# **TopicFM: Robust and Interpretable Topic-Assisted Feature Matching**

Khang Truong Giang<sup>1</sup>, Soohwan Song<sup>2\*</sup>, Sungho Jo<sup>1\*</sup>

 <sup>1</sup> School of Computing, KAIST, Daejeon 34141, Republic of Korea
 <sup>2</sup> Intelligent Robotics Research Division, ETRI, Daejeon 34129, Republic of Korea khangtg@kaist.ac.kr, soohwansong@etri.re.kr, shjo@kaist.ac.kr

#### Abstract

This study addresses an image-matching problem in challenging cases, such as large scene variations or textureless scenes. To gain robustness to such situations, most previous studies have attempted to encode the global contexts of a scene via graph neural networks or transformers. However, these contexts do not explicitly represent high-level contextual information, such as structural shapes or semantic instances; therefore, the encoded features are still not sufficiently discriminative in challenging scenes. We propose a novel image-matching method that applies a topic-modeling strategy to encode high-level contexts in images. The proposed method trains latent semantic instances called topics. It explicitly models an image as a multinomial distribution of topics, and then performs probabilistic feature matching. This approach improves the robustness of matching by focusing on the same semantic areas between the images. In addition, the inferred topics provide interpretability for matching the results, making our method explainable. Extensive experiments on outdoor and indoor datasets show that our method outperforms other state-of-the-art methods, particularly in challenging cases.

#### Introduction

Image matching is a long-standing problem in computer vision. It aims to find pixel-to-pixel correspondences across two or more images. Conventional image matching methods (Lowe 2004; Bay et al. 2008; Sattler, Leibe, and Kobbelt 2012) usually involve the following steps: i) local feature detection, ii) feature description, iii) matching, and iv) outlier rejection. These methods usually involve extracting sparse handcrafted local features (i.e., SIFT (Lowe 2004), SURF (Bay et al. 2008), or ORB (Rublee et al. 2011)) and matching them using a nearest neighbor search. Many recent studies have adopted convolutional neural networks (CNNs) to extract local features, which significantly outperform the conventional handcrafted features. However, such methods sometimes fail in challenging cases, such as illumination variations, repetitive structures, or low-texture conditions.

To address this issue, detector-free methods (Li et al. 2020; Rocco et al. 2018) have been proposed. These methods estimate dense feature maps without feature detection



Figure 1: The main idea of our human-friendly topicassisted feature matching, TopicFM. This method represents an image as a set of topics marked in different colors and quickly recognizes the same structures between an image pair. It then leverages the distinctive information of each topic to augment the pixel-level representation. As shown in the comparing illustrations above, TopicFM provides robust and accurate matching results, even for challenging scenes with large illumination and viewpoint variations.

and perform pixel-wise dense matching. Furthermore, a coarse-to-fine strategy has been applied to improve the computational efficiency. The strategy finds matches at a coarse level, and then refines the matches at a finer level. Such methods (Sun et al. 2021; Wang et al. 2022) produce a large number of matches, even for repetitive patterns and textureless scenes, thus achieving state-of-the-art performance.

However, detector-free methods still have some factors that degrade the matching performance. First, these methods cannot adequately incorporate the global context of a scene for feature matching. Several methods have attempted to implicitly capture global contextual information via transformers (Sun et al. 2021; Jiang et al. 2021) or patch-level matches (Zhou, Sattler, and Leal-Taixe 2021), but higherlevel contexts, such as semantic instances, should be effectively exploited to learn robust representations. Second, they exhaustively search for all features of the entire image area. Therefore, their matching performance is considerably low when there are limited covisible regions between images. Finally, these methods require intensive computation of dense matching, which increases runtime. Therefore, a more ef-

<sup>\*</sup>Corresponding authors

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ficient model is needed for real-time applications such as SLAM (Mur-Artal, Montiel, and Tardos 2015).

In this study, we propose a novel detector-free feature matching method, TopicFM, that encodes high-level contextual information on images based on a topic modeling strategy in data mining (Blei, Ng, and Jordan 2003; Yan et al. 2013). TopicFM models an image as a multinomial distribution over topics, where a topic represents a latent semantic instance such as an object or structural shape. TopicFM then performs probabilistic feature matching based on the distribution of the latent topics. It integrates topic information into local visual features to enhance their distinctiveness. Furthermore, it effectively matches features within overlapping regions between an image pair by estimating the covisible topics. Therefore, TopicFM provides robust and accurate feature-matching results, even for challenging scenes with large scale and viewpoint variations.

The proposed method also provides interpretability for matching results across topics. Fig. 1 illustrates the representative topics inferred from the image matching results. In Fig. 1, the image regions with the same object or structure are assigned to the same topic. Based on the sufficient high-level context information in the topic, TopicFM can learn discriminative features. Therefore, it is able to find the accurate dense correspondences in the same topic regions. This approach is similar to the human cognitive system, in which humans quickly recognize covisible regions based on semantic information and then search for matching points in these regions. By applying this top-down approach, our method successfully detected dense matching points in various challenging image conditions.

Going one step further, we designed an efficient endto-end network architecture to accelerate the computation. We adopted a coarse-to-fine framework and constructed lightweight networks for each stage. In particular, TopicFM only focuses on the same semantic areas between the images for learning features. Therefore, our method requires less computation compared to other methods (Sarlin et al. 2020; Sun et al. 2021; Wang et al. 2022) that apply the transformer to the whole domain.

The contributions of this study are as follows:

- We present a novel feature-matching method that fuses local context and high-level semantic information into latent features using a topic modeling strategy. This method produces accurate dense matches in challenging scenes by inferring covisible topics.
- We formulate the topic inference process as a learnable transformer module. These inferred topics can provide interpretability for matching results with humans.
- We design an efficient end-to-end network model to achieve real-time performance. This model processes image frames much faster than state-of-the-art methods such as Patch2Pix (Zhou, Sattler, and Leal-Taixe 2021) and LoFTR (Sun et al. 2021).
- We empirically evaluate the proposed method through extensive experiments. We also provide results on the interpretability of our topic models. Source code for the proposed method is publicly available.

