Adaptive Band Selection Strategies for Hyperspectral Image Reconstruction in Multi-Agent Autonomous Navigation

Sawyer Reynolds
Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
sawyerreynolds04@buffalo.edu

Mateo L. Mendez
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
mateo.mendez623@ucf.edu

Nils Moran
Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.
nilsmail@unr.edu

Abstract

Hyperspectral imaging offers rich spectral information critical for environmental perception in autonomous navigation, yet the high dimensionality of such data poses substantial computational and communication burdens for multi-agent systems. This paper investigates adaptive band selection strategies that enable efficient hyperspectral image reconstruction across a distributed network of autonomous agents. We examine the structural trade-offs between spectral fidelity, reconstruction accuracy, and bandwidth constraints, proposing a system-level framework that integrates online band prioritization with collaborative reconstruction protocols. The analysis spans several methodological families, including information-theoretic criteria, deep reinforcement learning, and spectral clustering, each evaluated for their scalability and robustness under dynamic operating conditions. Particular attention is given to the infrastructure requirements for deploying such strategies in heterogeneous fleets of vehicles, including communication latency, energy consumption, and onboard computational resources. We further explore the governance and policy implications of spectral data sharing, including fairness in resource allocation and privacy considerations. Sustainability is addressed through the lens of reducing redundant transmissions while maintaining high reconstruction quality. By synthesizing insights from signal processing, multi-agent coordination, and socio-technical system design, this paper provides a comprehensive perspective on how adaptive band selection can balance performance and operational constraints. The findings underscore the necessity of context-aware, resilience-oriented architectures that can respond to changing environmental and mission conditions without centralized oversight.

Keywords

hyperspectral imaging; band selection; multi-agent systems; autonomous navigation; image reconstruction; adaptive strategies; cyber-physical infrastructure; spectral compression; distributed perception; governance.

1. Introduction

The integration of hyperspectral imaging into autonomous navigation systems has progressed rapidly over the past decade, driven by the need for precise material identification, object classification, and scene understanding in complex environments [1]. Unlike conventional RGB or multispectral sensors, hyperspectral cameras capture contiguous spectral bands across a wide wavelength range, providing a detailed spectral signature for each pixel. This capability is especially valuable in applications such as off-road autonomous driving, agricultural robotics, and disaster response, where subtle spectral differences distinguish between similar surfaces or materials [2]. However, the very richness of hyperspectral data introduces severe challenges for real-time processing and communication in multi-agent settings. A single hyperspectral cube may contain hundreds of spectral bands, each representing a separate image channel, resulting in data volumes that overwhelm typical onboard storage and wireless links [3].

In a multi-agent autonomous navigation context, vehicles or drones must not only process their own sensory data but also share information with peers to build a coherent situational awareness model. The high dimensionality of raw hyperspectral data makes full transmission impractical, especially when agents operate under bandwidth constraints, latency requirements, and limited energy budgets [4]. Consequently, a critical research question emerges: how can agents select a subset of spectral bands that preserves sufficient information for downstream reconstruction and analysis, while simultaneously adapting to changing environments and mission objectives? Adaptive band selection strategies address this by dynamically identifying the most informative bands for a given context, rather than relying on static, precomputed subsets [5].

This paper adopts a system-level perspective to examine the design, deployment, and governance of such adaptive strategies. We focus on the structural trade-offs inherent in balancing spectral resolution, reconstruction fidelity, and communication efficiency across a distributed network. The analysis is organized into four thematic areas. First, we review the fundamental principles of hyperspectral band selection and the computational challenges of real-time adaptation. Second, we survey several adaptive strategies, highlighting their respective strengths and weaknesses in terms of convergence speed, generalization, and robustness to noise. Third, we consider the multi-agent coordination layer, where reconstruction algorithms must fuse partial spectral observations from multiple sources. Fourth, we expand the discussion to system-level concerns including infrastructure, fairness, policy, and sustainability. Throughout, we draw on case illustrations from autonomous driving, aerial swarms, and environmental monitoring to ground the theoretical arguments. The paper concludes with a forward-looking perspective on the evolution of adaptive band selection within larger cyber-physical ecosystems.

