

# ResFPN: Residual Skip Connections in Multi-Resolution Feature Pyramid Networks for Accurate Dense Pixel Matching

Rishav<sup>\*†‡</sup>, René Schuster<sup>\*†</sup>, Ramy Batraway<sup>†</sup>, Oliver Wasenmüller<sup>†§</sup> and Didier Stricker<sup>†¶</sup>

<sup>\*</sup>Equal contribution

<sup>†</sup>DFKI - German Research Center for Artificial Intelligence, {firstname.lastname}@dfki.de

<sup>‡</sup>Birla Institute of Technology and Science (BITS) Pilani, f2016108@pilani.bits-pilani.ac.in

<sup>§</sup>University of Applied Sciences Mannheim, o.wasenmueller@hs-mannheim.de

<sup>¶</sup>University of Kaiserslautern - TUK

**Abstract**—Dense pixel matching is required for many computer vision algorithms such as disparity, optical flow or scene flow estimation. Feature Pyramid Networks (FPN) have proven to be a suitable feature extractor for CNN-based dense matching tasks. FPN generates well localized and semantically strong features at multiple scales. However, the generic FPN is not utilizing its full potential, due to its reasonable but limited localization accuracy. Thus, we present ResFPN – a multi-resolution feature pyramid network with multiple residual skip connections, where at any scale, we leverage the information from higher resolution maps for stronger and better localized features. In our ablation study, we demonstrate the effectiveness of our novel architecture with clearly higher accuracy than FPN. In addition, we verify the superior accuracy of ResFPN in many different pixel matching applications on established datasets like KITTI, Sintel, and FlyingThings3D.

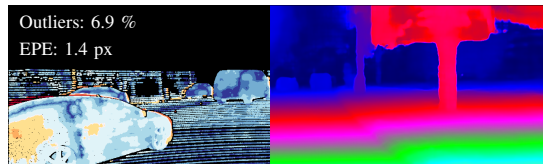
## I. INTRODUCTION

Dense pixel matching is the task to find pixel-wise correspondences across different images. It is one of the core challenges in computer vision and used for many algorithms such as optical flow, scene flow and disparity estimation. Traditionally, heuristic feature descriptors (e.g. SIFT [2] or CENSUS [3]) were used to represent every pixel via its surrounding. In recent years, especially CNN-based approaches, which were trained end-to-end, have achieved remarkable results for dense pixel matching [1], [4]–[6]. Within this category of algorithms, the feature representation turned out to be an essential factor for accurate matching [7]. The representation must be as characteristic as possible in order to be distinguishable. In addition, it must be as localizable as possible to allow for accurate matching and avoid small displacement mismatches. In the state-of-the-art, Feature Pyramid Networks (FPN) [8] seem to fulfill these properties best. FPN was originally proposed in the field of object detection, for which its localization is completely sufficient. However, the accuracy of the localization of FPN for dense pixel matching can be further improved.

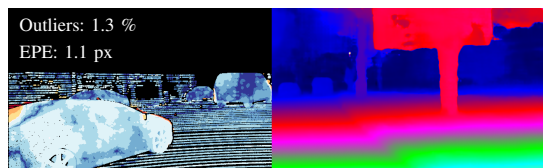
Thus, we present ResFPN which combines – compared to FPN – multiple feature representation of higher resolutions via residual skip connections. This is supposed to re-introduce



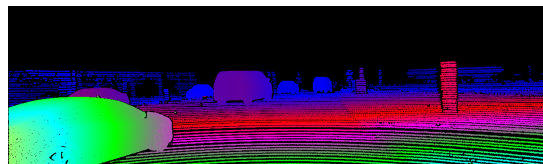
(a) Reference Image



(b) Disparity Result of PSMNet [1]



(c) Improved Result with our ResFPN



(d) Ground Truth

Fig. 1: ResFPN is a deep architecture to compute feature representations for dense matching. ResFPN applied together with SotA matching networks like PSMNet [1] preserves details better and is more robust under challenging conditions.

details for better localization in the final feature representation. Further, the residual skip connections can reduce the length of gradient paths during back-propagation to improve convergence [9]. We review our ResFPN in a comprehensive ablation study by validating each individual design decision in detail. In addition, we bring ResFPN into application for

the dense pixel matching tasks of optical flow, scene flow and disparity estimation. For these experiments, we utilize state-of-the-art algorithms and change nothing but the feature description. We confirm the superior accuracy across different algorithms as well as datasets, such as KITTI [10], Sintel [11] and FlyingThings3D [12].

## II. RELATED WORK

*Representations and Image Pyramids:* Feature maps (i.e. dense descriptors) are the basic cues for many computer vision tasks. A large number of methods show that a proper design of feature maps improves results especially for dense pixel-wise matching in terms of geometric reconstruction and motion estimation. Many approaches employ handcrafted designs like SIFT [2], HOG [13] or DAISY [14] features using image pyramid structure for seeking dense motion matches [15]–[17] or for scene flow estimation [18]. Pyramid feature representations use information from multiple scales for more improvement in terms of estimating correspondences. However, the advances of CNNs improve the robustness of feature maps against ill-conditioned environments, light or geometric changes compared to conventional solutions. In this context, many approaches aim to learn features [19], [20] for dense matching. These methods replace the conventional descriptors but they are not proven in end-to-end learning fashion for dense matching predictions. Our ResFPN is a flexible, modular network that can be plugged in as feature backbone for end-to-end matching networks.

*End-to-End Solutions using Feature Pyramids:* Early end-to-end learning solutions yielded impressive results based on encoder-decoder architectures, e.g. FlowNet [21], [22] for optical flow estimation. DispNet [12] extends the idea of FlowNet to disparity and scene flow estimation. The main idea of the encoder-decoder network is to aggregate the information from coarse-to-fine predictions, which is useful for large displacement predictions. However, it is a memory consuming approach and its computation is inefficient. SPyNet [23] is a lightweight model that aggregates information with a spatial pyramid network. Large motions can be handled with this approach. Compared to FlowNet, it is faster and yields better accuracy. PWC-Net [6] and LiteFlowNet [4] add warping and cost volume layers to the pyramid feature extractor which improves dense optical flow accuracy. PSMNet [1] uses a spatial pyramid pooling module to enlarge the receptive field of feature maps for stereo matching. Instead of using a generic CNN as feature extractor in PWC-Net [6], PWOC-3D [5] employs the FPN architecture [8] and utilizes those features for scene flow estimation with stereo images. Our ResFPN contributes to feature computation for many kinds of deep networks especially in the context of dense matching in a novel way.