## **Related Works**

Image Matching The standard pipeline for image matching (Ma et al. 2021) consists of four steps: feature detection, description, and matching, and outlier rejection. Traditional feature detection-and-description methods such as SIFT (Lowe 2004), SURF (Bay et al. 2008), and BRIEF (Calonder et al. 2010), although widely used in many applications, require a complicated selection of hyperparameters to achieve reliable performance (Efe, Ince, and Alatan 2021). Twelve years after SIFT, a fully learning-based architecture, LIFT (Yi et al. 2016), was proposed to address the hand-crafting issue of traditional approaches. Many studies (DeTone, Malisiewicz, and Rabinovich 2018; Ono et al. 2018; Dusmanu et al. 2019; Revaud et al. 2019; Bhowmik et al. 2020; Tyszkiewicz, Fua, and Trulls 2020) also proposed learning-based approaches, which have become dominant in feature detection and description. However, their methods mainly adopt standard CNNs to learn features from local context information, which is less effective when processing low-textured images.

To address this issue, some studies (Sun et al. 2021; Wang et al. 2022; Luo et al. 2019, 2020) have additionally considered global context information. ContextDesc (Luo et al. 2019) and ALSFeat (Luo et al. 2020) proposed a geometric context encoder using a large patch sampler and deformable CNN, respectively. LoFTR (Sun et al. 2021) applies transformers with self- and cross-attentions to extract dense feature maps. Although these methods are technically sound, they are unable to encode high-level contexts such as objects or structural shapes. They cannot explicitly represent hidden semantic structures in an image and lack interpretability. However, our method can capture latent semantic information via a topic modeling strategy; therefore, our matching results would be fairly interpretable.

Given two sets of features produced by the detection-anddescription methods, a basic feature-matching algorithm applies the nearest neighbor search (Muja and Lowe 2014) or ratio test (Lowe 2004) to find potential correspondences. Next, the matching outliers are rejected by RANSAC (Fischler and Bolles 1981), consensus- or motion-based heuristics (Lin et al. 2017b; Bian et al. 2017), or learning-based methods (Yi et al. 2018; Zhang et al. 2019). The outlier rejection performance relies heavily on the accuracy of the trained features. Recently, several studies (Sarlin et al. 2020; Chen et al. 2021; Shi et al. 2022) employed an attentional graph neural network (GNN) to enhance the quality of extracted features. These features were matched with an optimal transport layer (Cuturi 2013). As the performance of these methods depends on the features of the detector, these methods cannot guarantee robust and reliable performance.

Motivated by the above observation, several studies (Zhou, Sattler, and Leal-Taixe 2021; Sun et al. 2021; Wang et al. 2022; Jiang et al. 2021) have proposed an end-toend network architecture that performs image matching in a single forward pass instead of dividing separate steps. The network directly processed dense feature maps instead of extracting sparse feature points. Several studies applied a coarse-to-fine strategy to process the dense features of a high-resolution image efficiently. Patch2Pix detects coarse matches in low-resolution images and gradually refines them at higher resolutions. Similarly, other coarse-to-fine methods (Sun et al. 2021; Jiang et al. 2021; Wang et al. 2022) learn robust and distinctive features using transformers and achieve state-of-the-art performance. However, these methods remain inefficient when propagating global context information to the entire image region. We argue that the invisible regions between an image pair are redundant and may cause noise when learning the features with transformers. Therefore, we propose a topic modeling approach to utilize adequate context cues for learning representations.

Interpretable Image Matching The interpretability of vision models has recently been actively researched (Zhou et al. 2016; Selvaraju et al. 2017; Bau et al. 2018; Chefer, Gur, and Wolf 2021). It aims to explain a certain decision or prediction in image recognition (Williford, May, and Byrne 2020; Wang et al. 2021), or deep metric learning (Zhao et al. 2021). In image matching, the detector-based methods (Förstner, Dickscheid, and Schindler 2009; Lowe 2004) can estimate interpretable feature keypoints such as corners, blobs, or ridges. However, detected features do not represent spatial or semantic structures. Otherwise, existing end-toend methods only extract dense feature maps using the local context via CNNs (Zhou, Sattler, and Leal-Taixe 2021) or global context via transformers (Sun et al. 2021; Wang et al. 2022). However, these approaches cannot explicitly describe the details of the observed context information; therefore, their results lack interpretability.

The human cognitive system quickly recognizes covisible regions based on high-level contextual information, such as objects or structures. It then determines the matching points in the covisible regions. Inspired by this cognitive process, we designed an end-to-end model that is human-friendly. It categorizes local structures in images into different topics and uses only the information within topics to augment features. Moreover, our method performs interpretable matching by selecting important topics in the covisible regions of the two images. To the best of our knowledge, our method is the first to explicitly introduce interpretability to an image matching task.

**Semantic Segmentation** Various deep learning models for semantic segmentation have been introduced, such as fully convolutional networks (Long, Shelhamer, and Darrell 2015), encoder-decoder (Yuan, Chen, and Wang 2020), R-CNN-based (He et al. 2017), or attention-based models (Strudel et al. 2021). Unlike semantic segmentation, our topic modeling does not strictly detect semantic objects. However, it can effectively exploit the local structures or shapes, which benefits learning pixel-level representation for feature matching. Moreover, the topics can be trained in a self-supervised manner without requiring a large amount of labeled training data, as in semantic segmentation.

#### **Proposed Method**

### **Coarse-to-Fine Architecture**

This study addresses the feature-matching problem of an image pair. Let  $F^A$  and  $F^B$  be the feature maps extracted

from images  $I^A$  and  $I^B$ , respectively. Our objective is to find accurate and dense matching correspondences between two feature points,  $f_i^A \in F^A$  and  $f_j^B \in F^B$ . We employ a coarse-to-fine architecture (Sun et al. 2021) that trains a feature-matching network end-to-end. This architecture estimates coarse matches from low-resolution features and refines the matches to a finer level. This approach makes it possible to perform feature matching of high-resolution images in real time while preserving the pixel-level accuracy.

