World Knowledge as Indirect Supervision for Document Clustering

CHENGUANG WANG, Peking University

YANGQIU SONG, Hong Kong University of Science and Technology

DAN ROTH, University of Illinois at Urbana-Champaign

MING ZHANG, Peking University

JIAWEI HAN, University of Illinois at Urbana-Champaign

One of the key obstacles in making learning protocols realistic in applications is the need to supervise them, a costly process that often requires hiring domain experts. We consider the framework to use the world knowledge as indirect supervision. World knowledge is general-purpose knowledge, which is not designed for any specific domain. Then, the key challenges are how to adapt the world knowledge to domains and how to represent it for learning. In this article, we provide an example of using world knowledge for domain-dependent document clustering. We provide three ways to specify the world knowledge to domains by resolving the ambiguity of the entities and their types, and represent the data with world knowledge as a heterogeneous information network. Then, we propose a clustering algorithm that can cluster multiple types and incorporate the sub-type information as constraints. In the experiments, we use two existing knowledge bases as our sources of world knowledge. One is Freebase, which is collaboratively collected knowledge about entities and their organizations. The other is YAGO2, a knowledge base automatically extracted from Wikipedia and maps knowledge to the linguistic knowledge base, WordNet. Experimental results on two text benchmark datasets (20newsgroups and RCV1) show that incorporating world knowledge as indirect supervision can significantly outperform the state-of-the-art clustering algorithms as well as clustering algorithms enhanced with world knowledge features.

A preliminary version of this work appeared in the proceedings of KDD 2015 [Wang et al. 2015a]. This journal version has made several major improvements. First, we have proposed a new and general learning framework for machine learning with world knowledge as indirect supervision, where document clustering is a special case in the original paper. Second, in order to make our unsupervised semantic parsing method more understandable, we add several real cases from the original sentences to the resulting logic forms with all the necessary information. Third, we add details of the three semantic filtering methods and conduct

C. Wang gratefully acknowledges the support by the National Natural Science Foundation of China (NSFC Grant Number 61472006) and the National Basic Research Program (973 Program No. 2014CB340405). The research is also partially supported by the Army Research Laboratory (ARL) under agreement W911NF-09-2-0053, and by DARPA under agreement number FA8750-13-2-0008. The research is also partially sponsored by China National 973 project 2014CB340304, U.S. National Science Foundation IIS-1320617, IIS-1354329 and IIS 16-18481, HDTRA1-10-1-0120, and grant 1U54GM114838 awarded by NIGMS through funds provided by the trans-NIH Big Data to Knowledge (BD2K) initiative (www.bd2k.nih.gov), and MIAS, a DHS-IDS Center for Multimodal Information Access and Synthesis at UIUC. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied by these agencies or the U.S. Government.

Authors' addresses: C. Wang and M. Zhang (corresponding author), School of EECS, Peking University, Haidian, Beijing 100871, China; emails: {wangchenguang, mzhang_cs}@pku.edu.cn; Y. Song, Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong; email: yqsong@cse.ust.hk; D. Roth and J. Han, Department of Computer Science, University of Illinois at Urbana-Champaign, Champaign, IL 61820; emails: {danr, hanj}@illinois.edu.

Author's current address: C. Wang, IBM Research - Almaden, 650 Harry Road, San Jose, CA 95120; email: chenguang.wang@ibm.com.

Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2016 ACM 1556-4681/2016/12-ART13 \$15.00

DOI: <http://dx.doi.org/10.1145/2953881>

deep analysis of the three semantic filters, by using case studies to show why the conceptualization-based semantic filter can produce more accurate indirect supervision. Finally, in addition to the experiment on 20 newsgroup data and Freebase, we have extended the experiments on clustering results by using all the combinations of text (20 newsgroup, MCAT, CCAT, ECAT) and world knowledge sources (Freebase, YAGO2).

CCS Concepts: • **Theory of computation** → **Unsupervised learning and clustering**;

Additional Key Words and Phrases: World knowledge, heterogeneous information network, document clustering, knowledge base, knowledge graph

ACM Reference Format:

Chenguang Wang, Yangqiu Song, Dan Roth, Ming Zhang, and Jiawei Han. 2016. World knowledge as indirect supervision for document clustering. *ACM Trans. Knowl. Discov. Data* 11, 2, Article 13 (December 2016), 36 pages.

DOI: <http://dx.doi.org/10.1145/2953881>

1. INTRODUCTION

Machine learning algorithms have become pervasive in multiple domains, impacting a wide variety of applications. Nonetheless, a key obstacle in making learning protocols realistic in applications is the need to supervise them, a costly process that often requires hiring domain experts. In the past decades, machine learning community has elaborated to reduce the labeling work done by humans for supervised machine learning algorithms or to improve unsupervised learning with only minimum supervision. For example, semi-supervised learning [Chapelle et al. 2006] is proposed to use only partially labeled data and a lot of unlabeled data to perform learning with the hope that it can perform as good as fully supervised learning. Transfer learning [Pan and Yang 2010] uses the labeled data from other relevant domains to help the learning task in the target domain. However, there are still many cases that neither semi-supervised learning nor transfer learning can help. For example, in the era of big data, we can have a lot of textual information from different Web sites, e.g., blogs, forums, and mailing lists. It is impossible to ask humans to annotate all the required tasks. It is also difficult to find relevant labeled domains. Recognizing that some domains can be very specific and really need the domain experts to perform annotation, e.g., the medical domain publication classification, therefore we should consider a more general approach to further reducing the labeling cost for learning tasks in diverse domains.

Fortunately, with the proliferation of general-purpose knowledge bases (or knowledge graphs), e.g., Cyc project [Lenat and Guha 1989], Wikipedia, Freebase [Bollacker et al. 2008], KnowItAll [Etzioni et al. 2004], TextRunner [Banko et al. 2007], ReVerb [Fader et al. 2011], Ollie [Mausam et al. 2012], WikiTaxonomy [Ponzetto and Strube 2007], Probase [Wu et al. 2012], DBpedia [Auer et al. 2007], YAGO [Suchanek et al. 2007], NELL [Mitchell et al. 2015], and Knowledge Vault [Dong et al. 2014], we have an abundance of available world knowledge. We call these knowledge bases as world knowledge [Gabrilovich and Markovitch 2005], because they are universal knowledge that are either collaboratively annotated by human labelers or automatically extracted from big data. In general, world knowledge can be common sense knowledge, common knowledge, or domain-dependent knowledge. Common sense knowledge is the facts that an ordinary person is expected to know, but we seldom need to learn them on purpose. For example, we all know that a dog is an animal. Common knowledge is widely accessible knowledge. For example, the population of Illinois is around 12.8M. An ordinary person may not know the exact number of population, but we would find the answer easily from numerous sources. Domain knowledge can be very specific, such as the complete set of species of animals, or the taxonomic classification of trees.

When world knowledge is annotated or extracted, it is not collected for any specific domain. However, because we believe the facts in world knowledge bases are very useful and of high quality, we propose using them as supervision for many machine

learning problems. People have found it useful to use world knowledge as distant supervision for entity and relation extraction and embedding [Mintz et al. 2009; Wang et al. 2014b; Xu et al. 2014]. This is a direct use of the facts in world knowledge bases, where the entities in the knowledge bases are matched in the context regardless the ambiguity. A more interesting question is *Can we use the world knowledge to indirectly “supervise” more machine learning algorithms or applications?* For example, if we can use world knowledge as indirect supervision, then we can extend the knowledge about entities and relations to more generic text analytic problems, e.g., categorization and information retrieval.

Thus, we consider a general machine learning framework that can incorporate world knowledge into machine learning algorithms. In general, there are three challenges of incorporating world knowledge into machine learning algorithms: (1) *domain specification*, (2) *knowledge representation*, and (3) *propagation of indirect supervision*.

First, as mentioned, world knowledge is not designed for any specific domain. For example, when we want to cluster the documents about entertainment or sports, then the world knowledge about names of celebrities and athletes may help, whereas the terms used in science and technology may not be very useful. Thus, a key issue is how we should adapt world knowledge to the domain specific tasks. Second, when we have the world knowledge, we should consider how we can represent it for the domain-dependent tasks. An intuitive way is to use knowledge bases to generate features for machine learning algorithms [Gabrilovich and Markovitch 2005; Song et al. 2015]. We call these features “flat” because they do not consider the link information in the knowledge bases. However, most knowledge bases use a linked network to organize the knowledge. For example, a *CEO* is connected to an *IT company*, and the *IT company* is a *company*. Thus, the structure of the knowledge also provides rich information about the connections of entities and relations. Therefore, we should also carefully consider the best way to represent the world knowledge for machine learning algorithms. Third, given the world knowledge about entities and their relations, as well as the types of entities and relations, we should consider an effective algorithm that can propagate the knowledge about entity and relation categories to the categories of data that contain the entities and relations. This is a non-trivial task because we should consider both the data representation and the structural representation of the world knowledge.

In this article, we illustrate the framework of machine learning with world knowledge using a document clustering problem. We select two knowledge bases, i.e., Freebase and YAGO2, as the sources of world knowledge. Freebase [Bollacker et al. 2008] is a collaboratively collected knowledge base about entities and their organizations. YAGO2 [Suchanek et al. 2007] is a knowledge base automatically extracted from Wikipedia and maps the knowledge to the linguistic knowledge base, WordNet [Fellbaum 1998]. To adapt the world knowledge to domain specific tasks, we first use semantic parsing to ground any text to the knowledge bases [Berant et al. 2013]. We then apply entity frequency, document frequency, and conceptualization [Song et al. 2015] based semantic filters to resolve the ambiguity problem when adapting world knowledge to the domain tasks. After that, we have the documents as well as the extracted entities and their relations. Since the knowledge bases provide the entity types, the resulting data naturally form a heterogeneous information network (HIN) [Han et al. 2010]. We show an example of such HIN in Figure 1. The specified world knowledge, such as named entities (“Bush” and “Obama”) and their types (*Person*), as well as the documents and the words form the HIN. We then formulate the document clustering problem as an HIN partitioning problem, and provide a new algorithm to better perform clustering by incorporating the rich structural information as constraints in the HIN. For example, the HIN builds a link (a must-link constraint) between “Obama” of sub-type *Politician* in one document and “Bush” of sub-type *Politician* in another document. Such link

Fig. 1. Heterogeneous information network example. The network $\mathcal{G}$ contains five entity types: document, word, date, person, and location, which are represented with gray rectangle, gray round, green square, blue round, and yellow triangle, respectively.

and type information could be very useful if the target clustering domain is “Politics.” Then, we use the sub-type information as supervision information for the entities, and propagate the information to the documents through the network. Therefore, we call the sub-type information as indirect supervision of documents.

The main contributions of this work are highlighted as follows:

- We propose a new learning framework of machine learning with world knowledge as indirect supervision. We give the general data mining and machine learning framework, and use a specific problem of document clustering to illustrate the process.
- We propose to use semantic parsing and semantic filtering to specify world knowledge to the domain-dependent documents, and develop a new constrained HIN clustering algorithm to make better use of the structural information from the world knowledge for document clustering task.
- We conduct experiments on two benchmark datasets (20newsgroups and RCV1) to evaluate the clustering algorithm using HIN, compared with the state-of-the-art document clustering algorithms and clustering with “flat” world knowledge features. We show that our approach can be 13.3% better than the semi-supervised clustering algorithm incorporating 250K constraints which are generated by ground-truth labels.

This article is an extension of our previous work [Wang et al. 2015a]. We make the detailed algorithms clearer, and illustrate the effectiveness and efficiency of the algorithms with more extensive experiments.

The remainder of the article is organized as follows. Section 2 introduces the general learning framework of machine learning with world knowledge. Section 3 presents our world knowledge specification approach. The representation of world knowledge is introduced in Section 4. The model for document clustering with world knowledge is shown in Section 5. Experiments and results are discussed in Section 6. Section 7 discusses the related work, and we conclude this study in Section 8.

2. MACHINE LEARNING WITH WORLD KNOWLEDGE FRAMEWORK

In this section, we discuss the general framework on how we enable world knowledge to indirectly “supervise” machines, and give an overview on how to conduct document clustering in this framework. In general, performing machine learning with world knowledge, we should follow four steps.

