

# Cross-View Completion Models are Zero-shot Correspondence Estimators

## - Supplementary Material -

In this supplementary material, Section A outlines the implementation of zero-shot and learning-based methods, including a pipeline overview with notations, measurements, and a description of the learnable head for geometric matching and depth estimation, supported by equations and architectural illustrations. Section B details the experimental setup for these methods, highlighting the baselines and configurations for dynamic objects and noisy environments in learning-based depth estimation. Section C examines cross-view completion [93], focusing on layer-wise feature performance, the impact of inference resolution, pretrained weights, and additional ablation studies. Section D presents extended quantitative results for zero-shot and learning-based experiments, while Section E provides supplementary qualitative results to further validate our approach. Finally, Section F discusses the limitations of our proposed method.

### A. Implementation Details

#### A.1. Zero-shot Matching

For our main CroCo [93] zero-shot experiment reported in Table 1, the encoder is based on ViT-Large [17], consisting of 24 encoder blocks, while the decoder includes 12 decoder blocks, each incorporating both self-attention and cross-attention layers. We utilized pretrained weights from CroCo-v2 [93] trained on images with a resolution of  $224 \times 224$ .

As cross-view completion models are all pretrained on a resolution of 224, we resize all input source and target images,  $I_s, I_t \in \mathbb{R}^{224 \times 224 \times 3}$ , to fully leverage the pretrained model’s knowledge. Note that with models pretrained with different resolutions (e.g. 448), inference at higher resolutions is also possible. With a patch size of 16, the image is divided into 196 patches (or tokens). We then extract three components from CroCo: the encoder features  $D_s, D_t \in \mathbb{R}^{hw \times c}$ , decoder features  $D_{s \rightarrow t}, D_{t \rightarrow s} \in \mathbb{R}^{hw \times d}$ , and the cross-attention map  $C_{att} \in \mathbb{R}^{hw \times hw}$ , where  $hw = 14 \times 14 = 196$ ,  $c = 1024$  and  $d = 768$ . The cross-attention map is computed by averaging across the attention heads. To generate the final flow field  $F$ , we compute correlations for each component and apply a soft-argmax function with a temperature of  $\tau = 1e - 4$ . Finally,  $F$  is upsampled to match the original image size.

**Register tokens** We observe that the cross-attention map  $C^l$  often exhibits artifacts caused by the register token, a phenomenon linked to shortcut learning, similar to other Transformer-based models [15, 71]. To address this issue,




Figure 5. **Visualization of the attention map with and without the register token.** The initial cross-attention map of CroCo [93] often contains artifacts due to the register tokens as in (c). After correcting this, the proper attending point is identified as in (d).

we replace the attention values of the register token with the minimum attention value. This adjustment results in more accurate attended regions in the cross-attention map, as illustrated in Fig. 5.

#### A.2. Learning-based Geometric Matching

We conducted two types of learning-based geometric matching experiments, as described in our main paper: **ZeroCo-finetuned** and **ZeroCo-flow**. Both methods are based on the same network architecture used for zero-shot matching. In ZeroCo-finetuned, the cross-attention map is learned directly, whereas in ZeroCo-flow, the cross-attention map is further refined by a learnable head. The head used in our learning-based geometric matching comprises an aggregation module and an upsampling module, as shown in Fig. 6. The number of cross-attention layers is  $L = 12$ , and the aggregation module, repeated  $N_a = 4$  times, employs Swin-Transformer [58] and vanilla Transformer [85] blocks to aggregate the cross-attention map  $C_{att} \in \mathbb{R}^{hw \times hw}$  with the compressed decoder feature  $D'_{t \rightarrow s} \in \mathbb{R}^{hw \times d'}$ , where  $hw = 196$  and  $d' = 128$ . In the upsampling module, the 16-th and 8-th encoder features  $D_t^{16}, D_t^8 \in \mathbb{R}^{hw \times c}$  with  $c = 1024$  are used to guide the upsampling process, producing the final flow field  $F_{flow} \in \mathbb{R}^{224 \times 224 \times 2}$ . For ZeroCo-finetuned, we begin by averaging the  $L$  stacked cross-attention maps and combining them in a reciprocal manner. Next, we apply the soft-argmax operator [52] to transform them into a dense flow field, which is then upsampled to produce the final flow field  $F_{flow} \in \mathbb{R}^{224 \times 224 \times 2}$ .

**Flow head details.** Our geometric matching head consists of two modules: an aggregation module and an upsampling module. The aggregation module aggregates the

cross-attention map with the decoder features, while the up-sampling module enhances the geometric matching performance by upsampling low-resolution features.

The aggregation module refines inaccurate matches in the cross-attention map  $C_{\text{att}}$  by leveraging the decoder feature  $D_{t \rightarrow s} \in \mathbb{R}^{hw \times d}$  as guidance. First, we compress the decoder feature using linear projection and concatenate it with the cross-attention map. This combined representation is then processed by Swin-Transformer and vanilla Transformer blocks for spatial and multi-layer aggregation, respectively. The final refined cross-attention map  $C' \in \mathbb{R}^{hw \times hw}$  is then obtained with reciprocal summation:

$$C' = \mathcal{T}_c(C_{\text{att}}, \mathcal{P}(D_{t \rightarrow s})) + (\mathcal{T}_c(C_{\text{att,swap}}, \mathcal{P}(D_{s \rightarrow t})))^T$$

where  $\mathcal{T}_c$  denotes the successive transformer blocks for aggregation,  $C_{\text{att,swap}}$  represents the cross-attention map from swapped inputs and  $\mathcal{P}(\cdot)$  indicates linear projection, with  $h = 14$ ,  $w = 14$ ,  $d = 768$  and  $d' = 128$ .

The map is progressively upsampled to  $C'_{\text{up}} \in \mathbb{R}^{2hw \times 2hw}$  and  $C'' \in \mathbb{R}^{4hw \times 4hw}$  through two upsampling decoder layers, with each layer guided by the previous encoder features  $D_t^{16}$  and  $D_t^8$ , respectively. This process is illustrated as follows:

$$C'_{\text{up}} = \mathcal{U}^1(C', D_t^{16})$$

$$C'' = \mathcal{U}^2(C'_{\text{up}}, D_t^8)$$

Where  $\mathcal{U}^1$  and  $\mathcal{U}^2$  both comprise of multiple deconvolutional layers. The final flow field  $F_{\text{flow}} \in \mathbb{R}^{224 \times 224 \times 2}$  is obtained by applying soft-argmax followed by bilinear interpolation to  $C''$  as the following:

$$F_{\text{flow}} = \text{softargmax}(C'').$$

**Training details.** We adopt a two-stage training scheme, similar to conventional geometric matching works [44, 64, 82, 83]. In the first stage, we train on the DPED-CityScape-ADE dataset used in [82]. In the second stage, we extend the training by including the same dataset, augmented with COCO objects [84] and MegaDepth [56], enhancing generalizability to real-world scenarios. The model is optimized using the AdamW optimizer [50] and trained the model for 150 epochs in the first stage with a learning rate of  $1e-4$ . In the second stage, the model was trained for 50 epochs with a learning rate of  $5e-5$ . We applied multistep learning rate decay [66] in both stages and optimized only the learnable head while freezing all other components of the network.

### A.3. Learning-based Multi-frame Depth Estimation

For multi-frame depth estimation, the model receives two resized source and target images, as  $I_s, I_t \in \mathbb{R}^{192 \times 640 \times 3}$ .

Here, the source and target images correspond to the  $t$ -th and  $t - 1$ -th frames, respectively, sampled from a driving scene video. These frames are patchified (with a patch size of 16) and encoded by a transformer [17], resulting in encoded features  $D_s, D_t \in \mathbb{R}^{h \times w \times d}$ , which are then processed by a cross-attention layer to generate the cross-attention map  $C_{\text{att}} \in \mathbb{R}^{hw \times hw}$ , where  $h = 12$ ,  $w = 40$  and  $d = 768$ . The cross-attention map is obtained by averaging the attention heads. Similar to DPT [72], our learnable head employed for depth estimation receives four decoder features  $D_{t \rightarrow s}^l$ , and four cross-attention maps  $C_{\text{att}}^l$ , where  $l = 1, 4, 7, 11$ . Its objective is to aggregate the cross-attention map, guided by the decoder features, and upsample the depth map to capture fine details. Aggregation is performed by the transformer, while upsampling is performed by CNN layers as shown in Fig. 6. The final projection layer of our depth head outputs a depth map  $F_{\text{depth}} \in \mathbb{R}^{96 \times 320 \times 1}$ .

**Depth head details.** Analogous to previous geometric matching head, our depth head consists of two modules: the aggregation module and the upsampling module. The aggregation module aggregates the cross-attention map with the decoder features, while the upsampling module upsamples the features to capture fine details of the depth map.

The aggregation module aims to refine inaccurate matches in the cross-attention map and reduces its degree of freedom. Using a transformer [85], we first compress the initial cross-attention map  $C_{\text{att}}$  to obtain  $C'_{\text{att}} \in \mathbb{R}^{h \times w \times c}$  as the following:

$$C'_{\text{att}} = \mathcal{T}_{\text{att}}(C_{\text{att}}),$$

where  $\mathcal{T}_{\text{att}}(\cdot)$  denotes attention aggregation. Then we use the compressed target feature  $D'_{t \rightarrow s} \in \mathbb{R}^{h \times w \times d'}$  to further refine the compressed map as the following:

$$C' = \mathcal{T}_{\text{feat}}([C'_{\text{att}}, D'_{t \rightarrow s}]),$$

where  $C' \in \mathbb{R}^{h \times w \times hw}$  indicates the refined cross-attention map and  $\mathcal{T}_{\text{feat}}(\cdot)$  denotes the feature aggregation component composed of multiple convolutional layers. This whole process can be summarized as the following:

$$C' = \mathcal{T}_d(C_{\text{att}}, D_t) = \mathcal{T}_{\text{feat}}([\mathcal{T}_{\text{att}}(C_{\text{att}}), D'_{t \rightarrow s}]),$$

where  $\mathcal{T}_d(\cdot)$  denotes the aggregation module for depth estimation.

Next, the upsampling module aims to mitigate the limitations of low-resolution attention maps in fine depth prediction. Thus, we stack the aggregation modules in a pyramidal structure to hierarchically refine the cost volumes, similar to coarse-to-fine feature enhancement methods [62]. As illustrated in Fig. 6, we adopt a four-stage pyramidal process using aggregation modules [72], selecting four layers of decoder features  $D_{t \rightarrow s}^l$ , and cross-attention maps  $C_{\text{att}}^l$ , where




Figure 6. **Main architecture for our learning-based experiments.** (a) ZeroCo-flow and (b) ZeroCo-depth architectures extend the original zero-shot architecture by incorporating learnable heads, which consist of an aggregation and upsampling module, to effectively aggregate and refine the cross-attention map.

