

# TAG: Toward Accurate Social Media Content Tagging with a Concept Graph

Jiuding Yang\*

University of Alberta  
Edmonton, AB, Canada  
jiuding@ualberta.ca

Yakun Yu

University of Alberta  
Edmonton, AB, Canada  
yakun2@ualberta.ca

Linglong Kong

University of Alberta  
Edmonton, AB, Canada  
lkong@ualberta.ca

Weidong Guo\*

Platform and Content Group, Tencent  
Beijing, China  
weidongguo@tencent.com

Chaoyue Wang

Platform and Content Group, Tencent  
Beijing, China  
charieywang@tencent.com

Di Niu

University of Alberta  
Edmonton, AB, Canada  
dniu@ualberta.ca

Bang Liu\*†

Mila & DIRO, Université de Montréal  
Montréal, QC, Canada  
bang.liu@umontreal.ca

Jinwen Luo

Platform and Content Group, Tencent  
Beijing, China  
jamsluo@tencent.com

Zhen Wen

Platform and Content Group, Tencent  
Beijing, China  
zhenwen@tencent.com

## ABSTRACT

Although conceptualization has been widely studied in semantics and knowledge representation, it is still challenging to find the most accurate concept terms to tag fast-growing social media content. This is partly attributed to the fact that most traditional knowledge bases contain general terms of the world, such as trees and cars, which are not interesting to users, and do not have the defining power for social media content. Another reason is that the intricate use of tense, negation and grammar in social media content may change the logic or emphasis of the content, thus focusing on different main ideas. In this paper, we present TAG, a high-quality concept matching dataset consisting of 10,000 labeled pairs of fine-grained concepts and web-styled natural language sentences, mined from open-domain social media content. The concepts we provide are the trending terms on social media and have the right granularity to define user interests, e.g., highly educated actors instead of just actors. In the meantime, TAG offers a *concept graph* which interconnects these fine-grained concepts and entities to provide contextual information. We evaluate a wide range of neural text matching models as well as pre-trained language models for the concept matching task on TAG, and point out their insufficiency to tag social media content to characterize its main idea. We further propose a novel graph-graph matching framework that demonstrates superior abstraction and generalization performance

\*These authors contributed equally to this research.

†Canada CIFAR AI Chair.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

KDD '22, August 14–18, 2022, Washington, DC, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9385-0/22/08...\$15.00

<https://doi.org/10.1145/3534678.3539077>

by better utilizing both the structural information in the concept graph and logic interactions between semantic units in the natural language sentence via syntactic dependency parsing.

## CCS CONCEPTS

• **Computing methodologies** → **Information extraction.**

## KEYWORDS

datasets, concept-sentence matching, graph-graph matching, heterogeneous graph

### ACM Reference Format:

Jiuding Yang, Weidong Guo, Bang Liu, Yakun Yu, Chaoyue Wang, Jinwen Luo, Linglong Kong, Di Niu, and Zhen Wen. 2022. TAG: Toward Accurate Social Media Content Tagging with a Concept Graph. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22)*, August 14–18, 2022, Washington, DC, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3534678.3539077>

## 1 INTRODUCTION

Concepts are abstract ideas and notions about entities encountered in daily lives. While a common hypothesis in cognitive sciences, including linguistics, psychology and philosophy, is that all cognition must occur through concepts, their importance is also recognized by recent research in knowledge representation and natural language processing. Several recent efforts have focused on building a hierarchy, ontology or graph of concepts [20, 24, 46], which are key to document understanding [20, 21], search query understanding [17, 36], as well as recommender systems [20, 24].

Although many existing knowledge bases contain vast amounts of general notions for the world, e.g., trees, companies, cars, finding interesting concept terms that users would pay attention to among the exploding and fast-growing social media is becoming increasingly challenging, especially as fine-grained new concepts with defining powers are emerging on social media every day. For example, while people are familiar with notions like “best cars to



**Figure 1: An example of document retrieved by concept in the news feeds stream of QQ Browser.**

buy under 50k”, nowadays people search for “best electric cars under 50k”. Techniques have been developed to perform the query and the short text conceptualization [17, 37, 44] with topic models and probabilistic graphical models. However, these methods aim to tag short text or queries which usually consist of several words, and the focus is on word *disambiguation* in the limited context. For example, the query “most dangerous python in the world” implies the concept of a *snake* instead of a *programming language*. “Apple and iPad” implies a *firm* instead of *fruit*. Yet, what these techniques cannot solve is to find the defining concepts conforming to human intuition that just characterize the main idea behind the sentence. For example, “poisonous snakes” should be a more proper concept instead of just “snakes” in the above example.

Furthermore, social media content allows user-created content and usually implies more interesting information than what is available in Wikipedia or conventional knowledge bases, e.g., Probase [46], DBpedia [1], CNDBpedia [47], YAGO [39]. For example, a post about “Mr. Bean has a PhD degree” naturally reminds a human of *highly educated comedians* or *high IQ actors*, whereas a current knowledge base may only label it with “British actors” or “comedians”, which are correct but not matching the interests of social media viewers. Neither can they help improve the recommendation of similar items in mobile apps. The issue is that the “British” and “Comedian” parts of Mr. Bean are not what the sentence is trying to emphasize in the social media post. Unfortunately, terms like “highly educated actors”, although interesting to social media, cannot be found in most existing web-scale knowledge bases.

The Mr. Bean example points to another major issue of most existing knowledge representation methods. Although trying to find relevant concept tags, they do not attempt to understand the meaning of the natural language sentence and make sure the concept to be tagged conforms to the idea of the sentence. Taking the case in Figure 1 as an example, the text is talking about a recipe for cooking lotus root. The dish is delicious and can help you “get rid of hot-and-dry feeling in winter”, which means it is good for health.

However, a common pitfall is to label the content with “Delicious recipes”, only because it contains words like “crispy” and “fragrant” that are directly related to “Delicious recipes”, without respecting the meaning of this specific sentence. In fact, a correct tag to the document should be “Healthy recipes”, since the text emphasizes that the dish is healthy.

In this paper, we revisit text conceptualization beyond simple fact-based tagging, disambiguation and topic modeling, by focusing on finding the most accurate concept to characterize the main idea of a text snippet in open-domain social media. Such concept terms must be interesting to online viewers and are thus potentially highly diverse and involving social media slang. We ask the question—can machines comprehend the logical idea behind a piece of social media text and tag it with a *concept* that has just the right defining power and conforms to human intuition? According to Kant, “in order to make our mental images into concepts, one must thus be able to compare, reflect, and abstract.” (from Logic [16]) In other words, although it is easy to tag a sentence with relevant phrases, e.g., fruit, animal, trees, yet it is challenging to find the interesting concept at the right granularity that differentiates the sentence from other text mentioning the same entities.

