

Transfer and Alignment Network for Generalized Category Discovery

Wenbin An^{1,3}, Feng Tian^{2,3*}, Wenkai Shi^{1,3}, Yan Chen², Yaqiang Wu⁴, Qianying Wang^{4*}, Ping Chen⁵

¹School of Automation Science and Engineering, Xi'an Jiaotong University

²School of Computer Science and Technology, Xi'an Jiaotong University

³National Engineering Laboratory for Big Data Analytics

⁴Lenovo Research

⁵Department of Engineering, University of Massachusetts Boston

{wenbinan, kkkkkkai0611}@stu.xjtu.edu.cn, {fengtian, chenyan}@mail.xjtu.edu.cn

{wuyqe, wangqya}@lenovo.com, Ping.Chen@umb.edu

Abstract

Generalized Category Discovery (GCD) is a crucial real-world task that aims to recognize both known and novel categories from an unlabeled dataset by leveraging another labeled dataset with only known categories. Despite the improved performance on known categories, current methods perform poorly on novel categories. We attribute the poor performance to two reasons: *biased knowledge transfer* between labeled and unlabeled data and *noisy representation learning* on the unlabeled data. The former leads to unreliable estimation of learning targets for novel categories and the latter hinders models from learning discriminative features. To mitigate these two issues, we propose a *Transfer and Alignment Network* (TAN), which incorporates two *knowledge transfer mechanisms* to calibrate the biased knowledge and two *feature alignment mechanisms* to learn discriminative features. Specifically, we model different categories with prototypes and transfer the prototypes in labeled data to correct model bias towards known categories. On the one hand, we pull instances with known categories in unlabeled data closer to these prototypes to form more compact clusters and avoid boundary overlap between known and novel categories. On the other hand, we use these prototypes to calibrate noisy prototypes estimated from unlabeled data based on category similarities, which allows for more accurate estimation of prototypes for novel categories that can be used as reliable learning targets later. After knowledge transfer, we further propose two feature alignment mechanisms to acquire both instance- and category-level knowledge from unlabeled data by aligning instance features with both augmented features and the calibrated prototypes, which can boost model performance on both known and novel categories with less noise. Experiments on three benchmark datasets show that our model outperforms SOTA methods, especially on novel categories. Theoretical analysis is provided for an in-depth understanding of our model in general. Our code and data are available at <https://github.com/Lackel/TAN>.

Introduction

Despite remarkable breakthroughs achieved by modern deep learning systems, the majority of models are designed under a close-world setting, based on the assumption that training

*Corresponding author.




Figure 1: Average accuracy on three benchmark datasets compared to the SOTA method DPN. Accuracy for novel categories (Novel), known categories (Known) and their harmonic mean (H-score) over three runs are reported.

and test data are from the same set of pre-defined categories (Cao, Brbic, and Leskovec 2021). However, many practical problems such as intent detection (Lin, Xu, and Zhang 2020) and relation extraction (Hogan, Li, and Shang 2023) are open-world, where the well-trained models may encounter unlabeled data containing unseen novel categories. To meet the open-world demands, Generalized Category Discovery (GCD) was widely studied in both NLP (Zhang et al. 2021a; An et al. 2023c) and CV fields (Vaze et al. 2022; Wen, Zhao, and Qi 2022). GCD requires models to recognize both known and novel categories from encountered unlabelled data based on a set of labeled data with only known categories, which can adapt models to the increasing number of categories without any additional labeling cost.

Most existing works (Zhang et al. 2021b; An et al. 2023c; Shi et al. 2023) adopt a two-stage approach to address GCD: pre-training on labeled data and then transferring the pre-trained model for pseudo-label training on unlabeled data. Even though these methods have achieved good performance on known categories, they usually perform poorly on novel categories due to the lack of supervision (yellow bar in Fig. 1), which limits their applications in the real world. We attribute the poor performance to two reasons: **Biased knowledge transfer** and **Noisy representation learning**. First, models pretrained on labeled data with only known categories tend to be over-confident and biased towards known categories, so instances with novel categories in unlabeled data can be easily

misclassified into known categories (Fig. 2 (a) Top), where the biased knowledge transfer can lead to noisy estimation of learning targets for novel categories (e.g., biased category prototypes (An et al. 2023c) or unreliable pseudo labels (Zhang et al. 2021a)). Second, under the noisy learning targets, it is usually hard for current models to learn discriminative representations from unlabeled data effectively. Moreover, these methods solely focus on instance- (Vaze et al. 2022) or prototype-based discrimination (An et al. 2023c) with the noisy learning targets, but fail to combine them together to capture both instance- and category-level semantics, which can further disrupt the performance.

To mitigate above issues, we propose a *Transfer and Alignment Network* (TAN) to calibrate the biased knowledge and learn discriminative features. We first propose two knowledge transfer mechanisms to calibrate the biased knowledge caused by pre-training. Specifically, we model different categories with prototypes (Snell, Swersky, and Zemel 2017), then we leverage prototypes of known categories in labeled data as a prior to guide the training process on unlabeled data. On the one hand, we propose *Prototype-to-Instance Transfer* (**P2I Trans**) to cluster instances with known categories in unlabeled data around these prototypes to form compact clusters, which can make different categories discriminative and avoid boundary overlap between known and novel categories. On the other hand, inspired by the fact that similar categories may share some common features (e.g., cats and dogs), we propose *Prototype-to-Prototype Transfer* (**P2P Trans**) to transfer these prototypes to calibrate the noisy prototypes estimated from unlabeled data. To avoid negative transfer between dissimilar categories, we introduce semantic similarities between categories as weights and select only the most similar k prototypes for the prototype calibration. By transferring knowledge from known to novel categories, the calibrated prototypes can be used as reliable learning targets to guide the subsequent representation learning. Combining P2I and P2P Trans, TAN can learn clear decision boundaries for known categories and reliable prototypes for novel categories, which can help to alleviate the effects of biased knowledge transfer (Fig. 2 (a) Bottom).