$l = 1, 4, 7, 11$ . Also, resolutions are progressively adjusted using the Reassemble operation [72] and bilinear interpolation, enabling hierarchical feature and attention map aggregation into a refined cost volume:

$$C_{\text{ref}}^l = \mathcal{T}_d([\text{interp}(C_{\text{att}}^l + C_{\text{ref}}^{l-1})), D_{t \rightarrow s, \text{re}}^l])$$

$$D_{t \rightarrow s, \text{re}}^l = \text{Reassemble}_l(D_{t \rightarrow s}^l),$$

where  $l$  denotes the stage number,  $\text{interp}$  refers to bilinear interpolation,  $\text{Reassemble}_l$  indicates the Reassemble blocks for stage  $l$ , and  $C_{\text{ref}, l}$  represents the refined cost volume. The final refined cost volume input to the depth head has dimensions of  $C_{\text{ref}}^* \in \mathbb{R}^{4h \times 4w \times c}$ , resulting in a final depth map  $F_{\text{depth}} \in \mathbb{R}^{96 \times 320 \times 1}$  as the following:

$$F_{\text{depth}} = \mathcal{T}_{\text{depth}}(C_{\text{ref}}^*),$$

where  $\mathcal{T}_{\text{depth}}$  comprises multiple convolution layers and a final sigmoid layer.

**Training details.** We adopt the same augmentation scheme of previous depth estimation methods [27], our model is trained with input resolutions of  $192 \times 640$  for the KITTI dataset [25] and  $128 \times 416$  for the Cityscapes dataset [14]. Experiments were conducted using PyTorch [67] on an RTX 3090 GPU with a batch size of 8. We employ the Adam [50] optimizer, setting the learning rate as  $5e-6$  for the pretrained encoder and decoder, and  $5e-5$  for the remaining components. We incorporate static augmentation from ManyDepth [91] with a 50% ratio. For the photometric loss function, we set  $\lambda = 0.001$  and  $\alpha = 0.85$ . Lastly, we employ the same pose network used in previous depth estimation methods [27].

## B. Experimental Details

### B.1. Zero-shot Matching

**Datasets.** To validate the effectiveness of our zero-shot and learned geometric matching, we tested our approach

on the HPatches [3] and ETH3D [78] datasets. HPatches includes sequences of different views of the same scenes, with each sequence containing one source image and five target images, along with ground-truth optical flows. Later sequences (IV and V) present more challenging targets. ETH3D features 10 real-world 3D scene sequences with transformations beyond homographies. Following [82], we evaluate on sampled image pairs across varying intervals, covering 600K to 1,000K correspondences per interval. Larger intervals (e.g., 13 or 15) yield more challenging pairs. As per [61], we use the Average EndPoint Error (AEPE) as our metric, computed as the average Euclidean distance between estimated and ground-truth flow fields over all valid target pixels.

**Baselines.** We utilized DINOv2 [65], DIFT<sub>ad</sub> [80], DIFT<sub>sdm</sub> [80], and SD-DINO [99] as baseline models for comparison, all trained in a self-supervised manner without labels. For DINOv2, we used features extracted from the 11-th layer to construct matching costs, which are critical for achieving optimal performance. DIFT [80] employs a diffusion-based approach for zero-shot matching and is evaluated in two variants: DIFT<sub>ad</sub> and DIFT<sub>sdm</sub>. Both variants used the hyperparameter configurations originally proposed for DIFT, however, DIFT<sub>sdm</sub> exhibited instability at the original resolution due to its feature resolution being reduced by half. To enhance performance in dense matching, we adjusted the hyperparameters by setting the time step to 41, utilizing features from the 4-th layer, and employing an ensemble size of 2, which we denote as DIFT<sub>sdm</sub><sup>\*</sup>. For the HPatches-240 dataset, features were extracted at a resolution of  $224 \times 224$  for dense matching, while for the HPatches-original dataset, features were extracted at  $768 \times 768$ . SD-DINO [99], which also integrates a diffusion model with DINOv2 for zero-shot matching, was evaluated using the same hyperparameter configurations as its original implementation.

We applied L2-normalization to the features extracted from each image and constructed a correlation map to compute the matching cost. Following this, in alignment with our experimental setup, we applied a soft-argmax function with a temperature of  $\tau = 1e - 4$  to obtain the final flow field  $F$ , which was then upsampled to match the original image size.

## B.2. Learning-based Geometric Matching

**Baselines.** To ensure a fair comparison, we trained a dense matching model and evaluated it on the HPatches [3] and ETH3D [78] datasets, benchmarking against state-of-the-art models such as DiffMatch [64]. Many recent approaches, including GLU-Net [82], GoCor [81], and DiffMatch [64], were trained on DPED-Cityscapes-ADE dataset, while PDC-Net [83] and PDC-Net+ [84] extended their training to the MegaDepth [56] dataset, leveraging uncertainty modeling to address significant occlusions.

Additionally, we compared our results to models that implicitly or explicitly learn matching, such as CroCo-Flow [93], DUST3R [86], and MAST3R [53] which are fine-tuned from the pretrained CroCo-v2 [93]. CroCo-Flow was evaluated using DPED-Cityscapes-ADE dataset, while DUST3R and MAST3R were tested using pretrained models trained on more extensive datasets, including MegaDepth. DUST3R generates aligned point maps, enabling the use of point-cloud nearest neighbors (NN) to identify matching points, whereas MAST3R employs feature-based NN matching. For all evaluations, we performed dense matching without incorporating uncertainty or cycle-consistency checks during inference.

## B.3. Learning-based Multi-frame Depth Estimation

**Datasets.** To assess the performance of our self-supervised multi-frame depth estimation model, we evaluated it on the KITTI [25] and Cityscapes [14] datasets. For KITTI, we used the standard Eigen split [20] with preprocessing in [101], resulting in 39,810 training, 4,424 validation, and 697 test images. For Cityscapes, we used 69,731 preprocessed training images and 1,525 test images with SGM-derived disparity maps [39]. In both datasets, we evaluated pixels with ground-truth depth less than 80 meters, clipping predicted depths accordingly.

We used standard depth evaluation metrics [20] with median scaling following [101]. Lower error metric values (AbsRel, SqRel, RMSE, RMSElog) indicate better performance, while higher accuracy metric values ( $\delta_1$ ,  $\delta_2$ ,  $\delta_3$ ) are preferred. For noise robustness evaluation, we adopted metrics from RoboDepth [51], including mDEE (a combined AbsRel and  $\delta_1$ ) and mRR, which measures degradation under noisy conditions.

**Baselines.** To validate our performance, we conducted comparisons with models trained in a self-supervised manner using the same splits on KITTI [25] and Cityscapes [14] datasets. For monocular input, we evaluated against Monodepth2 [27], Packnet-SFM [30], MonoViT [100], GUDA [31], and RA-Depth [38]. Among multi-input models, we included those leveraging epipolar-based cost volumes to incorporate multi-view geometric information, such as ManyDepth [91], DynamicDepth [24], DepthFormer [32], MOVEDepth [87], and DualRefine [4], as well as attention-based approaches like Patil et al. [69] and TC-Depth [77]. Notably, most multi-frame depth estimation models struggle to handle dynamic scenes and often rely on knowledge distillation from Monodepth2 [27] to address these challenges.

**Experiments setting on dynamic objects.** To evaluate our model’s performance on dynamic objects, we followed the evaluation protocol of DynamicDepth [24]. Dynamic objects were identified using an off-the-shelf semantic segmentation model, EfficientPS [63]. Unlike instance-level masks or inter-frame correspondences, all dynamic category pixels were projected together in a single step. Specifically, we segmented and measured metrics within the regions corresponding to dynamic categories, including Vehicles, Persons, and Bikes.

**Experiments setting on noisy environments.** Building on the methodology of RoboDepth [51], which evaluated depth prediction performance in noisy, real-world scenarios, we also assessed the performance of multi-frame depth estimation under challenging conditions where epipolar-based cost volumes tend to struggle. We employed their evaluation metrics, including mDEE, a combined metric of AbsRel and  $\delta_1$ , as well as mRR, which measures the degradation in performance relative to noise-free conditions. These metrics are defined as follows:

$$\text{DEE} = \frac{\text{AbsRel} - \delta_1 + 1}{2} \quad (8)$$

$$\text{mDEE} = \frac{1}{N \cdot L} \sum_{i=1}^N \sum_{l=1}^L \text{DEE}_{i,l} \quad (9)$$

$$\text{RR}_i = \frac{\sum_{l=1}^L (1 - \text{DEE}_{i,l})}{L \times (1 - \text{DEE}_{\text{clean}})} \quad (10)$$

$$\text{mRR} = \frac{1}{N} \sum_{i=1}^N \text{RR}_i \quad (11)$$

To demonstrate the noise robustness of our framework, we evaluated it across five levels of severity using augmentations from the ‘Sensor & Movement’ and ‘Data & Processing’ categories. These augmentations included defocus



blur, glass blur, motion blur, zoom blur, elastic transformation, color quantization, gaussian noise, impulse noise, shot noise, ISO noise, pixelate, and JPEG compression.

