

Self-Supervised One-Step Diffusion Refinement for Snapshot Compressive Imaging

Shaoguang Huang¹, Yunzhen Wang¹, Haijin Zeng^{2*}, Hongyu Chen¹, Hongyan Zhang^{1*}

¹School of Computer Science, China University of Geosciences, Wuhan, China

²Department of Psychiatry, Harvard University, Cambridge, USA

{huangshaoguang, chenhongyu, zhanghongyan}@cug.edu.cn, yunzhen1110@163.com, haijin.zeng2018@gmail.com

Abstract

Snapshot compressive imaging (SCI) captures multispectral images (MSIs) using a single coded two-dimensional (2-D) measurement, but reconstructing high-fidelity MSIs from these compressed inputs remains a fundamentally ill-posed challenge. Recent diffusion-based methods improve quality but are limited by scarce MSI training data, domain shifts from RGB-pretrained models, and slow multi-step sampling. These drawbacks restrict their practicality in real-world applications. Unlike prior approaches that rely on expensive iterative refinement or subspace-based diffusion embeddings (e.g., DiffSCI, PSR-SCI)—we introduce a fundamentally different paradigm: a self-supervised One-Step Diffusion (OSD) framework designed specifically for SCI. The key novelty lies in using a single-step diffusion refiner to correct an initial reconstruction, eliminating iterative denoising entirely while preserving generative quality. Moreover, we adopt a self-supervised equivariant learning strategy to train both the predictor and refiner directly from raw 2-D measurements, enabling generalization to unseen domains without ground-truth MSI. To further address limited MSI data, we design a band-selection-driven distillation strategy that transfers core generative priors from large-scale RGB datasets, effectively bridging the domain gap. Extensive experiments confirm that our approach sets a new standard—yielding PSNR gains of 3.44dB, 1.61dB, and 0.28dB on the Harvard, NTIRE, and ICVL datasets respectively, while cutting reconstruction time from 8.9s to just 0.22s per image. These gains in efficiency and adaptability advance SCI reconstruction, enabling accurate and practical real-world deployment.

Code — <https://github.com/shhuang-1767/DiFASCI>

Introduction

Multispectral images (MSIs) record richer spectral information than RGB, allowing better material discrimination and supporting applications such as land cover classification (Liu et al. 2024), anomaly detection (Xu et al. 2015), and material identification (Hagen and Kudenov 2013). However, conventional MSI acquisition relies on spatial or temporal scanning and is slow. Driven by compressive sensing theory, snapshot compressive imaging (SCI) (Wagadarikar et al. 2009) addresses this by capturing a single

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1. Comparison with the latest diffusion-based SCI reconstruction methods on zero-shot datasets Harvard, NTIRE and ICVL in terms of (a) PSNR and (b) inference time.

coded 2D measurement in one shot, enabling efficient MSI acquisition via computational reconstruction.

Restoring a 3-D MSI from the 2-D measurement in SCI systems is an ill-posed inverse problem. Existing SCI restoration methods can be categorized into model-based and deep-learning-based approaches. The traditional model-based methods typically solve the inverse problem within a convex optimization framework with hand-crafted image priors. They enjoy good theory guarantees but require manual parameter tuning. Current deep-learning-based approaches include plug-and-play (PnP) based (Meng et al. 2021a), end-to-end (E2E) based (Zhang et al. 2024b) and deep unfolding-based methods (Zhang et al. 2024a). PnP methods integrate pre-trained denoising networks into the traditional optimization-based framework. Due to the fixed pre-trained denoising model, the performance of PnP methods is often limited. E2E methods learn a mapping function from 2-D measurement to 3-D MSI through supervised learning by using paired data. However, they often neglect the physical characteristics of SCI systems and thus are less interpretable and robust to different hardware systems. Deep unfolding-based methods unfold a convex optimization process by using multistage neural networks to restore MSI, resulting in an interpretable neural network architecture. However, in the inadequate sampling areas of 2-D measurement, existing methods often struggle to obtain accurate restoration with fine details (Zeng et al. 2025).

Diffusion model has yielded remarkable success in generating content from RGB images (Ho, Jain, and Abbeel

2020). Recent studies (Wu et al. 2025; Pan et al. 2024; Zeng et al. 2025) leverage its generative capacity to enhance the quality of MSI in SCI, achieving promising results (Zeng et al. 2024). DiffSCI (Pan et al. 2024) employs a pre-trained diffusion model on large RGB datasets as a generative denoiser within the PnP framework. Zeng et al. propose a diffusion-based two-stage framework where the high-frequency component of MSIs is refined by a fine-tuned diffusion model in the low dimensional subspace (Zeng et al. 2025). Another recent work (Wu et al. 2025) adopts a latent diffusion model to generate image priors, which are fed into a deep unfolding network for high-quality MSI reconstruction. Despite the encouraging results, the following issues remain: 1) The generation ability of the used diffusion model is limited due to the domain shift between RGB training data and MSI. Although the methods in (Zeng et al. 2025; Wu et al. 2025) fine-tune or retrain the diffusion models with paired measurements and MSIs, the scarcity of paired MSI data (e.g., merely a few thousand MSIs versus hundreds of millions of RGB images, and sometimes completely inaccessible) leads to suboptimal performance in the derived models’ generative capabilities. 2) The iterative nature of diffusion models necessitates significant sampling time, a limitation that becomes more pronounced in MSIs due to the high dimensionality of MSI.

To address these problems, we propose an efficient self-supervised one-step diffusion refinement framework to enhance the details of MSI. We leverage an existing SCI reconstruction network to generate an initial prediction of MSI from the input 2-D measurement and use a one-step diffusion model to generate an MSI residual for refinement. To ensure robust performance on unseen datasets, the SCI reconstruction network is co-trained with the one-step diffusion model. The coarse-to-fine strategy enables the use of a compact diffusion network by focusing on residual generation, which is significantly simpler than reconstructing the full MSI. To achieve fast one-step generation without compromising quality compared to multi-step diffusion models, we propose a spectral compression distillation loss for knowledge transfer from RGB images to MSI. Additionally, to address the lack of paired MSIs, we introduce a self-supervised learning paradigm based on the equivariant imaging framework, enabling network training solely with 2D measurements. The main contributions can be summarized as follows:

- 1) We propose an efficient self-supervised one-step diffusion refinement framework for SCI reconstruction. Our approach requires no paired MSIs, making it highly suitable for real-world applications where only 2D measurements are accessible. The framework also exhibits excellent generalization capabilities, as it can be flexibly combined with any E2E and deep unfolding-based SCI reconstruction networks.
- 2) We design a one-step diffusion adaptation (DiFA) module to enhance the details of the reconstructed MSIs from an initial predictor. It simplifies the whole modeling process while facilitating faster inference speed compared with the existing diffusion-based SCI reconstruction methods.

tion methods.