Toward solving this challenge, in this paper, we introduce **TAG**, a new dataset consisting of 10,000 concept-sentence pairs obtained from Chinese social media content. Each pair has a binary label indicating whether a (fine-grained) concept term can characterize the logical idea that a sentence wants to convey. The ontology that supplies these concept terms in our dataset is a concept graph, containing social media concepts such as “best mobile games for girls”, “top 10 kungfu stars”, created from vast search queries and search logs in QQ Browser, a social media app that serves more than 200 million active users, providing functions such as search, news feeds, etc. The concept graph organizes a large number of candidate concepts as well as entities in a big graph, where a higher-level concept is termed *superordinate* and a lower-level concept termed *subordinate*. In other words, different concepts and entities are linked by the *isA* relationship in the concept graph. For example, the concept “Marvel Superheroes” or “The Avengers” is instantiated by the entities like “Captain America” and “Black Widow”. Furthermore, “The Avengers” is a *subordinate* of “Marvel Superheroes”.

We evaluate a wide range of latest knowledge representation and neural text matching models on TAG, mainly including sequence-to-sequence text matching methods, e.g., ARC-I, ARC-II [13], BiMPM [42], MV-LSTM [41], MatchPyramid [28], CONV-KNRM [6], pre-trained language models, e.g., BERT [7], and graph embedding methods based on graph neural networks (GNNs) [18]. However, these state-of-the-art neural network models do not perform particularly well on this social media tagging task, due to their inability to either incorporate contextual information in the concept graph or accurately understand the fine logic conveyed by the sentence.

Motivated by the challenge of social media tagging in TAG, we further propose a novel **Graph-Graph matching** framework which converts a pair of concept and sentence to be matched into an aggregated *heterogeneous graph*, by converting the natural language sentence of interest into a graph via syntactical dependency parsing, while incorporating the surrounding connected terms via the concept graph. This allows us to reason over the constructed heterogeneous graph via Relational Graph Convolutional Network

[33] and acquires a deeper understanding of the logic interactions between terms in the sentence and their level of matching to the candidate concept.

Furthermore, the proposed Graph-Graph matching method can also tag a new concept which has never appeared in the dataset to a sentence, enabling the *zero-shot learning* ability. This is achieved by inferring the semantic meaning of a new concept based on its subordinate/superordinate relationships to existing concepts in the concept graph.

Extensive experiments on TAG have demonstrated the superior ability of the proposed model on the social media tagging task over a wide range of state-of-the-art neural matching and language model baselines in terms of accuracy, efficiency, and transferability to new concepts. To facilitate future research about confining and accurate conceptualization of open-domain social media content, we open source both the TAG dataset and the developed models at <https://github.com/XpastaX/TAG>. Our dataset is being constantly updated as new concepts and relationships are being discovered from online social media.

## 2 RELATED WORK

**Text tagging.** Text tagging aims to assign labels to text objects, including documents, sentences and phrases, to represent their main topics. Existing studies mainly focus on document tagging [4, 21], where a variety of approaches solve it as a content-based tag recommendation problem for information retrieval [5, 32, 38]. However, Chen et al. [4] uses predefined tags thus fails to consider the growing tag number in the real-world social media. Liu et al. [21] focuses on relation mining but fails to utilize the contextual information of its concept tags. Others [5, 32, 38] generate tags for a document by combining entities/keywords which can not represent the main idea of the document and lack a higher-level understanding of the text.

From the machine learning perspective, document tagging can be considered as a multi-label learning (MLL) problem, for which a lot of algorithms have been developed, including compressed-sensing based approach [12], ML-CSSP [2], LEML [48], FastXML [30], embedding-based DocTag2Vec [4], etc. Liu et al. [20, 21] propose a probabilistic inference-based approach and a matching-based approach to tag correlated concepts to a document through the key instances and their contextual words in the document. There are also studies about concept tagging for sentences, mainly through the perspective of a structured learning problem [10]. Algorithms have been developed from both generative and discriminative perspectives, such as LSTM-CRF [3, 14] and Seq2Seq models [40]. However, none of them takes the links between concepts into account, while TAG is the first dataset for tagging that considers them.

Lastly, text conceptualization [17, 36, 37, 44] aims to link a word or a short phrase to a set of concepts in an ontology or knowledge base, such as Probase [46]. However, the main tasks of those works focus on the disambiguation of words in the text. In contrast, the concept tagging task in TAG is different in that it focuses on understanding the main idea and semantic meaning of a sentence, when judging whether a concept phrase can be associated with the sentence. The solutions of the above works are mostly based on probabilistic inference and entity-based linking, and the knowledge

bases they utilized are mostly constructed from Wikipedia and web pages written in the author perspective which can not represent the interests of general users. In comparison, our proposed method matches the concept and the sentence by gathering the contextual information of the concepts mined from real-world search logs, which could best preserve the interest and the attention of the users.

**Neural Text matching.** A wide range of neural architectures have been designed for text matching. Representation-focused models, such as ARC-I [13] and BiMPM [42], transform a pair of text objects into context representation vectors through different encoders and output the matching result based on the context vectors [23, 26, 34, 41]. Interaction-focused models, such as MVLSTM [41], MatchPyramid [28] and ARC-II [13], extract the features of all pairwise interactions between words in the sentence pair, and aggregate the interaction features through deep neural networks to yield a matching result. Pre-trained language models such as BERT [7] can also be utilized for text matching through fine-tuning. Although neural text matching can characterize the semantic meaning of sentences to a certain degree, these models do not utilize the abundant relational and hierarchical information among concepts.

In contrast, the TAG dataset has called for solutions to a new challenge to the community. It asks the question of how to match a concept-sentence pair by learning from the comparisons between all semantic units in the given sentence and those in the given concept as well as in its subordinate/superordinate concepts. Another challenge is how to generalize a trained model to new concepts that have not appeared in training. There are also works utilizing graph structures [22, 27, 29] for text matching. However, by evaluating a wide range of techniques including GNN-based methods, we observe that the proposed concept tagging problem cannot readily be solved by combining existing off-the-shelf methods. Hence, we develop our Graph-Graph matching framework to tackle the above challenges.

**Concept graph construction.** The concept graph is constructed to empower models with the knowledge about the concepts in the real world. Existing works mainly focus on extracting the relations between entities in a text and concepts [15, 19, 43]. Instead, our main task is to conceptualize a document by tagging its sentences with the most semantically relevant concepts, aided by the concept graph which can best represent the attention and interest of users [20].

## 3 THE TAG DATASET

In this section, we formally introduce the concept tagging problem as a concept-sentence matching problem in the context of a concept graph. We describe the procedure in which the TAG dataset is constructed and analyze its unique characteristics.

