

Fine-grained Category Discovery under Coarse-grained supervision with Hierarchical Weighted Self-contrastive Learning

Wenbin An¹, Feng Tian², Ping Chen³, Siliang Tang⁴, Qinghua Zheng², QianYing Wang⁵

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

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

³ Department of Engineering, University of Massachusetts Boston

⁴ College of Computer Science and Technology, Zhejiang University ⁵ Lenovo Research
wenbinan@stu.xjtu.edu.cn, {fengtian, qhzheng}@mail.xjtu.edu.cn
ping.chen@umb.edu, siliang@zju.edu.cn, wangqya@lenovo.com

Abstract

Novel category discovery aims at adapting models trained on known categories to novel categories. Previous works only focus on the scenario where known and novel categories are of the same granularity. In this paper, we investigate a new practical scenario called Fine-grained Category Discovery under Coarse-grained supervision (FCDC). FCDC aims at discovering fine-grained categories with only coarse-grained labeled data, which can adapt models to categories of different granularity from known ones and reduce significant labeling cost. It is also a challenging task since supervised training on coarse-grained categories tends to focus on inter-class distance (distance between coarse-grained classes) but ignore intra-class distance (distance between fine-grained sub-classes) which is essential for separating fine-grained categories. Considering most current methods cannot transfer knowledge from coarse-grained level to fine-grained level, we propose a hierarchical weighted self-contrastive network by building a novel weighted self-contrastive module and combining it with supervised learning in a hierarchical manner. Extensive experiments on public datasets show both effectiveness and efficiency of our model over compared methods. Code and data are available at https://github.com/Lackel/Hierarchical_Weighted_SCL.

1 Introduction

Discovering novel categories based on some known categories has attracted much attention in both Natural Language Processing (Zhang et al., 2021; Zhao et al., 2021) and Computer Vision (Zhong et al., 2021; Han et al., 2019). Previous works assume that novel categories are of the same granularity (or of the same class hierarchy level) as known categories. However, in real-world scenarios, novel categories can be more fine-grained sub-categories of known ones (e.g., sports and tennis). A typical ap-




Figure 1: An example of proposed FCDC task (fine-grained label names need to be assigned by experts).

plication of this scenario is when data analysts want to perform more fine-grained analysis on data with only coarse-grained annotations, where re-labeling fine-grained categories can be time consuming and labour intensive. For example, in the intent detection field, discovering more fine-grained user intents can help to provide better services to customers, but labeling fine-grained intent categories is often much more difficult than labeling coarse-grained ones, since fine-grained annotation often requires higher expertise. To meet this requirement, we investigate a new scenario named Fine-grained Category Discovery under Coarse-grained supervision (FCDC). As shown in Figure 1, FCDC needs models to discover fine-grained categories (e.g., tennis and music) based only on coarse-grained (e.g., sports and arts) labeled data which are easier and cheaper to obtain.

In addition to being in line with above practical needs, FCDC is also a challenging task. Firstly, performing FCDC requires models to increase intra-class distance to ensure fine-grained separability with only coarse-grained supervision. However, coarse-grained classification only focuses on inter-class distance and does not care about intra-class distance (Bukchin et al., 2021), so samples with the same coarse-grained labels will be close to each other and hard to be separated in fine-grained feature space. Secondly, since fine-grained differentiation depends on correct coarse-grained classification, FCDC also requires models to con-

trol inter-class distance to ensure coarse-grained separability. Although increasing intra-class distance can contribute to fine sub-classes separability, it will also decrease inter-class distance, which can result in overlapping between different coarse-grained classes and therefore lead to misclassification. **So how to control and coordinate inter-class and intra-class distance to ensure both coarse-grained and fine-grained separability is the core challenge of FCDC.**

To address above challenges and transfer knowledge from coarse-grained level to fine-grained level, we propose a hierarchical weighted self-contrastive network. By performing different experiments on each layer of BERT, [Jawahar et al. \(2019\)](#) find bottom layers of BERT capture more surface features and top layers capture more high-level semantic features, which means BERT can extract features of different granularities from shallow to deep ([Xu et al., 2021](#)). Inspired by this phenomenon, the core motivation of our model is to learn coarse-grained knowledge by shallow layers of BERT and learn more fine-grained knowledge by the rest of deep layers hierarchically. This motivation is not only consistent with the feature extraction process of BERT, but also corresponding with the shallow-to-deep learning process of humans. Specifically, we use given coarse-grained labels to train shallow layers of BERT to learn some surface knowledge, then we propose a weighted self-contrastive module to train deep layers of BERT to learn more fine-grained knowledge based on the learned surface knowledge.

To ensure both coarse-grained and fine-grained separability, we further propose a weighted self-contrastive module to better coordinate inter-class and intra-class distance in the fine-grained feature space. Specifically, given a query sample, we firstly propose a weighting strategy by weighting different negative samples to control both inter-class and intra-class distance. Then we propose a self-contrastive strategy to generate positive samples to coordinate inter-class and intra-class distance to avoid the overlapping between different coarse-grained classes. We further verify effectiveness and efficiency of our model both theoretically (Section 3.2.4) and experimentally (Section 4.5).

