

MDFL: Multi-Domain Diffusion-Driven Feature Learning

Daixun Li, Weiyong Xie*, Jiaqing Zhang, Yunsong Li

State Key Laboratory of Integrated Services Networks, Xidian University, Xi'an 710071, China
ldx@stu.xidian.edu.cn, wyxie@xidian.edu.cn, jqzhang_2@stu.xidian.edu.cn, ysl@mail.xidian.edu.cn

Abstract

High-dimensional images, known for their rich semantic information, are widely applied in remote sensing and other fields. The spatial information in these images reflects the object's texture features, while the spectral information reveals the potential spectral representations across different bands. Currently, the understanding of high-dimensional images remains limited to a single-domain perspective with performance degradation. Motivated by the masking texture effect observed in the human visual system, we present a multi-domain diffusion-driven feature learning network (MDFL), a scheme to redefine the effective information domain that the model really focuses on. This method employs diffusion-based posterior sampling to explicitly consider joint information interactions between the high-dimensional manifold structures in the spectral, spatial, and frequency domains, thereby eliminating the influence of masking texture effects in visual models. Additionally, we introduce a feature reuse mechanism to gather deep and raw features of high-dimensional data. We demonstrate that MDFL significantly improves the feature extraction performance of high-dimensional data, thereby providing a powerful aid for revealing the intrinsic patterns and structures of such data. The experimental results on three multi-modal remote sensing datasets show that MDFL reaches an average overall accuracy of 98.25%, outperforming various state-of-the-art baseline schemes. Code available at <https://github.com/LDXDU/MDFL-AAAI-24>.

Introduction

Recently, remarkable strides have been achieved in precise classification within the realm of natural images, partially attributing to the distinctive representations found within the ImageNet dataset (Li et al. 2022a). These representations frequently encompass local structures, repetitive patterns, and hierarchical arrangements (Yao et al. 2019), rendering them advantages for deep learning models. However, progress in the domain of high-dimensional data analysis has been relatively sluggish in comparison to that in natural images due to several factors, including the absence of prior knowledge, sparse distributions, and low spatial resolutions (Jiang et al. 2021). Consequently, it becomes imperative to

*Correspond author.




Figure 1: The forward diffusion process in the spectral-spatial domain denoted by I_0 , the noised image at time step 0 and time step T denoted by I_T , and the feature visualization results of the hyperspectral image at different time steps.

delve into exploring intelligent interpretation algorithms that can be effectively applied to high-dimensional data, thereby improving their classification accuracy and effectiveness.

Multi-modal learning (MML), as a vital endeavor in high-dimensional data processing, plays a pivotal role in remote sensing applications (Zhao et al. 2023) by harnessing essential information from multiple sensor source images to enhance data acquisition (Feng et al. 2023; Hong et al. 2020). In scenarios where the availability of prior samples is limited, researchers have begun exploring fusion analysis approaches for high-dimensional data. Many of these approaches leverage Convolutional Neural Networks (CNNs) architectures for inter-layer fusion. For instance, Hong *et al.* (Hong et al. 2022) proposed a deep encoder-decoder network architecture for hyperspectral and LiDAR data classification, effectively capturing spectral and spatial information in a deep latent space. Subsequently, larger-scale models like transformers (Roy et al. 2023) and diffusion model (Zhao et al. 2023) have emerged. Typically, the optimization of these methods involves two distinct steps: extracting features from high-dimensional data using large-scale model architectures, and subsequently classifying the extracted features using a classifier. While these methods mitigate the curse of dimensionality, their feature extraction process does not adequately cater to the manifold structures of high-

dimensional data, and result in suboptimal solutions. Moreover, compared to natural images with feature domains in spatial, spectral, and frequency domains, high-dimensional data often exhibit more compact data relationships (Hu et al. 2022). Therefore, the majority of current techniques that exhibit exceptional performance in natural image scenarios are not applicable in high-dimensional contexts.

The aforementioned methodologies approach the fusion of multi-modal data as a mere feature extraction problem. Such methods fail to fully exploit the particularity of multi-modal fusion, which involves preserving the spectral, spatial, and frequency domains inherent in high-dimensional data. Thus, they treat the fusion process as a black-box deep learning problem. Motivated by the masking texture effect observed in the human visual system (Liu et al. 2021), the bottom-up attention driven by objective content can locate and analyze important target domains in detail, while other regions are only roughly analyzed or even ignored, this characteristic effectively enhances the focusing ability of image content and improves the performance accuracy of target recognition. However, excessive attention to a single target domain will lead to the loss of analytical capabilities for other effective information domains in the image (Liang et al. 2022), which is consistent with the texture masking effect (Song et al. 2020; Li et al. 2022b). This observation will be validated in the subsequent multi-domain learning methods. As shown in Figure 1, a spectral-spatial sample of multi-modal data is obtained in the preceding process of diffusion. However, with the increase of time steps, the texture occlusion effect caused by the digital domain information makes it difficult to extract effective features. For the multi-modal fusion problem, it is obvious that preserving the spatial, spectral, and frequency domains is the main objective. Therefore, deep learning methods should explicitly focus on this aspect, which inspires our proposed deep network called “MDFL”. To the best of our knowledge, this is the first discussion on diffusion characteristics in the frequency domain of high-dimensional data, and the first exploration of the intrinsic properties of high-dimensional data by combining spatial and spectral information:

