

GeoMoE: Divide-and-Conquer Motion Field Modeling with Mixture-of-Experts for Two-View Geometry

Jiajun Le, Jiayi Ma *

Electronic Information School, Wuhan University, Wuhan 430072, China
 jiajunle01@gmail.com, jyama2010@gmail.com

Abstract

Recent progress in two-view geometry increasingly emphasizes enforcing smoothness and global consistency priors when estimating motion fields between pairs of images. However, in complex real-world scenes, characterized by extreme viewpoint and scale changes as well as pronounced depth discontinuities, the motion field often exhibits diverse and heterogeneous motion patterns. Most existing methods lack targeted modeling strategies and fail to explicitly account for this variability, resulting in estimated motion fields that diverge from their true underlying structure and distribution. We observe that Mixture-of-Experts (MoE) can assign dedicated experts to motion sub-fields, enabling a divide-and-conquer strategy for heterogeneous motion patterns. Building on this insight, we re-architect motion field modeling in two-view geometry with GeoMoE, a streamlined framework. Specifically, we first devise a Probabilistic Prior-Guided Decomposition strategy that exploits inlier probability signals to perform a structure-aware decomposition of the motion field into heterogeneous sub-fields, sharply curbing outlier-induced bias. Next, we introduce an MoE-Enhanced Bi-Path Rectifier that enhances each sub-field along spatial-context and channel-semantic paths and routes it to a customized expert for targeted modeling, thereby decoupling heterogeneous motion regimes, suppressing cross-sub-field interference and representational entanglement, and yielding fine-grained motion-field rectification. With this minimalist design, GeoMoE outperforms prior state-of-the-art methods in relative pose and homography estimation and shows strong generalization.

Code — <https://github.com/JiajunLe/GeoMoE>

Introduction

Establishing precise and reliable correspondences between two images remains a critical and enduring challenge in computer vision, underpinning a variety of advanced vision tasks such as image stitching (Brown and Lowe 2007), Simultaneous Localization and Mapping (Mur-Artal, Martinez Montiel, and Tardos 2015), and Structure from Motion (Schonberger and Frahm 2016). Typically, these

*Corresponding author
 Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

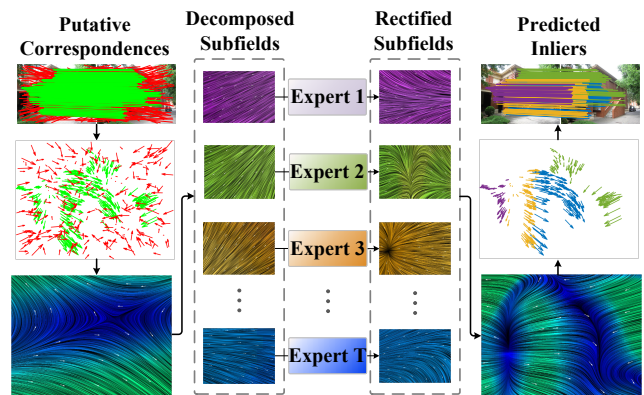


Figure 1: Illustration of MoE modeling heterogeneous motion patterns. Red lines denote outliers, while green lines represent inliers. Different colors in the filtered motion field indicate distinct motion sub-fields, each corresponding to a specific heterogeneous motion pattern.

correspondences are inferred based on the similarity between local descriptors (e.g., SIFT (Lowe 2004), SuperPoint (DeTone, Malisiewicz, and Rabinovich 2018)) extracted by detectors. Nevertheless, in complex, real-world scenarios characterized by significant viewpoint changes and illumination variations, the discriminative capability of these descriptors is often severely compromised, resulting in numerous incorrect matches (i.e., outliers). Consequently, such unreliable matches undermine the accuracy and robustness required for subsequent geometric estimation tasks. Therefore, effectively rejecting outliers while preserving accurate correspondences (i.e., inliers) has emerged as an essential research issue, which is the focus of this paper.

The rapid advances in deep learning over recent years have significantly accelerated research into two-view geometry problems. PointCN (Yi et al. 2018) serves as a flagship method, employing an MLP backbone to recast outlier rejection as a binary classification task distinguishing outliers from inliers. Successor methods derived from PointCN, such as OANet (Zhang et al. 2019), NCMNet (Liu and Yang 2023), U-Match (Li, Zhang, and Ma 2023), and BCLNet (Miao et al. 2024), have continued to demonstrate consistently improved performance. Another line of

work takes a more natural and structured perspective by leveraging consistency and smoothness priors that arise from the uniform deformation between image pairs. Recent methods, including LMCNet (Liu et al. 2021), ConvMatch++ (Zhang and Ma 2023), DeMatch (Zhang et al. 2024), and DeMo (Lu et al. 2025), explicitly exploit these priors, validating their effectiveness in improving model performance. Nevertheless, these methods are still subject to notable limitations. Specifically, LMCNet enforces motion consistency via Laplacian regularization but incurs high computational cost due to expensive matrix decomposition. ConvMatch++ structures unordered motion fields into image-like forms for CNN-based context extraction, but faces local over-smoothing and lacks global perceptual awareness. DeMo models global consistency within an RKHS framework (Aronszajn 1950; Baldassarre et al. 2012), while effectively preserving motion discontinuities while mitigating the risk of over-smoothing. Nonetheless, in realistic imaging conditions, extreme visual scenarios, drastic scale variations, and significant depth disparities often cause motion fields to exhibit multiple heterogeneous motion patterns as shown in Fig. 1. A straightforward strategy to address this is to decompose the motion field into sub-fields, as in DeMatch. However, DeMatch applies a uniform modeling strategy to all sub-fields, inevitably overlooking the intrinsic variability and distinct motion characteristics of each. Consequently, an unresolved key issue persists in the literature: *given the pronounced differences in both feature characteristics and distributions across motion sub-fields, how can we develop tailored, adaptive modeling strategies to effectively address their inherent heterogeneity?*

Rather than relying on complex and redundant hand-crafted designs, we advocate Mixture-of-Experts (MoE) (Jacobs et al. 1991) as a natural and efficient alternative. In this framework, a routing network dynamically analyzes the motion characteristics of each sub-field and assigns it to the most suitable expert for dedicated modeling. It allows experts to specialize in distinct motion patterns, preserving the independence of the modeling process and yielding more discriminative features for accurate sub-field rectification.