The main contributions of our work can be summarized as threefold:

- We propose to investigate a practical scenario called Fine-grained Category Discovery un-

der Coarse-grained supervision (FCDC), we further propose a hierarchical model to learn fine-grained knowledge from shallow to deep to facilitate the FCDC task.

- To better coordinate inter-class and intra-class distance, we propose a novel weighted self-contrastive module to ensure both coarse-grained and fine-grained separability.
- Extensive experiments on public datasets show that our model significantly advances best compared methods with a large margin and gets double training efficiency than state-of-the-art contrastive learning methods.

2 Related work

2.1 Contrastive learning

Contrastive Learning (CL) aims at grouping similar samples closer and separating dissimilar samples far from each other in a self-supervised way ([Jaiswal et al., 2021](#)), which has gained popularity in both Natural Language Processing (NLP) ([Wu et al., 2020](#)) and Computer Vision (CV) ([Chen et al., 2020](#)). A critical point for CL is to build high-quality positive and negative samples. The simplest way to construct negative samples is to use other in-batch data as negatives ([Chen et al., 2017](#)). Further, [He et al. \(2020\)](#) built a dynamic queue with momentum-updated encoder to keep consistency of representations of negatives. However, these methods consider all negatives equally important, which may lose discriminative information of different negatives. As for positive samples, in CV, one common way is taking two different transformations of the same image as the query and positive sample ([Dosovitskiy et al., 2014](#)). And in NLP, augmentation techniques such as word deletion ([Meng et al., 2021](#)), adversarial attack ([Yan et al., 2021](#)) and dropout ([Gao et al., 2021](#)) were proposed to generate positives. Although there are some recent works ([Bae et al., 2021](#)) using outputs from different levels of a network as positives, we have totally different motivations: they aim at providing more high-quality positives for representation learning while we aim at better adjusting intra-class and inter-class distance.

2.2 Novel Category Discovery

With data volume increases, novel categories especially novel fine-grained categories may be introduced into datasets ([Mekala et al., 2021](#)). To dis-

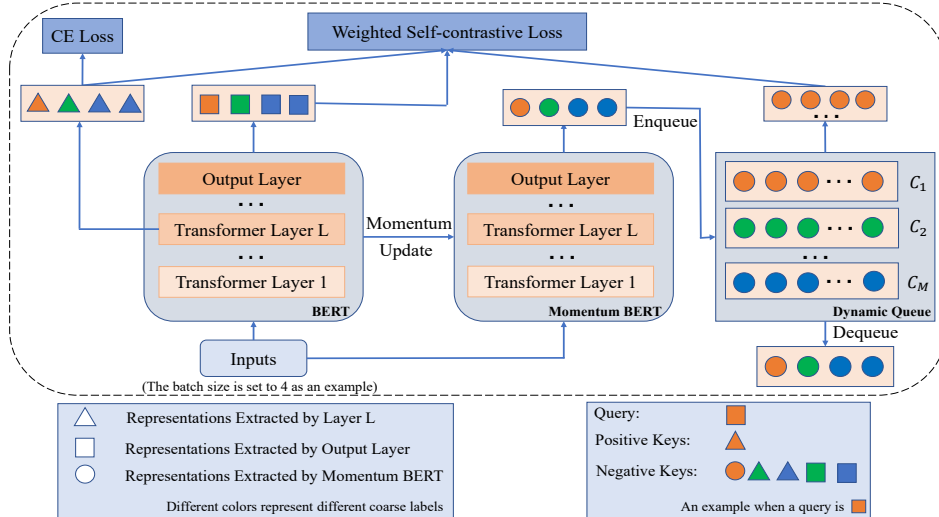


Figure 2: The overall architecture of our model. CE means Cross Entropy.

cover novel categories without human annotation, most previous work adopted clustering methods and transfer learning methods to generate pseudo labels for unlabeled data to train their models (Zhan et al., 2020). For example, Zhang et al. (2021) proposed an alignment strategy to perform Deep-Cluster (Caron et al., 2018) to discover novel categories. Ge et al. (2020) proposed a mutual mean teaching network to refine noisy pseudo labels to perform unsupervised person re-identification. Recently, Two similar tasks as FCDC are proposed. Bukchin et al. (2021) proposed to perform fine-grained image classification under coarse-grained supervision with angular contrastive learning, and they performed this task in a few-shot learning way which needs extra fine-grained labels for each categories. Mekala et al. (2021) proposed to perform fine-grained text classification with coarse-grained annotations, but they need extra fine-grained label hierarchy and corresponding label names to assist in the task. These two tasks both rely on extra fine-grained knowledge from human annotations, which is usually unavailable when novel categories appear in real-world applications. Comparatively, our FCDC is a category discovery task which does not require fine-grained knowledge and is more adapted to real world scenarios.