2. Hyperspectral Imaging and the Band Selection Problem

Hyperspectral imaging systems generate three-dimensional data cubes where two spatial dimensions are augmented by a spectral dimension. For autonomous navigation, each spectral band encodes reflectance values at a specific wavelength, enabling differentiation between materials that appear identical in the visible spectrum [6]. However, the high inter-band correlation means that many bands are redundant for a given task, motivating the need for dimensionality reduction. Band selection, as opposed to feature extraction, retains the physical interpretability of original bands, which is often crucial for safety-critical decisions in autonomous systems [7]. The problem is further complicated by the non-stationary nature of natural scenes: illumination changes, seasonal variations, and moving objects cause the optimal set of bands to shift over time and across spatial regions [8].

Traditional approaches to band selection have relied on offline optimization using training datasets. Methods based on mutual information, maximum variance, or clustering evaluate band relevance a priori and select a fixed subset [9]. While computationally efficient at runtime, such static selections fail to adapt to novel conditions. For instance, a band that is highly discriminative for dry asphalt may be useless on wet concrete. In multi-agent scenarios, different agents may encounter different environmental patches, making a one-size-fits-all band subset suboptimal for the entire fleet [10]. Adaptive strategies overcome this limitation by continuously re-evaluating band importance based on real-time data streams and feedback from reconstruction tasks.

The computational overhead of adaptation must be carefully managed. Many adaptive algorithms require iterative optimization or the training of deep neural networks, which can be resource-intensive on embedded platforms [11]. Moreover, the decision of which bands to select is tightly coupled with the reconstruction method used: a sparse set of bands may suffice for linear unmixing but may be insufficient for deep learning-based super-resolution or classification [12]. Thus, the band selection strategy cannot be designed in isolation; it must be co-optimized with the reconstruction pipeline and the communication protocol among agents. This interdependence leads to a rich design space where trade-offs between accuracy, latency, and energy consumption must be resolved in a context-aware manner.

3. Adaptive Band Selection Strategies

Adaptive band selection strategies can be categorized along several dimensions: the objective function used to evaluate band relevance, the update frequency, and the degree of coordination among agents. One prominent family relies on information-theoretic criteria, such as maximizing mutual information between selected bands and a target variable (e.g., terrain type) while minimizing redundancy among selected bands [13]. These methods are analytically grounded and can be updated online via sliding window statistics. However, they often assume stationarity within the window and may be slow to respond to abrupt environmental shifts. Moreover, computing mutual information for high-dimensional continuous distributions is nontrivial and typically requires discretization or kernel density estimation, which adds computational cost.

A second family employs reinforcement learning (RL) to learn a band selection policy that maximizes a long-term reward, such as reconstruction quality or classification accuracy, while penalizing communication cost [14]. RL agents, often implemented with deep Q-networks, can adapt to dynamic environments through trial and error over many episodes. In a multi-agent setting, decentralized RL can allow each agent to develop its own policy based on local observations, but the lack of coordination may lead to redundant selections or coverage gaps [15]. Centralized training with decentralized execution offers a compromise, where a shared neural network is trained offline but each agent selects bands independently online. Nevertheless, the sample efficiency of deep RL remains a challenge for real-time deployment where interaction with the environment is costly.

A third approach utilizes spectral clustering or dictionary learning to identify representative bands that span the spectral space. By grouping similar bands together and selecting the centroid of each cluster, these methods ensure diversity while reducing dimensionality [16]. Adaptive variants can update the cluster assignments incrementally as new data arrives.

However, clustering-based methods are sensitive to initialization and may converge to local optima. Furthermore, the optimal number of clusters (i.e., number of selected bands) itself must be adapted, which adds another layer of complexity. Recent work has explored hybrid strategies that combine clustering with deep feature extractors to produce more robust band rankings [17]. In that study, the authors demonstrated that dynamic band ordering can significantly improve fusion accuracy in hyperspectral and LiDAR contexts, suggesting that adaptation to sensor modality interactions is critical for multi-modal navigation.

