Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Adaptive Sensor Sampling and Data Center Offloading for Scalable Smart City GIS Models

[1]Mohamed Fahad A, [2]Housem Daaji, [3]Mohamed Elawadi Elsayed Eid*, [4]Amr Ezzat Abdou Shallal

[1,2,3,4]*IT-District Technology Management, KAFD Development & Management Company, Riyadh, Saudi Arabia*

ARTICLE INFO

Received: 02 Oct 2024

Revised: 17 Nov 2024

Accepted: 29 Nov 2024

ABSTRACT

Smart city geographic information system (GIS) platforms depend on dense networks of heterogeneous sensors, including traffic loops, CCTV, air-quality nodes, water-level gauges, and mobile crowd-sensed streams. These sensors keep a city's spatial digital twin continuously updated. But sending every raw reading to a central data center does not scale. It saturates uplink bandwidth. It inflates data center compute and storage costs. And it adds latency that undermines time-critical applications such as incident response and adaptive traffic control. This paper proposes an artificial intelligence (AI) driven framework for adaptive sensor sampling and edge-to-data-center offloading, guided explicitly by GIS context. The framework splits decision-making across three tiers: an edge sensing tier, a GIS context and policy tier, and a data center analytics tier. It formulates the sampling and offloading problem as a context-conditioned Markov decision process (MDP), in which spatial priors, such as zoning, points-of-interest density, historical incident maps, and proximity to critical infrastructure, shape the sampling rate, feature compression level, and offload priority assigned to each sensor stream. A reinforcement-learning (RL) policy is trained to maximize a reward that jointly captures event-detection accuracy, communication cost, energy consumption, and data-center processing load. We evaluate the framework through simulation across four canonical GIS zone types: residential, commercial, transit hub, and incident zone. The proposed GIS-context policy cuts uplink data volume by up to 82% relative to fixed-rate sampling. It keeps event-detection accuracy above 94% in the 55–75% bandwidth-reduction operating region, outperforming both static and threshold-based adaptive baselines. These results show that spatial context is a first-class signal for scalable smart-city sensing pipelines, not merely an annotation added after data collection.

Keywords: Edge computing, adaptive sampling, computation offloading, geographic information systems (GIS), smart cities, reinforcement learning, digital twin, IoT sensor networks.

I. INTRODUCTION

The smart city paradigm relies on the continuous fusion of sensor observations into a spatially referenced model of urban state, covering traffic flow, air quality, structural health, water infrastructure, and public safety events. This model is commonly realized as a GIS-based digital twin [1], [2]. The Internet of Things (IoT) has made ubiquitous urban sensing technically feasible. Municipalities now deploy tens of thousands of heterogeneous devices, spanning fixed environmental sensors, traffic cameras, vehicle-mounted probes, and citizen-carried mobile sensors [3], [4]. But the volume, velocity, and heterogeneity of this data create a fundamental scalability bottleneck. In a cloud-centric architecture, every sensor stream is transmitted continuously to a centralized data center for processing. This quickly exhausts wireless uplink capacity. It incurs unsustainable data-center ingestion and storage costs. And it introduces round-trip latency that is incompatible with applications such as adaptive signal control, structural failure warnings, or emergency dispatch, all of which require sub-second response [5], [6].

Edge computing has emerged as the main mitigation strategy. It relocates a portion of computation from the data center to devices and micro-servers physically close to the sensors [5], [7]. A large body of work has examined computation-offloading decisions from a communications and systems perspective, optimizing for latency, energy, and network cost under queueing, wireless channel, or task-graph models [8]– [11]. A separate body of work has examined adaptive or event-driven sampling. Here, sensors dynamically vary their acquisition rate based on signal

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

variance, energy budget, or learned reward signals, rather than sampling at a fixed rate [12]– [15]. What remains comparatively underexplored is the systematic use of geospatial context as an explicit input to the sampling and offloading decision itself. This is the same GIS context that the smart city platform is ultimately trying to update. A sensor at a residential intersection with low historical incident frequency does not carry the same information value per byte as a sensor at a transit hub during peak hours, or one inside an active incident zone. Treating all sensors identically wastes bandwidth on the former and risks under-sampling the latter.

