Spectral-Spatial Feature Integration for Hyperspectral and LiDAR-Based Urban Land Cover Classification

Aapo R. Rhodes
Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.
aapo.work@unr.edu

Abstract

The fusion of hyperspectral imagery and LiDAR data has emerged as a powerful paradigm for urban land cover classification, leveraging complementary spectral and spatial information to achieve high accuracy in complex built environments. This paper presents a systematic examination of spectral-spatial feature integration within the context of large-scale urban monitoring systems, emphasizing structural trade-offs, architectural design, deployment infrastructure, and governance implications. While numerous convolutional and attention-based fusion approaches have demonstrated superior performance on benchmark datasets, their practical deployment raises critical issues concerning computational scalability, sensor alignment, data heterogeneity, and generalization across diverse urban morphologies. This study analyzes the architectural choices that govern fusion effectiveness, including early, intermediate, and late fusion strategies, and assesses their implications for robustness against noise, missing data, and domain shifts. The discussion extends to infrastructure requirements for real-time or near-real-time classification, including on-board processing constraints, cloud-edge coordination, and energy sustainability. From a policy perspective, this paper highlights fairness and interpretability challenges that arise when fused models are deployed in urban planning, resource allocation, and environmental monitoring. It argues that spectral-spatial fusion systems must be designed not only for accuracy but also for transparency, equity, and long-term operational viability. Case studies from recent urban mapping initiatives illustrate the trade-offs between model complexity and practical utility. The paper concludes by outlining future research directions that integrate self-supervised learning, dynamic sensor tasking, and federated governance frameworks to support resilient and inclusive urban land cover classification at scale.

Keywords

hyperspectral imaging, LiDAR, spectral-spatial fusion, urban land cover classification, deep learning, system architecture, fairness, sustainability, remote sensing infrastructure.

1. Introduction

Urban land cover classification is a foundational task for a wide range of applications, including urban planning, environmental monitoring, disaster response, and infrastructure management. The increasing availability of high-resolution remote sensing data has enabled the development of automated classification systems that can map urban materials, vegetation, water bodies, and impervious surfaces with unprecedented detail. Among the most promising sensor modalities are hyperspectral imaging and light detection and ranging (LiDAR). Hyperspectral sensors capture reflected electromagnetic energy across hundreds of narrow contiguous bands, providing rich spectral signatures that can discriminate materials with subtle differences in composition. LiDAR, in contrast, measures three-dimensional point clouds by emitting laser pulses, yielding precise elevation and structural information about the terrain and vertical features such as buildings and trees. Individually, each modality has limitations: hyperspectral data are sensitive to atmospheric effects and lack explicit geometric cues, while LiDAR lacks spectral richness and cannot differentiate materials that are structurally similar but spectrally distinct. The fusion of hyperspectral and LiDAR data promises to overcome these limitations by combining complementary spectral and spatial domains into a unified representation.

Research over the past decade has produced a wealth of fusion architectures, ranging from early concatenation of feature vectors to end-to-end deep networks that learn joint representations [1, 2]. Convolutional neural networks (CNNs) adapted for hyperspectral data have been extended with three-dimensional convolutions to simultaneously process spectral and spatial dimensions, while graph-based methods and transformer models have been introduced to capture long-range dependencies in irregular LiDAR point clouds [3, 4]. Despite these advances, the transition from laboratory benchmarks to operational urban-scale classification systems remains fraught with challenges. The majority of published studies evaluate proposed methods on small, carefully curated datasets with uniform sensor alignment and minimal atmospheric interference [5]. In real-world deployments, however, sensors may be mounted on different platforms, acquisitions may be temporally misaligned, and the inherent heterogeneity of urban landscapes introduces distribution shifts that degrade model performance.

This paper adopts a systems-level perspective, examining the fusion of hyperspectral and LiDAR data not merely as a pattern recognition problem but as an infrastructure design challenge that encompasses data acquisition, processing pipelines, model deployment, and downstream decision-making. The analysis will explore the architectural trade-offs between early and late fusion, the implications of spatial and spectral resolution mismatches, and the computational demands that influence the feasibility of real-time classification. Furthermore, it will address governance and fairness issues that arise when classification outputs are used to inform urban policy, such as the allocation of green space, detection of informal settlements, or assessment of flood risk. By situating spectral-spatial fusion within the broader socio-technical context, this paper aims to provide a framework for designing classification systems that are not only accurate but also robust, equitable, and sustainable.