*Connecting Layers in Deep Neural Networks:* Traditional CNN architectures establish strictly sequential connections between layers [24]–[26]. Recently, more involved connections have been proposed. DenseNet [27] uses connections in a feed-forward fashion so that for each layer the feature maps of

all preceding layers are used as input to strengthen feature propagation. ResNet [28] and InceptionNet [29], [30] aim to improve deep networks through parallel shortcut connections.

Among modern architectures, Feature Pyramid Network (FPN) [8] leverages the concept of lateral connections for multi-level predictions based on features of multiple scales. Similar to the U-Net architecture [31], it fuses feature maps between the same levels of top-down and bottom-up paths using element-wise addition. In contrary, TDM [32] changes the lateral connections to convolutional layers and channel-wise concatenation with the output, which makes it computationally inefficient. Reverse Densely Connected Feature Pyramid Network [33] proposes to add reverse dense connections for the top-down module (decoder). Similarly, (A)RDFPN [34], [35] add dilated residual connections to the top-down stream of FPNs. The previous feature modules have been presented in the context of object detection. Recently, HRNet [36] has used multi-resolution feature maps to improve localization in the estimation of human poses.

Different from the aforementioned applications, our ResFPN uses the advantages of pyramidal networks to extract dense feature maps for dense matching tasks in terms of stereo matching, optical flow, and scene flow estimation. We utilize not only connections between similar levels of feature maps across bottom-up and top-down parts like FPN [8], but further enhance the spatial accuracy by adding new connections across high resolution feature maps of the bottom-up part and feature maps in the top-down part as shown in Fig. 2d.

## III. METHOD

ResFPN is a generic concept that can be applied in many different applications for different tasks. The general idea is to increase the number of lateral skip connections between encoder and decoder in feature pyramid networks in order to improve the spatial accuracy while maintaining high-level feature representations.

### A. Multiple Residual Skip Connections

Our work continues with the logical extension of regular lateral skip connections to further improve localization and feature abstraction in feature pyramid networks [8]. The reasoning is that additional connections from higher resolved levels of the encoder can benefit the final feature description (cf. Fig. 2). Further, more densely connected networks are assumed to have a better flow of gradients during training [27], [28] which improves convergence properties. Most recently, pyramidal feature extractors have also been shown to be more robust to adversarial attacks [37]. Moreover, the idea of ResFPN is independent of hyper-parameters of the pyramid like the number of levels, or the scale factor. It is applicable together with any building blocks for down-/up-sampling, like Residual [28], Dense [27], or Inception [29], [30] units. The idea of additional residual skip connections between encoder and decoder can be applied in all cases.

The theoretical idea of ResFPN can include any additional connection of layers in pyramid networks that goes beyond

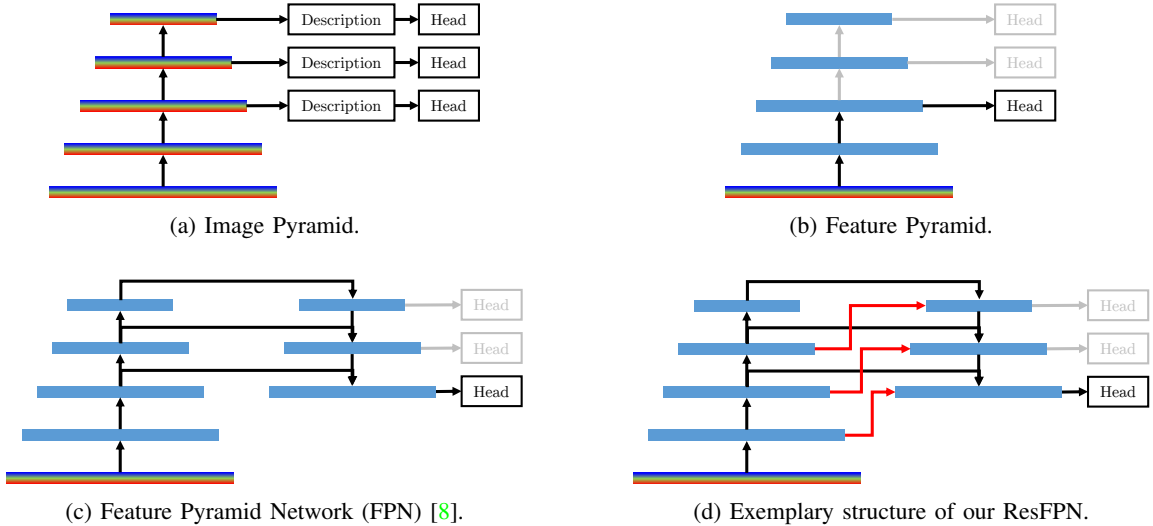


Fig. 2: Feature computation with different types of pyramids. (a) A simple image pyramid is used together with heuristic descriptors for multi-scale predictions. (b) Feature pyramids successively compress and encode the input image for multi-scale predictions. (c) Feature Pyramid Networks traverse the entire encoder and decode the representation until the required scale is reached. (d) Additional feature encodings of higher resolutions are combined during up-sampling in our ResFPN. Here, only a single additional connection per layer is visualized. Details about up-sampling and merging of ResFPN can be found in Fig. 3.

regular lateral skip connections, e.g. dense connections. However, for dense matching we argue that the set of possible connections can be restricted. More precisely, additional connections from lower resolved feature maps towards higher resolutions [33] are assumed to improve feature semantics only and do not contribute to the goal of better localization (they might even accomplish the opposite). As a result, we focus on (multiple) connections from higher resolution feature maps of the encoder to feature maps of the decoder (see Fig. 2d).

