

A Comprehensive Benchmark for Optical Remote Sensing Image Super-Resolution

Cesar Aybar^{ID}, David Montero^{ID}, *Student Member, IEEE*, Simon Donike^{ID},
Freddie Kalaitzis^{ID}, and Luis Gómez-Chova^{ID}, *Senior Member, IEEE*

Abstract—In recent years, there has been a growing interest in using image super-resolution (SR) techniques in remote sensing. These techniques aim to reconstruct high-resolution (HR) imagery from low-resolution (LR) sources. Despite the development of sophisticated SR methodologies, determining what constitutes “good” SR is still a matter of debate. Present-day literature often presents SR models through a strong computer vision perspective, heavily relying on synthetic datasets. Moreover, commonly used metrics often prioritize attributes that do not necessarily correspond to improvements in spatial resolution. To address this challenge, we present *OpenSR-test*, a comprehensive benchmark designed exclusively for evaluating SR of remote sensing images. Our framework incorporates specific quality metrics and curated cross-sensor datasets, each spanning various scale factors with consistent metadata. Utilizing *OpenSR-test*, we evaluate the state-of-the-art SR algorithms from a remote sensing perspective. The *OpenSR-test* framework and datasets are publicly available at <https://esaopnsr.github.io/opensr-test/>.

Index Terms—Benchmarking, datasets, deep learning, NAIP, Sentinel-2 (S2), SPOT, super-resolution (SR).

I. INTRODUCTION

OPTICAL remote sensing is a valuable source of information that predominantly uses satellite sensors to gather data of the Earth’s surface. The resolution of these sensors is determined by their ability to distinguish changes over time (temporal resolution), different wavelengths of the electromagnetic spectrum (spectral resolution), and spatially separated objects (spatial resolution). Depending on the application, remote sensing users may require higher resolution in one or more of these dimensions. For instance, in crop monitoring, a high temporal resolution may be preferred to track crop phenology and harvesting, while in the case of building delineation, a very high spatial resolution may be preferred

to capture the fine details of individual structures. However, remote sensing sensors do not always provide the required resolution due to technical and economic constraints.

In scenarios where spatial high-resolution (HR) images are not available, super-resolution (SR) algorithms emerge as a prominent solution [1]. SR is inherently an inverse problem, as it reconstructs the original HR image from its degraded low-resolution (LR) counterpart. Moreover, it is characterized as an ill-posed problem, given the one-to-many relation between LR and possible HR representations. Over recent years, substantial research efforts have been dedicated to the SR of natural images, particularly by employing deep-learning methods [1]. This letter has been directed toward developing more efficient architectures that integrate local and global features [2], obtaining more realistic degradation kernels [3], producing SR images over different scales [4], and proposing ingenious loss functions [5].

Despite significant advancements, a crucial question remains in the remote sensing community: can SR methods designed for natural images be effectively applied to satellite imagery? These SR methods use artificially degraded images through basic bicubic or bilinear interpolation as input. This approach circumvents the need for “real” LR images and results in LR–HR pairs that are spatially and spectrally consistent [3]. However, this approach may overlook the unique characteristics of the satellite sensors, potentially introducing a domain bias between synthetically generated and real LR observations. To mitigate this issue, recent advancements in SR methods have included simultaneous image fusion and SR processes [6] though these have shown limitations in preserving reflectance coherence, as recently noted in [5]. Alternatively, some studies have explored implementing harmonized cross-sensor SR approaches. Nevertheless, the preparation of training datasets for these approaches faces significant limitations due to the required preprocessing, including spatial alignment, removal of cloud-contaminated pixels, spatial coregistration, and radiometric adaptation [7], which ultimately constrains the dataset size and, consequently, the scalability of the model.

Another concern is the need for quality metrics explicitly designed for measuring the ground-resolved distance (GRD), a key indicator of spatial image resolution. While previous studies have often relied on pixelwise metrics, like peak signal-to-noise ratio (PSNR) and perceptual metrics, such as structural similarity index measure (SSIM) and learned perceptual image patch similarity (LPIPS) [8], these metrics have limitations. On the one hand, pixelwise metrics can

Manuscript received 16 February 2024; revised 7 May 2024; accepted 12 May 2024. Date of publication 15 May 2024; date of current version 22 May 2024. This work was supported by the European Space Agency (ESA, Φ-Lab, OpenSR Project). The work of Cesar Aybar was supported by the Consejo Nacional de Ciencia, Tecnología e Innovación Tecnológica (CONCYTEC), Peru (“Proyectos de Investigación Básica—2023-01” Program) under Grant PE501083135-2023-PROCIENCIA. The work of Luis Gómez-Chova was supported by the Spanish Ministry of Science and Innovation (funded by MCIN/AEI/10.13039/501100011033) under Project PID2019-109026RB-I00. (*Corresponding author: Cesar Aybar*.)