- (1) *Knowledge acquisition.* World knowledge acquisition is a challenging problem. There exist some world knowledge bases. They are either collaboratively constructed by humans (such as Cyc project [Lenat and Guha 1989], Wikipedia, Freebase [Bollacker et al. 2008]) or automatically extracted from big data (such as KnowItAll [Etzioni et al. 2004], TextRunner [Banko et al. 2007], ReVerb [Fader et al. 2011], Ollie [Mausam et al. 2012], WikiTaxonomy [Ponzetto and Strube 2007], Probase [Wu et al. 2012], DBpedia [Auer et al. 2007], YAGO [Suchanek et al. 2007], NELL [Mitchell et al. 2015], and Knowledge Vault [Dong et al. 2014]). Since we assume the world knowledge is given, we skip this step in this study. We select two knowledge bases in this article, which are Freebase and YAGO2. Different knowledge bases have different characteristics. For example, Freebase is collaboratively collected. It focuses on named entities and their organizations. YAGO2 is automatically extracted from Wikipedia and mapped to WordNet. It will be interesting to compare the effects of using different world knowledge bases.
- (2) *Data adaptation.* Given the world knowledge, it is not necessary to use the whole knowledge base to perform inference since not all the world knowledge is related to the specific domains. Therefore, we should consider specifying the world knowledge to the domain-dependent data, and adapting the world knowledge to better characterize the specific domains. Moreover, since the knowledge can be ambiguous without context, we should consider using domain-dependent data to find the best knowledge to use. For example, when a text mentions “apple,” it can refer to a company or a fruit. In the knowledge base, we have both. Therefore, we should choose the right one to use. Notice that the general data adaptation process contains the disambiguation phase. For example, in Section 3.2, we describe the way to data adaptation for documents including the entity disambiguation phase [Li et al. 2013] (semantic filtering procedure). The traditional entity disambiguation problem is focusing on leveraging the purely context information, such as the co-occurrence of words/phrases appearing in certain window of the entity to disambiguate the entity, while we are considering to use more information from the world knowledge base, such as the types of the near entities and relations to disambiguate the entities. Besides, we further explore the semantic context of the entity by considering the relations between entities in the world knowledge bases. Thus, disambiguation is more on entity side in the original raw documents, whereas data adaptation is more general.
- (3) *Data and knowledge representation.* Traditionally, machine learning algorithms use feature vectors to represent data. Some interesting algorithms can represent data as trees or graphs, and compute kernel based on trees and graphs for machine learning [Collins and Duffy 2001; Vishwanathan et al. 2010]. Given the specified knowledge we have as well as the domain-dependent data, we should use a better representation which considers the structure information of the linked knowledge rather than just considering the knowledge as flat features. Therefore, we propose to use a typed graph, which is called HIN to represent the data.
- (4) *Learning.* After we have the representation, we can design a learning algorithm for the domain-dependent task. The learning algorithm is dependent on the problem as well as the data and knowledge representation. We will show how to handle the HIN for our clustering problem. Particularly, by representing the world knowledge as HIN, we have type information for the named entities we detected from the documents. Moreover, the world knowledge also provides the sub-type information. The coarse-grained type information is denser than the fine-grained sub-type information. For example, we can have much more entities annotated as *Person* than *Politician*. Thus, we use the coarse-grained type information to construct the

Fig. 2. Major steps of incorporating world knowledge into machine learning algorithms.

Fig. 3. The overview of adapting machine learning with world knowledge framework to document clustering.

HIN, and use the fine-grained sub-type information as further supervision for the entities, which is then used as indirect supervision for the documents.

The illustration of these four steps is shown in Figure 2, where steps two and three are sometimes dependent. For example, we can do domain adaptation and knowledge representation simultaneously. The above four steps are general, which means they may apply to many applications. In the following sections, we demonstrate how to select the right knowledge to use and to represent this knowledge for the task of document clustering. After that, we will introduce the learning algorithm to perform better document clustering given the representation.

In Figure 3, we show the general overview of the framework illustrating the major procedures of how to adapt the framework to the document clustering task. Notice that each module of the general learning framework shown in Figure 2 is directly specified for the particular document clustering with world knowledge approach (Figure 3), e.g., data adaptation is specified as world knowledge specification. Generally, we first assume the knowledge acquisition is done, i.e., the world knowledge bases (e.g., Freebase) are given. Second, robust unsupervised semantic parsing (Section 3.1) and three alternative semantic filters (Section 3.2) (e.g., conceptualization-based semantic filtering) are introduced for data adaptation. Third, we generate the new representation of the documents by modeling the unstructured texts in the HIN with entities including documents themselves, words in the document set, and relevant named entities and relations with proper types from the knowledge bases. Finally, for the learning phase in the framework, based on the new HIN representation of the documents, we propose the

constrained HIN clustering model to take the type and link information obtained from the knowledge bases into consideration via constraints and probabilistic distributions. In simpler terms, the first three steps are performed to construct the document-based HIN. Then, new learning algorithms (e.g., HIN-based clustering algorithms) are proposed to handle the document clustering in the new network representation and make use of the relevant knowledge specified from the world knowledge, to improve the performance of the traditional clustering algorithms.

3. WORLD KNOWLEDGE SPECIFICATION

In this section, we propose a world knowledge specification approach to generate specified world knowledge given a set of domain-dependent documents. We first use semantic parsing to ground any text to the knowledge base, and then provide three semantic filtering approaches to avoid ambiguity of the extracted information.

3.1. Semantic Parsing

Semantic parsing is the task of mapping a piece of natural language text to a formal meaning representation [Mooney 2007]. This can support question answering by querying a knowledge base [Kwiatkowski et al. 2011]. Most previous semantic parsing algorithms or tools developed are for small-scale problems but with complicated logical forms. More recently, large-scale semantic parsing grounding to world knowledge bases has been investigated, e.g., using Freebase [Krishnamurthy and Mitchell 2012; Cai and Yates 2013; Kwiatkowski et al. 2013; Berant et al. 2013; Berant and Liang 2014; Yao and Durme 2014; Reddy et al. 2014] or ReVerb [Fader et al. 2013]. These methods use semantic parsing to answer questions with the world knowledge bases. Intuitively, they need to identify the candidate answers in the parsing results so that they can be ranked to answer the questions. Therefore, they either need the question-answer pairs as supervision, or need a large amount of resources as well as the questions as distant/weak supervision [Reddy et al. 2014]. Similar to them, we are also working with very large-scale world knowledge bases, but unlike them, we do not match question and answers. We have already got all the entities in the document that can be matched to the world knowledge bases. Our task is then to ground the text to the knowledge base entities and their relationships in the prescribed logical form. Because we do not have and do not necessarily have the question-answer pairs, our problem is a fully unsupervised problem. Then, the remaining problems are two-folds. First, we need to identify the relations between entities to map them to the knowledge base relations. Second, we need to resolve the ambiguity of the entities and relations.

We first introduce the problem formulation and then introduce how we perform unsupervised semantic parsing. Let $\mathcal{E}$ be a set of entities and $\mathcal{R}$ be a set of relations in the knowledge base. Then, the knowledge base $\mathcal{K}$ consists of triplets in the form of (e_1, r, e_2) , where $e_1, e_2 \in \mathcal{E}$ and $r \in \mathcal{R}$. We follow [Berant et al. 2013] to use a simple version of Lambda Dependency-Based Compositional Semantics (λ -DCS) [Liang 2013] as the logic language to query the knowledge base. We use λ -DCS because that it can generate logic forms simpler than lambda calculus forms. λ -DCS can reduce the number of variables by making existential quantification implicit. The logical form in simple λ -DCS is either in the form of unary (a subset of $\mathcal{E}$) or binary (a subset of $\mathcal{E} \times \mathcal{E}$). We briefly introduce the definition of basic λ -DCS logical forms x and the corresponding denotations $x_{\mathcal{K}}$ as below: (1) Unary base: an entity $e \in \mathcal{E}$ is a unary logic form (e.g., *Obama*) with $e_{\mathcal{K}} = \{e\}$; (2) Binary base: a relation $r \in \mathcal{R}$ is a binary logic form (e.g., *PresidentofCountry*) with $r_{\mathcal{K}} = \{(e_1, e_2) : (e_1, r, e_2) \in \mathcal{K}\}$; (3) Join: $b.u$ is a unary logic form, denoting a join and projection, where b is a binary and e is a unary. $b.u_{\mathcal{K}} = \{e_1 \in \mathcal{E} : \exists e_2. (e_1, e_2) \in b_{\mathcal{K}} \wedge e_2 \in u_{\mathcal{K}}\}$ (e.g., *PresidentofCountry.Obama*); (4) Intersection:

Fig. 4. Semantic parsing example. The figure shows a derivation d of the input text “Obama is president of United States of America.” and the sub-derivations. Each is labeled with composition rule (in blue) and logical form (in red). The derivation d ignores words “is” and “of.”

$u_1 \sqcap u_2$ (u_1 and u_2 are both unaries) denotes set intersection: $u_1 \sqcap u_2 \kappa = u_1 \kappa \cap u_2 \kappa$ (e.g., $Location.Olympics \sqcap PresidentofCountry.Obama$).

We generally adopt the semantic parsing framework proposed in Berant et al. [2013]. Given a piece of text s , the semantic parser produces a distribution over possible derivations $D(s)$. Each derivation $d \in D(s)$ is a tree that indicates applying a certain set of composition rules that ends in the root of the tree, i.e., the logical form $d.z$. We illustrate the semantic parsing process with an example as shown in Figure 4. In simpler terms, the semantic parsing process can be understood as the following. First, given a piece of text “Obama is president of United States of America,” it maps the entities as well as the relation phrases in the text to knowledge base. So “Obama” and “United States of America” are mapped to knowledge base, resulting in two unary logic forms $People.BarackObama$ and $Country.USA$, where $People$ and $Country$ are the type information in Freebase. The relation phrase “president” is mapped to a binary logic form $PresidentofCountry$. Notice that, the mapping process skips the words “is” and “of.” The mapping dictionary is constructed by aligning a large text corpus to the knowledge base. A phrase and a knowledge base entity or relation can be aligned if they co-occur with many of the same entities. We select two knowledge bases, i.e., Freebase and YAGO2. For Freebase, we just use the mapping already existing in the released tool shown in Berant et al. [2013]. For YAGO2, we follow [Berant et al. 2013] and download a subset of ClueWeb09¹ to find the new mapping for YAGO2 entities and relations. Second, it uses some rules (i.e., grammar) to combine the basic logic forms to generate the derivations, and rank the results (i.e., derivations). In detail, the semantic parsing framework constructs a set of derivations for each span of the input text. First, for each span of the input text, it generates the single-predicate derivations based on the lexicon mapping from text to knowledge base predicates (e.g., “president” maps to $PresidentofCountry$). According to the set of composition rules, given any logical form x_1 and x_2 that is constructed over span $[i_1, j_1]$ and $[i_2, j_2]$, we then generate the logic forms based on the span $[\min(i_1, i_2), \max(j_1, j_2)]$ as the following: $x_1 \sqcap x_2$ (intersection), $x_1.x_2$ (join), or $x_1 \sqcap r.x_2$ (bridging) for any relation $r \in \mathcal{R}$ (bridging operation is defined to generate additional predicates based on neighboring predicates). For the example shown in this paragraph, $People.BarackObama \sqcap President.USA$ is generated to represent its semantic meaning. Notice that, $President.USA$ is generated by joining the unary $Country.USA$ with the binary $PresidentofCountry$. Figure 5 shows

¹<http://www.lemurproject.org/clueweb09.php/>.

Fig. 5. A real example of the input documents and output logic forms for 20 newsgroups dataset. Left: a set of given documents; middle: semantic parser; right: resulting candidate logic forms for each document. We only show some examples of the parsed logic forms according to entities “John Smoltz,” “Braves,” and “Bob Horner” in the given documents for simplicity.

a real example in 20 newsgroups dataset. Given the documents on the left side, the semantic parsing model is performed to generate the logic form candidates on the right side. Notice that, we only show some examples of the parsed logic forms according to entities “John Smoltz,” “Braves,” and “Bob Horner” in the given documents for simplicity.

When there are more than one candidate semantic meanings (i.e., derivations) for a sentence, in Berant et al. [2013], they learn the ranks based on the annotated question–answer pairs. For our task, this annotation is not available. Therefore, instead of ranking or enumerating all the possible logic forms (which is found to be not feasible in limited time), we constrain the entities to be the maximum length spanning phrases recognized by a state-of-the-art named entity recognition tool [Ratinov and Roth 2009]. We then perform the two steps introduced above by using the maximum length spanning noun phrase as entities, and use the phrase between them in the text as relation phrase. We propose to use the following three semantic filtering methods to resolve the ambiguity problem.

3.2. Semantic Filtering

For each sentence in the given document, the output of semantic parsing is a set of derivations that represent the semantic meaning. However, the extracted entities (i.e., unaries in the resulting derivations) can be ambiguous. For example, “apple” may be associated with type *Company* or *Fruit*. Therefore, we should filter out the noisy entities and their types to ensure that the knowledge we have is good enough as indirect supervision for document clustering. We assume that in the domain specific tasks, given the context, the entities seldom have multiple mutually exclusive meanings. Given that we have the domain-dependent corpus containing the documents to be clustered, we are given a lot of evidence to disambiguate the entities. We propose the following three approaches to select the best knowledge to use for further learning process.

Frequency-based semantic filter (FBSF). For each entity e_i in document d_j , which can have multiple types choosing from all the types $t_1, \dots, t_{\bar{T}}$, where we assume there are in total $\bar{T}$ types. Then, we can use the frequency $n_{d_j}(e_i, t_k)$ of a type t_k for an entity e_i appearing in a document d_j as the criterion to decide whether the entity should be extracted for the domain specific task in a sentence. Then, we use a threshold to cut the entity types that appear less than the frequency. Here, we assume that the

most frequent type(s) of an entity appearing in the document is the correct semantic meaning(s) in the context.

