

# Deep Matching Prior: Test-Time Optimization for Dense Correspondence

Sunghwan Hong  
Korea University

sung\_hwan@korea.ac.kr

Seungryong Kim\*  
Korea University

seungryong\_kim@korea.ac.kr

## Abstract

Conventional techniques to establish dense correspondences across visually or semantically similar images focused on designing a task-specific matching prior, which is difficult to model in general. To overcome this, recent learning-based methods have attempted to learn a good matching prior within a model itself on large training data. The performance improvement was apparent, but the need for sufficient training data and intensive learning hinders their applicability. Moreover, using the fixed model at test time does not account for the fact that a pair of images may require their own prior, thus providing limited performance and poor generalization to unseen images.

In this paper, we show that an image pair-specific prior can be captured by solely optimizing the untrained matching networks on an input pair of images. Tailored for such test-time optimization for dense correspondence, we present a residual matching network and a confidence-aware contrastive loss to guarantee a meaningful convergence. Experiments demonstrate that our framework, dubbed Deep Matching Prior (DMP), is competitive, or even outperforms, against the latest learning-based methods on several benchmarks for geometric matching and semantic matching, even though it requires neither large training data nor intensive learning. With the networks pre-trained, DMP attains state-of-the-art performance on all benchmarks.

## 1. Introduction

Establishing dense correspondences across visually or semantically similar images facilitates a variety of computer vision applications [45, 44, 28, 37]. Unlike sparse correspondence [48, 5, 76] that detects sparse points and finds matches across them, dense correspondence [45, 6, 35, 37] aims at finding matches at each pixel and thus can benefit from prior<sup>1</sup> knowledge about matches among nearby pixels.

Typically, stereo matching [32, 6] and optical flow [24, 65] modeled the prior term that makes the correspondence

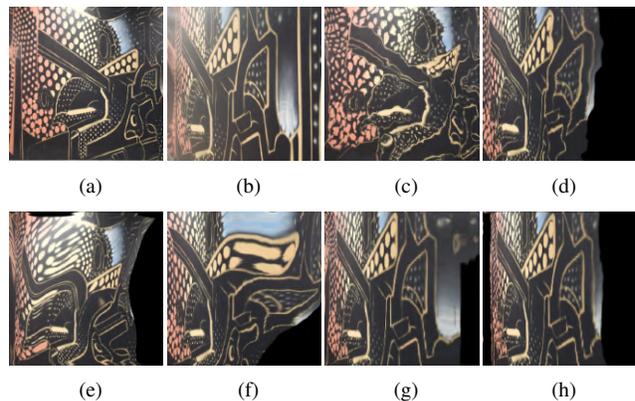


Figure 1. **Visualization of DMP results:** (a) source image, (b) target image, results of (c) winner-takes-all (WTA) and (d) learning-based method (e.g., GLU-Net [70]). As the iteration evolves (e), (f), (g), and (h), DMP with *untrained* networks estimates more optimal correspondence fields by optimizing the networks on a *single pair of images*, while learning-based methods utilize *pre-trained* and *fixed* networks at test time, thus providing limitation.

smooth while aligning discontinuities to image boundaries. Some methods [45, 67, 19, 38] for semantic matching also exploited this to regularize the correspondence within a local neighborhood. Although they can be formulated in various ways, these optimization-based methods [45, 67, 19, 38] formulate an objective function with explicitly defined matching data and prior terms and minimize the objective on a single image pair. These methods are capable of making corrections to the estimated correspondence during optimization, but they require a task-specific prior, which is complex to formulate.

Unlike these learning-free methods [45, 67, 19, 38], recent methods [55, 66, 61, 49, 70] cast this task as a learning problem, seeking to learn a matching model to directly regress the correspondence. The model, often implemented based on convolutional neural networks (CNNs) [64, 22], is generally trained on large datasets of image pairs, based on the belief that an optimal matching prior can be learned from such observations. As proven in literature [49, 70], these *learning*-based methods outperform traditional optimization-based methods [45, 19, 38],

\*Corresponding author

<sup>1</sup>It can also be called a smoothness or a regularizer in literature.

which can be attributed to their high capacity to learn a good matching prior. They, however, often require large training data with ground-truth correspondences, which are notoriously hard to collect, or intensive learning [49, 70], which hinders their applicability. In addition, a pair of images may require their own prior, and thus using *pre-trained* and *fixed* parameters at test time may provide limited and poor generalization performance to unseen image pairs.

In this paper, we show that, for the first time, the matching prior must not necessarily be learned from large training data; instead, it can be captured by optimizing the *untrained* matching networks on a *single pair of images*, proving that the structure of the networks is sufficient to capture a great deal of matching statistics. Our framework, dubbed Deep Matching Prior (DMP), requires neither large training sets nor intensive training, but it is competitive against learning-based methods [55, 66, 49, 61, 70, 62] or even outperforms, and does not suffer from the generalization issue.

Such a test-time optimization for dense correspondence, however, is extremely hard to converge due to the limited samples, high-dimensional search space, and non-convexity of the objective. To elevate the stability and boost the convergence, we propose a residual matching network, which exploits a distilled information from matching correlation that plays as a guidance for providing a good starting match and suppressing the possibility of getting stuck in local minima. We also propose a confidence-aware contrastive loss that only takes the matches with high confidence to eliminate the ambiguous matches. Fig. 1 visualizes the results of DMP in comparison with recent learning-based method [70].

The presented approach is evaluated on several benchmarks for dense correspondence and examined in an ablation study. The extensive experiments show that our model produces competitive results and even outperforms other learning-based methods [55, 66, 61, 49, 70], and once the networks are pre-trained, state-of-the-art performance is attained on all the benchmarks in experiments.

## 2. Related Works

**Dense correspondence.** Most early efforts for dense correspondence [45, 34, 67] have focused on designing a task-specific prior term in an objective function and improving optimization while employing hand-crafted features. Kim et al. [38] presented DCTM optimizer with a discontinuity-aware prior term to elevate geometric invariance.

