

GANTEE: Generative Adversarial Network for Taxonomy Entering Evaluation

Zhouhong Gu^{1,4}, Sihang Jiang¹, Jingping Liu^{2*}, Yanghua Xiao^{1,3*}, Hongwei Feng¹, Zhixu Li¹, Jiaqing Liang⁴, Jian Zhong⁵

¹Shanghai Key Laboratory of Data Science, School of Computer Science, Fudan University, China

²School of Information Science and Engineering, East China University of Science and Technology

³Fudan-Aishu Cognitive Intelligence Joint Research Center

⁴School of Data Science, Fudan University

⁵HUAWEI CBG Edu AI Lab

{zhgu20, shawyh, hwfeng, zhixuli, liangjiaqing}@fudan.edu.cn

tedsihangjiang@gmail.com, jingpingliu@ecust.edu.cn, zhongjian6@huawei.com

Abstract

Taxonomy is formulated as directed acyclic concepts graphs or trees that support many downstream tasks. Many new coming concepts need to be added to an existing taxonomy. The traditional taxonomy expansion task aims only at finding the best position for new coming concepts in the existing taxonomy. However, they have two drawbacks when being applied to the real-scenarios. The previous methods suffer from low-efficiency since they waste much time when most of the new coming concepts are indeed noisy concepts. They also suffer from low-effectiveness since they collect training samples only from the existing taxonomy, which limits the ability of the model to mine more hypernym-hyponym relationships among real concepts. This paper proposes a pluggable framework called Generative Adversarial Network for Taxonomy Entering Evaluation (GANTEE) to alleviate these drawbacks. A generative adversarial network is designed in this framework by discriminative models to alleviate the first drawback and the generative model to alleviate the second drawback. Two discriminators are used in GANTEE to provide long-term and short-term rewards, respectively. Moreover, to further improve the efficiency, pre-trained language models are used to retrieve the representation of the concepts quickly. The experiments on three real-world large-scale datasets with two different languages show that GANTEE improves the performance of the existing taxonomy expansion methods in both effectiveness and efficiency.

Introduction

Taxonomy is formulated as directed acyclic graphs or trees, which consist of hyponym-hypernym relations between concepts. One concept is a hypernym of another concept if the meaning of the former covers the latter (Sang 2007). A well-constructed taxonomy assist various downstream tasks, including web content tagging (Liu et al. 2020, 2019; Peng et al. 2019), personalized recommendation (Karamanolakis, Ma, and Dong 2020; Huang et al. 2019), query understanding (Yang, Zhang, and Han 2020) and so on.

Manually maintaining a taxonomy is labor-intensive and time-consuming. For example, millions of new concepts are

expected to be added to the taxonomy of one of the largest shopping platforms in the world each month (Jiang et al. 2019). So the task of automated taxonomy expansion is proposed (Jurgens and Pilehvar 2016), which aims to automatically assign an existing concept (anchor concept) as a hypernym concept to the newly input concept (query concept). These models successfully save much time from labor annotation in taxonomy maintenance. These methods predict the hypernym-hyponym relationships mainly based on the representation similarity, so they sample the training data based on the existing taxonomy to learn good representations for each concept.

However, we argue that these methods (Yu et al. 2020; Wang et al. 2022; Cheng et al. 2022; Shen et al. 2020; Mao et al. 2020; Wang et al. 2021; Zhang et al. 2021) have two drawbacks when being applied to real applications. **First**, these methods suffer from low effectiveness in real applications due to limited ability to represent the semantics of concepts. Only sampling from existing taxonomy provides limited data to train representations for zero-shot concepts, where thousands times queries of the number of concepts in existing taxonomy are all unseen to the models (Cheng et al. 2022). **Second**, these methods suffer from low efficiency in real applications because most of the query concepts are noisy. For instance, in a leading ordering take-outs online platform, only thousands of concepts out of billions of queries are finally added to their taxonomy (Cheng et al. 2022). Previous methods ignore the noisy concepts and waste much time finding the anchor concept for noisy query concepts.

In this paper, we propose a pluggable framework called Generative Adversarial Network for Taxonomy Entering Evaluation (GANTEE) to boost both the effectiveness and efficiency of the taxonomy expansion methods in real-scenarios as is shown in Figure 1. To improve the effectiveness, GANTEE proposes a generative model to generate more high-quality training data, which are required to train the high-quality representation for taxonomy expansion. To improve efficiency, GANTEE introduces a new task called taxonomy entering evaluation, which removes many noisy query concepts before finding the suitable anchor concept