After knowledge transfer, we further propose two feature alignment mechanisms to acquire both instance- and category-level knowledge from unlabeled data with less noise. First, we propose *Instance-to-Prototype Alignment* (**I2P Align**) to pull instance features closer to the corresponding calibrated prototypes to acquire category-level knowledge (i.e., common category semantics embedded in prototypes from multiple instances), so that these features can be discriminative and compact around the prototypes to form more distinguishable decision boundaries. Second, we propose *Instance-to-Instance Alignment* (**I2I Align**) to align instance features with their augmented features to acquire instance-level knowledge (i.e., specific instance semantics embedded in instance features), so that these features can be self-consistent and locally smooth for better representation learning (Chen et al. 2020). As shown in Figure 1, our model achieves the best performance on three benchmark datasets, and the improved performance on novel categories further validates the effectiveness of our model. Last but not least,

we theoretically justify the effectiveness of our model.

Our main contributions can be summarized as follows:

- We propose a *Transfer and Alignment Network* (TAN) to mitigate the performance gap between known and novel categories in GCD.
- We propose two knowledge transfer mechanisms to alleviate the effects of biased knowledge transfer and two feature alignment mechanisms to acquire both instance- and category-level knowledge with less noise.
- Extensive experiments show that our model outperforms SOTA methods, and theoretical analysis further validates the effectiveness of our model.

Related Work

Generalized Category Discovery

Generalized Category Discovery (GCD) is a practical and challenging task formalised by Vaze et al. (2022). Under the open-world setting, GCD assumes that the newly collected unlabeled data contain novel categories that have never been seen during training. Due to the lack of supervision for novel categories, previous methods mainly employed pseudo-labeling based methods (Han, Vedaldi, and Zisserman 2019; Lin, Xu, and Zhang 2020; Ge, Chen, and Li 2020) or self-supervised methods (Zhong et al. 2021; Wen, Zhao, and Qi 2022; Zhang et al. 2022) to learn from unlabeled data. For example, Caron et al. (2018) proposed to generate pseudo labels by clustering, Zhang et al. (2021a) and An et al. (2023a) improved this method by generating more robust and consistent pseudo labels. As for self-supervised methods, Vaze et al. (2022) proposed a semi-supervised k-means framework with contrastive learning to learn discriminative features, and An et al. (2023c) proposed a decoupled prototypical network to decouple known and novel categories from unlabeled data. Another line of work focuses on the discovery of novel fine-grained categories from coarsely-labeled data (An et al. 2022, 2023b; Ma et al. 2023), which can transfer knowledge between different label hierarchies. Despite the improved overall performance, most of these methods performed poorly on novel categories because of the biased knowledge transfer and noisy representation learning, which may limit their applications in the real world.

Transfer Learning

Transfer learning aims at transferring knowledge from the source domain to boost model performance on the target domain (Zhuang et al. 2020; Chen and Su 2023). GCD is related to transfer learning since we need to transfer knowledge from known to novel categories. Most of current methods are based on the pre-training and fine-tuning paradigm to transfer knowledge implicitly by initializing model parameters (Zhang et al. 2021a). For example, Vaze et al. (2022) proposed to use contrastive learning to pre-train their model, and Zhang et al. (2022) combined supervised learning and masked language modeling to initialize their model. However, we think this paradigm is sub-optimal for GCD because models pre-trained on labeled data tend to be biased towards known categories.

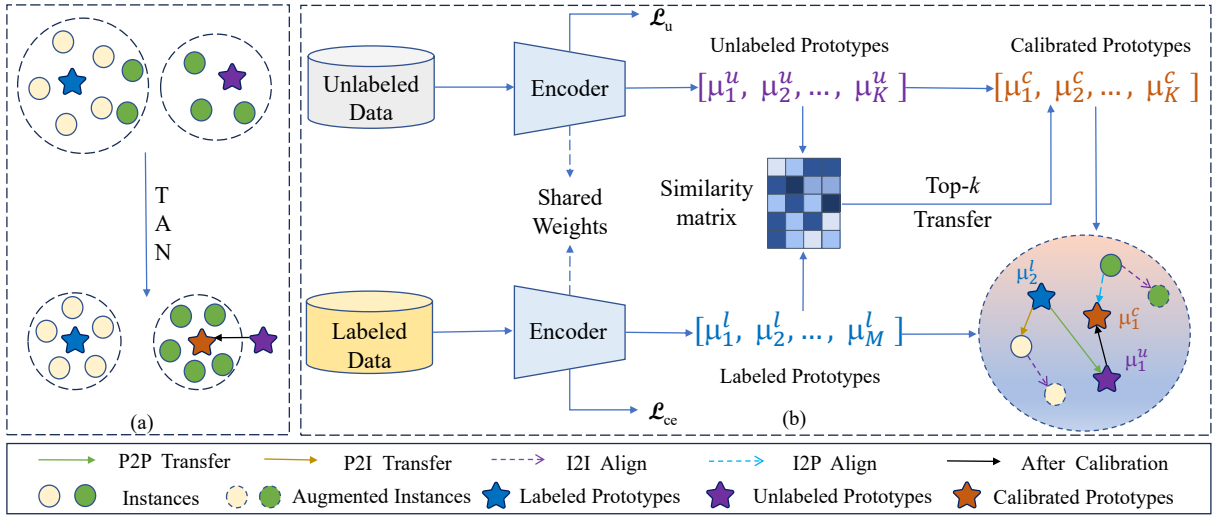


Figure 2: (a) An illustration of the biased knowledge transfer (Top) and the effectiveness of our knowledge transfer mechanisms (Bottom). (b) An overview of our model.