Along with these additional connections, novel questions arise. Higher resolution feature maps need to be adjusted to fit the spatial dimensions of the connected decoder layer. This can be done with any size-changing layer, e.g. strided convolution or pooling. Joining multiple feature maps into a single one requires a suitable strategy for merging. Commonly, either element-wise addition or concatenation is used. While the latter allows to maintain the separation of features, it also can lead to heavy computational loads for large and deep feature maps. Finally, one may ask which layers should be additionally connected. In theory, the more higher levels are used, the more the focus is shifted towards localization. On the other hand, a dense connection of every higher resolution to every lower one might be impractical. These questions are investigated in our ablation study in Section IV-A.

The final remark of the theoretical discussion of ResFPN is related to the spectrum of applications. We argue that ResFPN is especially powerful when used for deeper pyramids that realize a (coarse-to-fine, incremental) multi-level prediction at multiple scales. However, the application of ResFPN is not limited to this use case. It is also possible to use only a certain level of the decoder for a single final prediction. Our experiments (Section IV-B) cover a broad range of end-

to-end differentiable dense matching networks to demonstrate the flexibility of ResFPN.

### B. Feature Extraction Network

A ResFPN consists of  $l_d$  arbitrary down-sampling blocks of sub-sampling factor  $s$  (usually  $s = 2$ ), a bottleneck, and  $l_d \leq l_u$  up-sampling blocks using the same factor  $s$ . In regular FPNs [8], the up-sampling block merges the corresponding feature encoding of the target resolution with the up-sampled result to produce a refined feature map. In our extension of ResFPN, we additionally use  $h$  feature encodings of the next higher resolutions during merging. The feature encodings of higher resolutions need to be re-shaped to fit the spatial dimensions (and possibly the feature depth) of the target feature map. In theory, any re-sizing operation could be used for this task, e.g. strided convolution. We compare different strategies for re-shaping and merging in Section IV-A. Each (or one) of the feature maps of the decoder can then be used as features for the prediction.

One possible way to implement a ResFPN is described here. We base our architecture on the FPN in [5] which is an extension of the feature pyramid of [6]. That is, we use  $l_d = 6$  down-sampling blocks with a sub-sampling factor of  $s = 2$  to compute 6 feature maps, where the first one has  $1/2$  of the input resolution and the deepest encoding has  $1/64$  of the input resolution. This is followed by  $l_u = 4$  up-sampling blocks to reconstruct a feature map of  $1/4$  of the original image resolution. Higher resolutions are not required for most of the prediction heads in our experiments [1], [5], [6], but are possible. The down-sampling is performed by two  $3 \times 3$  convolutions, where the first one applies a stride of 2. For up-sampling, we apply a  $4 \times 4$  transposed convolution with stride 2, merge the up-sampled features with a regular skip

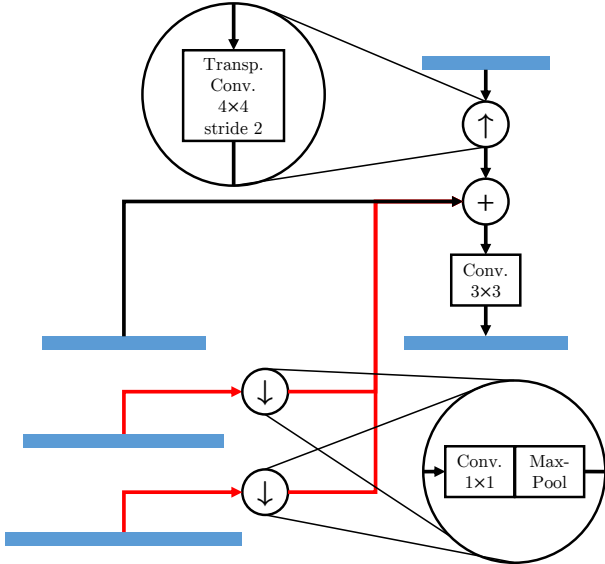


Fig. 3: A single up-sampling block in the decoder of our ResFPN combines four different resolutions. The previous lower resolution feature representation of the decoder is up-sampled with a transposed convolution, the equally resolved feature map from the encoder is connected through a classical skip connection, and two higher resolution feature encodings are additionally connected after down-sampling. For down-sampling,  $1 \times 1$  convolution and max-pooling are applied. Merging is performed by element-wise addition followed by convolution.

connection and  $h = 2$  additional lateral connections through element-wise addition, and then refine the fused features with a  $3 \times 3$  convolution. To align spatial size and feature depth for the merging of higher resolution feature encodings, we propose a  $1 \times 1$  convolution followed by max-pooling with a kernel size and stride of  $s \cdot \Delta l$ . Reshaping, merging, and refinement during up-sampling is illustrated in Fig. 3. In our experience, the combination of  $1 \times 1$  convolution with max-pooling is in the sweet spot of preserving spatial accuracy and computational efficiency, especially when feature depth is increased during the convolution (which is usually the case from higher to lower resolutions) (cf. Section IV-A). LeakyReLU activation [38] is used for all convolutions to introduce non-linearity into the model. The entire architecture of ResFPN with all details is given in Table I.

#### IV. EXPERIMENTS AND RESULTS

Our ResFPN is designed to extract features for dense matching such as stereo disparity, optical flow, or scene flow estimation. Our experiments cover end-to-end networks for all these matching tasks (cf. Section II). In particular, PWOC-3D [5] is used for scene flow estimation, PWCNet [6] and LiteFlowNet [4] represent optical flow estimators, and PSMNet [1] is the network used for disparity estimation.

The experiments consider three well established data sets. FlyingThings3D (FT3D) [12] is used in all cases for pre-

TABLE I: The detailed architecture of our ResFPN. (Up)Conv( $c, k, s, d$ ) and MaxPool( $k, s$ ) describe (transposed) convolution and max-pooling with  $c$  kernels, square kernel size  $k$ , stride  $s$ , and dilation rate  $d$ .

