

Propagating Confidences through CNNs for Sparse Data Regression

Abdelrahman Eldesokey¹
abdelrahman.eldeskey@liu.se

Michael Felsberg¹
michael.felsberg@liu.se

Fahad Shahbaz Khan^{1,2}
fahad.khan@liu.se

¹ Computer Vision Laboratory
Department of Electrical Engineering
Linköping University
Linköping, Sweden

² Inception Institute of Artificial
Intelligence
Abu Dhabi, UAE

Abstract

In most computer vision applications, convolutional neural networks (CNNs) operate on dense image data generated by ordinary cameras. Designing CNNs for sparse and irregularly spaced input data is still an open problem with numerous applications in autonomous driving, robotics, and surveillance. To tackle this challenging problem, we introduce an algebraically-constrained convolution layer for CNNs with sparse input and demonstrate its capabilities for the scene depth completion task. We propose novel strategies for determining the confidence from the convolution operation and propagating it to consecutive layers. Furthermore, we propose an objective function that simultaneously minimizes the data error while maximizing the output confidence. Comprehensive experiments are performed on the KITTI depth benchmark and the results clearly demonstrate that the proposed approach achieves superior performance while requiring three times fewer parameters than the state-of-the-art methods. Moreover, our approach produces a continuous pixel-wise confidence map enabling information fusion, state inference, and decision support.

1 Introduction

In recent years, machine learning methods have achieved significant successes in many computer vision applications, making use of data from monocular passive image sensors, such as grayscale, RGB, and thermal cameras. Typically, data generated by these image sensors are dense and most existing machine learning methods are designed to fully exploit this dense data in order to understand the scene content. Different to the aforementioned sensors, active sensors, such as LiDAR, RGB-D, and ToF cameras, produce sparse data. Here, the sparse output is caused by the acquisition process through active sensing compared to passively measuring light influx in conventional 2D sensors with dense output. The sparse output imposes additional challenges on the machine learning methods to infer the missing data and find an accurate reconstruction of the entire scene.

Sensors with sparse outputs are becoming increasingly popular and have numerous applications in autonomous driving, robotics, and surveillance due to their range measuring

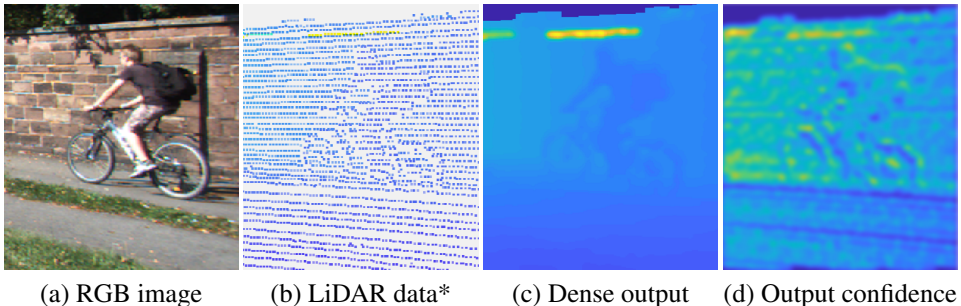


Figure 1: Depth map completion example. The RGB image (a) associated with a projected LiDAR point cloud in (b). The depth map completion output is shown in (c) together with pixel-wise confidence (d) (blue is low and yellow is high). Most existing deep learning methods struggle in scenarios as shown here due to the very high sparsity of the input data (95% of pixels are missing). [*The image is dilated for the sake of visibility]

capability. One fundamental task is *scene depth completion* that aims to reconstruct a full depth map from sparse input. Scene depth completion is a required processing step in, *e.g.* situation awareness and decision support. One of the key challenges when tackling the problem of scene depth completion is the handling of missing values while also differentiating them from the zero-valued regions. Besides, densifying the depth map, corresponding confidences are also desirable since they provide information about reliability of the output values. Such confidence maps are highly important for decision making in safety applications, *e.g.* obstacles detection in autonomous vehicles and robotics. Figure 1 shows an example of a depth completion task. Given the projected LiDAR point cloud, the objective is to densify the sparse depth map, either utilizing the RGB image (guided completion) or only using the projected point cloud (unguided completion). The output is a complete dense map together with pixel-wise output confidence.

