

## Revisiting Depth Completion from a Stereo Matching Perspective for Cross-domain Generalization

Luca Bartolomei<sup>\*,†</sup> Matteo Poggi<sup>\*,†</sup> Andrea Conti<sup>†</sup> Fabio Tosi<sup>†</sup> Stefano Mattoccia<sup>\*,†</sup>

<sup>\*</sup>Advanced Research Center on Electronic System (ARCES)

<sup>†</sup>Department of Computer Science and Engineering (DISI)

University of Bologna, Italy

{luca.bartolomei5, m.poggi, andrea.conti35, fabio.tosi5 stefano.mattoccia}@unibo.it

<https://vppdc.github.io/>




Figure 1. **Synth-to-real generalization.** Given an NYU Depth V2 [30] frame and 500 sparse depth points (a), our framework with RAFT-Stereo [24] trained only on the SceneFlow [28] synthetic dataset (e) outperforms the generalization capability of state-of-the-art depth completion networks NLSPN [31] (b), SpAgNet [14] (c), and CompletionFormer [60] (d) – all trained on the same synthetic dataset.

### Abstract

*This paper proposes a new framework for depth completion robust against domain-shifting issues. It exploits the generalization capability of modern stereo networks to face depth completion, by processing fictitious stereo pairs obtained through a virtual pattern projection paradigm. Any stereo network or traditional stereo matcher can be seamlessly plugged into our framework, allowing for the deployment of a virtual stereo setup that is future-proof against advancement in the stereo field. Exhaustive experiments on cross-domain generalization support our claims. Hence, we argue that our framework can help depth completion to reach new deployment scenarios.*

### 1. Introduction

Perceiving depth is of utmost importance for several computer vision tasks, such as autonomous driving, robotics, and augmented reality, to name a few of them. Different camera setups exist for this purpose in the literature, ranging from single to multiple imaging systems, and deep network approaches dominate these fields. An alternative strategy to infer depth in practical applications relies on deploying active depth sensors based on LiDAR or Time-of-Flight technologies, despite they feature much lower reso-

lution than conventional cameras. To overcome the limited depth density of active depth sensors, they are frequently coupled with higher-resolution imaging devices in stereo or monocular camera setups to achieve dense depth maps at higher resolution. The monocular setup, consisting of a conventional camera registered with an active depth sensor, has received much attention recently due to the minimalist setup needed and the outstanding results achieved by learning-based methods to tackle this task, known in the literature as depth completion. Despite these achievements, state-of-the-art networks struggle with out-of-domain data distribution, making their practical deployment challenging. In fact, these networks are tightly constrained to the depth sensor/data distribution and environments seen during training [18]. Nevertheless, increasing generalization robustness has been often overlooked and consequently scarcely investigated in the literature, with most studies limiting the training and testing of depth completion solutions over a single, very specific domain – either indoor [30] or outdoor [44].

Purposely, this paper aims to fill this gap by studying this issue deeply and tackling it from a different perspective. Inspired by the outstanding generalization capability of modern stereo networks [41, 49], we cast depth completion as if it were a stereo problem through a depth-based virtual pattern projection paradigm. In contrast to most depth completion approaches conceived to densify the input sparse depth

seeds according to the image content and monocular cues only, our strategy treats depth completion as a correspondence problem through existing stereo matchers. This is achieved by processing virtual stereo pairs characterized by less domain-specific features, enabling much higher robustness to out-of-domain issues.

Our main contributions can be summarized as follows:

- We cast depth completion as a virtual stereo correspondence problem, where two appropriately patterned virtual images enable us to face depth completion with robust stereo-matching algorithms or networks.
- Extensive experimental results with multiple datasets and networks demonstrate that our proposal vastly outperforms state-of-the-art concerning generalization capability, as shown in Fig. 1, and performs comparably in-domain.

## 2. Related Work

This section reviews the depth completion and stereo matching literature since both relate to our proposal.

**Depth Completion.** Given a sparse and potentially noisy depth map, gathered even from passive source [48], depth completion aims to predict missing values. The integration of an RGB camera provides a promising solution to address the challenges posed by pure sparse depth measurements, leading to more robust and reliable results. A variety of traditional methods tackle this task via interpolation [5], stochastic models [36], morphological operators [22] or geometric models [61]. Earlier deep learning approaches used convolutional neural networks (CNNs) [27], despite their suboptimal performance with sparse inputs. Some efforts, like [44], attempted to handle sparse measurements using sparsity-invariant CNNs. However, more recent advancements in guided spatial propagation-based networks have shown better results. For example, [26] proposed a network able to learn local affinities, which was further improved by [7] and then by [31]. [60] enhanced the latter approach using a jointly convolution-attention mechanism integrated into the encoder-decoder architecture. Finally, regardless of the works above, [11] exploits 3D geometric cues recasting depth completion as a stereo matching problem. Their complex framework consists of: i) an image inpainting network to synthesize the missing stereo image by filling the warped sparse RGB image using the reference RGB image as context; ii) a depth-guided stereo matching network originally proposed by [46]; iii) a depth refinement module. While our research and [11] share some similarities, it is important to highlight the key differences. In their work, the inpainting module generates an image that may be ambiguous in uniform regions and repetitive patterns. In contrast, our approach focuses on extracting valuable information from sparse depth points using a virtual stereo pair with highly distinctive patterns. Additionally, the synthesis process it-

self is learned and thus suffers from a lack of generalization across domains or when dealing with very sparse data. Instead, we leverage our virtual matching-friendly pattern to transform depth points into the image modality efficiently.

**Stereo Matching.** Stereo matching aims at reconstructing a 2.5D scene from a rectified pair of images. Traditional methods [38] rely on handcrafted algorithms, considering local [54] and global information [3, 21, 39, 45, 51]. [17] introduced a polynomial approximation method for efficient yet accurate stereo matching. Recently, deep-learning solutions [34] outperformed conventional stereo on standard benchmarks [55]. Initially, neural networks replaced some steps of the traditional pipeline. However, a paradigm shift has occurred with the emergence of end-to-end approaches. These approaches diverge based on their utilization of either a 2D encoder-decoder architecture [23, 28, 35, 40, 53] or a 3D architecture that relies on feature cost volumes [6, 8, 20, 37, 56]. More recent works [24, 49], instead, propose novel deep stereo networks that leverage the iterative refinement paradigm from the state-of-the-art optical flow network RAFT [41], or rely on Vision Transformers [50] to capture long-range contextual information. Nevertheless, deep stereo methods often encounter challenges when estimating depth in unseen scenarios, and different methods have been proposed to address this issue [1, 4, 12, 25, 43, 47, 57]. In tackling these challenges, other solutions have explored the integration of depth points obtained from external sensors (*e.g.* LiDAR) to enhance depth estimation either by concatenating them as input of CNN-based architectures [9, 10, 32, 46, 58] or by using them to guide the cost aggregation of existing cost volumes [19, 33, 46, 59]. In contrast, [2] augments the given stereo pair to enhance RGB images, providing more discriminative information to the network and making it easier to solve the correspondence problem. Guided by the insights of this latter work in the stereo domain, we improve and extend the paradigm to the depth completion task.