Method

In this section, we first formulate the Generalized Category Discovery (GCD) task. Then we introduce our *Transfer and Alignment Network* (TAN) in detail. Specifically, we first pre-train a feature encoder and learn category prototypes from both labeled and unlabeled data. Then we propose two knowledge transfer mechanisms to learn clear decision boundaries for known categories and reliable prototypes for novel categories, which can help to mitigate the effects of biased knowledge transfer. After knowledge transfer, we propose two feature alignment mechanisms to capture both instance- and category-level semantics to learn discriminative features. The framework of our model is shown in Fig. 2 (b).

Problem Formulation

Models trained on a labeled dataset $\mathcal{D}^l = \{(x_i, y_i) | y_i \in \mathcal{Y}_k\}$ can recognize pre-defined known categories \mathcal{Y}_k well. However, in the open world, the trained models may encounter unlabeled data $\mathcal{D}^u = \{x_i | y_i \in \{\mathcal{Y}_k, \mathcal{Y}_n\}\}$ that contain both known categories \mathcal{Y}_k and novel categories \mathcal{Y}_n , which can make the models fail. To cope with this limitation, Generalized Category Discovery (GCD) requires models to recognize both known and novel categories based on \mathcal{D}^l and \mathcal{D}^u , without any annotation for novel categories. We denote $M = |\mathcal{Y}_k|$ as the number of known categories and $K = |\mathcal{Y}_k| + |\mathcal{Y}_n|$ as the number of all categories. Finally, model performance will be measured on a testing set $\mathcal{D}^t = \{(x_i, y_i) | y_i \in \{\mathcal{Y}_k, \mathcal{Y}_n\}\}$.

Pre-training and Prototype Learning

We use the pre-trained BERT (Devlin et al. 2018) as our feature encoder F_θ to extract feature $z_i = F_\theta(x_i)$ for the input x_i . To adapt the pre-trained model to the downstream GCD task, we use cross-entropy loss on labeled data and masked language modeling (mlm) loss (Zhang et al. 2021b) on unlabeled data to pre-train F_θ , where the mlm loss can help

to learn general knowledge and reduce model bias towards known categories.

To transfer knowledge between different categories and acquire reliable learning targets for unlabeled data, we model categories with prototypes (Snell, Swersky, and Zemel 2017) and learn two sets of prototypes from the labeled and unlabeled dataset, respectively. For the labeled dataset, we take average of all instance features belonging to the same category as labeled prototypes $P^l = \{\mu_j^l\}_{j=1}^M$, where $\mu_j^l = \frac{1}{|\mathcal{D}_j^l|} \sum_{x_i \in \mathcal{D}_j^l} F_\theta(x_i)$ and \mathcal{D}_j^l is a set of labeled instances from the category j . For unlabeled data, we follow An et al. (2023c) to perform clustering and utilize cluster centers as estimated unlabeled prototypes $P^u = \{\mu_j^u\}_{j=1}^K$, where $\mu_j^u = \frac{1}{|\mathcal{D}_j^u|} \sum_{x_i \in \mathcal{D}_j^u} F_\theta(x_i)$ and \mathcal{D}_j^u is a set of unlabeled instances belonging to the cluster j .

Knowledge Transfer

To transfer knowledge between known and novel categories and alleviate the effects of biased knowledge transfer, we propose two knowledge transfer mechanisms called *Prototype-to-Prototype Transfer (P2P Trans)* and *Prototype-to-Instance Transfer (P2I Trans)*.

P2P Trans. Due to the lack of supervision and the model bias towards known categories, the estimated unlabeled prototypes can be noisy and biased, especially for novel categories. So directly performing prototypical learning (Snell, Swersky, and Zemel 2017) on these noisy prototypes can lead to inferior results. To mitigate this issue, we propose to transfer knowledge between categories and treat the labeled prototypes as unbiased estimation for known categories to calibrate the noisy unlabeled prototypes. To avoid negative transfer between dissimilar categories, we introduce semantic similarities between categories as weights and select only the most similar k prototypes for the prototype calibration. Specifically, we measure the semantic similarity of two categories based

on the Euclidean distance of the corresponding prototypes:

$$S_i = \{-\|\mu_i^u - \mu_j^l\|_2 \mid \mu_j^l \in P^l\} \quad (1)$$

where μ_i^u is the i -th unlabeled prototype and μ_j^l is the j -th labeled prototype. Then we select the top- k similar labeled prototypes as the transfer set S'_i for each unlabeled prototype:

$$T_i = \{j \mid S_{ij} \in \text{top}_k(S_i)\} \quad (2)$$

$$S'_i = \{S_{ij} \mid j \in T_i\} \quad (3)$$

where S_{ij} is the j -th element of S_i . Then we calculate the transfer weights by normalizing similarities in the transfer set:

$$w_i = \text{softmax}(S'_i / \sqrt{\text{dim}(z)}) \quad (4)$$

where $\text{dim}(z)$ is the feature dimension. Then we use the transfer set and weights to calibrate each unlabeled prototype. The calibrated prototypes $P^c = \{\mu_i^c\}_{i=1}^K$ can be estimated as follows:

$$\mu_i^c = \alpha \cdot \mu_i^u + (1 - \alpha) \cdot \sum_{j \in T_i} w_{ij} \cdot \mu_j^l \quad (5)$$

where α is a weighting factor, w_{ij} is the j -th element of w_i . By using category similarities as weights, we can transfer knowledge from labeled prototypes to calibrate the noisy unlabeled prototypes, where the calibrated prototypes can be used as reliable learning targets for unlabeled data later.