Similarly to other tasks, most recent works [37, 68, 73, 25, 57, 42, 43] leverage deep learning to achieve better performance for this task, at first replacing the hand-crafted features with the deep features [9, 36], and rapidly converging towards end-to-end networks embodying the entire pipeline [58, 37, 42]. In this context, most recent efforts have been focusing on how to better design the matching

networks and learn the networks without ground-truth correspondences. Rocco et al. [55, 56] proposed a geometric matching network, and their success inspired many variants that use local neighborhood consensus [58, 42], attention mechanism [23], or attribute transfer [37]. Similar to PWC-Net [65] designed specifically for optical flow estimation, PARN [27], DGC-Net [49], GLU-Net [70] and GOCor [69] formulate their networks in a pyramidal fashion. RANSAC-Flow [62] leverages RANSAC [14] within the network to improve the performance. On the other hand, to overcome the lack of ground-truth correspondences for training, synthetic supervision generated by applying random transformations [55, 70, 50, 69] and weak supervision in the form of image pairs based on the feature reconstruction [27, 56, 37] have been popularly used. However, these aforementioned works require either large labeled dataset or intensive training, which hinders their applicability.

**Self-supervised representation learning.** Popularized by [7, 21, 8, 15], contrastive learning seeks to learn feature representation in a self-supervised manner by minimizing the distance between two views augmented from an image and maximizing the distance to other images [18, 10, 17]. Several methods [59, 9, 38, 36, 72, 53] bring contrastive learning to the task of dense correspondence. A key benefit of self-supervised learning is that, because there is no reliance on labeled data, training needs not be limited to the training phase [12, 51, 63, 25, 40, 62]. CRW [25] and RANSAN-Flow [62] attempted to apply such a self-supervised adaptation at test time, but the performance was limited. The use of self-supervised adaptation for dense correspondence has never been thoroughly investigated and this paper is the first step in this direction.

**Test-time optimization.** Deep Image Prior (DIP) [71] initiated the trend that the low-level statistics in a single image can be captured by the structure of randomly-initialized CNNs, of which many variants were proposed, tailored to solve an inverse problem [13, 74, 63, 29]. GAN-inversion aims to invert an image back to latent code, and then reconstruct the image from the latent code to utilize the pre-trained generative prior [26, 1, 16, 30, 31, 77]. Deep Generative Prior (DGP) [52] extends this trend by, for the pre-trained generator, optimizing both latent codes and parameters. These aforementioned methods were tailored to capture an image prior, but there was no attempt to capture a matching prior, which is the topic of this paper.

## 3. Deep Matching Prior

### 3.1. Motivation

Let us denote a pair of images, i.e., source and target, as  $I^s$  and  $I^t$ , which may represent visually or semantically similar images. The objective of dense correspondence is

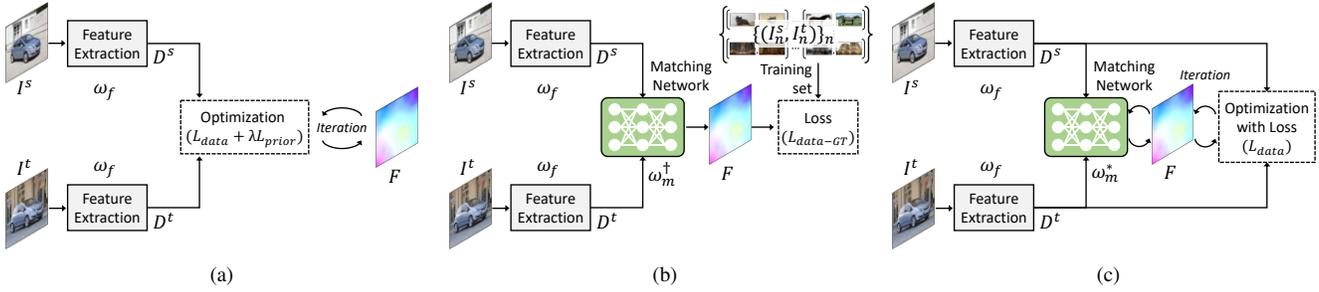


Figure 2. **Intuition of DMP:** (a) optimization-based methods [45, 67, 38] that formulate their objective function with data and prior terms, and minimize the energy function on a single image pair, (b) learning-based methods [49, 70, 62] that learn a matching prior from large training set of image pairs, and (c) our DMP, which takes the best of both approaches to estimate an image pair-specific matching prior.

to establish a dense correspondence field  $F(i)$  between two images that is defined for each pixel  $i$ , which warps  $I^s$  towards  $I^t$  so as to satisfy  $I^t(i) \approx I^s(i + F(i))$ .

To achieve this, the easiest solution may be, for a point in the source, to find the most similar one among all the points in the target with respect to a *data* term that accounts for matching evidence between features, defined such that

$$F^* = \operatorname{argmin}_F \mathcal{L}_{data}(D^s(F), D^t), \quad (1)$$

where a feature  $D$  is extracted with parameters  $\omega_f$  such that  $D = \mathcal{F}(I; \omega_f)$  and  $D^s(F)$  is an warped source feature with  $F$ . This solution  $F^*$ , so-called winner-takes-all (WTA) solution, is highly vulnerable to local minima or outliers.

To alleviate these, *optimization*-based methods [2, 45, 4, 67, 44, 38] typically formulate an objective function involving the data term and a *prior* term that favors similar correspondence fields among adjacent pixels with a balancing parameter  $\lambda$ , as shown in Fig. 2 (a), such that

$$F^* = \operatorname{argmin}_F \{ \mathcal{L}_{data}(D^s(F), D^t) + \lambda \mathcal{L}_{prior}(F) \}. \quad (2)$$

Traditional optimization-based methods [45, 67, 19, 45, 19, 38] model an explicit prior term such as total variation (TV) or discontinuity-aware smoothness and optimize the energy function on a given pair of images. Since this energy function is often non-convex, they use an iterative solver, e.g., gradient descent [39], to minimize this, thus they can benefit from an error feedback to find more optimal solution by making corrections to the estimated correspondence as the iteration evolves. However, they require a hand-designed task-specific prior, which is hard to design.

Unlike these optimization-based methods, recent *learning*-based methods [55, 66, 61, 49, 68, 70] learn a matching model to directly regress the correspondence such that  $F = \mathcal{F}(I^s, I^t; \omega_m)$  with parameters  $\omega_m$ , based on an assumption that an optimal prior could be learned within the model itself from massive training samples,  $\{(I_n^s, I_n^t)\}_{n \in \{1, \dots, N\}}$ , as shown in Fig. 2 (b). The model is often implemented in the form of convolutional neural networks (CNNs). During

the training phase, the parameters  $\omega_m^\dagger$  are first optimized, and the learned parameters are then used to regress a correspondence  $F^*$  through the networks with the learned prior at test time as follows:

$$\begin{aligned} \omega_m^\dagger &= \operatorname{argmin}_\omega \sum_n \mathcal{L}_{data-GT}(\mathcal{F}(I_n^s, I_n^t; \omega), \hat{F}_n), \\ F^* &= \mathcal{F}(I^s, I^t; \omega_m^\dagger), \end{aligned} \quad (3)$$

where  $\hat{F}_n$  denotes a ground-truth correspondence for  $n$ -th sample in the training set. As proven in literature [49, 70], these learning-based methods have shown better performance than the traditional optimization-based methods [45, 19, 38] thanks to their high capacity to learn the good matching prior. To optimize the parameters by minimizing the data term, e.g.,  $\mathcal{L}_{data-GT}(F, \hat{F}) = \|F - \hat{F}\|_1$ , they, however, require massive training image pairs with ground-truth correspondences, which are notoriously hard to collect. Some recent methods [55, 62] overcome this limitation by presenting an unsupervised loss defined only with source and target features, e.g.,  $\mathcal{L}_{data}(D^s(F), D^t)$ , or synthetic ground-truths, but these methods inherit the limitation of requiring an intensive training procedure and frequently fail to learn an image pair-specific prior, thus providing limited performance and generalization power.