This paper addresses that gap directly. We treat GIS context, including zoning classification, point-of-interest (POI) density, proximity to critical infrastructure, historical event frequency, and time-of-day/land-use interactions, as a first-class state variable in the sampling and offloading policy. This is different from treating it as a downstream annotation applied to already-collected data. Our contributions are threefold. First, we present a three-tier reference architecture (edge, GIS-context/policy, data center) that cleanly separates local sensing and inference from spatially informed decision-making, and from heavyweight city-scale analytics. Second, we formulate the joint sampling-rate and offload-priority decision as a context-conditioned MDP, and describe a reinforcement-learning policy whose reward balances event-detection accuracy, bandwidth cost, energy consumption, and data-center load. Third, we provide a simulation-based evaluation across representative GIS zone types, along with a bandwidth-accuracy trade-off analysis. This shows that GIS-aware adaptive offloading substantially outperforms both fixed-rate and generic threshold-based adaptive baselines, particularly in the operating region most relevant to municipal deployments (55–75% bandwidth reduction).

II. RELATED WORK

A. Edge and Cloud Computing for Smart Cities

Edge computing has been widely proposed as the architectural response to the latency, bandwidth, and privacy limitations of purely cloud-centric IoT systems [5], [7], [16]. Khan et al. [1] provide a comprehensive taxonomy of edge-computing-enabled smart city applications, spanning transportation, healthcare, and environmental monitoring. They identify resource heterogeneity and dynamic workload variability as persistent open challenges. Surveys on computation offloading [8], [9] classify offloading decisions along three axes: what to offload, where to offload it (edge, fog, or cloud), and when to offload. These are typically formulated as optimization problems over latency, energy, and monetary cost. Gharaibeh et al. [17] and Silva et al. [18] frame smart-city data management challenges more broadly. They note that heterogeneous, high-velocity sensor data must be triaged before it reaches back-end analytics, if city-scale systems are to remain economically viable. A more recent survey of the IoT-edge-cloud continuum [19] formalizes hybrid execution strategies, in which latency-sensitive tasks are handled at the edge while heavier analytics are offloaded to the cloud, echoing the tiered design adopted in this paper.

B. Adaptive and Context-Aware Sampling

A parallel line of research focuses on adaptive sampling, aiming to reduce the volume of data a sensor or sensor network needs to transmit in the first place. Early approaches used threshold-based sleep scheduling, exploiting spatiotemporal correlations between neighboring sensors to skip redundant measurements [20]. More recent work applies reinforcement learning directly to the sampling decision: Di Francesco Maesa et al. and related studies formulate information-driven adaptive sensing as an RL problem, in which an agent learns when to sample so as to maximize reconstruction accuracy under an energy or bandwidth budget [12]. Similar RL-based approaches have been used for sampling-parameter tuning in environmental IoT monitoring [13], for context caching in IoT platforms using actor-critic methods [14], and for AI/RL-based coverage optimization in wireless sensor networks [21]. Ghosh et al.'s UCB-based adaptive sensor-selection framework balances sensing quality against energy cost, while predicting inactive measurements from correlated active sensors. This shows that not every sensor needs to be queried at every timestep for accurate state reconstruction [22]. These works establish that adaptive sampling can substantially reduce data volume with limited accuracy loss. But the state representation driving the sampling decision is typically limited to signal statistics, such as variance and rate of change, or energy budget, without explicit spatial/GIS context.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

C. GIS Integration with IoT and Edge Systems

GIS has long served as the integrative substrate for urban sensor data, providing spatial referencing, visualization, and cross-domain analysis [23], [24]. Kamilari and Frank [23] discuss geospatial analysis as a natural complement to IoT data streams. Armstrong et al. [24] examine the challenges of coupling fast IoT data streams with geospatial data science pipelines. They note that most existing pipelines treat GIS as a post-hoc visualization or query layer, rather than as an input to real-time data-collection decisions. More recent work moves in this direction from the GIS side: deep reinforcement learning has been applied directly to spatial urban planning, generating and refining land-use and infrastructure layouts from geographic data [25], illustrating the growing appetite for RL-driven, spatially aware decision-making in urban systems, even though this line of work still operates on static geographic data rather than live IoT sensor streams. Spatio-temporal Bayesian learning has also been used for edge-resource planning in smart cities, using historical spatial and temporal patterns to anticipate where edge compute demand will concentrate [26]. Our work extends this direction. We use GIS layers not only to plan resource placement or visualize outcomes, but as a direct, low-latency input to the per-sensor sampling-rate and offload-priority decision at the edge.

