Heterogeneous Sensor Data Fusion with Gated Abundance Networks for Urban Change Detection

Otis Rhodes
Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.
rhodes1972@ku.edu

Jose Boung
School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
jyoung@uc.edu

Abstract

Urban environments are undergoing rapid transformation driven by population growth, infrastructure expansion, and climate variability, creating an urgent need for reliable change detection systems that can fuse heterogeneous sensor data. This paper introduces a conceptual architecture for heterogeneous sensor data fusion using Gated Abundance Networks, a framework designed to reconcile the disparate spectral, spatial, temporal, and structural characteristics of multi-source remote sensing data—including optical imagery, synthetic aperture radar, hyperspectral data, and LiDAR point clouds. The Gated Abundance Network leverages learnable gating mechanisms and abundance-based representations to adaptively weight contributions from each sensor modality, thereby enhancing detection accuracy while maintaining interpretability. Rather than focusing on algorithmic details, this paper provides a system-level analysis of the architectural trade-offs, deployment infrastructure, governance challenges, and sustainability implications of such a fusion framework for urban change detection. We examine how gated abundance modeling can mitigate issues of missing data, varying acquisition schedules, and domain shifts across sensors, and discuss the computational and energy costs associated with large-scale deployment. Furthermore, we address fairness and robustness concerns, particularly in heterogeneous urban landscapes where sensor coverage may be uneven. The paper also explores policy implications, including data sharing regulations, privacy safeguards, and the need for standardized evaluation benchmarks. Through cross-domain comparisons with other socio-technical systems, we highlight the importance of modular, scalable fusion architectures that can adapt to evolving urban sensing ecosystems. The findings underscore that successful deployment of Gated Abundance Networks depends not only on technical performance but also on institutional coordination, transparent governance, and long-term sustainability planning.

Keywords

sensor fusion, urban change detection, gated networks, heterogeneous data, remote sensing, socio-technical systems, governance, sustainability.

1. Introduction

Urban change detection has become a cornerstone application in remote sensing, enabling monitoring of land-use transitions, infrastructure development, disaster impacts, and environmental degradation. The increasing availability of data from diverse sensor platforms—satellite, aerial, and ground-based—offers unprecedented opportunities to capture the multi-faceted nature of urban dynamics. However, the heterogeneity in spectral resolution, spatial coverage, revisit frequency, and data quality introduces substantial challenges for fusion-based change detection systems. Traditional approaches often rely on simplistic concatenation or early fusion of sensor outputs, failing to account for the distinct uncertainty levels and complementary information carried by each modality. This paper proposes a system-level perspective on heterogeneous sensor data fusion through the lens of Gated Abundance Networks, a novel architectural paradigm that dynamically weights sensor contributions based on learned gating functions and abundance-based representations.

The concept of a "gated abundance network" emerges from the intersection of deep attention mechanisms and spectral unmixing theory. In this framework, each sensor modality is processed by a dedicated encoder that extracts feature abundance vectors—fractional estimates of land-cover components. A gating mechanism then learns to assign adaptive weights to these abundance vectors before fusion, enabling the network to suppress noisy or missing modalities and emphasize reliable ones. This approach is particularly valuable in urban settings where sensor coverage is often uneven; for example, optical imagery may be obstructed by clouds while synthetic aperture radar (SAR) remains unaffected, or LiDAR data may be sparse in dense building canopies. By incorporating gating, the network can dynamically adjust its reliance on each source, enhancing robustness and interpretability.

In this paper, we deliberately avoid detailed mathematical formulations or algorithmic pseudocode, focusing instead on the architectural, infrastructural, governance, and sustainability dimensions of deploying such a fusion system at scale. We draw on cross-domain analogies from autonomous systems, telecommunications, and healthcare informatics to situate the Gated Abundance Network within a broader socio-technical context. The discussion is structured around four main themes: system architecture and trade-offs, deployment infrastructure and operational challenges, robustness and fairness across heterogeneous urban landscapes, and policy implications for data governance and long-term sustainability. Through this analysis, we aim to provide a roadmap for researchers and practitioners seeking to design fusion systems that are not only technically effective but also institutionally viable and ethically sound.