We construct TAG dataset from the vast search logs and the online documents from QQ Browser. To emphasize the importance of zero-shot and few-shot learning in real-world scenarios, we gathered and developed solutions on a medium-sized dataset with a large number of concepts. Specifically, the dataset contains two parts: i) 10,000 concept-sentence pairs, each with a binary label indicating whether the concept can characterize the meaning of the sentence and thus be tagged onto the sentence; ii) a concept graph that contains a collection of phrases mined from vast amounts of

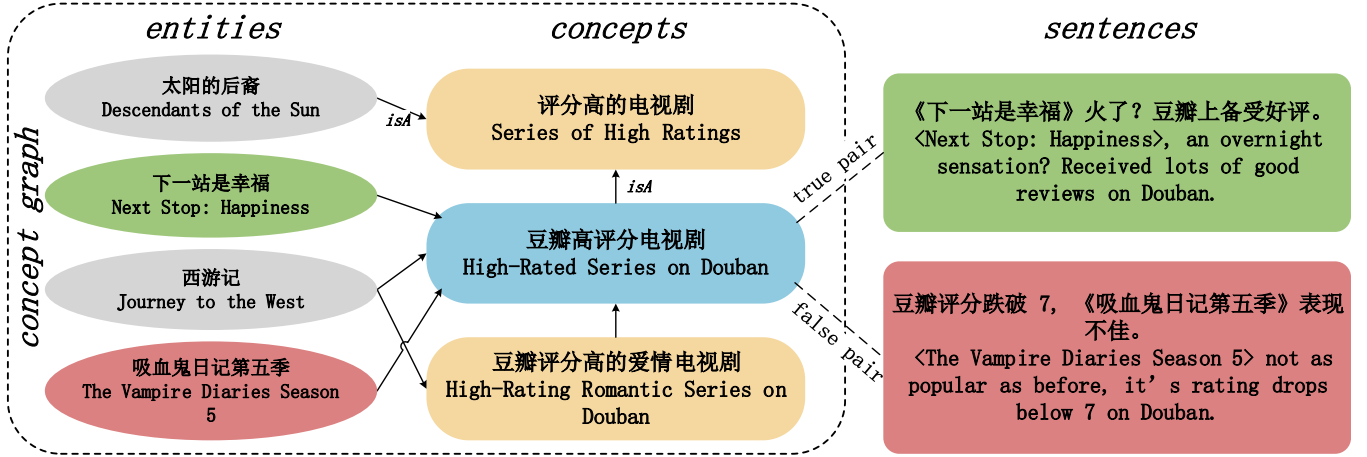


Figure 2: The TAG dataset consists of concept-sentence pairs with binary labels, depending on whether the concept can be concluded from the given sentence or not. The above is a snippet of the concept graph related to a concept “Highly-Rated Series on Douban” together with two sample pairs. *isA* means “is a subordinate of”.

online social media content, representing *concepts*, and *entities*, as well as the *isA* relationships between them.

The concepts are mined from search logs where popular queries satisfying certain rules serve as the concepts. The concept graph, as a user-centered ontology, is constructed from a web-scale corpus and is able to abstract entities and events into the higher-level concepts and topics based on user interests. Most existing taxonomy or knowledge bases, e.g., Probase, DBPedia, CNDBPedia [47], YAGO [39], utilize Wikipedia as their main source of content, which are not the best at tracking user-centered phrases on the online social media. Instead, the concept graph is mined from the *click graph*, which is a large bipartite graph formed by search queries and their corresponding clicked documents that can best represent the attention and interests of online users. More details regarding the construction of this concept graph can be found in our previous work [20]. In the concept graph we utilize, each concept is a phrase that represents a collection of entities sharing some common properties. The major relational connections in the concept graph are the *isA* (subordinate) relationships between the terms.

### 3.1 Task Description

The concept tagging task considered here is a challenging new problem that is rarely tackled in prior research. First, in TAG, each concept only appears in a few training sample pairs, leading to a few-shot learning problem, as opposed to traditional text classification where each class has many training samples. Second, each concept is usually composed of a few words and thus contains limited semantic information. Third, in real-world social media, new concepts constantly emerge on a daily basis, requiring the tagging model to generalize and transfer to new concepts that have not appeared in the training set.

We argue that a critical component to support the few-shot learning and transferability of a tagging model is to utilize the relational information in the concept graph, i.e., the comparisons among concepts and links between concepts and entities. To the best

of our knowledge, TAG is the first dataset specifically designed for online social media content tagging based on graphical knowledge representations. Existing datasets for document tagging, e.g., the UCCM dataset [21], do not have a concept graph component which lays the relational context around concepts and entities.

### 3.2 Dataset Construction

Table 1 presents the statistics of the TAG dataset. TAG contains 10,000 data examples with 230,504 relations and 5,438 concepts. Such a large amount of tags poses a big challenge to traditional labeling and text matching models. Figure 2 illustrates two concept-sentence pairs, a positive pair (matched) and a negative pair (unmatched). Two sentences on the right both contain some entities (titles of TV series) that belong to the concept “Highly-Rated Series on Douban” in the concept graph. However, the second sentence criticizes the Vampire Diaries and, together with the concept, should be labeled as a negative pair. We describe the detailed procedures of constructing the TAG dataset in the following:

**Candidate set retrieval.** Among all collected concepts, we sample 5438 of them from the “entertainment” domain, and extract their corresponding contexts to form our concept graph [20]. To prepare candidate concept-sentence pairs for human labeling, for each selected concept, we record the top-3 clicked documents retrieved by inputting the concept as a search query. For each document, we select the sentences which contain at least one entity that has *isA* relationship with the query concept. Note that even if a sentence contains an entity that belongs to a concept, this concept does not necessarily characterize the main idea of the sentence. Furthermore, many concepts share similar meanings as they are retrieved from similar search logs. For example, “Highly rated movies” has the same meaning as “Movies that receive a high rating”. Besides, many concepts are actually the combinations of other concepts. For example, the concept “Action movie with a high rating” is a combination of “Action movies” and “Highly rated movies”. To enhance the quality and cleanliness of the TAG dataset which aims to facilitate the

**Table 1: The statistics of the TAG dataset. *CG rel.* denotes the number of relations in the concept graph.**

	Total	Non-Overlap		Random	
		Train	Test	Train	Test
Concepts	5438	4326	1112	4660	1594
Sentences	8085	6780	1900	6733	1901
Positive	4152	3380	772	3332	820
Pairs	10000	8000	2000	8000	2000
CG rel.	230504	188312	42192	181461	49043

research in concept tagging, we further increase the diversity of the concepts in TAG. We use a filter to prevent semantically redundant concepts from being tagged to the same sentence. Specifically, for all the concepts paired to a sentence in the dataset, we rank them by the summation of frequencies of words and keep at most 4 concepts with the lowest aggregate word frequencies. This will also remove obvious matches to popular and easy concepts and increase the interestingness of the concepts.

