

Unsupervised Deep Learning-Based Pansharpening With Jointly Enhanced Spectral and Spatial Fidelity

Matteo Ciotola^{1b}, *Graduate Student Member, IEEE*, Giovanni Poggi^{1b}, *Member, IEEE*,
and Giuseppe Scarpa^{1b}, *Senior Member, IEEE*

Abstract—In latest years, deep learning (DL) has gained a leading role in the pansharpening of multiresolution images. Given the lack of ground truth data, most DL-based methods carry out supervised training in a reduced-resolution domain. However, models trained on downsized images tend to perform poorly on high-resolution target images. For this reason, several research groups are now turning to unsupervised training in the full-resolution domain, through the definition of appropriate loss functions and training paradigms. In this context, we have recently proposed a full-resolution training framework that can be applied to many existing architectures. Here, we propose a new DL-based pansharpening model that fully exploits the potential of this approach and provides cutting-edge performance. Besides architectural improvements with respect to previous work, such as the use of residual attention modules, the proposed model features a novel loss function that jointly promotes the spectral and spatial quality of the pansharpened data. In addition, thanks to a new fine-tuning strategy, it improves inference-time adaptation to target images. Experiments on a large variety of test images, performed in challenging scenarios, demonstrate that the proposed method compares favorably with the state-of-the-art both in terms of numerical results and visual output. The code is available online at <https://github.com/matciotola/Lambda-PNN>.

Index Terms—Deep learning (DL), image enhancement, image fusion, image registration, super resolution.

NOMENCLATURE

Symbol	Description
R	Resolution ratio.
B	Number of multispectral bands.
M, P	Original multispectral and panchromatic components.
\hat{M}	Pansharpened image.
$\hat{M}_{\downarrow}, \hat{M}_{\downarrow a}$	$(R \times)$ downscaled version of \hat{M} without/with alignment.
\tilde{M}	$(R \times)$ upscaled version of M .
$P^{\text{lp}}, \hat{M}^{\text{lp}}$	Low-pass filtered versions of P and \hat{M} .
$\mathcal{L}_{\lambda}, \mathcal{L}_S, \mathcal{L}$	Spectral, spatial and total loss.
$\langle \cdot \rangle$	Spatial and/or spectral average.

Manuscript received 29 June 2023; accepted 18 July 2023. Date of publication 1 August 2023; date of current version 9 August 2023. (*Corresponding author: Giovanni Poggi.*)

Matteo Ciotola and Giovanni Poggi are with the Department of Electrical Engineering and Information Technology, University of Naples Federico II, 80125 Naples, Italy (e-mail: matteo.ciotola@unina.it; poggi@unina.it).

Giuseppe Scarpa is with the Department of Engineering, Parthenope University of Naples, 80143 Naples, Italy (e-mail: giuseppe.scarpa@uniparthenope.it).

Digital Object Identifier 10.1109/TGRS.2023.3299356

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see <https://creativecommons.org/licenses/by-nc-nd/4.0/>

I. INTRODUCTION

DUE to physical and technological constraints, passive sensors are subject to a tradeoff between spectral and spatial resolution: the more spectral bands, the lower their spatial resolution. To overcome this limitation, multiresolution observation systems mount two sensors on the same flying platform, a multispectral (MS) sensor with high spectral resolution and a single-band panchromatic (PAN) sensor with high spatial resolution. The two pieces of information are then fused offline through the so-called pansharpening process [1] to recover the ideal full-resolution image.

Pansharpening has been a very active research field for more than 20 years. With the advances in technology and the increasing diffusion and use of remote sensing multiresolution data, a growing number of researchers have confronted the problem of pansharpening, proposing several ingenious approaches. According to the taxonomy proposed in [2], these can be roughly cast into four main categories: component substitution (CS) [3], [4], [5], [6], [7], [8], [9], multiresolution analysis (MRA) [10], [11], [12], [13], [14], [15], [16], variational optimization (VO) [17], [18], [19], and machine/deep learning (ML/DL) [20], [21], [22], [23], [24], [25], [26], [27], [28]. Both CS and MRA methods inject the PAN component into the resized (by interpolation) MS image, using, however, different schemes. The former performs the injection in a transform domain, obtained for example through principal component analysis (PCA) [29], Gram–Schmidt (GS) projection [4], and generalized intensity–hue–saturation (G-IHS) transform [30], where the “strongest” component is replaced by a suitably equalized version of the PAN image. The latter operate on “detail” (high-frequency) components, hence requiring multiresolution decomposition transforms, such as wavelets [10], [12], [31], [32] or Laplacian pyramids [11], [33], [34], [35], [36]. VO methods, instead, rely on the optimization of suitable acquisition or representation models, such as resolution downgrading models [17], sparse representations [18], total variation [19], and low-rank matrix representations [37].

In recent years, DL-based methods have taken center stage. The impressive performance observed in a number of computer vision and image processing problems, from denoising to super-resolution, to detection and classification, raised high expectations for similar improvements in remote sensing

applications. However, for pansharpening, this quantum leap is yet to come, even though the first DL-based method, PNN [20], dates back to 2016, and a plethora of other solutions have been proposed since then [21], [22], [23], [24], [25], [27], [38], [39], [40], [41], [42], [43], [44], [45], [46]. Two peculiar issues of pansharpening have long prevented reaping the full potential of DL: poor generalization ability and lack of ground truth data. Let us briefly analyze these problems.

In remote sensing, the wild variability of observed images, due to the diversity of sensors, scenes, and operating conditions, make it difficult to generalize data not seen in training. In computer vision, this problem is usually solved by increasing the training set and by using suitable forms of augmentation. Such solutions are hardly viable in remote sensing, due to the scarcity of high-quality training data (often proprietary) and the peculiarities of multiresolution imaging. Only recently, some field-specific forms of augmentation have been proposed [47] which appear to provide appreciable improvements. A more radical solution was proposed in [48]. The available pretrained model is adapted to the statistics of the target image by a few cycles of fine-tuning, trading off some processing time for a largely improved quality. This approach, called target-adaptive operating modality, solves the generalization issue at its root and is rapidly gaining popularity.

Turning to the second limiting issue, the lack of ground truth data prevents the use of straightforward and effective supervised training. To cope with this problem, early proposals [20], [21] resorted to supervised training in a reduced resolution domain. The network was trained on decimated versions of PAN and MS components, with the original MS acting as ground truth. The underlying assumption was that networks trained at low resolution would work equally well in the full-resolution domain. Experimental evidence, however, has clearly shown this scale-invariance hypothesis not to hold. Good performance at low resolution hardly scales to full resolution, and low-quality images are often obtained. Therefore, the idea of training pansharpening models on the original high-resolution data has lately gained momentum [49], [50]. With this approach, there is no scale mismatch problem. On the downside, lacking a ground truth, one must forsake supervised training and turn to alternative unsupervised solutions [49], [50], [51], [52], [53]. In particular, in [50], we developed a full-resolution training framework based on ad hoc losses, such that the quality of the fused image is assessed indirectly, by measuring how closely it matches the PAN in the spatial domain and the MS in the spectral domain. Early experiments show clear quality improvements for the resulting pansharpened images, fully supporting the validity of this choice.

This short analysis highlights recent important theoretical and practical advances in DL-based pansharpening. Building upon these ideas and tools, we propose here a new DL-based pansharpening method that further improves on the current state-of-the-art. First of all, we operate at full resolution, avoiding any processing of the original data that could impair image quality. We start from the full-resolution training framework developed in [50] and improve it in two ways: 1) by adopting a new spectral loss, with perceptually motivated reprojection

metrics in place of Euclidean norms and 2) by improving the accuracy of such metrics through loss-time realignment of spectral bands. Then, in the context of this improved training framework, we propose a new architecture that takes advantage of several promising and established tools, such as residual learning, and spatial and channel attention modules. To ensure good generalization to new data, we adopt the target-adaptive operating modality. However, we develop a new data preprocessing scheme which removes its computational bottlenecks, thus enabling its use in all real-world applications with off-the-shelf hardware.

We carried out a large number of experiments to validate our design choices and to assess the performance of the proposed method in comparison with state-of-the-art competitors, both model-based and data-driven. To this end, we considered a wide variety of datasets and the most challenging cross-dataset operating conditions. Numerical results show the proposed method to be always among the best performers, and often the best. Output images are characterized by high spatial and spectral fidelity, displaying a fine PAN-MS coregistration, obtained with no intervention on the part of the user. Our software tool is available online (upon publication) to enable in-depth analysis and further developments.

In summary, the main contributions of this work are as follows.

- 1) Novel attention-based residual learning architecture.
- 2) Full-resolution training with perceptually motivated spectral and spatial losses.
- 3) Automatic PAN-MS coregistration at subpixel precision.
- 4) Self-contained software module available online characterized by fast (almost size invariant) processing.
- 5) Thorough experimental validation on a large number of datasets and in challenging conditions.

In the rest of the manuscript, we account for related work (Section II), describe in detail the proposed method (Section III), carry out ablation studies and report on comparative experimental results (Section IV), and finally, draw conclusions (Section V).

II. RELATED WORK

In this section, we review the state-of-the-art on DL-based pansharpening. However, we neglect methods trained at reduced resolution, referring the reader to a recent review [54], and focus on those performing unsupervised training in the full resolution domain, more strictly related to our proposal. In addition, we describe the correlation-based spatial loss proposed in [50] and adopted here.

A. Unsupervised Methods in the Full Resolution Domain

Methods that work at full resolution overcome the scale mismatch problem, opening the way to superior performance. Lacking a ground truth, the central problem becomes the definition of a suitable unsupervised loss function that provide meaningful indications for network optimization. To the best of our knowledge, only four very recent papers [49], [50], [51], [52] have been proposed to date on this topic. All of them define the loss function as a weighted sum of spectral and spatial (sometimes called structural) terms, \mathcal{L}_λ , and \mathcal{L}_S ,

TABLE I

UNSUPERVISED LOSSES PROPOSED IN THE LITERATURE FOR HIGH-RESOLUTION PANSHARPENING $I = \sum_b \alpha_b \widehat{M}_b \simeq P$

Loss name	\mathcal{L}_λ	\mathcal{L}_S	more details
\mathcal{L}_{SSQ} [49]	$\ \widehat{M}^{\text{lp}} - \widetilde{M}\ _2 + [1 - \text{SSIM}(\widehat{M}^{\text{lp}}, \widetilde{M})]$	$\ P - I\ _2 + [1 - \text{SSIM}(P, I)]$	$\{\alpha_b\}$ estimated at reduced resolution; additional regularization term \mathcal{L}_{reg}
\mathcal{L}_{GDD} [51]	$\ \widehat{M}_\downarrow - M\ _2$	$\ \nabla P - \nabla I\ _1$	$\{\alpha_b\}$ learned in the training phase; \widehat{M}_\downarrow : bicubic downscaling of \widehat{M} .
\mathcal{L}_{PG} [52]	$\ \widehat{M}_\downarrow - M\ _2$	$\ \nabla P - \nabla I\ _2$	$\alpha_b = 1/B \forall b \rightarrow I = \langle \widehat{M} \rangle$; \widehat{M}_\downarrow : bilinear downscaling of \widehat{M} ; additional adversarial term \mathcal{L}_{adv} .
$\mathcal{L}_{\text{Z-PNN}}$ [50]	$\ \widehat{M}_\downarrow - M\ _1$	$\langle (1 - \rho^\sigma)u(\rho^{\text{max}} - \rho^\sigma) \rangle$	$\rho^\sigma = \text{corr}(P, \widehat{M})$, $\rho^{\text{max}} = \text{corr}(P^{\text{lp}}, \widehat{M})$; \widehat{M}_\downarrow : low-pass filtering and decimation of \widehat{M} ; $u(\cdot)$: unit-step function.

plus some possible additional terms, accounting for hybrid features, adversarial games, or simple weight regularization. Table I provides a synoptic view of such unsupervised losses, while Nomenclature lists the domain-specific symbols most frequently used in the following of this article.

Let us first consider the spectral loss term, \mathcal{L}_λ , whose goal is to quantify the consistency of the fused output image, \widehat{M} , with the $R \times R$ times smaller input MS, M . To compare them, these two images must be adjusted to the same size, which can be done by upscaling M [49] or downscaling \widehat{M} [50], [51], [52]. In both cases, there are implementation choices that impact heavily on the final result. Only Luo et al. [49] follow the first path. The MS is expanded, $M \rightarrow \widetilde{M}$, through interpolation (not specified in this article). Since this latter image lacks high-frequency details, also the model output is smoothed $\widehat{M} \rightarrow \widehat{M}^{\text{lp}}$, with a low-pass Gaussian-shaped filter. Finally, a combination of Euclidean distance and structural similarity (SSIM) index is used to compare them. The other path requires downgrading the model output, $\widehat{M} \rightarrow \widehat{M}_\downarrow$, so as to compare it with the original MS. A clear advantage of this second solution is that the spectral reference M is left unaltered. On the other hand, the procedure used to downgrade \widehat{M} also affects the accuracy of spectral consistency measurement. In [51] and [52], standard interpolators are used before decimation, bicubic, and bilinear, respectively. In our previous work [50], instead, a physically explainable process is implemented, with a Gaussian-shaped low-pass filter matching the modulation transfer function (MTF). In all cases, L_p norms are used to compute the loss, Euclidean in [51] and [52], L_1 in [50] to favor sharper edges and ensure faster convergence.

