

Enquire One’s Parent and Child Before Decision: Fully Exploit Hierarchical Structure for Self-Supervised Taxonomy Expansion

Suyuchen Wang*
suyuchen.wang@umontreal.ca
Mila & DIRO, Université de Montréal
Montréal, Québec, Canada

Ruihui Zhao
zacharyzhao@tencent.com
Tencent Jarvis Lab
Shenzhen, Guangdong, China

Xi Chen
jasonxchen@tencent.com
Tencent Jarvis Lab
Shenzhen, Guangdong, China

Yefeng Zheng
yefengzheng@tencent.com
Tencent Jarvis Lab
Shenzhen, Guangdong, China

Bang Liu†
bang.liu@umontreal.ca
Mila & DIRO, Université de Montréal
Montréal, Québec, Canada

ABSTRACT

Taxonomy is a hierarchically structured knowledge graph that plays a crucial role in machine intelligence. The taxonomy expansion task aims to find a position for a new term in an existing taxonomy to capture the emerging knowledge in the world and keep the taxonomy dynamically updated. Previous taxonomy expansion solutions neglect valuable information brought by the hierarchical structure and evaluate the correctness of merely an added edge, which downgrade the problem to node-pair scoring or mini-path classification. In this paper, we propose the Hierarchy Expansion Framework (HEF), which fully exploits the hierarchical structure’s properties to maximize the coherence of expanded taxonomy. HEF makes use of taxonomy’s hierarchical structure in multiple aspects: i) HEF utilizes subtrees containing most relevant nodes as self-supervision data for a complete comparison of parental and sibling relations; ii) HEF adopts a coherence modeling module to evaluate the coherence of a taxonomy’s subtree by integrating hypernymy relation detection and several tree-exclusive features; iii) HEF introduces the Fitting Score for position selection, which explicitly evaluates both path and level selections and takes full advantage of parental relations to interchange information for disambiguation and self-correction. Extensive experiments show that by better exploiting the hierarchical structure and optimizing taxonomy’s coherence, HEF vastly surpasses the prior state-of-the-art on three benchmark datasets by an average improvement of 46.7% in accuracy and 32.3% in mean reciprocal rank.

KEYWORDS

taxonomy expansion, self-supervised learning, hierarchical structure

1 INTRODUCTION

Taxonomy is a particular type of hierarchical knowledge graph that portrays the hypernym-hyponym relations or “is-A” relations of various concepts and entities. They have been adopted as the underlying infrastructure of a wide range of online services in various domains, such as product catalogs for e-commerce [14, 22], scientific indices like *MeSH* [19], and lexical databases like *WordNet* [25]. A well-constructed taxonomy can assist various downstream tasks,

*Work done during an internship at Tencent Jarvis Lab.

†Corresponding author.

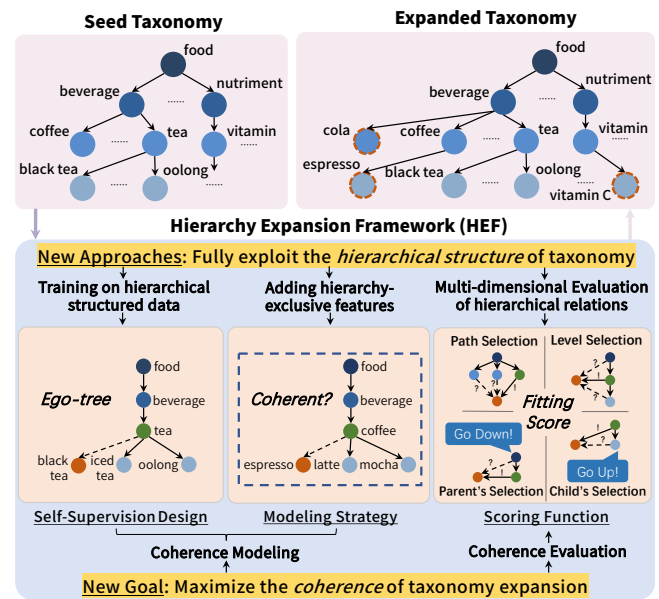


Figure 1: An illustration of the taxonomy expansion task and the contributions of the proposed HEF model.

including web content tagging [20, 27], web searching [46], personalized recommendation [10] and helping users achieve quick navigation on web applications [9]. Manually constructing and maintaining a taxonomy is laborious, expensive and time-consuming. It is also highly inefficient and detrimental for downstream tasks if we construct a taxonomy *from scratch* [7, 41] as long as the taxonomy has new terms to be added. A more realistic strategy is to insert new terms (“*query*”) into an existing taxonomy, i.e., the seed taxonomy, as a child of an existing node in the taxonomy (“*anchor*”) without modifying its original structure to best preserve its design. This problem is called *taxonomy expansion* [13].

Early taxonomy expansion approaches use terms that do not exist in the seed taxonomy and its best-suited position in the seed taxonomy as training data [12]. However, it suffers from the insufficiency of training data and the shortage of taxonomy structure supervision. More recent solutions adopt self-supervision and try to exploit the information of nodes in the seed taxonomy (seed

nodes) to perform node pair matching [34] or classification along mini paths in the taxonomy [47]. However, these approaches do not fully utilize the taxonomy’s hierarchical structure’s characteristics, and neglect the coherence of the extended taxonomy which oughts to be the core of the taxonomy expansion task. More specifically, existing approaches do not model a hierarchical structure identical to the taxonomy. Instead, they use ego-nets [34] or mini-paths [47] and feature few or no tree-exclusive information, making them unable to extract or learn the complete hierarchical design of a taxonomy. Besides, they do not consider the coherence of a taxonomy. They manage to find the most suitable node in a limited subgraph and only evaluate the correctness of a single edge instead of the expanded taxonomy, which downgrades the taxonomy expansion task to a hypernymy detection task. Lastly, their scoring approach regards the anchor node as an individual node without considering the hierarchical context information. However, the hierarchical structure provides multi-aspect criteria to evaluate a node, such as its path or level correctness. The structure also marks the nodes that are most likely to be wrongly chosen to be a parent in a specific parental relation.

To solve all the stated flaws in previous works, we propose the Hierarchy Expansion Framework (HEF), which aims to maximize the coherence of the expanded taxonomy instead of the fitness of a single edge by fully exploiting the hierarchical structure of a taxonomy for self-supervised training, as well as modeling and evaluating the structure of taxonomy. HEF’s designs and goals are illustrated in Fig. 1. Specifically, we make the following contributions.

Firstly, we design an innovative hierarchical data structure for self-supervision to mimic how humans construct a taxonomy. Relations in a taxonomy include hypernymy relations along a root-path and similarity among siblings. To find the most suitable parent node for the query term, human experts need to compare an anchor node with all its ancestors to distinguish the most appropriate one and compare the query with its potential siblings to testify their similarity. For example, to choose the parent for query “black tea” in the food taxonomy, the most appropriate anchor “tea” can only be selected by distinguishing from its ancestors “beverage” and “food”, which are all “black tea”’s hypernyms, as well as compare the query “black tea” with “tea”’s children like “iced tea” and “oolong” to guarantee similarity among siblings. Thus, we design a new structure called “ego-tree” for self-supervision, which contains all ancestors and sample of children of a node for taxonomy structure learning. Our ego-tree incorporates richer topological context information for attaching a query term to a candidate parent with minimal computation cost compared to previous approaches based on node pair matching or path information.

Secondly, we design a new modeling strategy to perform explicit ego-tree coherence modeling apart from the traditional node-pair hypernymy detection. Instead of merely modeling the correctness of the added edge, we adopt a more comprehensive approach to detect whether the anchor’s ego-tree after adding the query maintains the original design of the seed taxonomy. The design of taxonomy includes natural hypernymy relations, which needs the representation of node-pair relations and expert-curated level configurations, such as *species* must be placed in the eighth level of biological taxonomy, or adding one more adjective to a term means exactly one level higher in the e-commerce taxonomy. We adopt a coherence

modeling module to detect the two aspects of coherence: i) For natural hypernymy relations, we adopt a hypernymy detection module to represent the relation between the query and each node in the anchor’s ego-tree. ii) For expert-curated designs, we integrate hierarchy-exclusive features such as embeddings of a node’s absolute level and relative level to the anchor into the coherence modeling module.

Thirdly, we design a multi-dimensional evaluation to score the coherence of the expanded taxonomy. The hierarchical structure of taxonomy allows the model to evaluate the correctness of path selection and level selection separately and the parental relationships in a hierarchy not only allow the model to disambiguate the most similar terms but also enables the model to self-correct its level selection by deciding the current anchor’s granularity is too high or too low. We introduce the Fitting Score for the coherence evaluation of the expanded ego-tree by using a Pathfinder and a Stopper to score path correctness and level correctness, respectively. The Fitting Score calculation also disambiguates the most appropriate anchor from its parent and children and self-correct its level selection by bringing the level suggestion from the anchor’s parent and one of its children into consideration. The Fitting Score’s optimization adopts a self-supervised multi-task training paradigm for the Pathfinder and Stopper, which automatically generates training data from the seed taxonomy to utilize its information fully.