Cesar Aybar, Simon Donike, and Luis Gómez-Chova are with the Image Processing Laboratory (IPL), University of Valencia, 46010 Valencia, Spain (e-mail: cesar.aybar@uv.es).

David Montero is with the Remote Sensing Centre for Earth System Research (RSC4Earth), Leipzig University, 04103 Leipzig, Germany.

Freddie Kalaitzis is with the Department of Computer Science, University of Oxford, OX1 3QG Oxford, U.K.

Digital Object Identifier 10.1109/LGRS.2024.3401394




Fig. 1. Workflow for dataset generation: image selection (yellow blocks), harmonization (green blocks), and human inspection (purple block). For each LR–HR pair within each dataset, we report one spatial quality metric (QA_1) and one spectral quality metric (QA_2).

be too sensitive to spatial translations and local reflectance disparities unrelated to the SR problem [9]. On the other hand, perceptual metrics are not specifically intended for SR. They assess the general image quality based on human perception standards. Therefore, these metrics may not be the most accurate solution for domain-specific contexts like remote sensing, which requires specialized expertise.

This letter introduces a novel benchmark, *OpenSR-test*, for comprehensively evaluating SR techniques in optical remote sensing images. The benchmark presents two significant advancements over previous approaches [2], [10]. First, *OpenSR-test* comprises three distinct and carefully curated datasets, all compared against Sentinel-2 (S2), a reference in Earth observation. Second, we introduce a set of tailored quality metrics based on a remote sensing SR protocol. We demonstrate the utility of *OpenSR-test* by conducting detailed accuracy assessments of three SR methods and providing an in-depth analysis of the resulting metrics. Finally, we conclude by underscoring the importance that robust validation practices can play in shaping the future of SR in the remote sensing field.

II. CROSS-SENSOR MULTISCALE DATASETS

One of the primary goals of *OpenSR-test* is to prepare cross-sensor multiscale datasets with the maximum possible consistency between LR and HR image pairs. We selected S2 Level-2A (10-m bands) as the input LR imagery and three different preprocessed image sources to define the target HR domain: NAIP (2.5 m), SPOT (2.5 m), and VENμS (5 m). The entire process can be divided into three steps, as depicted in Fig. 1: selection of potential LR–HR pairs, harmonization, and visual inspection. For the LR–HR pair selection, we discarded all the LR–HR pairs that were not acquired on the same day to ensure similar atmospheric conditions. LR–HR pairs with cloud cover in the LR image were automatically discarded using a cloud detection algorithm trained on the CloudSEN12 dataset [11].

The harmonization phase comprises two key tasks: improving spatial collocation and radiometric matching. For spatial alignment assessment, ground control points between each LR and HR image pair are automatically identified using the LightGlue and DISK algorithms [12], [13]. Following the approach outlined in SEN2VENμS [7], we exclude any

TABLE I
DATASETS AVAILABLE IN *OpenSR-test*. THE HR AND LR PAIRS CONSIST OF FOUR BANDS: RED, GREEN, BLUE, AND NEAR INFRARED. THE WIDTH OF THE GAUSSIAN DEGRADATION KERNEL IS ESTIMATED BY COMPARING THE LR AND DOWNSAMPLED HR IMAGES

HR source imagery	Number of scenes	HR patch size (pix.)	Scaling factor	Gauss. kernel width (m)	Quality flags	
					$QA_1 \downarrow$	$QA_2 \downarrow$
NAIP [14]	30	1024	$\times 4$	5.50	0.33	1.30
SPOT 6/7 [6]	12	512	$\times 4$	5.75	0.44	1.21
VENμS [7]	59	512	$\times 2$	2.10	0.26	1.57

erroneous correspondence of points with a displacement greater than 30 m (3 pixels). The remaining points are used for calculating the mean square collocation error (quality flag QA_1 in Fig. 1). After collocating the LR–HR images, any image pair with a $QA_1 > 0.75$ is removed from the dataset. For radiometric harmonization, we apply histogram matching individually for each band and then compute the average spectral angle distance (SAD) between the LR–HR images (quality flag QA_2 in Fig. 1). LR–HR pairs with a $QA_2 > 5$ are not considered.

