

Permutation invariance and uncertainty in multitemporal image super-resolution

Original

Permutation invariance and uncertainty in multitemporal image super-resolution / Valsesia, D.; Magli, E.. - In: IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING. - ISSN 0196-2892. - STAMPA. - 60:(2022), pp. 1-12. [10.1109/TGRS.2021.3130673]

Availability:

This version is available at: 11583/2956129 since: 2022-02-22T12:01:45Z

Publisher:

Institute of Electrical and Electronics Engineers Inc.

Published

DOI:10.1109/TGRS.2021.3130673

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Permutation invariance and uncertainty in multitemporal image super-resolution

Diego Valsesia, *Member, IEEE* and Enrico Magli, *Fellow, IEEE*

Abstract—Recent advances have shown how deep neural networks can be extremely effective at super-resolving remote sensing imagery, starting from a multitemporal collection of low-resolution images. However, existing models have neglected the issue of temporal permutation, whereby the temporal ordering of the input images does not carry any relevant information for the super-resolution task and causes such models to be inefficient with the, often scarce, ground truth data that is available for training. Thus, models ought not to learn feature extractors that rely on temporal ordering. In this paper, we show how building a model that is fully invariant to temporal permutation significantly improves performance and data efficiency. Moreover, we study how to quantify the uncertainty of the super-resolved image so that the final user is informed on the local quality of the product. We show how uncertainty correlates with temporal variation in the series, and how quantifying it further improves model performance. Experiments on the Proba-V challenge dataset show significant improvements over the state of the art without the need for self-ensembling, as well as improved data efficiency, reaching the performance of the challenge winner with just 25% of the training data.

Keywords—multitemporal super-resolution, convolutional neural networks, self-attention, uncertainty estimation

I. INTRODUCTION

Mapping the Earth with high resolution imagery is critical for a wide range of applications including environmental monitoring, urban mapping, disaster assessment, military surveillance, and many more. At the same time, instruments onboard of satellites face constraints such as payload sizes, downlink bandwidth, etc. that can limit the spatial resolution of the images they acquire, or the temporal availability of high-resolution (HR) products. Super-resolution (SR) techniques address the problem of estimating HR images from one or more low-resolution (LR) images. The availability of multiple images of the same scene is particularly useful since small geometric displacements allow the images to carry complementary information, that, when suitably combined by means of SR methods, can significantly increase the spatial resolution. In the context of remote sensing, there are several ways to obtain multiple images of the same scene, e.g., they can be acquired by a spacecraft during multiple orbits, or by multiple satellites imaging the same scene at different times, or may be obtained at the same time with different sensors. The most challenging scenario is the multitemporal one, as the content of the scene may change due to a variety of reasons, such as change in illumination, occlusions due to clouds, human activity, etc.

Significant progress has recently been made on multitemporal image SR, also spurred by the Proba-V challenge by the European Space Agency [1], [2]. A curated dataset and a set of standard testing conditions allowed quick development of models tackling this difficult task and showed how deep learning models can effectively overcome difficulties such as unknown imaging model and temporal variation.

At the same time, we find that current models for multitemporal SR lack an important ingredient, which is the invariance to temporal ordering. Due to the unpredictability of change in the temporal series of LR images, no assumption can be made about patterns arising from a specific ordering of the images. In order to better understand this concept, let us compare it to a case where ordering matters, which is exploiting multiple spectral bands. In that case, each band has a specific physical meaning and ordering matters because there exist stable correlation patterns among bands. This is not the case when dealing with the temporal instead of spectral dimension, and any temporal permutation of the input LR images should always result in the same SR image. Therefore, we seek to build a model that is explicitly invariant to temporal permutation. Capturing this important prior allows to build a more robust and efficient model because it does not need to learn this property from the (possibly limited) data.

Moreover, in contrast with the LR images, for which established quality assessment criteria are typically available, the multitemporal SR process raises a problem related to the quality of the generated SR images, and, in particular, the degree of confidence to which a SR pixel has been reconstructed properly. If significant temporal variation is present, how can one trust the content of the SR scene to properly represent reality? We raise this issue for the first time in the context of deep multitemporal SR models and propose a technique that estimates an uncertainty value for each pixel in the SR image. We show that this uncertainty map correlates with temporal variations and with the true error signal. This uncertainty map can be made available to final users to judge the reliability of regions of the SR product.

Our main novel contributions in this paper can be thus summarized as:

- a new architecture for the multitemporal SR problem, called PIUnet (Permutation Invariance and Uncertainty network), which is invariant to temporal permutations and enables higher data efficiency, requiring smaller datasets for training;
- a method to estimate the aleatoric uncertainty of the SR image, consisting of an architectural design built in PIUnet and an ad-hoc training procedure;
- significant improvements over state-of-the-art on the

The authors are with Politecnico di Torino – Department of Electronics and Telecommunications, Italy. email: {name.surname}@polito.it.

Proba-V challenge dataset, in terms of both quality of the SR images and computational efficiency, since we do not require expensive temporal self-ensembles;

- a flexible model that can process an arbitrary number of input LR images in a stable manner and without ensembling or architecture redesigns.

II. RELATED WORK

Image super-resolution has received great attention and a broad literature is available. However, the majority of works focus on single-image super-resolution (SISR) and a comparatively smaller number address the more challenging multi-image setting, and even fewer consider realistic LR degradations and multi-temporal change.

SISR has been addressed by means of interpolation-based techniques, optimization-based methods [3]–[6] and, more recently, learning-based methods relying on deep neural networks [7]–[14]. While SISR is interesting because of the limited amount of available information to solve the HR reconstruction problem, MISR offers a unique set of challenges and requires solutions that go beyond simple extensions of SISR works.

The first work on MISR by Tsai and Huang [15] used a frequency-domain technique to combine multiple downsampled images with subpixel displacements. However, frequency-domain algorithms do not allow to easily incorporate prior knowledge about HR images, and thus several spatial-domain MISR techniques were proposed over the years, including nonuniform interpolation [16], iterative back-projection (IBP) [17], projection onto convex sets (POCS) [18], [19], sparse coding [20], [21], and other regularized methods [22]–[24]. In particular, IBP [17] enjoyed some success and works by improving the initial SR guess by back-projecting the difference between simulated LR images and actual LR images to the SR image, and iteratively attempting at inverting the forward imaging process. However, IBP is ultimately limited by the inability to deal with unknown or very difficult to model image degradation processes, as well as the difficulty in including image priors. Regularized methods generate the SR image by solving an optimization problem where a regularization cost can encode sophisticated image priors to improve performance. Among those, the bilateral total variation (BTV) method [22] exploits a combination of the total variation regularizer and the bilateral filter to create a robust edge-preserving prior.