D. Positioning of This Work

Relative to the surveyed literature, the framework proposed here differs in three respects. First, it treats static GIS layers (zoning, POI density, critical-infrastructure proximity) and dynamic GIS-derived signals (recent incident density, live congestion) as structured state features that feed directly into the sampling/offloading policy, rather than as metadata attached after ingestion. Second, it couples the sampling decision and the offload decision within a single policy, rather than treating them as sequential, independently optimized stages. This lets the system jointly trade off local sensing frequency against uplink volume. Third, it evaluates the policy specifically against zone-typed operating conditions that mirror realistic municipal deployments, namely residential, commercial, transit hub, and incident zone, rather than against a single aggregate workload.

III. PROPOSED SYSTEM ARCHITECTURE

Fig. 1 depicts the proposed three-tier architecture, consisting of an Edge Tier, a GIS Context and Policy Tier, and a Data Center Tier, connected by bidirectional feedback loops.

A. Edge Tier

Heterogeneous urban sensors, including CCTV, LiDAR, air-quality nodes, traffic counters, and water-level gauges, feed an edge gateway or micro-server co-located within the same street segment or district. The edge gateway performs lightweight local inference, such as object counting, simple anomaly scoring, and background-model differencing, to produce compact feature summaries rather than transmitting raw streams. A local buffer and event detector maintains short-horizon history and computes a change/anomaly score. This score triggers local actuation, such as traffic-signal adjustment or alerting, and tells the upstream sampling policy whether the current observation departs from the expected baseline for that location.

B. GIS Context and Policy Tier

This tier is the central contribution of the proposed framework. A GIS context engine maintains static spatial layers, such as zoning classification, road network, digital elevation model, and POI density, together with dynamic layers such as recent incident density, live congestion, and weather advisories. It exposes a compact context vector for each sensor location. This context vector, together with the edge-computed feature summary and event score, forms the state input to an adaptive sampling and offload policy. The policy is implemented as a reinforcement-learning agent operating over a context-conditioned MDP, detailed in Section IV. It outputs three things: (i) a target sampling rate for each sensor stream, (ii) a compression/feature-extraction level, and (iii) an offload priority, which a downstream task-prioritization and bandwidth scheduler uses to arbitrate among competing streams under a shared uplink budget. A feedback/reward module closes the loop. It computes accuracy, latency, energy, and cost signals, which are used both to update the policy and to push revised sampling thresholds back down to the edge tier.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

C. Data Center Tier

Only the samples selected by the policy are transmitted to the data center. There, an ingestion and stream-processing layer performs deduplication, format normalization, and quality checks before feeding a deep-analytics stage that updates the city-scale GIS digital twin. Processed and raw-selected data are persisted in a spatial data warehouse. City dashboards and planning APIs expose the updated model to downstream consumers, such as traffic-management centers or emergency-response coordination systems. Aggregate performance metrics computed at this tier, such as model drift and missed-event rate, are returned to the GIS context/policy tier as part of the reward signal. This enables continual policy refinement, without requiring the data center to micromanage individual sensor sampling rates.

