This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Self-Supervised Super-Resolution for Multi-Exposure Push-Frame Satellites

Ngoc Long Nguyen¹ Jérémy Anger^{1,2} Axel Davy¹ Pablo Arias¹ Gabriele Facciolo¹ ¹ Université Paris-Saclay, CNRS, ENS Paris-Saclay, Centre Borelli, France https://centreborelli.github.io/HDR-DSP-SR/

Abstract

Modern Earth observation satellites capture multiexposure bursts of push-frame images that can be superresolved via computational means. In this work, we propose a super-resolution method for such multi-exposure sequences, a problem that has received very little attention in the literature. The proposed method can handle the signal-dependent noise in the inputs, process sequences of any length, and be robust to inaccuracies in the exposure times. Furthermore, it can be trained end-to-end with selfsupervision, without requiring ground truth high resolution frames, which makes it especially suited to handle real data. Central to our method are three key contributions: i) a base-detail decomposition for handling errors in the exposure times, ii) a noise-level-aware feature encoding for improved fusion of frames with varying signal-to-noise ratio and iii) a permutation invariant fusion strategy by temporal pooling operators. We evaluate the proposed method on synthetic and real data and show that it outperforms by a significant margin existing single-exposure approaches that we adapted to the multi-exposure case.

1. Introduction

High resolution (HR) satellite imagery is a key element in a broad range of tasks, including human activity monitoring and disaster relief. Super-resolution by computational methods has recently been adopted [7, 41] by the remote sensing industry (Planet SkySat, Satellogic Aleph-1). By leveraging high frame rate low-resolution (LR) acquisitions, low-cost constellations can be effective competitors to more traditional high-cost satellites. In order to capture the full dynamic range of the scene, some satellites use exposure bracketing, resulting in sequences with varying exposures. While several works have addressed multi-image super-resolution (MISR) of single-exposure sequences, almost no previous work considers the multi-exposure case.

MISR techniques exploit the aliasing in several LR acquisitions to reconstruct a HR image. The maximum attainable resolution is capped by the spectral decay of the blur kernel resulting from the sensor's pixel integration and the camera optics, which imposes a frequency cutoff beyond which there is no usable high frequency information.

Aggregating many frames is also interesting as it allows significant noise reduction. If the LR frames are acquired with bracketed exposures, it is possible to integrate them in a super-resolved high dynamic range (HDR) image. Long exposures have higher signal-to-noise ratio (SNR) which helps reduce the noise in dark regions, whereas short exposures provide information in bright regions which can cause saturation with longer exposure times.

In this work, our goal is to perform joint super-resolution and denoising from a time series of bracketed satellite images. We focus on push-frame satellite sensors such as the SkySat constellation from Planet. We increase the resolution by a factor of two, which is the frequency cutoff of the combined optical and sensor's imaging system. The SkySat satellites [41] contain a full-frame sensor capable of capturing bursts of overlapping frames: a given point on the ground is seen in several consecutive images. However, our technique is general and can be applied to other satellites, or beyond satellite imagery to consumer cameras capable of multi-exposure burst or video acquisition.

Several methods have addressed either MISR or HDR imaging from multiple exposures, but their combination has received little attention. Existing works consider an ideal setup in which frames can be aligned with an affinity [7,53] or a homography [55], and the number of acquisitions is large enough to render the problem an overdetermined system of equations. Such motion models are good approximations for satellite bursts, but ignore parallax [8], which can be noticeable for mountains and tall buildings.

In the case of satellite imaging, push-frame cameras capable of capturing multi-exposure bursts are relatively recent, which explains why all previous works on MISR focus on the single-exposure case [7, 17, 40, 43], except for SkySat's proprietary method [41] producing the L1B product, whose details are not public. Deep learning methods currently outperform traditional model-based approaches [47]. In general, learning-based methods require large realistic datasets with ground truth to be trained, as



Figure 1. Super-resolution from a real multi-exposure sequence of 10 SkySat images. Top row: Original low resolution images with different exposures. Bottom row: Reconstructions from five methods, including ours trained with self-supervision (right).

methods trained on synthetic data [4] fail to generalize to real images [14]. One of such datasets is the PROBA-V dataset [37], acquired with a satellite equipped with two cameras of different resolutions. This dataset has fostered the publication of several deep learning approaches to satellite MISR [9,17,39]. However, the PROBA-V dataset is not appropriate for MISR of LR image bursts acquired at a high frame rate, as the PROBA-V sequences are multi-date and present significant content and illumination changes.

A promising direction is to use self-supervised learning techniques, which have been applied to video restoration tasks such as denoising and demosaicing [18–20, 51, 58], and recently to MISR [43]. These techniques benefit from the temporal redundancy in videos. Instead of using ground truth labels, one of the degraded frames in the input sequence is withheld from the network and used as label.

Our work builds upon *Deep Shift-and-Add* (DSA) [43], a self-supervised deep learning method for MISR of singleexposure bursts of satellite images. The model is trained without supervision by exploiting the frame redundancy.

Contributions. In this work, we propose *High Dynamic Range Deep Shift-and-Pool*, HDR-DSP a self-supervised method for joint super-resolution and denoising of bracketed satellite imagery. The method is able to handle timeseries with a variable number of frames and is robust to errors in the exposure times, as the ones provided in the metadata are often inaccurate. This makes our method directly applicable to real image data (see Figure 1). This is, to the best of our knowledge, the first multi-exposure MISR method for satellite imaging, and beyond satellite imagery, it is the first approach based on deep-learning.

Our contributions are the following:

Feature Shift-and-Pool. We propose a *shift-and-pool* module that merges features (computed by an encoder network on each input LR frame) into a HR feature map by temporal pooling using permutation invariant statistics: average, maximum, and standard deviation. This gives a rich fused representation which yields a substantial improvement over the average [43], in both single and multiple exposure cases. *Robustness to inaccurate exposure times via base-detail decomposition.* We propose normalizing the input frames and decomposing them into base and detail. The errors caused by the inaccurate exposure times affect mainly the base, whereas the detail containing the aliasing required for super-resolution can be safely processed by the network. Note that vignetting and stray light can also cause exposure issues that affect single and multi-exposure MISR alike.