## 3. Virtual Pattern Projection for Depth Completion (VPP4DC)

This section describes our approach for robustly facing depth completion, by deploying stereo matchers that process purely hallucinated image pairs generated according to the virtual pattern projection paradigm [2] – as if the scene were framed by two fictitious cameras and a pattern projector hitting the scene in sparse regions. The intuition behind this choice is to leverage the robustness of state-of-the-art stereo matchers [24, 49] at locating features across frames and domains to overcome the intrinsic limitations of conventional completion frameworks.




Figure 2. **Overview of the basic VPP4DC paradigm.** On the left, the proposed stereo setup we designed on top of the standard depth completion sensor suite enclosed in the green area. On the right, an outline of the proposed random projection that allows feeding a stereo matcher with a fictitious virtual patterned stereo pair and, optionally, an RGB image to tackle depth completion.

### 3.1. Virtual Stereo Setup

Given the standard setup for depth completion enclosed in the green area in Fig. 2 – consisting of a depth sensor (i) and an optional RGB camera (ii) – our proposal casts the task as a stereo correspondence problem using a virtual stereo setup with two fictitious cameras, one in the same position as the actual RGB device if present (ii), and the other (iii) at a distance  $b$ , *i.e.* the virtual stereo baseline. While the focal length  $f$  of the virtual cameras is constrained by the depth sensor (i) or the RGB camera (ii), the virtual stereo baseline  $b$  is a hyper-parameter.

We assume that the real RGB camera and the depth sensor are calibrated and we set the origin of the reference system in the camera. Therefore, we can project [38] sparse depth points  $\mathbf{Z}$  in the reference RGB camera view using the camera matrix  $\mathbf{K}_r$  and the roto-translation  $[\mathbf{R}_r|\mathbf{T}_r]$  between the depth sensor and the RGB camera:

$$\mathbf{Z}_r = \mathbf{K}_r [\mathbf{R}_r|\mathbf{T}_r] \mathbf{Z} \quad (1)$$

where  $\mathbf{Z}_r$  is the sparse depth map projected into the reference image plane. The proximity of the depth sensor and RGB camera can reduce occlusion issues [13] when projecting, although they cannot be entirely avoided – yet, can be easily identified and filtered out [13].

Then, we place an additional target virtual camera sharing the same intrinsics  $\mathbf{K}_r$  of the other virtual device at a horizontal distance to create a virtual baseline  $b$ . Although we will stick to this setup, it is worth noting that the target virtual camera is not constrained to the horizontal axis.

### 3.2. Virtual Pattern Projection

In the outlined setup, we aim to project onto the two fictitious cameras appropriate virtual patterns coherent with the 3D structure of the scene framed by the depth sensor, as if a projector were present in the setup [2]. At first, the sparse depth points are converted to the disparity domain using the parameters of the virtual stereo rig [38] as follows:

$$D_r = \frac{b \cdot f}{Z_r} \quad (2)$$

where  $Z_r$  is the sparse depth map aligned with the reference image,  $b$  is the virtual baseline, and  $f$  is the focal length of the virtual cameras (the same as the RGB camera).  $D_r$  is the sparse disparity map aligned with the reference virtual image  $I_r$  and the RGB image  $I$ .

Given our setup and the sparse depth points converted into disparity values, we can project the same pattern onto the fictitious reference  $I_r$  and target  $I_t$  cameras for each point  $(x, y)$  with an available disparity value  $d(x, y)$  in the reference image. It can be done by recalling [38] that with a calibrated stereo system, the disparity  $d(x, y)$  links one point  $I_r(x, y)$  in the reference image with the corresponding  $I_t(x', y)$  point in the target, with  $x' = x - d(x, y)$ . Once the two fictitious images have been generated, a stereo matcher processes them and produces a disparity map, that is then triangulated back into a densified depth map.

For projection: we manage real-valued disparities and occlusions, respectively by i) applying *weighted splatting* in the target image and ii) reprojecting the foreground pattern on occluded regions, as in [2]. Independently of the pattern choice, discussed next, the process outlined is feasible only for a subset of the image points, and we set other points to a constant color (e.g., black in all our experiments). Therefore, from a different point of view, each fictitious camera gathers sparse content coherent with the 3D structure of the scene only where a fictitious virtual pattern projector sends its rays. Regarding the virtually projected patterns, we outline the two following strategies.

**RGB Projection.** We project onto the two fictitious images the same content  $I(x, y)$  from the real camera, for each pixel with an available disparity value:

$$I_r(x, y) \leftarrow I(x, y), \quad I_t(x', y) \leftarrow I(x, y), \quad (3)$$

$$x' = x - d(x, y)$$

**Random Pattern Projection.** Instead of warping the image content, we project more *matching-friendly* patterns. Following [2], we project coherently distinctive patterns onto the two fictitious cameras:

$$I_r(x, y) \leftarrow \mathcal{P}, \quad I_t(x', y) \leftarrow \mathcal{P}, \quad (4)$$

$$x' = x - d(x, y)$$




Figure 3. **Adaptive random patches.** The patch adapts to the image content; we pick a random color for each projected point. The figure shows only the fictitious patterned image  $I_r$ .

where operator  $\mathcal{P}$  generates a random point-wise pattern applied coherently to both images. Fig. 2 shows the output of this virtual projection strategy as an intermediate output of the VPP4DC module.

On the one hand, compared to RGB Projection the random patterns are inherently less ambiguous by construction, for instance, in regions featuring a uniform texture. On the other hand, the sparse patterning prevents a complete awareness of the whole scene content for both strategies. However, such cue can be partially recovered from the RGB image if the stereo matcher can exploit image context [24, 49].

### 3.3. Additional Virtual Projection Strategies

We extend the strategy outlined so far to i) increase pattern density according to the RGB content and ii) handle issues regarding the horizontal field of view in the stereo system.