The goal of the spatial loss, \mathcal{L}_S , instead, is to measure how accurately the spatial structures of the high-resolution PAN are reproduced in the pansharpened image. The authors [49], [51], and [52] make the fundamental assumption that the PAN can be accurately estimated through a linear combination of the high-resolution spectral bands

$$I = \sum_{b=1}^B \alpha_b \widehat{M}_b \simeq P. \quad (1)$$

In other words, a space-invariant linear relationship is assumed to exist between the MS and PAN domains. Given this assumption, the problem reduces to computing the regression coefficient $\{\alpha_b\}$. They are estimated on a reduced resolution dataset in [49], learned during the training process in [51], and simply assumed to be all equal in [52]. The spatial loss is then computed by measuring the dissimilarity between I and P . In [49], a combination of Euclidean distance and SSIM is again proposed, while both [51] and [52], to better preserve high-frequency details, work on the gradients of I and P , using L_1 and L_2 norm, respectively.

Considering the limits and risks of a linear space-invariant model, already pointed out in [55], and in [50], we propose a radically different correlation-based spatial loss. Given its importance for this work, we describe it in detail below.

B. Correlation-Based Spatial Loss

Let X and Y be two image patches, and let σ_X^2 , σ_Y^2 , and σ_{XY} indicate their sample variances and covariance. Then, the correlation coefficient between X and Y is defined as

$$\rho_{XY} = \text{Corr}(X, Y) = \frac{\sigma_{XY}}{\sigma_X \sigma_Y}, \quad -1 \leq \rho_{XY} \leq 1. \quad (2)$$

The correlation coefficient indicates to what extent one image can be linearly predicted from the other, with $|\rho| = 1$ implying perfect predictability and $\rho = 0$ total incorrelation.

Now, the role of the spatial loss is to inject the detailed spatial structures of the PAN into the output image. Accordingly, one may be tempted to define a loss that forces the correlation between the PAN and each pansharpened band to be unitary. However, this would produce exact replicas of the PAN in each band, which is not our goal. We want to preserve the peculiar dynamics of each band, tolerating also the presence of local areas with low correlation. Therefore, we apply the basic idea with two corrections: 1) only *local* correlations are considered and 2) they are not forced to be unitary but only to reach a suitable reference level.

More formally, let $\widehat{M}^\sigma(i, j, b)$ be a $\sigma \times \sigma$ patch drawn from band b of the pansharpened image at spatial location

(i, j) , and $P^\sigma(i, j)$ the corresponding patch drawn from the PAN. We compute the correlation field

$$\rho^\sigma(i, j, b) = \text{Corr}(P^\sigma(i, j), \widehat{M}^\sigma(i, j, b)). \quad (3)$$

At the same time, from low-pass versions of the same quantities, P^{lp} and \widehat{M} , we compute a reference correlation field, $\rho^{\sigma, \text{max}}(i, j, b)$, which represents the local target for the correlation (the reader is referred to [50] for more details). Eventually, we define the local loss as

$$\ell^\sigma(i, j, b) = \begin{cases} 1 - \rho^\sigma(i, j, b), & \rho^\sigma < \rho^{\sigma, \text{max}} \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

and the overall spatial loss as its average on space and bands

$$\mathcal{L}_S = \langle \ell^\sigma(i, j, b) \rangle. \quad (5)$$

In practice, this loss pushes the local correlation to grow, but only until the conservative target value defined by the reference field is reached. Clearly, the size of the correlation window, σ , plays a critical role in performance. It should be large enough to allow reliable estimates over the local window, but not so large to lose resolution. Experiments in [50] show $\sigma = R$, the PAN-MS resolution ratio, to be a good compromise value. The reference field, instead, uses a window of size R^2 , since it is computed on smoother images.

III. PROPOSED METHOD

Our proposal builds upon the full-resolution training framework proposed in [50]. Here, we introduce a new spectral loss which, thanks to accurate coregistration, collaborates with the correlation-based spatial loss to provide enhanced spectral and spatial fidelity. In addition, we propose a fast training adaptation procedure and a deeper architecture including residual and attention modules.

A. Improved Spectral Loss

The L_1 and L_2 norms often used in the spectral losses are simple and popular measures of distortion but only loose proxies of the image quality as perceived by the end user. Several perceptual metrics have been proposed over the years that better fit the human visual system, such as the SSIM [56] or the universal image quality index (UIQI) [57]. In pansharpening, the use of L_p norms as quality indicators has been questioned for a long time and several alternative metrics have been proposed and are by now widely accepted.

A first popular measure of spectral quality is the *Erreur Relative Globale Adimensionnelle de Synthèse* (ERGAS) [58]. Given a B-band image, x , and its approximation, y , it is defined as

$$\text{ERGAS}(y, x) = \frac{100}{R} \sqrt{\frac{1}{B} \sum_{b=1}^B \frac{\|y_b - x_b\|^2}{\langle x_b \rangle^2}}. \quad (6)$$

Also widespread is the $Q2^n$ metric [59]. It generalizes to b-bands images the UIQI metric defined for single-band

images y and x , as

$$\text{UIQI}(y, x) = \left\langle \frac{\sigma_{XY}}{\sigma_x \sigma_y} \cdot \frac{2\sigma_x \sigma_y}{\sigma_x^2 + \sigma_y^2} \cdot \frac{2\mu_x \mu_y}{\mu_x^2 + \mu_y^2} \right\rangle \quad (7)$$

where μ_x , μ_y , σ_x , σ_y , and σ_{XY} are sample means, variances and covariance computed over a $W \times W$ window centered on corresponding pixels x and y of x and y , respectively. To extend UIQI to the B-band images, all statistics are redefined in vector/matrix form and properly combined [59], but the overall behavior of the metric remains the same. In the remote sensing field, there is wide agreement on the meaningfulness of both these metrics. Interestingly, they complement each other. In fact, while ERGAS is based on pixelwise measurements, $Q2^n$ relies on local statistics computed on relatively large windows.

Note that if only a reduced resolution target image, x' , is available, these metrics can still be used by properly downgrading the estimate y as well. For example, the spectral distortion index $D_\lambda^{(K)}$ proposed by Khan et al. [60] can be seen as an extension of $Q2^n$

$$D_\lambda^{(K)}(y_\downarrow, x') = 1 - Q2^n(y_\downarrow, x'). \quad (8)$$

Similar extensions have been also recently proposed [61] for ERGAS and other indexes.

Given the above considerations, we chose to define a new spectral loss based on these perceptually relevant metrics. In particular, to exploit the complementary features exhibited by ERGAS and $D_\lambda^{(K)}$, we consider their linear combination, with a weighting parameter, γ , to be set on the basis of experiments. In more detail, we use these metrics to compare the original MS component with the pansharpened image reprojected in the MS domain. This latter process consists in applying an MTF-matched Gaussian filter to all bands of \widehat{M} and decimating them. Eventually, the proposed spectral loss reads as

$$\mathcal{L}_\lambda = D_\lambda^{(K)}(\widehat{M}_\downarrow, M) + \gamma \text{ERGAS}(\widehat{M}_\downarrow, M). \quad (9)$$

Note that all terms are computed with respect to the original MS component, M , which is neither coregistered nor expanded or resampled, all operations that could degrade the quality of the reference.

B. Coregistration at Loss

For technological reasons, multiresolution images often present some misalignments, also in high-level products. Different spectral bands are shifted with respect to one another and to the PAN, with errors of up to a few PAN-scale pixels. Fig. 1 shows on the top an example of this phenomenon: band misalignment is clearly visible in the MS image, especially near geometric structures, where pixels with unnatural colors appear. Of course, to obtain a high-quality pansharpened image, such errors must be compensated for. A possible solution is to coregister the MS bands before proceeding with actual pansharpening. However, this operation modifies and possibly impairs the original MS reference data. In addition,

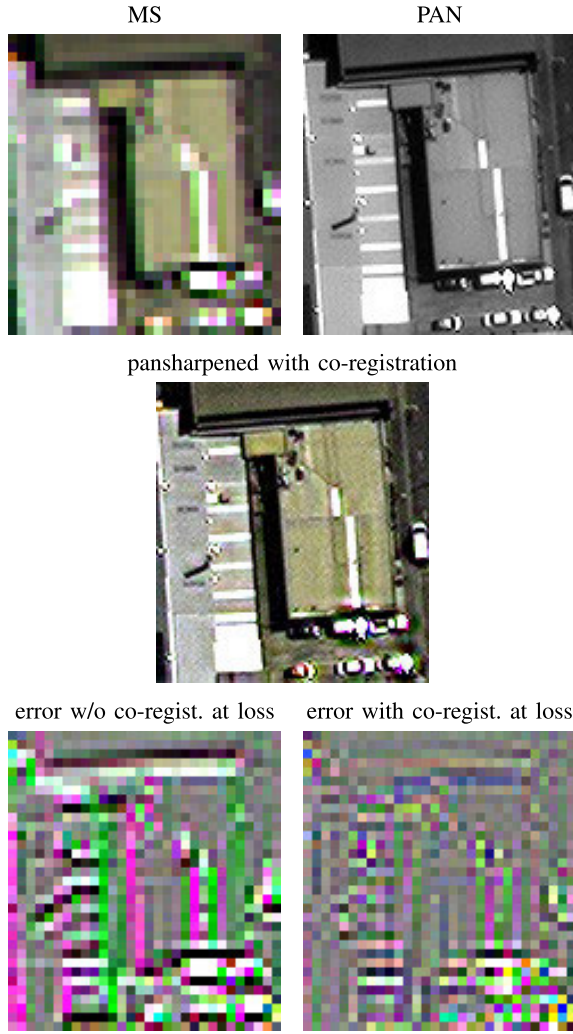


Fig. 1. Importance of coregistration at loss. (Top) Source data. The green/pink out-of-context pixels and strips in the false-color MS are due to band misalignment. (Middle) Pansharpened image with correct band coregistration. (Bottom) Difference (magnified for better visualization) between decimated pansharpened image and MS, $\widehat{M}_1 - M$, (left) without and (right) with coregistration at loss. Despite the good visual quality observed in the pansharpened image, large errors are observed with respect to the MS (e.g., pink/green strips near object boundaries) if loss-time band alignment is not performed. The corresponding spectral loss will be large, undermining correct model optimization.

it requires the active involvement of the user, not always appropriately skilled.

A better solution is to perform coregistration during pansharpening [50], [53], [62]. This is accomplished automatically in our full-resolution training framework, with no need for user intervention. In fact, to minimize the spatial loss, the local correlation between the PAN and each pansharpened band must be large, and this happens only when the spatial structures in all bands are correctly aligned. Going back to Fig. 1, the middle row shows the example image pansharpened by our method. Misalignment problems have been solved automatically and all bands appear to be correctly coregistered, greatly benefiting visual quality.

However, all our efforts would be undermined if we did not perform the so-called *coregistration at loss*. The example of

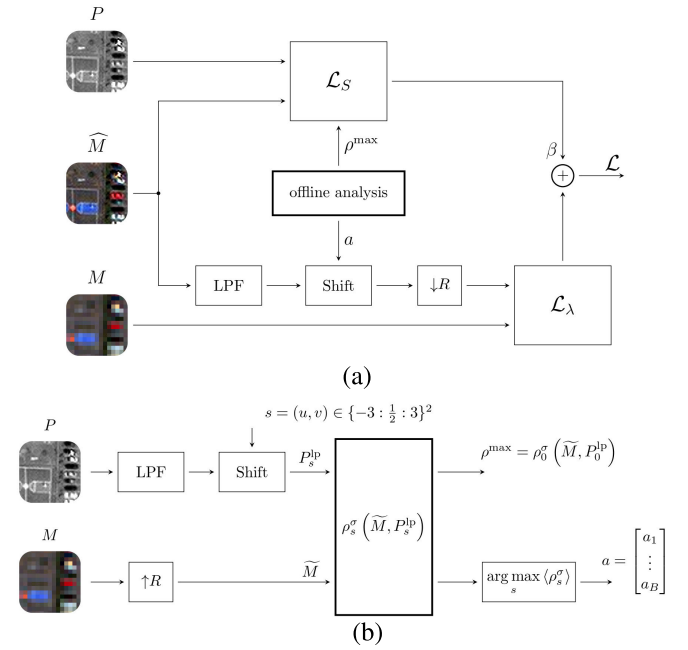


Fig. 2. Computation of JESSE loss. (a) Online operations: the two branches compute spatial and spectral loss components. (b) Offline correlation-based analysis provides the optimal band shifts and the reference correlation field.