Noise-level-aware detail encodings. The noise present in the LR images is signal-dependent, its variance being an affine function of the intensity. To deal with such noise, we provide the un-normalized LR images to the encoder in addition to the normalized detail components. This gives the encoder information about the noise level of each pixel, necessary for an optimal fusion.

Self-supervised loss with grid shifting. Using random shifts of the high-resolution grid, we make the self-supervised loss of [43] translation equivariant, leading to improved results.

We validate our contributions with an ablation study on a synthetic dataset (§5.2), designed to model the main characteristics of real bracketed SkySat sequences. Since there are no previous works on multi-exposure MISR, we compare against state-of-the-art single-exposure MISR methods which we adapt and retrain to multi-exposure inputs (§5.3).

We also introduce a dataset of 2500 multi-exposure real SkySat bursts (§5.4). The dataset only consists of noisy LR images, but we can nevertheless train our network on it, since it is self-supervised. Both on synthetic and real data, the proposed HDR-DSP method attains the best results by a significant margin *even though it is trained without high resolution ground truth data*. The dataset is available for download on the project website.

2. Related work

Most works on video and burst super-resolution focus on the single-exposure case [7, 12, 17, 34, 39, 43, 49, 52]. The problem of super-resolution from multi-exposure sequences has received much less attention. In [53] it is modeled as an overdetermined system and solved via a non regularized least-squares approach. An affine motion model and exact knowledge of the exposure times are assumed. The authors in [55] address the case in which the images have motion blur due to the camera shake. They also consider a static scene and do not consider noise. A related method for HDR imaging uses dual exposure sensors, which interlace two exposures in even and odd columns of the image [15, 26]. This can be seen as horizontally super-resolving the video.

Other works perform a related task: joint superresolution and reverse tone-mapping [30-32]. The difference with our problem is that the input video is a singleexposure LR video, and the goal is to artificially increase its dynamic range to adapt it to HDR screens.

Methods for HDR imaging from multiple exposures need to deal with the noise. Granados et al. [23] address the case of signal-dependent noise and propose a fixed point iteration of the MLE estimator which is close to the Cramer-Rao bound [3]. In these works, the denoising comes only from the temporal fusion. In [1, 2], this is incorporated in into spatio-temporal patch-based denoisers.

Our work can also be related to burst and video joint denoising and demosaicing [19,25,56], as demosaicing can be regarded as a super-resolution problem.

3. Observation model

We denote by \mathfrak{I}_t a dynamic infinite-resolution ideal scene. The camera on the satellite captures a sequence of mlow resolution images \bar{I}_i^{LR} with different exposures. For the i-th acquisition, the dynamic scene \mathfrak{I}_t is integrated during an exposure time e_i centered at t_i . Even if satellites travel at a very high speed relative to the ground, precise electrooptical image stabilization systems (with piezo-electric actuators [29, 33] or steering mirrors [46]) assure that the observed scene \mathfrak{I}_t is mostly constant during the exposure time (~2ms), which allows us to approximate the temporal integration with a product in our observation model

$$\bar{I}_i^{LR} = e_i \Pi_1 \left(\mathfrak{I}_{t_i} \ast k \right) + n_i = e_i \mathcal{I}_i^{LR} + n_i.$$
(1)

Here k is the Point Spread Function (PSF) modeling jointly optical blur and pixel integration, Π_1 is the bi-dimensional sampling operator due to the sensor array, \mathcal{I}_i^{LR} is the clean low-resolution image corresponding to an exposure of 1 unit of time and n_i denotes the noise. Throughout the text, calligraphic fonts \mathcal{I}_i denote noise-free images and regular fonts I_i noisy ones. A bar $\bar{I}_i = e_i I_i$ indicates that the image is multiplied by its exposure time (i.e. as it is acquired by the sensor), while its absence denotes images *normalized* to an exposure time of 1. We consider the *r*-th image \bar{I}_r^{LR} in the time series as the *reference*, and without loss of generality we assume its exposure time to be one, $e_r = 1$.

We model the noise as spatially independent, additive Gaussian noise with zero mean and signal-dependent variance $n_i(x) \sim \mathcal{N}(0, \sigma^2(\bar{\mathcal{I}}_i^{LR}(x)))$, where

$$\sigma^2(\bar{\mathcal{I}}_i^{LR}(x)) = ae_i \mathcal{I}_i^{LR}(x) + b, \qquad (2)$$

is an approximation of the Poisson shot noise plus Gaussian readout noise [21, 45], with parameters a and b.

Because of the spectral decay imposed by the pixel integration and optical blur (k), the images $\Im_{t_i} * k$ are band limited with a cutoff at about twice the sampling rate of the LR images for SkySat. Our goal is to increase the resolution by a factor 2 by estimating $\widehat{\mathcal{I}}_{r}^{HR}$, a non-aliased sampling of $\Im_{t_r} * k$ from several LR observations $\{\overline{I}_i^{LR}\}_{i=1}^m$ with varying exposures $\{e_i\}_{i=1}^m$. A sharp super-resolved image can then be recovered by partially deconvolving k.

In order for the method to be applicable in practice, it needs to handle time series with a variable number of frames m, and to be robust to inaccuracies in the exposure times e_i , as the exposure times in the image metadata are only a coarse approximation of the real ones.

4. Proposed method

Our method builds upon the DSA method for MISR introduced in [43], which can be regarded as a trainable generalization of the traditional shift-and-add (S&A) algorithms [6, 22, 24, 28, 38]. A *feature S&A* is used to fuse feature representations produced from the LR images by an encoder network. A motion estimation network computes the optical flows between each input LR frame and the reference frame. The output of the feature S&A is a high-resolution aggregated feature map, which is then decoded by another network to produce the output image.

The DSA method could be extended to multi-exposure sequences by applying it to the normalized images $I_i^{LR} = \bar{I}_i^{LR}/e_i$. This approach however is sub-optimal because it neglects the fact that the normalization alters the noise variance model, and fails if the reported exposure times are inaccurate, which is the case in practice.

