

# Semantic-embedded Unsupervised Spectral Reconstruction from Single RGB Images in the Wild

Zhiyu Zhu<sup>1</sup>, Hui Liu<sup>1</sup>, Junhui Hou<sup>1\*</sup>, Huanqiang Zeng<sup>2</sup>, Qingfu Zhang<sup>1</sup>

<sup>1</sup> City University of Hong Kong, <sup>2</sup> Huaqiao University

zhiyuzhu2-c@my.cityu.edu.hk, hliu99-c@my.cityu.edu.hk, jh.hou@cityu.edu.hk,  
zeng0043@hqu.edu.cn, qingfu.zhang@cityu.edu.hk.

## Abstract

*This paper investigates the problem of reconstructing hyperspectral (HS) images from single RGB images captured by commercial cameras, **without** using paired HS and RGB images during training. To tackle this challenge, we propose a new lightweight and end-to-end learning-based framework. Specifically, on the basis of the intrinsic imaging degradation model of RGB images from HS images, we progressively spread the differences between input RGB images and re-projected RGB images from recovered HS images via effective unsupervised camera spectral response function estimation. To enable the learning without paired ground-truth HS images as supervision, we adopt the adversarial learning manner and boost it with a simple yet effective  $\mathcal{L}_1$  gradient clipping scheme. Besides, we embed the semantic information of input RGB images to locally regularize the unsupervised learning, which is expected to promote pixels with identical semantics to have consistent spectral signatures. In addition to conducting quantitative experiments over two widely-used datasets for HS image reconstruction from synthetic RGB images, we also evaluate our method by applying recovered HS images from real RGB images to HS-based visual tracking. Extensive results show that our method significantly outperforms state-of-the-art unsupervised methods and even exceeds the latest supervised method under some settings. The source code is public available at <https://github.com/zbzhzhy/Unsupervised-Spectral-Reconstruction>.*

## 1. Introduction

As hyperspectral (HS) images delineate more reliable and accurate spectrum information of scenes than traditional RGB images, they facilitate many vision-based applications, such as visual tracking [44, 39], detection [26, 33],

and segmentation [34, 6]. However, it is costly to acquire HS images [50], which severely limits its wide deployment.

Without relying on expensive and specially designed hardware, recovering HS images from single RGB images via computational methods promises a convenient and low-cost manner, and recent deep neural network (DNN)-based methods have demonstrated the impressive abilities in solving such a highly ill-posed recovery problem. However, the majority of them have to be trained with paired RGB and HS images [25, 49, 45]. Unfortunately, it is non-trivial to collect a large number of such paired data via specially designed devices, e.g., well calibrated dual cameras. Although one can inversely synthesize RGB images from available HS images to form paired data for training, the huge gaps between the synthetic and real images may degrade the performance of a model trained with synthetic data when applied to real data.

To tackle the above issues, as illustrated in Fig. 1, we propose a lightweight, unsupervised, and end-to-end learning-based framework, which is capable of recovering HS images from single RGB images captured by commercial cameras *without* using paired RGB and HS images during training. Specifically, based on the low-dimensional property of camera spectral response functions (SRFs), we first propose a prior-driven method to estimate the underlying camera SRFs of input RGB images by extracting their deep features. Then, we propose an imaging degradation model-aware HS image generation approach, in which an HS image is generated in a coarse-to-fine manner by progressively spreading the information of the difference between input RGB images and re-projected RGB images obtained by integrating recovered HS images via the estimated camera SRFs. Such a generation manner well adapts to this unsupervised scenario. To enable the meaningful learning of our method under the lack of paired ground-truth HS images, we make types of efforts: (1) **globally**, we employ the popular and powerful adversarial learning to enforce the distribution of generated HS images to be close to that of real HS images, Moreover, we propose  $\mathcal{L}_1$  gradient clip-

\*Corresponding author. This work was supported by the Hong Kong Research Grants Council under grant CityU 11219019.

ping to stabilize and boost this learning manner; (2) **locally**, we embed the semantic information of scenes to encourage the pixels of reconstructed HS images with identical (resp. different) semantics to be similar (resp. dissimilar).

The main contributions of this paper are summarized as follows:

1. we propose an *unsupervised* HS image reconstruction framework from single RGB images in the wild;
2. we propose imaging degradation model-aware HS image generation and prior-driven unsupervised estimation of camera spectral response functions;
3. we propose to embed scenes’ semantic information efficiently to regularize the unsupervised learning;
4. we propose a simple yet effective  $\mathcal{L}_1$  gradient clipping strategy to stabilize and boost adversarial learning; and
5. we introduce visual tracking-based quality evaluation of HS images recovered from real RGB images.

## 2. Related Work

**Spectral Recovery from Single RGB Images.** Based on the assumption that HS images lie in a low-dimensional subspace, many traditional methods explore the map between RGB images and subspace coordinates. For example, Nguyen *et al.* [31] leveraged an RGB white-balancing algorithm to normalize the scene illumination for the scene reflectance recovery. Heikkinen *et al.* [17] utilized a scalar-valued Gaussian process regression with an-isotropic or combination kernels to estimate the spectral subspace coordinates. Arad *et al.* [?] proposed a sparse coding-based method, which learns an over-complete dictionary of HS images to describe the novel RGB images. Then Aeschbacher *et al.* [1] further improved it through introducing a shallow A+-based method [38].

Owing to the impressive representation ability of DNNs, DNN-based methods have been proposed to solve the challenging problem of HS image reconstruction from single RGB images. For example, Xiong *et al.* [45] proposed HSCNN-D for the reconstruction of HS images from RGB images and compressive HS imaging. Shi *et al.* [36] further improved HSCNN-D by introducing residual blocks and dense connections with a cross-scale fusion scheme to facilitate the feature extraction process. Fu *et al.* [13] explored non-negative structured information and utilized multiple sparse dictionaries to learn a compact basis representation for spectral reconstruction. To solve the uncertainty of camera SRFs, Berk *et al.* [21] trained a model to select different models, which are trained separately based on different SRFs. Li *et al.* [25] utilized both channel attention and spatial non-local attention for spectral reconstruction. Based on the assumption that pixels in an HS image belonging to different categories or spatial positions often require distinct mapping functions, Zhang *et al.* [49] proposed a pixel-aware deep function-mixture network, which learns recon-

struction functions with different receptive fields and linearly mixes them up according to pixel-level weights. Aitor *et al.* [2] applied a generative adversarial network to capture spatial semantics and map RGB to HS images. Yan *et al.* [46] introduced prior category information to generate distinct spectral data of objects via a U-Net-based architecture. Fu *et al.* [12] developed an SRF selection block to retrieve the optimal response function for HS image reconstruction. Galliani *et al.* [15] utilized a densely connected U-Net-based architecture for HS images reconstruction. Most of the aforementioned methods have to be trained with ground-truth HS images of input RGB images. However, such a heavy demand for paired HS and RGB images is hard to satisfy in practice. Although only a few works [49, 46] have been aware of the necessity of object context for HS image recovery, they could not explicitly utilize semantic information.

