchers and practitioners think about edge intelligence when resilience, rather than efficiency, is the dominant concern.

OVERVIEW

To ground this perspective more concretely, we now step back from motivating examples and articulate the system assumptions, conceptual distinctions, and design implications that underlie AEI for CISs.

Critical Infrastructure Systems

CISs comprise the physical and cyberphysical systems essential to societal stability, including energy grids, water and wastewater networks, transportation systems, emergency services, health care, communications, and financial infrastructure. Their uninterrupted operation underpins public safety, economic continuity, and national security. Despite differences across regions and sectors, CISs share several defining characteristics: They are foundational to daily life, highly interdependent, and increasingly digitized. Once largely standalone assets, modern CISs now operate as complex cyberphysical systems augmented with pervasive sensing, communication, and automated control. While these capabilities enable more efficient and adaptive operation, they also introduce new vulnerabilities, including cyberthreats, data overload, and fragility arising from tight interdependencies. In such environments, localized disturbances can propagate rapidly across subsystems, leading to cascading failures with potentially irreversible consequences.[1]

Why Conventional Edge AI Framing Is Misaligned With CISs

As intelligence increasingly moves closer to physical processes, Edge AI has become a common design paradigm for deploying learning-based perception, detection, and control on embedded devices, gateways, and edge servers. In this article, we treat Edge AI as a broad umbrella encompassing predictive, distributed, and multiagent approaches that perform inference near data sources and actuators.

Despite this breadth, the dominant framing of Edge AI remains centered on reactive or locally predictive decision making.[2] Prediction—when employed—is typically used to optimize task-level objectives, such as offloading efficiency, scheduling, caching, or resource allocation. While effective in many application domains, this framing implicitly assumes that disruptions are localized, system behavior degrades gradually, and coordination can be initiated after degradation becomes observable.

While Edge AI techniques have been applied to many individual CIS components, their dominant framing remains shaped by assumptions that do not consistently hold at the system level. In critical infrastructure settings, tight coupling, cross-domain interdependencies, and asymmetric consequences of delay mean that design choices appropriate for task-level or performance-driven edge applications can become brittle when scaled to infrastructure-wide operation. Table 1 distills this misalignment by contrasting common Edge AI evaluation and coordination lenses with the operational realities of CISs. The resulting gap is not one of missing techniques, but of how anticipatory information is interpreted, shared, and acted upon when resilience and containment are the overriding objectives.

TABLE 1. Framing mismatch between dominant Edge AI design and evaluation lenses and the operational realities of CISs.

Dominant Edge AI Framing	CIS Operational Reality	AEI Reframing
Edge AI formulations and evaluations emphasize task- or component-level performance objectives.	Infrastructure-wide operation is constrained by continuity, containment, and safety requirements.	Elevate infrastructure-wide objectives alongside, or over, isolated efficiency metrics.
Prediction is primarily used to optimize local or component-level decisions.	Early indicators of degradation are often weak, distributed, or ambiguous until escalation.	Use anticipatory signals as shared coordination inputs for early intervention.
Coordination is typically triggered by fault detection, alarms, or threshold violations.	Thresholds can be late, context-dependent, or misleading in tightly coupled systems.	Initiate coordinated preemptive actions prior to observable degradation.
Evaluation practice emphasizes average-case performance and localized fault models.	Failures can propagate across interdependent subsystems with cascading consequences.	Favor conservative, preemptive strategies to limit escalation.
Problem formulations often assume relatively stable operating regimes and objectives.	Operational priorities and constraints can shift rapidly during disruptions.	Incorporate uncertainty-aware, bounded decision making with safe fallbacks.

Anticipatory Edge Intelligence

The misalignment identified above motivates a reframing of edge intelligence that explicitly accounts for anticipation, coordination timing, and the collective behavior of interdependent components in CISs. We refer to this reframing as AEI. AEI is not a new algorithmic technique, but a coordination-centric design perspective that specifies how predictive information should be generated, exchanged, and acted upon at the edge to support resilience and containment, rather than task-level or average-case efficiency.

Under AEI, edge-resident components reason not only about current observations, but about near-future risks and expected degradation, and explicitly share such anticipatory signals with neighboring nodes. The purpose of this exchange is not to improve isolated local decisions, but to align pre-emptive actions across interdependent components before observable degradation propagates. In this sense, AEI shifts the role of prediction from a local optimization aid to a coordination mechanism that shapes when and how responses are initiated.