Finally, to ensure the optimal quality of the dataset, we carefully inspect the remaining LR–HR pairs through a manual review process. Any LR–HR pairs containing saturated or defective pixels, or other noticeable inconsistencies detected through visual inspection are excluded. The final datasets are accessible on HuggingFace at <https://huggingface.co/datasets/isp-uv-es/opensr-test>. Furthermore, for each LR–HR pair in the datasets, we report the spatial (QA_1) and spectral (QA_2) errors. The mean values of these quality metrics are presented in Table I for the three datasets (the lower the better).

III. PROTOCOL FOR REMOTE SENSING IMAGE SR EVALUATION

The second main goal of this work is to quantify the quality of SR images. Quantifying the quality of synthesized images is a well-known problem in remote sensing. Twenty-five years ago, Wald et al. [15] introduced a set of properties that an ideal pansharpening algorithm must satisfy to ensure its practical utility. Since then, numerous adaptations have been introduced, leveraging this protocol to tackle a comprehensive range of image enhancement issues, even transcending those related to multispectral data [16]. However, there is still no formal quality assessment protocol for SR in remote sensing.




Fig. 2. High-level summary of our workflow to estimate the consistency, synthesis, and correctness metrics. Users can choose from a wide range of distance metrics (blue blocks) according to their needs. The dotted lines represent the optional harmonization step conducted before the triple distance process (correctness). The red line ($d_{LR \leftrightarrow HR}$) is used to normalize the distances $d_{SR \leftrightarrow LR}$ and $d_{SR \leftrightarrow HR}$.

Wald’s protocol specifies two core quality properties that algorithms must satisfy: consistency and synthesis. Consistency entails that the LR image can be recovered by degrading the reconstructed (SR) image although the degradation strategy is usually unknown. Synthesis, on the other hand, states that a reconstructed (SR) image must preserve the characteristics of the original LR image at an increased spatial resolution. Inspired by Wald’s protocol, we propose a set of properties that can serve as guidelines for designing SR methods, metrics, and datasets, as presented in Section II.

SR images derived from LR images must satisfy the following three properties.

- 1) *Consistency Property*: Any SR image, when degraded (downsampled) to the original LR spatial resolution, SR_{down} , must maintain consistent reflectance values and spatial alignment with its LR counterpart. Testing this property can be challenging because the degradation model to generate LR (i.e., SR_{down}) from SR is typically unknown. However, assuming that the high-frequency characteristics between SR and HR should be similar, a degradation model learned from HR to LR might be employed to degrade SR.
- 2) *Synthesis Property*: Any SR image must improve the effective spatial resolution. In addition, the SR model must preserve the low-frequency details from the original LR image. If this condition is met, the metric of synthesis can be quantified by measuring the distance between the SR image and a reference image above the Nyquist frequency. The reference image can be obtained by simply upsampling the LR image (see Section IV). The larger the distance between the SR and the LR reference image, the more effective the gain in spatial resolution and detail representation.
- 3) *Correctness Property*: Any SR model must avoid hallucinations and omissions. Hallucinations are related to generating high-frequency features that do not align with the real ones present in the HR image. On the other hand, omissions denote the missing high-frequency features that the SR image failed to capture. One practical approach to evaluate the correctness property is to

measure the distances between the SR image to the HR and LR images, as we present in Section IV.

IV. SR QUALITY METRICS

The third objective of the *OpenSR-test* framework is to introduce specialized quality metrics that fulfill the described protocol for remote sensing image SR. The proposed quality metrics for SR can be computed at various aggregation levels: individual pixels, patches, or entire images; and are categorized into three distinct groups attending to the described *Consistency*, *Synthesis*, and *Correctness* properties (Fig. 2).