Name	Input	Layer	Output Shape
<i>input</i>	–	–	$H \times W \times 3$
<i>enc-1-1</i>	<i>input</i>	Conv(16,3,2,1)	$\frac{1}{2}H \times \frac{1}{2}W \times 16$
<i>enc-1-2</i>	<i>enc-1-1</i>	Conv(16,3,1,1)	$\frac{1}{2}H \times \frac{1}{2}W \times 16$
<i>enc-2-1</i>	<i>enc-1-2</i>	Conv(32,3,2,1)	$\frac{1}{4}H \times \frac{1}{4}W \times 32$
<i>enc-2-2</i>	<i>enc-2-1</i>	Conv(32,3,1,1)	$\frac{1}{4}H \times \frac{1}{4}W \times 32$
<i>enc-3-1</i>	<i>enc-2-2</i>	Conv(64,3,2,1)	$\frac{1}{8}H \times \frac{1}{8}W \times 64$
<i>enc-3-2</i>	<i>enc-3-1</i>	Conv(64,3,1,1)	$\frac{1}{8}H \times \frac{1}{8}W \times 64$
<i>enc-4-1</i>	<i>enc-3-2</i>	Conv(96,3,2,1)	$\frac{1}{16}H \times \frac{1}{16}W \times 96$
<i>enc-4-2</i>	<i>enc-4-1</i>	Conv(96,3,1,1)	$\frac{1}{16}H \times \frac{1}{16}W \times 96$
<i>enc-5-1</i>	<i>enc-4-2</i>	Conv(128,3,2,1)	$\frac{1}{32}H \times \frac{1}{32}W \times 128$
<i>enc-5-2</i>	<i>enc-5-1</i>	Conv(128,3,1,1)	$\frac{1}{32}H \times \frac{1}{32}W \times 128$
<i>enc-6-1</i>	<i>enc-5-2</i>	Conv(196,3,2,1)	$\frac{1}{64}H \times \frac{1}{64}W \times 196$
<i>enc-6-2</i>	<i>enc-6-1</i>	Conv(196,3,1,1)	$\frac{1}{64}H \times \frac{1}{64}W \times 196$
<i>bottleneck</i>	<i>enc-6-2</i>	Conv(196,1,1,1)	$\frac{1}{64}H \times \frac{1}{64}W \times 196$
<i>skip-5-6</i>	<i>enc-5-2</i>	Conv(196,1,1,1) MaxPool(2,2)	$\frac{1}{64}H \times \frac{1}{64}W \times 196$
<i>skip-4-6</i>	<i>enc-4-2</i>	Conv(196,1,1,1) MaxPool(4,4)	$\frac{1}{64}H \times \frac{1}{64}W \times 196$
<i>dec-6-2</i>	<i>bottleneck</i> + <i>skip-5-6</i> + <i>skip-4-6</i>	Conv(196,3,1,1)	$\frac{1}{64}H \times \frac{1}{64}W \times 196$
<i>dec-5-1</i>	<i>dec-6-2</i>	UpConv(128,4,2,1)	$\frac{1}{32}H \times \frac{1}{32}W \times 128$
<i>skip-4-5</i>	<i>enc-4-2</i>	Conv(128,1,1,1) MaxPool(2,2)	$\frac{1}{32}H \times \frac{1}{32}W \times 128$
<i>skip-3-5</i>	<i>enc-3-2</i>	Conv(128,1,1,1) MaxPool(4,4)	$\frac{1}{32}H \times \frac{1}{32}W \times 128$
<i>dec-5-2</i>	<i>dec-5-1</i> + <i>enc-5-2</i> + <i>skip-4-5</i> + <i>skip-3-5</i>	Conv(128,3,1,1)	$\frac{1}{32}H \times \frac{1}{32}W \times 128$
<i>dec-4-1</i>	<i>dec-5-2</i>	UpConv(96,4,2,1)	$\frac{1}{16}H \times \frac{1}{16}W \times 96$
<i>skip-3-4</i>	<i>enc-3-2</i>	Conv(96,1,1,1) MaxPool(2,2)	$\frac{1}{16}H \times \frac{1}{16}W \times 96$
<i>skip-2-4</i>	<i>enc-2-2</i>	Conv(96,1,1,1) MaxPool(4,4)	$\frac{1}{16}H \times \frac{1}{16}W \times 96$
<i>dec-4-2</i>	<i>dec-4-1</i> + <i>enc-4-2</i> + <i>skip-3-4</i> + <i>skip-2-4</i>	Conv(96,3,1,1)	$\frac{1}{16}H \times \frac{1}{16}W \times 96$
<i>dec-3-1</i>	<i>dec-4-2</i>	UpConv(64,4,2,1)	$\frac{1}{8}H \times \frac{1}{8}W \times 64$
<i>skip-2-3</i>	<i>enc-2-2</i>	Conv(64,1,1,1) MaxPool(2,2)	$\frac{1}{8}H \times \frac{1}{8}W \times 64$
<i>skip-1-3</i>	<i>enc-1-2</i>	Conv(64,1,1,1) MaxPool(4,4)	$\frac{1}{8}H \times \frac{1}{8}W \times 64$
<i>dec-3-2</i>	<i>dec-3-1</i> + <i>enc-3-2</i> + <i>skip-2-3</i> + <i>skip-1-3</i>	Conv(64,3,1,1)	$\frac{1}{8}H \times \frac{1}{8}W \times 64$
<i>dec-2-1</i>	<i>dec-3-2</i>	UpConv(32,4,2,1)	$\frac{1}{4}H \times \frac{1}{4}W \times 32$
<i>skip-1-2</i>	<i>enc-1-2</i>	Conv(32,1,1,1) MaxPool(2,2)	$\frac{1}{4}H \times \frac{1}{4}W \times 32$
<i>skip-0-2</i>	<i>input</i>	Conv(32,1,1,1) MaxPool(4,4)	$\frac{1}{4}H \times \frac{1}{4}W \times 32$
<i>dec-2-2</i>	<i>dec-2-1</i> + <i>enc-2-2</i> + <i>skip-1-2</i> + <i>skip-0-2</i>	Conv(32,3,1,1)	$\frac{1}{4}H \times \frac{1}{4}W \times 32$

training and evaluation. It provides dense scene flow ground truth and is thus also applicable for the training of optical flow or disparity networks. Further, we fine-tune networks on KITTI [10], [39] and Sintel [11]. The KITTI 2015 Scene

