

A Deep Primal-Dual Network for Guided Depth Super-Resolution

Gernot Riegler
riegler@icg.tugraz.at

David Ferstl
ferstl@icg.tugraz.at

Matthias Rütther
ruether@icg.tugraz.at

Horst Bischof
bischof@icg.tugraz.at

Institute for Computer Graphics and
Vision
Graz University of Technology
Austria

Abstract

In this paper we present a novel method to increase the spatial resolution of depth images. We combine a deep fully convolutional network with a non-local variational method in a *deep primal-dual network*. The joint network computes a noise-free, high-resolution estimate from a noisy, low-resolution input depth map. Additionally, a high-resolution intensity image is used to guide the reconstruction in the network. By unrolling the optimization steps of a first-order primal-dual algorithm and formulating it as a network, we can train our joint method end-to-end. This not only enables us to learn the weights of the fully convolutional network, but also to optimize all parameters of the variational method and its optimization procedure. The training of such a deep network requires a large dataset for supervision. Therefore, we generate high-quality depth maps and corresponding color images with a physically based renderer. In an exhaustive evaluation we show that our method outperforms the state-of-the-art on multiple benchmarks.

1 Introduction

In the last decade, a large range of affordable depth sensors became available on the mass market. This has pushed research to develop a variety of different applications based on these sensors. Especially active sensors based on structured light, or Time-of-Flight (ToF) measurements, enabled novel computer vision applications such as robot navigation [1], human pose estimation [13, 33], and hand pose estimation [35, 36]. Despite their success, these sensors suffer from a low spatial resolution and a high acquisition noise due to the physical limitations of the measurement principles. Even very recent ToF sensors have a spatial resolution of only 120×160 pixels [26]. Therefore, more and more approaches are proposed to improve the resolution and to suppress the noise of these depth cameras. Usually these depth cameras are equipped with an additional intensity camera of higher resolution. Hence, a very common practice [7, 25, 41] is to utilize the high resolution (HR) intensity image as guidance. These approaches build upon the observation that depth discontinuities

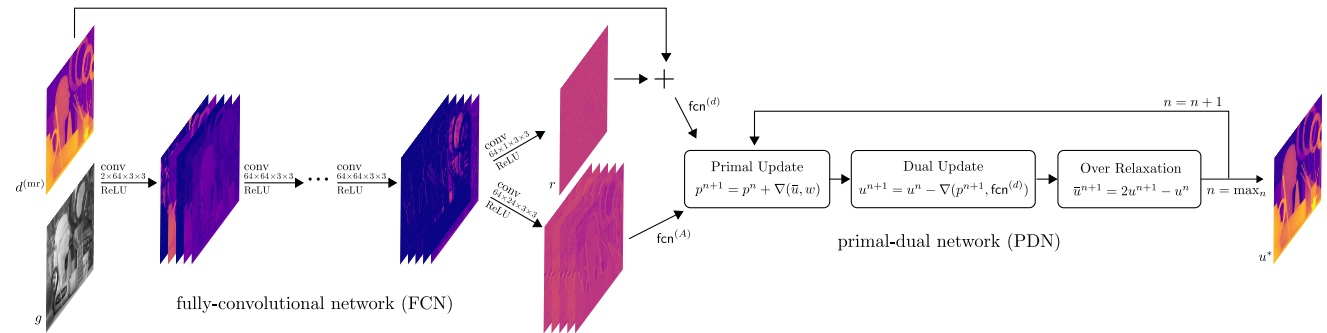


Figure 1: Our *deep primal-dual network* consists of two networks. A fully-convolutional network that computes a first HR estimate and weighting coefficients. Both outputs are then used in our primal-dual network, where we unroll the optimization steps of a non-local variational method that incorporates prior knowledge about the data modalities.

often occur at high intensity variations and that homogeneous areas in intensity images are also more likely to represent homogeneous areas in depth.

While the classical single image super-resolution for color images is dominated by machine learning approaches, *e.g.* [20, 30, 37], the field of depth super-resolution still mainly relies on Markov Random Field formulations [7], adaptive filters [41], or variational methods [10]. This is due to the lack of high quality training data in larger quantities, which is essential for large-scale machine learning methods. While it is quite easy to get a huge database of color image examples, *e.g.* from the web, there exists no equivalent source for depth data. One workaround [9, 23] is to densely reconstruct a 3D scene with KinectFusion [19] and facilitate these reconstructions as ground-truth. However, this also introduces artifacts in the training data, such as smoothed edges and the loss of fine details. Further, the scene preparation for reconstruction and the recording process itself are very time-consuming.

In this work we present a novel method based on machine learning for guided depth super-resolution, which combines the advantages of deep convolutional networks and variational methods. The training of this novel combination is enabled by creating a large corpus of high quality training data, which are automatically generated by rendering depth maps and corresponding color images from randomly placed and textured 3D objects in a virtual scene. This data is used to train our *deep primal-dual network* that maps low resolution (LR) and noisy depth maps to accurate HR estimates. The first part of the network consists of a series of fully convolutional layers to produce a guidance and rough super-resolved depth. This guidance and depth is used in a novel non-local variational model to optimize the final result. By unrolling the computation steps of a primal-dual algorithm [3] we formulate the variational optimization as a *primal-dual network*, where each numerical operation in this algorithm is defined as a layer in the network. In this way, our *deep primal-dual network* enables a joint optimization of all convolutional filter weights, the trade-off parameter of the variational cost function, and all hyper-parameters of the primal-dual algorithm.