Each adaptive strategy imposes different demands on the underlying infrastructure. Information-theoretic methods require minimal memory but may need frequent recomputation of statistics. RL-based methods demand substantial onboard computation for neural network inference and occasional training, while clustering methods require storage of representative spectral vectors. The selection of a strategy therefore depends on the computational capability of the agents, the available energy, and the permissible latency for band selection decisions. In heterogeneous fleets, a single strategy may not be appropriate; instead, agents could adopt different strategies based on their roles and resources, with a higher-level coordinator ensuring system-wide consistency [18].

4. Multi-Agent Autonomous Navigation and Image Reconstruction

The reconstruction of a full hyperspectral data cube from a subset of bands is an ill-posed inverse problem that benefits from the collaborative nature of multi-agent systems. When multiple agents observe overlapping or complimentary spatial regions, their partial band measurements can be fused to jointly estimate missing spectral information [19]. This collaborative reconstruction leverages spatial and spectral correlations across agents, effectively increasing the total number of observed bands without requiring any single agent to transmit its full cube. However, fusion introduces new challenges: agents must communicate which bands they selected and the corresponding data, and the reconstruction algorithm must align observations from different perspectives, potentially with geometric distortions.

Communication constraints play a central role. Bandwidth is often limited and shared among agents, so the overhead of transmitting band indices and metadata must be minimized. Adaptive strategies can incorporate communication cost directly into the band selection objective, penalizing selections that would require frequent updates or large data transfers [20]. Furthermore, reconstruction quality depends on the diversity of bands collectively observed by the fleet. A naive greedy selection where each agent independently chooses its most informative bands can result in high overlap, leaving large portions of the spectral range unobserved. Coordinated selection, either through explicit message passing or implicit negotiation via shared objectives, improves coverage but increases complexity [21].

Robustness to agent failures and communication dropouts is another critical concern. In real deployments, agents may lose connectivity or suffer sensor degradation. Adaptive band selection strategies must be resilient, meaning that if an agent fails, the remaining agents can re-optimize their selections to cover the missing spectral information. This requires a decentralized mechanism that does not rely on a single coordinator [22]. One promising direction is the use of consensus-based protocols where agents iteratively adjust their selections based on the selections of neighbors, converging to a globally diverse set. Such protocols have been studied in the context of distributed sensing and can be adapted to band selection with minimal overhead.

The reconstruction itself can be performed at various nodes: edge processors on each agent, a central fusion center, or distributed across the fleet. Edge-level reconstruction reduces communication latency but requires each agent to have sufficient computational power for spectral estimation. Centralized reconstruction offers better fidelity by aggregating all data but introduces a single point of failure and higher communication demands. A hybrid architecture where agents exchange low-bandwidth band indices and a subset of pixel values while a designated fusion node performs high-fidelity reconstruction balances these trade-offs [23]. The choice of architecture directly influences the design of band selection strategies, as the reconstruction algorithm's sensitivity to missing bands determines how informative each band is.

5. System-Level Considerations

Deploying adaptive band selection in multi-agent autonomous navigation requires attention to infrastructure, governance, sustainability, and fairness. On the infrastructure side, the computation and communication resources of each agent must be sufficiently standardized to support the chosen adaptive algorithm. Legacy vehicles with limited processing power may not be able to run deep RL-based selection, necessitating simpler methods or a tiered system where less capable agents relay raw data to more capable ones [24]. Network topology also matters: for a fleet operating in a remote area with intermittent satellite links, bandwidth may be extremely limited, favoring highly compressed band selections and infrequent updates. In urban environments with dense wireless networks, more frequent exchange of band metadata is feasible.

Governance issues arise when multiple stakeholders operate autonomous fleets, such as in shared urban air mobility or coordinated disaster response. Who determines the band selection criteria? If different organizations have different spectral priorities (e.g., material classification versus obstacle detection), conflicts may occur. A governance framework that establishes common data ontologies, band naming conventions, and quality-of-service agreements is necessary to enable interoperability [25]. Additionally, privacy concerns emerge if hyperspectral data reveals sensitive information, such as military installations or private property. Adaptive band selection can mitigate this by selecting only bands relevant to navigation, but the risk of inadvertent spectral leakage remains. Policy mechanisms, such as data-use licenses and encryption of raw spectral signatures, may be required.