Building upon this insight, we propose GeoMoE, a novel and streamlined network paradigm aimed at rebuilding motion field refinement for two-view geometry through the lens of Mixture-of-Experts. Specifically, we heuristically design a Probabilistic Prior-Guided Decomposition strategy that leverages inlier probabilities from the preceding layer as priors to guide a soft, structure-aware clustering of the motion field. This process decomposes the motion field into multiple sub-fields with strong scene-level consistency by concentrating on regions exhibiting similar motion patterns, while concurrently mitigating bias caused by uneven outliers. Upon completing the decomposition, we propose a novel MoE-Enhanced Bi-Path Rectifier, which first comprises two complementary paths: a spatial path that captures relative motion cues within the spatial context of each sub-field, and a channel path that enhances feature representations along the semantic dimension. A routing network then assigns each sub-field to the suitable expert based on its feature characteristics and distributions, enabling each expert to specialize in mod-

eling a specific motion pattern. This targeted expert assignment alleviates cross-sub-field information interference and representational entanglement, thereby yielding more precise tailored sub-field rectification. Through synergistic integration, the rectified sub-fields collectively reconstruct a global motion field with enriched detail, enhanced structural clarity, and closer alignment to real-world distributions.

Taken as a whole, as illustrated in Fig. 1, our method forms a unified and streamlined pipeline from decomposition to expert-guided modeling, fully leveraging MoE to address heterogeneous motion fields. Our main contributions are summarized as follows:

- We propose GeoMoE, a groundbreaking two-view geometry network that introduces MoE to enable fine-grained modeling and rectification of motion sub-fields.
- We develop a probabilistic prior-guided decomposition strategy that utilizes inlier probability as a prior to perform structure-aware motion field decomposition, while reducing bias induced by uneven outliers.
- We architect an innovative MoE-enhanced bi-path rectifier that reduces mutual interference and entanglement among sub-fields with diverse motion patterns, facilitating expert-driven, fine-grained modeling and rectification across sub-fields.
- Extensive experiments across multiple public datasets show that GeoMoE surpasses state-of-the-art methods, validating its effectiveness and strong generalizability.

Related Work

Learning-Based Outlier Rejection Methods

PointCN (Yi et al. 2018) pioneered the application of deep learning to the outlier rejection task by reframing it as an inlier/outlier classification problem modeled with MLPs. A series of subsequent works have focused on enhancing feature representations through increasingly sophisticated architectural designs. For example, OANet (Zhang et al. 2019) proposes differentiable pooling to capture local geometry, MS²DG-Net (Dai et al. 2022) introduces a graph-based representation in feature space to model inter-correspondence relationships, and U-Match (Li, Zhang, and Ma 2023) designs a U-shaped architecture to enable hierarchical and multi-scale local context learning. Furthermore, CLNet (Zhao et al. 2021) establishes a local-to-global consistency learning paradigm, progressively preserving reliable matches. This progressive filtering paradigm has been adopted by NCMNet (Liu and Yang 2023), BCLNet (Miao et al. 2024), and VSFormer (Liao et al. 2024).

From another perspective, ConvMatch++ (Zhang and Ma 2023) reinterprets two-view geometry as a motion field estimation problem, structuring unordered correspondences into fields and applying CNNs for local consistency. DeMatch (Zhang et al. 2024) alleviates computational burden by decomposing the global motion field into multiple implicit sub-fields, while DeMo (Lu et al. 2025) introduces learnable deep kernels in RKHS to capture global motion coherence. Nevertheless, these methods fail to account for multiple heterogeneous motion patterns induced by complex




Figure 2: Overall architecture of GeoMoE. Following initialization and the LOC module, high-dimensional motion features ${}^{\ell}F$ are decomposed into multiple motion sub-fields ${}^{\ell}Z'$ via the probabilistic prior-guided decomposition strategy. In the MoE-enhanced bi-path rectifier, each sub-field undergoes bi-path feature enhancement and subsequently is modeled by customized experts within F-MoE. The complete motion field ${}^{\ell+1}F$ is then reconstructed by a GAT module. Finally, the inlier probabilities ${}^{\ell}w$ predicted at each layer are utilized in loss computation.

scene variations, which pose a significant challenge to precise motion field modeling.

Mixture-of-Experts (MoE)

The Mixture-of-Experts (MoE) (Jacobs et al. 1991) framework leverages a routing network to dynamically activate specialized expert networks, enabling task-adaptive representation learning. It has shown strong effectiveness in both vision (Cai, Gu, and Zhang 2018) and language models (Lepikhin et al. 2020), particularly when integrated into Transformer (Fedus, Zoph, and Shazeer 2022) to boost capacity and efficiency. Recent work extends MoE to various tasks: it reconfigures U-Net skip connections for semantic decoding in medical segmentation (Luo et al. 2025); assigns correspondences to shared experts in point cloud registration (Huang et al. 2025) to enhance feature discriminability in overlapping regions; and improves robustness under cross-domain variations in correspondence pruning (Xia et al. 2025). In contrast, we introduce MoE into two-view geometry to address heterogeneous motion patterns arising from scene diversity and complexity at the underlying correspondence level.