Methods	Matching cost	Layers	HPatches-240						HPatches-Original					
			AEPE↓						AEPE↓					
			I	II	III	IV	V	Avg.	I	II	III	IV	V	Avg.
CroCo [93] Encoder	Correlation	0	77.68	81.24	80.58	81.38	84.60	81.10	331.33	347.28	348.05	354.11	363.28	348.81
		2	66.89	71.31	70.97	72.21	76.98	71.67	282.28	304.07	298.76	305.63	321.48	302.44
		4	59.57	64.55	64.22	66.41	71.11	65.17	250.96	272.93	271.94	283.00	301.25	276.02
		6	55.89	61.87	60.49	63.23	68.66	62.03	234.22	260.80	253.24	266.84	289.43	260.91
		8	49.46	55.76	54.82	57.03	62.82	55.98	208.88	237.07	229.45	241.63	268.51	237.11
		10	39.07	46.12	46.35	47.79	54.41	46.75	167.14	198.90	194.59	205.21	234.82	200.13
		12	34.85	41.88	42.19	44.29	50.36	42.71	149.11	181.04	176.83	189.94	216.16	182.62
		14	30.67	37.71	38.56	40.83	46.45	38.84	130.57	163.24	164.71	174.61	201.05	166.83
		16	28.73	35.89	37.21	39.56	44.92	37.26	123.70	156.69	158.48	166.67	194.58	160.02
		18	27.06	33.35	35.95	37.36	43.91	35.53	117.07	148.66	152.34	159.23	189.85	153.43
		20	26.06	32.96	35.50	37.16	43.31	35.00	113.62	145.45	150.61	157.60	189.93	151.44
		23	26.31	34.82	36.68	37.69	43.71	35.84	107.74	143.96	146.31	152.34	182.56	146.58
		0-23	39.69	47.28	47.35	48.64	54.63	47.52	168.56	201.52	197.82	206.63	234.83	201.87
CroCo [93] Decoder	Correlation	0	24.03	34.73	36.39	38.67	46.46	36.05	106.09	153.44	154.79	166.91	202.01	156.65
		1	22.82	35.18	35.82	40.43	40.30	36.73	99.02	153.25	154.57	175.31	215.75	159.58
		2	20.84	33.43	33.72	39.69	46.14	34.96	100.20	154.72	153.14	165.12	207.62	160.16
		3	22.72	36.93	36.54	42.39	53.36	38.39	98.65	160.79	157.14	186.83	236.06	167.89
		4	29.85	45.03	43.23	50.54	62.20	46.17	128.00	196.18	189.83	216.90	275.98	201.38
		5	37.31	53.51	51.08	57.05	69.58	53.71	157.78	229.95	219.40	246.55	297.62	230.26
		6	50.82	64.64	63.37	60.65	69.90	61.87	212.22	271.83	270.34	257.03	290.88	260.46
		7	63.51	79.03	76.00	74.26	81.53	74.86	270.40	335.62	324.02	311.68	339.78	316.30
		8	60.89	75.10	75.32	71.83	81.11	72.85	260.35	319.29	318.12	304.12	334.94	307.36
		9	64.55	76.34	76.85	72.39	80.32	74.09	277.35	327.92	323.53	305.80	337.55	314.43
		10	67.41	78.49	78.77	74.37	80.95	76.00	285.27	329.35	329.48	314.91	340.35	319.87
		11	66.89	75.14	72.98	73.67	77.91	73.32	280.18	311.57	305.92	306.77	323.13	305.51
		0-11	32.38	45.39	44.35	46.19	54.84	44.63	137.30	195.51	191.86	198.49	231.11	190.85
CroCo [93] Enc. + Dec.	Correlation	all	35.15	45.60	44.42	45.85	51.80	44.16	149.22	187.30	186.18	194.45	224.14	188.26
		23 / 3,7,11†	26.31	34.82	36.68	37.69	43.71	35.84	112.39	148.92	154.97	160.06	192.07	153.68
ZeroCo (Ours)	Cross-attention	0	9.83	19.68	22.37	27.82	31.11	22.16	41.419	85.40	92.59	116.39	137.12	94.58
		1	8.16	11.84	15.91	19.36	20.37	15.13	34.47	47.99	63.39	76.27	81.17	60.66
		2	7.70	11.28	15.32	17.61	19.04	14.19	31.33	44.89	59.52	69.76	78.65	56.83
		3	6.12	9.06	12.21	14.44	16.52	11.67	25.12	35.70	45.97	55.33	63.65	45.15
		4	6.19	9.81	12.87	14.29	16.42	11.92	25.14	37.76	49.62	53.97	63.51	46.00
		5	5.44	8.70	11.97	13.02	14.95	10.82	22.65	32.94	44.49	48.13	55.32	40.71
		6	6.42	9.14	13.40	13.42	15.87	11.65	25.18	34.58	50.56	52.51	61.38	44.84
		7	5.18	<u>7.13</u>	10.40	<u>11.38</u>	<u>12.77</u>	<u>9.37</u>	21.29	27.47	39.43	44.01	48.13	36.06
		8	<b>5.01</b>	<b>7.04</b>	<u>10.24</u>	11.54	<b>12.40</b>	<b>9.24</b>	<b>20.61</b>	<b>26.85</b>	<b>38.84</b>	<u>43.72</u>	<u>46.96</u>	<u>35.40</u>
		9	6.29	8.46	11.69	12.62	14.60	10.73	25.74	34.01	47.18	52.23	58.66	43.56
		10	5.47	7.43	10.40	11.96	13.28	9.71	22.23	29.33	40.90	46.36	52.19	38.20
		11	31.17	39.51	40.35	46.63	50.80	41.70	124.98	162.55	165.44	193.11	216.15	172.44
		0-11	<u>5.07</u>	7.16	<b>10.19</b>	<b>11.37</b>	13.26	9.41	<u>20.75</u>	<u>27.32</u>	<u>39.10</u>	<b>43.43</b>	<b>46.35</b>	<b>35.39</b>

Table 8. **Zero-shot matching results with different layers on HPatches [3].** The zero-shot performance of CroCo is evaluated across different layers by analyzing their cost volumes on HPatches-240 and HPatches-Original datasets, representing resolutions of  $240 \times 240$  and the original resolution, respectively. The best results are indicated in **bold**, while the second-best results are marked with underline. †: The encoder and decoder features are utilized in the CroCo-flow [93], DUS3R [86], and MAST3R [53].

Methods	Matching cost	Input Resolution	HPatches-Original					
			AEPE↓					
			I	II	III	IV	V	Avg.
CroCo [93] Encoder	Correlation	224 × 224	107.74	143.96	146.31	152.34	182.56	146.58
		448 × 448	133.71	174.30	173.62	189.65	213.98	177.05
		672 × 672	163.36	197.95	200.20	213.14	240.12	202.96
CroCo [93] Decoder	Correlation	224 × 224	106.09	153.44	154.79	166.91	202.01	156.65
		448 × 448	114.60	159.44	161.22	175.98	202.82	162.81
		672 × 672	147.90	186.95	189.64	203.14	224.58	190.44
<b>ZeroCo (Ours)</b>	Cross-attention	224 × 224	<b>20.75</b>	<b>27.32</b>	<b>39.10</b>	<b>43.43</b>	<b>46.35</b>	<b>35.39</b>
		448 × 448	<u>92.81</u>	173.12	184.96	226.13	256.65	186.73
		672 × 672	113.09	180.32	183.98	215.25	237.99	186.12

Table 9. **Zero-shot matching results at varying input resolutions on HPatches. [3].** The zero-shot performance of cross-view completion with different input resolutions is evaluated using cost volumes on HPatches-Original. The best results are highlighted in **bold**, while the second best results are marked with an underline. CroCo [93], trained exclusively at 224 × 224, achieves the best performance at this resolution. However, training on higher resolutions could potentially enhance its performance for larger input sizes.

Pretrained weights	Enc. / Dec.	Matching cost	HPatches-240						HPatches-Original					
			AEPE↓						AEPE↓					
			I	II	III	IV	V	Avg.	I	II	III	IV	V	Avg.
CroCo-v1 [92]	ViT-B / ViT-S	Encoder-Correlation	20.14	28.70	30.82	35.04	38.09	30.56	83.66	121.80	126.74	147.53	164.43	128.82
		Decoder-Correlation	25.21	40.35	42.06	49.68	54.71	42.40	106.72	174.56	179.56	213.87	236.85	182.31
		Cross-attention	22.77	38.04	40.05	46.04	50.13	39.41	96.27	165.30	169.30	196.17	216.41	168.69
CroCo-v2 [93]	ViT-B / ViT-S	Encoder-Correlation	16.09	22.55	25.51	27.45	31.86	24.69	69.46	96.43	104.09	115.62	134.44	104.01
		Decoder-Correlation	18.04	26.84	28.05	30.23	38.19	28.27	78.40	117.18	118.27	129.57	162.81	121.25
		Cross-attention	7.30	11.15	14.94	16.47	19.16	13.81	30.59	47.40	58.83	66.42	79.59	56.57
CroCo-v2 [93]	ViT-B / ViT-B	Encoder-Correlation	16.16	22.35	24.59	27.28	30.25	24.13	68.78	94.61	103.07	113.95	125.60	101.20
		Decoder-Correlation	18.89	25.76	28.90	30.76	37.36	28.33	82.98	113.79	122.72	129.42	160.06	121.80
		Cross-attention	<b>5.89</b>	<u>9.30</u>	<u>13.51</u>	<u>14.21</u>	<u>17.27</u>	<u>12.03</u>	<b>24.85</b>	<u>36.24</u>	<u>51.65</u>	<u>58.46</u>	<u>65.15</u>	<u>47.27</u>
CroCo-v2 [93]	ViT-L / ViT-B	Encoder-Correlation	24.62	32.92	34.75	35.80	41.69	33.96	105.02	141.72	146.45	151.03	182.55	145.36
		Decoder-Correlation	22.77	33.23	34.79	37.31	44.78	34.58	98.57	145.00	145.74	161.78	193.35	148.89
		Cross-attention	<u>6.09</u>	<b>8.47</b>	<b>11.87</b>	<b>12.89</b>	<b>13.83</b>	<b>10.63</b>	<u>25.10</u>	<b>33.03</b>	<b>45.05</b>	<b>50.26</b>	<b>54.24</b>	<b>41.54</b>

Table 10. **Zero-shot matching results with different pretrained weights on HPatches [3].** The zero-shot performance of CroCo with different pretrained weights is evaluated using cost volumes on HPatches-240 and HPatches-Original datasets, representing 240 × 240 and original resolutions, respectively. The best results are highlighted in **bold**, while the second best results are marked with an underline.

Computing Matching cost	Normalize	Reciprocity	Dense zoom-in	HPatches-240					
				AEPE↓					
				I	II	III	IV	V	Avg.
Query-Query	<b>X</b>	<b>X</b>	<b>X</b>	11.72	21.04	22.11	29.10	37.04	24.20
Key-Key	<b>X</b>	<b>X</b>	<b>X</b>	28.28	42.54	43.99	50.98	55.56	44.27
Value-Value	<b>X</b>	<b>X</b>	<b>X</b>	35.00	45.56	47.47	49.64	53.84	46.30
Query-Key	<b>X</b>	<b>X</b>	<b>X</b>	6.30	8.28	12.16	12.63	14.88	10.85
Query-Key	L2-Norm	<b>X</b>	<b>X</b>	5.90	8.07	11.65	12.69	14.98	10.66
Query-Key	Softmax	<b>X</b>	<b>X</b>	6.00	8.34	11.26	12.42	14.61	10.53
Query-Key	<b>X</b>	✓	<b>X</b>	5.93	8.13	11.56	12.46	13.12	10.24
Query-Key	L2-Norm	✓	<b>X</b>	5.93	8.19	11.50	12.51	13.54	10.34
Query-Key	Softmax	✓	<b>X</b>	6.06	8.35	11.44	12.85	13.52	10.45
<b>Query-Key (Ours)</b>	<b>X</b>	✓	✓	<b>5.07</b>	<b>7.16</b>	<b>10.19</b>	<b>11.37</b>	<b>13.26</b>	<b>9.41</b>

Table 11. **Ablation studies on HPatches-240 [3].** We conduct ablation studies on the HPatches-240.

## C. Analysis

### C.1. Layer Analysis of CVC Models

As with DIFT [80], we evaluated the performance of CroCo’s encoder, decoder, and attention map across individual layers on the HPatches [3] dataset to identify the layer yielding the best matching cost. Additionally, we assessed performance by averaging the matching costs obtained from all layers. We also tested the performance of features used in head training for models like CroCo-flow, DUST3R, and MAST3R, specifically the 23rd encoder feature and the 3rd, 7th, and 11th decoder features, by averaging their correlations.

The results, presented in Tab. 8, reveal that the attention map consistently outperformed the encoder and decoder in matching cost across all layers. The table further highlights that encoder layers tend to produce better matching costs in later layers, while decoder layers perform better in earlier layers. In contrast, the cross-attention mechanism achieved its best performance at intermediate layers (specifically the 7th and 8th). Notably, for datasets like Hpatches-Original, averaging the matching costs across layers yielded the best results. Furthermore, the features used for head training in existing models like CroCo-flow [93], DUST3R [86], and MAST3R [53] demonstrated descriptor qualities compared to other layer features. However, they exhibited inferior performance compared to the matching cost stored in the cross-attention maps. These results suggest that cross-attention maps encode superior geometric information compared to encoder or decoder features. Consistent with the findings from our learning-based model training, incorporating cross-attention into models like DUST3R and MAST3R would enable them to leverage the geometric information inherent in CroCo more effectively, surpassing the performance of models relying solely on feature-based approaches.

### C.2. Resolution Analysis of CVC Models

In our experiments, we also conducted zero-shot evaluations at the input resolution of  $224 \times 224$ , which is the resolution used during cross-view completion training. Additionally, we plan to extend this evaluation to higher resolutions, specifically  $2\times$  and  $3\times$  scales, at  $448 \times 448$  and  $672 \times 672$ , to measure the performance at these resolutions.