*Corresponding Authors

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

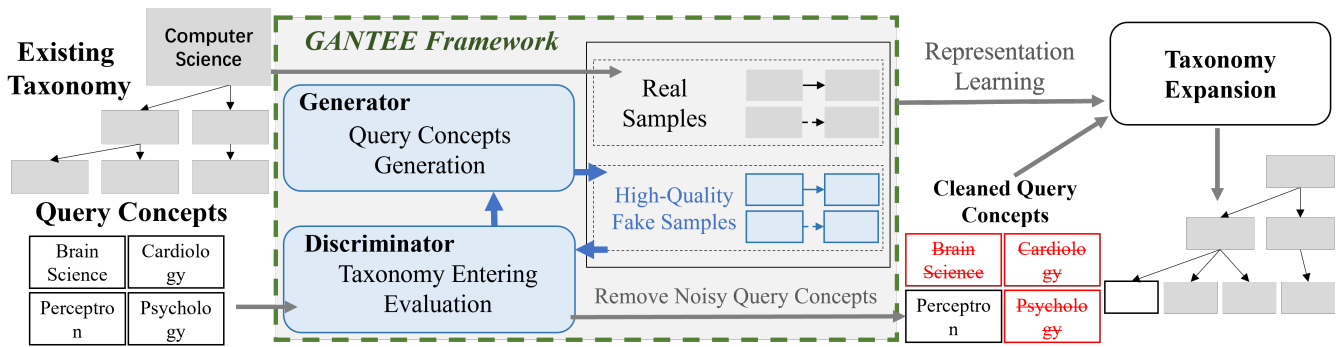


Figure 1: An example of the GANTEE framework. This framework is implemented as a plugin before taxonomy expansion. The demand for generating high-quality data and taxonomy entering evaluation is formulated as a generative adversarial network in GANTEE.

for them. Intuitively, the generative model and the taxonomy entering evaluation model are two adversarial models which can improve each other. The generative model produces fake samples similar to real ones to confuse the discriminative model. The taxonomy entering evaluation model learns to generate high-quality samples. Therefore, these two tasks can be formulated as a generative adversarial task (Goodfellow et al. 2014) to improve their performance further.

Specifically, to improve efficiency, the taxonomy entering evaluation model is expected to be more lightweight than the taxonomy expansion model. We introduce two mechanisms to ensure the efficiency of the taxonomy entering evaluation model. First, instead of determining where to extend the query concept, GANTEE introduces two more manageable tasks “isA Concept Access” and “isA Query Access” to determine whether the query concept is a concept and whether the query concept should be added to the taxonomy. Second, instead of training a new representation for emerging query concepts, we use pre-trained language models (PLMs) to directly obtain a representation of each concept based on the textual features of the concept. Although PLMs have limited ability in representing the fine-grained semantics based on conceptual text (Lauscher et al. 2019), the experimental results verify that PLMs are suitable to generate representation in “isA Concept Access” and “isA Query Access” task.

In the experiments, we benchmark the taxonomy entering evaluation task on three real-world taxonomies in two languages, English and Chinese. We reached the new SOTA in the experiments by improving all the metrics in only a third of the time of other methods in the prediction stage. Finally, extensive experiments have been conducted to study how the parameters affect the result of GANTEE, and provide an effective tuning scheme for the following researchers who will use our method in the future.

Contribution. To summarize, our major contributions include: (1) We propose a more manageable task called taxonomy entering evaluation which judges whether a query concept should be added to an existing taxonomy efficiently; (2) We propose a novel, effective method called GANTEE to generate confusing negative and fidelity positive data boosts the performance and efficiency for both traditional taxonomy expansion tasks and taxonomy entering evaluation tasks. (3) Extensive experiments have been conducted

to verify the superiority of GANTEE on three datasets in two languages. We also propose an effective tuning method based on the experiment results.

Problem Formalization

In this section, we first define the preliminaries used in this paper, then formally define the taxonomy entering evaluation task and traditional taxonomy expansion problems.

Preliminary and Problem Definition

Preliminary 1. (Taxonomy and Anchor Concept.) A taxonomy $\mathcal{T} = (\mathcal{A}, \mathcal{E})$ is a directed acyclic graph. Each node $a \in \mathcal{A}$ denotes an anchor concept and each directed edge $\langle a_p, a_c \rangle \in \mathcal{E}$ represents the hyponymy relation between the parent (hypernym) concept a_p and the child (hyponym) concept a_c .

Preliminary 2. (Query Concept.) Query concept is the concept prepared to be added to the existing taxonomy. The set of query concepts is represented as $\mathcal{Q} = \{q_1, q_2, \dots, q_m\}$.

Problem 1. Taxonomy entering Evaluation: In taxonomy entering evaluation task, the input is (1) an existing taxonomy $\mathcal{T}^0 = (\mathcal{A}^0, \mathcal{E}^0)$, (2) a set of query concepts \mathcal{Q} . The output of entering evaluation is a set of clean query concepts \mathcal{Q}^* , which is formalized into the following formula:

$$\mathcal{Q}^* = \mathcal{Q}, \forall q \in \mathcal{Q}^* \text{ s.t. } P(q|\mathcal{T}^0; \Theta_2) > \gamma. \quad (1)$$

Θ_2 denotes the parameters, and γ is a given threshold.