**Generative Adversarial Networks (GANs).** GANs have been widely adopted for image synthesis [32, 35] and reconstruction [41, 23]. With the emergency of GANs, the works on improving GAN’s stability and reducing the mode collapse have never stopped [5, 18, 28, 20, 37]. Among of them, some works solve the problems through replacing loss functions, e.g., Wasserstein GANs (W-GANs) [5], least squares GANs [27]. Others attempt to stabilize GANs’ training via regularizing discriminator. For example, Gulrajani *et al.* [16] proposed weight-clipping and gradient penalty to stabilize the training of W-GANs. Miyato *et al.* [28] proposed spectral normalization. However, these methods inevitably lose the discriminant ability, while enjoying the performance gains.

## 3. Proposed Method

### 3.1. Problem Statement and Overview

The physical degradation model of RGB images from HS images could be generally formulated as

$$\mathbf{X} = \mathbf{C}\mathbf{Y} + \mathbf{N}_y, \quad (1)$$

where  $\mathbf{X} \in \mathbb{R}^{3 \times hw}$  denotes the vectorial representation of an RGB image of spatial dimensions  $h \times w$ ,  $\mathbf{Y} \in \mathbb{R}^{s \times hw}$  is the corresponding HS image with  $s$  ( $s \gg 3$ ) spectral bands to be reconstructed,  $\mathbf{C} \in \mathbb{R}^{3 \times s}$  is the camera spectral response function (SRF), and  $\mathbf{N}_y \in \mathbb{R}^{3 \times hw}$  is the noise.

Apparently, one can design and train a DNN with paired RGB and HS images to learn the non-linear mapping from  $\mathbf{X}$  to  $\mathbf{Y}$ . However, it is costly, laborious and inconvenient to collect massive such image pairs for training in practice. Besides, due to different spatial resolution, field of views, and focal lengths, it is non-trivial to well register RGB and HS images captured by dual cameras. Alternatively, one may consider synthesizing RGB images from real HS images to form paired training data. However, the inevitable

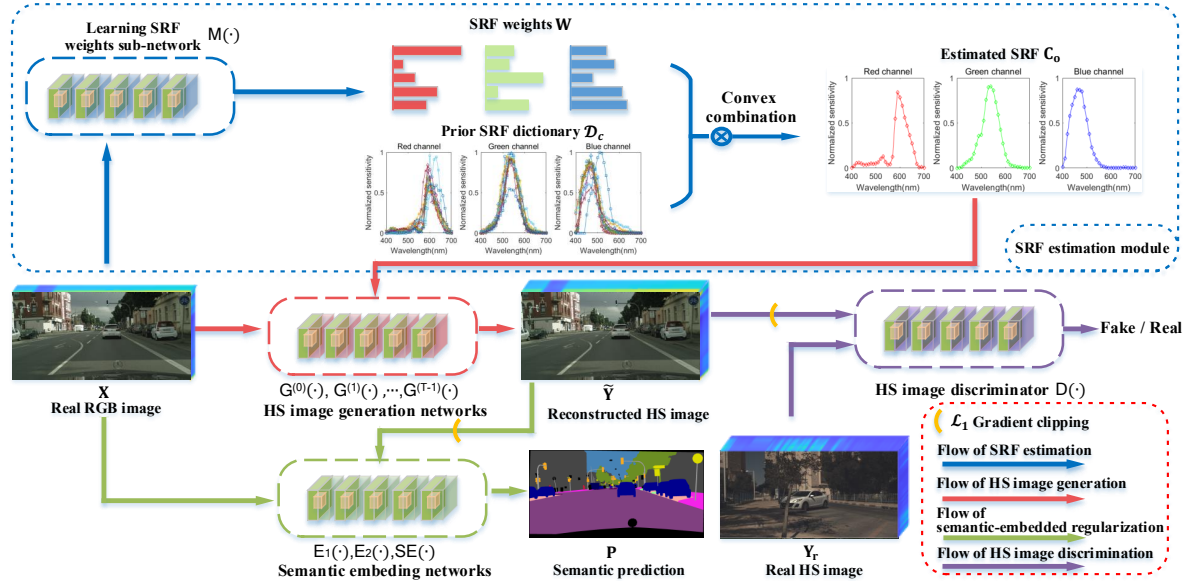


Figure 1. The flowchart of our **unsupervised** and **lightweight** learning-based framework. **Unpaired** semantically-labeled RGB images in the wild and real HS images are used for training. **Note** that the two modules, i.e., the semantic-embedded regularization and the  $\mathcal{L}_1$  gradient clipping boosted discriminator, are only used during training and not needed during testing. We refer readers to the *Supplementary Material* for the detailed network architecture of each module.

gap between real and synthetic RGB images may result in the DNN trained with synthetic data cannot be well generalized to real data.

To this end, as shown in Fig. 1, we propose an *unsupervised*, *lightweight*, and *end-to-end* learning-based HS image reconstruction framework from single RGB images captured by commercial cameras, where the corresponding ground-truth HS images of input RGB images are not required during training. Our framework is mainly composed of four modules:

(1) **Degradation model-aware HS image generation:** Let  $\hat{Y}$  be the recovered HS image from  $X$ , and  $\hat{X}$  be the re-projected RGB image by applying the camera SRF of  $X$  to  $\hat{Y}$ . According to Eq. (1), if  $\hat{Y}$  approximates the ground-truth one well, the difference between  $X$  and  $\hat{X}$  should be very small. Based on this intrinsic degradation relationship, we propose an HS image generation network, where a coarse HS image is first generated and then progressively refined by spreading useful information embedded in the difference between  $X$  and  $\hat{X}$ . The design of our generation network elegantly adapts to such an unsupervised scenario. See Section 3.2.

