AI-Enhanced Wide-Area Data Imaging via Massive Non-Orthogonal Direct Device-to-HAPS Transmission

Hyung-Joo Moon, Chan-Byoung Chae, Kai-Kit Wong, and Robert W. Heath, Jr.,

Abstract

Massive Aerial Processing for *X* (MAP-*X*) is an innovative framework for reconstructing spatially correlated ground data, such as environmental or industrial measurements distributed across a wide area, into data maps using a single high altitude pseudo-satellite (HAPS) and a large number of distributed sensors. With subframe-level data reconstruction, MAP-*X* provides a transformative solution for latency-sensitive IoT applications. This article explores two distinct approaches for AI integration in the post-processing stage of MAP-*X*. The model-driven pointwise estimation approach enables real-time, adaptive reconstruction through online training, while the end-to-end image reconstruction approach improves reconstruction accuracy through offline training with non-real-time data. Simulation results show that both approaches significantly outperform the conventional inverse discrete Fourier transform (IDFT)-based linear post-processing method. Furthermore, to enable AI-enhanced MAP-*X*, we propose a ground-HAPS cooperation framework, where terrestrial stations collect, process, and relay training data to the HAPS. With its enhanced capability in reconstructing field data, AI-enhanced MAP-*X* is applicable to various real-world use cases, including disaster response and network management.

Introduction

The exponential growth of connected devices has established the Internet of Things (IoT) as a foundational technology in modern digital infrastructure [1]. IoT networks now play a critical role in wide-area monitoring by leveraging spatiotemporal correlations in industrial and environmental data to facilitate real-time intelligence and decision-making. Existing IoT solutions, however, still face challenges in coverage, data acquisition speed, and connectivity, particularly in vast or remote areas. Addressing these limitations requires breakthrough wide-area sensing technologies that provide broader coverage, ultra-low-latency, and scalable support for massive device access.

Geographic information systems (GIS), which map spatially referenced information, are a key objective of wide-area sensing technologies [2]. Traditionally, GIS has been implemented through either wireless sensor networks (WSNs) or remote sensing. The conventional data communication-based approach reconstructs large-scale field data using methods such as kriging interpolation or spatial estimation from distributed sensor measurements [3, 4]. In contrast, remote sensing collects indirect environmental information via satellite or aerial imagery by capturing reflected sunlight or emitted radiation. Figure 1 presents a conceptual hierarchy of wide-area sensing approaches, with the proposed MAP-*X* framework positioned as a WSN-based system leveraging unlimited non-orthogonal access.

Each conventional approach has its own strengths and limitations. Satellite remote sensing provides broad coverage but often relies on indirect inference, limiting its ability to capture certain ground-level parameters [5]. UAV-based sensing and hybrid ground-air architectures offer deployment flexibility but are constrained by limited flight time, coordination overhead, and latency. Traditional WSNs enable targeted data acquisition through distributed low-power sensors, but their limited communication range typically necessitates multi-hop relay or scheduled access, which introduces latency and protocol complexity.

To increase access scalability in IoT networks, non-orthogonal communication schemes such as non-orthogonal multiple access (NOMA) and over-the-air computation (OAC)-based approaches have been proposed [6, 7]. However, these approaches require uplink channel estimation and phase-level synchronization, which limit the number of simultaneously supported devices and increase system complexity.

In contrast, MAP-*X* introduces a new paradigm that combines:

- Unlimited non-orthogonal access, enabling all devices to transmit simultaneously. This design leverages spatial redundancy for reliability and increases the expected signal power through random superposition.
- One-shot processing, completed within a few OFDM subframes, without requiring iterative

Hyung-Joo Moon is with the School of Integrated Technology, Yonsei University, South Korea and with the BK21 Graduate Program in Intelligent Semiconductor Technology, South Korea; Chan-Byoung Chae (corresponding author) is with the School of Integrated Technology, Yonsei University, South Korea; Kai-Kit Wong is with the Department of Electronic and Electrical Engineering, University College London, UK and with Yonsei Frontier Lab., Yonsei University, South Korea; Robert W. Heath, Jr. (corresponding author) is with the Department of Electrical and Computer Engineering, University of California at San Diego, USA.

Digital Object Identifier: 10.1109/MCOM.001.2500183

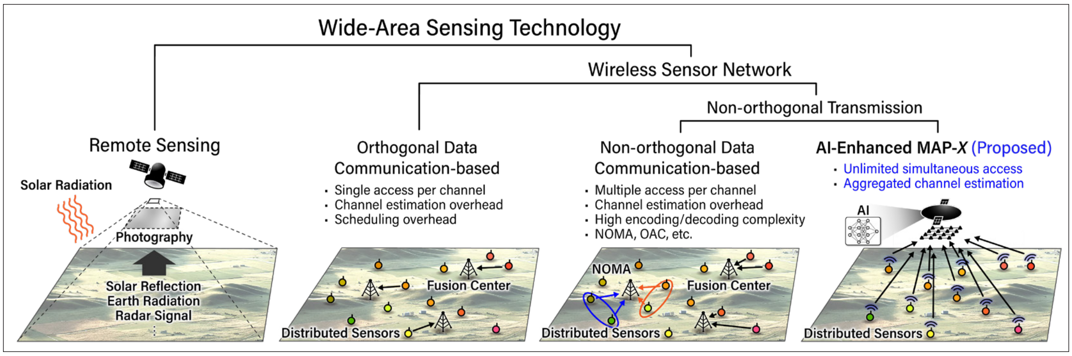

FIGURE 1. Comparison of GIS approaches, with MAP-X as a non-orthogonal WSN-based method.