Problem 2. Taxonomy Expansion: In taxonomy expansion task, the input is (1) an existing taxonomy $\mathcal{T}^0 = (\mathcal{A}^0, \mathcal{E}^0)$, (2) a set of clean query concepts \mathcal{Q}^* . The output of the taxonomy expansion is a best-expanded taxonomy \mathcal{T}^* with all the clean query concepts added in:

$$\begin{aligned} \mathcal{T}^* &= \arg \max_{\mathcal{T}} \mathbf{P}(\mathcal{T} | \mathcal{T}^0; \Theta) \\ &= \arg \max_{\mathcal{T}} \sum_{i=1}^{|\mathcal{A} \cup \mathcal{Q}^*|} \log \mathbf{P}(q_i | p(q_i); \Theta), \end{aligned} \quad (2)$$

where Θ denotes model parameters, and $p(q_i)$ denotes the parent of q_i in the taxonomy \mathcal{T} .

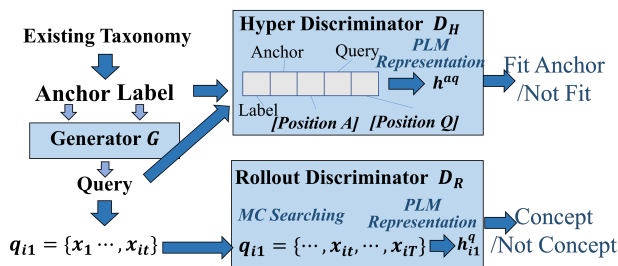


Figure 2: The framework of GANTEE. The generator of the GANTEE is on the left, and duel discriminator is on the right.

Framework of GANTEE

GANTEE consists of a generator G and two discriminators, which are hyper discriminator D_H and rollout discriminator D_R . G takes an anchor concept and a label that generates positive or negative samples as inputs and generates the corresponding query concept. D_H determines whether the given anchor concept is suitable for the generated query concept based on the given label, and D_R determines whether the generated text is a concept.

Specifically, when a new query concept is given in the predicting stage, GANTEE first determines if it is a concept by D_R . Then GANTEE determines whether the given concept has a hypernym-hyponym relation to the root concept by D_H . Finally, GANTEE output the confidence score of if the given query concept is suitable for this taxonomy.

The Generative Model for Query Concept Generating

We use Transformer (Vaswani et al. 2017) as the generative model. Transformer maps the input sequence x_1, \dots, x_n into a sequence of hidden states h_1, \dots, h_n . Specifically, a Transformer block consists of stacked self-attention layers with residual connections. Each self-attention layer receives n embeddings $\{h_i^{(l)}\}_{i=1}^n$ corresponding to unique input tokens, and outputs n hidden states $\{h_i^{(l+1)}\}_{i=1}^n$. The i -th token is mapped via linear transformations to a key k_i , query q_i , and value v_i . The i -th output of the self-attention layer is given by weighting the values v_j by the normalized dot product between the query q_i and other keys k_j :

$$h_i = \sum_{j=1}^n \text{softmax}(\langle q_i, k_j \rangle)_{j=1}^n \cdot v_j. \quad (3)$$

This allows the layer to assign ‘‘credit’’ by implicitly forming state-return associations via similarity of the query and key vectors (maximizing the dot product). It is worth noticing that most of the Transformer variants, such as BART (Lewis et al. 2019) or XLNet (Yang et al. 2019), and the RNN variants, such as the Gated Recurrent Unit (GRU) (Cho et al. 2014) or Long-Short Term Memory network (LSTM) (Hochreiter and Schmidhuber 1997), can be used as a generator in GANTEE.

We use GPT-2 (Radford et al. 2019) in this paper to give the best experimental results. GPT-2 modifies the transformer architecture with a causal self-attention mask to enable auto-regressive generation. It is one of the most famous

auto-regressive models, which best suits the fake query generation task.

As an auto-regressive generative model, GPT-2 could generate sequences based on all input sequences. So we formulate the input and the output of the GPT-2 in the following format:

$$GPT2(l_i \oplus a_i) = l_i \oplus a_i \oplus q_i. \quad (4)$$

$GPT2$ denotes the GPT-2 model, l_i denotes the given label, a_i denotes the text of the given anchor concept, q_i denotes the text of the given query concept, and \oplus denotes the concatenation of the tokens.

The Rollout Discriminative Model for Short Term Reward in Generating

The rollout discriminative model is used to judge whether the generated text is a concept before the generating process finish. Deep discriminative models such as deep neural network (DNN), convolutional neural network (CNN), and recurrent convolutional neural network (RCNN) have shown their effectiveness in complicated sequence classification tasks (Guo et al. 2018; Yu et al. 2017; Nie, Narodytska, and Patel 2018). In this paper, we choose LSTM (Hochreiter and Schmidhuber 1997) as our rollout discriminator. Since previous study (Guo et al. 2018) find that severe gradient vanishing occurs when the discriminator is too strict, LSTM is a lightweight model which is easy to do the training and will not provide too strict discriminative ability. LSTM will take in the input embedding x_1, \dots, x_t and use purpose-built memory cells to calculate the hidden state for each step h_1, \dots, h_t .

$$\{h_1, \dots, h_t\} = \text{lstm}(x_1, \dots, x_t), \quad (5)$$

where lstm denotes the model of LSTM. We take the final hidden state h_t and inject it into a linear layer to calculate the reward of the generated text:

$$r_r = W_r \cdot h_t + b_r, \quad (6)$$

where W_r and b_r are the parameters in a linear layer, the r_r is the confidence of whether the generated text is a concept.