2. Background and Related Work

The field of urban change detection has evolved from pixel-based difference methods to sophisticated deep learning frameworks that exploit spatial and temporal patterns [1][2]. Early work relied on optical image pairs, but the limitations of cloud cover and illumination variability motivated the integration of SAR data, which offers all-weather capability [7]. Subsequent advances in hyperspectral imaging provided rich spectral information for material discrimination, though at the cost of increased data dimensionality and processing complexity [5][9]. LiDAR data, with its precise height measurements, introduced a new dimension for structural change analysis, particularly in 3D urban modeling [4]. Fusing these heterogeneous sources has long been recognized as a promising direction, yet practical systems have struggled with misalignment, differing spatial resolutions, and asynchronous acquisition schedules [3][10].

Deep learning has revolutionized sensor fusion by enabling end-to-end learning of joint representations. Convolutional neural networks (CNNs) and transformer architectures have been adapted to multi-modal inputs, often using early, late, or intermediate fusion strategies [11][16][17]. The success of attention mechanisms in sequence modeling [18] has inspired gated fusion approaches, where a learned gate controls the flow of information from each modality. For instance, Hong et al. [6] demonstrated that multimodal deep learning outperforms single-modal methods for land-cover classification, particularly when diverse modalities are available. In the context of change detection, gated networks have shown resilience to missing data by suppressing noisy channels [15]. The concept of abundance estimation, originally developed for spectral unmixing, provides a physically interpretable representation that aligns well with urban component detection [9]. By combining gating with abundance vectors, the Gated Abundance Network inherits both adaptability and interpretability.

The required reference [14] (Yang et al., 2025) evaluated band ordering strategies in hyperspectral and LiDAR fusion, highlighting that the sequence in which spectral and height features are combined can significantly alter detection performance. This finding underscores the importance of the gating mechanism in dynamically ordering and weighting features rather than relying on a fixed fusion strategy. Meanwhile, work by Xiong et al. [13] on grounding text-to-image customization offers insights into how contextual cues can guide feature selection—a concept that parallels the use of gating in our framework. While these studies focus on specific algorithmic innovations, the present paper broadens the lens to encompass the systemic implications of deploying such fusion systems in real-world urban sensing environments.

3. System Architecture and Structural Trade-Offs

The core architecture of a Gated Abundance Network for urban change detection consists of four main components: modality-specific encoders, abundance estimation modules, a gating mechanism, and a fusion decoder that produces change maps. Each encoder is designed to handle the unique characteristics of its sensor—for example, a spectral-spatial CNN for optical and hyperspectral data, a polarimetric feature extractor for SAR, and a point cloud feature net for LiDAR. The abundance estimation module outputs a soft fractional assignment of each pixel or object to a predefined set of urban change classes (e.g., new construction, demolition, vegetation loss, pavement change). The gating mechanism learns a set of weights that are conditioned on both the input data and the current task context, effectively selecting which modalities to trust and how to combine their abundance vectors.

A major architectural trade-off lies in the granularity of the gating. Global gating, where a single set of weights applies to the entire scene, is computationally efficient but may overlook local variations in sensor reliability. For example, parts of an urban area affected by cloud cover may require lower optical weights, while clear regions can fully utilize optical abundance. Per-pixel or per-object gating increases flexibility but comes with higher memory and training cost. Another trade-off involves the choice between early fusion (combining raw features before abundance estimation) and late fusion (combining abundance vectors). The Gated Abundance Network adopts a mid-level fusion strategy: abundance vectors from each modality are gated and then concatenated before entering the decoder. This design preserves the interpretability of abundance while allowing the gating to modulate influence.