Using angle-of-arrival (AoA) domain processing, MAP-*X* superimposes signals from nearby devices while spatially distinguishing those from more distant sources.

scheduling or feedback overhead.

- Aggregated channel estimation, which replaces per-device channel estimation by capturing the collective spatial response during a shared reference transmission phase.
- AI-enhanced post-processing, enabling accurate field reconstruction from superimposed signals using neural models at a centralized high altitude pseudo-satellite (HAPS).

These characteristics make MAP-*X* particularly suited for time-sensitive applications such as disaster response, infrastructure monitoring, and large-scale network management.

MAP-*X* for Real-Time Wide-Area Data Imaging

In this section, we introduce the MAP-*X* framework by detailing the system model and the operational phases that enable real-time wide-area data imaging. We begin by outlining the system model, including the sensing target, devices, non-terrestrial station, terrestrial station, and environments. We then describe the three sequential communication phases of MAP-*X*. Finally, we discuss the key innovations in waveform design and transmission strategy that empower the system.

System Model of MAP-*X*

HAPS equipped with multiple-input multiple-output (MIMO) or reconfigurable intelligent surfaces (RIS) are expected to play a key role in future non-terrestrial networks (NTNs), providing active spatial processing at altitude [8, 9]. MAP-*X* leverages a non-orthogonal, direct device-to-HAPS transmission scheme for wide-area and low-latency GIS. In this framework, numerous ground devices simultaneously transmit unique multi-carrier radio-frequency (RF) waveforms over shared time-frequency resources. These randomly superimposed signals are captured by a HAPS equipped with a uniform planar array (UPA) antenna.

Using angle-of-arrival (AoA) domain processing, MAP-*X* superimposes signals from nearby devices while spatially distinguishing those from more distant sources. Given that environmental and industrial data typically exhibit high spatial correlation, adjacent devices transmit nearly identical data, leading to overlapping signals that reinforce each other. This increases the average signal-to-noise ratio (SNR), making direct transmission from low-power devices to the HAPS more efficient. Mathematical analysis in [10] confirms that effective data reconstruction is attainable using a simple inverse discrete Fourier transform (IDFT)-based linear post-processing method.

MAP-*X* can be considered a type of WSN framework; however, it differs significantly in its operational principles. Instead of orthogonal access and per-device data communication, MAP-*X* relies on a unique combination of non-orthogonal simultaneous transmission, spatial redundancy, and centralized aerial processing. Below, we outline the distinct characteristics required by each component in the MAP-*X* framework.

Sensing Target: MAP-*X* is specifically designed for spatially correlated data fields, such as environmental, industrial, or network metrics, where nearby sensors tend to measure similar values. Unlike conventional systems, it enables one-shot, wide-area reconstruction of such fields within milliseconds by exploiting spatial redundancy. This makes MAP-*X* suitable not only for slowly varying signals, but also for fast-changing, time-sensitive phenomena that require immediate situational awareness.

Devices: The key distinction in MAP-*X* is that all devices transmit simultaneously over shared time-frequency resources without establishing uplink connections. Each device encodes its sensing value into the amplitude of a predefined waveform, and applies a location-dependent phase modulation across subcarriers and OFDM symbols. These operations are performed autonomously and locally, based only on each device's location and sensing value. This eliminates the need for individual phase-level synchronization, feedback, or channel estimation.

Non-Terrestrial Station: MAP-*X* requires a high-altitude fusion center capable of resolving superimposed signals in the AoA domain. In this work, we adopt an aerostatic HAPS as the primary implementation due to its quasi-stationary positioning, suitable coverage range for MAP-*X*, and minimal Doppler shifts. Importantly, aerostatic HAPSs offer sufficient payload capacity to accommodate a large UPA and the associated RF chains needed for MAP-*X*'s AoA-domain signal separation. This aerial station plays a central role by receiving simultaneous transmissions from thousands of devices and reconstructing the spatial data field from a single aggregated reception, bypassing the need for per-link demodulation as in conventional WSNs.

Terrestrial Station: The terrestrial station is not involved in the real-time communication loop but plays a crucial role in supporting AI-based post-processing. It gathers ground-truth training data (either scalar measurements or full field maps) for training AI models used in reconstruction.