We conduct extensive evaluations based on three benchmark datasets to compare our method with state-of-the-art baseline approaches. The results suggest that the proposed HEF model significantly surpasses the previous solutions on all three datasets by an average improvement of 46.7% in accuracy and 32.3% in mean reciprocal rank. A series of ablation studies further demonstrate that HEF can effectively perform the taxonomy expansion task.

2 RELATED WORK

Taxonomy Construction. Taxonomy construction aims to create a tree-structured taxonomy with a set of terms (such as concepts and entities) from scratch, integrating hypernymy discovery and tree structure alignment. It can be further separated into two subdivisions. The first focuses on topic-based taxonomy, where each node is a cluster of several terms sharing the same topic [32, 48]. The other subdivision tackles the problem of term-based taxonomy construction, in which each node represents the term itself [3, 24, 35]. A typical pipeline for this task is to extract “is-A” relations with a hypernymy detection model first using either a pattern-based model [1, 8, 11, 28] or a distributional model [4, 18, 42, 45], then integrate and prune the mined hypernym-hyponym pairs into a single directed acyclic graph (DAG) or tree [7]. More recent solutions utilize hyperbolic embeddings [17] or transfer learning [31] to boost performance.

Taxonomy Expansion. In the taxonomy expansion task, an expert-curated seed taxonomy like MeSH [19] is provided as both the guidance and the base for adding new terms. The taxonomy expansion task is a ranking task to maximize a score of a node and its ground-truth parent in the taxonomy. Wang et al. [43] adopted Dirichlet distribution to model the parental relations. ETF [40] trained a learning-to-rank framework with handcrafted structural and semantic features. Arborist [23] calculated the ranking

score in a bi-linear form and adopted margin ranking loss. Taxo-Expan [34] modeled the anchor node by passing messages from its egonet instead of considering a single node, and scored by feeding a concatenation of egonet representation and query embedding to a feed-forward layer. STEAM [47] transformed the scoring task into a classification task on mini-paths and performed model ensemble of three sub-models processing distributional, contextual, and lexical-syntactic features, respectively. However, existing approaches mostly neglect the characteristics of taxonomy’s hierarchical structure and only evaluate the correctness of a single edge from anchor to query. On the contrary, our method utilizes the features and relations brought by the hierarchical structure and aims to enhance the expanded taxonomy’s overall coherence.

Modeling of Tree-Structured Data. Taxonomy expansion involves modeling a tree or graph structure. Plenty of works have been devoted to extending recurrent models to tree structures, like Tree-LSTM [38]. For explicit tree-structure modeling, previous approaches include modeling the likelihood of a Bayesian network [6, 43] or using graph neural net variants [34, 47]. Recently, Transformers [39] achieved state-of-the-art performance in the program translation task by designing a novel positional encoding related to paths in the tree [36] or merely transforming a tree to sequence by traversing its nodes [15]. In our work, we model tree-structure by a Transformer encoder, which, to the best of our knowledge, is the first to use the Transformer for taxonomy modeling. We adopt a more natural setting than [36] by using two different embeddings for a node’s absolute and relative level to denote positions.

3 PROBLEM DEFINITION

In this section, we provide the formal definition of the taxonomy expansion task and the explanation of key concepts that will occur in the following sections.

Definition and Concepts about Taxonomy. A taxonomy $\mathcal{T} = (\mathcal{N}, \mathcal{E})$ is an arborescence that presents hypernymy relations among a set of nodes. Each node $n \in \mathcal{N}$ represents a *term*, usually a concept mined from a large corpus online or an artificially extracted phrase. Each edge $\langle n_p, n_c \rangle \in \mathcal{E}$ points to a node from its most exact hypernym node, where n_p is n_c ’s *parent* node, and n_c is n_p ’s *child* node. Since hypernymy relation is transitive [29], such relation exists not only in node pairs connected by a single edge, but also in node pairs connected by a path in the taxonomy. Thus, for a node n in the taxonomy, its hypernym set and hyponym set consists of its *ancestors* \mathcal{A}_n , and its *descendants* \mathcal{D}_n respectively.

Definition of the Taxonomy Expansion Task. Given a *seed taxonomy* $\mathcal{T}^0 = (\mathcal{N}^0, \mathcal{E}^0)$ and the set of terms C to be added to the seed taxonomy, The model outputs the taxonomy $\mathcal{T} = (\mathcal{N}^0 \cup C, \mathcal{E}^0 \cup \mathcal{R})$, where \mathcal{R} is the newly added relations from seed nodes in \mathcal{N}^0 to new terms in C . More specifically, during the inference phase of a taxonomy expansion model, when given a *query* node $q \in C$, the model finds its best-suited parent node by iterating each node in the seed taxonomy as an *anchor* node $a \in \mathcal{N}^0$, calculating a score $f(a, q)$ representing the suitability for adding the edge $\langle a, q \rangle$, and deciding q ’s parent p_q in the taxonomy by $p_q = \arg \max_{a \in \mathcal{N}^0} f(a, q)$.

Accessible External Resources. As a term’s surface name is usually insufficient to convey the semantic information for hypernymy relationship detection, previous research usually utilizes

term definitions [13, 34] or related web pages [16, 43] to learn term representations. Besides, existing hypernymy detection solutions usually use large external corpora to discover lexical or syntactic patterns [37, 47]. As for the SemEval-2016 Task 13 datasets [2] used for our model’s evaluation, utilizing the WordNet [25] definitions is allowed by the original task, which guarantees a fair comparison with previous solutions.

4 THE HIERARCHY EXPANSION FRAMEWORK

In this section, we introduce the design of the Hierarchy Expansion Framework (HEF). An illustration of HEF is shown in Fig. 2. We first introduce the way HEF models the coherence of a tree structure, including two components for node pair hypernymy detection and ego-tree coherence modeling, respectively. Then, we discuss how HEF further exploits the hierarchical structure for multi-dimensional evaluation by the modules of Pathfinder and Stopper, and the self-supervised paradigm to train the model for the Fitting Score calculation.

4.1 Node Pair Hypernymy Detection

We first introduce the hypernymy detection module of HEF, which detects the hypernymy relationships between two terms. Unlike previous approaches that manually design a set of classical lexical-syntactic features, we accomplish the task more directly and automatically by expanding the surface names of terms to their descriptions and utilizing pre-trained language models to represent the relationship between two terms.

Given a seed term $n \in \mathcal{N}^0$ and a query term $q \in \mathcal{N}^0$ during training or $q \in C$ during inference, the hypernymy detection module outputs a representation $r_{n,q}$ suggesting how well these two terms form a hypernymy relation. Note that n might not be identical to the anchor a . Since the surface names of terms do not contain sufficient information for relation detection, we expand the surface names to their descriptions, enabling the model to better understand the semantic of new terms. We utilize the WordNet [25] concept definitions for completing this task. However, WordNet cannot explain all terms in a taxonomy due to its low coverage. Besides, many terms used in taxonomies are complex phrases like “*adaptation to climate change*” or “*bacon lettuce tomato sandwich*”. Therefore, we further develop a description generation algorithm $\text{descr}(\cdot)$, which generates meaningful and domain-related descriptions for a given term based on WordNet. Specifically, $\text{descr}(\cdot)$ is a dynamic programming algorithm that tends to integrate tokens into longer and explainable noun phrases. It describes each noun phrase by the most relative description to the taxonomy’s root’s surface name for domain relevance. The details are shown in Alg. 2 in the appendix. The input for hypernymy detection is organized as the input format of a Transformer:

$$D_{n,q} = [\text{<CLS>} \oplus \text{descr}(n) \oplus \text{<SEP>} \oplus \text{descr}(q) \oplus \text{<SEP>}],$$

where \oplus represents concatenation, <CLS> and <SEP> are the special token for classification and sentence separation in the Transformer architecture, respectively.