Fig. 1 will help us describe this issue. To compute the spectral loss, we must compare the decimated pansharpened image with the original MS. The former has been automatically coregistered while minimizing the spatial loss, but the latter presents band misalignment. Therefore, the difference image (bottom-left) will exhibit large errors and large spectral loss. So, while the spatial loss pushes toward band alignment, the spectral loss discourages it, a highly undesirable situation in which the two losses work against each other.

Fortunately, once identified, this annoying problem has a simple solution. We estimate in advance the global shift of each MS band with respect to the reference PAN. Then, at the moment of computing the spectral loss, and only to this aim, we temporarily shift back the bands of the pansharpened image to align them with the original MS before decimation. In practice, we reintroduce the band misalignment only to compute a meaningful spectral loss. The effect is visible in the bottom-right difference image. Large spectral errors are avoided and spatial and spectral losses are both small, in accordance with the good visual quality of the image.

C. Computing the JESSE Loss

We can now summarize the details of the total loss computation with the help of Fig. 2. The upper part (a) shows a high-level block diagram of online operations, where the top and bottom branches concern spatial and spectral losses, respectively. Note that both branches rely on products generated in advance through offline analyses: the reference correlation map, ρ^{\max} , required for the computation of \mathcal{L}_S , and the vector of bandwise alignment shifts, a , for the computation of \mathcal{L}_λ . This is a qualifying point of this approach. In fact, since ρ^{\max} and a are functions of P and M , both losses

eventually depend, either explicitly or implicitly, on both PAN and MS, that is,

$$\mathcal{L}_S = \mathcal{L}_S(\widehat{M}, P; \rho^{\max}) = \mathcal{L}_S^*(\widehat{M}, P; M) \quad (10)$$

$$\mathcal{L}_\lambda = \mathcal{L}_\lambda(\widehat{M}, M; a) = \mathcal{L}_\lambda^*(\widehat{M}, M; P). \quad (11)$$

Indeed, the correlation and alignment information, computed by exploiting all the available source data, put in tight relation the spectral and spatial losses, which do not fight with each other anymore but concur to provide the highest fidelity under all dimensions. Accordingly, we call this loss JESSE, after *Joint Enhancement of Spectral and Spatial fidelity*.

Note also that in the spectral loss defined in (9) \widehat{M}_\downarrow must be replaced by $\widehat{M}_{\downarrow a}$ to account for the band alignment in the computation of the distortion metrics and, accordingly, the total loss becomes

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_\lambda + \beta \mathcal{L}_S \\ &= D_\lambda^{(k)}(\widehat{M}_{\downarrow a}, M) + \gamma \text{ER GAS}(\widehat{M}_{\downarrow a}, M) \\ &\quad + \beta \langle (1 - \rho^\sigma) u(\rho^{\max} - \rho^\sigma) \rangle \end{aligned} \quad (12)$$

with $\rho^\sigma = \text{Corr}(P, \widehat{M})$.

Fig. 2(b) describes in detail the offline correlation-based analysis block. Since PAN and MS must be compared, they are brought in the same signal space. The MS is expanded by the resolution ratio R , by means of the 23-tap polynomial interpolator proposed in [11], to obtain its upsampled version, \widehat{M} . On the other side, the high-frequency content of P , not present in \widehat{M} , is removed through low-pass filtering. Then, we compute their correlation field using a window of size R^2 rather than R to match input images that have been smoothed. If no shift is applied, we obtain the reference field used to compute the spatial loss. However, in the presence of band misalignment, the PAN-MS matching may be improved by suitable shifts, and therefore, we look for the optimal shift vector a that maximizes the average correlation. In practice, we perform the search for the optimal shift band-by-band on a discrete grid, $\mathcal{S} = \{-3 : (1/2) : 3\}^2$, including all displacement with horizontal and vertical components going from -3 to $+3$ PAN-scale pixels at half-pixel steps. In formulas

$$a_b = \arg \max_{s \in \mathcal{S}} \langle \rho_s^\sigma \rangle = \arg \max_{s \in \mathcal{S}} \langle \text{Corr}(P_s^{\text{lp}}, \widehat{M}(b)) \rangle \quad (13)$$

with $\text{Corr}(\cdot, \cdot)$ computed only on valid pixels.

Eventually, the optimal shifts and the reference correlation field are provided as side input for the computation of the JESSE loss. It is worth underlining that this process does not involve the prediction \widehat{M} , hence its computational impact is negligible with respect to the cost of the fine-tuning cycles.

D. Fast Target Adaptation

A key asset of the proposed method is the use of a target-adaptive pansharpening modality, which adapts on-the-fly the pretrained model to the peculiar statistics of the target image. This inference-time fine-tuning process was first introduced in [48] in a supervised context, and then adapted in [50] for application in a full-resolution unsupervised context. In Fig. 3, we show its high-level workflow. The starting point is a DL model with initial weights $\phi^{(0)}$, obtained by offline pretraining on a suitably large dataset.

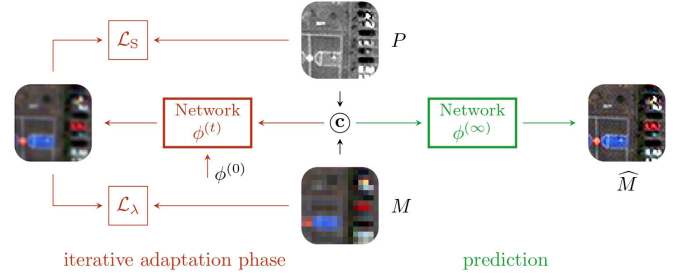


Fig. 3. Full-resolution target-adaptive inference scheme. The pretrained weights $\phi^{(0)}$ are iteratively fine-tuned (red blocks) to the target image. At convergence, the final weights $\phi^{(\infty)}$ are used for actual pansharpening (green block).

However, no matter how large and varied the training set, these weights will hardly fit the unique features of the target image, in terms of scene, illumination, and acquisition process, including possible band misalignments. So, to overcome this problem, the model undergoes an iterative tuning process (red blocks on the left of the figure) on the target image itself. The weights $\phi^{(t)}$ available at time t are used to pansharpen the source image. Then, the resulting loss is computed (the figure shows separate spectral and spatial terms as customary) and used to update the weights to the next values $\phi^{(t+1)}$. At the end of the process, the optimized weights, $\phi^{(\infty)}$, are used for actual pansharpening (green block on the right). Experiments show that target adaptation (TA) improves consistently the pretrained models, especially when the target image is not well represented by the training set.

In principle, the adaptation process should proceed until full convergence, but this would significantly delay real-time operations, and only a few tuning steps are carried out in practice. Moreover, the computational cost grows rapidly with image size and model capacity, preventing its application in important real-world cases. A first attempt to address this issue [63] provides only minor improvements. So, we propose here a new adaptation scheme that relies only on a fixed-size tuning sample drawn from the target image, thereby ensuring a low and almost size-invariant computational complexity.

The core of the proposed method is the selection of a suitable fine-tuning dataset, small enough to allow for fast online operations, yet so diverse to capture all the major features of the scene. This process is described in Fig. 4. First, the image (a) is partitioned into N_{tile} tiles of $c \times c$ -pixels. Then, the tiles (b) feed a CNN-based classifier, MobileNet v3-small [64], which extracts compact 576-component descriptors (c). These latter are further compacted (d) by means of PCA, retaining only the three principal components. The resulting vectors, regarded as points in a 3-D space (e), are processed by k -means clustering to obtain $N_{\text{clust}} \ll N_{\text{tile}}$ groups of similar vectors. The idea is that each cluster should be representative of a major land cover of the scene. Therefore, we eventually select a single feature vector per cluster (the median) to represent the whole group and include the associated tile in the fine-tuning dataset (f).

By acting on tile size and number of clusters, one can make this procedure more or less aggressive. Experiments described

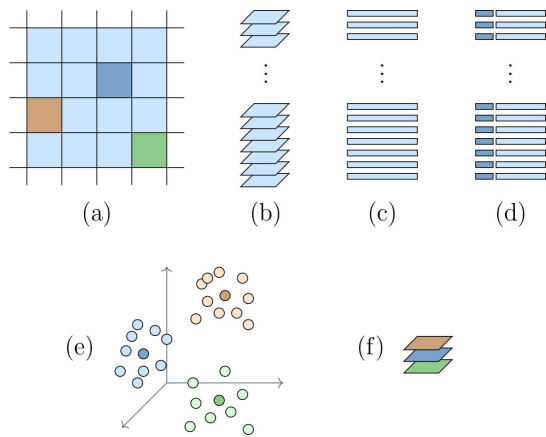


Fig. 4. Selection of the dataset for fast target adaptation. (a) Image tiling. (b) Initial dataset with all tiles. (c) Compact features extracted by the CNN. (d) PCA: the left part is kept. (e) k -means clustering and template selection. (f) Final dataset comprising only a few tiles representative of all land covers.

in Section IV show that the proposed fast procedure ensures excellent TA with processing times that are small, and almost invariant to image size.

E. λ -PNN: A Deeper Attention-Based Network Architecture

Thus far, we have improved the unsupervised training framework, especially from the point of view of spectral quality, and also the target-adaptive operating modality, which is now more accurate and efficient.

We now turn to propose a new model which takes advantage of the improved unsupervised training framework. The architecture is shown in Fig. 5(a). Before going into detail, we observe that, with seven convolutional layers and two attention modules, it is significantly deeper than our previous proposals. This is a direct consequence of the fast target-adaptive modality described above, which allows us to fine-tune a much heavier network than before without any major impact on the processing time.

Our architecture is of the residual type [65]: a global skip connection brings the resized MS directly to the output, where it is added to the output of the convolutional branch. In addition, also most of the convolutional blocks have a residual structure. This choice reflects the intuition that part of the desired image is already available (the low-frequency MS component), and only the high-frequency detail should be actually predicted. On the other hand, most DL models proposed for pansharpening in latest years [21], [22], [24], [28], [49], [50], rely on residual architectures, with advantages in terms of training speed and generalization properties.

In the convolutional branch, there are two 64-channel convolutional layers with ReLU activations, followed by two residual blocks Fig. 5(b), comprising again two 64-channel convolutional layers each and GeLU activation [66], and a final B -channel convolutional layer with linear activation feeding the summing node. In addition, there are two convolutional block attention modules (CBAMs) [67], these too in residual configuration. The R-CBAM modules aim of focusing attention on especially relevant portions of the input, both in space and along the channels. Their architecture is shown

TABLE II

WV3, WV2, AND GE1 DATASETS, WITH NUMBER OF CROPS FOR TRAINING, VALIDATION, AND TEST. ADELAIDE AND WASHINGTON, COURTESY OF DIGITALGLOBE[®]. FORTALEZA, MEXICO CITY, XI'AN, BERLIN, LONDON (WV2), ROME, NORIMBERGA, WATERFORD, AND GENOA (DIGITALGLOBE[®]) PROVIDED BY ESA. MUNICH, MIAMI, LONDON (GE1), AND TRENTON ARE PART OF PAIRMAX DATASET [68]

WorldView-3 (GSD at nadir: 0.31m)				
Dataset (PAN size)	Training (512×512)	Val. (512×512)	X-Val. (512×512)	Test (2048×2048)
Fortaleza	32	8	-	-
Mexico City	32	8	-	-
Xian	32	8	-	-
Adelaide	-	-	24	8
Munich (PairMax)	-	-	-	3

WorldView-2 (GSD at nadir: 0.46m)				
Dataset (PAN size)	Training (512×512)	Val. (512×512)	X-Val. (512×512)	Test (2048×2048)
Berlin	32	8	-	-
London	32	8	-	-
Rome	32	8	-	-
Washington	-	-	24	10
Miami (PairMax)	-	-	-	2

GeoEye-1 (GSD at nadir: 0.41m)				
Dataset (PAN size)	Training (512×512)	Val. (512×512)	X-Val. (512×512)	Test (2048×2048)
Norimberga	32	8	-	-
Rome	32	8	-	-
Waterford	32	8	-	-
Genoa	-	-	24	9
London (PairMax)	-	-	-	1
Trenton (PairMax)	-	-	-	1

in Fig. 5(c). In the channel attention module, global spatial pooling (max and average) is performed on two parallel paths followed by a shared multilayer perceptron. The resulting vectors, summed to one another and squashed on the 0–1 interval by a sigmoid, encode the channel importance. Accordingly, they multiply the 64-channel input stack emphasizing some channels more than others. The spatial attention module performs in a similar way, by exchanging the role of space and channel. After channelwise max and average pooling, the resulting feature maps are compacted to a single heatmap by a convolutional layer with sigmoid activation. This spatial attention map then multiplies the channel-emphasized input to put emphasis also on selected spatial sites. To help prevent vanishing gradients, the whole CBAM has a residual architecture, with the resulting feature stack summed to the input to provide the final output.

IV. EXPERIMENTAL ANALYSIS

In this section, after analyzing a sample experiment that provides more insight into the JESSE loss, we carry out the comparative performance assessment of the proposed method, studying numerical and visual results. Then, we show ablation studies that validate our design choices, and assess the efficiency of the new fast TA procedure.

