

# cube2net: Efficient Query-Specific Network Construction with Data Cube Organization

Carl Yang, Mengxiong Liu, Frank He, Jian Peng, Jiawei Han

University of Illinois, Urbana Champaign, 201 N Goodwin Ave, Urbana, Illinois 61801, USA

{jiayang3, mliu60, shibihe, jianpeng, hanj}@illinois.edu

**Abstract**—Networks are widely used to model objects with interactions and have enabled various downstream applications. However, in the real world, network mining is often done on particular query sets of objects, which does not require the construction and computation of networks including all objects in the datasets. In this work, for the first time, we propose to address the problem of *query-specific network construction*, to break the efficiency bottlenecks of existing network mining algorithms and facilitate various downstream tasks. To deal with real-world massive networks with complex attributes, we propose to leverage the well-developed data cube technology to organize network objects *w.r.t.* their essential attributes. An efficient reinforcement learning algorithm is then developed to automatically explore the data cube structures and construct the optimal query-specific networks. With extensive experiments of two classic network mining tasks on different real-world large datasets, we show that our proposed *cube2net* pipeline is general, and much more effective and efficient in query-specific network construction, compared with other methods without the leverage of data cube or reinforcement learning.

## I. INTRODUCTION

Networks provide a natural and generic way for modeling the interactions of objects, upon which various tasks can be performed, such as node classification [1], community detection [2] and link prediction [3], [4]. However, as real-world networks are becoming larger and more complex every day, various network mining algorithms need to be frequently developed or improved to scale up, but such innovations are often non-trivial, if not impossible. Moreover, the quality of networks taken by these algorithms is often questionable: Do the networks include all necessary information, and is every piece of information in the networks useful?

While existing network mining algorithms mostly focus on more complex models for better capturing of the given network structures [5], [6], [7], in this work, for the first time, we draw attention to the fact that network mining tasks are often specified on *particular sets of objects of interest*, which we call *queries*, and advocate for *query-specific network construction*, where the goal is to construct networks that are most *relevant* to the queries.

Under the philosophy of the well-developed technology of *data cube* for large-scale data management, massive real-world networks can be partitioned into small subnetworks residing in fine-grained multi-dimensional *cube cells* *w.r.t.* their essential node properties [8], [9]. Assuming the proper data cubes can be efficiently constructed for particular networks automatically

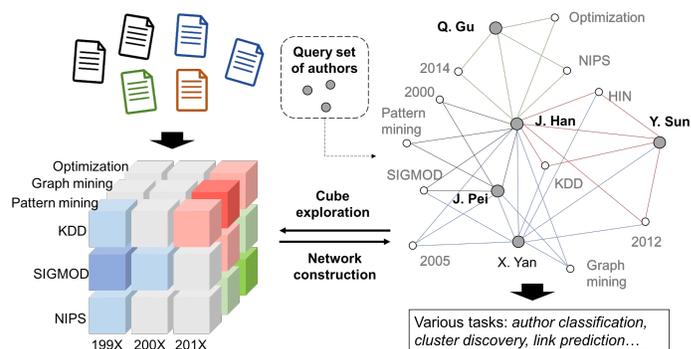


Fig. 1. *cube2net*: A running example on DBLP.

[10], [11], [12], [13], we can then clearly formulate the problem of this work as follows.

**Definition 1.1.** *Cube-Based Query-Specific Network Construction.* Given a massive network with objects organized in a data cube and a query set of objects, find a set of cells, so that objects in the cells are the most relevant to the query.

Figure 1 gives a toy example of query-specific network construction. Consider the massive author network of DBLP<sup>1</sup>. The task is to find pairs of close collaborators within a particular research group. Only retaining the co-author links within the group and ignoring all outside collaborations clearly leads to significant information loss, while incorporating all co-author links in the whole network is too costly and brings in lots of useless data. Based on the fact that the whole network can be partitioned into fine-grained multi-dimensional cube cells like  $\langle 200X, KDD, Graph\ mining \rangle$ ,  $\langle 201X, ACL, Text\ mining \rangle$ , we can look for a few subnetworks that are the most relevant to the considered group (*e.g.*, by looking at their overlap with the group), and leverage the union of them to serve as the query-specific network.

To identify the set of cells that contain subnetworks most relevant to the query, a straightforward method is exhaustive greedy search, by expanding the network with the most relevant cell incrementally. However, this method suffers from the following two drawbacks:

- 1) Since nodes in the real-networks can have various properties, there can be thousands or millions of cells, which makes exhaustive search very costly.

<sup>1</sup><http://dblp.uni-trier.de/>

- 2) The problem of looking for a set of most relevant cells is essentially a combinatorial problem over all cells, which can not be optimally solved by the greedy algorithm.

In this work, we propose and design *cube2net*, a simple and effective reinforcement learning algorithm over the data cube structures, to efficiently find a near optimal solution for the combinatorial problem of cube-based query-specific network construction. In our reinforcement learning framework, the state is represented by our novel designed continuous cell embedding vectors which capture the semantic proximities among cells in multiple dimensions, whereas the reward is designed to optimize the overall relevance between the set of selected subnetworks and the query set of objects. In this way, *cube2net* efficiently improves the utility estimation of various related cells regarding the relevance to the query by exploring each single cell, thus effectively approaching the optimal combination of relevant cells while avoiding the enumeration of all possible combinations.

The main contributions of this work are as follows:

- For real-world large-scale network data mining, we emphasize the lack of properly constructed networks and highlight the urge of query-specific network construction.
- We develop *cube2net*, based on the well-developed data cube technology and a novel simple reinforcement learning framework, to efficiently find the most relevant set of subnetworks from the cube structure for query-specific network construction.
- We conduct extensive experiments using two classic data mining tasks on different real-world massive networks to demonstrate the generality, effectiveness, and efficiency of *cube2net*.

## II. PROBLEM FORMULATION

### A. Preliminaries

1) *Data Cube Basics*: Data cube is widely used to organize multi-dimensional data, such as records in relational databases and documents in text collections [13], [9], [11], [12]. With well-designed cube structures, it can largely boost various downstream data analytics, mining and summarization tasks [8]. In the data cube, each object is assigned into a multi-dimensional *cell* which characterizes its properties from multiple aspects. We call each aspect as a *cube dimension* (to differentiate from vector dimension when necessary), which is formally defined as follows.

**Definition II.1.** *Cube Dimension.* A dimension in a data cube is defined as  $\mathcal{L}^p = \{l_1^p, l_2^p, \dots, l_{|\mathcal{L}^p|}^p\}$ , where  $l_i^p \in \mathcal{L}^p$  is a label in this dimension. For each particular dimension, labels are organized in either a flat or a hierarchical way.

For simplicity, we focus on flat labels in this work. Based on the notion of cube dimension, we formally define a data cube in the following.

