

PSReg: Prior-guided Sparse Mixture of Experts for Point Cloud Registration

Xiaoshui Huang^{1,5*}, Zhou Huang^{2,6*}, Yifan Zuo^{2,6†}, Yongshun Gong³,
Chengdong Zhang¹, Deyang Liu⁴, Yuming Fang^{2,6}

¹Shanghai Jiao Tong University,

²Jiangxi University of Finance and Economics,

³Shandong University,

⁴Anqing Normal University,

⁵Huaihua University,

⁶Jiangxi Provincial Key Laboratory of Multimedia Intelligent Processing

Abstract

The discriminative feature is crucial for point cloud registration. Recent methods improve the feature discriminative by distinguishing between non-overlapping and overlapping region points. However, they still face challenges in distinguishing the ambiguous structures in the overlapping regions. Therefore, the ambiguous features they extracted resulted in a significant number of outlier matches from overlapping regions. To solve this problem, we propose a prior-guided SMoE-based registration method to improve the feature distinctiveness by dispatching the potential correspondences to the same experts. Specifically, we propose a prior-guided SMoE module by fusing prior overlap and potential correspondence embeddings for routing, assigning tokens to the most suitable experts for processing. In addition, we propose a registration framework by a specific combination of Transformer layer and prior-guided SMoE module. The proposed method not only pays attention to the importance of locating the overlapping areas of point clouds, but also commits to finding more accurate correspondences in overlapping areas. Our extensive experiments demonstrate the effectiveness of our method, achieving state-of-the-art registration recall (95.7%/79.3%) on the 3DMatch/3DLoMatch benchmark. Moreover, we also test the performance on ModelNet40 and demonstrate excellent performance.

Introduction

Point cloud registration is fundamental in computer vision and widely applies in domains such as 3D reconstruction (Groß, Ošep, and Leibe 2019; Izadi et al. 2011), autonomous driving (Nagy and Benedek 2018), and augmented reality (Huang et al. 2016). The objective is to estimate the optimal rigid transformation between two point clouds, thereby establishing their correspondences within the same coordinate system. Currently, learning-based methods are leading the task of point cloud registration. However, successfully registering them remains challenging when the point clouds only partially overlap.

*These authors contributed equally.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

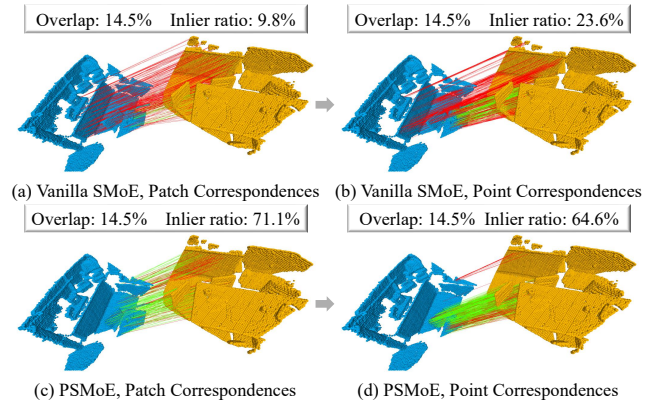


Figure 1: Performance gaps in applying vanilla SMoE and our PSMoE. Fig. (a) and (c) show coarse patch correspondences, (b) and (d) show fine point correspondences.

Learning-based works on partial overlap point cloud registration can be classified into two categories: non-Transformer-based methods and transformer-based methods. Non-Transformer-based methods (Zeng et al. 2017; Huang, Mei, and Zhang 2020; Choy, Park, and Koltun 2019; Ao et al. 2021) directly utilize existing networks for single-stage feature extraction and estimate the transformation matrix. However, they face challenges in integrating global information and feature alignment. In comparison, the Transformer-based methods (Wang and Solomon 2019; Huang et al. 2021, 2022a; Yu et al. 2021; Mei et al. 2022; Yew and Lee 2022; Qin et al. 2022) utilize a Transformer module to further enhance the features extracted by the existing backbones. Nevertheless, in scenarios with a large number of similar structures or textureless structures, these methods face the issue of ambiguous matching. Among them, PEAL (Yu et al. 2023) explicitly encodes intra-frame overlapping and non-overlapping regions using prior information, which alleviates the above ambiguous matching to a certain extent. However, if the overlapping region between the two point clouds contains a large amount of ambiguous structures (floor or wall without obvious texture), PEAL still faces difficulty. This is because it cannot distinguish the am-

biguous structures within the overlapping regions, resulting in a large number of outlier matches.

To address the aforementioned issues, we believe that applying sparse mixture of experts (SMoE) in Transformer is a feasible solution. The routing network of SMoE can cluster input tokens into different groups and then distribute these groups to different experts for processing. If the routing network can distinguish between overlapping and non-overlapping regions of tokens, and further capture the correspondence between tokens of two point clouds, then the multi-expert networks through SMoE can extract more discriminative features. To the best of our knowledge, we are pioneering the exploration of SMoE for addressing the point cloud registration task.

However, SMoE’s simple application does not fulfill our expectations. Due to insufficient interaction between the features of the source and target point clouds, and more importantly, the lack of relevant guidance, it’s challenging for SMoE’s routing network to differentiate tokens from overlapping and non-overlapping regions (Fig. 1 (a)). Inspired by PEAL (Yu et al. 2023), we propose **P**rior-guided routing **S**parse **M**ixture of **E**xperts (PSMoE), a novel module to address the aforementioned challenge. Specifically, we first calculate superpoint correspondences using transformations estimated by a SOTA pretrained model. Then, we design a **P**rior superpoint **C**orrespondences **E**ncoding (PCE) module to obtain prior embeddings. PCE module encodes matched superpoints with discrete ordered numbers, while non-matched superpoints are assigned the same number, all the numbers assigned to superpoints are encoded using an embedding layer. Finally, the routing network groups the original tokens based on prior embeddings and sends them to selected experts to extract distinctive features. Unlike PEAL, our PCE module not only encodes prior overlapping information but also further captures the potential correspondence of matching points. These prior correspondences guide the distribution of potential matching tokens to be processed by the same expert as much as possible, which facilitates feature alignment. Therefore, our method effectively locates overlapping areas and enhances matching accuracy within these regions (Fig. 1 (c)).

