

Guided Depth Super-Resolution by Deep Anisotropic Diffusion

Nando Metzger* Rodrigo Caye Daudt* Konrad Schindler
 Photogrammetry and Remote Sensing, ETH Zurich
 {metzger, rcayedaudt, schindler}@ethz.ch

Abstract

Performing super-resolution of a depth image using the guidance from an RGB image is a problem that concerns several fields, such as robotics, medical imaging, and remote sensing. While deep learning methods have achieved good results in this problem, recent work highlighted the value of combining modern methods with more formal frameworks. In this work, we propose a novel approach which combines guided anisotropic diffusion with a deep convolutional network and advances the state of the art for guided depth super-resolution. The edge transferring/enhancing properties of the diffusion are boosted by the contextual reasoning capabilities of modern networks, and a strict adjustment step guarantees perfect adherence to the source image. We achieve unprecedented results in three commonly used benchmarks for guided depth super-resolution. The performance gain compared to other methods is the largest at larger scales, such as $\times 32$ scaling. Code¹ for the proposed method is available to promote reproducibility of our results.

1. Introduction

It is a primordial need for visual data analysis to increase the resolution of images after they have been captured. In many fields one is faced with images that, for technical reasons, have too low resolutions for the intended purposes, *e.g.*, MRI scans in medical imaging [48], multi-spectral satellite images in Earth observation [22], thermal surveillance images [1] and depth images in robotics [9]. In some cases, an image of much higher resolution is available in a different imaging modality, which can act as a *guide* for super-resolving the low-resolution *source* image, by injecting the missing high-frequency content. For instance, in Earth observation, the guide is often a panchromatic image (hence the term "pan-sharpening"), whereas in robotics a conventional RGB image is often attached to the same

*Equal contribution.

¹ <https://github.com/prs-eth/Diffusion-Super-Resolution>

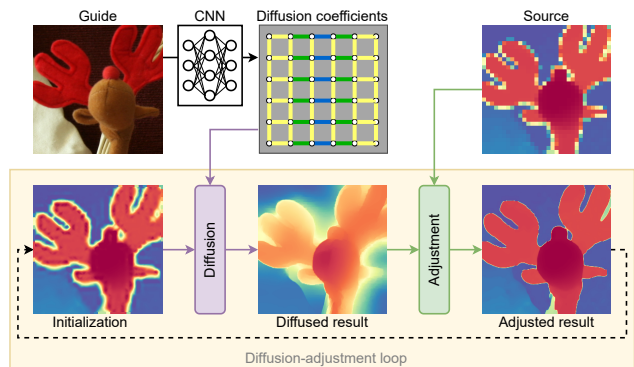


Figure 1. We super-resolve a low-resolution depth image by finding the equilibrium state of a constrained anisotropic diffusion process. Learned diffusion coefficients favor smooth depth within objects and suppress diffusion across discontinuities. They are derived from the guide with a neural feature extractor that is trained by back-propagating through the diffusion process.

platform as a TOF camera or laser scanner. In this paper, we focus on super-resolving depth images guided by RGB images, but the proposed framework is generic and can be adapted to other sensor combinations, too.

Research into guided super-resolution has a long history [16, 29]. The proposed solutions range from classical, entirely hand-crafted schemes [10] to fully learning-based methods [15], while some recent works have combined the two schools of thought, with promising results [5, 32]. Many classical methods boil down to an image-specific optimization problem that must be solved at inference time, which often makes them slow and memory-hungry. Moreover, they are limited to low-level image properties of the guide, such as color and contrast, and lack the high-level image understanding and contextual reasoning of modern neural networks. On the positive side, by design, they can not overfit the peculiarities of a training set and tend to generalize better. Recent work on guided super-resolution has focused on deep neural networks. Their superior ability to capture latent image structure has greatly advanced the state of the art over traditional, learning-free approaches. Still, these learning-based methods tend to struggle with sharp

discontinuities and often produce blurry edges in the super-resolved depth maps. Moreover, like many deep learning systems, they degrade – often substantially – when applied to images with even slightly different characteristics. Note also that standard feed-forward architectures cannot guarantee a consistent solution: feeding the source and guiding images through an encoder-decoder structure to obtain the super-resolved target will, by itself, not ensure that down-sampling the target will reproduce the source.

We propose a novel approach for guided depth super-resolution which combines the strengths of optimization-based and deep learning-based super-resolution. In short, our method is a combination of anisotropic diffusion (based on the discretized version of the heat equation) with deep feature learning (based on a convolutional backbone). The diffusion part resembles classical optimization approaches, solved via an iterative diffusion-adjustment loop. Every iteration consists of (1) an anisotropic diffusion step [2, 4, 23, 30], with diffusion weights driven by the guide in such a way that diffusion (i.e., smoothing) is low across high-contrast boundaries and high within homogeneous regions; and (2) an adjustment step that rescales the depth values such that they exactly match the low-resolution source when downsampled. To harness the unmatched ability of deep learning to extract informative image features, the diffusion weights are not computed from raw brightness values but are set by passing the guide through a (fully) convolutional feature extractor. An overview of the method is depicted in Fig. 1. The technical core of our method is the insight that such a feature extractor can be trained end-to-end to optimally fulfill the requirements of the subsequent optimization, by back-propagating gradients through the iterations of the diffusion loop. Despite its apparent simplicity, this hybrid approach delivers excellent super-resolution results. In our experiments, it consistently outperforms prior art on three different datasets, across a range of upsampling factors from $\times 4$ to $\times 32$. In our experiments, we compare it to six recent learning methods as well as five different learning-free methods. For completeness, we also include a learning-free version of our diffusion-adjustment scheme and show that it outperforms all other learning-free methods. Beyond the empirical performance, our method inherits several attractive properties from its ingredients: the diffusion-based optimization scheme ensures strict adherence to the depth values of the source, crisp edges, and a degree of interpretability; whereas deep learning equips the very local diffusion weights with large-scale context information, and offers a tractable, constant memory footprint at inference time. In summary, our contributions are:

1. We develop a hybrid framework for guided super-resolution that combines deep feature learning and anisotropic diffusion in an integrated, end-to-end trainable pipeline;

2. We provide an implementation of that scheme with constant memory demands, and with inference time that is constant for a given upsampling factor and scales linearly with the number of iterations;
3. We set a new state of the art for the Middlebury [38], NYUv2 [39] and DIML [20] datasets, for upsampling factors from $4\times$ to $32\times$, and provide empirical evidence that our method indeed guarantees exact consistency with the source image.