- 3) To leverage the prior knowledge from pre-trained multi-step diffusion models, we propose spectral compression distillation in RGB space, enabling fast one-step MSI generation without compromising quality.
- 4) Due to the lack of paired MSIs, recovering MSI is impossible solely from the 2-D measurement with a data fidelity loss. We develop an additional self-supervised training loss based on the transformation invariant assumption on MSI, allowing the model to be trained only with the compressed measurements.

Related Works

MSI Restoration Methods

Existing SCI reconstruction frameworks mainly include model-based and deep learning-based methods. Model-based methods (Zhang et al. 2019; Zeng et al. 2020) restore MSI within an optimization-based framework. To cope with the ill-posed inverse problem, they typically impose hand-crafted constraints such as sparsity, low-rank, graph Laplacian and total variation on the restored data by taking into account image priors. These methods are interpretable but often show low reconstruction quality and speed.

End-to-end and deep unfolding are two representative deep learning-based frameworks for SCI reconstruction. The end-to-end approaches (Fu et al. 2021; Xiong et al. 2017; Meng, Ma, and Yuan 2020; Miao et al. 2019; Cheng et al. 2022; Cai et al. 2022b,a; Zhang et al. 2024b) directly learn a mapping function on paired MSIs to restore MSIs from the 2-D compressed measurements. Deep unfolding methods (Zhang et al. 2021, 2022; Huang et al. 2021; Meng, Jalali, and Yuan 2020; Ma et al. 2019; Wang et al. 2020; Cai et al. 2022c; Li et al. 2023; Dong et al. 2023; Zhang et al. 2024a) unfold the iterative optimization-based reconstruction process into a multi-stage network, where each stage corresponds to one iteration of the optimization algorithm. However, these methods typically involve minimizing ℓ_1 or ℓ_2 based pixel loss and thus encounter the regression-to-mean issue, facing limitations in reconstructing high-frequency details (Wu et al. 2025).

Deep generative models such as the DDPM (Ho, Jain, and Abbeel 2020; Nichol and Dhariwal 2021) directly sample from the posterior, where a set of specific candidates can be generated from the learned posterior instead of calculating the average of all possible solutions, showing great potential for MSI reconstruction. Recently, attempts (Wu et al. 2025; Pan et al. 2024; Zeng et al. 2025) have been made to leverage the generative ability of diffusion models to enhance the reconstruction of MSI. Specifically, DiffSCI (Pan et al. 2024) employs a pre-trained RGB image denoising diffusion model as the denoiser in a PnP framework. Zeng et al. develop a two-stage SCI reconstruction method PSR-SCI (Zeng et al. 2025) where a diffusion model is used to enhance the high-frequency component of the MSI in a lower subspace derived with unmixing. Another recent work (Wu et al. 2025) integrates a latent diffusion model to generate clean image priors for a deep unfolding network, resulting in high-quality data reconstruction.

ResShift

ResShift (Yue, Wang, and Loy 2025) is a diffusion model starting with a prior distribution centered on the low-quality (LQ) image, which allows the high-quality (HQ) image to be iteratively recovered from its LQ counterpart instead of Gaussian white noise. The forward process can be expressed as follows:

$$q(x_t | x_0, y) = \mathcal{N}(x_t; x_0 + \eta_t(y - x), \kappa^2 \eta_t \mathbf{I}), \quad (1)$$

where y is a LR image, x_0 is a HR image, κ is a hyper-parameter for the noise variance and η_t is a shifting sequence, which is monotonically increasing and satisfied $\eta_1 \rightarrow 0$ and $\eta_T \rightarrow 1$. Based on this forward process, the reverse process will commence from the initial state with rich information in LQ images to perform denoising. The formula is as follows:

$$p_\theta(x_{t-1} | x_t, y) = \mathcal{N}\left(x_{t-1} \left| \frac{\eta_{t-1}}{\eta_t} x_t + \frac{\alpha_t}{\eta_t} x_0, \kappa^2 \frac{\eta_{t-1}}{\eta_t} \alpha_t \mathbf{I} \right.\right), \quad (2)$$

where $\alpha_t = \eta_t - \eta_{t-1}$ for $t > 1$ and $\alpha_1 = \eta_1$, and an initial state is $x_T = y + \kappa \sqrt{\eta_T} \epsilon$, where $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. In the reverse process in Eq. (2), x_0 is predicted by a trainable deep neural network f_θ . A non-Markovian reverse process exists which produces deterministic sampling and is reformulated as:

$$q(x_{t-1} | x_t, x_0, y) = \delta(k_t x_0 + m_t x_t + j_t y), \quad (3)$$

where δ is the unit impulse and k_t, x_t, j_t and m_t are derived from η_t . More details can be found in (Wang et al. 2024). Eq. (3) indicates a deterministic mapping from x_T to x_0 , expressed as $F_\theta(x_T, y)$.

Especially, the reverse process in one-step generation can be formulated as:

$$x_0 = f_\theta(x_T, y, T). \quad (4)$$

where the HQ image x_0 is generated by the trainable network f_θ directly, and T is the total timestep of ResShift.

The Proposed Method

Problem Definition

Degradation Model of CASSI. One of the representative SCI systems is the Coded Aperture Snapshot Spectral Compressive Imaging (CASSI) system (Gehm et al. 2007). It captures 3-D spectral data in a single snapshot by encoding spatial and spectral information through a coded aperture and reconstructing it from compressed 2-D measurements.

Let $\mathcal{X} \in \mathbb{R}^{H \times W \times B}$ denote a 3-D MSI and $\mathcal{Y} \in \mathbb{R}^{H \times W + d \times (B-1)}$ be the 2-D measurements, where H, W, d , and B are the MSI's height, width, shifting step, and total number of wavelengths, respectively. To obtain a concise equation for the imaging process, let $\mathbf{y} \in R^n$ with $n = H(W + d(B - 1))$ represent the vectorized measurement, $\mathbf{x} \in R^{nB}$ be the vectorized shifted MSI and $\mathcal{H} \in \mathbb{R}^{n \times nB}$ be the mask. Then, we obtain the imaging process in CASSI (Meng, Ma, and Yuan 2020; Ma et al. 2019):

$$\mathbf{y} = \mathcal{H}\mathbf{x} + \mathbf{n}, \quad (5)$$

where $\mathbf{n} \in R^n$ represents the noise on measurement.