2.3 Problem Formulation

Denote by $\mathcal{Y}_{coarse} = \{C_1, C_2, \dots, C_M\}$ a set of coarse-grained classes. The training set of FCDC is a set of texts $\mathcal{D}_{train} = \{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_N\}$ with their coarse-grained labels $\{c_1, c_2, \dots, c_N\}$, where $c_i \in \mathcal{Y}_{coarse}$. Different from previous tasks (Bukchin

et al., 2021; Mekala et al., 2021) where the fine-grained label set $\mathcal{Y}_{fine} = \{\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_K\}$ is already known, FCDC assumes that we do not have any prior knowledge about fine-grained labels. So FCDC requires models to perform clustering methods (e.g., K-Means) to discover fine-grained clusters \mathcal{Y}_{fine} with \mathcal{D}_{train} . Since performing clustering will assign each input with a specific cluster assignment, FCDC can also classify inputs into proper fine-grained categories $\{f_1, f_2, \dots, f_N\}$. Although the number of fine-grained clusters K can be estimated with various methods from the clustering area, we assume it is known in FCDC following previous similar works (Zhang et al., 2021; Bukchin et al., 2021) to make a fair comparison.

3 Proposed Approach

As shown in Figure 2, our model mainly contains three components: BERT, Dynamic Queue and Momentum BERT. BERT is used to extract both coarse-grained and fine-grained features. Dynamic Queue can store more negative samples grouping by their coarse-grained labels following Bukchin et al. (2021). Momentum BERT is used to update representations of samples in Dynamic Queue. Inspired by the "shallow to deep" learning process of humankind and the ability of pre-trained models to extract features from coarse-grained to fine-grained (Jawahar et al., 2019; Xu et al., 2021), a core motivation of our model is to learn fine-grained knowledge in a progressive way. Specifically, our model can learn coarse-grained knowledge at shallow layers under coarse-grained supervision and

learn more fine-grained knowledge at deep layers with the proposed weighted self-contrastive learning.

3.1 Supervised Learning

We firstly perform supervised learning on Transformer layer L of BERT to learn coarse-grained knowledge. Given the i -th document \mathcal{D}_i with its coarse-grained label c_i , we use all token embeddings from the L -th layer of BERT as its shallow features. Then we apply a mean-pooling layer to get its shallow feature representation h_i^L :

$$h_i^L = \text{mean-pooling}(\text{BERT}_L(\mathcal{D}_i)) \quad (1)$$

where $h_i^L \in \mathbb{R}^h$ is the hidden state of feature representations, h is the dimension of hidden representations. Then we perform supervised learning with cross entropy loss on coarse-grained labels to get supervised loss \mathcal{L}_{sup}^L at layer L :

$$z_i^L = \sigma(W_a h_i^L + b_a) \quad (2)$$

$$\mathcal{L}_{sup}^L = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp((z_i^L)^{c_i})}{\sum_{j=1}^M \exp((z_i^L)^j)} \quad (3)$$

where $z_i^L \in \mathbb{R}^M$ is the output logits, M is the number of coarse-grained classes. σ is the Tanh activation function, $W_a \in \mathbb{R}^{h \times M}$ and $b_a \in \mathbb{R}^M$ are learnable weights and bias terms, respectively. $(z_i^L)^j$ is the j -th element of output logits z_i^L .

3.2 Weighted Self-contrastive Learning

As shown in Figure 3, denote the coarse-grained inter-class and intra-class distance by d_{coarse} and d_{fine} , respectively. Supervised learning on coarse-grained labels can ensure $d_{coarse} \gg 0$ but will also make $d_{fine} \approx 0$, which can bring difficulties for fine-grained categorization. So how to increase d_{fine} to ensure separability of fine-grained subclasses is a severe challenge. Meanwhile, increasing d_{fine} without restraint will result in overlapping between different coarse-grained classes and therefore lead to misclassification. So how to constrain d_{fine} to ensure the proper classification on coarse-grained classes is the other challenge. In summary, our total goal can be described as:

$$0 \ll d_{fine} < d_{boundary} \ll d_{coarse} \quad (4)$$

where $d_{boundary}$ is a threshold to ensure that samples fall into proper coarse-grained classes.

To achieve above objectives, we propose a weighted self-contrastive module by introducing a




Figure 3: The effectiveness of our self-contrastive module, which can ensure both intra-class and inter-class distance.

novel generation strategy for positive samples and a weighting strategy for negative samples.