Model-based techniques such as the aforementioned ones are limited by the ability to accurately describe the forward imaging system and by the handcrafted prior used to capture the properties of real images. On the other hand, learning-based techniques directly use the data to overcome these modeling challenges and learn system models and image priors from observations. Deep neural networks are the tool of choice for such methods and recent years have shown progress in using deep learning for MISR. In particular, in the context of video SR [25], [26], convolutional neural networks (CNNs) have been developed to simultaneously perform motion compensation and SR frame generation. For instance, the dynamic upsampling filters (DUF) method proposed in [27]

estimates input-dependent interpolation filters for each pixel in the frames. Recently, the Proba-V SR challenge [1], [2], issued by the European Space Agency, stimulated research for MISR approaches in the remote sensing context. The new dataset is particularly interesting for the development of new methods, as it allows to deal with realistic image degradation, registration problems, and robustness to temporal variation in the scenes, both among LR images and between LR and HR. The challenge winner DeepSUM [28] proposed a modular CNN composed of a SISR part, a module performing dynamic registration from the feature space and a fusion module based on 3D convolution. The architecture was further improved in DeepSUM++ [29] by using non-local operations in the form of graph-convolutional layers. The challenge runner-up HighResNet [30] proposed a recursive fusion strategy, also including a dynamic registration module. The current state-of-the-art is represented by the RAMS model [31], which exploits feature attention at multiple stages. Finally, it is worth mentioning the related field of super-resolution from burst photography where recent progress [32], [33] has been made on novel architectures. However, research in this field is not directly comparable to the remote sensing case, due to the major focus on motion compensation and specificities of consumer cameras, such as demosaicing from the color filter array.

III. PROPOSED METHOD

A. Overview

We propose to tackle multitemporal image super-resolution by means of a novel neural network design, called PIUnet, addressing Permutation Invariance and Uncertainty estimation. An overview is presented in Fig. 1. The network input is a set with an arbitrary number of LR images. This is in contrast with other existing techniques which are designed for a fixed number of images. The LR images are assumed to be roughly registered with each other by means of preprocessing. The key features of PIUnet can be found in its invariance to permutations of the LR inputs and in the estimation of both the SR image and a corresponding pixel-by-pixel uncertainty value. This is achieved by means of two parallel heads which project the feature space built by the backbone to the two sets of information. Notice that the SR head uses pixel shuffling [34] as upsampling technique and only estimates the residual from bilinear upsampling and averaging of the LR inputs, while the uncertainty head does not use this global skip connection. The next sections describe in detail how to achieve permutation invariance and how to estimate the uncertainty map.

B. Invariance to temporal permutation

In this section, we are going to discuss the main modules that allow the model to be invariant to permutations of the LR images in the temporal dimension. Before that, we need to introduce some terminology from group theory, and, in particular, the notions of equivariant and invariant functions.

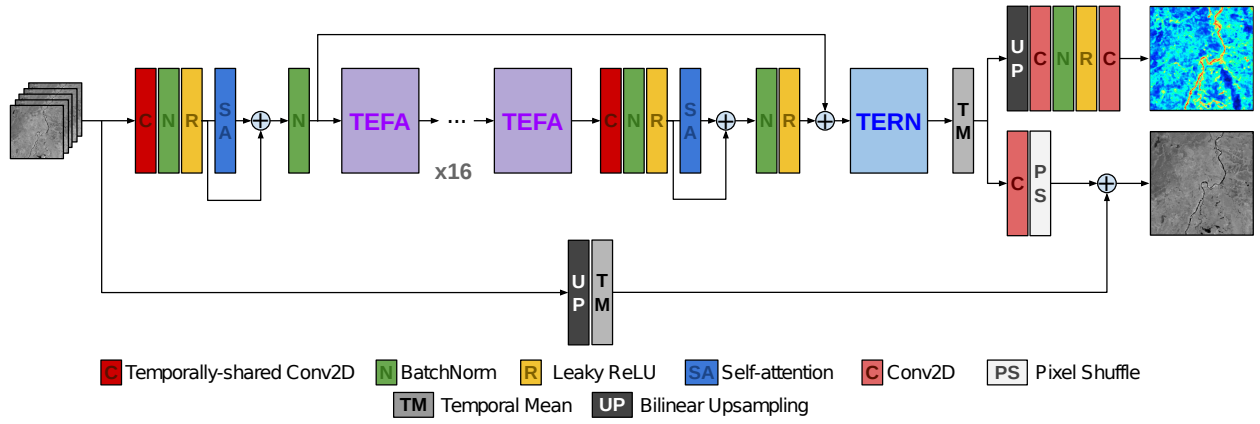


Fig. 1: PIUnet architecture. The model processes a stack of LR images and has two outputs, the top one being an uncertainty map and the bottom one the SR image.

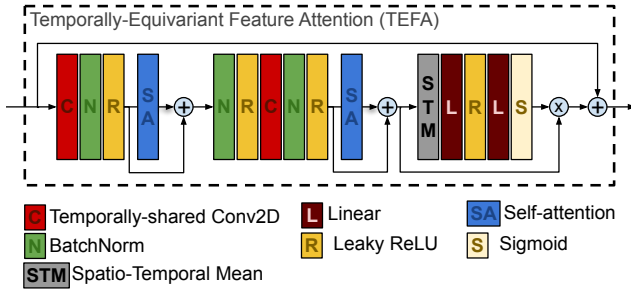


Fig. 2: Temporally-equivariant feature attention block.




Fig. 3: Temporally-equivariant RegNet.

Definition III.1 (Equivariance). A function $f : X \rightarrow Y$ is said to be equivariant to the actions g of a group \mathcal{G} if

$$f(g \circ x) = g \circ f(x) \quad \text{for all } x \in X, g \in \mathcal{G}. \quad (1)$$

Definition III.2 (Invariance). A function $f : X \rightarrow Y$ is said to be invariant to the actions g of a group \mathcal{G} if

$$f(g \circ x) = f(x) \quad \text{for all } x \in X, g \in \mathcal{G}. \quad (2)$$

In our case, we are dealing with the permutation group and its actions are all the possible temporal permutations of the input images. If we have an invariant function, the output will always be the same, no matter the permutation of the input, while for equivariant functions we will get an output that is exactly a permuted version of the output we would get without the input permutation.