The Hyper Discriminative Model for Long Term Reward in Generating

The hyper discriminative model is used to judge whether the generated concept has a hypernym relation to the given anchor concept after the generating process finish based on the given label. Such discriminative models offer a long-term reward for the generating process to guide the direction of the generation better. The existing models for taxonomy expansion can be used as the hyper discriminative model, which can detect whether there is a hypernym relation between two concepts.

However, these models rely heavily on the vector representation of the concepts, many studies such as STEAM (Yu et al. 2020) or HiExpan (Shen et al. 2018) use external information to get better representations for the concepts, and that would be time-consuming for the generating process since each time the generative model produce a new concept, it

will take time to get its detailed external information and get its vector representation.

Instead, we use Bert (Devlin et al. 2018) to retrieve the representation of concepts since it shows its ability in Dependency Judgement task (Shi and Demberg 2019), which is also good at understanding the hypernym (hyponym) relation between two concepts. We concatenate the label representation and the position embedding proposed by the TaxoExpan (Shen et al. 2020) to the Bert representation and put them into a linear layer to get the final predicted result.

Specifically, we get the representation of the concept by Bert in the following format:

$$\begin{aligned} x_a &= Bert(x_1^{c_0}, \dots, x_t^{c_0})[1], \\ x_q &= Bert(x_1, \dots, x_t)[1], \\ l &= EMB(positive) \text{ or } EMB(negative). \end{aligned} \quad (7)$$

Bert denotes the BERT model, *EMB* denotes the embedding model, and *positive* or *negative* denotes the input to the embedding model.

Then, we concatenate the label representation and position embedding to the representation of the anchor concept and the query concept and put them into a linear layer:

$$r_h = W_h \cdot (l + x_a + p_a + x_q + p_q) + b_h \quad (8)$$

Where W_h and b_h are the parameters in a linear layer, p_a and p_q are the position embedding. The r_h is the probability of the hypernym relation between the generated query concept and the anchor concept.

Training for GANTEE

GANTEE uses a Generative Adversarial Network (GAN) model to combine taxonomy entering evaluation and high-quality training samples generating. It needs a generative model to produce conceptual text and use discriminative models to offer supervised signals to the generative model. However, naive GAN is used for continuous data, not discrete tokens. The previous researchers (Bachman and Precup 2015; Bahdanau et al. 2016; Yu et al. 2017) propose considering the text generation procedure as a sequential decision-making process. The objective function of such generative models is formulated to generate a sequence from the start state s_0 to maximize its expected end reward via policy gradient algorithms (Sutton et al. 1999):

$$\begin{aligned} J(\theta) &= \mathbb{E}[R_T | s_0 = l \oplus a; \theta] \\ &= \sum_{y_1 \in Y} G(y_1 | s_0 = l \oplus a) \cdot Q_D^G(s_0 = l \oplus a, y_1), \end{aligned} \quad (9)$$

where $J(\theta)$ is the object function, R_T is the reward for a complete concept that has hypernym relation to the given anchor concept a and label l . Note that the reward is from both the hyper discriminator D_H and rollout discriminator D_R . G is the generative model, and $Q_D^G(s, a)$ is the action-value function of a sequence, which is calculated in the following format:

$$Q_D^G(a = y_t, s = Y_{1:t-1}, s_0) = D(s_0, Y_{1:t-1}) \quad (10)$$

Since the reward getting from an unfinished sequence often have a large bias over the finished one, the Monte Carlo

search is used to sample the unknown last $T - t$ tokens (Yu et al. 2017):

$$\{Y_{1:T}^1, \dots, Y_{1:T}^N\} = MC^G(Y_{1:t}; N) \quad (11)$$

We use rollout discriminator D_R to calculate the expected reward based on Monte Carlo search during the generative process for short-term reward and use hyper discriminator D_H to calculate the reward when the generating is done for long-term reward. The action-value function is transformed into the following format:

$$Q_D^G(a = y_t, s = Y_{1:t-1}, s_0) = \begin{cases} \frac{1}{N} \sum_{n=1}^N D_R(s_0, Y_{1:t-1}, Y_{1:T}^n \in MC^G(Y_{1:t}; N)) & \text{for } t < T \\ D_H(Y_{1:t}, s_0) & \text{for } t = T \end{cases} \quad (12)$$

Once we have a set of more realistic generated sequences, we will re-train the discriminator model as follows:

$$\min_D -\mathbb{E}_{Y \sim p_{data}} [\log D(Y)] - \mathbb{E}_{Y \sim G} [\log(1 - D(Y))] \quad (13)$$

Specifically, we pre-train the generator and the discriminator before the adversarial process begins. After the pre-training stage, the generator and discriminator have trained alternatively. As the generator is updated via some training epochs, the discriminators need to be re-trained periodically to keep a good pace with the generator.

When training the discriminator, positive samples are sampled from the existing taxonomy \mathcal{T}^0 based on the hypernym relation \mathcal{R} . In contrast, negative samples are the fake samples generated from our generator. To ensure the learning process is doing well, we use both generated samples and pure noisy samples as the negative query concepts to inject into the training process. The number of negative query concepts we use is the same as the size of positive query concepts.