The contribution of our work is three-fold and can be summarized as follows: (i) We extend our work of [31] by combining a deep fully convolutional network with a non-local primal-dual network that is trained end-to-end, shown in Section 3. Hence, we map a noisy, LR depth map along with a HR guidance image to an accurate HR estimate. (ii) We propose a framework based on the physically based Mitsuba renderer [39] to automatically generate high-quality depth maps with corresponding color images in large quantities which are used to train our model, shown in Section 4. (iii) The evaluations presented in Section 5 demonstrate the effectiveness of our method by outperforming state-of-the-art results on a set of standard synthetic and real-world benchmarks.

2 Related Work

Single-image super-resolution, *i.e.* enhancing the spatial resolution of an image, is a fundamental problem in low-level computer vision. It is inherently ill-posed, as several different HR images can map to the very same LR image. The field can be mainly divided into methods where an edge-preserving smoothness term is utilized [38], co-occurrences of patches within the same image are exploited [14], or, currently most successful, a mapping from LR to HR image patches is learned [8, 32, 37, 42].

Although the approaches for single-image super-resolution are quite general, models to increase the spatial resolution of depth maps differ. First, the modalities of depth maps are different than those in color images. While color images are characterized by high frequent textures and shading effects, depth data contains more noise, and consists of piece-wise affine regions and sharp edges at depth discontinuities. Second, training data for color images can be easily obtained, explaining the recent success of learning based approaches for single image super-resolution. Hence, several specific models have been proposed for depth super-resolution. In the seminal work of Diebel & Thrun [7] the super-resolution is formulated as a Markov Random Field (MRF) optimization, where the smoothness prior is weighted by the gradient magnitude of a guidance image. Park *et al.* [25] extend this MRF functional by incorporating a non-local means term to better preserve local structures of noisy data. In [41] Yang *et al.* propose an approach that builds upon the assumptions that surfaces are piecewise smooth and pixels with similar color have a similar depth. From this they derive a bilateral filter that is iteratively applied to the input depth map. Similarly, Chan *et al.* [4] present a modified bilateral filter to reduce artifacts in areas where a standard bilateral upsampling would cause a texture copy. A variational approach for guided depth super-resolution is proposed by Ferstl *et al.* [10]. They formulate the energy functional with an anisotropic Total Generalized Variation prior, which is weighted by the gradients in the guidance intensity image. Yang *et al.* [40] formulate the depth upsampling as a minimization of an adaptive color-guided auto-regressive model. One of the few learning based approaches is proposed by Kwon *et al.* [23]. They apply a multi-scale sparse coding approach to iteratively refine the LR depth map, and the HR data for training is acquired with KinectFusion [19].

While the approaches discussed above all utilize a HR guidance image, there exist also a few approaches that estimate the HR depth map without guidance. Aodha *et al.* [2] and Hornáček *et al.* [17] both utilize a MRF to fit HR candidate patches and differ in the search strategy to find similar patches. Aodha *et al.* exploit an external database of a few synthetic depth maps and Hornáček *et al.* search the 3D patches within the same depth map. A variational depth super-resolution model is presented by Ferstl *et al.* [11], where they use sparse coding to estimate the depth discontinuities in the HR depth map.

The integration of energy minimization models, like MRFs or variational methods, into deep networks recently gains a lot of interest. Chen *et al.* [5] show how to integrate a MRF on top of a deep network and train it with back-propagation. Similarly, Zheng *et al.* [43] unroll the computation steps of the mean field approximation [22] to optimize MRFs on top of a network for semantic segmentation. They show that the individual computation steps can be realised by operations in a convolutional network. One of the first integrations of variational models into convolutional networks is proposed by Ranftl & Pock [27] for foreground-background segmentation via implicit differentiation of the energy functional. Riegler *et al.* [29] use this formulation for guided depth denoising and super-resolution.

Our method is in the spirit of machine learning based approaches, especially it is related to the very deep network of [20]. To create a feasible amount of training data, we render

high quality depth maps and color images of randomly generated scenes in large quantities. Like [29], we combine the deep network with a variational approach, with the crucial difference that we do not implicitly differentiate the variational method, which drastically limits the choice of energy functionals, but unroll the steps of a fast optimization algorithm, as shown in [43] for MRFs. Finally, this work presents several improvements of our method presented in [31]. We demonstrate, how to incorporate an additional guidance image in our method and show that this is crucial for higher upsampling factors. Gathering training data for this scenario becomes also more difficult. We solve this problem by using a physically based renderer that produces high quality depth maps along with textured color images. Further, we evaluate different energy functionals for our method and show that a non-local Huber regularization term yields the best trade-off between accuracy and computational requirements for this task.

3 A Deep Primal-Dual Network

The proposed method consists of two main parts that are jointly trained end-to-end. The first one is a fully-convolutional network (FCN), which computes a HR estimate of the input depth map and input dependent weighting terms that are utilized in the subsequent primal-dual network (PDN). In the PDN we unroll the optimization procedure of a non-local variational model, namely of the first-order primal-dual algorithm [3]. The unrolling of the optimization steps enables us the integration of the variational model on top of the FCN and a joint training of both networks. A visual representation of our method is depicted in Fig. 1.