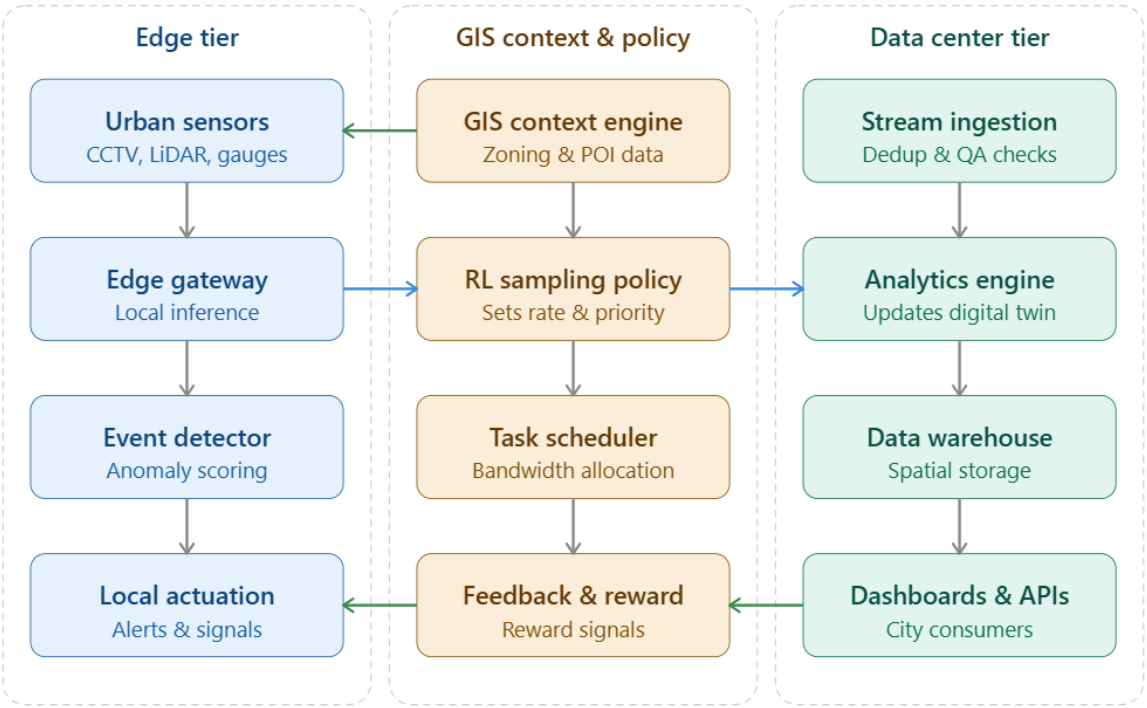


Fig. 1. Proposed three-tier architecture for GIS-guided adaptive sensor sampling and offloading, showing the Edge Tier, the GIS Context and Policy Tier, and the Data Center Tier connected by forward (feature/context/offload) and feedback (reward/policy-update) paths.

Interpretation: Raw sensor data is reduced to compact feature summaries at the Edge Tier, before it ever reaches the network. The GIS Context and Policy Tier fuse these summaries with static and dynamic spatial layers to produce a sampling rate and offload priority per stream, which the Data Center Tier consumes selectively. Two feedback paths keep the system adaptive, rather than relying on a one-time, hand-tuned configuration: context labels flow down to the edge, and reward/updated policy weights flow back from the data center.

IV. METHODOLOGY

A. Problem Formulation

We formulate the joint sampling-rate and offload-priority decision as a context-conditioned Markov decision process (MDP), defined by the tuple (S, A, P, R, γ). For each sensor i at discrete decision epoch t, the state $s_{i,t}$ has three groups of features. The first is local signal features: recent variance, rate of change, and the edge-computed anomaly score. The second is GIS-context features: zone class, POI density, distance to critical infrastructure, and historical incident frequency for the surrounding cell. The third is system features: current uplink utilization, remaining energy budget, and data-center backlog. The action $a_{i,t}$ is a discrete choice over a small set of operating points. Each one specifies a sampling interval, a feature-compression level, and an offload-priority tier: defer, edge-only, low-priority uplink, or high-priority uplink. The transition function P captures how the environment evolves, including traffic state, incident

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

occurrence, and sensor battery level, given the chosen sampling behavior. It is approximated during training via a discrete-event simulator, seeded with historical municipal sensor logs and synthetic incident injection.

B. Reward Function

The reward at each decision epoch combines four terms, as given in (1):

$$R_t = w_1 \cdot Acc_t - w_2 \cdot BW_t - w_3 \cdot E_t - w_4 \cdot L_t \quad (1)$$

where Acc_t is an estimate of event-detection accuracy achievable from the data actually collected. It is approximated by comparing the sub-sampled reconstruction against a full-rate reference during training. BW_t is the uplink bandwidth consumed, E_t is the energy expended by the sensor and edge gateway, and L_t is the additional processing load imposed on the data center. The weights w_1–w_4 are tunable per deployment. In our evaluation, we set them to reflect a municipal priority ordering, in which accuracy in high-criticality zones is weighted roughly twice as heavily as in low-criticality zones. We implement this by scaling w_1 as a function of the GIS zone-criticality feature, rather than using a single global weight. This design lets the same policy network behave conservatively, sampling near full rate, in an active incident zone, while behaving aggressively, with heavy subsampling, in a quiet residential zone, without requiring separate policies per zone type.