A. Datasets

To obtain a reliable wide-spectrum assessment, we carry out our experiments on a large variety of data, making also

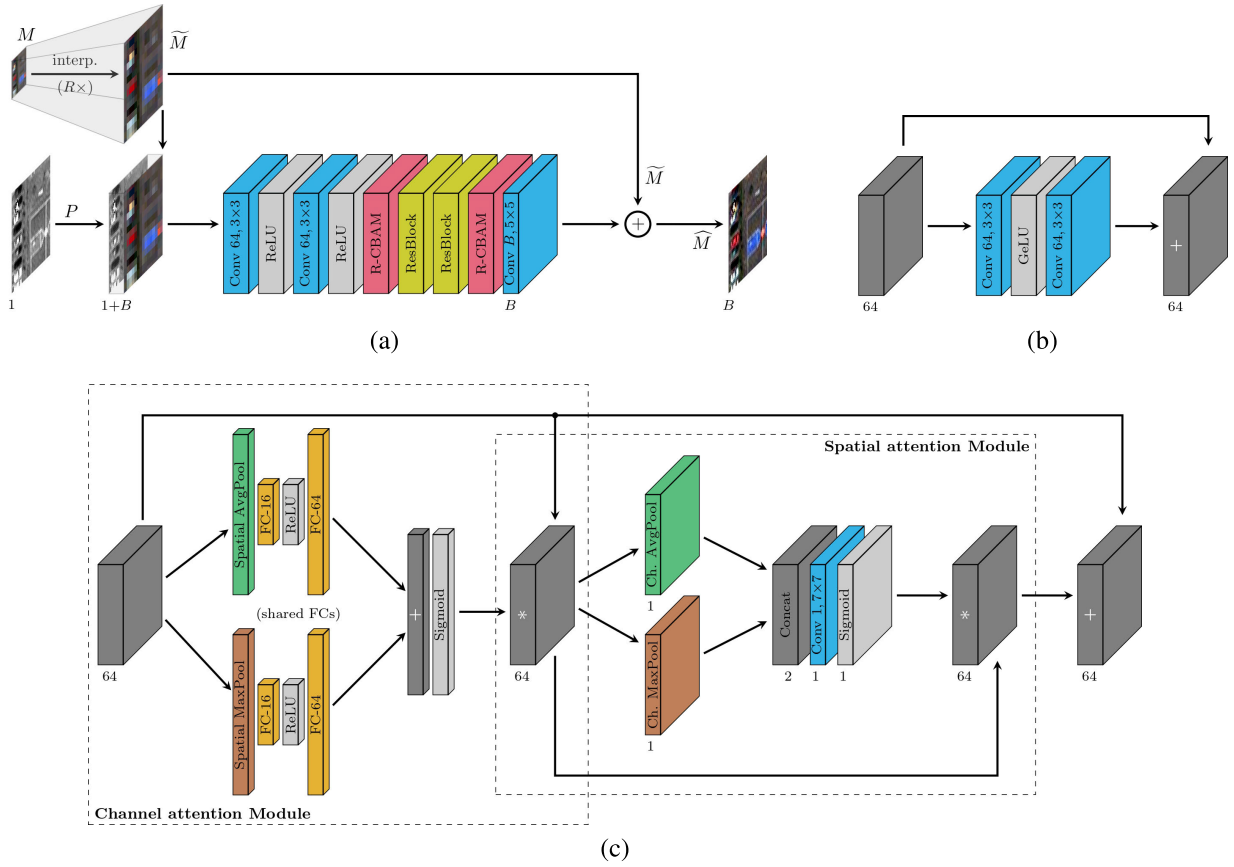


Fig. 5. λ -PNN pansharpening model. (a) Overall architecture. (b) ResBlock module. (c) R-CBAM module.

sure to test the generalization ability of competing methods. Therefore, we consider three distinct datasets, one for each of the sensors WorldView-3, WorldView-2, and GeoEye-1. For each dataset, we have several large images available (see Table II). Three of them are used exclusively for training and validation. Another image is never seen in training, but used to gather “cross-scenario” (change of place, date, daylight conditions, and sensing geometry) validation information and then also for testing. The last image is used exclusively for testing. This latter is the most challenging case since there is no link between the test image and the training process. Training and validation are carried out on 512×512 -pixel crops (PAN resolution) while there is no size constraint in the testing phase and we consider 2048×2048 crops. For all sensors, the PAN-MS resolution ratio is $R = 4$. In the following, we name datasets after the corresponding sensor: WV3, WV2, GE1, with the suffix Train, Val, X-Val, and Test when appropriate.

B. Reference Methods

We compare the performance of the proposed λ -PNN with a large number of reference methods. They are listed in Table III, grouped by their general approach: CS, MRA, VO, and ML, the latter trained at reduced or full resolution. Most of the methods are available in the toolboxes [2] and [54], from which we selected those that performed best in the experiments. Methods of the last group, instead, were downloaded

TABLE III

REFERENCE METHODS USED FOR COMPARATIVE ANALYSIS

Component Substitution (CS)

BT-H [14], BDSF [6], BDSF-PC [9], GS [4], GSA [5], PRACS [7]

Multiresolution Analysis (MRA)

AWLP [12], MTF-GLP [13], MTF-FS [16], MTF-HPM [13], MF [69]

Variational Optimization (VO)

FE-HPM [17], SR-D [18], TV [19]

Machine Learning (ML) Reduced Resolution

PNN [20], A-PNN [48], A-PNN-TA [48], BDPN [25], DiCNN [26], DRPNN [22], FusionNet [28], MSDCNN [24], PanNet [21]

Machine Learning (ML) Full Resolution

SSQ [49], GDD [51], PanGan [52], Z-PNN [50]

from the authors’ websites. For ML methods, together with the source code, the authors usually provide the weights obtained on their own training sets. In some cases (marked by an asterisk in the tables of results), weights were not available and we retrained the models on our dataset. Finally, we reimplemented and trained SSQ (marked with a double asterisk) because neither code nor weights were available.

C. Performance Metrics

In pansharpening, the lack of ground truths does not impact only model training but also performance assessment. All mea-

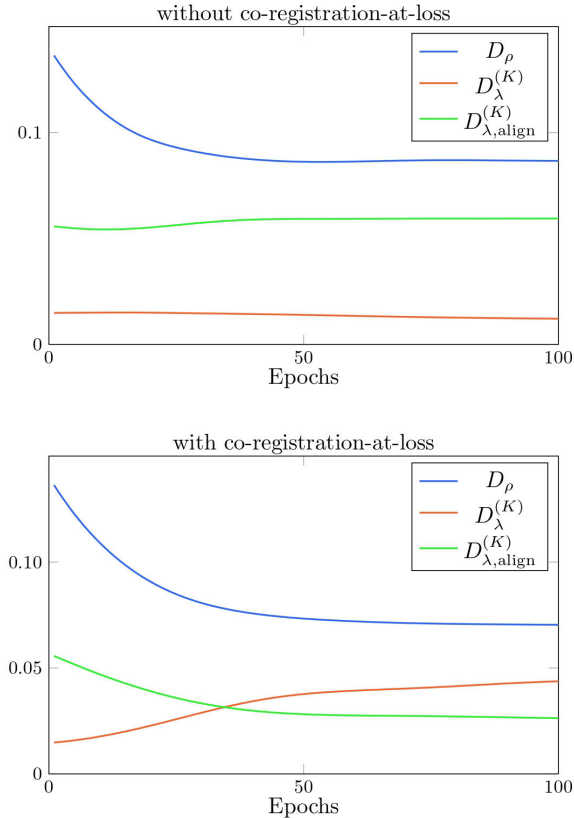


Fig. 6. Evolution of distortion metrics: D_ρ (blue), $D_\lambda^{(K)}$ (red), and $D_{\lambda,\text{align}}^{(K)}$ (green), with TA. (Top) Without realignment, TA is almost useless: $D_{\lambda,\text{align}}^{(K)}$ is always very high, indicating a large spectral distortion. (Bottom) With realignment, it becomes very effective: $D_{\lambda,\text{align}}^{(K)}$ reduces significantly and also D_ρ gets smaller than before.

asures of distortion are necessarily indirect and rely on explicit or implicit hypotheses that remain to be proved. Based on our reasoning and experience, we believe that the metrics used in λ -PNN to compute the total loss, that is, $D_{\lambda,\text{align}}^{(K)}$, R-ERGAS, and $D_\rho = \langle (1 - \rho^\sigma)u(\rho^{\max} - \rho^\sigma) \rangle$, provide a meaningful assessment of spectral and spatial fidelity. However, for the sake of completeness, we consider some more metrics largely used in the literature: Khan’s measure of spectral distortion without band alignment, $D_\lambda^{(K)}$, and two more spatial distortion metrics, D_S [70] and $D_S^{(R)}$ [71], for which description we refer the reader to the original papers.

D. On the Coherency of Spatial and Spectral Losses

To gain some more insight on the proposed JESSE loss, let us pansharpener with λ -PNN the image shown in Fig. 1, affected by obvious alignment problems. The network has been already pretrained on a large dataset and now must be fine-tuned on the target image. In Fig. 6, we show the evolution of spatial (D_ρ) and spectral ($D_{\lambda,\text{align}}^{(K)}$) distortion¹ as the adaptation proceeds.

In Fig. 6 (top), however, we use a “wrong” loss for adaptation, with the alignment vector set to $a = 0$, such that the pansharpener image is compared with the original MS without compensating the shifts. In other words, we replace

$D_{\lambda,\text{align}}^{(K)}$ with $D_\lambda^{(K)}$ in the loss. The effect of this wrong choice is to render TA almost ineffective. First of all, $D_{\lambda,\text{align}}^{(K)}$ (green curve) does not improve at all, it even worsens a little bit, because the network is trying instead to further improve $D_\lambda^{(K)}$ (red), which was already optimized offline. In addition, the spatial distortion D_ρ (blue) plateaus at 0.086 after a few cycles, because increasing the correlation with the PAN would align the bands and then increase the distortion with respect to the original MS, hence $D_\lambda^{(K)}$. In short, spatial and spectral loss components are fighting against each other.

In Fig. 6 (bottom), instead, the correct loss is used, with the optimal alignment vector a . Now, as the adaptation proceeds, the true measure of spectral distortion, $D_{\lambda,\text{align}}^{(K)}$, reduces constantly, going from 0.58 to 0.25, while $D_\lambda^{(K)}$ grows. In addition, also the spatial distortion, D_ρ , keeps lowering, reaching 0.071 after 100 cycles. Spectral and spatial losses provide coherent indications and reinforce one another, concurring with the joint enhancement of spectral and spatial fidelity.

E. Numerical Results

Tables IV–VI report numerical results for the WV3, WV2, and GE1 datasets, respectively. For each sensor there are two test datasets, e.g., Adelaide and Munich for WV3, for a total of six, and for each image, there are six columns of results, corresponding to three spectral and three spatial distortion metrics. If we consider only the metrics that contribute to the proposed JESSE loss, $D_{\lambda,\text{align}}^{(K)}$, R-ERGAS, and D_ρ , the proposed λ -PNN (last row) ranks always first (bold) or second (underlined) among all competing methods. This is a great achievement, and our goal from the beginning, because we think that these metrics are the most reliable indicators of quality. Moreover, the fact that λ -PNN succeeds in optimizing *both* spectral and spatial quality metrics further confirms that these latter provide coherent indications, with no domain mismatch. Critics will argue that this result is due, at least in part, to the alignment between the training and measurement phases. This is probably true, so we defer any further related consideration to visual inspection.

Nonetheless, even when we consider $D_\lambda^{(K)}$ (third column), λ -PNN keeps providing very good results, typically better than all competitors except for some MRA and TV methods. That is, spectral fidelity remains high despite the penalty given by the mismatch between MS and pansharpener output in the presence of imperfect coregistration. The situation is completely different when we consider D_S and $D_S^{(R)}$: according to these spatial distortion indicators, λ -PNN appears to be among the worst performers. We are not discouraged by these results, though, because we are skeptical about the ability of these indicators to really capture spatial quality, and report them for the sake of completeness. Indeed, unlike for *spectral* quality, the problem of no-reference *spatial* quality assessment is far from being solved. We refer the reader to our recent works [50], [61] for a deeper analysis of these metrics and, again, leave the final word to the visual inspection of images. However, it is worth noting that, on some images, the proposed λ -PNN is outperformed even by simple interpolation (EXP) according to these indicators. This speaks volumes about their

¹ERGAS is not shown to avoid cluttering the plots.