#### 3.3.1 Adaptive Patch-based Pattern Projection

The basic point-wise patterning strategy can be extended to increase the pattern density, as proposed in [2], at nearby points simply by assuming the same disparity value locally. However, it can lead to the degradation of fine details as they can be lost adopting this method. Purposely, we exploit the RGB image  $I$  to overcome this issue since it contains dense and meaningful cues about the scene. Specifically, we propose a heuristic inspired by the bilateral filter [42] to adapt the shape of the patch and handle overlapping patches.

Given a fixed size patch  $\mathcal{N}(x, y)$  centered on an available disparity point  $(x, y)$ , for each nearby point  $(x+x_w, y+y_w)$  within  $\mathcal{N}(x, y)$ , we estimate its *consistency* with the central point as:

$$W_{(x+x_w, y+y_w)} = e^{-\left(\frac{(x_w)^2 + (y_w)^2}{2\sigma_y^2} + \frac{(I(x+x_w, y+y_w) - I(x, y))^2}{2\sigma_i^2}\right)} \quad (5)$$

with  $\sigma_{xy}$ ,  $\sigma_i$  hyper-parameters. Then, we project a random value onto the two fictitious images only for those points with a similarity score  $W_{(x+x_w, y+y_w)}$  higher than a threshold  $t_{\text{adpt}}$ , a hyper-parameter of the method. Fig. 3 illustrates how the shape of the patch adapts to image content. Additionally, we update the upper threshold similarity scores in a data structure of the image size initialized to zero for




Figure 4. **Left padding on a fictitious patterned stereo pair.** The padding is visually highlighted using a teal color on  $I_r$  (left) and  $I_t$  (right). It allows us to project out-of-image warped points in  $I_t$ .

each available sparse disparity point. Hence, we can project the random pattern with the highest score for overlapping patches.

#### 3.3.2 Image Padding

As for any stereo setup, our made of two virtual cameras inherits a well-known issue: the cameras do not frame a completely overlapping portion of the scene. Specifically, in our setup depicted in Fig. 2, the left border of the reference image will not be visible in the target image. However, since we have complete control over image generation, we can easily eliminate this issue by extending the field of view of our fictitious cameras on the left side, by applying image padding to account for the largest warped point out of the image  $w_{\text{out}}$ . Accordingly, we can project virtual patterns that otherwise would pop out the left image border. Ultimately, the only trick needed is left cropping the output dense disparity map. Fig. 4 shows how padding works.

## 4. Experiments

We now present our experimental evaluation, including implementation details, datasets, and analysis of results.

### 4.1. Implementation Details

We implement RGB and Random pattern projection in Python, both with padding and adaptive patches. The hyper-parameters include  $\sigma_{xy} = 1$ ,  $\sigma_i = 1$ ,  $t_{\text{adpt}} = 0.001$  for adaptive patch. We conduct experiments with various stereo matchers: RAFT-Stereo [24], IGEV-Stereo [49], GMStereo [50], PSMNet [6], OpenCV’s SGM implementation [17, 29], and SDC [11]. We re-implemented SDC from scratch due to the unavailability of the source code. Depth refinement was omitted for a fairer comparison, although it has been shown to have a low impact on final accuracy [52]. Then, we compare VPP4DC in its best setting against conventional completion frameworks: NLSPN [31], CompletionFormer [60], and SpAgNet [14].

We will consider several cross-domain settings, either from synthetic to real data or across real datasets. Purposely, we will set the SceneFlow dataset [28] as the synthetic training dataset: for stereo matchers, we use the

Patch size	Hyperparameters		MAE (m)	
	Adaptive patch	Left Padding	NYU [30]	KITTI DC [44]
1 × 1	✗	✗	0.090	0.412
1 × 1	✗	✓	0.094	<b>0.408</b>
3 × 3	✗	✗	0.086	0.418
3 × 3	✗	✓	0.083	0.414
3 × 3	✓	✗	0.086	0.411
3 × 3	✓	✓	0.086	<b>0.406</b>
5 × 5	✗	✗	0.084	0.431
5 × 5	✗	✓	<u>0.079</u>	0.427
5 × 5	✓	✗	0.084	0.412
<b>5 × 5</b>	<b>✓</b>	<b>✓</b>	<b>0.080</b>	<b>0.409</b>
7 × 7	✗	✗	0.083	0.485
7 × 7	✗	✓	<b>0.078</b>	0.480
7 × 7	✓	✗	0.084	0.434
7 × 7	✓	✓	0.080	0.432
9 × 9	✗	✗	0.085	0.561
9 × 9	✗	✓	<b>0.078</b>	0.556
9 × 9	✓	✗	0.082	0.486
9 × 9	✓	✓	<b>0.078</b>	0.486

Table 1. **Hyper-parameters study.** Results on NYU and KITTI DC with a RAFT-Stereo model trained on synthetic data.

authors’ weights for IGEV-Stereo and GMStereo and re-trained RAFT-Stereo and PSMNet following [2], while we train from scratch completion frameworks by adapting their recommended settings to this specific dataset. Concerning generalization across real datasets, we perform either new training from scratch or fine-tuning of the synthetic pre-trained models. A detailed description of the training protocols, loss functions, and any hyper-parameter setting is reported in the supplementary material.

## 4.2. Evaluation Datasets & Protocol

Four datasets are used in our experiments, comprising scenes captured in both indoor and outdoor settings.

**KITTI DC [44].** The KITTI depth completion dataset depicts driving scenarios, captured at about  $1280 \times 384$  pixels resolution with sparse ground-truth depth maps collected using a LiDAR sensor. It provides 90K samples with RGB images, aligned sparse depth information, and semi-dense ground-truth data. We use the original training/validation splits, and each sample is top-cropped by 100px and then center-cropped to  $1216 \times 240$ . Raw LiDAR points are filtered according to distance from the minimum [13] within a local  $7 \times 7$  patch in any experiment.

**DDAD [16].** The DDAD dataset captures driving scenes using multiple synchronized cameras, measuring depths up to 250 meters. It includes about 12K training samples with RGB images and ground-truth depth maps aligned with each of the four cameras. The validation set comprises 3950 samples and ground-truth depth maps for each camera. We run experiments on validation samples from camera 1, where we sample about 20% depth points from ground-truth – filtered as done for KITTI raw LiDAR [13] – and evaluate at  $1936 \times 1216$  resolution.