**Labeling concept-sentence pairs.** All of our examples are from the “entertainment” domain, which is not deep or technical. Thus the annotation only requires common sense. We employ three annotators to ensure the quality of our dataset. Given the candidate pairs retrieved based on entities, each concept-sentence pair will be labeled with either 1 or 0 by two annotators (who are hired professional data labeling personnel yet totally ignorant and neutral to the research work being conducted), by solely examining the concept-sentence pair based on their own understanding without external information. If the results from the two annotators are different, the third annotator will be hired for a final check.

The criterion of the annotation is based on whether the concept can be concluded from the given sentence. In Figure 2, the sentence “<Next Stop: Happiness>, an overnight sensation? Received lots of good reviews on Douban” forms a positive example with the concept “High-rated series on Douabn” since it “received lots of good reviews on Douban”. In contrast, the sentence “<The Vampire Diaries Season 5> not as popular as before, its rating drops below 8 on Douban” is a negative pair with the concept since the series is “not popular as before”, which contradicts the concept.

**Post-processing.** After human labeling, we down-sample the negative pairs to ensure that the positive and negative examples are balanced. Furthermore, to simulate the real world scenario where the number of concepts keeps growing, we split the 10,000 sample pairs in TAG for training and testing under two different strategies: i) random splitting; ii) non-overlapped splitting, where the latter ensures the concepts that have appeared in the training set do not appear in the test set at all. The non-overlap scenario can be used to test the *zero-shot learning* capability of a model, i.e., the ability to generalize to new concepts it has never seen.

## 4 MODELS

In this section, we propose a novel and efficient solution for our concept-sentence tagging problem, which is a Graph-Graph matching model based on the representation of the given concept-sentence pair in a *heterogeneous graph*, and briefly describe our current

progress of the deployment. Before that, we first describe two existing neural text matching approaches that can also be applied to the problem, and show their limitations.

### 4.1 Seq-Seq and Graph-Seq Matching

There are two existing solutions that are intuitive and applicable to our tagging task. The first one is to build a sequence to sequence (Seq-Seq) matching model to leverage the semantic information contained in the concept-sentence pair. Specifically, we fine-tune the BERT [7] model to test sequence matching without any context information for the concept. We encode both the given concept and sentence by BERT, then either use the [CLS] tags or the embeddings of the sequences (BERT outputs one embedding for each Chinese character of the input text, and we encode those outputs by a Bi-LSTM layer) to perform binary classification. Denote the output embeddings of the sentence and the concept as  $V_S$  and  $V_T$ , respectively, we pass  $[|V_S - V_T|, V_S \circ V_T]$  into a Multi-layer perceptron (MLP) classifier to make the prediction.

Another natural solution to the proposed task is to build a Graph-Sequence (Graph-Seq) matching model to encode the contextual information of the given concept in the concept graph using a GNN. Specifically, on the graph side, the local context related to the concept is mapped into a directed graph with each node of the entities and the concepts represented by its word vectors (or the mean of the word vectors for the concepts formed by multiple words), and fed into a Relational Graph Convolutional Network (R-GCN) [33] to aggregate the contextual information in the concept graph. On the sequence side, we calculate the mean of the word vectors in the sentence as its embedding. Then, with the same operation performed in the Seq-Seq matching model, the node embedding of the concept is combined with the embedding of the sentence, and sent to an MLP classifier for prediction. We will introduce the details of the R-GCN model later in this section.

Both of the above two methods have some limitations. The Seq-Seq matching model fails to leverage the rich contextual information provided by the concept graph, while the Graph-Seq matching model lacks sufficient semantic support. Although the contextual information of the concept is utilized by the Graph-Seq model, the potential connection between the concept and the sentence is not explored. In other words, how the concept could relate to the sentence is not considered in the model.

### 4.2 Graph-Graph Matching

To incorporate the context information in the concept graph and model the interactions within a concept-sentence pair, we propose to form a *heterogeneous graph* representation of the given concept-sentence pair and combine it with R-GCN to estimate whether the concept is coherent with the main content of the sentence.

**Heterogeneous graph construction.** Figure 3 presents the heterogeneous graph structure of an example pair. For each concept-sentence pair, denote the sentence by  $S$  and the concepts in context by  $C = \{C_T, C_C\}$ , where  $C_T$  is the concept paired with the sentence (circled in red) and  $C_C$  (circled in black) represents other neighboring concepts in the concept graph. We further denote the set of entities associated with  $C$  as  $E_C$  and set of named entities in the



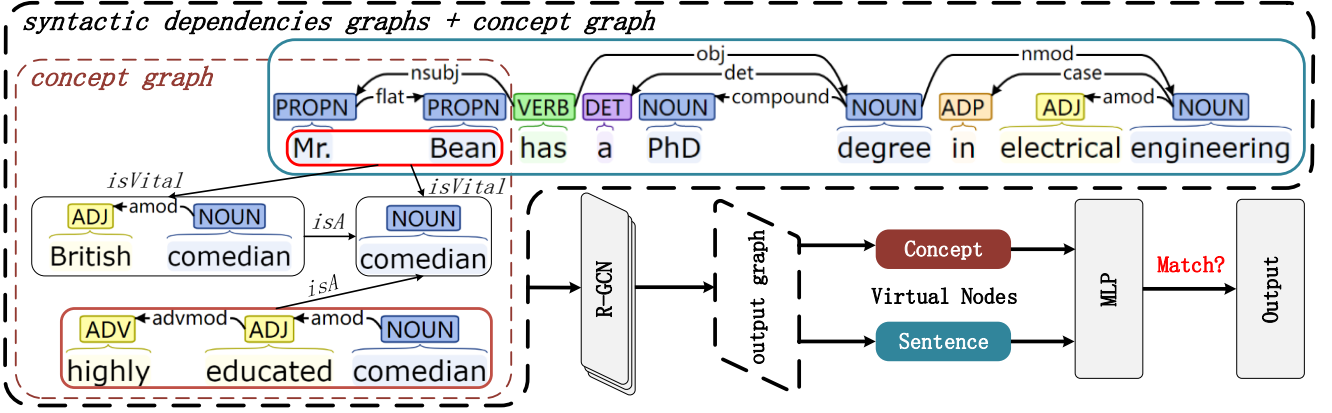


Figure 3: The procedure of predicting an example. Graph circled by the black dash line is the heterogeneous graph constructed for a concept-sentence pair, with the context information from the concept graph.

sentence  $S$  as  $E_S$ . Next, we introduce our procedure for constructing the heterogeneous graph based on Figure 4.

For each example, the first step is to obtain all syntactic dependencies in  $S$  and  $C$ . As presented in Figure 3, the sentence and the concepts will be segmented and fed into the Stanford CoreNLP model [25, 31] to generate a set of syntactic graphs in Figure 4. Each word is a node in these dependencies graphs either for the sentence or the concepts. To simplify the graphs, we fuse those syntactic graphs together by merging all nodes of the same word, as presented in Figure 4. The fused heterogeneous graph  $G_{syn}$  then stores all semantic information.

