

# Find Parent then Label Children: A Two-stage Taxonomy Completion Method with Pre-trained Language Model

Fei Xia<sup>1,2\*</sup>, Yixuan Weng<sup>1\*</sup>, Shizhu He<sup>1,2</sup>, Kang Liu<sup>1,2</sup> and Jun Zhao<sup>1,2</sup>

<sup>1</sup> National Laboratory of Pattern Recognition, Institute of Automation, CAS

<sup>2</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences

xiafei2020@ia.ac.cn, wengsyx@gmail.com

{shizhu.he, kliu, jzhao}@nlpr.ia.ac.cn

## Abstract

Taxonomies, which organize domain concepts into hierarchical structures, are crucial for building knowledge systems and downstream applications. As domain knowledge evolves, taxonomies need to be continuously updated to include new concepts. Previous approaches have mainly focused on adding concepts to the leaf nodes of the existing hierarchical tree, which does not fully utilize the taxonomy’s knowledge and is unable to update the original taxonomy structure (usually involving non-leaf nodes). In this paper, we propose a two-stage method called ATTEMPT for taxonomy completion. Our method inserts new concepts into the correct position by finding a parent node and labeling child nodes. Specifically, by combining local nodes with prompts to generate natural sentences, we take advantage of pre-trained language models for hypernym/hyponymy recognition. Experimental results on two public datasets (including six domains) show that ATTEMPT performs best on both taxonomy completion and extension tasks, surpassing existing methods.

## 1 Introduction

Taxonomies<sup>1</sup> are an important form of domain knowledge that organize concepts into hierarchical structures, representing “hypernym-hyponym” relationships among concepts in the form of trees or directed acyclic graphs (Shen et al., 2020). Taxonomies are essential components of knowledge systems such as ontologies and knowledge graph (Yu et al., 2020), and are widely used in various downstream applications, including search engineering (Yin and Shah, 2010), recommendation systems (Huang et al., 2019; Zhang et al.,

\* These authors contribute this work equally.

<sup>1</sup>In this paper, we mainly focus on the taxonomy represented as tree rather than directed acyclic graph, because trees are the mainstream form at present, such as the online catalog taxonomies of Amazon and Yelp.

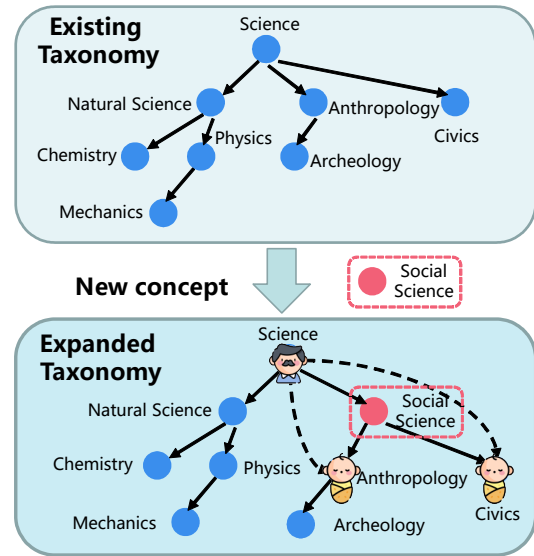


Figure 1: An example of taxonomy completion for a non-leaf node. The new concept “Social Science” needs to be inserted into the correct position in the existing taxonomy.

2014), and information filtering (Demeester et al., 2016).

As domain knowledge continues to evolve, especially with the rapid growth of web content, new concepts are constantly emerging. In order to stay current, original taxonomies must incorporate these new concepts and adapt their hierarchical relationships. For example, as shown in Figure 1, with the advancement of sociology and science, the concept of “Social Science” should be added to the science knowledge system, and the original structure should be adjusted accordingly.

However, existing taxonomies are primarily constructed by human experts (Shen et al., 2018). Manual extraction of domain concepts and detection of hierarchical relationships by domain experts is both time-consuming and labor-intensive, and may result in missing important concepts and relationships.

To extend existing taxonomies automatically, researchers have proposed the tasks of taxonomy

expansion (TE) and taxonomy completion (TC). Both tasks aim to append new nodes (concepts) to a given taxonomy. The main difference is that TE focuses on identifying the parent of a given node (usually a leaf node), while TC aims to identify both the parent and child nodes. As illustrated in Figure 1, TE would aim to identify the parent node of “Social Science”, while TC would also aim to identify the child nodes of “Social Science”.

Recently, researchers have been focusing on using pre-trained language models, such as BERT (Devlin et al., 2018), to improve the performance of taxonomy expansion (Liu et al., 2021; Takeoka et al., 2021). For example, TEMP (Liu et al., 2021) appends new concepts to leaf nodes and generates candidate taxonomy paths, then uses a pre-trained model for ranking and selecting the best path. Musubu (Takeoka et al., 2021) generates candidate “hypernym”-“new concept” pairs using Hearst patterns (Hearst, 1992), and relies on pre-trained knowledge to identify the optimal hypernym node for the new concept. These proposed models have greatly improved the effectiveness of taxonomy updates, thanks to the improved generalization performance of pre-trained language models (Liu et al., 2021).

Although current TE&TC methods have achieved good results, there are several main issues that need to be addressed. Firstly, existing TE methods struggle to extend non-leaf nodes or perform poorly in this task (Zhang et al., 2021). Secondly, while existing TC methods can extend both leaf and non-leaf nodes, they may be less effective in leaf node expansion than specialized TE methods (Liu et al., 2021), potentially due to a lack of sufficient utilization of knowledge. Furthermore, these methods often require large amounts of labeled samples or external resources, which are not always available (Takeoka et al., 2021). Lastly, current TC methods do not typically involve modifying the nodes of the original taxonomy system (all original parent-child relationships are preserved after adding nodes to the taxonomy). However, the insertion of new nodes can modify the relationship of the original nodes. For example, the insertion of “Social Science” in Figure 1 would change the relationship between “Science-Anthropology” from father-son to grandfather-grandson.