P2I Trans. Since the labeled prototypes are learned from ground-truth labels, they can be viewed as unbiased estimation for known categories. So we further utilize these labeled prototypes to guide the training process of unlabeled data with known categories, based on the pseudo labels from clustering. Specifically, we need to first match the prototypes for known categories in labeled and unlabeled data. Following the assumption that the closest prototypes in the feature space represent the same category, we can find the optimal match function \mathcal{P} through a bipartite matching algorithm introduced by An et al. (2023c), where μ_i^l and $\mu_{\mathcal{P}(i)}^u$ represent the same known category. Then we cluster unlabeled instances with known categories around the corresponding prototypes:

$$\mathcal{L}_{p2i} = \frac{1}{N} \sum_{i=1}^M \sum_{x_j \in \mathcal{D}_{\mathcal{P}(i)}^u} \|F_{\theta}(x_j) - \mu_i^l\|_2 \quad (6)$$

where N is the number of unlabeled data belonging to known categories based on clustering results. $\mathcal{D}_{\mathcal{P}(i)}^u$ is a set of unlabeled data belonging to the cluster $\mathcal{P}(i)$ (i.e., the i -th known category). In this way, we can form compact clusters and distinguishable decision boundaries for known categories.

Discussion. The knowledge transfer mechanisms are effective towards the problem of biased knowledge transfer in two aspects. First, we can calibrate the biased unlabeled prototypes and estimate more reliable prototypes through the P2P Trans (Sec. Prototype Calibration), which can help to correct instances that are biased to known categories. Second, we can form compact clusters and clear decision boundaries for known categories through the P2I Trans, which can also

help to correct the biased instances. Combining P2P and P2I Trans, our model can learn more reliable prototypes and more distinguishable decision boundaries for different categories, which can make them more discriminative and alleviate the effects of biased knowledge transfer (Fig. 2 (a) Bottom).

Feature Alignment

After knowledge transfer, we further propose two feature alignment mechanisms called *Instance-to-Prototype Alignment (I2P Align)* and *Instance-to-Instance Alignment (I2I Align)* to learn discriminative features from unlabeled data.

I2P Align. After P2P Trans, the calibrated prototypes are less noisy and can be treated as reliable learning targets for unlabeled data. So we propose to acquire category-level knowledge by pulling instance features of unlabeled data closer to the corresponding calibrated prototypes based on the clustering results, so that these features can be discriminative and compact around the prototypes to form more distinguishable decision boundaries for different categories:

$$\mathcal{L}_{i2p} = -\frac{1}{|\mathcal{D}^u|} \sum_{i=1}^K \sum_{x_j \in \mathcal{D}_i^u} \|F_{\theta}(x_j) - \mu_i^c\|_2 \quad (7)$$

where $|\mathcal{D}^u|$ is the number of unlabeled instances, \mathcal{D}_i^u is a set of unlabeled instances belonging to the cluster i .

I2I Align. In addition to I2P Align, we also propose Instance-to-Instance Alignment (I2I Align) to acquire instance-level knowledge from unlabeled data. Specifically, we use data augmentation methods to generate augmented instance x'_i for each instance x_i . Then we perform instance-wise contrastive learning (Oord, Li, and Vinyals 2018) to pull x_i closer to x'_i and push other instances away from x_i , which can help to learn self-consistent and locally-smooth features for better representation learning (Li et al. 2020):

$$\mathcal{L}_{i2i} = -\frac{1}{|\mathcal{D}^u|} \sum_{i=1}^{|\mathcal{D}^u|} \log \frac{\exp(z_i \cdot z'_i / \tau)}{\sum_{j=1}^{2B} \exp(z_i \cdot z'_j / \tau)} \quad (8)$$

where z'_i is the feature of x'_i , τ is a temperature hyper-parameter, and B is the batch size. We further add cross-entropy loss \mathcal{L}_u on unlabeled data to learn more discriminative features with cluster ids as pseudo labels (Caron et al. 2018). To avoid catastrophic forgetting for knowledge acquired from labeled data and mitigate the effects of label noise in unlabeled data (An et al. 2023a), we also add cross-entropy loss \mathcal{L}_{ce} on labeled data with ground-truth labels.

Overall Loss. The objective of our model is defined as:

$$\mathcal{L}_{TAN} = \mathcal{L}_{p2i} + \mathcal{L}_{i2p} + \mathcal{L}_{i2i} + \mathcal{L}_u + \beta \mathcal{L}_{ce} \quad (9)$$

where β is a weighting factor. By combining instance- and prototype-based learning, our model can capture both instance- and category-level semantics from unlabeled data, which can help to learn discriminative representations. In summary, our model can mitigate the problem of biased knowledge transfer by transferring knowledge from labeled to unlabeled data and alleviate the noisy representation learning problem by acquiring both instance- and category-level knowledge, which can help to learn clear decision boundaries for different categories and boost our model performance.