- We propose a diffusion-driven posterior sampling model MDFL, which formulates the MML problem as a unified framework for multi-domain joint feature learning in a manifold structure. It significantly improves feature extraction performance for high-dimensional data.
- A novel feature reuse mechanism is developed to integrate learned deep and shallow features. Consisting of two parallel attention modules, it can effectively aggregate cross-level features with minimal additional cost.
- We demonstrate the superiority of the proposed approach over the existing SOTA in multiple multimodal datasets. In addition, the effectiveness of feature reuse mechanism and multi-domain learning for high-dimensional data interpretation is verified by ablation experiments.

Related Works

The existing mainstream methods for multi-modal fusion mainly use CNNs and Transformers (Li et al. 2024),

but there are significant problems in the field of high-dimensional image processing. High-dimensional image semantic segmentation usually needs to consider the spatial information of the image, and the sparse distribution and adjacent space of ground objects have an important impact on the classification results. Traditional CNNs do not consider the spatial relationship between pixels, so their effectiveness in remote sensing image classification may be limited. The Transformer is mainly used to process sequence data, but for high-dimensional image semantic segmentation, spatial dependencies are more important, so Transformer is not the first choice.

The diffusion model is based on non-equilibrium thermodynamics and involves a Markov diffusion step chain that adds random noise to the data and learns the reverse diffusion process to generate desired data samples from the noise (Croitoru et al. 2023). Unlike VAE or flow models, diffusion models have a fixed process and high-dimensional latent variables (Nichol and Dhariwal 2021). Diffusion models are good at capturing fine details and realistic textures of high-dimensional images and have been proven to be very scalable to high-dimensional data. Moreover, diffusion models can effectively extract manifold features of different levels of noise steps. Diffusion models have been applied to various fields such as image super-resolution (Esser, Rombach, and Ommer 2021), semantic segmentation (Baranchuk et al. 2021), and classification (Han, Zheng, and Zhou 2022). Research on these can be categorized into three areas: effective sampling, improved likelihood estimation, and handling data with special structures (Yang et al. 2022). Effective sampling involves generating samples using iterative methods with a large number of evaluation steps. Karras *et al.* (Karras et al. 2022) determined optimal time discretization and applied high-order Runge-Kutta methods for sampling. They also evaluated different sampler schedules and analyzed the role of randomness in the sampling process. Remote sensing data, which include observations from satellites or sensors, often exhibit spatial correlations and geographic distributions. Zhou *et al.* (Zhou et al. 2023) designed a temporal leap feature library and dynamic feature fusion module to utilize rich temporal leap features and learn information-rich multi-temporal representations. Han *et al.* (Han, Zheng, and Zhou 2022) introduced the classification and regression diffusion model, which combines a denoising diffusion-driven generative model with a pretrained mean estimator for more accurate instance-level confidence evaluations in classification tasks.

High-dimensional hyperspectral images have been widely used as a representative type of data. In order to fully utilize the digital domain characteristics of hyperspectral data, diffusion-based classification models have been proposed. Although these methods have achieved satisfactory results, challenges still persist in the fusion of multi-modal data, such as the inability to capture the close relationships among multiple target domains. Additionally, current diffusion models used for remote sensing classification do not consider the manifold structure of high-dimensional data.

Methodology

Problem Formulation

The problem of high dimensional land cover classification can be defined as accurately assigning each pixel in a remote sensing image to different land cover classes. Given a high dimensional image I composed of p -th pixels, we focus on how to train different modalities of data $I^1, I^2 \in \mathbb{R}^{h \times w \times c}$ to achieve classification tasks. The h , w , and c denote the height, width, and the number of image channels, and $I^1(p)$ is represented as the p -th pixel pair of the first modality. The two modalities capture the same scene with the same label information, denoted as $L \in \mathbb{R}^{h \times w \times c}$ with $C \in N^*$ classification categories, to represent various land cover categories, such as buildings, roads, and fields. The objective of image classification is to develop a model $\Theta_{(I^1, I^2)}$, that can effectively map input images from various modalities to a novel representation $C_{\max}(I^1, I^2)$, thus, indicating the probability of each pixel being associated with different categories. By setting the maximum probability of distinct categories as the threshold (τ), a binary prediction map is obtained through hard classification, the values 1 and 0 in the map respectively represent the presence of the specific category and other categories, which is defined as:

$$\Theta_{(I^1, I^2)} = \begin{cases} 0, & \text{if } C_{\max}(I^1, I^2) < \tau, \\ 1, & \text{otherwise.} \end{cases} \quad (1)$$