Recently, deep learning, notably Convolutional Neural Networks (CNNs), have demonstrated great potential in solving a variety of computer vision tasks. Generally, CNNs are formed by several convolution, local normalization, and pooling layers. The final layers of CNNs are often fully connected (FC) where for classification problems, the last FC layer employs a softmax function to approximate the probability or confidence over the class memberships. Such confidences are often missing in regression settings, although both the regressed value *and* its confidence are required for numerous applications. For example, it is not only relevant to know how far away a potential obstacle is located, but also how reliable this information is. For the scene depth completion task, several deep regression networks have been proposed in the literature that introduce confidence measures [10, 11, 12, 13]. However, all these methods utilize the confidence as a binary valued mask to filter out the missing measurements. This strategy disregards valuable information available in the confidence maps. Different to these existing methods, we propose an approach that utilizes the signal confidence as a continuous measure for data uncertainty and propagates it through all layers.

In this paper, we propose an algebraically-constrained convolution operator for deep networks with sparse input to achieve a proper processing of confidences. The sparse input is equipped with confidences and the network is required to produce a dense output. We derive novel methods for determining the confidence from the convolution operation and propagating it to consecutive layers. To maintain the confidences within a valid range, we impose non-negativity constraints on the network weights during training. Further, we also introduce

an objective function that simultaneously minimizes the data error while maximizing the output confidence. Moreover, we demonstrate the significance of the proposed confidence measure by introducing a novel approach for performing scale-fusion based on confidences. Our proposed method achieves state-of-the-art results on the KITTI depth benchmark [10] while requiring only 480 parameters, which is three times fewer than state-of-the-art methods.

2 Related Work

Scene depth completion is a challenging problem that aims to construct a dense depth map given a sparse depth image. It shares similarities with image inpainting since both tasks require filling missing information/pixels in an image. In case of image inpainting, several approaches based on deep learning have been introduced recently, however restricted to binary masks. These masks define regions in the image where missing pixels have zero values and the remaining pixels have ones. Köhler *et al.* [9] showed quantitatively how incorporating those binary masks in training shallow networks leads to better results, even if the masks were not available during test time. Ren *et al.* [8] proposed a convolution operation based on Shepard interpolation [11] that also utilizes a binary mask to perform inpainting or super-resolution. They propagated the binary masks by convolving them with the same filters/weights as the data and thresholding insignificant values. Liu *et al.* [5] incorporated the use of binary masks in the U-Net architecture [9] for performing inpainting. The binary masks were propagated by setting the pixel at the filter origin to one if not all pixels within the filter support are unknown.

In case of scene depth completion, Uhrig *et al.* introduced the KITTI depth benchmark [10] which is a large-scale dataset for this task. They also proposed a method, where the convolution operations are weighted using the binary masks. The masks are propagated using the max pooling operation. In their work, they also investigated concatenating the binary mask to the input as an additional channel. Chodosh *et al.* [10] utilized compressed sensing to approach the sparsity problem for scene depth completion. A binary mask is employed to filter out the unmeasured values. Further, their method requires significantly fewer parameters compared to [10]. Ma and Karaman [6] proposed sparse-to-dense deep network which utilizes an RGB image and a randomly sampled set of sparse depth measurements to produce a dense depth map.