Then, the next step is to obtain \mathbf{x} given the obtained measurement \mathbf{y} and \mathcal{H} .

Challenges with Diffusion Model for SCI. Leveraging the generative power of diffusion models offers a promising solution to improve the detail of MSIs in SCI reconstruction. However, 1) directly integrating the existing pre-trained diffusion models, which are typically trained with RGB datasets, into the SCI reconstruction task is infeasible due to the band mismatch. 2) Training a diffusion model from scratch for MSI is challenging due to the limited MSI. The total amount of paired MSI in benchmark datasets is around only a thousand, which is far smaller than the RGB data. The difficulty intensifies when paired MSIs are entirely absent. 3) The inference process is time-consuming due to the high dimensionality of MSI and the substantial number of iterative sampling steps required.

The Self-supervised One-step Diffusion Refinement for SCI

Here, we develop a self-supervised one-step diffusion refinement framework as shown in Fig. 2 (a) where a core module DiFA is designed to obtain MSI with the input of \mathbf{y} and \mathcal{H} . As shown in Fig. 2 (b), we feed \mathbf{y} and \mathcal{H} into an initial predictor g_θ to obtain an initial prediction of MSI $\mathbf{x}_{\text{init}} = g_\theta(\mathbf{y}, \mathcal{H})$. Then, we refine \mathbf{x}_{init} by using a residual $\mathbf{r} = f_\theta(\mathbf{x}_{\text{init}})$ generated by a one-step diffusion model f_θ , leading to an enhanced MSI $\mathbf{x}^{(1)} = \mathbf{x}_{\text{init}} + \mathbf{r}$. The one-step diffusion model is trained from scratch by using only the measurement with self-supervised learning. Then, we compress the obtained MSI $\mathbf{x}^{(1)}$ into the RGB space and employ a distillation technique to leverage the prior knowledge of pre-trained multi-step diffusion model to guide the learning of the one-step diffusion model and the initial predictor. In response to the unavailability of paired MSIs in real applications, we develop a self-supervised loss that allows efficient model training with only 2-D measurement.

Efficient One-step Diffusion Adaption. Using a multi-step diffusion model for SCI reconstruction is time-consuming due to the high dimensionality of MSI and the large number of sampling steps. In addition, modeling the mapping process from 2-D measurement to 3-D MSI directly is challenging. Thus, we propose an efficient one-step diffusion adaption module DiFA, which first generates an initial prediction of MSI by an existing SCI reconstruction network and then refines it with details provided by a one-step diffusion network. Unlike the iterative noise prediction in the DDPM (Ho, Jain, and Abbeel 2020) reverse process, our approach employs a single-step diffusion to directly generate a clean image from an initial state, leading to a fast inference speed.

Let $\mathcal{F}_\theta = g_\theta(\cdot) + f_\theta(\cdot)$ denote the DiFA as shown in Fig. 2 (b), where $g_\theta(\cdot)$ and $f_\theta(\cdot)$ are the initial predictor and one-step diffusion model with θ represents the weights of neural networks. Note that our focus is not on the specific design of the initial predictor. Considering the availability of existing pre-trained SCI reconstruction neural networks, we can use an existing one as the initial predictor, which will be trained with self-supervised learning, for an initial guess of the restored MSI: $\mathbf{x}_{\text{init}} = g_\theta(\mathbf{y})$.