3.2.1 Negative Key Generation

Given the i -th document \mathcal{D}_i , we use all token embeddings from the output layer of BERT as its deep features. Then we apply a mean-pooling layer to get its deep feature representation $h_i^o \in \mathbb{R}^h$:

$$h_i^o = \text{mean-pooling}(\text{BERT}_o(\mathcal{D}_i)) \quad (5)$$

In-batch negative keys Given h_i^o with its coarse-grained label c_i as a query q , we treat shallow and deep features of other in-batch samples as its in-batch negative keys, where $k_-^{in}(i) = \{h_j^L, h_j^o\}_{j=1 \dots N, j \neq i}$. In this way, we can increase distance between different samples so that satisfying $d_{fine} \gg 0$ and $d_{coarse} \gg 0$. To satisfy $d_{coarse} \gg d_{fine}$, we propose a weighting strategy by giving more weights to samples with different coarse-grained labels as the query q to further increase their distance. So k_-^{in} can be divided into two groups according to the coarse-grained labels:

$$k_-^{diff}(i) = \{k \in k_-^{in}(i) : c_k \neq c_i\} \quad (6)$$

$$k_-^{same}(i) = \{k \in k_-^{in}(i) : c_k = c_i\} \quad (7)$$

Momentum negative keys To provide more negative keys, we build a momentum BERT and a set of dynamic queues $\{\mathcal{Q}_i\}_{i=1}^M$ to store previous samples grouped by their coarse-grained labels following Bukchin et al. (2021), where M is the number of coarse-grained classes. Specifically, given h_i^o with its coarse-grained label c_i as a query, we treat samples from the queue \mathcal{Q}_{c_i} as its momentum negative keys:

$$k_-^m(i) = \{k \in \mathcal{Q}_{c_i}\} \quad (8)$$

Feature representations of samples in dynamic queues are extracted by momentum BERT, and parameters of momentum BERT are updated in a momentum way (He et al., 2020). At the end of each iteration, the dynamic queues will be updated by adding novel samples and removing the

earliest samples. Since samples in $k_-^m(i)$ have the same coarse-grained labels as the query, they are much harder to be separated and beneficial to better representation learning.

The overall negative keys for the query h_i^o is :

$$k_-(i) = \{k_-^{diff}(i), k_-^{same}(i), k_-^m(i)\} \quad (9)$$

3.2.2 Positive Key Generation

By weighting different negative samples, we can satisfy the condition $0 \ll d_{fine} \ll d_{coarse}$. But increasing d_{fine} without restraint will violate the condition $d_{fine} < d_{boundary}$ and make some samples fall into incorrect coarse-grained classes. To solve this problem, we propose a self-contrastive strategy by treating shallow features of a query as its positive key. Specifically, given the deep feature representation h_i^o for document \mathcal{D}_i as a query, we treat h_i^L as its positive key:

$$k_+(i) = h_i^L \quad (10)$$

As shown in Figure 3, after supervised learning on coarse-grained labels at layer L , h_i^L can be very close to the class center of c_i , so pulling h_i^o close to h_i^L will also pull h_i^o close to the class center of c_i . In this way, we can increase d_{fine} with restraint and satisfy the condition $d_{fine} < d_{boundary}$ without computing the specific value of $d_{boundary}$. Another advantage of our self-contrastive strategy is that we can get double training efficiency than traditional data augmentation-based methods (Wu et al., 2020; Gao et al., 2021) since we only need to perform forward and backward propagation only once to get and update both queries and positive keys (Section 5.2).

3.2.3 Weighted Self-contrastive Loss

Given the query h_i^o with its positive key $k_+(i)$ and negative keys $k_-(i)$, the overall loss of our weighted self-contrastive module is:

$$\mathcal{L}_{cont} = \sum_{i=1}^N -\log \frac{e^{sim(h_i^o, h_i^L)/\tau}}{\sum_{l \in k_-(i)} \alpha_l \sum_{k \in l} e^{sim(h_i^o, h_k)/\tau}} \quad (11)$$

where $\alpha_l \in \{\alpha_{same}, \alpha_{diff}, \alpha_m\}$ are weighting factors for different negative keys, $sim(h_i, h_j)$ is cosine similarity $\frac{h_i^T h_j}{\|h_i\| \cdot \|h_j\|}$ and τ is a temperature hyperparameter.

By weighting different negative keys and selecting shallow features as the positive key, our model can satisfy the goal in Equation 4 and provide

conditions for subsequent fine-grained categorization.

3.2.4 Theoretical Analysis

The effectiveness of our weighted self-contrastive learning compared with traditional contrastive learning from the gradient perspective is analyzed below.

Self-contrastive Strategy Compared with traditional contrastive loss which only aims at grouping queries and their transformations closer, our self-contrastive strategy aims at pulling queries and their shallow features closer:

$$sim(h_i^o, h_i^L) := sim(h_i^o, h_i^L) + \frac{1}{\tau} \quad (12)$$

Since τ is positive, the positive similarity will increase and h_i^o will be grouped closer to h_i^L . After supervised learning on coarse-grained labels at layer L , h_i^L can be close to the class center of c_i , so pulling h_i^o closer to h_i^L will also pull h_i^o closer to the class center of c_i . So our Self-Contrastive strategy can guarantee queries fall into correct coarse-grained categories and get double training efficiency since we only need to perform forward and backward propagation only once to get and update both queries and positive keys.