Let $d^{(\text{lr})}$ be the LR input depth map with $d^{(\text{lr})} \in \mathbb{R}^{\rho^{-1}M \times \rho^{-1}N}$, where ρ is the scale factor. The only pre-processing step of our method is an upsampling of $d^{(\text{lr})}$ via bilinear interpolation to the target resolution. This yields the mid-resolution input depth map $d^{(\text{mr})} \in \mathbb{R}^{M \times N}$. As an additional input we have an intensity image g as guidance that is given in the target resolution, $g \in \mathbb{R}^{M \times N}$. For brevity we will condense the mid-resolution depth map and the guidance image to an input sample denoted as $s = (d^{(\text{mr})}, g) \in \mathbb{R}^{2 \times M \times N}$. To train our method, we require a dataset $\{(s_k, t_k)\}_{k=1}^K$ of K input samples s_k and corresponding HR depth maps as targets $t_k = d_k^{(\text{hr})} \in \mathbb{R}^{M \times N}$. The goal of the model training is to find the optimal parameters $w^* = (w_{\text{fcn}}^*, w_{\text{pdn}}^*)$ of our model $f = \text{pdn}(\text{fcn}(s; w_{\text{fcn}}); w_{\text{pdn}})$ that minimize a loss function L over all K training samples:

$$w^* = \arg \min_w \sum_{k=1}^K L(f(s_k, w), t_k). \quad (1)$$

3.1 Fully Convolution Network

Inspired by [8, 20] we use a deep convolutional network to compute an initial high-resolution depth estimate given a noisy, LR input depth map $d^{(\text{lr})}$ together with a corresponding guidance image g . The network consists of 10 convolutional layers and rectified linear units (ReLU) [24] as activation functions. For the convolutional layers we use 3×3 filter kernels for the benefits discussed in [34] and in each convolutional layer we employ 64 feature maps, which results in a receptive field of 21×21 pixels of our FCN.

An important aspect of this network is that it does not directly compute the high-resolution depth map $d^{(\text{hr})}$, but the residual $r = d^{(\text{hr})} - d^{(\text{mr})}$ to the mid-resolution input $d^{(\text{mr})}$ as shown

in Fig. 1. After addition of the mid-resolution depth to the residual, the network’s high-resolution estimate is given by $\text{fcn}^{(d)}(s_k, w_{\text{fcn}}) = d_k^{(\text{mr})} + r_k$. For an uncluttered notation we will denote $\text{fcn}^{(d)}(s_k, w_{\text{fcn}})$ simply as $\text{fcn}_{s_k}^{(d)}$. The calculation of a residual is especially beneficial for convolutional networks, since it omits the need for the intermediate layers to carry the input information through the whole network, as shown in [20, 32, 37]. This is also related to the recently proposed residual networks for image classification [16].

Additional to the residual output, our network computes weighting coefficients $\text{fcn}_{s_k}^{(A)}$ for the subsequent PDN. In short, they represent information about depth discontinuities in the HR domain, but we will discuss this in more detail in the next Section.

3.2 Primal-Dual Network

As we will show in our evaluations, a FCN already delivers quite satisfying HR estimates for smaller scaling factors. However, depth dependent noise is still apparent in homogeneous regions. In this case variational methods are an ideal solution, since they introduce prior knowledge about the data modalities. In our method we combine both a FCN and a variational method to estimate sharp and noise-free results all over the image. This Section gives the insights into the variational model, how we realize this model as a network and how we combine it with the FCN into our complete *deep primal-dual network*.

The cost function of a variational method typically consists of a data term D , which penalizes the deviation from the initial solution and a regularization term R , where we can formulate smoothness assumptions. Hence, our variational model is given by

$$u_k^* = \arg \min_u \lambda D(u, \text{fcn}_{s_k}^{(d)}) + R(u, \text{fcn}_{s_k}^{(A)}), \quad (2)$$

where D and R are parametrized by the outputs of the FCN, $\lambda \in \mathbb{R}^+$ steers the weighting between the two terms and u^* is the minimizer of the cost function. The data term in our model penalizes the deviations from the FCN depth output and is defined as

$$D(u, \text{fcn}_{s_k}^{(d)}) = \frac{1}{2} \int_{\Omega} (u(x) - \text{fcn}_{s_k}^{(d)}(x))^2 dx. \quad (3)$$

Most regularization terms are based on first order smoothness assumptions, *e.g.* the Total Variation (TV) semi norm, $R(u) = \int_{\Omega} \|\nabla u\|_1 dx$. Although the TV model is able to estimate sharp object discontinuities in the depth map, it has two major disadvantages: (i) the ℓ_1 norm favors piecewise constant solutions resulting in piecewise fronto-parallel depth reconstructions. (ii) the ∇ -operator is not suitable to preserve small scale structures because it only penalizes the forward differences to its direct neighbors.

In our work we model the regularization as TV in a “larger” (non-local) neighborhood \mathcal{N} and further choose a more robust norm. The idea of a non-local regularization [12] is to incorporate a low level segmentation process into the variational model. This non-local regularization is defined as

$$R(u) = \int_{\Omega} \int_{\mathcal{N}(x)} w(x, y) |u(x) - u(y)|_{\varepsilon} dx dy, \text{ where } |x|_{\varepsilon} = [x \leq \varepsilon] \frac{|x|^2}{2\varepsilon} + [x > \varepsilon] (|x| - \frac{\varepsilon}{2}), \quad (4)$$

where the operator $|\cdot|_{\varepsilon}$ denotes the Huber norm [18]. The parameter $\varepsilon \in \mathbb{R}^+$ defines the threshold between the quadratic ℓ_2 and linear ℓ_1 norm. In contrast to the TV this allows smooth depth reconstruction while preserving sharp edges. We further call it non-local Huber

(NLH) regularization. One crucial part of this non-local regularization is the weighting factor $w(x, y) \in \mathbb{R}^{|\mathcal{N}| \times \Omega}$, which sets the penalty influence of every pixel $y \in \mathcal{N}(x)$ to the center x . The support weight $w(x, y)$ combines the value-similarities and the spatial distances

$$w(x, y) = \exp\left(-\frac{\Delta_d}{\sigma_d} - \frac{\Delta_a}{\sigma_v}\right), \quad (5)$$

where Δ_d denotes the Euclidean proximity $\|x - y\|_2$, which means with increasing distance to x the influence of the penalty decreases. Δ_a denotes the Euclidean affinity for example to a given guidance image, *i.e.* with increasing homogeneity in the guidance also the regularization increases. The scalars $\sigma_v, \sigma_d \in \mathbb{R}^+$ define the influence of each term.