To address these issues, we propose **A Two-stage Taxonomy completion Method with Pre-**

**Trained Language model (ATTEMPT)**, which inserts new concepts into the correct position by identifying a parent node and labeling child nodes.

In the first stage of our proposed method, we use the “Taxonomy-path Prompt with Pre-trained model” (PPT) approach to take advantage of the local information of the taxonomy path and convert it into natural language using a prompt method, which helps to better utilize the implicit knowledge of the pre-trained model. Additionally, the pre-trained model’s extensive knowledge reserve allows us to avoid the need for external resources and large amounts of labeled data. In the second stage, we propose the “Multiple Nodes Labeling” (MNL) method, which jointly identifies each child node and better utilizes the interdependence between nodes, resulting in more accurate node type prediction (including father-son, sibling and other relationships). Additionally, MNL allows for modification of the original taxonomy nodes and simultaneous annotation of multiple child nodes.

We conduct detailed experiments on two public datasets (including six domains) to evaluate the effectiveness of our proposed method, ATTEMPT, in leaf and non-leaf node expansion. Specifically, for leaf nodes, our parent-finding method (PPT) outperforms the best baseline by 8.2% in accuracy. For non-leaf nodes, our children-finding method (MNL) improves by 21% and 20.3% respectively in accuracy and average F1 score, compared to a pair-wise classification method. On the overall task, our proposed method (ATTEMPT) outperforms other methods by 2.4% in average F1 score.

In summary, the main contributions of this paper include:

- The proposal of a two-stage taxonomy expansion method, ATTEMPT, that inserts new concepts into the correct position by identifying a parent node and labeling child nodes.
- The introduction of a multiple-nodes labeling method, MNL, for the children-finding stage, which allows for the label of zero to multiple children nodes of a given node simultaneously and modification of the original taxonomy nodes.
- The demonstration of the effectiveness of our approach through experiments on two public datasets (including six domains), with the best performance obtained in both non-leaf and leaf node expansion.

## 2 Related Work

Taxonomy construction aims to build a tree-structured taxonomy with a set of terms from scratch. Existing methods can be roughly divided into two categories. The first is an unsupervised method to construct the taxonomy based on clustering (Alfarone and Davis, 2015; Zhang et al., 2018a; Shang et al., 2020). The terms are grouped into a hierarchy based on hierarchical clustering or topic models (Downey et al., 2015). Each node of this taxonomy is a collection of topic-indicative terms, different from the taxonomy in this paper (each node represented by one individual term). The other approach constructs a taxonomy based on terms, where each node represents a term concept (Cocos et al., 2018; Dash et al., 2020). Hypernymy detection models are often used for this task. For example, pattern-based (Agichtein and Gravano, 2000; Jiang et al., 2017; Roller et al., 2018) or distributional models (Yin and Roth, 2018; Wang et al., 2019; Dash et al., 2020) extract the hypernymy for a given query node and then organize them into a tree structure.

Creating a taxonomy from scratch is labor-intensive. In many scenarios, such as e-commerce, some taxonomies may already be deployed in online systems, which involves the demand of taxonomy extension. QASSIT (Cleuziou and Moreno, 2016) is a semi-supervised vocabulary classification method, mainly based on genetic algorithms. The TAXI (Panchenko et al., 2016) system uses a taxonomy induction method based on lexico-syntactic patterns, substrings, and focused crawling. Later, TaxoGen (Zhang et al., 2018b) uses term embeddings and hierarchical clustering to construct topic taxonomies recursively. TEMP (Liu et al., 2021) is a self-supervised classification extension method that trains models with a new dynamic margin loss margin function.

Taxonomy completion (Zhang et al., 2021) is a recently proposed task that aims to find appropriate hypernyms-hyponyms for new nodes, not just hypernyms. GenTaxo (Zeng et al., 2021) gathers information from complex local structural information and learns to generate full names of concepts from corpora. TMN (Zhang et al., 2021) focuses on channel gating mechanisms and triplet matching networks. CoRel relies on concept learning and relation transferring to build a seed-oriented topic taxonomy.

But the above mentioned methods also have

some issues. The addition of new nodes may also lead to changes in the original taxonomy. The taxonomy completion task only finds the hyponyms of a given node, which cannot modify of the original taxonomy. GenTaxo (Zeng et al., 2021) requires a large amount of training data to learn enough information, and CoRel (Huang et al., 2020) focuses more on topic taxonomy than the taxonomy of individual terms. Other works such as CGExpan (Zhang et al., 2020) use the automatically generated class names and the class-guided entity selection module for entity expansion. However, CGExpan (Zhang et al., 2020) is more on the entity set than the tree taxonomy.

In addition, although the above methods can find both hypernyms and hyponyms of a given query node, they do not make sufficient use of the pre-trained model or do not use the pre-trained model at all (Zhang et al., 2021; Zeng et al., 2021). This may lead them to perform poorly on the hypernym recognition task, inferior to the specialized taxonomy extension methods of the pre-trained model (Liu et al., 2021). And most methods of taxonomy extension cannot perform well on the task of taxonomy completion (Zhang et al., 2021). We are dedicated to finding an approach that works in both tasks.

## 3 Method

Given an existing taxonomy  $T = (V, E)$  and a set of new terms  $V'$ , where  $V$  is a set of terms, and  $E$  is a set of "hyponym-hypernym" relationships between terms, the task of Taxonomy completion is to insert the new terms  $v' \in V'$  into the appropriate position of the existing taxonomy  $T$  one by one and extend them into a more complete taxonomy  $\tilde{T} = (\tilde{V}, \tilde{E})$ .

