

## Matching Semantically Similar Non-Identical Objects

Yusuke Marumo<sup>1†</sup> Kazuhiko Kawamoto<sup>2†</sup> Satomi Tanaka<sup>3‡</sup> Shigenobu Hirano<sup>4‡</sup> Hiroshi Kera<sup>5†\*</sup>

<sup>†</sup>Chiba University <sup>‡</sup>Ricoh Company

<sup>1</sup>yusuke.marumo@chiba-u.jp, <sup>2</sup>kawa@faculty.chiba-u.jp, <sup>3</sup>satomi.st.tanaka@jp.ricoh.com

<sup>4</sup>shigenobu.hirano@jp.ricoh.com, <sup>5</sup>kera@chiba-u.jp

### Abstract

*Not identical but similar objects are ubiquitous in our world, ranging from four-legged animals such as dogs and cats to cars of different models and flowers of various colors. This study addresses a novel task of matching such non-identical objects at the pixel level. We propose a weighting scheme of descriptors, Semantic Enhancement Weighting (SEW), that incorporates semantic information from object detectors into existing sparse feature matching methods, extending their targets from identical objects captured from different perspectives to semantically similar objects. The experiments show successful matching between non-identical objects in various cases, including in-class design variations, class discrepancy, and domain shifts (e.g., photo vs. drawing and image corruptions). The code is available at <https://github.com/Circ-Leaf/NIOM>.*

### 1. Introduction

We humans are able to find fine-grained and robust visual correspondences between two objects even when they are not necessarily identical or when they are captured in different situations. Thanks to this ability, for example, we can find out our missing cats in a gloomy backstreet by remembering them basking in the sun or assembling components to build a shelf based on their perspective drawing. In computer vision terms, this needs fine-grained matching of feature points between two objects, even with class discrepancy and domain shift (e.g., image corruptions).

The abovementioned matching has been partially addressed in the literature. In the feature matching [24, 37, 39, 69, 79], objects to be matched are assumed identical and captured from different perspectives. This assumption also enables supervised training using dense correspondence obtained from 3D reconstruction from image pairs. Recently, Transformer-based dense matching methods have shown great performance [8, 18, 30, 31, 64, 79, 95, 101, 103], and

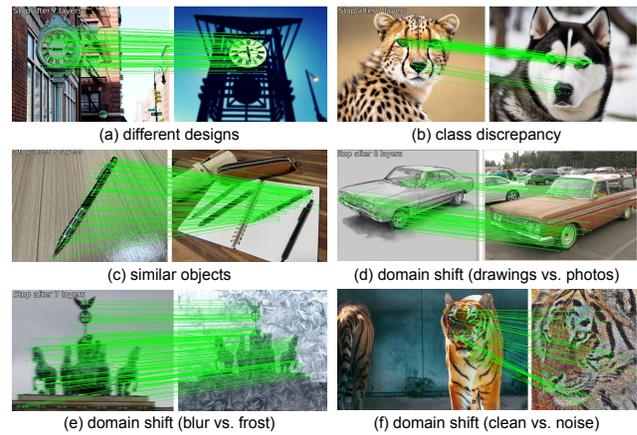


Figure 1. Non-identical object matching between various objects and image styles is achieved by our plug-and-play method with SuperPoint [16] keypoint detector and LightGlue [37] matcher.

there are also light-weight sparse matching methods, which run in real time with a slight cost in accuracy [24, 29, 43, 54, 65, 68, 69, 92]. In contrast, semantic correspondence models [34, 45, 62, 98, 99] focus on matching semantically similar parts between objects (e.g., tires of two bikes). They are trained on datasets with sparse part-to-part correspondence annotation and thus do not offer fine-grained matching.

In this study, we generalize feature matching and semantic correspondence to *non-identical object matching*, where objects to be matched are not necessarily identical (see Fig. 1). For example, the object pair can be street clocks of different designs or four-legged animals of different species. Further, the non-identity naturally requires robustness against image corruptions and domain shifts (e.g., photo vs. drawing) because two objects may be captured in different environments and domains. A crucial challenge of achieving non-identical object matching is that we cannot straightforwardly resort to the supervised learning framework because of the ambiguity of non-identity; there is no ground-truth matching between non-identical objects.

We propose a plug-and-play module, Semantic Enhance-

\*Corresponding Author.

ment Weighting (SEW), that extends off-the-shelf matchers to non-identical object matching. The sparse matching pipeline first detects keypoints in a pair of images and feeds their descriptors into a matcher. SEW is inserted before the matcher and weights descriptors using an object detector and Grad-CAM [73]. Due to task definition, traditional matching relies on low-level features (e.g., color, edges, texture), which are prone to ambiguity and noise (see Fig. 2). While these features are still vital for fine-grained matching, our high-level feature weighting improves robustness. Further, to handle the case with multiple objects in images, we propose Non-visual Object Pairing, which determines the pairs of similar objects to be matched between images and eliminates mismatching between irrelevant objects.

Our experiments demonstrate that our approach significantly boosts the robustness of sparse matching methods, including LightGlue [37] and GlueStick [54], in two tasks: non-identical object matching and robust image matching. For the former, we evaluate the matching results with various image pairs (a) of the same class, (b) with class discrepancy, (c) with domain shift, and (d) of the same appearance, and further provide a quantitative evaluation with a new annotation-free metric, *Triangular Matching Consistency* (TMC). While sparse matchers alone yield scattered mismatches, our method achieves fine-grained and consistent alignment, as reflected in higher TMC scores. For the latter, we evaluate the robustness of matching on the MegaDepth-1500 [79] dataset under common corruptions [27] in terms of the AUC score of standard relative pose estimation. Our method outperforms the state-of-the-art sparse matcher under various types of image corruptions. The average AUC of our method even exceeds that of dense matchers, LoFTR [79] and Efficient LoFTR [89], which pursue matching quality at the cost of speed.

To summarize, our contributions are as follows:

- We tackle non-identical object matching, a novel task of matching similar but not necessarily identical objects in images, even under image corruptions and/or across different domains, which can be considered a generalization of feature matching and semantic correspondence.
- We propose Semantic Enhancement Weighting (SEW), which enhances low-level features in a plug-and-play manner using semantic information and successfully extends various matchers to non-identical object matching.
- Our experiments show that under challenging conditions such as class discrepancy, image corruptions, and domain shifts, our method yields fine-grained and consistent correspondences, validated by TMC and pose estimation.

## 2. Related Work

Image matching [24, 37, 69, 79] is a fundamental task in computer vision with various applications [2, 14, 49, 51, 60, 71, 78, 80, 88, 93] that involves finding the fine-grained cor-

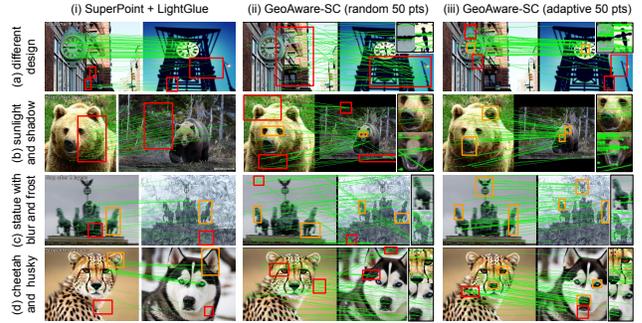


Figure 2. State-of-the-art models of sparse matching (LightGlue [37]) and semantic correspondence (GeoAware-SC [98]). The latter requires reference points in one image and finding their correspondence in the target image. To this end, 50 reference points are sampled randomly or based on the GradCAM [73] heatmap of YOLOv7 [87] (cf. Sec. D in the supplementary material). Orange boxes show close matches, which are correct at the part level, but not at the pixel level. This was typical for Geo-Aware-SC (e.g., matching between the nose tip of the cheetah and the nostrils of the husky in (d)). Further, Geo-Aware-SC was about ten times slower than LightGlue. (See Tab. B in the supplementary material).

respondence between two images. Many image matching methods have two stages. A keypoint detector first extracts keypoints with their positions and local feature descriptors, and then feature matchers find matches between them.