Invariance is a powerful property to exploit because it embeds robustness into the model, and allows to capture prior knowledge about the data properties, that would otherwise be very difficult to learn from the data. In our problem, multitemporal invariance is particularly desirable, as the SR image does not depend on any specific temporal ordering, and it would not make sense to assume that there exist ordered temporal patterns that always recur on every image set. Indeed, prior work on multitemporal super-resolution [31] attempted at capturing temporal invariance, but by means of data augmentation, where multiple permutations are fed to the model at training time, rather than by means of model operations. This augmentation approach is, at best, able to learn a weak invariance, as the model has to fully learn from the data that permutations correspond to the same output, rather than explicitly building this knowledge in the model operations. As an example, the fact that the training of the current state-of-the-art RAMS model [31] does not capture invariance is testified by the significant improvements obtained by test-time ensembling, where predictions corresponding to multiple temporal permutations are averaged.

Building a model out of invariant layers can be challenging because only a limited number of functions that are invariant to permutations exist, and since they are typically simple functions (e.g., the average) they may not be sufficiently expressive to build complex features. A simple technique to build an invariant model is, therefore, to use a sequence of equivariant operations, which are easier to design, followed by a global invariant function. A trivial way to do that could be to independently process all the LR images with the same neural network and then combine the results with an invariant

function, such as the mean. However, this is highly suboptimal because it does not exploit the correlation between images to build better feature spaces in the hidden layers. At the same time, using 2D convolution treating time as feature channels or 3D convolution would not provide equivariant operations, since the weights used by these layers assume a temporal ordering and any permutation would return a different result.

Therefore, we propose to use self-attention [35] in the temporal dimension as a permutation-equivariant operation that, at the same time, allows to effectively combine the information from the multiple time instants, exploiting their cross-correlations. We remark that, despite the name, self-attention is a significantly different operation from classic feature attention [36]. Self-attention projects its input feature vector to three different subspaces, using three learnable matrices, to generate the so-called key, query, value vectors and the cross-correlation matrix between key and query is used as transformation matrix to weigh the temporal components in the value vectors and generate the output. More formally, given the representation of a pixel $\mathbf{X} \in \mathbb{R}^{T \times F}$, characterized by F features and T temporal channels, the self-attention operation computes:

$$\mathbf{Q} = \mathbf{X}\mathbf{W}_q, \quad \mathbf{K} = \mathbf{X}\mathbf{W}_k, \quad \mathbf{V} = \mathbf{X}\mathbf{W}_v \quad (3)$$

$$\mathbf{Y} = \text{softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{T}} \right) \cdot \mathbf{V} = \mathbf{A}\mathbf{V} \quad (4)$$

where the softmax function is applied row-wise, and \mathbf{W}_q , \mathbf{W}_k , \mathbf{W}_v are learnable matrices. Notice how this operation may look like a classic linear layer in the sense that the input is transformed by a matrix $\mathbf{A} \in \mathbb{R}^{T \times T}$, mixing the temporal channels, but the key difference is that \mathbf{A} is computed as a function of the input itself, rather than being constituted of trainable values. It is easy to check that a permutation of the T temporal channels, results in a permutation of the columns of \mathbf{A} , ensuring the overall equivariance of the operation. This operation is performed for all pixels of all the images in our batch. We remark that self-attention has a quadratic complexity $O(T^2)$ in terms of computation and memory due to the $\mathbf{Q}\mathbf{K}^\top$ cross-correlations, but since we apply it to the temporal dimension, the value of T is typically fairly small (e.g., in the Proba-V dataset, $T = 9$), ensuring efficient implementations.

This operation provides us a building block to be used whenever we want to mix the temporal channels. Referring to Fig.1, we typically use 2D convolutions shared across the temporal channels to extract spatial features, and then use self-attention to temporally combine those features. Notice how sharing the 2D convolutions across the temporal dimension is crucial to maintaining equivariance to temporal permutation. Based on this idea, we also design a novel module to compute classic residual feature attention [36], called Temporally-Equivariant Feature Attention (TEFA), shown in Fig. 2, whose repetition serves as the backbone of our model. The classic residual feature attention [36] extracts a scalar value that is used to weigh each feature map. Our proposed TEFA module extends the idea to deal with the extra temporal dimension and computes attention scores to weigh the feature channels by extracting spatial and temporal features in an equivariant way, by means of the aforementioned shared 2D convolutions

and temporal self-attention, and averaging them over space and time.

We also propose a temporally-equivariant extension of the RegNet module presented in DeepSUM [28], called TERN and shown in Fig. 3. The goal of the original RegNet was to dynamically compute small $K \times K$ spatial kernels from the input features to be used as filters over the input itself. They served as adaptive filters that could implement registration filters, or, in general, spatial interpolators that could refine the registration of the multitemporal images, exploiting the powerful feature space of the network. For more details, we refer the reader to [28]. The original formulation was not perfectly equivariant as it relied on an explicit ordering where the first temporal image was taken as a reference and concatenated to each of the others to be processed in pairs as channels in a convolutional layer. In TERN, we overcome this limitation by exploiting self-attention to cross-correlate features over the temporal dimension and infer the values of the spatial kernels. Just like RegNet, TERN computes a different $K \times K$ spatial filter for each temporal image, while the filter is shared across multiple feature maps.

Finally, when we consider the proposed architecture from its input to the output of the TERN module, we can notice that it is equivariant to temporal permutations. In order to make the overall model invariant, we simply average the output of TERN along the temporal axis.

We also remark that the proposed model does not have any constraint on the temporal dimension, i.e., the same model could be used for any number of multitemporal images, which could be especially useful if, once deployed, fewer images were available. This is not the case for other existing methods, including DeepSUM [28] and RAMS [31], which have architectures that have hard constraints on using exactly 9 multitemporal images, while only HighResNet [30] allows this flexibility thanks to their recursive fusion approach.

C. Uncertainty estimation

In this section, we propose a technique to assess a measure of confidence on the super-resolved image. This is important from the perspective of a final user of the SR product, whom we can inform about the degree of confidence a certain region of an image has to have been super-resolved properly.

