

Learning Task-Specific Generalized Convolutions in the Permutohedral Lattice

Anne S. Wannenwetsch^{1*}[0000-0002-7016-3820],
Martin Kiefel²[0000-0001-9432-5428],
Peter V. Gehler²[0000-0002-5812-4052], Stefan Roth¹[0000-0001-9002-9832]

¹ TU Darmstadt, Germany ² Amazon, Germany

Abstract. Dense prediction tasks typically employ encoder-decoder architectures, but the prevalent convolutions in the decoder are not image-adaptive and can lead to boundary artifacts. Different generalized convolution operations have been introduced to counteract this. We go beyond these by leveraging guidance data to redefine their inherent notion of *proximity*. Our proposed network layer builds on the *permutohedral lattice*, which performs sparse convolutions in a high-dimensional space allowing for powerful non-local operations despite small filters. Multiple features with different characteristics span this permutohedral space. In contrast to prior work, we *learn* these features in a task-specific manner by generalizing the basic permutohedral operations to learnt feature representations. As the resulting objective is complex, a carefully designed framework and learning procedure are introduced, yielding rich feature embeddings in practice. We demonstrate the general applicability of our approach in different joint upsampling tasks. When adding our network layer to state-of-the-art networks for optical flow and semantic segmentation, boundary artifacts are removed and the accuracy is improved.

1 Introduction

Deep learning approaches are the backbone of many state-of-the-art methods across computer vision [7,36,39]. Convolutional neural networks (CNNs) are particularly common as they greatly lower the number of parameters compared to fully-connected networks and thus scale to practically relevant image sizes. While early CNNs employed large filters [28], it is now common to use small kernels stacked into deep networks [16,35]. Chaining several smaller filters requires fewer parameters for the same receptive field of a single large filter, and leads to more discriminative features by virtue of having more non-linearities [35].

While convolutions build a fundamental block of deep learning, they are not without drawbacks. First, they are not image-adaptive, *i.e.* content boundaries in a feature map are not respected but smoothed over. This is especially disadvantageous for dense prediction tasks, *e.g.* semantic segmentation or optical flow, leading to accuracy loss at boundaries [14,46]. Moreover, convolutions have

* This project was mainly done during an internship at Amazon, Germany.

a limited and predefined receptive field, which connects spatially close regions but cannot leverage similar, but more distant image structures. Here, a new definition of *pixel proximity* is needed that goes beyond two-dimensional (2D) spatial distance. For instance, image values themselves or abstract properties such as object classes could be used to define similarity in a more general setting.

Several methods have been proposed to counteract the named disadvantages. Sampling-based approaches [19,33] rearrange the image content but remain restricted to the 2D concept of proximity. Location specific networks [22,46] predict pixelwise filter kernels, but require many additional parameters. Other methods [8,40] determine neighboring pixels in an image-adaptive manner. However, the convolutional structure is fixed and only the position of neighbors is adjustable.

Image-adaptive filters, *e.g.* [15,41], have been used in traditional computer vision for years. The bilateral filter [41] adapts a Gaussian kernel according to the spatial distance and color difference of neighboring pixels. In [20,24], this concept is leveraged to construct bilateral convolution layers (BCLs) based on the *permutohedral lattice* [1] – a fast approximation of the bilateral filter. Filtering corresponds to a sparse convolution in a high-dimensional space, which is spanned by different features, *e.g.* spatial location and color. Jampani *et al.* [20,24] extend the Gaussian kernel to a general, image-adaptive convolution and learn the kernels from data. However, the features constituting the lattice space remain fixed. Feature parameters are not adjustable during training, which complicates integration into end-to-end learning. More importantly, relying on predefined features without further processing omits a possible source of improvements.

To counteract this disadvantage of BCLs, we present the *semantic lattice layer*. We rely on the permutohedral lattice as a backbone and show how to generalize its operations w.r.t. features with learnable parameters. The resulting computations are involved and may lead to practical challenges. We hence propose a specific setting in which basic features – as used in [20,24] – are processed by a CNN. This greatly simplifies the optimization since it allows to combine and especially refine features that are known to be beneficial for image-adaptive filtering. We further present various measures to avoid difficulties during learning. For instance, as the sparsity of the semantic lattice may avoid propagation of information if pixels are too distant, we restrict the output range of the embedded features. This rather simple measure has a large effect in practice.

Our setup enables us to learn meaningful feature embeddings from data. It allows to integrate feature parameters into training and effectively leverages guidance data to connect pixels due to their similar characteristics. As such, the semantic lattice is able to perform non-local operations while keeping the filter kernels and consequently the number of learnt parameters small and manageable.

We show the benefits of the semantic lattice in different areas for image-adaptive upsampling. For the task of color upsampling, the semantic lattice outperforms previous approaches by a large margin. We further replace bilinear upsampling in state-of-the-art networks for optical flow and semantic segmentation. Here, the semantic lattice leads to better aligned and crisper content boundaries and also improves the accuracy, especially at discontinuities.