TABLE II: Ablation study on our validation split of KITTI data for different numbers and kinds of residual connections with different strategies for merging. A simple FPN establishes only a single skip connection between layers of the same resolution. Our ResFPN adds two residual connections of higher resolutions (cf. Fig. 3). Results for scene flow estimation with PWOC-3D [5] validate that the setup of ResFPN yields the best results while at the same time increases the computational effort and network size only marginally.

	$h$	Re-shaping	Merging	FT3D [12]		KITTI [10]		Parameters $\times 10^6$	FLOPs $\times 10^{12}$
				>3px	EPE	>3px	EPE		
FPN [5]	0	–	addition	21.49	9.15	12.55	3.22	<b>8.05</b>	<b>6.07</b>
	1	$1 \times 1$ , max-pool	addition	20.95	8.28	11.37	3.09	8.09	6.50
	2	max-pool	concatenation	<u>19.90</u>	7.91	<u>11.21</u>	3.04	8.67	8.94
	2	$1 \times 1$ , max-pool	concatenation	21.16	8.34	11.83	<u>3.02</u>	9.03	12.09
	2	$3 \times 3$ , stride	addition	21.65	8.42	13.67	3.50	8.74	7.43
	2	$1 \times 1$ , bi-linear	addition	20.89	8.09	11.55	3.21	8.12	7.26
	2	max-pool, $1 \times 1$	addition	20.28	<u>7.67</u>	12.24	3.06	8.12	<u>6.24</u>
ResFPN	2	$1 \times 1$ , max-pool	addition	<b>18.91</b>	<b>7.19</b>	<b>10.63</b>	<b>2.98</b>	8.12	7.30

Flow data set also provides (sparse) labels for scene flow and can therefore be used for all evaluations. Sintel is a data set for optical flow and is thus used for experiments related to optical flow only. For validation and evaluation, the random split of [5] is used for KITTI, and we randomly sample 5 out of the 23 sequences for Sintel. These sequences are *alley\_2*, *ambush\_4*, *bamboo\_2*, *cave\_4*, and *market\_5*. For augmentation, we apply photometric transformations as in [5], [21] and temporal flipping for pre-training on FT3D. Unless mentioned otherwise, pre-training and fine-tuning are done with a batch size of 2 and 1, respectively.

The metrics being considered in the comparison are the end-point error (EPE) in pixels of the 1-, 2-, or 4-dimensional prediction and the KITTI outlier rate (>3px) of the respective task in percent [10]. For both, lower is better.

Using these setups, two sets of experiments are performed. First, we evaluate our design choices in Section IV-A and compare different ways to implement ResFPN. Second, we apply the features of ResFPN together with different end-to-end matching networks in Section IV-B.

### A. Design Decisions

There are multiple ways to implement the idea of ResFPN. In this section, we compare different entities of ResFPN and vary the number of additional skip connections  $h$ , the merging operation, and the method to adjust size and depth of the skip features. For those experiments, we use the up- and down-sampling blocks presented in Section III-B and Table I with the prediction head of PWOC-3D [5] for scene flow. The different variants are compared in Table II.

We vary the number of skip connections from 1 ( $h = 0$ , the original FPN) to 3 ( $h = 2$ ). More than three lateral connections were not realizable due to hardware constraints, yet we can clearly see that an increase of connections improves the final results. Furthermore, concatenation versus addition is tested. Since the concatenation is independent of the feature depths of the merging input, it is not necessary to reshape the depth of the additional skip features. However, when this step is omitted, performance decreases. Yet, if this step is included, it is not obvious what the output depth of the  $1 \times 1$  convolution should be. For the numbers reported in Table II,

we use the output depth of the up-sampled target feature map, i.e. the same number of output channels that is required for merging by element-wise addition. As a consequence, the computational effort is increased a lot.

Lastly, we change the re-shaping strategy to align spatial shapes, and in case of addition the depth of the skip feature maps. Our approach of  $1 \times 1$  convolution followed by max-pooling is opposed to strided convolution, convolution followed by bi-linear down-sampling, and max-pooling followed by convolution to show the importance of the order. Out of all strategies, our re-shaping approach with element-wise addition and  $1 + h = 3$  skip connections (visualized in Fig. 3) performs the best while, at the same time, is computationally affordable. The overhead of the additional residual connections in terms of numbers of parameters and floating point operations is negligibly small, but outlier rate and end-point error drop by 7 to 22 %. Note that the feature computation with either FPN or ResFPN requires less than 10 % of the entire floating point operations for the prediction of the scene flow with PWOC-3D [5]. In detail, inference with PWOC-3D for a single pair of stereo images on a GeForce GTX 1080 Ti requires about 0.2 s, i.e. feature computation with ResFPN for a single image ( $\sim 0.5$  MP) takes about 5 ms.

### B. Dense Matching with ResFPN

Four different end-to-end networks for scene flow, optical flow, and disparity estimation are used for dense matching. For our experiments, we replace nothing but the feature computation module with our ResFPN (and a simple FPN [8] for comparison). The predictions for the three different feature extractors are then compared. Evaluation is conducted on all mentioned data sets if ground truth for the respective task is available. Training schedules are as close as possible to the original, including multi-stage training if relevant, learning rate schedules, etc. Deviations are explicitly mentioned. Results for all networks on all data sets are presented in Table III.