2. Related Work

Learning-free methods. Early work on guided super-resolution consisted mostly of optimization methods. Several such methods [8, 25, 45] employ random field models to solve this problem. Other traditional methods rely on filters, such as the bilateral filter [46], the guided filter [11], the weighted median filter [24], the weighted mode filter [26], or the Static-Dynamic filter [10]. Liu *et al.* [23] showed an early, learning-free application of anisotropic diffusion for guided depth enhancement. De Lutio *et al.* [6] fit a pixel-wise MLP to map the guide to the target and obtain excellent results when compared to other learning-free methods, whereas Uezato *et al.* [41] adapt Deep Image Prior [42] to fuse the guide and the source images.

Learning-based methods. Other avenues have also been explored, such as the auto-regressive model proposed in [43]. More recent architectures [40, 47] have used successfully applied transformer modules to this problem. With the advent of deep learning in the past decade, several authors have proposed feedforward networks for upsampling depth images. *MSG-Net* [15] is a U-Net shaped architecture that embeds the source at its lowest scale and learns the residual errors of bicubic interpolations. Kim *et al.* [18] propose Deformable Kernel Networks (*DKN* and its efficient implementation *FDKN*), that learn sparse and spatially-invariant filter kernels. He *et al.* [13] employ a high-frequency guidance module to embed the guide details into the depth map. Wen *et al.* [44] use convolutional kernels of different sizes while also using data fidelity as a convergence criterion of their iterative refinement.

Hybrid methods. A final group of methods applies deep learning methods within formal frameworks that constrain their solution and improve the methods’ inductive biases for guided super-resolution. Riegler *et al.* [32] unroll the optimization steps of a first-order primal-dual algorithm into a neural network, such that they can train their deep feature extractor in an end-to-end manner. De Lutio *et al.* [5] use the implicit function theorem to propagate through a graph-based, MRF-style optimizer, combining deep learning within a framework inspired by traditional

approaches [8, 27, 34]. Work exploring this last family of algorithms has reported promising results from combining the contextual reasoning capabilities of convolutional neural networks (CNNs) with the explicit constraints of formal frameworks. Our work can also be regarded as a member of this group of methods.

3. Method

We first describe the diffusion-adjustment framework used for guided super-resolution, then we explain how deep learning is integrated into that framework. The latter part allows us to exploit high-level context in the guide image to find optimal coefficients for the diffusion operator.

3.1. Diffusion-adjustment

Anisotropic diffusion [30] is a form of iterative edge-aware filtering initially proposed with the aim of performing intra-region smoothing while avoiding inter-region smoothing. The filtering is done in a way analog to solving a discretized heat equation with anisotropic (*i.e.*, spatially varying) diffusion weights. The idea of computing these diffusion weights from a separate (coregistered) guide image has been explored in the context of edge enhancement [23] and semantic segmentation [2, 4].

We are given a source image $\mathbf{S} \in \mathbb{R}^{\frac{H}{s} \times \frac{W}{s}}$ and a guide $\mathbf{G} \in \mathbb{R}^{H \times W \times C}$, where $C = 3$ for RGB images or a larger number for deep features. The first step in our approach is to initialize $\mathbf{Y}_0 \in \mathbb{R}^{H \times W}$ with an upsampled version of \mathbf{S} . As we will show later, the exact initialization of \mathbf{Y}_0 has little impact on the final result. We can then define a diffusion step as:

$$\hat{\mathbf{y}}_t^p = \mathbf{y}_{t-1}^p + \lambda \cdot \sum_{n \in \mathcal{N}_4^p} (\mathbf{y}_{t-1}^n - \mathbf{y}_{t-1}^p) \cdot c(\mathbf{g}^p, \mathbf{g}^n), \quad (1)$$

where \mathbf{y}_t^p denotes the pixel value of \mathbf{Y}_t at location p (and similarly for \mathbf{g}^p). \mathcal{N}_4^p denotes the 4-neighborhood of pixel p . Note that this construction connects all pixels in the image into a (planar) graph. λ is a strictly positive hyperparameter that regulates the update and ensures stability. When using \mathcal{N}_4^p connectivity it should be set to $\lambda < \frac{1}{4}$. The function $c : (\mathbb{R}^C, \mathbb{R}^C) \rightarrow (0, 1)$ produces diffusion coefficients for neighboring pairs of pixels based on their values in the guide. Following [30], we use

$$c(\mathbf{g}^p, \mathbf{g}^n) = \frac{\kappa^2}{\kappa^2 + \|\mathbf{g}^p - \mathbf{g}^n\|_2^2} \quad (2)$$

where κ is a hyperparameter that regulates the sensitivity to gradients in \mathbf{G} . Note that c is symmetrical, $c(\mathbf{g}^p, \mathbf{g}^n) = c(\mathbf{g}^n, \mathbf{g}^p)$. Traditional (non-guided) anisotropic diffusion is a special case of the formulation above where $\mathbf{G} = \mathbf{Y}_{t-1}$.

When applied to a single image, anisotropic diffusion has edge-enhancing properties [30]. In guided diffusion,

where the diffusion weights are computed from a separate guide image, the diffusion process transfers edges from the guide to the target image [2, 4, 23]. This is the property that motivates its use in our guided depth super-resolution framework: the diffusion allows us to precisely recover depth discontinuities in the upsampling result.

By itself, diffusion does not take into account the constraint provided by the source \mathbf{S} . *I.e.*, with the machinery introduced so far \mathbf{Y}_t would approach a constant image with pixel values $\mu(\mathbf{Y}_0)$ as $t \rightarrow \infty$, losing all information. To tie the output to the source image \mathbf{S} , every diffusion step is followed by an adjustment step that restores compatibility with \mathbf{S} . Doing so guarantees that the output of every iteration, and therefore also the equilibrium $\lim_{t \rightarrow \infty} \mathbf{Y}_t$, is a valid upsampling of \mathbf{S} .