2 Related Work

Generalized Convolutions. We begin by reviewing work that generalizes convolution operations. Jaderberg *et al.* [19] introduce Spatial Transformers (STs), which transform feature maps depending on the data itself. Similar to warped convolutions [17], STs aim for invariance to certain transformations, *e.g.* rotation or scaling. [29] applies STs to allow for irregular patches in dense prediction tasks. In [33], saliency-based sampling emphasizes regions of high interest. These methods rearrange data in 2D space. In contrast, we can leverage additional feature dimensions to redefine the concept of pixel proximity.

Filter-weight networks [22] generate location-specific filters dependent on the input image. In [23], adaptive weights incorporate side information about the scene context. However, adaptive filters introduce many additional parameters in comparison to location-invariant networks and remain restricted to local transformations due to a fixed receptive field. Wu *et al.* [46] apply location-specific convolutions not only to the position itself but to several sampled neighboring regions, which extends their receptive field but requires additional computations.

Dilated convolutions use a fixed spacing between considered pixels to extend the spatial resolution [6,47]. It is possible to learn offsets for the input locations of each filter [21] or have them depend on the input and spatial location [8]. When using mixtures of Gaussians as filters, size and location of the receptive fields are learnable [40]. Structure-aware convolutions [5] use univariate functions as filters and are also applicable to non-Euclidean data. We do not learn individual neighborhoods for all filters but convolutions are instead performed consistently in a learnt feature space. Moreover, our convolution structure is not fixed; the number of neighbors is flexible and homogenous areas can be compressed.

Permutohedral Lattice. Adams *et al.* [1] propose the permutohedral lattice as a fast method for high-dimensional Gaussian filtering. It found widespread application, especially for fast inference in dense Conditional Random Fields (CRFs) [26,32,48] and upsampling or densification of data [10,34]. In contrast to our work, these approaches use fixed Gaussian filters and predefined features. [27] extends the fast inference method of [26] to learn parameters of dense CRFs, but the setup is restricted to Gaussian filters and customized to its application.

In [20,24], the high-dimensional filtering in permutohedral space is generalized by learnable convolution parameters. The proposed BCLs are beneficial in neural networks as they allow to redefine proximity of pixels w.r.t. different characteristics [12,20,30]. Moreover, BCLs can inherently cope with sparse data [24], *e.g.* in 3D point cloud processing [38]. Again, all methods rely on predefined features and thus restrict the flexibility of the generalized convolutions. We will show that a general setup with learnt features leads to better results in practice.

Another line of research aims for further speed-up of the permutohedral lattice. For instance, [9] proposes to encode its operations in a deep neural net.

Learnt Representations. Our embedding network aims to encode guidance data as discriminatively as possible for the task at hand. As such, our work is related to general embedding or metric learning; see [37,44] for an overview.

Image-Adaptive Filtering. We leverage additional properties for our generalized convolutions, which closely relates our approach to image-adaptive filtering methods such as bilateral filtering [41], non-local means [3], and guided image filtering [15]. All of these filters have been included into deep networks, *e.g.* for semantic segmentation [12,14], image processing [45], or video classification [42].

Only few approaches aim to learn guidance features for the filtering step in a general context. Harley *et al.* [14] propose segmentation-aware convolutions, which leverage image-adaptive masks from an embedding network. Object class labels are required for pre-training and large filter kernels increase the risk of overfitting. Deep joint image filtering (DJIF) [31] uses two individual networks to preprocess guidance and data features and subsequently merges the two branches for joint filtering. However, explicit knowledge about the relation between guidance and target data is not leveraged. Gharbi *et al.* [13] reproduce image enhancement operators with locally-affine models and upsample the low-resolution outputs guided by a learnt feature channel. Here, the offline learning of the models puts strong restrictions on the approximated operators. Deep guided filters [45] allow to learn a guidance image but restrict its dimensionality to a one-dimensional signal per output channel.

In contrast to previous work, the semantic lattice is applicable to a large variety of tasks and puts no restrictions on the guidance data. Moreover, the rich feature representations allow us to keep the applied filter kernels small.

3 The Semantic Lattice

To allow for the non-local combination of data, we build on the permutohedral lattice [1] to redefine the notion of *proximity* between the pixels of an image. The permutohedral lattice assumes that each input point is characterized by two properties – *features* and *data*. The feature value $\mathbf{f} \in \mathbb{R}^d$ indicates the location of the respective pixel in the d -dimensional permutohedral space, while the data value $\mathbf{v} \in \mathbb{R}^c$ describes the information stored at this location. In a first step, the data is projected into the lattice grid using the features to determine its position. Convolutions can then be performed in permutohedral space, considering a neighborhood defined by the feature dimensions. For instance, color values can be considered to connect visually similar areas and respect object boundaries [20,24]. This is in contrast to regular convolutions where the spatial location is used as the only feature to determine neighboring pixels. Finally, the convolved output is extracted at certain feature positions, which can but do not have to coincide with the input locations depending on the task at hand.

In the following, we introduce the permutohedral lattice and its properties more formally. We then extend the work in [20,24] to eliminate the usage of fixed, hand-crafted features. In particular, we show how to learn an appropriate feature space based on spatial positions as well as guidance data. As this approach allows to leverage semantically meaningful properties that go beyond the concept of predefined features, we refer to our proposed setup as the *semantic lattice*.