Experiments

There is an open question that whether the generated positive samples should be treated as positive training samples or negative training samples. Since the purpose of GAN is to train a generator that output a distribution similar to the real data, if all the data generated by the generator are regarded as negative samples, it will beyond the discriminative ability of the classifier in downstream applications. In this paper, we treat them as negative training samples in the first epoch of adversarial training and as positive training samples in the subsequent epochs of adversarial training.

Experimental Setup

Dataset We study the performance of GANTEE on three large-scale real-world taxonomies:

- **Microsoft Academic Graph on Field-of-Study (MAG-FoS):** This taxonomy (Sinha et al. 2015) consists of public Field-of-Study Taxonomy (FoS). It contains over 0.6 million scientific concepts and more than 0.7 million taxonomic relations in English.

Method	MAG-CS (10k-level)				MAG-FoS (100k-level)			CN-Probase (1000k-level)		
	MR	MRR@10	hit@5	time	MR	MRR@100	time	MR	MRR@1k	time
Close-Position	2.3k	0.137	0.023	-	23.2k	0.128	-	157.5k	0.159	-
Close-Neighbor	1.8k	0.153	0.043	-	15.6k	0.146	-	95.5k	0.188	-
BERT+MLP	2.5k	0.001	0.002	116.0 s	25.6k	0.023	5.4k s	159.1k	0.143	42.5k s
TaxoExpan	1.3k	0.319	0.157	471.4 s	3.7k	0.305	25.2k s	-	-	(200k+) s
TaxoExpan + GANTEE	1.2k	0.360	0.193	152.8 s	2.5k	0.382	6.1k s	6.6k	0.319	67.4k s

Table 1: Overall Performance on Taxonomy Expansion.

- **Microsoft Academic Graph on Computer-Science (MAG-CS):** Following the work of TaxoExpan (Shen et al. 2020), we construct MAG-CS based on the sub-graph of MAG-FoS related to the ‘‘Computer Science’’ domain.
- **CN-Probase:** This is an open chinese general concept taxonomy CN-Probase (Chen et al. 2019). It contains over 300 thousand concepts and more than 30 million ‘‘isA’’ relations between concepts and entities in Chinese.

Notice that there are many outlier concepts in the taxonomy, so we remove all concepts that do not relate to the original taxonomy. To ensure the validity and test data won’t be trained, we randomly mask 20% of leaf concepts (along with their relations) for validation and testing. Furthermore, we add noisy query concepts as 60% size of leaf concepts into the validation and test set to conduct the experiments as in the real-scenarios.

Evaluation Metric

- **Mean Rank (MR):** measures the average rank position of a query concept’s true parent among all candidates.
- **Mean Reciprocal Rank@k (MRR@k):** calculates the reciprocal rank of a query concept’s true parent.
- **hit@k:** is the number of query concepts whose parent is ranked in the top k positions, divided by the total number of queries.
- **time:** denotes the seconds used in the predicting stage of each model on each dataset.

Baseline Methods We compared our proposed GANTEE model with the following baseline approaches:

- **Closest-Position:** A rule-based method that chooses the anchor concept which have the most similar embedding with query concept as the query concept’s parent.
- **Closest-Neighbor:** Another rule-based method that scores each position based on its distance to the query concept plus the average distance between its children nodes and the query.
- **BERT+MLP:** This method utilizes BERT (Devlin et al. 2018) to perform hypernym detection, whose input is the text of the terms to get the representation from BERT.
- **Random:** This method is used only in the taxonomy entering evaluation task, which is random to determine if the query concept should be added to the existing taxonomy or not.

Method	MAG-CS		CN-Probase-Sampled	
	Acc	F1	Acc	F1
Closest-Position	21.5	13.9	15.4	11.5
Closest-Neighbor	28.8	17.8	19.6	13.7
BERT+MLP	57.6	46.1	52.3	43.8
Random	74.8	-	75.9	-
TaxoExpan	76.4	52.7	66.8	32.1
GANTEE	89.2	79.3	77.5	65.4

Table 2: Performance on Taxonomy Entering Evaluation.

Taxonomy + GANTEE	MAG-CS TE			MAG-CS TEE	
	MR	MRR	hit@5	Acc	F1
w/o hd	1235	0.357	0.186	78.5	62.9
w/o rd	1217	0.351	0.182	81.5	71.3
w/o D_ϕ	1250.1	0.348	0.180	71.4	51.9
-	1188.1	0.360	0.193	89.2	79.3

Table 3: Ablation study on GANTEE.

- **TaxoExpan** (Shen et al. 2020): One state-of-the-art taxonomy expansion framework which leverages position enhanced graph neural network to capture local information and InfoNCE (Oord, Li, and Vinyals 2018) loss for robust training.

Notably, except for the rule-based method Closest-Position and Closest-Neighbor, other baselines are learning-based methods and designed for one-to-one matching. There are other recently proposed taxonomy expansion methods, e.g., HiExpan (Shen et al. 2018) and STEAM (Yu et al. 2020). We do not include them as baselines because they leverage external sources, e.g., text corpus, to extract complicated features. In contrast, TaxoExpan and other baselines only take initial feature vectors as input.