The architecture must also accommodate asynchronous data collection. Urban change detection often relies on time series where sensor revisits are irregular. The gating mechanism can be extended to incorporate temporal attention, weighing not only across modalities but also across time steps. This temporal gating introduces additional complexity in model training, particularly when changes occur at different rates. A structural trade-off emerges between model depth and generalization: deeper fusion networks may overfit to specific sensor pairings, while shallow networks may underutilize complementary information. Cross-domain validation on multiple urban sites is essential to ensure that the learned gating strategies are not biased toward particular sensor configurations or seasonal patterns.

4. Deployment Infrastructure and Operational Challenges

The real-world deployment of a Gated Abundance Network for urban change detection requires a robust infrastructure that spans data acquisition, preprocessing, storage, model inference, and output dissemination. Satellite and aerial data streams are typically collected by different agencies or commercial providers, each with their own data formats, calibration standards, and update frequencies. An operational fusion system must therefore include a middleware layer that harmonizes coordinate systems, resamples spatial resolutions, and normalizes radiometric values. This preprocessing pipeline is non-trivial; errors in coregistration can propagate through the network and degrade performance [4]. The gating mechanism can partially compensate for misalignment by down-weighting misregistered channels, but it cannot fully correct systematic geometric errors.

Cloud-based computing platforms, such as those provided by major technology companies, offer scalable resources for processing large volumes of multi-temporal data. However, the latency and bandwidth requirements for streaming high-resolution hyperspectral or LiDAR data remain significant bottlenecks. Edge computing, where some inference is performed on-board satellites or drones, could reduce transmission loads but introduces constraints on model size and power consumption. The Gated Abundance Network, with its multiple encoders and gating layers, may be too heavy for real-time on-board processing; lightweight approximations or knowledge distillation would be necessary.

Operational sustainability also involves energy consumption. Training a deep multimodal network on urban-scale datasets can require thousands of GPU-hours, leading to a substantial carbon footprint. Once deployed, inference must be repeated frequently to capture temporal changes, especially in rapidly developing urban peripheries. Balancing the frequency of inference with energy constraints is a key operational challenge. Furthermore, the system must be maintained over time: sensor specifications change, new satellites are launched, and older instruments degrade. The gating mechanism can adapt to gradual sensor drift by relearning weights, but sudden sensor failures or decommissioning may require retraining on a reduced modality set. A modular architecture that allows easy addition or removal of sensor encoders is critical for long-term viability.

5. Robustness, Fairness, and Sustainability

Robustness in multi-sensor fusion is threatened by several sources of uncertainty, including sensor noise, atmospheric interference, seasonal variations, and adversarial examples. The Gated Abundance Network's gating mechanism inherently improves robustness by learning to suppress unreliable modalities. However, this resilience can be fragile if the gating is trained only on clean, well-aligned data. Domain adaptation techniques are necessary to ensure that the network generalizes to new urban areas with different sensor distributions [19]. Moreover, fairness concerns arise when sensor coverage is uneven across socioeconomic or geographic boundaries. Wealthier, centrally located urban districts may be imaged more frequently and with higher resolution than marginalized peripheral zones, leading to systematic biases in change detection accuracy. A gated network trained primarily on well-covered areas may underperform in underserved regions, potentially reinforcing existing inequalities in urban monitoring.

Sustainability encompasses not only environmental costs but also institutional longevity. Urban change detection systems are often funded by government agencies or research grants with fixed durations. Once funding ceases, the infrastructure may fall into disrepair, and the trained models become outdated. A sustainable approach involves designing the fusion framework as an open-source, community-maintained resource, with standardized benchmarks and continuous integration of new sensor types. The Gated Abundance Network can contribute to this by providing a flexible, plug-and-play architecture that accommodates evolving data sources without requiring complete redesign.