In the realm of *consistency*, our framework introduces three distinct metrics, which are delineated in purple in Fig. 2: *reflectance*, *spectral*, and *spatial* consistency. The core objective here is to evaluate the fidelity with which the SR image retains the intrinsic properties of the LR image. To facilitate this assessment, we downsample the SR image to the exact spatial resolution as the LR image, producing SR_{down} . We generate SR_{down} using bilinear interpolation by default, applying an antialiasing filter and then resampling. First, the *reflectance* metric is quantified using the mean absolute error. This metric focuses on determining the extent to which the SR process affects the reflectance values. Second, the *spectral* metric aims to ascertain the SR model’s capability in preserving the LR image’s spectral signatures, an essential factor in maintaining the spectral quality of the image. To compute this metric, we employ the SAD. Finally, the *spatial* consistency metric is derived by calculating the mean absolute error between the matching points identified by LightGlue in both SR and LR images. This measure is crucial for assessing the spatial alignment and structural integrity of the SR image relative to the LR image.

Before the synthesis and correctness estimation, the *OpenSR-test* framework performs two sequential steps.

- 1) The LR image is upsampled to match the HR spatial resolution, creating LR_{up} . Users can choose different methods, but bilinear interpolation is the default.
- 2) The SR image undergoes a harmonization process (SR_{harm}) to handle systematic errors identified during consistency inspection. The process includes correcting




Fig. 3. Visual representation of the three fundamental components within the correctness metric evaluation. The L1-norm defines the metric space. The first diagram in the figure delineates the improvement space (im), the second delineates the omission space (om), and the third depicts the hallucination space (ha). In all metrics, a value closer to zero indicates better accomplishment. Each diagram displays three stars representing the SR models' correctness evaluation. Refer to Section V for further information on the SR model setup. The gray areas depict regions where the Euclidean triangle inequality is not satisfied, i.e., where there are no possible solutions.

the reflectance values using histogram matching and performing spatial alignment using ground control points. These control points are obtained by applying LightGlue in the SR–HR pair.

With the three images standardized (i.e., three points in a high-dimensional space defining a triangle), we proceed to distance estimation, as depicted in Fig. 2. First, we calculate the distance between SR_{harm} and LR_{up} , and call it $d_{SR \leftrightarrow LR}$. This distance provides insights into the divergence of the SR image from the naive upsampled LR image (*omission*). Next, we calculate the distance between SR and HR, referred to as $d_{SR \leftrightarrow HR}$. This particular distance evaluates the extent to which the SR image aligns with the content of the reference HR image (*improvement*). Finally, we calculate the distance between LR_{up} and HR, referred to as $d_{LR \leftrightarrow HR}$. This distance indicates the amount of high-frequency information present in the HR image but absent in the LR image.

The default setting for distance calculations in the *OpenSR-test* is the L1-norm. However, users can choose from six distance metrics: L2, SAD, percentage bias (PBIAS), inverted peak signal-to-noise ratio (IPSNR), Kullback–Leibler divergence (KLD), and LPIPS. Choosing the right distance metric hinges on the user's end goal. For applications prioritizing reflectance preservation, the L1-norm is the recommended choice. On the other hand, for SR tasks with purely visual objectives, adopting LPIPS as the distance metric can lead to more accurate outcomes. We normalize $d_{SR \leftrightarrow LR}$ and $d_{SR \leftrightarrow HR}$ by the reference distance $d_{LR \leftrightarrow HR}$. This normalization results in relative distances, which we define as omission (d_{om}) and improvement (d_{im}) dimensions, respectively. The d_{om} dimension measures the total gain in high-frequency information and is used to quantify the synthesis property.

Regarding *correctness*, the primary objective is to classify and quantify how much of the high-frequency representation is an improvement (SR close to HR), omissions (SR close to LR), or hallucinations (SR far from both LR and HR). The *OpenSR-test* framework reports the percentage of pixels classified within these categories. Using the d_{om} and d_{im} , we estimate the omission (om), improvement (im), and hallucination (ha) scores using the following equations (Fig. 3):

$$H = d_{\text{im}} + d_{\text{om}} - 1 \quad (1)$$

$$\text{om} = d_{\text{om}} + d_{\text{im}}(1 - e^{-\gamma_{\text{om}}H}) \quad (2)$$

$$\text{im} = d_{\text{im}} + d_{\text{om}}(1 - e^{-\gamma_{\text{im}}H}) \quad (3)$$

$$\text{ha} = e^{-\gamma_{\text{ha}}d_{\text{im}}d_{\text{om}}} \quad (4)$$

The parameters γ_{om} , γ_{im} , and γ_{ha} govern the respective domains of the omission, improvement, and hallucination scores. We set γ to 0.8 for both the om and im scores, and to 0.4 for the ha score. These values were determined empirically based on the consensus of three remote sensing experts. The scores for om, im, and ha are used to categorize data units as hallucinations, omissions, or improvements based on the score that registers the minimum value. Refer to Fig. 3 for a visual representation of these dimensions.