### 3.2. Formulation

In this session, we argue that the matching prior does not necessarily need to be learned from an intensive learning; instead, an image pair-specific matching prior can be captured by solely minimizing the data term on a *single pair of images*, like what is done by traditional optimization-based methods [45, 38], with an *untrained* network for matching, as shown in Fig. 2 (c), which can be formulated as

$$\begin{aligned} \omega_m^* &= \operatorname{argmin}_\omega \mathcal{L}_{data}(D^s(\mathcal{F}(I^s, I^t; \omega)), D^t), \\ F^* &= \mathcal{F}(I^s, I^t; \omega_m^*), \end{aligned} \quad (4)$$

where  $\omega_m^*$  are parameters, over-fitted to a single pair of images to allow for generating  $F^*$  with a pair-specific matching prior, and can be obtained using an optimizer such as



Figure 3. **DMP convergence:** (From left to right) source image, target image, iterative evolution of warped images. Given a good initialization, which facilitates convergence and boosts matching performance, our approach successfully estimates the correspondence fields.

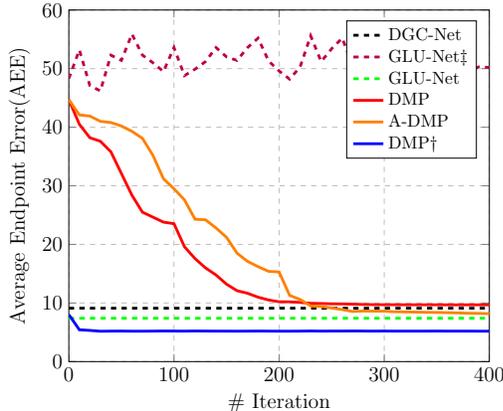


Figure 4. **Convergence analysis of DMP.** Starting from *untrained* network, our DMP converges to better correspondences by optimizing the network with a well-designed loss function on a single pair of images. Compared to the *learning*-based methods, such as DGC-Net [49] and GLU-Net [70] that are *pre-trained* with an intensive learning on large training datasets, our DMP has shown a competitive performance after convergence, which can be expedited with pre-trained initialization (denoted by DMP†). More details can be found in supplementary material.

gradient descent [39], starting from a *random* initialization. Our framework, dubbed Deep Matching Prior (DMP), requires neither massive training image pairs nor ground-truth correspondences that are major bottlenecks of existing learning-based methods [49, 70, 62], but only requires a single pair of images to be matched, which is competitive when compared to learning-based methods [49, 37, 70, 62] and even outperforms, as exemplified in Fig. 3 and Fig. 4. Thanks to its inherent error feedback nature, DMP does not suffer from a generalization issue for unseen image pairs, which existing learning-based methods [49, 37, 70, 62] frequently fail to avoid.

However, optimizing the existing matching network [49, 70, 62] on a single image pair with simple data term [45, 38] is extremely hard to converge due to the lack of samples, high-dimensional search space for dense correspondence and its non-convexity [38]. As shown in Fig. 4, when an existing matching network, e.g., GLU-Net [75], is optimized from a random initialization with a self-supervised loss, e.g., feature matching [62], denoted by GLU-Net†, it fails to find meaningful correspondences. To overcome these, we present a residual matching network and a confidence-aware contrastive loss tailored to boost the convergence and

matching performance. In the following, we describe the network architecture and loss function in detail.

### 3.3. Network Architecture

To guarantee a meaningful convergence during the optimization, a good initialization for correspondence  $F$  should be set, even though our network parameters  $\omega$  are randomly initialized. To achieve this, we formulate our model, consisting of *feature extraction* networks and *matching* networks, in a residual manner, as illustrated in Fig. 5.

**Feature extraction networks.** Our model accomplishes dense correspondence using deep features, e.g., VGG [64] or ResNet [22], which undergo  $l_2$  normalization. For the backbone feature, we exploit a model pre-trained on ImageNet [11], as done in almost most literature [58, 37, 49, 37, 70, 62]. The backbone model could be directly optimized in DMP framework, but we found that using an additional adaptation layer [41] to refine the backbone features and optimizing the layer only boosts the performance drastically. This is because it helps to focus on more learning matching networks at early training stages. The learnable parameters  $\omega_f$  in the feature extractor are thus from the adaptation layers. At the initialization, we zero-initialize the layer to make it behave like an identity.

**Matching networks.** Our matching networks consist of correlation map computation and inference modules. The correlation maps are first computed using the inner product between features as

$$\mathcal{C}(i, l) = D^s(i) \cdot D^t(l), \quad (5)$$

where  $l$  is defined within a search space in the target.

Based on the correlation  $\mathcal{C}$ , correspondence  $F_s$  is estimated through an inference module with parameters  $\omega_m$  as

$$F_s = \mathcal{F}(\mathcal{C}; \omega_m). \quad (6)$$

A randomly-initialized  $\omega_m$ , however, gives a noisy correspondence at the initial phase, making the optimization difficult. To overcome this, we enforce the inference module to estimate a residual of current best matches from given correlation volume, which is achieved by a soft-argmax operation [33] over  $\mathcal{C}$ , such that