The adjustment is done by simply re-scaling patches of $\hat{\mathbf{Y}}_t$ such that, when downsampled to the source resolution, they exactly match \mathbf{S} . More formally, the adjustment step can be written as

$$\mathbf{Y}_t = \hat{\mathbf{Y}}_t \cdot \underbrace{\text{up} \left(\frac{\mathbf{S}}{\text{down}(\hat{\mathbf{Y}}_t)} \right)}_{\mathbf{R}_t}, \quad (3)$$

where down denotes a linear downsampling operator, and up denotes nearest neighbor upsampling. \mathbf{r}_t and \mathbf{R}_t are the adjustment ratios at different scales. After this adjustment step, it is guaranteed that the target matches the source at a lower scale, *i.e.*, $\text{down}(\mathbf{Y}_t) = \mathbf{S}$.

An illustration of the method is displayed in Fig. 2. Furthermore, we show the evolution during the diffusion process for a 1D example in Fig. 3. Gradients in the diffused signal quickly dissipate where gradient in the guide are low, but diffusion (almost) stops at edges of the guide. Consistency with the source \mathbf{S} is preserved throughout the process.

3.2. Deep Feature-Guided Anisotropic Diffusion

Recent work has shown that the context features extracted by with CNNs over large receptive fields can massively boost super-resolution based on low-order graphs [5]. A main message of our paper is that this idea is even more powerful when combined with diffusion on the graph.

Let $F : \mathbb{R}^{H \times W \times C_1} \rightarrow \mathbb{R}^{H \times W \times C_2}$ be a neural feature extractor. In our experiments, F is a U-Net [33] with ResNet-50 [12] backbone pretrained on ImageNet [7] and $C_2 = 64$, but any other neural architecture could be used, as long as the spatial dimensions of the output match those of the input. While F could be applied directly to the RGB guide \mathbf{G} , we concatenate the upsampled source $\text{up}(\mathbf{S})$ as a fourth channel to support object separation based on coarse depth cues, so $C_1 = 4$.

Previous work that connected guided anisotropic diffusion with deep learning [2, 4] was restricted to using it only

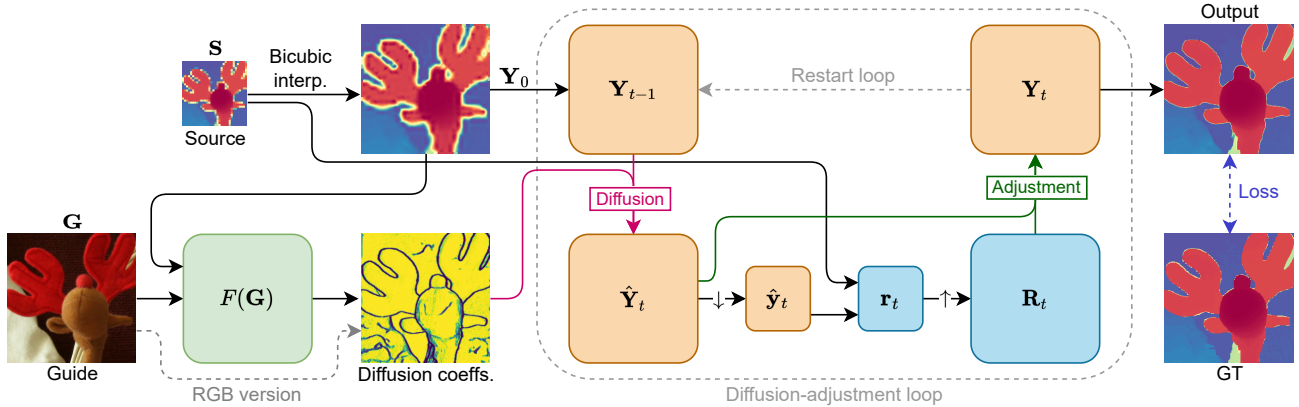


Figure 2. The diffusion step performs anisotropic diffusion on the depth image using diffusion weights given by a CNN. The adjustment step ensures that the diffused result matches the source image at its original resolution. Gradients from the loss function can be back-propagated back to the CNN, making the system end-to-end trainable.

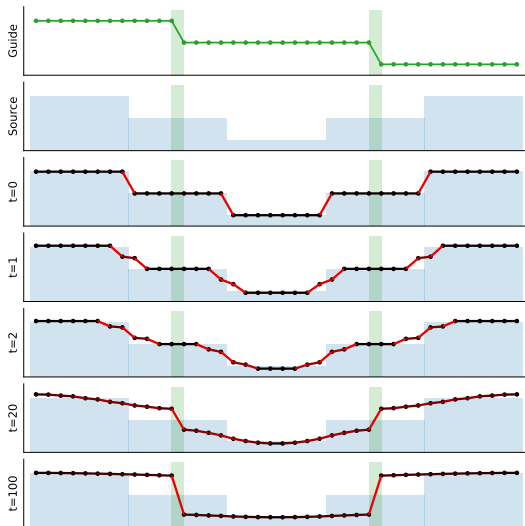


Figure 3. Diffusion diminishes gradients in the target signal (highlighted in red). Strong gradients in the guide signal (marked in green) impede diffusion and are thus transferred to the target. Additionally, the target is constrained to adhere to the low-resolution source (depicted in blue).

as a post-processing step at inference time, due to excessively high memory requirements, and a risk of vanishing/exploding gradients caused by the iterative nature of the algorithm. We found that it is in fact possible to propagate gradients through the diffusion process with the following scheme:

1. We compute the $c(\mathbf{g}^p, \mathbf{g}^n)$ only once before diffusion starts and freeze them. This is in contrast to earlier attempts that diffused the guide \mathbf{G} alongside \mathbf{Y} and had to update the $c(\mathbf{g}^p, \mathbf{g}^n)$ at every iteration [2, 4];
2. We back-propagate gradients only through the last N_{grad} diffusion-adjustment iterations (in practice

$N_{\text{grad}} \approx 10^3$). *I.e.*, during training the feature extractor receives training signals from the later stages of the diffusion process.

With these modifications, we can effectively back-propagate gradients all the way to the feature extractor F and train the entire pipeline end-to-end. Furthermore, note that κ in Eq. (2) is also trainable, and therefore does not need to be chosen manually.