As shown in Fig. 2, the hypernymy detection module utilizes a pre-trained DistilBERT [30], a lightweight variant of BERT [5], to

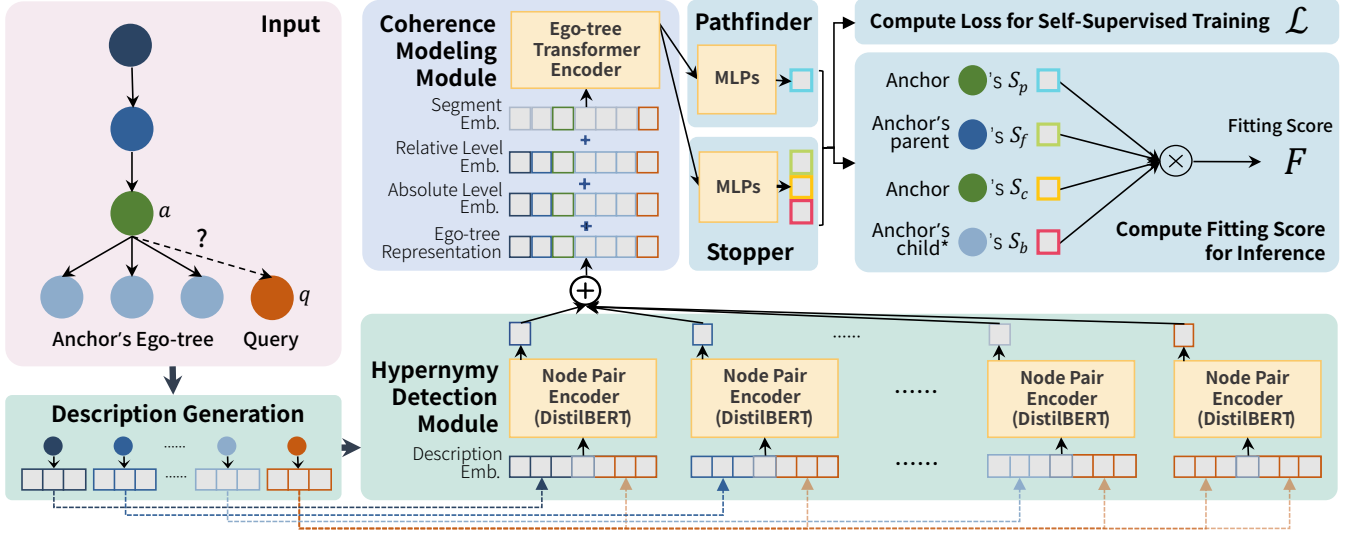


Figure 2: Illustration of the HEF model. Each circle denotes a seed node or a query node. The “Anchor’s child*” in Fitting Score calculation denotes the anchor’s child with maximum Pathfinder Score S_p .

learn the representations of cross-text relationships. Specifically, we first encode $D_{n,q}$ by $\text{DistilBERT}(\cdot)$ with positional encoding. Then we take the final layer representation of $\langle \text{CLS} \rangle$ as the representation of the node pair $\langle n, q \rangle$:

$$r_{n,q} = \text{DistilBERT}(D_{n,q})[0],$$

where index 0 represents the position of $\langle \text{CLS} \rangle$ ’s embedding.

4.2 Ego-Tree Coherence Modeling

We further design a coherence modeling module to evaluate the coherence of the tree structure after attaching the query term q into taxonomy \mathcal{T} as the anchor $a \in \mathcal{N}^0$ ’s child.

There are two different aspects for considering a taxonomy’s coherence: i) the natural hypernymy relations. Since a node’s ancestors in the taxonomy all hold hypernymy relations with it, an explicit comparison among a node’s ancestors is needed to distinguish the most appropriate one; ii) the expert-curated designs, which act as supplement information for maintaining the overall structure. Some taxonomies contain latent rules about a node’s absolute or relative levels in a taxonomy. For example, in the biological taxonomy, *kingdoms* and *species* are all placed in the second and eighth levels, respectively; in some e-commerce catalog taxonomies, terms that are one level higher than another term contain exactly one more adjective. Hence, the coherence modeling module needs to: i) model a subtree with the query as a node in it, rather than a single node pair, enabling the model to learn the design of a complete hierarchy; ii) add tree-exclusive features like absolute level or relative level compared to the anchor to assist learning the expert-curated designs of the taxonomy.

We design the *Ego-tree* \mathcal{H}_a , a novel contextual structure of an anchor a , which consists of all the ancestors and children of a (see Fig. 2). This structure contains all relevant nodes to both anchor and query, enabling the model to both compare all hypernymy relations along the root path and detect similarity among query

and its potential siblings with minimal computation cost:

$$\mathcal{H}_a = \mathcal{A}_a \cup \{a\} \cup \text{sample_child}(a), \quad (1)$$

where \mathcal{A}_a is all ancestors of a in the seed taxonomy \mathcal{T}^0 , and $\text{sample_child}(\cdot)$ means sampling at most three children of the anchor based on surface name similarity. The 3-children sampling is a trade-off between accuracy and speed, for three potential siblings are empirically enough for a comprehensive similarity comparison with the query (especially when these potential siblings are quite different) while decreasing the computation cost. Since this procedure is to leverage the similarity brought by a hierarchy’s sibling relations, sampling by surface name similarity is intuitive and cost-saving given that similar surface names usually indicate similar terms. The input of the coherence modeling module includes the anchor’s ego-tree \mathcal{H}_a and the query q as the anchor’s child in \mathcal{H}_a . For each node $n \in \mathcal{H}_a$, we represent the node pair $\langle n, q \rangle$ by the following representations:

- **Ego-tree representations.** The ego-tree representation $r_{n,q}$ is the output of the hypernymy detection module described in Sec. 4.1. It suggests the node pair’s relation.
- **Absolute level embedding.** The absolute level embedding $l_{n,q} = \text{AbsLvlEmb}(d_n)$, where d_n is the depth of n in the expanded taxonomy. When $n = q$, $l_{q,q} = \text{AbsLvlEmb}(d_a + 1)$. It assists the modeling of the expert-curated designs about granularity of a certain level.
- **Relative level embedding.** The relative level embedding $e_{n,q} = \text{RelLvlEmb}(d_n - d_q)$, where d_n is the depth of n in the expanded taxonomy. It assists the modeling of expert-curated designs about the cross-level comparison.
- **Segment embedding.** The segment embedding of $\langle n, q \rangle$ $g_{n,q} = \text{SegEmb}(\text{segment}(n))$ distinguishes anchor and query

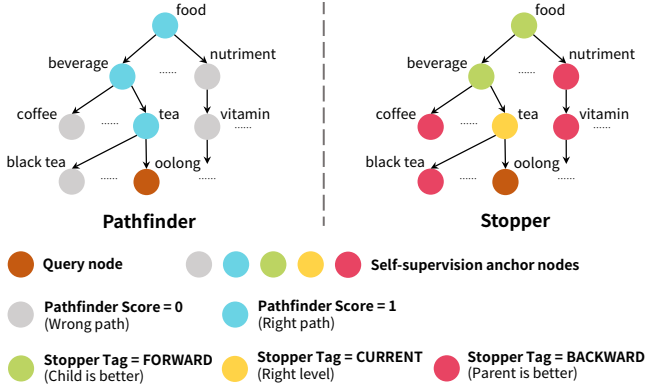


Figure 3: An illustration of the self-supervision data labels for Pathfinder and Stopper.

with other nodes in the ego-tree, where:

$$\text{segment}(n, q) = \begin{cases} 0, & \text{if } n \text{ is the anchor,} \\ 1, & \text{if } n \text{ is the query,} \\ 2, & \text{otherwise.} \end{cases}$$

The input of the coherence modeling module $R_{a,q} \in \mathbb{R}^{(|\mathcal{H}_a|+3) \times d}$ is the sum of the above embeddings calculated with the anchor’s ego-tree and the query, organized as the input of a Transformer:

$$R_{a,q} = \left[e_{\langle CLS \rangle} \oplus e_{\langle CLS \rangle} \oplus \bigoplus_{n \in \mathcal{H}_a \cup \{q\}} (r_{n,q} + l_{n,q} + e_{n,q} + g_{n,q}) \right], \quad (2)$$

where d is the dimension of embedding, $e_{\langle CLS \rangle}$ is a randomly initialized placeholder vector for obtaining the ego-tree’s path and level coherence representations, and \oplus denotes concatenation.

We implement the coherence modeling module using a Transformer encoder. Transformers are powerful to model sequences, but they lack positional information to process the relations among nodes in graphs. However, in a hierarchy like taxonomy, the level of nodes can be used as positional information, which simultaneously eliminates the positional difference of nodes on the same level. Transformers are also strong enough to integrate multiple-source information by adding their embeddings, thus they are quite suitable for modeling tree structures. In our HEF model, as shown in Fig. 2, by using two $\langle CLS \rangle$ s in the module’s input, we can obtain two different representations: $p_{a,q}$ representing the coherence of hypernymy relations (whether the *path* is correct), and $d_{a,q}$ representing the coherence of inter-level granularity (whether the *level* is correct), evaluating how well the query fits the current position in the taxonomy in both horizontal and vertical perspective:

$$p_{a,q} = \text{TransformerEncoder}(R_{a,q})[0]$$

$$d_{a,q} = \text{TransformerEncoder}(R_{a,q})[1],$$

where 0 and 1 are the position indexes of the two $\langle CLS \rangle$ s.

4.3 Fitting Score-based Training and Inference

The two representations $p_{a,q}$ and $d_{a,q}$ need to be transformed into scores indicating the fitness of placing the query q on a particular

path and a particular level. Thus, we propose the *Pathfinder* for path selection and the *Stopper* for level selection, as well as a new self-supervised learning algorithm for training and the Fitting Score calculation for inference.

Pathfinder. The Pathfinder detects whether the query is positioned on the right path. This module performs a binary classification using $p_{a,q}$. The Pathfinder Score $S_p = 1$ if and only if a and q are on the same root-path:

$$S_p(a, q) = \sigma(\mathbf{W}_{p2} \tanh(\mathbf{W}_{p1} p_{a,q} + b_{p1}) + b_{p2}), \quad (3)$$

where σ is the sigmoid function, and $\mathbf{W}_{p1}, \mathbf{W}_{p2}, b_{p1}, b_{p2}$ are trainable parameters for multi-layer perceptrons.