**Definition II.2.** *Data cube.* A data cube is defined as  $\mathcal{C} = \{\mathcal{L}^1, \dots, \mathcal{L}^P, \mathcal{D}\}$ , where  $\mathcal{L}^p$  is the  $p$ -th dimension, and  $\mathcal{D}$  is a collection of objects assigned to the cube cells. Each object

$d \in \mathcal{D}$  can be represented as  $\{l_d^1, \dots, l_d^P\}$ , where  $l_d^p$  is the label of  $d$  in dimension  $\mathcal{L}^p$ .  $l_d^p$  can also be a set of labels, allowing objects to reside in multiple cells at the same time.

Without loss of generality, we give a simple example of cube construction with the DBLP data, aiming to demonstrate the power of data cube for efficient data organization and exploration, which further facilitates query-specific network construction. In the meantime, we are aware that by no means this is the unique or best way of constructing a DBLP data cube, while we leave the exploration of more cube design choices as future work. A reasonable belief is that better cube designs can lead to more efficient and effective data organization and exploration.

Following the design of [12], we can create three cube dimensions based on paper attributes in DBLP:  $\mathcal{L}^{decade}$  derived from the numerical attribute *publication year*,  $\mathcal{L}^{venue}$  derived from the categorical attribute *publication venue*, and  $\mathcal{L}^{topic}$  derived from textual attributes like *paper titles and abstracts*. We assign labels of the *venue* dimension by *publication years*; *venue* labels are assigned by conference or journal names like KDD, TKDE; *topic* labels are assigned according to a latent phrase-based topic model by labels like *Neural networks*, *Feature selection*. To compute a useful topic model, we first apply the popular AutoPhrase tool [14] to extract quality phrases from the whole corpus and represent each paper as a bag of those quality phrases. Then we apply standard LDA and assign each paper with the highest weighted topic label. As a consequence, each paper is arranged into multi-dimensional cells like  $\langle 201X, KDD, Graph\ mining \rangle$ .

Each object (*e.g.*, author) can belong to multiple cells at the same time. The co-author links may also cross different cells. To enable fast network construction, we store all links on both end objects, and add them to the network only when both end objects are selected. As a consequence, a data cube does not partition the network into isolated parts, but rather provides a fine-grained multi-dimensional index over particular subnetworks *w.r.t.* the essential properties of objects in the dataset for efficient exploration and selection.

### B. Problem Definitions

In this work, given a massive real-world network  $\mathcal{N} = \{\mathcal{V}, \mathcal{E}, \mathcal{A}\}$ , where  $\mathcal{V}$  is the set of objects (*e.g.*, millions of authors on DBLP),  $\mathcal{E}$  is the set of links (*e.g.*, co-author links on DBLP), and  $\mathcal{A}$  is the set of object properties (*e.g.*, favorite venues, active years, research topics for authors on DBLP), we assume a proper data cube  $\mathcal{C}$  can be well constructed either automatically or by domain knowledge. Then the input and output of query-specific network construction can be defined as follows.

**Input:** A data cube  $\mathcal{C}$  that organizes and stores a massive real-world network  $\mathcal{N}$  *w.r.t.* its essential node properties  $\mathcal{A}$ ; a query set of objects  $\mathcal{Q} \in \mathcal{V}$ .

**Output:** A set of cube cells  $\mathcal{S}$ , which includes the most relevant subnetwork  $\mathcal{M} \in \mathcal{N}$  regarding  $\mathcal{Q}$ .

There are many possible ways to mathematically define the relevance between  $\mathcal{M}$  and  $\mathcal{Q}$ . For simplicity, in this work, we

intuitively require  $\mathcal{M}$  to be (1) *small* so that the downstream tasks can be solved efficiently, and (2) *complete* so that knowledge mining over  $\mathcal{Q}$  can be accurate and effective. Based on such philosophies, we formulate the relevance function as

$$\text{rel}(\mathcal{M}, \mathcal{Q}) = \frac{|\mathcal{M} \cap \mathcal{Q}|}{|\mathcal{M} \cup \mathcal{Q}|}, \quad (1)$$

where  $|\mathcal{M} \cap \mathcal{Q}|$  is large when  $\mathcal{M}$  is *complete* by covering more objects in  $\mathcal{Q}$ , and  $|\mathcal{M} \cup \mathcal{Q}|$  is small when  $\mathcal{M}$  is *small* and also overlaps much with  $\mathcal{Q}$ .  $r(\mathcal{M}, \mathcal{Q})$  reaches the optimal value 1 when the  $\mathcal{M}$  exactly covers all objects in  $\mathcal{Q}$ , which only happens when the union of objects in a certain set of cells are exactly the same as  $\mathcal{Q}$ .

### III. CUBE2NET

#### A. Motivations

The combinatorial property of query-specific network construction lies in the process of selecting a particular set of objects from the very large set of all objects to construct the query-relevant subnetwork. Without the guidance of node properties, such a process can only be done by link-based heuristic search algorithms [15], [16].

In this work, however, by leveraging the powerful technique of data cube, we can organize the massive network *w.r.t.* its essential node properties and search over the cube cells instead of individual objects to construct the query-specific network. However, even with such efficient data organization, we are still facing the challenge of selecting over a large number of cube cells. For example, on DBLP, after removing the empty ones, we still have about 80K cells, which leads to a total number of  $2^{80K}$  possible combinations.

Reinforcement learning has been intensively studied for solving complex planning problems with consecutive decision makings, such as robot control and human-computer games [17], [18]. Recently, there are several approaches to tackle the combinatorial optimization problems over network data by reinforcement learning [19], [20], which are shown even more efficient than other advanced neural network models [21]. Motivated by their success, we aim to leverage reinforcement learning to efficiently approximate the optimal combinatorial solution for query-specific network construction.

#### B. Framework Formulation

In this subsection, we demonstrate our reinforcement learning framework. Given a data cube  $\mathcal{C}$  with  $n$  cells, a partial solution is represented as  $\mathcal{S} = (c_1, \dots, c_{|\mathcal{S}|})$ ,  $c_i \in \mathcal{C}$ ;  $\bar{\mathcal{S}} = \mathcal{C} \setminus \mathcal{S}$  is the set of candidates to be selected conditional on  $\mathcal{S}$ .

In [20], an approximate solution is proposed by using a popular heuristic, namely a greedy algorithm. Their algorithm constructs a solution by sequentially and greedily adding nodes that maximize the evaluation function of a partial solution. The evaluation function is trained on small instances (*e.g.*, 50 nodes) with deep Q-learning and the solution is tested on larger instances (*e.g.*, 1000 nodes) for the performance of generalization. Despite their success on several graph combinatorial optimization tasks, this heuristic does not scale to our problem. In our scenario, the action space is of the size  $|\mathcal{C}|$ , which means

in every step of the deep Q-learning, thousands of selections over the cells need to be explored and exploited, making the learning process computationally intractable.