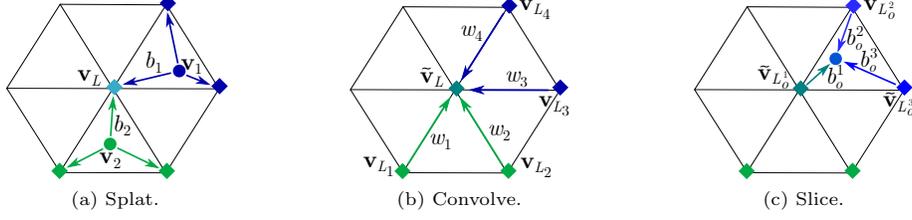


Fig. 1. Basic operations in the permutohedral lattice.

3.1 Permutohedral Lattice

Following [1], the permutohedral lattice is defined as the projection of the regular grid \mathbb{Z}^{d+1} onto the hyperplane $\mathcal{H} : \mathbf{h} \cdot \mathbf{1} = 0 \subseteq \mathbb{R}^{d+1}$. The projected grid points thus represent the corners of permutohedral simplices, which split the hyperplane \mathcal{H} into uniform cells. We refer to [1] for further details of the lattice structure.

The operation to read data into the lattice grid is denoted as *splatting*, see Fig. 1a. The feature vector \mathbf{f}_i is used to place an input point $\mathbf{i} = (\mathbf{f}_i, \mathbf{v}_i)$ into the permutohedral lattice. Then, the data value is splat onto the enclosing lattice points according to its barycentric coordinates. The data value of a lattice point L is given as

$$\mathbf{v}_L = \sum_{i \in \mathcal{I}(L)} b_i \cdot \mathbf{v}_i, \quad (1)$$

where b_i denote the barycentric coordinates of input points $\mathcal{I}(L)$ splatting on L .

The *convolution* step is subsequently performed on the permutohedral grid points considering corners in a neighborhood $\mathcal{N}(L)$, *c.f.* Fig. 1b. If a neighboring corner was not set during splatting, its value is assumed to be zero. Using a kernel $\mathbf{W} = (w_1, \dots, w_N)$, the convolution results in the updated lattice data

$$\tilde{\mathbf{v}}_L = \sum_{L_n \in \mathcal{N}(L)} w_n \cdot \mathbf{v}_{L_n}. \quad (2)$$

Fig. 1c illustrates the final *slicing* operation, which interpolates the data from lattice points to an output pixel \mathbf{o} . The value at pixel position \mathbf{f}_o is obtained as

$$\tilde{\mathbf{v}}_o = \sum_{k=1}^{d+1} b_o^k \cdot \tilde{\mathbf{v}}_{L_o^k}, \quad (3)$$

with enclosing simplex corners L_o^k and barycentric coordinates b_o^k , $1 \leq k \leq d+1$.

In [20,24], the permutohedral lattice is integrated into deep learning by providing partial derivatives of the permutohedral operations w.r.t. the input data \mathbf{v} and the kernel \mathbf{W} . As such, the original Gaussian kernel [1] is transformed into a general convolution with a flexible filter \mathbf{W} learnt from data. As the filter operation is performed in lattice space, the convolution respects the notion of proximity introduced by the features \mathbf{f} that span the permutohedral lattice.

3.2 Generalized Features in the Semantic Lattice

To define the permutohedral space, [20,24] resort to predefined features, which are usually taken as $\mathbf{f} = (x, y, r, g, b)$. Here, x and y describe the spatial x- and y-coordinates of a pixel, which are concatenated with corresponding RGB values.

Our semantic lattice instead aims to learn feature embeddings from data to leverage the full capacity of the lattice. To that end, we introduce generalized input features $\mathbf{f}(\mathbf{i}; \boldsymbol{\theta}_I)$ that depend on each pixel \mathbf{i} as well as a global set of parameters $\boldsymbol{\theta}_I$. The splatting operation in Eq. (1) then generalizes to

$$\mathbf{v}_L(\boldsymbol{\theta}_I) = \sum_{i \in \mathcal{I}(L; \boldsymbol{\theta}_I)} b_i(\boldsymbol{\theta}_I) \cdot \mathbf{v}_i, \quad (4)$$

since the splatting points as well as the corresponding barycentric coordinates depend on the feature values and thus also on $\boldsymbol{\theta}_I$. Due to the fixed lattice structure, the set of neighbors for the convolution remains unchanged. However, the data value \mathbf{v}_L at each lattice point depends on the inputs that splatted to this exact corner such that we rewrite the convolution in Eq. (2) as

$$\tilde{\mathbf{v}}_L(\boldsymbol{\theta}_I) = \sum_{L_n \in \mathcal{N}(L)} w_n \cdot \mathbf{v}_{L_n}(\boldsymbol{\theta}_I). \quad (5)$$