Our approach is different to the aforementioned methods in several aspects. Firstly, we treat the binary masks as continuous confidences instead of binary values. We further derive an algebraically-constrained deep convolution operator from the normalized convolution framework [9] that infers continuous output confidences. Secondly, different to [8], we enforce the trained filters to be positive to obtain a sound confidence. This is also different from [5, 10] that employed a constant averaging filter for the confidences, a strategy that assumes a uniform confidence distribution among all pixels, which is generally not the case in real-world data. Thirdly, we do not constrain the output confidences to be binary as in [5, 8, 10]. Instead, we propose a computational scheme that allows output confidences to be continuous *while* propagating confidence information from the input to the output. Finally, we demonstrate that utilizing normalized convolution to perform scale-fusion in multi-scale networks based on confidences outperforms the standard convolution used in, *e.g.* in U-Net [9, 9]. Moreover, our proposed approach requires remarkably fewer parameters compared to the aforementioned approaches, while achieving state-of-the-art results.

3 Our Approach

Here, we describe our approach by starting with a brief introduction to the normalized convolution framework. We then introduce an algebraically-constrained normalized convolution operator for CNNs and the propagation method for confidences. Finally, we describe our proposed network architecture and the loss function taking confidences into account.

3.1 Normalized Convolution

Assume a sparse signal/image \mathbf{F} with missing parts due to noise, acquisition process, pre-processing, or other system deficiencies. The missing parts of the signal are identified using a confidence mask \mathbf{C} , which has zeros or low values at missing/uncertain locations and ones otherwise. The signal is sampled and, at each sample point k , the neighborhood is represented as a finite dimensional vector $\mathbf{f}_k \in \mathbb{C}^n$ accompanied with a confidence vector \mathbf{c}_k of the same size, both assumed to be column vectors. Using the notation from [2], *normalized convolution* is defined at all locations k as (index k is omitted to reduce clutter):

$$\mathbf{r} = (\mathbf{B}^* \mathbf{D}_a \mathbf{D}_c \mathbf{B})^{-1} \mathbf{B}^* \mathbf{D}_a \mathbf{D}_c \mathbf{f} \quad , \quad (1)$$

where \mathbf{B} is a matrix which incorporates a set of basis functions $\{\mathbf{b}_i \in \mathbb{C}^n\}_1^m$ in its columns, \mathbf{D}_\times denotes a diagonal matrix with vectorized \times on the diagonal, \mathbf{a} is the applicability function which is a *non-negative* localization function for the basis \mathbf{B} , and \mathbf{r} holds the coefficients of the signal \mathbf{F} at location k projected onto the subspace spanned by \mathbf{B} .

The basis functions in \mathbf{B} could be polynomials or complex exponentials, but the simplest case is when $\mathbf{B} = \{\mathbf{b}_1 : \mathbf{b}_1 = \mathbf{1}\}$, and it becomes *normalized averaging*. In this case, the signal is mapped onto a constant, localized with the applicability function, and (1) simplifies to $\mathbf{r} = (\mathbf{1}^* \mathbf{D}_a \mathbf{D}_c \mathbf{1})^{-1} \mathbf{1}^* \mathbf{D}_a \mathbf{D}_c \mathbf{f}$, which can be formulated for the full signal \mathbf{F} and its confidence \mathbf{C} as ($*$ denotes convolution and \cdot point-wise multiplication):

$$\mathbf{r}[k] = \frac{\sum_i \mathbf{a}[i] \mathbf{F}[k-i] \mathbf{C}[k-i]}{\sum_i \mathbf{a}[i] \mathbf{C}[k-i]} = \frac{\mathbf{a} * (\mathbf{F} \cdot \mathbf{C})}{\mathbf{a} * \mathbf{C}}[k] \quad , \quad (2)$$

3.2 Training the Applicability

The appropriate choice of the applicability function is an open issue as it usually depends on the nature of the data. Therefore, methods for statically estimating the applicability function have been suggested [2], but we aim to learn \mathbf{a} as part of the training. This generalizes convolutional layers, as normalized averaging is equivalent to standard convolution in case of signals with constant confidence. As described above, the applicability function acts as a confidence or localization function for the basis and therefore it is essentially non-negative.