$$F = \mathcal{F}(\mathcal{C}; \omega_m) + \Phi(\mathcal{C}), \quad (7)$$

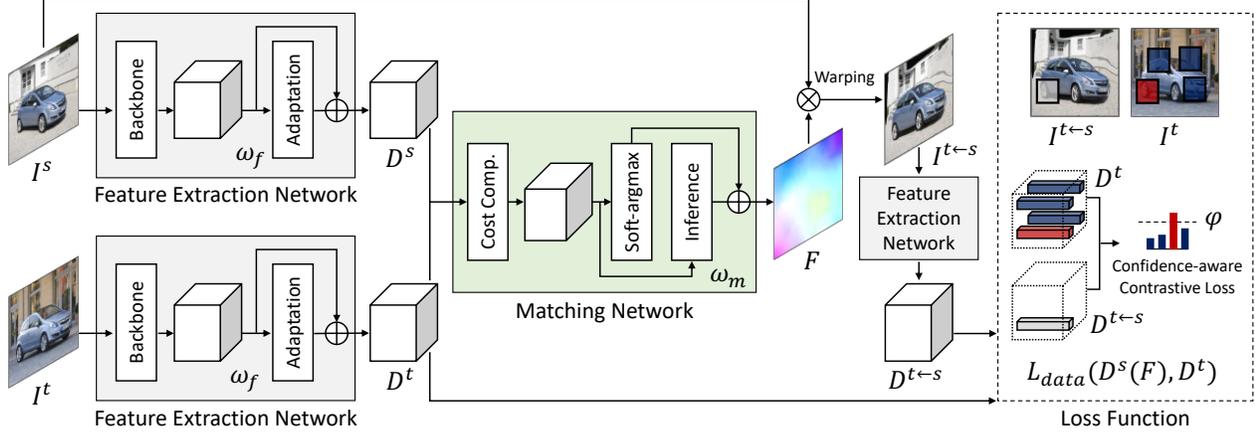


Figure 5. **Network configuration and loss function of DMP:** Our networks consist of feature extraction and matching networks, which are formulated in a residual manner to guarantee a good initialization for optimization. Note that a single-level version of the networks are illustrated for brevity, while the full model is formulated in a pyramidal fashion. A confidence-aware contrastive loss enables joint learning of feature extraction and matching networks by rejecting ambiguous matches while accepting confident matches through a thresholding.

where  $\Phi(\cdot)$  represents a soft-argmax operator. By setting the parameters of the last layer to generate zeros at the initialization, an optimizer can start at least the current best matches from given correlation volume. As evolving the iterations, the networks provide more regularized matching fields that encode an image pair-specific prior.

**Coarse-to-fine formulation.** Analogous to [66, 65, 27, 49, 70], we also utilize the coarse-to-fine approach through a pyramidal processing. Specifically, we exploit pyramidal features from coarse-to-fine levels to simultaneously provide robustness against deformations and improve fine-grained matching details by minimizing the data term at each layer. At the coarsest level, we consider a global correlation module, while a local correlation module is used at remaining levels. In addition to the use of soft-argmax for residual flow learning, at each level, previous level’s matches are up-sampled to play as a guidance. To handle any input resolution, we employ an adaptive resolution [70]. The details can be found in the supplementary material.

### 3.4. Loss function

Similarly to traditional optimization-based methods [45, 38] and unsupervised learning-based methods [56, 37], the data term in DMP is defined as the similarity function  $\mathcal{S}(\cdot, \cdot)$  of features, aggregated across the points  $i \in \{1, \dots, M\}$  as

$$\mathcal{L}_{\text{data}}(D^s(F), D^t) = -\frac{1}{M} \sum_i \mathcal{S}(D^{t \leftarrow s}(i), D^t(i)), \quad (8)$$

where  $D^{t \leftarrow s}$  represents an warped source feature with  $F$ . Although source feature can be directly warped, we extract the feature directly from warped source image  $I^{t \leftarrow s}$  with the final correspondence  $F$  at the pyramid that incorporates the information at all pyramidal layers, which demonstrates more stable convergence.

Perhaps the simplest similarity function is to use the inner product between  $l$ -2 normalized features as  $\mathcal{S}(D^{t \leftarrow s}(i), D^t(i)) = D^{t \leftarrow s}(i) \cdot D^t(i)$ , accounting for the intuition that the similarities between all pairs of the warped source and target features at the same locations should be maximized. Minimizing this, however, can induce erroneous solutions, e.g., constant features at all the points. To avoid such trivial solutions, we extend this similarity function in a contrastive learning fashion, aiming to maximize the similarities at the same locations while minimizing for the others, such that

$$\mathcal{S}_c = -\log \left( \frac{\exp(D^{t \leftarrow s}(i) \cdot D^t(i)/\tau)}{\sum_j \exp(D^{t \leftarrow s}(i) \cdot D^t(j)/\tau)} \right), \quad (9)$$

where  $j \in \{1, \dots, M\}$  and  $\tau$  is a temperature hyperparameter. By minimizing the loss function, both feature extraction and matching networks with parameters  $\omega_f$  and  $\omega_m$ , respectively, can be simultaneously optimized in a manner that the feature extractors embed the warped source and target features at same location to the same representation while the matching networks regress better correspondences that maximize the feature similarity.

Even though this loss helps to successfully avoid the trivial solution and allows joint learning of feature and matching networks, the erroneous initial estimates may be propagated, without a term to eliminate such estimates. Since our DMP optimizes the networks on a single pair of images, such adverse effects may be critical. To mitigate this, we present a confidence-aware contrastive loss that enables rejecting such ambiguous matches with a thresholding while accepting the confident matches, defined such that