Motivated by [19], our neural network architecture models a stochastic policy  $\pi(a | S, \mathcal{C})$ , where  $a$  is the action of selecting the next cell to be added from the candidate set  $\bar{\mathcal{S}} = \mathcal{C} \setminus \mathcal{S}$  and  $S$  is the state of the current partial solution  $\mathcal{S}$ . Our target is to learn the parameters of the policy network such that  $p(S | \mathcal{C})$  assigns a high probability to  $S$  that has high quality  $q(S)$  given all cells  $\mathcal{C}$ . We use the chain rule to factorize the probability of a solution as follows:

$$p(S | \mathcal{C}) = \prod_{i=1}^m \pi(S(i) | S(< i), \mathcal{C}), \quad (2)$$

where  $m$  is the size of the solution. Note that, we refer  $\pi(a | S)$  to  $\pi_{\Theta}(a | S, \mathcal{C})$ , where  $\Theta$  is the set of parameters inside the policy  $\pi$ .

1) *Representation*: As we just discussed, the number of possible states is exponential to the number of cells (*i.e.*,  $2^n$ ). It is impossible to model the states discretely with a look-up table. Therefore, we propose the novel technique of cell embedding, which captures the semantic proximities among cells. Particularly, for each cell  $c \in \mathcal{C}$ , a  $\kappa$ -dimensional embedding vector  $\mathbf{u}_c$  is computed to capture  $c$ 's multi-dimensional semantics regarding the essential properties of the objects it contains (more details are in *Appendix B*).

Now we study how to naturally leverage such cell embedding for the design of a reinforcement learning algorithm that efficiently explores the cube structures. Specifically, we need to design an appropriate reinforcement learning state representation to encode the currently selected cells in the subset  $\mathcal{S}$ . For model simplicity and easy optimization, we directly compute the state representation as a summation of the embeddings of selected cells in  $\mathcal{S}$ , *i.e.*,  $\mathbf{S} = \sum_{c \in \mathcal{S}} \mathbf{u}_c$ . With this simple form, we theoretically show that we allow the reinforcement learning agent to efficiently explore the cube structures, by simultaneously estimating the utility of semantically close cells at each state. Particularly, we prove the following theorem.

**Theorem III.1.** *Given our particular way of state representation, at each state  $S$ , the effect of the reinforcement learning agent in exploring a particular cell  $c \in \mathcal{C}$  is similar to that of exploring any cell  $c' \in \mathcal{N}_{\xi}(c)$ , when  $\xi$  is sufficiently small, in the sense that the output of the actor (critic) network is similar, *i.e.*,  $|\mathbf{f}(S, c) - \mathbf{f}(S, c')| \leq \eta\xi$ , and the gradients are similar too, *i.e.*,  $\|\nabla \mathbf{f}(S, c) - \nabla \mathbf{f}(S, c')\|_2^2 \leq \zeta\xi$ , where  $\mathbf{f} = \mu$  or  $\nu$ .*

Moreover, according to Theorem III.1, it is easy to arrive at the following Corollary.

**Corollary III.2.** *Given our particular way of state representation, the effect of the reinforcement agent of exploring a particular cell  $c \in \mathcal{C}$  in a particular state  $S$  is similar to that of exploring any cell  $c' \in \mathcal{N}_{\xi_1}(c)$  at any state  $S' \in \mathcal{N}_{\xi_2}(S)$ , when both  $\xi_1$  and  $\xi_2$  are sufficiently small.*

For detailed proofs, please refer to *Appendix A*.

Continue with our toy example on DBLP, where the deeply colored cells are directly sampled and estimated by the reinforcement learning agent during the exploration. According to Theorem III.1 and Corollary III.2, assuming the lightly colored cells are those semantically close to the deeply colored ones, their utilities regarding the quality function in similar states are also implicitly explored and evaluated. In this way, the reinforcement learning algorithm efficiently avoids the exhaustive search over all cube cells and their possible combinations.

In practice, we find that long trajectories leading to large constructed networks usually do not result in efficient model inference and satisfactory task performance. Therefore, we set  $\delta(S) = \{|\mathcal{S}| \leq m\}$ , which terminates the selection of more cells if the number of selected cells reaches  $m$ . In Section IV, we will study the impact of  $m$  on the performance of our algorithm.

2) *Reinforcement Learning Formulation*: The components in our reinforcement learning framework are defined as follows.

- 1) **State**: A state  $S$  corresponds to a set of cells  $\mathcal{S}$  we have selected. Based on previous discussion, a state is represented by a  $\kappa$ -dimensional vector  $\mathbf{S} = \sum_{c \in \mathcal{S}} \mathbf{u}_c$ .
- 2) **Action**: An action  $a$  is a cell  $c$  in  $\bar{\mathcal{S}} = \mathcal{C} \setminus \mathcal{S}$ . We will cast the details of the actions later.
- 3) **Reward**: The reward  $r$  of taking action  $a$  at state  $S$  is

$$r(S, a) = q(S') - q(S), \quad (3)$$

where  $S' := (S, c)$ . Given Eq. 3, maximizing the cumulative reward  $\sum_{i=1}^{n-1} r(S_i, c_i)$  is the same as maximizing the quality function  $q(S)$ .

- 4) **Transition**: The transition is deterministic by simply adding the cell  $c$  we have selected according to action  $a$  to the current state  $S$ . Thus, the next state  $S' := (S, c)$ .

According to our discussion in Section II.2, the quality function  $q$  is set to  $q(S) = \text{rel}(\mathcal{M}, \mathcal{Q})$ , where  $\mathcal{M}$  is the set of objects in the selected cells  $\mathcal{S}$  at state  $S$ .

Continue with our toy example in Figure 1, where  $\mathcal{Q}$  is the query set of authors including two data mining researchers, *e.g.*, Y. Sun and J. Pei. Assume the two authors are in the same cube cell <201X, KDD, Data mining>. Consider other authors in DBLP like X. Yan, who is also in the cell <201X, KDD, Data mining>, and Q. Gu, who is in another cell <201X, NIPS, Optimization>. Our quality function  $q(S)$  will prefer to include X. Yan into the constructed network, who is indeed more relevant to both Y. Sun and J. Pei, and such inclusion will likely benefit the proximity modeling among the particular authors in  $\mathcal{Q}$ .

Note that, although the toy example looks simple, when we consider  $\mathcal{Q}$  including thousands of objects distributed in multiple cells, as well as their relevant objects distributed in thousands of cells, the selection of the most relevant cells becomes a complex combinatorial problem. We stress that the main contribution of this work is to provide a general framework for efficient network construction, and it is trivial

to plug in different quality functions when certain properties of the constructed networks are required or preferred. In Section IV, we show the power of our framework for multiple network mining tasks with the simple quality function  $q$ , and we leave the exploration of more theoretically sound or task-specific ones to future work.

### C. Learning Algorithm

We illustrate how reinforcement learning is used to learn the policy parameters  $\Theta$ . It is impossible to directly apply the commonly used Q-learning to our task because finding the greedy policy requires optimization over  $a_t$  at every timestep  $t$ , which is too slow to be practical with large action spaces [22]. To this end, we adopt continuous policy gradient as our learning algorithm. We design our actions to be in the continuous cell embedding space. At each time, the policy outputs a continuous action vector and then we retrieve the closest cell by comparing the action vector with the cell embeddings.