Non-negative applicabilities are feasible to train in standard frameworks, since back-propagation is based on the chain rule and any differentiable function with non-negative co-domain can be plugged in. Thus, the function $\Gamma(\cdot)$, e.g. the softplus, is applied to the weights \mathbf{W} , and the gradients for the weight element $\mathbf{W}_{m,n}^l$ at the l^{th} convolution layer are calculated as:

$$\frac{\partial \mathbf{E}}{\partial \mathbf{W}_{m,n}^l} = \sum_{i,j} \frac{\partial \mathbf{E}}{\partial \mathbf{Z}_{i,j}^l} \cdot \frac{\partial \mathbf{Z}_{i,j}^l}{\partial \Gamma(\mathbf{W}_{m,n}^l)} \cdot \frac{\partial \Gamma(\mathbf{W}_{m,n}^l)}{\partial \mathbf{W}_{m,n}^l} \quad , \quad (3)$$

where \mathbf{E} is the loss between the output and the ground truth, $\mathbf{Z}_{i,j}^l$ is the output of the l^{th} layer at locations i, j that were convolved with the weight element $\mathbf{W}_{m,n}^l$. Accordingly, the forward pass for *normalized convolution* is defined as:

$$\mathbf{Z}_{i,j}^l = \frac{\sum_{m,n} \mathbf{Z}_{i+m,j+n}^{l-1} \mathbf{C}_{i+m,j+n}^{l-1} \Gamma(\mathbf{W}_{m,n}^l)}{\sum_{m,n} \mathbf{C}_{i+m,j+n}^{l-1} \Gamma(\mathbf{W}_{m,n}^l) + \varepsilon}, \quad (4)$$

where \mathbf{C}^{l-1} is the confidence from the previous layer, $\mathbf{W}_{m,n}^l$ is the applicability in this context and ε is a constant to prevent division by zero. Note that this is formally a correlation, as it is a common notation in CNNs.

3.3 Propagating Confidence

The main strength about the signal/confidence philosophy is the availability of confidence information apart from the signal. This is needed to be propagated through the network to output a pixel-wise confidence aside with the network prediction. In normalized convolution frameworks, Westelius [10] proposed a measure for propagating certainties:

$$\mathbf{C}_{\text{out}} = \left(\frac{\det \mathbf{G}}{\det \mathbf{G}_0} \right)^{\frac{1}{m}}, \quad (5)$$

where $\mathbf{G} = (\mathbf{B}^* \mathbf{D}_a \mathbf{D}_c \mathbf{B})$, $\mathbf{G}_0 = (\mathbf{B}^* \mathbf{D}_a \mathbf{B})$ and m is the number of basis functions.

This measure calculates a geometric ratio between the Grammian matrix \mathbf{G} in case of partial confidence and \mathbf{G}_0 in case of full confidence. Setting $\mathbf{B} = \{\mathbf{b}_1 : \mathbf{b}_1 = \mathbf{1}\}$, i.e., $m = 1$, we can utilize the already-computed term in (4) to propagate the confidence as follows:

$$\mathbf{C}_{i,j}^l = \frac{\sum_{m,n} \mathbf{C}_{i+m,j+n}^{l-1} \Gamma(\mathbf{W}_{m,n}^l) + \varepsilon}{\sum_{m,n} \Gamma(\mathbf{W}_{m,n}^l)}, \quad (6)$$

3.4 Loss Function

For scene depth completion task, we usually aim to minimize a norm, e.g. $l1$ or $l2$ norm, between the output from the network and the ground truth. In our proposed method, we use the Huber norm, which is a hybrid between the $l1$ and the $l2$ norm and it is defined as:

$$\|z - t\|_{\text{H}} = \begin{cases} 0.5(z - t)^2 & |z - t| < 1 \\ |z - t| - 0.5, & \text{otherwise} \end{cases} \quad (7)$$