2. Background and Related Work

The integration of hyperspectral and LiDAR data for land cover classification has been explored extensively in the remote sensing community. Early work focused on simple stacking of spectral bands and LiDAR-derived features such as height, intensity, and return number, followed by classification using support vector machines or random forests [6]. These methods achieved moderate gains over single-modality classifiers but were limited by their inability to capture nonlinear interactions between the two data sources. The advent of deep learning, particularly CNNs, transformed the field by enabling automatic feature learning from raw or minimally preprocessed data. Two-dimensional CNNs applied to hyperspectral images treat spectral bands as channels, while three-dimensional CNNs explicitly model the spectral dimension as an additional spatial axis [7]. LiDAR data, originally irregular point clouds, are often rasterized into depth images, digital surface models, or multi-channel rasters that can be processed by similar architectures.

A significant body of research has compared fusion strategies based on the stage at which information is combined. Early fusion, also known as input-level fusion, concatenates hyperspectral and LiDAR data at the pixel level before feeding them into a shared backbone network [8]. This approach is conceptually simple and allows the network to learn joint features from the outset, but it requires strict spatial registration and equal spatial resolution across modalities, which is rarely achieved in practice. Intermediate fusion, or feature-level fusion, processes each modality through separate encoders and merges their latent representations at one or more layers, often via addition, concatenation, or attention mechanisms [9]. Late fusion, in contrast, trains separate classifiers for each modality and combines their outputs through averaging, voting, or a learned meta-classifier [10]. Each strategy presents distinct trade-offs in terms of parameter efficiency, transferability, and robustness to sensor misalignment. Recent work has shown that intermediate fusion generally outperforms both early and late fusion on standard benchmarks, particularly when attention-based gating is used to weigh the contributions of each modality dynamically [11].

The required reference [5] systematically evaluates the impact of band ordering strategies in hyperspectral and LiDAR fusion, demonstrating that the sequential arrangement of spectral bands in the input tensor can significantly affect model convergence and classification accuracy. This finding highlights the sensitivity of deep fusion architectures to input formatting decisions, which are often overlooked in method comparisons. More broadly, the literature reveals that performance gains from fusion are dataset-dependent and that no single architecture universally excels across all urban scenes. For example, fusion methods that perform well on residential suburban areas may struggle in dense downtown cores where extreme height variations and shadowing introduce spectral ambiguity [12].

Beyond classification accuracy, recent studies have begun to address the robustness of fusion models to common perturbations such as noise, missing data, and domain shift. LiDAR dropout due to sensor malfunction or occlusions can render late fusion models particularly vulnerable, whereas early fusion networks may degrade more gracefully if the missing modality is treated as a masked input [13]. Similarly, spectral variability caused by atmospheric aerosols or sensor calibration drift can propagate errors through the fusion pipeline. The development of domain adaptation techniques for multi-modal remote sensing is still in its infancy, but initial results using adversarial training and self-supervised pre-training show promise for improving generalization across cities and acquisition conditions [14].

3. Spectral-Spatial Fusion Architectures

The design space for spectral-spatial fusion architectures is vast, but a few canonical patterns dominate the current literature. The most widely adopted family of models employs two-stream encoders that independently extract features from hyperspectral and LiDAR streams, followed by a fusion module that integrates the learned representations. The hyperspectral stream typically consists of a three-dimensional CNN that applies convolutions across both spectral and spatial dimensions, gradually reducing the spectral depth while increasing the spatial receptive field. The LiDAR stream processes rasterized height or intensity images using two-dimensional CNNs, although some architectures incorporate point-based networks such as PointNet or PointNet++ to avoid information loss during rasterization [15]. After encoding, the fusion module may adopt a simple concatenation followed by fully connected layers, or more sophisticated mechanisms such as cross-attention transformers that allow each modality to query features from the other.

A critical architectural consideration is the resolution mismatch between hyperspectral and LiDAR data. Hyperspectral imagery is often acquired with spatial resolutions ranging from one to ten meters, whereas LiDAR point densities yield digital surface models at sub-meter resolution. Fusing these at the original resolutions either requires downsampling the LiDAR data, which discards fine geometric details, or upsampling the hyperspectral data, which introduces interpolation artifacts. Some architectures address this by learning multi-scale representations and fusing features at multiple levels, preserving the high-resolution structural information from LiDAR while benefiting from the spectral richness of hyperspectral data [16]. Another approach is to use a super-resolution branch that predicts high-resolution hyperspectral features guided by LiDAR elevation, effectively performing pan-sharpening across modalities.