In traditional non-local methods the affinity Δ_a is given by an intensity image g and results in $\Delta_a = \|g(x) - g(y)\|_2$. It has been shown that this is beneficial since high gradients in the intensity image and high depth disparities are likely to co-occur. The main disadvantage of this approach is that textured surfaces violate this assumption which subsequently leads to erroneous results. Obviously, the optimal guidance would be the high resolution depth image $d^{(hr)}$. Therefore, we use our FCN to directly train for the optimal support weights w . Further, since the proximity Δ_d is constant we only have to train for the affinity term which is defined by the FCN output $\Delta_a = \text{fcn}^{(A)}$. In Fig. 2 the difference between the non-local weight from an intensity image and from our learned FCN guidance is shown.

To optimize the variational model (2) we use the primal-dual scheme as proposed in [3]. After discretization of the continuous image space on a Cartesian grid $\Omega \mapsto \mathbb{R}^{M \times N}$ the derived convex-concave saddle-point problem with dual variable p is given by

$$\min_{u \in \mathbb{R}^{M \times N}} \max_{p \in \mathcal{P}} \left\{ \sum_{x \in \mathbb{R}^{M \times N}} \sum_{y \in \mathcal{N}(x)} (u(x) - u(y)) p(x, y) + \frac{\lambda}{2} \|u - \text{fcn}^{(d)}\|_2^2 - \frac{\varepsilon}{2} \|p\|_2^2 \right\} \quad (6)$$

$$\text{s.t. } p \in \mathcal{P} = \{p: \mathbb{R}^{M \times N} \mapsto \mathbb{R} \mid |p(x, y)| \leq w(x, y), \forall x \in \mathbb{R}^{M \times N}, y \in \mathcal{N}(x)\}. \quad (7)$$

The iterations of the primal-dual scheme are then given by

$$\begin{cases} p^{n+1}(x, y) &= \max\left(-w(x, y), \min\left(w(x, y), \frac{p^n(x, y) + \sigma_p(\bar{u}^n(x) - \bar{u}^n(y))}{1 + \sigma_p \varepsilon}\right)\right) \\ u^{n+1}(x) &= \frac{u^n(x) - \tau_u \left(\sum_{y \in \mathcal{N}(x)} p^{n+1}(x, y) - p^{n+1}(y, x) + \lambda \text{fcn}^{(d)}(x)\right)}{1 + \tau_u \lambda} \\ \bar{u}^{n+1}(x) &= 2u^{n+1}(x) - u^n(x) \end{cases} \quad (8)$$

In traditional primal-dual optimization (8) is solved iteratively, the time-steps τ_u, σ_p are set to be Lipschitz continuous, and the parameters $\lambda, \sigma_d, \sigma_v, \varepsilon$ in the model are searched empirically. In contrast, we formulate the whole variational primal-dual optimization as our primal-dual network. Hence, each operation in the optimization (addition, multiplication, division, *etc.*) is defined as a network layer and a fixed number of iterations is *unrolled*, similar as in recurrent neural networks. Compared to standard primal-dual optimization our PDN has the advantages that it not only optimizes each parameter of the model in each iteration separately, but also trains separate time-steps for each iteration which are not tied to conservative Lipschitz boundaries.

3.3 Training

In general, we train our method by stochastic gradient descent with an additional momentum term. It is possible to randomly initialize the weights of our model and then train it from

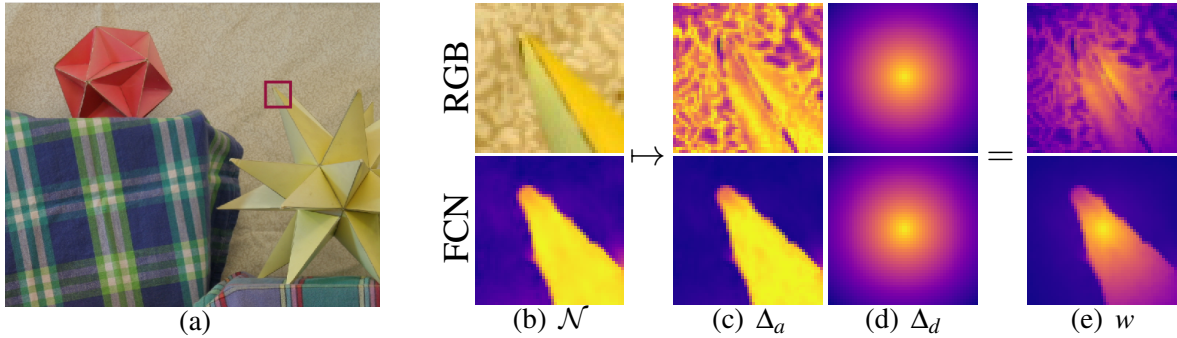


Figure 2: NL-support weights. In (a) the input image is shown. (b-e) depict the NL weight calculation. The 1st row shows the traditional calculation from the intensity image and the 2nd row shows the calculation from the trained FCN output. In detail (b) depicts the image neighborhood \mathcal{N} from which w is estimated. (c) shows the corresponding affinity part Δ_a and (d) the proximity part Δ_d . In (e) the final NL-weight matrix w is shown and the advantage of the fcn output is clearly visible.