The Huber norm helps preventing exploding gradients in case of highly sparse data, which stabilizes the convergence of the network. Nonetheless, our aim is not only to minimize the error norm between the output and the groundtruth, but also to increase the confidence of the output data. Thus, we propose a new loss which has a data term *and* a confidence term:

$$\mathbf{E}_{i,j} = \|\mathbf{Z}_{i,j}^L - \mathbf{T}_{i,j}\|_{\text{H}}, \quad \tilde{\mathbf{E}}_{i,j} = \mathbf{E}_{i,j} - \frac{1}{p} [\mathbf{C}_{i,j}^L - \mathbf{E}_{i,j} \mathbf{C}_{i,j}^L], \quad (8)$$

where $\mathbf{Z}_{i,j}^L$ is the data output from the final layer L , $\mathbf{C}_{i,j}^L$ is the corresponding confidence output, $\mathbf{T}_{i,j}$ is the ground truth, p is the epoch number, and $\|\cdot\|_{\text{H}}$ is the Huber norm. The main objective of the loss is to minimize the error of the data term and maximize the confidence of the output. Note that the third term in the loss prevents the confidence from growing indefinitely. We weight the confidence term using the reciprocal of the epoch number p to prevent it from dominating the loss function when the data error starts to converge.

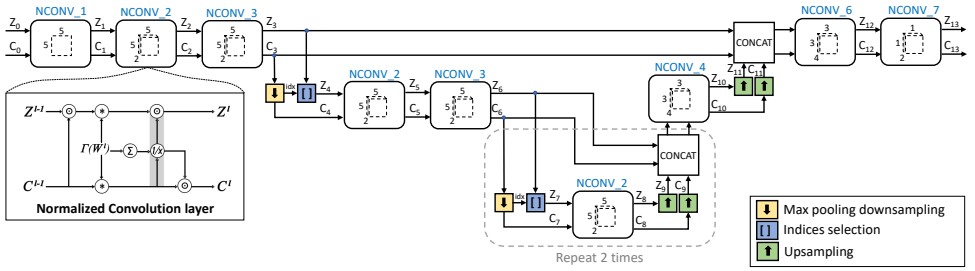


Figure 2: Our proposed multi-scale architecture for the task of scene depth completion which utilizes normalized convolution layers. Downsampling is performed using max pooling on confidence maps and the indices of the pooled pixels are used to select the same pixels from the feature maps. Different scales are fused by upsampling the coarser scale and concatenate it with the finer scale. A normalized convolution layer is then used fuse the feature maps based on the confidence information. Finally, a 1×1 normalized convolution layer is used to merge different channels into one channel dense output and output confidence map.

3.5 Network Architecture

Inspired by [9], we propose a hierarchical multi-scale architecture that shares the same weights between different scales, which leads to a very compact network as shown in Figure 2. Downsampling is performed using max pooling on the confidences and similar to [13] we keep the indices of the pooled pixels, which are then used to select the same pixels from the feature maps, *i.e.*, we keep the most confident feature map pixels. The downsampled confidences are divided by the Jacobian of the scaling to maintain absolute confidence levels. Scale fusion is performed by upsampling the coarser scale and concatenate it with the finer scale. We apply a normalized convolution operator on the concatenated feature map to allow the network to fuse different scales utilizing confidence information.

4 Experiments

4.1 Experimental Setup

Dataset: We evaluate our method on the KITTI depth benchmark [10] which consists of projected LiDAR point clouds. The resulting depth maps/images are very sparse (approximately 4% of pixels have values). The benchmark has 86,000 training images, 7,000 validation images and 1,000 test server images with no access to the ground truth. The ground truth has missing parts as it was matched with the stereo disparity to remove projected LiDAR outliers. We evaluate on the *full* validation set as in [10] and the test set.

Implementation details: All our experiments are performed on a workstation with Intel Xeon CPU (4 cores), 8 GB of RAM and NVIDIA GTX 1080 GPU with 8 GB of memory. NConv-HMS, NConv-1Scale(4ch), and NConv-SF-STD are trained with a batch size of 8, while NCONV-1-Scale(16ch) are trained with a batch size of 4. Our networks were trained on the first 10,000 out of 86,000 depth maps/images in the training set. We use the ADAM solver with default parameters except for the learning rate which we set to 0.01. ¹

¹The source code will be made publicly available.