TABLE IV
AVERAGE RESULTS ON WV3-TEST. (LEFT) ADELAIDE. (RIGHT) MUNICH (PAIRMAX)

Method	Adelaide						Munich (PairMax)					
	$D_{\lambda, \text{align}}^{(K)}$	R-ERGAS	$D_{\lambda}^{(K)}$	D_{ρ}	D_S	$D_S^{(R)}$	$D_{\lambda, \text{align}}^{(K)}$	R-ERGAS	$D_{\lambda}^{(K)}$	D_{ρ}	D_S	$D_S^{(R)}$
EXP	0.069	3.853	0.049	0.859	0.114	0.179	0.064	3.911	0.075	0.854	0.132	0.161
BT-H	0.074	3.685	0.153	0.071	0.123	0.001	0.079	4.161	0.156	0.049	0.122	0.001
BDS	0.100	4.392	0.175	0.110	0.047	0.016	0.120	5.596	0.204	0.075	0.049	0.052
BDS-PC	0.078	3.929	0.159	0.075	0.068	0.012	0.108	5.447	0.190	0.058	0.057	0.048
GS	0.113	4.370	0.199	0.086	0.085	<u>0.003</u>	0.128	5.012	0.213	0.072	0.083	<u>0.001</u>
GSA	0.066	3.698	0.142	<u>0.069</u>	0.112	0.004	0.075	4.416	0.164	0.051	0.093	0.001
PRACS	0.061	3.546	0.080	0.191	0.051	0.030	0.063	3.919	0.099	0.195	0.053	0.023
AWLP	0.049	3.237	<u>0.038</u>	0.093	0.069	0.059	0.043	3.078	<u>0.053</u>	0.079	0.076	0.039
MTF-GLP	0.049	3.259	0.043	0.073	0.100	0.055	0.042	2.948	0.056	0.056	0.098	0.035
MTF-GLP-FS	0.051	3.311	0.041	0.097	0.088	0.054	0.043	2.979	0.056	0.067	0.088	0.034
MTF-GLP-HPM	0.052	3.330	0.054	0.081	0.088	0.059	0.048	3.383	0.069	0.061	0.092	0.039
MF	0.046	3.094	0.058	0.093	0.093	0.058	0.042	3.037	0.065	0.078	0.085	0.051
FE-HPM	0.050	3.320	0.046	0.093	0.087	0.062	0.043	3.103	0.057	0.073	0.089	0.043
SR-D	0.054	3.542	0.023	0.301	0.032	0.133	<u>0.034</u>	<u>2.777</u>	0.040	0.186	0.070	0.086
TV	0.036	<u>2.646</u>	0.058	0.205	0.035	0.049	0.040	2.888	0.074	0.168	0.079	0.033
PNN	0.197	6.675	0.269	0.461	0.073	0.116	0.416	9.115	0.548	0.475	0.105	0.118
A-PNN	0.070	3.591	0.100	0.534	0.082	0.096	0.169	4.398	0.274	0.665	<u>0.187</u>	0.167
A-PNN-FT	0.060	3.489	0.068	0.332	<u>0.026</u>	0.069	0.085	3.521	0.121	0.335	0.029	0.084
BDPN	0.150	5.393	0.261	0.178	<u>0.078</u>	0.008	0.294	7.693	0.440	0.311	0.115	0.024
DiCNN	0.158	5.602	0.246	0.414	0.082	0.056	0.217	6.077	0.291	0.410	0.092	0.080
DRPNN	0.151	5.230	0.243	0.195	0.067	0.011	0.198	6.598	0.304	0.186	0.086	0.012
FusionNet	0.099	4.228	0.151	0.412	0.067	0.094	0.238	5.302	0.328	0.340	0.058	0.125
MSDCNN	0.157	5.450	0.243	0.226	0.098	0.013	0.392	6.664	0.514	0.348	0.098	0.023
PanNet	0.055	3.436	0.045	0.338	0.014	0.087	0.061	3.322	0.082	0.297	<u>0.035</u>	0.061
SSQ**	0.060	3.539	0.059	0.271	0.045	0.028	0.083	3.567	0.118	0.284	<u>0.073</u>	0.018
GDD*	0.310	9.731	0.377	0.662	0.111	0.118	0.304	8.693	0.400	0.587	0.095	0.104
PanGan*	0.143	4.825	0.284	0.130	0.079	0.048	0.630	17.589	0.773	0.132	0.117	0.070
Z-PNN	0.044	2.834	0.106	0.088	0.119	0.032	0.091	3.756	0.153	0.101	0.141	0.033
λ -PNN	0.021	1.978	0.095	0.044	0.083	0.076	0.031	2.526	0.066	0.033	0.140	0.079

TABLE V
AVERAGE RESULTS ON WV2-TEST. (LEFT) WASHINGTON. (RIGHT) MIAMI (PAIRMAX)

Method	Washington						Miami (PairMax)					
	$D_{\lambda, \text{align}}^{(K)}$	R-ERGAS	$D_{\lambda}^{(K)}$	D_{ρ}	D_S	$D_S^{(R)}$	$D_{\lambda, \text{align}}^{(K)}$	R-ERGAS	$D_{\lambda}^{(K)}$	D_{ρ}	D_S	$D_S^{(R)}$
EXP	0.042	2.238	0.044	0.808	0.069	0.151	0.056	3.663	0.057	0.738	0.065	0.135
BT-H	0.056	2.428	0.101	0.071	0.114	0.000	0.064	3.819	0.111	0.110	0.087	0.000
BDS	0.208	4.679	0.325	0.186	0.093	0.091	0.118	5.303	0.201	0.237	0.070	0.049
BDS-PC	0.107	3.335	0.181	0.088	0.040	0.055	0.095	4.742	0.161	0.138	0.030	0.046
GS	0.108	3.378	0.185	0.110	0.096	0.001	0.086	4.503	0.142	0.149	0.089	<u>0.001</u>
GSA	0.058	2.505	0.108	0.081	0.113	<u>0.001</u>	0.063	3.946	0.114	0.108	0.077	0.002
PRACS	0.042	2.214	0.056	0.279	0.034	0.052	0.056	3.620	0.074	0.229	0.036	0.017
AWLP	0.029	1.829	<u>0.033</u>	0.086	0.079	0.045	0.033	2.870	<u>0.036</u>	0.129	0.059	0.061
MTF-GLP	0.031	1.836	0.043	0.065	0.109	0.045	0.033	2.841	0.039	0.105	0.077	0.056
MTF-GLP-FS	0.031	1.853	0.038	0.098	0.099	0.038	0.035	2.909	0.039	0.129	0.069	0.053
MTF-GLP-HPM	0.057	2.375	0.071	0.076	0.094	0.047	0.038	2.951	0.058	0.129	0.060	0.060
MF	0.039	1.958	0.062	0.083	0.103	0.052	0.032	2.817	0.045	0.130	0.069	0.066
FE-HPM	0.040	2.068	0.049	0.082	0.100	0.050	0.033	2.893	0.042	0.122	0.073	0.058
SR-D	<u>0.025</u>	1.753	0.018	0.245	0.031	0.083	0.028	2.719	0.027	0.258	0.041	0.123
TV	0.047	<u>1.616</u>	0.086	0.241	0.025	0.042	0.016	2.063	0.039	0.274	0.023	0.051
PNN	0.073	2.349	0.110	0.249	0.029	0.011	0.055	3.575	0.083	0.335	0.027	0.015
A-PNN	0.047	2.097	0.072	0.419	0.040	0.036	0.044	3.302	0.058	0.508	0.052	0.046
A-PNN-FT	0.043	2.076	0.064	0.274	<u>0.025</u>	0.021	0.040	3.177	0.053	0.338	0.026	0.027
BDPN	0.093	2.649	0.150	0.195	0.047	0.013	0.079	4.548	0.140	0.215	0.046	0.013
DiCNN	0.067	2.492	0.115	0.302	0.045	0.026	0.091	5.258	0.166	0.413	0.042	0.051
DRPNN	0.063	2.258	0.105	0.233	0.035	0.013	0.054	3.512	0.091	0.279	<u>0.024</u>	0.016
FusionNet	0.057	2.170	0.091	0.305	0.039	0.028	0.046	3.194	0.080	0.367	0.028	0.054
MSDCNN	0.082	2.429	0.138	0.231	0.047	0.010	0.060	3.607	0.106	0.253	0.031	0.023
PanNet	0.041	1.839	0.060	0.345	0.031	0.046	0.030	2.704	0.042	0.368	0.036	0.072
SSQ**	0.043	1.938	0.064	0.302	0.035	0.026	0.033	2.883	0.043	0.386	0.030	0.038
GDD*	0.282	6.351	0.369	0.762	0.080	0.124	0.228	7.640	0.317	0.449	0.093	0.099
PanGan*	0.244	5.616	0.341	0.203	0.059	0.106	0.165	6.513	0.247	0.207	0.070	0.069
Z-PNN	0.047	1.924	0.094	<u>0.046</u>	0.119	0.021	0.050	3.155	0.095	<u>0.080</u>	0.095	0.026
λ -PNN	0.020	1.291	0.051	0.042	0.094	0.058	<u>0.024</u>	<u>2.246</u>	0.055	0.050	0.068	0.086

TABLE VI
AVERAGE RESULTS ON GEOEYE-1. (LEFT) GENOA. (RIGHT) LONDON + TRENTON (PAIRMAX)

Method	Genoa						London+Trenton (PairMax)					
	$D_{\lambda,align}^{(K)}$	R-ERGAS	$D_{\lambda}^{(K)}$	D_{ρ}	D_S	$D_S^{(R)}$	$D_{\lambda,align}^{(K)}$	R-ERGAS	$D_{\lambda}^{(K)}$	D_{ρ}	D_S	$D_S^{(R)}$
EXP	0.136	4.463	0.120	0.805	0.096	0.225	0.063	4.434	0.093	0.805	0.078	0.133
BT-H	0.158	4.579	0.261	0.070	0.091	0.012	0.100	5.240	0.179	0.057	0.080	0.004
BDS	0.166	4.497	0.282	0.082	0.078	0.015	0.137	6.230	0.231	0.065	0.044	0.038
BDS-PC	0.167	4.495	0.284	0.081	0.079	0.015	0.136	6.193	0.229	0.062	0.048	0.037
GS	0.186	4.615	0.296	0.092	0.088	0.002	0.108	5.543	0.183	0.057	0.074	0.002
GSA	0.161	4.446	0.263	0.052	0.152	<u>0.008</u>	0.104	5.306	0.185	0.039	0.106	0.002
PRACS	0.126	4.199	0.147	0.246	0.053	0.087	0.072	4.673	0.125	0.183	0.062	0.027
AWLP	0.098	3.816	<u>0.088</u>	0.082	0.110	0.108	0.041	3.704	0.059	0.098	0.058	0.051
MTF-GLP	0.092	3.696	0.092	0.065	0.130	0.109	0.037	3.498	0.055	0.078	0.066	0.057
MTF-GLP-FS	0.097	3.795	0.092	0.087	0.132	0.098	0.037	3.515	0.055	0.083	0.068	0.056
MTF-GLP-HPM	0.102	3.721	0.106	0.067	0.109	0.115	0.036	3.557	0.055	0.077	0.059	0.058
MF	0.091	3.542	0.111	0.077	0.134	0.096	0.041	3.697	0.062	0.101	0.055	0.065
FE-HPM	0.096	3.786	0.097	0.077	0.125	0.099	0.039	3.669	0.059	0.088	0.064	0.053
SR-D	0.087	3.718	0.056	0.268	0.062	0.305	0.023	3.135	0.035	0.202	0.078	0.191
TV	0.110	3.576	0.171	0.693	0.070	0.114	0.056	4.213	0.100	0.675	0.051	0.062
PNN	0.110	3.858	0.109	0.471	0.089	0.115	0.046	3.811	0.069	0.376	0.032	0.057
A-PNN	0.103	3.796	0.096	0.540	0.095	0.150	0.041	3.791	0.064	0.497	0.072	0.075
A-PNN-FT	0.106	3.840	0.099	0.372	0.067	0.130	0.044	3.932	0.067	0.264	0.030	0.059
BDPN*	0.182	4.559	0.287	0.346	0.065	0.016	0.103	5.685	0.188	0.334	0.041	0.009
DiCNN*	0.175	4.306	0.265	0.342	0.060	0.030	0.147	5.898	0.263	0.263	0.028	0.020
DRPNN*	0.144	3.669	0.240	0.614	0.148	0.075	0.083	4.702	0.159	0.502	0.038	0.041
FusionNet*	0.228	4.838	0.341	0.346	0.159	0.026	0.176	6.393	0.319	0.201	<u>0.036</u>	0.013
MSDCNN*	0.174	3.951	0.268	0.436	0.111	0.047	0.101	5.065	0.198	0.306	<u>0.036</u>	0.023
PanNet*	0.096	<u>3.190</u>	0.129	0.455	0.076	0.113	0.039	3.785	0.068	0.356	0.037	0.071
SSQ**	0.122	3.792	0.166	0.335	<u>0.057</u>	0.052	0.055	4.110	0.099	0.284	0.036	0.033
GDD*	0.399	8.146	0.470	0.564	0.097	0.201	0.282	10.338	0.386	0.643	0.100	0.170
PanGan*	0.263	5.840	0.387	0.178	0.076	0.040	0.194	8.347	0.349	0.107	0.091	0.029
Z-PNN	<u>0.083</u>	3.211	0.155	0.120	0.133	0.084	0.047	3.868	0.078	0.092	0.080	0.038
λ -PNN	0.043	2.220	0.134	<u>0.054</u>	0.080	0.265	<u>0.026</u>	<u>3.193</u>	<u>0.049</u>	<u>0.042</u>	0.095	0.178




Fig. 7. Example results for GE1 Trenton-PairMax test image. For the λ -PNN pancharpened image (bottom-left) we have $D_{\rho} = 0.042$, $D_S = 0.095$, and $D_S^{(R)} = 0.178$, while we have instead $D_{\rho} = 0.805$, $D_S = 0.078$, and $D_S^{(R)} = 0.133$ for the image obtained through simple interpolation of the MS (bottom-right). Both D_S and $D_S^{(R)}$ provide misleading indications on spatial quality.

reliability. In Fig. 7, we show a relevant example that illustrates this phenomenon and fully supports our standing.