C. Policy Learning

The policy $\pi(a|s)$ is represented by a compact neural network with two hidden layers and ReLU activations. It is trained using a proximal policy optimization (PPO)-style actor-critic update, chosen for its stability under the non-stationary reward landscape induced by time-varying incident activity. Training proceeds episodically over simulated 24-hour city cycles. The environment replays historical sensor traces annotated with GIS zone labels, and injects synthetic incidents at rates calibrated to observed municipal incident logs for each zone type. To keep inference feasible on constrained edge/micro-server hardware, the trained policy is distilled into a lightweight lookup-table/decision-tree approximation for deployment, following the general approach used in prior lightweight adaptive-sensing work [12]–[14]. Meanwhile, the full neural policy continues to be retrained periodically at the data center tier, using aggregated feedback.

D. Offload Scheduling

Multiple sensors within a district compete for a shared uplink budget. The offload-priority tiers output by the per-sensor policy are arbitrated by a bandwidth scheduler, using a weighted fair-queueing discipline. High-priority streams, whether incident-zone or anomaly-triggered, preempt low-priority background streams up to a configurable ceiling, preventing a single noisy sensor from starving the district's uplink capacity. This two-stage design, a per-sensor policy followed by district-level scheduling, mirrors the general edge/cloud hybrid execution strategies described in prior work [19]. But it replaces manually tuned offloading thresholds with a learned, GIS-conditioned policy.

V. EXPERIMENTAL SETUP

We evaluate the proposed framework using a discrete-event simulation representing a synthetic mid-sized city district, partitioned into four canonical GIS zone types: residential (low criticality), commercial (medium criticality), transit hub (high criticality), and a designated incident zone representing areas with elevated historical emergency-event density. Each zone hosts a mix of sensor types, namely traffic counters, air-quality nodes, and CCTV-derived event detectors, sampled at a nominal full rate of 1 Hz for CCTV-derived features, and 0.1 Hz for slower-changing environmental sensors. We compare three sampling/offloading policies: (i) a static fixed-rate baseline that always transmits at the nominal full rate; (ii) a threshold-based adaptive baseline that adjusts sampling rate solely from local signal variance, representative of conventional adaptive-sampling approaches [12], [20]; and (iii) the proposed GIS-context RL policy described in Section IV. Performance is measured along two axes: normalized uplink data volume relative to the fixed-rate baseline, and event-detection accuracy. The latter is computed as the fraction of injected synthetic incidents correctly flagged within a fixed detection-latency window, when analytics are run only on the data each policy actually transmits.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

VI. RESULTS AND DISCUSSION

Fig. 2(a) reports the normalized uplink data volume produced by each policy across the four GIS zone types. The static fixed-rate baseline transmits at 100% volume in every zone by construction, since it does not condition on context. The threshold-based adaptive baseline reduces volume substantially in low-activity zones, reaching 42% of full rate in residential areas, but continues to transmit at a high rate in high-criticality zones (95% in the incident zone). This is because local signal variance alone is a weak proxy for the value of that data to the downstream GIS model, when the surrounding context is already known to be volatile. The proposed GIS-context RL policy achieves the largest reduction in low-criticality zones, reaching 18% of full rate in residential areas, an 82% reduction relative to the fixed-rate baseline, while intentionally preserving near-full sampling in the incident zone (98%). This matches the intuitive requirement that data collection intensity should scale with spatial criticality, rather than with local signal statistics alone.

Fig. 2(b) examines the accuracy–bandwidth trade-off directly. We sweep a global aggressiveness parameter for both the uniform/random subsampling baseline and the GIS-context adaptive policy, and plot the resulting event-detection accuracy against the corresponding bandwidth reduction achieved. At modest bandwidth reductions (below 30%), the two curves are close, since neither policy needs to make difficult trade-offs. As bandwidth reduction increases into the 55–75% range, the region most representative of realistic municipal cost constraints, the two curves diverge. The GIS-context policy retains above 94% accuracy throughout this region, while uniform subsampling degrades to approximately 91–93% accuracy for the same bandwidth savings. The gap widens further at the most aggressive reduction levels (90%), where the GIS-context policy retains roughly 94.5% accuracy, compared with approximately 87.5% for uniform subsampling. This divergence matches the underlying mechanism: uniform subsampling discards information indiscriminately across zones, whereas the GIS-context policy concentrates the available uplink budget on locations where GIS priors indicate that events are more likely or more consequential, at the cost of coarser sampling in zones where the marginal value of additional data is low.