(2) **Prior-driven camera SRF estimation:** To realize the re-projection of the generation network, the camera SRFs of input RGB images have to be known. However, for RGB images in the wild, their SRFs are camera model-dependent and usually unavailable. Based on the fact that the SRFs of commercial cameras actually lie in a low-dimensional convex space (e.g., the first two principle components contain over 97 % of total variance of SRFs of 28 different camera models) [19], we propose a prior-driven

SRF estimation network, in which based on extracted deep features from input RGB images, SRFs are estimated as the convex combination of the atoms of a pre-defined SRF dictionary. See Section 3.3.

(3)  **$\mathcal{L}_1$  gradient clipping boosted adversarial learning:** To enable the learning of the generation network under the lack of paired ground-truth HS images, we adopt the adversarial learning manner, i.e., a discriminator is used to promote the distribution of generated HS images to approach that of real HS images. Besides, due to the inherent possibility of gradient explosion [37], it is non-trivial to train GANs [37, 48]. Inspired by traditional gradient clipping [7], we propose a simple yet effective strategy, namely  $\mathcal{L}_1$  gradient clipping, to stabilize the training process and likewise boost performance, in which the magnitude of the gradient matrix from the discriminator is adaptively clipped with reference to the gradient of  $\mathcal{L}_1$  loss function. See Section 3.4

(4) **Semantic-embedded regularization:** Owing to the rich spectrum information, the pixels of an HS image are discriminative [8], i.e., pixels corresponding to different (resp. identical) objects/materials always have distinct (resp. similar) spectral signatures [30, 10]. By leveraging this unique property, we propose a semantic-embedded regularization network to further regularize the solution space of the unsupervised generation network for meaningful reconstruction, i.e., during training, the generated HS images by the generation network are fed into a scene parsing/semantic segmentation network trained with the ground-truth semantic maps of the input RGB images to predict semantic maps. In contrast to the discriminator that *globally* regularizes the distribution of HS images, this




Figure 2. Illustration of the network architecture of the proposed degradation model-aware HS image generation.

module is able to *locally* regularize the pixels of generated HS images. See Section 3.5. Note that RGB image-based semantic segmentation is a popular image analysis task and there are many publicly available datasets, and thus our method will not introduce additional costs. Moreover, the emerging weakly-supervised scene parsing methods [42, 9] will lessen the requirements of the ground-truth semantic annotations.

Note that our framework is end-to-end trained with the objective functions in Section 3.6, and the two modules adversarial learning and semantic-embedded regularization are not needed during testing. In what follows, we will show the technical details of each module.

### 3.2. HS Image Generation

As shown in Fig. 2, this module generates an HS image from an input RGB image in a coarse-to-fine fashion via the following three steps:

**Coarse generation.** We first learn a coarse HS image denoted as  $\tilde{\mathbf{Y}}^{(1)}$  using an efficient sub-CNN  $G^{(0)}(\cdot)$ :

$$\tilde{\mathbf{Y}}^{(1)} = G^{(0)}(\mathbf{X}). \quad (2)$$

**Progressive refinement.** We then progressively refine  $\tilde{\mathbf{Y}}^{(1)}$  via a multi-stage structure, and at the  $t$ -th ( $1 \leq t \leq T-1$ ) stage, the refinement is carried out as

$$\tilde{\mathbf{Y}}^{(t+1)} = \tilde{\mathbf{Y}}^{(t)} + G^{(t)}(\mathbf{C}_o \tilde{\mathbf{Y}}^{(t)} - \mathbf{X}), \quad (3)$$

where  $G^{(t)}(\cdot)$  stands for the sub-CNN at the  $t$ -th stage, and  $\mathbf{C}_o \in \mathbb{R}^{3 \times s}$  is the predicted camera SRF explained in Section 3.3. Note that  $G^{(1)}(\cdot), \dots, G^{(T-1)}(\cdot)$  have the same network architecture equipped with different parameters, and the network architecture of  $G^{(0)}(\cdot)$  is different from that of  $G^{(t)}(\cdot)$  because  $G^{(t)}(\cdot)$  perceives error maps, while  $G^{(0)}(\cdot)$  takes RGB images as input.

**Non-negative thresholding.** Being aware that the pixel intensity of a physical meaningful HS image should be non-negative, we finally employ the LeakyReLU, which does not obstruct gradient propagation, to guarantee the non-negativity of generated HS images:

$$\tilde{\mathbf{Y}} = \text{LeakyReLU}(\tilde{\mathbf{Y}}^{(T)}). \quad (4)$$

where  $\tilde{\mathbf{Y}} \in \mathbb{R}^{s \times hw}$  is the generated HS image.

### 3.3. Camera SRF Estimation

Denote by  $\mathcal{D}_c = \{\mathbf{C}_1, \mathbf{C}_2, \dots, \mathbf{C}_N\}$  a pre-defined dictionary consisting of  $N = 28$  commonly-used camera SRFs

[19]. We first estimate an intermediate SRF  $\hat{\mathbf{C}} \in \mathbb{R}^{3 \times s}$  as

$$\hat{\mathbf{C}}(i) = \sum_{j=1}^N \hat{\mathbf{w}}_i(j) \times \mathbf{C}_j(i), \quad (i = 1, 2, 3),$$

$$\text{with } \hat{\mathbf{w}}_i = \mathcal{S}_{max}(\mathbf{W}(i)), \quad \mathbf{W} = M(\mathbf{X}), \quad (5)$$

where  $M(\cdot)$  is a sub-CNN extracting deep features from  $\mathbf{X}$  and outputs the weight matrix  $\mathbf{W} \in \mathbb{R}^{3 \times N}$ ;  $\hat{\mathbf{C}}(i)$ ,  $\mathbf{C}_j(i)$ , and  $\mathbf{W}(i)$  are the  $i$ -th row of  $\hat{\mathbf{C}}$ ,  $\mathbf{C}_j$ , and  $\mathbf{W}$ , respectively;  $\hat{\mathbf{w}}_i(j)$  is the  $j$ -th entry of  $\hat{\mathbf{w}}_i \in \mathbb{R}^{1 \times N}$ ;  $\mathcal{S}_{max}(\cdot)$  denotes the softmax operator.