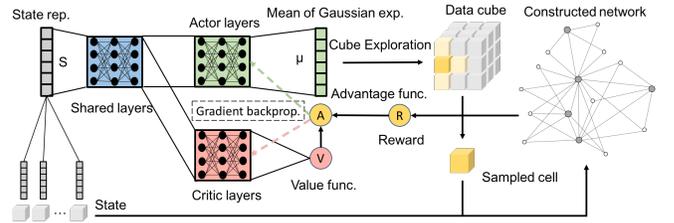


Fig. 2. The overall learning paradigm of cube2net.

Figure 2 illustrates our learning algorithm. More formally, our policy  $\pi(a | S)$  takes the state representation  $\mathbf{S}$  as the input and produces an action as the output. Instead of using Boltzmann exploration (*e.g.*, softmax) to choose a discrete action  $a = c$  from  $\bar{\mathcal{S}}$ , we project the actions to the  $\kappa$ -dimensional cell embedding space. Based on the property that semantically similar cells are close in the embedding space, a Gaussian exploration with mean  $\mu \in \mathbb{R}^\kappa$  and covariance matrix  $\Sigma \in \mathbb{R}^{\kappa \times \kappa}$  is used in our policy [22] as follows.

$$\pi(a | S, \{\mu, \Sigma\}) = \frac{1}{\sqrt{2\pi |\Sigma|}} \exp\left(-\frac{1}{2}(a - \mu)^T \Sigma^{-1} (a - \mu)\right), \quad (4)$$

$$\nabla_{\mu} \log \pi(a | S, \{\mu, \Sigma\}) = (a - \mu)^T \Sigma^{-1}, \quad (5)$$

$$\nabla_{\Sigma} \log \pi(a | S, \{\mu, \Sigma\}) = \frac{1}{2} (\Sigma^{-1} (a - \mu)(a - \mu)^T \Sigma^{-1} - \Sigma^{-1}). \quad (6)$$

A neural network  $\mu_{\Theta}(S)$  is used as our actor network to produce a Gaussian mean vector  $\mu$  as part of the action. A learnable parameter  $\sigma$  is used to model the Gaussian variance vector as  $\Sigma = \sigma^2 I$ .

1) *Policy Gradient*: Since our goal is to find a high-quality solution  $S$ , our training objective is

$$U(\Theta | C) = \mathbb{E}_{\tau \sim \pi_{\Theta}(S|C)} R(\tau), \quad (7)$$

where  $\tau$  denotes a state-action trajectory, *i.e.*,  $S_0, a_0, \dots, S_T, a_T$ . Each  $\tau$  corresponds to a selected set of cube cells.  $T(=m)$  is the length of the trajectory, and

$R(\tau) = \sum_{t=0}^T r(S_t, a_t)$  is the reward. Then we derive the gradient estimator based on policy gradient as

$$\nabla_{\Theta} U = \frac{1}{\alpha} \sum_{i=1}^{\alpha} \sum_{t=0}^{T-1} \nabla \log \pi_{\Theta}(a_t^{(i)} | S_t^{(i)}) \left( \sum_{k=t}^{T-1} r(S_k^{(i)}, a_k^{(i)}) - \nu_{\Theta}(S_k^{(i)}) \right), \quad (8)$$

where  $\alpha$  is the number of trajectories,  $\nu_{\Theta}(S)$  is the value function.

In a general form, we have the gradient estimator as

$$\hat{g} = \hat{\mathbb{E}}_t[\nabla_{\Theta} \log \pi_{\Theta}(a_t | S_t) \hat{A}_t], \quad (9)$$

where  $\hat{A}_t$  is the advantage function at timestep  $t$ .  $\hat{\mathbb{E}}_t$  denotes the empirical average over a mini-batch of samples in the algorithm that alternates between sampling and optimization using the proximal policy gradient algorithm [23].

To compute the advantage function  $\hat{A}_t$ , we use a T-step advantage function estimator as

$$\hat{A}_t = -\nu_{\Theta}(S_t) + r(S_t, a_t) + \gamma r(S_{t+1}, a_{t+1}) + \dots + \gamma^{T-t+1} r(S_{T-1}, a_{T-1}) + \gamma^{T-t} \nu_{\Theta}(S_T),$$

where  $\gamma$  is a clipping parameter.  $\nu_{\Theta}(S)$  is learned by minimizing the loss  $L^{\text{VF}}(\Theta) = \|\nu_{\Theta}(S_t) - \nu^{\text{target}}(S_t)\|_2^2$ , where  $V^{\text{target}}$  is the advantage function minus the value function.

2) *Neural Architecture*: Our actor network (policy network)  $\mu_{\Theta}(S)$  and critic network (value function)  $\nu_{\Theta}(S)$  shown in Figure 2 are both fully connected feedforward neural networks, each containing four layers including two hidden layers of size  $H$ , as well as the input and output layers. Sigmoid is used as the activation function for each layer, and the first hidden layer is shared between two networks. Both networks' inputs are  $\kappa$ -dimensional cell embeddings. The output of the actor network  $\mu_{\Theta}(S)$  is a  $\kappa$ -dimensional vector  $\mathbb{R}^{\kappa}$  and the output of critic network  $\nu_{\Theta}(S)$  is a real number.

As for model learning, in each iteration, our algorithm samples a set  $\Phi$  of  $\alpha$  trajectories (series of cells) of length  $T$  using the current policy  $\pi_{\Theta}(a | S)$ , and constructs the surrogate loss and value function loss based on  $\Phi$ . Mini-batch SGD is then used to optimize the objectives for  $\beta$  epochs, where the model parameters in  $\Theta$  are updated by Adam [24].

#### D. Computational Complexity

We theoretically analyze the complexity of our algorithm. Consider the problem of selecting  $m$  cells from a data cube  $\mathcal{C}$  with  $n$  cells ( $m \ll n$ ). During each step of model training, *cube2net* generates a target mean  $\mu_{\Theta}$  in constant time and then selects a cell from  $\mathcal{C}$  that is the closest to  $\mu_{\Theta}$ . Since computing the quality function and updating the neural network model based on particular trajectories take constant time and  $\text{argmax}$  takes logarithm time when parallelized on GPU, the overall complexity of training and planning with *cube2net* is  $O(\alpha\beta cm \log n)$ .  $\alpha, \beta$  are number of trajectories and epochs, which are both set to 40 in our experiments;  $c$  is the constant time in computing the quality function. The time for model inference is ignorable compared with model training.

As a comparison, a random algorithm takes  $O(1)$  time to select each cell and has an overall complexity of  $O(m)$ . A greedy algorithm at each of the  $m$  steps needs to consider all  $n$  actions and run the quality function on each of them to get the maximum, which has an overall complexity of  $O(cmn)$ . In practice, computing the quality function cannot be parallelized on GPU and always takes the most significant time. Also, the time is not truly constant—it largely favors *cube2net* as it constructs small networks by looking for the globally optimal combinations of relevant cells.