Based on this basic model, we propose the MDFL network architecture. As shown in Figure 2, the network consists of two branches that are multi-domain learning driven by a diffusion model. Thus, the fusion future graph in our work $C_{\max}(I^1, I^2)$ can be expressed as:

$$C_{\max}(I^1, I^2) = \Phi(I^1, I^2 | \alpha_1, \alpha_2), \quad (2)$$

where a nonlinear objective model $\Phi(\cdot)$ is used to transform the image space into the classification space, and α_1, α_2 represent the corresponding parameters of the two branches.

Network Architecture

In this section, the network architecture of the diffusion-driven multi-domain feature fusion learning method is introduced in Figure 2 to explicitly model the compact representation of each modal target domain for fusion classification of multi-modal data.

Diffusion-driven spectral-spatial feature learning In the realm of high-dimensional data analysis, a thorough investigation of the spectral and spatial domains is crucial for elucidating the diffusion characteristics. This section proposes a novel approach to learn diffusion-driven spectral-spatial features in a multi-domain posterior sampling framework. The proposed method utilizes a conditional generation module and a maximum likelihood estimation module, the two jointly analyzing the spectral, spatial and frequency domains. First, we add noise to the digital domain instance in the forward propagation stage and take the fusion of multiple noise scales as the input of the model. As shown in Figure 1, the diffusion process of the preceding term is inspired by the principle of non-equilibrium thermodynamics. As a Markov

chain, the one-step derivation formula can be obtained from the iterative formula:

$$\begin{cases} I_t = \sqrt{\bar{\alpha}_t} I_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \\ \bar{\alpha}_t = \prod_{i=1}^t \alpha_i, \end{cases} \quad (3)$$

where $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$ is the hyperparameter set with the Noise schedule, and $\epsilon \sim N(0, 1)$ is Gaussian noise. The noised maps are firstly concatenation in channel dimension as:

$$\tilde{I} = \text{Concat}[I_0, I_{50}, I_{100}, I_{200}, I_{400}]. \quad (4)$$

During training, a U-Net basic framework $f_\theta(z_t, t)$ is trained to predict \tilde{I} from z_t by minimizing the training objective with ℓ_2 loss:

$$\mathcal{L}_{\text{train}} = \frac{1}{2} \left\| f_\theta(z_t, t) - \tilde{I} \right\|_2,$$

based on Bayes' theorem, it is found that the posterior $q(z_{t-1} | z_t, z_0)$ is a Gaussian distribution as well:

$$q(z_{t-1} | z_t, z_0) = \mathcal{N}(z_{t-1}; \tilde{\mu}(z_t, z_0), \tilde{\beta}_t \tilde{I}), \quad (5)$$

where

$$\tilde{\mu}_t(z_t, z_0) = \frac{\sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)}{1 - \bar{\alpha}_t} z_0 + \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} z_t, \quad (6)$$

and

$$\tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} (1 - \alpha_t), \quad (7)$$

are mean and variance of this Gaussian distribution. We could get a sample from $q(z_0)$ by first sampling from $q(z_T)$ and running the reversing steps $q(z_{t-1} | z_t)$ until z_0 . Besides, the distribution of $q(z_T)$ is nearly an isotropic Gaussian distribution with a sufficiently large T and reasonable schedule of β_t ($\beta_t \rightarrow 0$), which makes it trivial to sample $z_T \sim \mathcal{N}(0, \tilde{I})$. Moreover, we could approximate $q(z_{t-1} | z_t)$ using a neural network, due to the fact that calculating $q(z_{t-1} | z_t)$ exactly should depend on the entire data distribution. The network is optimized to predict a mean μ_θ and a diagonal covariance matrix Σ_θ :

$$p_\theta(z_{t-1} | z_t) := \mathcal{N}(z_{t-1}; \mu_\theta(z_t, t), \Sigma_\theta(z_t, t)). \quad (8)$$

At inference stage, data sample z_0 is reconstructed from noise z_T with the model f_θ with an updating rule in an iterative way, i.e., $z_T \rightarrow z_{T-\Delta} \rightarrow \dots \rightarrow z_0$. At the same time, in order to capture the spatial and spectral residual information in the digital domain of the feature map, we add the spatial-spectral attention mechanism to the UNet basic framework. We then combine the attention mechanism and residual connection to enhance the representation ability of the feature map. Furthermore, we construct the transformation matrix (W_Q, W_K, W_V) before the multi-head self-attention (MSA) operation (Fan et al. 2021).