Stopper. The Stopper detects whether the query q is placed on the right level, i.e., under the most appropriate anchor a on a particular path. Selecting the right level is nonidentical to selecting the right path since levels are kept in order. The order of nodes on a path enables us to design a more representative module for further classifying whether the current level is too high (anchor a is a coarse-grained ancestor of q) or too low (a is a descendant of q). Thus, the Stopper module uses $d_{a,q}$ to perform a 3-class classification: searching for a better anchor node needs to go *Forward*, remain *Current*, or go *Backward*, in the taxonomy:

$$[S_f(a, q), S_c(a, q), S_b(a, q)] = \text{softmax}(\mathbf{W}_{s2} \tanh(\mathbf{W}_{s1} d_{a,q} + b_{s1}) + b_{s2}), \quad (4)$$

where $\mathbf{W}_{p1}, \mathbf{W}_{p2}, b_{p1}, b_{p2}$ are trainable parameters for multi-layer perceptrons. *Forward Score* S_f , *Current Score* S_c and *Backward Score* S_b are called *Stopper Scores*.

Self-Supervised Training. Training the HEF model needs data labels for both the Pathfinder and the Stopper. The tagging scheme is illustrated in Fig. 3. There are totally four kinds of Pathfinder-Stopper label combinations since Pathfinder Score is always 1 when Stopper Tag is *Forward* or *Current*. The training process of HEF is shown in Alg. 1. Specifically, we sample the ego-tree of all four types of nodes for a query: q ’s parent a , a ’s ancestors, a ’s descendants and other nodes, as a mini-batch for training the Pathfinder and Stopper simultaneously.

The optimization of Pathfinder and Stopper can be regarded as a multi-task learning process. The loss \mathcal{L}_q in Alg. 1 is a linear combination of the loss from Pathfinder and Stopper:

$$\mathcal{L}_q = -\eta \frac{1}{|\mathcal{X}_q|} \sum_{a \in \mathcal{X}_q} \text{BCELoss}(\hat{S}_p(a, q), S_p(a, q))$$

$$- (1 - \eta) \frac{1}{|\mathcal{X}_q|} \sum_{a \in \mathcal{X}_q} \sum_{k \in \{f, c, b\}} \hat{s}_k(a, q) \log s_k(a, q), \quad (5)$$

where $\text{BCELoss}(\cdot)$ denotes the binary cross entropy, and η is the weight of multi-task learning.

Fitting Score-based Inference. During inference, evaluation of an anchor-query pair $\langle a, q \rangle$ should consider both Pathfinder’s path evaluation and Stopper’s level evaluation. However, instead of merely using S_p and S_c , the multi-classifying Stopper also enables the HEF model to disambiguate the most suited anchor from its neighbors (its direct parent and children) and self-correct its level prediction by exchanging scores with its neighbors to find the best position for maintaining the taxonomy’s coherence. Thus, We

Algorithm 1 Self-Supervised Training Process of HEF.

```
1: procedure TRAINEPOCH( $\mathcal{T}^0, \Theta^0$ )
2:    $\Theta \leftarrow \Theta^0$ 
3:   for  $q \leftarrow \mathcal{N}^0 - \text{root}(\mathcal{T}^0)$  do  $\triangleright$  Root is not used as query
4:      $\mathcal{X}_q = \{\}$   $\triangleright$  Initialize anchor set
5:      $p \leftarrow \text{parent}(q)$   $\triangleright$  Reference node of labeling
6:      $\mathcal{X}_q \leftarrow \mathcal{X}_q \cup \{p\}$   $\triangleright$  Ground Truth Parent:  $S_p = 1, S_c = 1$ 
7:      $\mathcal{X}_q \leftarrow \mathcal{X}_q \cup \text{sample}(\mathcal{A}_p)$   $\triangleright$  Ancestors:  $S_p = 1, S_f = 1$ 
8:      $\mathcal{X}_q \leftarrow \mathcal{X}_q \cup \text{sample}(\mathcal{D}_p)$   $\triangleright$  Descendants:  $S_p = 1, S_b = 1$ 
9:      $\mathcal{X}_q \leftarrow \mathcal{X}_q \cup \text{sample}(\mathcal{N}^0 - \{p\} - \mathcal{A}_p - \mathcal{D}_p)$ 
10:     $\triangleright$  Other nodes:  $S_p = 0, S_b = 1$ 
11:    for  $a \leftarrow \mathcal{X}_q$  do
12:      Compute  $S_p(a, q)$  using Eqn. 3
13:      Compute  $S_f(a, q), S_c(a, q), S_b(a, q)$  using Eqn. 4
14:    end for
15:    Compute  $\mathcal{L}_q$  with  $S_p, S_f, S_c, S_b$  using Eqn. 5
16:     $\Theta \leftarrow \text{optimize}(\Theta, \mathcal{L}_q)$ 
17:  end for
18:  return  $\Theta$ 
19: end procedure
```

introduce the *Fitting Score* function during inference. For a new query term $q \in C$, we first obtain the Pathfinder Scores and Stopper Scores of all node pairs $\langle a, q \rangle, a \in \mathcal{N}^0$. For each anchor node a , we assign its Fitting Score by multiplying the following four items:

- **a 's Pathfinder Score:** $S_p(a, q)$, which suggests whether a is on the right path.
- **a 's parent's Forward Score:** $S_f(\text{parent}(a), q)$, which distinguishes a and a 's parent, and rectifies a 's Current Score. When a is the root node, we assign this item as a small number like $1e-4$ since the first level of taxonomy is likely to remain unchanged.
- **a 's Current Score:** $S_c(a, q)$, which suggests whether a is on the right level.
- **a 's child with maximum Pathfinder Score's Backward Score:** $S_b(c_a^*, q), c_a^* = \arg \max_{c_a \in \text{child}(a)} S_p(c_a, q)$, which distinguishes a and a 's children, and rectifies a 's Current Score. Since a might have multiple children, we pick the child with max Pathfinder Score, for larger S_p indicates a better hypernymy relation to q . When a is a leaf node, we assign this item as the proportion of leaf nodes in the seed taxonomy to keep its overall design.

The Fitting Score of $\langle a, q \rangle$ is given by:

$$F(a, q) = S_p(a, q) \cdot S_f(\text{parent}(a), q) \cdot S_c(a, q) \cdot S_b(c_a^*, q) \quad (6)$$
$$c_a^* = \arg \max_{c_a \in \text{child}(a)} S_p(c_a, q).$$

The Fitting Score can be computed using ordered S_p, S_f, S_c, S_b arrays and the seed taxonomy's adjacency matrix. Since a tree's adjacency matrix is sparse, the time complexity of Fitting Score computation is low. After calculating the Fitting Scores between all seed nodes and the query, we select the seed node with the highest Fitting Score as the query's parent in the expanded taxonomy:

$$\text{parent}(q) := \arg \max_{a \in \mathcal{N}^0} F(a, q). \quad (7)$$

Table 1: Statistics of datasets. $|N|$ and $|E_O|$ are the numbers of nodes and edges in the original datasets, respectively. D is the depth of the taxonomy. We adopt the spanning tree of each dataset, and $|E|$ is the number of remaining edges.

Dataset	$ N $	$ E_O $	$ E $	D
<i>SemEval16-Env</i>	261	261	260	6
<i>SemEval16-Sci</i>	429	452	428	8
<i>SemEval16-Food</i>	1486	1576	1485	8

5 EXPERIMENTS

In this section, we first introduce our experimental setups, including datasets, our implementation details, evaluation criteria, and a brief description of the compared baseline methods. Then, we provide extensive evaluation results for overall model performance, performance contribution brought by each design, and sensitivity analysis of the multi-task learning weight η in Equation 5. In-depth visualizations of hypernymy detection and coherence modeling modules are provided to analyze the model's inner behavior. We also provide a case study in the appendix.¹

5.1 Experimental Setup

5.1.1 Datasets. We evaluate HEF on three public benchmark datasets retrieved from SemEval-2016 Task 13 [2]. This task contains three taxonomies in the domain of Environment (*SemEval16-Env*), Science (*SemEval16-Sci*), and Food (*SemEval16-Food*), respectively. The statistics of the benchmark datasets are provided in Table 1. Note that the original dataset may not form a tree. In this case, we use a spanning tree of the taxonomy instead of the original graph to match the problem definition. The pruning process only removes less than 6% of the total edges, keeping the taxonomy's information and avoiding multiple ground truth parents for a single node.

Since HEF and the compared baselines [34, 47] are all limited to adding new terms without modifying the seed taxonomy, nodes in the test and validation set can only sample from leaf nodes to guarantee that the parents of test or validation nodes exist in the seed taxonomy. This is also the sampling strategy of TaxoExpan [34]. Following the previous state-of-the-art model STEAM [47], we exclude 20% of the nodes in each dataset, of which ten nodes of each dataset are separated as the validation set for early stopping, and the rest as the test set. The nodes not included in the validation set and test set are seed nodes for self-supervision in the training phase and potential anchor nodes in the inference phase. Note that pruning the dataset does not affect the node count, thus the scale of the dataset remains identical to our baselines' settings.