Fig. 2 depicts the proposed architecture for feature matching, which is composed of three steps: i) feature extraction, ii) coarse-level matching, and iii) fine-level refinement. The feature extraction step generates multiscale dense features through a UNet-like architecture (Lin et al. 2017a). Let  $\{F_c^A, F_c^B\}$  and  $\{F_f^A, F_f^B\}$  be pairs of coarse- and finer-lever feature maps of an image pair  $\{I^A, I^B\}$ , respectively. The coarse matching method estimates the matching probability distribution of  $\{F_c^A, F_c^B\}$  using a topic-assisted matching module, TopicFM. It then determines coarse correspondences based on the probability distribution (see the next section). The last stage refines the coarse matches to a finer level with high-resolution features  $\{F_f^A, F_f^B\}$ . We adopted the matching refinement method of LoFTR directly (Sun et al. 2021). For each coarse match (i, j), the method finds the best matching coordinate in  $F_f^B$  by measuring the similarities between a feature point  $F_{f,i}^{A} \in F_{f}^{A}$  for all features of the cropped patch at  $F_{f,j}^{B} \in F_{f}^{B}$ .

#### **Topic-Assisted Feature Matching**

**Probabilistic Feature Matching** The coarse feature maps  $F_c^A, F_c^B$  can be regarded as a bag-of-visual-words (Sivic and Zisserman 2003; Csurka et al. 2004), where each feature vector represents a visual word. Let  $m_{ij}$  be a random variable that indicates an event in which the  $i^{th}$  feature  $F_{c,i}^A$  is matched to the  $j^{th}$  feature  $F_{c,j}^B$ . Given two feature sets  $\{F_c^A, F_c^B\}$ , our goal is to estimate the match distribution of all possible matches  $M = \{m_{ij}\}$  (Bhowmik et al. 2020):

$$P(M \mid F_c^A, F_c^B) = \prod_{m_{ij} \in M} P\left(m_{ij} \mid F_c^A, F_c^B\right)$$
(1)

The matches with high match probability  $P(m_{ij}|F_c^A, F_c^B)$  are selected as the coarse correspondences. Existing methods (Bhowmik et al. 2020; Sun et al. 2021; Sarlin et al. 2020) directly infer the matching probabilities using Softmax (Bhowmik et al. 2020), Dual-Softmax (Sun et al. 2021), or optimal transport with Signkhorn regularization (Sarlin et al. 2020). Unlike these methods, TopicFM incorporates the latent distribution of topics to estimate the matching distribution.

To solve the matching problem of Eq. 1, our method infers a topic distribution for each feature point (Eq. 3). It then estimates a matching probability conditioned on topics for each matching candidate (Eq. 4). A sampling strategy is employed to calculate this probability (Eq. 6 and Eq. 7). Finally, our method selects the coarse matches from the candidates using probability thresholding.



Figure 2: Overview of the proposed architecture. (a) Our method first extracts multilevel feature maps. (b) Next, the method finds coarse matches from low-resolution features. It infers a topic distribution via a cross-attention layer with topic embeddings. It then samples topic labels of each feature point and augments the features with self/cross attention layers. The coarse matches are determined by estimating a matching probability with dual-softmax. (c) Finally, our method refines the coordinates inside the cropped patches at high resolution.

1

**Topic Inference via Transformers** We assume that the structural shapes or semantic instances of the images in a specific dataset can be categorized into K topics. Therefore, each image can be modeled as a multinomial distribution over K topics. The probability distribution of the topics was assigned to each feature point.

Let  $z_i$  and  $\theta_i$  be a topic indicator and topic distribution for feature  $F_i$ , respectively, where  $z_i \in \{1, \ldots, K\}$  and  $\theta_{i,k} = p(z_i = k | F)$  are the probabilities for assigning  $F_i$  to topic k. We represent topic k as an embedding vector,  $T_k$ , which is trainable. To estimate  $\theta_i$ , our method infers the local topic representations  $\hat{T}_k$  from the global representations  $T_k$  using transformers:

$$\hat{T}_k = \mathcal{CA}(T_k, F) \tag{2}$$

where  $CA(T_k, F)$  is the cross-attention layer between queries  $T_k$ , keys F, and values F. This function collects relevant information from an image of each topic. Finally, the topic probability  $\theta_{i,k}$  is defined as the distance between feature  $F_i$  and individual topics  $\hat{T}_k$  as follows:

$$\theta_{i,k} = \frac{\langle \hat{T}_k, F_i \rangle}{\sum_{h=1}^K \langle \hat{T}_h, F_i \rangle} \tag{3}$$

**Topic-Aware Feature Augmentation** This section describes the computation of Eq. 1 using inferred topics. We augment the features based on the high-level contexts of topics to enhance their distinctiveness. The augmented features are then used to estimate matching probability more precisely. Given a feature point pair  $(F_{c,i}^A, F_{c,j}^B)$ , we define an assigned topic of  $z_{ij}$  as a random variable  $z_{ij} \in \mathcal{Z} = \{1, 2, \ldots, K, NaN\}$ . If  $z_{ij} = k$  ( $k = 1, \ldots, K$ ), the pair belongs to the same topic k. Otherwise,  $z_{ij} = NaN$  indicates that  $F_{c,i}^A$  and  $F_{c,j}^B$  do not belong to the same topic; therefore, they are highly unmatchable.