Under the predict-and-refine pipeline, our OSD aims to

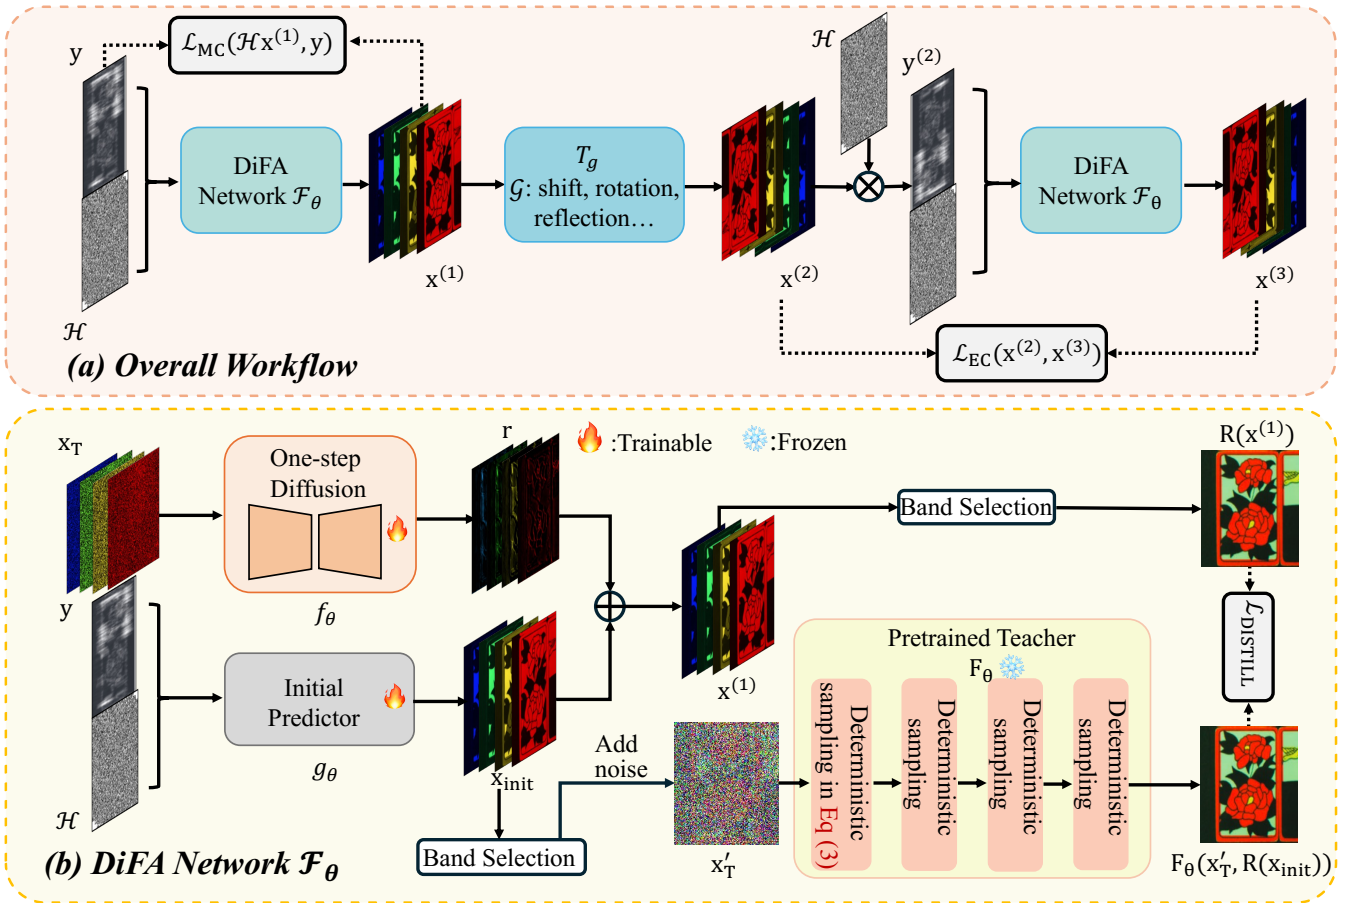


Figure 2. (a) Workflow for our self-supervised training strategy. The measurement y and mask \mathcal{H} are initially input into DiFA Network \mathcal{F}_θ , resulting in the recovered MSI $x^{(1)}$. Next, a series of transformations T_g containing shift, rotation, reflection, etc. are applied to $x^{(1)}$ to produce $x^{(2)}$. The MSI $x^{(2)}$ is then modulated again by the mask \mathcal{H} to obtain the compressed measurement $y^{(2)}$, which is finally input into \mathcal{F}_θ to obtain the re-reconstructed MSI $x^{(3)}$. (b) DiFA Network, on the one hand, the measurement y and mask \mathcal{H} are initially input into pre-trained reconstruction network g_θ to get the initial predictor x_{init} , the residual of MSI r is generated from noise x_T by one-step diffusion f_θ . Then, the refined image $x^{(1)}$ is obtained by adding the residual r to the initial prediction x_{init} . On the other hand, both x_{init} and $x^{(1)}$ are converted to RGB images through band selection $R(\cdot)$. A distillation loss is developed to leverage the prior knowledge of pre-trained multistep diffusion models.

estimate a 3-D residual r through a rough 3-D MSI, a similar task in the ResShift (Yue, Wang, and Loy 2025), which recovers a HR RGB image from a LQ RGB image. However, directly employing the ResShift as our OSD is infeasible because of the disparity of bands between RGB and MSI. Retraining one from scratch for MSI is challenging due to the limited paired MSIs. Here, we design a self-supervised training method by incorporating distillation and equivariant consistency to jointly learn the OSD and initial predictor with solely the 2-D measurement. Details of this strategy are described in the following sections.

Based on Eq. (4), we estimate the residual r as follows:

$$r = f_\theta(x_T, x_{init}, T), \quad (6)$$

where $x_T = x_{init} + \kappa\sqrt{\eta T}\epsilon$ is the initial state with $\epsilon \sim \mathcal{N}(0, 1)$ representing noise sampled from a standard normal distribution. We set $T = 4$ in our paper. The visual results

in Fig. 3 show that the refined MSI, produced by the initial predictor and OSD, exhibits superior visual quality with finer details and enhanced sharpness.

Spectral Compression Distillation. The one-step diffusion model is trained from scratch for the residual estimation without using any paired MSIs. While it shows faster inference speed, its generative capability might be sacrificed. To alleviate this problem, we design a distillation method to leverage the diffusion prior from pre-trained multi-stage diffusion models on RGB data to guide the training of the one-step diffusion model and the refinement of the initial predictor. However, due to the disparity of bands between MSI and RGB images, direct distillation is infeasible. To project the MSIs into the RGB space, we select three bands at 625.0 nm, 522.5 nm, and 476.5 nm corresponding to the red, green, and blue channels, respectively. This projection




Figure 3. The refined MSI, obtained with the initial and residual MSIs, exhibits superior visual quality with finer details compared to the baseline DAUHST.

process is denoted as $R(\cdot)$. Then, we perform distillation in the RGB space as follows:

$$\mathcal{L}_{DISTILL} = L_{MSE}(R(\mathbf{x}^{(1)}), F_{\theta}(\mathbf{x}'_T, R(\mathbf{x}_{init}))), \quad (7)$$

where F_{θ} is a pre-trained multi-step diffusion model (Yue, Wang, and Loy 2025) and $\mathbf{x}'_T = R(\mathbf{x}_{init}) + \kappa\sqrt{\eta_T}\epsilon$ with $\epsilon \sim \mathcal{N}(0, 1)$ is the initial state in the RGB space.

Equivariant Imaging Diffusion Consistency. Due to the costly nature of MSI acquisition and the lack of paired MSIs for training the one-step diffusion network in real scenarios, training a DiFA network for SCI reconstruction using only compressed measurement data becomes imperative.

Measurement Consistency Loss. Consider a naive unsupervised loss that only enforces measurement consistency:

$$\mathcal{L}_{MC} = \|\mathbf{y} - \mathcal{H}\mathcal{F}_{\theta}(\mathbf{y}, \mathcal{H})\|^2 = \|\mathbf{y} - \mathcal{H}\mathbf{x}^{(1)}\|^2, \quad (8)$$

where $\mathbf{x}^{(1)} = \mathcal{F}_{\theta}(\mathbf{y}, \mathcal{H})$ represents the image recovered by the DiFA Network as shown in Fig. 2 (b), and $\mathcal{H}\mathbf{x}^{(1)}$ represents the predicted measurements.

If the measurement process \mathcal{H}^T is incomplete, then even in the absence of noise, it is fundamentally impossible to learn a complete MSIs solely from the measurement \mathbf{y} , as there is no information in the null space of the measurement process \mathcal{H} . Thus, we need to learn more information beyond the range space of their inverse (Gleichman and Eldar 2011).

Equivariant Consistency. Recently, the EI framework (Chen, Tachella, and Davies 2021) showed that learning with only measurement data \mathbf{y} is possible with an additional transformation invariant assumption on the signal \mathcal{X} . That is, for a certain group of transformations (i.e., shifts, rotations, etc.) $\mathcal{G} = \{g_1, \dots, g_{|\mathcal{G}|}\}$ which are unitary matrices $T_g \in \mathcal{G}$, if $\forall \mathbf{x} \in \mathcal{X}$, we have $T_g\mathbf{x} \in \mathcal{X}$ for $\forall g \in \mathcal{G}$ and the sets $T_g\mathcal{X}$ and \mathcal{X} are the same.