We design a coarse-to-fine point cloud registration framework based on the PSMoE module, named **P**rior-guided **S**MoE-based **R**egistration (PSReg), to estimate the final transformation matrix. Additionally, we conduct extensive experiments on indoor benchmarks and synthetic datasets, demonstrating the superiority of our approach over previous methods. A comprehensive series of ablation studies validate the effectiveness of each module within our proposed methodology. Our contributions can be summarized as follows:

- We analyze the reasons for the poor performance of directly applying vanilla SMoE in point cloud registration, and propose a prior-guided SMoE (PSMoE) module to further improve the feature distinctiveness within the overlapping regions. To the best of our knowledge, we are the pioneers in exploring the SMoE in the point cloud registration task.

- We propose a registration method, PSReg, by combining both the Transformer and our PSMoE module.
- We conducted extensive experiments to verify the effectiveness of our methodology and explored the feasibility of applying SMoE to the point cloud registration task.

Related Work

Non-Transformer-based Registration

The overall idea of non-Transformer-based point cloud registration methods is typically using networks specifically designed for 3D data to extract features in a single stage, without further enhancing high-level features. These features are then utilized to establish correspondences and indirectly solve for the transformation matrix, or the transformation matrix is solved directly from the features. Based on the different types of input data, we can further broadly classify the existing methods into two categories. Firstly, methods based on structured data (Zeng et al. 2017; Elbaz, Avraham, and Fischer 2017; Choy, Park, and Koltun 2019; Ao et al. 2021), which typically involve projecting the point cloud into 2D images or constructing 3D volume, and directly applying 2D CNN or 3D CNN. Secondly, methods based on raw point clouds (Aoki et al. 2019; Du et al. 2019; Huang, Mei, and Zhang 2020; Bai et al. 2020; Xu et al. 2021; Huang et al. 2022b), which directly use the point cloud data as input. PointNet (Qi et al. 2017a) directly uses MLPs to extract point-wise features and addresses permutation invariance through max pooling. PointNet++ (Qi et al. 2017b) further introduces sampling and grouping layers to capture multi-scale contextual information. KPConv (Thomas et al. 2019) uses kernel points that carry convolution weights to simulate the kernel pixels in 2D convolution, thereby defining the convolution operation on raw point clouds. The above three backbone networks, which directly process raw point cloud data, are widely used in point cloud registration. However, the above methods face issues with information loss and challenges in capturing global contextual information.

Transformer-based Registration

Transformer-based registration methods typically involve a two-stage feature extraction process. Initially, high-level features are extracted by the existing 3D backbone, followed by further feature enhancement through Transformer. DCP (Wang and Solomon 2019) is the pioneer in utilizing Transformer to enhance feature extraction in point cloud registration, whereby unaligned point clouds are fed into a feature embedding module, followed by context aggregation executed by a Transformer encoder. REGTR (Yew and Lee 2022) replaces explicit feature matching and outlier filtering, typically performed by RANSAC, with an end-to-end Transformer framework to directly discover point cloud correspondences. Taking into consideration that the vanilla Transformer does not explicitly encode the geometric structure of point clouds, GeoTransformer (Qin et al. 2022) is engineered with a geometric self-attention module that expressly captures the internal geometric structures of point clouds. PEAL (Yu et al. 2023) introduces an overlapping

prior, utilizing one-way attention to relieving feature ambiguity. However, these methods are prone to failure when dealing with scenes containing a large number of ambiguous structures.

Sparse Mixture of Experts (SMoE)

SMoE is an efficient deep learning model architecture that utilizes a multi-expert system to handle complex tasks. In this structure, the network is composed of multiple experts (multiple FFNs) and a gating mechanism. The gating mechanism is responsible for deciding which experts to activate based on the input data, allowing the network to dynamically adjust its behavior for specific tasks, thereby enhancing processing speed and efficiency. The efficacy of SMoE has been extensively demonstrated across various tasks in NLP (Shazeer et al. 2017; Lepikhin et al. 2020; Zhou et al. 2022; Fedus, Zoph, and Shazeer 2022; Jiang et al. 2024) and CV (Gross, Ranzato, and Szlam 2017; Abbas and Andreopoulos 2020; Pavlitskaya et al. 2020; Riquelme et al. 2021).

However, research on SMoE in 3D vision remains unexplored. In fact, the SMoE is naturally suited for the task of partially overlapping point cloud registration. Partial overlapping point cloud registration focuses more on the overlapping regions of the point clouds. The gating mechanism of SMoE can sparsely activate a subset of experts to specifically handle data from these overlapping regions, which can enhance the uniqueness of the features. The aforementioned sparse activation method is theoretically feasible, but it is difficult to implement in practical applications. Therefore, to achieve controllable sparse activation, we propose a novel prior-guided routing mechanism that leverages SMoE for more effective discriminative feature extraction.

Method

Preliminary and Analysis

In this section, we will introduce the preliminaries of point cloud registration and sparse mixture of experts (SMoE).