Document frequency based semantic filter (DFBSF). Similar to the frequency-based method, we use the document frequency $\sum_j I_{d_j}(e_i, t_k)$ of a type t_k of an entity e_i as the criterion to find the most likely semantic meaning, where $I_{d_j}(e_i, t_k) = 1$ if e_i in d_j is with type t_k , otherwise $I_{d_j}(e_i, t_k) = 0$. Here, we assume that if an entity appears in multiple documents with the same type, then the type should be the correct semantic meaning in the whole document collection.

Conceptualization-based semantic filter (CBSF). Motivated by the approaches of conceptualization [Song et al. 2011, 2015] and entity disambiguation [Li et al. 2013], we represent each entity with a feature vector $\mathbf{t}_i = (t_1, \dots, t_{\bar{T}})^T$ of entity types, and use standard k -means to cluster the entities in a document. Suppose in one cluster we have a set of entities $\mathcal{E} = \{e_1, \dots, e_{N_{\mathcal{E}}}\}$. We then use the probabilistic conceptualization proposed in Song et al. [2011] to find the most likely entity types for the entities in the cluster. We make the naive Bayes assumption and use

$$P(t_k|\mathcal{E}) \propto P(t_k) \prod_{i=1}^{N_{\mathcal{E}}} P(e_i|t_k) \quad (1)$$

as the score of entity type t_k . Here, $P(e_i|t_k) = \frac{n(e_i, t_k)}{n(t_k)}$, where $n(e_i, t_k)$ is the co-occurrence count of entity type t_k and entity e_i in the knowledge base, and $n(t_k)$ is the overall number of entities with type t_k in the knowledge base. Besides, $P(t_k) = \frac{n(t_k)}{\sum_k n(t_k)}$. Note that, in this formulation, we have $n(e_i, t_k) \equiv 1$ for Freebase and YAGO since the evidence in Freebase and YAGO is deterministic. We also want to leverage the information provided by the document or corpus. Therefore, we can also replace $n(e_i, t_k)$ with the frequency $n_{d_j}(e_i, t_k)$ or document frequency $\sum_j I_{d_j}(e_i, t_k)$ used in the previous two methods. In this case, we mean that we construct a sub-knowledge base with the edges being weighted by the evidence shown in the document or the corpus. The probability $P(t_k|\mathcal{E})$ is used to rank the entity types and the largest ones are selected. In this case, different entities in a document can be used to disambiguate each other. For each cluster, only the common types are retained, and concepts with conflicts are filtered out. Here, we also assume that the type that can best fit the context is the correct semantic meaning. Different from the FBSF method, we also consider the entity cluster information. Therefore, CBSF will use more accurate context information about the entity types. Figure 6 shows the semantic filtering results based on the example shown in Figure 5. “John Smoltz” is actually a baseball player, and “Braves” means Atlanta Braves, which is a baseball team playing in the national league. Note that, based on CBSF, the entity disambiguation problem is resolved to some extent. For example, “Bob Horner” could help “John Smoltz” disambiguate the correct entity type to be *baseball_player* from *tv_actor*, by using the context information provided in the first two documents. Similarly, the context information in the first and third documents could help to disambiguate the entity type of “Braves” to be *baseball_team*.

4. WORLD KNOWLEDGE REPRESENTATION

The output of semantic parsing and semantic filtering is then the document associated with the entities, which are further associated with the types (or concepts, categories, the names can be different for different knowledge bases and relations). For example, in Freebase, we select the top-level named entity categories (i.e., domains) as the types, e.g., *Person*, *Location*, and *Organization*. In addition to the named entities, we also

Fig. 6. Illustration of semantic filtering results based on the example shown in Figure 5. Left: candidate logic forms according to each document; middle: semantic filter; right: filtered semantics. “John Smoltz” is actually a baseball player, and “Braves” means the Atlanta Braves, which is a baseball team playing in national league. By using CBSF, the entity ambiguity problem is resolved to some extent. For example, “Bob Horner” could help “John Smoltz” disambiguate the correct entity type to be *baseball_player* from *tv_actor*, by using the context information provided in the first two documents. Similarly, the context information in the first and the third documents could help to disambiguate the entity type of “Braves” to be *baseball_team*.

regard the document and word as two types. Then, we use an HIN to represent the data we get after semantic parsing and semantic filtering.

Definition 4.1. An HIN is a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with an entity-type mapping $\phi: \mathcal{V} \rightarrow \mathcal{A}$ and a relation-type mapping $\psi: \mathcal{E} \rightarrow \mathcal{R}$, where $\mathcal{V}$ denotes the entity set and $\mathcal{E}$ denotes the link set, $\mathcal{A}$ denotes the entity-type set and $\mathcal{R}$ denotes the relation-type set, and the number of entity types $|\mathcal{A}| > 1$ or the number of relation types $|\mathcal{R}| > 1$.

The network schema provides a high-level description of a given heterogeneous information network.

Definition 4.2. Given an HIN $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with the entity type mapping $\phi: \mathcal{V} \rightarrow \mathcal{A}$ and the relation type mapping $\psi: \mathcal{E} \rightarrow \mathcal{R}$, the *network schema* for network G , denoted as $\mathcal{T}_{\mathcal{G}} = (\mathcal{A}, \mathcal{R})$, is a graph with nodes as entity types from $\mathcal{A}$ and edges as relation types from $\mathcal{R}$.

Then, for our world knowledge dependent network, we use the network schema shown in Figure 7 to represent the data. The network contains multiple entity types: *document* $\mathcal{D}$, *word* $\mathcal{W}$, *named entities* $\{\mathcal{E}^i\}_{i=1}^T$, and a few relation types connecting the entity types. Notice that, we use “entity type” to represent the node type in HIN, as Definition 4.1 showed. We use “named entity type” to represent the type of the name mentioned in text (widely used in NLP community), e.g., person, location, and organization names. The entities in HIN do not have to be named entities, e.g., the categories of animals or diseases. We denote the document set as $\mathcal{D} = \{d_1, d_2, \dots, d_M\}$, where M is the size of $\mathcal{D}$, the word set as $\mathcal{W} = \{w_1, w_2, \dots, w_N\}$, where N is the size of $\mathcal{W}$, and the entity set as $\mathcal{E}^t = \{e_1^t, e_2^t, \dots, e_{V_t}^t\}$, where V_t is the size of $\mathcal{E}^t$. We have $t = 1, \dots, T$, where T is the total number of named entity types we find in the knowledge base. Note that if there are no named entities, then the network reduces to a bipartite graph containing only documents and words.

In Figure 8, we show a real example of the world knowledge dependent network. From the figure, we can see that two documents represented as gray rectangles are modeled in the HIN. Besides the two documents, we have words represented as gray rounds. The link between a word and the corresponding document indicates that the

Fig. 7. Heterogeneous information network schema. The specified knowledge is represented in the form of heterogeneous information network. The schema contains multiple entity types: document $\mathcal{D}$, word $\mathcal{W}$, named entities $\{\mathcal{E}^i\}_{i=1}^T$, and the relation types connecting the entity types.

Fig. 8. Document-based heterogeneous information network example. The network $\mathcal{G}$ contains five entity types: document, word, date, person, and location, which are represented with gray rectangle, gray round, green square, blue round, and yellow triangle, respectively.

document contains the word. We also have named entities that are associated with certain types (e.g., date, person, and location), relevant to the documents specified from the knowledge base. The link between a document and a named entity means that the document contains the named entity. The link between two named entities represents the relation between the named entities in the knowledge base generated by world knowledge specification. In summary, after performing world knowledge specification for documents, world knowledge representation aims to construct a document-based HIN that contains the link and type information that are useful to understand the documents, thus, leading to better text mining performance.

5. DOCUMENT CLUSTERING WITH WORLD KNOWLEDGE

In this subsection, we present our clustering algorithm using HIN, constructed from the domain-dependent documents and the world knowledge. Given the HIN, it is natural to perform HIN partitioning to obtain the document clusters. In addition to the HIN itself, let us revisit the structural information in a typical world knowledge base, e.g., Freebase. In the world knowledge base, the named entities are often organized in a hierarchy of categories. Although there is additional category information for each entity, we only use the top-level named entity types as the entity types in HIN.

For example, “Barack Obama” is a person, where person is the top-level category. In addition, he is the president of the “United States,” a politician, a celebrity, and the like. Another example is that “Google” is a software company, plus it has a CEO. This shows that the entities can have some attributes. We choose to use top-level entity types for the HIN schema since then we will have a relatively dense graph for each pairwise nodes in the network schema. The fine-grained named entity sub-types or the attributes are also very useful to identify the topics or the clusters of the documents. Therefore, in this section, we introduce how we incorporate the fine-grained level of named entity types as constraints in the HIN clustering algorithm.

5.1. Constrained Clustering Modeling

To formulate the clustering algorithm for the domain-dependent documents, we denote latent label sets of the documents as $\mathcal{L}_d = \{l_{d_1}, l_{d_2}, \dots, l_{d_M}\}$. We also denote $\mathcal{L}_w = \{l_{w_1}, l_{w_2}, \dots, l_{w_N}\}$ for words, and $\mathcal{L}_{e^t} = \{l_{e_1^t}, l_{e_2^t}, \dots, l_{e_{V_t}^t}\}$ for the t th named entities set. In general, we follow the framework of information-theoretic co-clustering (ITCC) [Dhillon et al. 2003] and constrained ITCC [Song et al. 2013; Wang et al. 2015c] to formulate our approach. Instead of only performing on the bipartite graph, we need to handle multi-type relational data, as well as more complicated constraints.

The original ITCC uses a variational function to approximate the joint probability of documents and words, which is

$$q(d_m, w_i) = p(\hat{d}_{k_d}, \hat{w}_{k_w})p(d_m|\hat{d}_{k_d})p(w_i|\hat{w}_{k_w}), \quad (2)$$

where $\hat{d}_{k_d}$ and $\hat{w}_{k_w}$ are cluster indicators to formulate the conditional probability, and k_d and k_w are the corresponding cluster indices. $q(d_m, w_i)$ is used to approximate $p(d_m, w_i)$ by minimizing the Kullback–Leibler (KL) divergence:

$$\begin{aligned} D_{KL}(p(\mathcal{D}, \mathcal{W})||q(\mathcal{D}, \mathcal{W})) \\ &= D_{KL}(p(\mathcal{D}, \mathcal{W}, \hat{\mathcal{D}}, \hat{\mathcal{W}})||q(\mathcal{D}, \mathcal{W}, \hat{\mathcal{D}}, \hat{\mathcal{W}})) \\ &= \sum_{k_d}^{K_d} \sum_{d_m: l_{d_m}=k_d} p(d_m) D_{KL}(p(\mathcal{W}|d_m)||p(\mathcal{W}|\hat{d}_{k_d})) \\ &= \sum_{k_w}^{K_w} \sum_{w_i: l_{w_i}=k_w} p(w_i) D_{KL}(p(\mathcal{D}|w_i)||p(\mathcal{D}|\hat{w}_{k_w})), \end{aligned} \quad (3)$$

where $\hat{\mathcal{D}}$ and $\hat{\mathcal{W}}$ are the cluster sets, $p(\mathcal{W}|\hat{d}_{k_d})$ denotes a multinomial distribution based on the probabilities

$$p(\mathcal{W}|\hat{d}_{k_d}) = (p(w_1|\hat{d}_{k_d}), \dots, p(w_N|\hat{d}_{k_d}))^T,$$

$$p(w_i|\hat{d}_{k_d}) = p(w_i|\hat{w}_{k_w})p(\hat{w}_{k_w}|\hat{d}_{k_d}), \text{ and}$$

$$p(w_i|\hat{w}_{k_w}) = p(w_i)/p(l_{w_i} = \hat{w}_{k_w}).$$

Symmetrically, we have

$$p(\mathcal{D}|\hat{w}_{k_w}) = (p(d_1|\hat{w}_{k_w}), \dots, p(d_M|\hat{w}_{k_w}))^T,$$

$$p(d_i|\hat{w}_{k_w}) = p(d_i|\hat{d}_{k_d})p(\hat{d}_{k_d}|\hat{w}_{k_w}), \text{ and}$$

$$p(d_i|\hat{d}_{k_d}) = p(d_i)/p(l_{d_i} = \hat{d}_{k_d}).$$

Table I. Notations for Clustering Algorithm

Meaning	Document	Word	Named Entity
Cluster Index	k_d	k_w	k_{e^t}
Cluster Indicator	$\hat{d}_{k_d}$	$\hat{w}_{k_w}$	$\hat{e}_{k_{e^t}}^t$
Data Indicator	d_m	w_i	e_i^t
Data Indicator Set	$\mathcal{D}$	$\mathcal{W}$	$\mathcal{E}^t$
Label	l_{d_m}	l_{w_i}	$l_{e_i^t}$
Label Indicator Set	$\mathcal{L}_d$	$\mathcal{L}_w$	$\mathcal{L}_{e^t}$

Moreover, $p(\hat{w}_{k_w}|\hat{d}_{k_d})$ and $p(\hat{d}_{k_d}|\hat{w}_{k_w})$ are computed based on the joint probability $q(\hat{d}_{k_d}, \hat{w}_{k_w}) = \sum_{l_{d_m}=k_d} \sum_{l_{w_i}=k_w} p(d_m, w_i)$.