$$\mathcal{S}_{\text{cac}} = -\log(\Psi(\mathcal{S}_c, \varphi)), \quad (10)$$

Methods	Pre-train	Test-opt.	HPatches (240 × 240)							HPatches						
			I	II	III	IV	V	Avg.	PCK	I	II	III	IV	V	Avg.	PCK
CNNGeo [55]	✓	✗	9.59	18.55	21.15	27.83	35.19	22.46	-	-	-	-	-	-	-	-
DGC-Net [49]	✓	✗	1.74	5.88	9.07	12.14	16.50	9.07	50.01	5.71	20.48	34.15	43.94	62.01	33.26	58.06
GLU-Net [70]	✓	✗	0.59	4.05	7.64	9.82	14.89	7.40	83.47	1.55	12.66	27.54	32.04	52.47	25.05	78.54
GLU-GOCor [69]	✓	✗	-	-	-	-	-	-	-	1.29	10.07	23.86	27.17	38.41	20.16	81.43
RANSAC-Flow [62]	✓	✗	0.51	2.36	2.91	4.41	<b>5.12</b>	3.06	-	-	-	-	-	-	-	-
	✗	✓	3.81	5.76	5.92	8.31	13.24	7.40	81.28	12.41	18.64	20.09	35.81	35.00	24.39	22.06
	✓	✓	1.29	3.52	4.01	6.88	10.27	5.19	88.23	10.25	15.22	15.19	30.43	48.09	23.84	26.21
<b>DMP</b>	✗	✓	1.21	5.12	12.31	13.68	16.12	9.69	79.21	3.21	15.54	32.54	38.62	63.43	30.64	63.21
<b>A-DMP</b>	✗	✓	1.31	4.81	10.21	10.69	13.88	8.18	82.35	3.42	14.21	29.90	32.82	55.8	27.22	71.64
<b>RANSAC-DMP</b>	✗	✓	0.53	<b>2.21</b>	2.76	4.62	5.14	3.05	96.28	4.32	11.21	22.80	31.34	33.64	20.65	75.35
<b>DMP†</b>	✓	✗	1.48	4.67	7.82	9.96	13.68	7.53	79.94	4.24	15.92	27.42	36.77	46.51	26.15	54.84
	✓	✓	0.77	3.36	5.22	7.32	9.38	5.21	90.89	2.41	9.88	20.64	28.21	30.15	18.23	81.74
<b>RANSAC-DMP†</b>	✓	✓	<b>0.48</b>	2.24	<b>2.41</b>	<b>4.32</b>	5.16	<b>2.92</b>	<b>97.52</b>	<b>3.57</b>	<b>8.59</b>	<b>10.18</b>	<b>21.21</b>	<b>23.81</b>	<b>13.47</b>	<b>87.62</b>

Table 1. **Quantitative evaluation on HPatches [3] dataset in terms of AEE and PCK.** Lower AEE and higher PCK (5-pixel (%)) are better. *Pre-train:* Pre-training, *Test-opt.:* Test-time optimization.

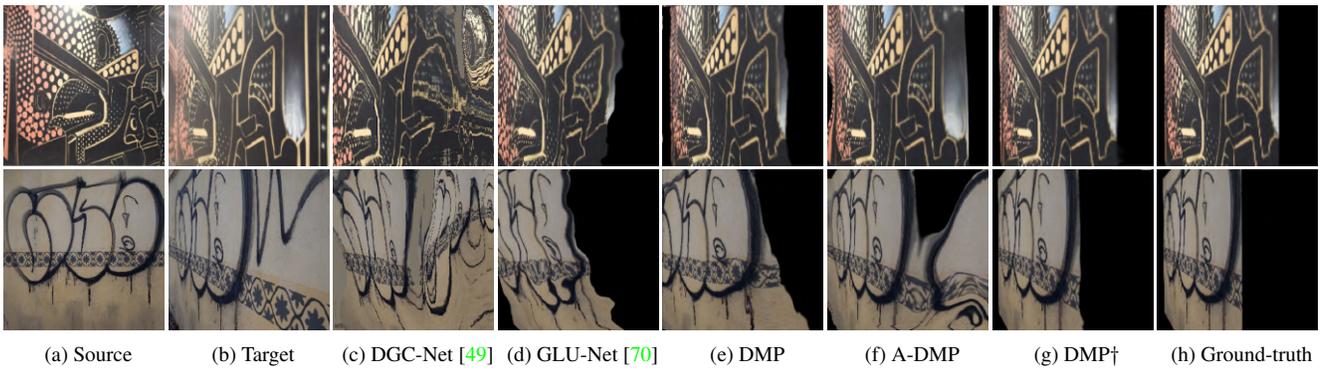


Figure 6. **Qualitative results on HPatches [3] dataset.** The source images were warped to the target images using correspondences.

where  $\varphi$  is a confidence hyperparameter;  $\Psi(x, \varphi)$  is a function designed to produce  $x$  if  $x \geq \varphi$  and 1 otherwise.

### 3.5. Extension

In this section, we introduce several variants of our DMP to expedite convergence and boost the performance.

**Test-time augmentation.** Since DMP solely depends on a single pair of images, the samples that capture matching distribution are rather limited, even though it surprisingly well learns from such limited samples, which inspires the enrichment through a test-time augmentation. We present Augmentation-DMP (A-DMP), where original target and a randomly-augmented target are fed concurrently to our model, and used to optimize the networks simultaneously. Given a pair of images having dramatically large geometric variations, the randomly-augmented target would be more similar to a source than original target at an iteration, generating the loss signal to accelerate the convergence. In addition, thanks to our confidence-aware loss, if augmented ones are more difficult to be matched, they are rejected. We apply one of spatial transformations at each iteration such as homography [12] or affine and thin plate spline (TPS) [55] to obtain the augmented targets.

**Pre-training.** Unsurprisingly, well-initialized parameters can guarantee better performance. We also show that by pre-training our networks, similarly to others [70, 69], denoted by  $\text{DMP}^\dagger$ , our DMP can start with a well learned initialization, thus yielding much better results. Note that [62] also attempted to pre-train and fine-tune their networks on a single pair of images, but as will be seen in our experiments, it failed to achieve the satisfactory performance, which proves that our well-designed networks and loss functions are essential for such a framework. In this work, we follow the training procedure identical to [70] and use the same dataset and hyperparameters.

## 4. Experiments

### 4.1. Implementation Details

For backbone feature extractor, we used VGG-16 [64] and ResNet-101 [22] pre-trained on ImageNet [11]. Specifically, for geometric matching, we used VGG-16, while for semantic matching, we additionally used ResNet-101. For the loss function, instead of using all the samples, we randomly sampled  $M = 256$  feature vectors at each iteration, considering the trade-off between performance and complexity. We set the maximum number of iterations as 2k and 0.3k for DMP (or A-DMP) and  $\text{DMP}^\dagger$ , respectively. In ad-

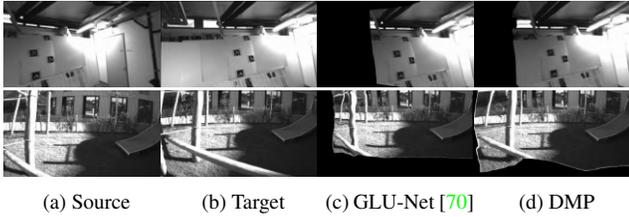


Figure 7. **Qualitative results on ETH3D [60] dataset.**

dition, we set the temperature  $\tau$  and confidence  $\varphi$  as 0.1 and 0.01, respectively, following ablation study in Sec. 4.5. We used Adam optimizer [39] fixed for all experiments, but the advanced optimizers may improve performance [54, 46].