Method

Problem Formulation

Following prior work (Liu et al. 2021), given a pair of images (I, I') , we use an off-the-shelf descriptor extraction method (Lowe 2004) to detect keypoints and compute descriptors. Based on descriptor similarity, a putative correspondence set $C = \{c_i | c_i = (x_i, x'_i)\} \in \mathbb{R}^{N \times 4}$ is constructed using a nearest neighbor (NN) matcher, where x_i and x'_i denote the keypoint coordinates normalized by the respective camera intrinsics. Each correspondence c_i defines a motion vector $m_i = (x_i, x'_i - x_i)$, which is then projected into a high-dimensional feature f_i to form the dense motion field $F = \{f_i\} \in \mathbb{R}^{N \times D}$. After the locality orthogonal context module, we apply a probabilistic prior-guided decomposition strategy to decompose the global motion field into multiple structure-aware sub-fields. Each sub-field, enhanced along spatial-context and channel-semantic paths, is subsequently routed by the MoE-enhanced bi-path rectifier

to a customized expert for dedicated modeling and refinement. The rectified sub-fields are then integrated to reconstruct a high-fidelity motion field with fine-grained structure. A lightweight MLP then predicts an inlier/outlier probability w_i for each correspondence. This cycle of local context aggregation, decomposition, refinement, and reconstruction is repeated across L sequential layers, progressively enhancing the motion field estimation. Finally, a weighted eight-point algorithm, guided by the predicted probabilities $w = \{w_i\} \in \mathbb{R}^{N \times 1}$, is used to estimate the essential matrix \hat{E} . The overall process is formalized as:

$$\begin{aligned} w &= g_{\phi}(C), \\ \hat{E} &= h(w, C), \end{aligned} \quad (1)$$

where g_{ϕ} denotes GeoMoE parameterized by ϕ , and h represents the robust weighted eight-point algorithm. As shown in Fig. 2, the overall architecture of GeoMoE is depicted, and we elaborate on its design in the following section.

Motion Initialization

Owing to their low dimensionality, motion vectors m_i lack the expressive power needed to capture deep features. To overcome this, inspired by prior work (Lu et al. 2025), we embed m_i into a high-dimensional space to obtain a richer representation $f_i \in \mathbb{R}^D$ that serves as input to the first layer:

$${}^0 f_i = \mathcal{E}(m_i), \quad i = 1, \dots, N. \quad (2)$$

Here, $\mathcal{E}(\cdot)$ denotes a transformation that elevates motion vectors from low- to high-dimensional space, implemented via a convolutional layer followed by ContextNorm (Yi et al. 2018). This embedding improves feature expressivity and lays a robust foundation for downstream processing.

Mixture-of-Experts for Motion Field (F-MoE)

The Mixture-of-Experts for sub-field modeling comprises a routing network and a set of feed-forward expert networks, as shown in Fig. 3. Given sub-field inputs f_i , the routing network dynamically activates the most relevant experts for specialized modeling. Unlike conventional fully-connected routing, we employ a sparse MLP to better reflect the sparsity and region-specific nature of sub-field rectification. The




Figure 3: Workflow diagram of F-MoE.

routing network is defined as:

$$\mathcal{R} = \text{top-k}(\text{softmax}(\text{MLP}(f_i))), \quad (3)$$

where the MLP shares parameters across all sub-fields, and the top-k operator selects the k most responsive experts based on activation scores. Each expert is a lightweight feed-forward network, tailored for dimensionality reduction and targeted motion field refinement. The final output is a weighted summation of the activated experts' outputs:

$$f'_i = \sum_{j \in \mathbb{N}_k} \mathcal{R}_j \cdot E_j(f_i), \quad (4)$$

where \mathbb{N}_k denote the indices of the top-k experts, E_j denotes the j -th expert, and \mathcal{R}_j is its routing weight.

Probabilistic Prior-Guided Decomposition

The motion field derived from putative correspondences often manifests a spectrum of heterogeneous dynamic patterns due to the complexity of real-world scenes, such as viewpoint changes, scale variations, and pronounced depth discontinuities. Instead of modeling the field holistically, prior work (Zhang et al. 2024) shows that grouping motion vectors with coherent geometric behavior yields a more efficient alternative. However, prior approaches that rely on learnable tokens to implicitly model motion patterns tend to neglect explicit spatial priors, i.e., the natural inductive bias that spatially adjacent motions are likely to share consistent transformations, resulting in physically ambiguous and semantically uninterpretable representations.

To address these limitations, we propose a Probabilistic Prior-Guided Decomposition (PPGD) strategy rooted in the structural regularities of the scene. Specifically, we first apply a pooling operation (Zhang et al. 2019) to coarsely partition the motion field ${}^\ell \mathbf{F} = \{f_i\}_{i=1}^N$ into M preliminary sub-fields ${}^\ell \mathbf{Z} \in \mathbb{R}^{M \times D}$:

$$\begin{aligned} {}^\ell \mathbf{Z} &= \mathbf{S} {}^\ell \mathbf{F}, \\ \mathbf{S} &= \text{softmax}(\text{MLP}({}^\ell \mathbf{F})). \end{aligned} \quad (5)$$

This operation is inherently structure-aware, leveraging localized contextual cues to inform sub-field assignment, as demonstrated in prior work (Zhang et al. 2019). To enhance the expressiveness and attenuate distortions caused by outliers, we incorporate an F-MoE module $\mathcal{M}(\cdot)$ that performs specialized refinement, guided by a probabilistic masking strategy using the previous-layer confidence scores ${}^{\ell-1} \mathbf{w} = \{w_i\} \in \mathbb{R}^N$:

$${}^\ell \mathbf{Z}' = \mathcal{G}(\mathcal{M}({}^\ell \mathbf{Z}), {}^{\ell-1} \mathbf{W} \odot {}^\ell \mathbf{F}), \quad (6)$$

where ${}^{\ell-1} \mathbf{W} \in \mathbb{R}^{N \times D}$ is a broadcasted version of ${}^{\ell-1} \mathbf{w}$, \odot denotes the Hadamard product, and $\mathcal{G}(\cdot)$ refers to the Graph Attention Network (Velickovic et al. 2017). For the initial layer, ${}^0 \mathbf{w}$ is initialized to an all-ones vector, implying no prior suppression. This prior-informed decomposition strategy is capable of yielding a motion field that is both physically plausible and structurally coherent.