$$Q = \tilde{F}W_Q, \quad K = \tilde{F}W_K, \quad V = \tilde{F}W_V. \quad (9)$$

The input intermediate feature \tilde{F} is partitioned into distinct Q, K , and V tensors. To calculate the attention weights, we initially perform a dot product operation between the

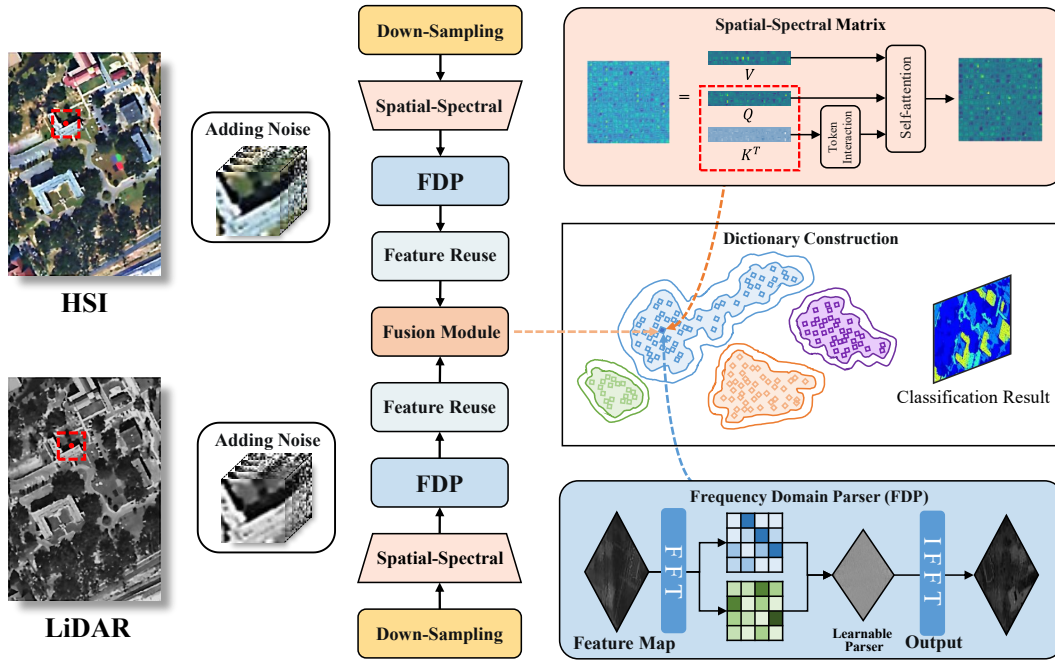


Figure 2: Overview of the multi-domain diffusion-driven feature learning framework, which combines spatial-spectral matrix and frequency domain parser for joint learning of multi-modal data. The MDFL framework is utilized to integrate the two modal data at the feature level, after which the extracted features are categorized into the downstream model.

query and key tensors. Subsequently, we scale this dot product by a factor of $1/\sqrt{c}$ to address the issue of vanishing gradients. Finally, the obtained dot product is passed through a softmax activation function, yielding normalized weights that reflect the relative significance of various elements within the input tensor.

$$A = \text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{c}}\right)V. \quad (10)$$

A linear layer is used to produce the output. Multi-head self-attention splits the queries, keys and values to h parts and performs the attention function in parallel, and then the output values of each head are concatenated and linearly projected to form the final output.

The resulting tensor is then reshaped to its original dimensions, providing a refined representation of the input tensor capturing relevant spatial and spectral features. This module enables our model to effectively capture and exploit the intrinsic interdependencies between elements in the input tensor, thereby facilitating enhanced feature extraction and improving the interpretation performance of high-dimensional data in the digital domain.

Frequency-aware discriminative feature learning The visual masking effect will be caused with the only digital domain interpretation of high-dimensional data. As suggested, the phase of a blurry image plays an important role in blurring, providing faithful information about motion patterns (Pan et al. 2019). The blurry naturally exists in remote sensing images whose resolution is constrained by imaging distance. We introduce the frequency domain to enhance the

detail information (such as texture and color information) in the different multi-modalities to make the object more discriminative. The pipeline and effect of frequency-aware discriminative feature learning is shown in Figure 4. Our main idea is to learn a parameterized filter by applying it to the Fourier-space features. Let $X(p) \in \mathbb{R}^{H \times W \times C}$ be the input feature matrix, where H , W , and C indicate the height, width, and channel of the of the feature respectively feature. Firstly, the 2D FFT is performed along the spatial dimensions, which can be represented as:

$$M = \mathcal{F}(X(p)) \in \mathbb{R}^{H \times W \times C}, \quad (11)$$

where $\mathcal{F}(\cdot)$ denotes the 2D FFT. We then modulate a convolution operation to obtain an attention map that can expose the importance of the different frequency components. In other words, the weight of convolution can be regarded as a learnable version of frequency filters widely utilized in digital image processing.