Motivated by ITCC, according to the network schema shown in Figure 7, our problem of HIN clustering is formulated as

$$\begin{aligned}
 \mathcal{J}_{\text{HINC}} = & D_{KL}(p(\mathcal{D}, \mathcal{W})||q(\mathcal{D}, \mathcal{W})) \\
 & + \sum_{t=1}^T D_{KL}(p(\mathcal{D}, \mathcal{E}^t)||q(\mathcal{D}, \mathcal{E}^t)) \\
 & + \sum_{t=1}^T \sum_{s=1}^T D_{KL}(p(\mathcal{E}^t, \mathcal{E}^s)||q(\mathcal{E}^t, \mathcal{E}^s)),
 \end{aligned} \tag{4}$$

where the equation aims to leverage the link and type information (documents, words, types of named entities) in the HIN shown in Figure 7 to estimate the cluster label of each document. The first part of the problem, $D_{KL}(p(\mathcal{D}, \mathcal{W})||q(\mathcal{D}, \mathcal{W}))$, according to Equation (3), is exactly the original ITCC on the document–word bipartite graph (the top portion of Figure 7 with two types of entities, document and word). It aims to use a variational function $q(\mathcal{D}, \mathcal{W})$ to approximate the joint probability of documents and words $p(\mathcal{D}, \mathcal{W})$ measured by KL divergence, to minimize the mutual information loss due to the mapping to generate the cluster indices of the documents and words as shown in Dhillon et al.’s [2003, Lemma 2.1]. Similarly, the second part $\sum_{t=1}^T D_{KL}(p(\mathcal{D}, \mathcal{E}^t)||q(\mathcal{D}, \mathcal{E}^t))$ means to use the variational function $q(\mathcal{D}, \mathcal{E}^t)$ to approximate the joint probability of documents and named entities that belong to every type t . This equals to perform ITCC on every document-named entity of type t bipartite graph (the middle portion of Figure 7). The third part is also defined similarly, which defines the variational function $q(\mathcal{E}^t, \mathcal{E}^s)$ to approximate the joint probability of named entities of every type t and named entities that belong to every type s . This is equivalent to performing ITCC on every named entity of type t -named entity of type s bipartite graph (the bottom portion of Figure 7). All the probabilities of the second and third parts can be similarly defined as the first part, document–word bipartite graph. We omit the detailed definitions for brevity. A summary of the notations is shown in Table I.

To incorporate the side information of the fine-grained named entity sub-types or the attributes as indirect supervision for document clustering, we define the constraints for the named entities we find after semantic parsing. We take the t th entity label set $\mathcal{E}^t$ as an example, and use must-links and cannot-links as the constraints. We denote the must-link set associated with $e_{i_1}^t$ as $\mathcal{M}_{e_{i_1}^t}$, and the cannot-link set as $\mathcal{C}_{e_{i_1}^t}$. The way how we build must-links and cannot-links is described in the experiment (Section 6.4.3). For must-links, the cost function is defined as

$$\begin{aligned}
 & V_{\mathcal{M}}(e_{i_1}^t, e_{i_2}^t \in \mathcal{M}_{e_{i_1}^t}) \\
 & = w_{\mathcal{M}} D_{KL}(p(\mathcal{D}|e_{i_1}^t)||p(\mathcal{D}|e_{i_2}^t)) \cdot \mathcal{I}_{l_{e_{i_1}^t}} \neq l_{e_{i_2}^t},
 \end{aligned} \tag{5}$$

where w_M is the weight for must-links, and $p(\mathcal{D}|e_{i_1}^t)$ denotes a multinomial distribution based on the probabilities $(p(d_1|e_{i_1}^t), \dots, p(d_M|e_{i_1}^t))^T$, and $\mathcal{I}_{true} = 1, \mathcal{I}_{false} = 0$. The above must-link cost function means that if the label of $e_{i_1}^t$ is not equal to the label of $e_{i_2}^t$, then we should take into account the cost function of how dissimilar the two entities $e_{i_1}^t$ and $e_{i_2}^t$ are. The dissimilarity is computed based on the probability of document $\mathcal{D}$ given the entities $e_{i_1}^t$ and $e_{i_2}^t$ as Equation (5). The more dissimilar the two entities are, the larger the cost is imposed. Please refer to the experimental section (Section 6.4) for details about the weight setting (w_M) for the must-links.

For cannot-links, the cost function is defined as

$$\begin{aligned} V_C(e_{i_1}^t, e_{i_2}^t \in \mathcal{C}_{e_{i_1}^t}) \\ = w_C(D_{max}^t - D_{KL}(p(\mathcal{D}|e_{i_1}^t)||p(\mathcal{D}|e_{i_2}^t))) \cdot \mathcal{I}_{e_{i_1}^t \neq e_{i_2}^t}, \end{aligned} \quad (6)$$

where w_C is the weight for cannot-links, and D_{max}^t is the maximum value for all the $D_{KL}(p(\mathcal{D}|e_{i_1}^t)||p(\mathcal{D}|e_{i_2}^t))$. The cannot-link cost function means that if the label of $e_{i_1}^t$ is equal to the label of $e_{i_2}^t$, then we should take into account the cost function of how similar they are. Also, please refer to the experimental section (Section 6.4) for details about how we set w_C for the cannot-links.

Integrating the constraints for $\mathcal{L}_{e^1}, \dots, \mathcal{L}_{e^T}$ to Equation (4), the objective function of constrained HIN clustering is

$$\begin{aligned} \mathcal{J}_{CHINC} = & D_{KL}(p(\mathcal{D}, \mathcal{W})||q(\mathcal{D}, \mathcal{W})) \\ & + \sum_{t=1}^T D_{KL}(p(\mathcal{D}, \mathcal{E}^t)||q(\mathcal{D}, \mathcal{E}^t)) \\ & + \sum_{t=1}^T \sum_{s=1}^T D_{KL}(p(\mathcal{E}^t, \mathcal{E}^s)||q(\mathcal{E}^t, \mathcal{E}^s)) \\ & + \sum_{t=1}^T \sum_{e_{i_1}^t=1} \sum_{e_{i_2}^t \in \mathcal{M}_{e_{i_1}^t}} V_M(e_{i_1}^t, e_{i_2}^t \in \mathcal{M}_{e_{i_1}^t}) \\ & + \sum_{t=1}^T \sum_{e_{i_1}^t=1} \sum_{e_{i_2}^t \in \mathcal{C}_{e_{i_1}^t}} V_C(e_{i_1}^t, e_{i_2}^t \in \mathcal{C}_{e_{i_1}^t}). \end{aligned} \quad (7)$$

From this objective function we can see that the must-links and cannot-links are imposed to the entities that the semantic parsing detects. Since the task is document clustering, the sub-types of entities serve as indirect supervision because they cannot directly affect the cluster labels of the documents. However, the constraints can affect the labels of entities, and then the labels of entities can be transferred to the document side to affect the labels of documents.

5.2. Alternating Optimization

Since global optimization of all the latent labels as well as the approximate function $q(\cdot, \cdot)$ is intractable, we perform an alternating optimization shown in Algorithm 1. We iterate the process to optimize the labels of documents, words, and entities. Meanwhile, we update the function $q(\cdot, \cdot)$ for the corresponding types.

ALGORITHM 1: Alternating Optimization for CHINC

Input: HIN defined ondocuments $\mathcal{D}$, words $\mathcal{W}$, and entities

$\mathcal{E}^t, t = 1, \dots, T$; Set maxIter and $\text{max}\delta$.

while $\text{iter} < \text{maxIter}$ and $\delta > \text{max}\delta$ **do**

$\mathcal{D}$ **Label Update:** minimize Equation (8) w.r.t. $\mathcal{L}_d$.

$\mathcal{D}$ **Model Update:** update $q(d_m, w_i)$ and $q(d_m, e_i^t)$.

for $t = 1, \dots, T$ **do**

$\mathcal{E}^t$ **Label Update:** minimize Equation (10) w.r.t. $\mathcal{L}_{e^t}$.

$\mathcal{E}^t$ **Model Update:** update $q(d_m, e_i^t)$ and $q(e_j^s, e_i^t)$.

end for

$\mathcal{D}$ **Label Update:** minimize Equation (8) w.r.t. $\mathcal{L}_d$.

$\mathcal{D}$ **Model Update:** update $q(d_m, w_i)$ and $q(d_m, e_i^t)$.

$\mathcal{W}$ **Label Update:** minimize Equation (9) w.r.t. $\mathcal{L}_w$.

$\mathcal{W}$ **Model Update:** update $q(d_m, w_i)$.

 Compute cost change δ using Equation (7).

end while

For example, to find label l_{d_m} of document d_m , we have

$$l_{d_m} = \arg \min_{l_{d_m}=k_d} D_{KL}(p(\mathcal{W}|d_m)||p(\mathcal{W}|\hat{d}_{k_d})) + \sum_{t=1}^T D_{KL}(p(\mathcal{E}^t|d_m)||p(\mathcal{E}^t|\hat{d}_{k_d})). \quad (8)$$

To find label l_{w_i} of word w_i , we have

$$l_{w_i} = \arg \min_{l_{w_i}=k_w} D_{KL}(p(\mathcal{D}|w_i)||p(\mathcal{D}|\hat{w}_{k_w})). \quad (9)$$

To find the label $l_{e_i^t}$, we use the iterated conditional mode (ICM) algorithm [Basu et al. 2004] to iteratively assign a label to the entity. We update one label $l_{e_i^t}$ at a time, and keep all the other labels fixed:

$$l_{e_i^t} = \arg \min_{l_{e_i^t}=k_{e^t}} D_{KL}(p(\mathcal{D}|e_i^t)||p(\mathcal{D}|\hat{e}_{k_{e^t}}^t)) + \sum_{s=1}^T D_{KL}(p(\mathcal{E}^s|e_i^t)||p(\mathcal{E}^s|\hat{e}_{k_{e^t}}^t)) + \sum_{\substack{e_{i'}^t \in \mathcal{M}_{e_i^t}; \\ l_{e_i^t} \neq l_{e_{i'}^t}}} w_{\mathcal{M}} D_{KL}(p(\mathcal{D}|e_i^t)||p(\mathcal{D}|e_{i'}^t)) + \sum_{\substack{e_{i'}^t \in \mathcal{C}_{e_i^t}; \\ l_{e_i^t} = l_{e_{i'}^t}}} w_{\mathcal{C}} (D_{max}^t - D_{KL}(p(\mathcal{D}|e_i^t)||p(\mathcal{D}|e_{i'}^t))). \quad (10)$$

To transfer the original objective function (7) to Equation (10), we should follow Equation (3), where we replace the document and word notations to the entity notations. To understand why Equation (3) holds, we suggest to refer to the original ITCC for detailed derivation [Dhillon et al. 2003].

Then, with the labels $\mathcal{L}_d$, $\mathcal{L}_{e^t}$, and $\mathcal{L}_w$ fixed, we update the model function $q(d_m, w_i)$, $q(d_m, e_i^t)$, and $q(e_j^s, e_i^t)$. The update of q is not influenced by the must-links and cannot-links. Thus, we can modify them the same as ITCC [Dhillon et al. 2003] and only show

the update of $q(d_m, e_i^t)$ here:

$$q(\hat{d}_{k_d}, \hat{e}_{k_e^t}^t) = \sum_{l_{d_m}=k_d} \sum_{l_{e_i^t}=k_e^t} p(d_m, e_i^t), \quad (11)$$

$$q(d_m | \hat{d}_{k_d}) = \frac{q(d_m)}{q(l_{d_m}=k_d)} [q(d_m | \hat{d}_{k_d}) = 0 \text{ if } l_{d_m} \neq k_d], \quad (12)$$

$$q(e_i^t | \hat{e}_{k_e^t}^t) = \frac{q(e_i^t)}{q(l_{e_i^t}=k_e^t)} [q(e_i^t | \hat{e}_{k_e^t}^t) = 0 \text{ if } l_{e_i^t} \neq k_e^t], \quad (13)$$

where $q(d_m) = \sum_{e_i^t} p(d_m, e_i^t)$, $q(e_i^t) = \sum_{d_m} p(d_m, e_i^t)$, $q(\hat{d}_{k_d}) = \sum_{k_e^t} p(\hat{d}_{k_d}, \hat{e}_{k_e^t}^t)$ and $q(\hat{e}_{k_e^t}^t) = \sum_{k_d} p(\hat{d}_{k_d}, \hat{e}_{k_e^t}^t)$.