To avoid ambiguity, AEI should be understood as a framing rather than a replacement for existing paradigms. It does not subsume predictive or proactive Edge AI, multiagent learning, or digital twin frameworks, nor does it prescribe specific learning models or control policies. Instead, AEI highlights a change in emphasis: from optimizing individual components to reasoning explicitly about anticipation, alignment, and action timing under uncertainty in tightly coupled systems.

AEI for CISs

In CIS settings, edge intelligence should be approached from a resilience-by-design perspective that begins with: 1) system-level objectives centered on continuity and containment rather than average-case efficiency, 2) anticipatory signals that capture short-horizon risk and expected degradation under uncertainty, and 3) coordinated pre-emptive actions aligned across interdependent edge nodes before observable degradation emerges. AEI articulates this starting point and organizes the associated research challenges in prediction, coordination, and assurance under a common framework. To illustrate how such anticipatory coordination may be operationalized, Figure 1 depicts a representative three-tier architecture linking device, edge, and cloud intelligence. This architecture is not proposed as a novel system design, but as an illustrative instantiation of how AEI principles can be realized in practice:

❯ The device layer collects high-frequency sensor data and executes rapid control actions.

❯ The edge layer performs short-horizon prediction, risk assessment, and peer coordination by exchanging compact foresight summaries.

❯ The cloud layer aggregates longer-term patterns, retrains models, and disseminates updated parameters to the edge.

Information flows upward as summarized forecasts, while refined models and policies flow downward, forming a predictive–adaptive loop that links edge immediacy with cloud-level awareness. Together, these tiers exemplify AEI's core principle of anticipatory coordination and provide a conceptual foundation for future research on synchronization, verification, and lightweight learning mechanisms for resilient CIS operation.

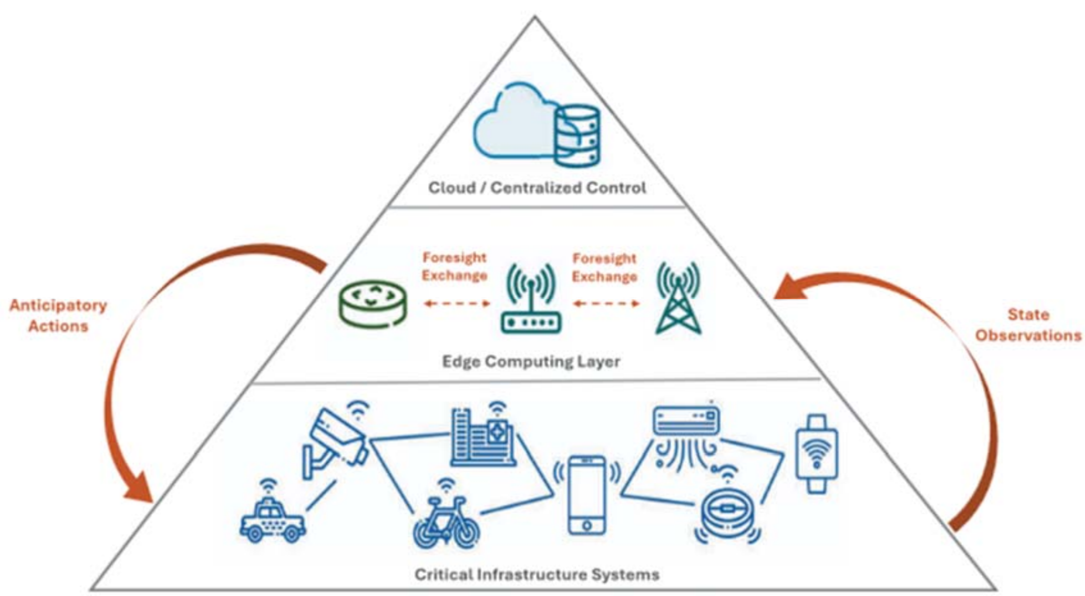

FIGURE 1. Illustrative three-tier architecture for AEI in CISs, showing the bidirectional flow of data and foresight.

CHALLENGES FACING MODERN CISs

CISs are essential yet increasingly vulnerable due to rising complexity, interdependence, and digital transformation. This section outlines key operational, technical, and systemic challenges to their resilience.

Challenge 1: Time Sensitivity and High Stakes

Many CIS services operate under stringent real-time constraints. In electrical grids, delays of even a few seconds in detecting voltage instabilities or frequency deviations can result in cascading blackouts.[3] In health care, wearable monitors and hospital sensors must flag anomalies, such as oxygen drops or arrhythmias, within milliseconds to enable life-saving interventions. Similarly, real-time decision making is essential in traffic control systems to prevent congestion or accidents during peak hours. The challenge is not only the speed of detection, but also the latency of response. Many existing systems rely on centralized data collection and decision pipelines, which introduce delays due to backhaul communication and remote processing. When seconds matter, reliance on faraway data centers can be a systemic vulnerability.