Note that, the novelty of our reinforcement learning algorithm lies in its unique leverage of cube structures and cell embedding, as well as the efficient Gaussian exploration policy. However, the training of it is rather routine, which as we will also show in Section IV, is very efficient and does not require heavy parameter tuning.

## IV. EXPERIMENTAL EVALUATION

In this section, we evaluate the effectiveness and efficiency of our proposed *cube2net* pipeline on two large real-world datasets from very different domains, *i.e.*, DBLP, an academic publication collection and Yelp<sup>2</sup>, a business review platform. We compare different network construction methods for two well-studied tasks on each of them—author clustering on DBLP and business-user link prediction on Yelp, to show the power and generality of *cube2net*.

### A. Author Clustering on DBLP

We first study the author networks of DBLP, which we have been using as the running example throughout this paper. Particularly, we focus on the problem of author clustering, since it is a well-studied problem with two sets of public standard evaluation labels [25], which are not directly captured by the network structures or attributes. Since the essential challenge of network mining is to capture object proximities (*e.g.*, using embeddings), we take an *embedding-clustering* approach and provide a comprehensive evaluation of the embedding quality regarding clustering. As shown in the experimental results of various embedding techniques [26], [27], [28], the quality of network embeddings are often consistent across different network mining tasks. Therefore, superior performances on author clustering with embedding can be a good indicator to the general high quality of constructed networks towards various network mining tasks.

#### 1) Experimental Settings:

**Dataset.** We use the public DBLP dataset V10<sup>3</sup> collected by [29]. The basic statistics of the dataset are shown in Table I.

Dataset	#papers	#authors	#links
DBLP	3,079,007	534,407	28,347,138
Dataset	#businesses	#customers	#links
Yelp	174,567	1,326,101	5,261,669

TABLE I

STATISTICS OF THE TWO PUBLIC DATASETS WE USE.

<sup>2</sup><https://www.yelp.com/>

<sup>3</sup><https://aminer.org/citation>

The DBLP dataset contains semi-structured scientific publications, with their corresponding authors, years, venues and textual contents. As discussed in Section II, a simple data cube with dimensions *year*, *venue* and *topic* is constructed, with each cell holding the corresponding *authors*.

For evaluating the task of author clustering, we use two sets of authors with ground-truth labels regarding research groups and research areas published by [25]. They are commonly used as ground-truth author labels for evaluating network classification and clustering algorithms on DBLP. The smaller set includes 4 labels of research groups on 116 authors, and the larger set includes 4 labels of research areas on 4,236 authors. We separately use them as the query sets of objects  $Q$  in two sets of experiments.

**Baselines.** We aim to improve network mining by constructing the proper network that contains only the query set of objects  $Q$  and the most relevant additional objects  $Q^+$ . Since we are the first to study the problem of query-specific network construction, we design a comprehensive group of baselines that can potentially choose a relevant set of objects  $Q^+$  to add to  $Q$  in the following.

- **NoCube:** Without a data cube, this method adds no additional object to the smallest network with only  $Q$  as nodes.
- **NoCube+:** Without a data cube, this method adds all objects that are directly linked with  $Q$  to the network.
- **NoCube++:** Without a data cube, this method adds all objects that are within two steps away from  $Q$  to the network.
- **MaxDisc** [15]: Without a data cube, this method leverages a heuristic search algorithm to find the largest network connected component that contains  $Q$ .
- **CubeRandom:** With a proper data cube, this method adds objects in  $m$  random cells to the network.
- **CubeGreedy:** With the same data cube, this method searches through all cells for  $m$  times, and greedily add the objects in one cell at each time to optimize the quality function of Eq. 1.

**Parameter settings.** When comparing *cube2net* with the baseline methods, we set its parameters to the following default values: For cube construction, we set the number of topics in LDA to 100 and we filter out cells with less than 10 objects; for reinforcement learning, we empirically set the size of hidden layers  $H$  to 256, the length of trajectories  $T$  (*i.e.*, the number of added cells  $m$ ) to 20, the sample size  $\alpha$  and the number of epochs  $\beta$  both to 40, and the clipping parameter of value function  $\gamma$  to 0.99. In addition to these default values, we also evaluate the effects of different parameters such as the length of trajectories  $T$  to study their impact on the performance of the three cube-based algorithms.

**Evaluation protocols.** We aim to show that *cube2net* is a general pipeline that breaks the bottleneck of network mining by constructing query-specific networks. Therefore, we evaluate the compared algorithms for network construction through two perspectives: effectiveness and efficiency.

We study the effectiveness by measuring the performance of downstream applications on the constructed network  $G$ . For author clustering, we firstly compute a network embedding  $E$  on  $G$ , which converts the network structure into distributed node representations. To provide a comprehensive evaluation, we use three popular algorithms, *i.e.*, *DeepWalk* [26], *LINE* [27] and *node2vec* [28] for this step. We then feed  $E$  into the standard  $K$ -means.

To evaluate the clustering results, we compute the *F1 similarity* (F1), *Jaccard similarity* (JC) and *Normalized Mutual Information* (NMI) *w.r.t.* ground-truth labels. Detailed definitions of the three metrics can be found in [30].

The efficiency of network construction is reflected by the *time* spent on constructing the network, as well as the *size* of the constructed network (in terms of #nodes and #edges), which can further influence the runtime of the network mining algorithms. All runtimes are measured on a server with one GeForce GTX TITAN X GPU and two Intel Xeon E5-2650V3 10-core 2.3GHz CPUs.

2) *Performance Comparison with Baselines:* We quantitatively evaluate *cube2net* against all baselines on clustering the query set of authors. Tables II and III comprehensively show the effectiveness and efficiency of compared algorithms.

The scores are deterministic for most algorithms and have very small standard deviations on the others except for *CubeRandom*. We run *CubeRandom* and *cube2net* for 10 times to take the average. The improvements of *cube2net* over other algorithms all passed the paired t-tests with significance value  $p < 0.01$ .

As for effectiveness, the network constructed by *cube2net* is able to best facilitate network mining around the query, particularly, author network clustering in this experiment, under the three metrics computed on standard  $K$ -means clustering results after network embedding. As can be observed, (1) blindly adding neighbors into the network without a cube organization or randomly adding cells can hurt the task performance; (2) with a proper cube structure, greedily adding cells *w.r.t.* our quality function can significantly boost the task performance; however, (3) the performance of *CubeGreedy* is still inferior to *cube2net*, which confirms our arguments that the task of network construction is essentially a combinatorial problem, which requires a globally optimal solution that can be efficiently achieved only by reinforcement learning.