MoE-Enhanced Bi-Path Rectifier

Amid substantial heterogeneity in motion patterns across sub-fields, prior approaches (Zhang et al. 2024) that adopt a shared architecture for all sub-fields inherently fail to capture the diverse statistical characteristics and spatial dynamics specific to each region. This uniform treatment overlooks crucial structural discrepancies, resulting in constrained representational capacity and limited modeling fidelity.

To address this, we propose the MoE-enhanced Bi-Path Rectifier (MBPR), a specialized rectification module that routes each sub-field to a set of dedicated expert networks for adaptive refinement and motion field rectification, enabling more tailored and expressive modeling. Specifically, we begin by enriching sub-field representations through a bi-path encoding strategy: the spatial branch models relative motion via graph-based aggregation, while the channel branch extracts high-level semantics by applying global average pooling followed by channel-wise modulation. The two branches are defined as:

$$\begin{aligned} {}^\ell \mathbf{Z}_s &= \mathcal{G}({}^\ell \mathbf{Z}', {}^\ell \mathbf{Z}'), \\ {}^\ell \mathbf{Z}_c &= \text{MLP}(\text{AvgPool}({}^\ell \mathbf{Z}')) {}^\ell \mathbf{Z}'. \end{aligned} \quad (7)$$

The outputs are concatenated and fused via a lightweight MLP to obtain the final bi-path enhanced representation:

$${}^\ell \mathbf{Z}_{bi} = \text{MLP}({}^\ell \mathbf{Z}_c \parallel {}^\ell \mathbf{Z}_s). \quad (8)$$

Next, each enhanced sub-field is dynamically assigned to a subset of expert networks via a sparsely-activated routing mechanism:

$${}^\ell \mathbf{Z}_{moe} = \mathcal{M}({}^\ell \mathbf{Z}_{bi}), \quad (9)$$

where $\mathcal{M}(\cdot)$ denotes the F-MoE module that adaptively assigns sub-fields to specialized experts according to their individual characteristics. Finally, the rectified sub-fields are globally aggregated to reconstruct the updated motion field:

$${}^{\ell+1} \mathbf{F} = \mathcal{G}({}^\ell \mathbf{F}, {}^\ell \mathbf{Z}_{moe}). \quad (10)$$

This staged rectification pipeline progressively refines the global motion representation with enhanced structural clarity, higher spatial fidelity, and stronger alignment to the true underlying motion distribution.

Local Orthogonal Context (LOC)

Before motion field decomposition and sub-field modeling, we incorporate orthogonal contextual cues within KNN-based local neighborhoods, following (Lu et al. 2025). This preparatory step facilitates the detection of local motion discontinuities while suppressing excessive smoothness artifacts. Concretely, each motion vector f_i undergoes dimensionality reduction. A KNN-based neighborhood $\{f_{i,j}\}_{j=0}^K$

Method	YFCC100M (%↑)			SUN3D (%↑)		
	@5°	@10°	@20°	@5°	@10°	@20°
RANSAC	3.47	9.10	18.60	1.04	3.43	8.75
MAGSAC++	1.51	4.67	11.95	0.88	2.99	8.42
NG-RANSAC	16.47	28.38	42.03	4.22	10.87	21.39
GMS	- / 13.29	- / 24.38	- / 37.83	- / 4.12	- / 10.53	- / 20.82
VFC	- / 17.43	- / 29.98	- / 43.00	- / 5.26	- / 13.05	- / 24.84
PointCN	10.16 / 26.73	24.43 / 44.01	43.31 / 60.49	3.05 / 6.09	10.00 / 15.43	24.06 / 29.74
OANet	15.92 / 27.26	35.93 / 45.93	57.11 / 63.17	5.93 / 6.78	16.91 / 17.10	34.32 / 32.41
NCMNet	26.89 / 32.30	46.19 / 52.29	64.21 / 69.65	6.31 / 7.10	16.84 / 18.56	33.11 / 35.58
UMatch	30.93 / 33.92	52.17 / 53.09	69.75 / 69.45	8.03 / 7.01	20.83 / 17.79	38.75 / 33.57
MS ² DG-Net	20.61 / 31.55	42.90 / 50.94	64.26 / 68.34	5.89 / 6.95	16.85 / 17.66	34.25 / 33.38
DeMatch	30.89 / 32.98	52.67 / 52.37	70.33 / 69.01	<u>9.31</u> / 7.44	23.10 / 18.66	41.55 / 34.78
DeMo	<u>32.57</u> / 34.63	<u>54.82</u> / <u>54.42</u>	<u>72.51</u> / <u>71.15</u>	<u>9.31</u> / <u>7.48</u>	<u>23.18</u> / <u>18.85</u>	<u>41.75</u> / <u>35.31</u>
GeoMoE	34.21 / <u>34.52</u>	56.25 / <u>54.72</u>	73.53 / <u>71.63</u>	9.98 / <u>7.54</u>	24.05 / <u>18.87</u>	42.71 / 35.22

Table 1: Quantitative results of relative pose estimation. AUC is reported under 5°, 10°, and 20° thresholds with the weighted eight-point algorithm / RANSAC. Best and second-best results are shown in bold and underlined, respectively.

is then constructed to capture its local structure, and pairwise differences $\mathbf{h}_{i,j} = \mathbf{f}_i - \mathbf{f}_{i,j}$ are computed. To jointly capture spatial correlations and feature-level semantics, we utilize a two-stage aggregation scheme composed of lightweight MLPs:

$$\begin{aligned} \mathbf{H}_f &= \text{MLP}(\mathbf{H}) + \mathbf{H}, \\ \mathbf{F}_{local} &= \text{MLP}(\text{MLP}(\mathbf{H}_f^T) + \mathbf{H}_f^T). \end{aligned} \quad (11)$$

The final output \mathbf{F}_{local} is projected into a higher-dimensional space, equipping it with enriched local representations for downstream motion modeling.