Second, we create *virtual nodes* for each element in  $S \cup C$  in the graph and establish connections between those virtual nodes and  $G_{syn}$ . Specifically, we link the concept nodes and all of its words together with *isA* relation, then we link the sentence node and all of its named entities together with *isNamedEntity* relation. With such connections, the virtual nodes perform like hubs, which aggregate all semantic information among  $G_{syn}$ .

Finally, to utilize the concept graph, for each entity in  $E_C \cap E_S$  and each concept in  $C$ , we map their context relations into the heterogeneous graph (red dash circle in Figure 3). Moreover, to emphasize the entities shared by  $S$  and  $C$ , we assign the links between such entities and  $C$  with a special relation named *isVital*. Since the concept-sentence pair is retrieved by those shared entities, such links are expected to help distinguish a successful retrieval.

There are in total 36 types of dependencies and 3 types of manually introduced relations in the heterogeneous graph. Note that the relations in the graph are all directed. Thus, some nodes may not have outward links to pass their messages to others. To fix this, for each relation in the heterogeneous graph, we introduce an opposite relation to ensure information transmission between the nodes, which doubles the number of relations to 78. This fix is also taken in the Graph-Seq matching model.

**Relevance estimation.** As shown in Figure 3, we apply R-GCN on the heterogeneous graph of the concept-sentence pair to encode the virtual nodes of  $S$  as  $V_S$  and  $C$  as  $V_T$ . Then we pass  $[|V_S - V_T|, V_S \circ V_T]$  into an MLP classifier to make a binary decision on

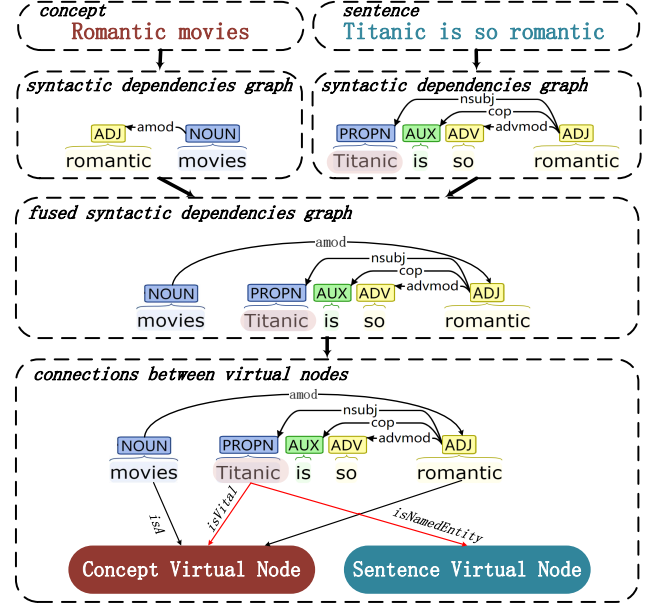


Figure 4: The process of connecting the concept and the sentence. “Titanic” is the entity shared between the concept graph and the sentence.

the concept-sentence pair. Note that the concept and the sentence are embedded together, which means the same concepts will have different embedding if paired with a different sentence. The insight is that the concept and the sentence will exchange their information to output the embedding specialized for the pair, which could lead to a better prediction.

The embedding of each node in the heterogeneous graph is formed by the following features: the word vector, part-of-speech (POS) tag, named entity recognition (NER) tag, and the source of the node. Each part will be tagged with “None” if not applicable.

Specifically, we employ the word vectors from [35], which contains a large number of domain-specific words, entities, slangs, and fresh words that are of high frequencies in our TAG dataset. For each virtual node in the heterogeneous graph, we use the mean of its word vectors as the embedding. For the source tag, we use “SENTENCE”, “CONCEPT” and “BOTH” to represent where the node comes from.

**Relational GCN.** We now briefly describe the R-GCN used in Graph-Graph and Graph-Seq matching models. Denote a directed graph as  $G = (V, E, R)$  with labeled edges  $e_{vw} = (v, r, w) \in E$ , where  $r \in R$  is a type of relation and  $v, w \in V$  are two nodes connected by  $r$ . In the message-passing framework [9] of graph neural networks, the hidden states  $h_v^l$  of each node  $v \in G$  at layer  $l$  are updated based on messages  $m_v^{l+1}$  according to:

$$m_v^{l+1} = \sum_{w \in N(v)} M_l(h_v^l, h_w^l, e_{vw}), \quad (1)$$

$$h_v^{l+1} = U_l(h_v^l, m_v^{l+1}), \quad (2)$$

where  $N(v)$  are the neighbors of  $v$  in  $G$ ,  $M_l$  and  $U_l$  are the message function and vertex update function at layer  $l$ , respectively.

The message passing function of Relational Graph Convolutional Networks is defined as:

$$h_v^{l+1} = \sigma \left( \sum_{r \in R} \sum_{w \in N^r(v)} \frac{1}{c_{vw}} W_r^l h_w^l + W_0^l h_v^l \right), \quad (3)$$

where  $N^r(v)$  is the set of neighbors under relation  $r \in R$ .  $\sigma(\cdot)$  is an element-wise activation function such as  $\text{ReLU}(\cdot) = \max(0, \cdot)$ .  $W_r^l$  and  $W_0^l$  are learned weight matrices.  $c_{vw}$  is a problem-specific normalization constant that can be learned or pre-defined (e.g.,  $c_{vw} = |N^r(v)|$ ). We can see that R-GCN learns relation-specific transformation matrices to take the type and direction of each edge into account.

Moreover, R-GCN exploits basis or block-diagonal decomposition to regularize the number of weights. For basis decomposition, each weight matrix  $W_r^l \in \mathbb{R}^{d^{l+1} \times d^l}$  is decomposed as  $W_r^l = \sum_{b=1}^B a_{rb}^l V_b^l$ , where  $V_b^l \in \mathbb{R}^{d^{l+1} \times d^l}$  are base weight matrices. In this way, only the coefficients  $a_{rb}^l$  depend on  $r$ . For block-diagonal decomposition,  $W_r^l$  is defined through the direct sum over a set of low-dimensional matrices as  $W_r^l = \bigoplus_{b=1}^B Q_{br}^l$ , where  $W_r^l = \text{diag}(Q_{1r}^l, Q_{2r}^l, \dots, Q_{Br}^l)$  is a block-diagonal matrix with  $Q_{br}^l \in \mathbb{R}^{(d^{l+1}/B) \times (d^l/B)}$ . The basis function decomposition introduces weight sharing between different relation types, while the block decomposition applies sparsity constraint on the weight matrices.