Point Cloud Registration. Given a pair of partially overlapping point clouds, namely the source point cloud $\mathcal{P} = \{p_i \in \mathbb{R}^3 | i = 1, \dots, N\}$ and the target point cloud $\mathcal{Q} = \{q_j \in \mathbb{R}^3 | j = 1, \dots, M\}$, where N and M are the number of points in point clouds \mathcal{P} and \mathcal{Q} . The goal of point cloud registration is to predict the rotation matrix $R \in SO(3)$ and translation vector $t \in \mathbb{R}^3$ to align \mathcal{P} and \mathcal{Q} under the same coordinate system. Our approach needs to be combined with Transformer, while Transformer-based registration methods typically employ an encoder (existing 3D feature extraction networks) for downsampling the raw point clouds and extracting relevant neighborhood features before applying the Transformer. We use $\hat{\mathcal{P}} = \{\hat{p}_i \in \mathbb{R}^3 | i = 1, \dots, N'\}$ and $\hat{\mathcal{Q}} = \{\hat{q}_j \in \mathbb{R}^3 | j = 1, \dots, M'\}$ to represent downsampled points (superpoints), and the corresponding superpoint features are represented by $\hat{\mathcal{F}}^{\mathcal{P}} \in \mathbb{R}^{N' \times d}$ and $\hat{\mathcal{F}}^{\mathcal{Q}} \in \mathbb{R}^{M' \times d}$.

Vanilla SMoE. The vanilla SMoE (Riquelme et al. 2021) is a neural network layer comprising self-attention mechanisms and multiple feed-forward expert networks. Given an




Figure 2: The visualization of the routing of tokens (patches) in the last SMoE block, where the color of the point cloud patches indicates the selected experts.

input point cloud token x , it is sent to the routing network $\mathcal{G}(\cdot)$, which delegates the token to several highly relevant experts for processing. Typically, the routing network is a fully connected layer, represented as follows:

$$\mathcal{G} = \text{top-}k(\text{softmax}(W_g x)), \quad (1)$$

Where W_g represents the trainable parameters in the router neural network and $\text{top-}k(\cdot)$ select the k experts with the largest gating values. In our approach, $k = 1$. The final output of the SMoE layer will be the summary of the weighted active expert features:

$$y = \mathcal{G}_i \cdot E_i(x), \quad (2)$$

where E_i represents the neural network of expert i , and \mathcal{G}_i represents the corresponding gate value.

Analyze Vanilla SMoE. The multi-expert network of SMoE can focus on different features or patterns, which should help enhance the uniqueness of superpoint features in different regions. To analyze the ability of SMoE in the point cloud registration, we directly use vanilla SMoE on the selected baseline model (Yu et al. 2023). Specifically, we replace the FFN layer in the self-attention block of the baseline with the SMoE layer. This SMoE with the tokens choice routing policy, following the default configurations in (Fedus, Zoph, and Shazeer 2022). To intuitively observe the routing results of vanilla SMoE, we visualized the distribution of tokens in the last SMoE block. As illustrated in Fig. 2 (b), the tokens do not exhibit a clear clustering phenomenon, indicating that the vanilla SMoE does not handle the expert subsets for tokens in overlapping and non-overlapping regions separately.

Analysis: The primary reason for the failure of vanilla SMoE in locating overlapping regions is that the routing network is typically just a fully connected layer. Without explicit guiding signals, its expressive capability struggles to capture the overlapping information between point clouds.

PSMoE Module

To solve the above-mentioned problem and further alleviate matching ambiguities in overlapping areas, we propose **Prior-guided Sparse MoE (PSMoE)**, as shown in Fig. 3, which offers explicit prior guidance. Fig. 2 (c) shows the token distribution in our method. Due to the introduction of prior information, the routing network can distinguish between overlapping and non-overlapping patches. Therefore, providing clear guidance signals can assist the routing network in better identifying which experts are more

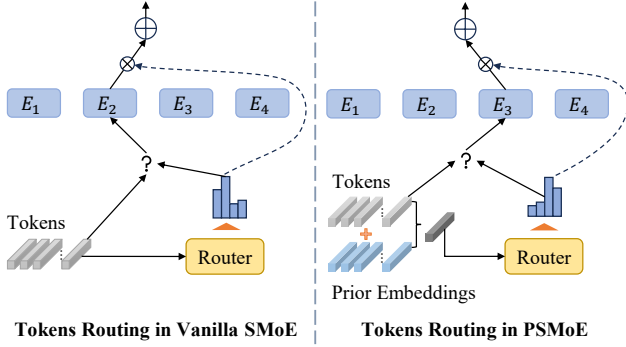


Figure 3: Difference between vanilla SMoE and PSMoE.

suitable for specific inputs. Next, we will sequentially introduce our approach from three aspects: prior superpoint correspondence prediction, prior superpoint correspondence encoding, and prior-guided routing.

Prior Superpoint Correspondence Prediction. Superpoint correspondences not only provide overlapping information but also directly associate matching superpoints in two point clouds. These explicit signals can effectively guide the routing network as priors to activate proper expert extract discriminative point-wise features. However, it is a chicken-and-egg problem between superpoint correspondences and discriminative point-wise feature extraction. To solve this problem, inspired by the approach of (Yu et al. 2023), we calculate the *prior superpoint correspondences* by using a pretrained SOTA registration model. Specifically, we employ GeoTransformer to predict the prior transformation matrix $T_p = \{R_p, t_p\}$, and use T_p to compute the overlap ratio matrix $\hat{O} \in \mathbb{R}^{N' \times M'}$ between the superpoints (patches) of \hat{P} and \hat{Q} . Then, we select entries in \hat{O} with an overlap ratio greater than τ_o to serve as prior superpoint correspondences $\bar{C} = (\mathcal{P}', \mathcal{Q}')$, and the corresponding prior overlap ratios are represented as \bar{O} .

Prior Superpoint Correspondence Encoding (PCE). To aid the routing network in discerning the precise locations of tokens within overlapping regions and to effectively capture the token correspondences in these putative overlap areas, our aim is to partition all tokens into L (the number of experts) clusters. These clusters can be broadly classified into two categories: the non-anchor cluster (a singular cluster where the majority of tokens reside in non-overlapping zones) and anchor clusters (comprising $L - 1$ clusters primarily housing tokens within overlapping regions). Ideally, each sub-cluster within the anchor clusters should encompass tokens with potential matches between the two sets of data points.