Weighting Strategy Since negatives with the same coarse-grained labels as queries have larger gradients (Wang and Liu, 2021), traditional contrastive loss will push these negatives farther from queries than those with different coarse-grained labels as queries, which leads to $d_{coarse} < d_{fine}$ and is opposite of what we expect to solve the FCDC task. To mitigate this limitation, we propose a weighting strategy to give more weights to samples with different coarse-grained labels as the query to further increase their distance:

$$sim(h_i^o, h_j^o) := sim(h_i^o, h_j^o) - \alpha_l \cdot \mathcal{P}_{i,j} \quad (13)$$

$$\mathcal{P}_{i,j} = \frac{1}{\tau} \cdot \frac{e^{sim(h_i^o, h_j^o)/\tau}}{\sum_{l \in k_-(i)} \alpha_l \sum_{k \in l} e^{sim(h_i^o, h_k)/\tau}} \quad (14)$$

By increasing the weighting factor α_l for negatives with different coarse-grained labels as queries, the corresponding similarity will decrease faster. So negatives with different coarse-grained labels from queries will be pushed farther than those with the same coarse-grained labels as queries, which can guarantee $d_{fine} < d_{coarse}$ for the FCDC task.

Table 1: Statistics of datasets. # indicates the number of samples. $|\mathcal{C}|$, $|\mathcal{F}|$ means the number of coarse-grained and fine-grained classes, respectively.

Dataset	$ \mathcal{C} $	$ \mathcal{F} $	# Train	# Test
CLINC	10	150	18,000	1,000
WOS	7	33	8,362	2,420
HWU64	18	64	8,954	1,031

3.3 Overall Loss

To further guarantee samples to be classified into proper coarse-grained categories, we also add supervised learning on coarse-grained labels at the output layer. So the overall loss for our hierarchical weighted self-contrastive network is:

$$\mathcal{L} = \mathcal{L}_{sup}^o + \gamma_1 \mathcal{L}_{sup}^L + \gamma_2 \mathcal{L}_{cont} \quad (15)$$

where \mathcal{L}_{sup}^o is the cross entropy loss at the output layer. γ_1 and γ_2 are weighting factors.

After representation learning, we simply perform the non-parametric clustering method K-Means to discover fine-grained categories based on features extracted by the output layer of BERT.

4 Experiments

4.1 Datasets

To evaluate effectiveness of our model, we conduct experiments on three public datasets. Statistics of three datasets can be found in Table 1.

CLINC is an intent classification dataset released by Larson et al. (2019).

Web of Science (WOS) is a paper classification dataset released by Kowsari et al. (2017).

HWU64 is a personal assistant query classification dataset released by Liu et al. (2021).

4.2 Implementation Details

We use the pre-trained BERT model (bert-base-uncased) implemented by Pytorch (Wolf et al., 2020) as our backbone and adopt most of its suggested hyper-parameters. We use the cuml library (Raschka et al., 2020) to perform K-Means on GPU to speed up calculations. We use the AdamW optimizer with 0.01 weight decay. Gradient clipping is also used with the norm 1.0. For hyper-parameters, temperature τ is set to 0.1, layer L is set to 8, and the weighting factors α_l for $\{k_-^{diff}(i), k_-^{same}(i), k_-^m(i)\}$ are set to $\{1.4, 1.0, 1.0\}$, weighting factors $\{\gamma_1, \gamma_2\}$ are set to $\{0.001, 0.008\}$. The training batch size is set to 128, and

the testing batch size is set to 64. The momentum queue size for each coarse-grained category is set to 128, and the momentum factor for Momentum BERT is set to 0.9. The hidden dimension h is 768, the learning rate is set to $5e^{-5}$, the dropout rate is set to 0.1. The training epoch is set to 20.

For a fair comparison, we use the same BERT model as ours to extract features for all compared methods and adopt hyper-parameters in their original papers.

4.3 Compared Methods

Baselines We perform FCDC with BERT in unsupervised and coarse-supervised way as baselines.

Self-supervised Methods DeepCluster (Caron et al., 2018) and DeepAligned (Zhang et al., 2021) are self-supervised methods using self-training techniques and achieve state-of-the-art results in many category discovery tasks. Ancor (Bukchin et al., 2021) is a self-supervised method designed for few-shot fine-grained classification with coarse-grained labels. SimCSE (Gao et al., 2021) and Delete One Word (Wu et al., 2020) are contrastive learning methods in NLP with different data augmentation techniques.

Self-supervised + Cross Entropy To investigate the influence of coarse-grained supervision on compared models, we further add cross entropy loss on coarse-grained labels \mathcal{L}_{sup}^o to their loss function.

4.4 Evaluation Metrics

We use fine-grained labels as ground truth to evaluate model performance on testing sets. Since no fine-grained knowledge is available for the FCDC task, we need to perform clustering to discover fine-grained categories. Clustering performance can reflect the quality of discovered fine-grained clusters (more compact clusters usually mean better discovered categories). And classification performance can reflect the semantic overlap between discovered clusters and real categories.