5.1.2 Baselines for Comparison. We compare our proposed HEF model with the following baseline approaches:

- **BERT+MLP:** This method utilizes BERT [5] to perform hypernym detection. This model's input is the term's surface name, and the representation of BERT's classification token $\langle CLS \rangle$ is fed into a feed-forward layer to score whether the first sequence is the ground-truth parent.

¹The code will be available at <https://github.com/sheryc/HEF>.

Table 2: Comparison of the proposed method against state-of-the-art methods. All metrics are presented in percentages (%). Best results for each metric of each dataset are marked in bold. Reported performance is the average of three runs using different random seeds. The MRR of TAXI [26] is inaccessible since it outputs the whole taxonomy instead of node rankings. The performance of baseline methods are retrieved from [47].

Dataset Metric	<i>SemEval16-Env</i>			<i>SemEval16-Sci</i>			<i>SemEval16-Food</i>		
	Acc	MRR	Wu&P	Acc	MRR	Wu&P	Acc	MRR	Wu&P
BERT+MLP	11.1	21.5	47.9	11.5	15.7	43.6	10.5	14.9	47.0
TAXI [26]	16.7	-	44.7	13.0	-	32.9	18.2	-	39.2
HypeNet [37]	16.7	23.7	55.8	15.4	22.6	50.7	20.5	27.3	63.2
TaxoExpan [34]	11.1	32.3	54.8	27.8	44.8	57.6	27.6	40.5	54.2
STEAM [47]	36.1	46.9	69.6	36.5	48.3	68.2	34.2	43.4	67.0
HEF	55.3	65.3	71.4	53.6	62.7	75.6	47.9	55.5	73.5

- **HypeNet** [37]: HypeNet is an LSTM-based hypernym extraction model that scores a term pair by representing node paths in the dependency tree.
- **TAXI** [26]: TAXI was the top solution of SemEval-2016 Task 13. It explicitly splits the task into a pipeline of hypernym detection using substring matching and pattern extraction, and hypernym pruning to avoid multiple parents.
- **TaxoExpan** [34]: TaxoExpan is a self-supervised taxonomy expansion model. The anchor’s representation is modeled by a graph network of its Egonet with consideration of relative levels, and the parental relationship is scored by a feed-forward layer. BERT embedding is used as its input instead of the model’s original configuration.
- **STEAM** [47]: STEAM is the state-of-the-art self-supervised taxonomy expansion model, which scores parental relations by ensembling three classifiers considering graph, contextual, and hand-crafted lexical-syntactic features, respectively.

5.1.3 Implementation Details. In our setting, the coherence modeling module is a 3-layer, 6-head, 768-dimensional Transformer encoder initialized from Gaussian distribution $\mathcal{N}(0, 0.02)$. The first hidden layers of Pathfinder and Stopper are both 300-dimensional. The input to the hypernymy detection module is either truncated or padded to a length of 64. Each training step contains a set of 32 query-ego-tree pairs of 32 query nodes using gradient accumulation, with each query-ego-tree pair set containing one ground-truth parent ($S_p = 1, S_c = 1$), at most 6 ground-truth parent’s ancestors ($S_p = 1, S_f = 1$), at most 8 ground-truth parent’s descendants ($S_p = 1, S_b = 1$), and at least 16 other nodes ($S_p = 0, S_b = 1$). The hyperparameters above are empirically set, since our algorithm is not sensitive to the setting of splits. Each dataset is trained for 150 epochs. In a single epoch, each seed node is trained as the query exactly once. AdamW [21] is used for optimization with ϵ set to 1×10^{-6} . A linear warm-up is adopted with the learning rate linearly rise from 0 to 5e-5 in the first 10% of total training steps and linearly drop to 0 at the end of 150 epochs. The multi-task learning weight η is set to 0.9. After each epoch, we validate the model and save the model with the best performance on the validation set. These hyperparameters are tuned on on *SemEval16-Env*’s validation set, and are used across all datasets and experiments unless specified in the ablation studies or sensitivity analysis.

5.1.4 Evaluation Metrics. Assume $k := |C|$ to be the term count of the test set, $\{p_1, p_2, \dots, p_k\}$ to be the predicted parents for test set queries, and $\{\hat{p}_1, \hat{p}_2, \dots, \hat{p}_k\}$ to be the ground truth parents accordingly. Following the previous solutions [23, 40, 47], we adopt the following three metrics as evaluation criteria.

- **Accuracy (Acc):** It measures the proportion that the predicted parent for each node in the test set exactly matches the ground truth parent:

$$\text{Acc} = \text{Hit@1} = \frac{1}{k} \sum_{i=1}^k \mathbb{I}(p_i = \hat{p}_i),$$

where $\mathbb{I}(\cdot)$ denotes the indicator function,

- **Mean Reciprocal Rank (MRR):** It calculates the average reciprocal rank of each test node’s ground truth parent:

$$\text{MRR} = \frac{1}{k} \sum_{i=1}^k \frac{1}{\text{rank}(\hat{p}_i)},$$

- **Wu & Palmer Similarity (Wu&P)** [44]: It is a tree-based measurement that judges how close the predicted and ground truth parents are in the seed taxonomy:

$$\text{Wu\&P} = \frac{1}{k} \sum_{i=1}^k \frac{2 \times \text{depth}(\text{LCA}(p_i, \hat{p}_i))}{\text{depth}(p_i) + \text{depth}(\hat{p}_i)},$$

where $\text{depth}(\cdot)$ is the node’s depth in the seed taxonomy and $\text{LCA}(\cdot, \cdot)$ is the least common ancestor of two nodes.

5.2 Main Results

The performance of HEF is shown in Table 2. HEF achieves the best performance on all datasets and surpasses previous state-of-the-art models with a significant improvement on all metrics.

From the table, we get an overview of how taxonomy expansion models evolve chronologically. The solution of BERT+MLP does not utilize any structural and lexical-syntactic features of terms, and the insufficiency of information attributes to its poor results. Models of the first generation like TAXI and HypeNet utilize lexical, syntactic, or contextual information to achieve better results, mainly for the task’s hypernymy detection part. However, these two models do not utilize any of the structural information of taxonomy; hence they are unable to maintain the taxonomy’s structural design. Models of

Table 3: Ablation experiment results on the *SemEval16-Env* dataset. All metrics are presented in percentages (%). For each ablation experiment setting, only the best result is reported.

Abl. Type	Setting	Acc	MRR	Wu&P
Original	HEF	55.3	65.3	71.4
Dataflows	- WordNet Descriptions	41.5	55.3	62.6
	- Ego-tree + Egonet	45.3	60.6	69.9
	- Relative Level Emb.	49.1	59.2	60.9
	- Absolute Level Emb.	49.1	60.6	68.4
Scoring Function	Stopper Only	52.8	62.5	68.7
	Pathfinder + Current Only	50.9	62.1	66.8
	Current Only	41.5	54.7	58.6

the second generation, like TaxoExpan and STEAM, inherit hand-crafted lexical-syntactic features for detecting hypernymy relations. They also utilize the structural information by self-supervision from seed taxonomy and graph neural networks on small subgraphs of the taxonomy. However, they neglect the hierarchical structure of taxonomies, and they do not consider the coherence of the whole expanded taxonomy. Thus, their usage of structural information is only an improvement for performing hypernymy detection rather than taxonomy expansion.

HEF further improves both the previous two generations’ strength by proposing a new approach that better fits the taxonomy expansion task. Moreover, it introduces a new goal for the task: to best preserve the taxonomy’s coherence after expansion. We propose the description generation algorithm to generate accurate and domain-specific descriptions for complex and rare terms, to incorporate lexical-syntactic features for hypernymy detection. Aided by DistilBERT’s power of sentence-pair representation, HEF can mine hypernymy features more automatically and accurately. HEF also aims to fully exploit the information brought by the taxonomy’s hierarchical structure to boost performance. HEF uses ego-trees to perform thorough comparison along root path and among siblings, injects tree-exclusive features to assist modeling the expert-curated taxonomy designs and explicitly evaluates both path and level for the anchor node as well as its parent and child. Experiment results suggest that these designs are capable of modeling and maximizing the coherence of taxonomy in different aspects, which results in a vast performance increase in the taxonomy expansion task.

5.3 Ablation Studies

We discuss how exploiting different characteristics of taxonomy’s hierarchical structure brings performance increase by a series of ablation studies. We substitute some designs of HEF in dataflow and score function to a vanilla setting and rerun the experiments. The results of the ablation studies are shown in Table 3.

- **- WordNet Descriptions:** WordNet descriptions are substituted with the term’s surface name as the hypernymy detection module’s input.
- **- Ego-tree + Egonet:** the Egonet from TaxoExpan [34] is used instead of the ego-tree for modeling the tree structure.