Algorithm 1 summarizes the main steps in the procedure. The objective function (7) with our alternating update monotonically decreases to a local optimum. This is because the ICM algorithm decreases the non-negative objective function (7) to a local optimum given a fixed q function. Then, the update of q is monotonically decreasing as guaranteed by the theorem proven in Song et al. [2013]. Besides, the original proof of the decrease of q is shown in Dhillon et al. [2003, Theorem 4.1].

The time complexity of Algorithm 1 is $O(n_{\mathcal{D}, \mathcal{W}} \cdot (K_d + K_w) + \sum_{t=1}^T n_{\mathcal{D}, \mathcal{E}^t} \cdot (K_d + K_{e^t}) + \sum_{t=1}^T \sum_{s=1}^T (n_{\mathcal{E}^t, \mathcal{E}^s} + (n_c * \text{iter}_{ICM})) \cdot (K_{e^t} + K_{e^s})) \cdot \text{iter}_{AO}$, where $n_{\cdot}$ is the total number of non-zero elements in the corresponding co-occurrence matrix, n_c is the number of constraints, iter_{ICM} is the number of ICM iterations, K_d , K_w , and K_{e^t} are respectively the number of document clusters, word clusters, and entity clusters of type t , and iter_{AO} is the number of the alternating optimization iterations.

Discussion. The major factors contributing to the time complexity of Algorithm 1 are the numbers of non-zero elements in the corresponding matrices. We have about 20,000 documents and 60,000 words in 20NG. From Figure 10, we find around 20,000 entities specified from Freebase with 79 types of entities. In contrast, in 20NG, the number of document clusters, word clusters, and named entity clusters is empirically set to be 20, 40, and 79, according to the number of categories of documents, twice the number of document clusters following Song et al. [2013], and the total number of top-level named entity types in Freebase, respectively. Compared to the number of non-zero elements in the matrix (e.g., hundreds of thousands), the number of relevant clusters can be ignored. In such situation, if the network schema contains more types of named entities, the running time of the algorithm will not increase that much, compared to the impacts caused by the size of document sets and vocabulary, as well as the named entity sets. There could be lots of ways to improve the performance of machine learning algorithms in large-scale datasets, e.g., distributed computation. However, as this is beyond the scope of this article, we leave it for future work.

6. EXPERIMENTS

In this section, we show the experimental results to demonstrate the effectiveness and efficiency of our approach on document clustering with world knowledge as indirect supervision.

6.1. Datasets

We use the following two benchmark datasets to evaluate domain-dependent document clustering. For both datasets, we assume the numbers of document clusters are given.

Table II. RCV1 Dataset Statistics

	#(Categories)	#(Leaf Categories)	#(Documents)
MCAT	9	7	44,033
CCAT	31	26	47,494
ECAT	23	18	19,813

#(Categories) is the number of all categories; #(Leaf Categories) is the number of leaf categories; #(Documents) is the number of documents.

20Newsgroups (20NG): The 20newsgroups dataset [Lang 1995] contains about 20,000 newsgroups documents evenly distributed across 20 newsgroups.² We use all the 20 groups as 20 classes.

RCV1: The RCV1 dataset is a dataset containing manually labeled newswire stories from Reuters Ltd. [Lewis et al. 2004]. The news documents are categorized with respect to three controlled vocabularies: industries, topics, and regions. There are 103 categories including all nodes except for root in the hierarchy. The maximum depth is four, and 82 nodes are leaves. We select top categories MCAT (Markets), CCAT (Corporate/Industrial), and ECAT (Economics) in one portion of the test partition to form three clustering tasks. The three clustering tasks are summarized in Table II. We use the original source of this data, and use the leaf categories in each task as the ground-truth classes.

6.2. World Knowledge Bases

Then, we introduce the knowledge bases we use.

Freebase: Freebase³ is a publicly available knowledge base consisting of entities and relations collaboratively collected by its community members. Now, it contains more than 2 billion relation expressions between 40 million entities. We convert a logical form generated by our unsupervised semantic parser of the world knowledge specification approach introduced in Section 3 into a SPARQL query and execute it on our copy of Freebase using the Virtuoso engine.

YAGO2: YAGO2⁴ is also a semantic knowledge base, derived from Wikipedia, WordNet, and GeoNames. Currently, YAGO2 has knowledge of more than 10 million entities (such as persons, organizations, cities, etc.) and contains more than 120 million facts about these entities. Similar to Freebase, we also convert a logical form into a SPARQL query and execute it on our copy of YAGO2 using the Virtuoso engine.

In Table III, we show some statistics about Freebase and YAGO2.

Note that in most knowledge bases, such as Freebase and YAGO2, entities types are often organized in a hierarchical manner. For example, *Politician* is a sub-type of *Person*. *University* is a sub-type of *Organization*. All the types or attributes share a common root, called *Object*. Figure 9 depicts an example of hierarchy of types. In general, we use the highest level under the root object as the entity types (e.g., *Person*) as specified world knowledge incorporated in the HIN, and the direct children (e.g., *Politician*) as entity constraints. In the following experiments, we select *Person*, *Organization*, and *Location* as the three entity types in the HIN, because they are popular in both Freebase and YAGO2.

Discussion. In all the following experiments, we conduct experiments based on different combinations of world knowledge bases (i.e., Freebase and YAGO2) introduced in this section with the four datasets (i.e., 20NG, MCAT, CCAT, and ECAT) described

²<http://qwone.com/~jason/20Newsgroups/>.

³<https://developers.google.com/freebase/>.

⁴<http://www.mpi-inf.mpg.de/departments/databases-and-information-systems/research/yago-naga/yago/>.

Table III. Statistics of Freebase and YAGO2

Name	Freebase	YAGO2
#(Entity Types)	1,500 ^a	350,000
#(Entity Instances)	40 million	10 million
#(Relation Types)	35,000	100
#(Relation Instances)	2 billion	120 million

#(Entity Types) is the number of entity types; #(Entity Instances) is the number of entity instances; #(Relation Types) is the number of relation types; #(Relation Instances) is the number of relation instances.

^aThe number of 1,500 types is reported in Dong et al. [2014]. In our downloaded dump of Freebase, we found 79 domains, 2,232 types, and 6,635 properties.

Fig. 9. Hierarchy of entity types.

in Section 6.1. In total, we illustrate eight different world knowledge base and document dataset combinations for both world knowledge specification and the clustering algorithms. For example, we show eight clustering results of our clustering algorithm named constrained heterogeneous information network clustering (CHINC) in Table VI(b). As shown in the results, for instance, the clustering result of CHINC on 20NG dataset using Freebase as world knowledge source is 0.631.

6.3. Effectiveness of World Knowledge Specification

Before applying the specified world knowledge to downstream text analytics tasks, such as document clustering in our case, we need to evaluate whether our world knowledge specification approach could produce the correct specified world knowledge.

In order to test the effectiveness of our world knowledge specification approach, we first sample 200 documents from 20 newsgroups, i.e., 10 documents from each category. Second, we split the documents into sentences. After post-processing, 3,232 sentences are generated for human evaluation. Third, we use our world knowledge specification approach in Section 3 with three different semantic filtering modules to generate the specified world knowledge for each sentence, which consists of relation triplets in the form of (e_1, r, e_2) with the type information. Notice that we use Freebase as the world knowledge source. For FBSF, in each document, we decide the type of an entity in that document by choosing the one with the largest frequency in the document according to Section 3.2; For DFBSF, similar to that of FBSF, the type with the largest document frequency is selected as the correct type of the entity in all the documents; For CBSF, we set the number of entity type clusters as 79, which is the number of top types (in Freebase, i.e., domains) as shown in Table III, since we assume the document set including all the possible entity types in the world knowledge base. Afterwards, we ask three annotators to label the specified world knowledge according two criteria:

Table IV. Precision of Different Semantic Filtering Results

Semantic filter	FBSF	DFBSF	CBSF
Precision	0.751	0.890	0.916

FBSF represents frequency-based semantic filter; DFBSF represents document frequency-based semantic filter; CBSF represents conceptualization-based semantic filter.

Table V. Error Analysis of Specified World Knowledge Generated by the World Knowledge Specification Approach with Three Different Semantic Filters

Type of error	Example sentence	Number and percentage of errors		
		FBSF (805)	DFBSF (359)	CBSF (272)
Entity recognition	"Einstein's theory of relativity explained mercury's motion."	179 (22.2%)	129 (35.9%)	105 (38.6%)
Entity disambiguation	"Bill said all this to make the point that Christianity is eminently."	537 (66.7%)	182 (50.7%)	130 (47.8%)
Subordinate clause	"Bruce S. Winters, worked at United States Technologies Research Center, bought a Ford."	89 (11.1%)	48 (13.4%)	37 (13.6%)

FBSF represents frequency-based semantic filter; DFBSF represents document frequency-based semantic filter; CBSF represents conceptualization-based semantic filter.

(1) whether the boundaries of e_1 and e_2 are correctly recognized or not; (2) whether the entity type of e_1 and e_2 are correct or not. It is annotated as correct if both (1) and (2) are satisfied. We check the mutual agreement of the human annotation, which is around 91.3% accuracy.

We then test the precision of three different specified world knowledge generated by the corresponding semantic filtering method. The results are shown in Table IV. From the results we can see that CBSF outperforms the other two ways to generate the correct semantic meaning. The main reason is that conceptualization-based method is able to use the context information to help judge the real semantic of the text rather than only taking the statistics of the data into account. Here, we only care about precision because we wish to use world knowledge as indirect supervision. The recall will not be very important.

Error analysis. To further investigate what triggers the errors in our semantic parsing and semantic filtering pipelines, we analyze the cause of errors for the incorrect specified world knowledge. The errors are collected from the error cases based on the annotation generated by the three annotators. Then, we ask the annotators to try to classify the errors. Finally, we summarize the following three error categories as shown in Table V.

Entity recognition: In semantic parsing, entities can be extracted incorrectly. Long entities are composed of multiple simple entities. For example, "Einstein's theory of relativity" may be extracted as "Einstein" and "theory of relativity." Paraphrasing and misspelling entities cause their textual expressions to deviate from any knowledge base entries. Idiomatic expressions are incorrectly picked up as entities. Using a larger mapping from text to knowledge base phrases, or paraphrasing techniques will help avoid some errors. However, this is out of the scope of this article.

Entity disambiguation: Selecting an incorrect entity out of multiple matching candidates causes this error, e.g., "Bill" in our example sentence can be "Bill Clinton" or "Bill Gates." Primarily due to two reasons: first, entity disambiguation is a tough research

problem in NLP community. Second, the type information of relations is not sufficient to further prune out mismatching entities during the semantic filtering process. Notice that, entity disambiguation is the major cause of the errors. By using CBSF, the number of incorrect entities caused by disambiguation can be dramatically reduced.

Subordinate clause: Semantic parsing sometimes produces wrong relation phrases in the subordinate clauses. For example, in the example wrong case shown in Table V, it takes the relation phrase “worked at” meaning the working place of “Bruce S. Winters,” ignores the phrase “bought,” which could be more informative for the target clustering domain. This could be resolved by adding more concrete rules in the semantic parsing grammar.

Discussion. Notice that both 20NG and RCV1 have relatively larger number of named entities specified from knowledge bases; thus, we show more significantly improved clustering results. When encountering the text with little named entities, our algorithm would be very similar to that of the original CITCC since the schema of the document-based HIN constructed in Sections 3 and 4 would tend to be similar to the schema with only two types of entities, i.e., documents and entities. Besides, CITCC only uses the constraints built upon the words and documents to cluster the documents, whereas our approach explores the constraints built based on not only documents and words, but also the named entities from the knowledge base.

In the following experiments, we use the world knowledge specification approach with CBSF, because it performs the best among the three semantic filtering methods.

6.4. Clustering Result

In this experiment, we compare the performance of our model, CHINC, with several representative clustering algorithms such as k -means, ITCC [Dhillon et al. 2003] and CITCC [Song et al. 2013]. The parameters used in CHINC to control the constraints are w_M and w_C . We set them as $\frac{1}{\sum_{t=1}^T |\mathcal{E}^t|}$ following the rules tested in Song et al. [2013], where $|\mathcal{E}^t|$ is the number of entities in $\mathcal{E}^t$.⁵ We also denote our algorithm without constraints as HINC. For both CHINC and HINC, we set the numbers of document clusters, word clusters, and named entity clusters according to the numbers of categories of the document set, twice the number of document clusters following [Song et al. 2013], and the total number of top-level named entity types in the world knowledge base, respectively. The number of constraints used in each combination of document set and knowledge base is the same, which is equal to 3×10^8 . The constraints are randomly selected from all the constraints generated by the method introduced in Section 6.4.3. “FB” and “YG” represent two different world knowledge sources, Freebase and YAGO2, respectively. We re-implement all the above clustering algorithms. Notice that, for CITCC, we follow Song et al. [2013] to generate and add constraints for documents and words. We also use the specified world knowledge as features to enhance the k -means and ITCC. The feature settings are defined as below:

- BOW: Traditional bag-of-words model with the tf-idf weighting mechanism.
- BOW+FB: BOW integrated with additional features from entities in specified world knowledge of Freebase.
- BOW+YG: BOW integrated with additional features from entities in specified world knowledge of YAGO2.