$$M' = W \otimes M, \quad (12)$$

where W denotes the trainable weights in the frequency domain. It serves the purpose of regulating the noise at high frequencies while simultaneously learning the output of the spectral-spatial convolution from the preceding layer of the network, considering its spectral characteristics. And \otimes represents the element-wise product. We reverse back to the spatial domain by adopting inverse FFT:

$$X' = \mathcal{F}^{-1}(M'). \quad (13)$$

Different from the spectral-spatial domain, FDFL makes globally adjustments to the components of the specific frequencies to enhance the discriminative information. It can




Figure 3: The graph comparison between single step training and multi-step fusion training under different synchronization lengths is provided.




Figure 4: The frequency domain parser visualizations.

be learned to constrain the different frequency components for adaptive integration.

Feature reuse module In this framework, all network layers share an objective, which is to extract diverse types of features to achieve accurate recognition. The downstream model is expected to leverage the deep features and shallow features from the upstream model U-Net to optimize efficiency. To actualize this concept, we introduce a feature reuse mechanism that combines the inherent and semantic information of an image. Specifically, we propose two parallel attention modules, wherein each module modulates the current layer to attain high-level features. This procedure can be described as follows:

$$\begin{aligned} X_{\text{FRM}} &= \text{FRM}(X_{\text{low}}, X_{\text{deep}}) \\ &= X_{\text{low}} + X_{\text{deep}}, \end{aligned} \quad (14)$$

where $\text{FRM}(\cdot)$ represents the mapping function learned by the feature reuse module. X_{low} and X_{deep} represent the shallow and deep features in the U-Net encoder, respectively. X_{low} and X_{deep} are the attention features in shallow layer and deep layer, which can be obtained as follows:

$$X_{\text{deep}} = \text{Sigmoid}(F_d(X_{\text{low}})) F_b(X_{\text{deep}}) + F_b(X_{\text{low}}), \quad (15)$$

$$X_{\text{low}} = \text{Sigmoid}(F_l(X_{\text{low}})) F_b(X_{\text{deep}}) + F_b(X_{\text{low}}), \quad (16)$$

$F_b(\cdot)$ is suitable for the 1×1 convolutional layer structure constraining high-frequency noise, $F_d(\cdot)$ represents the 1×1 deformable convolutional layer used to extract the deep layer, $F_l(\cdot)$ represents the 1×1 variability convolutional

layer used to extract the shallow layer. Therefore, the embedded X_{FRM} is injected into the downstream model to provide prior knowledge to recognize the input image, which can be expressed by the following equation:

$$X(p) = \text{MLP}(\text{FRM}(\text{FRM}(X_0^1, X_1^1), \text{FRM}(X_0^2, X_1^2))). \quad (17)$$

The $l \in \{1, \dots, L\}$, and the original and deep information of the single modality are first fused by feature reuse, and then the FRM results of the two modalities are fused in the same way for the classification of downstream models.

Experiments and Analysis

Experimental Settings

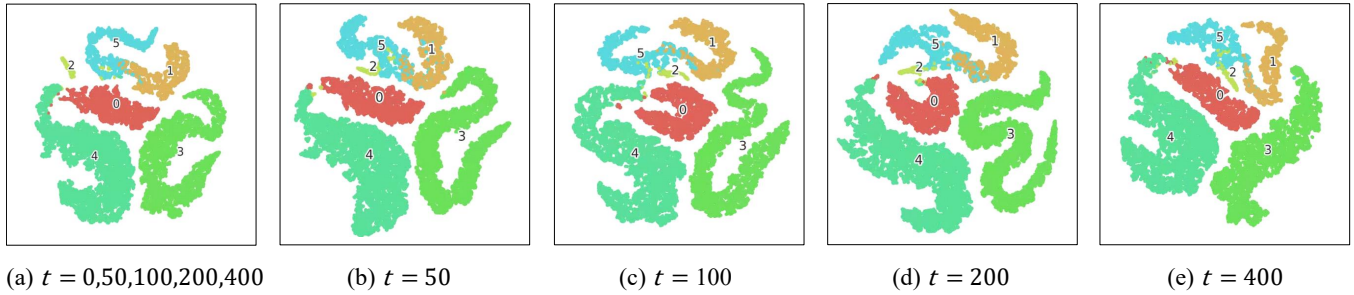
Datasets To validate the effectiveness of the proposed method in analyzing high-dimensional data, three remote sensing multimodal datasets with hyperspectra, namely the Houston2013 dataset, the Trento dataset, and the MUUFL dataset, are selected for verification of the proposed classification model.

To assess the performance of the proposed method in classifying test images, three metrics are used: overall accuracy (OA), average accuracy (AA), and kappa (κ) coefficient. OA measures the ratio of correctly classified test samples to the total number of test samples. AA represents the average accuracy across all classes. The κ coefficient measures the agreement between the classification maps generated by the model and the true values provided.