Finally, the set of lattice points surrounding an output pixel \mathbf{o} and its barycentric coordinates are again dependent on its features $\mathbf{f}(\mathbf{o}; \boldsymbol{\theta}_O)$, which are parametrized by a set $\boldsymbol{\theta}_O$. This definition results in a generalized slicing operation given as

$$\tilde{\mathbf{v}}_o(\boldsymbol{\theta}_I, \boldsymbol{\theta}_O) = \sum_{k=1}^{d+1} b_o^k(\boldsymbol{\theta}_O) \cdot \tilde{\mathbf{v}}_{L_o^k(\boldsymbol{\theta}_O)}(\boldsymbol{\theta}_I). \quad (6)$$

As operations in the lattice require specific computations, common automatic differentiation packages cannot be easily applied. Instead, we rely on customized functions for the above generalized operations as well as their parameter gradients. The derivatives then allow us to apply gradient based optimizers to learn task-specific feature representations $\mathbf{f}(\mathbf{i}; \boldsymbol{\theta}_I)$ and $\mathbf{f}(\mathbf{o}; \boldsymbol{\theta}_O)$ from data.

However, the nested occurrence of the parameter sets $\boldsymbol{\theta}_I$ and $\boldsymbol{\theta}_O$ already suggests that learning these generalized features may not be straightforward. Reconsidering the generalized operations in Eqs. (4) – (6), we observe that information between input and output pixels only propagates via a set of lattice corners defined by the neighborhood size of the convolution step. It is thus essential that input and output feature positions are sufficiently close in lattice space when starting the learning process. Otherwise, the loss gradient does not affect the input feature parameters $\boldsymbol{\theta}_I$ and no learning occurs.

To avoid this situation, we propose a specific framework as illustrated in Fig. 2 for the sample task of color upsampling. For given input and output points \mathbf{p} , we first generate several *basic features* $\mathbf{f}_B(\mathbf{p}) \in \mathbb{R}^{d'}$, *i.e.* hand-crafted features that we assume to be helpful for the task of interest. In the example case, the spatial location of each pixel and the corresponding grayscale image are chosen as basic

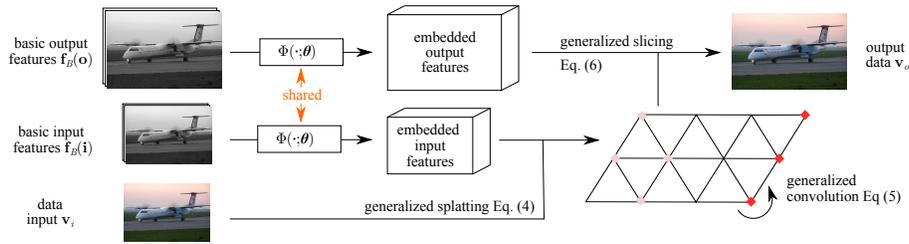


Fig. 2. Visualization of generalized feature learning in the semantic lattice; illustrated for the task of guided color upsampling.

features. Here, the additional grayscale information needs to be available for the input as well as output pixels. We denote it more generally as *guidance data* in the following. Then, a parametric function $\Phi : \mathbb{R}^{d'} \rightarrow \mathbb{R}^d$ is defined, which takes the basic features \mathbf{f}_B as an input and returns a learnt feature embedding $\Phi(\mathbf{f}_B(\mathbf{p}); \boldsymbol{\theta})$ in \mathbb{R}^d . The parameter set $\boldsymbol{\theta}$ is shared across input and output points, *i.e.* $\boldsymbol{\theta} = \boldsymbol{\theta}_I = \boldsymbol{\theta}_O$, to ensure the necessary consistency of the feature embedding. We propose to use a multi-layer CNN for the embedding function Φ and refer to this network in the following as *feature or embedding network*. The parameter set $\boldsymbol{\theta}$ thus corresponds to the network weights. Embedded features as well as data inputs \mathbf{v}_i are then used for the generalized operations in Eqs. (4) – (6).

As there are no restriction on the basic features, the semantic lattice is able to learn different kinds of non-local operations. Due to the characteristics of the permutohedral lattice, input and output positions are rearranged according to the learnt features and spatially far pixels are connected if they share the same characteristics. With this redefinition of proximity, the number of weight parameters remains limited while the semantic lattice yet operates globally.

4 Training the Semantic Lattice

The training and setup of the semantic lattice requires careful consideration. We detail this in the following and provide an experimental analysis in Sec. 5.2.

4.1 Training Procedure

Feature Scaling. As the size of the permutohedral lattice cells is fixed, a scaling factor applied to the individual features determines the importance of the different dimensions as well as the number of pixels that fall into one lattice cell. These scaling factors thus constitute important hyperparameters. Following [20,24], we determine factors for our basic features via grid search. While with the semantic lattice it is possible to refine these factors in end-to-end training through backpropagation, we found little benefit in our experiments. Hence, we use the scaled features from the grid search as inputs to our embedding network.

Data Centering. If predefined features are used to map to the lattice space, the output is largely invariant to a global translation of the features. In contrast,

we found feature network training to be more stable with zero mean input. We thus subtract the dataset mean from the basic features before feature scaling.