Loss Functions

During supervised training, we adopt the widely used binary cross-entropy loss for inlier/outlier classification, together with a regression loss for essential matrix estimation. Importantly, to address the often-overlooked issue of expert imbalance in the Mixture-of-Experts (MoE) framework (Fedus, Zoph, and Shazeer 2022), we design a load-balancing regularization term:

$$\mathcal{L}_{load} = \frac{1}{T} \sum_{t=1}^T \left(\frac{1}{N} \sum_{i=1}^N \ell_{\mathbf{r}_{i,t}} \right)^2, \quad (12)$$

where $\ell_{\mathbf{r}_{i,t}}$ denotes the activation probability of the t -th expert for the i -th sub-field, and T is the total number of experts. This regularization encourages balanced expert utilization, promoting more stable training and better generalization. The overall loss function is defined as:

$$\mathcal{L} = \sum_{\ell=0}^{L-1} \mathcal{L}_{cls}(\mathbf{w}, \ell \mathbf{w}) + \mu \mathcal{L}_{reg}(\mathbf{E}, \ell \hat{\mathbf{E}}) + \frac{\beta}{L} \mathcal{L}_{load}, \quad (13)$$

where β and μ are weighting coefficients that balance the contributions of each term.

Implementation Details

In our implementation, GeoMoE is unrolled for $L = 8$ layers, with motions represented in a 128-dimensional space (i.e., $D = 128$). For each image pair, up to $N = 2000$ putative matches are obtained using SIFT followed by NN matcher. The motion field is decomposed into 48 sub-fields. Each F-MoE comprises 4 experts, and every sub-field is

routed to 2 of them for parallel processing. For loss balancing, the regularization weight μ is initially set to 0 and increased to 0.5 after 20k iterations. The weight β for expert load balancing remains fixed at 0.01.

Experiments

Relative Pose Estimation

Many vision applications require estimating the relative camera pose, i.e., rotation and translation, between images, thereby imposing stringent demands on high-quality inliers.

Comparisons with Outlier Rejection Methods *Datasets and Evaluation Protocols:* Following the standard evaluation protocol (Zhang et al. 2019), we conduct relative pose estimation on the outdoor dataset YFCC100M (Thomee et al. 2016) and the indoor dataset SUN3D (Xiao, Owens, and Torralba 2013). We report the Area Under the cumulative error Curve (AUC) for the maximum angular error in both rotation and translation under multiple thresholds.

Baselines: We benchmark our method against a broad range of baselines, including both classical methods (GMS (Bian et al. 2017), and VFC (Ma et al. 2014)) and deep learning-based approaches (PointCN (Yi et al. 2018), OANet (Zhang et al. 2019), NCMNet (Liu and Yang 2023), U-Match (Li, Zhang, and Ma 2023), MS²DG-Net (Dai et al. 2022), DeMatch (Zhang et al. 2024), and DeMo (Lu et al. 2025)). In addition, we evaluate against the RANSAC family of methods, namely RANSAC (Fischler and Bolles 1981), MAGSAC++ (Barath et al. 2020), and NG-RANSAC (Brachmann and Rother 2019).

Results: All quantitative evaluation results are reported in Table 1. As shown, GeoMoE consistently delivers state-of-the-art performance across most evaluation thresholds in both indoor and outdoor scenarios, underscoring its empirical effectiveness and clear advantage over existing methods. Notably, under the 5° threshold, our method achieves a 10.75% improvement over DeMatch when integrated with the weighted eight-point algorithm. This result highlights the effectiveness and strong potential of employing MoE to assign individual sub-fields to customized experts for fine-

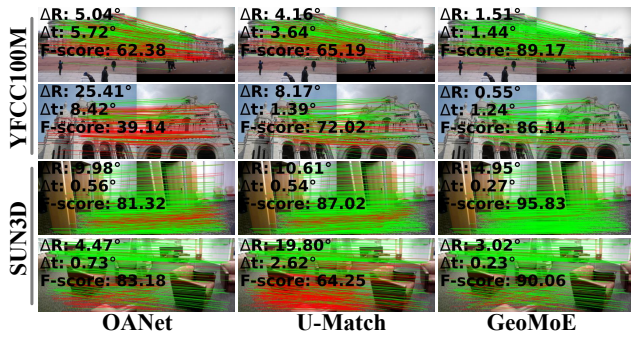


Figure 4: Qualitative comparison of outlier rejection. From left to right: matching results of OANet, U-Match, and GeoMoE. Green and red lines denote correct and incorrect matches, respectively. Estimated relative poses are shown in the top-left corner. Zoom in for better visibility.

Matcher	Filter	YFCC100M (outdoor) (% \uparrow)		
		AUC@5 $^\circ$	AUC@10 $^\circ$	AUC@20 $^\circ$
SuperPoint+NN	/	6.83	15.99	29.09
	GeoMoE	29.61	48.70	66.06
SuperPoint+SuperGlue	/	38.24	58.47	74.64
	GeoMoE	39.24	59.78	75.70
LoFTR	/	40.42	60.97	76.77
	GeoMoE	41.66	62.58	78.16

Table 2: Compatibility with learning-based matchers on YFCC100M. AUC at angular error thresholds of 5 $^\circ$, 10 $^\circ$, and 20 $^\circ$ for relative pose estimation is reported.

grained motion modeling. In addition, qualitative results in Fig. 4 show that GeoMoE can effectively suppress outliers and preserve inliers under challenging conditions such as large viewpoint changes and significant depth disparities.