Figure 2 provided illustrates the overall structure of the ATTEMPT method, which is broken down into two main stages: the parent finding stage and the children finding stage. These two stages work together to identify the relationships between terms in the taxonomy, specifically determining the parent and children of a given term.

### 3.1 Stage one: Parent Finding

The first stage of the process is to identify the parent node of a given node in the taxonomy. For example, finding the parent node "science" for the node "social science" in Figure 2.

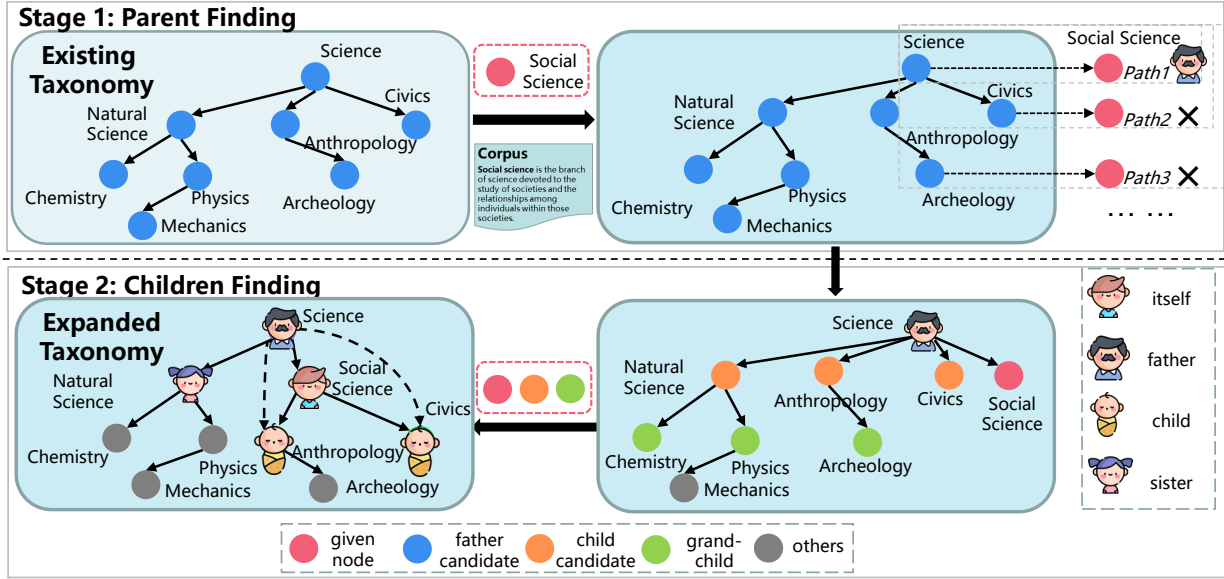


Figure 2: An overview of the proposed method ATTEMPT. ATTEMPT consists of two stages. Given an existing taxonomy and the new concept term (“Social Science”), the first stage is to find the correct parent (“Science”) and the second stage is to find all possible children (“Anthropology” and “Civics”).

### 3.1.1 TEMP

The TEMP method (Liu et al., 2021) is the first approach to use pre-trained contextual encoders as the core component for taxonomy extension. The pre-trained contextual embeddings are useful for capturing relationships between terms because they have been trained on a large corpus. TEMP predicts the location of new concepts by ranking the generated taxonomy paths. A taxonomy path of a new term (ND) in the tree-structured taxonomy is the unique path from that term to the root of the taxonomy. The taxonomy path is represented as  $P = [ROOT, N_1, N_2, \dots, N_D]$ , where  $D$  is the depth of the ND and  $ROOT$  is the root of the taxonomy. In the taxonomy,  $N_{i-1}$  is the parent of  $N_i$ . TEMP generates taxonomy paths for each term, then adds the new term to be expanded to the end of each path to form new paths. Finally, the new paths are ranked and the highest-scoring path is chosen as the correct parent term.

Equation 1 describes how TEMP uses a contextual encoder to return a sequence of vectors, given a term’s definition  $S$  and an arbitrary taxonomy path  $P$ .

$$\text{Encoder}(S, P) = v_{[\text{CLS}]}, v_1, \dots, v_{[\text{SEP}]}, v_{p_d}, \dots, v_{\text{root}} \quad (1)$$

The TEMP method, which uses pre-trained contextual encoders to model taxonomy paths, has been an inspiration for our work. However, TEMP

also has some limitations. One of the main limitations is that it can only expand new leaf nodes. Additionally, TEMP has some issues such as:

1) Limited use of local information - although TEMP uses paths to narrow the search range within the taxonomy tree, the problem of too long paths can still arise. In such cases, distant relationships may have a limited impact on the determination of leaf nodes.

2) Inadequate utilization of pre-trained model - TEMP only connects the nodes of the path using special tokens such as  $[\text{SEP}]$  or  $[\text{UNK}]$ , which does not fully leverage the knowledge encoded by the pre-trained language model.