At training time we use a random number of iterations without gradient before computing iterations with gradient. We define N_{pre} as the maximum number of iterations without gradient updates at training time and N_{grad} as the number of iterations with gradients. We found that this randomization helps to speed up the convergence at inference time (see supplementary material). For inference we use a constant $N = N_{\text{pre}} + N_{\text{grad}}$ iterations. The rationale behind using the last N_{grad} iterations for the computation of gradients is that we are interested in the steady-state equilibrium of the system, and therefore need a high number of iterations. We study the effect of N_{pre} and N_{grad} and the associated computational costs in Section 4, and find that there is little gain for $N_{\text{grad}} > 1024$, a setting at which the system can still be trained on a single GPU. Finally, we point out that while the algorithm needs a relatively large number of iterations (we find 8000 to be a suitable number), the diffusion and adjustment operators are extremely lightweight and parallelizable, and as a result, our algorithm is still faster than other optimization approaches.

4. Experiments

We evaluate our method on three different RGB-D datasets commonly used for super-resolution, namely Middlebury 2005-2014 [14, 35–38], NYUv2 [39], and DIML [3, 19–21]. We closely follow the setup of [5], including the

	×4	×8	×16	×32
	MSE / MAE	MSE / MAE	MSE / MAE	MSE / MAE
Middlebury				
MSG	4.13 / 0.22	10.5 / 0.43	34.2 / 1.06	92.9 / 2.55
DKN	4.29 / 0.18	11.2 / 0.38	47.6 / 1.42	119 / 3.35
FDKN	3.60 / 0.16	10.4 / 0.37	38.5 / 1.18	112 / 3.23
PMBA	4.72 / 0.25	9.48 / 0.38	30.6 / 0.89	175 / 5.15
FDSR	7.72 / 0.35	23.2 / 0.69	55.4 / 1.51	179 / 4.11
LGR	3.04 / 0.13	7.26 / 0.24	24.7 / 0.67	63.4 / 1.75
DADA	2.52 / 0.11	5.63 / 0.20	15.6 / 0.47	47.6 / 1.35
NYUv2				
MSG	6.85 / 0.81	24.1 / 1.66	84.5 / 3.35	262 / 6.70
DKN	11.4 / 1.03	29.8 / 1.82	115 / 4.01	364 / 8.33
FDKN	8.07 / 0.85	29.9 / 1.80	113 / 3.95	365 / 8.39
PMBA	10.8 / 0.93	31.5 / 1.79	84.9 / 3.26	449 / 9.11
FDSR	10.5 / 0.94	35.4 / 1.96	179 / 4.68	771 / 11.8
LGR	6.45 / 0.73	19.6 / 1.42	67.5 / 2.90	253 / 6.50
DADA	4.87 / 0.64	17.1 / 1.33	59.2 / 2.65	223 / 5.76
DIML				
MSG	1.73 / 0.22	4.13 / 0.40	13.0 / 0.93	55.0 / 2.56
DKN	3.47 / 0.33	5.47 / 0.45	19.3 / 1.20	91.5 / 3.75
FDKN	2.2 / 0.23	5.95 / 0.47	20.8 / 1.24	89.9 / 3.75
PMBA	3.05 / 0.31	5.87 / 0.47	13.8 / 0.87	55.1 / 2.30
FDSR	2.75 / 0.29	8.40 / 0.66	32.9 / 1.66	124 / 4.39
LGR	1.68 / 0.20	3.51 / 0.31	9.45 / 0.68	57.0 / 2.47
DADA	1.30 / 0.17	2.87 / 0.27	7.75 / 0.60	38.6 / 1.90

Table 1. Performance comparison of learning-based methods in terms of MSE (in cm^2) and MAE (in cm). DADA consistently outperforms all other methods, especially at large scaling factors. We report the variability of our method in the supplementary material.

train/validation/test splits, data preprocessing and usage of evaluation metrics.

Middlebury [14, 35–38] consists of 50 photogrammetrically created high-resolution depth maps and their associated RGB images from the years from 2005 to 2014. Five samples are reserved for validation and testing each, while the rest is used for training. The depth maps contain data gaps, which are also present in the source images. This dataset provides the most accurate ground truth among the considered datasets.

NYUv2 [39] was captured with a Microsoft Kinect depth camera. It consists of a total of 1449 RGB-D images, from which 849 are used for training, and 300 each for validation and testing.

DIML [3, 19–21] contains 2 million RGB-D samples, from which we only use the high quality indoor samples that were acquired with a Microsoft Kinect depth camera. 1440 are training samples, 169 are validation samples and 503 are test samples.

4.1. Experimental Setup

We compare our deep anisotropic diffusion-adjustment network (*DADA*) against a broad, representative range of

	×4	×8	×16	×32
	MSE / MAE	MSE / MAE	MSE / MAE	MSE / MAE
Middlebury				
Bicubic	13.3 / 0.55	30.0 / 1.10	68.5 / 2.13	143 / 3.87
GF	33.3 / 1.27	40.5 / 1.49	67.4 / 2.21	134 / 3.82
SD	24.9 / 0.46	82.5 / 0.86	511 / 1.73	4062 / 3.37
PixT	39.8 / 0.79	32.7 / 0.82	41.5 / 1.24	107 / 2.71
LGR [†]	14.8 / 0.42	68.3 / 0.83	297 / 1.69	897 / 3.31
DADA[†]	11.1 / 0.40	18.9 / 0.70	36.7 / 1.30	84.1 / 2.58
NYUv2				
Bicubic	31.9 / 1.59	89.9 / 3.17	242 / 6.01	588 / 10.5
GF	114 / 3.91	142 / 4.47	249 / 6.34	556 / 10.4
SD	36.0 / 1.31	105 / 2.57	533 / 5.07	3246 / 10.4
PixT	112 / 3.61	122 / 3.86	219 / 5.40	759 / 11.6
LGR [†]	19.0 / 1.11	68.4 / 2.30	264 / 4.56	790 / 9.31
DADA[†]	18.1 / 1.05	49.9 / 2.05	125 / 3.88	328 / 7.50
DIML				
Bicubic	10.4 / 0.63	28.6 / 1.32	73.2 / 2.68	187 / 5.37
GF	25.6 / 1.45	34.1 / 1.77	66.3 / 2.74	165 / 5.18
SD	10.5 / 0.40	44.9 / 0.83	411 / 1.91	5905 / 4.45
PixT	20.7 / 1.15	23.0 / 1.26	39.3 / 1.78	141 / 4.19
LGR [†]	7.02 / 0.35	15.2 / 0.67	133 / 1.72	815 / 3.98
DADA[†]	4.40 / 0.28	9.35 / 0.51	21.3 / 1.15	72.6 / 3.02