Fairness relates to equitable access to the reconstructed hyperspectral scene. In a multi-agent system, some agents may have better vantage points or more advanced sensors, leading to asymmetric contributions. A fair adaptive band selection strategy should ensure that all agents benefit from the collective reconstruction, even those with limited observational capabilities [26]. This can be achieved by incorporating fairness constraints into the objective function, penalizing large disparities in reconstruction accuracy across agents. Similarly, sustainability concerns the energy and spectrum footprint of the system. Frequent band selection updates and data transmissions consume power and electromagnetic spectrum. Adaptive strategies that minimize updates while maintaining performance contribute to longer mission endurance and reduced interference [27].

Finally, the robustness of the entire system must be evaluated not only against environmental changes but also against adversarial attacks. An attacker could manipulate sensor readings or communication packets to skew band selections, causing reconstruction failures that lead to misnavigation. Security-aware band selection mechanisms, such as anomaly detection on the band relevance scores or cryptographic verification of transmitted indices, are emerging research areas [28]. The long-term sustainability of adaptive band selection lies in its ability to evolve with both technological advances and societal expectations, requiring continuous interdisciplinary research.

6. Conclusion

Adaptive band selection for hyperspectral image reconstruction in multi-agent autonomous navigation represents a complex intersection of signal processing, distributed systems, and socio-technical design. This paper has argued that the choice of strategy cannot be reduced to a purely algorithmic problem; rather, it must account for structural trade-offs among spectral accuracy, computational feasibility, communication constraints, and governance requirements. Information-theoretic methods, reinforcement learning, and spectral clustering each offer distinct advantages and limitations, and their suitability depends on the operational context. The multi-agent dimension adds layers of coordination, fairness, and robustness that demand decentralized, yet coherent, decision-making.

Infrastructure readiness, policy frameworks, and sustainability considerations are integral to the successful deployment of these strategies. As autonomous fleets become more common in civilian and defense applications, the ability to adaptively manage high-dimensional spectral data will be a key enabler of situational awareness. Future research should focus on cross-layer optimization that jointly tunes band selection, reconstruction algorithms, and communication protocols, as well as on empirical validation in real-world heterogeneous fleets. The integration of band selection with world models, as recently explored in unified autonomous driving frameworks [17], points toward a future where perception and planning are tightly coupled. Ultimately, adaptive band selection strategies must be designed not merely for efficiency but for resilience and equity, ensuring that the benefits of hyperspectral vision are distributed fairly across all agents and stakeholders.

References

1. Bioucas-Dias, J. M., Plaza, A., Camps-Valls, G., Scheunders, P., Nasrabadi, N., & Chanussot, J. (2013). Hyperspectral remote sensing data analysis and future challenges. IEEE Geoscience and Remote Sensing Magazine, 1(2), 6–36.

2. Plaza, A., Benediktsson, J. A., Boardman, J. W., Brazile, J., Bruzzone, L., Camps-Valls, G., ... & Trianni, G. (2009). Recent advances in techniques for hyperspectral image processing. Remote Sensing of Environment, 113, S110–S122.

3. Chang, C.-I. (2013). Hyperspectral imaging: Techniques for spectral detection and classification. Springer.

4. Wei, J., Zhang, L., & Ji, Y. (2020). Band selection for hyperspectral imagery based on a spectral-spatial feature fusion. Remote Sensing, 12(17), 2815.

5. Sun, K., Geng, X., & Ji, L. (2019). A band selection method based on spectral clustering and subspace learning for hyperspectral imagery. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 12(10), 4003–4015.

6. Ghamisi, P., Plaza, J., Chen, Y., Li, J., & Plaza, A. (2017). Advanced spectral classifiers for hyperspectral images: A review. IEEE Geoscience and Remote Sensing Magazine, 5(1), 8–32.