Explicit Spatial Features. As the embedding network combines basic features with various scale factors, we find that a random initialization may lead to a poor initial accuracy. While the feature network is able to recover a reasonable embedding starting from a random initialization, we observe long training times in practice as well as occasional convergence to poor local minima (inferior to the scaled basic features themselves). We find that this is mainly caused by the absence of reliable spatial coordinate features in the initial embedding network. Hence, we do not input the spatial coordinates into the embedding network, and instead explicitly concatenate the scaled spatial features and the learnt feature embedding to jointly define the lattice space for the subsequent convolutions.

Normalization Weights. The number of points per lattice cell can vary considerably, resulting in differing ranges of absolute data values at corner points (Eq. 4). Moreover, the flexible structure of the lattice results in a variable number of non-zero neighbors for the convolution in Eq. (5). For this reason, computations in permutohedral space require a normalization step on the slice result in Eq. (6). We divide by a normalization value, which is obtained by performing all lattice operations with a placeholder input with the same features as the regular input and $\mathbf{v}_i = 1$. This implies that an all-one input remains unchanged by the lattice operations. Note that this normalization becomes invalid as convolution weights turn negative. [20,24] resolve this by introducing a separate set of convolution weights for the normalization. They rely on a fixed Gaussian filter, which reduces the flexibility. In contrast, we explicitly learn separate convolution weights for the normalization step and constrain them to be non-negative.

Learning Rates. If the feature network and permutohedral kernels are learnt simultaneously, individual learning rates are applied to both parameter sets.

4.2 Architecture

We use a CNN with 3×3 filters and leakyReLU activations as our feature embedding network. The non-linearities are omitted after the last convolution to allow for positive and negative features.¹ We experimented with ResNet-like feature networks [16], but observed little benefit. In contrast, we found it essential to add a batch normalization layer [18] at the end of the embedding network, *c.f.* Sec. 5.2. This can be understood as follows: Even in our carefully designed semantic lattice, it is possible that no data is splat to the lattice cells surrounding a specific output location. In such a case, the slice operation returns zero and the corresponding gradient with respect to the output location turns zero as well. Without further gradient signals from such pixels, learning keeps pushing more pixels into this disadvantageous state and the accuracy starts to degrade. Consequently, it is necessary to restrict the output range of the feature network, which batch normalization admits. While other normalization methods

¹ Details of the network architecture are provided in the supplemental material.

Table 1. PSNR for *color upsampling* on the Pascal VOC 2012 Segmentation test set.

	PSNR [dB]		PSNR [dB]
Semantic lattice (<i>scaled basic features</i>)	36.55	Nearest neighbors	22.17
Semantic lattice (<i>learnt kernels</i>)	36.62	Bicubic upsampling	23.45
Semantic lattice (<i>learnt embedding</i>)	36.81	DGF [45]	35.17
Semantic lattice (<i>both learnt</i>)	36.83	DJIF [31]	23.99

are possible, *e.g.* a simple min-max normalization, they show no clear benefit over batch normalization, which is commonly available in deep learning libraries.

In permutohedral space, we use a single kernel per input channel with a neighborhood of size one, *c.f.* supplemental material. For upsampling tasks as in Sec. 5.1, transitions of the sparse inputs between lattice cells may cause sudden changes of training loss. For this reason, we apply a nearest neighbor upsampling to the low-resolution guidance and input data before feeding them into the feature network and lattice, respectively. The spatial features are adapted to the upsampled input, which spreads the data more evenly over the lattice and leads to more reliable gradients w.r.t. the features.

5 Experiments

5.1 Color Upsampling

Guided upsampling is a common application of image-adaptive filters [2,25,31,45]. Here, guidance data is available at a higher resolution than the data of interest. This is particularly interesting if sensor data is available at different resolutions.

We evaluate our approach on the the task of joint color upsampling in which a grayscale image guides the upsampling of a low-resolution color image. Following [20], we use images of the Pascal VOC Segmentation splits [11] for training, validation, and test from which we removed grayscale images for fair comparison. Bilinear interpolation is used to downsample color and grayscale images by $4\times$.

The semantic lattice is applied to learn the offset between grayscale images and the RGB data. We use spatial coordinates and grayscale values as basic features. The semantic lattice is trained for 100 epochs on random crops of size 200×272 with learning rates of 0.001 and 0.01 for the feature network and permutohedral kernels, respectively. For comparison, we also train the deep guided filter (DGF) [45] in the same setting for 150 epochs using their procedure for image processing tasks. Again, the DGF predicts the offset between RGB and grayscale images as this slightly improves the results. We also compare with Deep Joint Image Filtering (DJIF) [31] by applying their residual network trained for the task of depth upsampling to predict color offsets.

Table 1 summarizes color upsampling results on Pascal VOC Segmentation test. When learning the permutohedral kernels (*learnt kernels*), we observe only a small benefit in comparison to the usage of a Gaussian filter (*scaled basic*

features). In contrast, our learnt feature embedding (*learnt embedding*) leads to a significant improvement, highlighting the importance of using task-specific features. Combining both leads to another (minor) gain. Overall, we outperform the baselines of nearest neighbor and bicubic upsampling as well as related work [31,45] by a large margin. Please see supplemental material for visualizations.