## 4.2. Experimental Settings

In this section, we conduct comprehensive experiments for geometric matching and semantic matching, by evaluating our approach through comparisons to state-of-the-art methods including, CNNGeo [55], PARN [27], NC-Net [58], DGC-Net [49], SAM-Net [37], GLU-Net [70], GOCor [69], RANSAC-Flow [62], and SCOT [47]. For the geometric matching task, we evaluate our method on Hpatches [3] and ETH3D [60] datasets, while for the semantic matching task, we evaluate our method on TSS [67] and PF-PASCAL [19] datasets. We also include the results of our variants for all the experiments and analyze the influence of different components of our method.

## 4.3. Geometric Matching Results

**HPatches.** We first evaluate our method on Hpatches [3] which consists of images with different views of the same scenes. Each sequence contains a source and five target images with different viewpoints with corresponding ground-truth flows. We use images of high resolutions ranging from  $450 \times 600$  to  $1,613 \times 1,210$ , whereas the down-scaled images ( $240 \times 240$ ) are also used for the evaluation as in [49]. For the evaluation metric, we use the Average Endpoint Error (AEE), computed by averaging the euclidean distance between the ground-truth and estimated flow, and Percentage of Correct Keypoints (PCK), computed as the ratio of estimated keypoints within the threshold from ground-truths to the total number of keypoints.

Table 1 summarizes the quantitative results, and Fig. 6 visualizes the qualitative results. We also provide an ablation study for RANSAC-Flow [62] to validate the effect of test-time optimization, which deteriorates performance as studied in [62]. However, DMP is competitive against the state-of-the-art learning-based methods [49, 70, 62] thanks to the feedback signals during the test-time optimization. Leveraging RANSAC [14] within our network allows outstanding performance, and with the networks pre-trained, DMP† even outperforms GOCor [69] by a large margin.

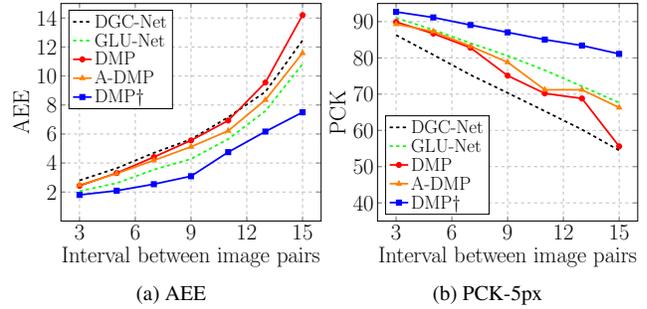


Figure 8. **Quantitative results on ETH3D [60] dataset.** AEE and PCK are computed on image pairs sampled at different intervals.

Methods	Backbone feature	TSS				PF-PA.
		FG	JO	PA	Avg.	
CNNGeo [55]	ResNet-101	90.3	76.4	56.5	74.3	69.5
PARN [27]	VGG-16	87.6	71.6	68.8	76.0	49.1
NC-Net [58]	ResNet-101	94.5	81.4	57.1	77.7	78.9
SAM-Net [37]	VGG-16	94.4	75.5	78.3	82.7	80.2
SCOT [47]	ResNet-101	95.3	81.3	57.7	78.1	88.8
GLU-Net [70]	VGG-16	93.2	73.3	71.1	79.2	79.7
GOCor [69]	VGG-16	95.0	78.9	81.3	85.1	-
<b>DMP</b>	VGG-16	93.9	77.2	72.1	81.7	81.5
	ResNet-101	94.6	78.1	74.2	82.3	85.3
<b>A-DMP</b>	VGG-16	94.1	80.1	74.3	82.8	82.4
	ResNet-101	94.4	81.4	75.2	86.7	85.4
<b>DMP†</b>	VGG-16	96.3	83.1	80.1	86.5	86.2
	ResNet-101	<b>96.7</b>	<b>89.2</b>	<b>82.3</b>	<b>89.4</b>	<b>89.1</b>

Table 2. **Quantitative evaluation on TSS [67] and PF-PASCAL (PF-PA.) [19] benchmark.** Higher PCK is better. *FG: FG3DCar, JO: JODS, PA: PASCAL datasets.*

**ETH3D.** We further evaluate our method on ETH3D [60] dataset. Unlike Hpatches [3], ETH3D consists of real 3D scenes, where the image transformations are not constrained to homography transformation, thus more challenging. We follow the protocol of [70], where we sample the image pairs at different intervals to evaluate on varying magnitude of geometric transformations. We evaluate on 7 intervals in total, each interval containing approximately 500 image pairs, and employ the standard evaluation metrics, AEE and PCK. Fig. 7 and Fig. 8 show quantitative and quantitative results. Our methods yield highly competitive results, and DMP†, significantly outperforming other methods, attains state-of-the-art performance.

## 4.4. Semantic Matching Results

**TSS.** For semantic matching, we evaluate our method on TSS [67] benchmark, which contains 3 groups (FG3DCar, JODS, and PASCAL) of 400 image pairs of 7 object categories, foreground masks and dense ground-truth flow. We employ the PCK to measure the precision. For TSS, we set the threshold as  $\alpha = 0.05$ . Following [70], we address the reflections by inferring the flow on both the source and flipped target and the source and original target. We then output the flow field with smaller mean horizontal magnitude. Table 2

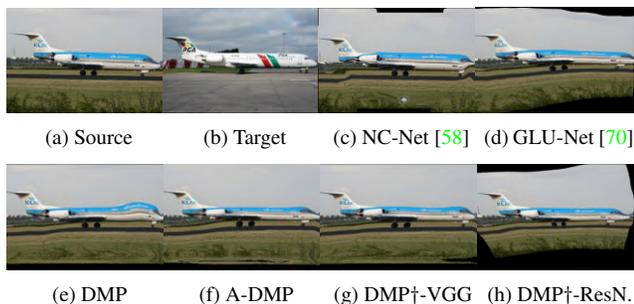


Figure 9. Qualitative results on PF-PASCAL [19] benchmark.

summarizes the PCK values. As shown in the results, our method produces highly competitive results compared to other learning-based methods. Interestingly, DMP<sup>†</sup> already attains state-of-the-art performance with VGG-16 [64] as feature backbone, but using the ResNet-101 [22] further boosts the performance. Qualitative results can be found in supplementary material.

**PF-PASCAL.** We further evaluate on PF-PASCAL [19] dataset, containing 1,351 image pairs over 20 object categories with keypoint annotations. Following the common practice [58, 20], we use PCK for the evaluation metric, and the results are reported with the PCK threshold  $\alpha = 0.1$ . Table 2 shows the quantitative results and Fig. 9 visualizes the qualitative comparisons. Similar to the experiments on TSS [67], our DMP provides the highest matching accuracy.