	MAE [m]	RMSE [m]	MRE	$\delta < 1.01$	$\delta < 1.01^2$	$\delta < 1.01^3$	#Params	Output Conf.
CNN [10]	0.78	2.97	-	-	-	-	2.5×10^4	No
CNN+mask [10]	0.79	2.24	-	-	-	-	2.5×10^4	No
SparseConv [10]	0.58	1.80	0.035	0.33	0.65	0.82	2.5×10^4	No
Sparse-To-Dense [6]	0.70	1.68	0.039	0.21	0.41	0.59	3.4×10^6	No
DCCS-1-Layer [10]	0.83	2.77	0.054	0.30	0.47	0.59	1.0×10^3	No
DCCS-2-Layers [10]	0.47	1.45	0.028	0.41	0.68	0.80	1.8×10^3	No
DCCS-3-Layers [10]	0.43	1.35	0.024	0.48	0.73	0.83	1.7×10^3	No
NConv-1-Scale(16ch)	0.40	1.58	0.022	0.60	0.81	0.88	2.5×10^4	Yes
NConv-1-Scale(4ch)	0.42	1.59	0.022	0.59	0.80	0.88	2.0×10^3	Yes
NConv-HMS	0.38	1.37	0.021	0.60	0.81	0.89	4.8×10^2	Yes
NConv-SF-STD	0.53	3.0	0.037	0.59	0.80	0.88	4.8×10^2	No

Table 1: Evaluation results on the validation set. The results for CNN and CNN+mask are taken from [10], SparseConv, Sparse-To-Dense and DCCS are from [10]. Our multi-scale architecture NConv-HMS outperforms all other method in all evaluation metrics except for RMSE, where it is slightly inferior to DCCS-3-Layers.

Evaluation metrics: For comparison, we use the same evaluation metrics as defined in [10, 6]: *Mean Absolute Error* (MAE) which is an unbiased error metric, *Root Mean Square Error* (RMSE) which penalizes large errors, *Mean Absolute Relative Error* (MRE) is a ratio between the error magnitude and the groundtruth value, and *Inliers Ratio* (δ_i) which is the percentage of pixels having relative error less than a specific threshold to the power of i . As in [10], we use a challenging threshold value of $\delta = 1.01$.

4.2 Quantitative Comparisons

We compare our method with state-of-the-art methods in the literature: Sparsity Invariant Convolution (SparseConv) [10], Deep Convolutional Compressed Sensing (DCCS) [10], and Sparse-To-Dense [6] approaches. As mentioned earlier, the Sparsity Invariant Convolution method applies a constrained convolution operation using binary masks. The DCCS approach [10] employs compressed sensing and Alternating Direction Neural Networks (ADNNs) to create a deep auto-encoder that constructs a dense output. The Sparse-To-Dense method [6] utilizes a ResNet architecture to encode the sparse LiDAR point clouds and RGB images and then decode a dense output

Impact of continuous confidences: To evaluate the impact of employing our proposed confidence scheme, we evaluate a single-scale architecture as described in [10]. This architecture consists of 6 normalized convolution layers with filter sizes of 11×11 , 7×7 , 5×5 , 3×3 , 3×3 and 1×1 respectively with *16 channels* each and we denote it as NConv-1-Scale(16ch). To further demonstrate the efficiency of our approach, we evaluate the same architecture with 4 channels only and we denote it as NConv-1-Scale(4ch). Table 1 shows the results for both experiments as well as other methods in comparison. Our single-scale architecture NConv-1-Scale(16ch) achieves superior results in terms of MAE, MRE and δ_i compared to all other methods. This demonstrates the advantage of our proposed confidences scheme compared to SparseConv [10]. Moreover, our compact architecture NConv-1-Scale(4ch) maintains the performance while requiring remarkably fewer parameters. However, DCCS-2-Layers and DCCS-3-Layers achieve better RMSE than our proposed single-