We focus on characterizing the *aleatoric* uncertainty [37] on the SR image. This kind of uncertainty is due to stochastic perturbations in the input data, typically noise, or, in this case, also temporal variations of the scene. Characterizing aleatoric uncertainty allows us to determine whether a portion of the image was poorly super-resolved due to variability in the input LR images.

We use a heteroscedastic model for aleatoric uncertainty, essentially modeling each SR pixel as a random variable whose distribution can change on a pixel-by-pixel basis. The parameters of these distributions are then directly estimated by the neural network as its outputs. Training then uses the negative log likelihood (NLL) as loss function to be minimized. Most of the works on regression problems with neural networks [38] model the samples as Normal random variables, thus the

corresponding NLL can be seen as a generalization of the mean squared error (MSE) loss. However, it has been observed [39] that, in many image restoration problems, the L1 loss typically outperforms the MSE loss, even if the evaluation metric is the tightly related PSNR. This has also been observed in the context of multitemporal super-resolution [31] and could be explained by the robustness of the L1 metric to outliers. We therefore seek an extension of the L1 loss to train our network, and this leads to modeling the pixel distribution as a Laplacian:

$$p(x_i) = \frac{1}{2\beta_i} \exp\left(-\frac{|x_i - \mu_i|}{\beta_i}\right) \quad (5)$$

$$\mathbb{E}[x_i] = \mu_i, \quad \text{Var}[x_i] = 2\beta_i^2. \quad (6)$$

The goal of the neural network is to output μ_i , which will be the pixels in our SR image, and β_i which is proportional to the standard deviation and will therefore be our aleatoric uncertainty. This is done with two parallel heads, as shown in Fig.1, splitting after the temporal averaging operation. In practice, $\delta = \log \beta$ is estimated for numerical stability, resulting in the following loss function:

$$\begin{aligned} L &= -\frac{1}{NB} \sum_{b,i} \log p(x_i) \\ &= \frac{1}{NB} \sum_b \left[\sum_i \left(\delta_i^{(b)} + e^{-\delta_i^{(b)}} |x_i^{\text{HR}(b)} - \mu_i^{(b)}| \right) \right] \end{aligned} \quad (7)$$

for $i = 1, \dots, N$ pixels and $b = 1, \dots, B$ images.

It is often the case that the HR ground truth is not registered with the SR image. In that case, it is common practice [28], [30], [31] to use the minimum value of L computed for all possible registration shifts as loss function.

We remark that estimating the uncertainty and using Eq. (7) as loss function actually serves a dual purpose. Not only it provides information on reliability of the SR image, but it also improves model performance with respect to the L1 loss. In fact, the contribution of the variance serves as a regularizer against excessively confident predictions, leading to higher quality solutions, as shown in the experiments in Sec. IV-D.

IV. EXPERIMENTAL RESULTS AND DISCUSSIONS

In this section, we discuss the experimental performance of the proposed method by focusing on the Proba-V dataset from the corresponding SR challenge. We present comparisons with the state-of-the-art models, including the most recent ones developed after the challenge. We also analyze the properties of the proposed method in terms of label efficiency, the impact of the Laplacian NLL loss, and the information provided by the uncertainty map. Code is available online: <https://github.com/diegovalsesia/piunet>.

A. Experimental setting

Our experiments employ the Proba-V super-resolution dataset [1], released by the European Space Agency in the context of a challenge [2]. The unique feature of this dataset is the availability of both LR and HR captured by the same

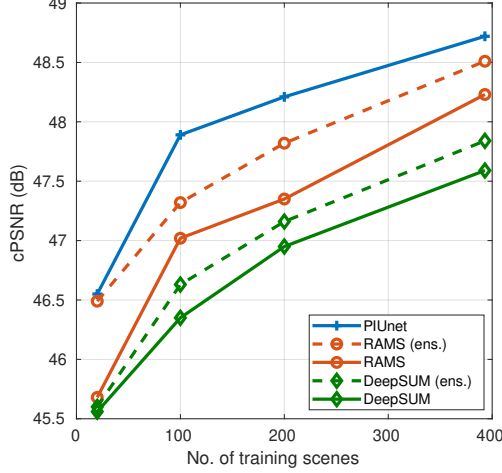
satellite. In fact, Proba-V is able to acquire LR images at 300m resolution daily, but also HR images at 100m resolution with a longer revisit time of 5 days. The availability of real acquisitions for both LR and HR images makes for an interesting case study for SR techniques, since it avoids synthetic degradation of the HR data to generate the LR ones, which often results in simplistic degradations and distorts model performance. This is also well-suited for supervised learning techniques, which are able to fully learn an inverse to the complex and unknown degradation mapping. However, the higher revisit time for HR data means that only a limited amount of images could be available, once cloud cover and limits to temporal change are taken into account, so it is important to study efficient models that can perform well with few data. The dataset provides single-band Level 2A (radiometrically and geometrically corrected Top-of-Atmosphere reflectance) images divided in a near-infrared (NIR) and visible (RED) categories. At least 9 LR images, acquired over the course of 30 days, are available for each HR scene. A total of 396 NIR scenes and 415 RED scenes are available for training, with additional 170 NIR and 176 RED scenes available for validation with known ground truth. The pixel size of the LR images is 128×128 , while the HR images are 384×384 .

We preprocess the data by selecting only the LR images having cloud coverage lower than 15% according to the provided clearance masks. We fix the number of used LR images per scene to 9 to match the standard setting used in other works. However, notice that the proposed model does not have any constraint on the size of the temporal dimension. We also normalize the images by subtracting the average intensity over the training set and dividing by the standard deviation. LR images are registered to each other by means of cross-correlation, assuming a translational model. Training uses LR patches of size 32×32 extracted from all spatial locations in the available images and augmented with rotations, while testing directly processes the full pixel size, since the model is fully-convolutional. Separate models are trained for NIR and RED.