Challenge 2: Fragmentation and Heterogeneity

CISs are inherently multivendor, multiprotocol environments, composed of diverse subsystems installed over decades and rarely designed with interoperability in mind. A single city's infrastructure might include legacy supervisory control and data acquisition systems in water management, proprietary telemetry in electricity distribution, and GPS-based traffic sensors in transportation, each with different formats, latencies, and reliability assumptions.[4] This heterogeneity creates silos that inhibit shared awareness and coordinated control. In disaster scenarios and modernization efforts alike, fragmented systems have led to misaligned responses, noninteroperable components, and even service disruptions when devices from different vendors fail to communicate. Even when systems are digitized, data cannot be easily exchanged or fused to inform unified decisions. Fragmentation also increases maintenance overhead and complicates modernization efforts, particularly when integrating newer intelligent components with legacy infrastructure.

Challenge 3: Data Gaps and Inconsistencies

Despite growing sensor deployment, many CIS domains still suffer from sparse, noisy, or low-resolution data.[5] In rural or underserved areas, sensor coverage may be minimal or unreliable. Legacy systems often lack telemetry altogether, relying on manual logging or batch uploads that introduce delays and inconsistencies. Additionally, data quality is not always guaranteed. Sensors may drift out of calibration, communication links may drop intermittently, and recorded values may include missing or erroneous entries. These issues are not just theoretical: During the early phases of COVID-19, many public health and hospital systems struggled to track intensive care unit availability and oxygen supplies in real time due to fragmented and manually reported data.[6] Similarly, in water networks, inaccurate or infrequent flow and pressure measurements can delay leak detection or trigger false alarms in anomaly detection systems. Such gaps undermine confidence in analytics pipelines and hinder early-warning systems, which depend on clean, continuous data to detect anomalies and predict future states.

Challenge 4: Static Systems in a Dynamic World

CISs are increasingly affected by changing external conditions: climate variability, urban growth, population shifts, energy transitions, and evolving policy mandates. However, many CIS control systems still operate based on fixed rules, predefined thresholds, or models trained on historical data.

As a result, they struggle to adapt when patterns shift or new stresses emerge. For example, during the 2021 Texas winter storm, rigid electricity market rules and demand forecasting models failed to anticipate the scale of residential heating loads under extreme cold, contributing to widespread blackouts.[7] Similarly, energy forecasting tools may underperform as electric vehicle adoption and rooftop solar generation accelerate, introducing bidirectional power flows and demand volatility that were not considered in legacy grid models. The inability to adapt undermines responsiveness and can lead to degraded service or safety margins.

Challenge 5: Limited Autonomy and Centralized Dependence

Traditional infrastructure management often relies on centralized control rooms and human-in-the-loop decision making. While this structure allows oversight, it introduces bottlenecks, especially when systems are geographically dispersed or network connectivity is limited. In disaster scenarios, centralized control can become a single point of failure: Even when local infrastructure components remain operational, they may be rendered inactive due to their dependence on remote coordination. For example, substations or distribution nodes that could function independently may instead default to passive modes because they lack local intelligence or authority to act.[8] Without autonomous control at the edge, actionable data go unused and critical services remain disrupted, limiting both resilience and operational continuity.

Challenge 6: Expanding Security and Trust Risks in AEI

As CISs evolve toward distributed intelligence, their security and trust surface expands. Legacy devices often remain unprotected, while AEI introduces vulnerabilities tied to autonomous foresight-driven agents. Local models may be corrupted through adversarial or data-poisoning attacks or diverge due to bias and inconsistent training data, producing conflicting predictions. Manipulating a single node's foresight can cascade across peers, amplifying disruption. Robust coordination and verification are essential: Agents must cross-validate forecasts, detect anomalies in shared models, and sustain consensus under uncertainty, integrating cybersecurity with AI-aware resilience in distributed infrastructures.