### 3.1.2 PPT: Taxonomy-path Prompt with Pre-trained model

To address the limitations of the TEMP method, we proposed PPT (A Taxonomy Expansion Method Based on Taxonomy Path Prompt and Pre-Trained Model). Our approach includes a few improvements:

Utilization of local information - Instead of using the entire taxonomy path, we use the local information nodes  $l_p$  closest to the nodes. For example, in Figure 2, for the node “Archeology”, the local information nodes would be “Archeology” and “Anthropology”. When the depth of the taxonomy path is less than two, we take only one node.

$$l_p = \text{local}(P) = \{N_{D-1}, N_D\} \quad (2)$$

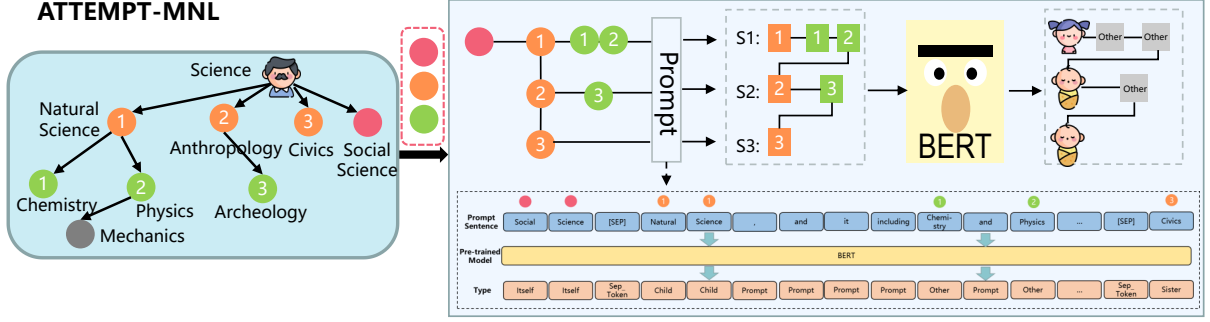


Figure 3: An overview of the MNL (multiple nodes labeling). We concatenate the current node (red), parent node, candidate children nodes (orange) and their children nodes (green) together to form a unified prompt, and then input into BERT. Finally, we realize the prediction of various types of nodes through joint labeling.

Improved pre-trained model utilization - We form a set of taxonomy path points  $P_{SocialScience} = (\text{Archeology-Antropology-Social Science})$  by combining the local information points of each node and the node Social Science to be extended. We then generate the appropriate natural language  $S_{Gen}$  using a prompt function.

$$S_{Gen}(q, l_p) = Prompt(q, l_p) \quad (3)$$

where  $q$  is the node to be expanded and  $Prompt$  is a function to generate natural language from prompts. For example,  $S_{Gen}(q, l_p) = \text{"Social Science including Anthropology, and Anthropology including Archeology"}$ . We feed this generated language  $S_{Gen}$  into the pre-trained model, rank the results in the same way as TEMP, and use the highest score as the parent node of the given node.

$$Encoder(S_{Gen}) = v'_{[CLS]}, v'_1, \dots, v'_w \quad (4)$$

The encoder results are as above, where  $w$  is the number of output vectors. We trained the model with Margin Ranking Loss (MRL), which is defined as follows:

$$\mathcal{L} = \sum_{P \in \mathcal{P}^+} \sum_{P' \in \mathcal{P}^-} \max(0, f(P') - f(P) + \gamma(P, P')) \quad (5)$$

where  $\mathcal{P}^+$  is the set of taxonomy-paths in the taxonomy,  $\mathcal{P}^-$  is the set of negative samples, and  $\gamma(P, P')$  is a function designed for the margin between positive and negative taxonomy-paths. To capture the semantic similarity of different taxonomy-paths, we follow TEMP to set a dynamic margin function based on the semantic similarity as follows:

$$\gamma(P, P') = \left( \frac{|P \cup P'|}{|P \cap P'|} - 1 \right) * k \quad (6)$$

where  $k$  is a parameter used to adjust margins (usually between 0.1 and 1).

### 3.2 Stage two: Children Finding

The second stage of ATTEMPT is to identify all the children nodes of a given node in the taxonomy. For example, finding the children nodes "Anthropology" and "Civics" for the node "Social Science", as shown in Figure 2. We propose two methods for this stage: PWC and MNL.

#### 3.2.1 PWC: Pair-wise Classification

In the second stage, we identify all the child nodes of a given term. To do this, we form pairs of possible "hypernym-hyponym" term pairs from the node to be expanded (red node) and each candidate child node (orange node, child of the parent identified in the first stage). These term pairs are connected with the special token  $[SEP]$  and fed into a pre-trained language model such as BERT. An example can be seen in Figure 2, where the node to be classified is "Social Science" and the orange candidate child nodes are "natural science," "anthropology" and "civics."

We use the pre-trained model to perform binary classification to determine whether the term pairs have a "hypernym-hyponym" relationship or not. The traditional cross-entropy function is used as the loss function to train the classification model. This method is simple, because the pre-trained model has been trained on a large corpus already and it can identify whether the term pairs have a hierarchical relationship or not. This method is called Pair-wise classification.

Dataset		Environment			Food			Science		
Metric		Wu&P	MRR	Acc	Wu&P	MRR	Acc	Wu&P	MRR	Acc
Leaf Nodes	BERT+MLP	47.9	21.5	11.1	47.0	14.9	10.5	43.6	15.7	11.5
	TaxoExpan	54.8	32.3	11.1	54.2	40.5	27.6	57.6	44.8	27.8
	STEAM	69.6	46.9	36.1	67.0	43.4	34.2	68.2	48.3	36.5
	TMN	54.0	43.6	35.0	65.9	47.2	34.7	75.9	53.2	41.9
	TEMP-BERT	75.9	62.0	49.0	78.3	57.1	45.2	84.6	64.6	54.4
	ATTEMPT-PPT(ours)	<b>82.3</b>	<b>75.1</b>	<b>65.4</b>	<b>78.4</b>	<b>58.1</b>	<b>46.5</b>	<b>86.6</b>	<b>70.7</b>	<b>61.2</b>
All Nodes	TEMP-BERT	64.8	33.5	15.6	84.3	41.1	37.4	88.0	52.9	43.4
	ATTEMPT-PPT(ours)	<b>81.1</b>	<b>52.5</b>	<b>37.5</b>	84.3	<b>42.6</b>	<b>39.0</b>	<b>90.9</b>	<b>60.0</b>	<b>52.8</b>

Table 1: Baseline comparison on the three datasets in stage one (in %)

### 3.2.2 MNL: Multiple Nodes Labeling

MNL is a new approach that addresses the problem of identifying multiple children of a given node in the taxonomy. There are two main challenges: determining whether a node has children and how many children it has, and identifying as many children as possible if there are multiple children.