Compatibility with Learning-based Matchers As a general-purpose outlier rejection method, GeoMoE can accept an arbitrary set of putative correspondences and effectively filter out false matches. Leveraging this flexibility, we integrate GeoMoE with several learning-based matchers on YFCC100M to assess its compatibility and potential to further improve matching accuracy. Specifically, we adopt the learning-based descriptor SuperPoint (DeTone, Malisiewicz, and Rabinovich 2018), paired with NN Matcher and the sparse matcher SuperGlue (Sarlin et al. 2020), and the detector-free matcher LoFTR (Sun et al. 2021). All models are evaluated using official pretrained weights. For each image pair, $2k$ keypoints are extracted using SuperPoint. Table 2 presents relative pose estimation results with and without GeoMoE. As shown, GeoMoE consistently improves the pose estimation accuracy of all matchers across all evaluation thresholds, demonstrating strong compatibility and generalization capability.

Homography Estimation

Datasets and Evaluation Protocols We conducted homography estimation experiments on the HPatches (Balntas et al. 2017) benchmark, which consists of 116 image

Method	HPatches (% \uparrow)		
	Acc.@3px	Acc.@5px	Acc.@10px
PointCN	38.97 / 68.97	51.38 / 83.10	65.52 / <u>92.59</u>
OANet	39.83 / 68.62	52.41 / 82.59	63.10 / 91.90
LMCNet	47.93 / 72.24	58.62 / 85.34	69.83 / <u>92.59</u>
UMatch	47.07 / 70.34	57.41 / 85.00	70.00 / 91.72
DeMatch	46.90 / 70.69	60.52 / 82.76	71.55 / <u>92.59</u>
DeMo	53.79 / 73.10	63.79 / 85.52	75.69 / <u>92.59</u>
GeoMoE	53.83 / 74.02	64.92 / 85.59	<u>74.60</u> / 95.20

Table 3: Comparison of homography estimation results. The estimation accuracy (Acc.) at different pixel thresholds is reported using DLT / RANSAC as post-processing methods.

Method	3DMatch (Zeng et al. 2017)		
	RR (% \uparrow)	Prec (% \uparrow)	F-score (% \uparrow)
RANSAC	66.77	65.87	62.24
PointCN	73.66	64.60	61.49
OANet	76.90	67.91	68.58
DeMatch	77.29	68.74	69.82
U-Match	76.30	67.61	68.64
DeMo	77.10	68.94	69.58
PointDSC	<u>77.50</u>	<u>69.34</u>	70.68
GeoMoE	78.33	69.65	<u>70.65</u>

Table 4: Results of point cloud registration on 3DMatch.

sequences, each comprising one reference image and five query images with corresponding ground-truth homography matrices. For each image, up to $4k$ keypoints are extracted, and putative correspondences are established via an NN matcher. Homographies are estimated using the Direct Linear Transform (DLT) algorithm and RANSAC as post-processing. We assess performance by computing the average corner error from warping with estimated and ground-truth homographies, and report the percentage of correct estimates under 3, 5, and 10 pixel thresholds.

Results Table 3 presents the quantitative results of homography estimation on the HPatches dataset. GeoMoE consistently achieves top performance across most threshold settings and post-processing strategies, clearly demonstrating its effectiveness in outlier rejection.

Point Cloud Registration

Point cloud registration aims to estimate the optimal rigid transformation between two point clouds.

Datasets and Evaluation Protocols We perform point cloud registration experiments on the 3DMatch dataset (Zeng et al. 2017). Due to the higher dimensionality of point cloud data, all methods, except PointDSC (Bai et al. 2021), are retrained on 3DMatch. For evaluation, we follow the registration recall (RR) metric from PointDSC, which considers an alignment successful if the rotation error is below 20 $^\circ$ and the translation error is less than 30cm. We additionally report inlier precision (Prec) and F-score to quantify the effectiveness of outlier filtering.

Method	YFCC100M (% \uparrow)			SUN3D (% \uparrow)			
	RootSIFT	LIFT	SuperPoint	SIFT	RootSIFT	LIFT	SuperPoint
PointCN	24.71	15.31	14.94	1.56	1.71	1.99	3.52
OANet	36.46	28.83	20.54	3.37	3.59	3.20	3.39
UMatch	52.49	38.86	25.90	7.63	7.66	4.84	3.45
BCLNet	47.04	37.88	26.48	5.33	5.50	3.11	1.96
DeMatch	53.32	42.27	<u>29.84</u>	7.12	7.31	6.35	<u>5.18</u>
DeMo	<u>55.56</u>	<u>44.23</u>	<u>28.45</u>	<u>8.22</u>	<u>8.53</u>	<u>6.89</u>	5.01
GeoMoE	56.74	45.00	30.00	8.52	8.84	7.39	6.11

Table 5: Generalization ability. AUC@10 $^\circ$ on YFCC100M is reported using the weighted eight-point algorithm.

Method	Size(M) (\downarrow)	Flops(G) (\downarrow)	YFCC100M		
			Mem(MB)	Time(ms)	AUC@5 $^\circ$ (\uparrow)
LMCNet	0.926	2.323	121.38	301.54	22.35
DeMatch	5.853	2.346	109.21	46.16	30.89
U-Match	7.764	3.740	144.14	65.34	30.93
BCLNet	4.460	8.971	141.01	59.25	27.04
DeMo	4.520	3.612	90.22	54.78	32.57
GeoMoE	5.552	3.106	106.73	56.14	34.21

Table 6: Comparative results of computational usage.