Tab. 9 indicate that increasing the resolution degrades the performance of not only the encoder-decoder correlation map but also the cross-attention map. This phenomenon is likely due to CroCo being trained exclusively at a resolution of  $224 \times 224$ , where the cross-attention map achieves its best performance. Consequently, we hypothesize that if CroCo were trained at higher resolutions, it would produce effective cross-attention maps even at those resolutions, leading to finer zero-shot matching results.

Methods	Time(ms)	Memories(MB)
DINOv2 [65]	<b>56.74</b>	<b>336.82</b>
DIFT <sub>SD</sub> [80]	345.91	4,888.75
DIFT <sub>ADM</sub> [80]	174.31	2,118.89
SD-DINO [99]	5,016.78	9,095.14
<b>ZeroCo w/o dense zoom-in</b>	<u>100.52</u>	<u>1,607.81</u>
<b>ZeroCo (ZeroCo-finetuned)</b>	2,648.01	1,608.00

Table 12. **Evaluation results of memory and time complexity on HPatches-240 [3].** We conduct all experiments on a single RTX 3090 GPU. The best results are highlighted in **bold**, and the second-best results are marked with underline.

Methods	DINO	DIFT <sub>ADM</sub>	SD-DINO	Encoder-corr.	Decoder-corr.	Ours
AEPE↓	85.04	72.65	<u>62.16</u>	102.11	93.75	<b>47.63</b>

Table 13. **Zero-shot matching on MegaDepth [56].**

### C.3. Different CVC Models

CroCo provides four types of pretrained weights, including CroCo-v1 [92] and CroCo-v2 [93] with different encoder-decoder sizes. We analyzed the zero-shot matching performance of these weights.

As shown in Tab. 10, the cross-attention maps of CroCo-v2 models consistently capture better geometric information than the encoder-decoder correlation maps. Additionally, larger model sizes demonstrated improved matching performance. Interestingly, while CroCo-v1 was trained using the same method, its cross-attention maps showed performance comparable to correlation maps. This difference likely stems from CroCo-v2’s use of a retrieval-based approach with more diverse and extensive training data, enabling better learning of cross-attention maps.

### C.4. Ablation Study

We also conducted additional ablation studies, as detailed in Tab. 11, examining the impact of various factors, including the method for constructing matching costs based on cross-attention, normalization, reciprocity, and dense zoom-in.

Initially, we computed matching costs by multiplying the query-key matrices before applying softmax. However, with reciprocity-based inference, we obtain query, key, and value features for both the target and source images, allowing us to calculate correlations between them to construct matching costs. Interestingly, as shown in the results, the original query-key method yielded the best performance. For normalization, excluding reciprocity, both L2-normalization and softmax improved performance. However, applying normalization after combining costs through reciprocity resulted in reduced performance. Additionally, incorporating dense zoom-in provided slight performance improvements. Our final model omits normalization while utilizing reciprocity and dense zoom-in, achieving the best overall results.



Methods	Training dataset	Matching cost	HPatches-240						HPatches-Original					
			AEPE↓						AEPE↓					
			I	II	III	IV	V	Avg.	I	II	III	IV	V	Avg.
DUST3R [86] Encoder	MIX <sub>8</sub>	Correlation	18.43	27.62	28.98	32.89	40.99	29.78	79.54	123.12	123.23	137.96	170.49	126.87
DUST3R [86] Decoder	MIX <sub>8</sub>	Correlation	17.01	27.42	28.03	32.49	40.14	29.02	72.22	118.66	121.22	137.65	169.66	123.88
DUST3R [86]	MIX <sub>8</sub>	Cross-attention	6.49	9.00	10.64	14.04	16.78	11.39	27.75	38.28	43.12	52.97	63.84	45.19
MASt3R [53] Encoder	MIX <sub>14</sub>	Correlation	24.30	31.97	33.91	35.52	42.47	33.63	105.25	133.68	141.91	149.18	178.33	131.67
MASt3R [53] Decoder	MIX <sub>14</sub>	Correlation	22.23	29.34	31.28	33.40	40.61	31.37	94.06	125.98	131.58	144.19	170.76	133.31
MASt3R [53]	MIX <sub>14</sub>	Cross-attention	5.65	6.41	8.14	8.16	8.52	7.38	24.00	27.08	32.16	33.72	38.91	31.17
CroCo-flow [93]	MIX <sub>4</sub>	-	5.21	19.47	21.70	21.44	28.82	19.33	27.33	94.37	93.51	107.60	147.43	94.05
CroCo-flow [93]	MIX <sub>4</sub> , $\mathcal{D}$	-	<b>0.43</b>	<u>2.64</u>	7.73	8.79	10.63	6.05	<b>1.32</b>	<b>7.80</b>	24.88	30.53	38.75	20.66
DUST3R [86]	MIX <sub>8</sub>	Point-Corr.	4.18	6.32	10.15	11.35	15.32	9.46	10.11	13.19	<b>15.28</b>	<u>18.44</u>	30.78	17.56
MASt3R [53]	MIX <sub>14</sub>	Feature-Corr.	1.31	<b>1.36</b>	<b>4.01</b>	<b>2.06</b>	<b>4.34</b>	<b>2.62</b>	13.07	11.58	20.74	<b>15.18</b>	24.49	17.01
<b>ZeroCo-finetuned (Ours)</b>	$\mathcal{D}^*$ , $\mathcal{M}$	Cross-attention	1.86	4.03	6.03	6.63	7.79	5.26	5.15	11.55	18.59	20.86	<u>22.72</u>	<u>15.77</u>
<b>ZeroCo-flow (Ours)</b>	$\mathcal{D}^*$ , $\mathcal{M}$	Refined cost	<u>0.49</u>	2.73	<u>5.46</u>	<u>6.40</u>	<u>6.25</u>	<u>4.27</u>	<u>1.51</u>	<u>9.09</u>	<u>15.62</u>	21.07	<b>20.73</b>	<b>13.61</b>

Table 14. **Dense geometric matching results on different resolutions of HPatches [3].** A higher scene label such as V corresponds to a more challenging setting with extreme geometric deformation. The best result is highlighted in **bold**, and the second-best result is marked with underline. The following notations are used in the table:  $\mathcal{D}$ : DPED-CityScapes-ADE,  $\mathcal{D}^*$ : COCO-augmented DPED-CityScapes-ADE,  $\mathcal{M}$ : Megadepth, MIX<sub>4</sub>: 4 mixed dataset used in CroCo-flow [93], MIX<sub>8</sub>: 8 mixed dataset used in DUST3R [86], MIX<sub>14</sub>: 14 mixed dataset used in MASt3R [53], Corr.: Correlation.

## C.5. Time and Memory Complexity

We also compared the time complexity and GPU memory usage of our model against other zero-shot matching models. For this analysis, we averaged the inference time and GPU memory consumption per image pair on the HPatches-240 dataset, with all models operating at a resolution of  $224 \times 224$ . All experiments were conducted on a single NVIDIA RTX 3090 GPU.

As shown in Tab. 12, diffusion-based models [80, 99] exhibit high memory and time complexity due to their large model sizes. Notably, SD-DINO [99] requires even more resources as it combines stable diffusion and DINOv2 features. In contrast, models like DINOv2 [65], which use only a ViT encoder, benefit from faster speeds and lower memory usage but suffer from inferior performance. Our model uses slightly more resources than DINOv2, yet it remains significantly lighter and faster than diffusion-based models while delivering superior performance.

## D. Additional Quantitative Results

### D.1. Zero-shot Geometry Matching

**Megadepth results.** To evaluate the generalizability of ZeroCo, we conducted experiments on MegaDepth [56], a larger and more realistic real-world dataset, with the results presented in Table 13. The results show that ZeroCo outperforms not only existing zero-shot models such as DINOv2 [65], DIFT [80], and SD-DINO [99], but also the zero-shot performance obtained by using correlation maps constructed from CroCo’s encoder and decoder features as matching costs.

### D.2. Learning-based Geometry Matching

**HPatches results.** We compared models trained on extended datasets, including CroCo-flow [93], DUST3R [86],

Methods	AEPE↓					
	I	II	III	IV	V	Avg.
Encoder-corr.	5.06	13.95	26.00	32.87	36.77	22.93
Decoder-corr.	5.52	13.91	27.11	33.36	36.05	23.19
CroCo-flow (decoder)	5.56	15.42	26.43	33.58	38.37	23.88
CroCo-flow (query)	4.60	12.78	26.37	32.19	37.33	22.66
CroCo-flow (key)	65.40	96.04	113.69	125.47	140.87	108.29
ZeroCo-flow w/o reciprocity	2.73	10.25	20.99	27.31	27.01	17.66
ZeroCo-flow w/o $\mathcal{T}$	2.45	10.15	19.89	25.40	24.86	16.55
ZeroCo-flow w/o $\mathcal{U}$	3.29	10.66	21.48	24.82	27.94	17.64
<b>ZeroCo-flow</b>	<b>1.51</b>	<b>9.09</b>	<b>15.62</b>	<b>21.07</b>	<b>20.73</b>	<b>13.61</b>

Table 15. **Ablation studies of learning-based matching on HPatches-Original dataset.**

and MASt3R [53], all leveraging pretrained CroCo. For CroCo-flow, comparisons were made between models trained on flow datasets and DPED-Cityscapes-ADE datasets, while DUST3R and MASt3R were evaluated on their respective mixed datasets. Additionally, we analyzed the geometric understanding encoded in the encoder, decoder, and cross-attention maps of DUST3R and MASt3R.

Tab. 14 reveals that CroCo-flow, trained on the flow dataset, underperformed due to its focus on small displacements and the limited geometric understanding of its encoder and decoder, even with additional dense-matching datasets. In contrast, DUST3R and MASt3R showed improved geometric comprehension in encoder, decoder, and cross-attention features with larger datasets. Notably, MASt3R excelled in cross-attention maps, highlighting the structural advantage of cross-attention for capturing geometric information.

**Ablation studies.** To highlight the importance of each component in our learning-based method and to underscore the significance of employing cross-attention, we provide ablation studies in Table 15. First, we observed a per-

Methods	Additional network	Test frames	AbsRel↓	SqRel↓	RMSE↓	RMSElog↓	$\delta_1\uparrow$	$\delta_2\uparrow$	$\delta_3\uparrow$
Struct2Depth 2 [7]	-	1	0.145	1.737	7.280	0.205	0.813	0.942	0.976
Monodepth2 [27]	-	1	0.129	1.569	6.876	0.187	0.849	0.957	0.983
Videos in the Wild [28]	-	1	0.127	1.330	6.960	0.195	0.830	0.947	0.981
Li et al. [54]	-	1	0.119	1.290	6.980	0.190	0.846	0.952	0.982
Struct2Depth 2 [7]	$\mathcal{M}, \mathcal{F}$	3 (-1, 0, 1)	0.151	2.492	7.024	0.202	0.826	0.937	0.972
ManyDepth [91]	$\mathcal{M}$	2 (-1, 0)	0.114	1.193	6.223	0.170	<u>0.875</u>	0.967	<u>0.989</u>
DynamicDepth <sup>†</sup> [24]	$\mathcal{M}, \mathcal{S}$	2 (-1, 0)	<b>0.104</b>	<b>1.009</b>	<b>5.991</b>	<b>0.150</b>	<b>0.889</b>	0.972	<b>0.991</b>
<b>ZeroCo-depth (Ours)</b>	-	2 (-1, 0)	<u>0.105</u>	<u>1.050</u>	<u>6.117</u>	<u>0.162</u>	<b>0.889</b>	<b>0.973</b>	<b>0.991</b>

Table 16. **Depth estimation results on Cityscapes [14]**. We compare our model with previous single- and multi-frame depth estimation networks. For our baseline, we adopted the CroCo-stereo architecture [93] and trained it using a self-supervised depth learning manner. The best is in **bold**, and the second-best is underlined. <sup>†</sup>: our reproduced results from the official repository,  $\mathcal{M}$ : monocular depth [27],  $\mathcal{S}$ : segmentation, and  $\mathcal{F}$ : flow network.

formance drop when replacing the cross-attention map in ZeroCo-flow with either the encoder correlation map or the decoder correlation map. This suggests that the cross-attention map, which forms a cost volume with more accurate correspondences during matching, is more beneficial for the matching process.

