Academic Journal of Applied Sciences
ISSN: 3105-0549 (Print), ISSN: 3105-0557 (Online) | Vol. 2, No. 2, 2026

Cross-Domain Classification for Hyperspectral Imagery: A Comprehensive Review of Domain Adaptation, Domain Generalization, and Multimodal Approaches

Jiashuo Hou

School of Surveying and Land Information Engineering, Henan Polytechnic University, Jiaozuo 454000, China

Abstract: Cross-domain hyperspectral image (HSI) classification has emerged as a critical challenge in remote sensing, driven by the inherent spectral variability across different acquisition conditions, sensors, and geographical regions. Traditional supervised classification methods rely heavily on abundant labeled data, which is expensive and time-consuming to obtain for each new scene. This review provides a comprehensive synthesis of 23 representative works published between 2015 and 2025, organized into three methodological paradigms: (1) domain adaptation (DA)—13 methods that align source and target distributions through statistical matching, adversarial learning, contrastive learning, or active sample selection; (2) domain generalization (DG)—5 methods that learn domain-invariant representations without target data exposure; and (3) cross-modality and graph-based approaches—5 methods that leverage multimodal data or capture nonlocal topological structures. Our cross-cutting synthesis yields four key findings: (i) feature-level alignment alone is insufficient—class-conditional, decision-boundary, and topological alignment each provide complementary signals; (ii) contrastive learning has emerged as a unifying paradigm since 2023, adopted by 5 of the 10 most recent works; (iii) only 2 of 23 methods support open-set transfer, revealing a critical gap between research assumptions and real-world deployment; and (iv) domain generalization methods have narrowed the accuracy gap with DA while offering zero-shot deployment capability. We identify seven research gaps, each assessed for urgency and impact, with foundation model adaptation, benchmark standardization, and open-set domain adaptation rated as near-term, high-impact priorities. This review provides researchers with a structured taxonomy, a 23-method comparative matrix across six design dimensions, and a prioritized roadmap for future investigations in cross-domain remote sensing image classification.

Keywords: Hyperspectral image classification, domain adaptation, domain generalization, cross-scene classification, transfer learning, contrastive learning, graph neural networks, multimodal learning.

1. Introduction

1.1. Topic and Rationale

Hyperspectral imaging captures rich spectral information across hundreds of contiguous narrow bands, enabling fine-grained discrimination of land cover materials that is unattainable with conventional multispectral or RGB imagery. This capability has driven widespread adoption of hyperspectral image (HSI) classification in applications ranging from precision agriculture and mineral exploration to urban planning and environmental monitoring [1, 11].

However, a fundamental challenge undermines the practical deployment of HSI classification systems: spectral domain shift. When a classifier trained on labeled data from one geographical scene (the source domain) is applied to a different scene (the target domain), classification accuracy degrades substantially due to variations in atmospheric conditions, illumination geometry, sensor characteristics, and seasonal phenology. This phenomenon—where the same material exhibits different spectral signatures across acquisition conditions—has been formally characterized as the "hyperspectral heterospectra" problem [3].

The traditional solution—collecting and manually labeling training data for each new scene—is prohibitively expensive. A single HSI scene may contain millions of pixels, and expert annotation requires field surveys and spectroradiometer measurements. The increasing temporal frequency and spatial coverage of hyperspectral acquisitions (from satellites such as PRISMA, EnMAP, and Gaofen-5) have rendered per-scene labeling economically infeasible, motivating the urgent development of cross-domain classification methods that can transfer knowledge from previously labeled scenes to new, unlabeled acquisitions.

1.2. Scope and Boundaries

This review focuses on cross-domain classification methods for hyperspectral imagery, covering the period from 2015 to 2025. We concentrate on three complementary research paradigms:

Domain Adaptation (DA): methods that assume access to both labeled source domain data and (typically unlabeled) target domain data during training, aiming to align their distributions while preserving discriminative information.

Domain Generalization (DG): methods that train exclusively on source domain data and directly generalize to unseen target domains without any target data exposure, suitable for real-time or one-shot deployment scenarios.

Cross-Modality and Graph-Based Approaches: methods that exploit complementary information from heterogeneous data modalities (multispectral, LiDAR, SAR) or capture nonlocal topological structures through graph representations.

We exclude classical single-scene supervised classification methods, physics-based atmospheric correction approaches, and pure data augmentation strategies that do not explicitly model domain shift. We also do not cover active learning methods in isolation, though we discuss their integration with domain adaptation.

1.3. Review Methodology

The literature corpus for this review comprises 23 journal articles published in major remote sensing, computer vision, and machine learning venues including IEEE Transactions on Geoscience and Remote Sensing, IEEE Transactions on Neural Networks and Learning Systems, IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Transactions on Image Processing, ISPRS Journal of Photogrammetry and Remote Sensing, Pattern Recognition, and Neural Networks. Articles were collected through systematic database searches on Semantic Scholar, IEEE Xplore, and Crossref using combinations of the keywords: "domain adaptation," "domain generalization," "cross-scene," "cross-domain," "hyperspectral image classification," "transfer learning," and "contrastive learning."

1.4. Organization of the Review

The remainder of this review is organized as follows. Section 2 covers domain adaptation methods, subdivided into statistical alignment, adversarial learning, contrastive learning, and active-learning-integrated approaches. Section 3 addresses domain generalization methods, organized into representation-learning approaches (conditional alignment, semantic guidance, frequency decomposition) and domain-expansion approaches. Section 4 discusses cross-modality learning, graph-based methods, and general cross-domain approaches developed in the broader machine learning community. Section 5 provides a cross-cutting synthesis with comparative analysis, identifying convergent findings, divergent debates, and methodological observations. Section 6 identifies seven research gaps with urgency and impact assessments and proposes a corresponding research agenda. Section 7 concludes with key takeaways and practical implications.

1.5. Methodological Taxonomy

Before proceeding to the detailed review, we present a taxonomy that organizes the cross-domain HSI classification literature along two axes: the transfer paradigm (domain adaptation vs. domain generalization vs. cross-modality) and the technical mechanism (statistical alignment, adversarial learning, contrastive learning, generative expansion, graph reasoning, and attention-based fusion). Figure 1 provides a visual overview.

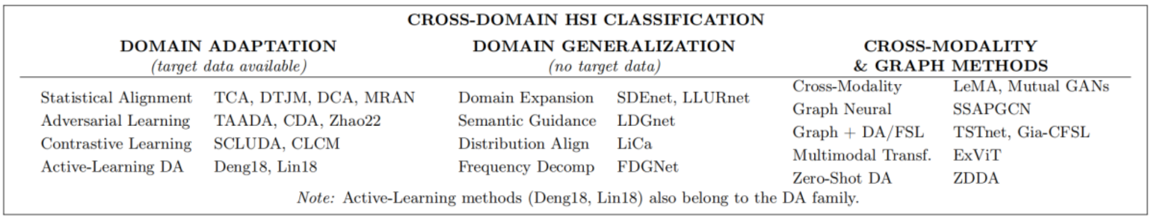

Figure 1. Taxonomic organization of cross-domain HSI classification methods reviewed in this paper. Methods are grouped by transfer paradigm (columns) and technical mechanism (rows), with active-learning-integrated approaches forming a cross-cutting category

Table 1 provides a chronological overview of all 23 reviewed works with their core technical contributions.