	SparseConv [10]	NN+CNN [10]	DCCS-3-Layers [10]	NConv-HMS (Ours)
MAE [m]	0.48	0.41	0.44	0.37
RMSE [m]	1.60	1.41	1.32	1.29

Table 2: Quantitative results on the test set. All the results are taken from the online KITTI depth benchmark [10]. Our method outperforms all published methods on the benchmark.

scale architecture, which we attribute to the insufficient receptive field of the network.

Multi-scale architecture: To address the problem of the limited receptive field of our single-scale architecture, we incorporate a multi-scale architecture inspired by [9]. We further maintain the low number of parameters by sharing the weights/filters between different scales. The multi-scale architecture is illustrated in Figure 2 and denoted as NConv-HMS. Table 1 provides the comparison between NConv-HMS and existing methods. Our NConv-HMS achieves better results compared to the single-scale architectures with respect to all the evaluation metrics. The RMSE is the most significantly reduced measure and becomes almost the same as for DCCS-3-Layers. Note also that the number of parameters was reduced to 480, which is remarkably fewer than all other methods in comparison.

Impact of proposed scale-fusion scheme: A common approach to perform multi-scale fusion is to upsample the coarser scales, concatenate it with the finer scale and then use a convolution layer to learn the proper fusion as in [6, 9]. Instead, we perform scale-fusion using a normalized convolution layer which takes into account the confidence information embedded in different scales. We evaluate both approaches in our multi-scale architecture and our confidence-based approach NConv-HMS significantly outperforms the standard fusion approach NConv-SF-STD as shown in Table 1. This clearly demonstrates the significance of utilizing confidence information for selecting the most confident data within the network.

Comparison on the test set: Here, we evaluate on the test set, which can only be performed on the benchmark server. Table 2 shows the error metrics for state-of-the-art methods published in the literature that are based on deep learning. SparseConv [10] performs significantly better on the test set than the validation set, while DCCS-3-Layers maintains its performance. NN+CNN corresponds to performing nearest-neighbor filling for missing pixels and then train a CNN with the same architecture as [10] to enhance the output. Our approach outperforms all published state-of-the-art methods on the test set. Contrary to the validation set, our approach outperforms DCCS-3-Layers on the test set.

4.3 Qualitative Analysis

To further analyze the impact of the proposed contributions, we perform a qualitative study on the KITTI depth benchmark [10]. Figure 3 shows examples for performing depth scene completion on two images from the benchmark. The input are projected LiDAR point clouds that are highly sparse. The ground truth images are not completely dense due to the strict outlier filtering adopted by [10]. These missing data impose a big challenge on methods to learn a good representation. As shown in the figure, our multi-scale architecture performs very well on densifying the sparse input. Moreover, the output confidences from our method provide indication about how reliable the output depth maps are. At locations where neither input points nor groundtruth information is available, e.g. behind the cyclists or below the billboard, the output confidence is very low. Further, the results show that regions in the center of the scene tend to have high confidence due to the high point cloud density in the




Figure 3: Examples of scene depth completion using our multi-scale architecture on the KITTI depth benchmark [10]. *First* row are the sparse projected LiDAR point clouds, *second* row are the ground truth images, *third* row are the dense outputs from our method, and the *last* row are the output confidence maps. Our method performs favorably on densifying the sparse input, while providing a confidence map that indicates the output reliability.

input. This demonstrates that our method for confidence propagation enables the network to learn the prominence of different regions with respect to the groundtruth.