Implement details These experiments are conducted on a machine equipped with an NVIDIA A100 Tensor Core GPU. The training samples are randomly cropped to a size of 7×7 . The Adam optimizer is employed with an initial learning rate set to $1e-3$, and a weight decay of $5e-3$ was applied. The training process spans 1000 epochs. In addition, a step scheduler with a step size of 50 and gamma value of 0.9 is utilized. The batch size is set to 64. Additionally, for incorporating noise, a total step size of 500 was used, and the values of $t = 0, 50, 100, 200, 400$ are selected.

Ablation Study

In this section, we conduct an ablation study to assess the individual contributions of different components in the pro-

Figure 5: t -SNE effect display for different time steps.

	Houston2013			Trento			MUUFL		
Method	OA(%)	AA(%)	$\kappa(\times 100)$	OA(%)	AA(%)	$\kappa(\times 100)$	OA(%)	AA(%)	$\kappa(\times 100)$
RNN	62.61	62.24	59.64	96.43	92.38	95.21	88.79	75.84	85.18
Cross	91.84	92.70	91.16	97.82	97.35	97.09	87.29	63.81	82.75
CNN-2D	92.30	92.69	91.65	98.65	94.58	98.19	91.88	78.44	89.22
CALC	88.97	90.78	88.06	94.62	91.33	92.81	93.94	74.09	92.00
ViT	85.05	86.83	83.84	96.47	94.56	95.28	92.15	78.50	89.56
MFT	89.80	91.54	88.93	98.32	95.98	97.75	94.34	81.48	92.51
MDFL	99.16	96.72	94.09	99.43	96.13	94.41	96.96	82.63	94.63

Table 1: OA, AA and κ on the Houston2013 dataset, Trento dataset and MUUFL dataset by considering HSI and LiDAR data.

posed MDFL to MML. Specifically, five scenarios are designed: (A) a comparison of noise with various synchronization lengths selected by us against single-step training and multi-step fusion training; (B) the omission of frequency domain analysis during the training of diffusion models; (C) the exclusion of the original information in upstream model training from participating in feature reuse; (D) the consideration of the complete proposal model, MDFL.

Description	OA(%)	AA(%)	$\kappa(\times 100)$
(B)	97.75	95.62	93.24
(C)	98.10	90.44	89.39
(D)	98.51	91.82	94.37

Table 2: Average results of ablation study for MDFL conduct on 3 datasets. The best result is highlighted.

As shown in Figure 3, we observe with the line plots of three datasets that single-step training performs worse than multi-step fusion training in the process of forward diffusion, thereby validating the effectiveness of multi-step noise fusion. To further validate the benefits of multi-step time step fusion to classification, we demonstrate this with t-SNE maps at different time steps, as shown in Figure 5. It can be observed that (a) exhibits excellent clustering performance, with class 2 clearly separated from other classes in the two-

dimensional projection space. However, in (b), (c), (d), and (e), class 2 shows incorrect associations with other classes in the space, with confirms the effectiveness of multi-step fusion. Additionally, Table 2 demonstrates that the average OA precision of MDFL across the three datasets is 98.51%. In (B), by removing the joint analysis in the frequency domain, we observe a decrease in the results to 97.75%, resulting in a performance loss of 0.76%, which indicates the importance of multi-domain joint learning. The effectiveness of feature reuse in high-dimensional manifold structures is validated in (C). Therefore, in high-dimensional feature extraction, multi-objective domain joint learning and feature reuse modules can effectively enhance its performance.

Comparisons with Previous Methods

The accuracy performance of the proposed model as well as other models on the Houston2013, Trento, and MUUFL datasets is presented in Table 1. The best results are denoted in bold. In selecting comparison methods, we consider classic deep learning techniques such as RNN (Cho et al. 2014), mainstream RS multi-modal methods including Cross (Hong et al. 2020), CNN-2D and CALC (Ding et al. 2022), and transformer methods such as ViT (Dosovitskiy et al. 2021) and MFT (Roy et al. 2023).

Our evaluation demonstrates that the proposed method attains the highest OA, AA, and κ scores on most classification tasks, thus surpassing other methods. Specifically, on the Houston2013 dataset, MDFL exhibits superior accu-

racy across all categories, outperforming mainstream methods in terms of these coefficients. Although traditional methods perform better in terms of accuracy when combining HSI and LiDAR data, our model still presents significant advancements compared to other models. In comparison to transformer architecture, MDFL demonstrates a 9.36% increase in OA, a 5.18% increase in AA, and a 5.16% increase in κ compared MFT. When compared to Cross, MFT increases by 7.32% in OA, 4.02% in AA, and 2.93% in κ . Overall, MDFL achieves comprehensive performance improvements across three multimodal datasets and offers an effective solution for feature extraction from high-dimensional data.