Table 1. Chronological Summary of Reviewed Methods

Year	Method	Paradigm	Core Mechanism	Venue
2015	SSTCA (Matasci et al.)	DA	MMD + manifold regularization	IEEE TGRS
2018	AMKDA (Deng et al.)	DA+AL	Multi-kernel MMD + margin sampling	Pattern Recognition
2018	AL-DTL (Lin et al.)	DA+AL	Deep transfer + active salient selection	IEEE JSTARS
2019	DTJM (Peng et al.)	DA	Joint MMD + instance reweighting	IEEE GRSL
2019	MRAN (Zhu et al.)	DA	Multi-representation MMD alignment	Neural Networks
2019	LeMA (Hong et al.)	Cross-Mod	Learnable manifold alignment	ISPRS JPRS
2021	CDA (Liu et al.)	DA	Class-wise adversarial + weighted MMD	IEEE TGRS
2021	DCA (Zhang et al.)	DA	Cooperative subspace+distribution alignment	IEEE TGRS
2021	ZDDA (Kutbi et al.)	DA	Zero-shot common representation	IEEE TPAMI
2021	Mutual GANs (Hong et al.)	Cross-Mod	Self-GAN + mutual-GAN modules	IEEE TGRS
2022	TAADA (Huang et al.)	DA	Two-branch attention adversarial	IEEE TGRS
2022	MLDB-UDA (Zhao C.\ et al.)	DA	Multi-level feature + decision boundary UDA	IEEE TGRS
2023	SSAPGCN (Ding et al.)	Graph	Unified adaptive probability GCN	IEEE TNNLS
2023	SCLUDA (Li et al.)	DA	Supervised contrastive + OCC loss	IEEE TGRS
2023	CLCM (Ning et al.)	DA	Instance contrastive + category matching	IEEE TGRS
2023	LiCa (Gao et al.)	DG	Conditional alignment + error bounds	IEEE TGRS
2023	SDEnet (Zhang et al.)	DG	Domain expansion + contrastive learning	IEEE TIP
2023	LDGnet (Zhang et al.)	DG	Language-guided visual-linguistic alignment	IEEE TGRS
2023	TSTnet (Zhang et al.)	DA+Graph	Graph optimal transport + topological alignment	IEEE TNNLS
2023	LLURnet (Zhao H.\ et al.)	DG	Locally linear unbiased randomization	IEEE TGRS
2023	ExViT (Yao et al.)	Cross-Mod	Vision transformer + cross-modal attention	IEEE TGRS
2024	Gia-CFSL (Zhang et al.)	DA+FSL	Graph aggregation + few-shot alignment	IEEE TNNLS
2025	FDGNet (Qin et al.)	DG	Frequency disentanglement + data geometry	IEEE TNNLS

2. Domain Adaptation for Hyperspectral Image Classification

Domain adaptation (DA) aims to improve the classification performance on a target domain by leveraging labeled data from a related but distributionally different source domain. Unlike standard transfer learning, DA explicitly models and reduces the discrepancy between source and target feature distributions. In the HSI classification context, DA methods address the fundamental challenge that spectral signatures of identical materials shift across scenes due to environmental and sensor-induced variability. As shown in Table 1, DA represents the largest category of reviewed work (13 of 23 papers), spanning statistical alignment, adversarial learning, contrastive learning, and active-learning-integrated approaches.

2.1. Statistical Distribution Alignment Methods

The earliest and most theoretically grounded DA approaches for HSI classification rely on statistical measures of distribution discrepancy, primarily maximum mean discrepancy (MMD) and its variants, to explicitly align source and target feature distributions in a learned embedding space.

Transfer Component Analysis (TCA). Matasci et al. (2015) provided one of the first systematic investigations of DA for remote sensing image classification. They adapted Transfer Component Analysis—originally developed for general machine learning tasks—to the cross-scene HSI classification problem. TCA learns a feature transformation that minimizes the MMD between source and target domain representations in a reproducing kernel Hilbert space (RKHS). The semi-supervised variant (SSTCA) further incorporates label information from the source domain to preserve class discriminability. Experimental evaluations on multi- and hyperspectral acquisitions demonstrated that SSTCA achieved remarkable cross-image classification performance compared to both linear and kernel-based feature extraction baselines. This work established the feasibility of statistical domain alignment for remote sensing and identified two critical objectives that subsequent methods would build upon: alignment with source labels and preservation of local data structure. However, TCA's reliance on a fixed kernel function limits its adaptability—the kernel must be chosen a priori and cannot adapt to the specific characteristics of HSI data (e.g., narrow spectral bands with strong correlation). This limitation motivated the learned-graph and deep adversarial approaches developed in subsequent years (Sections 2.2 and 4.1).

Discriminative Transfer Joint Matching (DTJM). Building on the statistical alignment paradigm, Peng et al. (2019) proposed DTJM, which simultaneously addresses three objectives in a unified optimization framework: (1) feature matching via empirical MMD minimization in kernel PCA space, (2) instance reweighting through an $\ell_{2,1}$ -norm regularized embedding matrix that identifies and down-weights irrelevant source samples, and (3) manifold regularization that preserves the local geometric structure of both domains while maximizing the statistical dependence between the learned embedding and source labels. Experiments on benchmark HSI datasets demonstrated DTJM's superiority over feature extraction techniques operating without domain alignment, highlighting the importance of joint instance- and feature-level adaptation.

Class-Wise Distribution Adaptation (CDA). Liu et al. (2021) advanced the MMD-based paradigm by introducing class-wise adversarial adaptation. The CDA network employs adversarial learning between a feature extractor and multiple per-class domain discriminators to generate domain-invariant features at the class level. A key innovation is the integration of a probability-prediction-based MMD method that achieves superior feature alignment by weighting the MMD computation with the classifier's confidence scores. This class-wise distribution adaptation, operationalized through both adversarial training and class-weighted MMD, enables unsupervised classification of target images without requiring any target domain labels. Experimental validation using Hyperion and AVIRIS hyperspectral data confirmed CDA's effectiveness.

Discriminative Cooperative Alignment (DCA). Zhang et al. (2021) identified a fundamental limitation of existing subspace-based DA methods: the assumption that a shared subspace exists between source and target domains. When marginal distributions are substantially different, this assumption breaks down. DCA addresses this through cooperative alignment of both subspace geometry and data distribution. The framework performs geometrical alignment of subspaces and statistical alignment of distributions simultaneously, while preserving class-discriminative information in the original feature space. A reconstruction constraint enhances the robustness of the subspace projection against overfitting to noise. Cross-scene experiments on three HSI datasets confirmed DCA's significant improvement over state-of-the-art domain-adaptive methods, with the cooperative strategy proving more robust than purely geometrical or purely statistical alignment approaches.