To better exploit multiple exposures, we propose two modifications: (1) A base-detail decomposition, which provides robustness to errors in the exposure times; (2) An encoding of the images that is made dependent on the noise variance, which allows the encoder to weight different contributions according to their signal-to-noise ratio. In addition, we also propose a new feature pooling fusion intended to capture a richer picture of the encoded features, leading to a substantial improvement in reconstruction quality, both for single and multiple exposure cases. The resulting net-



Figure 2. Overview of our proposed multi-exposure super-resolution network architecture HDR-DSP at inference time.

work can be trained end-to-end with self-supervision, i.e. without requiring ground truth.

4.1. Architecture

Figure 2 shows a diagram of our proposed architecture which takes as input a sequence of multi-exposed LR images $\{\bar{I}_i^{LR}\}_{i=1}^m$ along with the corresponding exposure times e_i and produces one super-resolved image $\widehat{\mathcal{I}}_r^{HR}$. The input LR images are first normalized to unit exposure time. The normalized LR images $\{I_i^{LR}\}_{i=1}^m$ are then decomposed into base $\{B_i^{LR}\}$ and detail $\{D_i^{LR}\}$ components. The bases contain the low frequencies. We align and average them to reduce the low frequency noise and upsample the result using bilinear zooming to produce the HR base component. The LR detail images are fed to a shared convolutional En*coder* network that outputs a feature representation of each LR image. The features are then merged into a HR feature map by our shift-and-pool block (FSP), which aligns the LR features into the HR grid of the reference frame, and applies different pooling operations. The pooled features are then concatenated and fed to a *Decoder* CNN module that produces the HR detail image. The final HR image is obtained by adding the HR base and detail $\widehat{\mathcal{I}}_{r}^{HR} = \widehat{\mathcal{B}}_{r}^{HR} + \widehat{\mathcal{D}}_{r}^{HR}$.

The trainable modules of the proposed architecture (shown in red in Figure 2) include the Motion Estimator, the Encoder, and the Decoder.

Base-Detail decomposition. As mentioned above, normalizing a sequence of the frames \bar{I}_i^{LR} by their reported exposures e_i does not result in stable intensity levels across the sequence. This can be due to small errors in e_i . However, uncorrected vignetting or stray light also contribute the same effect, even in single-exposure imagery.

The nature of the super-resolution task makes it very sensitive to these exposure fluctuations. The shift-and-add operation would merge the LR features into an incoherent high-resolution feature map, making the task of the decoder more difficult, resulting in loss of details or high-frequency artifacts (see Figure 3). Refining the initial e_i could limit this problem. But this entails its own challenges, especially if one also considers vignetting and stray light sources.

Instead, in this paper we propose a more robust and simple alternative, which is based on a base-detail decomposition [44] of the normalized LR images defined as follows

$$B_i^{LR} = I_i^{LR} * G, \qquad D_i^{LR} = I_i^{LR} - B_i^{LR}, \quad (3)$$

for i = 1, ..., m. Here G is a Gaussian kernel of standard deviation 1. We then process independently the details $\{D_i^{LR}\}$ and the bases $\{B_i^{LR}\}$ to produce the corresponding high resolution estimates $\widehat{\mathcal{D}}_r^{HR}$ and $\widehat{\mathcal{B}}_r^{HR}$. This decomposition is linear and does not affect the super-resolution since the alias is preserved in the detail components $\{D_i^{LR}\}$.

As the detail images span a smaller intensity range than the complete image I_i^{LR} , an error δ in the exposure time results in a small deviation in the detail and a large one in the base: $\delta B_i^{LR} + \delta D_i^{LR} = \delta I_i^{LR}$. The small error in the detail can be handled by a super-resolution method.

On the other hand, the base images do not need to be super-resolved, but still need to be denoised. In this work we propose a simple processing that aligns and averages the bases and upsamples the result. To fully exploit the high signal-to-noise ratio of longer exposures, the average is weighted by the exposure times e_i

$$B^{HR} = \text{Zoom}\left(\frac{\sum_{i} e_i \operatorname{Warp}(B_i^{LR})}{\sum_{i} e_i}\right).$$
 (4)

This weighting is an approximation of the ML estimator of Granados et al. [23] (details in the supplementary material).

Base and detail decompositions have been used in superresolution networks [27, 31] to focus the network capacity on the details. In our case, the decomposition also provides robustness to errors in the radiometric normalization.

Motion Estimator. We follow the works of [43,49] to build a network (with the same hourglass architecture) that esti-



Figure 3. High frequency artifacts in a reconstruction from a real SkySat sequence (using DSA [43]) with exposure time errors (left). HDR-DSP with the proposed base-detail (BD) decomposition does not present artifacts (right).

mates the optical flows between the normalized LR frames $\{I_i^{LR}\}_{i=1}^m$ and the normalized reference frame I_r^{LR}

$$F_{i \to r} = \mathbf{MotionEst}(I_i^{LR}, I_r^{LR}; \Theta_{\mathbf{M}}) \in [-R, R]^{H \times W \times 2},$$
(5)

where $\Theta_{\mathbf{M}}$ denotes the network parameters. A small Gaussian filter ($\sigma = 1$) is applied to the input images to reduce the alias [43, 54]. The network is trained with a maximum motion range of $[-R, R]^2$ (with R = 5 pixels). The training was adapted to better handle the noise difference due to the multi-exposure setting (see §4.2).

Noise-level-aware detail encodings. The Encoder module generates relevant features J_i^{LR} for each normalized LR detail image D_i^{LR} in the sequence

$$J_i^{LR} = \mathbf{Encoder}(D_i^{LR}, \bar{I}_i^{LR}; \Theta_{\mathbf{E}}) \in \mathbb{R}^{H \times W \times N}, \quad (6)$$

where $\Theta_{\mathbf{E}}$ is the set of parameters of the encoder and N = 64 is the number of produced features. The network architecture is detailed in the supplementary material.