Besides, real RGB images by commercial cameras are usually adjusted to be visually-pleasing via white balance. That is, the red, green, and blue three channels are separately re-scaled to have the same average value via the widely-used Gray World white balance algorithm [24]. Considering its pixel position-independent property, we mimic this operation by re-scaling the SRF:

$$\mathbf{C}_o(i) = \frac{\sum_{j=1}^N \tilde{\mathbf{W}}(i, j)}{\sum_{j=1}^N \sum_{k=1}^3 \tilde{\mathbf{W}}(k, j)/3} \times \hat{\mathbf{C}}(i), \quad (i = 1, 2, 3),$$

$$\text{with } \tilde{\mathbf{W}} = \text{Sigmoid}(\mathbf{W}), \quad (6)$$

where  $\mathbf{C}_o(i)$  is the  $i$ -th row of the finally estimated camera SRF  $\mathbf{C}_o \in \mathbb{R}^{3 \times s}$ ,  $\text{Sigmoid}(\cdot)$  is the sigmoid function, and  $\tilde{\mathbf{W}}(k, j)$  denotes the  $(k, j)$ -th entry of matrix  $\tilde{\mathbf{W}} \in \mathbb{R}^{3 \times N}$ . The use of the sigmoid function is able to avoid the cancellation when computing the the summation of positive and negative values, and also maintains the original order.

### 3.4. Adversarial Learning with Gradient Clipping

It is commonly-known that training a good GAN is challenging, and the poor convergence may severely compromise the final performance. Instead of using existing regularization methods [28, 5, 16], which may compromise the discriminant ability of the discriminator, we propose  $\mathcal{L}_1$  gradient clipping to stabilize the training process and likewise boost performance. Specifically, the existing gradient clipping [7], which was originally designed for solving the gradient issue of training RNNs, re-scales the intermediate gradient to a certain level, i.e., ,

$$\mathcal{C}(\mathbf{G}, \tau) = \text{Min} \left( 1, \frac{\tau}{\|\mathbf{G}\|_F} \right) \cdot \mathbf{G}, \quad (7)$$

where  $\mathcal{C}(\cdot, \cdot)$  is the gradient clipping function for training,  $\tau$  is the gradient threshold,  $\mathbf{G} \in \mathbb{R}^{c \times hw}$  is the gradient matrix

to be regularized, and  $\|\cdot\|_F$  is the Frobenius norm of a matrix. Although gradient clipping can regularize the gradient magnitude without changing discriminator parameters and weakening its discriminant ability, the threshold  $\tau$  has to be carefully set. Based on the observation that the classic  $\mathcal{L}_1$  loss function can produce a gradient with a fixed  $\mathcal{F}$ -norm, which only varying with the input size, we modify Eq. (7) and propose  $\mathcal{L}_1$  gradient clipping, which re-scales the gradient matrix as  $\mathcal{C}(\mathbf{G}, \frac{1}{\sqrt{c \times hw}})$ . We apply this  $\mathcal{L}_1$  gradient clipping to the gradient matrix back-propagated from the discriminator to the HS image generation network. We experimentally validated the **impressive** performance of such a simple strategy in Table 5.

### 3.5. Semantic-embedded HS image Regularization

We utilize existing scene parsing datasets as the source of input RGB images during training, which contain semantically-labeled RGB images in pixel-wise. As it is usually complex and time-consuming to train a scene parsing network, we propose the following simple manner: (1) first utilizing a pre-trained scene parsing network [40] as the backbone, denoted as  $E_1(\cdot)$ , to extract features denoted as  $\Phi_X$  from  $\mathbf{X}$  and a plain sub-CNN, denoted as  $E_2(\cdot)$ , to extract features denoted as  $\Phi_Y$  from  $\tilde{\mathbf{Y}}$ ; and (2) then fusing  $\Phi_X$  and  $\Phi_Y$  via another sub-CNN, denoted as  $SE(\cdot, \cdot)$ , to predict the semantic map  $\mathbf{P} \in \mathbb{R}^{M \times hw}$  with  $M$  being the total number of semantic classes. It is expected that via back-propagation, the generation network will be promoted to produce  $\tilde{\mathbf{Y}}$  which could be distinguished by the semantic-embedded network trained with ground-truth semantic maps of  $\mathbf{X}$ . We also apply  $\mathcal{L}_1$  gradient clipping to stabilize the fluctuated gradient from the semantic-embedded network.

### 3.6. Objective Functions

Analogy to most GAN-based frameworks, we alternately train the generator (including HS image generation, semantic-embedded regularization, and camera SRF estimation) and the discriminator. The objective function of the generator is

$$\mathcal{L}_G = \mathcal{B}(D(\tilde{\mathbf{Y}}/\bar{y}), 1) + \lambda_b \mathcal{S}_{bp}(\mathbf{C}_o, \tilde{\mathbf{Y}}, \mathbf{X}) + \lambda_s \mathcal{S}_1(\tilde{\mathbf{Y}}) + \lambda_p \mathcal{S}_{pos}(\tilde{\mathbf{Y}}) + \lambda_{sem} \mathcal{S}_{sem}(\mathbf{P}, \mathbf{L}), \quad (8)$$

where  $\lambda_b$ ,  $\lambda_s$ ,  $\lambda_p$ , and  $\lambda_{sem}$  are non-negative parameters to balance different terms, which are empirically set to  $1e^2$ ,  $1e^1$ ,  $1e^{-2}$ , and  $1e^0$ , respectively;  $\bar{y}$  is the mean value of  $\tilde{\mathbf{Y}}$ ;  $\mathcal{B}(\cdot)$  is the binary cross-entropy loss;  $D(\cdot)$  stands for the discriminator; the  $\mathcal{S}_{bp}(\cdot, \cdot, \cdot)$  is the back-projection constraint,  $\mathcal{S}_1(\cdot)$  is the first-order smoothness constraint,  $\mathcal{S}_{pos}(\cdot)$  is the positive-definite constraint, and  $\mathcal{S}_{sem}(\cdot, \cdot)$  is the loss for the semantics embedding regularization, where are respectively defined as

$$\mathcal{S}_{bp}(\mathbf{C}_o, \tilde{\mathbf{Y}}, \mathbf{X}) = \frac{1}{3 \times hw} \left\| \mathbf{C}_o \tilde{\mathbf{Y}} - \mathbf{X} \right\|_F^2, \quad (9)$$