We define  $z_{ij}$  as a latent variable for computing the

matching distribution in Eq. 1 as follows:

$$\log P\left(M \mid F_c^A, F_c^B\right) = \sum_{m_{ij} \in M} \log P\left(m_{ij} \mid F_c^A, F_c^B\right)$$
$$= \sum_{m_{ij} \in M} \log \sum_{k \in \mathcal{Z}} P\left(m_{ij}, z_{ij} = k \mid F_c^A, F_c^B\right) \quad (4)$$

To compute Eq. 4, we approximated this equation with an evidence lower bound (ELBO):

$$\mathcal{L}_{ELBO} = \sum_{m_{ij}} \sum_{k \in \mathcal{Z}} P\left(z_{ij} = k \mid F_c\right) \log P\left(m_{ij} \mid z_{ij}, F_c\right)$$
$$= \sum_{m_{ij}} E_{p(z_{ij})} \log P\left(m_{ij} \mid z_{ij}, F_c^A, F_c^B\right) \quad (5)$$

where  $P(m_{ij}|z_{ij}, F_c^A, F_c^B)$  refers to the matching probability conditioned on topic  $z_{ij}$ . Eq. 5 can be estimated by applying Monte-Carlo (MC) sampling, as follows:

$$\mathcal{L}_{ELBO} = \sum_{m_{ij} \in M} \frac{1}{S} \sum_{s=1}^{S} \log P\left(m_{ij} \mid z_{ij}^{(s)}, F_c^A, F_c^B\right)$$
(6)  
$$z_{ij}^{(s)} \sim P\left(z_{ij} \mid F_c^A, F_c^B\right)$$
(7)

where S is the number of samples  $(S \ll K)$ . This sampling approach improves computational efficiency because it is unnecessary to iterate all K topics to compute the expectation in Eq. 5. Finally, the problem is reduced to the computation of the topic distribution  $P(z_{ij} | F_c^A, F_c^B)$  and the conditional matching distribution  $P(m_{ij} | z_{ij}^{(s)}, F_c^A, F_c^B)$ .

**Topic Distribution** We estimate the distribution of  $z_{ij}$  by factorizing it into two distributions of  $z_i$  and  $z_j$  as follows:

$$P(z_{ij} = k \mid F_c^A, F_c^B) =$$

$$P(z_i = k \mid F_c^A)P(z_j = k \mid F_c^B) = \theta_{i,k}^A \theta_{j,k}^B \quad (8)$$

where  $\theta_{i,k}^A$ ,  $\theta_{j,k}^B$  are computed using Eq. 2 and Eq. 3. This represents the probability of assigning feature pair  $\{F_{c,i}^A, F_{c,j}^B\}$  to a specific topic  $k \in \{1, \ldots, K\}$ . The probability of being in at least one topic is calculated as follows:

$$P(z_{ij} \in \{1, \dots, K\} \mid F_c^A, F_c^B) = \sum_{k=1}^K \theta_{i,k}^A \theta_{j,k}^B$$
(9)

Otherwise, the probability of not being on the same topic is calculated by

$$P(z_{ij} = NaN \mid .) = 1 - \sum_{k=1}^{K} P(z_{ij} = k \mid .)$$
  
=  $1 - \sum_{k=1}^{K} \theta_{i,k}^{A} \theta_{j,k}^{B}$  (10)

In summary, the topic distribution for each pair of features was determined as follows:

$$P(z_{ij} = k \mid F_c) = \begin{cases} \theta_{i,k}^A \theta_{j,k}^B & k \in \{1...K\} \\ 1 - \sum_{k=1}^K \theta_{i,k}^A \theta_{j,k}^B & k = NaN \end{cases}$$
(11)

We can sample  $z_{ij}^{(s)}$  from this distribution by sampling  $z_i^{(s)}$  and  $z_j^{(s)}$  from  $\theta_i^A$  and  $\theta_j^B$  separately based on the independent and identically distributed (i.i.d.) assumption:

$$z_{ij}^{(s)} = \begin{cases} k & \text{if } z_i^{(s)} = z_j^{(s)} = k\\ NaN & \text{if } z_i^{(s)} \neq z_j^{(s)} \end{cases}$$
(12)

**Conditional Matching Distribution** After sampling, we classified a pair of features into topics. Let  $F_c^{A,\tilde{k}} \subset F_c^A$  and  $F_c^{B,\tilde{k}} \subset F_c^B$  be a set of features sampled with topic  $\tilde{k} = z_{ij}^{(s)}$ . These features are augmented to improve their distinctiveness by applying self- and cross-attentions (SA and CA) of the transformer (Sarlin et al. 2020; Sun et al. 2021):  $\hat{F}_{c,i}^{A,\tilde{k}} \leftarrow S\mathcal{A}\left(F_{c,i}^{A,\tilde{k}}, F_c^{A,\tilde{k}}\right), \quad \hat{F}_{c,j}^{B,\tilde{k}} \leftarrow S\mathcal{A}\left(F_{c,j}^{B,\tilde{k}}, F_c^{B,\tilde{k}}\right)$  $\hat{F}_{c,i}^{A,\tilde{k}} \leftarrow C\mathcal{A}\left(F_{c,i}^{A,\tilde{k}}, F_c^{B,\tilde{k}}\right), \quad \hat{F}_{c,j}^{B,\tilde{k}} \leftarrow C\mathcal{A}\left(F_{c,j}^{B,\tilde{k}}, F_c^{A,\tilde{k}}\right)$ This augmentation learns powerful representation by con-

sidering adequate context information inside the topic  $\tilde{k}$ . Finally, the matching probability conditioned on topic  $z_{ij}^{(s)}$  in Eq. 6 is determined by computing the feature distance and normalizing it with a dual-softmax (Sun et al. 2021):

$$P(m_{ij} \mid z_{ij}^{(s)} = \tilde{k}, F_c^A, F_c^B) = DS\left(\langle \hat{F}_{c,i}^{A,\tilde{k}}, \hat{F}_{c,j}^{B,\tilde{k}} \rangle\right)$$
(13)

To reduce redundant computation, we only augmented the features with covisible topics. Covisible topics were determined by comparing the topic distributions of the two images. The topic distribution in an image is estimated by aggregating the distributions of all features:

$$\theta_k^A \propto \sum_{i=1}^{|F_c^A|} \theta_{i,k}^A, \quad \theta_k^B \propto \sum_{j=1}^{|F_c^B|} \theta_{j,k}^B$$
(14)

where  $\propto$  denotes the normalization operator. We then calculated the covisible probability by multiplying the two topic distributions as  $\theta_k^{Vis} = \theta_k^A \theta_k^B$ . Finally, the most important topics were selected as the covisible topics for feature augmentation based on probability.