**Keypoint detectors.** Classical detectors such as SIFT [40] and ORB [67] rely on hand-crafted features. Recent learning-based methods, including SuperPoint [16] and DISK [84], achieve higher robustness, with lightweight [21] and unsupervised variants [25] also proposed.

**Feature matchers.** Feature matchers [90, 96] find a fine-grained correspondence between images. The feature matchers are categorized into two types based on the density of matching points: sparse matchers and dense matchers. Sparse matchers (e.g., SuperGlue [69], which combines Transformer attention with optimal transport; LightGlue [37], which introduces a hierarchical matching structure for real-time accuracy; GlueStick [54], which exploits line segments as structural features) use limited keypoints for efficiency. In contrast, dense matchers (e.g., LoFTR [79], DKM [17], ROMA [18]) pursue matching accuracy rather than speed by matching all pixels.

**Semantic correspondence models.** Semantic Correspondence [3, 9, 23, 45] is a task of identifying a semantically similar point in the target image from the given point in the source image, with benchmarks such as SPair-71k [48] and PF-PASCAL [22]. Recent models include GeoAware-SC [98], which integrates geometric cues with pre-trained features, and unsupervised methods without an-

notations [26]. It is also worth noting that semantic correspondence models run significantly slower than sparse matching methods because feature extraction depends on large generative models, such as Stable Diffusion [66]. For example, it runs roughly 10 times slower in Fig. 2.

This study addresses non-identical object matching, unifying the fine-grained alignment of feature matching and the high-level matching of semantic correspondence. This makes supervised training technically challenging but also offers unique applications (cf. Sec. 3.2). In another view, we aim for matching that runs in real-time and is robust to object identity variations, corruptions, and domain shifts. This covers broader robustness than prior studies [5, 63, 91], including DarkFeat [24] and SAM [43] for lighting changes. See Sec. H in the supplementary for detailed related work.

### 3. Non-Identical Object Matching

We introduce non-identical object matching, matching between unnecessarily identical objects in two images. For example, a dog and a cat are not identical but still have similar structures, such as two eyes, four legs, and one tail. The same mass-produced products, such as cars of the same model, have the same appearance but are still not identical (e.g., owned by different people) and can appear in images with largely different backgrounds and lighting conditions.

#### 3.1. Task formulation

We introduce a more formal setup of non-identical object matching. To this end, we first present fine-level image matching (i.e., dense pixel-wise matching, as opposed to coarse-level image matching tasks such as semantic correspondence) as a general task and then discuss feature matching and non-identical object matching as special cases.

**Problem 1 (fine-level image matching)** *Let  $\mathcal{X}$  be the image domain. Given two images  $\mathbf{x}_A, \mathbf{x}_B \in \mathcal{X}$ , let  $D_A = \{(\mathbf{p}_i^A, \mathbf{d}_i^A)\}_{i=1}^{n_A}$  be  $n_A$  keypoints in image  $\mathbf{x}_A$ , where  $\mathbf{p}_i^A \in [0, 1]^2$  and  $\mathbf{d}_i^A \in [0, 1]^d$  are its position and descriptor of the  $i$ -th keypoint. We define  $D_B = \{(\mathbf{p}_i^B, \mathbf{d}_i^B)\}_{i=1}^{n_B}$  for image  $\mathbf{x}_B$  similarly. Let  $\mathcal{L}(\cdot)$  be a (conceptual) matching loss function. The task can be formalized as follows.*

$$\min \mathcal{L}(\{m_{kl}\}_{k=1, \dots, n_A, l=1, \dots, n_B}) \quad (1)$$

$$\text{s.t. } m_{kl} \in [0, 1], \forall k \in \{1, \dots, n_A\}, l \in \{1, \dots, n_B\} \quad (2)$$

$$\sum_{l=1}^{n_A} m_{kl} \leq 1, \forall k \in \{1, \dots, n_A\} \quad (3)$$

$$\sum_{k=1}^{n_B} m_{kl} \leq 1, \forall l \in \{1, \dots, n_B\}. \quad (4)$$

The matching coefficient  $m_{kl} = 1$  indicates that the  $k$ -th keypoint in image  $\mathbf{x}_A$  and the  $l$ -th keypoint in image  $\mathbf{x}_B$



Figure 3. Non-identical object warping. Homography is estimated from fine-grained matching of city clocks with different styles.

is matched and otherwise  $m_{kl} = 0$ . Each keypoint in one image has at most one matched point in the other image. This condition is implemented by the last two constraints.

In the feature matching task, where an identical object is captured from different angles in two images, matching loss  $\mathcal{L}(\{m_{kl}\}_{k,l})$  measures the discrepancy between a camera pose computed the matching  $\{m_{kl}\}_{k,l}$  and the ground truth.

Our task, non-identical object matching covers a broader class of matching. Given two images  $\mathbf{x}_A, \mathbf{x}_B \in \mathcal{X}$  that contain (potentially non-identical) objects to be matched, image  $\mathbf{x}_I$  can be characterized by several attributes  $(o_I, y_I, P_I, D_I)$  for  $I \in \{A, B\}$ . If  $o_A = o_B$ , then the objects to be matched are identical. If  $o_A \equiv o_B$ , they are the same in appearance but not necessarily identical (e.g., two cars of the same model). If  $y_A = y_B$ , they belong to the same class (e.g., the car class). If  $P_A = P_B$ , then they are captured from the same camera pose. If  $D_A = D_B$ , they are captured in the same domain (e.g., the drawing domain).

The classical feature matching corresponds to the case with  $o_A = o_B$  (consequently,  $o_A \equiv o_B$  and  $y_A = y_B$ ),  $P_A \neq P_B$ , and  $D_A = D_B$ . In words, it considers identical objects captured from different camera poses in the same domain. In contrast, non-identical object matching encompasses  $y_A \neq y_B$  (consequently,  $o_A \neq o_B$  and  $o_A \not\equiv o_B$ ),  $P_A \neq P_B$ , and  $D_A \neq D_B$ . Note that we practically assume  $o_A$  and  $o_B$  to be *similar* in some sense. Below, we highlight two important special cases, both assuming  $P_A \neq P_B$ .

**Class discrepancy ( $y_A \neq y_B$ ).** Four-legged animals, such as tigers and cats, belong to different classes but have similar structures (two eyes, four legs, and a tail). Despite differences in color, texture, and shape, we humans can still find part-by-part matches. However, it is not trivial to design matching methods to find fine-grained correspondence while being robust to detail differences.

**Domain shift ( $o_A \equiv o_B, D_A \neq D_B$ ).** We humans can know the flower in the painting on the wall depicts that in the garden. This is a robust fine-grained matching of an identical object in different domains. Typical domain shifts include photos vs. drawings and clean vs. corrupted images. For instance, if  $\mathbf{x}_A$  is a photo and  $\mathbf{x}_B$  is a painting, the latter lacks colors, textures, and shadows, causing its descriptors to differ significantly. Without sufficient robustness, this discrepancy leads to missing correspondences.

### 3.2. Distinction from prior tasks and applications

Non-identical object matching relates to feature matching (i.e., fine-level perspective alignment) and semantic correspondence (i.e., coarse-level part correspondences between similar objects). However, achieving fine-grained and semantic alignment of different objects simultaneously is technically non-trivial and has interesting applications.

Non-identical object matching is challenging because it does not readily fit into supervised training. There is no available dataset with fine-grained human annotation of matching between distinct objects, and the construction would be laborious. In contrast, feature matching relies on supervised training with ground-truth homography from 3D reconstruction, while semantic correspondence relies on sparse part-level annotations. Yet its standard metric, PCK [22, 48], evaluates only coarse part matches and is therefore inadequate for fine-grained alignment.