Challenge 7: Rising Expectations for Resilience and Sustainability

Beyond operational safety, CISs are now expected to deliver environmental sustainability, cost-efficiency, and continuous service, even in the face of unpredictable events. As urban populations grow and infrastructure ages, service providers must navigate increasingly complex performance tradeoffs: how to meet rising demand, reduce emissions, manage distributed energy sources, and prepare for extreme weather events. These pressures are no longer hypothetical: Prolonged heatwaves in Europe have strained electricity grids due to concurrent spikes in cooling demand and reduced renewable generation,[9] while record flooding events have overwhelmed stormwater and wastewater systems in cities, like New York and Jakarta.[10] Meeting these expectations requires more than incremental improvements. It calls for infrastructure that is adaptive, proactive, and capable of operating under uncertainty. Current systems, designed for stability and predictability, are ill-equipped for the emerging era of complexity and change.

Challenge 8: Latency–Computation Tradeoffs at the Edge

Edge intelligence is often motivated by reduced communication latency, yet in CIS deployments end-to-end decision time is shaped by more than data transmission alone. Sensing, feature extraction, inference, communication, and actuation all contribute to delay, and on resource-constrained edge devices the cost of nontrivial predictive models can rival or exceed network latency, particularly under bursty workloads or when competing with safety-critical processes. These constraints weaken the assumption that moving intelligence closer to the physical system inherently yields faster or more reliable decisions. In tightly coupled infrastructures, reacting once degradation becomes observable may already be too late, regardless of where inference is performed. This tension highlights a core challenge for AEI: Anticipation must be valuable not because it is faster, but because it is earlier. Exploiting this temporal shift raises open questions around when prediction creates decision slack, how to bound inference cost using lightweight or approximate models, and how to schedule anticipatory computation so its overhead is justified by the benefit of pre-emptive, coordinated action.

Challenge 9: Out-of-Distribution Events and Nonstationarity in Disasters

Many AEI use cases arise during extreme or rare events, such as natural disasters, cyber-physical attacks, or cascading infrastructure failures, which are inherently out-of-distribution relative to historical training data. In such regimes, learning-based models may extrapolate unreliably, and online adaptation can become unstable as system dynamics and objectives shift rapidly. Errors in anticipatory prediction may further propagate across interdependent components, amplifying rather than mitigating risk. AEI therefore cannot assume that online learning will converge or remain robust under disruption. Instead, it emphasizes mechanisms for safe operation under deep uncertainty, including uncertainty-aware prediction, bounded adaptation, explicit confidence thresholds for anticipatory action, fallback to reactive control, and runtime assurance to limit the impact of incorrect foresight.

RESEARCH OPPORTUNITIES IN AEI FOR CISs

While CISs face mounting challenges, AEI provides a promising foundation for resilience. The following opportunities illustrate how anticipatory intelligence can address current limitations while identifying open research directions.

Opportunity 1: Timely Intervention in Safety-Critical Systems

In domains where milliseconds determine stability or failure, delayed detection or response can escalate into crises. AEI enables predictive intelligence to operate at the edge, supporting near-instantaneous foresight and action. Power grids already use predictive analytics to anticipate voltage collapse and reroute loads before failure,[11] while health-care systems, such as the Technology Integrated Health Management project, detect early signs of deterioration.[12] AEI advances these efforts by relocating decision making from centralized servers to substations, roadside units, or wearable devices, reducing latency and enabling autonomous operation when the cloud is unreachable. Reliable real-time intervention must still address model drift, false positives, and resource constraints. Predictive errors can propagate rapidly across interconnected systems, risking unsafe actions. Future work should focus on model adaptation, context-aware confidence estimation, and verifiable decision logic to ensure anticipatory actions remain accurate and safe.

Opportunity 2: Distributed Fault Detection and Recovery

Centralized control centers for fault detection and recovery are often too slow and vulnerable for large or remote networks. AEI decentralizes this process by enabling nodes to detect and respond locally while coordinating recovery collaboratively in near real time. Examples, such as distributed digital twins and edge analytics in smart cities, show how exchanging state information among nodes supports global awareness.[13] A water node can coordinate with nearby pumps to reroute flow, while a transport hub can pre-empt congestion through predictive scheduling. However, distributed recovery must maintain temporal and causal consistency despite latency and partial observability, limit message exchange to prevent congestion, and preserve trust so compromised nodes cannot mislead peers. Future work should explore lightweight consensus, fault-tolerant coordination, and distributed assurance frameworks to sustain accuracy and trust under uncertainty.

Opportunity 3: Privacy-Aware, Personalized Intelligence

CISs increasingly rely on sensitive data, such as health metrics or household energy use, to deliver adaptive services, raising privacy and ethical concerns. AEI embeds predictive foresight within edge nodes, allowing systems to anticipate needs while keeping raw data local. Forecasting future states from local trends can reduce continuous monitoring; for instance, a smart building can infer occupancy to preadjust ventilation, or a wearable can predict distress signals and respond locally without cloud uploads. However, predictive foresight may expose behavioral traits or amplify bias, and personalized models risk overfitting on small, sensitive datasets. Future work should design privacy-preserving foresight, differential learning, and adaptive anonymization to retain accuracy while ensuring confidentiality and fairness.