The un-normalized low resolution frames \bar{I}_i^{LR} are also fed to the encoder. This is motivated by the fact that the maximum likelihood fusion of noisy acquisitions into a (HDR) image is a weighted average, where the weights are the inverse of the noise variances [3,23]. In the proposed architecture, the normalized details D_i^{LR} are fused to produce a high resolution detail \hat{D}_r^{HR} . The noisy un-normalized images are unbiased estimators of an affine function of the noise variances $\sigma^2(\tilde{I}_i^{LR})/a - b/a$, thus they provide to the encoder the information required to compute the optimal fusion weights. The resulting features J_i^{LR} are then aggregated via a set of pooling operations, without any particular handling related to different source exposures.

Feature Pooling. We propose the Feature Shift-and-Pool block (FSP) which maps the LR features into their positions on the reference HR grid and pools them. First the features are "splatted" bilinearly onto the HR grid by the SPMC module [52]. Each LR frame is upscaled by introducing zeros between samples and motion compensated following the flows $F_{i \rightarrow r}$. This is differentiable with respect to

the intensities and the optical flows. Each splatted pixel is assigned a bilinear weight depending on the fractional part of its position in the HR grid. See [43, 52] for details.

This results in a set of aligned sparse HR feature maps

$$J_i^{HR} = \operatorname{SPMC}(J_i^{LR}, \{F_{i \to r}\}) \in \mathbb{R}^{sH \times sW \times N}, \quad (7)$$

and the corresponding bilinear splatting weights $W_i^{HR} =$ SPMC $(1, \{F_{i \to r}\})$. The upscaling factor s is set to 2.

As in [43], we use a weighted average pooling in the temporal direction (8). In addition, we propose computing the standard deviation and the max (9):

$$J_A^{HR} = (\sum_i J_i^{HR}) (\sum_i W_i^{HR})^{-1},$$
(8)

$$J_M^{HR} = \max_i J_i^{HR}, \qquad J_S^{HR} = \operatorname{std}_i J_i^{HR}.$$
(9)

Note that this block does not have any trainable parameters, a trainable layer may attain a similar performance at a much higher computational cost (see the supplementary material).

These feature pooling operations render the architecture invariant to permutations of the input frames [5]. The key idea is that through end-to-end training, the encoder network will learn to output features for which the pooling is meaningful. Therefore, it is essential that the pooling operation is capable of passing all the necessary information to the decoder. Indeed, average pooling captures a consensus of the features, which amounts to a temporal denoising. But in aliased image sequences, it is common to come across features that are only visible in a single frame. Thus, the idea of the max-pooling operation is to preserve these unique features that would otherwise be lost in the average. The standard deviation pooling completes the picture by measuring the point-wise variability of the features.

The pooled features are independent of the number of processed frames. But this information is important as the decoder may interpret features resulting from aggregating many images differently than those resulting from just a few. For this reason, the aggregation weights $W^{HR} = \sum_i W_i^{HR}$ are also concatenated with the pooled features. As we will see in §5.2, incorporating W^{HR} improves the network ability to handle a variable number of input frames. **Decoder.** The Decoder network reconstructs the HR detail image $\hat{\mathcal{D}}_r^{HR}$ from the pooled features

$$\widehat{\mathcal{D}}_{r}^{HR} = \mathbf{Decoder}(J_{A}^{HR}, J_{M}^{HR}, J_{S}^{HR}, W^{HR}; \Theta_{\mathbf{D}}) \in \mathbb{R}^{sH \times sW}$$
(10)

where $\Theta_{\mathbf{D}}$ denotes the set of parameters of the decoder. The architecture is detailed in the supplementary material.

4.2. Self-supervised learning

To train the HDR-DSP detail fusion network, we adapt the fully self-supervised framework of [43], which requires no ground truth HR images. During training, the LR frames are randomly selected and for every sequence, one frame is set apart as the reference I_r^{LR} . Then, all the other LR images in each sequence are registered against the reference using the **MotionEst** network yielding the flows $F_{i \rightarrow r}$. The reference frame serves as the target for the self-supervised training similarly to noise-to-noise [20, 35]. The procedure relies on the minimization of a reconstruction loss in the LR domain plus a motion estimation loss to ensure accurate alignment of the frames. The losses and the proposed adaptations are detailed in the following paragraphs.

Self-supervised SR loss. The self-supervised loss forces the network to produce an HR detail \hat{D}_r^{HR} such that when subsampled, it coincides (modulo the noise) with the withheld target detail D_r^{LR}

$$\ell_{self}(\widehat{D}_{r}^{HR}, D_{r}^{LR}) = \|\Pi_{2}(\widehat{D}_{r}^{HR} * k) - D_{r}^{LR}\|_{1}, \quad (11)$$

where $\widehat{\mathcal{D}}_{r}^{HR} = \operatorname{Net}(\{D_{i}^{LR}\}_{i \neq r}, \{\overline{I}_{i}^{LR}\}_{i=1}^{m})$ is the SR output, and Π_{2} is the subsampling operator that takes one pixel over two in each direction. As in [43] we include the convolution kernel k in the loss. This forces the network to produce a deconvolved HR image that once convolved with k and subsampled matches the optical blur present in D_{r}^{LR} .

During training, the LR reference is only used in the motion estimator to compute the optical flows, but it is not fused into the HR result to avoid unwanted trivial solutions [11, 18, 43]. At inference time we use the reference as this leads to improved results [43].

Grid shifting. The self-supervised loss (11) downsamples the super-resolved detail to compare it with the reference LR detail. But since the downsampling is fixed, only the sampled positions intervene in the loss, which breaks the translation equivariance of the method. To avoid this issue, during training we augment the data by adding to the estimated optical flows a random shift of 0.5ϵ in each dimension ($\epsilon \in \{0, 1\}$). As a result, the super-resolved image is shifted by ϵ , which is easily compensated before computing the loss. This yields an improvement in PSNR of 0.2dB.

Motion estimation loss. The motion estimator is trained with unsupervised learning as in [57]. The loss consists of a warping term and a regularization term. We observed that the optical flow is very sensitive to the intensity fluctuations between frames (as in our normalized LR frames I_i^{LR}), which result in imprecise alignments. To prevent this issue we compute the warping loss on the details rather than on the images, which is common in traditional optical flow [36, 49]. The loss is computed for each flow $F_{i\to r}$ estimated by the **MotionEst** module

$$\ell_{me}(\{F_{i\to r}\}_{i=1}^{m}) = \lambda_1 TV(F_{i\to r}) + \sum_i \|\mathbf{Detail} \left(I_i^{LR} - \mathbf{Pullback}(I_r^{LR}, F_{i\to r})\right)\|_1, \quad (12)$$

where **Pullback** computes a bicubic warping of I_r^{LR} according to a flow, **Detail** applies a high-pass filter, TV is the finite difference discretization of the classic Total Variation regularizer [48], and $\lambda_1 = 0.003$ is a hyperparameter controlling the regularization strength.