F. Visual Results

As already mentioned, finding a reliable measure of pancharpened image quality is still under intense research and there is no consensus, to date, on which metric or combination of metrics best fits the judgment and needs of end users. Therefore, we rely on visual inspection to confirm (or possibly disconfirm) our conclusions. An expert viewer can still spot spectral and spatial artifacts better than compact (averaged) indicators can.

Figs. 8–10 show visual results for the WV3, WV2, and GE1 datasets, respectively. For each sensor, two crops are selected, one for each dataset, e.g., Adelaide and Munich-PairMax for WV3. For each selected crop, we show the original MS and PAN, on the left, followed by the pancharpened outputs of the proposed method, λ -PNN, and the six strongest competitors, one per each distortion metric. For example, next to λ -PNN, we show the image generated by the method that performs best (except for λ -PNN itself) under the $D_{\lambda,align}^{(K)}$ distortion metric. This choice allows us to limit the number of images to inspect, but also to assess the proposed method against the most competitive reference methods. Note that the “challenger” under a given metric changes from image to image. For example, for the WV3 test set and $D_{\lambda,align}^{(K)}$, it is TV for Adelaide but SR-D for Munich. Of course, if a method is the best competitor under more than one metric, we avoid showing the same image multiple times and use the next in the list. MS images are shown using a red–green–blue (RGB bands) composition. In theory, artifacts may happen to occur *only*

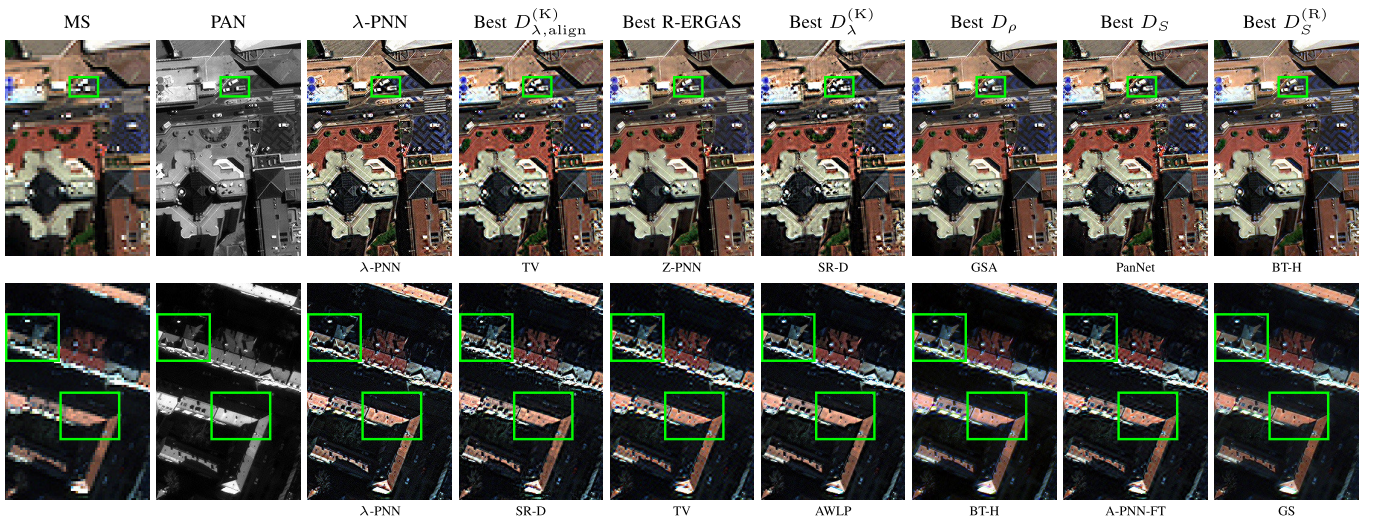


Fig. 8. Samples from WV3 (top) Adelaide and (bottom) Munich-PairMax. (Left to right) MS, PAN, λ -PNN, and best references.

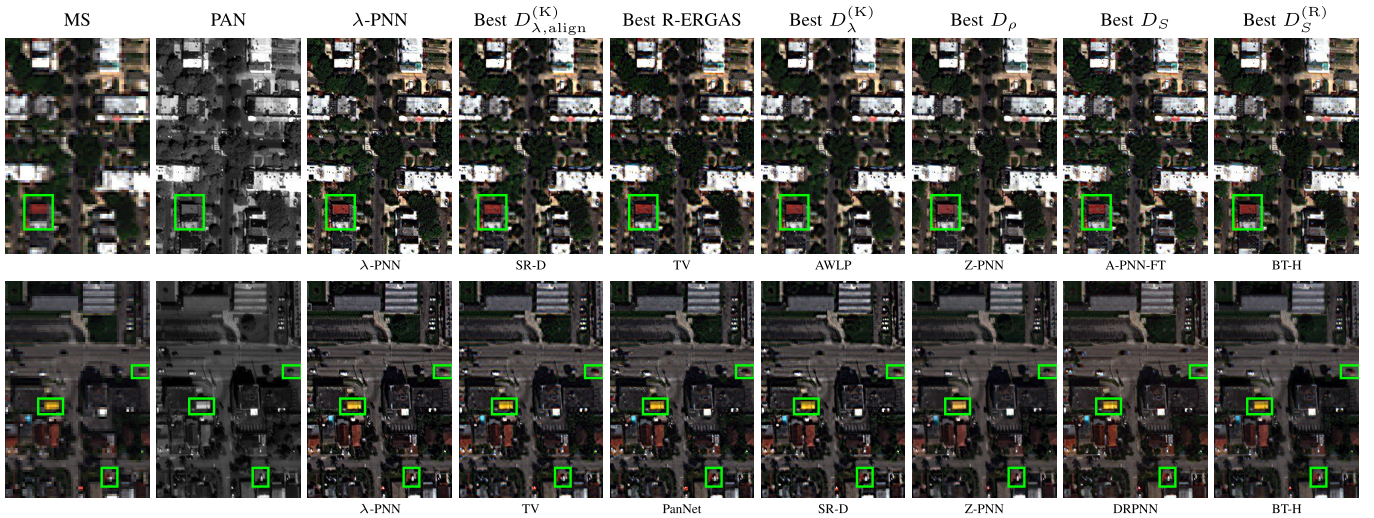


Fig. 9. Samples from WV2 (top) Washington and (bottom) Miami-PairMax. (Left to right) MS, PAN, λ -PNN, and best references.

in bands not shown, thus evading visual inspection. However, we checked several images, characterized by both low and high spectral distortion, and never spotted such phenomena.

To begin, we study the absolute quality of the images pan-sharpened by λ -PNN, that is, whether they can be satisfactory for the end user, no matter how they compare with competitors. To this end, we analyze PAN, MS, and the λ -PNN image together, looking for possible spectral or spatial artifacts. With fastidious scrutiny, we singled out a few geometric errors (marked by green boxes for easier inspection), like in the upper-most truck in Adelaide, or in the facade of the bottom-right building in Genoa. Likewise, minor color deviations can sometimes be observed with respect to the MS, e.g., some very small cars in Miami. All in all, these are minor exceptions to the general rule of high spatial and spectral fidelity, with textures accurately copied from PAN and colors closely matching those of MS. Nonetheless, we notice an annoying quasi-periodic pattern in some surfaces supposed to be flat, like the streets of Miami. However, a close inspection

reveals that such patterns were already present in the PAN image. λ -PNN, by forcing a strong correlation with the PAN, replicates such patterns which are made more visible by the injection of color. In general, the quality of λ -PNN images appears to be fully satisfactory and to meet the original goal of high joint spatial-spectral fidelity.

Let us now proceed to a comparative analysis, focusing on the products that appear to be the most competitive under one or more distortion metrics. First, we consider the methods that provide the least spectral distortion under the $D_{\lambda,align}^{(K)}$, R-ERGAS, and $D_{\lambda}^{(K)}$ metrics. For WV3 Adelaide these are TV, MF, and SR-D. They all show a significant loss of resolution with respect to λ -PNN and to the PAN, somewhat milder for SR-D which, however, exhibits spatial artifacts. On the other hand, for all of them, D_{ρ} is much larger than for λ -PNN signaling a limited correlation with the PAN. What is observed for Adelaide, repeats over and over as we go through the other test images, and always in agreement with the numerical results. In all cases, good spectral quality is obtained at the price of

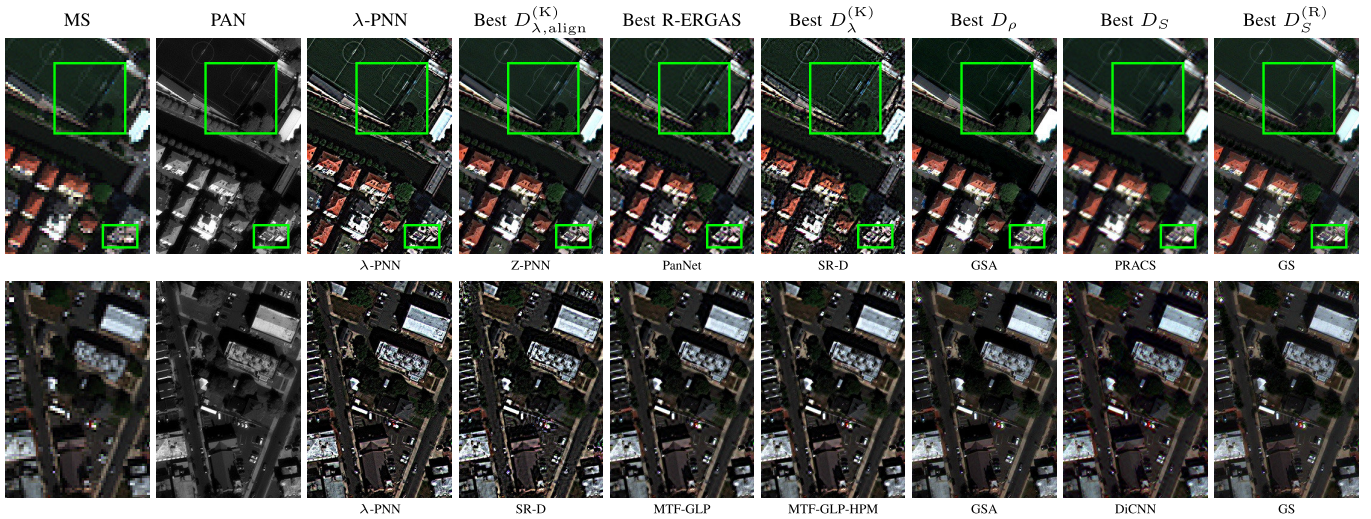


Fig. 10. Samples from GE1 (top) Genoa and (bottom) Trenton-PairMax. (Left to right) MS, PAN, λ -PNN, and best references.




Fig. 11. Loss curves observed during TA to a GE1 genoa image. All models pretrained on the GE1 training set.

reduced resolution or spatial artifacts, Therefore, it is fair to conclude that only the proposed λ -PNN seems able to consistently provide *both* good spectral and good spatial quality.

Similar conclusions can be drawn by looking at the best competitors under the spatial distortion metrics. In this case, we discard right away the last two columns because, as already noted, the D_S and $D_S^{(R)}$ indicators are not very reliable and often point to images with unsatisfactory spatial quality. Consider for example the PRACS image obtained for Genoa (Fig. 10), which has the best D_S but shows a dramatic loss of resolution. We focus instead on the best D_{ρ} images, which are always characterized by good spatial quality. Spectral quality is generally good for them, but deeper scrutiny reveals several problematic details. A few examples (marked again by green boxes): BT-H (Munich) is generally bluish, especially in high-contrast areas; Z-PNN (Washington and Miami) has rooftops with desaturated colors and a yellow rooftop (Miami) with a plain wrong hue; the football field of GSA (Genoa) is too dark. Again, λ -PNN appears to provide the fewest spatial and spectral artifacts.

G. Ablation Studies: Architecture

To gain further insight into the effectiveness of our design choices, we carry out some ablation studies, starting with

TABLE VII
 λ -PNN VERSIONS CONSIDERED IN THE ABLATION STUDY, FROM BASELINE, WITH JUST THREE CONVOLUTIONAL LAYERS, TO FULL FLEDGED

architecture	ConvLayers	ResBlocks	CBAM	R-CBAM
λ -PNN/a	✓	–	–	–
λ -PNN/b	✓	✓	–	–
λ -PNN/c	✓	✓	✓	–
full λ -PNN	✓	✓	–	✓

architecture. To this end, we compare four versions of the proposed model (see Table VII), starting with a basic CNN and gradually adding new layers until we get the full λ -PNN architecture of Fig. 5(a). The basic model comprises just three convolutional layers, then, in λ -PNN/b the two ResBlocks (red boxes are added, followed by the two attention layers, conventional (CBAM) in λ -PNN/c or residual (R-CBAM) in full λ -PNN.