Method	BANKING			StackOverflow			CLINC		
	H-score	Known	Novel	H-score	Known	Novel	H-score	Known	Novel
DeepCluster	13.97	13.94	13.99	19.10	18.22	14.80	26.48	27.34	25.67
DCN	16.33	18.94	14.35	29.22	28.94	29.51	29.20	30.00	28.45
DEC	17.82	20.36	15.84	25.99	26.20	25.78	19.78	20.18	19.40
KM-BERT	21.08	21.48	20.70	16.93	16.67	17.20	34.05	34.98	33.16
KM-GloVe	29.25	29.11	29.39	28.32	28.60	28.05	51.62	51.74	51.50
AG-GloVe	30.47	29.69	31.29	29.95	28.49	31.56	44.16	45.17	43.20
SAE	37.77	38.29	37.27	62.65	57.36	69.02	45.74	47.35	44.24
Simple	40.52	49.96	34.08	57.53	57.87	57.20	62.76	70.60	56.49
Semi-DC	47.40	53.37	42.63	64.90	63.57	61.20	73.41	75.60	71.34
Self-Labeling	48.19	61.64	39.56	59.99	78.53	48.53	61.29	80.06	49.65
CDAC+	50.28	55.42	46.01	75.78	77.51	74.13	69.42	70.08	68.77
DTC	52.13	59.98	46.10	63.22	80.93	51.87	68.71	82.34	58.95
Semi-KM	54.83	73.62	43.68	61.43	81.02	49.47	70.98	89.03	59.01
DAC	54.98	69.60	45.44	63.64	76.13	54.67	78.77	89.10	70.59
GCD	55.78	75.16	44.34	64.63	82.00	53.33	63.08	89.64	48.66
PTJN	60.69	77.20	50.00	77.48	72.80	82.80	83.34	91.79	76.32
DPN	60.73	80.93	48.60	83.13	85.29	81.07	84.56	92.97	77.54
TAN (OC)	65.49	81.21	54.87	85.20	84.80	85.60	86.20	92.74	80.53
TAN (Ours)	66.70	81.97	56.23	86.64	86.36	86.93	87.02	93.39	81.46
Improvement	+5.97	+1.04	+6.23	+3.51	+1.07	+4.13	+2.46	+0.42	+3.92

Table 1: Average results (%) over 3 runs on the testing set. OC means over clustering.

Theoretical Analysis

We formalize the error bound of our model with the theory of unsupervised domain adaptation (Pan et al. 2019; Ben-David et al. 2010). Training on both labeled and unlabeled data, the classification error of GCD can be written as the linear weighted sum of errors on labeled and unlabeled data:

$$\epsilon(h) = \gamma \cdot \epsilon_u(h, \hat{y}^u) + (1 - \gamma) \cdot \epsilon_l(h, y^l) \quad (10)$$

where h is a hypothesis, γ is a weighting factor, $\epsilon_u(h, \hat{y}^u) = E_{x \sim \mathcal{D}^u} |h(x) - \hat{y}^u|$ and $\epsilon_l(h, y^l) = E_{x \sim \mathcal{D}^l} |h(x) - y^l|$ represent the error over the sample distribution of unlabeled data \mathcal{D}^u with pseudo labels \hat{y}^u and labeled data \mathcal{D}^l with ground-truth labels y^l , respectively. Then we want to analyse how close the error $\epsilon(h)$ is to an oracle error $\epsilon_u(h, y^u)$ that evaluates the model learned on the unlabeled data with ground truth labels y^u . Following the analysis in Pan et al. (2019), difference between the two losses can be bounded by the following Lemma.

Lemma 1 *Let h be a hypothesis in class \mathcal{H} . Then*

$$|\epsilon(h) - \epsilon_u(h, y^u)| \leq (1 - \gamma) \left(\frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^l, \mathcal{D}^u) + \lambda \right) + \gamma \rho \quad (11)$$

where $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^l, \mathcal{D}^u) = 2 \sup_{h, h' \in \mathcal{H}} |\epsilon_u(h, h') - \epsilon_l(h, h')|$ measures the domain discrepancy between the labeled and unlabeled data in the hypothesis space \mathcal{H} . $\lambda = \epsilon_l(h^*, y^l) + \epsilon_u(h^*, y^u)$ is the total error on the labeled and unlabeled data with the joint optimal hypothesis h^* . And ρ denotes the ratio of false pseudo labels for unlabeled data.

From the Lemma 1 we can see that the bound are decided by three terms. First, the term λ is negligibly small

with the joint optimal hypothesis h^* and ground-truth labels, which can be disregard. Second, for the discrepancy term $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^l, \mathcal{D}^u)$ that can be quantified by category-level discrepancy of prototypes (Pan et al. 2019), the main discrepancy of the labeled and unlabeled data is from the differences between known and novel categories. Our **knowledge transfer mechanisms** can mitigate this discrepancy by transferring knowledge from known to novel categories, where the P2P Trans can help to estimate more reliable prototypes for novel categories and mitigate the discrepancy between known and novel categories. Third, the ratio of false pseudo labels ρ can be gradually reduced by learning more discriminative representations during training. Our **feature alignment mechanisms** can capture both instance- and category-level semantics to learn more discriminative features, which can reduce the noise of pseudo labels (Sec. Accuracy of Pseudo Labels). In summary, our knowledge transfer and feature alignment mechanisms can help to tighten the bound in Lemma 1, which can prove the effectiveness of our model theoretically.