Training. The self-supervised training of HDR-DSP is done in two stages. We first pretrain the motion estimator on the simulated data to ensure that it produces accurate flows. Then, we train the entire system end-to-end with the pre-trained **MotionEst** using the self-supervised loss ($\lambda_2 = 3$)

$$loss = \ell_{self} + \lambda_2 \ell_{me}.$$
 (13)

Other training details are in the supplementary material.

5. Experiments

For our experiments, we use real multi-exposure pushframe images (L1A) acquired by SkySat satellites [41]. For the quantitative evaluations we also simulated a multiexposure and a single-exposure datasets from L1B products (super-resolved products by Planet with a factor of 1.25).

5.1. Simulated multi-exposure dataset

The two simulated datasets were generated from 1371 crops of L1B products (1096 train, 200 test, 75 val). First, we generate the noise-free LR images normalized to an exposure time of 1. Random subpixel translations of $\{\Delta_i\}_{i=1}^m$ are applied to the ground truth followed by ×2 subsampling

$$\mathcal{I}_{r}^{LR} = \Pi_{2}(\mathcal{I}^{HR}),$$

$$\mathcal{I}_{i}^{LR} = \Pi_{2}(\mathrm{Shift}_{\Delta_{i}}(\mathcal{I}^{HR})), \qquad i \neq r$$
(14)

where Π_2 is the subsampling operator. The exposure times are simulated as $e_i = \alpha^{c_i}$, where $c_i \in \{-5, ..., 5\}$, and $\alpha =$ uniform(1.2, 1.4). The noises $n_i = \sqrt{ae_i \mathcal{I}_i^{LR} + b} \mathcal{N}(0, 1)$ are then added to all the un-normalized frames to produce the noisy multi-exposure sequence $\bar{I}_i^{LR} = e_i \mathcal{I}_i^{LR} + n_i$. The constants a = 0.119, b = 12.050 were estimated from real SkySat images with the Ponomarenko noise curve estimation method [16, 45]. The single-exposure dataset is generated in the same manner but with all $e_i = 1$. To simulate the exposure inaccuracies, during training and testing the e_i values are contaminated with noise within a range of 5%.

We use a PSNR score in our evaluation. The SkySat L1A images have a dynamic range of 12 bits, but we observed that the peak signal is at about 3400 DN. Therefore, our PSNR is normalized with a peak of 3400. We denote PSNR ME (resp. PSNR SE) as the average PSNRs computed on all the multi-exposure (resp. single-exposure) test sequences.

5.2. Ablation study

We study in Table 1 the importance of the base-detail decomposition. We consider simulated multi-exposure (ME)

Table 1. Handling of multi-exposure sequences with base-detail decomposition (BD) and using the un-normalized LR frames I_i^{LR} as an additional encoder input.

Methods (all HDR-DSP)	full	w/o BD	w/o BD (trained SE)	w/o LR
PSNR(dB) ME	54.70	53.76	52.91	53.94
PSNR(dB) SE	54.72	54.16	54.54	54.16

Table 2. Feature pooling choice. Using average (A), maximum (M), and standard deviation (S) pooling improves the results.

Features	AMS (HDR-DSP)	AS	AM	А
PSNR(dB) ME	54.70	54.46	54.44	54.17
PSNR(dB) SE	54.72	54.47	54.48	54.20

Table 3. Handling variable number of frames (PSNR ME (dB)).

Methods (all HDR-DSP)	full	w/o W^{HR}	HDR-DSP 4	HDR-DSP 14
4 frames	52.81	52.60	52.69	51.31
14 frames	55.85	55.59	54.26	55.53
variable n frames	54.70	54.45	53.85	54.07

and single-exposure (SE) sequences presenting small exposure errors that match the ones observed in real sequences. If we train HDR-DSP without the proposed base-detail (w/o BD), the performance drops noticeably, which is also visible on real sequences (Figure 3). Even when training specifically for a single-exposure setting, as with DSA [43], the performance with base-detail is superior. In addition, we can see that removing the un-normalized LR frame from the encoder inputs (w/o LR) leads to a large performance drop for both single- and multi-exposure.

The experiment shown in Table 2 studies the impact of using multiple feature pooling strategies: average, maximum, and standard deviation. It shows that using the three greatly improves the results: about 0.5dB with respect to just using average. We observed that not including the average among the pooling strategies yields much worse results.

The aggregation weight feature W^{HR} was added to improve the handling by the decoder of sequences with variable number of input frames. The results in Table 3 confirm the importance of providing these weights. We also compare with networks trained for a fixed number of frames (HDR-DSP 4 and 14) and observe that in this case the performance drops even when testing for those specific configurations. We conclude that the weights become useless if the training does not consider a variable number of frames.

Lastly, removing the grid shifting (§4.2) from the training also reduces the PSNR ME: from 54.70 to 54.49dB.

5.3. Comparison with the state-of-the-art

We compare our self-supervised network on the simulated dataset against state-of-the-art MISR methods for

Table 4. PSNR ME (dB) over the synthetic test set with 15 images in the case of 0%, 5% and 20% exposure time errors.

Methods	RAMS	ME S&A	HR-net	BD-ACT	DSA	HDR-DSP
0% exp. error	52.05	53.33	54.30	54.24	55.55	56.00
5% exp. error	51.84	52.43	54.22	54.23	54.99	55.99
20% exp. error	49.95	49.19	53.82	54.20	54.30	55.90

satellite images: *DSA* [43], *HighRes-net* (HR-net) [17], *RAMS* [50], and *ACT* [7]. A weighted *Shift-and-add* [38] with bicubic splatting adapted to multi-exposure sequences (ME S&A) serves as the baseline. HR-net and RAMS are two supervised networks designed to perform super-resolution of multi-temporal PROBA-V satellite images. In the context of push-frame satellites, we use the reference-aware version [42] of HR-net and RAMS rather than the original approaches, as they achieve higher quality results. DSA and ACT are two state-of-the-art super-resolution methods for SkySat imagery. ACT also serves as a proxy for comparison with other interpolation-based methods from the literature [56].