Figure 3: Using a physically based renderer we automatically generate 3D scenes of random objects varying in position, size and texture. We also randomly change the lighting directions and intensities. Each sample consists of a noisy LR depth, a HR guidance and a HR target.

scratch. However, in practice we observed faster convergence and increased accuracy, if we pre-train the FCN in advance. Therefore, we train the FCN for 25 epochs with a constant learning rate of 10^{-3} and momentum term set to 0.9 minimizing

$$\sum_{k=1}^K \sum_{x \in \mathbb{R}^{M \times N}} \|\text{fcn}_{s_k}^{(d)}(x) - t_k(x)\|_2^2 + \sum_{y \in \mathcal{N}(x)} \|\text{fcn}_{s_k}^{(A)}(x, y) - (t_k(x) - t_k(y))\|_\varepsilon. \quad (9)$$

After the pre-training step, we plug 20 iterations of our PDN on top of the FCN and train both networks jointly for 10 epochs at a learning rate of 10^{-4} , minimizing the Euclidean loss. We note that the loss function can easily be changed, to evaluate different metrics than the Root Mean Squared Error (RMSE). In this joint training the parameters of the FCN adapt to the PDN, *i.e.* the outputs $\text{fcn}^{(d)}$ and $\text{fcn}^{(A)}$ get optimized to increase the overall accuracy. Additionally, all parameters of the PDN improve as well. This includes the trade-off parameter λ and all hyper-parameters of the primal-dual algorithm. Especially, the parameters get tuned for each iteration individually, yielding an optimal convergence for a fixed number of PDN iterations.

4 Training Data

In this Section we show how we automatically generate our training data. Each training sample (s_k, t_k) is generated using the open source Mitsuba Render Software [39]. In this physically based renderer a scene is defined by placing objects, light sources and sensors freely in an environment defined by a configuration file. Using this file, the renderer generates an intensity- and a depth-map of the scene in definable quality and size.

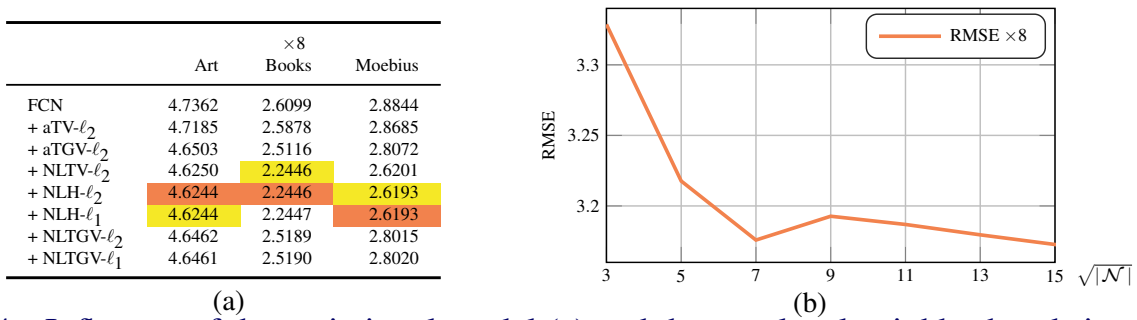


Figure 4: Influence of the variational model (a) and the non-local neighborhood size \mathcal{N} (b) on the RMSE. Best results highlighted in orange and second best in yellow.

	×2			×4			×8			×16		
	Art	Books	Moebius	Art	Books	Moebius	Art	Books	Moebius	Art	Books	Moebius
NN	6.55	6.16	6.59	7.48	6.31	6.78	9.02	6.62	7.00	11.45	7.33	7.52
Bilinear	4.58	3.95	4.20	5.62	4.31	4.56	7.14	4.71	4.87	9.72	5.38	5.43
Yang <i>et al.</i> [41]	3.01	1.87	1.92	4.02	2.38	2.42	4.99	2.88	2.98	7.85	4.27	4.40
He <i>et al.</i> [15]	3.55	2.37	2.48	4.41	2.74	2.83	5.72	3.42	3.57	8.49	4.53	4.58
Diebel & Thrun [7]	3.49	2.06	2.13	4.51	3.00	3.11	6.39	4.05	4.18	9.39	5.13	5.17
Chan <i>et al.</i> [4]	3.44	2.09	2.08	4.46	2.77	2.76	6.12	3.78	3.87	8.68	5.45	5.57
Park <i>et al.</i> [25]	3.76	1.95	1.96	4.56	2.61	2.51	5.93	3.31	3.22	9.32	4.85	4.48
Ferstl <i>et al.</i> [10]	3.19	1.52	1.47	4.06	2.21	2.03	5.08	2.47	2.58	7.61	3.54	3.50
FCN($d^{(mr)}$)	1.83	1.10	1.26	3.03	1.73	1.99	5.39	2.65	3.08	9.31	4.34	4.40
FCN+NLH($d^{(mr)}$)	2.10	1.25	1.38	2.95	1.63	1.88	5.31	2.40	2.91	9.29	4.08	4.18
FCN-PDN($d^{(mr)}$)	1.81	1.05	1.21	2.85	1.53	1.74	5.20	2.26	2.68	8.68	3.70	3.99
FCN(s)	1.99	1.20	1.37	3.25	1.78	1.96	4.74	2.61	2.88	7.80	4.08	4.16
FCN+NLH(s)	2.00	1.18	1.31	3.26	1.62	1.83	4.63	2.25	2.62	7.60	3.59	3.84
FCN-PDN(s)	1.87	1.01	1.16	3.11	1.56	1.68	4.48	2.24	2.48	7.35	3.46	3.62

Table 1: Quantitative results on noisy Middlebury data: We present our results on the disparity maps of the noisy Middlebury dataset [25] as RMSE of the disparity values. Best results highlighted in orange and second best in yellow.