Environments: MAP-*X* is robust to spatial inhomogeneity and irregular terrain [10]. Unlike con-

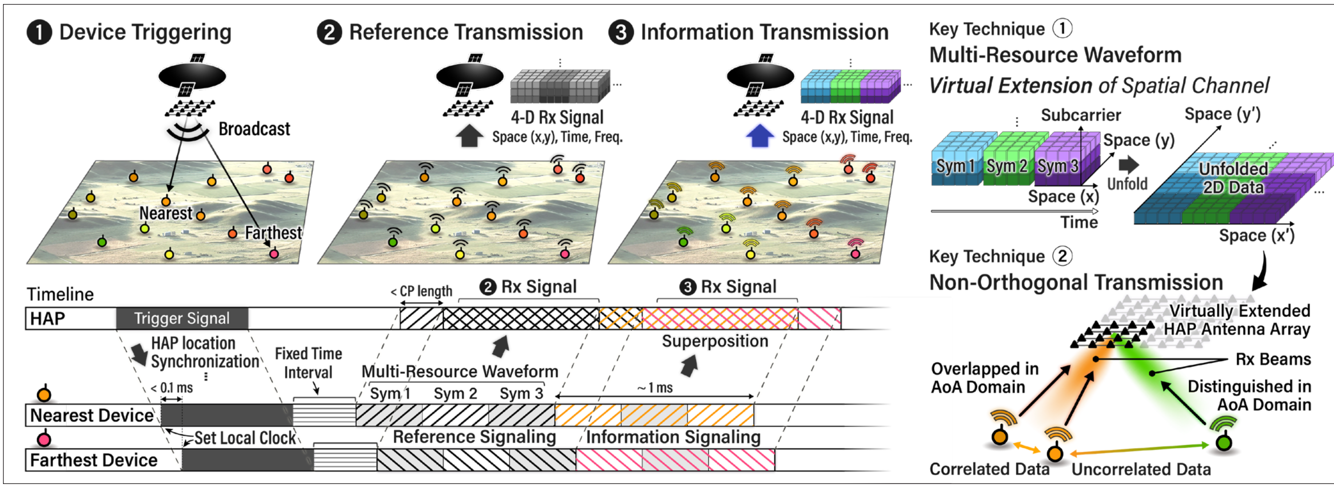

FIGURE 2. Three transmission phases of MAP-X and the two key techniques supporting its operation.

ventional WSNs that require structured routing or relay communications, MAP-*X* performs well with non-uniform device distributions. The system experiences only minor performance degradation in non-line-of-sight (NLoS) conditions, primarily due to SNR reduction. The system requires channel coherence on the order of several milliseconds, which is a realistic assumption given the low mobility of sensing devices and the pseudo-stationary position of the HAPS. Moreover, because MAP-*X* does not require channel estimation for each individual device-to-HAPS link, it avoids performance issues related to channel estimation errors or feedback overhead.

Three Transmission Phases of MAP-*X*

MAP-*X* is an efficient data collection and field reconstruction strategy within a WSN, distinguished by its use of non-orthogonal, simultaneous uplink transmissions and centralized aerial signal fusion. Unlike conventional WSNs that require individual connections and channel estimation per device, MAP-*X* exploits spatial signal structure and timing alignment to estimate data collectively. Figure 2 illustrates the three key phases of MAP-*X* operation.

Device Triggering: The MAP-*X* process begins with the HAPS broadcasting a trigger signal to all sensing devices. This signal includes metadata such as the type of data to be transmitted, the HAPS's location, and a reference delay corresponding to the nearest device. Additionally, the trigger signal synchronizes the devices' local oscillators and clocks, ensuring time alignment based on their distance from the HAPS. For instance, the nearest device maintains the earliest clock timing, while the farthest device operates on the latest clock timing.

Reference Transmission: Following synchronization, each device generates a waveform according to a shared protocol. Using a uniquely designed OFDM-based waveform structure, all devices apply a predefined phase modulation across time-frequency resource elements (i.e., across OFDM subcarriers and symbols), while fixing the symbol amplitude to a known constant (e.g., 1). The resulting signal received at the HAPS is the random superposition of all device transmissions. Through receiver beamforming using the HAPS's UPA antenna, this superimposed signal can be spatially separated into narrow angle-of-arrival AoA bins, each representing the aggregated channel response of devices within a specific angular region. This process replaces the need for per-device uplink channel estimation, as required in conventional communication-based WSNs.

Information Transmission: Immediately after the reference transmission, each device transmits again using the same waveform and phase modulation pattern. However, this time the symbol amplitudes encode the devices' sensing data, mapped to positive real values using a predefined encoding function. Because neighboring devices observe similar environmental conditions, the amplitude values tend to have high local correlation. Since both transmissions occur within the channel coherence time, the HAPS can associate the two received signals, reference and information, as corresponding measurements from the same spatial region. By dividing the information signal by the reference signal in each AoA-domain bin, the aggregated channel effects are canceled, and the original data value for that spatial region is recovered with signal noise.

Key Techniques Underlying the Transmission Phases

The main innovation of MAP-*X* lies in its novel waveform design and non-orthogonal transmission strategy. The mathematical foundations and performance analysis of these techniques are detailed in [10]. This section focuses on the technical implications and practical effects of these advancements.