We employ the widely used normalized mutual information (NMI) [Strehl and Ghosh 2003] as the evaluation measure. The NMI score is 1 if the clustering results match the

⁵In Song et al. [2013], the experiment shows the parameter study on the weights of the constraints varying from $1E-8$ to 100, and conclude this is one of the best settings.

Table VI. Performance of Different Clustering Algorithms on 20NG and RCV1 data

(a) Performance of k -means, ITCC, and CITCC with different features.

Features Data	k -means			ITCC			CITCC
	BOW	BOW +FB	BOW +YG	BOW	BOW +FB	BOW +YG	BOW
20NG	0.429	0.447	0.437	0.501	0.525	0.513	0.569
MCAT	0.549	0.575	0.559	0.604	0.630	0.619	0.652
CCAT	0.403	0.419	0.410	0.472	0.494	0.481	0.535
ECAT	0.417	0.436	0.424	0.493	0.516	0.505	0.562

(b) Performance of HINC and CHINC with different world knowledge sources.

Features Data	HINC		CHINC	
	FB	YG	FB	YG
20NG	0.571 (+0.4%)	0.541 (−4.9%)	0.631 (+10.9%)	0.600 (+5.5%)
MCAT	0.645 (−1.1%)	0.625 (−4.1%)	0.698 (+7.1%)	0.685 (+5.1%)
CCAT	0.542 (+1.3%)	0.515 (−3.7%)	0.606 (+13.3%)	0.574 (+7.3%)
ECAT	0.561 (−0.2%)	0.530 (−5.7%)	0.624 (+11.0%)	0.588 (+4.6%)

CHINC is our proposed method. BOW, FB (Freebase), or YG (YAGO2) represent bag-of-words features, the entities generated by our world knowledge specification approach based on Freebase or YAGO2, respectively. We compared all the numbers of HINC and CHINC with CITCC, which is the strongest baseline. The percentages in the brackets are the relative numbers compared to CITCC. CITCC uses 250K constraints generated based on ground-truth labels of documents. With t-test, we find the improvements are significant at 0.05 significance level.

category labels perfectly and 0 if the clusters are obtained from a random partition. In general, the larger the scores, the better the clustering results.

In Table VI, we show the performance of all the clustering algorithms with different experimental settings. The NMI is the average NMI of five random trials per experiment setting. Overall, among all the methods we test, CHINC consistently performs the best among all the clustering methods. We can see that HINC+FB and HINC+YG perform better than ITCC with BOW+FB or BOW+YG features, respectively. This means that by using the structural information provided by the world knowledge, we can further improve the clustering results. In addition, the algorithms with Freebase consistently outperform the ones with YAGO2 since Freebase has much more facts compared with YAGO2 as shown in Table III; besides, one can see in Figure 10 that Freebase could consistently specify more entities than YAGO2 does from all of the document datasets. CITCC is the strongest baseline clustering algorithm, because it uses the ground-truth constraints derived from category labels based on the human knowledge. We use 250K constraints to perform CITCC. As shown in Table VI, HINC performs competitive with the CITCC. CHINC significantly outperforms CITCC. This shows that by automatically using world knowledge, it has the potential to perform better than the algorithm with the specific domain knowledge.

Discussion. Based on the results, we can see that, when performing existing clustering algorithms with world knowledge, it is better to choose the world knowledge sources including relatively larger number of instances of entities and relations. Since, in general, the more entities and relations one knowledge base has, the bigger possibility that more useful entities and relations could be parsed out, and thus impacting the final clustering result. Table VI shows such difference between Freebase and YAGO2 on the document clustering task, which marches the difference between the number of specified entities in Figure 10. Besides, it will be interesting to explore the clustering validation methods [Luo et al. 2009; Wu et al. 2009; Liu et al. 2013] to further demonstrate the reliability of the clustering results.

Fig. 10. Statistics of the number of entities in different document datasets with different world knowledge sources.

6.4.1. Analysis of Number of Entity Clusters. We also evaluate the effect of varying the number of entity clusters of each entity type in CHINC on the document clustering task. Figure 11(a) shows the results of clustering with different numbers of entity clusters of each entity type on “CHINC + Freebase” for the 20NG dataset. The number of entity clusters varies from 2 to 128. The default number of iterations is set as 20, which will be discussed in Section 6.4.2. When testing the effect of the number of entity clusters of one entity type, the numbers of entity clusters of the other two entity types are fixed as twice as the number of document clusters, which are 40 and 40 in 20NG, respectively. It is shown that for this dataset, more entity clusters may not result in improved document clustering results when a sufficient number of entity clusters are reached. For example, as shown in Figure 11(a), after reaching 32, the NMI scores of CHINC actually decrease when the numbers of entity clusters further increases. One can also find the effects of the numbers of entity clusters on the clustering performance with the other document dataset and knowledge base combinations in Figures 11(b)–(d) for Freebase and Figures 12(a)–(d) for YAGO2. From the results, we can conclude that there exist certain values of the number of entity clusters leading to the best clustering performance.

6.4.2. Analysis of Number of Iterations in Alternating Optimization. We evaluate the impact of the number of iterations of the alternating optimization (Algorithm 1) on CHINC in relation to the execution time of the optimization algorithm as well as the clustering performance. We increase the number of iterations from 1 to 80. For example, for each number of iterations, we run CHINC five trials, and the average execution time and NMI are summarized in Figures 13 and 14. From the result, one can conclude that the larger the number of iterations, the more significant the improvement on clustering performance. This improvement eventually drops, tapers out, and becomes stable. The reason is that with the increase of the number of iterations, the alternating optimization algorithm comes to convergence. However, the execution time still increases in a nearly linear manner. For example, as shown in Figure 13(a), after reaching 20, the performance stays stable. Thus, we set the number of iterations as 20 in the remaining experiments with the consideration of both performance and efficiency. As shown in Figures 13(b) and 14(d), we set the number of iterations to 20 when conducting

Fig. 11. Effect of number of entity clusters of each entity type on document clustering on different datasets with Freebase as world knowledge source.

experiments on the other combinations of document datasets and world knowledge bases.

6.4.3. Analysis of Specified World Knowledge-Based Constraints. Rather than using human knowledge as constraints, we use the specified world knowledge automatically generated by our approach as constraints in CHINC. Based on the specified world knowledge, it is straightforward to design constraints for entities.

Entity constraints. (1) *Must-links*. If two entities belong to the same entity sub-type, we add a must-link. For example, the entity sub-types of “Obama” and “Bush” are both *Politician*, we thus build a must-link between them. (2) *Cannot-links*. If two entities belong to different entity sub-types, we add a cannot-link. For example, the entity sub-types of “Obama” and “United States” are *Politician* and *Country*, respectively. In this case, we add a cannot-link to them.

We then test the performance of our proposed CHINC by using the specified world knowledge as constraints described above. We show the experiments on all of the different combinations of datasets and world knowledge sources in Figures 15 and 16. Each x -axis represents the number of entity type constraints used in each experiment,

Fig. 12. Effect of number of entity clusters of each entity type on document clustering on different datasets with YAGO2 as world knowledge source.

and y-axis is the average NMI of five random trials. For example, the constraints derived from entity types #1, #2, and #3 are eventually added to CHINC as shown in Figures 15(a), (b), and (c) respectively, when using Freebase as world knowledge and testing on 20NG dataset. We can see that CHINC outperforms the best clustering algorithm with the human knowledge as shown in Table VI (CITCC:0.569) with even no constraints (HINC:0.571). By adding more and more constraints, the clustering result of CHINC is significantly better. So CHINC is able to use information in world knowledge specified in the HIN, and the entity sub-type information can be transferred to the document side. The results show the power of modeling data as heterogeneous information networks, as well as the high quality of constraints derived from world knowledge.

From Figures 15 and 16, by increasing the number of constraints, we find that the average execution time of five trials increases linearly, and the clustering performance measured by NMI is increasing as mentioned before. For example, Figure 15(c) shows the effects of the constraints of all the three entity types on the clustering performance as well as the execution time, when Freebase is used as world knowledge and CHINC is tested on 20NG dataset. After the number of constraints reaches around 50M, the

Fig. 13. Analysis of number of iterations in alternating optimization algorithm on different datasets with Freebase as world knowledge source. Left y-axis: average NMI; Right y-axis: average execution time (ms).

Fig. 14. Analysis of number of iterations in alternating optimization algorithm on different datasets with YAGO2 as world knowledge source. Left y-axis: average NMI. Right y-axis: average execution time (ms).

Fig. 15. Effects of entity constraints and world knowledge source (Freebase). Left y-axis: average NMI. Right y-axis: average execution time (ms).

(a) Constraints of type #1 of "CHINC + YAGO2" for 20NG.

(b) Constraints of types #1+#2 of "CHINC + YAGO2" for 20NG.

(c) Constraints of types #1+#2+#3 of "CHINC + YAGO2" for 20NG.

(d) Constraints of type #1 of "CHINC + YAGO2" for MCAT.

(e) Constraints of types #1+#2 of "CHINC + YAGO2" for MCAT.

(f) Constraints of types #1+#2+#3 of "CHINC + YAGO2" for MCAT.

(g) Constraints of type #1 of "CHINC + YAGO2" for CCAT.

(h) Constraints of types #1+#2 of "CHINC + YAGO2" for CCAT.

(i) Constraints of types #1+#2+#3 of "CHINC + YAGO2" for CCAT.

(j) Constraints of type #1 of "CHINC + YAGO2" for ECAT.

(k) Constraints of types #1+#2 of "CHINC + YAGO2" for ECAT.

(l) Constraints of types #1+#2+#3 of "CHINC + YAGO2" for ECAT.

Fig. 16. Effects of entity constraints and world knowledge source (YAGO2). Left y-axis: average NMI. Right y-axis: average execution time (ms).

Fig. 17. Analysis of the efficiency of our algorithm on different document datasets with different world knowledge sources.

increase of performance drops and stays stable. At this point, the execution time is around 1.2M (ms). In Figure 17, we can see the similar results on the other combinations of document datasets and knowledge bases. We also find that the average execution time of our algorithm with Freebase as world knowledge source is greater than that with YAGO2. As shown in Figure 10, the reason is that each document datasets with Freebase could be specified much more entities than that with YAGO2. From the results, we can see that our algorithm is scalable to use the large scale specified world knowledge as constraints, and cluster large amounts of documents.

7. RELATED WORK

In this section, we review the related work on document clustering, machine learning with domain and world knowledge, as well as heterogeneous information networks.

7.1. Document Clustering

Document clustering has been studied for many years. We can use traditional one-dimensional clustering algorithms (e.g., k -means) to cluster the documents. If we treat the document and corresponding words as a bipartite graph, we can use co-clustering algorithms [Dhillon et al. 2003] to cluster the documents. Moreover, with the help of labeled seed documents, semi-supervised clustering can be used [Basu et al. 2002]. When the seeds are not available, we can use side information as constraints to guide clustering algorithms [Basu et al. 2004]. When the supervision from target domain is not available, we can also perform transfer learning to transfer the labeled information from other domains to the target domain [Dai et al. 2007a; Wang et al. 2009]. All the above clustering algorithms with supervision need domain or relevant domain knowledge. When there are diverse domains and the supervision is needed, they will still be very costly to ask a lot of different domain experts to label.

7.2. Machine Learning with Domain Knowledge

The general idea of incorporating domain knowledge into machine learning algorithms has already been studied extensively in natural language processing community. Chang et al. [2012] presented constrained conditional models (CCMs) that allow to inject high-level knowledge as soft constraints into linear models, such as Hidden Markov Models and Structured Perceptron, for lots of natural language processing tasks, including

information extraction [Roth and Yih 2004, 2007], semantic role labeling [Roth and Yih 2005; Punyakanok et al. 2008], summarization [Clarke and Lapata 2006], and co-reference resolution [Denis et al. 2007]. Posterior Regularization (PR) [Ganchev et al. 2010] works on incorporating indirect supervision via constraints on posterior distributions of probabilistic models with latent variables. The key difference between CCM-based learning and PR is that CCMs allow the use of hard constraints, whereas PR uses expectation constraints. Samdani et al. [2012] have proposed a unified Expectation-Maximization algorithm that can cover CCM-based learning and PR. Besides, lots of models have been studied to incorporate the domain knowledge for better performance. Markov Logic Network (MLN) [Richardson and Domingos 2006] is proposed to integrate first-order logic with Markov Random Field. A combination of the Bayesian network model with a collection of deterministic constraints is presented in Dechter and Mateescu [2004]. However, all of the mentioned works use domain knowledge to improve the performance of the corresponding machine learning algorithms. Different from their work, we explore a more general learning framework with world knowledge. Given domain-dependent data, we aim to automatically generate domain knowledge by specifying the world knowledge, and represent the domain knowledge in a unified format for more general machine learning. It will be interesting to use our learning framework in more learning models, to conduct empirical experiments to compare to CCM- and PR-based models in various domain-dependent tasks.