Result Visualization

The results are depicted in Figure 6, Figure 7, and Figure 8, illustrating the macroscopic performance evaluation of the classification graphs produced by MDFL. To achieve this, we employ a visualization technique that assigns a unique color to each class. MDFL reconstructs high-dimensional features by integrating the joint information from various domains, including spectral, spatial, and frequency domains. This approach effectively reduces the granularity of textures in the classification graphs, resulting in a more diverse and finely-detailed representation. Overall, MDFL is well-suited for generating classification maps with enhanced performance and intricate details, thus, it is particularly suitable for land use and scene classification applications.

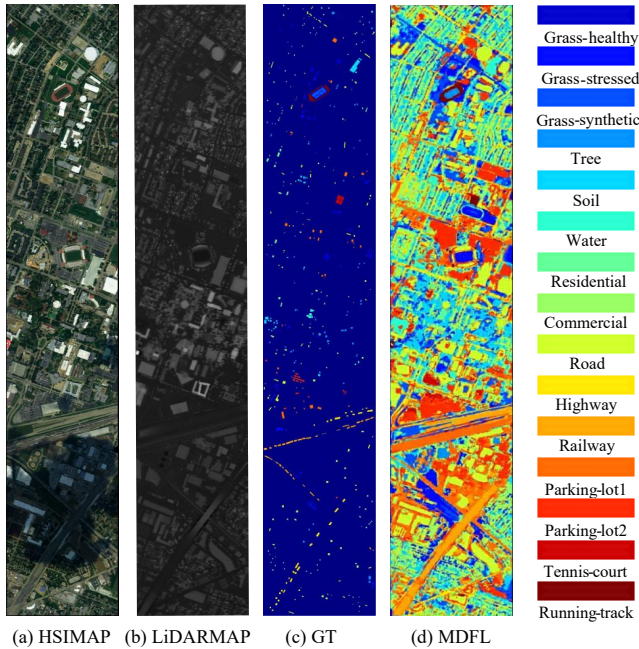


Figure 6: Visualization of false-color HSI and LiDAR images using MDFL based on the Houston2013 dataset.

Conclusion

In the context of high-dimensional image feature extraction, a significant challenge lies in developing a unified MML

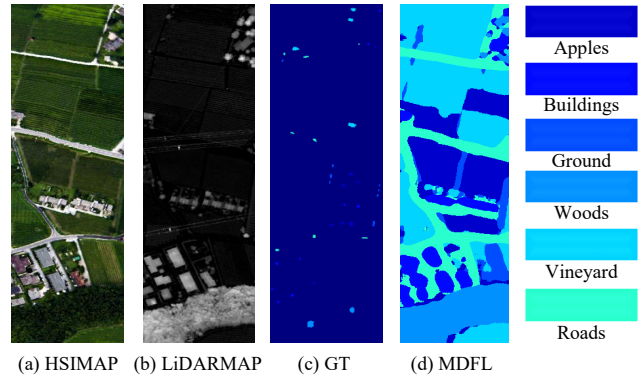


Figure 7: Visualization of false-color HSI and LiDAR images using MDFL based on the Trento dataset.




Figure 8: Visualization of false-color HSI and LiDAR images using MDFL based on the MUUFL dataset.

model capable of facilitating joint information exchange and feature extraction across multiple domains. This problem has not been fully addressed within the previous framework that primarily focuses on a single spatial domain. Taking inspiration from the masking texture effect observed in the human visual system, we propose a novel approach for joint feature learning in high-dimensional multimodal images. This approach encompasses spectral domain, spatial domain, and frequency domain to comprehensively capture the intricate features. Additionally, we consider the manifold structure characteristics of high-dimensional data and introduce a feature reuse mechanism to aggregate deep and primitive features from multiple modes. With multi-domain joint learning, our method not only reveals the underlying patterns and structures of high-dimensional data, but also maximizes the performance of feature extraction. Extensive experiments have been conducted to validate the efficacy of our proposed method, MDFL, in tackling MML problems. The results consistently demonstrate that MDFL outperforms existing methods, showcasing superior performance and highlighting its potential in addressing the challenges of high-dimensional data analysis.