Consequently, we have devised a PCE module to explicitly encode prior superpoint correspondences, as illustrated in Fig. 4. Specifically, we first encode the prior superpoint correspondences \bar{C} using a discrete and ordered sequence \mathcal{Z}_A , with $\bar{c} = |\mathcal{Z}_A| = |\bar{C}|$. We then employ the same numerical value \mathcal{Z}_{NA} to encode unmatched superpoints in \mathcal{P} and \mathcal{Q} . Finally, we obtain a set of natural numbers, denoted




Figure 4: Diagram of PCE module.

as $\mathcal{N} = \{i \in \mathbb{N} \mid i = 0, \dots, \bar{c}\}$:

$$\mathcal{N} = \text{Concat}[\mathcal{Z}_{NA}, \mathcal{Z}_A], \quad (3)$$

Considering the continuous nature of the sinusoidal embedding function (Vaswani et al. 2017), we utilize this function and MLP to form an embedding layer to compute the prior embeddings $\mathcal{F}^{\mathcal{N}} \in \mathbb{R}^{(\bar{c}+1) \times d}$ for \mathcal{N} :

$$\begin{cases} \Gamma_{i,2k}^{\mathcal{N}} = \sin\left(\frac{i}{10000^{2k/d}}\right) \\ \Gamma_{i,2k+1}^{\mathcal{N}} = \cos\left(\frac{i}{10000^{2k/d}}\right) \end{cases}, \quad \mathcal{F}^{\mathcal{N}} = \text{MLP}(\Gamma^{\mathcal{N}}), \quad (4)$$

Based on \bar{C} , we select the corresponding superpoints from \mathcal{P} to form a set of anchor superpoints $\mathcal{A} \subset \hat{\mathcal{P}}$. However, superpoints in \mathcal{A} may correspond to multiple superpoints in \mathcal{Q} that meet the criteria (overlap ratio greater than τ_o), which are denoted as \mathcal{A}' . Therefore, to ensure the uniqueness of the prior embedding corresponding to \mathcal{A}' , we perform a weighted sum based on the overlap ratio. For each superpoint in \mathcal{A} , we search for the entries corresponding to this superpoint from \mathcal{P}' in \bar{C} and the corresponding index is represented as \mathcal{X}_i . Ultimately, the prior embeddings $\check{\mathcal{F}}^{\mathcal{P}} \in \mathbb{R}^{|\hat{\mathcal{P}}| \times d}$ corresponding to the superpoints in $\hat{\mathcal{P}}$ is denoted as:

$$\check{\mathcal{F}}_i^{\mathcal{P}} = \begin{cases} \text{softmax}(\bar{O}(\mathcal{X}_i)) \cdot \mathcal{F}_{\mathcal{Z}_A}^{\mathcal{N}}(\mathcal{X}_i), & |\mathcal{X}_i| > 1, \\ \mathcal{F}_{\mathcal{Z}_A}^{\mathcal{N}}(\mathcal{X}_i), & |\mathcal{X}_i| = 1, \\ \mathcal{F}_{\mathcal{Z}_{NA}}^{\mathcal{N}}, & \text{otherwise.} \end{cases} \quad (5)$$

Where $\mathcal{F}_{\mathcal{Z}_A}^{\mathcal{N}}$ and $\mathcal{F}_{\mathcal{Z}_{NA}}^{\mathcal{N}}$ respectively represent the embeddings corresponding to \mathcal{Z}_A and \mathcal{Z}_{NA} . The prior embeddings $\check{\mathcal{F}}^{\mathcal{Q}} \in \mathbb{R}^{|\hat{\mathcal{Q}}| \times d}$ corresponding to $\hat{\mathcal{Q}}$ are obtained in the same manner.

Prior-guided Routing. In Fig. 3, it is illustrated that to ensure prior embeddings effectively guide the routing process and simplify model design, we adopt a straightforward yet efficient approach. This involves directly integrating them with prior information before distributing the input tokens through the routing network. More precisely, We sum the prior embeddings encoded by the PCE module with the original tokens, and giving them equal weight. Subsequently, we

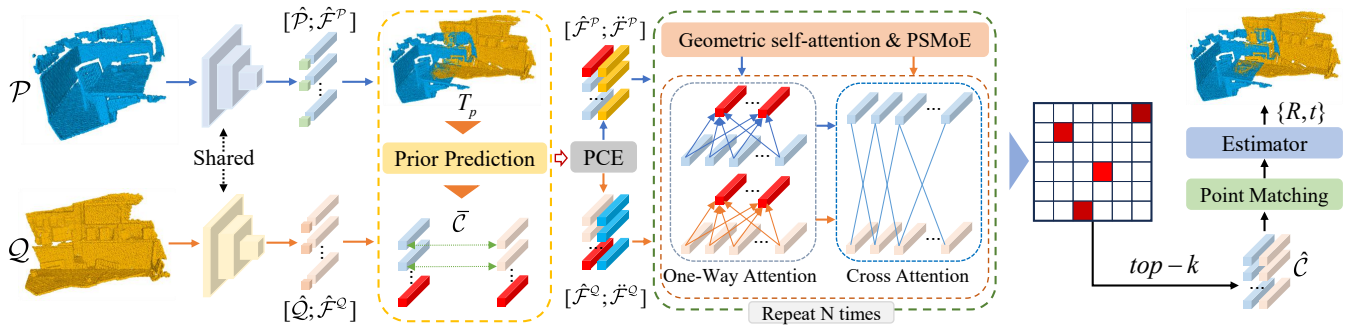


Figure 5: The framework of PSReg. The red blocks represent superpoints located in the putative non-overlapping areas.

conduct routing using these composite tokens enriched with prior information, as detailed below:

$$\mathcal{G}^P = \text{top-}k(\text{softmax}(W_g(X^P + \tilde{F}^P))). \quad (6)$$

Where X^P denote the tokens in \mathcal{P} , and \tilde{F}^P represent their corresponding prior embeddings. Under the guidance of prior information, distribute the original tokens to the appropriate experts for processing. We obtain the final output of the PSMoE block for tokens in \mathcal{P} according to Equation 2, i.e., $G_i^P \cdot E_i(x^p)$. Tokens in \mathcal{Q} are processed using the same method.