OFDM-Based Multi-Resource Waveform: One of the primary challenges in aerial platforms is the physical limitation on the size of the antenna array, imposed by weight and energy constraints. To address this, MAP-*X* employs an innovative waveform that uses additional time-frequency resource elements to effectively "extend" the spatial domain resolution. In both reference and information transmission steps, devices modulate OFDM symbol phases, which are computed using a predefined formula based on their precise location relative to the HAPS. This allows the received symbols in the time-frequency domain to be unfolded into a virtually extended spatial domain, as illustrated in Fig. 2. As a result, this technique simulates the effect of an extended antenna array, significantly enhancing the resolution of AoA domain

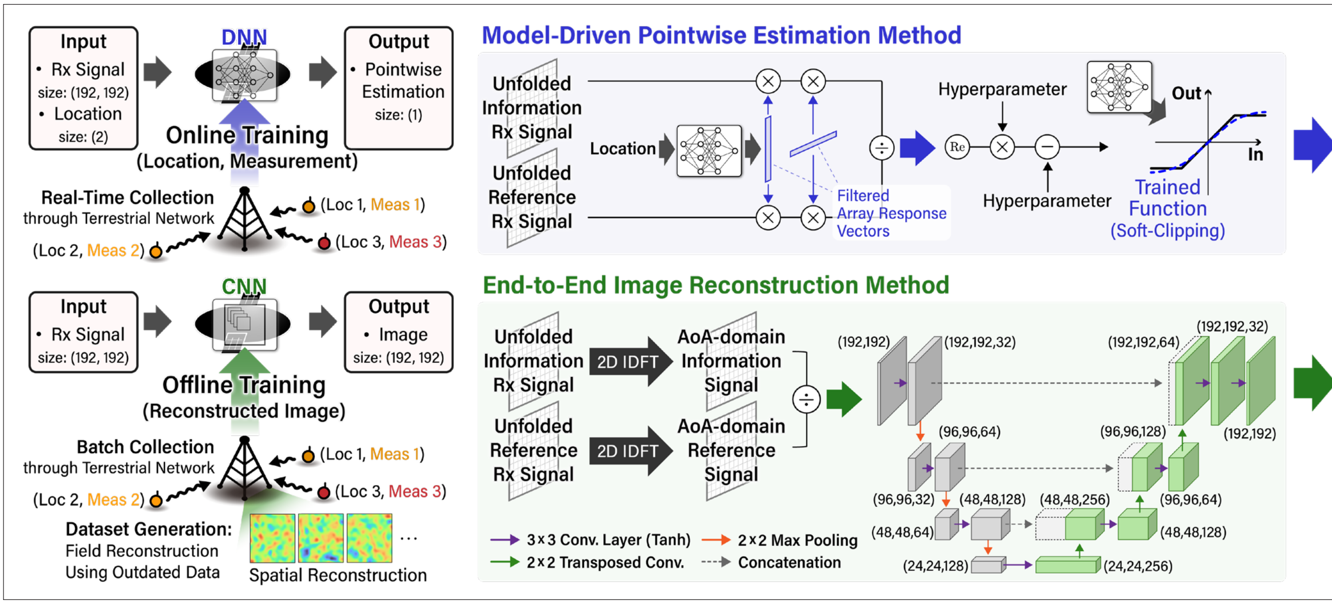

FIGURE 3. Training and inference pipelines for the two deep learning-based post-processing methods. In the pointwise estimation method, only the blue modules are implemented using lightweight AI models.

channelization. Narrowing the AoA domain channel increases the uniformity of data from a single AoA region, thereby maximizing the effectiveness of the superposition-based transmission scheme.

Overlapping Signals via Non-Orthogonal Transmission: As outlined in the three transmission phases of MAP-*X*, all devices transmit signals over the same multiple time-frequency resources, leading to signal superposition at the receiver. A related technique, OAC, intentionally exploits signal overlap from different devices to perform nomographic operations [7]. MAP-*X* also leverages signal superposition to ensure that information from closely located devices is effectively averaged, while signals from spatially distant devices remain fully distinguishable in the AoA domain, as illustrated in Fig. 2. Unlike OAC, which requires uplink channel state information at the transmitters to achieve phase and magnitude alignment at the receiver, MAP-*X* avoids this requirement entirely, as it does not rely on such alignment. This non-orthogonal transmission strategy offers several key advantages:

- Overlapping data from adjacent devices exploits spatial correlation in the raw data, improving accuracy.
- Non-orthogonal transmission significantly shortens transmission time compared to data gathering through orthogonal data transmission.
- The SNR enhancement from signal overlap enables direct device-to-HAPS transmission, mitigating the high pathloss.
- Under the channel coherence time assumption, MAP-*X* operates without requiring channel knowledge or individual data decoding, making it resilient to incoherent signal superposition.

AI-Enhanced Post-Processing for MAP-*X*

Upon completing the three transmission phases of MAP-*X*, the HAPS fusion center acquires two received symbol tensors: one from the reference transmission and another from the information transmission (Fig. 2). These tensors are four-dimensional, spanning the spatial *x*- and *y*-axes (antenna array coordinates), time (symbols), and frequency (subcarriers). Based on our waveform design, the time-frequency structure can be unfolded into a 2D virtually extended spatial domain. According to the mathematical derivation in [10], field reconstruction can be achieved by applying a 2D IDFT to the unfolded data, followed by element-wise division and clipping. This linear reconstruction process, which relies solely on the received signals, serves as the foundation for the advanced deep learning-based methods discussed next.