$$\mathcal{S}_1(\tilde{\mathbf{Y}}) = \frac{1}{(s-1) \times hw} \sum_{i=1}^{s-1} \left\| \tilde{\mathbf{Y}}(i) - \tilde{\mathbf{Y}}(i+1) \right\|_1, \quad (10)$$

$$\mathcal{S}_{pos}(\tilde{\mathbf{Y}}) = \left\| \tilde{\mathbf{Y}} - \left| \tilde{\mathbf{Y}} \right| \right\|_F^2, \quad (11)$$

$$\mathcal{S}_{sem}(\mathbf{P}, \mathbf{L}) = \frac{1}{hw} \sum_{n=1}^{hw} -\mathbf{L}(n) \log(\mathbf{P}(n)), \quad (12)$$

where  $|\cdot|$  computes the absolute value of the input in element-wise,  $\|\cdot\|_1$  is the  $\ell_1$ -norm of a vector,  $\mathbf{P}(n) \in \mathbb{R}^{M \times 1}$  is the semantic prediction of the  $n$ -th pixel, and  $\mathbf{L}(n) \in \mathbb{R}^{M \times 1}$  is the ground-truth of  $\mathbf{P}(n)$ . The objective function of the discriminator is

$$\mathcal{L}_D = \mathcal{B}(D(\tilde{\mathbf{Y}}/\bar{y}), 0) + \mathcal{B}(D(\mathbf{Y}_r/\bar{y}_r), 1), \quad (13)$$

where  $\mathbf{Y}_r \in \mathbb{R}^{s \times hw}$  is a real HS image,  $\bar{y}_r$  is the mean value of  $\mathbf{Y}_r$ . We normalize the  $\tilde{\mathbf{Y}}$  and  $\mathbf{Y}_r$  respectively with their mean values to focus the reconstruction process much on the spectral angular of HS images. Besides, the proposed  $\mathcal{L}_1$  gradient clipping operation in Section 3.4 will be used to regularize the backward gradients.

## 4. Experiments

**Real HS image dataset.** The ICVL [3] and NTIRE'20 [4] datasets contain HS images of 31 spectral bands covering from 400 to 700 nm with interval of 10nm captured under outdoor daylight illumination. There are 201 HS images of spatial dimensions  $1392 \times 1300$  and 460 HS images of spatial dimensions  $482 \times 512$  in the ICVL and NTIRE'20 datasets, respectively.

**Semantically-labeled real RGB image datasets.** The PASCAL-Context [29] and Cityscapes [11] datasets contain RGB images with pixel-wise categorical labels for semantic embedding. PASCAL-Context contains 59 semantic classes, which is divided into 4,998/5,105 images for training and testing, respectively. The Cityscapes dataset is officially divided into 2,975/500/1525 images for training, validation, and testing, respectively.

**RGB and HS video dataset for visual tracking.** The hyperspectral object tracking challenge (HOTAC) dataset [44] contains 50 RGB videos with each about 425 frames, where the bounding-boxes of target objects are given. Besides, HOTAC provides 50 HS videos with target objects labeled with bounding-boxes, and the HS videos capture approximately the same scenes as the RGB videos<sup>1</sup>.

**Compared methods.** We compared with 5 methods, including 1 baseline, i.e., the bicubic interpolation (BI)

<sup>1</sup>Although the RGB and HS videos have been registered, there is still a large disparity between them.




Figure 3. Visualization of the error maps of different methods on 6 testing RGB images (the most left column) in the ICVL (the top three rows in 440, 590, and 640 nm) and NTIRE’20 (the bottom three rows in 490, 640, and 690 nm) datasets. Additionally, we also draw the spectral signatures of 6 pixels respectively selected from the 6 testing images the selected pixels are marked by green rectangle in the RGB images.

over the spectral dimension, and 4 most recent DNN-based methods, i.e., CHS-GAN [2], DBF [14], HSCNN-D [36], and FM-Net [49]. Note that CHS-GAN, HSCNN-D, and FM-Net are supervised and have to be trained with paired HS and RGB images, and DBF is unsupervised.

#### 4.1. Evaluation on Synthetic RGB Images

**Experiment settings.** Although there is a gap between synthetic and real data, we still conducted experiments on synthetic data as done in existing works [49], [36] to have a quantitative and intuitive understanding of our framework.

We selected the last 20 HS images of ICVL to form the test dataset, and the remaining ones as the training dataset. All the HS images of NTIRE’20 were used as the test dataset. For fair comparisons, we applied the same dictionary of camera SRFs in Section 3.3 to the training HS images to synthesize RGB images, generating paired training data, which were used to train HSCNN-D [36], FM-Net [49], and CHS-GAN [2] with the paired synthetic RGB and HS images. Following the original implementation [14], we first trained DBF to learn a basis function and then trained it on the synthetic RGB images with the learnt basis function. The proposed method was trained with PASCAL-Context and the aforementioned training HS images from ICVL as RGB and HS image sources, respectively, with the batch size of 24. The number of stages  $T$  was set to 4. We adopted the ADAM [22] optimizer for optimization of the generation network, the discriminator, and the cam-

era SRF estimation network. The exponential decay rates were set as  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$  for the first and second moment estimates, respectively. We also adopted the TTUR scheme [18] to set the learning rate of the SRF estimation network, HS image generation network, and the HS image discriminator to  $1e^{-4}$ ,  $1e^{-4}$ , and  $5e^{-4}$ , respectively. Meanwhile, for the semantic embedding module, we fixed  $E_1(\cdot)$  and trained  $E_2(\cdot)$  and  $SE(\cdot)$  using the SGD optimizer with the base learning rate, the momentum, and the weight decay equal to  $1e^{-3}$ , 0.9, and  $1e^{-4}$ , respectively. The poly learning rate policy with the power of 0.9 is used for dropping the learning rate [40]. The proposed method is trained 50 epochs with 208 iterations each epoch. We tested all the methods with the same protocols on the synthetic RGB images generated from the widely-used camera SRF— Nikon D700. To quantitatively compare different methods, we adopted 3 commonly-used metrics, i.e., Peak Signal-to-Noise Ratio (PSNR), Average Structural Similarity Index (ASSIM) [43], and Spectral Angle Mapper (SAM) [47]. We reported all the methods with the training results of the last epoch.