5.2 Validation of Architectural Choices

We now compare different settings for feature learning to validate our proposed lattice setup. Table 2 summarizes results obtained with fixed Gaussian kernels. First, we train an embedding network without batch normalization to evaluate the importance of restricting its output range. We observe a significant drop in PSNR with a result only slightly better than that with scaled basic features. This is due to the fact that input and output locations do not necessarily coincide, which may lead to empty cells without gradients, *c.f.* Sec. 4.2. If the output range of the network is restricted, the number of such pixels can be kept small.

Next, we validate feeding our embedding network with guidance data and concatenating its output with spatial features. We first learn the scale factor of x- and y-coordinates jointly with the embedding. As this yields a negligible improvement over our baseline, we generally do not refine the scale factors. However, note that bigger benefits may be obtained from scale refinement if the initial scale factors are estimated only coarsely. In a second experiment, we apply the embedding network to all features, *i.e.* spatial coordinates and grayscale values. The network learns reasonable features from random initialization, but the PSNR is clearly lower than our baseline despite training $9\times$ longer.

Finally, we evaluate our new normalization approach and learn kernels using scaled basic features. Applying a fixed Gaussian filter rather than a flexible, positive kernel for normalization reduces the PSNR by 0.04dB.

Table 2. Validation of architectural choices for color upsampling on Pascal VOC test.

	PSNR [dB]
Baseline (<i>learnt embedding</i>)	36.81
W/o batch normalization layer	36.61
Learnt spatial scale factor	36.82
Spatial features embedded	36.55
Baseline (<i>learnt kernels</i>)	36.62
Gaussian normalization	36.58

5.3 Dense Prediction Tasks

We next apply our semantic lattice in deep networks for challenging dense prediction tasks, where networks typically operate on downsampled images and use bilinear upsampling as a last step, *e.g.* [7,39].

Optical Flow. We first consider optical flow for which PWC-Net [39] performs competitively on different benchmarks, *e.g.* [4]. However, the calculated flow looks blurry and boundary details are oversmoothed, see Fig. 3. We attribute this to the non-adaptive upsampling that enlarges the estimated flow by $\sim 4\times$.

To obtain sharper and more detailed flow, we replace the bilinear upsampling with a single convolution in the semantic lattice. As basis, we use the so-called

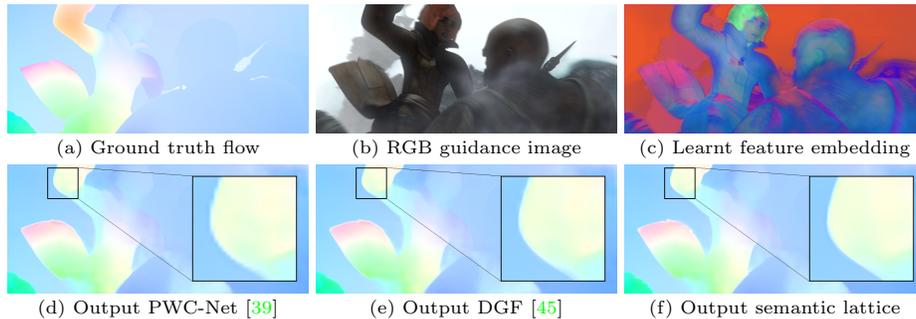


Fig. 3. Flow fields and corresponding guidance data for a Sintel sequence.

Table 3. Average-end-point error (AEE) and boundary AEE (bAEE) on our Sintel test split (ours) and the official Sintel test set (off.). Semantic lattice abbreviated to SL.

	clean (ours)		final (ours)		clean (off.)	final (off.)
	AEE	bAEE	AEE	bAEE	AEE	AEE
SL (<i>scaled basic features</i>)	1.27	7.84	1.66	8.65	–	–
SL (<i>learnt embedding</i>)	1.26	7.50	1.66	8.61	–	–
SL (<i>both learnt</i>)	1.25	7.49	1.65	8.56	3.84	4.89
PWC-Net [39]	1.30	8.52	1.67	8.98	3.90	4.90
DGF [45]	1.29	8.31	1.67	8.91	–	–

PWC-Net.ROB model trained on a variety of datasets. For fair comparison, we do not backpropagate into the network itself but only update the parameters of the semantic lattice, since bilinear upsampling cannot benefit from learning. Spatial coordinates and color values of the first image are leveraged as basic features. The high-resolution guidance image is equally used for input and output features. Our setup is trained on the Sintel dataset [4], which we split randomly into 862 training, 80 validation, and 99 test images. We use the average end-point error (AEE) as loss function and train all configurations for 100 epochs on random 281×512 crops. Learning rates are set to $1e-3$ and $1e-7$ for embedding parameters and permutohedral kernels, respectively. We again compare our approach to DGF [45], which we trained for 500 epochs using their setup for computer vision tasks and hyperparameters tuned on validation.