PSReg

Algorithm Design. To solve the point cloud registration task, we propose a **P**rior-guided **S**MoE-based **R**egistration algorithm (PSReg). Our approach builds upon (Yu et al. 2023), and the pipeline is illustrated in Fig. 5. The framework adopts a coarse-to-fine paradigm. Initially, we utilize the KPConv-FPN backbone (Lin et al. 2017; Thomas et al. 2019) network to downsample the original point clouds, obtaining superpoints and their corresponding features. Discriminative superpoint features are crucial for high-quality point cloud registration, as the accuracy of correspondences in the fine stage depends on superpoint correspondences. Therefore, our method primarily focuses on feature extraction in the coarse stage.

We integrate our PSMoE with Transformer, leveraging PSMoE’s multi-expert network to extract highly discriminative superpoint features for precise superpoint correspondence identification. Subsequently, the point matching module deduces exact point correspondences, followed by the application of local-to-global registration (LGR) (Qin et al. 2022) for the final transformation estimation.

Optimization. The optimization of PSReg is to minimize two correspondence supervision losses ($\mathcal{L}_c, \mathcal{L}_f$) and a load balancing loss (\mathcal{L}_g). The final loss function is :

$$\mathcal{L} = \mathcal{L}_c + \mathcal{L}_f + \mathcal{L}_g. \quad (7)$$

The coarse correspondence loss \mathcal{L}_c uses the overlap-aware circle loss (Qin et al. 2022), which focuses more on the positive samples with high overlap. The fine correspondence loss \mathcal{L}_f uses the negative log-likelihood loss. To encourage a balanced load across experts, we add a load balancing loss \mathcal{L}_g , which is the same as in (Fedus, Zoph, and Shazeer 2022).

Samples	3DMatch					3DLoMatch				
	5000	2500	1000	500	250	5000	2500	1000	500	250
<i>Feature Matching Recall (%)</i> ↑										
FCGF	97.4	97.3	97.0	96.7	96.6	76.6	75.4	74.2	71.7	67.3
D3Feat	95.6	95.4	94.5	94.1	93.1	67.3	66.7	67.0	66.7	66.5
SpinNet	97.6	97.2	96.8	95.5	94.3	75.3	74.9	72.5	70.0	63.6
YOHO	98.2	97.6	97.5	97.7	96.0	79.4	78.1	76.3	73.8	69.1
Predator	96.6	96.6	96.5	96.3	96.5	78.6	77.4	76.3	75.7	75.3
CoFiNet	98.1	98.3	98.1	98.2	98.3	83.1	83.5	83.3	83.1	82.6
GeoTr.	97.9	97.9	97.9	97.9	97.6	88.3	88.6	88.8	88.6	88.3
OIF-PCR	98.1	98.1	97.9	98.4	98.4	84.6	85.2	85.5	86.6	87.0
PEAL	98.4	98.4	98.4	98.4	98.4	87.7	87.8	87.8	88.0	87.4
SIRA-PCR	98.2	98.4	98.4	98.5	98.5	88.8	89.0	88.9	88.6	87.7
PSReg	98.6	98.6	98.6	98.6	98.6	86.4	86.6	87.1	87.4	87.1
<i>Inlier Ratio (%)</i> ↑										
FCGF	56.8	54.1	48.7	42.5	34.1	21.4	20.0	17.2	14.8	11.6
D3Feat	39.0	38.8	40.4	41.5	41.8	13.2	13.1	14.0	14.6	15.0
SpinNet	47.5	44.7	39.4	33.9	27.6	20.5	19.0	16.3	13.8	11.1
YOHO	64.4	60.7	55.7	46.4	41.2	25.9	23.3	22.6	18.2	15.0
Predator	58.0	58.4	57.1	54.1	49.3	26.7	28.1	28.3	27.5	25.8
CoFiNet	49.8	51.2	51.9	52.2	52.2	24.4	25.9	26.7	26.8	26.9
GeoTr.	71.9	75.2	76.0	82.2	85.1	43.5	45.3	46.2	52.9	57.7
OIF-PCR	62.3	65.2	66.8	67.1	67.5	27.5	30.0	31.2	32.6	33.1
PEAL	74.6	81.1	85.8	87.7	88.9	48.1	53.5	59.9	62.4	64.4
SIRA-PCR	70.8	78.3	83.7	85.9	87.4	43.3	49.0	55.9	58.8	60.7
PSReg	75.8	82.4	87.1	88.9	90.0	49.9	55.5	61.9	64.5	66.3
<i>Registration Recall (%)</i> ↑										
FCGF	85.1	84.7	83.3	81.6	71.4	40.1	41.7	38.2	35.4	26.8
D3Feat	81.6	84.5	83.4	82.4	77.9	37.2	42.7	46.9	43.8	39.1
SpinNet	88.6	86.6	85.5	83.5	70.2	59.8	54.9	48.3	39.8	26.8
YOHO	90.8	90.3	89.1	88.6	84.5	65.2	65.5	63.2	56.5	48.0
Predator	89.0	89.9	90.6	88.5	86.6	59.8	61.2	62.4	60.8	58.1
CoFiNet	89.3	88.9	88.4	87.4	87.0	67.5	66.2	64.2	63.1	61.0
GeoTr.	92.0	91.8	91.8	91.4	91.2	75.0	74.8	74.2	74.1	73.5
OIF-PCR	92.4	91.9	91.8	92.1	91.2	76.1	75.4	75.1	74.4	73.6
PEAL	94.4	94.3	94.0	93.8	93.8	78.9	78.6	77.7	76.9	77.0
SIRA-PCR	93.6	93.9	93.9	92.7	92.4	73.5	73.9	73.0	73.4	71.1
PSReg	95.7	94.9	95.1	95.0	95.2	79.3	79.3	78.7	78.7	78.4

Table 1: Quantitative results on 3DMatch and 3DLoMatch with different numbers of samples. GeoTr. denotes Geo-Transformer.