The applications of non-identical object matching are unique and cannot be attained by feature matching and semantic correspondence. A fine-grained matching with robustness to class discrepancy allows transferring fine-grained landmark annotations. For example, computer vision for animals works on datasets targeting a single animal species, such as facial datasets of dogs [46], cats [47], and cattle [12]. Non-identical object matching allows transferring their annotations to other animals. Its robustness to domain shift can also compare the sketch or blueprint of a product to a real one, which helps workers assemble parts and build it up. Figure 3 illustrates a simple example, *non-identical object warping*, that is not covered by prior tasks.

## 4. Proposed Method

We propose a plug-and-play module, Semantic Enhancement Weighting (SEW), that can be inserted into various sparse matchers, such as SuperGlue [69], LightGlue [37], and more [7, 54]. Sparse matchers consist of a pair of a keypoint detector and a feature matcher. Our module processes the keypoints before feeding them to the feature matcher. Particularly, given a pair of images  $\mathbf{x}_A, \mathbf{x}_B \in \mathcal{X}$ , The keypoint detector outputs collections of keypoints  $D_A = \{(\mathbf{p}_i^A, \mathbf{d}_i^A)\}_i, D_B = \{(\mathbf{p}_i^B, \mathbf{d}_i^B)\}_i$ . The proposed module then computes weighted descriptors  $\tilde{\mathbf{d}}_i^A, \tilde{\mathbf{d}}_i^B$  for each  $\mathbf{d}_i^A, \mathbf{d}_i^B$ , respectively. Lastly, the feature matcher performs matching using  $\tilde{D}_A = \{(\mathbf{p}_i^A, \tilde{\mathbf{d}}_i^A)\}_i, \tilde{D}_B = \{(\mathbf{p}_i^B, \tilde{\mathbf{d}}_i^B)\}_i$ .

As will be shown in Sec. 5, this simple pipeline works surprisingly well and adapts pretrained sparse matching models without any training, addressing the technical challenge in non-identical object matching discussed above.

### 4.1. Semantic Enhancement Weighting

The proposed weighting module uses an object detector and a visual explanation model. This study used YOLOv7 [87]

and Grad-CAM [73]. Let  $\mathbf{x} \in \mathcal{X}$  be an image, and let  $D = \{(\mathbf{p}_i, \mathbf{d}_i)\}_{i=1}^n$  be the collection of  $n$  keypoints of  $\mathbf{x}$ .

First, the object detector is applied to  $\mathbf{x}$  and detects the objects with their classes. For each of the detected objects (say, the  $k$ -th object), the visual explanation model gives pixel-wise class-activation scores or a heatmap  $H_k$ . Let  $H_{k, \mathbf{p}_i} \in [0, 1]$  be the heatmap score of  $H_k$  at  $\mathbf{p}_i$ . Then, the scores at  $\mathbf{p}_i$  are aggregated as  $H_{\mathbf{p}_i} := \max_k H_{k, \mathbf{p}_i}$ <sup>1</sup>. Finally, the weighted descriptor  $\tilde{\mathbf{d}}_i$  is computed by

$$\tilde{\mathbf{d}}_i = \frac{1 + H_{\mathbf{p}_i}}{\max_{j \in \{1, \dots, n\}} (1 + H_{\mathbf{p}_j})} \cdot \mathbf{d}_i \in [0.5, 1]^d. \quad (5)$$

We empirically found that the addition of base value 1 and the normalization by the maximum value is important. For example, the simplest approach  $\tilde{\mathbf{d}}_i = H_{\mathbf{p}_i} \mathbf{d}_i \in [0, 1]^d$  ends up with a scarcity of matching because of the heatmap scores on the keypoints on background tend to be too low, i.e.,  $H_{\mathbf{p}_i} \approx 0$ . Although such keypoints will not likely appear in the matching results, they seem to help the keypoints on objects to be matched. If one instead uses  $\tilde{\mathbf{d}}_i = (1 + H_{\mathbf{p}_i}) \mathbf{d}_i \in [1, 2]^d$ , this also affects negatively. We consider that this is because of the softmax operations used in the subsequent feature matching. We next discuss how the weighting of descriptors impacts the matching.

### 4.2. Effect of weights in feature matcher

We now illustrate how our weighting module affects the subsequent feature matching. We take LightGlue [37] as an example, but a similar discussion holds for other feature matchers. For simplicity, let  $\alpha_i$  be the weight in Eq. (5), i.e.,  $\tilde{\mathbf{d}}_i = \alpha_i \mathbf{d}_i$ . The matcher takes as input the keypoints  $\tilde{D}_A = \{(\mathbf{p}_i^A, \tilde{\mathbf{d}}_i^A)\}_i, \tilde{D}_B = \{(\mathbf{p}_i^B, \tilde{\mathbf{d}}_i^B)\}_i$ . The vectors  $\tilde{\mathbf{d}}_i^A = \alpha_i^A \mathbf{d}_i^A, \tilde{\mathbf{d}}_i^B = \alpha_i^B \mathbf{d}_i^B$  are the  $i$ -th weighted descriptors of image  $\mathbf{x}_A, \mathbf{x}_B$ , respectively. Below, symbols with and without  $\tilde{\cdot}$  present those with and without weighting.

In the feature matcher, the descriptors are first converted to key and query vectors using matrices  $W_q, W_k \in \mathbb{R}^{d \times d}$ .

$$\tilde{\mathbf{q}}_i^I = W_q \tilde{\mathbf{d}}_i^I = \alpha_i W_q \mathbf{d}_i^I = \alpha_i \mathbf{q}_i^I \quad (6)$$

$$\tilde{\mathbf{k}}_i^I = W_k \tilde{\mathbf{d}}_i^I = \alpha_i W_k \mathbf{d}_i^I = \alpha_i \mathbf{k}_i^I, \quad (7)$$

where  $I \in \{A, B\}$ . Then, the self-attention score  $a_{ij}^I$  between  $i$ -th and  $j$ -th keypoints of image  $\mathbf{x}_I$  is calculated as

$$\tilde{a}_{ij}^I = \tilde{\mathbf{q}}_i^{I\top} \mathcal{R} (\mathbf{p}_j^I - \mathbf{p}_i^I) \tilde{\mathbf{k}}_j^I \quad (8)$$

$$= (\alpha_i^I \alpha_j^I) \cdot (\mathbf{q}_i^{I\top} \mathcal{R} (\mathbf{p}_j^I - \mathbf{p}_i^I) \mathbf{k}_j^I) \quad (9)$$

$$= (\alpha_i^I \alpha_j^I) \cdot a_{ij}^I, \quad (10)$$

<sup>1</sup>The max operation outperformed other candidates including mean, sigmoid, and softmax operations in our preliminary experiments.

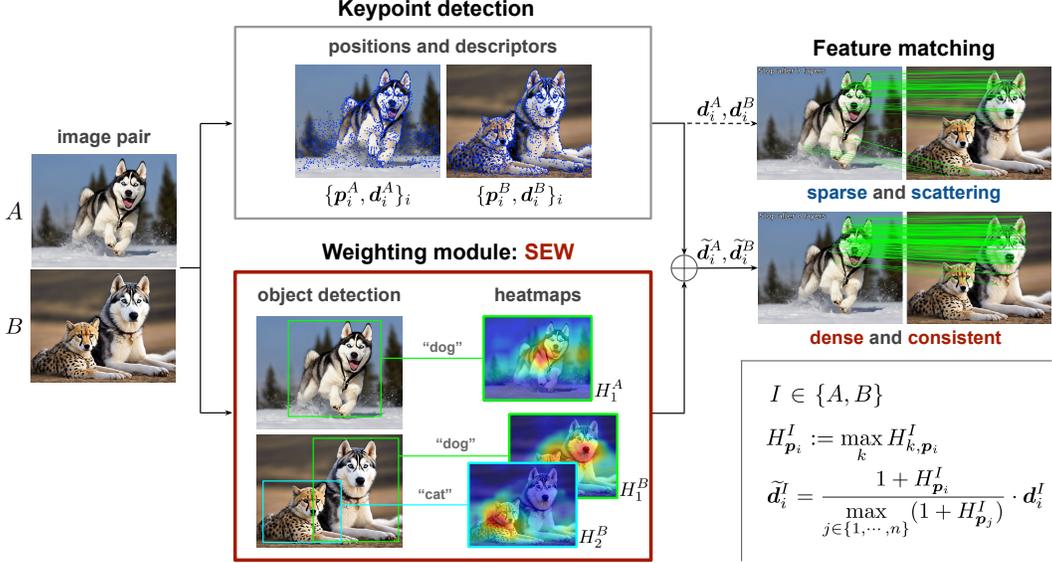


Figure 4. Pipeline of the matching. The keypoint detection and feature matching are done by off-the-shelf models. The proposed plug-and-play module, SEW, computes the heatmap scores of objects and weights the descriptors with this semantic information.