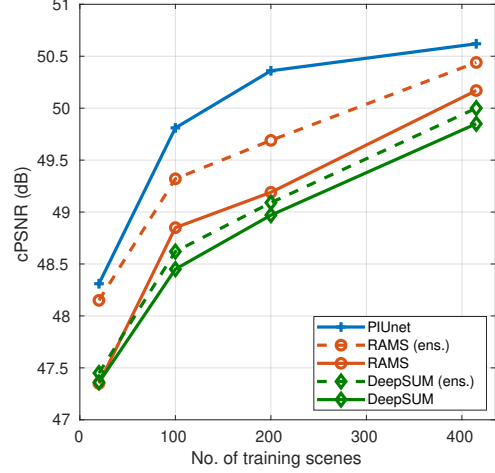
A feature size $F = 42$ is used throughout the model, except for the linear layers in TEFA which form a bottleneck reducing to $F = 5$ features before returning to 42. A total of 16 TEFA modules and 1 TERN module are used. The spatial kernels computed by TERN have size 5×5 . The total number of trainable parameters is slightly under a million, and it is comparable with existing works. Training minimizes the NLL loss (Eq. (7)) for approximately 500 epochs, using the Adam optimizer [40]. The NLL loss is made insensitive to misregistration between SR and HR and to absolute image brightness using the same approach as described in [28]:

TABLE I: Quantitative performance - cPSNR (dB) and cSSIM

	Bicubic	IBP [17]	BTV [22]	DUF [27]	HighResNet [30]	DeepSUM (ens.) [28]	DeepSUM++ (ens.) [29]	RAMS [31]	RAMS [31] (ens.)	PIUnet
NIR cPSNR	45.44	45.96	45.93	47.06	47.55	47.84	47.93	48.23	48.51	48.72
NIR cSSIM	0.9771	0.9778	0.9794	0.9842	0.9855	0.9858	0.9862	0.9875	0.9880	0.9883
RED cPSNR	47.34	48.21	48.12	49.36	49.75	50.00	50.08	50.17	50.44	50.62
RED cSSIM	0.9840	0.9865	0.9861	0.9842	0.9904	0.9908	0.9912	0.9913	0.9917	0.9921



(a) NIR



(b) RED

Fig. 4: Label efficiency. Performance on the validation set as function of the number of HR scenes available for training. Dashed lines represent temporal ensembles.




Fig. 5: NIR validation *imgset0792*. Left to right: HR, SR, SR uncertainty, ground truth absolute error.

$$L = \min_{u,v \in [0,6]} L_{u,v} \quad (8)$$

$$L_{u,v} = \frac{1}{NB} \sum_j \left[\sum_i \left(\delta_i^{(j)} + e^{-\delta_i^{(j)}} |x_i^{\text{HR}(u,v,j)} - \hat{\mu}_i^{(j)}| \right) \right] \quad (9)$$

$$\hat{\mu}_i^{(j)} = \mu_i^{(j)} + b^{(j)} \quad (10)$$

$$b^{(j)} = \frac{1}{\|m_i^{(j)}\|_1} \sum_j \left[x_i^{\text{HR}(u,v,j)} m_i^{(j)} - \mu_i^{(j)} m_i^{(j)} \right] \quad (11)$$

where $m_i^{(j)}$ is the clearance mask for image j , and u, v indicate the amount of horizontal and vertical shift applied to the HR image. The evaluation metric is the corrected PSNR, as used in earlier works, which accounts for shifts between the SR and HR images, and is insensitive to absolute brightness:

$$\text{cPSNR} = \max_{u,v \in [0,6]} 10 \log_{10} \frac{(2^{16} - 1)^2}{\text{MSE}_{u,v}} \quad (12)$$

$$\text{MSE}_{u,v} = \frac{\|\mathbf{x}^{\text{HR}(u,v)} \odot \mathbf{m} - (\mathbf{x}^{\text{SR}} \odot \mathbf{m} + \mathbf{b} \odot \mathbf{m})\|_2^2}{\|\mathbf{m}\|_1}, \quad (13)$$

with \odot denoting elementwise product.

The learning rate is 10^{-4} and it is reduced to $2 \cdot 10^{-5}$ for the final epochs. A batch size equal to 24 was used. Training required approximately 2 days on a Titan RTX GPU.

B. Comparison with state of the art

In this section, we compare PIUnet with a number of approaches, representing the state of the art for the multitemporal SR task. We include both model-based techniques and recent neural network methods to ensure a complete overview. The baseline (“Bicubic”) consists in bicubic upsampling, followed by pixel-domain registration and temporal averaging. For what concerns model-based approaches, we consider IBP [17] and BTV [22]. IBP takes as input an initial guess corresponding to our Bicubic baseline and the spatial shifts related to the LR images using phase correlation algorithm. BTV also uses the same initial guess for the bilateral total variation minimization problem. Regarding neural network approaches, DUF [27] is a technique from the video SR literature, using dynamic filters, a concept similar to what we use in TERN, but on a pixel-by-pixel basis; HighResNet [30] is the runner-up in the ESA challenge, and has a unique setting in which images in the LR set are recursively fused; DeepSUM [28] is the challenge winner and DeepSUM++ [29] an evolution with non-local operators based on graph convolution; RAMS [31] is the current state-of-the-art, heavily exploiting feature attention mechanisms.

Table I show quantitative results in terms of cPSNR and SSIM for the various methods. Notice that the table reports “(ens.)” for results from published material obtained with self-ensembling techniques. In particular, DeepSUM (ens.) used ensembles over five subsets of 9 images, while RAMS (ens.) averaged the results of 20 temporal permutations of the input images. Highlighting the use of self-ensembling is important as it increases computational complexity, requiring more processing time and/or memory, with respect to a single-shot model. The results show that PIUnet outperforms the other methods, and, interestingly, even the ensembled version of RAMS.

We are also interested in testing the label efficiency of the methods that require training, i.e., what performance can be achieved if the training set only has a limited number of image with HR ground truth. The results are reported in Fig. 4 which shows that PIUnet has significantly higher efficiency, by providing higher quality SR images even with constrained training sets. In particular, notice how just 100 scenes are required to reach approximately the same performance as DeepSUM (ens.).

Qualitative results are reported in Figs. 5, 6, 7. In particular, in Fig. 5, notice how the output uncertainty map estimated by the PIUnet neural network correlates with the absolute error between the SR and HR images. This point is further quantitatively discussed in Sec. IV-D. Figs. 6, 7 show NIR and RED scenes from the validation set and highlight increased sharpness and higher fidelity with respect to the HR image achieved by the proposed method. In particular, the absolute error with respect to the HR ground truth is reported in Figs. 8 and 9 for RAMS and PIUnet. It can be seen that the error of the

TABLE II: Impact of TERN (cPSNR)

	TERN	TEFA
NIR	48.72 dB	48.43 dB
RED	50.62 dB	50.55 dB

proposed method is generally lower on some high-frequency image features. Fig. 8 highlights the lower error achieved by PIUnet on the main river segment and on the terrain in the bottom part of the image. In Fig. 9, it is evident that PIUnet is able to more faithfully recover the river branch near the bottom of the image which exhibits high errors in RAMS.