*Stereo Disparity:* PSMNet [1] is used to compute stereo disparity. This network predicts single scale dense stereo displacements at  $1/4$  resolution, i.e. only the output of *dec-2-2* from our ResFPN is used for the prediction. For the comparison between baseline and ResFPN, we replace the

TABLE III: Comparison of feature extractors. For different prediction networks on different data sets, we evaluate the original network and a version where nothing but the feature module is changed to our improved ResFPN. To validate if the additional lateral connections in ResFPN are the reason for the improvement, we also compare to a simple FPN [8].

Prediction Head	FT3D [12]				KITTI [10]				Sintel [11]									
	Original >3px	EPE	FPN >3px	EPE	ResFPN >3px	EPE	Original >3px	EPE	FPN >3px	EPE	ResFPN >3px	EPE	Original >3px	EPE	FPN >3px	EPE	ResFPN >3px	EPE
PWOC-3D [5]	–	–	21.5	9.2	<b>18.9</b>	<b>7.2</b>	–	–	12.6	3.2	<b>10.6</b>	<b>3.0</b>	–	–	–	–	–	–
PWCNet [6]	19.9	8.5	19.4	8.4	<b>18.7</b>	<b>8.2</b>	15.6	3.7	14.6	3.3	<b>13.9</b>	<b>3.2</b>	20.2	6.0	19.6	5.7	<b>18.5</b>	<b>5.7</b>
LiteFlowNet [4]	23.1	9.8	22.8	9.9	<b>20.9</b>	<b>9.0</b>	18.0	3.7	18.0	3.6	<b>16.4</b>	<b>3.5</b>	20.7	5.7	19.6	5.7	<b>18.3</b>	<b>5.6</b>
PSMNet [1]	16.0	5.3	<b>10.9</b>	5.2	<b>10.9</b>	<b>4.9</b>	3.0	<b>1.0</b>	2.6	<b>1.0</b>	<b>2.2</b>	<b>1.0</b>	–	–	–	–	–	–

CNN module for feature extraction (see Table 1 in [1]) with our ResFPN. To smooth the interface between our code and the *SPP* module of [1], we pass the used feature representation through a  $1 \times 1$  convolution with 128 output channels to match the input shapes between baseline and ResFPN. For any training of PSMNet, we use a batch-size of 3 for pre-training. For the training of PSMNet together with ResFPN, we reduce the entire learning rate schedule by factor 10, because the additional skip connections affect the flow of gradients and thus can influence stability.

Our results show a significant reduction of outliers ( $>3\text{px}$ ) for both stereo data sets when using ResFPN. End-point errors on FT3D are also reduced. ResFPN also outperforms the simple FPN with a single lateral skip connection only.

*Optical Flow:* PWCNet [6] and LiteFlowNet [4] are used for estimation of optical flow. For our experiments with PWCNet, we use the exact ResFPN as described in Section III-B and Table I. For LiteFlowNet, we demonstrate the flexibility of ResFPN and test a version that is closer to the original feature computation module of LiteFlowNet. We still apply the concept of multiple residual skip connections in a pyramidal encoder-decoder network using the up-sampling concept shown in Fig. 3, but we change the hyper-parameters to fit the settings of the encoder of LiteFlowNet [4]. In detail, the feature encoder is formed by the input image, a first feature representation at full resolution, and then 5 additional down-sampled feature maps. This setup reaches a minimal resolution of  $1/32$  with feature depths of 3, 32, 32, 64, 96, 128, 192 for the 7 parts of the encoder (including the image itself). For the prediction, multiple scales are used iteratively until  $1/2$  of the input resolution is reached. This is different from all other networks, where the final resolution for prediction is  $1/4$ .

For both optical flow networks, the results improve on all data sets when features from our ResFPN are used. This holds for both metrics, outlier rate and average end-point error. ResFPN also outperforms a simple FPN [8] in all our optical flow experiments. A visual comparison of the results of the baselines and ResFPN is given in Fig. 4 for exemplary images from KITTI and Sintel. It is evident that not only localization of features was improved to capture more details during matching. Moreover, ResFPN shows an increased robustness compared with its competitors in general. In the first sample from Sintel for example, the relatively small, badly illuminated character is outlined much better when ResFPN features are used for the matching, even if the overall results for this

TABLE IV: Evaluation in boundary regions of objects on KITTI [10].  $d_n$  defines the average end-point error for areas around object boundaries of  $n$  pixels width.

Predictor	Original				with ResFPN			
	$d_3$	$d_5$	$d_{10}$	$d_{20}$	$d_3$	$d_5$	$d_{10}$	$d_{20}$
PWOC [5]	10.75	10.23	8.24	6.54	10.18	9.87	8.30	6.80
PWC [6]	10.34	9.79	8.25	6.94	9.46	9.06	7.91	6.84
LFN [4]	13.72	12.57	10.13	8.74	12.36	11.43	9.15	7.72
PSM [1]	2.07	2.31	2.12	1.70	2.77	2.46	1.79	1.35

frame are slightly worse. On a global scale, especially for large displacements or occluded areas, ResFPN outperforms the baseline (e.g. in the last example of Fig. 4).

*Scene Flow:* For estimation of scene flow with PWOC-3D [5], our original design of ResFPN is applied again. The major differences here are that four instead of two images are processed for matching with ResFPN and that the baseline is already using a FPN with lateral skip connections [5], [8]. Therefore, this experiment has the strongest baseline. Still, ResFPN achieves a considerable reduction of outliers of about 15 % and cuts the end-point errors by  $\sim 6\%$  and  $\sim 22\%$  for KITTI and FT3D, respectively.

In summary, using ResFPN for feature computation in end-to-end matching networks reduces outlier rates and end-point errors (or maintains them) in all our experiments. The better localized features preserve details during matching and produce more consistent and smooth results in comparison to simple Feature Pyramids (FP) and basic Feature Pyramid Networks (FPN). ResFPN could achieve this for prediction networks with very different characteristics, e.g. single and multi-scale estimation, different encoding (down-sampling) blocks, and different final resolutions. More results and visualizations are provided within our supplementary video.