In addition, we conducted ablation studies on CroCo-flow, which originally utilizes CroCo for matching, to investigate whether individual components of the cross-attention map—such as the query and key—contribute to performance. The results show that using decoder features, query, or key individually yields inferior performance compared to directly using the full cross-attention map.

Furthermore, we performed ablations on each component of ZeroCo-Flow, including reciprocity, the aggregation module, and the upsampling module. The absence of each component led to a clear degradation in performance, confirming the necessity of each for optimal results.

### D.3. Learning-based Depth Estimation

**Cityscape results.** Alongside KITTI [25] results, we present Cityscapes [14] results, showing our model’s competitive performance against state-of-the-art models in Tab. 16. Unlike methods like Monodepth2 [27] that rely on additional networks or labels, our approach harnesses the geometric quality of cross-attention alone. This demonstrates that a well-constructed cost volume can achieve strong performance in multi-frame depth estimation.

**Ablation studies.** In Table 17, we present ablation studies on multi-frame depth estimation using the KITTI dataset [25]. First, we evaluated the performance change when replacing the cross-attention map with the encoder and decoder correlation maps from CroCo. The results demonstrate that, similar to matching, the rich correspondence information embedded in the cross-attention map is also beneficial for depth estimation.

We further investigated the impact of using individual

Methods	AbsRel↓	SqRel↓	RMSE↓	$\delta_1\uparrow$
Encoder-corr.	0.098	0.722	4.326	0.908
Decoder-corr.	0.098	0.723	4.342	0.907
CroCo-flow (decoder)	0.102	0.767	4.351	0.903
CroCo-flow (query)	0.096	0.724	4.351	0.910
CroCo-flow (key)	0.097	0.748	4.383	0.912
ZeroCo-depth w/o $\mathcal{F}_{feat}$	0.150	1.235	5.515	0.808
ZeroCo-depth w/o $\mathcal{F}_{attn}$	0.095	0.693	4.251	0.908
ZeroCo-depth	<b>0.090</b>	<b>0.637</b>	<b>4.128</b>	<b>0.915</b>

Table 17. **Ablation studies of depth estimation on KITTI [25]**.

components of CroCo-flow. Specifically, training the model using only the query and key features, excluding the decoder features. The results again indicate that directly utilizing the cross-attention map yields better performance than using its individual components.

Additionally, we conducted ablations on each component involved in the depth estimation pipeline. The findings highlight the importance of both the feature representations and the attention map, confirming that both are essential for accurate depth estimation.

## E. Additional Qualitative Results

The additional qualitative results on zero-shot matching, as well as learning-based geometric matching and depth estimation, highlight the effectiveness of the Cross-View Completion (CVC) pretext task [93], which enables the cross-attention map to learn geometric relationships that can either be used directly or fine-tuned for optimal performance in downstream tasks.

### E.1. Zero-shot Matching

Fig. 7, 8, and 9 showcase correlations from encoder and decoder features and cross-attention maps for target and source images. Blue points mark the query point in the target image, with the corresponding red point indicating the highest attention value based on matching costs.

Fig. 10 and 11 present additional visualizations of warped images, leveraging the cross-attention maps from the successful prior works DUS3R [86] and MAST3R [53]. As both models were also pretrained on CVC, their attention maps effectively demonstrate their effectiveness for zero-shot matching.

### E.2. Learning-based Geometric Matching

Fig. 12 presents visualizations of warped images using the dense output flow from learning-based geometric matching methods. In comparison to recent dense matching models that were not pretrained with CVC, (e) ZeroCo-flow exhibits superior performance. This is due to the geometric knowledge acquired through CVC, which enables the model to achieve improved matching and warping results, as seen in the first and third rows.

### E.3. Learning-based Multi-frame Depth Estimation

Fig. 13 and 14 display multi-frame depth estimation results from our model and state-of-the-art models on the KITTI [25] and Cityscapes [14] datasets. In contrast to previous multi-frame depth estimation methods that use epipolar-based cost volumes to explicitly model matching costs, our cross-attention map acts as a full cost volume, which more effectively captures geometric relationships and matching costs, allowing for better detection of dynamic objects such as cars and pedestrians.



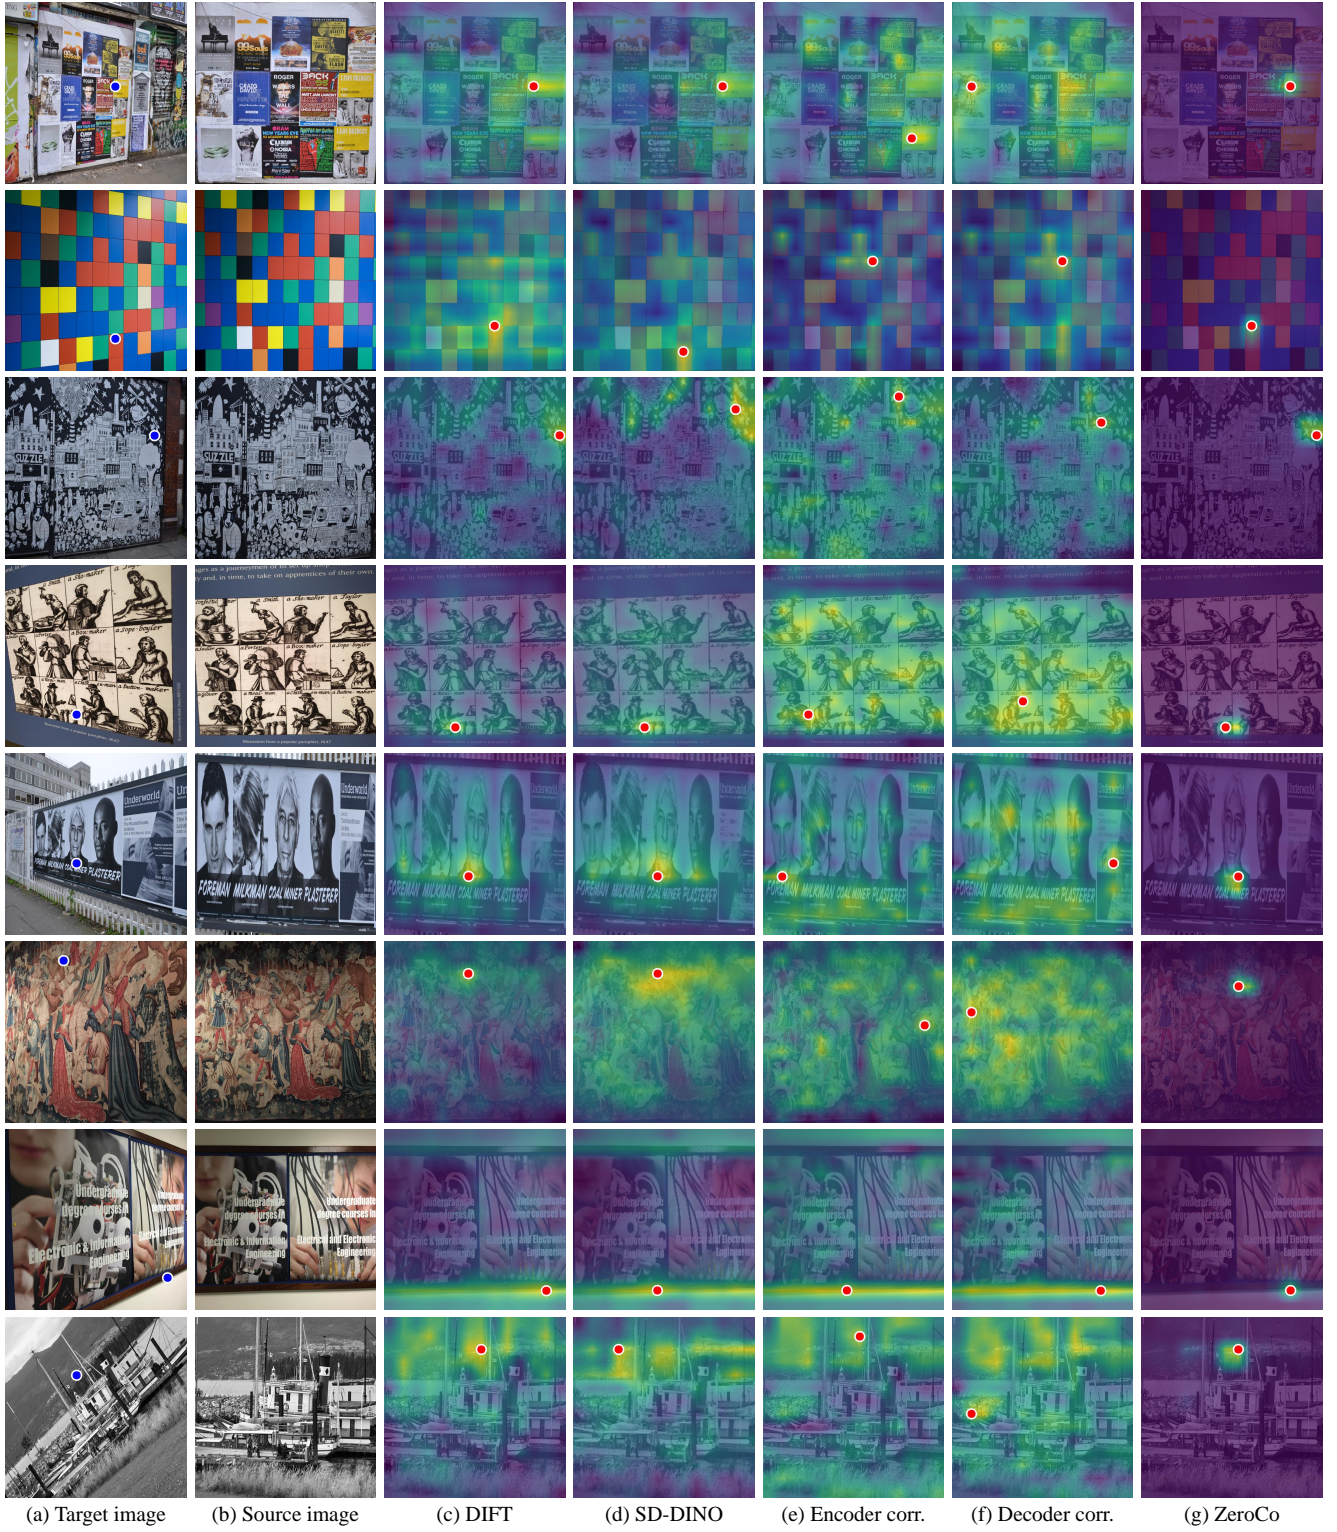


Figure 7. Visualization of matching costs in previous zero-shot matching methods [80, 99], encoder and decoder features within cross-view completion models, and our ZeroCo.