To evaluate clustering performance, we use two broadly used external evaluation metrics. Adjusted Rand Index (ARI) is used to evaluate the degree of agreement between cluster assignments and ground truth. And Normalized Mutual Information (NMI) is used to evaluate the mutual information between cluster assignments and ground truth. To evaluate classification performance, we use metric Accuracy (ACC), which is obtained from Hungarian algorithm (Kuhn, 1955) to align cluster assignments and ground truth.

Table 2: Model comparison results (%) on fine-grained categories. Average results over 5 runs are reported. '+ CE' means adding coarse-grained supervision with cross entropy loss. We have also performed statistical significance test and all the p-values are less than 10^{-6} , which means our improvement is significant.

Methods	CLINC			WOS			HWU64		
	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI
Unsupervised	33.38	16.42	63.46	32.32	18.21	47.12	33.66	16.88	56.78
Coarse Supervised	45.91	32.27	75.04	39.42	33.67	61.60	42.41	33.74	71.57
DeepCluster	26.40	12.51	61.26	29.17	18.05	43.34	29.74	13.98	53.27
DeepAligned	29.16	14.15	62.78	28.47	15.94	43.52	29.14	12.89	52.99
SimCSE	40.22	23.57	69.02	25.87	13.03	38.53	24.48	8.42	46.94
Anchor	45.60	33.11	75.23	41.20	37.00	65.42	37.34	34.75	74.99
Delete One Word	47.11	31.28	73.39	24.50	11.68	35.47	21.30	6.52	44.13
DeepCluster + CE	30.28	13.56	62.38	38.76	35.21	60.30	41.73	27.81	66.81
DeepAligned + CE	42.09	28.09	72.78	39.42	33.67	61.60	42.19	28.15	66.50
Anchor + CE	44.44	31.50	74.67	39.34	26.14	54.35	32.90	30.71	74.73
Delete One Word + CE	47.87	33.79	76.25	41.53	33.78	61.01	35.13	31.84	74.88
SimCSE + CE	52.53	37.03	77.39	41.28	34.47	61.62	34.04	31.81	74.86
Ours	74.67	64.77	89.14	63.64	51.55	72.46	58.45	48.20	78.66

4.5 Main Results

Model performance on fine-grained categories are reported in Table 2. From the results we can draw following conclusions. Our model significantly outperforms other compared methods across all datasets. We contribute reasons of better performance of our model to following two points. Firstly, we propose a hierarchical architecture to learn fine-grained knowledge from shallow to deep, which is consistent with the feature extraction process of BERT and the shallow-to-deep learning process of humans. Secondly, we propose a weighted self-contrastive module to coordinate inter-class and intra-class distance so that we can better learn both coarse-grained and fine-grained knowledge. Self-training methods perform badly on all datasets and evaluation metrics since they rely on abundant labeled data to generate high-quality pseudo labels for unlabeled data. Contrastive learning methods perform better than self-training methods since they do not need fine-grained labels to initialize their models. However, their performance is still much worse than ours since they cannot fully utilize given coarse-grained labels to control inter-class and intra-class distance between samples. We also find that model performance of most compared methods increases with the addition of coarse-grained supervision, which means coarse-grained supervision can boost model performance on fine-grained tasks.

Our model performance on coarse-grained

Table 3: Classification accuracy (%) on coarse-grained categories on test sets.

Model	CLINC	WOS
Coarse Supervised	98.58	91.86
Ours	98.71	91.45

Table 4: Results (%) of different model variants. '-' means that we remove the component from our model.

Model	ACC	ARI	NMI
ALL	74.67	64.77	89.14
- Momentum	73.38	64.47	88.91
- \mathcal{L}_{sup}^L	72.93	63.72	88.06
- Weighting	71.75	62.99	88.47
- Self-Contrast	53.21	40.05	75.36

categories are reported in Table 3. From the table we can see that our model gets similar classification accuracy as the upper-bound coarse-supervised BERT, which means that our model can control not only intra-class distance to ensure fine-grained separability, but also inter-class distance to ensure coarse-grained variability.

5 Discussion

5.1 Ablation Study

To investigate contributions of different components to our model, we compare the performance of our model with its variants on the CLINC dataset. As shown in Table 4, removing different compo-

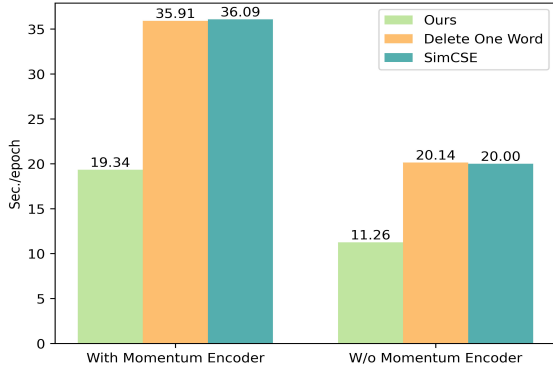


Figure 4: Training efficiency comparison.

nents will affect model performance more or less, which indicates the effectiveness of different components of our model. Removing Momentum Encoder has minimal impact, since our model is insensitive to the number of negative samples. Removing weighting strategy or cross entropy loss at shallow layers also hurt model performance since they can help to learn coarse-grained knowledge and lay foundation for learning fine-grained knowledge. Above all, removing self-contrastive strategy results in a significant decrease, since it is responsible for controlling intra-class and inter-class distance.