References

Baranchuk, D.; Rubachev, I.; Voynov, A.; Khulkov, V.; and Babenko, A. 2021. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*.
 Cho, K.; Van Merriënboer, B.; Bahdanau, D.; and Ben-

- gio, Y. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Croitoru, F.-A.; Hondru, V.; Ionescu, R. T.; and Shah, M. 2023. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*.
- Ding, K.; Lu, T.; Fu, W.; Li, S.; and Ma, F. 2022. Global-Local Transformer Network for HSI and LiDAR Data Joint Classification. *IEEE Transactions on Geoscience and Remote Sensing (TGRS)*, 60: 1–13.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations (ICLR)*.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12873–12883.
- Fan, H.; Xiong, B.; Mangalam, K.; Li, Y.; Yan, Z.; Malik, J.; and Feichtenhofer, C. 2021. Multiscale vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6824–6835.
- Feng, Z.; Song, L.; Yang, S.; Zhang, X.; and Jiao, L. 2023. Cross-Modal Contrastive Learning for Remote Sensing Image Classification. *IEEE Transactions on Geoscience and Remote Sensing (TGRS)*.
- Han, X.; Zheng, H.; and Zhou, M. 2022. Card: Classification and regression diffusion models. *Advances in Neural Information Processing Systems (NIPS)*, 35: 18100–18115.
- Hong, D.; Gao, L.; Hang, R.; Zhang, B.; and Chanussot, J. 2022. Deep Encoder-Decoder Networks for Classification of Hyperspectral and LiDAR Data. *IEEE Geoscience and Remote Sensing Letters (GRSL)*, 19: 1–5.
- Hong, D.; Gao, L.; Yokoya, N.; Yao, J.; Chanussot, J.; Du, Q.; and Zhang, B. 2020. More diverse means better: Multimodal deep learning meets remote-sensing imagery classification. *IEEE Transactions on Geoscience and Remote Sensing (TGRS)*, 59(5): 4340–4354.
- Hu, X.; Cai, Y.; Lin, J.; Wang, H.; Yuan, X.; Zhang, Y.; Timofte, R.; and Van Gool, L. 2022. Hdnet: High-resolution dual-domain learning for spectral compressive imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17542–17551.
- Jiang, K.; Xie, W.; Lei, J.; Jiang, T.; and Li, Y. 2021. LREN: Low-rank embedded network for sample-free hyperspectral anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 35, 4139–4146.
- Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems (NIPS)*, 35: 26565–26577.
- Li, D.; Ling, H.; Kim, S. W.; Kreis, K.; Fidler, S.; and Torralba, A. 2022a. BigDatasetGAN: Synthesizing ImageNet with pixel-wise annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 21330–21340.
- Li, D.; Xie, W.; Li, Y.; and Fang, L. 2024. FedFusion: Manifold-Driven Federated Learning for Multi-Satellite and Multi-Modality Fusion. *IEEE Transactions on Geoscience and Remote Sensing (TGRS)*, 62: 1–13.
- Li, W.; Lin, Z.; Zhou, K.; Qi, L.; Wang, Y.; and Jia, J. 2022b. Mat: Mask-aware transformer for large hole image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10758–10768.
- Liang, J.; Hu, D.; Feng, J.; and He, R. 2022. Dine: Domain adaptation from single and multiple black-box predictors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8003–8013.
- Liu, S.; Wang, S.; Liu, X.; Gandomi, A. H.; Daneshmand, M.; Muhammad, K.; and De Albuquerque, V. H. C. 2021. Human memory update strategy: a multi-layer template update mechanism for remote visual monitoring. *IEEE Transactions on Multimedia (TM)*, 23: 2188–2198.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning (ICML)*, 8162–8171. PMLR.
- Pan, L.; Hartley, R.; Liu, M.; and Dai, Y. 2019. Phase-only image based kernel estimation for single image blind deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6034–6043.
- Roy, S. K.; Deria, A.; Hong, D.; Rasti, B.; Plaza, A.; and Chanussot, J. 2023. Multimodal fusion transformer for remote sensing image classification. *IEEE Transactions on Geoscience and Remote Sensing (TGRS)*.
- Song, K.; Wei, X.-S.; Shu, X.; Song, R.-J.; and Lu, J. 2020. Bi-modal progressive mask attention for fine-grained recognition. *IEEE Transactions on Image Processing (TIP)*, 29: 7006–7018.
- Yang, L.; Zhang, Z.; Song, Y.; Hong, S.; Xu, R.; Zhao, Y.; Shao, Y.; Zhang, W.; Cui, B.; and Yang, M.-H. 2022. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796*.
- Yao, T.; Pan, Y.; Li, Y.; and Mei, T. 2019. Hierarchy parsing for image captioning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2621–2629.
- Zhao, Z.; Bai, H.; Zhu, Y.; Zhang, J.; Xu, S.; Zhang, Y.; Zhang, K.; Meng, D.; Timofte, R.; and Van Gool, L. 2023. DDFM: denoising diffusion model for multi-modality image fusion. *arXiv preprint arXiv:2303.06840*.
- Zhou, J.; Sheng, J.; Fan, J.; Ye, P.; He, T.; Wang, B.; and Chen, T. 2023. When Hyperspectral Image Classification Meets Diffusion Models: An Unsupervised Feature Learning Framework. *arXiv preprint arXiv:2306.08964*.