Model-Driven Pointwise Estimation Method

The model-driven pointwise estimation method is rooted in a closed-form estimator and enhances it with two learnable modules: a spatial filter and a soft-clipping function, as illustrated in Fig. 3. These modules are introduced to better adapt to variations in signal characteristics and noise across operating environments. Because these functions are relatively simple and mathematically grounded, they can be implemented with compact DNNs, making this approach ideal for real-time adaptation and lightweight deployment.

This method focuses on estimating a single data value at a given spatial location. Therefore, its input consists of the received signal and the coordinate of the target location, and its output is the estimated data value. As shown in Fig. 3, the signal is first transformed into the AoA domain. The trained spatial filter selectively shapes this domain, and the soft-clipping function mitigates outlier distortion caused by symbol division.

In our simulations, the spatial filter module is implemented using a four-layer MLP with dimensions 2, 96, 192, and 192 (ReLU activations), while the soft-clipping module uses a four-layer MLP with dimensions 1, 3, 3, and 1 (Tanh activation). Due to its compact structure and low data requirements, this method enables online training. The HAPS collects received signals autonomously, while the required ground-truth data (location-measurement pairs) are gathered by terrestrial stations via sidelink communication and forwarded to the HAPS for training. Since the

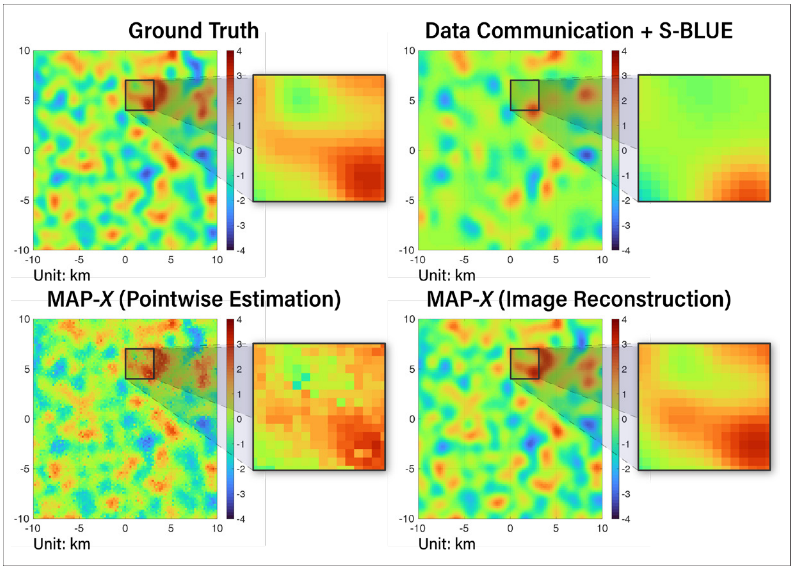

FIGURE 4. Image reconstruction results of the conventional data communication-based method and deep learning-aided MAP-X with equivalent resource consumption.

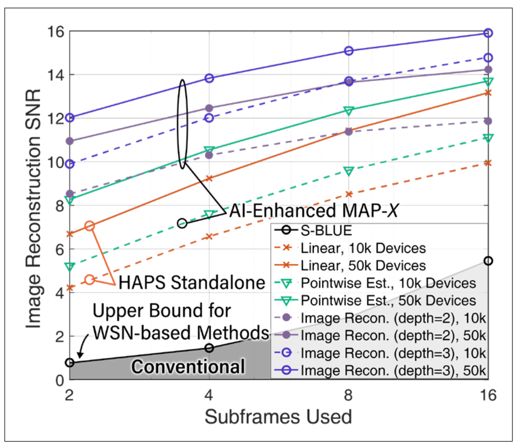

FIGURE 5. Comparison of image reconstruction SNR performance between the conventional data communication-based method, the HAPS standalone method, and the methods incorporating deep learning models. The black solid line represents the upper bound performance of conventional distributed sensing with orthogonal transmission.

base estimator performs reasonably well without training, only limited data are needed to fine-tune the filter and clipping functions, making this method adaptable in real-time.

End-to-End Image Reconstruction Method

The end-to-end image reconstruction method takes a fully data-driven approach to spatial field estimation. Unlike the model-driven pointwise estimation method, which outputs individual point estimates, this method generates an entire 2D map from the received signal. It uses a U-Net CNN architecture that is trained to learn the full transformation from the AoA-domain input to the ground-truth spatial field, including denoising and nonlinear distortions. The processing begins with a 2D IDFT and division to obtain a coarse reconstruction in the AoA domain. This intermediate output is then passed through the U-Net, which performs nonlinear transformation and refinement to produce a high-precision field map. The model architecture and signal flow are shown in Fig. 3.