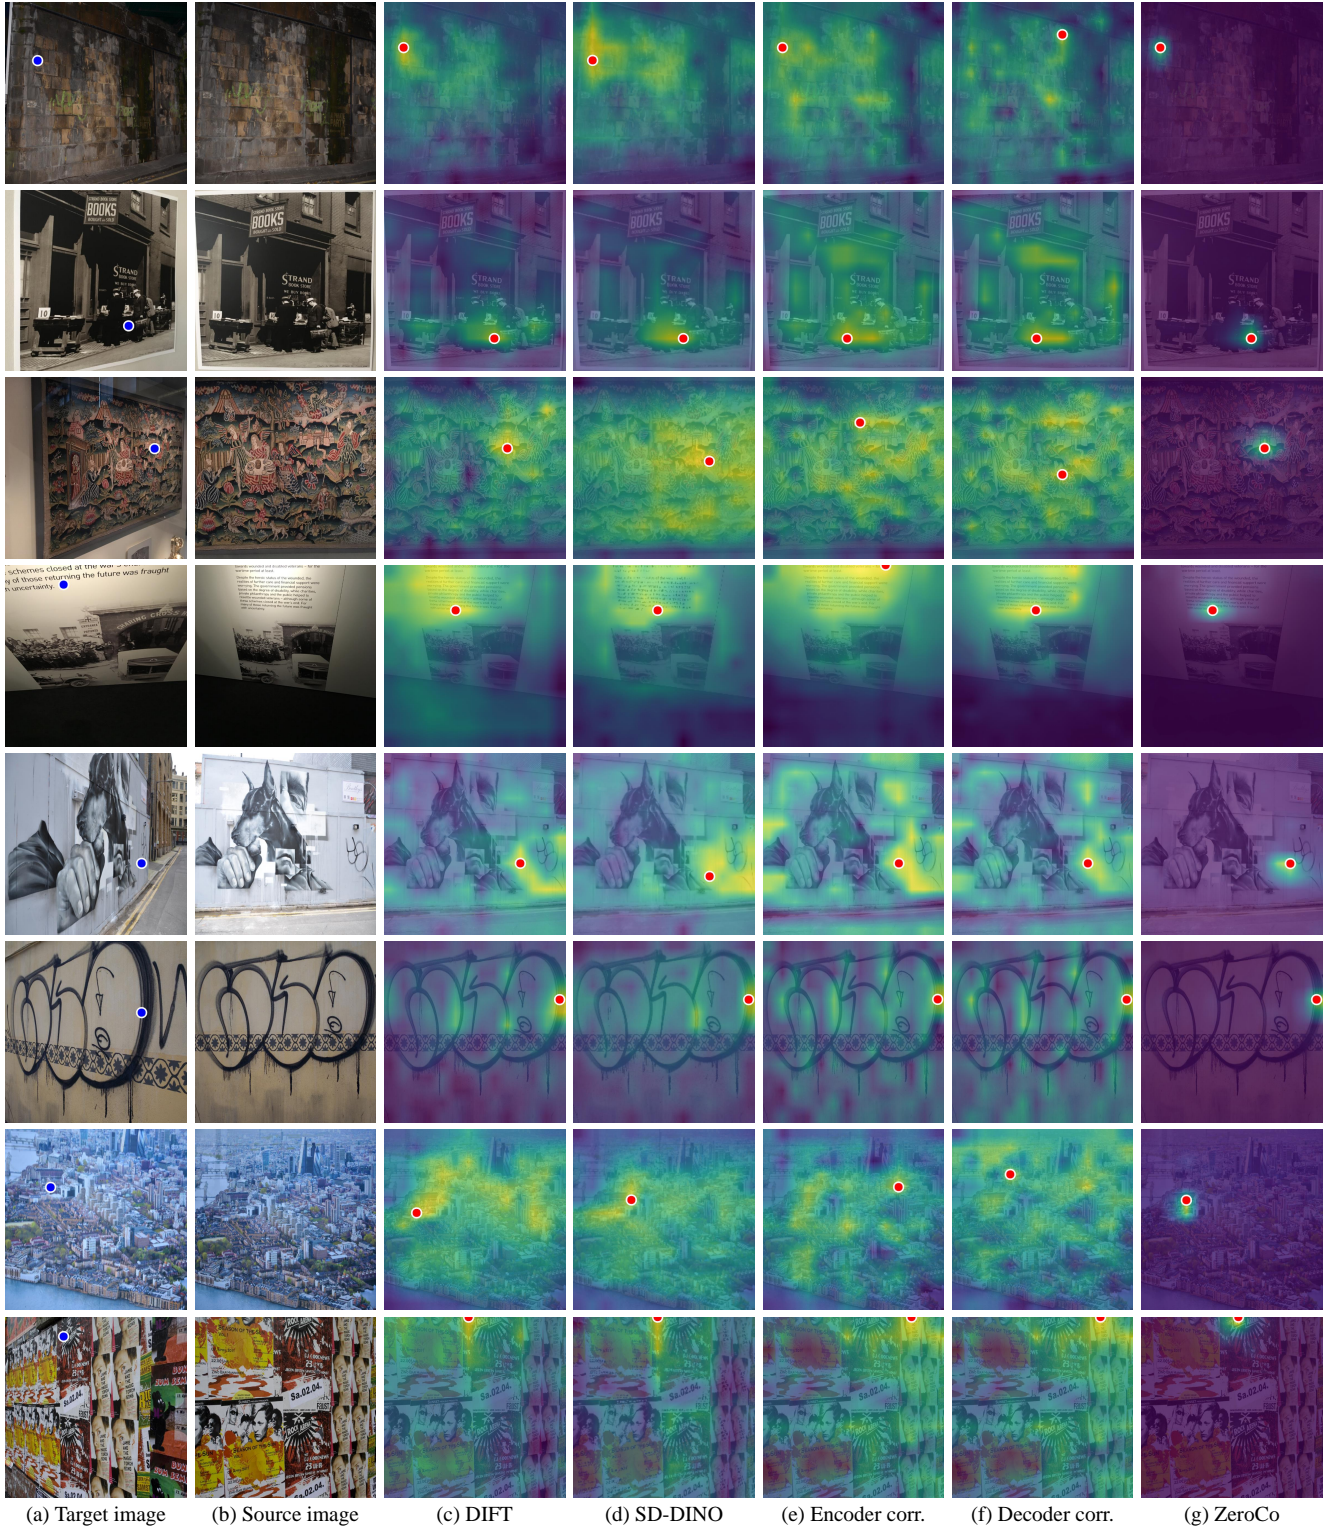


Figure 8. Visualization of matching costs in previous zero-shot matching methods [80, 99], encoder and decoder features within cross-view completion models, and our ZeroCo.