**Quantitative and visual results.** Table 1 lists the results of different methods on ICVL, where it can be seen that our method significantly outperforms the unsupervised DBF and achieves comparable performance to the supervised HSCNN-D and CHS-GAN, which are credited to the carefully designed framework. But the gap between our method and FM-Net — the most recent supervised method

Table 1. Quantitative comparisons of different methods on the ICVL dataset. “↑ (resp. ↓)” indicates that the larger (resp. smaller), the better. For # Params and # FLOPs, the smaller, the more compact and efficient.

Methods	Paired data	Real RGB	# Params	# FLOPs	Running time (ms)	PSNR ↑	ASSIM ↑	SAM ↓
BI	×	✓	—	—	—	25.96	0.911	19.18
HSCNN-D [36]	✓	×	3.61 M	5.22 T	5184.75ms	38.15	0.992	3.28
FM-Net [49]	✓	×	11.79 M	17.07 T	4780.41ms	40.26	0.993	2.65
CHS-GAN [2]	✓	×	1.09M	0.91T	505.70ms	38.03	0.988	3.34
DBF [14]	×	×	1.13M	1.02T	531.25ms	30.75	0.940	14.64
Ours	×	✓	0.58M	0.93T	561.90 ms	37.83	0.991	4.01

Table 2. Quantitative comparisons of different methods on the NTIRE’20 dataset.

Methods	Paired data	Real RGB	# Params	# FLOPs	Running time	PSNR ↑	ASSIM ↑	SAM ↓
BI	×	✓	—	—	—	25.91	0.876	23.00
HSCNN-D [36]	✓	×	3.61 M	0.890 T	1003.38 ms	31.45	0.919	10.76
FM-Net [49]	✓	×	11.79 M	2.955 T	1066.40 ms	34.25	0.968	9.11
CHS-GAN [2]	✓	×	1.09 M	0.15T	87.74 ms	30.63	0.906	15.38
DBF [14]	×	×	1.13 M	0.17T	85.91 ms	31.51	0.891	16.81
Ours	×	✓	0.58 M	0.16T	134.81 ms	37.53	0.979	5.97

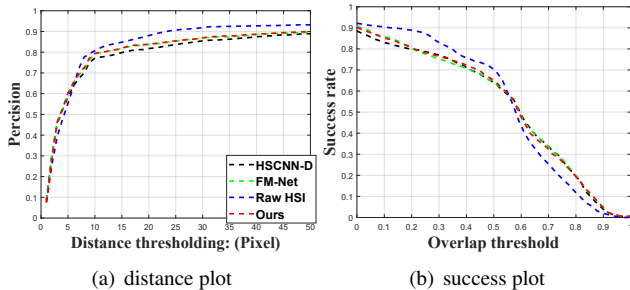


Figure 4. Comparison of different methods evaluated with the visual tracking task.

also indicates the potential of this topic. However, as listed in Table 2, our method even exceeds FM-Net on NTIRE’20, which somewhat demonstrates better generalization ability. As visualized in Fig. 3, we can see that the error between the reconstructed spectral band and the corresponding ground-truth one is much smaller than those of compared methods. Such an advantage may be benefit from the semantic information of the PASCAL-Context dataset containing the class of sky. Besides, the spectral signatures of selected pixels are closer to ground-truth ones. Finally, the proposed method consumes much less time than HSCNN-D and FM-Net.

## 4.2. Visual Tracking-based Evaluation on Real RGB Images

For real RGB images, the well registered ground-truth HS images are usually unavailable, making it impossible to quantitatively evaluate different methods, as done in Section 4.1. Besides, due to the high spectral resolution, it is also difficult to visually compare the results of different methods. Considering that the rich spectrum information makes pixels of HS images discriminative, which will be benefit high-level vision tasks, we introduced the HS video-based visual tracking experiment to evaluate different HS image

reconstruction methods. It is expected a better method will generate HS images that are closer to unknown ground-truth ones, which then produce high tracking accuracy.

**Experiment settings.** We compared with HSCNN-D and FM-Net, which always achieve the top two performance among all the compared methods in the previous scenario. For all methods, we directly adopted the trained models in Section 4.1 **without** additional training to reconstruct the 10 randomly selected RGB videos from HOTAC into HS videos, which were then fed to a pre-trained HS video-based tracker named MHT [44] to realize tracking. Besides, we also performed the MHT tracker on the 10 HS videos provided by HOTAC, corresponding to the 10 selected RGB videos. The tracking performance of such a setting, denoted as “Raw”, could be thought of as the upper-bound of the reconstruction methods. We quantitatively measured the tracking performance by using Average Precision (AP) among all the distance threshold, AUC, and location error, three commonly-used metrics in visual tracking.

**Quantitative results.** Table 4 and Fig. 4 show the results, where it can be seen that our method even produces better tracking performance than the two supervised methods FM-Net and HSCNN-D, and as expected, the performance of all the three computation-based reconstruction methods is lower than that of “Raw”. Please refer to our *Video Demo* for more visual results.

## 4.3. Ablation Studies

**Regularization terms.** We experimentally validated the effectiveness of the smooth constraint  $\mathcal{S}_1(\cdot)$ , the positive constraint  $\mathcal{S}_{pos}(\cdot)$ , the back-projection constraint  $\mathcal{S}_{bp}(\cdot, \cdot, \cdot)$ ,  $\mathcal{L}_1$  gradient clipping, and the semantic-embedded module. As shown in Table 3, we can see that our framework benefits from each of these components. Among them, the back-projection regularization, semantic embedding module and  $\mathcal{L}_1$  gradient clipping have the pivotal influence. Besides, only using traditional regularization terms

Table 3. Results of ablation studies on the ICVL dataset.  $\times$  indicates that the corresponding component was removed when training testing the proposed method. The bottom row corresponds to the complete model.

Adversarial learning	$\mathcal{S}_{bp}$	$\mathcal{S}_{pos}$	$\mathcal{S}_1$	$\mathcal{L}_1$	Gradient clipping	$\mathcal{S}_{sem}$	PSNR $\uparrow$	ASSIM $\uparrow$	SAM $\downarrow$
$\checkmark$	$\times$	$\times$	$\times$		$\times$	$\times$	14.47	0.121	63.34
$\checkmark$	$\checkmark$	$\times$	$\times$		$\times$	$\times$	20.48	0.781	35.70
$\checkmark$	$\checkmark$	$\checkmark$	$\times$		$\times$	$\times$	21.72	0.783	30.64
$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$		$\times$	$\times$	26.11	0.882	14.44
$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$		$\checkmark$	$\times$	34.51	0.983	7.04
$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$		$\checkmark$	$\checkmark$	37.83	0.991	4.01

Table 4. Comparison of different methods on visual tracking.