To train this model, a complete dataset of received signals and corresponding field maps is required. These ground-truth maps are generated offline by terrestrial stations using local sensor measurements and high-quality spatial interpolation (e.g., kriging or covariance-based estimation). This dataset is then used to train the model, either on the HAPS or at a terrestrial station, with the final model deployed to the HAPS for inference. Due to its large parameter count and the difficulty of acquiring high-quality, full-resolution ground-truth maps in real-time, this method is restricted to offline training. However, once trained, it achieves superior reconstruction accuracy and denoising capability compared to the model-driven pointwise method.

Capabilities of AI-Enhanced MAP-*X*

We compare MAP-*X*, integrated with linear, model-driven pointwise estimation, and end-to-end image reconstruction post-processing strategies, against the conventional orthogonal data communication-based data reconstruction method. In the orthogonal data communication-based approach, a ground fusion center assigns each device a unique time-frequency resource element to transmit its location and measurement as a single symbol, without requiring prior signaling. The collected data are then used to reconstruct the continuous 2D field using the S-BLUE method, an optimal linear estimator that assumes perfect knowledge of the first- and second-order statistics of the spatial signal [4]. Based on this knowledge, S-BLUE computes optimal coefficients to fuse the individual observations into a field estimate. As a result, the orthogonal data communication-based method represents a highly idealized baseline. In contrast, MAP-*X* requires no prior statistical knowledge of the data, making it significantly more practical and scalable for real-world deployments.

In the simulation, we consider a total of 50,000 devices distributed over a 40 km × 40 km area, resulting in a device density of 31.25 devices per km^2. The raw target data in the simulation are modeled as a 2D Gaussian random field with normalized variance, filtered by a Gaussian function [10]. This filtering yields a correlation coefficient between two data points that decays proportionally to a Gaussian function of their separation distance. The HAPS is equipped with a 16 × 16 antenna array, and each subframe consists of 12 symbols and 12 subcarriers. Each device has a single isotropic antenna; the transmit signal is centered at 2.5 GHz, with a transmission power of 0 dBm. The direct device-to-HAPS transmission channel is modeled as Rician fading.

Image Reconstruction Results: Figure 4 presents the field reconstruction results produced by MAP-*X* using deep learning-based post-processing methods. In our simulations, four pairs of subframes are used to generate four independent estimates of the ground-truth data map, which are then averaged to produce a single, high-accuracy result. For a fair comparison, eight subframes are allocated to the orthogonal data communication-based data collection method. For the model-driven pointwise estimation method, values across the entire field are obtained by repeating the orthogonal data communication-based approach. However, due to the division step in the estimation process, it exhibits a heavy-tailed

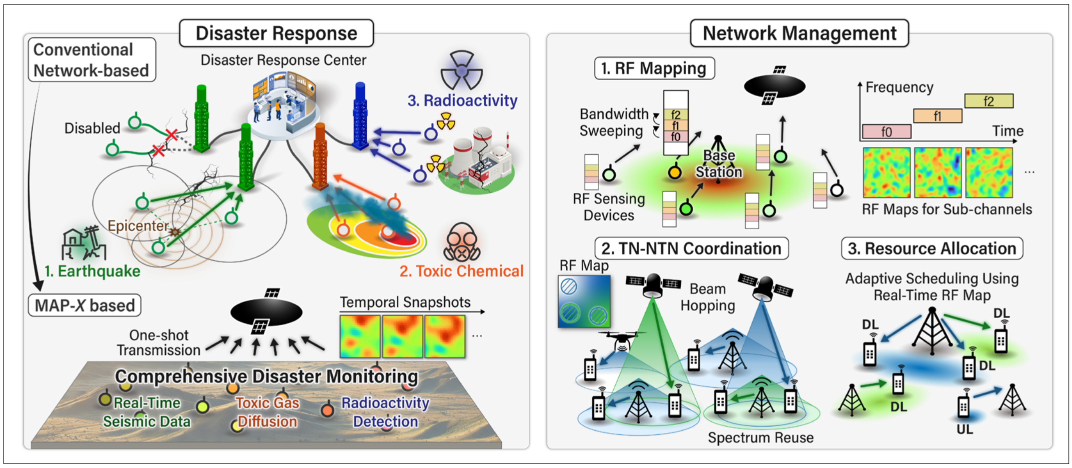

FIGURE 6. Potential use cases of MAP-*X*.

In disaster scenarios such as earthquakes, toxic chemical leaks, or radioactive incidents, rapid sensing is critical for damage mitigation.

error distribution that results in salt-and-pepper noise. This noise can be effectively mitigated by the CNN model, which not only improves estimation accuracy but also recovers finer signal details. Overall, these results demonstrate that MAP-*X* can accurately reconstruct field data, making it suitable for a wide range of applications, as further discussed next.