Results The quantitative results of point cloud registration are presented in Table 4. GeoMoE achieves superior overall performance, particularly excelling in registration recall, the primary evaluation metric, and notably surpasses PointDSC, a method specifically tailored for point cloud registration. These results highlight the strong cross-task generalization capability of GeoMoE, enabled by its motion field decomposition and expert-driven modeling strategy.

Analysis

Generalization Ability Generalization performance serves as a key indicator of the robustness of a model, gauging its ability to adapt across varying datasets and descriptor types. To evaluate this, we test the model trained with SIFT on YFCC100M under alternative descriptors (RootSIFT (Arandjelović and Zisserman 2012), LIFT (Yi et al. 2016), SuperPoint) across both YFCC100M and SUN3D. Except for SuperPoint, which yields $1k$ keypoints, all descriptors extract $2k$ keypoints. As shown in Table 5, GeoMoE consistently delivers state-of-the-art results across all combinations, highlighting the remarkable adaptability and the superior generalization of its expert-driven sub-field modeling, tailored to capture localized motion dynamics.

Computation Usage We systematically evaluate the computational cost of GeoMoE in comparison to existing baselines, in terms of model parameter size, FLOPs, peak memory usage (Mem), and average inference time per image on the YFCC100M dataset. As shown in Table 6, despite incorporating a Mixture-of-Experts architecture, well known for increasing parameter counts, GeoMoE maintains a compact computational profile while achieving state-of-the-art performance. These results demonstrate that GeoMoE is both computationally efficient and deployment-friendly.

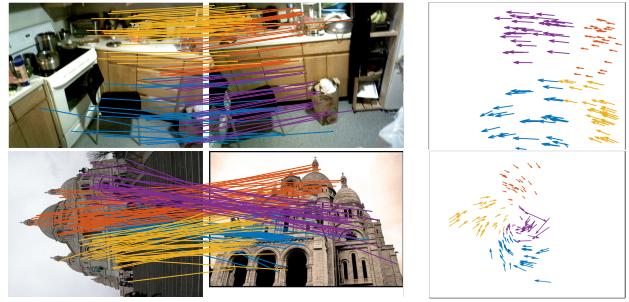


Figure 5: Visualization of expert assignments in F-MoE.

Num.	Prob.	S-Path.	C-Path.	F-MoE	@5 $^\circ$	@10 $^\circ$	@20 $^\circ$
(i)					31.69	53.65	71.24
(ii)	✓				32.37	53.65	71.16
(iii)	✓			✓	32.78	54.60	72.16
(iv)	✓	✓		✓	33.14	55.31	72.89
(v)	✓	✓	✓	✓	32.59	54.92	72.58
(vi)	✓	✓	✓	✓	34.21	56.25	73.53

Table 7: Ablation study. AUC on YFCC100M with the weighted eight-point algorithm are shown. “Prob.” denotes probability injection in PPGD. “S-Path” and “C-Path” represent the spatial and channel paths in MBPR, respectively.

Visualization of MoE Behavior Across Sub-Fields To more clearly elucidate the functionality of F-MoE, Fig. 5 visualizes the correspondence between each expert in the final network layer and its representative subregion, along with the associated motion vectors. For each expert, a sub-field is randomly selected from its response region according to the expert response probabilities \mathcal{R} defined in Eq. 3. Distinct colors indicate assignment to different customized experts. As illustrated, each sub-field exhibits a characteristic motion pattern, attending to different objects (e.g., first row) or spatial regions (e.g., second row), and is delegated to the most suitable expert for tailored modeling and refinement.

Ablation Studies We conduct ablation experiments on the YFCC100M dataset, with results reported in Table 7. The last row represents the complete GeoMoE model. As observed from rows (i)-(v), incorporating probability injection during the decomposition process, bi-path enhancement, and F-MoE for sub-field-specific modeling in MBPR each contributes to incremental performance gains, collectively validating the effectiveness of our designs.

Conclusion

We propose GeoMoE, a unified two-view geometry framework that employs a MoE paradigm to model diverse motion patterns, particularly under extreme viewpoints, large-scale variations, and significant depth disparities. Guided by probabilistic priors, GeoMoE decomposes the motion field into sub-fields and assigns each to specialized experts for precise modeling, achieving state-of-the-art performance with a compact design and high computational efficiency.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant no. 62276192.