Opportunity 4: Robust Operation Under Connectivity Constraints

Many CISs face intermittent connectivity, whether in rural areas or during disasters, such as floods and wildfires. Reactive systems often fail when disconnected, but AEI enhances resilience by embedding predictive foresight within edge nodes so they can operate autonomously. Flood sensors could project river levels and adjust gates, while agricultural drones beyond coverage could forecast microclimate shifts and adapt flight paths. By anticipating short-term risks, isolated subsystems can maintain continuity until cloud access is restored. However, prolonged disconnection may cause model drift, desynchronization, and inconsistent decisions when connectivity resumes. Future work should investigate temporal model stabilization, version-aware synchronization, and self-healing coordination to align local and global states safely.

Opportunity 5: Proactive Resource Optimization

Reactive resource management often leads to inefficiencies, such as electricity overproduction, water waste, and traffic congestion. AEI enables foresight-driven allocation of resources based on predicted demand and risk. Smart grids can forecast consumption and shift loads before peaks; water utilities can anticipate leaks or contamination and schedule maintenance; and transport systems can reroute flows preemptively.[14] This proactive coordination improves efficiency, reduces emissions, and extends infrastructure life. However, dependable optimization at the edge is difficult: Predictive accuracy degrades under sparse or noisy data, and distributed foresight must align across heterogeneous nodes with limited bandwidth and compute. Overconfidence can destabilize control. Future work should pursue uncertainty-aware prediction–control loops, adaptive online learning, and multiagent optimization to sustain stability, efficiency, and safety at scale.

Opportunity 6: Adaptive Response to Environmental and Operational Change

Rapidly evolving conditions—such as climate variability, population growth, aging assets, and regulatory change—make static control systems ineffective. AEI enables adaptive foresight by detecting gradual shifts, such as temperature trends or demand peaks, and adjusting operations proactively. It can balance renewable fluctuations, modify irrigation schedules, or anticipate public transit surges. The rise of wildfire-related precautionary power shutoffs highlights the need for predictive adaptability.[15] Adaptive AEI requires continual learning with bounded updates, drift detection through uncertainty metrics, and runtime verification for stability and compliance. Future work should design safe learning mechanisms that let AEI systems evolve with changing environments while maintaining resilience and trust.

Opportunity 7: Toward Self-Managing, Autonomous Infrastructure

As infrastructures scale, continuous human oversight becomes impractical. AEI provides the foundation for self-monitoring and self-configuring systems capable of autonomous operation. Energy networks can reroute power flows pre-emptively, hospitals can adjust heating ventilation and air conditioning and power provisioning based on occupancy forecasts, and smart water grids can optimize delivery in real time. Autonomous vertical farms operating fully on edge infrastructure already demonstrate the feasibility of decentralized, foresight-driven control.[16] Achieving trustworthy autonomy, however, requires stability, interpretability, and alignment between local actions and global policies. Future work should explore hierarchical autonomy, runtime assurance, explainability, and human-in-the-loop governance to keep AEI systems transparent, accountable, and safe.

Opportunity 8: Anticipation With Bounded Risk and Temporal Slack

AEI reframes how prediction, computation, and control are coupled in time-critical infrastructure systems. Rather than optimizing inference latency at the moment of action, AEI exploits short-horizon anticipation to create temporal slack, enabling decisions to be precomputed, staged, or validated ahead of degradation. This is critical in CISs, where inference cost on resource-constrained edge devices, resource contention, and rapidly evolving conditions can undermine reactive or just-in-time control. Acting earlier enables conservative decision making, cross-node coordination, and safety validation even when predictive models are approximate. Temporal slack can be operationalized through anytime inference and early-exit architectures, which produce coarse predictions quickly and refine them only when time and confidence permit,[17] while runtime monitors or control barriers validate staged actions before execution. Future work should focus on latency-aware anticipation, bounded-complexity models, and scheduling mechanisms that trade computational cost against the value of early, coordinated action.