Transparency and explainability are also dimensions of fairness. Many deep fusion models operate as black boxes, making it difficult for urban planners or policymakers to understand why a particular change was flagged. The abundance-based representation of the Gated Abundance Network offers a degree of interpretability: each modality's fractional estimates can be visualized, and the gating weights reveal which sensors contributed most to a given detection. This transparency fosters trust and enables human oversight, particularly in high-stakes decisions such as post-disaster damage assessment or land use regulation enforcement.

6. Case Illustrations and Cross-Domain Comparisons

To illustrate the systemic considerations, we examine two hypothetical deployment scenarios. In the first scenario, a metropolitan planning authority seeks to monitor informal settlement expansion on the urban fringe. Optical imagery is available at weekly intervals but is frequently obstructed by aerosol pollution; SAR provides cloud-free coverage every three weeks; and LiDAR is acquired annually. A Gated Abundance Network, trained on a multi-year time series, learns to rely heavily on SAR for short-term change detection between cloud-free optical windows, while LiDAR data calibrates the vertical growth patterns. The gating weights reveal that SAR abundance is most activated during monsoon seasons, while optical abundance dominates in clear winter months. This dynamic adaptivity reduces false positives from cloud shadows and improves detection of new structures. However, the network's performance in low-income neighborhoods is worse due to fewer ground-truth labels; fairness requires active collection of validation data in those areas, which the planning authority must budget for.

In the second scenario, a humanitarian organization deploys a portable fusion system for post-earthquake building damage assessment using drone-based optical and LiDAR data. The gating mechanism allows the system to operate even if the LiDAR sensor malfunctions mid-mission; the network gracefully degrades to optical-only inference with a drop in accuracy but still produces usable damage maps. This robustness is critical in disaster contexts where sensor failure is common. Cross-domain comparisons with autonomous vehicle sensor fusion show analogous strategies: vehicles combine camera, LiDAR, and radar with learned gating to navigate in adverse weather [12]. The parallels suggest that gated abundance networks could be standardized across domains, enabling technology transfer between urban monitoring and robotics.

7. Policy Implications and Future Directions

The deployment of heterogeneous sensor fusion systems raises a host of policy issues. Data sharing agreements between public agencies and private satellite operators must define ownership, access rights, and usage restrictions. Privacy is a growing concern, particularly when high-resolution optical imagery can reveal individual activities. Gated abundance networks, by operating on aggregated abundance vectors rather than raw pixels, may offer a degree of privacy preservation, but the risk of reconstruction attacks remains. Policymakers need to establish guidelines for anonymization and differential privacy in urban change detection.

Standardization is another policy priority. Different cities use different sensor networks, and the lack of common benchmarks hampers cross-site comparisons and model transfer. International bodies such as the Committee on Earth Observation Satellites could promote a unified evaluation protocol for change detection fusion systems, including metrics for fairness and robustness. Additionally, funding agencies should incentivize long-term maintenance and open-source contributions rather than one-off prototype deployments.

Future research directions include developing self-supervised learning methods for pretraining the abundance estimators without extensive labeled data, incorporating social and demographic data as additional modalities, and designing federated learning frameworks that allow multiple cities to collaboratively train a Gated Abundance Network without sharing sensitive imagery. The integration of causal inference to distinguish between natural and human-induced changes could further enhance the system's analytical power. Ultimately, the success of such fusion architectures depends on careful alignment between technical design and the institutional ecosystems in which they operate.

8. Conclusion

Heterogeneous sensor data fusion for urban change detection is a complex systems challenge that transcends algorithmic innovation. This paper has presented the Gated Abundance Network as a conceptual framework that combines learnable gating with abundance-based representations to achieve robust, interpretable, and adaptive fusion. Rather than focusing on performance numbers, we have examined the architectural trade-offs, deployment infrastructure, robustness and fairness concerns, and policy implications that must be addressed for real-world viability. The required reference [14] and related studies underscore the importance of careful fusion strategy design, while cross-domain comparisons highlight transferable lessons. As urban sensing ecosystems continue to expand, the need for scalable, sustainable, and equitable fusion systems will only intensify. The Gated Abundance Network offers a promising direction, but its ultimate impact will be determined by the degree to which researchers, engineers, policymakers, and communities collaborate to embed it within a robust socio-technical framework.