Parameter Setting For learning-based methods, we use SGD optimizer with initial learning rate 0.0001 and ReduceLROnPlateau¹ scheduler with ten patience epochs. During model training, the batch size and negative sample size are set to 16 and 256 in the overall performance experiments, respectively. We set the epochs to be ten and use two layers of Graph Attention Layer with 8 and 1 attention head and 100 dimensions of the position dim.

Experimental Results

Overall Performance on Taxonomy Expansion As our proposal is a plugin framework for improving both the ef-

¹https://pytorch.org/docs/stable/optim.html/#torch.optim.lr_scheduler.ReduceLROnPlateau




Figure 3: The changing of the results on different metrics with different negative size sampling during the taxonomy expansion model training stage and different number of generated high-quality samples by GANTEE.




Figure 4: The changing of the losses and results on different metrics with the adversarial training goes on. In these figures, red lines denote “loss”, blue line denotes the metrics, label “N” means taxonomy expansion without GANTEE, label “P” means the stage of pretraining generative model, and label “A” means the adversarial training stage.

fectiveness and efficiency of the taxonomy expansion, we are curious how GANTEE will affect the performance of the previous methods on the taxonomy expansion task. Thus, we use GANTEE on the most prevailing State-Of-The-Art (SOTA) method TaxoExpan (Shen et al. 2020) on the taxonomy expansion task. The results are presented in Table 1.

Notably, since the time expense of the rule-based method only depends on the computational complexity of the algorithms or matrix operations, it is not recorded, for it can be neglected most of the time in our experiments. And the time expense of TaxoExpan on CN-Probbase is too large, and we do not finish this experiment in days, so the expected time expense is recorded in the table.

Effectiveness Performance: From the results, we can see that all of the rule-based methods did a poor job in most metrics, which demonstrates that the semantics of concepts in real-scenarios are complicated, so simple rule-based methods have difficulty performing well in the taxonomy expansion task. However, BERT+MLP has also done a poor job on most metrics, demonstrating that simply using the representation of PLM is not suitable for taxonomy expansion tasks. With our GANTEE, TaxoExpan can reach a new SOTA in all three datasets. Additionally, the larger scale of the taxonomy, the more improvement is made through GANTEE.

Efficiency Performance: BERT+MLP method is proven itself to be an efficient method, and we think it is because this method also uses PLMs to fetch the representation of each concept, which is a trade to the final effectiveness of the taxonomy expansion task. Moreover, when TaxoExpan is implemented after our proposed method, a new SOTA is reached within one-third of the predicting time on three datasets on average.

Performance on Taxonomy Entering Evaluation Table 2 presents the results of all compared methods on the

three datasets. The rule-based methods and BERT+MLP did a poor job on taxonomy entering evaluation task, which is much worse than the Random. The TaxoExpan has limited ability to distinguish the noisy query concept, for it outperforms four of the other baseline methods. GANTEE outperforms all the other methods by a large margin on all datasets in taxonomy entering evaluation task. Furthermore, since the PLM representation of the English word is better than the Chinese word, it is easier for the discriminator to distinguish the noisy query concepts in English than in Chinese, which makes all the metrics in English higher than in Chinese.

Ablation Study The detail of ablation studies is presented in this subsection. We remove three of the essential designs in the GANTEE. In “w/o hd”, we remove the hyper discriminator during the adversarial generating process. In “w/o rd”, we remove the rollout discriminator during the adversarial generating process. In “w/o D_ϕ ”, we remove all the adversarial processes and only use the generator to generate samples. The results are shown in Table 3

We compare the removal of the hyper discriminative model D_H and rollout discriminative model D_R together. We can see that without the rollout discriminative model, all metrics have a more significant decline than without the hyper discriminative model in the taxonomy expansion task. However, when the hyper discriminative model is removed, there will be a more significant decline in the metrics in the taxonomy entering evaluation task than removing the rollout discriminative. We conclude that the rollout discriminative model generates better conceptual texts, and the hyper discriminative is used to strictly denoise all the noisy query concepts.

And without all of the discriminative models, which means only using the generative model to learn from the query conceptual text, a significant decline is witnessed in

all metrics. We conclude there are two main reasons. The first reason for this significant decline is because there has no model trained for taxonomy entering evaluation specifically, so the metrics of taxonomy entering evaluation task is decreasing. The second reason for this significant decline is because there have no supervised signals provided during the generating process, and the GPT-2 is hard to capture the semantic of the hypernym-hyponym relationship in the existing taxonomy, so the result of the taxonomy expansion task is decreasing.

Analysis of GANTEE The Table 3 denotes how different the negative size during training and a different number of the high-quality generated training sampled by GANTEE affect the final result of the taxonomy expansion task. The upper right corner of the figure is significantly darker than the other areas. So we conclude that when generating more high-quality training samples by GANTEE, the taxonomy expansion task will improve performance by using a bigger negative size of the taxonomy expansion in its training stage.