V. RESULTS

A. Experimental Setup

We compare three pretrained SR models using different datasets and architectures to showcase the effectiveness of our benchmark. One of the models, *SR4RS* [17], employed S2 images, while the other two models, *diffuser* and *SuperImage*, were trained using natural images. The evaluation metrics reported in our comparison only considered the red, green, and blue bands to ensure consistency across all SR model configurations.

Despite the extensive literature on remote sensing SR, our experiment did not include other SR models because only *SR4RS* provides open access to both the weight parameters and the code of their trained models. Pretrained models available in the *SUPER-RES* [5] and *WorldStrat* [6] projects were omitted as they focus on multi-image SR, whereas *OpenSR-test* is designed explicitly for single-image SR.

B. Overall Performance Comparison

The comparative performance analysis of the three SR models, presented in Table II, offers a detailed assessment across the *OpenSR-test* datasets, including NAIP, SPOT, and VEN μ S.

From a *consistency* perspective, the *SuperImage* and *SR4RS* instances display an acceptable balance in preserving spatial, reflectance, and spectral characteristics across all evaluated

TABLE II

BENCHMARK COMPARISON OF SR MODELS. DOWNWARD ARROWS ($\downarrow$) DENOTE METRICS IN WHICH LOWER VALUES ARE PREFERABLE, AND UPWARD ARROWS ($\uparrow$) INDICATE METRICS IN WHICH HIGHER VALUES REFLECT BETTER PERFORMANCE

		Consistency			Synthesis	Correctness		
		reflectance $\downarrow$	spectral $\downarrow$	spatial $\downarrow$	high-frequency $\uparrow$	ha $\downarrow$	om $\downarrow$	im $\uparrow$
NAIP	SuperImage	0.008	7.286	0.131	0.003	0.117	0.784	0.098
	SR4RS	0.016	3.471	1.156	0.010	0.869	0.077	0.054
	diffusers	0.463	12.437	2.88	0.013	0.905	0.055	0.040
SPOT	SuperImage	0.009	3.512	0.062	0.006	0.160	0.794	0.046
	SR4RS	0.039	3.232	1.151	0.023	0.834	0.115	0.051
	diffusers	0.417	11.730	0.817	0.014	0.686	0.251	0.063
VEN μ S	SuperImage	0.009	8.687	0.099	0.003	0.403	0.380	0.217
	SR4RS	0.014	3.394	1.122	0.012	0.971	0.017	0.012
	diffusers	0.467	13.303	0.806	0.009	0.933	0.043	0.024

datasets. This indicates its suitability for remote-sensing downstream tasks. On the other hand, the *diffusers* model exhibits significant alterations in reflectance and spectral values, along with spatial displacements greater than 1 pixel at LR (10 m), which may suggest a degree of local geometric distortion.

From a *synthesis* perspective, *SR4RS* and *diffusers* models exhibit close scores. This outcome aligns with expectations, as generative models like *SR4RS* and *diffusers* tend to introduce more high-frequency energy into images compared to discriminative models such as *SuperImage*. Notably, *SuperImage* introduces approximately two times less high-frequency information, indicating a more conservative approach.

Lastly, in terms of *correctness*, the generative models (*SR4RS* and *diffuser*) demonstrate a higher incidence of hallucinations across all datasets (Fig. 3). These models tend to introduce artifacts that, while visually appealing, do not accurately correlate with the HR data. Conversely, the *SuperImage* model tends more toward the omission space (Fig. 3), implying that the added high-frequencies are similar to the LR_{up} image. While this approach results in fewer hallucinations, it may also lead to the exclusion of finer details, which is the primary motivation for utilizing an SR algorithm. A tradeoff between hallucinations and omissions highlights the importance of a balanced approach in remote sensing SR algorithms.