Error analysis: As discussed earlier, our single-scale architecture suffers from a limited receptive field and fails to predict values for regions above the horizon in some images. This leads to a significant increase in the RMSE. We addressed this problem by adopting a multi-scale architecture to cover the whole receptive field. This allows our method to perform well on the whole validation set. For the case of the multi-scale architecture, the error is mainly distributed along sharp edges and upon the horizon. This is likely due to the absence of structural information that could be found in RGB images. Figure 4 shows an example of where the largest errors of our method are located. Obviously, those errors are distributed along the vehicles edges and close to the horizon. This problem could be addressed by incorporating prior knowledge about the structure of the scene from the RGB image.

5 Conclusion

In this paper, we proposed an algebraically-constrained convolution layer for CNNs to tackle the issue of sparse and irregularly spaced input data. Unlike previous works, we treated the input masks as continuous confidences instead of binary values and equip the sparse input with confidences. We further derived novel methods for determining the confidence from the convolution operation and propagating it to consecutive layers. A non-negativity constraints on the network weights is imposed to maintain the confidences within a valid range. Moreover, we introduced an objective function that simultaneously minimizes the data error while maximizing the output confidence. Comprehensive experiments are performed on the KITTI depth benchmark for scene depth completion. The results show that our approach achieves




Figure 4: An example of error analysis for our proposed method on KITTI Depth benchmark [10]. *Top-left* is the input RGB image, *top-right* is the projected LiDAR point cloud, *bottom-left* is the output from our method and *bottom-right* is the error map in logarithmic scale. The error is mainly distributed along edges and close to the horizon.

superior performance while requiring significantly fewer parameters. Finally, the continuous pixel-wise confidence map produced by our approach is shown to produce reasonable results enabling proper information fusion, state inference, and decision support.

6 Acknowledgments

This research is funded by Vinnova through grant CYCLA, the Swedish Research Council through a framework grant for the project Energy Minimization for Computational Cameras (2014-6227), CENIIT grant (18.14) and VR starting grant (2016-05543).

References

- [1] Nathaniel Chodosh, Chaoyang Wang, and Simon Lucey. Deep Convolutional Compressed Sensing for LiDAR Depth Completion. mar 2018. URL <http://arxiv.org/abs/1803.08949>.
- [2] Gunnar Farneböck. *Polynomial expansion for orientation and motion estimation*. PhD thesis, Linköping University Electronic Press, 2002.
- [3] Hans Knutsson and C-F Westin. Normalized and differential convolution. In *Computer Vision and Pattern Recognition, 1993. Proceedings CVPR'93., 1993 IEEE Computer Society Conference on*, pages 515–523. IEEE, 1993.
- [4] Rolf Köhler, Christian Schuler, Bernhard Schölkopf, and Stefan Harmeling. Mask-specific inpainting with deep neural networks. In *German Conference on Pattern Recognition*, pages 523–534. Springer, 2014.
- [5] Guilin Liu, Fitsum A. Reda, Kevin J. Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image Inpainting for Irregular Holes Using Partial Convolutions. apr 2018. URL <http://arxiv.org/abs/1804.07723>.
- [6] Fangchang Ma and Sertac Karaman. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. *arXiv preprint arXiv:1709.07492*, 2017.
- [7] Matthias Mühlich and Rudol Mester. A statistical extension of normalized convolution and its usage for image interpolation and filtering. In *EUSIPCO*, 2004.
- [8] Jimmy SJ Ren, Li Xu, Qiong Yan, and Wenxiu Sun. Shepard convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 901–909, 2015.
- [9] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. pages 234–241. Springer, Cham, oct 2015. doi: 10.1007/978-3-319-24574-4_28. URL http://link.springer.com/10.1007/978-3-319-24574-4_{_}28.
- [10] Donald Shepard. A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 1968 23rd ACM national conference*, pages 517–524. ACM, 1968.
- [11] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity Invariant CNNs. aug 2017. URL <http://arxiv.org/abs/1708.06500>.
- [12] Carl-Johan Westelius. *Focus of attention and gaze control for robot vision*. PhD thesis, Linköping University, Computer Vision, The Institute of Technology, 1995.
- [13] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.