Table 2. Performance comparison of learning-free methods in terms of MSE (in cm^2) and MAE (in cm). Our diffusion-adjustment scheme, without feature learning, outperforms other learning-free methods. [†] marks learning-free variants of hybrid methods.

guided super-resolution methods, both learning-based and learning-free. We also include simple non-guided bicubic upsampling (*Bicubic*) [17] as a baseline and sanity check. The learning-based methods we consider are *MSG-Net* (*MSG*) [15], Deformable Kernel Network (*DKN*) [18], Fast Deformable Kernel Network (*FDKN*) [18], Fast Depth Super-Resolution (*FDSR*) [13], PMBANet (*PMBA*) [47] and Learned Graph Regularizer (*LGR*) [5]. Learning-free methods in our evaluation are the guided filtering (*GF*) [11], Static/Dynamic filtering (*SD*) [10], Pixtransform (*PixT*) [6], and a version of *LGR* [5] that is based on raw RGB values in the guide, rather than learned features. In much the same way, we also run our diffusion-adjustment scheme with diffusion weights derived from raw RGB values. For upsampling factors of $\times 4$ to $\times 16$, the scores are taken directly from [5], for scale $\times 32$ we have generated all results ourselves, following the setup described in [5] and using their open-source code base. As error metrics, we show both the mean squared error (MSE) and the mean absolute error (MAE) of the predicted depth images.

Our experiments were conducted using PyTorch [28] and we train all methods, including our own, with the L_1 loss function. Further details regarding hyper-parameters for training will be provided in the supplementary material and code to ensure reproducibility. For the learning-free variant of *DADA* we set $\kappa = 0.03$ (see Eq. (2)).

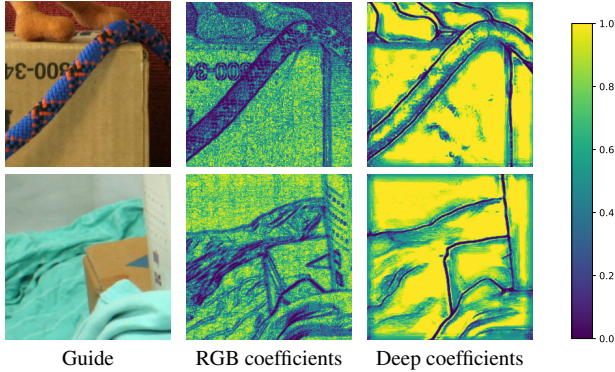


Figure 4. Diffusion coefficients calculated from raw RGB values and from deep features. The CNN focuses on object discontinuities rather than texture edges. Note the text on the box, top left.

4.2. Results

Table 1 shows quantitative results on all three datasets for the learned methods, while Tab. 2 shows the results for the non-learned ones. Our proposed method obtains the best performance in terms of both MSE and MAE on all three datasets, and across the full range of upsampling factors, with one single exception (at factor $\times 16$ on Middlebury, it is 2nd-best). The results underline the general trend towards learned, high-level feature representations: generally speaking, the learned methods clearly outperform the learning-free ones. In particular, for LGR as well as for our DADA, the versions with learned features easily beat their RGB-based counterparts. Having said that, we note that many learning-based methods seem to struggle with large-scale differences between the source and the guide (respectively, the target): the gap between learned and non-learned methods shrinks with increasing upsampling factors, and in fact, the performance of several dedicated super-resolution methods (learning-free ones, but also learned ones) falls below that of naïve bicubic upsampling.

RMSE curves for different scaling factors using the Middlebury dataset can be seen in Fig. 7. The plots for the other datasets follow the same trends and will be included in the supplementary material. As seen in these results, the performances of several of the considered methods degrade severely for larger scaling factors. In fact, for scale $\times 32$ many of the methods perform worse than simple bicubic interpolation. The results also show that for the lowest scaling factor ($\times 4$) many learned methods achieve similar performance, suggesting that guided super-resolution for moderate scale factors is reaching saturation, at least w.r.t. the currently available training and test data.

Figure 5 contains results from several of the compared methods on all three datasets with scaling factors of $\times 16$ and $\times 32$. The residual images make it clear that DADA is superior to all other methods, including LGR – the previous

		DIML			Middlebury		
		MSE	MAE	MSE [‡]	MSE	MAE	MSE [‡]
$\times 8$	MSG	5.76	0.51	6.16	11.0	0.54	5.01
	FDKN	6.74	0.53	0.20	10.0	0.43	0.20
	PMBA	7.35	0.59	0.04	9.62	0.46	0.06
	FDSR	7.73	0.74	0.45	18.4	0.73	7.20
	LGR	4.95	0.40	0.002	8.25	0.35	0.001
	DADA	4.58	0.38	0.0	7.82	0.35	0.0
$\times 32$	MSG	60.2	2.82	0.59	68.5	2.31	0.46
	FDKN	97.2	3.91	1.09	98.3	3.05	0.68
	PMBA	312	4.96	1.24	118	3.71	6.62
	FDSR	239	6.45	3.70	177	4.67	2.47
	LGR	64.5	2.83	0.001	70.7	2.28	0.008
	DADA	56.9	2.35	0.0	52.3	1.67	0.0

Table 3. Cross-dataset generalization performance. All learned methods were trained on NYUv2, then tested on DIML and Middlebury, with scaling factors of $\times 8$ and $\times 32$. Errors are in cm^2 for the MSE and in cm for the MAE. The low-resolution MSE[‡], *i.e.*, after downsampling the predicted target, indicates inconsistency with the forward model (linear downsampling).

best. Strict constraint to the source is especially advantageous in regions with no depth discontinuities. DADA is nevertheless able to produce sharp edges where necessary. Encircled depth discontinuities seem to be the most challenging cases for all methods. Although DADA still outperforms other methods, those situations still offer room for improvement.

Our quantitative and qualitative results both show the superiority of the results DADA yields when compared to all other methods. Our results also lend further support to the findings of LGR [5] that theoretically founded, optimization-based methods still offer excellent performance within-domain and especially out-of-domain if (and only if) they are integrated with the large receptive field and superior perceptual grouping abilities of modern, deep feature extractors. Having said that, we find that DADA surpasses LGR in terms of both reconstruction error and computational cost (memory consumption as well as runtime), especially at high upsampling factors.

In Figure 4 we show a pair of diffusion coefficients derived from deep features and contrast them with coefficients that are directly derived from the RGB values. The learned coefficients are more robust to printed text and logos, image noise, and textured surfaces, whereas the unlearned ones are overly sensitive to those.