2.2. Adversarial Domain Adaptation

Adversarial DA methods adopt the generative adversarial network (GAN) paradigm, pitting a feature extractor (generator) against a domain discriminator that attempts to distinguish source from target domain representations. The adversarial objective drives the extractor to produce features that are both discriminative for the classification task and invariant to domain identity.

Two-Branch Attention Adversarial DA (TAADA). Huang et al. (2022) observed that most deep HSI DA approaches directly apply deep networks without explicitly exploiting the rich spectral and spatial information encoded in hyperspectral data. TAADA addresses this gap through a two-branch feature extraction (TBFE) subnetwork that separately processes spectral and spatial dimensions using attention mechanisms. The generator produces attention-based spectral-spatial features, while a discriminator based on dual classifiers with FC-BN-ReLU-Dropout structure engages in adversarial learning. The adversarial interplay simultaneously improves discriminative feature extraction and cross-domain classification, with the network learning to adjust source-target distributions to extract domain-invariant features. Three cross-scene HSI classification tasks validated TAADA's superiority over existing DA methods, with the attention mechanisms proving particularly effective for capturing the fine-grained spectral-spatial patterns characteristic of HSI data.

Multi-Level Features and Decision Boundaries UDA. Zhao et al. (2022) addressed the often-overlooked issue of decision boundary alignment in DA for HSI classification. While most methods focus exclusively on feature-level distribution alignment, Zhao et al. argued that domain shift also manifests at the classifier decision boundary. Their unsupervised DA method jointly aligns features at multiple network levels while explicitly regularizing the decision boundaries learned from source and target domains. This multi-granularity approach captures domain discrepancies at different semantic levels—from low-level spectral features to high-level semantic representations—while the decision boundary alignment ensures that the classifier's separation surfaces generalize across domains.

2.3. Contrastive Learning-Based Domain Adaptation

The most recent paradigm shift in HSI DA has been the integration of contrastive learning—a self-supervised representation learning technique that has revolutionized computer vision. Rather than merely reducing inter-domain distance, contrastive DA methods explicitly enforce intra-class compactness and inter-class separability within each domain.

Supervised Contrastive Learning-Based UDA (SCLUDA). Li et al. (2023) identified a critical shortcoming of existing DA methods: while they align inter-domain distributions, they often neglect the intra-domain separability of source and target data. SCLUDA introduces supervised contrastive learning in both source and target domains, pulling samples from the same category closer while pushing samples from different categories apart. A novel domain similarity loss cast as a one-class classification (OCC) problem reduces inter-domain discrepancy without collapsing class structure. Furthermore, a confidence-learning-based sample selection strategy identifies high-confidence pseudo-labeled samples from the target domain to fine-tune the adaptation model, progressively enhancing target domain discrimination. The source code is publicly available, facilitating reproducibility.

Contrastive Learning with Category Matching (CLCM). Ning et al. (2023) independently explored whether contrastive learning—whose core principle of pulling similar samples together while pushing dissimilar ones apart is conceptually analogous to domain alignment—could achieve effective cross-scene HSI classification. CLCM operates at the instance level, treating category information as a structural prior in feature space. Each source sample serves as an anchor; positive matches are target samples of the same (pseudo-labeled) category, and negative matches are samples of different categories. Instance-level discriminative embeddings emerge through the attraction of positive pairs and repulsion of negative pairs across domains. To enhance contrastive quality, CLCM incorporates spectral-spatial feature extraction for more accurate semantic representation and screens high-confidence target samples for iterative network refinement. Results on three DA tasks demonstrated the effectiveness and feature discriminativeness of contrastive learning for cross-scene HSI classification, providing new methodological directions beyond traditional MMD and adversarial approaches.

2.4. Active Learning-Integrated Domain Adaptation

A complementary line of research integrates active learning with domain adaptation to achieve maximal classification performance with minimal target domain annotation effort.

Active Multi-Kernel Domain Adaptation. Deng et al. (2018) proposed one of the earliest frameworks combining active learning with domain adaptation for HSI classification. The core idea is to retrain a multi-kernel classifier using available labeled samples from the source domain while actively querying the most informative samples in the target domain for manual annotation. The method adaptively combines multiple kernels to form a DA classifier that minimizes the distribution bias between domains. An inner active updating process sequentially expands the training set with margin sampling strategy, gradually converging to satisfactory classification performance. Experiments on two popular HSI datasets demonstrated the approach's effectiveness, establishing the active DA paradigm as a practical solution when a small annotation budget is available for the target domain.

Active-Learning-Incorporated Deep Transfer Learning. Lin et al. (2018) concurrently introduced two novel contributions: (1) for the first time in HSI transfer learning, active learning initializes salient samples on the source data to determine what knowledge is worth transferring, and (2) higher-level features are constructed and connected across source and target domains to further bridge cross-domain disparity. Unlike most DA methods that require prior knowledge about the target domain, Lin et al.'s framework assumes no such knowledge and works for both homogeneous (same sensor) and heterogeneous (different sensors) HSI data. Experiments on three real-world HSIs confirmed effectiveness, with the salient sample selection mechanism proving critical for transfer quality [9].

3. Domain Generalization for Hyperspectral Image Classification

While domain adaptation methods assume access to target domain data (labeled or unlabeled) during training, domain generalization (DG) addresses the more challenging and practically relevant scenario where the target domain is entirely unseen during model development. A DG model, trained exclusively on one or more source domains, is expected to generalize directly to previously unobserved target domains. For HSI classification, DG is particularly attractive for real-time applications (e.g., disaster response, onboard satellite processing) where retraining on each new acquisition is infeasible. Five of the 23 reviewed works (Table 1) fall under the DG paradigm, which we organize here into two broad strategies: representation-learning approaches that seek domain-invariant features through distribution alignment, semantic guidance, or frequency decomposition; and domain-expansion approaches that generate synthetic intermediate domains to simulate the diversity a model might encounter at test time.

3.1. Representation-Learning Approaches: Alignment, Semantics, and Frequency

Rather than generating new domains, representation-learning DG methods aim to discover feature spaces or representational modalities that are inherently robust to domain variation.