5.2 Training Efficiency

In this section, we compare the training efficiency of our model with contrastive methods SimCSE and Delete One Word on the CLINC dataset. We test all methods using the BERT base model trained on the same hardware platform (an AMD EPYC CPU 7702 and a RTX 3090 GPU) with the batch size 128. Average results over 100 epochs are shown in Figure 4. Compared with SimCSE and Delete One Word, our model gets double training efficiency both when adding or removing Momentum Encoder, which benefits from our self-contrastive strategy. Specifically, our model utilizes shallow features of queries as positive keys, which only needs to perform forward and backward propagation once to get and update both queries and positive keys.

5.3 Visualization

We further visualize the learned embeddings of our model and SimCSE using t-SNE on the CLINC dataset in Figure 5. Our model can separate different coarse-grained categories with a larger margin than SimCSE (Top in Figure 5), which benefits




Figure 5: TSNE visualization of learned embeddings. **Top**: coarse-grained categories. **Bottom**: fine-grained categories of one arbitrary coarse-grained category. **Left**: Ours. **Right**: SimCSE + CE.

from our strategy of combining supervised learning and contrastive learning in a hierarchical way. Furthermore, our model can also separate different fine-grained categories with a larger margin (Bottom in Figure 5), which benefits from the weighted self-contrastive module. In summary, our model can better control both inter-class and intra-class distance between samples to facilitate the FCDC task than traditional contrastive learning methods.

5.4 Choices of L and weighting factors

Effect of Shallow Layer L The influence of the choice of shallow layer L on model performance is shown in Figure 6. Our model achieves the best performance when $L=8$. In this way, our model can learn coarse-grained knowledge at shallow layers ($L<8$) and provide enough model capacity to learn fine-grained knowledge at deeper layers ($L>8$), which is consistent with the feature extraction process of BERT (Jawahar et al., 2019).

Effect of Weighting Factors We investigate the influence of the ratio $\beta = \alpha_{diff}/\alpha_{same}$ in Figure 7 (We fixed $\alpha_m = 1$ since it has little influence). As analyzed in Section 3.2.4, by giving more weights to negatives with different coarse-grained labels as queries ($\beta > 1$), our weighting strategy can keep these negatives further away from queries and guarantee $d_{fine} < d_{coarse}$. On the contrary, when $\beta < 1$, negatives with the same coarse-grained labels as queries will be further away from the queries, which can hurt our model performance.

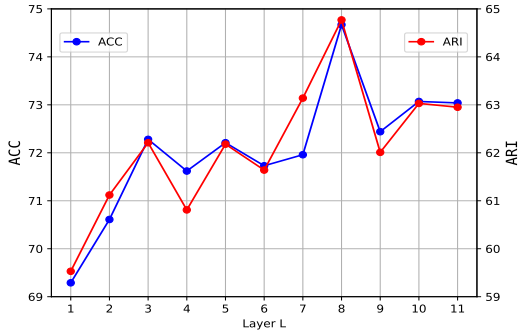


Figure 6: Effect of shallow layer L .

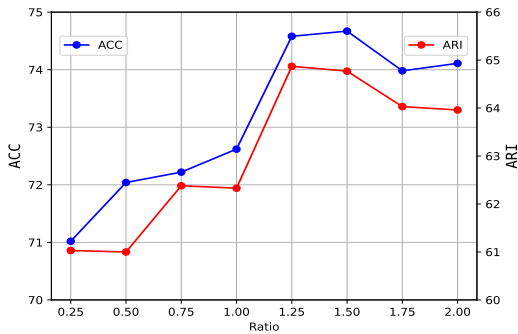


Figure 7: Effect of ratio $\beta = \alpha_{diff} / \alpha_{same}$.

6 Conclusion

In this paper, we investigate a novel task named Fine-grained Category Discovery under Coarse-grained supervision (FCDC), which can reduce significant labeling cost and adapt models to novel categories of different granularity from known ones. We further propose a hierarchical weighted self-contrastive model to approach the FCDC task by better controlling intra-class and inter-class distance. By performing supervised and contrastive learning on shallow and deep layers of pre-trained models, our model can learn fine-grained knowledge from shallow to deep with only coarse-grained supervision. Extensive experiments on public datasets show that our approach is more effective and efficient than compared methods.

Limitations

The limitations of our method lies in two aspects. Firstly, following previous works, we need to know the number of fine-grained clusters K as prior knowledge, which is usually difficult to get in real-world scenarios. Secondly, our method cannot predict semantic meanings (e.g., label names) of discovered fine-grained categories, which is also an unexplored question in the field of novel category

discovery.