**NYU Depth V2 [30].** The NYU Depth V2 dataset comprises 464 indoor scenes captured using a Kinect sensor. Each sample was downsampled to  $320 \times 240$  and center-

Model	MAE(m)					
	RGB projection		Synthesis Net [11]		Random projection	
	NYU	KITTI DC	NYU	KITTI DC	NYU	KITTI DC
OpenCV-SGM [17]	0.152	0.955	2.125	13.938	0.150	0.817
RAFT-Stereo [24]	<b>0.096</b>	<b>0.463</b>	1.211	7.937	<b>0.080</b>	<b>0.409</b>
IGEV-Stereo [49]	0.119	0.609	0.908	6.178	0.107	0.604
GMStereo [50]	0.251	0.756	0.914	4.913	0.264	0.881
PSMNet [6]	0.175	1.076	1.288	6.104	0.166	1.009
SDC [11]	0.134	0.678	<b>0.293</b>	<b>0.830</b>	0.132	0.613

Table 2. **VPP4DC with off-the-shelf stereo networks.** Results on NYU and KITTI DC. Networks trained on SceneFlow [28].

cropped to  $304 \times 228$ . As for the previous datasets, we used the original train/validation split for training and testing.

**VOID [48].** The VOID dataset provides synchronized color images at  $640 \times 480$  and sparse depth maps of indoor and outdoor scenes. It includes about 1500 depth points obtained via XIVO [15] together with dense ground-truth depth sensed by an active stereo camera. Of the total 56 sequences, 8 are used for testing purposes defining three benchmarks: VOID150 (150 points per frame), VOID500 (500 points), and VOID1500 (1500 points).

**Evaluation Protocol.** We assess the performance of all methods using root mean square error (RMSE) and mean absolute error (MAE). Specifically, we evaluate our computed disparity maps converted into depth maps, for each pixel with valid ground-truth depth values. In the remainder, we will highlight the best results in a given configuration using **bold** font and indicate the absolute bests in **red**.

## 4.3. Ablation Study

We assess the impact of various design choices in VPP4DC, by validating over 100 images sampled from each of the KITTI DC, DDAD, NYU, and VOID training sets.

**Hyper-Parameters.** Tab. 1 collects the MAE achieved on the NYU and KITTI DC datasets by varying different hyper-parameters of the VPP4DC framework while maintaining the stereo model unchanged – i.e., RAFT-Stereo [24]. The virtual baseline is fixed at 0.15 and 0.54 for NYU and KITTI DC respectively. We conduct experiments with the virtual pattern at pixel level or patches – either without or with the adaptive pattern – without or with left padding. By applying padding to the pointwise pattern improves the results on KITTI DC, but leads to a small drop in accuracy on NYU. Switching to patchwise patterns yields consistent improvements on NYU, mainly due to the much sparser sets of depth points used on this benchmark compared to the raw LiDAR scans used in KITTI. However, it has a negative impact on KITTI, except when employing small –  $3 \times 3$  or  $5 \times 5$  – adaptive patches. With patches, left padding is generally beneficial on both datasets. Based on this study, we can conclude that using a  $7 \times 7$  adaptive patch yields the best results on NYU, while a  $3 \times 3$  adaptive patch is optimal for KITTI DC, with padding being used for both. Therefore, for the subsequent evaluations, we choose a configuration yielding a good trade-off on both datasets – i.e., a  $5 \times 5$

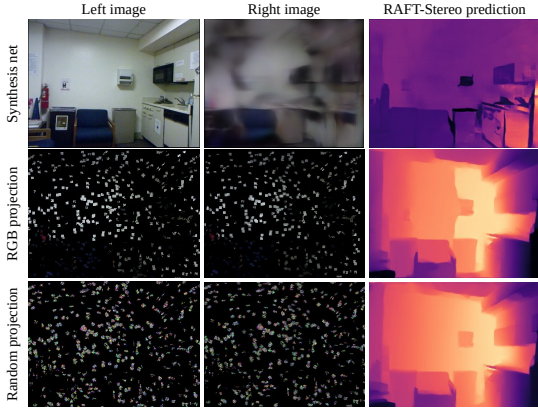


Figure 5. **Comparison of generated stereo pairs.** From top to bottom: stereo pairs by [11], RGB pattern projection and Random pattern projection. For each, we report the output of RAFT-Stereo.




Figure 6. **Qualitative generalization analysis of synthesizer network [11].** The network cannot handle high level details well.

adaptive patch with left padding, in orange in the table.

**Stereo Models.** In Tab. 2, we evaluate VPP4DC in combination with various stereo matchers. The virtual baseline is fixed at 0.15 and 0.54 for NYU and KITTI DC respectively. At the very top, we report the results obtained using the OpenCV SGM implementation [17] as a reference, followed by recent stereo networks and SDC [11] trained on the SceneFlow dataset [28]. On three different columns, we apply different strategies for generating the stereo images processed by the different methods:  $5 \times 5$  RGB projection, the use of an image synthesis model as proposed in [11], and  $5 \times 5$  Random projection. Starting from the left, we observe that PSMNet and GMStereo demonstrate lower accuracy compared to other networks and the SGM algorithm itself. We ascribe the better accuracy achieved by RAFT-Stereo and IGEV-Stereo to the use of a contextual network within their architecture. However, SDC compensates the lack of this latter by processing the sparse depth points as a direct input to the network itself. By generating a dense right image through the synthesis network and processing it together with the real left image, the accuracy of any method drops, with SDC being the best-performing model under this setting. We attribute this outcome to the lack of high-frequency details in the synthetically generated right




Figure 7. **Virtual baseline analysis.** VPP4DC performance with different virtual baselines.