LiCa: Conditional Distribution Alignment with Theoretical Guarantees. Gao et al. (2023) provided a formal mathematical treatment of the HSI domain generalization problem by characterizing hyperspectral heterospectra through the lens of conditional probability distributions. The key insight is that spectra of the same material across different acquisition conditions can be treated as samples from different domains, and the goal is to align the conditional distributions $P(X|Y)$ across domains. The LiCa block introduces two novel loss functions: (1) cross-domain conditional alignment, which directly matches the spectral conditional distributions of different domains for each class, and (2) cross-domain entropy (CdE), which quantifies the heterogeneity of HSI data across domains. Importantly, Gao et al. provided theoretical grounding by deriving an upper bound on the classification error in any unseen target domain, connecting the empirical alignment objectives to formal generalization guarantees—this makes LiCa one of the few theoretically grounded methods in the entire cross-domain HSI literature. Comparative experiments against pixel-wise classification methods demonstrated LiCa's superior generalization, establishing conditional distribution alignment as a principled approach to HSI domain generalization.

LDGnet: Language as a Cross-Domain Invariant. Zhang et al. (2023b) introduced a paradigm-shifting approach by leveraging linguistic knowledge as a cross-domain invariant representation. The key observation is that while spectral signatures of "forest" or "water" vary across domains, the semantic concept encoded in language remains stable. LDGnet employs a dual-stream architecture with an image encoder and a text encoder that extracts visual and linguistic features, respectively. Two levels of text representation—coarse-grained (class-level descriptions) and fine-grained (attribute-level descriptions)—capture complementary semantic information. The linguistic features serve as a cross-domain shared semantic space, with visual-linguistic alignment achieved through supervised contrastive learning. This language-guided approach represents a novel direction in HSI DG, connecting remote sensing with the broader vision-language pretraining paradigm. Notably, LDGnet is the only method in the reviewed corpus that exploits vision-language alignment, marking an underexplored but promising research direction.

FDGNet: Frequency-Domain Disentanglement. Qin et al. (2025) proposed the most recent advance in HSI domain generalization by operating in the frequency domain rather than the conventional spatial-spectral domain. FDGNet is motivated by the observation that different frequency components of HSI data encode different types of domain information: low-frequency components tend to capture domain-invariant structural information (e.g., material composition), while high-frequency components encode domain-specific details (e.g., sensor noise, atmospheric effects). By explicitly disentangling these frequency components, FDGNet learns representations that isolate domain-invariant information. The data geometry component further ensures that the learned feature manifold preserves the intrinsic structure of HSI data across domains. This frequency-domain perspective introduces a new dimension to cross-domain HSI analysis—previous methods operate entirely in the spatial-spectral domain and implicitly conflate frequency components with varying degrees of domain sensitivity. FDGNet and LiCa, while operating in fundamentally different representational spaces (frequency vs. probability), share the core principle that domain invariance can be achieved through explicit decomposition rather than implicit adversarial learning [14].

3.2. Domain-Expansion Approaches: Learning from Synthesized Diversity

A distinctive family of DG methods for HSI classification generates synthetic intermediate domains to bridge the gap between source and unseen target distributions. The underlying hypothesis is that exposing a model to diverse domain variations during training—even if artificially generated—will improve its ability to generalize to genuinely novel domains.

SDEnet: Single-Source Domain Expansion Network. Zhang et al. (2023a) developed SDEnet as the first domain-expansion method for HSI classification. Using generative adversarial learning, SDEnet trains a generator comprising a semantic encoder and a morph encoder designed with an encoder-randomization-decoder architecture. Spatial randomization and spectral randomization inject controlled variability into the generated intermediate domains, while implicit morphological knowledge serves as domain-invariant information during expansion. The discriminator employs supervised contrastive learning to learn class-wise domain-invariant representations that pull intra-class samples of source and extended domains together. Adversarial training optimizes the generator to separate intra-class samples, creating a diverse yet class-consistent set of training domains. Experiments on two HSI datasets and one MSI dataset confirmed SDEnet's superiority, with the domain expansion strategy proving effective even from a single source domain—a finding that challenges the intuitive assumption that multi-source diversity is necessary for effective DG.

LLURnet: Unbiased Randomization with Manifold Preservation. Zhao et al. (2023) independently explored the domain expansion strategy with a complementary focus on unbiased feature generation. While SDEnet emphasizes diversity through randomization, LLURnet emphasizes fidelity—the generated domain variants must lie on the natural HSI data manifold. LLURnet employs a symmetric encoder-decoder structure to implicitly construct local joint features under style randomization, generating expanded domain variants without introducing systematic biases. Supervised contrastive learning prevents duplicate augmentation collapse—a common failure mode in generation-based DG where the generator produces trivial variations. An adversarial penalty term, formed by inter-class and intra-class contrastive regularization in the discriminator, balances domain-specific and domain-invariant feature learning. The "locally linear" constraint ensures that generated features respect the intrinsic geometric structure of HSI data. The complementary philosophies of SDEnet (maximize diversity) and LLURnet (preserve fidelity) together define the design space for domain-expansion DG: the optimal balance between diversity and fidelity remains an open question, and no existing work systematically explores this trade-off.

4. Cross-Modality Learning and Graph-Based Approaches

Beyond the core DA and DG paradigms, two complementary research directions have significantly advanced cross-domain HSI classification: leveraging heterogeneous multimodal data and exploiting graph-structured representations.

4.1. Cross-Modality Learning

Cross-modality methods address scenarios where different modalities (e.g., hyperspectral and multispectral) are available for the source and target domains, requiring alignment not only of distributions but also of fundamentally different feature spaces.

Learnable Manifold Alignment (LeMA). Hong et al. (2019) tackled the problem of using limited highly-discriminative hyperspectral data to improve classification with abundant but poorly-discriminative multispectral data. Traditional semi-supervised manifold alignment methods underperform in this setting due to the severe modality gap. LeMA learns a joint graph structure directly from data—rather than relying on a fixed kernel-defined graph—enabling adaptive capture of cross-modality data distribution. Graph-based label propagation with the learned graph finds more accurate decision boundaries. An alternating direction method of multipliers (ADMM) optimization strategy efficiently solves the proposed model. Experiments on two hyperspectral-multispectral datasets demonstrated LeMA's superiority, establishing the learned-graph paradigm as more effective than fixed-kernel manifold alignment for cross-modality remote sensing.

Multimodal GANs for Crossmodal Segmentation. Hong et al. (2021) addressed cross-modality learning through GAN-based architectures, introducing two novel plug-and-play units: (1) a self-GAN module that learns perturbation-insensitive feature representations by generating and discriminating augmented versions of the input, and (2) a mutual-GAN module that eliminates the gap between heterogeneous modalities (HS and MS) through adversarial cross-modality generation. A patchwise progressive training strategy enables effective learning with limited training samples—a common challenge in remote sensing where multimodal annotations are scarce. The self-GAN provides robustness to within-modality perturbations, while the mutual-GAN handles the more severe cross-modality distribution gap.