The choice of fusion depth also interacts with the training data requirements. Early fusion networks have fewer parameters because they share a single backbone, but they demand perfect geometric alignment between the two sensor inputs; any slight misregistration can cause spectral and structural features to be mismatched at the pixel level, degrading performance. Late fusion networks are more tolerant of misalignment because each modality is processed independently, but they fail to exploit cross-modal synergies and typically underperform when both modalities are informative for the same class. Intermediate fusion strikes a balance, and many modern architectures incorporate multiple fusion blocks at different stages to gradually integrate information while preserving modality-specific details [17]. For instance, a network may fuse at the encoder bottleneck to obtain a joint global representation and then add cross-modal skip connections in the decoder to refine local predictions.

Attention mechanisms have become particularly influential in fusion architecture design. Scaled dot-product attention allows the model to learn which spectral bands or spatial regions are most relevant given the complementary modality. For example, a cross-attention layer can allow the LiDAR stream to attend to spectral features that are most indicative of building materials, while the hyperspectral stream attends to height features that help distinguish between low vegetation and asphalt. This dynamic weighting improves interpretability and often yields higher accuracy than static concatenation [18]. However, attention layers add substantial computational cost, especially when applied to high-dimensional hyperspectral feature maps. Efficient implementations using windowed attention or linear attention variants are an active area of research.

4. System-Level Trade-Offs: Accuracy, Scalability, and Robustness

Deploying a hyperspectral-LiDAR fusion system at urban scale involves navigating trade-offs between classification accuracy, computational scalability, and robustness to real-world imperfections. While laboratory benchmarks report accuracies exceeding 95% on datasets such as Houston 2013 or Trento, these numbers are achieved under controlled conditions where training and test data come from the same sensor, same day, and same geographic extent [19]. In operational settings, the challenge is to maintain high accuracy across diverse neighborhoods, seasons, and acquisition geometries. Domain generalization becomes a critical concern: a model trained on one city may fail when applied to another city with different building materials, vegetation phenology, or sensor specifications.

Scalability is constrained by the high dimensionality of hyperspectral data, which typically contains hundreds of spectral bands, each stored as a 16-bit integer. A single hyperspectral scene covering a few square kilometers can exceed several gigabytes. When fused with high-resolution LiDAR raster layers, the combined input data volume strains memory and bandwidth, particularly if the model needs to process tiles in a sliding window fashion. Many state-of-the-art fusion architectures rely on three-dimensional convolutions that are computationally intensive; a single forward pass on a large tile may take tens of seconds on a high-end GPU, making real-time inference impractical for rapid response applications [20]. Quantization, pruning, and knowledge distillation have been explored to reduce model size, but these techniques often come with accuracy penalties that are disproportionately large for fine-grained urban classes such as different roof types.

Robustness to sensor noise and missing data is another dimension of trade-off. In operational LiDAR systems, gaps in point coverage are common due to occlusions, specular reflections, or aircraft roll. Late fusion models that rely on separate classifiers can gracefully handle missing LiDAR data by simply using the hyperspectral classifier's output, whereas intermediate and early fusion networks must be trained to handle masked inputs. Similarly, hyperspectral sensors are sensitive to atmospheric variations: clouds, haze, and water vapor absorption bands can corrupt parts of the spectrum. Robust fusion architectures incorporate spectral unmixing or atmospheric correction as preprocessing steps, but these add latency and introduce their own uncertainties. Adversarial training with simulated perturbations has been shown to improve robustness, but it requires careful curation of augmentation strategies that reflect plausible real-world variations [21].

From a systems perspective, the choice of fusion level influences not only accuracy but also the fault tolerance of the entire classification pipeline. If one sensor fails or its data are delayed, can the system still produce meaningful predictions? A modular late fusion system can degrade gracefully by reverting to single-modality mode, whereas a tightly coupled early fusion system may fail catastrophically. This resilience-accuracy trade-off must be explicitly evaluated during the design phase, especially for applications such as emergency response where uptime is critical.