Methods	AP (%) $\uparrow$	AUC $\uparrow$	Location error $\downarrow$
HSCNN-D [36]	78.65	0.5235	17.4358
FM-Net [49]	80.03	0.5268	16.3825
Ours	80.23	0.5272	16.4884
Raw	83.21	0.5360	14.7835

Table 5. Comparison of different regularization methods for boosting the training of GANs.

Methods	PSNR $\uparrow$	SSIM $\uparrow$	SAM $\downarrow$
1-GP [16]	28.15	0.944	12.60
0-GP [37]	33.33	0.984	6.29
$\mathcal{L}_1$ -GC (ours)	37.83	0.991	4.01

Table 6. Performance of our method trained with various RGB and HS image sources. The test data are 20 samples from ICVL.

RGB source	HS source	PSNR $\uparrow$	ASSIM $\uparrow$	SAM $\downarrow$
Cityscapes	NTIRE'20	34.08	0.9838	6.61
Cityscapes	ICVL	37.68	0.988	5.03
PASCAL-Context	NTIRE'20	35.87	0.9837	6.24
PASCAL-Context	ICVL	37.83	0.991	4.01

such as  $\mathcal{S}_{bp}(\cdot, \cdot, \cdot)$ ,  $\mathcal{S}_{pos}(\cdot)$ , and  $\mathcal{S}_1(\cdot)$ , the method produces very limited performance comparable to that of Bicubic interpolation, indicating that the unsupervised HS image reconstruction actually is a very challenging task.

**Camera SRF estimation.** Even with supervision, camera SRF estimation of RGB images is a quite challenging task [21]. To evaluate the effectiveness of our camera SRF estimation module, we randomly chose ten SRFs to generate synthetic RGB images from real HS images, which were then fed into our estimation module to estimate the SRFs. We calculated the precision of estimated SRFs using the AUC metric. As shown in Fig. 5, the average precision of our estimation module achieves 77.79%. Note that our SRF estimation module is *unsupervised*, i.e., the ground-truth camera SRFs were not employed as supervision during training.

**$\mathcal{L}_1$  gradient clipping.** To show the advantage of the proposed  $\mathcal{L}_1$  gradient clipping, we trained our framework with different gradient regularization methods [37, 16], where we only changed the gradient regularization terms and re-




Figure 5. Reconstruction accuracy of our unsupervised prior-driven camera SRF estimation module.

mained the other modules unchanged. As listed in Table 5, the proposed  $\mathcal{L}_1$  gradient clipping exceeds the other gradient clipping methods to a large extent, demonstrating its superiority.

**Various training datasets.** To further explore the performance of our method on different datasets, we conducted experiments by training our method with different image datasets as training data. We tested all the trained models with the 20 samples of ICVL dataset. As listed in Table 6, we can see that for the RGB image source, the performance of using PASCAL-Context is better than that of using Cityscapes, which may be credited to the more diverse samples contained in the PASCAL-Context dataset.

## 5. Conclusion

We have presented a new unsupervised, lightweight, and end-to-end learning-based framework for the reconstruction of HS images from single RGB images in the wild. Our framework does not need paired RGB and HS images for training. Owing the well-motivated reconstruction modules including prior-driven camera SRF estimation and imaging degradation model-aware generation, as well as the effective semantic embedding and  $\mathcal{L}_1$  gradient clipping boosted adversarial learning for regularizing the learning of the framework without referencing ground-truth HS images, our framework achieves impressive performance, and especially our method even exceeds the most recent supervised reconstruction method under some settings.



## References

- [1] Jonas Aeschbacher, Jiqing Wu, and Radu Timofte. In defense of shallow learned spectral reconstruction from rgb images. In *Proc. of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 471–479, 2017.
- [2] Aitor Alvarez-Gila, Joost Van De Weijer, and Estibaliz Garrote. Adversarial networks for spatial context-aware spectral image reconstruction from rgb. In *Proc. of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 480–490, 2017.
- [3] Boaz Arad and Ohad Ben-Shahar. Sparse recovery of hyperspectral signal from natural rgb images. In *Proc. of the European Conference on Computer Vision*, pages 19–34. Springer, 2016.
- [4] Boaz Arad, Radu Timofte, Ohad Ben-Shahar, Yi-Tun Lin, and Graham D Finlayson. Ntire 2020 challenge on spectral reconstruction from an rgb image. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 446–447, 2020.
- [5] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223. PMLR, 2017.
- [6] Hakan Aytaylan and Seniha Esen Yuksel. Semantic segmentation of hyperspectral images with the fusion of lidar data. In *IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 2522–2525, 2016.
- [7] Yoshua Bengio, Ian Goodfellow, and Aaron Courville. *Deep learning*, volume 1. MIT press Massachusetts, USA:, 2017.
- [8] Gustavo Camps-Valls and Lorenzo Bruzzone. Kernel-based methods for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 43(6):1351–1362, 2005.
- [9] Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, and Ming-Hsuan Yang. Weakly-supervised semantic segmentation via sub-category exploration. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8991–9000, 2020.
- [10] Yi Chen, Nasser M Nasrabadi, and Trac D Tran. Hyperspectral image classification via kernel sparse representation. *IEEE Transactions on Geoscience and Remote Sensing*, 51(1):217–231, 2012.
- [11] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016.
- [12] Ying Fu, Tao Zhang, Yinqiang Zheng, Debing Zhang, and Hua Huang. Joint camera spectral sensitivity selection and hyperspectral image recovery. In *Proc. of the European Conference on Computer Vision*, pages 788–804, 2018.
- [13] Ying Fu, Yongrong Zheng, Lin Zhang, and Hua Huang. Spectral reflectance recovery from a single rgb image. *IEEE Transactions on Computational Imaging*, 4(3):382–394, 2018.
- [14] Biebele Joslyn Fubara, Mohamed Sedky, and David Dyke. Rgb to spectral reconstruction via learned basis functions and weights. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 480–481, 2020.
- [15] Silvano Galliani, Charis Lanaras, Dimitrios Marmanis, Emmanuel Baltsavias, and Konrad Schindler. Learned spectral super-resolution. *arXiv preprint arXiv:1703.09470*, 2017.
- [16] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Proc. of the Advances in Neural Information Processing Systems*, volume 30, 2017.
- [17] Ville Heikkinen. Spectral reflectance estimation using gaussian processes and combination kernels. *IEEE Transactions on Image Processing*, 27(7):3358–3373, 2018.
- [18] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proc. of the International Conference on Neural Information Processing Systems*, pages 6629–6640, 2017.
- [19] Jun Jiang, Dengyu Liu, Jinwei Gu, and Sabine Süsstrunk. What is the space of spectral sensitivity functions for digital color cameras? In *Proc. of the IEEE Workshop on Applications of Computer Vision*, pages 168–179, 2013.
- [20] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *Proc. of the International Conference on Learning Representations*, 2018.
- [21] Berk Kaya, Yigit Baran Can, and Radu Timofte. Towards spectral estimation from a single rgb image in the wild. In *Proc. of the IEEE/CVF International Conference on Computer Vision Workshop*, pages 3546–3555, 2019.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proc. of the International Conference on Learning Representations*, 2015.
- [23] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiří Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8183–8192, 2018.
- [24] Edmund Y Lam. Image restoration in digital photography. *IEEE Transactions on Consumer Electronics*, 49(2):269–274, 2003.
- [25] Jiaojiao Li, Chaoxiong Wu, Rui Song, Yunsong Li, and Fei Liu. Adaptive weighted attention network with camera spectral sensitivity prior for spectral reconstruction from rgb images. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 462–463, 2020.
- [26] Konstantinos Makantasis, Konstantinos Karantzas, Anastasios Doulamis, and Konstantinos Loupos. Deep learning-based man-made object detection from hyperspectral data. In *International Symposium on Visual Computing*, pages 717–727. Springer, 2015.
- [27] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proc. of the IEEE/CVF Interna-*