We adapt these methods to multi-exposure sequences. The deep learning approaches are fed with the normalized input images, whereas for ACT method we apply the same base-detail decomposition described in §4.1 and use ACT to restore the details (denoted BD-ACT). The registration step of ME S&A, BD-ACT, and RAMS are done with the inverse compositional algorithm [10, 13], which is robust to noise and brightness changes. The motion estimator of DSA is also trained with the loss on the details (§4.2).

Table 4 shows a quantitative comparison of the methods over the test set in the case of adding exposure time errors of 5% (as during training) and 20%. These errors are estimated from SkySat data (exposures ranging from 0.5 to 4.5 ms); see the supplementary material for details. Note that even with exact exposure times (row 0%), vignetting or stray light effects still justify the use of the proposed basedetail decomposition. Our self-supervised network ranks first in all cases with a significant gain of more than 1dB over all others (see Figure 4). Interestingly, the performance of most methods degrades quickly for large inaccuracy in exposure times. Only the methods using the base-detail decomposition (BD-ACT and ours) are robust to these inaccuracies. Note that HDR-DSP has never seen errors of 20% during training.

5.4. Results on real data

The proposed self-supervised training allows to train HDR-DSP on real multi-exposure sequences taken from SkySat satellites. From the L1A product of Planet SkySat, we extracted 2500 sequences (128×128 pixels) pre-registered up-to an integer translation. Out of 2500 sequences, 300 are used for testing. Each sequence contains



Figure 4. Super-resolution from a synthetic multi-exposure sequence (5% exp. error) of 15 aliased LR images. Methods are trained on a synthetic dataset and receive as inputs the normalized ME images except BD-ACT and HDR-DSP, which use the base-detail decomposition.



Figure 5. Super-resolution from a real multi-exposure sequence of 9 SkySat images. The first line corresponds to 4 normalized LR images in that sequence with different exposure times. The second line shows the reconstructions by Planet (L1B), DSA, BD-ACT and our method HDR-DSP.

from 4 to 15 frames. In about 75% of the sequences the exposure time varies within each sequence and we used the exposure time information provided in the metadata.

Figure 5 compares HDR-DSP against Planet L1B, DSA, and BD-ACT. The top row shows four normalized frames of the sequence, where we can notice the dependence of the noise level on the exposure time. The method used in the Planet L1B product is unknown. It super-resolves by a factor of 1.25 but contains noticeable artifacts and lacks fine details. The result from DSA exhibits a high-frequency pattern due to the imprecise exposure times. BD-ACT is able to cope with the exposure changes thanks to the base-detail decomposition, but the result is still very noisy. In contrast, HDR-DSP shows a clean and detailed reconstruction.

Figure 1 also shows a multi-exposure LR sequence along with the results from ME S&A, Planet L1B, ACT, DSA and HDR-DSP. Comparing HDR-DSP with DSA, we see that the former provides a cleaner result thanks to the basedetail decomposition and the proposed improvements over the DSA architecture and training procedure, which is also observed in the synthetic experiments.

6. Conclusion and limitations

The proposed HDR-DSP method is able to reconstruct high-quality results from multi-exposure bursts, providing fine details, low-noise, and high dynamic range. The proposed base-detail processing allows robustness to errors in the exposure time that are common in practice. In addition, a significant performance improvement is obtained by making the image encoding dependent on the noise variance, and using a new feature pooling designed to capture richer representations. Thanks to its fully self-supervised training, the method requires no ground truth and can thus be applied on real data. We show its effectiveness by training a model that super-resolves multi-exposure SkySat L1A acquisitions, leading to a substantial resolution gain with respect to the state-of-the-art.

Limitations. The context of remote sensing allows one to make additional assumptions that do not hold in more general settings: 1. The considered noise levels are away from the challenging photon-limited regime; 2. Motion and occlusions are much easier to handle. In particular, the latter point should be improved to apply this method to video or burst super-resolution. Besides, the proposed method does not handle saturation. This will be studied in future work.

Acknowledgments. Work supported by a grant from Région Île-de-France. This work was performed using HPC resources from GENCI–IDRIS (grants 2022-AD011012453R1 and 2022-AD011012458R1) and from the "Mésocentre" computing center of CentraleSupélec and ENS Paris-Saclay supported by CNRS and Région Île-de-France (http://mesocentre.centralesupelec.fr/). We thank Planet for providing the L1A SkySat images.

References

- [1] Cecilia Aguerrebere, Andrés Almansa, Julie Delon, Yann Gousseau, and Pablo Musé. A bayesian hyperprior approach for joint image denoising and interpolation, with an application to hdr imaging. *IEEE Transactions on Computational Imaging*, 3(4):633–646, 2017. 3
- [2] Cecilia Aguerrebere, Julie Delon, Yann Gousseau, and Pablo Muse. Simultaneous hdr image reconstruction and denoising for dynamic scenes. In *IEEE International Conference* on Computational Photography (ICCP), pages 1–11. IEEE, 2013. 3
- [3] Cecilia Aguerrebere, Julie Delon, Yann Gousseau, and Pablo Musé. Best algorithms for hdr image generation. a study of performance bounds. *SIAM Journal on Imaging Sciences*, 7(1):1–34, 2014. 3, 5
- [4] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pages 126–135, 2017. 2
- [5] Miika Aittala and Frédo Durand. Burst image deblurring using permutation invariant convolutional neural networks. In Proceedings of the European Conference on Computer Vision (ECCV), pages 731–747, 2018. 5
- [6] Mohammad S. Alam, John G. Bognar, Russell C. Hardie, and Brian J. Yasuda. Infrared image registration and high-resolution reconstruction using multiple translationally shifted aliased video frames. *IEEE Transactions on instrumentation and measurement*, 49(5):915–923, 2000. 3
- [7] Jérémy Anger, Thibaud Ehret, Carlo de Franchis, and Gabriele Facciolo. Fast and accurate multi-frame superresolution of satellite images. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences*, 5(1), 2020. 1, 2, 3, 7
- [8] Jérémy Anger, Thibaud Ehret, and Gabriele Facciolo. Parallax estimation for push-frame satellite imagery: application to super-resolution and 3d surface modeling from skysat products. In 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS, pages 2679–2682. IEEE, 2021. 1
- [9] Md Rifat Arefin, Vincent Michalski, Pierre-Luc St-Charles, Alfredo Kalaitzis, Sookyung Kim, Samira E. Kahou, and Yoshua Bengio. Multi-image super-resolution for remote sensing using deep recurrent networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pages 206–207, 2020. 2
- [10] Simon Baker and Iain Matthews. Equivalence and efficiency of image alignment algorithms. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 1, pages I–I. IEEE, 2001. 7
- [11] Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. In *International Conference on Machine Learning*, pages 524–533. PMLR, 2019. 6
- [12] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Deep burst super-resolution. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9209–9218, 2021. 3