Besides, transfer learning [Pan and Yang 2010] is another direction on leveraging domain knowledge for better machine learning. The main idea of transfer learning is to leverage the domain knowledge in source domain to help the learning tasks in the target domain. The key intuition is that the labeled data is relatively easier to collect in the source domain than in the target domain. Unsupervised transfer learning [Dai et al. 2008] has been adapted for developing new clustering algorithms. Self-taught clustering [Dai et al. 2008] is an instance of unsupervised transfer learning, which aims at clustering a small collection of unlabeled data in the target domain with the help of a large amount of unlabeled data in the source domain. Besides unsupervised transfer learning, there are inductive transfer learning [Dai et al. 2007b; Mihalkova et al. 2007] and transductive transfer learning [Arnold et al. 2007; Ling et al. 2008] for the problem of regression and classification. Recently, source-free transfer learning [Lu et al. 2014] and selective transfer learning [Lu et al. 2013] have been proposed to address the text classification and cross domain recommendation problems, and shown improved results respectively. Lifelong learning [Eaton and Ruvolo 2013; Chen and Liu 2014] is also a framework of machine learning with domain knowledge. Lifelong learning test [Li and Yang 2015] has been proposed to take both the current performance and the performance growth rate into consideration. However, both transfer learning and lifelong learning are focusing on leveraging domain knowledge rather than world knowledge, whereas our framework is focusing on first how to specify the world knowledge to automatically generate domain knowledge and then modeling the learning task(s) with the help of the specified world knowledge. Again, it will be of great interests to build an end-to-end machine learning system by using the automatically generated domain knowledge while conducting transfer learning or lifelong learning.

7.3. Machine Learning with World Knowledge

Most of the existing usage of world knowledge is to enrich the features beyond bag-of-words representation of documents. For example, by using the linguistic knowledge base WordNet to resolve synonyms and introduce WordNet concepts, the quality of document clustering can be improved [Hotho et al. 2003]. The first paper using the term “world knowledge” [Gabrilovich and Markovitch 2005] extends the bag-of-words features with the categories in Open Directory Project (ODP), and shows that it can help improve text classification with additional knowledge. Following this, by mapping

the text to the semantic space provided by Wikipedia pages or other ontologies, it has been proven to be useful for short text classification [Gabrilovich and Markovitch 2006, 2007, 2009], clustering [Hu et al. 2008, 2009a, 2009b; Fodeh et al. 2011], and information retrieval [Egozi et al. 2011]. After that, a bunch of research also uses another knowledge base of taxonomy, Probase [Wu et al. 2012], to enrich the features of short text or keywords for classification [Wang et al. 2014a], clustering [Song et al. 2011, 2015], information retrieval [Hua et al. 2013; Song et al. 2014], or mines knowledge from text for information retrieval [Wang et al. 2013]. All of the above approaches just consider to use world knowledge as a source of features. However, the knowledge in the knowledge bases indeed has annotations of types, categories, and so on. Thus, it can be more effective to consider this information as “supervision” to supervise other machine learning algorithms and tasks. Along this research direction, recent work [Wang et al. 2015b] applies our world knowledge enabled machine learning framework for the document similarity computation task.

Distant supervision uses the knowledge of entities and their relationships from world knowledge bases, e.g., Freebase, as supervision for the task of entity and relation extraction [Mintz et al. 2009; Wang et al. 2014b; Xu et al. 2014]. It considers to use knowledge supervision to extract more entities and relations from new text or to generate a better embedding of entities and relations. Thus, the application of direct supervision is limited to entities and relations themselves.

Song et al. [2013] consider using fully unsupervised method to generate constraints of words using an external general-purpose knowledge base, WordNet. This can be regarded as an initial attempt to use general knowledge as indirect supervision to help clustering. However, the knowledge from WordNet is mostly linguistically related. It lacks of the information about named entities and their types. Moreover, their approach is still a simple application of constrained co-clustering, where it misses the rich structural information in the knowledge base.

7.4. Heterogeneous Information Network

An HIN is defined as a graph of multi-typed entities and relations [Han et al. 2010; Sun and Han 2012]. Different from traditional graphs, HIN incorporates the type information which can be useful to identify the semantic meaning of the paths in the graph [Sun et al. 2011b]. This is a good property to perform graph search and matching [He and Singh 2006; Yan et al. 2004; Zhang et al. 2007]. Original HINs are developed for the applications of scientific publication network analysis [Zhao et al. 2009; Sun et al. 2011a, 2011b, 2012]. Then, social network analysis also leverages this representation for user similarity and link prediction [Kong et al. 2013; Zhang et al. 2013, 2014]. Seamlessly, we can see that the knowledge in world knowledge bases, e.g., Freebase and YAGO2, can be naturally represented as an HIN since the entities and relations in the knowledge base are all typed. We introduce this representation to knowledge-based analysis, and show that it can be very useful for our document clustering task. Note that there is also a series of methods called multi-type relational data clustering [Long et al. 2006, 2007] and collective matrix factorization [Singh and Gordon 2008; Nickel et al. 2011; Bouchard et al. 2013; Klami et al. 2013]. Although they require the data to be structural beforehand (e.g., providing information of authors, and co-authors), our method only needs the input of raw documents. In addition to the multi-type relational information, we also incorporate the type information provided by the knowledge base as constraints to further improve the clustering results.

8. CONCLUSION

In this article, we study a novel problem of machine learning with world knowledge. Particularly, we take document clustering as an example and show how to use world knowledge as indirect supervision to improve the clustering results. To use the world

knowledge, we show how to adapt the world knowledge to the domain-dependent tasks by using semantic parsing and semantic filtering. Then, we represent the data as a heterogeneous information network, and use a constrained network clustering algorithm to obtain the document clusters. We demonstrate the effectiveness and efficiency of our approach on two real datasets along with two popular knowledge bases. In the future, we plan to use world knowledge to help more text mining and text analytic tasks, such as text classification and information retrieval.

REFERENCES

- Andrew Arnold, Ramesh Nallapati, and William W. Cohen. 2007. A comparative study of methods for transductive transfer learning. In *Proceedings of ICDM Workshop on Mining and Management of Biological Data (ICDM'07)*. 77–82.
- Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. *DBpedia: A Nucleus for a Web of Open Data*. Springer.
- Michele Banko, Michael J. Cafarella, Stephen Soderland, Matthew Broadhead, and Oren Etzioni. 2007. Open information extraction from the web. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'07)*. 2670–2676.
- Sugato Basu, Arindam Banerjee, and Raymond J. Mooney. 2002. Semi-supervised clustering by seeding. In *Proceedings of International Conference on Machine Learning (ICML'02)*. 27–34.
- Sugato Basu, Mikhail Bilenko, and Raymond J. Mooney. 2004. A probabilistic framework for semi-supervised clustering. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'04)*. 59–68.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP'13)*. 1533–1544.
- Jonathan Berant and Percy Liang. 2014. Semantic parsing via paraphrasing. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL'14)*. 1415–1425.
- Kurt D. Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: A collaboratively created graph database for structuring human knowledge. In *Proceedings of SIGMOD Conference*. 1247–1250.
- Guillaume Bouchard, Dawei Yin, and Shengbo Guo. 2013. Convex collective matrix factorization. In *Proceedings of International Conference on Artificial Intelligence and Statistics (AISTATS'13)*. 144–152.
- Qingqing Cai and Alexander Yates. 2013. Large-scale semantic parsing via schema matching and lexicon extension. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL'13)*. 423–433.
- Ming-Wei Chang, Lev Ratinov, and Dan Roth. 2012. Structured learning with constrained conditional models. *Machine Learning* 88, 3 (2012), 399–431.
- O. Chapelle, B. Schölkopf, and A. Zien (Eds.). 2006. *Semi-Supervised Learning*. MIT Press.
- Zhiyuan Chen and Bing Liu. 2014. Topic modeling using topics from many domains, lifelong learning and big data. In *Proceedings of International Conference on Machine Learning (ICML'14)*. 703–711.
- James Clarke and Mirella Lapata. 2006. Constraint-based sentence compression an integer programming approach. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL'06)*. 144–151.
- Michael Collins and Nigel Duffy. 2001. Convolution kernels for natural language. *NIPS*. 625–632.
- Wenyuan Dai, Gui-Rong Xue, Qiang Yang, and Yong Yu. 2007a. Co-clustering based classification for out-of-domain documents. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'07)*. 210–219.
- Wenyuan Dai, Qiang Yang, Gui-Rong Xue, and Yong Yu. 2007b. Boosting for transfer learning. In *Proceedings of the 24th International Conference on Machine Learning*. ACM, 193–200.
- Wenyuan Dai, Qiang Yang, Gui-Rong Xue, and Yong Yu. 2008. Self-taught clustering. In *Proceedings of International Conference on Machine Learning (ICML'08)*. 200–207.
- Rina Dechter and Robert Mateescu. 2004. Mixtures of deterministic-probabilistic networks and their AND/OR search space. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (AUAI'04)*. 120–129.
- Pascal Denis, Jason Baldridge, and others. 2007. Joint determination of anaphoricity and coreference resolution using integer programming. In *Proceedings of NAACL Conference*. 236–243.

- Inderjit S. Dhillon, Subramanyam Mallela, and Dharmendra S. Modha. 2003. Information-theoretic co-clustering. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'03)*. 89–98.
- Xin Dong, Evgeniy Gabrilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmman, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'14)*. 601–610.
- Eric Eaton and Paul L. Ruvolo. 2013. ELLA: An efficient lifelong learning algorithm. In *Proceedings of the 30th International Conference on Machine Learning (ICML'13)*. 507–515.
- Ofer Egozi, Shaul Markovitch, and Evgeniy Gabrilovich. 2011. Concept-based information retrieval using explicit semantic analysis. *ACM Transactions on Information Systems* 29, 2 (2011), 8.
- Oren Etzioni, Michael Cafarella, and Doug Downey. 2004. WebScale information extraction in KnowItAll (preliminary results). In *Proceedings of International Conference on World Wide Web (WWW'04)*. 100–110.
- Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. Identifying relations for open information extraction. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP'11)*. 1535–1545.
- Anthony Fader, Luke S. Zettlemoyer, and Oren Etzioni. 2013. Paraphrase-driven learning for open question answering. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL'13)*. 1608–1618.
- Christiane Fellbaum (Ed.). 1998. *WordNet: An Electronic Lexical Database*. MIT Press.
- Samah Fodeh, Bill Punch, and Pang-Ning Tan. 2011. On ontology-driven document clustering using core semantic features. *Knowledge and Information Systems* 28, 2 (2011), 395–421.
- Evgeniy Gabrilovich and Shaul Markovitch. 2005. Feature generation for text categorization using world knowledge. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'05)*. 1048–1053.
- Evgeniy Gabrilovich and Shaul Markovitch. 2006. Overcoming the brittleness bottleneck using Wikipedia: Enhancing text categorization with encyclopedic knowledge. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'06)*. 1301–1306.
- Evgeniy Gabrilovich and Shaul Markovitch. 2007. Computing semantic relatedness using Wikipedia-based explicit semantic analysis. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'07)*. 1606–1611.
- Evgeniy Gabrilovich and Shaul Markovitch. 2009. Wikipedia-based semantic interpretation for natural language processing. *Journal of Artificial Intelligence Research* 34 (2009), 443–498.
- Kuzman Ganchev, Joao Graça, Jennifer Gillenwater, and Ben Taskar. 2010. Posterior regularization for structured latent variable models. *Journal of Machine Learning Research* 11 (2010), 2001–2049.
- Jiawei Han, Yizhou Sun, Xifeng Yan, and Philip S. Yu. 2010. Mining knowledge from databases: An information network analysis approach. In *Proceedings of SIGMOD Conference*. 1251–1252.
- Huahai He and Ambuj K. Singh. 2006. Closure-tree: An index structure for graph queries. In *Proceedings of IEEE International Conference on Data Engineering (ICDE'06)*. IEEE, 38–38.
- Andreas Hotho, Steffen Staab, and Gerd Stumme. 2003. Ontologies improve text document clustering. In *Proceedings of IEEE International Conference on Data Mining (ICDM'03)*. 541–544.
- Jian Hu, Lujun Fang, Yang Cao, Hua-Jun Zeng, Hua Li, Qiang Yang, and Zheng Chen. 2008. Enhancing text clustering by leveraging Wikipedia semantics. In *Proceedings of ACM SIGIR International Conference on Research and Development in Information Retrieval (SIGIR'08)*. 179–186.
- Xia Hu, Nan Sun, Chao Zhang, and Tat-Seng Chua. 2009a. Exploiting internal and external semantics for the clustering of short texts using world knowledge. In *Proceedings of ACM International Conference on Information and Knowledge (CIKM'09)*. 919–928.
- Xiaohua Hu, Xiaodan Zhang, Caimei Lu, E. K. Park, and Xiaohua Zhou. 2009b. Exploiting Wikipedia as external knowledge for document clustering. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'09)*. 389–396.
- Wen Hua, Yangqiu Song, Haixun Wang, and Xiaofang Zhou. 2013. Identifying users' topical tasks in web search. In *Proceedings of ACM International Conference on Web Search and Data Mining (WSDM'13)*. 93–102.
- Arto Klami, Guillaume Bouchard, and Abhishek Tripathi. 2013. Group-sparse embeddings in collective matrix factorization. *Technical Report, ArXiv*.
- Xiangnan Kong, Jiawei Zhang, and Philip S. Yu. 2013. Inferring anchor links across multiple heterogeneous social networks. In *Proceedings of ACM International Conference on Information and Knowledge (CIKM'13)*. 179–188.