In our case, the automatic dataset generation is scripted by randomly placing different objects (cubes, spheres and planes) in varying poses and dimensions in the scene. Further, the objects are randomly textured using samples from the publicly available *Describable Textures Dataset* [6]. The light intensity and position is also slightly varied during the data generation. We depict two such generated samples in Fig. 3. The output intensity image is used as HR guidance image g , the clean depth output defines the HR target depth t_k , and by downsampling t_k and adding noise we generate the LR depth input $d^{(lr)}$.

5 Evaluation

In the following Section we present a comprehensive evaluation of our *deep primal-dual network*. First, we demonstrate the influence of different energy functionals and the non-local window size on our PDN. Then, we compare our method to state-of-the-art approaches for guided depth super-resolution on the Middlebury dataset as proposed by Park *et al.* [25]. Finally, we present our results on the challenging ToFMark benchmark [10] for real Time-of-Flight data.

5.1 Influence of Energy Functional and Non-Local Window Sizes

In this evaluation we show the influence of the variational model and non-local window size on the accuracy of our model. First, we optimize a variety of different variational models without joint training on top of the FCN output. The RMSE accuracy is shown in Fig. 4(a) evaluated on the noisy Middlebury data [25] ($\times 8$) for 20 iterations. We compare two local

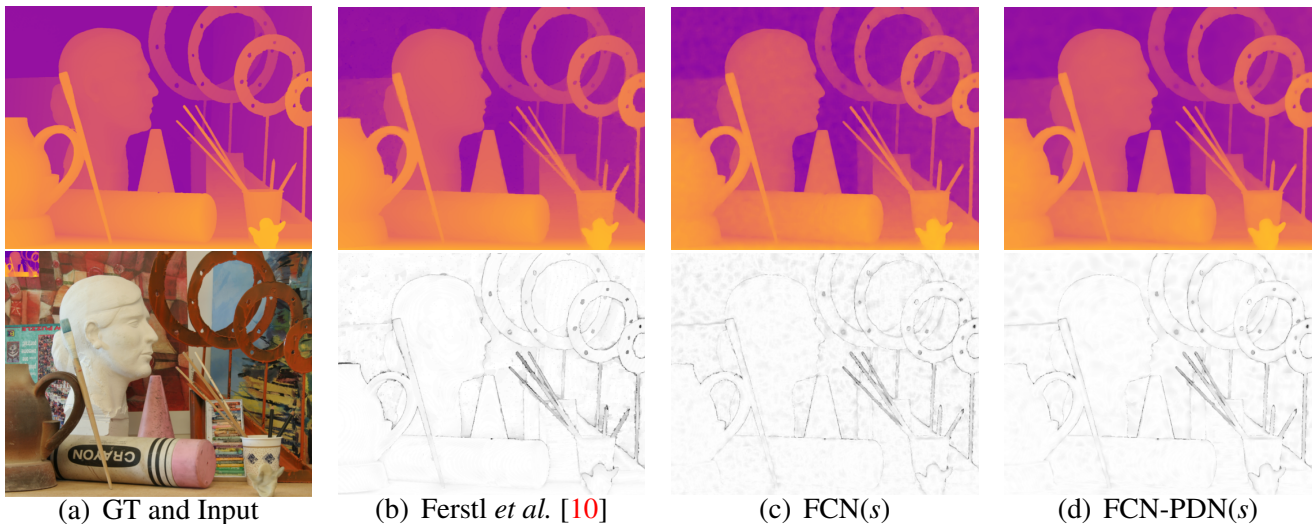


Figure 5: Qualitative results for the image *Art* from the noisy Middlebury dataset [25] and a scale factor of $\times 8$. The first image in (a) shows the ground-truth HR depth and the second image depicts the input sample. In (b)-(c) we present in the first row the HR estimates of a state-of-the-art method, as well as our results, and in the second row we show the corresponding error maps.

models, the anisotropic TV and the anisotropic Total Generalized Variation (TGV) with ℓ_2 data term [10], and the non-local models with TV, Huber and the recently proposed non-local TGV regularization [28] with ℓ_1 and ℓ_2 data term. The neighborhood size \mathcal{N} is set to 7×7 . We can observe that all variational models increase the final accuracy. While the influence of the data penalization is not very significant, the non-local regularization has a superior performance over the local models. Overall, the non-local Huber regularization gives the best results. Second, we evaluate the influence of the non-local window size on the accuracy in Fig. 4(b). The error decreases with a larger neighborhood \mathcal{N} , but also the computational complexity and memory requirements increase dramatically. Hence, we use a 7×7 window size since it provides the best trade-off between accuracy and computational resources.