Experiments

We validate the efficacy of our approach on real-world datasets such as 3DMatch/3DLoMatch and synthetic datasets like ModelNet/ModelLoNet. Furthermore, we have performed comprehensive ablation experiments.

3DMatch & 3DLoMatch Benchmarks

Dataset. The 3DMatch (Zeng et al. 2017) comprises data collected from 62 indoor scenes, with 46 scenes designated for training, 8 for validation, and 8 for testing. We utilize data processed by (Huang et al. 2021). The 3DMatch bench-

mark includes pairs of point clouds with an overlap rate of more than 30%, while the 3DLoMatch benchmark only includes pairs of scans with an overlap rate of 10%-30%.

Metrics. We follow (Huang et al. 2021), using the following three metrics to evaluate the performance of registration: Registration Recall (RR), Feature Matching Recall (FMR) and Inlier Ratio (IR).

Registration Results. In Tab. 1, we compare our method against deep learning-based baselines. Specifically, we evaluated four methods based on local descriptors: FCGF (Choy, Park, and Koltun 2019), D3Feat (Bai et al. 2020), SpinNet (Ao et al. 2021), and YOHO (Wang et al. 2022), along with six Transformer-based methods: Predator (Huang et al. 2021), CoFiNet (Yu et al. 2021), GeoTransformer (Qin et al. 2022), OIF-PCR (Yang et al. 2022), PEAL (Yu et al. 2023) and SIRA-PCR (Chen et al. 2023). Following (Huang et al. 2021), we report the performance of each method at five different numbers of correspondences (5000, 2500, 1000, 500, 250). As PEAL only reports results based on a 2D prior, while other baselines do not need 2D information, we train PEAL using a 3D prior as the input from scratch for a fair comparison. In addition, we adopted the same iterative update strategy, with 6 iterations.

RR directly reflects the final performance of the registration. As shown in Tab. 1, our method achieved the highest RR across all sampled correspondences on 3DMatch and 3DLoMatch. Compared to the baseline GeoTransformer, our approach has improved by no less than 3.3% on 3DMatch and no less than 5.7% on 3DLoMatch. Under the condition of introducing the same prior information, our method has achieved a performance improvement of up to 1.5% on 3DMatch and up to 1.8% on 3DLoMatch compared to PEAL. Notably, our approach maintains a high RR even when the number of sampling correspondences is reduced. Especially when the number of sampling correspondences on 3DMatch is 250, we have achieved a RR of 95.2%, which is higher than the RR of other baseline methods when the number of sampling correspondences is 5000.

As for IR, our method significantly outperforms other baseline methods, demonstrating that under the strategy of prior-guided routing, the multi-expert network of PSMoE significantly enhances the distinguishability of features, thereby obtaining reliable correspondences. For FMR, our method does not perform as well as some baseline methods on 3DLoMatch. We speculate that this is because our method is more sensitive to the quality of the priors. It can significantly improve the IR in most point cloud pairs, but in some point cloud pairs where the priors cannot provide enough correct superpoint correspondences, our method would reduce their IR, resulting in a slightly poorer FMR.

ModelNet & ModelLoNet Benchmarks

Dataset. ModelNet40 (Wu et al. 2015) consists of CAD models of 12,311 objects from 40 different categories. We follow the data settings in (Yew and Lee 2020; Huang et al. 2021), where the point clouds are sampled randomly from mesh faces of the CAD models, cropped and subsampled. A total of 5,112 samples are used for training, 1,202 samples for validation, and 1,266 samples for testing. Following to

Methods	ModelNet			ModelLoNet		
	RRE	RTE	CD	RRE	RTE	CD
DCP-v2	11.975	0.171	0.0117	16.501	0.300	0.0268
RPM-Net	1.712	0.018	0.00085	7.342	0.124	0.0050
Predator	1.739	0.019	0.00089	5.235	0.132	0.0083
LGR						
GeoTr.	2.284	0.023	0.00106	4.176	0.098	0.0047
PEAL	2.070	0.021	0.00100	4.092	0.108	0.0057
PSReg	2.146	0.022	0.00103	3.952	0.096	0.0046
RANSAC						
GeoTr.	1.875	0.018	0.00089	3.801	0.093	0.0047
PEAL	1.760	0.018	0.00088	3.524	0.103	0.0056
PSReg	1.745	0.017	0.00086	3.498	0.091	0.0043

Table 2: Evaluation results on ModelNet40 dataset.

(Huang et al. 2021), we evaluated two partially overlapping settings: the average overlap rate for ModelNet is 73.5%, and for ModelLoNet, it is 53.6%. We trained only on ModelNet and directly generalized to ModelLoNet.

Metrics. Following (Yew and Lee 2020; Huang et al. 2021), we report the Relative Rotation Error (RRE) that evaluates the error between estimated and ground truth rotation matrices, and Relative Translation Error (RTE) that measures the error between estimated and ground truth translation vectors, as well as the Chamfer distance (CD) between the registered scans.

Registration Results. We compare the proposed PSReg with five SOTA baselines, including two end-to-end methods (Wang and Solomon 2019; Yew and Lee 2020), one correspondence-based method (Huang et al. 2021), and two coarse-to-fine registration approaches (Qin et al. 2022; Yu et al. 2023). For the coarse-to-fine methods, we report the metrics based on both the LGR estimator and the RANSAC estimator, with the number of RANSAC iterations set to 50K. The detailed comparison results are shown in Tab. 2. As mentioned in (Huang et al. 2021), many end-to-end methods are specifically tuned for ModelNet40, and RPM-Net (Yew and Lee 2020) also utilizes surface normal information. Although our method’s RRE and CD on ModelNet are slightly lower than RPM-Net (Yew and Lee 2020), it surpasses other baseline methods on the more challenging ModelLoNet in terms of the RRE metric, even without using RANSAC.