The Table 4 denotes how the final result of the taxonomy expansion task will be affected when the GANTEE stop training in a different training stage. In these figures, we can intuitively conclude that with the stage of the training going deeper, the performance of the final result will not be better because the line in MRR, hit@1, and hit@3 are not consistently decreasing, and the line in MR is not continuously increasing. We conclude that when the loss no longer drops dramatically, it is better to stop the training of the GANTEE to receive a better performance in the taxonomy expansion task. We find out that the goal of the generative model in GANTEE is to confuse the discriminator, and when the distribution of the generated text is close to the distribution of the real conceptual text, the discriminator will not be able to identify the samples effectively. And the gradients are not effectively obtained, causing the discriminator become less effective and the low quality of the generator.

Case Analysis We present three good cases and two typical bad cases of GANTEE in Tabel 4

The result shows that GANTEE has the ability in generating a new concept in a domain-specific field. For example, “Deep Q-Learning”, “Computer Vision Algorithm” and “Deep Graphics Processing Unit” are all correct concepts and have true hypernym-hyponym relationships to the given anchor concepts based on the given label.

However, we find GANTEE has limited ability in generating positive query concepts for too fine-grained or too coarse-grained anchor concepts. We conclude that GANTEE has difficulty in capturing the semantics of too fine-grained anchor concepts which are seldom seen in the training process of GANTEE “Deep Q-Learning Algorithm” to “Deep Algorithm”. And GANTEE tends to assign any concept to the too coarse-grained anchor concepts like “Deep Graphics Processing Unit” to “Deep Learning”.

Related Work

The previous methods for taxonomy expansion task can be divided into two categories. The studies in the first category

Samples Generated by GANTEE
Given Label: Positive Given Anchor Concept: Deep Learning Algorithm Generated Text: Deep Q-Learning ✓
Given Label: Positive Given Anchor Concept: Computer Vision Generated Text: Computer Vision Algorithm ✓
Given Label: Negative Given Anchor Concept: Deep Learning Algorithm Generated Text: Deep Graphics Processing Unit ✓
Given Label: Positive Given Anchor Concept: Deep Q-Learning Generated Text: Deep Algorithm ✗
Given Label: Positive Given Anchor Concept: Deep Learning Generated Text: Deep Graphics Processing Unit ✗

Table 4: Case study of the samples generated by GANTEE. tend to construct a better model to predict the hypernym-hyponym relationships of the concepts. Position embedding (Shen et al. 2020), concept sorting mechanism (Song et al. 2021) and multilevel position matching (Wang et al. 2022) are proposed to better understand the semantics of the query concepts and anchor concepts. The second category focus on learning good representations for all concepts. Open source (Yu et al. 2020), user-clicked-logs (Cheng et al. 2022) and parent-child information (Wang et al. 2021, 2022) are all used by the previous methods to learn a better representation for concepts. Besides, there exist many studies use PLMs during the representation learning stage (Cheng et al. 2022; Wang et al. 2022; Takeoka, Akimoto, and Oyamada 2021), however, these methods only use the PLMs as a external information provider to learn a better representation.

Conclusion

This paper studies the drawbacks of the existing taxonomy expansion methods when applied to the real-scenarios. We propose a pluggable framework called GANTEE to alleviate these drawbacks. Extensive experiments have been conducted to demonstrate that GANTEE can boost the efficiency and effectiveness of the taxonomy expansion task. Interesting future work includes lower the representation learning expense or learning representation for multi-model queries in the task of taxonomy expansion.

Acknowledgments

This work was supported by Shanghai Municipal Special Fund for Promoting High-quality Development of Industries (No. 2021-GZL-RGZN-01018), Shanghai Sailing Program (No. 23YF1409400), National Key Research and Development Project (No.2020AAA0109302), Scientific and technological innovation 2030 - major project of new generation artificial intelligence (No.2020AAA0109300), National Natural Science Foundation of China (No.62072323), Shanghai Science and Technology Innovation Action Plan (No. 22511104700). National Natural Science Foundation of China (No. 62102095).