5.2 Noisy Middlebury

In the following experiment we evaluate our method on the noisy Middlebury dataset as proposed by Park *et al.* [25]. According to [25] we interpret the disparity values as depth. The disparity maps are corrupted by multiplicative Gaussian noise $\eta(x) = \mathcal{N}(0, 651 \cdot d^{(\text{lr})}(x)^{-1})$. The same noise is added to our training data. In Tab. 1 we compare our method to standard interpolation methods and a variety of state-of-the-art methods for guided depth super-resolution. Further, we compare our method once trained solely on the depth maps as input ($d^{(\text{mr})}$) and once with the additional guidance image as input (s). In this comparison we also show the results of the FCN output only (FCN), the results after applying the variational NLH- ℓ_2 on top of the FCN (FCN+NLH), and after joint training of our *deep primal-dual network* (FCN-PDN). For smaller upsampling factors ($\times 2$, $\times 4$) the FCN alone already outperforms all other state-of-the-art methods on this dataset, and our complete method after joint training performs best. At smaller upsampling factors the additional guidance input is not beneficial to the accuracy, but this changes drastically for higher upsampling factors ($\times 8$, $\times 16$). There, we can observe a significant boost in performance by adding the guidance input. In those cases the PDN clearly improves the results over the FCN alone. In Fig. 5 we show a example of the qualitative results. We refer to the supplemental material for more visualizations.

	Books	Devil	Shark
NN	30.46	27.53	38.21
Bilinear	29.11	25.34	36.34
Kopf <i>et al.</i> [21]	27.82	24.30	34.79
He <i>et al.</i> [15]	27.11	23.45	33.26
Ferstl <i>et al.</i> [10]	24.00	23.19	29.89
FCN-PDN ($d^{(mr)}$ & g)	23.74	20.47	28.81

Figure 6: Quantitative and qualitative results on the ToFMark [10] benchmark. In (a) we present our quantitative results as RMSE in *mm*. Best results highlighted in orange and second best in yellow. In (b) and (c) we show the results of the FCN and the full model, respectively. For comparison, we also show in (d) the ground-truth HR depth.

5.3 ToFMark

In our final evaluation we compare our method on the challenging real-world ToFMark dataset [10]. The dataset consists of three different scenes. For each scene it provides a noisy, LR ToF image, a HR depth map, generated with a structure light scanner, and a HR intensity image. The intensity image and the HR depth map are in the same camera coordinate system, however, the HR depth map is given in its own system. Therefore, the depth pixels are mapped to the LR coordinate system of the intensity image via the provided projection matrix. This yields a sparse depth map which we fill with bilinear interpolation to generate the mid-resolution input. In the training data we simulated this projection. First, the HR training depth maps are mapped in the LR ToF coordinate system via the inverse projection matrix. Since multiple HR points can map onto the same LR pixel, we compute the mean over the corresponding depth values. Second, we apply depth dependent noise on the LR depth from which the mid-resolution input is generated. In Tab. 6a we compare the results of our method with state-of-the-art guided depth super-resolution methods, where we can observe a significant improvement in terms of the root mean squared error over previous approaches. A qualitative result is depicted in Fig. 6b-d. Again, we refer to the supplemental material for more visualizations.

6 Conclusion

We presented a novel method that combines the advantages of deep fully convolutional networks and variational methods for guided depth super-resolution. We formulated the non-local variational model as a network which is placed on top of a fully convolutional network by unrolling the optimization steps of a primal-dual algorithm. In a complete end-to-end training our *deep primal-dual network* is able to learn an efficient parameterization of the model including the convolutional filters, and all hyper-parameter and step-sizes of the variational optimization. We created the necessary training data with a physically based renderer in high quality and large quantities. In our evaluations we have shown that this novel combination significantly outperforms state-of-the-art results on different synthetic and real-world benchmarks.

Acknowledgments This work was supported by *Infineon Technologies Austria AG* and the Austrian Research Promotion Agency (FFG) under the *FIT-IT Bridge* program, project #838513 (TOFUSION).

References

- [1] Sergio Almansa-Valverde, José Carlos Castillo, and Antonio Fernández-Caballero. Mobile robot map building from time-of-flight camera. *Expert Systems with Applications*, 39(10):8835–8843, 2012.
- [2] Oisín Mac Aodha, Neill D.F. Campbell, Arun Nair, and Gabriel J. Brostow. Patch Based Synthesis for Single Depth Image Super-Resolution. In *European Conference on Computer Vision (ECCV)*, 2012.
- [3] Antonin Chambolle and Thomas Pock. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011.
- [4] Derek Chan, Hylke Buisman, Christian Theobalt, and Sebastian Thrun. A Noise-aware Filter for Real-time Depth Upsampling. In *European Conference on Computer Vision Workshops (ECCVW)*, 2008.
- [5] Liang-Chieh Chen, Alexander G. Schwing, Alan L. Yuille, and Raquel Urtasun. Learning Deep Structured Models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- [6] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi. Describing textures in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [7] James Diebel and Sebastian Thrun. An Application of Markov Random Fields to Range Sensing. In *Proceedings of Conference on Neural Information Processing Systems (NIPS)*, 2005.
- [8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a Deep Convolutional Network for Image Super-Resolution. In *European Conference on Computer Vision (ECCV)*, 2014.
- [9] Sean Fanello, Cem Keskin, Pushmeet Kohli, Shahram Izadi, Jamie Shotton, Antonio Criminisi, Ugo Pattacini, and Tim Paek. Filter Forests for Learning Data-Dependent Convolutional Kernels. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [10] David Ferstl, Christian Reinbacher, René Ranftl, Matthias Rüther, and Horst Bischof. Image Guided Depth Upsampling using Anisotropic Total Generalized Variation. In *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [11] David Ferstl, Matthias Rüther, and Horst Bischof. Variational Depth Superresolution using Example-Based Edge Representations. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [12] G. Gilboa and S. Osher. Nonlocal Operators with Applications to Image Processing. *Multiscale Modeling and Simulation*, 7(3):1005–1028, 2009.
- [13] Ross Girshick, Jamie Shotton, Pushmeet Kohli, Antonio Criminisi, and Andrew W. Fitzgibbon. Efficient Regression of General-Activity Human Poses from Depth Images. In *IEEE International Conference on Computer Vision (ICCV)*, 2011.