Table 3 shows results on our own test split of Sintel clean and final as well as on the official test images of the benchmark. Our proposed semantic lattice layer leads to a moderate AEE improvement on both sets. This is to be expected as our experimental setup can only refine the flow estimates. However, sharper flow boundaries are clearly visible when considering the results in Fig. 3. As the AEE is known to be insensitive towards boundary accuracy, we also evaluate a boundary average end-point error (bAEE). It focuses on accuracy close to motion discontinuities, which are determined from ground truth flow by applying a threshold to the flow gradient norm, *c.f.* [43]. As the varying motion ranges

require different thresholds [4,43], we follow Weinzaepfel *et al.* [43] and generate multiple masks using values in $\{1, 3, 7, 10\}$. These masks are subsequently dilated with a structuring element of size 3. We finally calculate the bAEE by evaluating flow on boundary regions only and averaging over the different boundary masks. Our proposed approach shows a clear benefit for boundary regions, improving the bAEE much more significantly than DGF [45] on Sintel clean and final.

Semantic Segmentation. We finally consider the task of semantic segmentation and replace the bilinear upsampling of the recent DeepLabv3+ [7] with our semantic lattice. Again, we only update parameters of the semantic lattice and keep the remaining network fixed to an Xception65 model trained on COCO and Pascal VOC augmented, *c.f.* [7]. Basic features and the setup of our lattice layer remain the same as for optical flow. We train the semantic lattice with random crops of size 200×272 on the training set of Pascal VOC 2012 [11], which we further split into training and validation. The embedding network is trained for 25 epochs with a learning rate of $1e-3$, which we reduce by $10\times$ for the remaining 75 epochs. The learning rate for permutohedral kernels is fixed to $1e-8$. DGF is trained as for optical flow with hyperparameters used in [45].

While the semantic lattice without learnt embedding performs slightly worse than the original implementation, the full semantic lattice and DGF outperform DeepLabv3+. Table 4 summarizes results on Pascal VOC 2012 validation; see supplemental for visualizations. The overall improvement is rather small, which we attribute mainly to DeepLabv3+ being highly engineered, with particular focus on the decoder (unlike the previous DeepLabv3). Nevertheless, image-adaptive filters may benefit further from jointly training with the entire network.

Table 4. Mean intersection over union (mIoU) for *semantic segmentation* on our Pascal VOC 2012 test set.

	mIoU
Semantic lattice (<i>scaled basic features</i>)	82.17%
Semantic lattice (<i>learnt embedding</i>)	82.24%
Semantic lattice (<i>both learnt</i>)	82.25%
DeepLabv3+ [7]	82.20%
DGF [45]	82.26%

6 Conclusion

We introduced the semantic lattice layer, a task-specific, generalized convolution. Our approach is built on the permutohedral lattice that rearranges input data according to different features and thus performs non-local operations with small filter kernels. First, we generalized the operations in permutohedral space to feature representations that can be learnt from data. We then showed how rich feature embeddings can be learnt in practice and validated the proposed architecture. When applying the semantic lattice to color upsampling, learning task-specific features showed a clear benefit. Adding the semantic lattice to decoders in deep neural networks for optical flow and semantic segmentation allowed to reduce boundary artifacts and improved the accuracy for both tasks.

References

1. Adams, A., Baek, J., Davis, M.A.: Fast high-dimensional filtering using the permutohedral lattice. *Comput. Graph. Forum* **29**(2) (2010) 2, 3, 4, 5
2. Barron, J.T., Poole, B.: The fast bilateral solver. In: *ECCV (2016)* 9
3. Buades, A., Coll, B., Morel, J.M.: A non-local algorithm for image denoising. In: *CVPR (2005)* 4
4. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: *ECCV (2012)* 10, 11, 12
5. Chang, J., Gu, J., Wang, L., Meng, G., Xiang, S., Pan, C.: Structure-aware convolutional neural network. In: *NeurIPS*2018* 3
6. Chen, L., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Semantic image segmentation with deep convolutional nets and fully connected CRFs. In: *ICLR (2015)* 3
7. Chen, L., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: *ECCV (2018)* 1, 10, 12
8. Dai, J., Qi, H., Xiong, Y., Li, Y., Zhang, G., Hu, H., Wei, Y.: Deformable convolutional networks. In: *ICCV (2017)* 2, 3
9. Dai, L., Tang, L., Xie, Y., Tang, J.: Designing by training: Acceleration neural network for fast high-dimensional convolution. In: *NeurIPS*2018* 3
10. Dolson, J., Baek, J., Plogemann, C., Thrun, S.: Upsampling range data in dynamic environments. In: *CVPR (2010)* 3
11. Everingham, M., Eslami, S.M.A., Gool, L.V., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL visual object classes challenge: A retrospective. *Int. J. Comput. Vision* **111**(1), 98–136 (2015) 9, 12
12. Gadde, R., Jampani, V., Kiefel, M., Kappler, D., Gehler, P.: Superpixel convolutional networks using bilateral inceptions. In: *ECCV (2016)* 3, 4
13. Gharbi, M., Chen, J., Barron, J.T., Hasinoff, S.W., Durand, F.: Deep bilateral learning for real-time image enhancement. *SIGGRAPH (2017)* 4
14. Harley, A.W., Derpanis, K.G., Kokkinos, I.: Segmentation-aware convolutional networks using local attention masks. In: *ICCV (2017)* 1, 4
15. He, K., Sun, J., Tang, X.: Guided image filtering. In: *ECCV (2010)* 2, 4
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR (2017)* 1, 8
17. Henriques, J.F., Vedaldi, A.: Warped convolutions: Efficient invariance to spatial transformations. In: *ICML (2017)* 3
18. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *ICML (2015)* 8
19. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: *NIPS*2015* 2, 3
20. Jampani, V., Kiefel, M., Gehler, P.V.: Learning sparse high dimensional filters: Image filtering, dense CRFs and bilateral neural networks. In: *CVPR (2016)* 2, 3, 4, 5, 6, 7, 8, 9
21. Jeon, Y., Kim, J.: Active convolution: Learning the shape of convolution for image classification. In: *CVPR (2017)* 3
22. Jia, X., Brabandere, B.D., Tuytelaars, T., Gool, L.V.: Dynamic filter networks. In: *NIPS*2016* 2, 3
23. Kang, D., Dhar, D., Chan, A.B.: Incorporating side information by adaptive convolution. In: *NIPS*2017* 3