To address these challenges, we first determine whether the given node is a leaf node (has no children) and if so, the second stage ends. If there are multiple children, we treat this as a multiple-choice problem and model it as a sequential labeling task. As shown in Figure 3, we extract the possible siblings, children, and grandchildren (orange and green nodes) of the given node to make use of local information. We then use a prompt function to convert these three types of nodes into natural language (e.g., “Natural Science - Chemistry, Physics” is converted to “Natural Science, and it including Chemistry and Physics”).

We concatenate the node to be expanded “Social Science” with all the sentences generated by the prompt, and then feed this into the pre-trained model. Since the model was trained on a large corpus of natural language, the input of natural language is consistent with the pre-training phase, which helps to fully utilize the hidden information of the model and correctly identify the contextual relationships. The addition of local information provides additional context to the model, which allows it to make more accurate predictions about the children of the given term.

## 4 Experiments

In this section, we first describe the experimental setup and implementation details in Section 4.1 and Section 4.2. We then present the results of

our experiments in Section 4.3, including a comparison of our approach to the baseline method. To further understand the contribution of different components of our approach, we conduct ablation experiments in Section 4.4 to investigate the effectiveness of using local information and prompts in ATTEMPT.

### 4.1 Experimental Setup

**Datasets.** We conducted experiments on two datasets that include six domains and two types of nodes. The first dataset is the Semeval-2016 task 13 dataset, which was used to evaluate the performance of expanding leaf nodes in stage one. We compared our method to previous approaches such as TEMP (Liu et al., 2021) and STEAM (Yu et al., 2020), which have also been tested on this dataset for leaf node expansion.

To evaluate the expansion of non-leaf nodes, we constructed a new dataset based on Semeval, as there are limited previous datasets that are relevant to this task. This dataset was specifically designed for the purpose of non-leaf node expansion and evaluation.

The following is a description of the two datasets: 1) We used the dataset from Semeval-2016 task 13<sup>2</sup>, which contains three English datasets for the environment, science, and food domains. We followed the setup as in (Yu et al., 2020) and used the randomly-grown taxonomies for self-supervised learning, and sampled 20% of the leaf nodes for testing. We used this dataset to compare our method with other taxonomy extension methods for leaf nodes. 2) As there is limited data available for non-leaf node expansion, we re-constructed the original data. We defined nodes with one parent and no children as leaf nodes and

<sup>2</sup><https://alt.qcri.org/semeval2016/task13/>

Dataset		Environment		Science		Food	
Metric		acc	Avg(F1)	acc	Avg(F1)	acc	Avg(F1)
Non-Leaf	PWC	21.5	50.8	20.0	54.3	12.5	32.1
Nodes	ATTEMPT-MNL(ours)	<b>46.2</b>	<b>68.7</b>	<b>33.3</b>	<b>59.8</b>	<b>37.5</b>	<b>69.6</b>
All	PWC	47.1	56.3	58.1	<b>74.65</b>	60.3	74.8
Nodes	ATTEMPT-MNL(ours)	<b>64.7</b>	<b>82.8</b>	<b>61.3</b>	74.09	<b>73.3</b>	<b>84.5</b>

Table 2: Baseline comparison on the three datasets in stage two (in %)

Dataset	Environment		Science		Food	
Metric	acc	Avg(F1)	acc	Avg(F1)	acc	Avg(F1)
Baseline	11.8	14.7	22.5	23.3	25.9	30.9
ATTEMPT(ours)	<b>12.0</b>	<b>15.0</b>	<b>22.6</b>	<b>28.1</b>	<b>26.7</b>	<b>32.9</b>

Table 3: Comparison of the baseline method and ATTEMPT in the overall process(in %)

nodes with one parent and at least one child as non-leaf nodes. More details about the dataset are provided in Appendix A.

**Metrics.** For the parent finding process in stage 1, we followed the evaluation strategy of (Yu et al., 2020) using Accuracy, Mean reciprocal rank (MRR), and Wu & Palmer similarity (Wu&P) to evaluate our methods. Accuracy (ACC) measures the count of parent or child nodes that are accurately predicted. MRR calculates the average of reciprocal ranks of the true taxonomy path. Wu&P measures the semantic similarity between the predicted taxonomy path and the truth taxonomy-path.

For stage two, we proposed two metrics for evaluating the effectiveness of this phase. One is ACC, which represents whether all children can be found or not. The second one is Avg F1, which can further evaluate how many children are found for a given node.  $\text{Avg(F1)} = \frac{1}{n} \sum_{i=1}^n \text{F1}$

**Compared Methods.** We compare with the following methods:

- **BERT+MLP** The method extracts terms embeddings from BERT and then feeds them into a multilayer perceptron (MLP) to predict their relationship.

- **TEMP** (Liu et al., 2021) One state-of-the-art taxonomy expansion framework which predicts new concepts’ position by ranking the generated taxonomy paths. The first method that employs pre-trained contextual encoders in taxonomy construction and hypernym detection problems.

- **STEAM** (Yu et al., 2020) A taxonomy expansion framework that leverages natural supervision in the existing taxonomy for expansion.

- **TaxoExpan** (Shen et al., 2020) A self-supervised method for encoding local structures in seed taxonomy using location-enhanced graph

neural networks.

- **TMN** (Zhang et al., 2021) A Triplet Matching Network (TMN) that finds suitable hypernym, hyponym word pairs for a given query concept.