With the invariance assumption, the following equations should be met in our method:

$$\mathcal{F}_{\theta}(\mathcal{H}T_g\mathbf{x}) = T_g\mathcal{F}_{\theta}(\mathcal{H}\mathbf{x}), \quad (9)$$

for $\forall g \in \mathcal{G}$ and $\forall \mathbf{x} \in \mathcal{X}$. This indicates that the composition $\mathcal{F}_{\theta} \circ \mathcal{H}$ should be transformation invariant.

After obtaining the estimated MSI $\mathbf{x}^{(1)} = \mathcal{F}_{\theta}(\mathbf{y}, \mathcal{H})$, based on the transformation invariant property, we obtain

$\mathbf{x}^{(2)} = T_g\mathbf{x}^{(1)}$ and subsequently feed it to \mathcal{H} and \mathcal{F}_{θ} as illustrated in Fig. 2 (a), resulting in a recovered MSI:

$$\mathbf{x}^{(3)} = \mathcal{F}_{\theta}(\mathcal{H}\mathbf{x}^{(2)}) = \mathcal{F}_{\theta}(\mathcal{H}T_g\mathbf{x}^{(1)}), \quad (10)$$

which is the estimation of $\mathbf{x}^{(2)}$. Thus, our equivariant consistency (EC) loss is formulated as:

$$\mathcal{L}_{EC} = \|\mathbf{x}^{(2)} - \mathbf{x}^{(3)}\|^2 = \|T_g\mathcal{F}_{\theta}(\mathbf{y}) - \mathcal{F}_{\theta}(\mathcal{H}(T_g\mathcal{F}_{\theta}(\mathbf{y})))\|^2. \quad (11)$$

The EC loss in Eq. (11) incorporates the transformation invariant prior information of \mathcal{X} , allowing us to learn additional information that is beyond the range space of \mathcal{H}^T , which is impossible by using the MC loss alone in Eq. (8).

The Overall Training Loss. By combining the aforementioned losses, the overall objective function \mathcal{L}_{Total} is:

$$\mathcal{L}_{Total} = \mathcal{L}_{MC} + \alpha\mathcal{L}_{EC} + \beta\mathcal{L}_{DISTILL}, \quad (12)$$

where α , and β are hyper-parameters. The pre-trained network g_{θ} and the one-step diffusion f_{θ} are jointly trained by minimizing Eq. (12) and the teacher network F_{θ} is frozen.

Experiments

Experiment Setting

Simulated Dataset. Three benchmark MSI datasets NTIRE (Arad et al. 2022), ICVL (Arad and Ben-Shahar 2016), and Harvard (Choi et al. 2017) are used here. Following (Meng, Ma, and Yuan 2020), we obtain the simulated 2-D compressed measurement \mathbf{y} for the three datasets. In the experiments, the number of bands is 28 and the image size is 256×256 for NTIRE, ICVL, and Harvard.

Real Dataset. Five real MSIs collected by the CASSI system (Meng, Ma, and Yuan 2020) are used in the experiments.

Implementation Detail. The number of training steps is set to 50,000. And we set $\alpha = 1$, and $\beta = 0.001$ in Eq. (12).

The denoising network f_{θ} for one-step diffusion is a U-Net incorporating several Swin Transformer blocks (Yue, Wang, and Loy 2025). An Nvidia RTX 4090 GPU is used for model training. Our method is implemented by PyTorch.

Compared Methods. We compare our method with 15 state-of-the-art methods, including the end-to-end networks (λ -Net (Miao et al. 2019), TSA-Net (Meng, Ma, and Yuan 2020), BiSRNet (Cai et al. 2023), CST-L (Cai et al. 2022a), MST-L (Cai et al. 2022b), HDNet (Hu et al. 2022) and SSR-L (Zhang et al. 2024b)), the PnP networks (DIP-MSI (Meng et al. 2021b)), the deep unfolding methods (ADMM-Net (Ma et al. 2019), DAUHST-9stg (Cai et al. 2022c), PADUT-3stg (Li et al. 2023), RDLUF-MixS2-9stg (Dong et al. 2023) and DPU-9stg (Zhang et al. 2024a)) and diffusion-based methods (DiffSCI (Pan et al. 2024), LADE-10stg (Wu et al. 2025) and PSR-SCI-D (Zeng et al. 2025)). Since our method is flexible and can incorporate existing end-to-end and deep unfolding SCI reconstruction neural networks as the initial predictor, we employ the MST-L, HDNet, SSR-L, DAUHST-9stg, PADUT-3stg, and DPU-9stg networks that are trained on CAVE dataset as our initial predictor and refer to our methods as MST-DiFA, HDNet-DiFA, SSR-DiFA, DAUHST-DiFA, PADUT-DiFA and DPU-DiFA, respectively. Note that these pre-trained networks were not trained on the test datasets (ICVL,

Methods	Category	Reference	ICVL	NTIRE	Harvard
λ -Net (Miao et al. 2019)	CNN	ICCV 2019	26.62 / 0.763	24.20 / 0.656	26.84 / 0.658
ADMM-Net (Ma et al. 2019)	Deep Unfolding	ICCV 2019	26.44 / 0.775	25.47 / 0.748	28.22 / 0.741
TSA-Net (Meng, Ma, and Yuan 2020)	CNN	ECCV 2020	24.15 / 0.592	21.63 / 0.577	19.92 / 0.408
DIP-HSI (Meng et al. 2021b)	PnP	ICCV 2021	31.53 / 0.832	28.71 / 0.829	27.56 / 0.712
CST-L+ (Cai et al. 2022a)	Transformer	ECCV 2022	25.68 / 0.773	24.39 / 0.768	27.75 / 0.724
BiSRNet (Cai et al. 2023)	BNN	NeurIPS 2023	24.71 / 0.707	22.49 / 0.637	25.59 / 0.624
RDLUF-MixS2-9stg (Dong et al. 2023)	Deep Unfolding	CVPR 2023	35.89 / 0.945	32.17 / 0.909	31.04 / 0.820
MST-L (Cai et al. 2022b)	Transformer	CVPR 2022	27.38 / 0.804	25.24 / 0.775	27.49 / 0.706
MST-DiFA	Diffusion	Ours	32.81 / 0.894	30.86 / 0.880	30.31 / 0.796
HDNet (Hu et al. 2022)	Transformer	CVPR 2022	25.71 / 0.749	24.16 / 0.749	26.99 / 0.687
HDNet-DiFA	Diffusion	Ours	30.62 / 0.822	28.08 / 0.784	28.15 / 0.700
DAUHST-9stg (Cai et al. 2022c)	Deep Unfolding	NuerIPS 2022	30.21 / 0.886	29.08 / 0.878	30.61 / 0.809
DAUHST-DiFA	Diffusion	Ours	37.19 / 0.952	34.71 / 0.939	31.93 / 0.837
PADUT-3stg (Li et al. 2023)	Deep Unfolding	ICCV 2023	29.43 / 0.866	27.74 / 0.839	28.94 / 0.768
PADUT-DiFA	Diffusion	Ours	35.07 / 0.919	32.51 / 0.897	30.33 / 0.793
DPU-9stg (Zhang et al. 2024a)	Deep Unfolding	CVPR 2024	33.76 / 0.934	32.28 / 0.920	31.64 / 0.829
DPU-DiFA	Diffusion	Ours	36.39 / 0.918	33.25 / 0.927	32.15 / 0.849
SSR-L (Zhang et al. 2024b)	Transformer	CVPR 2024	36.27 / 0.947	34.10 / 0.938	31.34 / 0.823
SSR-DiFA	Diffusion	Ours	37.42 / 0.941	37.14 / 0.964	32.46 / 0.857