Experience suggests that deep learning methods tend to (over-)fit to dataset-specific characteristics and generalize poorly from one dataset to another. For LGR, which like our method explicitly enforces consistency between the (down-sampled) target and the source, the authors in fact demonstrate improved generalization compared to purely learned frameworks. They attribute this behavior to the consistency constraint, which is data-independent and therefore not affected by domain gaps [5]. We have also explored the gen-

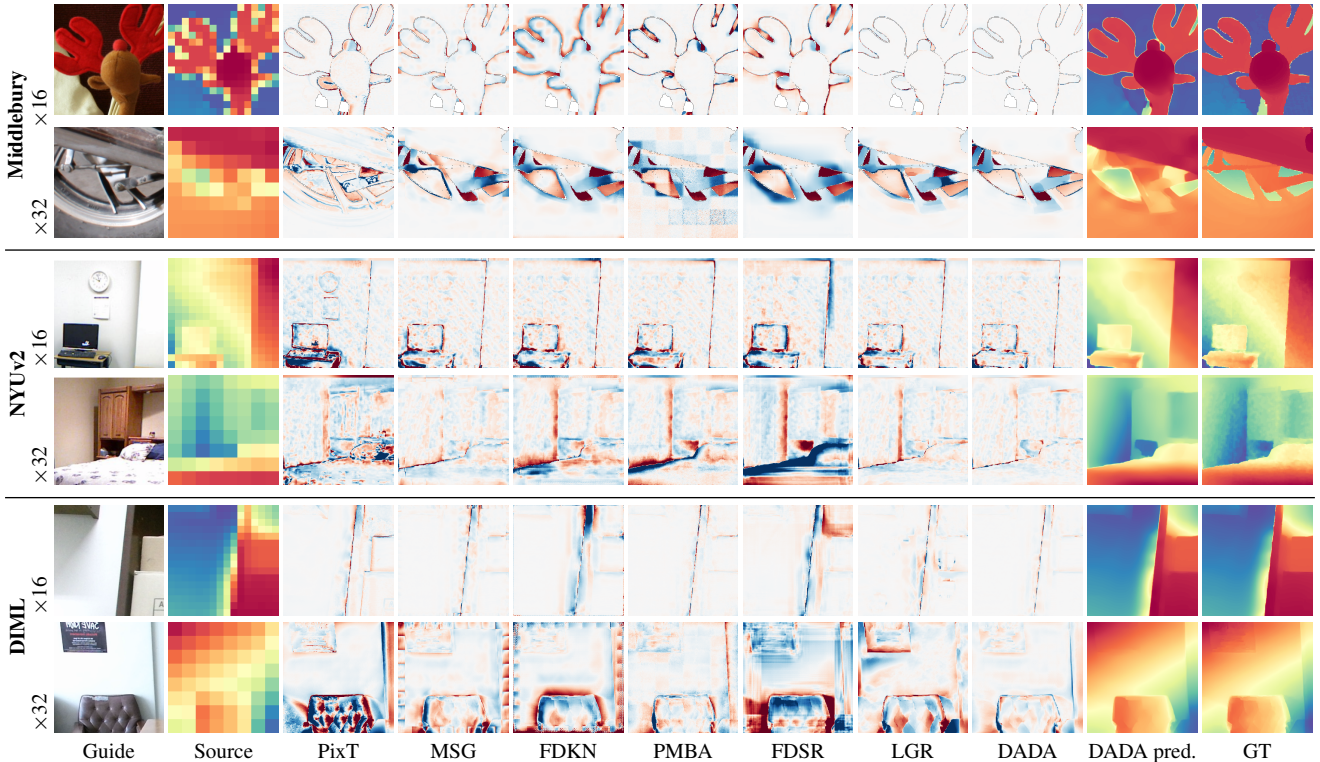


Figure 5. Error maps for different guided super-resolution methods. Blue denotes under-estimated depth, red denotes over-estimation. All errors maps in a row have the same color scale. The last two columns juxtapose our predictions and the ground truth.

eralization of our method across datasets, and report the results in Table 3. For these experiments, we train on the NYUv2 dataset for a super-resolution factor of $\times 8$ and $\times 32$, and test the resulting network on Middlebury and DIML. Also for this task, our method outperforms all competitors, including LGR, which ranks second.

4.3. Ablations

In this section, we experimentally study the effects of different parameters and modules on the results of our super-resolution method. We run these experiments on the Middlebury dataset with upsampling factor $\times 8$, which we found to be a representative setting. The memory requirements reported below always refer to a training batch consisting of 8 images of size 256×256 .

Initialization. Our default initialization scheme for the super-resolved image (bicubic upsampling of the guide) and for the feature extractor (Imagenet pretraining) are not the only options. Compared to the base case ($MSE = 5.67$, $MAE = 0.199$), we explore the two simplifications: (1) random initialization of the ResNet-50 backbone weights instead of ImageNet pretraining (which yields ($MSE = 5.92$, $MAE = 0.211$); and (2) constant initialization of \mathbf{Y}_0 rather than bicubic interpolation of \mathbf{S} (which yields ($MSE = 5.64$, $MAE = 0.200$). In both cases, we ob-

serve only tiny differences w.r.t. the base case. With regard to the backbone weights, this suggests that the available amount of training data is sufficient to learn a suitable feature representation. For the initialization of the target image \mathbf{Y}_0 , it is not surprising that this has almost no effect on the result. It is easy to see that the first diffusion step will not alter the constant input and that the output \mathbf{Y}_1 of the first adjustment step will then be the nearest-neighbor upsampling of \mathbf{S} . Meaning that after a single iteration of the diffusion-adjustment loop, the constant initialization has largely caught up with the seemingly more sophisticated bicubic one (and in fronto-parallel areas perhaps even has overtaken it).

Number of training iterations with gradient. As explained above, due to memory constraints we limit the number of diffusion iterations that propagate gradients into the feature extractor. In Figure 6a, we explore how this affects feature learning. Memory requirements increase with N_{grad} , ranging from 11 GB for $N_{\text{grad}} = 32$ up to 23 GB for $N_{\text{grad}} = 1024$. We find that larger values of N_{grad} generally lead to better performance, likely due to the gradients being more representative of the equilibrium point of the diffusion-adjustment process.