## 4.2 Implementation Details

We present the PPT method for the first stage of leaf node expansion, which is based on TEMP (TEMPS’ code link <sup>3</sup>). We use BERT (bert-base-uncased) as the pre-trained language model and split the terms into 10% for validation and 10% for testing. To expand the full type of nodes, both leaf and non-leaf, we use the new data introduced previously and select the same number of leaf and non-leaf nodes as the test set. We use the default optimal hyperparameters of the original TEMP authors and experiment with different learning rates to obtain the best performance. We also use multiple prompts (see Appendix C) according to the settings of Musubu (Takeoka et al., 2021), and take the average result as the experimental result. To reduce the impact of randomness, we repeat the experiment three times.

For the MNL method in stage two, we connect the nodes to be expanded (red), the candidate child nodes (orange), and the child nodes of the candidate nodes (green) and generate natural language by way of prompt. The generated natural language is fed into the pre-training model and labelled. We label the real children of a given node as 1, the sibling nodes as 0, and ignore the computational loss for all the rest of the nodes. In addition, if a term has multiple tokens and one of the tokens is marked as one by the model, we mark all those tokens as child nodes. See Appendix B for more details.

<sup>3</sup><https://github.com/liu-zichen/TEMP>

Dataset	Environment			Science			Food		
Metric	Wu&P	MRR	Acc	Wu&P	MRR	Acc	Wu&P	MRR	Acc
no path nodes	81.8	73.2	59.6	81.0	67.2	56.5	78.5	57.5	44.1
no prompt	79.0	64.8	51.9	84.1	67.0	56.5	78.4	56.9	45.1
ATTEMPT-PPT(ours)	<b>82.3</b>	<b>75.1</b>	<b>65.4</b>	<b>86.6</b>	<b>70.7</b>	<b>61.2</b>	<b>78.7</b>	<b>58.1</b>	<b>46.5</b>

Table 4: Results of ablation experiments on the three dataset in stage one (in %)

Dataset	Environment		Science		Food	
Metric	acc	Avg(F1)	acc	Avg(F1)	acc	Avg(F1)
no prompt	52.9	53.7	<b>62.0</b>	<b>75.0</b>	52.7	63.4
no grandchild	52.9	52.9	51.6	62.3	51.1	59.8
ATTEMPT-MNL(ours)	<b>64.7</b>	<b>82.8</b>	61.3	74.2	<b>73.3</b>	<b>84.5</b>

Table 5: Results of ablation experiments on the three dataset in stage two (in %)

### 4.3 Experimental Results

As shown in Table 1, our method PPT outperforms the existing TEMP model significantly on both leaf and non-leaf nodes. For leaf nodes, we improved the TEMP model by 8.2%, 6.7%, and 2.8% on Acc, MRR, and Wu&P, respectively. For all types of nodes, the improvement is 11.0%, 9.2%, and 6.4%, respectively.

The comparison results of the two methods tested in the child discovery phase are presented in Table 2. For leaf nodes, the MNL method improves Acc and Avg(F1) by 21% and 20.3%, respectively, compared to the pair-wise classification method over the three benchmark datasets. For all type nodes, the improvement is 11.3% and 11.9%, respectively.

Table 3 presents the comparison results between the baseline method and our ATTEMPT. The baseline method achieves 14.7%, 23.3%, and 30.9% in Avg(F1) metrics for the three datasets of environment, science, and food, respectively. Our ATTEMPT method improved the Avg(F1) by an average of 2.4% over the baseline. The low results in Table 3 are due to the challenging nature of the task. To obtain the correct parent node, all child nodes must be successfully identified. This highlights the potential for further improvement.

### 4.4 Ablation Studies

To verify local information and prompt effectiveness, we compare and test the changes in experimental results with/without these two types of information on both stages.

**Local Information** As shown in Table 4, after removing the path nodes, the PPT method in stage 1 decreases on average by 4.3%, 2%, and

2.1% on acc, mrr, wu&p, respectively, on the three datasets. Table 5 also shows that the MNL method decreases by 14.6% and 22.1% on average on accuracy and average F1 score, respectively, after removing the grandchild node information in the child finding stage. We found that local information is essential in both the first and second phases, particularly in the second child lookup phase. Removing local information brings about a significant performance degradation, which may be attributed to our method’s modelling of relationships. The individual nodes are closely associated in our MNL method.

**Prompt** In Table 4, the PPT method with prompt removal decreased in acc, mrr, wu&p by 6.5%, 5.1%, and 2.0% on average, respectively. Meanwhile, in the second stage, the MNL method decreased 8.6% and 15.6% for accuracy and average F1 metrics, respectively, after prompt removal. The scientific data in the second stage showed a slight performance improvement after prompt removal, which we speculate may be due to insufficient data and pre-trained corpus. Overall, the prompt is essential for the parent finding process in the first stage and the child finding process in the second stage.

## 5 Conclusion

This paper proposes a two-stage taxonomy completion method based on pre-trained Language models (ATTEMPT), which effectively inserts the new concept in the correct position by finding a parent node then labeling children nodes. In addition, we use prompt to generate natural language information suitable for the pre-trained model further to improve the effectiveness of parent node



recognition and children labeling for the given node. Our experiments on two types and three domains with six datasets show that our method can enhance the effectiveness of locating the position of a given node in existing taxonomies. Furthermore, the efficacy of local information and prompts in ATTEMPT is also demonstrated by ablation experiments. In conclusion, our proposed ATTEMPT method is an effective approach for taxonomy completion, and it can be further improved with more comprehensive datasets.