Table 1. PSNR/SSIM of different methods on three zero-shot datasets. The average results of ten images are reported.

Methods	DiffSCI	LADE-10stg	PSR-SCI-D	SSR-DiFA
Reference	CVPR 2024	ECCV 2024	ICLR 2025	Ours
Inference Complexity	50 84.54s	8 1.10s	50 8.90s	1 0.22s
Quality Index	ICVL NTIRE Harvard	33.02/0.868 33.58/0.923 28.02/0.739	37.14/0.918 35.53/0.948 29.02/0.728	37.42/0.941 37.14/0.964 32.46/0.857

Table 2. Inference time (a 256×256 MSI) and average reconstruction quality (PSNR/SSIM over 10 images) of diffusion-based methods on zero-shot datasets.

NTIRE, and Harvard), and all experiments are conducted in the zero-shot setting.

Evaluation Metrics. We evaluate the performance of different methods with the peak signal-to-noise ratio (PSNR) and the structural similarity index metrics (SSIM).

Quantitative Results

We show the quantitative results of different methods on NTIRE, ICVL and Harvard datasets in Tab. 1 and 2. Our method SSR-DiFA achieves the best performance in terms of PSNR and SSIM on all three datasets. It is observed in Tab. 1 that all our methods outperform their counterparts. This indicates that our method is highly generalized to varying initial predictor.

In Tab. 2, it is observed that compared with the latest SOTA methods DiffSCI, LADE-10stg and PSR-SCI-D, which are the current leading diffusion-based methods in SCI, we obtain the best performance on all three datasets, demonstrating the effectiveness of our design. Compared with LADE-10stg and PSR-SCI-D, which retrain or fine-tune pre-trained diffusion models with paired MSIs, we train a one-step diffusion model from scratch in a self-supervised manner, which only uses 2-D measurement for training, which are more applicable in real applications. Moreover, our method yields the fastest inference speed of 0.22 seconds among the diffusion-based methods, demonstrating the efficiency of our approach.

Case	DiFA	MC	EC	DSTILL	HDNet-DiFA	MST-DiFA	PADUT-DiFA	DAUHST-DiFA
(a)	×	×	×	×	24.16 / 0.749	25.24 / 0.775	27.74 / 0.839	29.08 / 0.878
(b)	✓	×	×	×	26.41 / 0.765	26.97 / 0.800	28.23 / 0.848	29.75 / 0.889
(c)	✓	✓	×	×	26.82 / 0.779	27.05 / 0.801	28.60 / 0.851	29.81 / 0.890
(d)	✓	✓	✓	×	27.95 / 0.782	30.52 / 0.868	32.44 / 0.892	34.50 / 0.937
(e)	✓	✓	✓	✓	28.08 / 0.784	30.86 / 0.880	32.51 / 0.897	34.71 / 0.939
Methods	HDNet-DiFA	MST-DiFA	PADUT-DiFA	DAUHST-DiFA				
Baseline	24.16 / 0.749	25.24 / 0.775	27.74 / 0.839	29.08 / 0.878				
w/o OSD	27.42 / 0.810	28.83 / 0.834	31.96 / 0.903	33.11 / 0.914				
with OSD	28.08 / 0.784	30.86 / 0.880	32.51 / 0.897	34.71 / 0.939				

Table 3. Ablation study on different components (top) and OSD (bottom) on zero-shot dataset NTIRE.

Qualitative Results

We show the visual results of different methods on the simulated datasets NTIRE, ICVL, and Harvard in Fig. 4. We can see that our DiFA provides better visual results with more details and fewer artifacts than the original methods. Moreover, we show the spectral curves of different methods on the right of Fig. 4. It is observed that our DiFA strategy effectively improves the spectral reconstruction accuracy, demonstrating the superior spectral reconstruction ability of DiFA. In addition, we show the visual results of MST, DAUHST and our methods MST-DiFA and DAUHST-DiFA on the real data set in Fig. 5. Compared with the original methods, our approaches obtain improved visual quality as demonstrated in the region of strawberry.

Experimental analysis

Ablation Study. There are three key components in our method, including DiFA, the EC loss \mathcal{L}_{EC} and the distillation loss $\mathcal{L}_{DISTILL}$. In addition, we investigate the influence of the data fidelity MSE loss \mathcal{L}_{MC} . We remove different components from our methods and conduct experiments on NTIRE with four different initial predictors HDNet, MST, PADUT and DAUHST. The results are shown in Tab. 3 top. Case (a) represents the original method. It is observed that compared with the baselines, i.e., case (a), all other cases obtain improved performance. Our final methods