### **Implementation Details**

**Efficient Model Design** To achieve a fast computation, we designed an efficient lightweight network for each coarse-to-fine step. For feature extraction, we applied a standard UNet instead of ResUnet, as in other methods (Zhou, Sattler, and Leal-Taixe 2021; Sun et al. 2021). In the coarse matching step, TopicFM uses a single block of self/cross-attention and shares it across topics to extract the features. This operation is applied only to covisible topics; therefore, it is more efficient than methods that use a multi-block transformer (Sun et al. 2021; Wang et al. 2022). Finally, in the fine matching step, our method applies only a cross-attention layer instead of both self- and cross-attention, as in LoFTR

**Training Loss** The loss function is defined as  $\mathcal{L} = \mathcal{L}_f + \mathcal{L}_c$ , where  $\mathcal{L}_f$  and  $\mathcal{L}_c$  are fine- and coarse-level losses, respectively. We directly adopted the fine-level loss  $\mathcal{L}_f$  of LoFTR (Sun et al. 2021). It considers  $l_2$  loss of fine-level matches with the total variance on a cropped patch.

For coarse-level loss  $\mathcal{L}_c$ , we define a new loss function considering the topic model. Given a set of ground truth matches  $\mathcal{M}_c$  at a coarse level, we label each ground truth pair as one. The loss for the positive samples has the following form.

$$\mathcal{L}_{c}^{pos} = -\sum_{m_{ij} \in \mathcal{M}_{c}} \left( E_{p(z_{ij})} \log P(m_{ij} \mid z_{ij}, F_{c}^{A}, F_{c}^{B}) + \log \sum_{k=1}^{K} \theta_{i,k}^{A} \theta_{j,k}^{B} \right)$$
(15)

where the first term represents the ELBO loss estimated by Eqs. 6 and 7, and the second term is used to enforce the pair on the same topic, which is derived from Eq. 9.

We also needed to add a negative loss to prevent the assignment of all features to a single topic. For each ground truth match  $m_{ij}$ , we sampled N unmatched pairs  $\{m_{in}\}_{n=1}^{N}$  and then defined the negative loss using Eq. 10:

$$\mathcal{L}_{c}^{neg} = -\sum_{m_{ij}} \left( \frac{1}{N} \sum_{n=1}^{N} \log \left( 1 - \sum_{k=1}^{K} \theta_{i,k}^{A} \theta_{n,k}^{B} \right) \right) \quad (16)$$

The final coarse-level loss involves these positive and negative terms,  $\mathcal{L}_c = \mathcal{L}_c^{pos} + \mathcal{L}_c^{neg}$ 

### Experiments

#### **Settings and Datasets**

**Training** We trained the proposed network model on the MegaDepth dataset (Li and Snavely 2018), in which the highest dimension of the image was resized to 800. Compared with state-of-the-art transformer-based models (Sarlin et al. 2020; Sun et al. 2021) (e.g., LoFTR (2021) requires approximately 19GB of GPU), our model is much more efficient. Therefore, we used only four GPUs with 11GB of memory to train the model with a batch size of 4. We implemented our network model in PyTorch, with an initial learning rate of 0.01. For the network hyperparameters, we set

Method	Homo. Est. AUC (%)			#M
	3px	5px	10px	
<b>D2Net</b> (2019) + <b>NN</b>	23.2	35.9	53.6	0.2K
<b>R2D2</b> (Revaud et al. 2019) + <b>NN</b>	50.6	63.9	76.8	0.5K
DISK (2020) + NN	52.3	64.9	78.9	1.1K
<b>SP</b> (2018) + <b>SuperGlue</b> (2020)	53.9	68.4	81.7	0.6K
Sparse-NCNet (2020)	48.9	54.2	67.1	1.0K
<b>DRC-Net</b> (Li et al. 2020)	50.6	56.2	68.3	1.0K
Patch2Pix (2021)	59.3	70.6	81.2	0.7K
LoFTR (Sun et al. 2021)	65.9	75.6	84.6	1.0K
TopicFM (Ours)	67.3	77.0	85.7	1.0K

Table 1: Evaluation of homography estimation on HPatches (Balntas et al. 2017). We compute AUC metrics following Sun et al. (2021). #M denotes the number of estimated matches

Method	Relative Pose Estimation (MegaDepth / ScanNet)			
	$5^{o}$	$10^{o}$	$20^{o}$	
SP + SuperGlue	42.2/16.16	61.2/33.81	76.0/ <b>51.84</b>	
<b>DRC-Net</b> * (2020)	27.0/7.69	43.0/17.93	58.3/30.49	
Patch2Pix* (2021)	41.4/9.59	56.3/20.23	68.3/32.63	
LoFTR* (2021)	52.8/16.88	69.2/33.62	81.2/50.62	
MatchFormer <sup>*</sup> (2022)	52.9/-	69.4/-	82.0/-	
TopicFM* (ours)	54.1/17.34	70.1/34.54	<u>81.6/50.91</u>	

Table 2: Evaluation of relative pose estimation on MegaDepth and ScanNet. We use models trained only on MegaDepth for the coarse-to-fine methods denoted by \*

the number of topics K to 100, threshold of coarse match selection  $\tau$  to 0.2, and number of covisible topics for feature augmentation  $K_{co}$  to 6.

We evaluated the image-matching performance on three application tasks: i) homography estimation, ii) relative pose estimation, and iii) visual localization. All of these experiments used the pre-trained model of MegaDepth without fine-tuning. However, some hyperparameters, including  $\tau$  and  $K_{co}$  can be modified during testing.

#### **Benchmark Performance**

**Homography Estimation** The homography matrix between two images can be estimated by matching correspondences using the algorithm (Hartley and Zisserman 2003). We used the HPatches dataset (Balntas et al. 2017) to estimate the homography matrices. For each image pair, we first warped the four corners of the first image to the second image based on the estimated and ground-truth homographies. We then computed the corner error between the two warped versions (DeTone, Malisiewicz, and Rabinovich 2018). The error was measured using the AUC metric with thresholds of 3, 5, and 10 pixels (Sarlin et al. 2020). To report the results, we followed the same setup as in LoFTR. Table 1 shows the homography estimation performance of our method and the state-of-the-art methods. Our method generally outperformed other methods, demonstrating its effectiveness



Figure 3: Qualitative comparison between our method and other coarse-to-fine methods Patch2Pix and LoFTR. Our method can produce a high number of accurate correspondences in challenging conditions such as large relative viewpoints (MegaDepth) or untextured scenes (ScanNet).