Opportunity 9: Learning and Coordination Under Deep Uncertainty

In CISs, the scenarios where anticipatory intelligence is most needed—natural disasters, cascading failures, and large-scale disruptions—are rare, high-impact, and fundamentally out-of-distribution relative to historical data. AEI therefore cannot assume reliable generalization, stable convergence, or effective continuous online learning. Instead, it motivates learning and coordination mechanisms that explicitly account for uncertainty, partial correctness, and model failure. While individual events may not repeat, prior CIS disruptions can inform anticipation by revealing structural vulnerabilities, failure propagation pathways, timing constraints, and safe response envelopes. Anticipatory coordination should be paired with explicit confidence estimation, bounded adaptation, fallback to reactive control, and runtime assurance mechanisms that constrain actions when predictive reliability degrades. Future work should focus on uncertainty-aware forecasting, conservative or risk-sensitive learning,[18] verification and runtime assurance, and governance mechanisms that treat deep uncertainty as a first-class design constraint.

ILLUSTRATIVE SCENARIO: AEI FOR DISASTER-RESILIENT COMMUNICATION

In large-scale disasters, such as earthquakes or flash floods, communication infrastructure is often among the first services to fail.[19] Base stations collapse, backhaul links degrade, and coordination erodes as demand peaks. Centralized communication architectures—dependent on stable connectivity and human intervention—offer limited resilience under such conditions. AEI offers an alternative perspective in which distributed nodes anticipate disruption, reorganize proactively, and coordinate through lightweight, decentralized intelligence at the edge. The following scenario, as shown in Figure 2, illustrates how AEI principles may be instantiated in a disaster-response communication network composed of heterogeneous, cooperative agents.

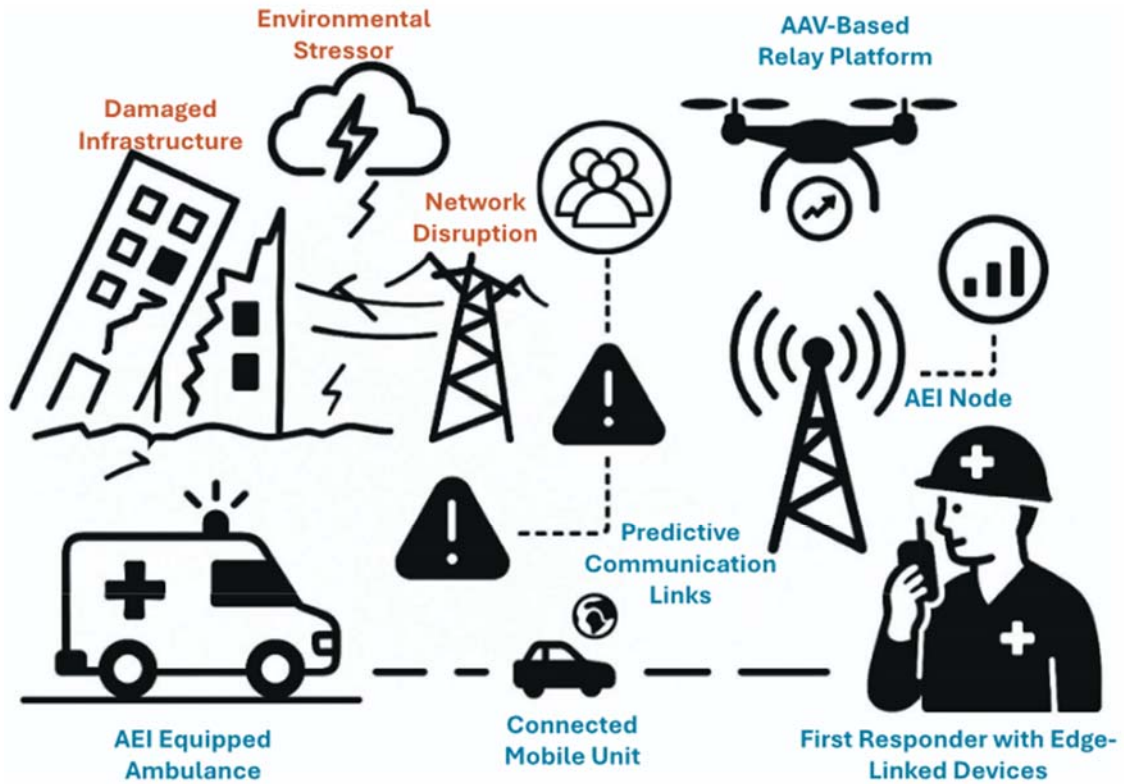

FIGURE 2. AEI-enabled disaster response framework for foresight-driven coordination

Here, AEI operates as a coordinated ecosystem of predictive agents that sustain connectivity and resilience under extreme conditions. Each component contributes a distinct capability within the network.