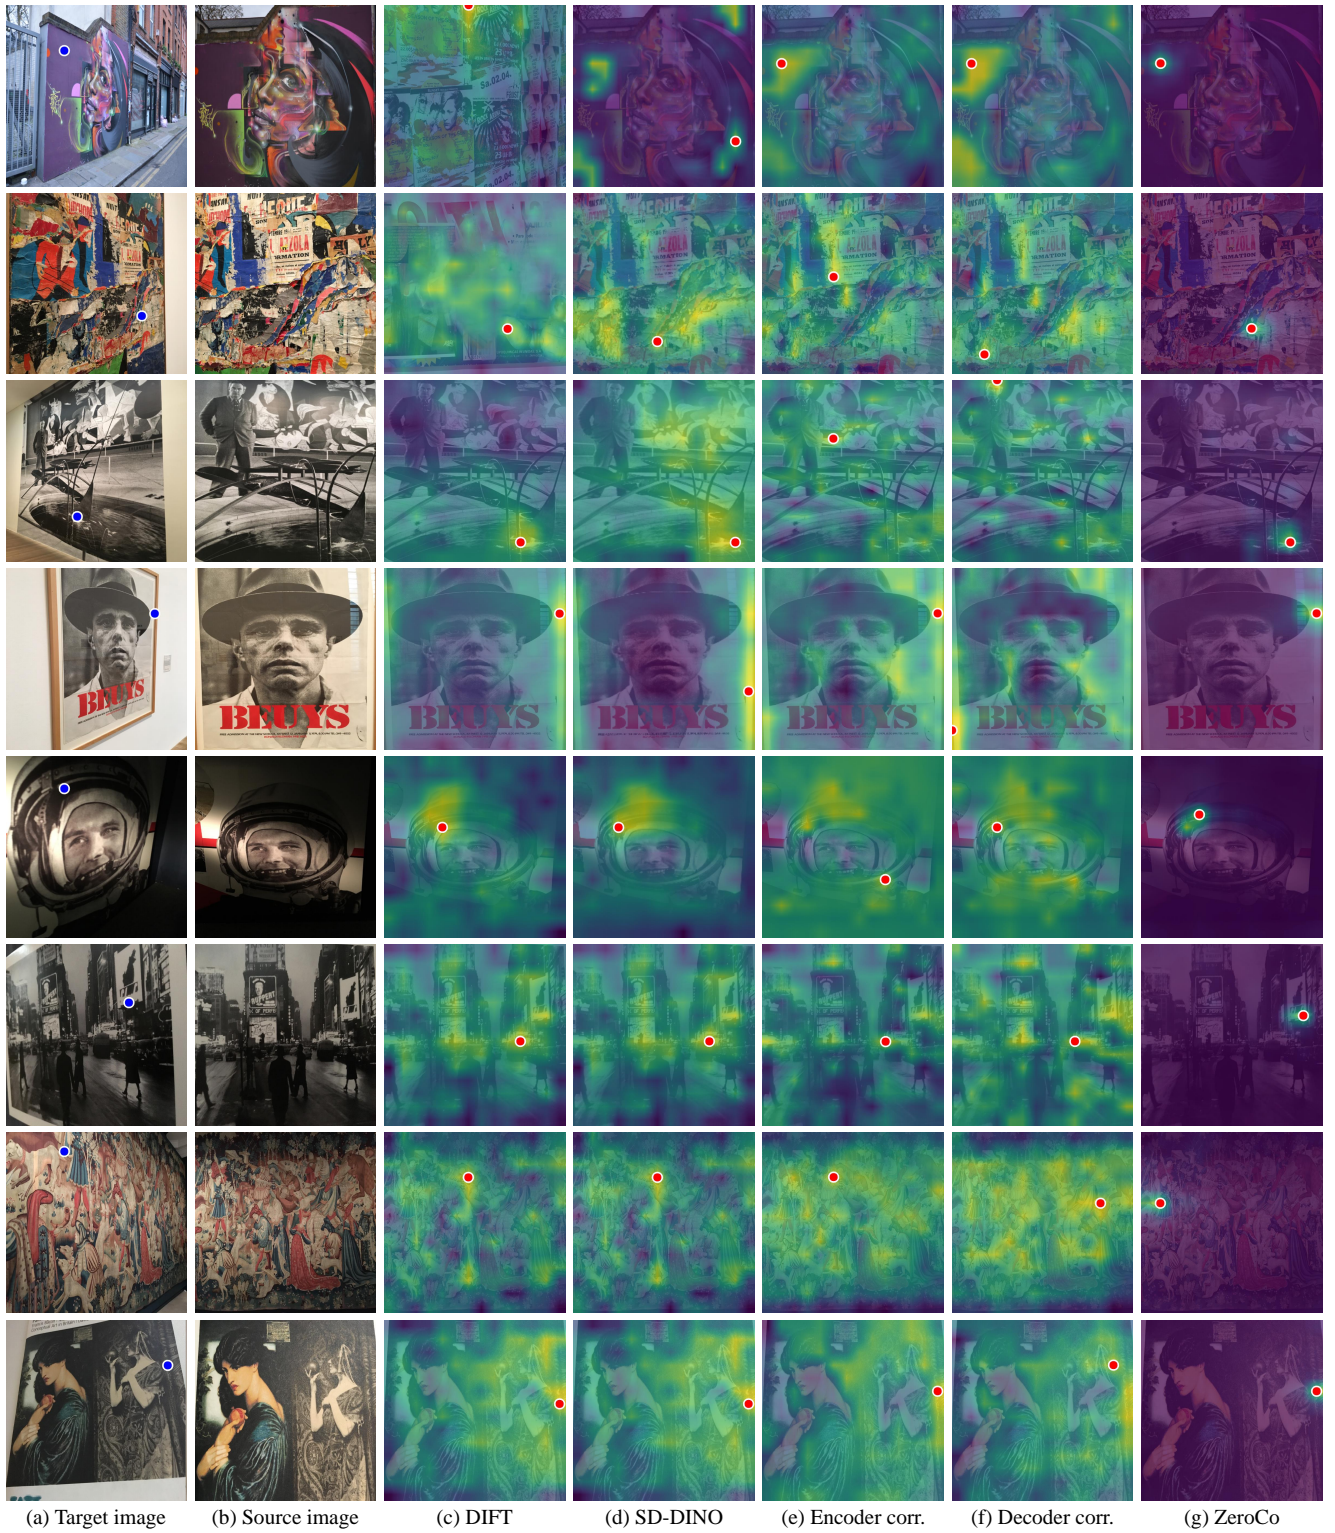



Figure 9. Visualization of matching costs in previous zero-shot matching methods [80, 99], encoder and decoder features within cross-view completion models, and our ZeroCo.





Figure 10. **Visualization of warped images using cross-attention maps.** Based on our findings, we used the **cross-attention** maps from CroCo-v2 [93], DUS3R [86], and MAST3R [53] to warp the source image to the respective target image, which shows the effectiveness of the cross-attention maps in various cross-view completion-based models for dense correspondence.





(a) Target image

(b) Source image

(c) CroCo-v2

(d) DUS3R

(e) MAST3R

(f) Ground-truth

Figure 11. **Visualization of warped images using cross-attention maps.** Based on our findings, we used the **cross-attention** maps from CroCo-v2 [93], DUS3R [86], and MAST3R [53] to warp the source image to the respective target image, which shows the effectiveness of the cross-attention maps in various cross-view completion-based models for dense correspondence.





Figure 12. **Visualization of warped images using estimated dense correspondence.** We used the output flow from GLU-Net-GOcor [81], PDCNet+ [84], DiffMatch [64], and our ZeroCo-flow to warp the source image to the respective target image.



Figure 13. **Qualitative results for multi-frame depth estimation on the KITTI [25] dataset.** We compare our ZeroCo-depth with multi-view depth estimation models that leverage epipolar-based cost volumes [4, 91] and demonstrate improved depth prediction performance for dynamic objects through a full cost volume represented by a cross-attention map.



Figure 14. **Qualitative results for multi-frame depth estimation on the Cityscapes [14] dataset.** We compare our ZeroCo-depth with multi-view depth estimation models that leverage epipolar-based cost volumes [24, 91] and demonstrate improved depth prediction performance for dynamic objects through a full cost volume represented by a cross-attention map.

## **F. Limitation**

Our method may face challenges with semantic object correspondence tasks [9, 10, 49], which involve additional complexities such as intra-class variations and background clutter in image pairs. However, this limitation could potentially be addressed by training cross-view completion models on semantically similar object pairs. Additionally, our method is currently constrained to two-view inputs, and extending it to handle a large number of view inputs would require further implementation efforts and considerations. Finally, the method may encounter difficulties when applied to extremely high-resolution images.

## References

- [1] Roman Bachmann, David Mizrahi, Andrei Atanov, and Amir Zamir. Multima: Multi-modal multi-task masked autoencoders. In *European Conference on Computer Vision*, pages 348–367. Springer, 2022. 2
- [2] Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. Multi-view depth estimation by fusing single-view depth probability with multi-view geometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2842–2851, 2022. 3
- [3] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, and Krystian Mikolajczyk. Hpatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5173–5182, 2017. 6, 8, 16, 17, 19, 20, 21, 22
- [4] Antyanta Bangunharcana, Ahmed Magd, and Kyung-Soo Kim. Dualrefine: Self-supervised depth and pose estimation through iterative epipolar sampling and refinement toward equilibrium. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 726–738, 2023. 3, 7, 8, 17, 31
- [5] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv*, 2021. 3
- [6] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33: 9912–9924, 2020. 2
- [7] Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Unsupervised monocular depth and ego-motion learning with structure and semantics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 23
- [8] Junda Cheng, Wei Yin, Kaixuan Wang, Xiaozhi Chen, Shijie Wang, and Xin Yang. Adaptive fusion of single-view and multi-view depth for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10138–10147, 2024. 3
- [9] Seokju Cho, Sunghwan Hong, Sangryul Jeon, Yunsung Lee, Kwanghoon Sohn, and Seungryong Kim. Cats: Cost aggregation transformers for visual correspondence. *Advances in Neural Information Processing Systems*, 34: 9011–9023, 2021. 5, 33
- [10] Seokju Cho, Sunghwan Hong, and Seungryong Kim. Cats++: Boosting cost aggregation with convolutions and transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(6):7174–7194, 2022. 33
- [11] Seokju Cho, Heeseong Shin, Sunghwan Hong, Anurag Arnab, Paul Hongsuck Seo, and Seungryong Kim. Catseg: Cost aggregation for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4113–4123, 2024. 5
- [12] Hyesong Choi, Hunsang Lee, Seyoung Joung, Hyejin Park, Jiyeong Kim, and Dongbo Min. Emerging property of masked token for effective pre-training. In *European Conference on Computer Vision*, pages 272–289. Springer, 2025. 2
- [13] Hyesong Choi, Hyejin Park, Kwang Moo Yi, Sungmin Cha, and Dongbo Min. Saliency-based adaptive masking: revisiting token dynamics for enhanced pre-training. In *European Conference on Computer Vision*, pages 343–359. Springer, 2025. 2
- [14] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 8, 16, 17, 23, 24, 32
- [15] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*, 2023. 14
- [16] Jacob Devlin. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [17] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3, 5, 14, 15
- [18] Johan Edstedt, Ioannis Athanasiadis, Mårten Wadenbäck, and Michael Felsberg. Dkm: Dense kernelized feature matching for geometry estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17765–17775, 2023. 2, 3
- [19] Johan Edstedt, Qiyu Sun, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. Roma: Robust dense feature matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19790–19800, 2024. 2, 3
- [20] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE international conference on computer vision*, pages 2650–2658, 2015. 7, 17
- [21] Mohamed El Banani, Amit Raj, Kevis-Kokitsi Maninis, Abhishek Kar, Yuanzhen Li, Michael Rubinstein, Deqing Sun, Leonidas Guibas, Justin Johnson, and Varun Jampani. Probing the 3d awareness of visual foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21795–21806, 2024. 5
- [22] José M Fácil, Alejo Concha, Luis Montesano, and Javier Civera. Single-view and multi-view depth fusion. *IEEE Robotics and Automation Letters*, 2(4):1994–2001, 2017. 3
- [23] Yuxin Fang, Li Dong, Hangbo Bao, Xinggang Wang, and Furu Wei. Corrupted image modeling for self-supervised visual pre-training. *arXiv preprint arXiv:2202.03382*, 2022. 3
- [24] Ziyue Feng, Liang Yang, Longlong Jing, Haiyan Wang, YingLi Tian, and Bing Li. Disentangling object motion and occlusion for unsupervised multi-frame monocular depth. In *European Conference on Computer Vision*, pages 228–244. Springer, 2022. 3, 7, 8, 17, 23, 32



- [25] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. [7](#), [8](#), [16](#), [17](#), [23](#), [24](#), [31](#)
- [26] Clément Godard, Oisín Mac Aodha, and Gabriel J Brostow. Unsupervised monocular depth estimation with left-right consistency. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 270–279, 2017. [2](#)
- [27] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3828–3838, 2019. [5](#), [7](#), [8](#), [16](#), [17](#), [23](#)
- [28] Ariel Gordon, Hanhan Li, Rico Jonschkowski, and Anelia Angelova. Depth from videos in the wild: Unsupervised monocular depth learning from unknown cameras. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8977–8986, 2019. [2](#), [23](#)
- [29] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. [2](#)
- [30] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2485–2494, 2020. [2](#), [7](#), [17](#)
- [31] Vitor Guizilini, Jie Li, Rares Ambrus, and Adrien Gaidon. Geometric unsupervised domain adaptation for semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8537–8547, 2021. [7](#), [17](#)
- [32] Vitor Guizilini, Rares Ambrus, Dian Chen, Sergey Zakharov, and Adrien Gaidon. Multi-frame self-supervised depth with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 160–170, 2022. [3](#), [5](#), [7](#), [17](#)
- [33] Agrim Gupta, Jiajun Wu, Jia Deng, and Fei-Fei Li. Siamese masked autoencoders. *Advances in Neural Information Processing Systems*, 36:40676–40693, 2023. [1](#), [2](#)
- [34] Guangxing Han and Ser-Nam Lim. Few-shot object detection with foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28608–28618, 2024. [1](#)
- [35] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [4](#)
- [36] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [1](#), [2](#)
- [37] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. [1](#), [2](#), [3](#)
- [38] Mu He, Le Hui, Yikai Bian, Jian Ren, Jin Xie, and Jian Yang. Ra-depth: Resolution adaptive self-supervised monocular depth estimation. In *European Conference on Computer Vision*, pages 565–581. Springer, 2022. [7](#), [17](#)
- [39] Heiko Hirschmüller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on pattern analysis and machine intelligence*, 30(2):328–341, 2007. [17](#)
- [40] Sunghwan Hong and Seungryong Kim. Deep matching prior: Test-time optimization for dense correspondence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9907–9917, 2021. [6](#), [8](#)
- [41] Sunghwan Hong, Seokju Cho, Jisu Nam, Stephen Lin, and Seungryong Kim. Cost aggregation with 4d convolutional swin transformer for few-shot segmentation. In *European Conference on Computer Vision*, pages 108–126. Springer, 2022. [3](#)
- [42] Sunghwan Hong, Jisu Nam, Seokju Cho, Susung Hong, Sangryul Jeon, Dongbo Min, and Seungryong Kim. Neural matching fields: Implicit representation of matching fields for visual correspondence. *Advances in Neural Information Processing Systems*, 35:13512–13526, 2022. [3](#)
- [43] Sunghwan Hong, Jaewoo Jung, Heeseong Shin, Jiaolong Yang, Seungryong Kim, and Chong Luo. Unifying correspondence, pose and nerf for pose-free novel view synthesis from stereo pairs. *arXiv preprint arXiv:2312.07246*, 2023. [8](#)
- [44] Sunghwan Hong, Seokju Cho, Seungryong Kim, and Stephen Lin. Unifying feature and cost aggregation with transformers for semantic and visual correspondence. In *The Twelfth International Conference on Learning Representations*, 2024. [2](#), [4](#), [15](#)
- [45] Sunghwan Hong, Jaewoo Jung, Heeseong Shin, Jiaolong Yang, Seungryong Kim, and Chong Luo. Unifying correspondence pose and nerf for generalized pose-free novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20196–20206, 2024. [3](#)
- [46] Adrian Johnston and Gustavo Carneiro. Self-supervised monocular trained depth estimation using self-attention and discrete disparity volume. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4756–4765, 2020. [2](#)
- [47] Rico Jonschkowski, Austin Stone, Jonathan T. Barron, Ariel Gordon, Kurt Konolige, and Anelia Angelova. What matters in unsupervised optical flow, 2020. [2](#), [3](#)
- [48] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9492–9502, 2024. [3](#)
- [49] Seungryong Kim, Stephen Lin, Sang Ryul Jeon, Dongbo Min, and Kwanghoon Sohn. Recurrent transformer net-

works for semantic correspondence. *Advances in neural information processing systems*, 31, 2018. 5, 33