References

1. Zhang, L., Zhang, L., & Du, B. (2016). Deep learning for remote sensing data: A technical tutorial on the state of the art. IEEE Geoscience and Remote Sensing Magazine, 4(2), 22–40.

2. Zhu, X. X., Tuia, D., Mou, L., Xia, G. S., Zhang, L., Xu, F., & Fraundorfer, F. (2017). Deep learning in remote sensing: A comprehensive review and list of resources. IEEE Geoscience and Remote Sensing Magazine, 5(4), 8–36.

3. Ghassemian, H. (2016). A review of remote sensing image fusion methods. Information Fusion, 32, 75–89.

4. Schmitt, M., & Zhu, X. X. (2016). Data fusion and remote sensing: An ever-growing relationship. IEEE Geoscience and Remote Sensing Magazine, 4(4), 6–23.

5. Ghamisi, P., Yokoya, N., Li, J., Liao, W., Liu, S., Plaza, J., ... & Benediktsson, J. A. (2019). Advances in hyperspectral image and signal processing: A comprehensive overview of the state of the art. IEEE Geoscience and Remote Sensing Magazine, 7(2), 48–85.

6. Hong, D., Gao, L., Yokoya, N., Yao, J., Chanussot, J., Du, Q., & Zhang, B. (2019). More diverse means better: Multimodal deep learning meets remote-sensing imagery classification. IEEE Transactions on Geoscience and Remote Sensing, 59(5), 4340–4354.

7. Wagner, W. (2016). A review of the use of synthetic aperture radar images for urban change detection. Remote Sensing, 8(4), 258.

8. Taubenböck, H., et al. (2012). The impact of urbanization on the urban heat island in Beijing using remote sensing data. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 5(3), 878–888.

9. Benediktsson, J. A., Palmason, J. A., & Sveinsson, J. R. (2005). Classification of hyperspectral data from urban areas based on extended morphological profiles. IEEE Transactions on Geoscience and Remote Sensing, 43(3), 480–491.

10. Huang, B., Zhang, H., & Song, H. (2019). Spatiotemporal urban change detection from remote sensing images: A review. International Journal of Remote Sensing, 40(9), 3331–3365.

11. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems, 25, 1097–1105.

12. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.

13. Xiong, Z., Xiong, W., Shi, J., Zhang, H., Song, Y., & Jacobs, N. (2024). Groundingbooth: Grounding text-to-image customization. arXiv preprint arXiv:2409.08520.

14. Yang, J. X., Wang, J., Li, Z., Sui, C., Long, Z., & Zhou, J. (2025). HSLiNets: Evaluating Band Ordering Strategies in Hyperspectral and LiDAR Fusion. IEEE Geoscience and Remote Sensing Letters.

15. Diakogiannis, F. I., Waldner, F., Caccetta, P., & Wu, C. (2020). ResUNet-a: A deep learning framework for semantic segmentation of remotely sensed data. ISPRS Journal of Photogrammetry and Remote Sensing, 162, 94–114.

16. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. Medical Image Computing and Computer-Assisted Intervention (MICCAI), 234–241.

17. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778.

18. Vaswani, A., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30.

19. Ravanbakhsh, M., et al. (2019). Open-set domain adaptation for remote sensing image classification. IEEE Transactions on Geoscience and Remote Sensing, 57(8), 5649–5664.

20. Tuia, D., et al. (2020). Advances in deep learning for remote sensing: From deep learning to explainable AI. ISPRS Journal of Photogrammetry and Remote Sensing, 169, 178–191.