Network	Test domain	RMSE (m)	MAE (m)
NLSPN [31]	NYU [30]	0.716	0.440
SpAgNet [14]	NYU [30]	0.292	0.158
CompletionFormer [60]	NYU [30]	0.374	0.186
VPP4DC (ours)	NYU [30]	<b>0.247</b>	<b>0.077</b>
NLSPN [31]	VOID500 [48]	2.394	0.972
SpAgNet [14]	VOID500 [48]	0.782	0.366
CompletionFormer [60]	VOID500 [48]	2.617	1.290
VPP4DC (ours)	VOID500 [48]	<b>0.614</b>	<b>0.188</b>
NLSPN [31]	KITTI DC [44]	2.076	1.335
SpAgNet [14]	KITTI DC [44]	1.788	0.518
CompletionFormer [60]	KITTI DC [44]	1.935	0.952
VPP4DC (ours)	KITTI DC [44]	<b>1.609</b>	<b>0.413</b>
NLSPN [31]	DDAD [16]	11.612	3.498
SpAgNet [14]	DDAD [16]	13.236	4.578
CompletionFormer [60]	DDAD [16]	9.959	2.518
VPP4DC (ours)	DDAD [16]	<b>7.303</b>	<b>1.529</b>

Table 3. **Synthetic-to-real generalization.** All networks are trained on SceneFlow [28] and tested on real datasets.

view due to the very low density of the depth points used to guide the synthesis process. Finally, using the random pattern enables any model to achieve the lowest errors, with RAFT-Stereo proving to be the architecture better suited to exploit it. Thus, we use Random pattern projection coupled with RAFT-Stereo in our framework for the remaining experiments. We refer to it as VPP4DC.

Fig. 5 shows a qualitative example of the results obtained by RAFT-Stereo using virtual stereo pairs generated through the three approaches. Using the left and synthesized right images yields poor results due to the lack of details in the generated view. We ascribe this behaviour to the lack of generalization capabilities by the synthesis network: Fig. 6 shows how images generated in the training domain appear detailed, with their quality degrading when applied to the KITTI DC dataset. On the contrary, Fig. 5 shows that RGB projection produces qualitatively good results already, with Random projection further facilitating the recovery of fine structures in the final prediction.

**Virtual Baseline.** Fig. 7 plots how the virtual baseline length impacts the accuracy of VPP4DC. For each baseline  $b$ , we plot the relative accuracy as the minimum MAE achieved across all the baselines over the MAE computed with baseline  $b$ . In general, a tiny baseline dampens the resolution at farther distances, reducing accuracy in those regions. In contrast, a wide baseline produces larger disparity values, potentially falling outside the disparity distribution observed during training by most stereo networks [28]. We can observe similar trends across driving datasets, such as

Network	Train on KITTI DC [44]				Train on NYU [30]			
	DDAD [16]		VOID500 [48]		DDAD [16]		VOID500 [48]	
	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)
NLSPN [31]	11.646	4.621	5.627	4.196	20.180	6.882	0.802	0.381
SpAgNet [14]	18.247	9.130	1.153	0.485	36.728	21.732	<b>0.752</b>	0.326
CompletionFormer [60]	<b>9.606</b>	3.328	11.640	9.856	16.479	5.649	0.821	0.385
VPP4DC (ours)	10.247	<b>2.290</b>	<b>0.934</b>	<b>0.356</b>	<b>9.246</b>	<b>3.001</b>	0.840	<b>0.307</b>

(a)

Network	Train on SceneFlow [28] + KITTI DC [44]				Train on SceneFlow [28] + NYU [30]			
	DDAD [16]		VOID500 [48]		DDAD [16]		VOID500 [48]	
	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)
NLSPN [31]	9.231	2.498	2.426	0.886	40.221	18.487	0.783	0.301
SpAgNet [14]	17.540	9.202	0.878	0.458	36.878	21.808	0.765	0.342
CompletionFormer [60]	9.471	3.607	3.418	2.294	35.590	18.928	0.929	0.429
VPP4DC (ours)	<b>7.048</b>	<b>1.580</b>	<b>0.582</b>	<b>0.187</b>	<b>6.781</b>	<b>1.344</b>	<b>0.652</b>	<b>0.187</b>

(b)

Table 4. **Real-to-real generalization.** Results in different train/test scenarios, without (a) and with (b) pre-training on SceneFlow [28].

Network	Train on SceneFlow					
	VOID150 [48]		VOID500 [48]		VOID1500 [48]	
	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)
NLSPN [31]	4.989	2.523	2.394	0.972	1.353	0.427
SpAgNet [14]	0.901	0.452	0.782	0.366	0.729	0.270
CompletionFormer [60]	4.131	2.369	2.617	1.290	1.684	0.642
VPP4DC (ours)	<b>0.748</b>	<b>0.245</b>	<b>0.614</b>	<b>0.188</b>	<b>0.606</b>	<b>0.166</b>

Table 5. **Density generalization.** Results in different train/test scenarios, without (a) and with (b) pre-training on SceneFlow [28].

DDAD and KITTI DC, where the optimal baseline is 0.30 and 0.50m, respectively, with larger baselines yielding similar results for KITTI – up to 0.80m – while increasing the error on DDAD. We ascribe this to the different distribution of depth values observed in the two datasets, showing a higher percentage of pixels at closer distances on DDAD. On datasets featuring a much lower depth range, such as NYU and VOID, the best results are achieved with shorter baselines, specifically 0.15m and 0.10m, respectively. From now on, we will use these optimal virtual baselines for the corresponding datasets.

#### 4.4. Synthetic-to-Real Generalization

We now assess the robustness of state-of-the-art depth completion architectures and our VPP4DC framework to harsh domain shifts, starting from a synthetic-to-real transfer setup. Specifically, we train VPP4DC and completion models [14, 31, 60] on the SceneFlow dataset and evaluate their accuracy on the four real benchmarks – NYU, VOID500, KITTI DC, and DDAD. Three are the main challenges for these approaches: i) the highly diverse image content depicted in the various datasets, ii) the significantly different distributions of sparse depth points, and iii) the varying image resolutions, in particular on DDAD.

Tab. 3 collects the outcome of this experiment. NLSPN often suffers the most from the domain shift, followed by CompletionFormer, partially compensating for it thanks to its global receptive field. SpAgNet proves to be highly robust by its design, specifically suited for handling different input distributions. Nevertheless, it faces significant challenges in adapting to DDAD, indicating limitations in handling high resolution. Finally, VPP4DC consistently achieves the highest accuracy, making it the preferred option for addressing synthetic-to-real generalization for depth completion.

#### 4.5. Real-to-Real Generalization

The synthetic-to-real experiment already highlights the potentially superior robustness of VPP4DC compared to well-established completion frameworks across different datasets. To confirm this, we conduct further experiments by training each method on real datasets and evaluating their accuracy in unseen, different environments. Purposely, we choose NYU and KITTI DC as the training datasets, as they have historically been the most commonly used in the depth completion literature. We then test the models on VOID500 and DDAD, creating the KITTI-to-DDAD and NYU-to-VOID500 benchmarks, where the training and testing domains are quite similar. Additionally, we establish the KITTI-to-VOID500 and NYU-to-DDAD ones, where the domain shift is more severe. Tab. 4 summarizes the results of this experiment, conducted under two settings.