- **- Relative Level Emb.:** The relative level embedding for the coherence modeling module is removed.
- **- Absolute Level Emb.:** The absolute level embedding for the coherence modeling module is removed.
- **Stopper Only:** Only the Stopper Scores are used for Fitting Score calculation. More specifically, $\eta = 0$, and the Fitting Score in Equation 6 becomes:

$$F(a, q) = S_f(\text{parent}(a), q) \cdot S_c(a, q) \cdot S_b(c_a^*, q),$$

$$c_a^* = \arg \max_{c_a \in \text{child}(a)} S_p(c_a, q).$$

- **Pathfinder + Current Only:** Only the Pathfinder Score and the Current Score are used for Fitting Score calculation. More specifically, the Fitting Score in Equation 6 and the loss in Equation 5 become:

$$F(a, q) = S_p(a, q) \cdot S_c(a, q),$$

$$\mathcal{L}_q = -\eta \frac{1}{|\mathcal{X}_q|} \sum_{a \in \mathcal{X}_q} \text{BCELoss}(\hat{S}_p(a, q), S_p(a, q))$$

$$- (1 - \eta) \frac{1}{|\mathcal{X}_q|} \sum_{a \in \mathcal{X}_q} \text{BCELoss}(\hat{S}_c(a, q), S_c(a, q)).$$

- **Current Only:** Only the Current Score is used for Fitting Score calculation. This is the scoring strategy identical to prior arts [34, 47]. More specifically, the Fitting Score in Equation 6 and the loss in Equation 5 become:

$$F(a, q) = S_c(a, q),$$

$$\mathcal{L}_q = -\frac{1}{|\mathcal{X}_q|} \sum_{a \in \mathcal{X}_q} \text{BCELoss}(\hat{S}_c(a, q), S_c(a, q)).$$

We notice that by changing the design of dataflows, the performance of the HEF model suffers from various deteriorations. Substituting WordNet descriptions with a term’s surface name surprisingly remains a relatively high performance, which might attribute to the representation power of the DistilBERT model. Using Egonets rather than ego-trees for coherence modeling also affects the performance. Although egonets can capture the local structure of taxonomy, ego-trees are more capable of modeling the complete construction of a hierarchy. For the introduction of level embeddings, the results show that removing one of the two level embeddings for the coherence modeling module hurts the learning of taxonomy’s design. This is in accordance with the previous research about the importance of using the information of both absolute and relative positions in Transformers [33] and confirms our assumption that taxonomies have intrinsic designs about both absolute and relative levels.

Changes to the score function bring a smaller negative impact on the model compared to the dataflow changes, except for the setting of using merely Current Score. When using only the Current Score, the model loses the ability to disambiguate with its neighbors and the capacity of directly choosing the right path, downgrading the problem to be a series of independent scoring problems like the previous solutions. Adding Backward Score and Forward Score into Fitting Score calculation allows the model to distinguish the ground truth from its neighbors, bringing a boost to accuracy. However, without the Pathfinder, the “Stopper Only” setting only explicitly

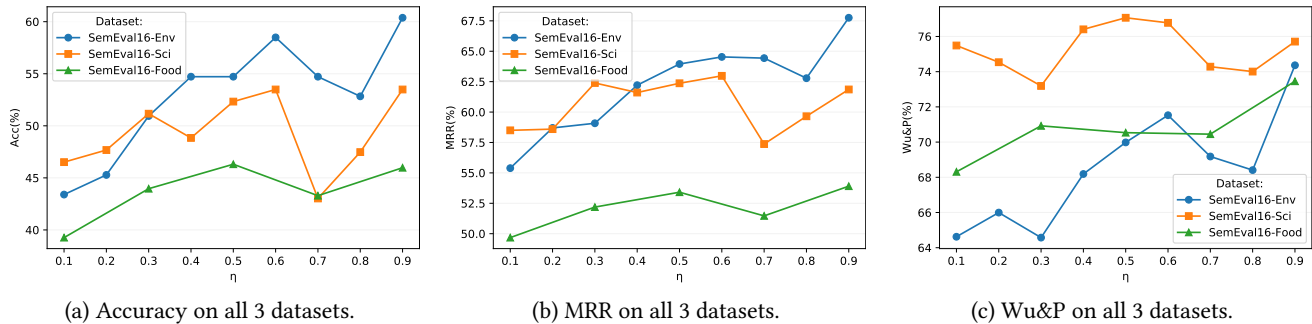


Figure 4: Sensitivity analysis of model performance under different multi-task learning weight η .

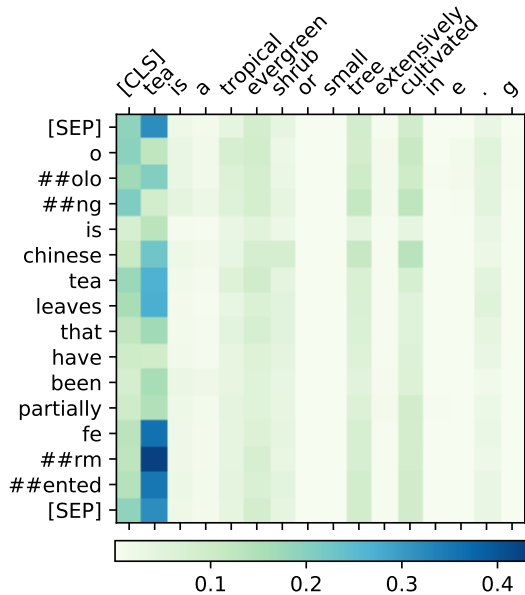


Figure 5: Illustration of one self-attention head in the last layer of the hypernymy detection module, showing how the hypernymy detection module detects hypernymy relations. In this example, the seed node is "tea", and the query is "oolong".

focuses on choosing the right level without considering the path and is inferior to the original HEF model.

However, we observe that although changing several designs of dataflow or scoring function deteriorates the performance, our method can still surpass the previous state-of-the-art in Acc and MRR, suggesting that the HEF model introduces improvements in multiple aspects, which also testifies that maximizing the taxonomy’s coherence is a better goal for the taxonomy expansion task.

5.4 Impact of Multi-Task Learning Weight η

In this section, we discuss the impact of η in Equation 5 through a sensitivity analysis. Since η controls the proportion of loss from the path-oriented Pathfinder and the level-oriented Stopper, this

hyperparameter affects HEF’s tendency to prioritize path or level selection. The results on all three datasets are shown in Fig. 4.

From the result, we notice that η cannot be set too low, which means that explicit path selection contributes a lot to the model’s performance. This is in accordance with the fact that taxonomies are based on hypernymy relations and selecting the right path is the essential guidance for anchor selection. A better setting of η is [0.4, 0.6]. In this setting, the model tends to balance path and level selections, which results in better performance. Surprisingly, setting η to a high value like 0.9 also brings a performance boost, and sometimes even achieves the best result. This phenomenon consistently exists when changing random seeds. However, setting η to 1 means using merely Pathfinder, which cannot distinguish the ground truth from other nodes and breaks the model. This discovery further testifies the importance of explicitly evaluating path selection in the taxonomy expansion task.

5.5 Visualization of Self-Attentions

5.5.1 Node Pair Hypernymy Detection Module. To illustrate how the hypernymy detection module works, we show the weights of one of the attention heads of the hypernymy detection module’s last Transformer encoder layer in Fig. 5.

By expanding a term to its description, the model is capable of understanding the term "oolong" by its description, which cannot be achieved by constructing rule-based lexical-syntactic features since "oolong" and "tea" have no similarity in their surface names. Furthermore, by adopting the pretrained DistilBERT, the hypernymy detection module can also discover more latent patterns such as the relation between "leaves" and "tree", allowing the model to discover more in-depth hypernymy relations.

5.5.2 Ego-Tree Coherence Modeling Module. To illustrate how the coherence modeling module compares the nodes in the ego-tree to maintain the taxonomy’s coherence, we present the weights of an attention head of the module’s first Transformer encoder layer in Fig. 6. Since the last layer’s attention mostly focuses on the anchor node ("herb"), the first layer can better illustrate the model’s comparison among ego-tree nodes.

Based on our observation, the two <CLS>s are capable of finding the best-suited parent node in the ego-tree even if it is not the anchor. Since the coherence modeling module utilizes ego-trees for hierarchy modeling, the coherence modeling module can compare

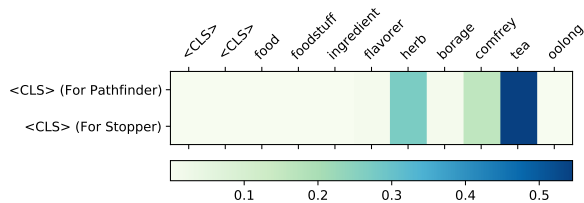


Figure 6: Illustration of one self-attention head in the first layer of the coherence modeling module, showing how the coherence modeling module finds the most fitted node in an ego-tree. In this example, the anchor is “herb”, the query is “oolong”, and the query’s ground truth parent is “tea”.

a node with all its ancestors and its children to find the most suited anchor, which makes the model more robust. Besides, the coherence modeling module is also able to assign a lower attention weight to the best-suited parent’s parent node when it is on the right path, suggesting that the coherence modeling module can achieve both path-wise and level-wise comparison.