C. Analysis and Ablation experiments

1) *Number of input images:* We remark that PIUnet is capable of processing an arbitrary number of input images without resorting to self-ensembling techniques, which is a major limitation of state-of-the-art approaches like DeepSUM [28] and RAMS [31]. Moreover, those approaches cannot process a set of images which is smaller than nine without modifications to the architectures, due to the assumptions made on sizes in unpadded convolutional layers. It can be expected that PIUnet reaches optimal performance close to the number of images used during training. However, we expect a graceful behaviour as function of T . This is shown in Fig. 10, where we can see that the best performance is achieved when the number of test images is close to (but not always at) the number of training images. The figure also shows that finetuning for a few iterations with the target number of images can improve performance, at the cost of degrading it for significantly different input sizes. Building a model with a broader optimality window so that finetuning is not required can be an interesting direction for future work.

2) *Impact of TERN module:* We evaluate the impact of TERN to the final performance of the model, and report it in Table II. This experiments replaces the TERN module with an extra TEFA module to keep the number of parameters approximately constant. We can notice that TERN provides significant performance improvements, especially on the NIR band.

3) *Computational complexity:* Table III reports some results on computational complexity of the proposed method with respect to state-of-the-art. We can see that PIUnet is significantly faster and with lower memory requirements than DeepSUM. The baseline version of RAMS is slightly faster and with comparable memory requirements. However, as shown in the previous section, we are able to outperform the temporally ensembled version of RAMS, which has a significant penalty in terms of complexity, as time or memory roughly scale linearly with the number of permutations in the ensemble. In particular, we report results for two ways of running the ensemble, i.e., serially, thus trading time for memory, or in parallel, trading memory for time.

D. Uncertainty estimation

In this experiment, we first show that the two-headed model using the NLL loss improves performance over the classic L1




Fig. 6: NIR validation *imgset0792* detail. Left: HR image. Right from top left to bottom right: one among LR images, bicubic (47.71 dB/0.9874), IBP (48.46 dB/0.9892), DeepSUM (50.82 dB/0.9933), RAMS (51.24 dB/0.9939), PIUnet (51.81 dB/0.9946).




Fig. 7: RED validation *imgset0353* detail. Left: HR image. Right from top left to bottom right: one among LR images, bicubic (47.34 dB/0.9882), IBP (47.57 dB/0.9889), DeepSUM (48.78 dB/0.9913), RAMS (50.22 dB/0.9938), PIUnet (51.78 dB/0.9946).

TABLE III: Computational complexity

	runtime	memory
DeepSUM	484 ms	5420 MB
RAMS	102 ms	1250 MB
RAMS (ens.)	1642 ms / 1075 ms	1250 MB / 5340 MB
PIUnet	181 ms	1280 MB

TABLE IV: Training loss comparison (cPSNR)

	L1 loss	NLL loss
NIR	48.41 dB	48.72 dB
RED	50.53 dB	50.62 dB

loss used by models that only estimate the SR image. Table IV shows the results obtained with the two losses for the same network architecture. We can see that the information on the variance of the super-resolved pixel helps convergence by regularizing high-confidence predictions and results in superior

performance. This would already be a significant reason to use such loss, but an additional benefit is the availability of the uncertainty values on a pixel-by-pixel basis for the SR product.

We can show how this uncertainty map can be used as guide to predict the unreliability of the SR image in areas with significant temporal variation. Fig. 11 shows an example




Fig. 8: NIR validation *imgset0792* absolute error. Left: RAMS error. Right: PIUnet error. Notice the lower error on the main river and the terrain in the bottom part of the image.




Fig. 9: RED validation *imgset0353* absolute error. Left: RAMS error. Right: PIUnet error. Notice how the river branch at the bottom is more faithfully recovered by PIUnet.

of this concept. We focus on a crop of a scene displaying some temporal variation in the river bed and surrounding areas. Notice how the variance of the SR pixels is higher in the corresponding areas, highlighting to the user that the SR image might be less reliable due to the high content variability in the available images.

A quantitative way to determine whether uncertainty is correctly estimated is to check how it correlates with the true error map. This can be done by means of sparsification plots [41]. These plots are generated by first sorting all the pixels by decreasing uncertainty and then removing a progressively larger fraction of those with high uncertainty. If the uncertainty estimate correlates with the error signal, a quality metric measured on the remaining pixels should show improvements as more pixels are removed. Viceversa, if the uncertainty values were random, the curve would be a constant, as it does not provide informative ways of removing pixels to improve quality. Fig. 12 shows the sparsification curve obtained for the proposed method, and compared with random selection and with an oracle that uses the true error signal to sort the pixels (for all variants, the brightness bias and shifts are computed before sparsification and kept fixed). The curves indeed confirm that the uncertainty map carries meaningful information. Future techniques addressing uncertainty estimation can compare sparsification curves to show improvements

on the task.

E. Experiments on Sentinel-2

The Proba-V dataset is an excellent case study that models complex degradations between LR and HR data. However, its limited spatial resolution might not allow to fully understand model performance at resolutions more commonly used by modern satellites. In this section, we train and test PIUnet and state-of-the-art methods on the Sentinel-2 images from the multitemporal dataset assembled for the “Enhanced Sentinel 2 Agriculture” challenge [42]. This dataset is focused on super-resolved segmentation where multiple $10m$ images are used to estimate $2.5m$ segmentation maps. Since segmentation is beyond the scope of this paper, we re-purpose the data to synthetically create a $10m$ HR ground truth and $40m$ multitemporal LR images via bicubic downsampling. Notice we choose a scaling factor of 4 to also test generalization beyond the factor of 3 used on the Proba-V data. The HR image is chosen as a cloudless image in the temporal series, disjoint from the images used in the LR series. In order to keep the single-band framework used in Proba-V we only use Band 3 (green). How to optimally exploit spectral correlation in presence of multispectral imagers is an interesting subject for future work. Testing is done on 5 out of the 100 scenes available. Training follows the same protocol as the one in Sec. IV-A, except for the lower number of epochs, set to 100 due to the smaller dataset. It must be remarked that synthetic data generation via downsampling can be a limiting factor, as it has simplistic assumptions on the degradation kernel and training on such data often results in sub-optimal real-world performance [43].