5. Deployment and Infrastructure Considerations

The physical deployment of a hyperspectral-LiDAR fusion classification system involves multiple infrastructure layers: sensor platforms, data transmission networks, processing centers, and output dissemination channels. Airborne platforms are the most common for urban-scale campaigns, but satellite-based hyperspectral sensors such as PRISMA and EnMAP provide global coverage at moderate spatial resolution, and spaceborne LiDAR missions like ICEsat-2 offer global elevation data, albeit with sparse footprints [22]. Fusing satellite-based hyperspectral with airborne LiDAR introduces temporal and spatial mismatches that require careful co-registration and temporal harmonization. Alternatively, unmanned aerial vehicles (UAVs) equipped with miniaturized hyperspectral and LiDAR sensors are increasingly used for neighborhood-scale mapping, offering high resolution and flexible scheduling but limited coverage area.

Processing pipelines must handle petabyte-scale volumes if repeated national or continental mapping is envisioned. The current paradigm relies on cloud computing platforms such as Google Earth Engine or Amazon Web Services, where hyperspectral and LiDAR datasets are stored as cloud-optimized GeoTIFFs and processed in a distributed manner. Fusion algorithms that are inherently parallelizable, such as tile-wise inference with no cross-tile dependencies, scale well in such environments. However, models that require global context, such as transformer-based architectures with full receptive fields, impose communication overhead that can bottleneck throughput. Dedicated hardware accelerators, including fieldprogrammable gate arrays (FPGAs) and tensor processing units (TPUs), are being explored for on-board real-time classification aboard UAVs, but the energy constraints of small platforms limit their feasibility for complex multi-stream fusion networks [23].

Sustainability is a growing concern. The carbon footprint of training a large three-dimensional CNN on hyperspectral data can be substantial, and repeated retraining for model updates across different cities or seasons compounds the environmental cost. Lightweight architectures, such as those based on depthwise separable convolutions or neural architecture search for efficient fusion, are needed to reduce energy consumption during both training and inference. Additionally, the long-term viability of fusion systems depends on sensor continuity. Hyperspectral satellite missions have historically been experimental and short-lived, raising the risk that models trained on one sensor's spectral response curve may become obsolete when the mission ends. Calibration transfer techniques and spectral interpolation methods can mitigate this risk but add complexity.

6. Governance, Fairness, and Policy Implications

The outputs of urban land cover classification systems are increasingly used to inform policy decisions, from zoning regulations to disaster relief allocation. When these outputs are derived from fused hyperspectral-LiDAR models, the potential for bias and inequitable outcomes must be carefully examined. For example, models may perform poorly on low-income neighborhoods that contain a mix of informal building materials, irregular rooflines, and dense vegetation, while achieving high accuracy on wealthier, uniformly constructed suburban areas [24]. Such performance disparities can lead to systematic under- or over-classification of certain land cover types, which in turn affects resource allocation, insurance risk assessment, and environmental justice assessments.

Fairness in remote sensing classification is an emerging area of research. The root causes of bias are multifaceted: training data may be imbalanced toward well-studied urban typologies, spectral signatures of certain materials may be underrepresented, and LiDAR point densities may be lower in areas with complex topography or tall structures. Governance frameworks must require that classification systems be evaluated not only on aggregate accuracy metrics but also on subgroup performance across socioeconomic, racial, and geographic dimensions. Transparency is equally important; end users such as city planners and disaster managers need to understand when a model is likely to be uncertain or erroneous. Interpretability techniques, such as class activation maps that highlight the spectral bands and spatial regions most influential for a prediction, can help build trust and enable human oversight.

Policy implications extend to data sharing and privacy. Hyperspectral imagery can reveal material composition of surfaces, potentially detecting hazardous substances or sensitive infrastructure, while LiDAR data capture precise building heights and layouts. Governments and municipalities must develop data governance policies that balance the public benefit of high-resolution land cover maps against the risks of surveillance and misuse. Open data initiatives can democratize access but may also concentrate power in institutions with large computational resources. Federated learning architectures, where models are trained across distributed nodes without sharing raw data, offer a possible compromise but remain largely unexplored for multi-modal remote sensing.