where  $\mathcal{R} : \mathbb{R}^2 \rightarrow \mathbb{R}^{d \times d}$  denotes the rotary-encoding [76]. Therefore, our weighting by heatmap scores leads to the weighting of attention scores. The self-attention scores reflect the internal relationship of objects and background, thereby extracting better feature representations in the subsequent steps of the feature matcher. Roughly speaking, noting that  $\alpha_i^I \in [0.5, 1.0]$  for all  $i$  and  $I \in \{A, B\}$ , we have  $\alpha_i^I \alpha_j^I \approx 1$  between the keypoints on objects,  $\alpha_i^I \alpha_j^I \approx 0.5$  between one on object and the other on background, and  $\alpha_i^I \alpha_j^I \approx 0.25$  for between those on the background. This is reasonable because we want nice feature representations for object keypoints for nice matching. Nevertheless, such nice representations should be obtained by taking into account the background conditions, e.g., to compensate for the overall lighting conditions. If one uses  $\alpha_i^I = H_{p_i}$  as discussed earlier, i.e.,  $H_{p_i} \approx 0$  for background keypoints, the feature representations cannot take into account the background.

Similarly, our weighting affects the cross-attention score  $\alpha_{ij}^{AB}$  between the  $i$ -th keypoint of image  $x_A$  and the  $j$ -th keypoint of image  $x_B$ , computed from the key vectors.

$$\tilde{\alpha}_{ij}^{AB} = \tilde{\mathbf{k}}_i^{A\top} \tilde{\mathbf{k}}_j^B = (\alpha_i^A \alpha_j^B) \cdot \alpha_{ij}^{AB}. \quad (11)$$

As discussed earlier,  $\alpha_i^A \alpha_j^B$  is large between keypoints of the objects, encouraging matching between objects.

### 4.3. Non-visual Object Pairing

We next propose Non-visual Object Pairing for the case where images contain multiple objects. In such a case, it is essential to determine which object in one image should be matched with another object in the other image, and the

weighting module should exploit the heatmaps of them and exclude those of irrelevant objects. Otherwise, the local spurious similarity of objects causes scattered and inconsistent matching; see Fig. 5 (bottom right), where a nose of a dog and an ear of a cat match.

A pair of *similar* objects for matching can be determined using visual or non-visual features. For the former, we can use the feature vectors extracted from an intermediate layer of the object detector. This can be done simultaneously during the heatmap computation process. For the latter, text embedding (e.g., by the CLIP text encoder) of object labels may be a reasonable choice. The text embedding vectors can be pre-computed for all the classes if one pursues fast matching. The visual and non-visual features are complementary, and the choice should depend on the use case. For example, given (*lesser panda*, *raccoon*, *giant panda*), the visually similar pair is (*lesser panda*, *raccoon*), while (*lesser panda*, *giant panda*) is more conceptually similar.

We here provide concrete procedures in the Non-visual Object Pairing based on non-visual features. Suppose that an image  $x_A$  contains  $n_A$  objects, and let  $Y_A = \{y_i^A\}_{i=1}^{n_A}$  be their class labels. We define  $Y_B = \{y_i^B\}_{i=1}^{n_B}$  similarly for the other image  $x_B$ . Let  $\mathcal{T}$  be a pretrained text encoder that maps a label  $y$  to an embedding vector  $\mathcal{T}(y)$ . Particularly, we used the CLIP text encoder [59]. The adapter determines the object pair  $(i, j)$  to be matched based on the cosine similarity between the embedding vectors as follows.

$$(y_i^A, y_j^B) = \operatorname{argmax}_{y^A \in Y_A, y^B \in Y_B} \frac{\mathcal{T}(y^A)^\top \mathcal{T}(y^B)}{\|\mathcal{T}(y^A)\| \|\mathcal{T}(y^B)\|}, \quad (12)$$

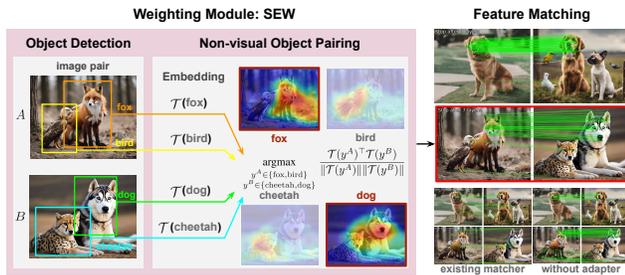


Figure 5. Non-visual Object Pairing converts class labels from an object detector into embeddings using the CLIP [59] text encoder and compares them with cosine similarity to identify semantically similar object pairs. The GradCAM [73] heatmap of object detector YOLOv7 [87] generated for the chosen object pair is used for our weighting method, enhancing matching with multiple objects.

where  $\|\cdot\|$  denotes the  $L_2$  norm of vector. Then, the weighting module only uses the Grad-CAM heatmaps of the  $i$ -th object in  $x_A$  and the  $j$ -th object in  $x_B$ . The process is almost the same when visual features are used; the text embedding vectors are replaced by the feature vectors from the intermediate layer of the object detector. Finally, the weighting module uses the corresponding heatmaps for weighting the descriptors. Figure 5 shows this pipeline.

## 5. Experiments

In this section, we demonstrate the effectiveness of the proposed method in non-identical object matching tasks. Section 5.1 presents various qualitative comparisons from state-of-the-art sparse feature matching methods. Section 5.2 then introduces *Triangular Matching Consistency* (TMC), a new annotation-free metric that allows us to quantitatively evaluate non-identical object matching itself. Section 5.3 presents quantitative evaluations. In the latter experiment, as there are no available datasets providing fine-grained matching between non-identical objects, we examine the matching results between identical objects under 15 types of corruptions, which measures the robustness to domain shifts. We emphasize that while semantic correspondence is a conceptually similar task to ours, it does not target fine-grained matching, and thus, our experiment setup (i.e., datasets, baseline methods, and metrics) is based on feature matching literature.

**Setup.** The sparse matching pipeline involves a keypoint detector and a feature matcher. For the former, we adopted pre-trained SuperPoint [16], extracting 2,048 keypoints. For the latter, the extracted keypoints are matched using either pretrained SGMNet [7], GlueStick [54], and the state-of-the-art method, LightGlue [37], with and without the proposed method. We additionally tested a detector-free dense matcher, LoFTR [79] and Efficient LoFTR [89]. We downloaded the pretrained models from their official repos-

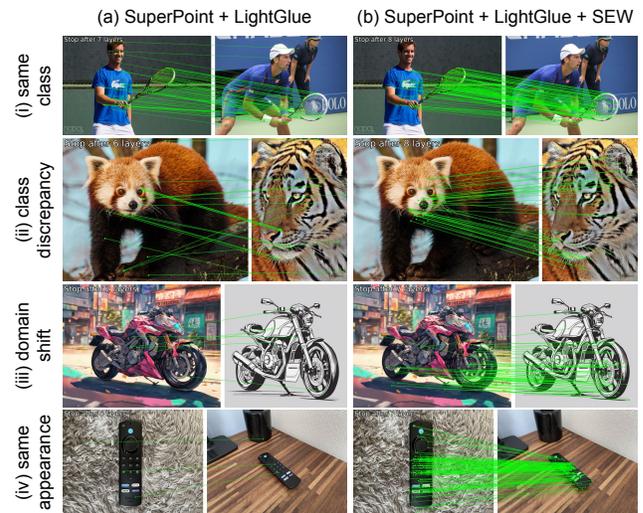


Figure 6. Typical categories of non-identical object matching. (a) A combination of SuperPoint [16] and LightGlue [37] only finds a small number of correspondences, and many of them are incorrect. (b) Our method SEW significantly improves the matching.

itories (details provided in the supplementary material). All experiments used NVIDIA GeForce RTX 3090 GPU and 24GB memory.