Experiments

Experimental Setup

Datasets. We validate the effectiveness of our model on three benchmark datasets. **BANKING** is an intent detection dataset in the bank domain (Casanueva et al. 2020). **StackOverflow** is a question classification dataset processed by Xu et al. (2015). **CLINC** is a text classification dataset from diverse domains (Larson et al. 2019). For each dataset, we randomly select 25% categories as novel categories and 10% data as labeled data.

Variant	H-score	Known	Novel
TAN	66.70	81.97	56.23
w/o <i>P2I Trans</i>	65.46	80.46	55.18
w/o <i>P2P Trans</i>	63.45	82.54	51.53
w/o \mathcal{L}_{ce}	62.70	81.45	50.96
w/o <i>I2I Align</i>	61.26	80.19	49.56
w/o \mathcal{L}_u	60.11	82.50	47.28
w/o <i>I2P Align</i>	59.10	84.51	45.44

Table 2: Ablation study with different model variants.

Comparison with SOTA. We compare the proposed model with various baselines and SOTA methods.

Unsupervised Models. (1) DeepCluster: Deep Clustering (Caron et al. 2018). (2) DCN: Deep Clustering Network (Yang et al. 2017). (3) DEC: Deep Embedding Clustering (Xie, Girshick, and Farhadi 2016). (4) KM-BERT: KMeans with BERT embeddings (Devlin et al. 2018). (5) KM-GloVe: KMeans (MacQueen et al. 1967) with GloVe embeddings (Pennington, Socher, and Manning 2014). (6) AG-GloVe: Agglomerative Clustering (Gowda and Krishna 1978) with GloVe embeddings. (7) SAE: Stacked Auto Encoder.

Semi-supervised Models. (1) Simple: A Simple Parametric model (Wen, Zhao, and Qi 2022). (2) Semi-DC: Deep Clustering (Caron et al. 2018) pretrained on labeled data. (3) Self-Labeling: Self-Labeling Framework (Yu et al. 2022). (4) CDAC+: Constrained Adaptive Clustering (Lin, Xu, and Zhang 2020). (5) DTC: Deep Transfer Clustering (Han, Vedaldi, and Zisserman 2019). (6) Semi-KM: KMeans with BERT pretrained on labeled data. (7) DAC: Deep Aligned Clustering (Zhang et al. 2021a). (8) GCD: Label Assignment with Semi-supervised KMeans (Vaze et al. 2022). (9) PTJN: Robust Pseudo-label Training (An et al. 2023a). (10) DPN: Decoupled Prototypical Network (An et al. 2023c).

Evaluation Metrics. We measure model performance with clustering accuracy on the testing set with Hungarian algorithm (Kuhn 1955). (1) **H-score**: harmonic mean of the accuracy for known and novel categories, which can avoid evaluation bias towards known categories (Saito and Saenko 2021). (2) **Known**: accuracy for instances with known categories. (3) **Novel**: accuracy for instances with novel categories.

Implementation Details. We use the pretrained bert-base-uncased model (Wolf et al. 2019) and adopt its suggested hyper-parameters. We only fine-tune the last four Transformer layers with AdamW optimizer. Early stopping is used during pretraining with wait patience 20. For hyper-parameters, k is set to 5, α is set to 0.8, β is set to 100 and τ is set to 0.07. Training epochs for StackOverflow, BANKING and CLINC dataset are set to $\{10, 20, 20\}$. The learning rate for pretraining and training is set to $5e^{-5}$ and $1e^{-5}$, respectively. For masked language modeling, the mask probability is set to 0.15 following previous works. And SimCSE (Gao, Yao, and Chen 2021) is used to generate augmented instances.

Query Category	Selected Categories
Play_Music Timezone	Play_Music, Next_Song, Whisper Timezone, Travel_Alert, Time
Timer	Alarm, Reminder_Update, Time
Book_Flight	Flight_Status, Book_Hotel, Car_Rental

Table 3: Examples of the top-3 similar categories.

Results and Discussion

Main Results. We show the results in Table 1. From the results we can get following observations. First, our model gets the best performance on all evaluation metrics and datasets, which can show the effectiveness of our model. Second, our model achieves the best results on H-score (average **3.98%** improvement), which means that our model can better balance the performance on known and novel categories and alleviate the effects of model bias towards known categories. Third, our model achieves the best performance on accuracy for known categories (average **0.84%** improvement). Thanks to the knowledge transfer and feature alignment mechanisms, our model can form compact clusters and discriminative decision boundaries for known categories, which means that our model can boost model performance on novel categories without sacrificing the model performance on known categories. Last but not least, our model achieves the best performance on accuracy for novel categories (average **4.76%** improvement). We attribute the significant improvement to following reasons. First, our knowledge transfer mechanisms (*P2I Trans* and *P2P Trans*) can help to alleviate the effects of biased knowledge transfer by calibrating the noisy prototypes and forming clear decision boundaries for both known and novel categories, where the calibrated prototypes can be used as reliable learning targets for the subsequent training. Second, our feature alignment mechanisms (*I2P Align* and *I2I Align*) can alleviate the effects of noisy representation learning by acquiring instance- and category-level knowledge simultaneously. And under the guidance of the calibrated prototypes, our model can learn discriminative features to form compact clusters with less noise.