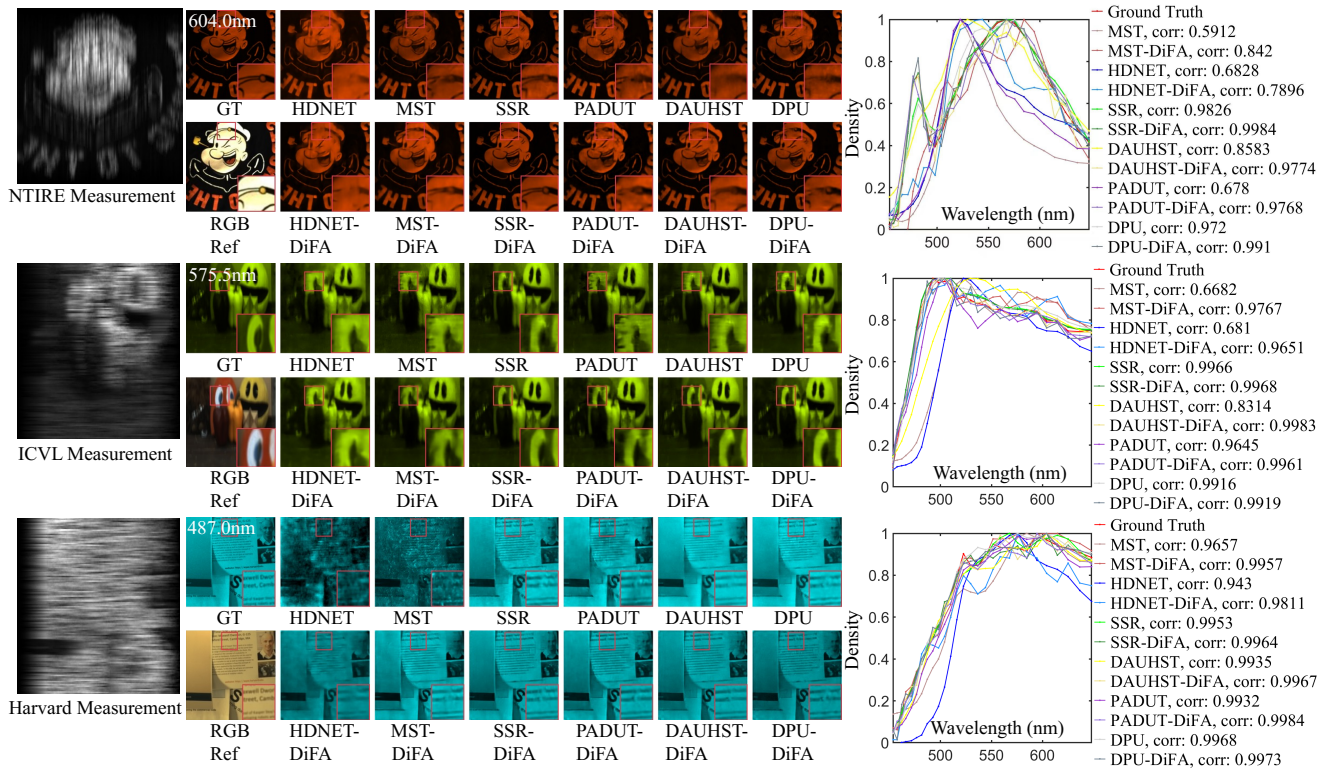


Figure 4. Simulated MSI reconstruction comparisons on zero-shot datasets NTIRE, ICVL and Harvard. The right shows the reconstructed spectral curves corresponding to the selected region.




Figure 5. Real MSI reconstruction results on Scene 4 with 6 spectral channels. Please zoom in for a better view.

in case (e) with all components yield the best performance, demonstrating the effectiveness of each design.

The Influence of the One-step Diffusion in DiFA. The OSD model in DiFA is used to estimate the details of MSI for the refinement of the initial predicted MSI. To verify its effectiveness, we conduct experiments on zero-shot dataset NTIRE by removing the one-step diffusion from our approach. Our four methods with different initial predictors are tested. The results are shown in Tab. 3 bottom. It is observed that without using OSD, the PSNR performance is dropped. For instance, the PSNR is reduced by 2.03dB for MST-DiFA. It should be noted that compared with the base-

lines HDNet, MST-L, PADUT-3stg and DAUHST-9stg, our reduced methods without OSD still obtain significant performance improvement.

Conclusion

In this paper, we propose a novel self-supervised one-step diffusion refinement framework for SCI, which reconstructs a high-quality MSI with an initial predictor and an OSD model. The OSD model provides details for the rough prediction obtained by the initial predictor, leading to a high-quality MSI with fine details. We design a distillation loss in the RGB space for knowledge transfer from the RGB domain to the MSI domain and devise a self-supervised learning method based on a transformation invariant property, which allows effective model training with only 2-D measurement. Our method achieves state-of-the-art performance on both simulated and real data, offering faster inference and excellent generalization compared to the latest diffusion-based methods in the field.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under grants 42301425 and T2525018, in part by the China Postdoctoral Science Foundation (2023M743299), in part by the ‘‘CUG Scholar’’ Scientific Research Funds at China University of Geosciences (Wuhan) under grant 2022164, and in part by the Fundamen-

tal Research Funds for the Central Universities, China University of Geosciences (Wuhan) under grant CUG240628.