In Fig. 11, we show the evolution of the three selected metrics during the TA phase. Models were all pretrained in the same conditions on the GE1 training set and then adapted to the target Genoa image (GE1 X-Val dataset). The initial point of each curve corresponds to the distortion observed with pretrained weights, which is then gradually




Fig. 12. Sample from WV3 X-Val (Adelaide). (Left to right) MS, PAN, and pansharpened by λ -PNN with various losses.

reduced with TA. The basic model (λ -PNN/a, blue curves), is clearly unable to exploit the potential of full-resolution training. Both spatial and spectral distortions are the largest and, what is worse, they do not decrease significantly with fine-tuning. We conjecture that with just three layers, the composite receptive field is too small to capture necessary dependencies. All deeper architectures perform much better. By including residual blocks (λ -PNN/b, orange) the spatial distortion D_ρ is more than halved, even with the pretrained weights. However, the spectral distortion remains relatively large and reduces somewhat only after intense fine-tuning. The further inclusion of attention mechanisms seems to unlock the potential of full-resolution training. With CBAM blocks (λ -PNN/c, green) and even more with R-CBAM (λ -PNN, red) a major boost in spectral quality is obtained. Eventually, λ -PNN reduces all distortion metrics with respect to the basic reference, by as much as 20% (R-ERGAS), 40% ($D_{\lambda, \text{align}}^{(K)}$), and 60% (D_ρ).

It is also worth underlining that TA plays a pivotal role in achieving such good results, since it keeps reducing the spectral distortion during the whole process and, apparently, could provide further gains with more epochs. All this is possible thanks to the proposed fast TA procedure. Otherwise, with large images and large models, only a few epochs would have been affordable.

H. Ablation Studies: Loss Function

We now turn to assess the impact of the proposed JESSE loss on the good performance of λ -PNN. Therefore, we now keep the architecture fixed and change only the loss function. For each of these losses, we train a new λ -PNN on the VW3 training set. Numerical results on the WV3 X-Val dataset (24 tiles from the Adelaide image) are reported in Table VIII, considering only the $D_\lambda^{(K)}$ and D_ρ metrics for brevity. In the upper part of the table, we report results for the JESSE loss itself (first row) and some ablated variants: without the $D_{\lambda, \text{align}}^{(K)}$ spectral (sub)term, without the R-ERGAS spectral (sub)term, with $D_S^{(R)}$ in place of D_ρ as the spatial term. In the lower part, instead, we show the results obtained when the JESSE loss is replaced altogether by one of the loss functions proposed for other pansharpening methods that rely on unsupervised full-resolution training, SSQ, GDD, PG, and Z-PNN. Together with the numerical results, we rely also on a visual example. In Fig. 12, for a small test crop we show, as usual, MS and PAN original data, followed by the images pansharpened by λ -PNN trained with each of the aforementioned losses. Note that the automatic alignment mechanism is deactivated in all cases to ensure a fair comparison.

TABLE VIII
AVERAGE RESULTS ON WV3 X-VAL (ADELAIDE)

Loss function			$D_\lambda^{(K)}$	D_ρ
1)		JESSE as is	0.041	0.040
2)	variants of JESSE	w/o $D_{\lambda, \text{align}}^{(K)}$	0.055	0.034
3)		w/o ERGAS	0.071	0.028
4)		$D_S^{(R)} \rightarrow D_\rho$	0.013	0.768
5)	from other papers	\mathcal{L}_{SSQ} [49]	0.065	0.268
6)		\mathcal{L}_{GDD} [51]	0.067	0.961
7)		\mathcal{L}_{PG} [52]	0.100	0.110
8)		$\mathcal{L}_{\text{Z-PNN}}$ [50]	0.045	0.044

By keeping the JESSE loss, but removing one of the spectral subterms (second and third rows), the balance between spectral and spatial fidelity changes somewhat in favor of the latter, but not dramatically so. However, the pansharpened samples show some spectral aberrations. This is especially marked when the ERGAS term is removed, with clear hue distortions (e.g., the yellow rooftop on the top-right). Instead, replacing the correlation-based metric, D_ρ , with the popular $D_S^{(R)}$ metric in the spatial loss term (row 4), has a catastrophic impact on image quality, as clear from both the numbers and the sample image. On the other hand, this is only a further confirmation of what was observed before.

Let us now turn to consider completely different loss functions, copied by recently published papers. We keep using the λ -PNN architecture and replace only spatial and spectral loss terms with those proposed in the references. With this study, we want to understand whether the good performance of λ -PNN is mainly due to its improved architecture or the loss function plays a fundamental role in it. Both the numbers reported in the bottom part of the table, (rows 5–8) and the pansharpened images provide unambiguous indications. With all other losses, except for Z-PNN, a significant quality impairment is observed, especially in the spatial resolution, as also confirmed by the D_ρ indicator, sometimes very high. With the Z-PNN loss, instead, the quality impairment is minimal both visually and according to the quality indicators (about 10% worse). On the other hand, this is the ancestor of the JESSE loss and exhibits already most of its good qualities.

I. Ablation Studies: Target Adaptation

In our experiments, we used always the proposed fast TA procedure. For the large 2048×2048 -pixel test images, based on preliminary experiments, we used 16 tiles of 256×256 pixels. In Fig. 13, we show one such image together with 16 tiles selected for fine-tuning, representative of the different contexts and spatial layouts observed in the scene. For the same image,




Fig. 13. Example of tile selection for fast target-adaptation.

TABLE IX

RESULTS OF SAMPLE EXPERIMENTS WITH CONVENTIONAL AND FAST TA ON A 2048×2048 TEST IMAGE

Target Adaptation	Time (s)	$D_{\lambda, \text{align}}^{(K)}$	D_{ρ}
fast ($256 \times 256, 16$)	37.6	0.0141	0.0381
conventional	824.6	0.0137	0.0396
no t.a.	0.0	0.0193	0.0503

tile size	# tiles	Time (s)	$D_{\lambda, \text{align}}^{(K)}$	D_{ρ}
512×512	4	68.6	0.0140	0.0386
	16	37.6	0.0141	0.0381
256×256	8	24.2	0.0144	0.0376
	4	18.0	0.0143	0.0382
128×128	64	31.9	0.0149	0.0354
	32	18.5	0.0145	0.0371
	16	12.1	0.0150	0.0361
	8	7.7	0.0148	0.0372
	4	5.7	0.0155	0.0379

in Table IX, we report numerical results observed with λ -PNN pansharpening. In the upper part of the table, we compare fast TA with conventional TA carried out on the whole image, and in the case where no TA is used. For each solution, we report the processing time and the spectral and spatial distortions, measured by $D_{\lambda, \text{align}}^{(K)}$ and D_{ρ} . It is clear that the fast procedure has little or no impact on quality indicators while sharply reducing the processing time, from almost 15 min to about 40 s, certainly viable for real-world applications. Of course, this extra time is still much larger than the 2.2 s inference time. However, the choice of not performing TA at all is not appealing, as this would significantly worsen both spectral and spatial indicators.

In the lower part of the table, we study how performance depends on the choice of parameters, showing a number of combinations of tile size (from 512×512 to 128×128 pixels) and number of tiles (from 64 to 4) going from more conservative to more aggressive. It appears that, barring the extreme case of just four small tiles, distortion indicators are not significantly affected by fast TA, with some impairments observed mainly on spectral quality. Eventually, we selected the parameters that ensure the least increase in spectral dis-

ortion, despite the larger processing time. In summary, with a judicious choice of the tuning dataset, TA becomes not only effective, ensuring good generalization, but also efficient enough to allow its application with any image and model.

V. CONCLUSION

We proposed a new DL-based pansharpening model, λ -PNN, with unsupervised training on full-resolution data. Besides architectural improvements, with spatial and spectral attention modules, a key asset of the proposal is the JESSE loss, which jointly promotes the spectral and spatial fidelity of the output images to the available reference data. Moreover, we proposed a fast TA procedure to ensure good generalization ability in all practical applications. To assess the proposed method we performed a large number of experiments on images acquired by various sensors, obtaining always excellent numerical results and convincing quality of output images.


We believe that full-resolution unsupervised training offers the best opportunity to unlock the full potential of DL in pansharpening. However, despite the good results obtained with the proposed method, there are still problems to be solved and room for improvement. An immediate goal is to characterize and ultimately eliminate the anomalous patterns observed in some cases. More ambitious and long-term objectives concern the extension of the method to some challenging cases: 1) the joint coregistration and pansharpening of images with moving objects, e.g., vehicles, which cause long-range local shifts and 2) the pansharpening of MS or hyperspectral bands weakly correlated with the PAN, for which a new suitable spatial loss term needs to be defined.

REFERENCES

- [1] G. Vivone et al., "A critical comparison among pansharpening algorithms," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 5, pp. 2565–2586, May 2015.
- [2] G. Vivone et al., "A new benchmark based on recent advances in multispectral pansharpening: Revisiting pansharpening with classical and emerging pansharpening methods," *IEEE Geosci. Remote Sens. Mag.*, vol. 9, no. 1, pp. 53–81, Mar. 2021.
- [3] V. K. Shettigara, "A generalized component substitution technique for spatial enhancement of multispectral images using a higher resolution data set," *Photogram. Eng. Remote Sens.*, vol. 58, no. 5, pp. 561–567, 1992.
- [4] C. Laben and B. Brower, "Process for enhancing the spatial resolution of multispectral imagery using pan-sharpening," U.S. Patent 6011 875, Jan. 4, 2000.
- [5] B. Aiazzi, S. Baronti, and M. Selva, "Improving component substitution pansharpening through multivariate regression of MS+Pan data," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 10, pp. 3230–3239, Oct. 2007.
- [6] A. Garzelli, F. Nencini, and L. Capobianco, "Optimal MMSE pan sharpening of very high resolution multispectral images," *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 1, pp. 228–236, Jan. 2008.
- [7] J. Choi, K. Yu, and Y. Kim, "A new adaptive component-substitution-based satellite image fusion by using partial replacement," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 1, pp. 295–309, Jan. 2011.
- [8] A. Garzelli, "Pansharpening of multispectral images based on nonlocal parameter optimization," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 4, pp. 2096–2107, Apr. 2015.
- [9] G. Vivone, "Robust band-dependent spatial-detail approaches for panchromatic sharpening," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 9, pp. 6421–6433, Sep. 2019.
- [10] T. Ranchin and L. Wald, "Fusion of high spatial and spectral resolution images: The ARSIS concept and its implementation," *Photogramm. Eng. Remote Sens.*, vol. 66, no. 1, pp. 49–61, Jan. 2000.