Acknowledgements

This work was supported by National Key Research and Development Program of China (2020AAA0108800), National Natural Science Foundation of China (62137002, 61721002, 61937001, 61877048, 62177038, 62277042). Innovation Research Team of Ministry of Education (IRT_17R86), Project of China Knowledge Centre for Engineering Science and Technology. MoE-CMCC “Artificial Intelligence” Project (MCM20190701), Project of Chinese academy of engineering “The Online and Offline Mixed Educational Service System for ‘The Belt and Road’ Training in MOOC China”. “LENOVO-XJTU” Intelligent Industry Joint Laboratory Project.

References

- Sangmin Bae, Sungnyun Kim, Jongwoo Ko, Gihun Lee, Seungjong Noh, and Se-Young Yun. 2021. Self-contrastive learning. *arXiv preprint arXiv:2106.15499*.
- Guy Bukchin, Eli Schwartz, Kate Saenko, Ori Shahar, Rogerio Feris, Raja Giryes, and Leonid Karlinsky. 2021. Fine-grained angular contrastive learning with coarse labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8730–8740.
- Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. 2018. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 132–149.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.
- Ting Chen, Yizhou Sun, Yue Shi, and Liangjie Hong. 2017. On sampling strategies for neural network-based collaborative filtering. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 767–776.
- Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. 2014. Discriminative unsupervised feature learning with convolutional neural networks. *Advances in neural information processing systems*, 27:766–774.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.

- Yixiao Ge, Dapeng Chen, and Hongsheng Li. 2020. Mutual mean-teaching: Pseudo label refinery for unsupervised domain adaptation on person re-identification. *arXiv preprint arXiv:2001.01526*.
- Kai Han, Andrea Vedaldi, and Andrew Zisserman. 2019. Learning to discover novel visual categories via deep transfer clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8401–8409.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738.
- Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. 2021. A survey on contrastive self-supervised learning. *Technologies*, 9:2.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. What does bert learn about the structure of language? In *ACL 2019-57th Annual Meeting of the Association for Computational Linguistics*.
- Kamran Kowsari, Donald E Brown, Mojtaba Heidarysafa, Kiana Jafari Meimandi, , Matthew S Gerber, and Laura E Barnes. 2017. Hdltext: Hierarchical deep learning for text classification. In *Machine Learning and Applications (ICMLA), 2017 16th IEEE International Conference on*. IEEE.
- Harold W Kuhn. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2:83–97.
- Stefan Larson, Anish Mahendran, Joseph J Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K Kummerfeld, Kevin Leach, Michael A Laurenzano, Lingjia Tang, et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. *arXiv preprint arXiv:1909.02027*.
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2021. Benchmarking natural language understanding services for building conversational agents. In *Increasing Naturalness and Flexibility in Spoken Dialogue Interaction*, pages 165–183. Springer.
- Dheeraj Mekala, Varun Gangal, and Jingbo Shang. 2021. Coarse2fine: Fine-grained text classification on coarsely-grained annotated data. *arXiv preprint arXiv:2109.10856*.
- Yu Meng, Chenyan Xiong, Payal Bajaj, Saurabh Tiwary, Paul Bennett, Jiawei Han, and Xia Song. 2021. Coco-lm: Correcting and contrasting text sequences for language model pretraining. *arXiv preprint arXiv:2102.08473*.
- Sebastian Raschka, Joshua Patterson, and Corey Nolet. 2020. Machine learning in python: Main developments and technology trends in data science, machine learning, and artificial intelligence. *arXiv preprint arXiv:2002.04803*.
- Feng Wang and Huaping Liu. 2021. Understanding the behaviour of contrastive loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2495–2504.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zhuofeng Wu, Sinong Wang, Jiatao Gu, Madian Khabsa, Fei Sun, and Hao Ma. 2020. Clear: Contrastive learning for sentence representation. *arXiv preprint arXiv:2012.15466*.
- Keyang Xu, Tongzheng Ren, Shikun Zhang, Yihao Feng, and Caiming Xiong. 2021. Unsupervised out-of-domain detection via pre-trained transformers. *arXiv preprint arXiv:2106.00948*.
- Yuanmeng Yan, Rumei Li, Sirui Wang, Fuzheng Zhang, Wei Wu, and Weiran Xu. 2021. Consert: A contrastive framework for self-supervised sentence representation transfer. *arXiv preprint arXiv:2105.11741*.
- Xiaohang Zhan, Jiahao Xie, Ziwei Liu, Yew-Soon Ong, and Chen Change Loy. 2020. Online deep clustering for unsupervised representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6688–6697.
- Hanlei Zhang, Hua Xu, Ting-En Lin, and Rui Lyu. 2021. Discovering new intents with deep aligned clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Jun Zhao, Tao Gui, Qi Zhang, and Yaqian Zhou. 2021. A relation-oriented clustering method for open relation extraction. *arXiv preprint arXiv:2109.07205*.
- Zhun Zhong, Enrico Fini, Subhankar Roy, Zhiming Luo, Elisa Ricci, and Nicu Sebe. 2021. Neighborhood contrastive learning for novel class discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10867–10875.