24. Kiefel, M., Jampani, V., Gehler, P.V.: Permutohedral lattice CNNs. In: ICLR Workshop Track (2016) [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
25. Kopf, J., Cohen, M.F., Lischinski, D., Uyttendaele, M.: Joint bilateral upsampling. *ACM T. Graphics* **26**(3) (2007) [9](#)
26. Krähenbühl, P., Koltun, V.: Efficient inference in fully connected CRFs with Gaussian edge potentials. In: NIPS*2011 [3](#)
27. Krähenbühl, P., Koltun, V.: Parameter learning and convergent inference for dense random fields. In: ICML (2013) [3](#)
28. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: NIPS*2012 [1](#)
29. Li, J., Chen, Y., Cai, L., Davidson, I., Ji, S.: Dense transformer networks. arXiv:1705.08881 [cs.CV] (2017) [3](#)
30. Li, S., Seybold, B., Vorobyov, A., Lei, X., Kuo, C.J.: Unsupervised video object segmentation with motion-based bilateral networks. In: ECCV (2018) [3](#)
31. Li, Y., Huang, J.B., Ahuja, N., Yang, M.H.: Joint image filtering with deep convolutional networks. In: IEEE T. Pattern Anal. Mach. Intell. (2019) [4](#), [9](#), [10](#)
32. Perazzi, F., Krähenbühl, P., Pritch, Y., Hornung, A.: Saliency filters: Contrast based filtering for salient region detection. In: CVPR (2012) [3](#)
33. Recasens, A., Kellnhofer, P., Stent, S., Matusik, W., Torralba, A.: Learning to zoom: A saliency-based sampling layer for neural networks. In: ECCV (2018) [2](#), [3](#)
34. Russell, C., Yu, R., Agapito, L.: Video pop-up: Monocular 3d reconstruction of dynamic scenes. In: ECCV (2014) [3](#)
35. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015) [1](#)
36. Singh, B., Najibi, M., Davis, L.S.: Sniper: Efficient multi-scale training. In: NeurIPS*2018 [1](#)
37. Song, H.O., Xiang, Y., Jegelka, S., Savarese, S.: Deep metric learning via lifted structured feature embedding. In: CVPR (2016) [3](#)
38. Su, H., Jampani, V., Sun, D., Maji, S., Kalogerakis, E., Yang, M.H., Kautz, J.: Splatnet: Sparse lattice networks for point cloud processing. In: CVPR (2018) [3](#)
39. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In: CVPR (2018) [1](#), [10](#), [11](#)
40. Tabernik, D., Kristan, M., Leonardis, A.: Spatially-adaptive filter units for deep neural networks. In: CVPR (2018) [2](#), [3](#)
41. Tomasi, C., Manduchi, R.: Bilateral filtering for gray and color images. In: ICCV (1998) [2](#), [4](#)
42. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: CVPR (2018) [4](#)
43. Weinzaepfel, P., Revaud, J., Harchaoui, Z., Schmid, C.: Learning to detect motion boundaries. In: CVPR (2015) [11](#), [12](#)
44. Wu, C.Y., Manmatha, R., Smola, A.J., Krähenbühl, P.: Sampling matters in deep embedding learning. In: ICCV (2017) [3](#)
45. Wu, H., Zheng, S., Zhang, J., Huang, K.: Fast end-to-end trainable guided filter. In: CVPR (2018) [4](#), [9](#), [10](#), [11](#), [12](#)
46. Wu, J., Li, D., Yang, Y., Bajaj, C., Ji, X.: Dynamic filtering with large sampling field for convnets. In: ECCV (2018) [1](#), [2](#), [3](#)
47. Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. In: ICLR (2016) [3](#)
48. Zhang, Z., Fidler, S., Urtasun, R.: Instance-level segmentation for autonomous driving with deep densely connected MRFs. In: CVPR (2016) [3](#)