#### 4.5. Ablation Study

We show an ablation analysis on both architecture and loss function in our model. We measure the AEE over all the scenes of the HPatches [3], and each ablation experiment is conducted under the same experimental setting. We further report an analysis on computational complexity.

**Network architecture.** Table 3 shows the comparisons with different architectural components in our model. The baseline only consists backbone feature and inference module, optimized using contrastive loss without confidence threshold. Since optimization-based model is highly sensitive to initialization, exploiting residual correspondence module in matching networks dramatically helps the convergence and boosts performance in comparison to the baseline. We found that jointly exploiting both adaptation feature and residual correspondence showed apparent improvements. In addition, our model was not sensitive to the number of samples chosen for the loss computation.

**Loss function.** As shown in Table 3, the performance improvement by confidence threshold was apparent. In addition, Fig. 10 plots the accuracy as varying both the temperature ( $\tau$ ) and confidence ( $\varphi$ ). We found that the temperature controlling the sharpness of the softmax function is highly influential to the convergence and performance of

# samples ( $M$ )	Adaptation feature	Residual correspond.	Confidence ( $\varphi$ )	AEE
256	✗	✗	✗	33.1
256	✓	✗	✗	32.8
256	✗	✓	✗	18.6
256	✓	✓	✗	12.3
128	✓	✓	✓	9.84
256	✓	✓	✓	<b>9.69</b>
512	✓	✓	✓	9.67

Table 3. Ablation study of DMP.

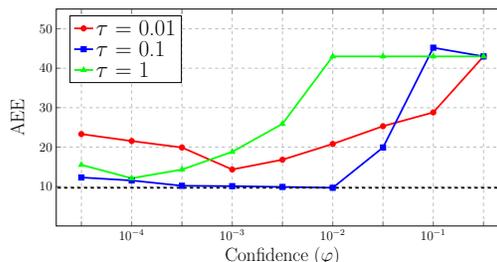


Figure 10. Ablation study on temperature and confidence.

our model. We thus found from extensive experiments that setting the temperature and confidence as 0.1 and 0.01, respectively, reports the best results. If the confidence is set too high, the loss signal does not occur and learning stops.

**Computational complexity.** In experiments, we found that given good initialization, DMP, including all the variants, converges within 100 iterations. However, given difficult image pairs to be matched, DMP struggles to find the good correspondences, and we thus let it iterate 2k times. The average runtime of 100 iterations of DMP and DMP<sup>†</sup> is 3-5 seconds, while A-DMP is 6-10 seconds on a single GPU Geforce GTX 2080 Ti. More details can be found in supplementary material. A natural next step, which we leave for future work, is to reduce the runtime for the optimization by advanced techniques, e.g., early-stopping.

## 5. Conclusion

In this paper, we have proven that, for the first time, a matching prior must not necessarily be learned from large training data and have introduced DMP that captures an image pair-specific matching prior by optimizing the untrained matching networks on a single pair of images. Tailored for such test-time optimization for matching, we have developed a residual matching network and a confidence-aware contrastive loss. Our experiments have shown that although our framework requires neither a large training data nor an intensive learning, it is competitive against the latest learning-based methods and even outperforms.

**Acknowledgments** This research was supported by the MSIT, Korea, under the ICT Creative Consilience program (IITP-2021-2020-0-01819) and Regional Strategic Industry Convergence Security Core Talent Training Business (IITP-2019-0-01343) supervised by the IITP and National Research Foundation of Korea (NRF-2021R1C1C1006897).

## References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *ICCV*, 2019. 2
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In *ACM SIGGRAPH*. 2005. 3
- [3] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, and Krystian Mikolajczyk. Hpatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *CVPR*, 2017. 6, 7, 8
- [4] Michael Bleyer, Christoph Rhemann, and Carsten Rother. Patchmatch stereo-stereo matching with slanted support windows. In *BMVC*, 2011. 3
- [5] Michael Calonder, Vincent Lepetit, Christoph Strecha, and Pascal Fua. Brief: Binary robust independent elementary features. In *ECCV*, 2010. 1
- [6] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *CVPR*, 2018. 1
- [7] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv:2002.05709*, 2020. 2
- [8] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv:2003.04297*, 2020. 2
- [9] Christopher B Choy, JunYoung Gwak, Silvio Savarese, and Manmohan Chandraker. Universal correspondence network. *NeurIPS*, 29:2414–2422, 2016. 2
- [10] Bo Dai and Dahua Lin. Contrastive learning for image captioning. In *NeurIPS*, 2017. 2
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 4, 6
- [12] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Deep image homography estimation. *arXiv:1606.03798*, 2016. 2, 6
- [13] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *ECCV*, 2014. 2
- [14] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 2, 7
- [15] Jean-Bastien Grill, Florian Strub, Florent Althé, 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. In *NeurIPS*, 2020. 2
- [16] Jinjin Gu, Yujun Shen, and Bolei Zhou. Image processing using multi-code gan prior. In *CVPR*, 2020. 2
- [17] Tanmay Gupta, Arash Vahdat, Gal Chechik, Xiaodong Yang, Jan Kautz, and Derek Hoiem. Contrastive learning for weakly supervised phrase grounding. *arXiv:2006.09920*, 2020. 2
- [18] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006. 2
- [19] Bumsub Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal flow. In *CVPR*, 2016. 1, 3, 7, 8
- [20] Kai Han, Rafael S Rezende, Bumsub Ham, Kwan-Yee K Wong, Minsu Cho, Cordelia Schmid, and Jean Ponce. Ssnet: Learning semantic correspondence. In *ICCV*, 2017. 8
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 2
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 4, 6, 8
- [23] Shuaiyi Huang, Qiuyue Wang, Songyang Zhang, Shipeng Yan, and Xuming He. Dynamic context correspondence network for semantic alignment. In *ICCV*, 2019. 2
- [24] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, 2017. 1
- [25] Allan Jabri, Andrew Owens, and Alexei Efros. Space-time correspondence as a contrastive random walk. *NeurIPS*, 2020. 2
- [26] Ali Jahanian, Lucy Chai, and Phillip Isola. On the “steerability” of generative adversarial networks. In *ICLR*, 2020. 2
- [27] Sangryul Jeon, Seungryong Kim, Dongbo Min, and Kwanghoon Sohn. Parn: Pyramidal affine regression networks for dense semantic correspondence. In *ECCV*, 2018. 2, 5, 7
- [28] Sangryul Jeon, Dongbo Min, Seungryong Kim, and Kwanghoon Sohn. Joint learning of semantic alignment and object landmark detection. In *ICCV*, 2019. 1
- [29] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 2
- [30] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 2
- [31] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 2
- [32] 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 *ICCV*, 2017. 1
- [33] 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 *ICCV*, 2017. 4
- [34] Jaechul Kim, Ce Liu, Fei Sha, and Kristen Grauman. Deformable spatial pyramid matching for fast dense correspondences. In *CVPR*, 2013. 2
- [35] Seungryong Kim, Stephen Lin, Sang Ryul Jeon, Dongbo Min, and Kwanghoon Sohn. Recurrent transformer networks for semantic correspondence. In *NeurIPS*, 2018. 1