- [13] Thibaud Briand, Gabriele Facciolo, and Javier Sánchez. Improvements of the Inverse Compositional Algorithm for Parametric Motion Estimation. *IPOL*, 8:435–464, 2018. 7
- [14] Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3086–3095, 2019. 2
- [15] Uğur Çoğalan, Mojtaba Bemana, Karol Myszkowski, Hans-Peter Seidel, and Tobias Ritschel. Hdr denoising and deblurring by learning spatio-temporal distortion models. arXiv preprint arXiv:2012.12009, 2020. 3
- [16] Miguel Colom and Antoni Buades. Analysis and extension of the ponomarenko et al. method, estimating a noise curve from a single image. *Image Processing On Line*, 3:173–197, 2013. 6
- [17] Michel Deudon, Alfredo Kalaitzis, Israel Goytom, Md Rifat Arefin, Zhichao Lin, Kris Sankaran, Vincent Michalski, Samira E Kahou, Julien Cornebise, and Yoshua Bengio. Highres-net: Recursive fusion for multi-frame superresolution of satellite imagery. arXiv 2020. arXiv preprint arXiv:2002.06460, 2020. 1, 2, 3, 7
- [18] Valéry Dewil, Jérémy Anger, Axel Davy, Thibaud Ehret, Gabriele Facciolo, and Pablo Arias. Self-supervised training for blind multi-frame video denoising. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2724–2734, 2021. 2, 6
- [19] Thibaud Ehret, Axel Davy, Pablo Arias, and Gabriele Facciolo. Joint demosaicing and denoising by overfitting of bursts of raw images. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 3
- [20] Thibaud Ehret, Axel Davy, Jean-Michel Morel, Gabriele Facciolo, and Pablo Arias. Model-blind video denoising via frame-to-frame training. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2, 6
- [21] Alessandro Foi, Mejdi Trimeche, Vladimir Katkovnik, and Karen Egiazarian. Practical poissonian-gaussian noise modeling and fitting for single-image raw-data. *IEEE Trans. Im*age Process., 17(10):1737–1754, 2008. 3
- [22] Andrew Fruchter and Richard Hook. Drizzle: A method for the linear reconstruction of undersampled images. *Publications of the Astronomical Society of the Pacific*, 114(792):144, 2002. 3
- [23] Miguel Granados, Boris Ajdin, Michael Wand, Christian Theobalt, Hans-Peter Seidel, and Hendrik PA Lensch. Optimal hdr reconstruction with linear digital cameras. In 2010 IEEE computer society conference on computer vision and pattern recognition, pages 215–222. IEEE, 2010. 3, 4, 5
- [24] Thomas J. Grycewicz, Stephen A. Cota, Terrence S. Lomheim, and Linda S. Kalman. Focal plane resolution and overlapped array TDI imaging. In *Remote Sensing System Engineering*, volume 7087, page 708704. International Society for Optics and Photonics, 2008. 3
- [25] Samuel W. Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T. Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. Burst photography for high dynamic range and

low-light imaging on mobile cameras. ACM Transactions on Graphics (ToG), 35(6):1–12, 2016. 3

- [26] Felix Heide, Markus Steinberger, Yun-Ta Tsai, Mushfiqur Rouf, Dawid Pajak, Dikpal Reddy, Orazio Gallo, Jing Liu, Wolfgang Heidrich, Karen Egiazarian, et al. Flexisp: A flexible camera image processing framework. ACM Transactions on Graphics (ToG), 33(6):1–13, 2014. 3
- [27] Takashi Isobe, Xu Jia, Shuhang Gu, Songjiang Li, Shengjin Wang, and Qi Tian. Video Super-Resolution with Recurrent Structure-Detail Network. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2020. 4
- [28] Yunwei Jia. Method and apparatus for super-resolution of images, Nov. 6 2012. US Patent 8,306,121. 3
- [29] Emiliano Kargieman, Gerado Gabriel Richarte, and Juan Manuel Vuletich. Imaging device for scenes in apparent motion. U.S. Patent 9813601B2, issued November 7, 2017.
- [30] Soo Ye Kim and Munchurl Kim. A multi-purpose convolutional neural network for simultaneous super-resolution and high dynamic range image reconstruction. In *Asian Conference on Computer Vision*, pages 379–394. Springer, 2018.
 3
- [31] Soo Ye Kim, Jihyong Oh, and Munchurl Kim. Deep SR-ITM: Joint learning of super-resolution and inverse tonemapping for 4K UHD HDR applications. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision, pages 3116–3125, 2019. 3, 4
- [32] Soo Ye Kim, Jihyong Oh, and Munchurl Kim. Jsi-gan: Ganbased joint super-resolution and inverse tone-mapping with pixel-wise task-specific filters for uhd hdr video. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11287–11295, 2020. 3
- [33] Mary Knapp, Sara Seager, Brice-Olivier Demory, Akshata Krishnamurthy, Matthew W. Smith, Christopher M. Pong, Vanessa P. Bailey, Amanda Donner, Peter Di Pasquale, Brian Campuzano, Colin Smith, Jason Luu, Alessandra Babuscia, Robert L. Bocchino, Jr., Jessica Loveland, Cody Colley, Tobias Gedenk, Tejas Kulkarni, Kyle Hughes, Mary White, Joel Krajewski, and Lorraine Fesq. Demonstrating Highprecision Photometry with a CubeSat: ASTERIA Observations of 55 Cancri e. *The Astronomical Journal*, 160(1):23, jun 2020. 3
- [34] Bruno Lecouat, Jean Ponce, and Julien Mairal. Lucaskanade reloaded: End-to-end super-resolution from raw image bursts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2370–2379, 2021. 3
- [35] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning Image Restoration without Clean Data. In 35th International Conference on Machine Learning, ICML 2018, 2018. 6
- [36] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Selflow: Self-supervised learning of optical flow. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2019. 6