### 5.1. Matching Non-Identical Objects

Here, we perform the matching of non-identical objects and evaluate them qualitatively.

**Datasets.** We collected various images with objects from the COCO [36] and ImageNet [15] datasets. To prepare the drawing version of several images, we used the DreamStudio<sup>2</sup>, which is backed by the Stable Diffusion [66].

**Robust matching across non-identical objects.** We examined four cases: (i) same class (i.e.,  $y_A = y_B, D_A = D_B$ ), (ii) class discrepancy (i.e.,  $y_A \neq y_B, D_A = D_B$ ), (iii) domain shift (i.e.,  $o_A \equiv o_B, D_A \neq D_B$ ), and (iv) same appearance (i.e.,  $o_A = o_B$ ). For all the cases, we have a perspective difference (i.e.,  $P_A \neq P_B$ ). Due to the page number restriction, we only present a limited number of matching examples here, but more results can be found in Sec. A in the supplementary material. Figure 6 shows the matching results with and without the proposed method in the four cases. As shown in Fig. 6 (a), SuperPoint and LightGlue alone fail to distinguish objects from the background or match corresponding animal structures. In contrast, Fig. 6 (b) shows that the proposed method significantly reduces unreasonable matches (e.g., human and racket, animal eye and mouth) while increasing correct part-to-part matchings. Particularly, it (i) increased racket matches while reducing mismatches between people,

<sup>2</sup><https://beta.dreamstudio.ai/generate>



Figure 7. The results with different matchers [7, 37, 54]. SuperPoint [16] is used for the keypoint detection. The introduction of SEW increases the matching of eyes and noses between the cheetah and husky and decreases mismatching in the background.

(ii) aligned corresponding animal parts, (iii) matched front wheels despite domain gaps, and (iv) significantly increased matches despite background and angle variations. Notably, structural features, such as eyes and noses, are successfully matched between different classes. We consider that the heatmaps of these animals gave similar importance to the eyes or mouth, thus increasing the number of matches between the same parts of the body. In addition, the results for the same class and appearance showed a decrease in mismatches with the background, which means that the background and foreground had relatively lower similarity.

**Different matchers.** Figure 7 presents the results with other matchers, showing that our method boosts all of them to match corresponding parts between different animals.

## 5.2. Triangular Matching Consistency

Direct quantitative evaluation of non-identical object matching is difficult, as fine-grained ground-truth correspondences between non-identical objects are inherently ambiguous. We therefore introduce *Triangular Matching Consistency* (TMC), an annotation-free metric. TMC measures the agreement between direct matching ( $A, C$ ) and one-hop matching via ( $A, B$ ) and ( $B, C$ ), where  $A, B, C$  are semantically related images. By selecting ( $A, C$ ) to be highly similar, their direct matches provide a reliable reference, and a robust our task method should approximate them even through the one-hop path. TMC reports error magnitude (RMSE), accuracy (PCK), and coverage (Recall). See Sec. E and Fig. C for the details.

**Dataset.** We build an image set of 1,500 triplets for TMC. We use animal images from the Wild category of the Animal Faces dataset [11]. This dataset contains diverse animal face images, including dogs, cats, lions, tigers, and more. For each image  $A_i$ , we generate a new view  $C_i$  of the same object by applying a  $0.9\times$  pinch-out geometric transformation. Then, another species image  $A_j$  ( $j \neq i$ ) is used as  $B$ , forming  $(A, B, C) = (A_i, A_j, C_i)$ . Here, ( $A, C$ ) are identical-instance pairs with viewpoint change, while  $B$  provides a semantically similar but different object. This design allows us to use ( $A, C$ ) as a reliable reference while introduc-

Table 1. Triangular Matching Consistency (TMC) results for the wild category in the Animal Faces dataset [11]. Sparse matchers (LightGlue (LG), GlueStick (GS)) with SuperPoint; dense matchers (LoFTR, ELoFTR) are shown for reference. Metrics include RMSE (lower is better), PCK, and Recall at distance thresholds of 0.01, 0.05, and 0.10 of the image size. Our method SEW shows large improvements across all metrics. See Sec. F for details.

Matching Method		RMSE↓	PCK↑			Recall↑		
			0.01	0.05	0.10	0.01	0.05	0.10
Sparse Matcher (with SuperPoint)	LG	0.14	54.0	68.9	68.2	70.2	76.9	81.4
	LG+SEW	<b>0.07</b>	<b>73.8</b>	<b>87.8</b>	<b>92.6</b>	<b>73.8</b>	<b>87.7</b>	<b>92.9</b>
	GS	0.16	49.0	62.1	62.3	63.8	71.3	73.0
	GS+SEW	0.08	68.8	81.0	86.6	67.3	82.1	86.3
Dense Matcher (detector-free)	LoFTR	0.19	52.7	61.9	66.8	57.3	62.1	65.7
	ELoFTR	0.14	64.5	73.7	76.7	65.9	73.3	76.9

ing non-identical variation through  $B$ , which is essential for evaluating our new task.

**Results.** Table 1 summarizes the evaluation. RMSE measures geometric error, and our module reduces it to about half of each baseline. It even outperforms dense detector-free methods. PCK measures the proportion of matches within a tolerance. It improves at all thresholds, with the largest gains at broader ones and clear benefits even under strict settings. This shows that our method enhances consistent matching between corresponding object parts. Recall measures how well one-hop correspondences cover the  $A \leftrightarrow C$  baseline, and it also increases with our module. Across all three metrics, it indicates that our method enhances the matching consistency of non-identical objects.

## 5.3. Robust Image Matching

We next quantitatively evaluate the proposed method on a standard image matching task with common corruptions. To our knowledge, image matching methods have only been evaluated clean images in literature, except for a few specialized to low-lighting [24] and thermal imaging [83]. We address two cases: (i) both input images are corrupted, and (ii) only one of them is corrupted. The former shows robustness to corruptions, while the latter shows it to domain shifts (i.e., clean image domain vs. corrupted image domain).

**Dataset.** We used the MegaDepth-1500 [79] test set because it contains many objects, making it suitable for a fair evaluation of the non-identical object matching task. This is not the case with other standard datasets [4, 13, 70, 72, 82]. MegaDepth [35] dataset consists of 3D reconstructed data from one million images for two popular phototourism destinations, with camera poses and depth maps computed via COLMAP [71] and refined as ground truth. Common Corruptions [27] simulates the 15 types of corruptions that occur when images are captured. The severity level of corruption is set to 5, which is the most challenging condition.

Table 2. The pose accuracy (AUC) at the maximum angular error of 20° of the relative pose estimation from image pairs MegaDepth-1500 [79] under common corruptions [27]. The left side shows both images corrupted with the same type of corruption, while the right side shows only one of the input images corrupted and the other as a clean image. Our method SEW largely improves LightGlue (LG) [37] and GlueStick (GS) [54] for most categories and the average AUC. The dense matchers, LoFTR [79] and Efficient LoFTR (ELoFTR) [89], are shown for reference. Our method improves the robustness of sparse matchers with a slight reduction in clean accuracy and runtime.