- Jayant Krishnamurthy and Tom M. Mitchell. 2012. Weakly supervised training of semantic parsers. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language (EMNLP-CoNLL'12)*. 754–765.
- Tom Kwiatkowski, Eunsol Choi, Yoav Artzi, and Luke S. Zettlemoyer. 2013. Scaling semantic parsers with on-the-fly ontology matching. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP'13)*. 1545–1556.
- Tom Kwiatkowski, Luke S. Zettlemoyer, Sharon Goldwater, and Mark Steedman. 2011. Lexical generalization in CCG grammar induction for semantic parsing. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP'11)*. 1512–1523.
- Ken Lang. 1995. NewsWeeder: Learning to filter netnews. In *Proceedings of International Conference on Machine Learning (ICML'95)*. 331–339.
- Douglas B. Lenat and R. V. Guha. 1989. *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Addison-Wesley.
- David D. Lewis, Yiming Yang, Tony G. Rose, and Fan Li. 2004. Rcv1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research* 5 (2004), 361–397.
- Lianghao Li and Qiang Yang. 2015. Lifelong machine learning test. In *Proceedings of the AAAI Workshop*.
- Yang Li, Chi Wang, Fangqiu Han, Jiawei Han, Dan Roth, and Xifeng Yan. 2013. Mining evidences for named entity disambiguation. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'13)*. 1070–1078.
- Percy Liang. 2013. Lambda dependency-based compositional semantics. *Technical Report, ArXiv*.
- Xiao Ling, Wenyuan Dai, Gui-Rong Xue, Qiang Yang, and Yong Yu. 2008. Spectral domain-transfer learning. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'08)*. 488–496.
- Yanchi Liu, Zhongmou Li, Hui Xiong, Xuedong Gao, Junjie Wu, and Sen Wu. 2013. Understanding and enhancement of internal clustering validation measures. *IEEE Transactions on Cybernetics* 43, 3 (2013), 982–994.
- Bo Long, Zhongfei (Mark) Zhang, Xiaoyun Wú, and Philip S. Yu. 2006. Spectral clustering for multi-type relational data. In *Proceedings of International Conference on Machine Learning (ICML'06)*. 585–592.
- Bo Long, Zhongfei Mark Zhang, and Philip S. Yu. 2007. A probabilistic framework for relational clustering. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'07)*. 470–479.
- Zhongqi Lu, Weike Pan, Evan Wei Xiang, Qiang Yang, Lili Zhao, and ErHeng Zhong. 2013. Selective transfer learning for cross domain recommendation. In *Proceedings of SIAM International Conference on Data Mining (SDM'13)*. 641–649.
- Zhongqi Lu, Yin Zhu, Sinno Jialin Pan, Evan Wei Xiang, Yujing Wang, and Qiang Yang. 2014. Source free transfer learning for text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'14)*. 122–128.
- Ping Luo, Hui Xiong, Guoxing Zhan, Junjie Wu, and Zhongzhi Shi. 2009. Information-theoretic distance measures for clustering validation: Generalization and normalization. *IEEE Transactions on Knowledge and Data Engineering* 21, 9 (2009), 1249–1262.
- Mausam, Michael Schmitz, Stephen Soderland, Robert Bart, and Oren Etzioni. 2012. Open language learning for information extraction. In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language (EMNLP-CoNLL'12)*. 523–534.
- Lilyana Mihalkova, Tuyen Huynh, and Raymond J. Mooney. 2007. Mapping and revising Markov logic networks for transfer learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'07)*, Vol. 7. 608–614.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the ACL/AFNLP Conference*. 1003–1011.
- Tom M. Mitchell, William W. Cohen, Estevam R. Hruschka Jr., Partha Pratim Talukdar, Justin Betteridge, Andrew Carlson, Bhavana Dalvi Mishra, Matthew Gardner, Bryan Kisiel, Jayant Krishnamurthy, Ni Lao, Kathryn Mazaitis, Thahir Mohamed, Ndapandula Nakashole, Emmanouil Antonios Platanios, Alan Ritter, Mehdi Samadi, Burr Settles, Richard C. Wang, Derry Tanti Wijaya, Abhinav Gupta, Xinlei Chen, Abulhair Saparov, Malcolm Greaves, and Joel Welling. 2015. Never-ending learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'15)*. 2302–2310.
- Raymond J. Mooney. 2007. Learning for semantic parsing. In *Proceedings of Conference on Intelligent Text Processing and Computational Linguistics (CICLing'07)*. 311–324.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A three-way model for collective learning on multi-relational data. In *Proceedings of International Conference on Machine Learning (ICML'11)*. 809–816.

- Sinno Jialin Pan and Qiang Yang. 2010. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* 22, 10 (2010), 1345–1359.
- Simone Paolo Ponzetto and Michael Strube. 2007. Deriving a large-scale taxonomy from Wikipedia. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'07)*. 1440–1445.
- Vasin Punyakanok, Dan Roth, and Wen-tau Yih. 2008. The importance of syntactic parsing and inference in semantic role labeling. *Computational Linguistics* 34, 2 (2008), 257–287.
- Lev Ratinov and Dan Roth. 2009. Design challenges and misconceptions in named entity recognition. In *Proceedings of Conference on Computational Natural Language Learning (CoNLL'08)*. 147–155.
- Siva Reddy, Mirella Lapata, and Mark Steedman. 2014. Large-scale semantic parsing without question-answer pairs. *Transactions of the Association for Computational Linguistics* 2 (2014), 377–392.
- Matthew Richardson and Pedro Domingos. 2006. Markov logic networks. *Machine Learning* 62, 1–2 (2006), 107–136.
- Dan Roth and Wen-tau Yih. 2004. A linear programming formulation for global inference in natural language tasks. In *Proceedings of Conference on Computational Natural Language Learning (CoNLL'08)*. 1–8.
- Dan Roth and Wen-tau Yih. 2005. Integer linear programming inference for conditional random fields. In *Proceedings of International Conference on Machine Learning (ICML'05)*. 736–743.
- Dan Roth and Wen-tau Yih. 2007. Global inference for entity and relation identification via a linear programming formulation. *Introduction to Statistical Relational Learning* (2007), MIT Press, 553–580.
- Rajhans Samdani, Ming-Wei Chang, and Dan Roth. 2012. Unified expectation maximization. In *Proceedings of the Conference of the North American North American Chapter of the Association for Computational Linguistics (NAACL'12)*. 688–698.
- Ajit P. Singh and Geoffrey J. Gordon. 2008. Relational learning via collective matrix factorization. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'08)*. 650–658.
- Yangqiu Song, Shixia Liu, Xueqing Liu, and Haixun Wang. 2015. Automatic taxonomy construction from keywords via scalable Bayesian rose trees. *IEEE Transactions on Knowledge and Data Engineering* 27, 7 (2015), 1861–1874.
- Yangqiu Song, Shimei Pan, Shixia Liu, Furu Wei, M. X. Zhou, and Weihong Qian. 2013. Constrained text coclustering with supervised and unsupervised constraints. *IEEE Transactions on Knowledge and Data Engineering* 25, 6 (2013), 1227–1239.
- Yangqiu Song, Haixun Wang, Weizhu Chen, and Shusen Wang. 2014. Transfer understanding from head queries to tail queries. In *Proceedings of ACM International Conference on Information and Knowledge (CIKM'14)*. 1299–1308.
- Yangqiu Song, Haixun Wang, Zhongyuan Wang, Hongsong Li, and Weizhu Chen. 2011. Short text conceptualization using a probabilistic knowledgebase. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'11)*. 2330–2336.
- Yangqiu Song, Shusen Wang, and Haixun Wang. 2015. Open domain short text conceptualization: A generative + descriptive modeling approach. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'15)*. 3820–3826.
- Alexander Strehl and Joydeep Ghosh. 2003. Cluster ensembles—A knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research* 3 (2003), 583–617.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: A core of semantic knowledge. In *Proceedings of International Conference on World Wide Web (WWW'07)*. 697–706.
- Yizhou Sun, Rick Barber, Manish Gupta, Charu C. Aggarwal, and Jiawei Han. 2011a. Co-author relationship prediction in heterogeneous bibliographic networks. In *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks (ASONAM'11)*. IEEE, 121–128.
- Yizhou Sun and Jiawei Han. 2012. Mining heterogeneous information networks: Principles and methodologies. *Synthesis Lectures on Data Mining and Knowledge Discovery* 3, 2 (2012), 1–159.
- Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S. Yu, and Tianyi Wu. 2011b. PathSim: Meta path-based top-k similarity search in heterogeneous information networks. In *Proceedings of the VLDB Endowment (PVLDB'11)*. 992–1003.
- Yizhou Sun, Brandon Norick, Jiawei Han, Xifeng Yan, Philip S. Yu, and Xiao Yu. 2012. Integrating meta-path selection with user-guided object clustering in heterogeneous information networks. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'12)*. 1348–1356.
- S. V. N. Vishwanathan, Nicol N. Schraudolph, Risi Kondor, and Karsten M. Borgwardt. 2010. Graph kernels. *Journal of Machine Learning Research* 11 (Aug. 2010), 1201–1242.
- Chenguang Wang, Nan Duan, Ming Zhou, and Ming Zhang. 2013. Paraphrasing adaptation for web search ranking. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL'13)*. 41–46.

- Chenguang Wang, Yangqiu Song, Ahmed El-Kishky, Dan Roth, Ming Zhang, and Jiawei Han. 2015a. Incorporating world knowledge to document clustering via heterogeneous information networks. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'15)*. 1215–1224.
- Chenguang Wang, Yangqiu Song, Haoran Li, Ming Zhang, and Jiawei Han. 2015b. KnowSim: A document similarity measure on structured heterogeneous information networks. In *Proceedings of IEEE International Conference on Data Mining (ICDM'15)*. 1015–1020.
- Chenguang Wang, Yangqiu Song, Dan Roth, Chi Wang, Jiawei Han, Heng Ji, and Ming Zhang. 2015c. Constrained information-theoretic tripartite graph clustering to identify semantically similar relations. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'15)*. 3882–3889.
- Zheng Wang, Yangqiu Song, and Changshui Zhang. 2009. Knowledge transfer on hybrid graph. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI'09)*. 1291–1296.
- Zhongyuan Wang, Fang Wang, Ji-Rong Wen, and Zhoujun Li. 2014a. Concept-based short text classification and ranking. In *Proceedings of ACM International Conference on Information and Knowledge (CIKM'14)*. 1069–1078.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014b. Knowledge graph and text jointly embedding. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP'14)*. 1591–1601.
- Junjie Wu, Hui Xiong, and Jian Chen. 2009. Adapting the right measures for k-means clustering. In *Proceedings of ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'09)*. 877–886.
- Wentao Wu, Hongsong Li, Haixun Wang, and Kenny Q. Zhu. 2012. Probase: A probabilistic taxonomy for text understanding. In *Proceedings of SIGMOD Conference*. 481–492.
- Chang Xu, Yalong Bai, Jiang Bian, Bin Gao, Gang Wang, Xiaoguang Liu, and Tie-Yan Liu. 2014. Rc-net: A general framework for incorporating knowledge into word representations. In *Proceedings of ACM International Conference on Information and Knowledge (CIKM'14)*. 1219–1228.
- Xifeng Yan, Philip S. Yu, and Jiawei Han. 2004. Graph indexing: A frequent structure-based approach. In *Proceedings of SIGMOD Conference*. ACM, 335–346.
- Xuchen Yao and Benjamin Van Durme. 2014. Information extraction over structured data: Question answering with freebase. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL'14)*. 956–966.
- Jiawei Zhang, Xiangnan Kong, and Philip S. Yu. 2013. Predicting social links for new users across aligned heterogeneous social networks. In *Proceedings of IEEE International Conference on Data Mining (ICDM'13)*. 1289–1294.
- Jiawei Zhang, Xiangnan Kong, and Philip S. Yu. 2014. Transferring heterogeneous links across location-based social networks. In *Proceedings of the ACM International Conference on Web Search and Data Mining (WSDM'14)*. 303–312.
- Shijie Zhang, Meng Hu, and Jiong Yang. 2007. Treepi: A novel graph indexing method. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE'07)*. 966–975.
- Peixiang Zhao, Jiawei Han, and Yizhou Sun. 2009. P-rank: A comprehensive structural similarity measure over information networks. In *Proceedings of ACM International Conference on Information and Knowledge (CIKM'09)*. 553–562.

Received October 2015; revised March 2016; accepted June 2016