Table V presents the performance achieved PIUnet and state-of-the-art models on this Sentinel-2 dataset. The models are tested for various numbers of multitemporal images to show performance scaling when more images are available. The T images are taken as the ones with least cloud coverage among those available. It can be noticed that PIUnet is still highly effective despite the significantly higher resolution with respect to Proba-V setting, and significantly improves over state-of-the-art.

Fig. 13 shows a qualitative comparison on a test image from this dataset. It can be noticed that details of the scene are recovered in a significantly sharper way by PIUnet.

Finally, we show how transfer learning from a model pre-trained on the Proba-V data is beneficial, as it allows to achieve better performance with respect to a random initialization. Table ?? shows the results on the Sentinel 2 data for a scaling factor of 3 and $T = 9$ multitemporal images. This is expected since the extraction of low-level features is similar between the two datasets, despite the differences in resolution, so finetuning the pretrained model uses the Sentinel 2 data more efficiently.

V. CONCLUSION

We showed the importance of invariance to temporal permutation when building deep models for multitemporal image super-resolution. Building equivariant layers by means of self-attention and using a globally invariant operation close to the

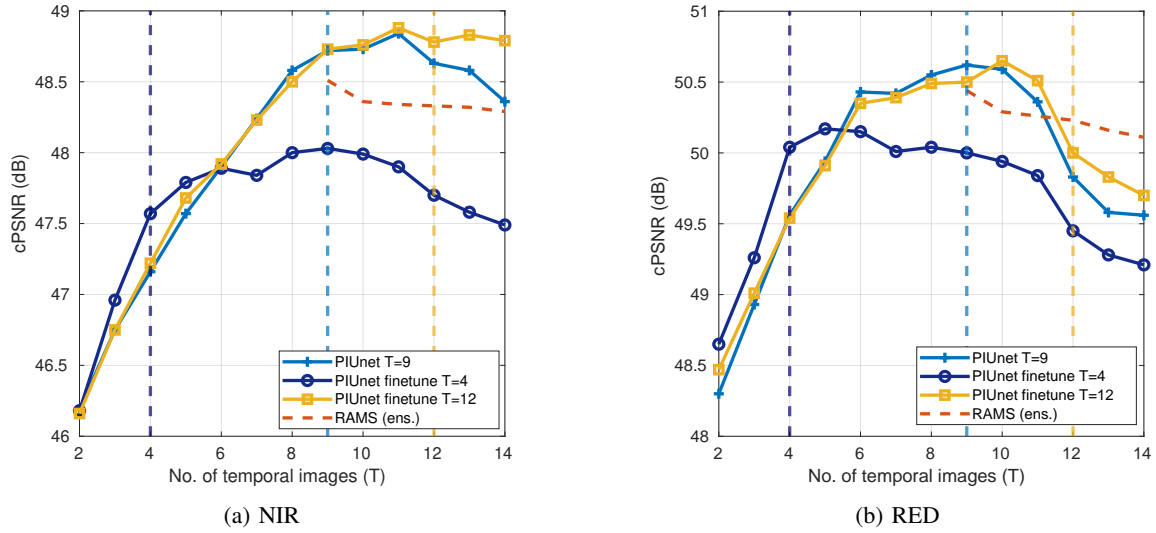


Fig. 10: Performance on the validation set as function of number of input LR images. Notice how the trained model gracefully handles any input size, but optimal performance may require finetuning.




Fig. 11: Temporal variation is captured by the SR uncertainty map. NIR validation *imgset0975*. From left to right: four of the LR images, SR image, SR uncertainty map, HR image.

TABLE V: SR performance on Sentinel 2 data (x4 SR) - cPSNR (dB) and cSSIM

		Bicubic+Mean	RAMS (ens.)	PIUnet
$T = 5$	cPSNR	38.95	40.15	41.01
	cSSIM	0.8781	0.9192	0.9243
$T = 10$	cPSNR	39.16	42.02	43.94
	cSSIM	0.8866	0.9512	0.9689
$T = 15$	cPSNR	40.22	42.51	44.87
	cSSIM	0.9187	0.9641	0.9754

TABLE VI: Transfer learning performance on Sentinel 2 data (x3 SR) - cPSNR (dB) and cSSIM

	Random init.	Proba-V pretrain
cPSNR	46.78	46.91
cSSIM	0.9803	0.9812

output allows to significantly improve performance over this challenging task and removes the computationally-expensive self-ensembling operation. We also showed how to estimate

the uncertainty of the SR image, so that the final user can be guided in trusting the generated product, and that this further improves model performance.

REFERENCES

- [1] M. Märtens, D. Izzo, A. Krzic, and D. Cox, “Super-resolution of proba-v images using convolutional neural networks,” *Astrodynamics*, vol. 3, no. 4, pp. 387–402, 2019.
- [2] Kelvins - ESA’s Advanced Concepts, “PROBA-V Super Resolution,” <https://kelvins.esa.int/proba-v-super-resolution>.
- [3] M. K. Ng, H. Shen, E. Y. Lam, and L. Zhang, “A total variation regularization based super-resolution reconstruction algorithm for digital video,” *EURASIP Journal on Advances in Signal Processing (JASP)*, vol. 2007, no. 1, p. 074585, 2007.
- [4] L. Wang, S. Xiang, G. Meng, H. Wu, and C. Pan, “Edge-directed single-image super-resolution via adaptive gradient magnitude self-interpolation,” *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, vol. 23, no. 8, pp. 1289–1299, Aug 2013.
- [5] K. Zhang, X. Gao, D. Tao, and X. Li, “Single image super-resolution with non-local means and steering kernel regression,” *IEEE Transactions on Image Processing (TIP)*, vol. 21, no. 11, pp. 4544–4556, Nov 2012.




Fig. 12: Sparsification curves for the proposed method against random and oracle. The improvement over random shows that the output uncertainty map actually correlates with the error signal and provides valuable information on the local quality of the SR product.




Fig. 13: Test on Sentinel 2 data (*eopatch-924*). From left to right: HR image, bicubic+mean (45.61 dB/0.9641), RAMS (52.43 dB/0.9966), PIUnet (53.15 dB/0.9972).