Figure 6. Plots for ablation experiments that explore the effects of different numbers of iterations in the diffusion-adjustment loop. We generally see an increase in performance for larger numbers of iterations, albeit at higher computational cost.

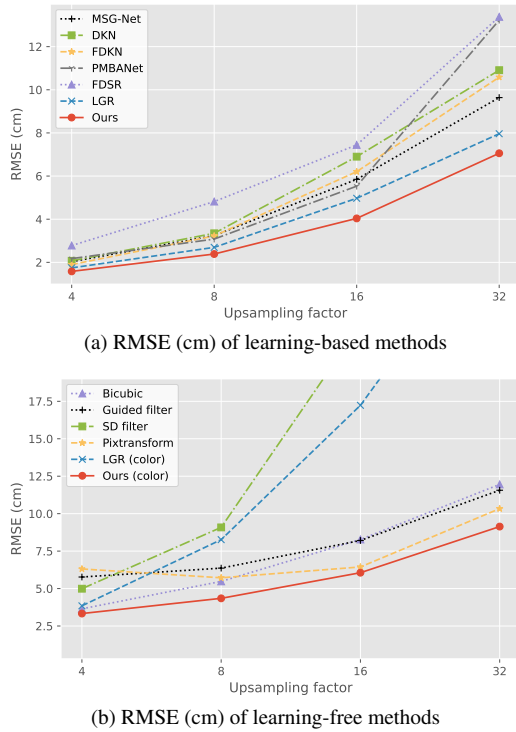


Figure 7. RMSE values at different upsampling factors, for the Middlebury dataset. Our method consistently achieves the lowest error.

Number of training iterations without gradient. A main hyper-parameter of the (asymptotic) diffusion process is the total number of iterations. For a fixed setting of N_{grad} this is, in our scheme, governed by the number of preceding gradient-free iterations, N_{pre} . We vary that value and once again we find that, as expected, larger values generally lead to better results, likely because the gradients are computed from a state closer to the equilibrium point. We see little further gain after $N_{\text{pre}} = 8000$.

Number of iterations at test time. Once trained, we are free to decide for how long we run the diffusion process at

inference time. In Figure 6c we explore how the number of iterations at test time impacts the results. Once more, higher iteration counts N lead to better results, as the process converges further towards the equilibrium point $\lim_{t \rightarrow \infty} \mathbf{Y}_t$. We see little improvement for $N > 8000$. As computing time (also shown in the graph) scales linearly with N , we see no reason to go beyond $N = 8000$.

5. Conclusion

We have presented DADA, a simple yet extremely effective approach to guided super-resolution of depth images. The method links guided anisotropic diffusion with deep learning. The diffusion part offers a computationally efficient optimization framework to optimally align the prediction with the individual input images at test time and explicitly constrains the results to comply with the source image. The deep learning-based feature extractor harnesses the representation power of CNNs to optimally inform the diffusion process, by injecting context cues collected over large receptive fields into the diffusion weights. In our experiments, DADA reaches top performance on three popular datasets and outperforms the prior state-of-the-art, including learning-based, optimization-based as well as hybrid methods. We also found that DADA is fairly robust in terms of engineering details: it is barely affected by the choice of pretraining and initialization scheme and is relatively tolerant against hyper-parameter changes.

An interesting direction for future work will be to formulate the diffusion-adjustment iteration as an ordinary differential equation. This could potentially make it possible to back-propagate through the entire diffusion process with the adjoint sensitivity method [31], rather than unrolling a limited number of iterations. At a conceptual level, we hope that our work motivates further research about the combination of classical optimization-based computer vision and modern, deep learning-based image representations. In line with others [5,32], our work suggests that such hybrid methods hold great potential, not only for super-resolution, but also for other low-level vision tasks, and perhaps beyond.