On top (a), where the models are trained on either KITTI DC or NYU only, we observe that NLSPN and CompletionFormer underperform when trained on datasets significantly different from the testing domain, i.e., in the KITTI-to-VOID500 setup or the NYU-to-DDAD. SpAgNet performs the worst on DDAD while it shows considerable robustness in the KITTI-to-DDAD setting and achieves the absolute best RMSE on NYU-to-VOID500. VPP4DC demonstrates strong generalization either when trained on KITTI DC or NYUv2 solely, yielding the lowest error in most cases.

At the bottom (b), we repeat the same evaluation, this time using models being pre-trained on SceneFlow – firstly, we pre-train the models on SceneFlow, where the completion networks handle the completion task and VPP4DC focuses on stereo matching. Then, we fine-tune these pre-trained models on either KITTI DC or NYU for evaluation. Overall, pre-training appears to be beneficial for models that are subsequently trained on KITTI DC, as it improves their accuracy in most cases. Conversely, it has a negative impact – negligible or not – when the models are eventually trained on NYU, except for VPP4DC. Indeed, our method consistently achieves the best results when empowered by stereo pre-training – which is crucial for properly learning how to match images, either real or generated through our pattern – largely reducing RMSE and MAE on VOID500 especially.

Network	Train on KITTI DC [44]						Train on NYU [30]					
	VOID150 [48]		VOID500 [48]		VOID1500 [48]		VOID150 [48]		VOID500 [48]		VOID1500 [48]	
	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)
NLSPN [31]	6.587	4.933	5.627	4.196	4.784	3.591	0.963	0.492	0.802	0.381	0.737	0.298
SpAgNet [14]	1.240	0.533	1.154	0.485	1.108	0.446	<b>0.866</b>	0.408	<b>0.752</b>	0.326	<b>0.706</b>	<b>0.244</b>
CompletionFormer [60]	15.543	13.490	11.640	9.856	8.804	7.446	0.956	0.487	0.821	0.385	0.726	0.261
VPP4DC (ours)	<b>1.118</b>	<b>0.507</b>	<b>0.934</b>	<b>0.356</b>	<b>0.789</b>	<b>0.244</b>	0.960	<b>0.397</b>	0.840	<b>0.307</b>	0.800	0.253

(a)

Network	Train on SceneFlow [28] + KITTI DC [44]						Train on SceneFlow [28] + NYU [30]					
	VOID150 [48]		VOID500 [48]		VOID1500 [48]		VOID150 [48]		VOID500 [48]		VOID1500 [48]	
	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)
NLSPN [31]	4.616	1.991	2.426	0.886	1.349	0.448	1.138	0.535	0.783	0.301	0.668	0.210
SpAgNet [14]	0.959	0.503	0.878	0.458	0.875	0.420	0.874	0.418	0.766	0.342	0.718	0.253
CompletionFormer [60]	5.396	3.736	3.418	2.294	2.446	1.608	1.162	0.628	0.929	0.429	0.800	0.287
VPP4DC (ours)	<b>0.690</b>	<b>0.247</b>	<b>0.582</b>	<b>0.187</b>	<b>0.543</b>	<b>0.148</b>	<b>0.865</b>	<b>0.259</b>	<b>0.652</b>	<b>0.187</b>	<b>0.595</b>	<b>0.157</b>

(b)

Table 6. **Density generalization.** Results in different train/test scenarios, without (a) and with (b) pre-training on SceneFlow [28].

Network	NYU		KITTI DC	
	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)
SCPU [22]	0.544	0.240	1.591	0.329
NLSPN [31]	0.092	0.035	<b>0.772</b>	<b>0.197</b>
SpAgNet [14]	0.114	0.045	0.856	0.222
CompletionFormer [60]	<b>0.090</b>	<b>0.035</b>	0.849	0.216
VPP4DC (ours)	0.117	0.044	0.999	0.269

(a)

Network	NYU		KITTI DC	
	RMSE (m)	MAE (m)	RMSE (m)	MAE (m)
SCPU [22]	0.544	0.240	1.591	0.329
NLSPN [31]	0.123	0.051	1.129	0.353
SpAgNet [14]	0.104	0.043	<b>0.832</b>	<b>0.216</b>
CompletionFormer [60]	<b>0.102</b>	<b>0.042</b>	1.090	0.259
VPP4DC (ours)	0.119	0.048	1.095	0.304

(b)

Table 7. **In-domain performance.** Results without (a) and with (b) pre-training on SceneFlow [28].

## 4.6. Density Generalization

We investigate model robustness to varying input depth point densities across domains using the VOID150, VOID500, and VOID1500 benchmarks.

**Synthetic-to-Real.** We start by evaluating networks trained on the SceneFlow dataset [28]. Tab. 5 summarizes the outcome of this experiment. Among the established completion frameworks, SpAgNet demonstrates outstanding robustness to the different densities, while NLSPN and CompletionFormer suffer about  $2\times$  increase of RMSE and MAE from VOID1500 to VOID500 and from this latter to VOID150. Once again, VPP4DC surpasses all other models, affirming its robustness in this aspect too.

**Real-to-Real.** Next, we study the impact of density when the networks have been trained on KITTI DC or NYU datasets. Tab. 6 collects the results on the three VOID benchmarks. On top (a), we report the results for the models trained directly on real data. SpAgNet confirms its robustness to the varying input density. Moreover, when trained on a similar domain, such as the one of the NYU dataset, it often achieves the best results. It also outperforms VPP4DC, which proves to be more effective when generalizing from KITTI DC to VOID. At the bottom (b), we evaluate models pre-trained on SceneFlow. We observe how this pre-training strategy improves the results when moving from KITTI DC to VOID splits – i.e., very different do-

main – while it has a negligible or negative impact when moving from NYU to VOID – i.e., more similar ones. On the contrary, VPP4DC consistently benefits from stereo pre-training on SceneFlow, often halving the errors and emerging as the most robust framework to varying input density.

## 4.7. In-Domain Performance

We conclude by evaluating the accuracy achieved by VPP4DC when training and testing on the same domains. Tab. 7 reports the outcome of this final experiment on two sub-tables. On top of both, we include the results achieved by a hand-crafted completion solution [22] as a reference.