Common Corruptions	AUC@20° with pairs of <b>corrupted</b> images						AUC@20° with pairs of <b>clean and corrupted</b> images					
	<i>keypoint detector : SuperPoint</i>				<i>dense matcher</i>		<i>keypoint detector : SuperPoint</i>				<i>dense matcher</i>	
	LG	LG+SEW	GS	GS+SEW	LoFTR	ELoFTR	LG	LG+SEW	GS	GS+SEW	LoFTR	ELoFTR
None (Clean)	<b>80.61</b>	78.42	<b>78.04</b>	75.08	80.93	83.48	<b>80.61</b>	78.42	<b>78.04</b>	75.08	80.93	83.48
Gaussian Noise	43.09	<b>53.27</b>	45.67	<b>52.24</b>	34.97	35.71	27.54	<b>41.01</b>	41.93	<b>44.98</b>	33.02	36.56
Shot Noise	43.41	<b>53.82</b>	46.10	<b>51.72</b>	40.00	40.36	32.10	<b>42.35</b>	41.75	<b>43.82</b>	36.53	39.52
Impulse Noise	44.98	<b>50.11</b>	44.91	<b>50.05</b>	37.73	39.89	35.95	<b>43.67</b>	40.97	<b>42.91</b>	34.92	38.63
Defocus Blur	32.69	<b>48.35</b>	45.58	<b>49.23</b>	57.22	54.97	18.25	<b>32.13</b>	23.16	<b>30.10</b>	49.40	41.97
Frosted Glass Blur	33.37	<b>47.88</b>	47.40	<b>49.94</b>	54.48	47.09	34.25	<b>44.80</b>	42.96	<b>46.21</b>	54.53	47.52
Motion Blur	42.12	<b>53.63</b>	44.76	<b>51.84</b>	55.60	48.12	50.07	<b>53.10</b>	44.27	<b>52.69</b>	52.81	45.10
Zoom Blur	24.40	<b>31.21</b>	22.87	<b>30.97</b>	23.08	22.54	34.67	<b>40.04</b>	31.18	<b>37.11</b>	31.09	28.49
Snow	<b>31.51</b>	30.33	<b>25.26</b>	23.21	31.85	37.89	56.24	<b>58.15</b>	47.22	<b>48.53</b>	49.15	54.56
Frost	<b>32.24</b>	31.71	<b>26.39</b>	24.22	19.23	25.49	62.20	<b>64.87</b>	54.34	<b>57.26</b>	43.17	45.18
Fog	70.99	<b>73.49</b>	71.19	<b>73.89</b>	67.10	67.74	76.96	<b>78.40</b>	<b>75.18</b>	74.87	74.06	75.36
Brightness	<b>75.48</b>	75.08	<b>73.03</b>	70.84	75.45	79.83	<b>77.14</b>	76.98	<b>76.78</b>	75.06	78.41	81.26
Contrast	<b>39.50</b>	38.47	39.47	<b>42.61</b>	53.22	59.78	43.09	<b>45.17</b>	40.44	<b>43.53</b>	39.66	52.76
Elastic Transform	54.78	<b>66.21</b>	60.37	<b>64.36</b>	58.56	60.50	64.93	<b>68.34</b>	68.02	<b>69.92</b>	68.95	71.49
Pixelate	67.94	<b>68.24</b>	<b>65.01</b>	64.81	74.77	77.27	66.80	<b>68.43</b>	<b>67.16</b>	65.46	76.21	78.45
JPEG Compression	29.54	<b>36.97</b>	39.66	<b>44.70</b>	59.05	59.80	47.41	<b>53.82</b>	<b>60.39</b>	60.12	68.14	71.70
Average	44.40	<b>50.58</b>	46.51	<b>49.64</b>	49.49	50.47	48.51	<b>54.08</b>	50.38	<b>52.84</b>	52.67	53.90
Time [ms/pair]	46.67	68.43	105.72	128.18	312.54	191.17	45.98	68.37	105.34	127.86	305.01	192.53

**Metric.** We follow the standard evaluation of image matching based on relative pose estimation (e.g., see [37]). From the matching results, the relative pose between two images is computed with RANSAC [20] as follows. First, an essential matrix representing the positional transformation between the images is calculated, which is then decomposed into a rotation and a translation. The camera pose error is computed by the maximum angular error, and the area under the cumulative error curve (AUC) at 20° is calculated.

**Robustness against common corruptions.** The left side of Tab. 2 shows the AUC of matching when input images are both corrupted. The results of the dense matching methods (i.e., LoFTR and Efficient LoFTR) are given for reference. As the results show, the proposed method makes the LightGlue and GlueStick robust against most types of common corruptions (roughly 5% to 10% increase) with a slight increase in runtime. The average AUC even exceeds that of Efficient LoFTR. The matching methods benefit from the proposed method, particularly when images are corrupted by noise and blur. We observe AUC decrease in a few types of corruptions, such as Snow and Frost, by the introduction of the proposed method. These corruptions introduce edgy geometric patterns, and the matcher finds correspondence between them. See Sec. B.2 in the supplementary material for more detailed comparisons.

**Robustness against environmental changes.** As the objects to be matched are non-identical, they can appear in

largely different environments. The right side of Tab. 2 shows the AUC when only one of the input images is corrupted and the other is clean. As in the left side of Tab. 2, the proposed method boosts the LightGlue and GlueStick for most cases and the average case. Interestingly, the results are better for almost all weather categories. Only the brightness category shows a decrease in accuracy, but the drop is slight. Weather categories such as Snow and Frost add strong geometry to the overall image, causing mismatches as occlusions. Heatmap weighting is considered to structurally complement the areas hidden by occlusion. Overall, the proposed method enhances robustness to environmental changes, with only a slight drop in precision on clean images. See Sec. B.3 in the supplementary for details.

## 6. Conclusion

We addressed a novel task of matching non-identical objects, which are only semantically similar, not necessarily identical. Due to the non-identity, two objects may be captured in different environments and domains. Thus, robust matching under image corruptions and domain shifts is also encompassed. Our method exploits object detection and visual explanation as high-level features, thereby focusing matching on objects. The experiments showed the effectiveness of our method in non-identical object matching using TMC, a new annotation-free metric, and relative pose estimation under various image corruptions.

## Acknowledgement

Kazuhiko Kawamoto and Hiroshi Kera were supported by JSPS KAKENHI Grant Numbers JP23K24914 and JP22K17962, respectively. They were also supported by Ricoh Company, Ltd.