Latency vs. Accuracy Performance Analysis: Figure 5 compares the performance of MAP-*X* with conventional data communication-based field reconstruction. The *x*-axis represents the number of subframes used to reconstruct a single data map. S-BLUE utilizes additional subframes to gather more device data, while MAP-*X* employs a scalable approach, averaging the results of multiple MAP-*X* processes. As shown in Fig. 5, the upper bound performance of the traditional data communication-based method is significantly lower than the performance achieved by MAP-*X*. The DNN trained through online training improves performance, while the CNN, which requires offline training, achieves much higher accuracy under the same latency conditions. The pointwise estimation method is preferable for scenarios with rapidly changing signal characteristics and limited training data, whereas the image reconstruction method is more effective when target data exhibit strong temporal correlation and ample datasets are available from the terrestrial station.

Strengths of MAP-X and Future Applications

As demonstrated in the previous sections, MAP-*X* offers a new paradigm for large-scale data acquisition by leveraging simultaneous, non-orthogonal transmissions and centralized aerial fusion. One of its most critical strengths is its ability to reconstruct high-precision data maps within a few milliseconds, significantly faster than conventional WSN-based approaches that rely on multi-hop transmission, protocol handshaking, or sequential scheduling to gather data from individual links [11]. The MAP-*X* pipeline consists of only three tightly coupled stages: a one-shot device triggering signal, reference and information transmissions over two OFDM subframes, and lightweight post-processing at the HAPS. This architecture enables data collection and field reconstruction at latency levels on the order of few milliseconds, regardless of the number of participating devices.

This ultra-low-latency performance is a key differentiator in scenarios where rapid, wide-area sensing is critical. Environmental and industrial monitoring systems, for example, require frequent data snapshots of specific physical phenomena, often under constraints of unreliable connectivity or large sensing coverage. In such contexts, MAP-*X* functions as an extremely efficient GIS, capable of producing high-precision spatial maps from superimposed signals, without the need for per-device channel estimation or connection setup. This is particularly advantageous when sensing targets exhibit strong spatial correlation but limited temporal variation, as high-resolution data can be reconstructed from a single transmission round.

Latency-Critical Applications of MAP-X

The ultra-low-latency characteristics of MAP-*X* make it particularly effective in scenarios that demand fast, wide-area sensing. These include time-sensitive applications such as disaster response, real-time infrastructure monitoring, and network state management, where conventional WSNs suffer from protocol overhead, scheduling delays, or connectivity limitations. The potential use cases of MAP-*X* are illustrated in Fig. 6.

In disaster scenarios such as earthquakes, toxic chemical leaks, or radioactive incidents, rapid sensing is critical for damage mitigation. Conventional systems struggle with disrupted infrastructure, sparse coverage, or delayed data aggregation. MAP-*X*, by contrast, can quickly reconstruct spatial fields from distributed ground sensors without relying on individual links or relays. For earthquakes, this enables rapid estimation of ground motion patterns [12]; for chemical or radiation leaks, it supports real-time safety zone mapping where traditional remote sensing cannot detect invisible or non-radiative hazards [3, 13]. Its ability to aggregate field-level measurements from thousands of nodes into precise spatial maps is particularly useful for emergency response and forecasting.

Beyond disaster response, MAP-*X* can serve as a wide-area observatory for next-generation wireless network management. It enables subframe-level RF map generation to support real-time interference coordination, spectrum sharing in TN-NTN integration, and dynamic scheduling in multi-cell systems [14, 15]. While traditional RF mapping relies on long-term averaging or dedicated probes, MAP-*X* can generate low-latency, high-resolution spatial maps [14]. This allows base stations or satellites to adapt scheduling, beamforming, or spectrum allocation decisions based on timely environmental awareness. The combination of fast sensing, massive access, and AI-enhanced reconstruction gives MAP-*X* a strategic role in both emergency and infrastructure-level intelligence.

Future Research Directions and Challenges

Future work should focus on optimizing MAP-*X* under practical system constraints. This includes refining performance under realistic wireless channels, HAPS mobility, and device limitations such as imperfect synchronization and hardware nonidealities. Furthermore, validation through hardware-in-the-loop testing or prototyping is needed to assess feasibility in real-world deployments.

Another key direction is enhancing the AI pipeline. In addition to the AI-based post-processing introduced in this article, AI-based pre-processing can also be applied at the data level to compress multi-dimensional sensor readings into a lower-dimensional latent representation before transmission. This approach is particularly effective in MAP-*X*, where each device transmits only one scalar value per process. Optimizing the cooperation between ground stations and HAPS for efficient data relay and model adaptation will be essential to scaling MAP-*X* across various use cases.

Conclusion

In this article, we introduced the integration of AI with Massive Aerial Processing for *X* (MAP-*X*), an innovative wide-area data reconstruction framework. We compared two deep learning approaches, model-driven pointwise estimation and end-to-end image reconstruction methods, built on top of the mathematically-driven linear model. While the pointwise estimation method enables fast adaptation to environmental changes through online training, the image reconstruction method achieves higher data reconstruction accuracy by leveraging long-term data collection at terrestrial stations. These AI-enhanced MAP-*X* systems have the potential to revolutionize data collection and decision-making in latency-critical IoT applications, enabling faster and more accurate responses across various real-world scenarios.