Ablation Study. We inspect the contribution of different components to our model on the BANKING dataset in Table 2. First, removing different components from TAN degrades model performance on novel categories and H-score, which can show the effectiveness of different components towards mitigating the model bias to known categories and boosting model performance on novel categories. Removing components related to representation learning (*I2I Align*, *I2P Align* and \mathcal{L}_u) has the greatest impact on the performance, which means that learning discriminative features is crucial for novel categories. Second, removing *P2I Trans* and *I2I Align* degrades model performance on known categories since they are responsible for learning compact and discriminative clusters for known categories. Even though removing some components (*P2P Trans*, *I2P Align* and \mathcal{L}_u) can improve the performance on known categories, they can also greatly exacerbate the model bias and degrade the performance on




Figure 3: Average distance between the ground-truth prototypes and the prototypes before and after calibration.

Dataset	BANKING	StackOverflow	CLINC
Ground Truth	77	20	150
Ours	74	19	148

Table 4: Estimation of the number of categories.

novel categories. In summary, our model can balance model performance on known and novel categories by mitigating the model bias and learning discriminative features.

Prototype Calibration. We investigate the effectiveness of our Prototype Calibration (*P2P Trans*) mechanism by answering the following questions. **(1) Can prototype distances measure semantic similarities between categories?** In Table 3, We show the top-3 selected known categories in Eq. (3) for both known (Top) and novel (Bottom) categories, based on the prototype distance. From the results we can see that the selected categories are highly relevant to query categories, which means that modeling categories with prototypes can preserve semantic similarities between categories. And by measuring distances between prototypes, we can transfer knowledge between similar categories and calibrate the noisy prototypes. **(2) Can prototype calibration help to learn better prototypes?** In Fig. 3, we compare the average distance between the ground-truth prototypes and the prototypes before and after calibration. We can see that our model can learn prototypes that are closer to the ground-truth prototypes after calibration, which means that the prototype calibration can help to estimate more accurate and reliable prototypes.

Real-world Applications. In the real world, the number of categories K is usually unknown. We show the robustness of our model towards this real-world setting from two aspects. **(1) Number of categories estimation.** We report the results on estimating the number of categories with the filtering algorithm (Zhang et al. 2021a) in Table 4. We can see that our estimations come very close to the ground-truth number of categories, which can show the effectiveness of our model. **(2) Over Clustering.** To investigate the sensitivity of our model to the number of categories, we over-estimate the number of categories used for training and testing by a factor of one point two. As shown in Table 1, our model (TAN (OC)) gets close performance even without knowing the ground-truth number of categories, which can show the robustness

Method	CLINC	StackOverflow	BANKING
PTJN	80.23	72.32	67.50
DPN	80.89	77.16	68.34
TAN (Ours)	82.67	81.37	69.05

Table 5: Accuracy of pseudo labels.




Figure 4: The t-SNE visualization of embeddings.

of our model towards the real-world settings.

Accuracy of Pseudo Labels. We report the accuracy of pseudo labels generated by different models for unlabeled data in Table 5. Our model gets the highest accuracy, so the ratio of false pseudo labels ρ in the Sec. Theoretical Analysis can be controlled by our model, which can verify the validity of our theoretical analysis.

Visualization. We visualize the learned embeddings of our model before and after training with t-SNE in Fig. 4. From the figure we can see that novel categories are mixed together before training. And the clusters are more distinguishable after training, especially for novel categories, which indicates that our model can learn discriminative features and form distinguishable decision boundaries for different categories.

Conclusion

In this paper, we propose *Transfer and Alignment Network* for GCD, which incorporates two knowledge transfer mechanisms to mitigate the effects of biased knowledge transfer and two feature alignment mechanisms to learn discriminative features with less noise. By modeling different categories with prototypes and transferring knowledge from labeled to unlabeled data, our model can calibrate the noisy prototypes for novel categories and learn more discriminative clusters for known categories, which can help to mitigate the model bias towards known categories. After knowledge transfer, our model can acquire both instance- and category-level knowledge by aligning instance features with both augmented features and the calibrated prototypes, which can help to learn more discriminative features and form more distinguishable decision boundaries for different categories. Experimental results on three benchmark datasets show that our model outperforms SOTA methods, especially for model performance on novel categories. And the theoretical analysis further justifies the effectiveness of our model.

Acknowledgments

This work was supported by National Key Research and Development Program of China (2022ZD0117102), National Natural Science Foundation of China (62293551, 62177038, 62277042, 62137002, 61721002, 61937001, 62377038). Innovation Research Team of Ministry of Education (IRT_17R86), Project of China Knowledge Centre for Engineering Science and Technology, "LENOVO-XJTU" Intelligent Industry Joint Laboratory Project.