Method	<b>Day</b>	<b>Night</b>	over-
ISDE (2020)	0.23,10 )/(0.3	74 2/96 0/07 4	<u>20 0</u>
ISKF (2020)	87.1/94.7/98.5	/4.5/80.9/9/.4	09.0
$\mathbf{KAPTURE} + \mathbf{R2D2}$	90.0/96.2/99.5	72.3/86.4/97.9	90.4
+ APGeM (2020)	<u>&gt; 010</u> / > 012/>>10		2011
SP + SuperGlue	89.8/ <u>96.1/99.4</u>	77.0/ <u>90.6</u> / <b>100</b>	<u>92.1</u>
Patch2Pix (2021)	86.4/93.0/97.5	72.3/88.5/97.9	89.2
LoFTR (2021)	88.7/95.6/99.0	78.5/ <u>90.6</u> /99.0	91.9
TopicFM (Ours)	<b>90.2</b> /95.9/98.9	<u>77.5</u> / <b>91.1</b> / <u>99.5</u>	92.2

Table 3: Evaluation of visual localization on Aachen Day-Night v1.1 (Zhang, Sattler, and Scaramuzza 2021). We report the results using HLoc pipeline (Sarlin et al. 2019)

**Relative Pose Estimation** To evaluate the imagematching performance, we measured the accuracy of the transformation matrix between the two images. We tested outdoor (MegaDepth (Li and Snavely 2018)) and indoor (ScanNet (Dai et al. 2017)) datasets. Each test set includes 1500 image pairs of images. We set the image resolution to  $640 \times 480$  for ScanNet and resized the highest dimension of the image to 1200 for MegaDepth. Similar to (Sarlin et al. 2020; Sun et al. 2021), we measured the area under the cumulative curve (AUC) of the pose estimation error at thresholds of  $\{5^o, 10^o, 20^o\}$ .

Table 2 shows the AUC results for both MegaDepth and ScanNet datasets. To make a fair comparison on ScanNet, we used models trained only on MegaDepth for all the coarse-to-fine methods. As shown in Table 2, our method performed better than the other coarse-to-fine baselines for all evaluation metrics. Compared with SuperPoint (SP) (DeTone, Malisiewicz, and Rabinovich 2018) + SuperGlue (Sarlin et al. 2020), our method had a worse performance only at 20° of AUC on the ScanNet. The main reason for this is that SuperGlue is trained directly on the ScanNet. However, TopicFM was still better than SP+SuperGlue. We provide a detailed comparison with additional baselines for ScanNet in the Supplementary Material.



Figure 4: Topic visualization across images and datasets. Our method can model a specific kind of structure by a topic that then supports the matching process effectively, as described in the method section.

Method	DUC1	DUC2	over-
Wiethou	$(0.25, 10^{\circ})/(0.5, 10^{\circ})/(1.0, 10^{\circ})$		all
<b>ISRF</b> (2020)	39.4/58.1/70.2	41.2/61.1/69.5	56.6
<b>KAPTURE</b> (2020) +	41.4/60.1/73.7	47.3/67.2/73.3	60.5
<b>R2D2</b> (2019)	11. 00000000000	11.5/07.2/75.5	00.5
SP + SuperGlue	<u>49.0</u> /68.7/80.8	53.4/ <b>77.1</b> /82.4	68.6
Patch2Pix (2021)	44.4/66.7/78.3	49.6/64.9/72.5	62.7
LoFTR (2021)	47.5/72.2/84.8	<u>54.2/74.8</u> / <b>85.5</b>	<u>69.8</u>
<b>CoTR</b> (2021)	41.9/61.1/73.2	42.7/67.9/75.6	60.4
MatchFormer (2022)	46.5/ <u>73.2</u> / <u>85.9</u>	<b>55.7</b> /71.8/81.7	69.1
TopicFM (Ours)	52.0/74.7/87.4	53.4/ <u>74.8/83.2</u>	70.9

Table 4: Visual localization on InLoc dataset (Taira et al. 2018) using HLoc pipleline. We achieve best performance in overall.

**Visual Localization** Unlike relative pose estimation, visual localization aims to estimate a camera pose for each image in a global coordinate system; however, it involves several steps. First, the pipeline builds a 3D structure of the scene from a set of database images. Next, given an input query image, it registers this image into the database and finds a set of 2D-3D matches that are then used to output the pose of the query image. Finding correspondences plays an important role in these steps. Therefore, we plugged the matching method into a visual localization pipeline to evaluate the matching performance. Following Patch2Pix, we use a full localization pipeline with HLoc (Sarlin et al. 2019). The benchmark datasets were the Aachen Day-Night v1.1 containing outdoor images and the InLoc dataset with indoor scenes.

Tables 3 and 4 present the results for the Aachen v1.1 (Zhang, Sattler, and Scaramuzza 2021) and InLoc (Taira et al. 2018) datasets, respectively. Our method achieved competitive performance on both benchmarks compared with state-of-the-art baselines. As shown in Table 3, Top-icFM had a similar overall performance to SP+SuperGlue. SP and SuperGlue are trained by leveraging different types of datasets with various shapes and scenes, such as MS-COCO 2014 (Lin et al. 2014) (SP), synthetic shapes (SP), and MegaDepth (SuperGlue). Compared with the second-

best LoFTR method, our overall result was slightly better. The main reason for achieving a satisfactory performance of LoFTR is that it was fine-tuned by augmenting the color images of MegaDepth to fit the nighttime images. In contrast to all the aforementioned setups, our method uses only a unified model trained on MegaDepth. This demonstrated the robustness of the proposed architecture. Similarly, for the InLoc evaluation shown in Table 4, our method is better for all baselines on the DUC1 set with a large margin, although it is worse on the DUC2 set. However, we still achieved the best performance on average.