Extended Vision Transformer (ExViT). Yao et al. (2023) extended the vision transformer architecture to multimodal remote sensing classification with minimal architectural modifications. Processing multimodal image patches through parallel branches of position-shared ViTs extended with separable convolution modules provides an economical solution that leverages both spatial and modality-specific channel information. A cross-modality attention (CMA) module fuses tokenized embeddings from heterogeneous modalities (e.g., HS and LiDAR, or HS and SAR) by exploiting pixel-level spatial correlation. A full token-based decision-level fusion module produces the final classification. ExViT outperformed both CNN-based and transformer-based competitors on Houston2013 (HS+LiDAR) and Berlin (HS+SAR) benchmark datasets, demonstrating that vision transformers can effectively bridge modality gaps when equipped with appropriate cross-modal attention mechanisms.

4.2. Graph-Based Methods

Graph-based approaches model HSI pixels or regions as nodes in a graph, where edges encode spectral-spatial similarity relationships. This representation naturally captures nonlocal dependencies that convolutional approaches—limited by their local receptive fields—cannot.

SSAPGCN: Spatial-Spectral Unified Adaptive Probability GCN. Ding et al. (2023) addressed two limitations of existing GCN-based HSI classification methods: inflexible initial graph structures and the separation between graph construction and feature extraction. The SSAPGCN combines spectral information and spatial coordinates to obtain a spatial-spectral adaptive probability graph structure that captures probabilistic connectivity between homogeneous HSI pixels. Critically, the graph structure and GCN are unified into a single framework that simultaneously learns both the graph and output features through feedback—the output features inform graph refinement, and the refined graph improves features in a closed loop. Evaluation on four public HSI datasets with two-layer SSAPGCN demonstrated superiority over various classification methods, particularly with small training sample sizes where graph-based information propagation compensates for label scarcity.

TSTnet: Topological Structure and Semantic Information Transfer. Zhang et al. (2023c) bridged domain adaptation with graph representation learning through TSTnet. The key insight is that CNNs only model local spatial relationships, missing nonlocal topological relationships that better represent the underlying data structure of HSI. TSTnet employs graph optimal transmission (GOT) to align topological relationships across domains, complementing the distribution alignment achieved through MMD. Subgraphs from source and target domains are dynamically constructed based on CNN features, leveraging CNN discriminative capacity while gaining GCN's nonlocal reasoning ability. A consistency constraint integrates CNN and GCN outputs to ensure coherent predictions. This topological perspective adds a dimension to domain alignment—beyond aligning feature distributions, TSTnet aligns the relational structures between classes.

Gia-CFSL: Graph Information Aggregation Cross-Domain Few-Shot Learning. Zhang et al. (2024) combined few-shot learning (FSL) with graph-based domain alignment for the challenging scenario where the target domain contains new (unseen) classes. Gia-CFSL conducts FSL episodic training with all source domain labels and only a few target domain labels. Two specialized blocks handle intra-domain and cross-domain information: the Intra-domain Distribution Extraction (IDE) block characterizes and aggregates nonlocal relationships within each domain, while the Cross-domain Similarity Aware (CSA) block captures feature and distribution similarities between domains. Feature-level and distribution-level cross-domain graph alignments mitigate domain shift impact on FSL. By operating in a few-shot setting with graph-structured information aggregation, Gia-CFSL addresses a practical gap in existing DA methods that assume identical label sets across domains.

4.3. Cross-Domain Methods from the Broader Machine Learning Community

While Sections 2 and 3 covered methods developed specifically for HSI data, two influential DA methods from the general computer vision and machine learning literature merit discussion. Although neither was originally designed for hyperspectral data, both ZDDA and MRAN embody design principles—zero-shot common representation learning and multi-branch feature alignment, respectively—that have influenced subsequent HSI-specific DA methods (e.g., the multi-representation ideas in TAADA's two-branch architecture and the zero-shot aspirations of SDEnet's domain expansion).

Zero-shot Deep Domain Adaptation (ZDDA). Kutbi et al. (2021) proposed ZDDA to eliminate the need for task-relevant target domain training data entirely—a more extreme setting than unsupervised DA where at least unlabeled target data is available. ZDDA uses paired dual-domain task-irrelevant data to learn common representations for source and target domains. Two variants extend applicability:

ZDDA-C for classification and ZDDA-ML for metric learning tasks. Notably, ZDDA-C handles open-set classification where target domain classes are not fully known during training, addressing a practical limitation of closed-set DA that assumes identical label spaces. As confirmed in Table 2, ZDDA is one of only two reviewed methods supporting open-set transfer, alongside Gia-CFSL (2024). The zero-shot setting represents the ultimate cross-domain transfer challenge—no target data whatsoever—and is increasingly relevant for HSI applications where task-relevant target data may be entirely unavailable, such as classifying a newly imaged region with no historical ground truth.

Multi-Representation Adaptation Network (MRAN). Zhu et al. (2019) observed that single-structure feature extractors capture only partial information from input data (e.g., subset of saturation, brightness, and hue properties in visual imagery). MRAN uses an Inception Adaptation Module (IAM)—a hybrid multi-branch architecture—to extract multiple complementary representations from different receptive fields and processing pathways. Multi-representation alignment captures information from diverse aspects of the data, and an extended MMD formulation computes the adaptation loss across all representation branches simultaneously. The modular design (IAM can be inserted into most feed-forward models) makes MRAN a versatile enhancement to existing architectures. While MRAN was validated on natural image benchmarks rather than HSI data, its core insight—that domain alignment benefits from processing complementary representations in parallel—is conceptually echoed in HSI-specific designs such as TAADA's separate spectral and spatial branches [6, 23].

5. Cross-Cutting Synthesis

5.1. Convergent Findings

Several consistent findings emerge across the reviewed literature:

The insufficiency of feature-level alignment alone. Multiple works [8, 16, 21] independently concluded that purely aligning marginal feature distributions is insufficient for robust cross-domain HSI classification. Class-conditional alignment [3, 10], decision boundary alignment [21], and topological structure alignment [19] each provide complementary adaptation signals that significantly improve upon feature-only alignment.

Contrastive learning as a unifying paradigm. Since 2023, contrastive learning has been independently adopted by multiple research groups for both DA [8, 12] and DG [3, 17, 22] in HSI classification. The conceptual alignment between contrastive learning's push-pull mechanism and domain alignment objectives has proven productive, with contrastive methods consistently outperforming earlier MMD-only and adversarial-only approaches.

Spatial-spectral joint processing is essential. Methods that explicitly model both spatial and spectral dimensions [2, 6, 12, 19] consistently outperform those treating HSI data as 1D spectral vectors. The spatial context provides structural information that is partially domain-invariant—the spatial arrangement of land cover categories follows physical constraints that transcend acquisition conditions.

The value of theoretical grounding. Works that provide theoretical analysis [3, 13, 16] offer more interpretable and trustworthy results than purely empirical contributions. Theory-driven method design tends to produce more robust solutions that generalize beyond the specific experimental settings.