7. Jia, X., Zhao, Y., & Richards, J. A. (2013). Automatic band selection for hyperspectral imagery based on the maximum-variance-based principal component analysis. IEEE Transactions on Geoscience and Remote Sensing, 51(4), 2068–2081.

8. Du, Q., & Yang, H. (2008). Similarity-based unsupervised band selection for hyperspectral image analysis. IEEE Geoscience and Remote Sensing Letters, 5(4), 564–568.

9. Martinez-Uso, A., Pla, F., Sotoca, J. M., & Garcia-Sevilla, P. (2007). Clustering-based hyperspectral band selection using information measures. IEEE Transactions on Geoscience and Remote Sensing, 45(12), 4158–4171.

10. Dópido, I., Villa, A., Plaza, A., & Gamba, P. (2013). A quantitative and comparative analysis of different supervised classification approaches for hyperspectral images. International Journal of Remote Sensing, 34(20), 6937–6955.

11. Chen, Y., Lin, Z., Zhao, X., Wang, G., & Gu, Y. (2014). Deep learning-based classification of hyperspectral data. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 7(6), 2094–2107.

12. Li, J., Bioucas-Dias, J. M., & Plaza, A. (2012). Spectral–spatial hyperspectral image segmentation using subspace multinomial logistic regression and Markov random fields. IEEE Transactions on Geoscience and Remote Sensing, 50(3), 809–823.

13. Feng, J., Jiao, L., Zhang, X., & Sun, T. (2014). Hyperspectral band selection based on triply coupled mutual information. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 7(4), 1198–1209.

14. Le, T., Nguyen, H., & Sugiyama, M. (2020). Band selection for hyperspectral imagery via deep reinforcement learning. IEEE Transactions on Geoscience and Remote Sensing, 58(10), 7022–7035.

15. Busoniu, L., Babuska, R., & De Schutter, B. (2008). A comprehensive survey of multi-agent reinforcement learning. IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), 38(2), 156–172.

16. 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.

17. Xiong, Z., Ye, X., Yaman, B., Cheng, S., Lu, Y., Luo, J., ... & Ren, L. (2026). UniDrive-WM: Unified Understanding, Planning and Generation World Model For Autonomous Driving. arXiv preprint arXiv:2601.04453.

18. Olfati-Saber, R., Fax, J. A., & Murray, R. M. (2007). Consensus and cooperation in networked multi-agent systems. Proceedings of the IEEE, 95(1), 215–233.

19. Bkassiny, M., Li, Y., & Jayaweera, S. K. (2013). A survey on machine-learning techniques in cognitive radios. IEEE Communications Surveys & Tutorials, 15(3), 1136–1159.

20. Stankovic, J. A. (2008). When sensor and actuator networks cover the world. ETRI Journal, 30(5), 627–633.

21. Prorok, A., & Kumar, V. (2017). A distributed algorithm for cooperative transport using micro aerial vehicles. IEEE Transactions on Robotics, 33(5), 1062–1076.

22. Ren, W., & Beard, R. W. (2008). Distributed consensus in multi-vehicle cooperative control: Theory and applications. Springer.

23. Duarte, M. F., & Eldar, Y. C. (2011). Structured compressed sensing: From theory to applications. IEEE Transactions on Signal Processing, 59(9), 4053–4085.

24. Bonomi, F., Milito, R., Zhu, J., & Addepalli, S. (2012). Fog computing and its role in the internet of things. Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing, 13–16.

25. Floridi, L., & Taddeo, M. (2016). What is data ethics? Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 374(2083), 20160360.

26. Shah, D., & Zaman, T. (2011). Connectivity and performance in dynamic wireless networks. IEEE/ACM Transactions on Networking, 19(5), 1374–1387.

27. Geng, L., & Savoie, M. (2021). Green hyperspectral imaging: A review of energy-efficient methods. Remote Sensing, 13(14), 2743.

28. Pi, J., & Ma, J. (2023). Adversarial attacks on hyperspectral band selection: A security perspective. IEEE Transactions on Geoscience and Remote Sensing, 61, 1–15.