- [50] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 15, 16
- [51] Lingdong Kong, Shaoyuan Xie, Hanjiang Hu, Lai Xing Ng, Benoit R. Cottureau, and Wei Tsang Ooi. Robodepth: Robust out-of-distribution depth estimation under corruptions. In *Advances in Neural Information Processing Systems*, 2023. 8, 17
- [52] Junhyup Lee, Dohyung Kim, Jean Ponce, and Bumsuh Ham. Sfnet: Learning object-aware semantic correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2278–2287, 2019. 14
- [53] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. *arXiv preprint arXiv:2406.09756*, 2024. 2, 4, 6, 8, 17, 19, 21, 22, 24, 28, 29
- [54] Hanhan Li, Ariel Gordon, Hang Zhao, Vincent Casser, and Anelia Angelova. Unsupervised monocular depth learning in dynamic scenes. In *Conference on Robot Learning*, pages 1908–1917. PMLR, 2021. 2, 23
- [55] Rui Li, Dong Gong, Wei Yin, Hao Chen, Yu Zhu, Kaixuan Wang, Xiaozhi Chen, Jinqiu Sun, and Yanning Zhang. Learning to fuse monocular and multi-view cues for multi-frame depth estimation in dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21539–21548, 2023. 3
- [56] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2041–2050, 2018. 15, 17, 21, 22
- [57] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Self-low: Self-supervised learning of optical flow. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4571–4580, 2019. 2, 3
- [58] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 14
- [59] Ao Luo, Xin Li, Fan Yang, Jiangyu Liu, Haoqiang Fan, and Shuaicheng Liu. Flowdiffuser: Advancing optical flow estimation with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19167–19176, 2024. 1
- [60] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 2
- [61] Iaroslav Melekhov, Aleksei Tiulpin, Torsten Sattler, Marc Pollefeys, Esa Rahtu, and Juho Kannala. Dgc-net: Dense geometric correspondence network. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1034–1042. IEEE, 2019. 2, 6, 16
- [62] Juhong Min, Dahyun Kang, and Minsu Cho. Hypercorrelation squeeze for few-shot segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 15
- [63] Rohit Mohan and Abhinav Valada. Efficientpts: Efficient panoptic segmentation. *International Journal of Computer Vision*, 129(5):1551–1579, 2021. 17
- [64] Jisu Nam, Gyuseong Lee, Sunwoo Kim, Hyeonsu Kim, Hyungwon Cho, Seyeon Kim, and Seungryong Kim. Diffusion model for dense matching. In *The Twelfth International Conference on Learning Representations*. 6, 8, 15, 17, 30
- [65] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023. 1, 5, 6, 7, 16, 21, 22
- [66] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 15
- [67] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 16
- [68] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016. 3
- [69] Vaishakh Patil, Wouter Van Gansbeke, Dengxin Dai, and Luc Van Gool. Don’t forget the past: Recurrent depth estimation from monocular video. *IEEE Robotics and Automation Letters*, 5(4):6813–6820, 2020. 7, 17
- [70] Luigi Piccinelli, Yung-Hsu Yang, Christos Sakaridis, Mattia Segu, Siyuan Li, Luc Van Gool, and Fisher Yu. Unidepth: Universal monocular metric depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10106–10116, 2024. 3
- [71] Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024. 14
- [72] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12179–12188, 2021. 5, 15, 16
- [73] Zhe Ren, Junchi Yan, Bingbing Ni, Bin Liu, Xiaokang Yang, and Hongyuan Zha. Unsupervised deep learning for optical flow estimation. In *Proceedings of the AAAI conference on artificial intelligence*, 2017. 2

- [74] Jerome Revaud, Yohann Cabon, Romain Brégier, JongMin Lee, and Philippe Weinzaepfel. Sacreg: Scene-agnostic coordinate regression for visual localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 688–698, 2024. [1](#), [8](#)
- [75] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. *Advances in neural information processing systems*, 31, 2018. [5](#)
- [76] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. [7](#)
- [77] Patrick Ruhkamp, Daoyi Gao, Hanzhi Chen, Nassir Navab, and Benjamin Busam. Attention meets geometry: Geometry guided spatial-temporal attention for consistent self-supervised monocular depth estimation. In *2021 International Conference on 3D Vision (3DV)*, pages 837–847. IEEE, 2021. [7](#), [17](#)
- [78] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3260–3269, 2017. [6](#), [7](#), [8](#), [16](#), [17](#)
- [79] Xi Shen, François Darmon, Alexei A Efros, and Mathieu Aubry. Ransac-flow: generic two-stage image alignment. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 618–637. Springer, 2020. [2](#)
- [80] Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan. Emergent correspondence from image diffusion. *Advances in Neural Information Processing Systems*, 36:1363–1389, 2023. [5](#), [6](#), [7](#), [16](#), [21](#), [22](#), [25](#), [26](#), [27](#)
- [81] Prune Truong, Martin Danelljan, Luc V Gool, and Radu Timofte. Gocor: Bringing globally optimized correspondence volumes into your neural network. *Advances in Neural Information Processing Systems*, 33:14278–14290, 2020. [2](#), [6](#), [8](#), [17](#), [30](#)
- [82] Prune Truong, Martin Danelljan, and Radu Timofte. Glunet: Global-local universal network for dense flow and correspondences. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6258–6268, 2020. [2](#), [3](#), [5](#), [6](#), [15](#), [16](#), [17](#)
- [83] Prune Truong, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning accurate dense correspondences and when to trust them. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5714–5724, 2021. [2](#), [6](#), [15](#), [17](#)
- [84] Prune Truong, Martin Danelljan, Radu Timofte, and Luc Van Gool. Pdc-net+: Enhanced probabilistic dense correspondence network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8):10247–10266, 2023. [2](#), [3](#), [5](#), [6](#), [8](#), [15](#), [17](#), [30](#)
- [85] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. [4](#), [5](#), [14](#), [15](#)
- [86] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. [1](#), [2](#), [4](#), [6](#), [8](#), [17](#), [19](#), [21](#), [22](#), [24](#), [28](#), [29](#)
- [87] Xiaofeng Wang, Zheng Zhu, Guan Huang, Xu Chi, Yun Ye, Ziwei Chen, and Xingang Wang. Crafting monocular cues and velocity guidance for self-supervised multi-frame depth learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2689–2697, 2023. [3](#), [7](#), [17](#)
- [88] Yang Wang, Peng Wang, Zhenheng Yang, Chenxu Luo, Yi Yang, and Wei Xu. Unos: Unified unsupervised optical-flow and stereo-depth estimation by watching videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8071–8081, 2019. [2](#), [4](#)
- [89] Yihan Wang, Lahav Lipson, and Jia Deng. Sea-raft: Simple, efficient, accurate raft for optical flow. In *European Conference on Computer Vision*, pages 36–54. Springer, 2025. [1](#)
- [90] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. [2](#), [4](#)
- [91] Jamie Watson, Oisín Mac Aodha, Victor Prisacariu, Gabriel Brostow, and Michael Firman. The temporal opportunist: Self-supervised multi-frame monocular depth. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1164–1174, 2021. [2](#), [3](#), [5](#), [7](#), [8](#), [16](#), [17](#), [23](#), [31](#), [32](#)
- [92] Philippe Weinzaepfel, Vincent Leroy, Thomas Lucas, Romain Brégier, Yohann Cabon, Vaibhav Arora, Leonid Antsfeld, Boris Chidlovskii, Gabriela Csurka, and Jérôme Revaud. Croco: Self-supervised pre-training for 3d vision tasks by cross-view completion. *Advances in Neural Information Processing Systems*, 35:3502–3516, 2022. [1](#), [2](#), [3](#), [4](#), [20](#), [21](#)
- [93] Philippe Weinzaepfel, Thomas Lucas, Vincent Leroy, Yohann Cabon, Vaibhav Arora, Romain Brégier, Gabriela Csurka, Leonid Antsfeld, Boris Chidlovskii, and Jérôme Revaud. Croco v2: Improved cross-view completion pre-training for stereo matching and optical flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17969–17980, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#), [14](#), [17](#), [19](#), [20](#), [21](#), [22](#), [23](#), [24](#), [28](#), [29](#)
- [94] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatofghi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8121–8130, 2022. [4](#)
- [95] Lihe Yang, Bingyi Kang, Zilong Huang, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10371–10381, 2024. [3](#)
- [96] Zhenpei Yang, Zhile Ren, Qi Shan, and Qixing Huang. Mvs2d: Efficient multi-view stereo via attention-driven 2d

- convolutions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8574–8584, 2022. [3](#)
- [97] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9043–9053, 2023. [3](#)
- [98] Jason J Yu, Adam W Harley, and Konstantinos G Derpanis. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part III 14*, pages 3–10. Springer, 2016. [2](#), [4](#)
- [99] Junyi Zhang, Charles Herrmann, Junhwa Hur, Luisa Polania Cabrera, Varun Jampani, Deqing Sun, and Ming-Hsuan Yang. A tale of two features: Stable diffusion complements dino for zero-shot semantic correspondence. *Advances in Neural Information Processing Systems*, 36, 2024. [5](#), [6](#), [7](#), [16](#), [21](#), [22](#), [25](#), [26](#), [27](#)
- [100] Chaoqiang Zhao, Youmin Zhang, Matteo Poggi, Fabio Tosi, Xianda Guo, Zheng Zhu, Guan Huang, Yang Tang, and Stefano Mattoccia. Monovit: Self-supervised monocular depth estimation with a vision transformer. In *2022 international conference on 3D vision (3DV)*, pages 668–678. IEEE, 2022. [7](#), [17](#)
- [101] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1851–1858, 2017. [17](#)
- [102] Chaoyang Zhu and Long Chen. A survey on open-vocabulary detection and segmentation: Past, present, and future. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. [1](#)