### **Interpretability Visualization**

We visualized the inferred topics to demonstrate the interpretability of the proposed model. As shown in Fig. 4, our method can partition the contents of an image into different types of spatial structures, in which the same semantic instances are assigned to the same topic. For instance, the topic "human" is marked in green color in the first image pair of MegaDepth and Aachen; the "tree" is marked in orange, and the "ground" is in blue. Different parts of a building, such as roofs, windows, and pillars, are separated into different topics. This phenomenon was repeated across images of MegaDepth and Aachen Day-Night, demonstrating the effectiveness of our topic modeling and inference modules. Notably, as illustrated in the third image pair of the first two rows in Fig. 4, our method focuses on the covisible structures in the same topic (marked with color) and ignores the non-overlapping information (marked without color). Although TopicFM was trained on the outdoor dataset MegaDepth, it could still generalize well on the indoor dataset ScanNet, as shown in the last row of Fig. 4.

#### Conclusion

We introduced a novel architecture using latent semantic modeling for image matching. Our method can learn a powerful representation without high computational power by leveraging adequate context information in latent topics. As a result, the proposed method is robust, interpretable, and efficient compared with state-of-the-art methods.

### Acknowledgments

This work was supported by the National Research Foundation of Korea (NRF) funded by the Ministry of Education under Grant 2016R1D1A1B01013573; and the Industrial Strategic Technology Development Program (No. 20007058, Development of safe and comfortable human augmentation hybrid robot suit) funded by the Ministry of Trade, Industry, & Energy (MOTIE, Korea).

### References

Balntas, V.; Lenc, K.; Vedaldi, A.; and Mikolajczyk, K. 2017. HPatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *CVPR*, 5173–5182.

Bau, D.; Zhu, J.-Y.; Strobelt, H.; Zhou, B.; Tenenbaum, J. B.; Freeman, W. T.; and Torralba, A. 2018. Gan dissection: Visualizing and understanding generative adversarial networks. *arXiv preprint arXiv:1811.10597*.

Bay, H.; Ess, A.; Tuytelaars, T.; and Van Gool, L. 2008. Speeded-up robust features (SURF). *Computer vision and image understanding*, 110(3): 346–359.

Bhowmik, A.; Gumhold, S.; Rother, C.; and Brachmann, E. 2020. Reinforced feature points: Optimizing feature detection and description for a high-level task. In *CVPR*, 4948–4957.

Bian, J.; Lin, W.-Y.; Matsushita, Y.; Yeung, S.-K.; Nguyen, T.-D.; and Cheng, M.-M. 2017. Gms: Grid-based motion statistics for fast, ultra-robust feature correspondence. In *CVPR*, 4181–4190.

Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan): 993–1022.

Calonder, M.; Lepetit, V.; Strecha, C.; and Fua, P. 2010. Brief: Binary robust independent elementary features. In *ECCV*, 778–792. Springer.

Chefer, H.; Gur, S.; and Wolf, L. 2021. Transformer interpretability beyond attention visualization. In *CVPR*, 782– 791.

Chen, H.; Luo, Z.; Zhang, J.; Zhou, L.; Bai, X.; Hu, Z.; Tai, C.-L.; and Quan, L. 2021. Learning to match features with seeded graph matching network. In *ICCV*, 6301–6310.

Csurka, G.; Dance, C.; Fan, L.; Willamowski, J.; and Bray, C. 2004. Visual categorization with bags of keypoints. In *Workshop on statistical learning in computer vision, ECCV*, volume 1, 1–2. Prague.

Cuturi, M. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *NeurIPS*, 26.

Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 5828–5839.

DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2018. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 224–236.

Dusmanu, M.; Rocco, I.; Pajdla, T.; Pollefeys, M.; Sivic, J.; Torii, A.; and Sattler, T. 2019. D2-net: A trainable cnn for joint description and detection of local features. In CVPR, 8092–8101.

Efe, U.; Ince, K. G.; and Alatan, A. A. 2021. Effect of Parameter Optimization on Classical and Learning-based Image Matching Methods. In *CVPR*, 2506–2513.

Fischler, M. A.; and Bolles, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6): 381–395.

Förstner, W.; Dickscheid, T.; and Schindler, F. 2009. Detecting interpretable and accurate scale-invariant keypoints. In *ICCV*, 2256–2263. IEEE.

Hartley, R.; and Zisserman, A. 2003. *Multiple view geometry in computer vision*. Cambridge university press.

He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *CVPR*, 2961–2969.

Humenberger, M.; Cabon, Y.; Guerin, N.; Morat, J.; Revaud, J.; Rerole, P.; Pion, N.; de Souza, C.; Leroy, V.; and Csurka, G. 2020. Robust image retrieval-based visual localization using kapture. *arXiv preprint arXiv:2007.13867*.

Jiang, W.; Trulls, E.; Hosang, J.; Tagliasacchi, A.; and Yi, K. M. 2021. Cotr: Correspondence transformer for matching across images. In *CVPR*, 6207–6217.

Li, X.; Han, K.; Li, S.; and Prisacariu, V. 2020. Dualresolution correspondence networks. *NeurIPS*, 33: 17346– 17357.

Li, Z.; and Snavely, N. 2018. Megadepth: Learning singleview depth prediction from internet photos. In *CVPR*, 2041– 2050.

Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017a. Feature pyramid networks for object detection. In *CVPR*, 2117–2125.

Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*, 740–755. Springer.

Lin, W.-Y.; Wang, F.; Cheng, M.-M.; Yeung, S.-K.; Torr, P. H.; Do, M. N.; and Lu, J. 2017b. CODE: Coherence based decision boundaries for feature correspondence. *TPAMI*, 40(1): 34–47.

Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*, 3431–3440.

Lowe, D. G. 2004. Distinctive image features from scaleinvariant keypoints. *International journal of computer vision*, 60(2): 91–110.

Luo, Z.; Shen, T.; Zhou, L.; Zhang, J.; Yao, Y.; Li, S.; Fang, T.; and Quan, L. 2019. Contextdesc: Local descriptor augmentation with cross-modality context. In *CVPR*, 2527–2536.