Qualitative Results

Qualitative comparison with other SOTA methods is shown in Fig. 6. It is worth noting that in the 3DLoMatch example, the point cloud pair only has partially overlapping walls. In such scenarios with a large number of similar structures, our method successfully achieves registration, whereas PEAL fails. From the quantitative and qualitative results, it can be seen that our proposed PSReg significantly enhances the distinguishability of features, which helps infer reliable superpoint correspondences in low-overlap point cloud registration, thereby achieving precise alignment.




Figure 6: Visual comparison examples on 3DMatch and 3DLoMatch datasets.

Ablation Studies

In this section, following (Qin et al. 2022), we conduct extensive ablation experiments based on the LGR estimator. We report the FMR and IR of all dense point correspondences, as well as RR.

Threshold (τ_o)	3DMatch			3DLoMatch		
	FMR	IR	RR	FMR	IR	RR
0.0	98.6	74.7	95.0	86.9	49.7	78.5
0.1	98.3	74.4	93.8	86.2	49.6	78.0
0.3	98.3	74.5	93.1	86.3	49.0	77.5

Table 3: Ablation experiments on overlap threshold.

To evaluate the impact of the selection of prior superpoint correspondences, we tested various overlap thresholds, as shown in Tab. 3. The experimental results indicate that more prior superpoint correspondences are more conducive to registration. Hence, we ultimately set $\tau_o = 0$.

Model	3DMatch			3DLoMatch		
	FMR	IR	RR	FMR	IR	RR
(a) w/o SMoE	98.4	73.2	94.2	88.0	47.1	77.8
(b) w/ Vanilla SMoE	98.5	74.0	94.2	86.8	48.8	78.1
(c) w/ PSMoE (BC)	98.4	73.9	94.6	87.8	49.5	79.2
(d) w/ PSMoE (OC)	98.6	74.7	95.0	86.9	49.7	78.5

Table 4: Ablation experiments on PSReg. 'BC' means binary coding, and 'OC' means ordered coding.

In Tab. 4, we compare different configurations of PSReg, including: (a) the baseline (Yu et al. 2023) we chose does not integrate SMoE, (b) using the vanilla SMoE (Fedus, Zoph, and Shazeer 2022) on the baseline. It can be observed from Tab. 4 that the integration of SMoE offers negligible improvement in registration performance. In addition, we also compared different encoding methods in the PCE module. According to the prior superpoint correspondences, we classify superpoints into matched and unmatched categories, encoding them with 0 and 1 respectively, which is binary coding. The ordered coding is the method shown in Fig. 4. Tab. 4 presents the results of different encoding methods: (c) the

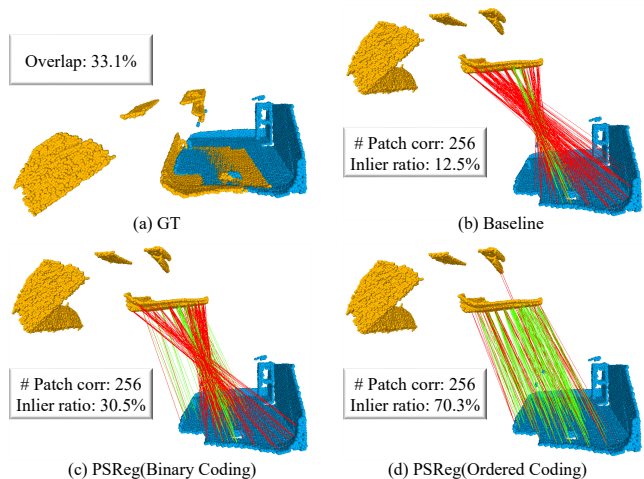


Figure 7: Visualization of superpoint correspondences with different coding methods in the PCE module.

PCE module uses binary coding, (d) the PCE module uses ordered coding. Both encoding methods effectively improve registration performance, particularly with the ordered coding method significantly enhancing the IR.

Fig. 7 visualizes the correspondence of superpoints, it can be clearly observed that the ordered coding method achieves more accurate correspondences in overlapping areas compared to the binary coding method and the baseline. This is because the ordered coding method can distribute potentially matching tokens to the same expert as much as possible, which helps to infer more accurate correspondences.

Conclusion

In this paper, we propose the prior-guided SMoE (PSMoE) to differentiate between the ambiguous structures in the overlapping regions. Based on the PSMoE, we propose the first SMoE-based registration framework. Through extensive experiments, our method demonstrates a significant enhancement in feature discriminability, leading to a notable improvement in inlier ratio. Our method explores a new direction by utilizing a multi-expert neural network to enhance feature discriminability for point cloud registration.

Acknowledgements

This work is supported in part by National Natural Science Foundation of China (62271237, 62132006), Jiangxi Provincial Natural Science Foundation of China (20242BAB26014) and Science and Technology Department of Jiangxi Province (20223AEI91002).