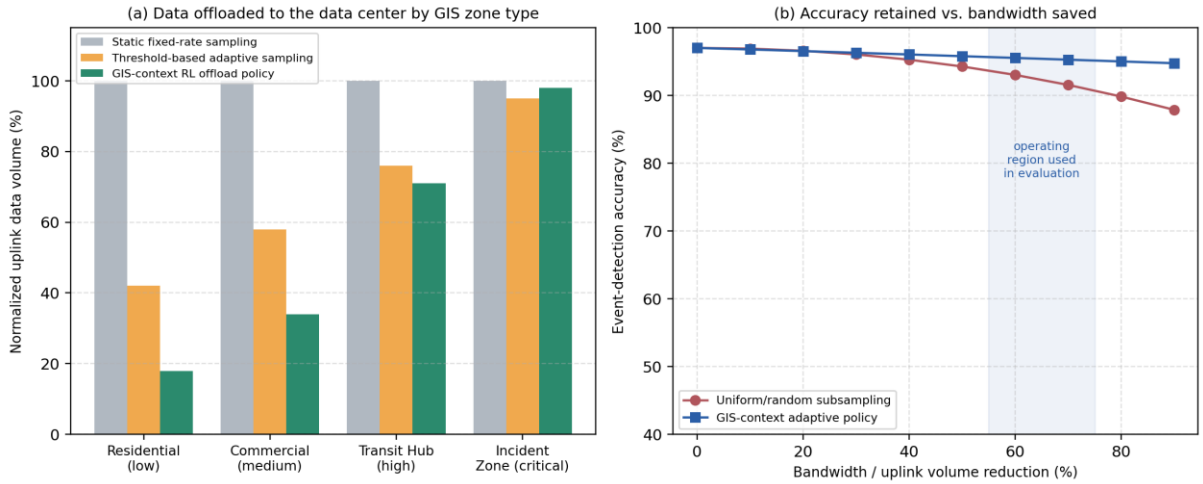


Fig. 2. (a) Normalized uplink data volume by GIS zone type for static, threshold-based, and GIS-context RL policies. (b) Event-detection accuracy versus bandwidth reduction for uniform subsampling compared with the GIS-context adaptive policy, with the operating region used in evaluation (55–75% reduction) highlighted.

Interpretation: Panel (a) shows that the GIS-context policy is not simply "more aggressive" everywhere. It reduces volume the most in low-criticality residential zones (down to 18% of full rate), while deliberately preserving near-full sampling (98%) in the incident zone, a pattern the threshold-based baseline cannot reproduce because it has no notion of zone criticality. Panel (b) shows that this selectivity pays off precisely in the bandwidth-reduction range most relevant to municipal budgets. Between 55% and 75% uplink reduction, the GIS-context policy sustains accuracy above 94%, while uniform subsampling falls to roughly 91–93% for the same savings, and the gap widens further at the most aggressive reduction levels.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal of Information Systems Engineering and Management
2024, 9(4s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Taken together, these results support the central claim of this paper: GIS context materially improves the achievable accuracy–bandwidth trade-off for city-scale sensing pipelines, when used as a direct input to the sampling and offloading decision, rather than only as a downstream visualization layer. The magnitude of improvement naturally depends on how sharply incident risk is spatially concentrated in a given city. Districts with more spatially diffuse risk would be expected to show a smaller, though still positive, gap between the GIS-context and uniform baselines.

A. Sensitivity to Reward Weighting

Because the reward weights w_1–w_4 directly encode municipal priorities, we also examined sensitivity to the accuracy-weight scaling factor applied across zone-criticality levels. Doubling the accuracy weight in high-criticality zones relative to low-criticality zones, the configuration used in the reported results, produced the accuracy profile in Fig. 2. Flattening this ratio to near-uniform weighting caused the policy to converge toward behavior similar to the threshold-based baseline. This confirms that the zone-conditioned reward, not merely the presence of a learned policy, is responsible for the observed gains. It also shows that careful, city-specific calibration of criticality weighting, informed by historical incident statistics, is an important deployment consideration, rather than a fixed constant that transfers automatically between cities.