Note that, we observe that blindly enriching the network with neighbor nodes does not always hurt the performance. To better understand it, we look into the data and find that links among the 116 labeled authors are much denser than the 4,236 authors (*i.e.*, 4.77% *v.s.* 0.05%). It indicates adding neighbors are more helpful when the existing interactions among the query set itself are more sparse. Particularly, we find that among the 4,236 authors, 930 were dangling in the original network, but can be bridged into the constructed network with other queried authors by adding their direct neighbors.

As for efficiency, (1) without cube organization, the network can easily get too large, which requires significant network construction time, and leads to long runtimes of network

Net. Con. Algorithm	Effectiveness on the smaller set of authors								
	F1			JC			NMI		
	DeepWalk	LINE	node2vec	DeepWalk	LINE	node2vec	DeepWalk	LINE	node2vec
NoCube	0.7034	0.5620	0.6195	0.4828	0.4621	0.5179	0.4013	0.3642	0.3976
NoCube+	0.6559	0.5263	0.5812	0.4359	0.4178	0.4763	0.2928	0.2814	0.3058
NoCube++	0.6794	0.5247	0.5874	0.4678	0.4155	0.4616	0.3799	0.3152	0.3654
MaxDisc	0.7162	0.5423	0.6299	0.4983	0.4437	0.5305	0.4052	0.3246	0.4122
CubeRandom	0.5839	0.5087	0.5748	0.4681	0.4194	0.5069	0.3693	0.3105	0.3930
CubeGreedy	0.7445	0.5988	0.6432	0.5492	0.4850	0.5343	0.3981	0.3369	0.4086
cube2net	<b>0.7628</b>	<b>0.6295</b>	<b>0.6913</b>	<b>0.5720</b>	<b>0.5312</b>	<b>0.5834</b>	<b>0.4196</b>	<b>0.3784</b>	<b>0.4517</b>
Net. Con. Algorithm	Effectiveness on the larger set of authors								
	F1			JC			NMI		
	DeepWalk	LINE	node2vec	DeepWalk	LINE	node2vec	DeepWalk	LINE	node2vec
NoCube	0.4336	0.3044	0.3708	0.2541	0.1759	0.2195	0.0809	0.0426	0.0439
NoCube+	0.5207	0.3212	0.5113	0.3022	0.1773	0.3362	0.1105	0.0454	0.0996
NoCube++	0.5515	0.3261	0.4984	0.3989	0.1725	0.3465	0.2174	0.0436	0.1086
MaxDisc	0.5859	0.3282	0.5619	0.4249	0.1950	0.4071	0.2120	0.0484	0.2278
CubeRandom	0.4018	0.2972	0.3416	0.2158	0.1543	0.1766	0.0620	0.4021	0.0376
CubeGreedy	0.6125	0.3509	0.5768	0.4526	0.1954	0.4186	0.2513	0.0595	0.2224
cube2net	<b>0.6447</b>	<b>0.3718</b>	<b>0.6214</b>	<b>0.4865</b>	<b>0.2288</b>	<b>0.4623</b>	<b>0.2597</b>	<b>0.0646</b>	<b>0.2418</b>

TABLE II

EFFECTIVENESS OF QUERY-SPECIFIC NETWORK CONSTRUCTION FOR CLUSTERING QUERY SET OF AUTHORS.

Net. Con. Algorithm	Efficiency on the smaller set of authors					
	Computation Time			Network Size		
	DeepWalk	LINE	node2vec	Net. Con.	#nodes	#edges
NoCube	1s	1s	1s	2s	116	318
NoCube+	25s	58s	65s	5s	3,853	32,882
NoCube++	732s	652s	904s	62s	55,724	540,780
MaxDisc	1s	2s	3s	794s	140	440
CubeRandom	56s	59s	62s	4s	2,512	24,578
CubeGreedy	28s	43s	32s	3,082s	1,464	14,892
cube2net	7s	11s	10s	296s	526	2,953
Net. Con. Algorithm	Efficiency on the larger set of authors					
	Computation Time			Network Size		
	DeepWalk	LINE	node2vec	Net. Con.	#nodes	#edges
NoCube	11s	74s	23s	3s	4,236	4,678
NoCube+	270s	127s	2,147s	84s	74,459	120,803
NoCube++	2,450s	1,514s	8,485s	1,128s	434,941	1,372,892
MaxDisc	18s	78s	113s	3,390s	4,718	7,004
CubeRandom	104s	108s	430s	4s	21,126	31,481
CubeGreedy	62s	96s	208s	4,194s	10,046	16,259
cube2net	23s	83s	92s	314s	6,842	12,497

TABLE III

EFFICIENCY OF QUERY-SPECIFIC NETWORK CONSTRUCTION FOR CLUSTERING QUERY SET OF AUTHORS.

mining algorithms; (2) the sizes of constructed networks are much more controllable with a proper cube organization, because we can easily set the number of cells to add; (3) even with a proper cube, greedily searching the cube at each step to select the proper cells is extremely time-consuming—on the contrary, *cube2net* efficiently explores the cube structures with reinforcement learning and is able to find the particular set of cells to construct the most relevant subnetwork, which also makes the downstream network mining more efficient.

Comparing the results on the two sets of query objects of different sizes, we further find that, (1) as the query set of authors becomes larger, blindly bringing in neighbors leads to much larger networks, and subsequently much longer runtimes of network mining algorithms. Such low efficiency is exactly what we want to avoid by aiming at query-specific network construction in this work; (2) when the query set becomes much larger, the runtimes of the cube-based algorithms only increase a little, since they still work on the same well organized cube structure by evaluating the utility of cells rather than individual nodes, indicating the power of the data cube organization.

Note that, since *cube2net* needs to learn the policy every time given a new set of objects, the network construction time we report for *cube2net* is a sum of three runtimes: model

training, model inference and network link retrieval. In this case, the comparison to other algorithms is fair, whose network construction time is a sum of node selection and link retrieval.

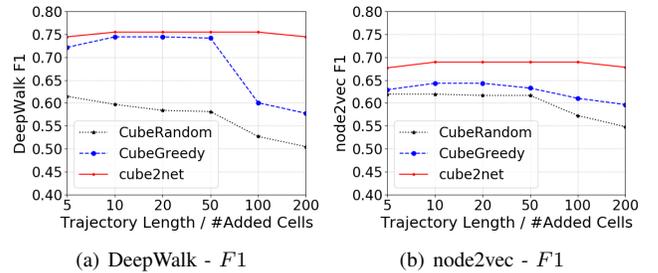


Fig. 3. Comparing different trajectory lengths.

In Figure 3, we also show the performances of the three cube-based algorithms as we consider different numbers of cells to add into the network with the small set of 116 labeled authors. Although particular network mining algorithms might prefer different network structures, the performances of three compared network construction algorithms vary in very similar trends: (1) *cube2net* maintains the best and most robust performance in all cases; (2) *CubeGreedy* performs closely with *cube2net* when the number of added cells is small, but gets worse as more cells are considered and the combinatorial property of the problem emerges; (3) *CubeRandom* keeps

getting worse as more random cells are added into the network.