On top (a), each model is trained directly on the training set of the target dataset – NYU or KITTI DC. Notably, the well-established completion frameworks outperform VPP4DC concerning specialization within a single domain. At the bottom (b), we show the performance obtained by pre-training the networks on SceneFlow. Except for SpAgNet, the pre-training impacts negatively on the accuracy over the very same domain. Indeed, the absolute best results are obtained by CompletionFormer and NLSPN, on NYU and KITTI DC, respectively, when trained only on the target domain training set.

Despite being slightly less effective at specializing in single domains, VPP4DC still achieves competitive results, not far from the state-of-the-art. Nonetheless, we believe this small gap is a moderate price to pay for achieving much stronger generalization across very different domains. This makes VPP4DC better suited for in-the-wild deployment and a first significant step towards studying depth completion solutions with such capabilities.

## 5. Conclusion

This paper proposed a novel paradigm to achieve depth completion robust against the training environment. Combining a state-of-the-art stereo network and our framework achieves competitive results on the same domain evaluation and greatly outperforms other task-specific networks in cross-domain generalization. We believe this can pave the way for deploying depth completion in countless and exciting practical application contexts.



## References

- [1] Filippo Aleotti, Fabio Tosi, Pierluigi Zama Ramirez, Matteo Poggi, Samuele Salti, Stefano Mattoccia, and Luigi Di Stefano. Neural disparity refinement for arbitrary resolution stereo. In *2021 International Conference on 3D Vision (3DV)*, pages 207–217. IEEE, 2021. 2
- [2] Luca Bartolomei, Matteo Poggi, Fabio Tosi, Andrea Conti, and Stefano Mattoccia. Active stereo without pattern projector. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 18470–18482, 2023. 2, 3, 4, 5
- [3] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on pattern analysis and machine intelligence*, 23(11):1222–1239, 2001. 2
- [4] Changjiang Cai, Matteo Poggi, Stefano Mattoccia, and Philippos Mordohai. Matching-space stereo networks for cross-domain generalization. In *2020 International Conference on 3D Vision (3DV)*, pages 364–373, 2020. 2
- [5] Massimo Camplani and Luis Salgado. Efficient spatio-temporal hole filling strategy for kinect depth maps. In *Three-dimensional image processing (3DIP) and applications II*, pages 127–136. SPIE, 2012. 2
- [6] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5410–5418, 2018. 2, 4, 5
- [7] Xinjing Cheng, Peng Wang, and Ruigang Yang. Depth estimation via affinity learned with convolutional spatial propagation network. In *Proceedings of the European conference on computer vision (ECCV)*, pages 103–119, 2018. 2
- [8] Xinjing Cheng, Peng Wang, and Ruigang Yang. Learning depth with convolutional spatial propagation network. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2361–2379, 2019. 2
- [9] Xuelian Cheng, Yiran Zhong, Yuchao Dai, Pan Ji, and Hongdong Li. Noise-aware unsupervised deep lidar-stereo fusion. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2019. 2
- [10] Jaesung Choe, Kyungdon Joo, Tooba Imtiaz, and In So Kweon. Volumetric propagation network: Stereo-lidar fusion for long-range depth estimation. *IEEE Robotics and Automation Letters*, 6(3):4672–4679, 2021. 2
- [11] Keunhoon Choi, Somi Jeong, Youngjung Kim, and Kwanghoon Sohn. Stereo-augmented depth completion from a single rgb-lidar image. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13641–13647, 2021. 2, 4, 5, 6
- [12] WeiQin Chuah, Ruwan Tennakoon, Reza Hoseinnezhad, Alireza Bab-Hadiashar, and David Suter. Itsa: An information-theoretic approach to automatic shortcut avoidance and domain generalization in stereo matching networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13022–13032, 2022. 2
- [13] Andrea Conti, Matteo Poggi, Filippo Aleotti, and Stefano Mattoccia. Unsupervised confidence for lidar depth maps and applications. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022. IROS. 3, 5
- [14] Andrea Conti, Matteo Poggi, and Stefano Mattoccia. Sparsity agnostic depth completion. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 5871–5880, 2023. 1, 4, 6, 7, 8
- [15] X. Fei, A. Wong, and S. Soatto. Geo-supervised visual depth prediction. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2019. 5
- [16] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Ravnatos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5, 6, 7
- [17] Heiko Hirschmuller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on pattern analysis and machine intelligence*, 30(2):328–341, 2007. 2, 4, 5, 6
- [18] Junjie Hu, Chenyu Bao, Mete Ozay, Chenyou Fan, Qing Gao, Honghai Liu, and Tin Lun Lam. Deep depth completion from extremely sparse data: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7):8244–8264, 2023. 1
- [19] Yu-Kai Huang, Yueh-Cheng Liu, Tsung-Han Wu, Hung-Ting Su, Yu-Cheng Chang, Tsung-Lin Tsou, Yu-An Wang, and Winston H Hsu. S3: Learnable sparse signal superdensity for guided depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16706–16716, 2021. 2
- [20] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *The IEEE International Conference on Computer Vision (ICCV)*, 2017. 2
- [21] Vladimir Kolmogorov and Ramin Zabih. What energy functions can be minimized via graph cuts? *IEEE transactions on pattern analysis and machine intelligence*, 26(2):147–159, 2004. 2
- [22] Jason Ku, Ali Harakeh, and Steven L. Waslander. In defense of classical image processing: Fast depth completion on the cpu. In *2018 15th Conference on Computer and Robot Vision (CRV)*, pages 16–22, 2018. 2, 8
- [23] Zhengfa Liang, Yiliu Feng, Yulan Guo, Hengzhu Liu, Wei Chen, Linbo Qiao, Li Zhou, and Jianfeng Zhang. Learning for disparity estimation through feature constancy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [24] Lahav Lipson, Zachary Teed, and Jia Deng. Raft-stereo: Multilevel recurrent field transforms for stereo matching. In *International Conference on 3D Vision (3DV)*, 2021. 1, 2, 4, 5
- [25] Biyang Liu, Huimin Yu, and Guodong Qi. Graftnet: Towards domain generalized stereo matching with a broad-spectrum and task-oriented feature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13012–13021, 2022. 2