**Deployment in QQ Browser.** Our current document tagging method is based on the term-matching between the concepts and the documents [20]. However, with the growing number of concepts and documents, it is harder and harder for traditional term-matching to satisfy the need for tagging accuracy, and a more effective tagging method is in need. Thus, we reform the document tagging task into the concept-sentence tagging style (TAG), and develop the proposed Graph-Graph matching method on TAG.

We have constructed a concept graph with over 440k existing concepts covering 38 categories, e.g., entertainment, technology, sports, etc. A real-time extraction method is further deployed to integrate the new concepts into our concept graph. To build the concept-sentence pairs proposed in TAG, we also deployed a key

sentence extraction method which now is processing over one million newly uploaded documents every day. We are re-tagging the existing documents with the concepts in the entertainment domain for experiments, and will soon deploy the real-time Graph-Graph matching method to tag all existing and new documents with the concepts from all categories.

## 5 EXPERIMENTS

In this section, we evaluate the proposed concept tagging method on TAG with a number of state-of-the-art text matching baselines.

**Baselines and Metrics.** We compare the proposed method with 6 existing baseline methods, including ARC-I, ARC-II [13], BiMPM [42], MV-LSTM [41], MatchPyramid [28], CONV-KNRM [6], and 3 designed baseline methods: Graph-Seq matching, BERT-Siamese and BERT-CLS. The 6 existing baselines and the BERT model all perform semantic matching of a given concept-sentence pair, while the Graph-Graph matching and Graph-Seq matching models also leverage the contextual information in the concept graph. We use the Chinese word vector proposed in [35] as the common input to all the models except BERT.

The performance of each model is evaluated on the two versions of TAG, including random and non-overlapped versions, as shown in Table 1. We utilize 1,000 sample split from the training set for validation. Additionally, for the *non-overlapped* training set, we will ensure the concepts in the validation set are different from those remains in the training set as well. The evaluation metrics are the F1 score and the accuracy of the prediction.

**Implementation details.** The 6 existing baselines are implemented by MatchZoo [11]. The BERT is from Transformers [45]. And our graph models are developed with Pytorch Geometric [8]. We have manually tuned the hyper-parameters of all the models. For the *Graph-Graph matching* model, we set the number of R-GCN layers to 3 and the number of bases to 14. The initial learning rate is  $1e^{-4}$  with a cosine decay scheduler. The model is trained for 20 epochs with 10% steps to warm up. For the *Graph-Seq matching* model, We use a 3-layer and 2-base R-GCN on the graph side, since we only have *isA* and its opposite relation. The learning rate and scheduler are the same as those in Graph-Graph matching. The number of epochs is 50. For *BERT*, we fine-tune the model with a batch size of 8 and a learning rate of  $1e^{-6}$  for 30 epochs with 10% steps to warm up. For each model, the epoch with the best validation F1 score is used for final evaluation.

### 5.1 Results and Analysis

Table 2 summarizes the performance comparison of all the methods on the two versions of TAG. Each score in the table is an average of 5 runs. We can observe that the proposed Graph-Graph matching method significantly outperforms all other baselines. Comparing the Graph-Seq matching model with the 6 existing baselines (i.e., from ARC-I to CONV-KNRM), the Graph-Seq matching model outperforms all of them. That confirms our hypothesis that the relational information in concept hierarchy offers additional help to the prediction. Furthermore, the proposed Graph-Graph matching model significantly outperforms the Graph-Seq matching model, since the syntactic graph provides richer semantic information than encoding a concept/sentence by taking the mean of word vectors.

**Table 2: Experimental results.**

Model	Overlapped		Non-overlapped	
	F1	Acc	F1	Acc
ARC-I	64.97	71.48	56.51	65.99
ARC-II	66.39	70.50	56.11	64.59
BiMPM	62.50	69.66	58.46	66.62
MV-LSTM	66.60	71.10	59.61	64.19
MatchPyramid	63.70	71.32	56.87	67.06
CONV-KNRM	66.93	73.31	60.59	64.93
G-S	68.26	73.84	61.89	67.58
BERT-Siamese	69.91	74.64	65.15	70.43
BERT-CLS	70.14	74.49	65.21	70.32
G-G w/o Syn	69.14	74.61	63.98	69.72
G-G w/o CG	74.38	79.06	66.73	72.98
G-G	<b>76.04</b>	<b>80.26</b>	<b>69.77</b>	<b>75.72</b>

BERT outperforms all other baseline methods because it has learned lots of prior knowledge through pre-training. Our proposed model outperforms the BERT methods with the 40 times smaller model size: both of the BERT implementations have more than 110 million parameters while our proposed method contains less than 2.66 million parameters. This is because the training corpus of Chinese BERT is mainly collected from official sites (e.g. Wikipedia) which are usually written in an objective aspect, while in real case application, most online texts are written in user-aspect containing complicated sentiments and willing, and usually come with fresh words or slangs. Different from BERT, with proper word segmentation, the proposed method can effectively utilize the word vectors trained with a large amount of user-aspect corpus, which contains lots of slang and self-made words. Therefore, it can better understand the text. Furthermore, BERT does not utilize the context information in the concept graph either. Compared with BERT, the syntactic graph provides great advantages to semantic aggregation. With the additional context contained in the concept graph, the model can characterize the semantics of the new concepts more accurately, leading to better performance.

Our model also gives better results on Non-overlapped data because each new concept will be linked with the sentence through the shared entities retrieved from the concept graph, e.g., “Titanic” in Figure 4. The contextual information in the concept graph can further help in understanding the new concept. By leveraging both the advantages, our Graph-Graph matching model enriches the information contained in the output embedding, which helps its MLP classifier performs better prediction on zero-shot learning.

## 5.2 Ablation Study

Furthermore, to measure the impacts of the syntactic dependencies graph (SDG) and the concept graph (CG) to our proposed model, we perform ablation studies on our proposed method by removing the SDG and the CG, respectively. As shown in Table 2, “G-G w/o Syn” represents the model without syntactic dependencies graph, while “G-G w/o CG” represents the model without the concept graph. Both variations yield lower performance compared to the original one. Our proposed method of extracting information from

SDG reaches a higher performance than BERT. The “G-G w/o Syn” model can also reach a comparable performance with a 100 times smaller model size than BERT methods.

Comparing the impacts of CG and SDG on the proposed model, SDG has a greater overall improvement on the performance of the model. Together with the word vectors, our proposed method can extract rich semantic information from the syntactic dependencies. The performance of “G-G w/o CG” model decreases less than “G-G w/o Syn”, this is because the SDG is the backbone of extracting semantic information, which is most important for understanding a sentence. Instead, CG mainly contributes to providing contextual information to enhance the understanding of a sentence, especially for zero-shot learning. By putting CG back to the “G-G w/o CG” model, the F1 score increases 3.04% for the non-overlapped dataset while increasing 1.66% for the random dataset. This fact further shows the importance of CG in TAG: the contextual information it provides can help the model better understand the new concept, which leads to higher performance on the zero-shot tagging task.