- [11] B. Aiazzi, L. Alparone, S. Baronti, and A. Garzelli, "Context-driven fusion of high spatial and spectral resolution images based on over-sampled multiresolution analysis," *IEEE Trans. Geosci. Remote Sens.*, vol. 40, no. 10, pp. 2300–2312, Jan. 2002.
- [12] X. Otazu, M. Gonzalez-Audicana, O. Fors, and J. Nunez, "Introduction of sensor spectral response into image fusion methods. Application to wavelet-based methods," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 10, pp. 2376–2385, Oct. 2005.
- [13] L. Alparone, A. Garzelli, and G. Vivone, "Intersensor statistical matching for pansharpening: Theoretical issues and practical solutions," *IEEE Trans. Geosci. Remote Sens.*, vol. 55, no. 8, pp. 4682–4695, Aug. 2017.
- [14] S. Lolli, L. Alparone, A. Garzelli, and G. Vivone, "Haze correction for contrast-based multispectral pansharpening," *IEEE Geosci. Remote Sens. Lett.*, vol. 14, no. 12, pp. 2255–2259, Dec. 2017.
- [15] G. Vivone, R. Restaino, and J. Chanussot, "A regression-based high-pass modulation pansharpening approach," *IEEE Trans. Geosci. Remote Sens.*, vol. 56, no. 2, pp. 984–996, Feb. 2018.
- [16] G. Vivone, R. Restaino, and J. Chanussot, "Full scale regression-based injection coefficients for panchromatic sharpening," *IEEE Trans. Image Process.*, vol. 27, no. 7, pp. 3418–3431, Jul. 2018.
- [17] G. Vivone et al., "Pansharpening based on semiblind deconvolution," *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 4, pp. 1997–2010, Apr. 2015.
- [18] M. R. Vicinanza, R. Restaino, G. Vivone, M. D. Mura, and J. Chanussot, "A pansharpening method based on the sparse representation of injected details," *IEEE Geosci. Remote Sens. Lett.*, vol. 12, no. 1, pp. 180–184, Jan. 2015.
- [19] F. Palsson, J. R. Sveinsson, and M. O. Ulfarsson, "A new pansharpening algorithm based on total variation," *IEEE Geosci. Remote Sens. Lett.*, vol. 11, no. 1, pp. 318–322, Jan. 2014.
- [20] G. Masi, D. Cozzolino, L. Verdoliva, and G. Scarpa, "Pansharpening by convolutional neural networks," *Remote Sens.*, vol. 8, no. 7, p. 594, Jul. 2016. [Online]. Available: <http://www.mdpi.com/2072-4292/8/7/594>
- [21] J. Yang, X. Fu, Y. Hu, Y. Huang, X. Ding, and J. Paisley, "PanNet: A deep network architecture for pan-sharpening," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 1753–1761.
- [22] Y. Wei, Q. Yuan, H. Shen, and L. Zhang, "Boosting the accuracy of multispectral image pansharpening by learning a deep residual network," *IEEE Geosci. Remote Sens. Lett.*, vol. 14, no. 10, pp. 1795–1799, Oct. 2017.
- [23] G. Masi, D. Cozzolino, L. Verdoliva, and G. Scarpa, "CNN-based pansharpening of multi-resolution remote-sensing images," in *Proc. Joint Urban Remote Sens. Event (JURSE)*, Dubai, United Arab Emirates, 2017, pp. 1–4, doi: [10.1109/JURSE.2017.7924534](https://doi.org/10.1109/JURSE.2017.7924534).
- [24] Q. Yuan, Y. Wei, X. Meng, H. Shen, and L. Zhang, "A multiscale and multidepth convolutional neural network for remote sensing imagery pan-sharpening," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 11, no. 3, pp. 978–989, Mar. 2018.
- [25] Y. Zhang, C. Liu, M. Sun, and Y. Ou, "Pan-sharpening using an efficient bidirectional pyramid network," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 8, pp. 5549–5563, Aug. 2019.
- [26] L. He et al., "Pansharpening via detail injection based convolutional neural networks," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 12, no. 4, pp. 1188–1204, Apr. 2019.
- [27] S. Vitale, G. Ferraioli, and G. Scarpa, "A CNN-based model for pansharpening of WorldView-3 images," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2018, pp. 5108–5111.
- [28] L.-J. Deng, G. Vivone, C. Jin, and J. Chanussot, "Detail injection-based deep convolutional neural networks for pansharpening," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 8, pp. 6995–7010, Aug. 2021.
- [29] P. S. Chavez Jr., and A. Y. Kwarteng, "Extracting spectral contrast in Landsat thematic mapper image data using selective principal component analysis," *Photogramm. Eng. Remote Sens.*, vol. 55, no. 3, pp. 339–348, 1989.
- [30] T.-M. Tu, P. S. Huang, C.-L. Hung, and C.-P. Chang, "A fast intensity–hue–saturation fusion technique with spectral adjustment for IKONOS imagery," *IEEE Geosci. Remote Sens. Lett.*, vol. 1, no. 4, pp. 309–312, Oct. 2004.
- [31] J. Nunez, X. Otazu, O. Fors, A. Prades, V. Pala, and R. Arbiol, "Multiresolution-based image fusion with additive wavelet decomposition," *IEEE Trans. Geosci. Remote Sens.*, vol. 37, no. 3, pp. 1204–1211, May 1999.
- [32] M. M. Khan, J. Chanussot, L. Condat, and A. Montanvert, "Indusion: Fusion of multispectral and panchromatic images using the induction scaling technique," *IEEE Geosci. Remote Sens. Lett.*, vol. 5, no. 1, pp. 98–102, Jan. 2008.
- [33] B. Aiazzi, L. Alparone, S. Baronti, A. Garzelli, and M. Selva, "An MTF-based spectral distortion minimizing model for pan-sharpening of very high resolution multispectral images of urban areas," in *Proc. 2nd GRSS/ISPRS Joint Workshop Remote Sens. Data Fusion Over Urban Areas*, May 2003, pp. 90–94.
- [34] B. Aiazzi, L. Alparone, S. Baronti, A. Garzelli, and M. Selva, "MTF-tailored multiscale fusion of high-resolution MS and pan imagery," *Photogramm. Eng. Remote Sens.*, vol. 72, no. 5, pp. 591–596, May 2006.
- [35] J. Lee and C. Lee, "Fast and efficient panchromatic sharpening," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 1, pp. 155–163, Jan. 2010.
- [36] R. Restaino, M. Dalla Mura, G. Vivone, and J. Chanussot, "Context-adaptive pansharpening based on image segmentation," *IEEE Trans. Geosci. Remote Sens.*, vol. 55, no. 2, pp. 753–766, Feb. 2017.
- [37] F. Palsson, M. O. Ulfarsson, and J. R. Sveinsson, "Model-based reduced-rank pansharpening," *IEEE Geosci. Remote Sens. Lett.*, vol. 17, no. 4, pp. 656–660, Apr. 2020.
- [38] Y. Wei and Q. Yuan, "Deep residual learning for remote sensed imagery pansharpening," in *Proc. Int. Workshop Remote Sens. With Intell. Process. (RSIP)*, May 2017, pp. 1–4.
- [39] Y. Rao, L. He, and J. Zhu, "A residual convolutional neural network for pan-sharpening," in *Proc. Int. Workshop Remote Sens. With Intell. Process. (RSIP)*, May 2017, pp. 1–4.
- [40] A. Azarang and H. Ghassemian, "A new pansharpening method using multi resolution analysis framework and deep neural networks," in *Proc. 3rd Int. Conf. Pattern Recognit. Image Anal. (IPRIA)*, Apr. 2017, pp. 1–6.
- [41] X. Liu, Y. Wang, and Q. Liu, "PSGAN: A generative adversarial network for remote sensing image pan-sharpening," in *Proc. 25th IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2018, pp. 873–877.
- [42] Z. Shao and J. Cai, "Remote sensing image fusion with deep convolutional neural network," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 11, no. 5, pp. 1656–1669, May 2018.
- [43] W. Dong, T. Zhang, J. Qu, S. Xiao, J. Liang, and Y. Li, "Laplacian pyramid dense network for hyperspectral pansharpening," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5507113.
- [44] W. Dong, S. Hou, S. Xiao, J. Qu, Q. Du, and Y. Li, "Generative dual-adversarial network with spectral fidelity and spatial enhancement for hyperspectral pansharpening," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 12, pp. 7303–7317, Dec. 2022.
- [45] M. Gong, H. Zhang, H. Xu, X. Tian, and J. Ma, "Multipatch progressive pansharpening with knowledge distillation," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5401115.
- [46] M. Gong, J. Ma, H. Xu, X. Tian, and X.-P. Zhang, "D2TNet: A ConvLSTM network with dual-direction transfer for pan-sharpening," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5409114.
- [47] L. Chen, G. Vivone, Z. Nie, J. Chanussot, and X. Yang, "Spatial data augmentation: Improving the generalization of neural networks for pansharpening," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5401711.
- [48] G. Scarpa, S. Vitale, and D. Cozzolino, "Target-adaptive CNN-based pansharpening," *IEEE Trans. Geosci. Remote Sens.*, vol. 56, no. 9, pp. 5443–5457, Sep. 2018.
- [49] S. Luo, S. Zhou, Y. Feng, and J. Xie, "Pansharpening via unsupervised convolutional neural networks," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 13, pp. 4295–4310, 2020.
- [50] M. Ciotola, S. Vitale, A. Mazza, G. Poggi, and G. Scarpa, "Pansharpening by convolutional neural networks in the full resolution framework," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5408717.
- [51] T. Uezato, D. Hong, N. Yokoya, and W. He, "Guided deep decoder: Unsupervised image pair fusion," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2020, pp. 87–102.
- [52] J. Ma, W. Yu, C. Chen, P. Liang, X. Guo, and J. Jiang, "PanGAN: An unsupervised pan-sharpening method for remote sensing image fusion," *Inf. Fusion*, vol. 62, pp. 110–120, Oct. 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1566253520302591>
- [53] S. Seo et al., "UPSNet: Unsupervised pan-sharpening network with registration learning between panchromatic and multispectral images," *IEEE Access*, vol. 8, pp. 201199–201217, 2020.
- [54] L.-J. Deng et al., "Machine learning in pansharpening: A benchmark, from shallow to deep networks," *IEEE Geosci. Remote Sens. Mag.*, vol. 10, no. 3, pp. 279–315, Sep. 2022.

- [55] C. Thomas, T. Ranchin, L. Wald, and J. Chanussot, "Synthesis of multispectral images to high spatial resolution: A critical review of fusion methods based on remote sensing physics," *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 5, pp. 1301–1312, May 2008.
- [56] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [57] Z. Wang and A. C. Bovik, "A universal image quality index," *IEEE Signal Process. Lett.*, vol. 9, no. 3, pp. 81–84, Mar. 2002.
- [58] L. Wald, "Data fusion: Definitions and architectures—fusion of images of different spatial resolutions," École des Mines de Paris, Paris, France, Tech. Rep., 2002.
- [59] A. Garzelli and F. Nencini, "Hypercomplex quality assessment of multi/hyperspectral images," *IEEE Geosci. Remote Sens. Lett.*, vol. 6, no. 4, pp. 662–665, Oct. 2009.
- [60] M. M. Khan, L. Alparone, and J. Chanussot, "Pansharpening quality assessment using the modulation transfer functions of instruments," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 11, pp. 3880–3891, Nov. 2009.
- [61] G. Scarpa and M. Ciotola, "Full-resolution quality assessment for pansharpening," *Remote Sens.*, vol. 14, no. 8, p. 1808, Apr. 2022. [Online]. Available: <https://www.mdpi.com/2072-4292/14/8/1808>
- [62] J. Lee, S. Seo, and M. Kim, "SIPSA-Net: Shift-invariant pan sharpening with moving object alignment for satellite imagery," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 10161–10169.
- [63] M. Ciotola and G. Scarpa, "Fast full-resolution target-adaptive CNN-based pansharpening framework," *Remote Sens.*, vol. 15, no. 2, p. 319, Jan. 2023. [Online]. Available: <https://www.mdpi.com/2072-4292/15/2/319>
- [64] A. Howard et al., "Searching for MobileNetV3," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 1314–1324.
- [65] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [66] D. Hendrycks and K. Gimpel, "Gaussian error linear units (GELUs)," 2016, *arXiv:1606.08415*.
- [67] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional block attention module," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 3–19.
- [68] G. Vivone, M. D. Mura, A. Garzelli, and F. Pacifici, "A benchmarking protocol for pansharpening: Dataset, preprocessing, and quality assessment," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 14, pp. 6102–6118, 2021.
- [69] R. Restaino, G. Vivone, M. D. Mura, and J. Chanussot, "Fusion of multispectral and panchromatic images based on morphological operators," *IEEE Trans. Image Process.*, vol. 25, no. 6, pp. 2882–2895, Jun. 2016.
- [70] L. Alparone, B. Aiuzzi, S. Baronti, A. Garzelli, F. Nencini, and M. Selva, "Multispectral and panchromatic data fusion assessment without reference," *Photogramm. Eng. Remote Sens.*, vol. 74, no. 2, pp. 193–200, Feb. 2008.
- [71] L. Alparone, A. Garzelli, and G. Vivone, "Spatial consistency for full-scale assessment of pansharpening," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2018, pp. 5132–5134.



Matteo Ciotola (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. degrees (cum laude) in automation engineering from the University of Naples Federico II, Naples, Italy, in 2018 and 2020, respectively, where he is currently pursuing the Ph.D. degree.

In 2020, he held a traineeship with the Centre de Coopération Internationale en Recherche Agronomique pour le Développement (CIRAD), Montpellier, France, where he conducted his research at the Maison de la Télédétection.

From 2022 to 2023, he was a Visiting Scholar with the Université Bretagne Sud, Vannes, France, where he fulfilled his research with the Environment Observation With Complex Imagery (OBELIX) Research Team, Naples. He is a member of the GRIP Research Team. His main research interests focus on data fusion of remotely sensed images, in particular super-resolution and pansharpening, through deep learning algorithms.



Giovanni Poggi (Member, IEEE) is a Full Professor of telecommunications at the University of Naples Federico II, Naples, Italy. His research interests are in statistical image processing, including compression, restoration, segmentation, classification, and fusion, with application to remote-sensing (both optical and SAR images) digital forensics, and biometry.

Prof. Poggi has been an Associate Editor of *IEEE TRANSACTIONS ON IMAGE PROCESSING* and *Signal Processing* (Elsevier).



Giuseppe Scarpa (Senior Member, IEEE) is an Associate Professor of Telecommunications at the Parthenope University of Naples. His research activity concerns image segmentation, texture modeling and classification, object detection, pansharpening, feature extraction, data fusion, SAR despeckling, image coregistration, and deep learning, with applications in the remote sensing domain.

Prof. Scarpa has been a Senior Area Editor of *IEEE SIGNAL PROCESSING LETTERS*. He has served as a Guest Editor for several special issues of the *Remote Sensing* (MDPI) journal.