Unmanned Aerial Vehicle-Based Relay Platforms

In AEI-enabled architectures, autonomous aerial vehicles (AAVs) act as agile relays that restore and extend connectivity when ground infrastructure fails. Each AAV carries lightweight edge computing and predictive models that forecast demand and reposition ahead of anticipated coverage gaps. Flight paths and resource allocation may be adapted using reinforcement learning under multi-objective criteria, such as signal quality, energy consumption, and user priority. AAVs can further anticipate bandwidth surges via contextual bandit models that adjust channels and power dynamically. By caching mission-critical content at predicted hotspots, AEI-enabled AAVs shift from reactive relays to foresight-driven coordinators that support decentralized and resilient communication under stress.

Edge Nodes and First Responders

First responders carry rugged, edge-enabled devices that perform local inference to detect hazards, stress indicators, or network faults without relying on cloud connectivity. These devices exchange compact summaries over short-range links and coordinate through graph-based methods, such as message-passing networks or graph neural networks (GNNs), to maintain situational awareness. Federated and transfer learning enable continual adaptation while preserving data locality. While GNNs support cooperative perception, deployment on constrained devices remains challenging, motivating research into lightweight architectures, model compression, and edge–cloud co-inference to balance scalability, energy efficiency, and responsiveness.

AEI-Equipped Ambulances

Ambulances operate as mobile edge nodes that unify sensing, prediction, and coordination. Real-time telemetry—such as speed, location, and patient vitals—feeds onboard models that forecast travel times and patient condition to support anticipatory triage and routing. Clustering methods, including online DB-SCAN or mobility-aware Voronoi tessellation, identify emergent hotspots or coverage gaps, enabling predictive repositioning and bandwidth planning with AAVs and responders. Compact neural models assess vitals for early escalation, allowing ambulances to function as predictive agents that enhance both medical and communication resilience.

Predictive Communication Links

AEI introduces predictive communication links that anticipate, rather than react to, channel degradation and mobility. Each node—AAV, vehicle, or fixed unit—runs lightweight forecasting using techniques, such as contextual bandits, Gaussian process regression, or meta-learning to estimate link stability and demand. Inspired by emerging 6G work on predictive link control using in-band signal features,[20] AEI adapts these ideas for resource-constrained, latency-critical environments. In this scenario, AAVs anticipate ambulance movement toward low-coverage zones and pre-emptively reserve spectrum or form multihop relays ahead of arrival.

Connected Mobile Units

The connected mobile unit in Figure 2 represents an opportunistic sensing and relay agent embedded in everyday traffic. These vehicles collect data on road blockages, crowd density, and interference, and use local models, such as long short-term memories or convolutional neural networks to forecast disruptions. When static relays fail, mobile units act as temporary gateways coordinated via multiagent learning to improve coverage. By combining sensing, prediction, and relaying, such units contribute mobility-aware foresight to the AEI mesh and enable more adaptive, fine-grained response in dynamic environments.

CONCLUSION

CISs increasingly operate under conditions where delays, partial information, and localized failures can rapidly escalate into system-wide disruption. This position article argued for AEI as a reframing of edge-enabled intelligence for CISs, one that elevates anticipation, coordination timing, and system-level resilience as primary design concerns. Rather than proposing new algorithms or architectures, AEI provides a conceptual lens for reasoning about how short-horizon foresight can be generated, exchanged, and acted upon at the edge to support a pre-emptive, coordinated response under uncertainty. By examining key challenges and articulating emerging opportunities, this work highlights how anticipation can create temporal slack, enable bounded operation under deep uncertainty, and complement existing reactive mechanisms. Advancing AEI will require interdisciplinary research spanning systems, networking, learning, and control, with particular emphasis on assurance, coordination, and safe operation in nonstationary environments.

References

1. S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, “Catastrophic cascade of failures in interdependent networks,” *Nature*, vol. 464, no. 7291, pp. 1025–1028, 2010, doi: 10.1038/nature08932.