Finally, a particular characteristic can be noted in the VEN μ S dataset, where all algorithms tend to show higher hallucination values (Table II). Visual analysis indicates that the VEN μ S imagery appears smoother when compared to the SR results of the three SR networks. This observation is supported by comparing the degradation kernel widths (Table I); specifically, the kernel width of VEN μ S is less than half that of NAIP and SPOT, i.e., the effective spatial resolution of VEN μ S images appears to be worse than 5 m. Since the correctness metrics are based on comparing distances, if the reference HR is inadequate (i.e., smoother than expected), the results will always lean toward hallucination.

VI. CONCLUSION

This letter presents the *OpenSR-test*, a novel and comprehensive benchmark tailored for SR of optical remote sensing images. Designed around three foundational properties, the *OpenSR-test* provides a more precise interpretation of SR performance than traditional metrics and provides three curated cross-sensor datasets. While the current datasets offer valuable insights, their small size limits the generalization of results across regional or global scales. The *OpenSR-test* framework

can be seamlessly extended to tackle other remote sensing image synthesis problems such as cloud removal or image dehazing. We hope this work will encourage further research in the field of SR performance assessment.

REFERENCES

- [1] P. Wang, B. Bayram, and E. Sertel, "A comprehensive review on deep learning based remote sensing image super-resolution methods," *Earth-Sci. Rev.*, vol. 232, Sep. 2022, Art. no. 104110.
- [2] C. Tuna, G. Unal, and E. Sertel, "Single-frame super resolution of remote-sensing images by convolutional neural networks," *Int. J. Remote Sens.*, vol. 39, no. 8, pp. 2463–2479, Apr. 2018.
- [3] R. Dong, L. Mou, L. Zhang, H. Fu, and X. X. Zhu, "Real-world remote sensing image super-resolution via a practical degradation model and a kernel-aware network," *ISPRS J. Photogramm. Remote Sens.*, vol. 191, pp. 155–170, Sep. 2022.
- [4] H. Wu, N. Ni, and L. Zhang, "Learning dynamic scale awareness and global implicit functions for continuous-scale super-resolution of remote sensing images," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5602315.
- [5] P. Wolters, F. Bastani, and A. Kembhavi, "Zooming out on zooming in: Advancing super-resolution for remote sensing," 2023, *arXiv:2311.18082*.
- [6] J. Cornebise, I. Oršolić, and F. Kalaitzis, "Open high-resolution satellite imagery: The WorldStrat dataset—With application to super-resolution," 2022, *arXiv:2207.06418*.
- [7] J. Michel, J. Vinasco-Salinas, J. Inglada, and O. Hagolle, "SEN2VEN μ S, a dataset for the training of Sentinel-2 super-resolution algorithms," *Data*, vol. 7, no. 7, p. 96, Jul. 2022.
- [8] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 586–595.
- [9] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? A new look at signal fidelity measures," *IEEE Signal Process. Mag.*, vol. 26, no. 1, pp. 98–117, Jan. 2009.
- [10] Y. Wang et al., "Remote sensing image super-resolution and object detection: Benchmark and state of the art," *Expert Syst. Appl.*, vol. 197, Jul. 2022, Art. no. 116793.
- [11] C. Aybar et al., "CloudSEN12, a global dataset for semantic understanding of cloud and cloud shadow in Sentinel-2," *Sci. Data*, vol. 9, no. 1, p. 782, Dec. 2022.
- [12] P. Lindenberger, P.-E. Sarlin, and M. Pollefeys, "LightGlue: Local feature matching at light speed," 2023, *arXiv:2306.13643*.
- [13] M. Tyszkiewicz, P. Fua, and E. Trulls, "DISK: Learning local features with policy gradient," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 14254–14265.
- [14] NAIP: National Agriculture Imagery Program, USDA-FSA-APFO Geospatial Enterprise Oper., USA, 2023.
- [15] L. Wald, T. Ranchin, and M. Mangolini, "Fusion of satellite images of different spatial resolutions: Assessing the quality of resulting images," *Photogramm. Eng. Remote Sens.*, vol. 63, no. 6, pp. 691–699, 1997.
- [16] W. G. C. Bandara and V. M. Patel, "HyperTransformer: A textural and spectral feature fusion transformer for pansharpening," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2022, pp. 1767–1777.
- [17] R. Cresson, "SR4RS: A tool for super resolution of remote sensing images," *J. Open Res. Softw.*, vol. 10, no. 1, p. 1, Mar. 2022.