From the results above, we can conclude that both the contextual information in CG and the syntactic dependencies contribute to improving the semantic understanding of the model. The contextual information in CG of TAG shows excellent potential on the zero-shot task. Our proposed method well unitizes both of the advantages to achieve superior performance.

## 5.3 Error Analysis

To analyze the weaknesses of the proposed model and the BERT model, we further gathered all the incorrect predictions of the non-overlapped dataset. As shown in Table 3, we can generally categorize the reason for the errors into three main types: lack of reasoning ability, failed Named Entity Recognition, and complex context. Our model is less prone to the first type of errors in general, and the last two types are only common in the BERT model.

**Reasoning ability.** Words may have implicit information that is hard to extract. In the first example in Table 3, the concept is about the **movie** of “Resident Evil 7”, while the sentence is talking about the **game** “Resident Evil 7”. The annotators label the example as “False”, because “speed run” means clearing a game as fast as possible and “Zoe ending” suggests multiple endings exist, while the movies usually tell stories with single-endings. However, the models may lack the related knowledge to perform such reasoning and make the correct prediction. As shown in this example, BERT incorrectly predicts the label as “True” since it failed to distinguish “Resident Evil 7” in the sentence from “Movie Resident Evil 7”, while our proposed method successfully recognizes “Resident Evil 7” in the sentence as a video game and predict the example as “False”. For all 14 concept-sentence pairs with the concept “Movie Resident Evil 7” in the test set, BERT made 13 false predictions while our model made only 4 false predictions. However, by comparing the difference between the entities in the concept graph and the sentence, e.g. the entity “Zoe” is related to the game “Resident Evil 7” while not exists in the movie, we may further improve the performance of our proposed model.

**Named Entity Recognition.** Limited by the nature of the Chinese language, BERT will encode a Chinese sentence in characters. This could cause problems when it comes to embedding an entity,



**Table 3: Typical types of errors in concept matching for social media content.**

Error Types	Label	Concept	Sentence
Reasoning Ability	False	Movie Resident Evil 7	“Resident Evil 7”, 45min speed run video of Zoe ending.
Named Entity Recognition	True	Previous Oscars winners	On the 14th Oscars in 1942, Gary Cooper was nominated for the second time and finally won the award.
Complex Context	True	TV series about hospitals	Yesterday, a reader who has been following the TV series “Behind” called the hotline of our newspaper and said: The current hospital is indeed like that, and the plot is very real.

which usually is a combination of multiple characters. Take the second sentence in Table 3 as an example, Gary Copper is the name of a person, which in Chinese is “加里·库珀”, formed by 4 unrelated Chinese characters since it is a transliteration of a name. BERT is not able to distinguish such an entity and the irrelevant characters make it more difficult to understand the sentence. It then outputs an inaccurate embedding of the sentence, which leads to a false prediction. In contrast, by leveraging the NER technique, our proposed model can distinguish these entities, extract their meaning of them from word vectors, and output the correct embedding.

**Complex context.** Long sentences may have complex and rich information which could be hard for the models to understand. As an example, in the third sentence of Table 3, the context information includes: “a reader called the hotline of the newspaper”, “the reader has been following the TV series Behind”, and “the plot about the hospital in Behind is true”. Based on the information above, we can draw the key “Behind is a TV series about hospitals” and match the sentence with the concept “TV series about hospitals”. However, when encoding the sentence, BERT will not focus on “Behind is a TV series about hospitals”. The key information is then “diluted” in the output embedding, which results in the mismatch between the concept and the sentence. In contrast, our proposed method successfully established the connections between the concept and the sentence as described in Section 4 and integrated the connections in encoding. Specifically, in this example, the concept and the sentence are connected by the words “TV series”, “Behind” and “hospital”, through which the information will be passed to encode the concept and the sentence. Such a mechanism emphasizes the significance of the concept during the encoding process, and results in a better prediction in complex context situations.

## 6 CONCLUSION

In this paper, we present TAG, a high quality labeled dataset of 10,000 concept-sentence pairs to facilitate research on fine-grained concept tagging for social media content understanding in the presence of a concept graph. By evaluating a wide range of state-of-the-art neural text matching methods, we observe that the proposed concept tagging problem becomes a challenging task, as existing sequence-to-sequence or graph embedding methods either fail to leverage the context information in the concept graph or cannot

accurately grasp the logic meaning conveyed in the natural language sentence. The need to generalize and transfer models to new concepts further increases the challenge. We further design a novel framework to convert text to a graph via dependency parsing such that Graph-Graph matching can be conducted between the sentence and the concept in the context of the concept graph. Extensive experiments have demonstrated the superior ability of the proposed method to utilize contextual information in the concept graph, to reason about the logical match between a concept and a sentence, and to generalize and transfer to new concepts.

## ACKNOWLEDGMENTS

The authors would like to thank Dr. De-Chuan Zhan for his helpful advice and comments on various technical issues investigated in the paper.

## REFERENCES

- [1] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. *The semantic web* (2007), 722–735.
- [2] Wei Bi and James Kwok. 2013. Efficient multi-label classification with many labels. In *ICML*. 405–413.
- [3] Raghavendra Chalapathy, Ehsan Zare Borzeshi, and Massimo Piccardi. 2016. Bidirectional LSTM-CRF for clinical concept extraction. *arXiv preprint arXiv:1611.08373* (2016).
- [4] Sheng Chen, Akshay Soni, Aasish Pappu, and Yashar Mehdad. 2017. Doctag2vec: An embedding based multi-label learning approach for document tagging. *arXiv preprint arXiv:1707.04596* (2017).
- [5] Paul-Alexandru Chirita, Stefania Costache, Wolfgang Nejdl, and Siegfried Handschuh. 2007. P-tag: large scale automatic generation of personalized annotation tags for the web. In *Proc. WWW*. 845–854.
- [6] Zhuyun Dai, Chenyan Xiong, Jamie Callan, and Zhiyuan Liu. 2018. Convolutional neural networks for soft-matching n-grams in ad-hoc search. In *Proc. ACM WSDM*. 126–134.
- [7] 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* (2018).
- [8] Matthias Fey and Jan E. Lenssen. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLR Workshop on Representation Learning on Graphs and Manifolds*.
- [9] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural message passing for quantum chemistry. In *Proc. ICML-Volume 70*. JMLR. org, 1263–1272.
- [10] Jacopo Gobbi, Evgeny Stepanov, and Giuseppe Riccardi. 2018. Concept Tagging for Natural Language Understanding: Two Decadelong Algorithm Development. *arXiv preprint arXiv:1807.10661* (2018).
- [11] Jiafeng Guo, Yixing Fan, Xiang Ji, and Xueqi Cheng. 2019. MatchZoo: A Learning, Practicing, and Developing System for Neural Text Matching. In *Proc. ACM SIGIR (Paris, France) (SIGIR’19)*. ACM, New York, NY, USA, 1297–1300.