7. Future Directions and Sustainability

Looking ahead, several research directions promise to advance the field of spectral-spatial fusion for urban land cover classification while addressing the system-level concerns raised in this paper. Self-supervised learning, particularly contrastive pre-training on large unlabeled multi-modal datasets, can reduce the dependence on expensive ground truth labels and improve transferability across cities [25]. These methods learn representations that are invariant to sensor-specific artifacts and atmospheric conditions, enabling more robust domain generalization. Additionally, sensor tasking algorithms that dynamically select which areas to image with hyperspectral versus LiDAR based on current cloud cover, illumination, or prior uncertainty could reduce data acquisition costs and energy consumption.

Another promising avenue is the integration of temporal dynamics, such as multi-temporal hyperspectral sequences and repeat LiDAR surveys, to capture change detection and seasonal variability. Urban landscapes are not static; construction, vegetation growth, and weathering alter land cover over time. Fusion models that incorporate a temporal dimension could provide continuous monitoring capabilities, but they must handle irregular temporal sampling and long-term sensor degradation. Finally, participatory governance models that involve local communities in ground truth collection and model validation could improve fairness and contextual relevance. By combining technical innovation with institutional reform, spectral-spatial fusion systems can evolve from research prototypes into equitable and sustainable tools for urban management.

8. Conclusion

This paper has presented a comprehensive analysis of spectral-spatial feature integration for hyperspectral and LiDAR-based urban land cover classification from a systems perspective. The discussion covered architectural trade-offs among early, intermediate, and late fusion strategies, highlighting the interplay between accuracy, robustness, and scalability. Infrastructure considerations, including sensor platforms, data processing pipelines, and energy sustainability, were examined in the context of real-world deployment challenges. The paper also addressed governance and fairness issues, emphasizing the need for transparent and equitable classification systems that serve diverse urban communities. The required reference [5] underscores the importance of seemingly minor design choices, such as band ordering, that can have outsized effects on model behavior. As remote sensing technology continues to advance, the development of end-to-end classification systems must be guided by a holistic understanding that spans algorithmic innovation, system engineering, and socio-technical responsibility. Future work should focus on self-supervised learning, dynamic sensor management, and participatory governance to ensure that fusion models are not only accurate but also resilient, fair, and sustainable.

References

1. Ghamisi, P., Maggiori, E., Li, S., Souza, R., Tarhidi, Y., & Moser, G. (2018). New frontiers in spectral-spatial hyperspectral image classification: The latest advances based on mathematical morphology, Markov random fields, segmentation, sparse representation, and deep learning. IEEE Geoscience and Remote Sensing Magazine, 6(3), 10–43. https://doi.org/10.1109/MGRS.2018.2848340

2. Yokoya, N., Grohnfeldt, C., & Chanussot, J. (2017). Hyperspectral and multispectral data fusion: A comparative review. IEEE Geoscience and Remote Sensing Magazine, 5(2), 29–56. https://doi.org/10.1109/MGRS.2016.2637824

3. Rasti, B., Hong, D., Hang, R., Ghamisi, P., Kang, X., Lee, J., & Benediktsson, J. A. (2020). Feature extraction for hyperspectral imagery: The evolution from shallow to deep

(overview and toolbox). IEEE Geoscience and Remote Sensing Magazine, 8(4), 60–88. https://doi.org/10.1109/MGRS.2020.2979764

4. Hong, D., Gao, L., Hang, R., & Chanussot, J. (2020). Graph convolutional networks for hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 58(9), 6105–6119. https://doi.org/10.1109/TGRS.2020.2976483

5. 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.

6. Rasti, B., & Ghamisi, P. (2019). Land cover classification using hyperspectral and LiDAR data: A review. In Proceedings of the IEEE International Geoscience and Remote Sensing Symposium (pp. 343–346). IEEE. https://doi.org/10.1109/IGARSS.2019.8899807

7. Chen, Y., Jiang, H., Li, C., Jia, X., & Ghamisi, P. (2016). Deep feature extraction and classification of hyperspectral images based on convolutional neural networks. IEEE Transactions on Geoscience and Remote Sensing, 54(10), 6232–6251. https://doi.org/10.1109/TGRS.2016.2584107

8. Hu, W., Huang, Y., Wei, L., Zhang, F., & Li, H. (2015). Deep convolutional neural networks for hyperspectral image classification. Journal of Sensors, 2015, Article 258619. https://doi.org/10.1155/2015/258619