5.2. Divergent Findings and Debates

Domain adaptation vs. domain generalization: which paradigm? The literature is divided on whether DA (requiring target domain data) or DG (training without target data) is the more practical paradigm. DA proponents argue that target domain data is almost always available (even if unlabeled) and that ignoring it sacrifices achievable accuracy. DG proponents [3, 17, 22] counter that requiring target data for retraining undermines the fundamental goal of rapid, zero-shot deployment. The empirical evidence is mixed and resists direct quantification because no standardized benchmark compares DA and DG methods under identical conditions. Within individual papers that include both paradigms as baselines, DG methods typically trail DA by 2–8 percentage points in overall accuracy when target data is available, but the operational benefit—no retraining, no target data collection—can outweigh this gap in time-sensitive deployments. This trade-off cannot be resolved empirically alone; the optimal choice depends on the deployment constraints (latency requirements, computational budget, target data availability), and we consider the absence of a standardized DA-vs-DG comparison protocol a significant methodological gap (see Gap 3, Section 6.1).

Closed-set vs. open-set assumptions. Most DA methods assume identical label sets across domains (closed-set). However, real-world HSI scenes often contain land cover classes unique to specific geographical regions. Kutbi et al. (2021) and Zhang et al. (2024) have begun addressing open-set and partial-set transfer, but these settings remain underexplored relative to their practical importance.

Single-source vs. multi-source generalization. The majority of HSI DG methods [3, 17, 22] operate from a single source domain. While practically convenient, theoretical results from the broader DG literature suggest that multi-source training with diverse domains yields stronger generalization guarantees. Whether domain expansion techniques [17] can effectively substitute for genuine multi-source diversity is an open empirical question.

5.3. Comparative Analysis of Method Characteristics

To facilitate method selection for practitioners and highlight design trade-offs, Table 2 provides a comparative summary of key methodological properties across all reviewed works.

Table 2. Comparative Summary of Methodological Properties

Method	Year	TargetData	Open-Set	Spatial	Code	Complexity
SSTCA	2015	Unsup.	--	--	--	$O(n^2)$ kernel
AMKDA	2018	Fewlabeled	--	--	--	$O(K \cdot n^2)$
AL-DTL	2018	Unsup.	--	Yes	--	$O(n^2)$ deep
DTJM	2019	Unsup.	--	--	--	$O(n^2)$ MMD
MRAN	2019	Unsup.	--	--	--	$O(B \cdot n^2)$
LeMA	2019	Unsup.	--	Lim.	--	$O(n^2)$ ADMM
CDA	2021	Unsup.	--	--	--	O(2n) adv.
DCA	2021	Unsup.	--	--	--	$O(n^2)$ subsp.
ZDDA	2021	Zero-shot	Yes	--	--	$O(n^2)$ pair
MutualGANs	2021	Unsup.	--	Yes	--	O(2n) dual
TAADA	2022	Unsup.	--	Yes	--	O(2n) adv.
MLDB-UDA	2022	Unsup.	--	--	--	O(2n) adv.
SSAPGCN	2023	Single-sc.	N/A	Yes	--	$O(N^2)$ graph
SCLUDA	2023	Unsup.	--	--	Yes	$O(B \cdot n)$ cont.
CLCM	2023	Unsup.	--	Yes	--	\$O(B·n) cont.
LiCa	2023	No (DG)	--	--	--	O(n) cond.
SDEnet	2023	No (DG)	--	--	Yes	O(2n) GAN
LDGnet	2023	No (DG)	--	--	Yes	O(2n) dual
TSTnet	2023	Unsup.	--	Yes	Yes	$O(N^2)$ GOT
LLURnet	2023	No (DG)	--	--	Yes	O(2n) GAN
ExViT	2023	Single-sc.	N/A	Yes	Yes	$O(N^2)$ attn.
Gia-CFSL	2024	Few-shot	Yes	Yes	Yes	$O(N^2)$ graph
FDGNet	2025	No (DG)	--	--	--	O(nlogn) FFT

Notes: Target Data: "Unsup." = unsupervised (unlabeled target data), "Single-sc." = single-scene (not cross-domain). Complexity: n = pixels, N = graph nodes, K = kernels, B = batch/branch count. O(2n) = adversarial with alternating optimization. $O(N^2)$ = pairwise similarity dominated. "–" = not supported.

Several patterns emerge from Table 2. First, open-set capability is rare: only ZDDA (2021) and Gia-CFSL (2024) explicitly handle scenarios where source and target label sets differ—a striking gap given that real-world deployment routinely encounters unseen land cover classes. Second, source code availability is concentrated in the 2023–2025 period, reflecting a positive but recent shift toward reproducible research. Third, computational complexity varies by orders of magnitude: frequency-domain (FDGNet) and conditional-alignment (LiCa) methods scale near-linearly, while graph-based (GCN, graph optimal transport) and kernel-based methods scale quadratically, potentially limiting their applicability to large-scale HSI scenes.

5.4. Methodological Observations

Dataset standardization is lacking. The reviewed works evaluate on different combinations of HSI datasets (Pavia, Houston, Salinas, Indian Pines, HyRANK, etc.) with varying experimental protocols (training ratios, number of runs, evaluation metrics). This heterogeneity complicates direct method comparison and may contribute to the "every method outperforms baselines" pattern observed across the literature. The community would benefit from standardized cross-domain HSI benchmarks analogous to Office-31 and DomainNet in computer vision.

Reproducibility is improving. Encouragingly, recent works [8, 12, 17-20, 22] have released source code, marking a positive shift toward reproducible research. This trend should be accelerated through shared evaluation frameworks.

Computational cost is rarely reported. While classification accuracy is the dominant evaluation metric, the computational requirements (training time, GPU memory, inference latency) of different methods vary substantially. Adversarial methods [6, 21] require training domain discriminators; contrastive methods [8, 12] require large batch sizes for effective negative sampling; graph methods [2, 19] incur quadratic complexity in the number of nodes. The trade-off between accuracy and computational efficiency deserves explicit consideration.

The dominant journals are concentrated. The reviewed literature is heavily concentrated in IEEE journals—IEEE TGRS (11 papers), IEEE TNNLS (4 papers), IEEE TPAMI (1 paper), IEEE TIP (1 paper), IEEE GRSL (1 paper)—with only a few appearances in Pattern Recognition, ISPRS, Neural Networks, and IEEE JSTARS. This concentration suggests limited cross-pollination with the broader computer vision DA/DG community, whose methods (e.g., self-training, mean teacher, augmentation-based DG) have seen limited adaptation to the HSI domain.

6. Research Gaps and Future Directions

6.1. Identified Gaps

We present seven research gaps with an assessment of their urgency (near-term: 1–3 years; medium-term: 3–5 years) and potential impact (high/medium) on the field. The assessment reflects the authors' judgment based on the synthesis of the reviewed literature and identified trends.