References

- Arandjelović, R.; and Zisserman, A. 2012. Three things everyone should know to improve object retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2911–2918.
- Aronszajn, N. 1950. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68(3): 337–404.
- Bai, X.; Luo, Z.; Zhou, L.; Chen, H.; Li, L.; Hu, Z.; Fu, H.; and Tai, C.-L. 2021. PointDSC: Robust Point Cloud Registration Using Deep Spatial Consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 15859–15869.
- Baldassarre, L.; Rosasco, L.; Barla, A.; and Verri, A. 2012. Multi-output learning via spectral filtering. *Machine Learning*, 87: 259–301.
- Balntas, V.; Lenc, K.; Vedaldi, A.; and Mikolajczyk, K. 2017. HPatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5173–5182.
- Barath, D.; Nuskova, J.; Ivashechkin, M.; and Matas, J. 2020. MAGSAC++: A Fast, Reliable and Accurate Robust Estimator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1304–1312.
- 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4181–4190.
- Brachmann, E.; and Rother, C. 2019. Neural-Guided RANSAC: Learning Where to Sample Model Hypotheses. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 4322–4331.
- Brown, M.; and Lowe, D. G. 2007. Automatic Panoramic Image Stitching Using Invariant Features. *International Journal of Computer Vision*, 74(1): 59–73.
- Cai, J.; Gu, S.; and Zhang, L. 2018. Learning a Deep Single Image Contrast Enhancer from Multi-Exposure Images. *IEEE Transactions on Image Processing*, 27(4): 2049–2062.
- Dai, L.; Liu, Y.; Ma, J.; Wei, L.; Lai, T.; Yang, C.; and Chen, R. 2022. MS2DG-Net: Progressive Correspondence Learning via Multiple Sparse Semantics Dynamic Graph. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8973–8982.
- 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 (CVPRW)*, 224–236.
- Fedus, W.; Zoph, B.; and Shazeer, N. 2022. Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity. *Journal of Machine Learning Research*, 23(120): 1–39.
- 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.
- Huang, X.; Huang, Z.; Zuo, Y.; Gong, Y.; Zhang, C.; Liu, D.; and Fang, Y. 2025. PSReg: Prior-Guided Sparse Mixture of Experts for Point Cloud Registration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 3788–3796.
- Jacobs, R. A.; Jordan, M. I.; Nowlan, S. J.; and Hinton, G. E. 1991. Adaptive Mixtures of Local Experts. *Neural Computation*, 3(1): 79–87.
- Lepikhin, D.; Lee, H.; Xu, Y.; Chen, D.; Firat, O.; Huang, Y.; Krikun, M.; Shazeer, N.; and Chen, Z. 2020. GShard: Scaling Giant Models with Conditional Computation and Automatic Sharding. *arXiv preprint*.
- Li, Z.; Zhang, S.; and Ma, J. 2023. U-Match: Two-view Correspondence Learning with Hierarchy-aware Local Context Aggregation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1169–1176.
- Liao, T.; Zhang, X.; Zhao, L.; Wang, T.; and Xiao, G. 2024. VSFormer: Visual-Spatial Fusion Transformer for Correspondence Pruning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 3369–3377.
- Liu, X.; and Yang, J. 2023. Progressive neighbor consistency mining for correspondence pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9527–9537.
- Liu, Y.; Liu, L.; Lin, C.; Dong, Z.; and Wang, W. 2021. Learnable motion coherence for correspondence pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3237–3246.
- Lowe, D. G. 2004. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2): 91–110.
- Lu, Y.; Le, J.; Li, Z.; Yuan, Y.; and Ma, J. 2025. Deep Motion Field Consensus with Learnable Kernels for Two-View Correspondence Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 5829–5837.
- Luo, Z.; Zhu, X.; Zhang, L.; and Sun, B. 2025. Rethinking U-Net: Task-Adaptive Mixture of Skip Connections for Enhanced Medical Image Segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 5874–5882.
- Ma, J.; Zhao, J.; Tian, J.; Yuille, A. L.; and Tu, Z. 2014. Robust point matching via vector field consensus. *IEEE Transactions on Image Processing*, 23(4): 1706–1721.
- Miao, X.; Xiao, G.; Wang, S.; and Yu, J. 2024. BCLNet: Bilateral Consensus Learning for Two-View Correspondence Pruning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4225–4232.

- Mur-Artal, R.; Martinez Montiel, J. M.; and Tardos, J. D. 2015. ORB-SLAM: A Versatile and Accurate Monocular SLAM System. *IEEE Transactions on Robotics*, 31(5): 1147–1163.
- Sarlin, P.-E.; DeTone, D.; Malisiewicz, T.; and Rabinovich, A. 2020. SuperGlue: Learning Feature Matching with Graph Neural Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4938–4947.
- Schonberger, J. L.; and Frahm, J.-M. 2016. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4104–4113.
- Sun, J.; Shen, Z.; Wang, Y.; Bao, H.; and Zhou, X. 2021. LoFTR: Detector-Free Local Feature Matching with Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8922–8931.
- Thomee, B.; Shamma, D. A.; Friedland, G.; Elizalde, B.; Ni, K.; Poland, D.; Borth, D.; and Li, L.-J. 2016. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2): 64–73.
- Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph Attention Networks. *arXiv preprint*.
- Xia, P.; Liao, T.; Zhu, W.; Zhao, D.; Ke, J.; Zhang, K.; Lu, T.; and Wang, T. 2025. CorrMoE: Mixture of Experts with De-Stylization Learning for Cross-Scene and Cross-Domain Correspondence Pruning. *arXiv preprint*.
- Xiao, J.; Owens, A.; and Torralba, A. 2013. Sun3d: A database of big spaces reconstructed using sfm and object labels. In *Proceedings of the IEEE International Conference on Computer Vision*, 1625–1632.
- Yi, K. M.; Trulls, E.; Lepetit, V.; and Fua, P. 2016. Lift: Learned invariant feature transform. In *Proceedings of the European Conference on Computer Vision*, 467–483.
- Yi, K. M.; Trulls, E.; Ono, Y.; Lepetit, V.; Salzmann, M.; and Fua, P. 2018. Learning to find good correspondences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2666–2674.
- Zeng, A.; Song, S.; Nießner, M.; Fisher, M.; Xiao, J.; and Funkhouser, T. 2017. 3DMatch: Learning Local Geometric Descriptors from RGB-D Reconstructions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1802–1811.
- 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 *Proceedings of the IEEE International Conference on Computer Vision*, 5845–5854.
- Zhang, S.; Li, Z.; Gao, Y.; and Ma, J. 2024. DeMatch: Deep Decomposition of Motion Field for Two-View Correspondence Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 20278–20287.
- Zhang, S.; and Ma, J. 2023. ConvMatch: Rethinking Network Design for Two-View Correspondence Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5): 2920–2935.
- Zhao, C.; Ge, Y.; Zhu, F.; Zhao, R.; Li, H.; and Salzmann, M. 2021. Progressive correspondence pruning by consensus learning. In *Proceedings of the IEEE International Conference on Computer Vision*, 6464–6473.