References

- [1] Feras Almasri and Olivier Debeir. RGB guided thermal super-resolution enhancement. *Int'l Conf on Cloud Computing Technologies and Applications*, 2018. 1
- [2] Rodrigo Caye Daudt, Bertrand Le Saux, Alexandre Boulch, and Yann Gousseau. Guided anisotropic diffusion and iterative learning for weakly supervised change detection. *CVPR Workshops*, 2019. 2, 3, 4
- [3] Jaehoon Cho, Dongbo Min, Youngjung Kim, and Kwanghoon Sohn. Deep monocular depth estimation leveraging a large-scale outdoor stereo dataset. *Expert Systems with Applications*, 178:114877, 2021. 4, 5
- [4] Rodrigo Caye Daudt, Bertrand Le Saux, Alexandre Boulch, and Yann Gousseau. Weakly supervised change detection using guided anisotropic diffusion. *Machine Learning*, pages 1–27, 2021. 2, 3, 4
- [5] Riccardo de Lutio, Alexander Becker, Stefano D’Aronco, Stefania Russo, Jan D Wegner, and Konrad Schindler. Learning graph regularisation for guided super-resolution. *CVPR*, 2022. 1, 2, 3, 4, 5, 6, 8
- [6] Riccardo de Lutio, Stefano D’Aronco, Jan Dirk Wegner, and Konrad Schindler. Guided super-resolution as pixel-to-pixel transformation. *ICCV*, 2019. 2, 5
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. *CVPR*, 2009. 3
- [8] James Diebel and Sebastian Thrun. An application of Markov random fields to range sensing. *NIPS*, 2005. 2, 3
- [9] Ivan Eichhardt, Dmitry Chetverikov, and Zsolt Janko. Image-guided ToF depth upsampling: a survey. *Machine Vision and Applications*, 28(3):267–282, 2017. 1
- [10] Bumsub Ham, Minsu Cho, and Jean Ponce. Robust guided image filtering using nonconvex potentials. *IEEE TPAMI*, 40(1):192–207, 2017. 1, 2, 5
- [11] Kaiming He, Jian Sun, and Xiaoou Tang. Guided image filtering. *IEEE TPAMI*, 35(6):1397–1409, 2012. 2, 5
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CVPR*, 2016. 3
- [13] Lingzhi He, Hongguang Zhu, Feng Li, Huihui Bai, Runmin Cong, Chunjie Zhang, Chunyu Lin, Meiqin Liu, and Yao Zhao. Towards fast and accurate real-world depth super-resolution: Benchmark dataset and baseline. *CVPR*, 2021. 2, 5
- [14] Heiko Hirschmüller and Daniel Scharstein. Evaluation of cost functions for stereo matching. *CVPR*, 2007. 4, 5
- [15] Tak-Wai Hui, Chen Change Loy, and Xiaoou Tang. Depth map super-resolution by deep multi-scale guidance. *ECCV*, 2016. 1, 2, 5
- [16] David Izraelevitz. Model-based multispectral sharpening. *Proceedings of SPIE*, 1994. 1
- [17] Robert Keys. Cubic convolution interpolation for digital image processing. *IEEE T Acoustics, Speech, and Signal processing*, 29(6):1153–1160, 1981. 5
- [18] Beomjun Kim, Jean Ponce, and Bumsub Ham. Deformable kernel networks for joint image filtering. *IJCV*, 129(2):579–600, 2021. 2, 5
- [19] Sunok Kim, Dongbo Min, Bumsub Ham, Seungryong Kim, and Kwanghoon Sohn. Deep stereo confidence prediction for depth estimation. *ICIP*, 2017. 4, 5
- [20] Youngjung Kim, Bumsub Ham, Changjae Oh, and Kwanghoon Sohn. Structure selective depth superresolution for RGB-D cameras. *IEEE TIP*, 25(11):5227–5238, 2016. 2, 4, 5
- [21] Youngjung Kim, Hyunjoo Jung, Dongbo Min, and Kwanghoon Sohn. Deep monocular depth estimation via integration of global and local predictions. *IEEE TIP*, 27(8):4131–4144, 2018. 4, 5
- [22] Charis Lanaras, José Bioucas-Dias, Silvano Galliani, Emmanuel Baltsavias, and Konrad Schindler. Super-resolution of Sentinel-2 images: Learning a globally applicable deep neural network. *ISPRS J Photogrammetry and Remote Sensing*, 146:305–319, 2018. 1
- [23] Junyi Liu and Xiaojin Gong. Guided depth enhancement via anisotropic diffusion. *Pacific-Rim Conf on Multimedia*, 2013. 2, 3
- [24] Ziyang Ma, Kaiming He, Yichen Wei, Jian Sun, and Enhua Wu. Constant time weighted median filtering for stereo matching and beyond. *ICCV*, 2013. 2
- [25] Oisín Mac Aodha, Neill D. F. Campbell, Arun Nair, and Gabriel J. Brostow. Patch based synthesis for single depth image super-resolution. *ECCV*, 2012. 2
- [26] Dongbo Min, Jiangbo Lu, and Minh N Do. Depth video enhancement based on weighted mode filtering. *IEEE TIP*, 21(3):1176–1190, 2011. 2
- [27] Min-Gyu Park and Kuk-Jin Yoon. As-planar-as-possible depth map estimation. *Computer Vision and Image Understanding*, 181:50–59, 2019. 3
- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. *NeurIPS*, 2019. 5
- [29] Tim J. Patterson, Michael E. Bullock, and Alan D. Wada. Multispectral band sharpening using pseudoinverse estimation and fuzzy reasoning. *Proceedings of SPIE*, 1992. 1
- [30] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE TPAMI*, 12(7):629–639, 1990. 2, 3
- [31] Lev S. Pontryagin. *Mathematical theory of optimal processes*. CRC Press, 1987. 8
- [32] Gernot Riegler, David Ferstl, Matthias Rüther, and Horst Bischof. A deep primal-dual network for guided depth super-resolution. *BMVC*, 2016. 1, 2, 8
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *MICCAI*, 2015. 3
- [34] Mattia Rossi, Mireille El Gheche, Andreas Kuhn, and Pascal Frossard. Joint graph-based depth refinement and normal estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12154–12163, 2020. 3
- [35] Daniel Scharstein, Heiko Hirschmüller, York Kitajima, Greg Krathwohl, Nera Nešić, Xi Wang, and Porter West-

- ling. High-resolution stereo datasets with subpixel-accurate ground truth. *GCPR*, 2014. 4, 5
- [36] Daniel Scharstein and Chris Pal. Learning conditional random fields for stereo. *CVPR*, 2007. 4, 5
- [37] Daniel Scharstein and Richard Szeliski. High-accuracy stereo depth maps using structured light. *CVPR*, 2001. 4, 5
- [38] Daniel Scharstein, Richard Szeliski, and Ramin Zabih. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *IJCV*, 47(1):7–42, 2002. 2, 4, 5
- [39] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *IEEE TPAMI*, 2017. 2, 4, 5
- [40] Xibin Song, Yuchao Dai, Dingfu Zhou, Liu Liu, Wei Li, Hongdong Li, and Ruigang Yang. Channel attention based iterative residual learning for depth map super-resolution. *CVPR*, 2020. 2
- [41] Tatsumi Uezato, Danfeng Hong, Naoto Yokoya, and Wei He. Guided deep decoder: Unsupervised image pair fusion. *ECCV*, 2020. 2
- [42] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. *CVPR*, 2018. 2
- [43] Jin Wang, Wei Xu, Jian-Feng Cai, Qing Zhu, Yunhui Shi, and Baocai Yin. Multi-direction dictionary learning based depth map super-resolution with autoregressive modeling. *IEEE T Multimedia*, 22(6):1470–1484, 2019. 2
- [44] Yang Wen, Bin Sheng, Ping Li, Weiyao Lin, and David Dagan Feng. Deep color guided coarse-to-fine convolutional network cascade for depth image super-resolution. *IEEE TIP*, 28(2):994–1006, 2018. 2
- [45] Jun Xie, Rogerio Schmidt Feris, and Ming-Ting Sun. Edge-guided single depth image super resolution. *IEEE TIP*, 25(1):428–438, 2015. 2
- [46] Qingxiong Yang, Ruigang Yang, James Davis, and David Nistér. Spatial-depth super resolution for range images. *CVPR*, 2007. 2
- [47] Xinchun Ye, Baoli Sun, Zhihui Wang, Jingyu Yang, Rui Xu, Haojie Li, and Baopu Li. PMBANet: Progressive multi-branch aggregation network for scene depth super-resolution. *IEEE TIP*, 2020. 2, 5
- [48] Yongqin Zhang, Feng Shi, Jian Cheng, Li Wang, Pew-Thian Yap, and Dinggang Shen. Longitudinally guided super-resolution of neonatal brain magnetic resonance images. *IEEE T Cybernetics*, 49(2):662–674, 2018. 1