Gap 1: Foundation model adaptation for cross-domain HSI. Urgency: Near-term | Impact: High

The remote sensing community has begun exploring foundation models (e.g., pretrained vision transformers, CLIP-based models, GeoFMs), but their systematic adaptation to cross-domain HSI classification remains unexplored. LDGnet [18] is the only reviewed work leveraging pretrained language models, and ExViT [15] adapts vision transformers but trains from scratch. This gap is particularly urgent given the rapid pace of foundation model development in both computer vision (DINOv3, SAM-2) and remote sensing (SpectralGPT, GeoChat). The primary barrier is the scarcity of large-scale, diverse, publicly available HSI training corpora comparable to ImageNet or LAION.

Gap 2: Theoretical generalization bounds specific to HSI data.

Urgency: Medium-term | Impact: High

Only Gao et al. (2023) provide theoretical analysis (error bounds for DG), and only in the context of conditional distribution alignment. Generalization theory for HSI domain adaptation—accounting for the unique properties of hyperspectral data: high dimensionality (~200 bands), strong spectral correlation, and spatial autocorrelation—is largely absent. Deriving bounds that account for the spectral correlation structure—which fundamentally differs from the i.i.d. assumptions of standard generalization theory—represents a significant theoretical challenge.

Gap 3: Benchmark standardization and rigorous comparison.

Urgency: Near-term | Impact: High

As analyzed in Section 5.4, the absence of standardized cross-domain HSI benchmarks impedes objective assessment of methodological progress. This gap is rated near-term and high-impact because the infrastructure costs are modest (existing public HSI datasets are sufficient) while the benefit—enabling rigorous comparison and preventing the "every method outperforms baselines" pattern—is immediate. The CV community's DomainBed and WILDS benchmarks provide proven templates.

Gap 4: Open-set and universal domain adaptation.

Urgency: Near-term | Impact: High

Only two reviewed works [7, 20] address settings where source and target label sets differ, as confirmed in Table 2. Real-world HSI deployment scenarios routinely encounter new land cover classes not present in historical labeled data—for example, urban expansion introducing "construction site" in a target scene where the source domain only contained "built-up" and "vegetation." Open-set DA (target has unknown classes), partial DA (source has classes absent in target), and universal DA (both have private classes) represent critical practical settings.

Gap 5: Computational efficiency and onboard deployment.

Urgency: Medium-term | Impact: Medium

No reviewed work explicitly addresses the computational constraints of satellite or UAV onboard processing. As shown in Table 2, methods span a wide complexity spectrum—from O(nlogn) (FDGNet) to $O(N^2)$ (graph-based)—but none are evaluated under realistic edge-computing constraints. Lightweight cross-domain methods suitable for onboard inference represent a bridging problem between the remote sensing and TinyML communities.

Gap 6: Multimodal domain adaptation and generalization.

Urgency: Medium-term | Impact: Medium

While cross-modality learning [4, 5, 15] addresses alignment between heterogeneous modalities, the integration of multimodal fusion with domain adaptation/generalization is nascent. Most DA/DG methods operate on single-modality HSI data (18 of 23 reviewed works), yet modern remote sensing increasingly involves multiple colocated sensors (HSI+MSI+LiDAR+SAR). Joint multimodal-domain adaptation—where complementary modalities serve as auxiliary sources of domain-invariant information—represents a natural extension of both research threads but requires multimodal benchmark datasets that are currently scarce.

Gap 7: Temporal domain shift.

Urgency: Near-term | Impact: Medium

All reviewed works address spatial cross-scene domain shift (different geographical locations), but temporal domain shift (same location, different time) introduces additional challenges—phenological changes, land cover transitions, seasonal spectral variations—that are underexplored. The key methodological challenge is distinguishing domain shift caused by acquisition conditions (to be removed via DA/DG) from genuine land cover change (to be preserved).

6.2. Proposed Research Agenda

Develop HSI-specific foundation models pretrained on large-scale, diverse hyperspectral archives, and systematically evaluate their zero-shot and few-shot cross-domain transfer capabilities.

Establish community-standard cross-domain HSI benchmarks with predefined source-target splits, multi-source configurations, and standardized evaluation metrics (per-class accuracy, macro-F1, domain generalization gap).

Derive theory-driven domain adaptation bounds that account for HSI-specific properties: spectral dimensionality, spatial correlation structure, and class-conditional distribution geometry.

Extend open-set and universal DA/DG to HSI, developing methods that gracefully handle new classes in target domains without catastrophic forgetting of source knowledge.

Design lightweight cross-domain architectures targeting onboard deployment, using techniques such as knowledge distillation, neural architecture search, and quantization adapted to domain-shifted settings.

Integrate multimodal fusion with domain generalization, leveraging complementary modalities as auxiliary sources of domain-invariant information.

Investigate temporal domain shift, developing methods that distinguish domain shift caused by acquisition conditions (to be removed) from genuine land cover change (to be preserved).

7. Conclusion

7.1. Key Takeaways

This review has examined 23 representative works in cross-domain hyperspectral image classification, spanning domain adaptation, domain generalization, cross-modality learning, and graph-based approaches, published between 2015 and 2025. Three principal conclusions emerge:

First, the field has undergone a clear methodological evolution: from statistical alignment (MMD, TCA: 2015–2019) through adversarial learning (GAN-based DA: 2020–2022) to contrastive learning and domain generalization (2023–2025). Each paradigm shift has brought not only improved accuracy but also conceptual advances in understanding the nature of domain shift in hyperspectral data.

Second, domain generalization—learning models that transfer to unseen target domains without target data exposure—has emerged as a practical alternative to domain adaptation. DG methods using domain expansion (SDEnet, LLURnet), language-guided alignment (LDGnet), and frequency disentanglement (FDGNet) have significantly narrowed the performance gap with DA methods, while offering the operational advantage of zero-shot deployment.

Third, the integration of complementary perspectives—graph-based topological reasoning, multimodal sensor fusion, linguistic semantic priors, and frequency-domain analysis—is yielding richer and more robust cross-domain representations than purely spectral-spatial approaches. The most successful methods exploit information beyond the conventional spatial-spectral feature space.

7.2. Implications for Practice and Policy

For practitioners developing operational HSI classification systems, this review suggests several actionable recommendations. (1) When target domain data is available (even unlabeled), contrastive learning-based DA methods (SCLUDA, CLCM) currently offer the best accuracy-efficiency trade-off. (2) For rapid deployment scenarios without target data access, domain expansion DG methods (SDEnet, LLURnet) provide the most practical solution from a single source domain. (3) Graph-based methods should be considered when nonlocal spatial relationships are important (e.g., large homogeneous regions) but their computational cost may prohibit large-scale application. (4) The release of source code and pretrained models should be standard practice to enable reproducible comparison and practical deployment.

For the broader remote sensing research community, standardized cross-domain benchmarks and rigorous theoretical analysis remain the most pressing needs to advance the field beyond incremental empirical improvements toward principled, generalizable solutions for the fundamental challenge of spectral domain shift.