## References

- [1] Samira Abnar and Willem Zuidema. Quantifying Attention Flow in Transformers. In *Association for Computational Linguistics*, 2020. 6
- [2] Jonathan Astermark, Anders Heyden, and Viktor Larsson. Dense Match Summarization for Faster Two-view Estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2
- [3] Mehmet Aygun and Oisin Mac Aodha. Demystifying Unsupervised Semantic Correspondence Estimation. In *European Conference on Computer Vision (ECCV)*, 2022. 2, 7
- [4] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 7
- [5] Felipe Cadar, Guilherme Potje, Renato Martins, Cédric Demonceaux, and Erickson R. Nascimento. Leveraging Semantic Cues from Foundation Vision Models for Enhanced Local Feature Correspondence. In *Asian Conference on Computer Vision (ACCV)*, 2024. 3
- [6] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-CAM++: Generalized Gradient-Based Visual Explanations for Deep Convolutional Networks. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2018. 6
- [7] Chen, Hongkai, Luo, Zixin, Zhang, Jiahui, Zhou, Lei, Bai, Xuyang, Hu, Zeyu, Tai, Chiew-Lan, Quan, and Long. Learning to Match Features with Seeded Graph Matching Network. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 4, 6, 7
- [8] Hongkai Chen, Zixin Luo, Lei Zhou, Yurun Tian, Mingmin Zhen, Tian Fang, David Mckinnon, Yanghai Tsin, and Long Quan. ASpanFormer: Detector-Free Image Matching with Adaptive Span Transformer. In *European Conference on Computer Vision (ECCV)*, 2022. 1, 6
- [9] Seokju Cho, Sunghwan Hong, Sangryul Jeon, Yunsung Lee, Kwanghoon Sohn, and Seungryong Kim. CATs: Cost Aggregation Transformers for Visual Correspondence. In *Neural Information Processing Systems (NeurIPS)*, 2021. 2, 7
- [10] Seokju Cho, Sunghwan Hong, and Seungryong Kim. CATs++: Boosting Cost Aggregation With Convolutions and Transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 7
- [11] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. StarGAN v2: Diverse Image Synthesis for Multiple Domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 7, 6
- [12] Ethan Coffman, Reagan Clark, Nhat-Tan Bui, Trong Thang Pham, Beth Kegley, Jeremy G. Powell, Jiangchao Zhao, and Ngan Le. Cattleface-rgb: Rgb-t cattle facial landmark benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024. 4
- [13] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 7
- [14] Davison, Andrew J., Reid, Ian D., Molton, Nicholas D., Stasse, and Olivier. MonoSLAM: Real-Time Single Camera SLAM. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2007. 2
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 6, 1, 8, 9, 10
- [16] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperPoint: Self-Supervised Interest Point Detection and Description. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 6, 7, 8, 9, 10, 11, 12
- [17] 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 (CVPR)*, 2023. 2
- [18] Johan Edstedt, Qiyu Sun, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. RoMa: Robust Dense Feature Matching. *arXiv preprint arXiv:2305.15404*, 2023. 1, 2, 6
- [19] Johan Edstedt, Georg Bökman, and Zhenjun Zhao. DeDoDe v2: Analyzing and Improving the DeDoDe Keypoint Detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024. 6
- [20] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. In *Communications of the ACM*, 1981. 8, 4
- [21] Pierre Gleize, Weiyao Wang, and Matt Feiszli. SiLK – Simple Learned Keypoints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2, 6
- [22] Bumsub Ham, Minsu Cho, Cordelia Schmid, and Jean Ponce. Proposal Flow: Semantic Correspondences from Object Proposals. In *arXiv:1703.07144*, 2017. 2, 4, 5, 7
- [23] Kai Han, Rafael S. Rezende, Bumsub Ham, Kwan-Yee K. Wong, Minsu Cho, Cordelia Schmid, and Jean Ponce. SCNet: Learning Semantic Correspondence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 2, 7

- [24] Yuze He, Yubin Hu, Wang Zhao, Jisheng Li, Yong-Jin Liu, Yuxing Han, and Jiangtao Wen. DarkFeat: Noise-Robust Feature Detector and Descriptor for Extremely Low-Light RAW Images. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2023. 1, 2, 3, 7
- [25] Hedlin, Eric, Sharma, Gopal, Mahajan, Shweta, He, Xingzhe, Isack, Hossam, Rhodin, Abhishek Kar Helge, Tagliasacchi, Andrea, Yi, and Kwang Moo. Unsupervised Keypoints from Pretrained Diffusion Models. *arXiv preprint arXiv:2312.00065*, 2023. 2, 6
- [26] Eric Hedlin, Gopal Sharma, Shweta Mahajan, Hosam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. Unsupervised Semantic Correspondence Using Stable Diffusion. In *Neural Information Processing Systems (NeurIPS)*, 2023. 3, 7
- [27] Dan Hendrycks and Thomas Dietterich. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. 2, 7, 8, 1, 10
- [28] Sunghwan Hong, Seokju Cho, Seungryong Kim, and Stephen Lin. Unifying Feature and Cost Aggregation with Transformers for Semantic and Visual Correspondence. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. 7
- [29] Hanwen Jiang, Arjun Karapur, Bingyi Cao, Qixing Huang, and Andre Araujo. OmniGlue: Generalizable Feature Matching with Foundation Model Guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1
- [30] Wei Jiang, Eduard Trulls, Jan Hosang, Andrea Tagliasacchi, and Kwang Moo Yi. COTR: Correspondence Transformer for Matching Across Images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 1, 6
- [31] Dongki Jung, Jaehoon Choi, Yonghan Lee, Somi Jeong, Taejae Lee, Dinesh Manocha, and Suyong Yeon. EDM: Equirectangular Projection-Oriented Dense Kernelized Feature Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 1
- [32] Seungwook Kim, Juhong Min, and Minsu Cho. Efficient Semantic Matching with Hypercolumn Correlation. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2023. 7
- [33] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment Anything. *arXiv:2304.02643*, 2023. 6
- [34] Xinghui Li, Jingyi Lu, Kai Han, and Victor Prisacariu. SD4Match: Learning to Prompt Stable Diffusion Model for Semantic Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1
- [35] Zhengqi Li and Noah Snavely. MegaDepth: Learning Single-View Depth Prediction from Internet Photos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 7, 4
- [36] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common Objects in Context. *CoRR*, 2014. 6, 1, 8, 9, 10
- [37] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. LightGlue: Local Feature Matching at Light Speed. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 1, 2, 4, 6, 7, 8, 9, 10, 11, 12, 13, 14
- [38] Xiangzeng Liu, Chi Wang, Guanglu Shi, Xiaodong Zhang, Qiguang Miao, and Miao Fan. SGAD: Semantic and Geometric-aware Descriptor for Local Feature Matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 6
- [39] Yepeng Liu, Wenpeng Lai, Zhou Zhao, Yuxuan Xiong, Jinchi Zhu, Jun Cheng, and Yongchao Xu. LiftFeat: 3D Geometry-Aware Local Feature Matching. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2025. 1
- [40] David G Lowe. Distinctive image features from scaleinvariant keypoints. In *International Journal of Computer Vision (IJCV)*, 2004. 2, 6
- [41] Xiaoyong Lu and Songlin Du. Raising the Ceiling: Conflict-Free Local Feature Matching with Dynamic View Switching. In *European Conference on Computer Vision (ECCV)*, 2024. 6
- [42] Xiaoyong Lu and Songlin Du. JamMa: Ultra-lightweight Local Feature Matching with Joint Mamba. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 6
- [43] Xiaoyong Lu, Yaping Yan, Tong Wei, and Songlin Du. Scene-Aware Feature Matching, 2023. 1, 3
- [44] Lundberg, Scott M., Lee, and Su-In. A Unified Approach to Interpreting Model Predictions. In *Neural Information Processing Systems (NeurIPS)*, 2017. 6
- [45] Grace Luo, Lisa Dunlap, Dong Huk Park, Aleksander Holynski, and Trevor Darrell. Diffusion Hyperfeatures: Searching Through Time and Space for Semantic Correspondence. In *Neural Information Processing Systems (NeurIPS)*, 2023. 1, 2, 7
- [46] George Martvel, Greta Abele, Annika Bremhorst, Chiara Canori, Nareed Farhat, Giulia Pedretti, Ilan Shimshoni, and Anna Zamansky. Dogflw: Dog facial landmarks in the wild dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024. 4
- [47] George Martvel, Nareed Farhat, Ilan Shimshoni, and Anna Zamansky. Catflw: Cat facial landmarks in the wild dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024. 4
- [48] Juhong Min, Jongmin Lee, Jean Ponce, and Minsu Cho. Hyperpixel Flow: Semantic Correspondence with Multi-layer Neural Features. In *Proceedings of the IEEE/CVF In-*