- [36] Seungryong Kim, Dongbo Min, Bumsub Ham, Sangryul Jeon, Stephen Lin, and Kwanghoon Sohn. Fcss: Fully convolutional self-similarity for dense semantic correspondence. In *CVPR*, 2017. 2
- [37] Seungryong Kim, Dongbo Min, Somi Jeong, Sunok Kim, Sangryul Jeon, and Kwanghoon Sohn. Semantic attribute matching networks. In *CVPR*, 2019. 1, 2, 4, 5, 7
- [38] Seungryong Kim, Dongbo Min, Stephen Lin, and Kwanghoon Sohn. Dctm: Discrete-continuous transformation matching for semantic flow. In *ICCV*, 2017. 1, 2, 3, 4, 5
- [39] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014. 3, 4, 7
- [40] Zihang Lai, Erika Lu, and Weidi Xie. Mast: A memory-augmented self-supervised tracker. In *CVPR*, 2020. 2
- [41] Junghyup Lee, Dohyung Kim, Jean Ponce, and Bumsub Ham. Sfnets: Learning object-aware semantic correspondence. In *CVPR*, 2019. 4
- [42] Shuda Li, Kai Han, Theo W Costain, Henry Howard-Jenkins, and Victor Prisacariu. Correspondence networks with adaptive neighbourhood consensus. In *CVPR*, 2020. 2
- [43] Xinghui Li, Kai Han, Shuda Li, and Victor Adrian Prisacariu. Dual-resolution correspondence networks. *arXiv:2006.08844*, 2020. 2
- [44] Jing Liao, Yuan Yao, Lu Yuan, Gang Hua, and Sing Bing Kang. Visual attribute transfer through deep image analogy. *arXiv:1705.01088*, 2017. 1, 3
- [45] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):978–994, 2010. 1, 2, 3, 4, 5
- [46] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. *arXiv preprint arXiv:1908.03265*, 2019. 7
- [47] Yanbin Liu, Linchao Zhu, Makoto Yamada, and Yi Yang. Semantic correspondence as an optimal transport problem. In *CVPR*, 2020. 7
- [48] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 1
- [49] Iaroslav Melekhov, Aleksei Tiulpin, Torsten Sattler, Marc Pollefeys, Esa Rahtu, and Juho Kannala. Dgc-net: Dense geometric correspondence network. In *WACV*, 2019. 1, 2, 3, 4, 5, 6, 7
- [50] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Learning to compose hypercolumns for visual correspondence. In *ECCV*, 2020. 2
- [51] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2
- [52] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. In *ECCV*, 2020. 2
- [53] Pedro O Pinheiro, Amjad Almahairi, Ryan Y Benmaleck, Florian Golemo, and Aaron Courville. Unsupervised learning of dense visual representations. *arXiv:2011.05499*, 2020. 2
- [54] Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*, 2019. 7
- [55] Ignacio Rocco, Relja Arandjelovic, and Josef Sivic. Convolutional neural network architecture for geometric matching. In *CVPR*, 2017. 1, 2, 3, 6, 7
- [56] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. End-to-end weakly-supervised semantic alignment. In *CVPR*, 2018. 2, 5
- [57] Ignacio Rocco, Relja Arandjelović, and Josef Sivic. Efficient neighbourhood consensus networks via submanifold sparse convolutions. In *ECCV*, 2020. 2
- [58] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. In *NeurIPS*, 2018. 2, 4, 7, 8
- [59] Tanner Schmidt, Richard Newcombe, and Dieter Fox. Self-supervised visual descriptor learning for dense correspondence. *IEEE Robotics and Automation Letters*, 2(2):420–427, 2016. 2
- [60] 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 *CVPR*, 2017. 7
- [61] Tianwei Shen, Lei Zhou, Zixin Luo, Yao Yao, Shiwei Li, Jiahui Zhang, Tian Fang, and Long Quan. Self-supervised learning of depth and motion under photometric inconsistency. In *ICCV Workshops*, 2019. 1, 2, 3
- [62] Xi Shen, François Darmon, Alexei A Efros, and Mathieu Aubry. Ransac-flow: generic two-stage image alignment. *arXiv:2004.01526*, 2020. 2, 3, 4, 6, 7
- [63] Assaf Shocher, Nadav Cohen, and Michal Irani. “zero-shot” super-resolution using deep internal learning. In *CVPR*, 2018. 2
- [64] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014. 1, 4, 6, 8
- [65] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018. 1, 2, 5
- [66] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *CVPR*, 2018. 1, 2, 3, 5
- [67] Tatsunori Tanai, Sudipta N Sinha, and Yoichi Sato. Joint recovery of dense correspondence and cosegmentation in two images. In *CVPR*, 2016. 1, 2, 3, 7, 8
- [68] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020. 2, 3
- [69] Prune Truong, Martin Danelljan, Luc V Gool, and Radu Timofte. Gocor: Bringing globally optimized correspondence volumes into your neural network. In *NeurIPS*, 2020. 2, 6, 7
- [70] Prune Truong, Martin Danelljan, and Radu Timofte. Glunet: Global-local universal network for dense flow and correspondences. In *CVPR*, 2020. 1, 2, 3, 4, 5, 6, 7, 8

- [71] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *CVPR*, 2018. 2
- [72] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. *arXiv:2011.09157*, 2020. 2
- [73] Olivia Wiles, Sebastien Ehrhardt, and Andrew Zisserman. D2d: Learning to find good correspondences for image matching and manipulation. *arXiv:2007.08480*, 2020. 2
- [74] Li Xu, Jimmy S Ren, Ce Liu, and Jiaya Jia. Deep convolutional neural network for image deconvolution. In *NeurIPS*, 2014. 2
- [75] Gengshan Yang and Deva Ramanan. Volumetric correspondence networks for optical flow. In *NeurIPS*, 2019. 4
- [76] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. Lift: Learned invariant feature transform. In *ECCV*, 2016. 1
- [77] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *ECCV*, 2020. 2