### C. Improved Localization

The previous section confirms that matching with ResFPN yields an overall better result on various domains with all kinds of networks. However, one of our major claims is improved localization by utilization of multiple higher resolution feature maps. Therefore, a final experiment to validate this claim is conducted. Towards this end, we make use of the object masks provided by the KITTI data set [10] to repeat the previous experiment on boundary regions only. The average end-point error for different maximum distances to object boundaries is evaluated and reported in Table IV.

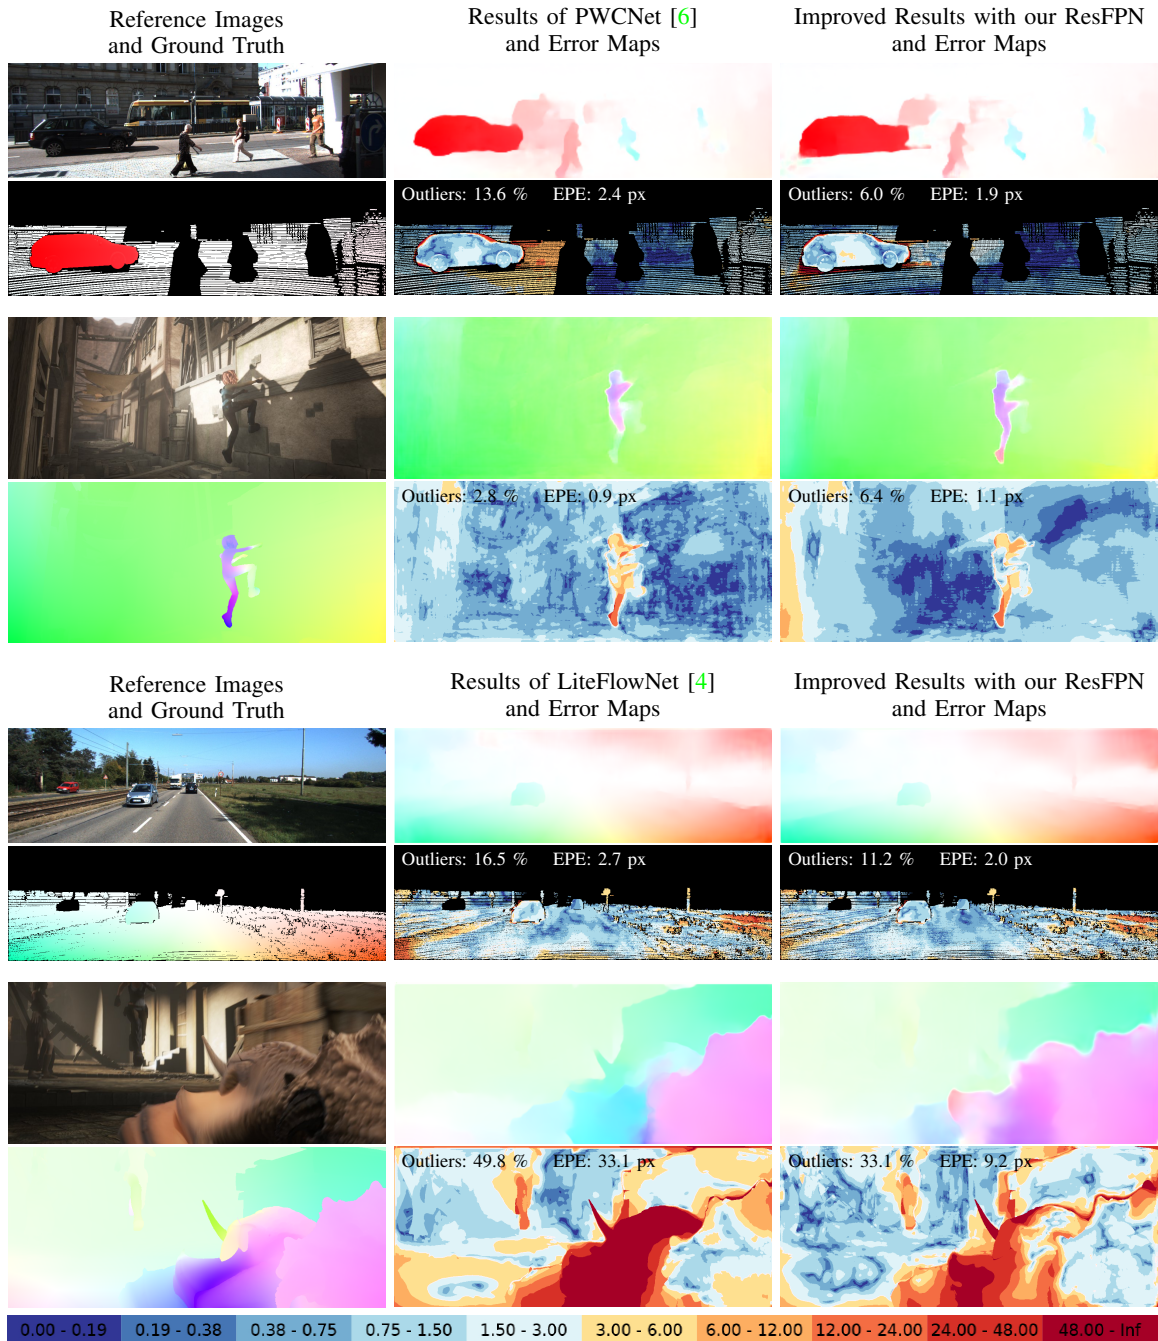


Fig. 4: Some examples of how ResFPN improves optical flow prediction on KITTI and Sintel. Please note the subtle differences around objects, e.g. vehicles. More visual comparisons for all networks and data sets are presented in our supplementary video.

The numbers indicate that results obtained from our feature module are better at discontinuities around objects in most cases. Except for very narrow evaluation regions for predictions with PSMNet [1] and wide boundaries for scene flow prediction with PWOC-3D [5], ResFPN reduces the error in these difficult image regions.