References

[1] Deng, C., Liu, X., Li, C., & Tao, D. (2018). Active multi-kernel domain adaptation for hyperspectral image classification. Pattern Recognition, 77, 306–315. https://doi.org/10.1016/j.patcog.2017.12.006

[2] Ding, Y., Chong, Y., Pan, S., Wang, Y., & Nie, C. (2023). Spatial–spectral unified adaptive probability graph convolutional networks for hyperspectral image classification. IEEE Transactions on Neural Networks and Learning Systems, 34(7), 3653–3667. https://doi.org/10.1109/TNNLS.2022.3141286

[3] Gao, Z., Pan, B., Xu, T., Li, T., & Shi, Z. (2023). LiCa: Label-indicate-conditional-alignment domain generalization for pixel-wise hyperspectral imagery classification. IEEE Transactions on Geoscience and Remote Sensing, 61, 5513516. https://doi.org/10.1109/TGRS.2023.3284098

[4] Hong, D., Yokoya, N., Ge, N., Chanussot, J., & Zhu, X. X. (2019). Learnable manifold alignment (LeMA): A semi-supervised cross-modality learning framework for land cover and land use classification. ISPRS Journal of Photogrammetry and Remote Sensing, 147, 193–205. https://doi.org/10.1016/j.isprsjprs.2018.11.014

[5] Hong, D., Yao, J., Meng, D., Xu, Z., & Chanussot, J. (2021). Multimodal GANs: Toward crossmodal hyperspectral–multispectral image segmentation. IEEE Transactions on Geoscience and Remote Sensing, 59(6), 5103–5113. https://doi.org/10.1109/TGRS.2020.3032203

[6] Huang, Y., Peng, J., Sun, W., Chen, N., Du, Q., Ning, Y., & Su, H. (2022). Two-branch attention adversarial domain adaptation network for hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 60, 5536914. https://doi.org/10.1109/TGRS.2022.3188944

[7] Kutbi, M., Peng, K.-C., & Wu, Z. (2021). Zero-shot deep domain adaptation with common representation learning. IEEE Transactions on Pattern Analysis and Machine Intelligence, 44(9), 5430–5446. https://doi.org/10.1109/TPAMI.2021.3060090

[8] Li, Z., Xu, Q., Ma, L., Fang, Z., Wang, Y., He, W., & Du, Q. (2023). Supervised contrastive learning-based unsupervised domain adaptation for hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 61, 5528617. https://doi.org/10.1109/TGRS.2023.3296323

[9] Lin, J., Zhao, L., Li, S., Ward, R., & Wang, Z. J. (2018). Active-learning-incorporated deep transfer learning for hyperspectral image classification. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 11(11), 4249–4262. https://doi.org/10.1109/JSTARS.2018.2868438

[10] Liu, Z., Ma, L., & Du, Q. (2021). Class-wise distribution adaptation for unsupervised classification of hyperspectral remote sensing images. IEEE Transactions on Geoscience and Remote Sensing, 59(1), 508–521. https://doi.org/10.1109/TGRS.2020.2994443

[11] Matasci, G., Volpi, M., Kanevski, M., Bruzzone, L., & Tuia, D. (2015). Semisupervised transfer component analysis for domain adaptation in remote sensing image classification. IEEE Transactions on Geoscience and Remote Sensing, 53(7), 3550–3564. https://doi.org/10.1109/TGRS.2014.2380378

[12] Ning, Y., Peng, J., Liu, Q., Huang, Y., Sun, W., & Du, Q. (2023). Contrastive learning based on category matching for domain adaptation in hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 61, 5522915. https://doi.org/10.1109/TGRS.2023.3269127

[13] Peng, J., Sun, W., Ma, L., & Du, Q. (2019). Discriminative transfer joint matching for domain adaptation in hyperspectral image classification. IEEE Geoscience and Remote Sensing Letters, 16(6), 953–957. https://doi.org/10.1109/LGRS.2019.2893141

[14] Qin, B., Feng, S., Zhao, C., Xi, B., Li, W., & Tao, R. (2025). FDGNet: Frequency disentanglement and data geometry for domain generalization in cross-scene hyperspectral image classification. IEEE Transactions on Neural Networks and Learning Systems, 36(5), 7171–7185. https://doi.org/10.1109/TNNLS.2024.3484707

[15] Yao, J., Zhang, B., Li, C., Hong, D., & Chanussot, J. (2023). Extended vision transformer (ExViT) for land use and land cover classification: A multimodal deep learning framework. IEEE Transactions on Geoscience and Remote Sensing, 61, 5517915. https://doi.org/10.1109/TGRS.2023.3260295

[16] Zhang, Y., Li, W., Tao, R., Peng, J., Du, Q., & Cai, Z. (2021). Cross-scene hyperspectral image classification with discriminative cooperative alignment. IEEE Transactions on Geoscience and Remote Sensing, 59(9), 7646–7660. https://doi.org/10.1109/TGRS.2021.3060137

[17] Zhang, Y., Li, W., Sun, W., Tao, R., & Du, Q. (2023). Single-source domain expansion network for cross-scene hyperspectral image classification. IEEE Transactions on Image Processing, 32, 1499–1514. https://doi.org/10.1109/TIP.2023.3241811

[18] Zhang, Y., Zhang, M., Li, W., Wang, S., & Tao, R. (2023). Language-aware domain generalization network for cross-scene hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 61, 5502812. https://doi.org/10.1109/TGRS.2023.3244601

[19] Zhang, Y., Li, W., Zhang, M., Qu, Y., Tao, R., & Qi, H. (2023). Topological structure and semantic information transfer network for cross-scene hyperspectral image classification. IEEE Transactions on Neural Networks and Learning Systems,

34(6), 3112–3125. https://doi.org/10.1109/TNNLS.2022.3141964

[20] Zhang, Y., Li, W., Zhang, M., Wang, S., Tao, R., & Du, Q. (2024). Graph information aggregation cross-domain few-shot learning for hyperspectral image classification. IEEE Transactions on Neural Networks and Learning Systems, 35(4), 5712–5726. https://doi.org/10.1109/TNNLS.2023.3314276

[21] Zhao, C., Qin, B., Feng, S., Zhu, W., Zhang, L., & Ren, J. (2022). An unsupervised domain adaptation method towards multi-level features and decision boundaries for cross-scene hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 60, 5545916. https://doi.org/10.1109/TGRS.2022.3189122

[22] Zhao, H., Zhang, J., Lin, L., Wang, J., Gao, S., & Zhang, Z. (2023). Locally linear unbiased randomization network for cross-scene hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 61, 5527215. https://doi.org/10.1109/TGRS.2023.3294820

[23] Zhu, Y., Zhuang, F., Wang, J., Chen, Z., Shi, Z., & Wu, W. (2019). Multi-representation adaptation network for cross-domain image classification. Neural Networks, 119, 214–221. https://doi.org/10.1016/j.neunet.2019.08.011