## Limitations

Since ATTEMPT uses the pre-trained language model to complete the taxonomy, the expansion effect is limited by the model. Generally, pre-trained models with more knowledge scales are better (e.g., BERT-Large V.S. BERT-Base-uncased). However, our paper focuses on how to fully use the knowledge of the pre-trained model rather than verifying whether more knowledge scales better or not. Based on the above, this paper does not conduct more related research (in fact, TEMP (Liu et al., 2021) has been compared and reached similar conclusions). In addition, the selection of prompts will also affect the expansion effect. For the convenience of comparison, we have selected several basic prompts (the same as Musubu (Takeoka et al., 2021)) for experimentation. In future work, we plan to study how to construct or select better prompts for classification expansion. We do not consider the situation of multi-parent nodes according to the TEMP (Liu et al., 2021) settings. And according to our statistics, there are only a few multi-parent nodes in the Semeval-2016(task-13) datasets (1/3843). We will continue investigating how to make better use of the pre-trained model knowledge to solve the taxonomy completion problem.

## Acknowledgements

This work was supported by the National Key Research and Development Program of China (No. 2020AAA0106400), the National Natural Science Foundation of China (No.61976211) and the Strategic Priority Research Program of Chinese Academy of Sciences (Grant No. XDA27020200). This research work was supported by the Youth Innovation Promotion Association CAS, Yunnan Provincial Major Science and Technology Special Plan Projects (No.202202AD080004) and CCF-

DiDi GAIA Collaborative Research Funds for Young Scholars.

## References

- Eugene Agichtein and Luis Gravano. 2000. Snowball: Extracting relations from large plain-text collections. In *Proceedings of the fifth ACM conference on Digital libraries*, pages 85–94.
- Daniele Alfarone and Jesse Davis. 2015. Unsupervised learning of an is-a taxonomy from a limited domain-specific corpus. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- Guillaume Cleuziou and Jose G. Moreno. 2016. QAS-SIT at SemEval-2016 task 13: On the integration of semantic vectors in pretopological spaces for lexical taxonomy acquisition. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1315–1319, San Diego, California. Association for Computational Linguistics.
- Anne Cocos, Marianna Apidianaki, and Chris Callison-Burch. 2018. Comparing constraints for taxonomic organization. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 323–333.
- Sarthak Dash, Md Faisal Mahbub Chowdhury, Alfio Gliozzo, Nandana Mihindukulasooriya, and Nicolas Rodolfo Fauceglia. 2020. Hypernym detection using strict partial order networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7626–7633.
- Thomas Demeester, Tim Rocktäschel, and Sebastian Riedel. 2016. Lifted rule injection for relation embeddings. *arXiv preprint arXiv:1606.08359*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Doug Downey, Chandra Bhagavatula, and Yi Yang. 2015. Efficient methods for inferring large sparse topic hierarchies. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 774–784.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Marti A. Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *International Conference on Computational Linguistics*.

- Jiaxin Huang, Yiqing Xie, Yu Meng, Yunyi Zhang, and Jiawei Han. 2020. Corel: Seed-guided topical taxonomy construction by concept learning and relation transferring. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1928–1936.
- Jin Huang, Zhaochun Ren, Wayne Xin Zhao, Gaole He, Ji-Rong Wen, and Dong Daxiang. 2019. Taxonomy-aware multi-hop reasoning networks for sequential recommendation. In *Web Search and Data Mining*.
- Meng Jiang, Jingbo Shang, Taylor Cassidy, Xiang Ren, Lance M Kaplan, Timothy P Hanratty, and Jiawei Han. 2017. Metapad: Meta pattern discovery from massive text corpora. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 877–886.
- Zichen Liu, Hongyuan Xu, Yanlong Wen, Ning Jiang, HaiYing Wu, and Xiaojie Yuan. 2021. TEMP: Taxonomy expansion with dynamic margin loss through taxonomy-paths. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3854–3863, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2018. Decoupled weight decay regularization. In *International Conference on Learning Representations*.
- Alexander Panchenko, Stefano Faralli, Eugen Ruppert, Steffen Remus, Hubert Naets, Cédric Fairon, Simone Paolo Ponzetto, and Chris Biemann. 2016. TAXI at SemEval-2016 task 13: a taxonomy induction method based on lexico-syntactic patterns, substrings and focused crawling. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 1320–1327, San Diego, California. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. *Pytorch: An imperative style, high-performance deep learning library*. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Stephen Roller, Douwe Kiela, and Maximilian Nickel. 2018. Hearst patterns revisited: Automatic hypernym detection from large text corpora. *arXiv preprint arXiv:1806.03191*.
- Jingbo Shang, Xinyang Zhang, Liyuan Liu, Sha Li, and Jiawei Han. 2020. Nettaxo: Automated topic taxonomy construction from text-rich network. In *Proceedings of The Web Conference 2020*, pages 1908–1919.
- Jiaming Shen, Zhihong Shen, Chenyan Xiong, Chi Wang, Kuansan Wang, and Jiawei Han. 2020. *Taxo-Expan: Self-Supervised Taxonomy Expansion with Position-Enhanced Graph Neural Network*, page 486–497. Association for Computing Machinery, New York, NY, USA.
- Jiaming Shen, Zeqiu Wu, Dongming Lei, Chao Zhang, Xiang Ren, Michelle Vanni, Brian M. Sadler, and Jiawei Han. 2018. Hiexpan: Task-guided taxonomy construction by hierarchical tree expansion. In *Knowledge Discovery and Data Mining*.
- Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958.
- Kunihiro Takeoka, Kosuke Akimoto, and Masafumi Oyamada. 2021. Low-resource taxonomy enrichment with pretrained language models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2747–2758, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Chengyu Wang, Yan Fan, Xiaofeng He, and Aoying Zhou. 2019. A family of fuzzy orthogonal projection models for monolingual and cross-lingual hypernymy prediction. In *The World Wide Web Conference*, pages 1965–1976.
- Wenpeng Yin and Dan Roth. 2018. Term definitions help hypernymy detection. *arXiv preprint arXiv:1806.04532*.
- Xiaoxin Yin and Sarthak Shah. 2010. Building taxonomy of web search intents for name entity queries. In *Proceedings of the 19th international conference on World wide web*, pages 1001–1010.
- Yue Yu, Yinghao Li, Jiaming Shen, Hao Feng, Jiemeng Sun, and Chao Zhang. 2020. *STEAM: Self-Supervised Taxonomy Expansion with Mini-Paths*, page 1026–1035. Association for Computing Machinery, New York, NY, USA.
- Qingkai Zeng, Jinfeng Lin, Wenhao Yu, Jane Cleland-Huang, and Meng Jiang. 2021. Enhancing taxonomy completion with concept generation via fusing relational representations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2104–2113.
- Chao Zhang, Fangbo Tao, Xiushi Chen, Jiaming Shen, Meng Jiang, Brian Sadler, Michelle Vanni, and Jiawei Han. 2018a. Taxogen: Unsupervised topic taxonomy construction by adaptive term embedding and clustering. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2701–2709.
- Chao Zhang, Fangbo Tao, Xiushi Chen, Jiaming Shen, Meng Jiang, Brian M. Sadler, Michelle Vanni, and Jiawei Han. 2018b. Taxogen: Unsupervised topic