-
- [14] Daniel Glasner, Shai Bagon, and Michal Irani. Super-Resolution from Single Image. In *IEEE International Conference on Computer Vision (ICCV)*, 2009.
- [15] Kaiming He, Jian Sun, and Xiaoou Tang. Guided Image Filtering. In *European Conference on Computer Vision (ECCV)*, 2010.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [17] Michael Hornáček, Christoph Rhemann, Margrit Gelautz, and Carsten Rother. Depth Super Resolution by Rigid Body Self-Similarity in 3D. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [18] Peter J. Huber. Robust regression: Asymptotics, conjectures and monte carlo. *Annal. of Stat.*, 1(5):799–821, 1973.
- [19] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, and Andrew Fitzgibbon. KinectFusion: Real-time 3D Reconstruction and Interaction Using a Moving Depth Camera. In *ACM Symposium on User Interface Software and Technology*, 2011.
- [20] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate Image Super-Resolution Using Very Deep Convolutional Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [21] Johannes Kopf, Michael F. Cohen, Dani Lischinski, and Matthew Uyttendaele. Joint Bilateral Upsampling. *ACM Transactions on Graphics (TOG)*, 26(3):96, 2007.
- [22] Philipp Krähenbühl and Vladlen Koltun. Efficient Inference in Fully Connected CRFs with Gaussian Edge Potentials. In *Proceedings of Conference on Neural Information Processing Systems (NIPS)*, 2012.
- [23] HyeokHyen Kwon, Yu-Wing Tai, and Stephen Lin. Data-Driven Depth Map Refinement via Multi-scale Sparse Representations. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [24] Vinod Nair and Geoffrey E. Hinton. Rectified Linear Units Improve Restricted Boltzmann Machines. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2010.
- [25] Jaesik Park, Hyeonwoo Kim, Yu-Wing Tai, Michael S. Brown, and In-So Kweon. High Quality Depth Map Upsampling for 3D-TOF Cameras. In *IEEE International Conference on Computer Vision (ICCV)*, 2011.
- [26] *Camboard Pico^s*. PMD Technologies. Germany.
- [27] René Ranftl and Thomas Pock. A Deep Variational Model for Image Segmentation. In *German Conference on Pattern Recognition (GCPR)*, 2014.

- [28] René Ranftl, Kristian Bredies, and Thomas Pock. Non-local total generalized variation for optical flow estimation. In *European Conference on Computer Vision (ECCV)*, 2014.
- [29] Gernot Riegler, René Ranftl, Matthias Rüther, and Horst Bischof. Joint Training of an Convolutional Neural Net and a Global Regression Model. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2015.
- [30] Gernot Riegler, Samuel Schulter, Matthias Rüther, and Horst Bischof. Conditioned Regression Models for Non-Blind Single Image Super-Resolution. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [31] Gernot Riegler, Matthias Rüther, and Horst Bischof. ATGV-Net: Accurate Depth Super-Resolution. In *European Conference on Computer Vision (ECCV)*, 2016.
- [32] Samuel Schulter, Christian Leistner, and Horst Bischof. Fast and Accurate Image Upscaling with Super-Resolution Forests. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [33] Jamie Shotton, Toby Sharp, Alex Kipman, Andrew Fitzgibbon, Mark Finocchio, Andrew Blake, Mat Cook, and Richard Moore. Real-time Human Pose Recognition in Parts from Single Depth Images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [34] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. 2015.
- [35] Danhang Tang, Tsz-Ho Yu, and Tae-Kyun Kim. Real-Time Articulated Hand Pose Estimation Using Semi-supervised Transductive Regression Forests. In *IEEE International Conference on Computer Vision (ICCV)*, 2013.
- [36] Danhang Tang, Hyung Jin Chang, Alykhan Tejani, and Tae-Kyun Kim. Latent Regression Forest: Structured Estimation of 3D Articulated Hand Posture. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [37] Radu Timofte, Vincent De Smet, and Luc Van Gool. A+: Adjusted Anchored Neighborhood Regression for Fast Super-Resolution. In *Asian Conference on Computer Vision (ACCV)*, 2014.
- [38] Markus Unger, Thomas Pock, Manuel Werlberger, and Horst Bischof. A convex approach for variational super-resolution. In *German Conference on Pattern Recognition (GCPR)*, 2010.
- [39] Jakob Wenzel. Mitsuba Renderer, 2010. URL <http://www.mitsuba-renderer.org>.
- [40] Jingyu Yang, Xinchun Ye, Kun Li, Chunping Hou, and Yao Wang. Color-Guided Depth Recovery From RGB-D Data Using an Adaptive Autoregressive Model. *IEEE Transactions on Image Processing*, 23(8):3443–3458, 2014.
- [41] Qingxiong Yang, Ruigang Yang, James Davis, and David Nistér. Spatial-Depth Super Resolution for Range Images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.

- [42] Roman Zeyde, Michael Elad, and Matan Protter. On Single Image Scale-Up Using Sparse-Representations. In *Curves and Surfaces*, 2010.
- [43] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip Torr. Conditional Random Fields as Recurrent Neural Networks. In *IEEE International Conference on Computer Vision (ICCV)*, 2015.