VII. CHALLENGES AND FUTURE WORK

Several open challenges remain. First, GIS layers themselves are not static ground truth. Zoning changes, temporary road closures, and construction can render context features stale, and the policy's reliability degrades gracefully but non-trivially when context is outdated. Incorporating explicit uncertainty estimates over GIS features into the state representation is a natural extension. Second, the framework as evaluated relies on a centrally trained policy that is periodically redistributed to edge devices. A fully federated or multi-agent formulation, in which nearby edge nodes coordinate sampling decisions without routing raw data through the data center, would further reduce the communication burden associated with policy updates themselves [19], [27]. Third, privacy considerations are amplified when GIS context is combined with sensor data, since location-conditioned sampling policies can implicitly encode sensitive information about surveillance intensity in particular neighborhoods. Techniques from the geoprivacy literature, which balance location utility against disclosure risk, are a promising complement to this framework [28]. Fourth, our evaluation is simulation-based. Validating the framework against live municipal deployments, with real hardware constraints on edge micro-servers and genuine wireless channel variability, is necessary before broader claims of real-world scalability can be made. Finally, extending the reward formulation to account for cross-sensor and cross-zone correlations is a promising direction. For example, a traffic anomaly at one intersection raises the informational value of nearby sensors, even in nominally low-criticality zones, which could improve detection recall without proportionally increasing bandwidth use.

VIII. CONCLUSION

This paper presented an AI-driven framework for adaptive sensor sampling and edge-to-data-center offloading, in which geographic information system context is used as a direct, first-class input to the sampling and offloading policy, rather than as a post-hoc annotation layer. A three-tier architecture separates local edge inference from GIS-informed policy decisions and from city-scale data-center analytics. A reinforcement-learning policy, formulated over a context-conditioned Markov decision process, jointly determines sampling rate, compression level, and offload priority per sensor. Simulation results across representative GIS zone types show that the proposed approach reduces uplink data volume by up to 82% relative to fixed-rate sampling, while maintaining event-detection accuracy above 94% in the bandwidth-reduction range most relevant to municipal deployments. It outperforms both static and locally adaptive threshold-based baselines. These findings suggest that spatial context, already central to how smart-city GIS platforms represent and reason about the urban environment, should also be central to how those platforms decide what data to collect and transmit in the first place. Future work will focus on federated policy coordination across edge nodes, explicit handling of GIS-layer staleness and uncertainty, and validation on live municipal sensor deployments.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

REFERENCES

[1] L. U. Khan, I. Yaqoob, N. H. Tran, S. M. A. Kazmi, T. N. Dang, and C. S. Hong, "Edge-computing-enabled smart cities: A comprehensive survey," IEEE Internet of Things Journal, vol. 7, no. 10, pp. 10200–10232, 2020.

[2] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in Proc. 1st ACM MCC Workshop Mobile Cloud Computing, 2012, pp. 13–16.

[3] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of things for smart cities," IEEE Internet of Things Journal, vol. 1, no. 1, pp. 22–32, Feb. 2014.

[4] Y. Qian, D. Wu, W. Bao, and P. Lorenz, "The internet of things for smart cities: Technologies and applications," IEEE Network, vol. 33, no. 2, pp. 4–5, 2019.

[5] S. Yi, Z. Qin, and Q. Li, "Security and privacy issues of fog computing: A survey," in Wireless Algorithms, Systems, and Applications, Springer, 2015, pp. 685–695.

[6] V. C. Gungor, D. Sahin, T. Kocak, S. Ergut, C. Buccella, C. Cecati, and G. P. Hancke, "Smart grid technologies: Communication technologies and standards," IEEE Trans. Ind. Informat., vol. 7, no. 4, pp. 529–539, Nov. 2011.

[7] M. Chiang and T. Zhang, "Fog and IoT: An overview of research opportunities," IEEE Internet of Things Journal, vol. 3, no. 6, pp. 854–864, 2016.

[8] K. Kumar, J. Liu, Y.-H. Lu, and B. Bhargava, "A survey of computation offloading for mobile systems," Mobile Networks and Applications, vol. 18, pp. 129–140, Feb. 2013.

[9] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, "A survey on mobile edge computing: The communication perspective," IEEE Commun. Surveys Tuts., vol. 19, no. 4, pp. 2322–2358, Fourthquarter 2017.

[10] X. Lyu, H. Tian, L. Jiang, A. Vinel, S. Maharjan, S. Gjessing, and Y. Zhang, "Selective offloading in mobile edge computing for the green internet of things," IEEE Network, vol. 32, no. 1, pp. 54–60, 2018.