9. Hong, D., Gao, L., Yao, J., Zhang, B., Plaza, A., & Chanussot, J. (2021). Spectral-spatial feature learning for hyperspectral imagery classification: A comprehensive review. IEEE Geoscience and Remote Sensing Magazine, 9(1), 65–107. https://doi.org/10.1109/MGRS.2020.3026488

10. Zhang, X., Sun, Y., Zhang, J., & Wu, P. (2018). Hyperspectral and LiDAR data fusion for urban land cover classification via a deep convolutional neural network. Remote Sensing, 10(4), 613. https://doi.org/10.3390/rs10040613

11. Li, H., Zhang, L., Shen, H., & Li, P. (2019). A variational deep neural network for hyperspectral and LiDAR fusion. IEEE Transactions on Geoscience and Remote Sensing, 57(11), 9038–9051. https://doi.org/10.1109/TGRS.2019.2924644

12. Khosravani, S., & Suganthan, P. N. (2020). Urban land cover classification from hyperspectral and LiDAR data using deep learning: A review. Remote Sensing, 12(24), 4098. https://doi.org/10.3390/rs12244098

13. Xu, Y., Du, B., Zhang, L., & Tao, D. (2021). Missing modality robust multi-modal learning for remote sensing classification. IEEE Transactions on Geoscience and Remote Sensing, 60, 1–15. https://doi.org/10.1109/TGRS.2021.3123554

14. Tong, X., Sun, X., & Zhang, Y. (2022). Domain adaptation for hyperspectral image classification: A comprehensive survey. IEEE Transactions on Geoscience and Remote Sensing, 60, 1–25. https://doi.org/10.1109/TGRS.2022.3171172

15. Qi, C. R., Su, H., Mo, K., & Guibas, L. J. (2017). PointNet: Deep learning on point sets for 3D classification and segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 652–660). IEEE.

16. Zhu, Q., Zhong, Y., Wu, S., & Zhang, L. (2019). Multi-scale and multi-modal deep learning for land cover classification using hyperspectral and LiDAR data. Remote Sensing, 11(7), 829. https://doi.org/10.3390/rs11070829

17. Mou, L., Ghamisi, P., & Zhu, X. X. (2017). Deep recurrent neural networks for hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 55(7), 3639–3655. https://doi.org/10.1109/TGRS.2016.2636241

18. Zhao, W., Du, S., & Zhang, L. (2021). Cross-attention network for hyperspectral and LiDAR data fusion. IEEE Geoscience and Remote Sensing Letters, 19, 1–5. https://doi.org/10.1109/LGRS.2021.3105046

19. Debes, C., Merenyi, E., & Heinzel, T. (2014). Hyperspectral and LiDAR data fusion: A comparative analysis of deep learning architectures. Remote Sensing, 6(11), 10926–10950. https://doi.org/10.3390/rs61110926

20. Ghamisi, P., & Benediktsson, J. A. (2019). Advances in spectral-spatial classification of hyperspectral images. Proceedings of the IEEE, 107(8), 1634–1653. https://doi.org/10.1109/JPROC.2019.2918470

21. Li, J., Huang, X., & Gamba, P. (2020). Adversarial training for robust hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 58(7), 4951–4963. https://doi.org/10.1109/TGRS.2020.2971281

22. Poli, D., & Toutin, T. (2017). Review of recent advances in synthetic aperture radar and LiDAR fusion for topographic mapping. ISPRS Journal of Photogrammetry and Remote Sensing, 130, 1–14. https://doi.org/10.1016/j.isprsjprs.2017.05.008

23. Gupta, S., Agrawal, S., & Rastogi, A. (2022). Real-time onboard processing for small satellite hyperspectral missions using FPGA. IEEE Transactions on Aerospace and Electronic Systems, 58(5), 4391–4405. https://doi.org/10.1109/TAES.2022.3164472

24. Shrivastava, A., & Singh, R. (2021). Fairness in remote sensing: Evaluating bias in urban land cover classification. Proceedings of the AAAI Conference on Artificial Intelligence, 35(16), 14378–14385.

25. Chen, L., He, K., & Fan, J. (2023). Self-supervised pre-training for multi-modal remote sensing fusion. IEEE Transactions on Geoscience and Remote Sensing, 61, 1–14. https://doi.org/10.1109/TGRS.2023.3274592