taxonomy construction by adaptive term embedding and clustering. In *Knowledge Discovery and Data Mining*.

Jieyu Zhang, Xiangchen Song, Ying Zeng, Jiaze Chen, Jiaming Shen, Yuning Mao, and Lei Li. 2021. Taxonomy completion via triplet matching network. In *AAAI*.

Yuchen Zhang, Amr Ahmed, Vanja Josifovski, and Alexander Smola. 2014. Taxonomy discovery for personalized recommendation. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 243–252.

Yunyi Zhang, Jiaming Shen, Jingbo Shang, and Jiawei Han. 2020. Empower entity set expansion via language model probing. *arXiv preprint arXiv:2004.13897*.

## A Dataset

The original dataset and our reconstructed dataset statistics are in Table 6 and 7.

Dataset	Environment	Science	food
$ \mathcal{N} $	261	429	1486
$ \mathcal{L} $	200	306	1161
$ \mathcal{D} $	6	8	8
$ \Delta_{leaf} $	3.79	5.17	5.37
$ \Delta_{non} $	2.97	4.41	4.78
train(leaf/non)	148/59	226/112	893/284
test(leaf/non)	52/0	85/0	297/0

Table 6: Statistics of the original taxonomy datasets for evaluation.  $|\mathcal{N}|$  and  $|\mathcal{L}|$  are the number of nodes and leaf nodes in the taxonomy.  $|\mathcal{D}|$ ,  $|\Delta_{leaf}|$  and  $|\Delta_{non}|$  indicate the depth of the taxonomy and the average depth of leaf nodes and non-leaf nodes respectively. Train(leaf/non) and test(leaf/non) represents the proportion of leaf and non-leaf nodes in train data and test data.

Dataset	Environment	Science	food
$ \mathcal{N} $	261	429	1486
$ \mathcal{L} $	200	306	1161
$ \mathcal{D} $	6	8	8
$ \Delta_{leaf} $	3.79	5.17	5.37
$ \Delta_{non} $	2.96	4.41	4.05
train(leaf/non)	187/46	292/94	1124/214
test(leaf/non)	13/13	21/21	74/74

Table 7: Statistics of the new taxonomy datasets for evaluation.

To prevent the test data from being leaked during training and to thoroughly test the generalization ability of the model when encountering unseen data, we split each original taxonomy tree into two subtrees, one for training and one for testing. For example, the left subtree of the scientific taxonomy in Figure 2, natural science and its children, is used as the test subtree, and the rest is used for training. Specifically, we select the subtree with 20% of the number of nodes of the current taxonomy tree as the subtree for testing and

ignore too many leaf nodes to ensure the ratio of leaf nodes to non-leaf nodes is 1:1. Too many leaf nodes will make the child finding stage degenerate into an expansion of leaf nodes, and the model will be easily overfitting. And too few leaf nodes will make the test inadequate, so we use equal leaf and non-leaf node data as the test.

In the training and testing phases, we dig out the node to be expanded in the current taxonomy tree, and if the node has N children, these N children are reassigned to the original parent of the node to be expanded as child nodes. We ignore the case of double parent nodes because their existence is too rare. Only one node in the three datasets containing more than 2000 nodes in our experiments has a dual-parent node. We will consider this case further in our future work.

## B Implementation Details

For the fairness of the experiment, we follow the setting of TEMP (Liu et al., 2021). We use 10% terms for validating and 10% for testing. For each benchmark, we try various learning rates and report the best performance. We use multiple prompts to experiment and select the average result as the experimental result. We repeated the experiment three times to reduce the impact of randomness. We train the model using the Pytorch<sup>4</sup> (Paszke et al., 2019) on the NVIDIA RTX3090 GPU. For all methods, the bert-base-uncased<sup>5</sup> model are chosen for feature extraction. The pre-trained contextual encoders are of base size with 12 layers. We use the AdamW (Loshchilov and Hutter, 2018) as the optimizer with the warm-up (He et al., 2016), and fine tune the whole model with a learning rate of 2e-5. The dropout (Srivastava et al., 2014) of 0.1 is applied to prevent overfitting.

## C Prompt Details

Name	Prompt
Such-as	Y such as X and Z
One-of	X is one of Y, and Z is one of Y
Especiallly	Y, especiallly X and Z
Is-a	X is a Y, and Z is a Y
Including	Y including X and Z

Table 8: List of prompts used in the experiments. Y denotes a parent term of a term X and Z

<sup>4</sup><https://pytorch.org>

<sup>5</sup><https://huggingface.co/bert-base-uncased>