- tional Conference on Computer Vision*, pages 2794–2802, 2017.
- [28] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *Proc. of the International Conference on Learning Representations*, 2018.
- [29] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 891–898, 2014.
- [30] L. Mou, P. Ghamisi, and X. X. Zhu. Deep recurrent neural networks for hyperspectral image classification. *IEEE Transactions on Geoscience and Remote Sensing*, 55(7):3639–3655, 2017.
- [31] Rang MH Nguyen, Dilip K Prasad, and Michael S Brown. Training-based spectral reconstruction from a single rgb image. In *Proc. of the European Conference on Computer Vision*, pages 186–201. Springer, 2014.
- [32] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *Proc. of the International Conference on Machine Learning*, pages 2642–2651. PMLR, 2017.
- [33] Thuy T Pham, MA Takalkar, M Xu, Dinh Thai Hoang, HA Truong, Eryk Dutkiewicz, and S Perry. Airborne object detection using hyperspectral imaging: Deep learning review. In *Proc. of the International Conference on Computational Science and Its Applications*, pages 306–321. Springer, 2019.
- [34] Daniele Ravi, Himar Fabelo, Gustavo Marrero Callic, and Guang-Zhong Yang. Manifold embedding and semantic segmentation for intraoperative guidance with hyperspectral brain imaging. *IEEE Transactions on Medical Imaging*, 36(9):1845–1857, 2017.
- [35] Krishna Regmi and Ali Borji. Cross-view image synthesis using conditional gans. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3501–3510, 2018.
- [36] Zhan Shi, Chang Chen, Zhiwei Xiong, Dong Liu, and Feng Wu. Hscnn+: Advanced cnn-based hyperspectral recovery from rgb images. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 939–947, 2018.
- [37] Hoang Thanh-Tung, Svetha Venkatesh, and Truyen Tran. Improving generalization and stability of generative adversarial networks. In *Proc. of the International Conference on Learning Representations*, 2019.
- [38] Radu Timofte, Vincent De Smet, and Luc Van Gool. A+: Adjusted anchored neighborhood regression for fast super-resolution. In *Proc. of the Asian Conference on Computer Vision*, pages 111–126. Springer, 2014.
- [39] Hien Van Nguyen, Amit Banerjee, and Rama Chellappa. Tracking via object reflectance using a hyperspectral video camera. In *Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition-Workshops*, pages 44–51. IEEE, 2010.
- [40] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [41] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proc. of the European Conference on Computer Vision Workshops*, pages 63–79, 2018.
- [42] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12275–12284, 2020.
- [43] Zhou Wang and Alan C Bovik. A universal image quality index. *IEEE Signal Processing Letters*, 9(3):81–84, 2002.
- [44] Fengchao Xiong, Jun Zhou, and Yuntao Qian. Material based object tracking in hyperspectral videos. *IEEE Transactions on Image Processing*, 29:3719–3733, 2020.
- [45] Zhiwei Xiong, Zhan Shi, Huiqun Li, Lizhi Wang, Dong Liu, and Feng Wu. Hscnn: Cnn-based hyperspectral image recovery from spectrally undersampled projections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 518–525, 2017.
- [46] Longbin Yan, Xiuheng Wang, Min Zhao, Maboud Kaloorazi, Jie Chen, and Susanto Rahardja. Reconstruction of hyperspectral data from rgb images with prior category information. *IEEE Transactions on Computational Imaging*, 6:1070–1081, 2020.
- [47] Roberta H Yuhas, Alexander FH Goetz, and Joe W Boardman. Discrimination among semi-arid landscape endmembers using the spectral angle mapper (sam) algorithm. In *Proc. Summaries 3rd Annu. JPL Airborne Geosci. Workshop*, volume 1, pages 147–149, 1992.
- [48] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *Proc. of the International Conference on Machine Learning*, pages 7354–7363. PMLR, 2019.
- [49] Lei Zhang, Zhiqiang Lang, Peng Wang, Wei Wei, Shengcai Liao, Ling Shao, and Yanning Zhang. Pixel-aware deep function-mixture network for spectral super-resolution. In *Proc. of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12821–12828, 2020.
- [50] Zhiyu Zhu, Junhui Hou, Jie Chen, Huanqiang Zeng, and Jiantao Zhou. Hyperspectral image super-resolution via deep progressive zero-centric residual learning. *IEEE Transactions on Image Processing*, 30:1423–1438, 2020.