References

- Arad, B.; and Ben-Shahar, O. 2016. Sparse recovery of hyperspectral signal from natural RGB images. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VII 14*, 19–34. Springer.
- Arad, B.; Timofte, R.; Yahel, R.; Morag, N.; Bernat, A.; Cai, Y.; Lin, J.; Lin, Z.; Wang, H.; Zhang, Y.; et al. 2022. Ntire 2022 spectral recovery challenge and data set. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 863–881.
- Cai, Y.; Lin, J.; Hu, X.; Wang, H.; Yuan, X.; Zhang, Y.; Timofte, R.; and Van Gool, L. 2022a. Coarse-to-fine sparse transformer for hyperspectral image reconstruction. In *European Conference on Computer Vision*, 686–704. Springer.
- Cai, Y.; Lin, J.; Hu, X.; Wang, H.; Yuan, X.; Zhang, Y.; Timofte, R.; and Van Gool, L. 2022b. Mask-guided spectral-wise transformer for efficient hyperspectral image reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17502–17511.
- Cai, Y.; Lin, J.; Wang, H.; Yuan, X.; Ding, H.; Zhang, Y.; Timofte, R.; and Gool, L. V. 2022c. Degradation-aware unfolding half-shuffle transformer for spectral compressive imaging. *Advances in Neural Information Processing Systems*, 35: 37749–37761.
- Cai, Y.; Zheng, Y.; Lin, J.; Yuan, X.; Zhang, Y.; and Wang, H. 2023. Binarized spectral compressive imaging. *Advances in Neural Information Processing Systems*, 36: 38335–38346.
- Chen, D.; Tachella, J.; and Davies, M. E. 2021. Equivariant imaging: Learning beyond the range space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4379–4388.
- Cheng, Z.; Chen, B.; Lu, R.; Wang, Z.; Zhang, H.; Meng, Z.; and Yuan, X. 2022. Recurrent neural networks for snapshot compressive imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 2264–2281.
- Choi, I.; Kim, M.; Gutierrez, D.; Jeon, D.; and Nam, G. 2017. High-quality hyperspectral reconstruction using a spectral prior. Technical report.
- Dong, Y.; Gao, D.; Qiu, T.; Li, Y.; Yang, M.; and Shi, G. 2023. Residual degradation learning unfolding framework with mixing priors across spectral and spatial for compressive spectral imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22262–22271.
- Fu, Y.; Zhang, T.; Wang, L.; and Huang, H. 2021. Coded hyperspectral image reconstruction using deep external and internal learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7): 3404–3420.
- Gehm, M. E.; John, R.; Brady, D. J.; Willett, R. M.; and Schulz, T. J. 2007. Single-shot compressive spectral imaging with a dual-disperser architecture. *Optics Express*, 15(21): 14013–14027.
- Gleichman, S.; and Eldar, Y. C. 2011. Blind compressed sensing. *IEEE Transactions on Information Theory*, 57(10): 6958–6975.
- Hagen, N.; and Kudenov, M. W. 2013. Review of snapshot spectral imaging technologies. *Optical Engineering*, 52(9): 090901–090901.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- 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*, 17542–17551.
- Huang, T.; Dong, W.; Yuan, X.; Wu, J.; and Shi, G. 2021. Deep gaussian scale mixture prior for spectral compressive imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16216–16225.
- Li, M.; Fu, Y.; Liu, J.; and Zhang, Y. 2023. Pixel adaptive deep unfolding transformer for hyperspectral image reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12959–12968.
- Liu, R.; Luo, T.; Huang, S.; Wu, Y.; Jiang, Z.; and Zhang, H. 2024. CrossMatch: Cross-View Matching for Semi-Supervised Remote Sensing Image Segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–15.
- Ma, J.; Liu, X.-Y.; Shou, Z.; and Yuan, X. 2019. Deep tensor admm-net for snapshot compressive imaging. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10223–10232.
- Meng, Z.; Jalali, S.; and Yuan, X. 2020. Gap-net for snapshot compressive imaging. *arXiv preprint arXiv:2012.08364*.
- Meng, Z.; Ma, J.; and Yuan, X. 2020. End-to-end low cost compressive spectral imaging with spatial-spectral self-attention. In *European Conference on Computer Vision*, 187–204. Springer.
- Meng, Z.; Yu, Z.; Xu, K.; and Yuan, X. 2021a. Self-supervised neural networks for spectral snapshot compressive imaging. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2622–2631.
- Meng, Z.; Yu, Z.; Xu, K.; and Yuan, X. 2021b. Self-supervised neural networks for spectral snapshot compressive imaging. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2622–2631.
- Miao, X.; Yuan, X.; Pu, Y.; and Athitsos, V. 2019. l-net: Reconstruct hyperspectral images from a snapshot measurement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4059–4069.
- Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171. PMLR.
- Pan, Z.; Zeng, H.; Cao, J.; Zhang, K.; and Chen, Y. 2024. DiffSCI: Zero-Shot Snapshot Compressive Imaging via Iterative Spectral Diffusion Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25297–25306.

- Wagadarikar, A. A.; Pitsianis, N. P.; Sun, X.; and Brady, D. J. 2009. Video rate spectral imaging using a coded aperture snapshot spectral imager. *Optics Express*, 17(8): 6368–6388.
- Wang, L.; Sun, C.; Zhang, M.; Fu, Y.; and Huang, H. 2020. Dnu: Deep non-local unrolling for computational spectral imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1661–1671.
- Wang, Y.; Yang, W.; Chen, X.; Wang, Y.; Guo, L.; Chau, L.-P.; Liu, Z.; Qiao, Y.; Kot, A. C.; and Wen, B. 2024. SinSR: diffusion-based image super-resolution in a single step. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25796–25805.
- Wu, Z.; Lu, R.; Fu, Y.; and Yuan, X. 2025. Latent Diffusion Prior Enhanced Deep Unfolding for Snapshot Spectral Compressive Imaging. In *European Conference on Computer Vision*, 164–181. Springer.
- Xiong, Z.; Shi, Z.; Li, H.; Wang, L.; Liu, D.; and Wu, F. 2017. Hscnn: Cnn-based hyperspectral image recovery from spectrally undersampled projections. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 518–525.
- Xu, Y.; Wu, Z.; Li, J.; Plaza, A.; and Wei, Z. 2015. Anomaly detection in hyperspectral images based on low-rank and sparse representation. *IEEE Transactions on Geoscience and Remote Sensing*, 54(4): 1990–2000.
- Yue, Z.; Wang, J.; and Loy, C. C. 2025. Efficient diffusion model for image restoration by residual shifting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(1): 116–130.
- Zeng, H.; Cao, J.; Zhang, K.; Chen, Y.; Luong, H.; and Philips, W. 2024. Unmixing diffusion for self-supervised hyperspectral image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27820–27830.
- Zeng, H.; Sun, B.; Chen, Y.; Su, J.; and Xu, Y. 2025. Spectral Compressive Imaging via Unmixing-driven Subspace Diffusion Refinement. In *The Thirteenth International Conference on Learning Representations*.
- Zeng, H.; Xie, X.; Cui, H.; Yin, H.; and Ning, J. 2020. Hyperspectral image restoration via global L 1-2 spatial-spectral total variation regularized local low-rank tensor recovery. *IEEE transactions on geoscience and remote sensing*, 59(4): 3309–3325.
- Zhang, J.; Zeng, H.; Cao, J.; Chen, Y.; Yu, D.; and Zhao, Y.-P. 2024a. Dual Prior Unfolding for Snapshot Compressive Imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25742–25752.
- Zhang, J.; Zeng, H.; Chen, Y.; Yu, D.; and Zhao, Y.-P. 2024b. Improving Spectral Snapshot Reconstruction with Spectral-Spatial Rectification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25817–25826.
- Zhang, S.; Wang, L.; Fu, Y.; Zhong, X.; and Huang, H. 2019. Computational hyperspectral imaging based on dimension-discriminative low-rank tensor recovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10183–10192.
- Zhang, S.; Wang, L.; Zhang, L.; and Huang, H. 2021. Learning tensor low-rank prior for hyperspectral image reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12006–12015.
- Zhang, X.; Zhang, Y.; Xiong, R.; Sun, Q.; and Zhang, J. 2022. Herosnet: Hyperspectral explicable reconstruction and optimal sampling deep network for snapshot compressive imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17532–17541.