- [6] K. Zhang, X. Gao, D. Tao, and X. Li, "Single image super-resolution with multiscale similarity learning," *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*, vol. 24, pp. 1648–1659, 2013.
- [7] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE Transactions on Image Processing (TIP)*, vol. PP, 08 2016.
- [8] D. Liu, B. Wen, Y. Fan, C. C. Loy, and T. S. Huang, "Non-local recurrent network for image restoration," in *Advances in Neural Information Processing Systems*, 2018, pp. 1673–1682.
- [9] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *European Conference on Computer Vision (ECCV)*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham: Springer International Publishing, 2014, pp. 184–199.
- [10] J. Kim, J. K. Lee, and K. M. Lee, "Accurate image super-resolution using very deep convolutional networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 1646–1654.
- [11] —, "Deeply-recursive convolutional network for image super-resolution," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 1637–1645.
- [12] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1874–1883, 2016.
- [13] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, "Enhanced deep residual networks for single image super-resolution," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1132–1140, 2017.
- [14] Y. Zhang, Y. Tian, Y. Kong, B. Zhong, and Y. Fu, "Residual dense network for image super-resolution," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2472–2481, 2018.
- [15] R. Tsai and T. Huang, "Multiframe image restoration and registration," in *Advances in Computer Vision and Image Processing*, vol. 1. JAI Press: Greenwich, CT, USA, 1984, pp. 317–339.
- [16] S. Lertrattanapanich and N. K. Bose, "High resolution image formation from low resolution frames using delaunay triangulation," *IEEE Trans-*

- actions on Image Processing (TIP), vol. 11, no. 12, pp. 1427–1441, Dec 2002.
- [17] M. Irani and S. Peleg, “Improving resolution by image registration,” *Graphical Models and Image Processing (CVGIP)*, vol. 53, no. 3, pp. 231–239, 1991.
- [18] H. Stark and P. Oskoui, “High-resolution image recovery from image-plane arrays, using convex projections,” *Journal of the Optical Society of America*, vol. 6, no. 11, pp. 1715–1726, Nov 1989.
- [19] A. J. Patti, M. Ibrahim Sezan, and A. Murat Tekalp, “High-resolution image reconstruction from a low-resolution image sequence in the presence of time-varying motion blur,” in *Proceedings of 1st International Conference on Image Processing (ICIP)*, vol. 1, Nov 1994, pp. 343–347 vol.1.
- [20] T. Kato, H. Hino, and N. Murata, “Double sparsity for multi-frame super resolution,” *Neurocomputing*, vol. 240, no. C, pp. 115–126, May 2017.
- [21] K. Toshiyuki, H. Hideitsu, and M. Noboru, “Multi-frame image super resolution based on sparse coding,” *Neural Networks*, vol. 66, pp. 64–78, 2015.
- [22] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar, “Fast and robust multiframe super resolution,” *IEEE Transactions on Image Processing (TIP)*, vol. 13, no. 10, pp. 1327–1344, Oct 2004.
- [23] H. Takeda, S. Farsiu, and P. Milanfar, “Kernel regression for image processing and reconstruction,” *IEEE Transactions on Image Processing (TIP)*, vol. 16, no. 2, pp. 349–366, Feb 2007.
- [24] H. Shen, M. K. Ng, P. Li, and L. Zhang, “Super-resolution reconstruction algorithm to modis remote sensing images,” *The Computer Journal*, vol. 52, pp. 90–100, 01 2009.
- [25] A. Kappeler, S. Yoo, Q. Dai, and A. K. Katsaggelos, “Video super-resolution with convolutional neural networks,” *IEEE Transactions on Computational Imaging (TCI)*, vol. 2, no. 2, pp. 109–122, June 2016.
- [26] J. Caballero, C. Ledig, A. P. Aitken, A. Acosta, J. Totz, Z. Wang, and W. Shi, “Real-time video super-resolution with spatio-temporal networks and motion compensation,” *CoRR*, vol. abs/1611.05250, 2016.
- [27] Y. Jo, S. Wug Oh, J. Kang, and S. Joo Kim, “Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [28] A. B. Molini, D. Valsesia, G. Fracastoro, and E. Magli, “Deepsum: Deep neural network for super-resolution of unregistered multitemporal images,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 5, pp. 3644–3656, 2019.
- [29] —, “Deepsum++: Non-local deep neural network for super-resolution of unregistered multitemporal images,” *arXiv preprint arXiv:2001.06342*, 2020.
- [30] M. R. Arefin, V. Michalski, P. St-Charles, A. Kalaitzis, S. Kim, S. E. Kahou, and Y. Bengio, “Multi-image super-resolution for remote sensing using deep recurrent networks,” in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2020, Seattle, WA, USA, June 14-19, 2020*. IEEE, 2020, pp. 816–825.
- [31] F. Salvetti, V. Mazzia, A. Khaliq, and M. Chiaberge, “Multi-image super resolution of remotely sensed images using residual attention deep neural networks,” *Remote Sensing*, vol. 12, no. 14, 2020. [Online]. Available: <https://www.mdpi.com/2072-4292/12/14/2207>
- [32] G. Bhat, M. Danelljan, L. Van Gool, and R. Timofte, “Deep burst super-resolution,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9209–9218.
- [33] Z. Luo, L. Yu, X. Mo, Y. Li, L. Jia, H. Fan, J. Sun, and S. Liu, “Ebsr: Feature enhanced burst super-resolution with deformable alignment,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2021, pp. 471–478.
- [34] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, “Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1874–1883.
- [35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 6000–6010.
- [36] Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong, and Y. Fu, “Image super-resolution using very deep residual channel attention networks,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 286–301.
- [37] A. Kendall and Y. Gal, “What uncertainties do we need in bayesian deep learning for computer vision?” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 5580–5590.
- [38] A. Kendall, Y. Gal, and R. Cipolla, “Multi-task learning using uncertainty to weigh losses for scene geometry and semantics,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7482–7491.
- [39] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, “Loss functions for image restoration with neural networks,” *IEEE Transactions on Computational Imaging*, vol. 3, no. 1, pp. 47–57, 2017.
- [40] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [41] O. Mac Aodha, A. Humayun, M. Pollefeys, and G. J. Brostow, “Learning a confidence measure for optical flow,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 5, pp. 1107–1120, 2012.
- [42] “AI4EO Enhanced Sentinel 2 Agriculture,” <https://platform.ai4eo.eu/enhanced-sentinel2-agriculture>, accessed: 2010-10-13.
- [43] S. Bell-Kligler, A. Shocher, and M. Irani, “Blind super-resolution kernel estimation using an internal-gan,” in *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019, pp. 284–293.