References

- An, W.; Tian, F.; Chen, P.; Tang, S.; Zheng, Q.; and Wang, Q. 2022. Fine-grained Category Discovery under Coarse-grained supervision with Hierarchical Weighted Self-contrastive Learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 1314–1323.
- An, W.; Tian, F.; Chen, P.; Zheng, Q.; and Ding, W. 2023a. New User Intent Discovery with Robust Pseudo Label Training and Source Domain Joint-training. *IEEE Intelligent Systems*.
- An, W.; Tian, F.; Shi, W.; Chen, Y.; Zheng, Q.; Wang, Q.; and Chen, P. 2023b. DNA: Denoised Neighborhood Aggregation for Fine-grained Category Discovery. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 12292–12302.
- An, W.; Tian, F.; Zheng, Q.; Ding, W.; Wang, Q.; and Chen, P. 2023c. Generalized category discovery with decoupled prototypical network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 12527–12535.
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine learning*, 79: 151–175.
- Cao, K.; Brbic, M.; and Leskovec, J. 2021. Open-world semi-supervised learning. *arXiv preprint arXiv:2102.03526*.
- Caron, M.; Bojanowski, P.; Joulin, A.; and Douze, M. 2018. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 132–149.
- Casanueva, I.; Temčin, T.; Gerz, D.; Henderson, M.; and Vulić, I. 2020. Efficient intent detection with dual sentence encoders. *arXiv preprint arXiv:2003.04807*.
- Chen, J.; and Su, B. 2023. Transfer Knowledge from Head to Tail: Uncertainty Calibration under Long-tailed Distribution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19978–19987.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Gao, T.; Yao, X.; and Chen, D. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- Ge, Y.; Chen, D.; and Li, H. 2020. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. *arXiv preprint arXiv:2001.01526*.
- Gowda, K. C.; and Krishna, G. 1978. Agglomerative clustering using the concept of mutual nearest neighbourhood. *Pattern recognition*, 10(2): 105–112.
- Han, K.; Vedaldi, A.; and Zisserman, A. 2019. Learning to discover novel visual categories via deep transfer clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8401–8409.
- Hogan, W.; Li, J.; and Shang, J. 2023. Open-world Semi-supervised Generalized Relation Discovery Aligned in a Real-world Setting. *arXiv preprint arXiv:2305.13533*.
- Kuhn, H. W. 1955. The Hungarian method for the assignment problem. *Naval research logistics quarterly*, 2: 83–97.
- Larson, S.; Mahendran, A.; Peper, J. J.; Clarke, C.; Lee, A.; Hill, P.; Kummerfeld, J. K.; Leach, K.; Laurenzano, M. A.; Tang, L.; et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. *arXiv preprint arXiv:1909.02027*.
- Li, J.; Zhou, P.; Xiong, C.; and Hoi, S. C. 2020. Prototypical contrastive learning of unsupervised representations. *arXiv preprint arXiv:2005.04966*.
- Lin, T.-E.; Xu, H.; and Zhang, H. 2020. Discovering new intents via constrained deep adaptive clustering with cluster refinement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 8360–8367.
- Ma, R.; Lin, Z.; Chen, X.; Zhou, X.; Wang, J.; Gui, T.; Zhang, Q.; Gao, X.; and Chen, Y. W. 2023. Coarse-to-fine few-shot learning for named entity recognition. In *Findings of the Association for Computational Linguistics: ACL 2023*, 4115–4129.
- MacQueen, J.; et al. 1967. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, 281–297. Oakland, CA, USA.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Pan, Y.; Yao, T.; Li, Y.; Wang, Y.; Ngo, C.-W.; and Mei, T. 2019. Transferrable prototypical networks for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2239–2247.
- Pennington, J.; Socher, R.; and Manning, C. D. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Saito, K.; and Saenko, K. 2021. Ovanet: One-vs-all network for universal domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9000–9009.
- Shi, W.; An, W.; Tian, F.; Zheng, Q.; Wang, Q.; and Chen, P. 2023. A Diffusion Weighted Graph Framework for New Intent Discovery. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 8033–8042.

- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30.
- Vaze, S.; Han, K.; Vedaldi, A.; and Zisserman, A. 2022. Generalized category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7492–7501.
- Wen, X.; Zhao, B.; and Qi, X. 2022. A Simple Parametric Classification Baseline for Generalized Category Discovery. *arXiv preprint arXiv:2211.11727*.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2019. HuggingFace’s Transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Xie, J.; Girshick, R.; and Farhadi, A. 2016. Unsupervised deep embedding for clustering analysis. In *International conference on machine learning*, 478–487. PMLR.
- Xu, J.; Wang, P.; Tian, G.; Xu, B.; Zhao, J.; Wang, F.; and Hao, H. 2015. Short text clustering via convolutional neural networks. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, 62–69.
- Yang, B.; Fu, X.; Sidiropoulos, N. D.; and Hong, M. 2017. Towards k-means-friendly spaces: Simultaneous deep learning and clustering. In *international conference on machine learning*, 3861–3870. PMLR.
- Yu, Q.; Ikami, D.; Irie, G.; and Aizawa, K. 2022. Self-Labeling Framework for Novel Category Discovery over Domains. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Zhang, H.; Xu, H.; Lin, T.-E.; and Lyu, R. 2021a. Discovering New Intents with Deep Aligned Clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Zhang, J.; Bui, T.; Yoon, S.; Chen, X.; Liu, Z.; Xia, C.; Tran, Q. H.; Chang, W.; and Yu, P. 2021b. Few-shot intent detection via contrastive pre-training and fine-tuning. *arXiv preprint arXiv:2109.06349*.
- Zhang, Y.; Zhang, H.; Zhan, L.-M.; Wu, X.-M.; and Lam, A. 2022. New intent discovery with pre-training and contrastive learning. *arXiv preprint arXiv:2205.12914*.
- Zhong, Z.; Fini, E.; Roy, S.; Luo, Z.; Ricci, E.; and Sebe, N. 2021. Neighborhood contrastive learning for novel class discovery. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10867–10875.
- Zhuang, F.; Qi, Z.; Duan, K.; Xi, D.; Zhu, Y.; Zhu, H.; Xiong, H.; and He, Q. 2020. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1): 43–76.