Acknowledgments

This work was supported by Institute of Information & communications Technology Planning &Evaluation (IITP) grant funded by the MSIT (2022-0-00704 and RS-2024-00428780). The work of R. Heath was supported in part by the National Science Foundation under Grant No.NSF-ECCS-2435261, NSF-CNS-2433782, funds from federal agency and industry partners as specified in the Resilient & Intelligent NextG Systems (RINGS) program, as well as by industry affiliates of the Center for Wireless Communications at UC San Diego. The work of K. K. Wong was supported by the Engineering and Physical Sciences Research Council (EPSRC) under Grant EP/W026813/1.

References

[1] A. Ikpehai *et al.*, "Low-Power Wide Area Network Technologies for Internet-of-Things: A Comparative Review," *IEEE Internet Things J.*, vol. 6, no. 2, Apr. 2019, pp. 2225–40.
[2] Q. Meng *et al.*, "Assessment of Regression Kriging for Spatial Interpolation – Comparisons of seven GIS interpolation methods," *Cartogr. Geographic Inf. Sci.*, vol. 40, no. 1, Jan. 2013, pp. 28–39.
[3] L. Yi *et al.*, "Localized Confident Information Coverage Hole Detection in Internet of Things for Radioactive Pollution Monitoring," *IEEE Access*, vol. 5, Sep. 2017, pp. 18,665–74.
[4] P. Zhang *et al.*, "Spatial Field Reconstruction and Sensor Selection in Heterogeneous Sensor Networks with Stochastic Energy Harvesting," *IEEE Trans. Signal Process.*, vol. 66, no. 9, May 2018, pp. 2245–57.
[5] M. Fiore, "Full network sensing: Architecting 6G beyond communications," *IEEE Network*, vol. 37, no. 3, May/June 2023, pp. 232–39.
[6] D. Zhai *et al.*, "Energy-Saving Deployment Optimization and Resource Management for UAV-Assisted Wireless Sensor Networks with NOMA," *IEEE Trans. Veh. Technol.*, vol. 71, no. 6, 2022, pp. 6609–23.
[7] G. Zhu *et al.*, "Over-the-Air Computing for Wireless Data Aggregation in massive IoT," *IEEE Wireless Commun.*, vol. 28, no. 4, Sep. 2021, pp. 57–65.
[8] A. Azizi and A. Farhang, "RIS Meets Aerodynamic HAPS: A Multiobjective Optimization Approach," *IEEE Wireless Commun. Lett.*, vol. 12, no. 11, Nov. 2023, pp. 1851–55.
[9] A. Azizi, M. A. Kishk, and A. Farhang, "Exploring the impact of HAPS-RIS on UAV-Based Networks: A Novel Network Architecture," 2025; https://arxiv.org/abs/2409.17817
[10] H.-J. Moon *et al.*, "MAP-X: Massive Field Data Processing for Real-Time Wide-Area Mapping Using High-Altitude Platforms with MIMO," *IEEE Trans. Wireless Commun.*, to appear.
[11] M. Doudou, D. Djenouri, and N. Badache, "Survey on Latency Issues of Asynchronous MAC Protocols in Delay-Sensitive Wireless Sensor Networks," *IEEE Commun. Surveys Tuts.*, vol. 15, no. 2, 2nd Quarter, 2013, pp. 528–50.
[12] G. Liu *et al.*, "Volcanic Earthquake Timing Using Wireless Sensor Networks," *Proc. ISPN*, 2013.
[13] M. Faruolo *et al.*, "A Daytime Multisensor Satellite System for Global Gas Flaring Monitoring," *IEEE Trans. Geosci. and Remote Sens.*, vol. 60, Jan. 2022, pp. 1–17.
[14] Y. Li *et al.*, "Geo2SigMap: High-Fidelity RF Signal Mapping Using Geographic Databases," *Proc. IEEE DySPAN*, May 2024, pp. 277–85.
[15] H.-J. Moon and C.-B. Chae, "Cooperative Ground-Satellite Scheduling and Power Allocation for Urban Air Mobility Networks," *IEEE JSAC*, vol. 43, no. 1, Jan. 2025, pp. 218–33.

Biographies

Hyung-Joo Moon [GStM] (moonhj@yonsei.ac.kr) is currently pursuing a Ph.D. degree in the School of Integrated Technology, Yonsei University, South Korea. He is also affiliated with the BK21 Graduate Program in Intelligent Semiconductor Technology. His research interest is emerging technologies in 6G non-terrestrial networks (NTN).

Chan-Byoung Chae [F] is (cbchae@yonsei.ac.kr) an Underwood Distinguished Professor and Lee Youn Jae Fellow with the School of Integrated Technology, Yonsei University, South Korea. His research interests include emerging technologies in 6G and cloud systems.

Kai-Kit Wong [F] (kai-kit.wong@ucl.ac.uk) is a Professor in the Department of Electronic and Electrical Engineering, University College London (UCL). His research interest is wireless communications theory.

Robert W. Heath Jr. [F] (rwheathjr@ucsd.edu) is a Professor in the Department of ECE at the University of California, San Diego. He is an elected member of the National Academy of Engineering. His research interests include wireless communications, MIMO, and millimeter-wave systems.