Authors	Cells selected for the given set of authors			
	Rank	Decade	Venue	Topic
J. Han	1	200X	KDD	Recommender systems
D. Blei	2	200X	AAAI	Knowledge discovery
D. Roth	3	200X	ICIC	Multi-task learning
J. Leskovec	4	200X	ICML	Statistic models
...	5	200X	EMNLP	Language models

TABLE IV  
EXAMPLE CELLS SELECTED BY *cube2net* ON DBLP.

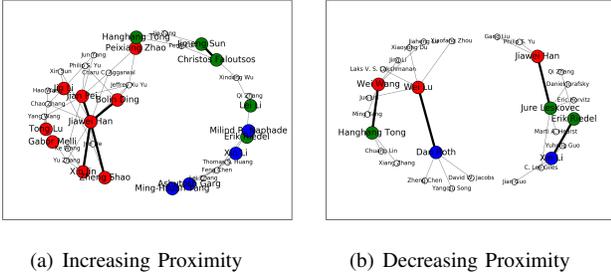


Fig. 4. Close-ups of the networks constructed by *cube2net* (better viewed with zoom-in).

3) *Case Study*: As a side product of *cube2net*, we are able to gain valuable insight into the query set of objects by looking at the selected cells. Table IV shows the first few cells selected by *cube2net* on DBLP for the small set of 116 authors from four well-known research groups led by authors in the left column. In the table, the rows are intentionally misaligned because there exists no one-to-one correspondence between individual authors and cells. However, all cells selected by *cube2net* are indeed highly relevant to the given set of authors, and they together provide a multi-dimensional view of the whole query set. By looking at the *decade* dimension, we can see that the majority authors in the set are most active in the 2000s, and *cube2net* does select the most relevant cells in this dimension. However, since the authors have quite different research focuses, the selected labels in the *venue* dimension are quite heterogeneous, which helps drag away authors from different research groups, such as those active in data mining v.s. machine learning. The *topic* dimension further describes the exact research problems that mostly appeal as well as differentiate authors in  $\mathcal{Q}$ .

In Figure 4, we focus on particular parts of the network constructed around the small set of 116 labeled authors to gain insight into how *cube2net* can construct high-quality networks that facilitate downstream tasks. Colored nodes are those in  $\mathcal{Q}$ , where different colors denote the ground-truth research group labels, and white nodes are those within the cells selected by *cube2net*. The links among nodes in the given set are thicker than those between them and the added nodes. As we can see, in Figure 4 (a), some authors in the same research groups do not have direct co-authorships, but they are drawn closer by the additional authors who co-author with both of them. Moreover, in Figure 4 (b), some authors in different research groups have co-authorships, but they are drawn further away by their different other co-authors. *cube2net* can efficiently find those helpful additional authors, and adding them into the

network can effectively improve modeling of the proximities among the query set of authors.

### B. Link Prediction On Yelp

Although we focus on the author networks of DBLP when developing *cube2net*, the concepts and techniques are general and applicable to various real-world networks. To demonstrate this, we conduct experiments on the widely used Yelp dataset. Instead of clustering, we focus on the more popular and well-studied problem of place recommendation. Moreover, since the constructed networks are business-user bipartite networks, instead of the homogeneous network embedding algorithms, we apply standard matrix completion algorithms to evaluate their quality towards link prediction. Superior performance in these experiments can further indicate the general and robust advantages of *cube2net* across different datasets and tasks.

#### 1) Experimental Settings:

**Dataset.** We use the public Yelp dataset from the Yelp Challenge Round 11<sup>4</sup>. Its basic statistics are also in Table I.

The dataset contains semi-structured data around businesses and users, such as locations, reviews, check-ins, etc. To organize the complex multi-dimensional data, we leverage quality attributes of businesses like *city* and *category*, and aggregate the *review texts* of each business to model them similarly as for paper contents in DBLP. Thus, we are able to build a simple data cube with dimensions *city*, *category* and *topic*, with each cell holding the corresponding businesses and users. For cell embedding, since all labels are common phrases, we leverage the same pre-trained word embedding table and general cell embedding approach as used for DBLP.

To further demonstrate the power of *cube2net*, we focus on the well studied task of place recommendation on Yelp, which is essentially a link prediction problem on a bipartite network of businesses and users as objects, and reviews as links between them. Based on additional attributes of the businesses, we generate two different query sets of objects and hide some of the links among them for performance evaluation. Specifically, the first set (ILKID) includes 198 businesses randomly selected from the total 414 that are within the state of Illinois and are Good For Kids, and the second set (NVOUT) includes 2,982 businesses randomly selected from the total 6,020 that are within the state of Nevada and offer Takeouts. Both sets also include 80% of all users that have left reviews for the corresponding businesses (i.e., 528 users for ILKID and 20,968 users for NVOUT). 20% of the links generated by the latest reviews are hidden for evaluation (i.e., 161 for ILKID and 6,269 for NVOUT). For each  $\mathcal{Q}$  of the two query sets, our task is thus to efficiently bring in relevant businesses and users from the large remaining part of the data, to construct a bipartite network around  $\mathcal{Q}$ , and improve the performance of link prediction in  $\mathcal{Q}$ .

**Evaluation protocols.** To evaluate the quality of constructed networks on the particular task of business-user link prediction

<sup>4</sup><https://www.yelp.com/dataset>

on Yelp, we chose a popular low-rank matrix completion algorithm called *Soft-Impute* [31] with open source implementations<sup>5</sup>, due to its relatively high and stable performance and efficiency. Upon the bipartite networks constructed by different algorithms in comparison, *Soft-Impute* is applied to recover the missing links, which are then evaluated on the held out links in the query sets.

To evaluate the link prediction performance, we compute precision and recall at  $K$ , as commonly done in related works [32]. We also record the runtime of network construction and *Soft-Impute* on the same server, as well as the size of the constructed networks.

**Parameter settings.** Since the textual contents in reviews are simpler than in papers, we set the number of *topics* to 50, and we also filter out cells with less than 10 objects. For the reinforcement learning model, we use all the same default parameter settings as in Section IV.1, except for the length of trajectories  $m$ : We use 50 for ILKID and 100 for NVOU. This also confirms that *cube2net* is not sensitive to parameter settings and can easily maintain robust performance across domains, datasets, and tasks.

2) *Performance Comparison with Baselines:* Figure 5 demonstrates the quality of constructed query-specific networks on Yelp.

In the figures, similar merits of *cube2net* as discussed on DBLP can be observed, *i.e.*, it is able to construct networks that facilitate downstream tasks with both effectiveness and efficiency. Particularly, the networks constructed by *cube2net* lead to the best precision and recall for business-user link prediction among all compared algorithms. Moreover, the networks constructed by *cube2net* have the smallest size, except for *MaxDisc*, which only considers the connected network components. To compute such a dense and connected network component, *MaxDisc* spends a lot of time searching over the node neighborhoods, but still leaves a lot of nodes dangling outside of the constructed network, resulting in quite poor performance. While NVOU (2,982 businesses + 20,968 users) is much larger than ILKID (198 businesses + 161 users), the results follow quite similar trends. Such results strongly indicate the general and robust advantages of *cube2net*. Finally, for different queries on the corresponding datasets, we use almost the same parameters (except for the trajectory lengths), which confirms the generalization and easy training of *cube2net*.