- [26] Sifei Liu, Shalini De Mello, Jinwei Gu, Guangyu Zhong, Ming-Hsuan Yang, and Jan Kautz. Learning affinity via spatial propagation networks. *Advances in Neural Information Processing Systems*, 30, 2017. 2
- [27] Fangchang Ma and Sertac Karaman. Sparse-to-dense: Depth prediction from sparse depth samples and a single image. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 4796–4803. IEEE, 2018. 2
- [28] N. Mayer, E. Ilg, P. Häusser, 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 *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. arXiv:1512.02134. 1, 2, 4, 5, 6, 7, 8
- [29] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 4
- [30] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012. 1, 5, 6, 7, 8
- [31] Jinsun Park, Kyungdon Joo, Zhe Hu, Chi-Kuei Liu, and In So Kweon. Non-local spatial propagation network for depth completion. In *Proc. of European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 4, 6, 7, 8
- [32] Kihong Park, Seungryong Kim, and Kwanghoon Sohn. High-precision depth estimation with the 3d lidar and stereo fusion. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2156–2163. IEEE, 2018. 2
- [33] Matteo Poggi, Davide Pallotti, Fabio Tosi, and Stefano Mattoccia. Guided stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 979–988, 2019. 2
- [34] Matteo Poggi, Fabio Tosi, Konstantinos Batsos, Philippos Mordohai, and Stefano Mattoccia. On the synergies between machine learning and binocular stereo for depth estimation from images: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):5314–5334, 2022. 2
- [35] Tonmoy Saikia, Yassine Marrakchi, Arber Zela, Frank Hutter, and Thomas Brox. Autodispnet: Improving disparity estimation with automl. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1812–1823, 2019. 2
- [36] Ju Shen and Sen-Ching S Cheung. Layer depth denoising and completion for structured-light rgb-d cameras. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1187–1194, 2013. 2
- [37] Zhelun Shen, Yuchao Dai, and Zhibo Rao. Cfnet: Cascade and fused cost volume for robust stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13906–13915, 2021. 2
- [38] Richard Szeliski. *Computer Vision - Algorithms and Applications, Second Edition*. Springer, 2022. 2, 3
- [39] Tatsunori Tanai, Yasuyuki Matsushita, and Takeshi Naemura. Graph cut based continuous stereo matching using locally shared labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1613–1620, 2014. 2
- [40] Vladimir Tankovich, Christian Hane, Yinda Zhang, Adarsh Kowdle, Sean Fanello, and Sofien Bouaziz. Hitnet: Hierarchical iterative tile refinement network for real-time stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14362–14372, 2021. 2
- [41] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, pages 402–419. Springer, 2020. 1, 2
- [42] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Computer Vision, 1998. Sixth International Conference on*, pages 839–846, 1998. 4
- [43] Fabio Tosi, Alessio Tonioni, Daniele De Gregorio, and Matteo Poggi. Nerf-supervised deep stereo. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 855–866, 2023. 2
- [44] Jonas Uhrig, Nick Schneider, Lukas Schneider, Uwe Franke, Thomas Brox, and Andreas Geiger. Sparsity invariant cnns. In *International Conference on 3D Vision (3DV)*, 2017. 1, 2, 5, 6, 7, 8
- [45] Olga Veksler. Stereo correspondence by dynamic programming on a tree. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, pages 384–390. IEEE, 2005. 2
- [46] Tsun-Hsuan Wang, Hou-Ning Hu, Chieh Hubert Lin, Yi-Hsuan Tsai, Wei-Chen Chiu, and Min Sun. 3d lidar and stereo fusion using stereo matching network with conditional cost volume normalization. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5895–5902. IEEE, 2019. 2
- [47] Jamie Watson, Oisín Mac Aodha, Daniyar Turmukhambetov, Gabriel J. Brostow, and Michael Firman. Learning stereo from single images. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [48] Alex Wong, Xiaohan Fei, Stephanie Tsuei, and Stefano Soatto. Unsupervised depth completion from visual inertial odometry. *IEEE Robotics and Automation Letters*, 5(2): 1899–1906, 2020. 2, 5, 6, 7, 8
- [49] Gangwei Xu, Xianqi Wang, Xiaohuan Ding, and Xin Yang. Iterative geometry encoding volume for stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21919–21928, 2023. 1, 2, 4, 5
- [50] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatoughi, Fisher Yu, Dacheng Tao, and Andreas Geiger. Unifying flow, stereo and depth estimation. *arXiv preprint arXiv:2211.05783*, 2022. 2, 4, 5
- [51] Qingxiang Yang, Liang Wang, and Narendra Ahuja. A constant-space belief propagation algorithm for stereo matching. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1458–1465. IEEE, 2010. 2
- [52] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. *European Conference on Computer Vision (ECCV)*, 2018. 4
- [53] Zhichao Yin, Trevor Darrell, and Fisher Yu. Hierarchical discrete distribution decomposition for match density estimation.

- tion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6044–6053, 2019. [2](#)
- [54] Ramin Zabih and John Woodfill. Non-parametric local transforms for computing visual correspondence. In *Third European Conference on Computer Vision (Vol. II)*, pages 151–158, Secaucus, NJ, USA, 1994. Springer-Verlag New York, Inc. [2](#)
- [55] Jure Zbontar, Yann LeCun, et al. Stereo matching by training a convolutional neural network to compare image patches. *J. Mach. Learn. Res.*, 17(1):2287–2318, 2016. [2](#)
- [56] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. GA-Net: Guided aggregation net for end-to-end stereo matching. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#)
- [57] Feihu Zhang, Xiaojuan Qi, Ruigang Yang, Victor Prisacariu, Benjamin Wah, and Philip Torr. Domain-invariant stereo matching networks. In *Europe Conference on Computer Vision (ECCV)*, 2020. [2](#)
- [58] Junming Zhang, Manikandasriram Srinivasan Ramanagopal, Ram Vasudevan, and Matthew Johnson-Roberson. Listereo: Generate dense depth maps from lidar and stereo imagery. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7829–7836. IEEE, 2020. [2](#)
- [59] Yongjun Zhang, Siyuan Zou, Xinyi Liu, Xu Huang, Yi Wan, and Yongxiang Yao. Lidar-guided stereo matching with a spatial consistency constraint. *ISPRS Journal of Photogrammetry and Remote Sensing*, 183:164–177, 2022. [2](#)
- [60] Youmin Zhang, Xianda Guo, Matteo Poggi, Zheng Zhu, Guan Huang, and Stefano Mattoccia. Completionformer: Depth completion with convolutions and vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18527–18536, 2023. [1](#), [2](#), [4](#), [6](#), [7](#), [8](#)
- [61] Yiming Zhao, Lin Bai, Ziming Zhang, and Xinming Huang. A surface geometry model for lidar depth completion. *IEEE Robotics and Automation Letters*, 6(3):4457–4464, 2021. [2](#)