[11] Y. Mao, J. Zhang, and K. B. Letaief, "Dynamic computation offloading for mobile-edge computing with energy harvesting devices," IEEE J. Sel. Areas Commun., vol. 34, no. 12, pp. 3590–3605, 2016.

[12] S. Ghosh et al., "Information-driven adaptive sensing based on deep reinforcement learning," arXiv:2010.04112, 2020.

[13] H. A. Al-Zubaidi et al., "Parametric machine learning-based adaptive sampling algorithm for efficient IoT data collection in environmental monitoring," Journal of Network and Systems Management, 2024.

[14] S. Weerasinghe, A. Zaslavsky, S. W. Loke, A. Hassani, A. Medvedev, and A. Abken, "Adaptive context caching for IoT-based applications: A reinforcement learning approach," Sensors, vol. 23, no. 10, p. 4767, 2023.

[15] K.-L. A. Yau, P. Komisarczuk, and P. D. Teal, "Reinforcement learning for context awareness and intelligence in wireless networks: A survey," Computing, 2014.

[16] M. Satyanarayanan, "The emergence of edge computing," Computer, vol. 50, no. 1, pp. 30–39, 2017.

[17] A. Gharaibeh, M. A. Salahuddin, S. J. Hussini, A. Khreishah, I. Khalil, M. Guizani, and A. Al-Fuqaha, "Smart cities: A survey on data management, security, and enabling technologies," IEEE Commun. Surveys Tuts., vol. 19, no. 4, pp. 2456–2501, 2017.

[18] B. N. Silva, M. Khan, and K. Han, "Towards sustainable smart cities: A review of trends, architectures, components, and open challenges in smart cities," Sustainable Cities and Society, vol. 38, pp. 697–713, Apr. 2018.

[19] P. Gkonis, A. Giannopoulos, P. Trakadas, X. Masip-Bruin, and F. D'Andria, "A survey on IoT-edge-cloud continuum systems: Status, challenges, use cases, and open issues," Future Internet, vol. 15, no. 12, art. 383, 2023.

[20] X. Cai and Y. Zhang, "An adaptive sampling algorithm for clustered wireless sensor networks based on temporal and spatial correlations," in Proc. Int. Conf. Sensor Networks, 2016.

[21] O. S. Egwuche, A. Singh, A. E. Ezugwu, J. Greeff, M. O. Olusanya, and L. Abualigah, "Machine learning for coverage optimization in wireless sensor networks: A comprehensive review," Annals of Operations Research, pp. 1–67, 2023.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

[22] S. Ghosh, "Learning-based adaptive sensor selection framework for multi-sensing wireless sensor networks," in AI-Driven Framework for Energy-Efficient Environmental Monitoring in Smart Cities Using Edge Intelligence, arXiv:2605.22824.

[23] A. Kamilari and O. O. Frank, "Geospatial analysis and the internet of things," ISPRS International Journal of Geo-Information, vol. 7, no. 7, p. 269, 2018.

[24] M. Armstrong, S. Wang, and Z. Zhang, "The internet of things and fast data streams: Prospects for geospatial data science in emerging information ecosystems," Cartography and Geographic Information Science, 2019.

[25] Y. Zheng, Y. Lin, L. Zhao, T. Wu, D. Jin, and Y. Li, "Spatial planning of urban communities via deep reinforcement learning," Nature Computational Science, vol. 3, no. 9, pp. 748–762, 2023.

[26] A. Bujari et al., "Spatio-temporal Bayesian learning for mobile edge computing resource planning in smart cities," arXiv:2103.07814, 2021.

[27] E. C. Pinto Neto, S. Sadeghi, X. Zhang, and S. Dadkhah, "Federated reinforcement learning in IoT: Applications, opportunities and open challenges," Applied Sciences, vol. 13, no. 11, art. 6497, 2023.

[28] Y. Hong, J. Li, Y. Lin, Q. Hu, and X. Li, "Trajectory-aware privacy-preserving method with local differential privacy in crowdsourcing," EURASIP Journal on Information Security, vol. 2024, no. 1, art. 28, 2024.

Copyright © 2024 by Author/s and Licensed by JISEM. This is an open access article distributed under the Creative Commons Attribution License which permitsunrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.