- [12] Daniel J Hsu, Sham M Kakade, John Langford, and Tong Zhang. 2009. Multi-label prediction via compressed sensing. In *NIPS*. 772–780.
- [13] Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. In *NIPS*. 2042–2050.
- [14] Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991* (2015).
- [15] Lei Ji, Yujing Wang, Botian Shi, Dawei Zhang, Zhongyuan Wang, and Jun Yan. 2019. Microsoft concept graph: Mining semantic concepts for short text understanding. *Data Intelligence* 1, 3 (2019), 238–270.
- [16] Immanuel Kant. [n.d.]. *Logic*. Dover Publications.
- [17] Dongwoo Kim, Haixun Wang, and Alice Oh. 2013. Context-dependent conceptualization. In *IJCAI*.
- [18] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [19] Jiaqing Liang, Yanghua Xiao, Haixun Wang, Yi Zhang, and Wei Wang. 2017. Probase+: Inferring missing links in conceptual taxonomies. *IEEE TKDE* 29, 6 (2017), 1281–1295.
- [20] Bang Liu, Weidong Guo, Di Niu, Jinwen Luo, Chaoyue Wang, Zhen Wen, and Yu Xu. 2020. GIAN: Scalable Creation of a Web-scale Ontology. In *Proc. ACM SIGMOD*. 393–409.
- [21] 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 *Proc. ACM SIGKDD*. 1831–1841.
- [22] Bang Liu, Di Niu, Haojie Wei, Jinghong Lin, Yancheng He, Kunfeng Lai, and Yu Xu. 2019. Matching Article Pairs with Graphical Decomposition and Convolutions. In *Proc. ACL*. 6284–6294.
- [23] Bang Liu, Ting Zhang, Fred X Han, Di Niu, Kunfeng Lai, and Yu Xu. 2018. Matching natural language sentences with hierarchical sentence factorization. In *Proc. WWW. IW3C2*, 1237–1246.
- [24] Xusheng Luo, Luxin Liu, Yonghua Yang, Le Bo, Yuanpeng Cao, Jinhang Wu, Qiang Li, Keping Yang, and Kenny Q Zhu. 2020. AliCoCo: Alibaba E-commerce Cognitive Concept Net. *arXiv preprint arXiv:2003.13230* (2020).
- [25] Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proc. ACL: system demonstrations*. 55–60.
- [26] Jonas Mueller and Aditya Thyagarajan. 2016. Siamese recurrent architectures for learning sentence similarity. In *AAAI Conference on Artificial Intelligence*.
- [27] Giannis Nikolentzos, Polykarpos Meladianos, François Rousseau, Yannis Stavrakas, and Michalis Vazirgiannis. 2017. Shortest-path graph kernels for document similarity. In *Proc. EMNLP*. 1890–1900.
- [28] Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text Matching as Image Recognition. In *AAAI*. 2793–2799.
- [29] Christian Paul, Achim Rettinger, Aditya Mogadala, Craig A Knoblock, and Pedro Szekely. 2016. Efficient graph-based document similarity. In *ESWC*. Springer, 334–349.
- [30] Yashoteja Prabhu and Manik Varma. 2014. Fastxml: A fast, accurate and stable tree-classifier for extreme multi-label learning. In *Proc. ACM SIGKDD*. 263–272.
- [31] Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. Stanza: A Python Natural Language Processing Toolkit for Many Human Languages. *arXiv preprint arXiv:2003.07082* (2020).
- [32] Steffen Rendle, Leandro Balby Marinho, Alexandros Nanopoulos, and Lars Schmidt-Thieme. 2009. Learning optimal ranking with tensor factorization for tag recommendation. In *Proc. ACM SIGKDD*. 727–736.
- [33] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *ESWC*. Springer, 593–607.
- [34] Aliaksei Severyn and Alessandro Moschitti. 2015. Learning to rank short text pairs with convolutional deep neural networks. In *Proc. ACM SIGIR*. ACM, 373–382.
- [35] Yan Song, Shuming Shi, Jing Li, and Haisong Zhang. 2018. Directional Skip-Gram: Explicitly Distinguishing Left and Right Context for Word Embeddings. In *Proc. NAACL (Short Papers)*. 175–180.
- [36] Yangqiu Song, Haixun Wang, Zhongyuan Wang, Hongsong Li, and Weizhu Chen. 2011. Short text conceptualization using a probabilistic knowledgebase. In *IJCAI*.
- [37] Yangqiu Song, Shusen Wang, and Haixun Wang. 2015. Open domain short text conceptualization: A generative+ descriptive modeling approach. In *IJCAI*.
- [38] Yang Song, Lu Zhang, and C Lee Giles. 2011. Automatic tag recommendation algorithms for social recommender systems. *ACM TWEB* 5, 1 (2011), 1–31.
- [39] Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proc. WWW*. 697–706.
- [40] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NIPS*. 3104–3112.
- [41] Shengxian Wan, Yanyan Lan, Jiafeng Guo, Jun Xu, Liang Pang, and Xueqi Cheng. 2016. A Deep Architecture for Semantic Matching with Multiple Positional Sentence Representations. In *AAAI*, Vol. 16. 2835–2841.
- [42] Zhiguo Wang, Wael Hamza, and Radu Florian. 2017. Bilateral multi-perspective matching for natural language sentences. *arXiv preprint arXiv:1702.03814* (2017).
- [43] Zhongyuan Wang and Haixun Wang. 2016. Understanding short texts. (2016).
- [44] Zhongyuan Wang, Kejun Zhao, Haixun Wang, Xiaofeng Meng, and Ji-Rong Wen. 2015. Query understanding through knowledge-based conceptualization. In *IJCAI*.
- [45] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace's Transformers: State-of-the-art Natural Language Processing. *ArXiv abs/1910.03771* (2019).
- [46] Wentao Wu, Hongsong Li, Haixun Wang, and Kenny Q Zhu. 2012. Probase: A probabilistic taxonomy for text understanding. In *Proc. ACM SIGMOD*. 481–492.
- [47] Bo Xu, Yong Xu, Jiaqing Liang, Chenhao Xie, Bin Liang, Wanyun Cui, and Yanghua Xiao. 2017. CN-DBpedia: A never-ending Chinese knowledge extraction system. In *IEA/AIE*. Springer, 428–438.
- [48] Hsiang-Fu Yu, Prateek Jain, Purushottam Kar, and Inderjit Dhillon. 2014. Large-scale multi-label learning with missing labels. In *ICML*. 593–601.