Luo, Z.; Zhou, L.; Bai, X.; Chen, H.; Zhang, J.; Yao, Y.; Li, S.; Fang, T.; and Quan, L. 2020. Aslfeat: Learning local features of accurate shape and localization. In *CVPR*, 6589–6598.

Ma, J.; Jiang, X.; Fan, A.; Jiang, J.; and Yan, J. 2021. Image matching from handcrafted to deep features: A survey. *International Journal of Computer Vision*, 129(1): 23–79.

Melekhov, I.; Brostow, G. J.; Kannala, J.; and Turmukhambetov, D. 2020. Image stylization for robust features. *arXiv* preprint arXiv:2008.06959.

Muja, M.; and Lowe, D. G. 2014. Scalable nearest neighbor algorithms for high dimensional data. *TPAMI*, 36(11): 2227–2240.

Mur-Artal, R.; Montiel, J. M. M.; and Tardos, J. D. 2015. ORB-SLAM: a versatile and accurate monocular SLAM system. *IEEE transactions on robotics*, 31(5): 1147–1163.

Ono, Y.; Trulls, E.; Fua, P.; and Yi, K. M. 2018. LF-Net: Learning local features from images. *NeurIPS*, 31.

Revaud, J.; Weinzaepfel, P.; De Souza, C.; Pion, N.; Csurka, G.; Cabon, Y.; and Humenberger, M. 2019. R2D2: repeatable and reliable detector and descriptor. *arXiv preprint arXiv:1906.06195*.

Rocco, I.; Arandjelović, R.; and Sivic, J. 2020. Efficient neighbourhood consensus networks via submanifold sparse convolutions. In *ECCV*, 605–621. Springer.

Rocco, I.; Cimpoi, M.; Arandjelović, R.; Torii, A.; Pajdla, T.; and Sivic, J. 2018. Neighbourhood consensus networks. *NeurIPS*, 31.

Rublee, E.; Rabaud, V.; Konolige, K.; and Bradski, G. 2011. ORB: An efficient alternative to SIFT or SURF. In *ICCV*, 2564–2571. Ieee.

Sarlin, P.-E.; Cadena, C.; Siegwart, R.; and Dymczyk, M. 2019. From coarse to fine: Robust hierarchical localization at large scale. In *CVPR*, 12716–12725.

Sarlin, P.-E.; DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2020. Superglue: Learning feature matching with graph neural networks. In *CVPR*, 4938–4947.

Sattler, T.; Leibe, B.; and Kobbelt, L. 2012. Improving image-based localization by active correspondence search. In *ECCV*, 752–765. Springer.

Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *ICCV*, 618–626.

Shi, Y.; Cai, J.-X.; Shavit, Y.; Mu, T.-J.; Feng, W.; and Zhang, K. 2022. ClusterGNN: Cluster-based Coarse-to-Fine Graph Neural Network for Efficient Feature Matching. In *CVPR*, 12517–12526.

Sivic, J.; and Zisserman, A. 2003. Video Google: A text retrieval approach to object matching in videos. In *ICCV*, volume 3, 1470–1470. IEEE Computer Society.

Strudel, R.; Garcia, R.; Laptev, I.; and Schmid, C. 2021. Segmenter: Transformer for semantic segmentation. In *ICCV*, 7262–7272.

Sun, J.; Shen, Z.; Wang, Y.; Bao, H.; and Zhou, X. 2021. LoFTR: Detector-free local feature matching with transformers. In *CVPR*, 8922–8931.

Taira, H.; Okutomi, M.; Sattler, T.; Cimpoi, M.; Pollefeys, M.; Sivic, J.; Pajdla, T.; and Torii, A. 2018. InLoc: Indoor

visual localization with dense matching and view synthesis. In *CVPR*, 7199–7209.

Tyszkiewicz, M.; Fua, P.; and Trulls, E. 2020. DISK: Learning local features with policy gradient. *NeurIPS*, 33: 14254– 14265.

Wang, J.; Liu, H.; Wang, X.; and Jing, L. 2021. Interpretable image recognition by constructing transparent embedding space. In *ICCV*, 895–904.

Wang, Q.; Zhang, J.; Yang, K.; Peng, K.; and Stiefelhagen, R. 2022. MatchFormer: Interleaving Attention in Transformers for Feature Matching. *arXiv preprint arXiv:2203.09645*.

Williford, J. R.; May, B. B.; and Byrne, J. 2020. Explainable face recognition. In *ECCV*, 248–263. Springer.

Yan, X.; Guo, J.; Lan, Y.; and Cheng, X. 2013. A biterm topic model for short texts. In *Proceedings of the 22nd international conference on World Wide Web*, 1445–1456.

Yi, K. M.; Trulls, E.; Lepetit, V.; and Fua, P. 2016. Lift: Learned invariant feature transform. In *ECCV*, 467–483. Springer.

Yi, K. M.; Trulls, E.; Ono, Y.; Lepetit, V.; Salzmann, M.; and Fua, P. 2018. Learning to find good correspondences. In *CVPR*, 2666–2674.

Yuan, Y.; Chen, X.; and Wang, J. 2020. Object-contextual representations for semantic segmentation. In *ECCV*, 173–190. Springer.

Zhang, J.; Sun, D.; Luo, Z.; Yao, A.; Zhou, L.; Shen, T.; Chen, Y.; Quan, L.; and Liao, H. 2019. Learning two-view correspondences and geometry using order-aware network. In *ICCV*, 5845–5854.

Zhang, Z.; Sattler, T.; and Scaramuzza, D. 2021. Reference pose generation for long-term visual localization via learned features and view synthesis. *International Journal of Computer Vision*, 129(4): 821–844.

Zhao, W.; Rao, Y.; Wang, Z.; Lu, J.; and Zhou, J. 2021. Towards interpretable deep metric learning with structural matching. In *ICCV*, 9887–9896.

Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *CVPR*, 2921–2929.

Zhou, Q.; Sattler, T.; and Leal-Taixe, L. 2021. Patch2pix: Epipolar-guided pixel-level correspondences. In *CVPR*, 4669–4678.