- [37] Marcus Märtens, Dario Izzo, Andrej Krzic, and Daniël Cox. Super-resolution of PROBA-V images using convolutional neural networks. *Astrodynamics*, 3(4):387–402, 2019. 2
- [38] Maria Teresa Merino and Jorge Nunez. Super-resolution of remotely sensed images with variable-pixel linear reconstruction. *IEEE TGRS*, 45(5):1446–1457, 2007. 3, 7
- [39] Andrea Bordone Molini, Diego Valsesia, Giulia Fracastoro, and Enrico Magli. Deepsum: Deep neural network for super-resolution of unregistered multitemporal images. *IEEE Transactions on Geoscience and Remote Sensing*, 58(5):3644–3656, 2019. 2, 3
- [40] Andrea Bordone Molini, Diego Valsesia, Giulia Fracastoro, and Enrico Magli. Deepsum++: Non-local deep neural network for super-resolution of unregistered multitemporal images. In IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium, pages 609–612, 2020. 1
- [41] Kiran Murthy, Michael Shearn, Byron D. Smiley, Alexandra H. Chau, Josh Levine, and Dirk Robinson. SkySat-1: very high-resolution imagery from a small satellite. In Sensors, Systems, and Next-Generation Satellites XVIII, volume 9241, page 92411E. International Society for Optics and Photonics, 2014. 1, 2, 6
- [42] Ngoc Long Nguyen, Jérémy Anger, Axel Davy, Pablo Arias, and Gabriele Facciolo. PROBA-V-REF: Repurposing the PROBA-V Challenge for Reference-Aware Super Resolution. In 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS, pages 3881–3884. IEEE, jul 2021. 7
- [43] Ngoc Long Nguyen, Jérémy Anger, Axel Davy, Pablo Arias, and Gabriele Facciolo. Self-supervised multi-image superresolution for push-frame satellite images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1121–1131, June 2021. 1, 2, 3, 4, 5, 6, 7
- [44] Joan M. Ogden, Edward H. Adelson, James R. Bergen, and Peter J. Burt. Pyramid-based computer graphics. *RCA Engineer*, 30(5):4–15, 1985. 4
- [45] Nikolay N. Ponomarenko, Vladimir V. Lukin, M.S. Zriakhov, Arto Kaarna, and Jaakko Astola. An automatic approach to lossy compression of aviris images. In 2007 IEEE International Geoscience and Remote Sensing Symposium, pages 472–475. IEEE, 2007. 3, 6
- [46] Dirk Robinson, Jonathan Dyer, Joshua Levine, Brendan Hermalyn, Ronny Votel, and Matt William Messana. Controlling a line of sight angle of an imaging platform. U.S. Patent 10432866B2, issued October 1, 2019. 3
- [47] G Rohith and Lakshmi Sutha Kumar. Paradigm shifts in super-resolution techniques for remote sensing applications. *The Visual Computer*, 37(7):1965–2008, 2021.
- [48] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992. 6
- [49] Mehdi SM Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 6626–6634, 2018. 3, 4, 6
- [50] Francesco Salvetti, Vittorio Mazzia, Aleem Khaliq, and Marcello Chiaberge. Multi-image super resolution of remotely

sensed images using residual attention deep neural networks. *Remote Sensing*, 12(14):2207, 2020. 7

- [51] Dev Yashpal Sheth, Sreyas Mohan, Joshua L. Vincent, Ramon Manzorro, Peter A. Crozier, Mitesh M. Khapra, Eero P. Simoncelli, and Carlos Fernandez-Granda. Unsupervised deep video denoising. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1759– 1768, 2021. 2
- [52] Xin Tao, Hongyun Gao, Renjie Liao, Jue Wang, and Jiaya Jia. Detail-revealing deep video super-resolution. In Proceedings of the IEEE International Conference on Computer Vision, pages 4472–4480, 2017. 3, 5
- [53] Yann Traonmilin and Cecilia Aguerrebere. Simultaneous high dynamic range and superresolution imaging without regularization. *SIAM Journal on Imaging Sciences*, 7(3):1624–1644, 2014. 1, 3
- [54] Patrick Vandewalle, Luciano Sbaiz, Joos Vandewalle, and Martin Vetterli. Super-resolution from unregistered and totally aliased signals using subspace methods. *IEEE Transactions on Signal Processing*, 2007. 5
- [55] Subeesh Vasu, Abhijeet Shenoi, and A.N. Rajagopazan. Joint hdr and super-resolution imaging in motion blur. In 2018 25th IEEE International Conference on Image Processing (ICIP), pages 2885–2889. IEEE, 2018. 1, 3
- [56] Bartlomiej Wronski, Ignacio Garcia-Dorado, Manfred Ernst, Damien Kelly, Michael Krainin, Chia-Kai Liang, Marc Levoy, and Peyman Milanfar. Handheld multi-frame superresolution. ACM Transactions on Graphics (TOG), 38(4):1– 18, 2019. 3, 7
- [57] Jason J. Yu, Adam W. Harley, and Konstantinos G. Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *European Conference on Computer Vision*, pages 3–10. Springer, 2016.
 6
- [58] Songhyun Yu, Bumjun Park, Junwoo Park, and Jechang Jeong. Joint learning of blind video denoising and optical flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 500–501, 2020. 2