6 CONCLUSION

We proposed HEF, a self-supervised taxonomy expansion model that fully exploits the hierarchical structure of a taxonomy for better hierarchical structure modeling and taxonomy coherence maintenance. Compared to previous methods that evaluate the anchor by merely a new edge in a normal graph neglecting the tree structure of taxonomy, we used extensive experiments to prove that, evaluating a tree structure for coherence maintenance, and mining multiple tree-exclusive features in the taxonomy, including hypernymy relations from parent-child relations, term similarity from sibling relations, absolute and relative levels, path+level based multi-dimensional evaluation, and disambiguation based on parent-current-child chains all brought performance boost. This indicates the importance of using the information of tree structure for the taxonomy expansion task. We also proposed a framework for injecting these features, introduced our implementation of the framework, and surpassed the previous state-of-the-art. We believe that these novel designs and their motivations will not only benefit the taxonomy expansion task, but also be influential for all tasks involving hierarchical or tree structure modeling and evaluation. Future works include how to model and utilize these or new tree-exclusive features to boost other taxonomy-related tasks, and better implementation of each module in HEF.

ACKNOWLEDGMENTS

Thanks to everyone who helped me with this paper in the Tencent Jarvis Lab, my family, and my loved one.

REFERENCES

- [1] Eugene Agichtein and Luis Gravano. 2000. Snowball: Extracting Relations From Large Plain-Text Collections. In *Proceedings of JCDL*. 85–94.
- [2] Georgeta Bordea, Els Lefever, and Paul Buitelaar. 2016. SemEval-2016 Task 13: Taxonomy Extraction Evaluation (TExEval-2). In *Proceedings of the SemEval-2016*. 1081–1091.
- [3] Anne Cocos, Marianna Apidianaki, and Chris Callison-Burch. 2018. Comparing Constraints for Taxonomic Organization. In *Proceedings of NAACL*. 323–333.

- [4] Sarthak Dash, Md Faisal Mahub Chowdhury, Alfio Gliozzo, Nandana Mihindukulasooriya, and Nicolas Rodolfo Faceglia. 2020. Hypernym Detection Using Strict Partial Order Networks. In *Proceedings of AAAI*. 7626–7633.
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL*. 4171–4186.
- [6] Trevor Fountain and Mirella Lapata. 2012. Taxonomy Induction Using Hierarchical Random Graphs. In *Proceedings of NAACL*. 466–476.
- [7] Amit Gupta, Rémi Lebret, Hamza Harkous, and Karl Aberer. 2017. Taxonomy Induction Using Hypernym Subsequences. In *Proceedings of CIKM*. 1329–1338.
- [8] Marti A. Hearst. 1992. Automatic Acquisition of Hyponyms From Large Text Corpora. In *Proceedings of ACL*. 539–545.
- [9] Wen Hua, Zhongyuan Wang, Haixun Wang, Kai Zheng, and Xiaofang Zhou. 2017. Understand Short Texts by Harvesting and Analyzing Semantic Knowledge. *IEEE Transactions on Knowledge and Data Engineering* (2017), 499–512.
- [10] Jin Huang, Zhaochun Ren, Wayne Xin Zhao, Gaole He, Ji-Rong Wen, and Daxiang Dong. 2019. Taxonomy-Aware Multi-Hop Reasoning Networks for Sequential Recommendation. In *Proceedings of WSDM*. 573–581.
- [11] 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 KDD*. 877–886.
- [12] David Jurgens and Mohammad Taher Pilehvar. 2015. Reserating the Awesometastic: An Automatic Extension of the WordNet Taxonomy for Novel Terms. In *Proceedings of NAACL*. 1459–1465.
- [13] David Jurgens and Mohammad Taher Pilehvar. 2016. SemEval-2016 Task 14: Semantic Taxonomy Enrichment. In *Proceedings of the SemEval-2016*. 1092–1102.
- [14] Giannis Karamanolakis, Jun Ma, and Xin Luna Dong. 2020. TXtract: Taxonomy-Aware Knowledge Extraction for Thousands of Product Categories. (2020).
- [15] Seohyun Kim, Jinman Zhao, Yuchi Tian, and Satish Chandra. 2020. Code Prediction by Feeding Trees to Transformers. (2020).
- [16] Zornitsa Kozareva and Eduard Hovy. 2010. A Semi-Supervised Method to Learn and Construct Taxonomies Using the Web. In *Proceedings of EMNLP*. 1110–1118.
- [17] Matthew Le, Stephen Roller, Laetitia Papaxanthos, Douwe Kiela, and Maximilian Nickel. 2019. Inferring Concept Hierarchies From Text Corpora via Hyperbolic Embeddings. In *Proceedings of ACL*. 3231–3241.
- [18] Dekang Lin. 1998. An Information-Theoretic Definition of Similarity. In *Proceedings of ICML*. 296–304.
- [19] Carolyn E. Lipscomb. 2000. Medical Subject Headings (MeSH). *Bulletin of the Medical Library Association* (2000), 265–266.
- [20] Bang Liu, Weidong Guo, Di Niu, Chaoyue Wang, Shunnan Xu, Jinghong Lin, Kunfeng Lai, and Yu Xu. 2019. A User-Centered Concept Mining System for Query and Document Understanding at Tencent. In *Proceedings of KDD*. 1831–1841.
- [21] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. (2019).
- [22] Xusheng Luo, Luxin Liu, Yonghua Yang, Le Bo, Yuanpeng Cao, Jinghang Wu, Qiang Li, Keping Yang, and Kenny Q. Zhu. 2020. AliCoCo: Alibaba E-Commerce Cognitive Concept Net. In *Proceedings of SIGMOD*. 313–327.
- [23] Emaad Manzoor, Rui Li, Dhananjay Shrivastava, and Jure Leskovec. 2020. Expanding Taxonomies With Implicit Edge Semantics. In *Proceedings of TheWebConf*. 2044–2054.
- [24] Yuning Mao, Xiang Ren, Jiaming Shen, Xiaotao Gu, and Jiawei Han. 2018. End-To-End Reinforcement Learning for Automatic Taxonomy Induction. In *Proceedings of ACL*. 2462–2472.
- [25] George A. Miller. 1995. WordNet: A Lexical Database for English. *Commun. ACM* (1995), 39–41.
- [26] Alexander Panchenko, Stefano Faralli, Eugen Ruppert, Steffen Remus, Hubert Naets, Cédric Faron, 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 SemEval-2016*. 1320–1327.
- [27] Hao Peng, Jianxin Li, Senzhang Wang, Lihong Wang, Qiran Gong, Renyu Yang, Bo Li, Philip Yu, and Lifang He. 2019. Hierarchical Taxonomy-Aware and Attentional Graph Capsule RCNNs for Large-Scale Multi-Label Text Classification. *IEEE Transactions on Knowledge and Data Engineering* (2019).
- [28] Stephen Roller, Douwe Kiela, and Maximilian Nickel. 2018. Hearst Patterns Revisited: Automatic Hypernym Detection From Large Text Corpora. In *Proceedings of ACL*. 358–363.
- [29] Erik Tjong Kim Sang. 2007. Extracting Hypernym Pairs From the Web. In *Proceedings of ACL*. Association for Computational Linguistics, 165–168.
- [30] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. (2020).
- [31] Chao Shang, Sarthak Dash, Md. Faisal Mahub Chowdhury, Nandana Mihindukulasooriya, and Alfio Gliozzo. 2020. Taxonomy Construction of Unseen Domains via Graph-Based Cross-Domain Knowledge Transfer. In *Proceedings of ACL*. 2198–2208.
- [32] Jingbo Shang, Xinyang Zhang, Liyuan Liu, Sha Li, and Jiawei Han. 2020. Net-Taxo: Automated Topic Taxonomy Construction From Text-Rich Network. In *Proceedings of TheWebConf*. 1908–1919.