- ternational Conference on Computer Vision (ICCV)*, 2019. 2, 4, 7
- [49] Guillaume Le Moing, Jean Ponce, and Cordelia Schmid. Dense Optical Tracking: Connecting the Dots. *arXiv preprint arXiv:2312.00786*, 2023. 2
- [50] James Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 1957. 4
- [51] Raúl Mur-Artal, J. M. M. Montiel, and Juan D. Tardós. ORB-SLAM: A Versatile and Accurate Monocular SLAM System. In *IEEE Transactions on Robotics*, 2015. 2
- [52] Junjie Ni, Guofeng Zhang, Guanglin Li, Yijin Li, Xinyang Liu, Zhaoyang Huang, and Hujun Bao. ETO:Efficient Transformer-based Local Feature Matching by Organizing Multiple Homography Hypotheses. In *European Conference on Computer Vision (ECCV)*, 2024. 6
- [53] 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. 3, 6, 7
- [54] Rémi Pautrat, Iago Suárez, Yifan Yu, Marc Pollefeys, and Viktor Larsson. GlueStick: Robust Image Matching by Sticking Points and Lines Together. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 1, 2, 4, 6, 7, 8, 13, 14
- [55] William Peebles, Jun-Yan Zhu, Richard Zhang, Antonio Torralba, Alexei Efros, and Eli Shechtman. GAN-Supervised Dense Visual Alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 7
- [56] Gabriel Peyré and Marco Cuturi. Computational Optimal Transport. In *Foundations and Trends® in Machine Learning*, 2019. 6
- [57] Guilherme Potje, Felipe Cadar, André Araujo, Renato Martins, and Erickson R. Nascimento. XFeat: Accelerated Features for Lightweight Image Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 6
- [58] Qiyang Qian, Hansheng Chen, Masayoshi Tomizuka, Kurt Keutzer, Qianqian Wang, and Chenfeng Xu. Bridging Viewpoint Gaps: Geometric Reasoning Boosts Semantic Correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 7
- [59] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning (ICML)*, 2021. 5, 6
- [60] Umer Rafi, Andreas Doering, Bastian Leibe, and Juergen Gall. Self-supervised Keypoint Correspondences for Multi-person Pose Estimation and Tracking in Videos. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [61] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Juntao Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chaoyuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. SAM 2: Segment Anything in Images and Videos. *arXiv preprint arXiv:2408.00714*, 2024. 6
- [62] I. Rocco, R. Arandjelović, and J. Sivic. Convolutional neural network architecture for geometric matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1
- [63] I. Rocco, R. Arandjelović, and J. Sivic. End-to-end weakly-supervised semantic alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [64] I. Rocco, R. Arandjelović, and J. Sivic. Efficient Neighbourhood Consensus Networks via Submanifold Sparse Convolutions. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 6
- [65] Barbara Roessle and Matthias Nießner. End2End Multi-View Feature Matching with Differentiable Pose Optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 1
- [66] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models, 2021. 3, 6, 1, 7, 10
- [67] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2011. 2, 6
- [68] Kihwan Ryoo, Hyungtae Lim, and Hyun Myung. MambaGlue: Fast and Robust Local Feature Matching With Mamba. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2025. 1
- [69] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning Feature Matching with Graph Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 4, 6
- [70] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, and Tomas Pajdla. Benchmarking 6DOF Outdoor Visual Localization in Changing Conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 7
- [71] Johannes Lutz Schonberger and Jan-Michael Frahm. Structure-From-Motion Revisited. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 7
- [72] Johannes Lutz Schönberger, Hans Hardmeier, Torsten Sattler, and Marc Pollefeys. Comparative Evaluation of Hand-Crafted and Learned Local Features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 7

- [73] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 2, 4, 6, 3
- [74] Ayush Shrivastava and Andrew Owens. Self-Supervised Spatial Correspondence Across Modalities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 7
- [75] Aleksandar Shtedritski, Andrea Vedaldi, and Christian Rupprecht. Learning Universal Semantic Correspondences with No Supervision and Automatic Data Curation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 2023. 7
- [76] Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. RoFormer: Enhanced Transformer with Rotary Position Embedding. In *arXiv preprint arXiv:2104.09864*, 2021. 5
- [77] Kosuke Sumiyasu, Kazuhiko Kawamoto, and Hiroshi Kera. Identifying Important Group of Pixels using Interactions, 2024. 6
- [78] Sun, Jiaming, Wang, Zihao, Zhang, Siyu, He, Xingyi, Zhao, Hongcheng, Zhang, Guofeng, Zhou, and XiaoWei. OnePose: One-Shot Object Pose Estimation without CAD Models. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [79] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and XiaoWei Zhou. LoFTR: Detector-Free Local Feature Matching with Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2, 6, 7, 8, 4, 13, 14
- [80] Jiaming Sun, Zihao Wang, Siyu Zhang, Xingyi He, Hongcheng Zhao, Guofeng Zhang, and XiaoWei Zhou. OnePose: One-Shot Object Pose Estimation without CAD Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [81] Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan. Emergent Correspondence from Image Diffusion. In *Neural Information Processing Systems (NeurIPS)*, 2023. 7
- [82] Bart Thomee, David A. Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. YFCC100M: The new data in multimedia research. *Communications of the ACM*, 2016. 7
- [83] Önder Tuzcuoğlu, Aybora Köksal, Buğra Sofu, Sinan Kalkan, and A Aydin Alatan. XoFTR: Cross-modal Feature Matching Transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 7
- [84] Michał J. Tyszkiewicz, Pascal Fua, and Eduard Trulls. DISK: Learning local features with policy gradient. In *Neural Information Processing Systems (NeurIPS)*, 2020. 2, 6
- [85] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Neural Information Processing Systems (NeurIPS)*, 2017. 6
- [86] Wang, Xinjiang, Liu, Zeyu, Hu, Yu, Xi, Wei, Yu, Wenxian, Zou, and Danping. FeatureBooster: Boosting Feature Descriptors with a Lightweight Neural Network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 6
- [87] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2, 4, 6, 3
- [88] Qianqian Wang, Yen-Yu Chang, Ruojin Cai, Zhengqi Li, Bharath Hariharan, Aleksander Holynski, and Noah Snavely. Tracking Everything Everywhere All at Once. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2
- [89] Yifan Wang, Xingyi He, Sida Peng, Dongli Tan, and Xiaowei Zhou. Efficient LoFTR: Semi-Dense Local Feature Matching with Sparse-Like Speed. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 6, 8, 13, 14
- [90] Shibiao Xu, Shunpeng Chen, Rongtao Xu, Changwei Wang, Peng Lu, and Li Guo. Local Feature Matching Using Deep Learning: A Survey, 2024. 2, 6
- [91] Fei Xue, Ignas Budvytis, and Roberto Cipolla. SFD2: Semantic-guided Feature Detection and Description. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [92] Fei Xue, Ignas Budvytis, and Roberto Cipolla. IMP: Iterative Matching and Pose Estimation with Adaptive Pooling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1
- [93] Fei Xue, Sven Elfle, Laura Leal-Taixé, and Qunjie Zhou. MATCHA: Towards Matching Anything. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 2
- [94] Toshinori Yamauchi, Hiroshi Kera, and Kazuhiko Kawamoto. Explaining Object Detectors via Collective Contribution of Pixels, 2024. 6
- [95] Jiahuan Yu, Jiahao Chang, Jianfeng He, Tianzhu Zhang, and Feng Wu. Adaptive Spot-Guided Transformer for Consistent Local Feature Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 6
- [96] Liao Yun, Di Yide, Zhu Kaijun, Zhou Hao, Lu Mingyu, Zhang Yijia, Qing Duan, and Liu Junhui. Local feature matching from detector-based to detector-free: a survey. *Applied Intelligence*, 2024. 2, 6
- [97] 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. In *Neural Information Processing Systems (NeurIPS)*, 2023. 7
- [98] Junyi Zhang, Charles Herrmann, Junhwa Hur, Eric Chen, Varun Jampani, Deqing Sun, and Ming-Hsuan Yang. Telling Left from Right: Identifying Geometry-Aware Semantic Correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1, 2, 3, 4, 7

- [99] Kaiyan Zhang, Xinghui Li, Jingyi Lu, and Kai Han. Semantic Correspondence: Unified Benchmarking and a Strong Baseline, 2025. [1](#)
- [100] Yesheng Zhang and Xu Zhao. MESA: Matching Everything by Segmenting Anything. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [6](#)
- [101] Yesheng Zhang, Xu Zhao, and Dahong Qian. Searching from Area to Point: A Hierarchical Framework for Semantic-Geometric Combined Feature Matching. In *arXiv preprint arXiv:2305.00194*, 2023. [1](#), [6](#)
- [102] Zhou, Bolei, Khosla, Aditya, Lapedriza, Agata, Oliva, Aude, Torralba, and Antonio. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [6](#)
- [103] Qunjie Zhou, Torsten Sattler, and Laura Leal-Taixe. Patch2Pix: Epipolar-Guided Pixel-Level Correspondences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [1](#), [6](#)