References

- Bachman, P.; and Precup, D. 2015. Data generation as sequential decision making. *Advances in Neural Information Processing Systems*, 28.
- Bahdanau, D.; Brakel, P.; Xu, K.; Goyal, A.; Lowe, R.; Pineau, J.; Courville, A.; and Bengio, Y. 2016. An actor-critic algorithm for sequence prediction. *arXiv preprint arXiv:1607.07086*.
- Chen, J.; Wang, A.; Chen, J.; Xiao, Y.; Chu, Z.; Liu, J.; Liang, J.; and Wang, W. 2019. CN-Probase: a data-driven approach for large-scale Chinese taxonomy construction. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, 1706–1709. IEEE.
- Cheng, S.; Gu, Z.; Liu, B.; Xie, R.; Wu, W.; and Xiao, Y. 2022. Learning What You Need from What You Did: Product Taxonomy Expansion with User Behaviors Supervision. *arXiv preprint arXiv:2203.14921*.
- Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- 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*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Guo, J.; Lu, S.; Cai, H.; Zhang, W.; Yu, Y.; and Wang, J. 2018. Long text generation via adversarial training with leaked information. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Huang, J.; Ren, Z.; Zhao, W. X.; He, G.; Wen, J.-R.; and Dong, D. 2019. Taxonomy-aware multi-hop reasoning networks for sequential recommendation. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 573–581.
- Jiang, S.; Liang, J.; Xiao, Y.; Tang, H.; Huang, H.; and Tan, J. 2019. Towards the completion of a domain-specific knowledge base with emerging query terms. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, 1430–1441. IEEE.
- Jurgens, D.; and Pilehvar, M. T. 2016. Semeval-2016 task 14: Semantic taxonomy enrichment. In *Proceedings of the 10th international workshop on semantic evaluation (SemEval-2016)*, 1092–1102.
- Karamanolakis, G.; Ma, J.; and Dong, X. L. 2020. Txtract: Taxonomy-aware knowledge extraction for thousands of product categories. *arXiv preprint arXiv:2004.13852*.
- Lauscher, A.; Vulić, I.; Ponti, E. M.; Korhonen, A.; and Glavaš, G. 2019. Specializing unsupervised pretraining models for word-level semantic similarity. *arXiv preprint arXiv:1909.02339*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Liu, B.; Guo, W.; Niu, D.; Luo, J.; Wang, C.; Wen, Z.; and Xu, Y. 2020. GIANT: scalable creation of a web-scale ontology. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, 393–409.
- Liu, B.; Guo, W.; Niu, D.; Wang, C.; Xu, S.; Lin, J.; Lai, K.; and Xu, Y. 2019. A user-centered concept mining system for query and document understanding at tencent. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1831–1841.
- Mao, Y.; Zhao, T.; Kan, A.; Zhang, C.; Dong, X. L.; Faloutsos, C.; and Han, J. 2020. Octet: Online Catalog Taxonomy Enrichment with Self-Supervision. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2247–2257.
- Nie, W.; Narodytska, N.; and Patel, A. 2018. Relgan: Relational generative adversarial networks for text generation. In *International conference on learning representations*.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Peng, H.; Li, J.; Wang, S.; Wang, L.; Gong, Q.; Yang, R.; Li, B.; Philip, S. Y.; and He, L. 2019. Hierarchical taxonomy-aware and attentional graph capsule RCNNs for large-scale multi-label text classification. *IEEE Transactions on Knowledge and Data Engineering*, 33(6): 2505–2519.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Sang, E. T. K. 2007. Extracting hypernym pairs from the web. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, 165–168.
- Shen, J.; Shen, Z.; Xiong, C.; Wang, C.; Wang, K.; and Han, J. 2020. TaxoExpan: self-supervised taxonomy expansion with position-enhanced graph neural network. In *Proceedings of The Web Conference 2020*, 486–497.

- Shen, J.; Wu, Z.; Lei, D.; Zhang, C.; Ren, X.; Vanni, M. T.; Sadler, B. M.; and Han, J. 2018. Hiexpan: Task-guided taxonomy construction by hierarchical tree expansion. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2180–2189.
- Shi, W.; and Demberg, V. 2019. Next sentence prediction helps implicit discourse relation classification within and across domains. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, 5790–5796.
- Sinha, A.; Shen, Z.; Song, Y.; Ma, H.; Eide, D.; Hsu, B.-J.; and Wang, K. 2015. An overview of microsoft academic service (mas) and applications. In *Proceedings of the 24th international conference on world wide web*, 243–246.
- Song, X.; Shen, J.; Zhang, J.; and Han, J. 2021. Who Should Go First? A Self-Supervised Concept Sorting Model for Improving Taxonomy Expansion. *arXiv preprint arXiv:2104.03682*.
- Sutton, R. S.; McAllester, D.; Singh, S.; and Mansour, Y. 1999. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12.
- Takeoka, K.; Akimoto, K.; and Oyamada, M. 2021. Low-resource taxonomy enrichment with pretrained language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2747–2758.
- 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, S.; Zhao, R.; Chen, X.; Zheng, Y.; and Liu, B. 2021. Enquire One’s Parent and Child Before Decision: Fully Exploit Hierarchical Structure for Self-Supervised Taxonomy Expansion. *arXiv preprint arXiv:2101.11268*.
- Wang, S.; Zhao, R.; Zheng, Y.; and Liu, B. 2022. QEN: Applicable Taxonomy Completion via Evaluating Full Taxonomic Relations. In *Proceedings of the ACM Web Conference 2022*, 1008–1017.
- Yang, C.; Zhang, J.; and Han, J. 2020. Co-embedding network nodes and hierarchical labels with taxonomy based generative adversarial networks. *ICDM*.
- Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R. R.; and Le, Q. V. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32.
- Yu, L.; Zhang, W.; Wang, J.; and Yu, Y. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Yu, Y.; Li, Y.; Shen, J.; Feng, H.; Sun, J.; and Zhang, C. 2020. STEAM: Self-Supervised Taxonomy Expansion with Mini-Paths. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1026–1035.
- Zhang, J.; Song, X.; Zeng, Y.; Chen, J.; Shen, J.; Mao, Y.; and Li, L. 2021. Taxonomy Completion via Triplet Matching Network. *arXiv preprint arXiv:2101.01896*.