## V. CONCLUSION

In this paper, we presented ResFPN – a multi-resolution feature pyramid network with residual skip connections. With

this novel design we were able to significantly improve the representativity and localization of feature description for end-to-end learned dense pixel matching tasks. We validated our design in a comprehensive ablation study. In various experiments, we showed that ResFPN achieves significant improvements in application for optical flow, scene flow and disparity estimation. These improvements have been confirmed for a wide range of state-of-the-art methods over a large number of renowned data sets.

As future work, we plan to explicitly consider further input modalities like LiDAR [40] or radar [41] in the design of ResFPN. The additional 3D information plays an essential role for various applications. Furthermore, we want to improve ResFPN with respect to its robustness against disturbances in the input data [37].

## REFERENCES

- [1] J.-R. Chang and Y.-S. Chen, "Pyramid stereo matching network," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 3, 4, 5, 6, 7
- [2] D. G. Lowe, "Object recognition from local scale-invariant features," in *International Conference on Computer Vision (ICCV)*, 1999. 1, 2
- [3] R. Zabih and J. Woodfill, "Non-parametric local transforms for computing visual correspondence," in *European Conference on Computer Vision (ECCV)*, 1994. 1
- [4] T.-W. Hui, X. Tang, and C. Change Loy, "LiteFlowNet: A lightweight convolutional neural network for optical flow estimation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 4, 6, 7
- [5] R. Saxena, R. Schuster, O. Wasenmüller, and D. Stricker, "PWOC-3D: Deep occlusion-aware end-to-end scene flow estimation," in *Intelligent Vehicles Symposium (IV)*, 2019. 1, 2, 3, 4, 5, 6, 7
- [6] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz, "PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 3, 4, 6, 7
- [7] C. Bailer, K. Varanasi, and D. Stricker, "CNN-based patch matching for optical flow with thresholded hinge embedding loss," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [8] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 3, 5, 6
- [9] L. Zhu, R. Deng, M. Maire, Z. Deng, G. Mori, and P. Tan, "Sparsely aggregated convolutional networks," in *European Conference on Computer Vision (ECCV)*, 2018. 1
- [10] M. Menze and A. Geiger, "Object scene flow for autonomous vehicles," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 4, 5, 6
- [11] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black, "A naturalistic open source movie for optical flow evaluation," in *European Conference on Computer Vision (ECCV)*, 2012. 2, 4, 6
- [12] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, "A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 4, 5, 6
- [13] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005. 2
- [14] E. Tola, V. Lepetit, and P. Fua, "Daisy: An efficient dense descriptor applied to wide-baseline stereo," *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2009. 2
- [15] C. Bailer, B. Taetz, and D. Stricker, "Flow fields: Dense correspondence fields for highly accurate large displacement optical flow estimation," in *International Conference on Computer Vision (ICCV)*, 2015. 2
- [16] Y. Hu, R. Song, and Y. Li, "Efficient coarse-to-fine patchmatch for large displacement optical flow," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [17] L. Xu, J. Jia, and Y. Matsushita, "Motion detail preserving optical flow estimation," *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2011. 2
- [18] R. Schuster, O. Wasenmüller, G. Kusch, C. Bailer, and D. Stricker, "Sceneflowfields: Dense interpolation of sparse scene flow correspondences," in *Winter Conference on Applications of Computer Vision (WACV)*, 2018. 2
- [19] C. B. Choy, J. Gwak, S. Savarese, and M. Chandraker, "Universal correspondence network," in *Advances in Neural Information Processing Systems (NIPS)*, 2016. 2
- [20] R. Schuster, O. Wasenmüller, C. Unger, and D. Stricker, "SDC – Stacked Dilated Convolution: A unified descriptor network for dense matching tasks," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [21] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox, "Flownet: Learning optical flow with convolutional networks," in *International Conference on Computer Vision (ICCV)*, 2015. 2, 5
- [22] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "Flownet 2.0: Evolution of optical flow estimation with deep networks," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [23] A. Ranjan and M. J. Black, "Optical flow estimation using a spatial pyramid network," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [24] Y. LeCun, L. Bottou, Y. Bengio, P. Haffner *et al.*, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, 1998. 2
- [25] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems (NIPS)*, 2012. 2
- [26] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *International Conference on Learning Representations (ICLR)*, 2015. 2
- [27] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [28] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [29] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, Inception-ResNet and the impact of residual connections on learning," in *Conference on Artificial Intelligence (AAAI)*, 2017. 2
- [30] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [31] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-assisted Intervention (MICCAI)*, 2015. 2
- [32] A. Shrivastava, R. Sukthankar, J. Malik, and A. Gupta, "Beyond skip connections: Top-down modulation for object detection," *arXiv preprint arXiv:1612.06851*, 2016. 2
- [33] Y. Xin, S. Wang, L. Li, W. Zhang, and Q. Huang, "Reverse densely connected feature pyramid network for object detection," in *Asian Conference on Computer Vision (ACCV)*, 2018. 2, 3
- [34] X. Zhao, W. Li, Y. Zhang, F. Zhang, S. Chang, and Z. Feng, "Residual dilation based feature pyramid network," in *International Conference on Image Processing (ICIP)*, 2019. 2
- [35] X. Zhao, W. Li, Y. Zhang, S. Chang, Z. Feng, and P. Zhang, "Aggregated residual dilation-based feature pyramid network for object detection," *IEEE Access*, 2019. 2
- [36] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep high-resolution representation learning for human pose estimation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [37] A. Ranjan, J. Janai, A. Geiger, and M. J. Black, "Attacking optical flow," in *International Conference on Computer Vision (ICCV)*, 2019. 2, 8
- [38] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *International Conference on Machine Learning (ICML)*, 2013. 4
- [39] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 4
- [40] R. Battrawy, R. Schuster, O. Wasenmüller, Q. Rao, and D. Stricker, "LiDAR-Flow: Dense scene flow estimation from sparse lidar and stereo images," in *International Conference on Intelligent Robots and Systems (IROS)*, 2019. 8
- [41] M. Meyer and G. Kusch, "Deep learning based 3d object detection for automotive radar and camera," in *European Radar Conference (EuRAD)*, 2019. 8