References

- Abbas, A.; and Andreopoulos, Y. 2020. Biased mixtures of experts: Enabling computer vision inference under data transfer limitations. *IEEE Transactions on Image Processing*, 29: 7656–7667.
- Ao, S.; Hu, Q.; Yang, B.; Markham, A.; and Guo, Y. 2021. Spinnet: Learning a general surface descriptor for 3d point cloud registration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11753–11762.
- Aoki, Y.; Goforth, H.; Srivatsan, R. A.; and Lucey, S. 2019. Pointnetlk: Robust & efficient point cloud registration using pointnet. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7163–7172.
- Bai, X.; Luo, Z.; Zhou, L.; Fu, H.; Quan, L.; and Tai, C.-L. 2020. D3feat: Joint learning of dense detection and description of 3d local features. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6359–6367.
- Chen, S.; Xu, H.; Li, R.; Liu, G.; Fu, C.-W.; and Liu, S. 2023. SIRA-PCR: Sim-to-Real Adaptation for 3D Point Cloud Registration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14394–14405.
- Choy, C.; Park, J.; and Koltun, V. 2019. Fully convolutional geometric features. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8958–8966.
- Du, A.; Huang, X.; Zhang, J.; Yao, L.; and Wu, Q. 2019. Kpsnet: Keypoint detection and feature extraction for point cloud registration. In *2019 IEEE International Conference on Image Processing (ICIP)*, 2576–2580. IEEE.
- Elbaz, G.; Avraham, T.; and Fischer, A. 2017. 3D point cloud registration for localization using a deep neural network auto-encoder. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4631–4640.
- 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.
- Groß, J.; Ošep, A.; and Leibe, B. 2019. Alignnet-3d: Fast point cloud registration of partially observed objects. In *2019 International conference on 3d vision (3DV)*, 623–632. IEEE.
- Gross, S.; Ranzato, M.; and Szlam, A. 2017. Hard mixtures of experts for large scale weakly supervised vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6865–6873.
- Huang, S.; Gojcic, Z.; Usvyatsov, M.; Wieser, A.; and Schindler, K. 2021. Predator: Registration of 3d point clouds with low overlap. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 4267–4276.
- Huang, X.; Fan, L.; Zhang, J.; Wu, Q.; and Yuan, C. 2016. Real time complete dense depth reconstruction for a monocular camera. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 32–37.
- Huang, X.; Mei, G.; and Zhang, J. 2020. Feature-metric registration: A fast semi-supervised approach for robust point cloud registration without correspondences. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11366–11374.
- Huang, X.; Qu, W.; Zuo, Y.; Fang, Y.; and Zhao, X. 2022a. IMFNet: Interpretable multimodal fusion for point cloud registration. *IEEE Robotics and Automation Letters*, 7(4): 12323–12330.
- Huang, X.; Wang, Y.; Li, S.; Mei, G.; Xu, Z.; Wang, Y.; Zhang, J.; and Bennamoun, M. 2022b. Robust real-world point cloud registration by inlier detection. *Computer Vision and Image Understanding*, 224: 103556.
- Izadi, S.; Newcombe, R. A.; Kim, D.; Hilliges, O.; Molyneaux, D.; Hodges, S.; Kohli, P.; Shotton, J.; Davison, A. J.; and Fitzgibbon, A. 2011. Kinectfusion: real-time dynamic 3d surface reconstruction and interaction. In *ACM SIGGRAPH 2011 Talks*, 1–1.
- Jiang, A. Q.; Sablayrolles, A.; Roux, A.; Mensch, A.; Savary, B.; Bamford, C.; Chaplot, D. S.; Casas, D. d. l.; Hanna, E. B.; Bressand, F.; et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- 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 arXiv:2006.16668*.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Mei, G.; Huang, X.; Zhang, J.; and Wu, Q. 2022. Overlap-guided coarse-to-fine correspondence prediction for point cloud registration. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Nagy, B.; and Benedek, C. 2018. Real-time point cloud alignment for vehicle localization in a high resolution 3D map. In *Proceedings of the european conference on computer vision (ECCV) workshops*, 0–0.
- Pavlitckaya, S.; Hubschneider, C.; Weber, M.; Moritz, R.; Huger, F.; Schlicht, P.; and Zollner, M. 2020. Using mixture of expert models to gain insights into semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 342–343.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.

- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30.
- Qin, Z.; Yu, H.; Wang, C.; Guo, Y.; Peng, Y.; and Xu, K. 2022. Geometric transformer for fast and robust point cloud registration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11143–11152.
- Riquelme, C.; Puigcerver, J.; Mustafa, B.; Neumann, M.; Jenatton, R.; Susano Pinto, A.; Keysers, D.; and Houlsby, N. 2021. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems*, 34: 8583–8595.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Thomas, H.; Qi, C. R.; Deschaud, J.-E.; Marcotegui, B.; Goulette, F.; and Guibas, L. J. 2019. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6411–6420.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, H.; Liu, Y.; Dong, Z.; and Wang, W. 2022. You only hypothesize once: Point cloud registration with rotation-equivariant descriptors. In *Proceedings of the 30th ACM International Conference on Multimedia*, 1630–1641.
- Wang, Y.; and Solomon, J. M. 2019. Deep closest point: Learning representations for point cloud registration. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3523–3532.
- Wu, Z.; Song, S.; Khosla, A.; Yu, F.; Zhang, L.; Tang, X.; and Xiao, J. 2015. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1912–1920.
- Xu, H.; Liu, S.; Wang, G.; Liu, G.; and Zeng, B. 2021. Omnet: Learning overlapping mask for partial-to-partial point cloud registration. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3132–3141.
- Yang, F.; Guo, L.; Chen, Z.; and Tao, W. 2022. One-inlier is first: Towards efficient position encoding for point cloud registration. *Advances in Neural Information Processing Systems*, 35: 6982–6995.
- Yew, Z. J.; and Lee, G. H. 2020. Rpm-net: Robust point matching using learned features. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11824–11833.
- Yew, Z. J.; and Lee, G. H. 2022. Regtr: End-to-end point cloud correspondences with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6677–6686.
- Yu, H.; Li, F.; Saleh, M.; Busam, B.; and Ilic, S. 2021. Cofinet: Reliable coarse-to-fine correspondences for robust pointcloud registration. *Advances in Neural Information Processing Systems*, 34: 23872–23884.
- Yu, J.; Ren, L.; Zhang, Y.; Zhou, W.; Lin, L.; and Dai, G. 2023. PEAL: Prior-embedded explicit attention learning for low-overlap point cloud registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17702–17711.
- 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*, 1802–1811.
- Zhou, Y.; Lei, T.; Liu, H.; Du, N.; Huang, Y.; Zhao, V.; Dai, A. M.; Le, Q. V.; Laudon, J.; et al. 2022. Mixture-of-experts with expert choice routing. *Advances in Neural Information Processing Systems*, 35: 7103–7114.