2. S. S. Gill et al., “Edge AI: A taxonomy, systematic review and future directions,” *Cluster Comput.*, vol. 28, no. 1, Art. no. 18, 2024, doi: 10.1007/s10586-024-04686-y.
3. L. Dueñas-Osorio and S. M. Vemuru, “Cascading failures in complex infrastructure systems,” *Struct. Saf.*, vol. 31, no. 2, pp. 157–167, 2009, doi: 10.1016/j.strusafe.2008.06.007.
4. B. Ahlgren, M. Hidell, and E. C.-H. Ngai, “Internet of things for smart cities: Interoperability and open data,” *IEEE Internet Comput.*, vol. 20, no. 6, pp. 52–56, Nov./Dec.2016, doi: 10.1109/MIC.2016.124.
5. M. A. Jamshed, K. Ali, Q. H. Abbasi, M. A. Imran, and M. Ur-Rehman, “Challenges, applications, and future of wireless sensors in internet of things: A review,” *IEEE Sensors J.*, vol. 22, no. 6, pp. 5482–5494, Mar.2022, doi: 10.1109/JSEN.2022.3148128.
6. R. Badker et al., “Challenges in reported COVID-19 data: Best practices and recommendations for future epidemics,” *BMJ Global Health*, vol. 6, no. 5, Art. no. e005542, 2021, doi: 10.1136/bmjgh-2021-005542.
7. J. W. Busby et al., “Cascading risks: Understanding the 2021 winter blackout in Texas,” *Energy Res. Social Sci.*, vol. 77, Jul.2021, Art. no. 102106, doi: 10.1016/j.erss.2021.102106.
8. M. Gallucci, “Rebuilding Puerto Rico’s grid,” *IEEE Spectr.*, vol. 55, no. 5, pp. 30–38, May2018, doi: 10.1109/MSPEC.2018.8352572.
9. K. van der Wiel et al., “The influence of weather regimes on European renewable energy production and demand,” *Environ. Res. Lett.*, vol. 14, no. 9, 2019, Art. no. 094010, doi: 10.1088/1748-9326/ab38d3.
10. E. C. O’Donnell and C. R. Thorne, “Drivers of future urban flood risk,” *Philos. Trans. R. Soc. London, Ser. A*, vol. 378, no. 2181, 2020, Art. no. 20190216.
11. R. Eskandarpour and A. Khodaei, “Machine learning based power grid outage prediction in response to extreme events,” *IEEE Trans. Power Syst.*, vol. 32, no. 4, pp. 3315–3316, Jul.2017, doi: 10.1109/TPWRS.2016.2631895.
12. S. Enshaeifar et al., “The internet of things for dementia care,” *IEEE Internet Comput.*, vol. 22, no. 1, pp. 8–17, Jan./Feb.2018, doi: 10.1109/MIC.2018.112102418.
13. B. Ketzler, V. Naserentin, F. Latino, C. Zangelidis, L. Thuvander, and A. Logg, “Digital twins for cities: A state of the art review,” *Built Environ.*, vol. 46, no. 4, pp. 547–573, 2020, doi: 10.2148/benv.46.4.547.
14. S. Kulkarni et al., “Enabling a decentralized smart grid using autonomous edge control devices,” *IEEE Internet Things J.*, vol. 6, no. 5, pp. 7406–7419, Oct.2019, doi: 10.1109/JIOT.2019.2898837.
15. J. W. Muhs, M. Parvania, and M. Shahidehpour, “Wildfire risk mitigation: A paradigm shift in power systems planning and operation,” *IEEE Open Access J. Power Energy*, vol. 7, pp. 366–375, 2020, doi: 10.1109/OAJPE.2020.3030023.
16. A. Jandl, P. A. Frangoudis, and S. Dustdar, “Edge-based autonomous management of vertical farms,” *IEEE Internet Comput.*, vol. 26, no. 1, pp. 68–75, Jan./Feb.2022, doi: 10.1109/MIC.2021.3129271.
17. S. Teerapittayanon, B. McDanel, and H. T. Kung, “BranchyNet: Fast inference via early exiting from deep neural networks,” in *Proc. 23rd Int. Conf. Pattern Recognit. (ICPR)*, 2016, pp. 2464–2469, doi: 10.1109/ICPR.2016.7900006.
18. A. Wachi and Y. Sui, “Safe reinforcement learning in constrained Markov decision processes,” in *Proc. 37th Int. Conf. Mach. Learn. (ICML)*, 2020, pp. 9797–9806.
19. M. Matracia, N. Saeed, M. A. Kishk, and M.-S. Alouini, “Post-disaster communications: Enabling technologies, architectures, and open challenges,” *IEEE Open J. Commun. Soc.*, vol. 3, pp. 1177–1205, 2022, doi: 10.1109/OJCOMS.2022.3192040.

20. S. Wu, C. Chakrabarti, and A. Alkhateeb, “Proactively predicting dynamic 6G link blockages using LiDAR and in-band signatures,” *IEEE Open J. Commun. Soc.*, vol. 4, pp. 392–412, 2023, doi: 10.1109/OJCOMS.2023.3239434.