- [33] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. Self-Attention With Relative Position Representations. In *Proceedings of NAACL*. 464–468.
- [34] Jiaming Shen, Zhihong Shen, Chenyan Xiong, Chi Wang, Kuansan Wang, and Jiawei Han. 2020. TaxoExpan: Self-Supervised Taxonomy Expansion With Position-Enhanced Graph Neural Network. In *Proceedings of TheWebConf*. 486–497.
- [35] Jiaming Shen, Zeqiu Wu, Dongming Lei, Chao Zhang, Xiang Ren, Michelle T. Vanni, Brian M. Sadler, and Jiawei Han. 2018. HiExpan: Task-Guided Taxonomy Construction by Hierarchical Tree Expansion. In *Proceedings of KDD*. 2180–2189.
- [36] Vighnesh Shiv and Chris Quirk. 2019. Novel Positional Encodings to Enable Tree-Based Transformers. In *Advances in Neural Information Processing Systems* 32. 12081–12091.
- [37] Vered Shwartz, Yoav Goldberg, and Ido Dagan. 2016. Improving Hypernymy Detection With an Integrated Path-Based and Distributional Method. In *Proceedings of ACL*. Association for Computational Linguistics, 2389–2398.
- [38] Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks. In *Proceedings of ACL/IJCNLP*. 1556–1566.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Advances in Neural Information Processing Systems* 30. 5998–6008.
- [40] Nikhita Vedula, Patrick K. Nicholson, Deepak Ajwani, Sourav Dutta, Alessandra Sala, and Srinivasan Parthasarathy. 2018. Enriching Taxonomies With Functional Domain Knowledge. In *Proceedings of SIGIR*. 745–754.
- [41] Paola Velardi, Stefano Faralli, and Roberto Navigli. 2013. OntoLearn Reloaded: A Graph-Based Algorithm for Taxonomy Induction. *Computational Linguistics* (2013), 665–707.
- [42] 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 *Proceedings of WWW*. 1965–1976.
- [43] Jinghui Wang, Changsung Kang, Yi Chang, and Jiawei Han. 2014. A Hierarchical Dirichlet Model for Taxonomy Expansion for Search Engines. In *Proceedings of WWW*. 961–970.
- [44] Zhibiao Wu and Martha Palmer. 1994. Verbs Semantics and Lexical Selection. In *Proceedings of ACL*. 133–138.
- [45] Wenpeng Yin and Dan Roth. 2018. Term Definitions Help Hypernymy Detection. In *Proceedings of SEM*. 203–213.
- [46] Xiaoxin Yin and Sarthak Shah. 2010. Building Taxonomy of Web Search Intents for Name Entity Queries. In *Proceedings of WWW*. 1001–1010.
- [47] Yue Yu, Yinghao Li, Jiaming Shen, Hao Feng, Jimeng Sun, and Chao Zhang. 2020. STEAM: Self-Supervised Taxonomy Expansion With Mini-Paths. In *Proceedings of KDD*. 1026–1035.
- [48] Chao Zhang, Fangbo Tao, Xiuxi Chen, Jiaming Shen, Meng Jiang, Brian Sadler, Michelle Vanni, and Jiawei Han. 2018. TaxoGen: Unsupervised Topic Taxonomy Construction by Adaptive Term Embedding and Clustering. In *Proceedings of KDD*. 2701–2709.

A CASE STUDY

To understand how different Fitting Score components contribute to HEF’s performance, we conduct a case study on the *SemEval16-Food* dataset and show the detailed results in Table 4.

The first two rows of Table 4 shows two cases where HEF successfully predicts the query’s parent. We can see that the Pathfinder Score and the three Stopper Scores all contribute to the correct selection, which testifies the effectiveness of the Fitting Score design.

The last two rows of Table 4 provide situations when HEF fails to select the correct parent. In the third row, “*bourguignon*” is described as “reduced red wine”, thus the model attaches it to the node “*wine*”. However, “*bourguignon*” is also a sauce for cooking beef. Such ambiguity consequently affects the meaning of a term by assigning an incorrect description, which hurts the model’s performance. In the last row, although “*hot fudge sauce*”’s description contains “chocolate sauce”, the node “*chocolate sauce*” still gets a low Current Score. In HEF, the design of Stopper Scores enables the model to self-correct the occasionally wrong Current Scores by assigning larger Forward Score from a node’s parent and larger Backward Score from one of the node’s children. However, since “*chocolate sauce*” is a leaf node, its child’s Backward Score is assigned to be the proportion of leaf nodes in the seed taxonomy, which is

0.07 in the *SemEval16-Food* dataset. This indicates that future work includes designing a more reasonable Backward Score function for leaf nodes to improve the model’s robustness.

B DESCRIPTION GENERATION ALGORITHM

Algorithm 2 shows the description generation algorithm `descr(·)` used in HEF’s hypernymy detection module. `descr(·)` utilizes WordNet descriptions to generate domain-related term descriptions by dynamic programming. In this algorithm, `WordNetNounDescr(·)` means the set of a concept’s noun descriptions from WordNet [25], and `CosSimilarity(t, nroot)` means calculating the average token cosine similarity of word vectors between a candidate description t and the surface name of a taxonomy’s root term n_{root} .

Algorithm 2 Description Generation Algorithm for the Hypernymy Detection Module

```

1: procedure DESCR( $n$ )                                ▶ Input: term  $n$ 
2:    $N \leftarrow \text{split}(n)$ 
3:   for  $i \leftarrow 0, \dots, \text{length}(N)$  do
4:      $S[i] = 0$                                         ▶ Initialize score array
5:      $C[i] = 0$                                         ▶ Initialize splitting positions
6:   end for
7:   for  $i \leftarrow 0, \dots, \text{length}(N) - 1$  do
8:     for  $j \leftarrow 0, \dots, i$  do
9:       if WordNetNounDescr(N[j : i + 1]) > 0 then
10:         $s_{ij} = (i - j + 1)^2 + 1$  ▶ Prefer longer concepts
11:      else
12:         $s_{ij} = 1$ 
13:      end if
14:      if  $S[j] + s_{ij} > S[i + 1]$  then
15:         $S[i + 1] \leftarrow S[j] + s_{ij}$  ▶ Save max score
16:         $C[i] = j$  ▶ Save splitting position
17:      end if
18:    end for
19:  end for
20:   $D \leftarrow \text{“”}$  ▶ Initialize description
21:   $p \leftarrow \text{length}(N)$  ▶ Generate split pointer
22:  while  $p \neq -1$  do
23:     $D_{WN} = \text{WordNetNounDescr}(N[C[p] : p + 1])$ 
24:    if len(DWN) > 0 then ▶ Noun or noun phrase
25:       $d \leftarrow \arg \max_{t \in D_{WN}} \text{CosSimilarity}(t, n_{root})$ 
26:    else ▶ Prep. or adj.
27:       $d \leftarrow \text{join}(N[C[p] : p + 1])$ 
28:    end if
29:     $D \leftarrow d + D$  ▶ Put new description in the front
30:     $p \leftarrow C[p] - 1$  ▶ Go to next split
31:  end while
32: end procedure

```

Table 4: Examples of HEF’s prediction, with detailed Fitting Score composition and comparison between the ground truth and the predicted parent. Scores in this table correspond to the node in the same tabular cell with the score.

Query (q)	Ground Truth (\hat{p})	Scores	Prediction (p)	Scores
q : paddy is rice in the husk either gathered or still in the field	\hat{p} : rice is grains used as food either unpolished or more often polished	$S_p = 0.9997$ $S_c = 0.4599$	p : rice is grains used as food either unpolished or more often polished	$S_p = 0.9997$ $S_c = 0.4599$
	parent(\hat{p}): starches is a commercial preparation of starch that is used to stiffen textile fabrics in laundering	$S_f = 0.9755$	parent(p): starches is a commercial preparation of starch that is used to stiffen textile fabrics in laundering	$S_f = 0.9755$
$F(\hat{p}, q) = 0.4483$ \hat{p} 's Ranking: 1	$c_{\hat{p}}^*$: white rice is having husk or outer brown layers removed	$S_b = 0.9995$	c_p^* : white rice is having husk or outer brown layers removed	$S_b = 0.9995$
q : fish meal is ground dried fish used as fertilizer and as feed for domestic livestock	\hat{p} : feed is food for domestic livestock	$S_p = 0.9993$ $S_c = 0.3169$	p : feed is food for domestic livestock	$S_p = 0.9993$ $S_c = 0.3169$
	parent(\hat{p}): food is any substance that can be metabolized by an animal to give energy and build tissue	$S_f = 0.9984$	parent(p): food is any substance that can be metabolized by an animal to give energy and build tissue	$S_f = 0.9984$
$F(\hat{p}, q) = 0.3158$ \hat{p} 's Ranking: 1	$c_{\hat{p}}^*$: mash is mixture of ground animal feeds	$S_b = 0.9988$	c_p^* : mash is mixture of ground animal feeds	$S_b = 0.9988$
q : bourguignon is reduced red wine with onions and parsley and thyme and butter	\hat{p} : sauce is flavorful relish or dressing or topping served as an accompaniment...	$S_p = 0.0002$ $S_c = 0.0001$	p : wine is a red as dark as red wine	$S_p = 0.9997$ $S_c = 0.1399$
	parent(\hat{p}): condiment is a preparation (a sauce or relish or spice) to enhance flavor or enjoyment	$S_f = 0.0004$	parent(p): alcohol is any of a series of volatile hydroxyl compounds that are made from hydrocarbons by distillation	$S_f = 0.9812$
$F(\hat{p}, q) = 1e - 11$ \hat{p} 's Ranking: 328	$c_{\hat{p}}^*$: bercy is butter creamed with white wine and shallots and parsley	$S_b = 0.9997$	c_p^* : red wine is wine having a red color derived from skins of dark-colored grapes	$S_b = 0.8784$
q : hot fudge sauce is hot thick chocolate sauce served hot	\hat{p} : chocolate sauce is sauce made with unsweetened chocolate or cocoa...	$S_p = 0.9471$ $S_c = 9e - 5$	p : sauce is flavorful relish or dressing or topping served as an accompaniment...	$S_p = 0.9995$ $S_c = 0.0172$
	parent(\hat{p}): sauce is flavorful relish or dressing or topping served as an accompaniment...	$S_f = 0.9617$	parent(p): condiment is a preparation (a sauce or relish or spice) to enhance flavor or enjoyment	$S_f = 0.9888$
$F(\hat{p}, q) = 6e - 6$ \hat{p} 's Ranking: 20	$c_{\hat{p}}^*$: None	$S_b = 0.0700$	c_p^* : lyonnaise sauce is brown sauce with sauteed chopped onions and parsley...	$S_b = 0.9995$