Note that, compared with DBLP, on Yelp, the node attributes, network properties, downstream tasks and evaluation protocols are all quite different. Therefore, the consistent and significant gains of *cube2net* over all compared algorithms clearly indicate its general power in query-specific network construction to facilitate various network mining algorithms towards different downstream tasks.

## V. RELATED WORK

### A. Network Mining

Various network mining algorithms have been developed in the past decade, such as information propagation models [33], stochastic block models [34] and generative probabilistic models [35]. Recently, unsupervised network embedding based on the advances in neural language models like [36] has been extremely popular [26], [27], [28]. The idea is to compute a distributed representation of nodes, which captures their network proximities, regarding both neighborhood similarity and structural similarity [5], [7], [6], [37], [38].

However, almost all existing network mining methods are limited to a fixed set of objects. Among them, most methods assume the network structures to be given [39], [40], [41], [42], [43], [44], while some others try to infer links among fixed sets of objects [45], [46], [47]. Some recent methods consider dynamic networks with changing sets of objects [48], [49], [50], but they do not actively choose which objects to include into the network. For the particular problem of subnetwork construction, some heuristic search algorithms have been developed to find interesting subnetworks from seed nodes, but they only work for particular network properties and do not consider node attributes [15], [16]. Regarding scalability, commercial graph databases like neo4j [51] can handle massive networks efficiently, but they only support fixed sets of operations and do not provide seamless combination to arbitrary graph mining algorithms.

In this work, we consider a more realistic situation where each network mining task is performed on a relatively small query set of objects [10], and focus on efficiently leveraging the appropriate portion of the whole massive network to facilitate query-specific knowledge discovery. To the best of our knowledge, this is first research effort to tackle network mining through this novel perspective.

### B. Reinforcement Learning

Reinforcement learning studies the problem of automated decision making by learning policies to take actions and maximize a reward signal from the environment. Some recent significant progress has been made by combining advances in deep learning for sensory processing with reinforcement learning [52], resulting in the super-human performance on high-dimensional game controls [17].

As close to our work, several approaches have been proposed to solve the combinatorial optimization problems with reinforcement learning. [21] blends an attention mechanism into a sequence-to-sequence structure, in order to learn a conditional probability as a solution to graph mining tasks such as the TSP (Traveling Salesman Problem). They rely on significant training data and can hardly generalize to different tasks or larger networks. Policy gradient [18] is used to learn the conditional probability policy in [19], while [20] learns a greedy heuristic using deep Q-learning [53]. In general, they aim to learn policies that generalize to unseen instances from a certain data distribution. These methods have achieved notable

<sup>5</sup><https://github.com/iskandr/fancyimpute>

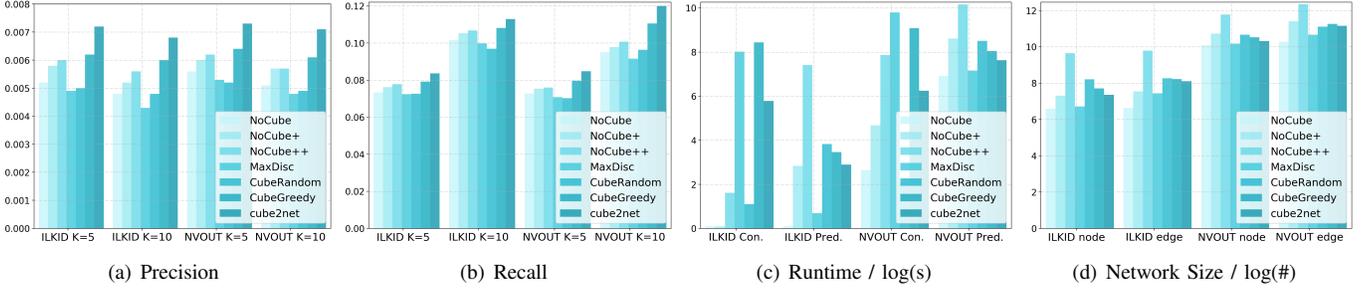


Fig. 5. Performance of network construction for link prediction on the ILKID Yelp subset.

successes on some classical graph combinatorial optimization problems. However, when it comes to our scenario of network construction with large action spaces, these algorithms are no longer applicable.

In this work, we connect continuous control policy gradient with powerful data cube organization through novel cell embedding. To the best of our knowledge, this is the first research effort to leverage reinforcement learning for efficient data cube exploration and quality network construction.

## VI. CONCLUSIONS

Network mining has been intensively studied by a wide research community. In this work, we stress on the novel angle of query-specific network construction, which aims to break the efficiency bottleneck of existing network mining algorithms and facilitate various downstream applications on particular query sets of objects. To achieve this goal, we propose to leverage the power of data cube, which significantly benefits the organization of massive real-world networks. Upon that, a novel reinforcement learning framework is designed to automatically explore the data cube structures and construct the most relevant query-specific networks. We demonstrate the effectiveness and efficiency of *cube2net* as a universal network data preprocessor in improving various network mining algorithms for different tasks, with a simple design of data cube and reinforcement learning, while many potential improvements can be explored in future works.

## ACKNOWLEDGEMENTS

Research was sponsored in part by U.S. Army Research Lab. under Cooperative Agreement No. W911NF-09-2-0053 (NSCTA), DARPA under Agreement No. W911NF-17-C-0099, National Science Foundation IIS 16-18481, IIS 17-04532, and IIS-17-41317, DTRA HDTRA11810026, and grant 1U54GM114838 awarded by NIGMS through funds provided by the trans-NIH Big Data to Knowledge (BD2K) initiative ([www.bd2k.nih.gov](http://www.bd2k.nih.gov)).

## APPENDIX A: Proof of Theorem III.1

The  $\xi$  neighborhood  $\mathcal{N}_\xi(c)$  of cell  $c$  is a  $\kappa$ -dimensional ball around  $\mathbf{u}_c$ . We firstly write out our actor function  $\mathbf{f}(S, c) = \mu(S, c)$  for the current particular state  $S$  and cell  $c$  being explored. The situation for the critic function  $\nu(S, c)$  is exactly

the same. According to our neural architecture described in Section III, we have

$$\begin{aligned} \mathbf{f}(S, c) &= \mathbf{g}(\mathbf{S} + \mathbf{u}_c) \\ &= \mathbf{h}^Q(\mathbf{h}^{Q-1}(\dots \mathbf{h}^1(\mathbf{S} + \mathbf{u}_c) \dots)), \end{aligned} \quad (10)$$

where

$$\mathbf{h}^q(\mathbf{u}) = \text{Sigmoid}(\mathbf{W}^q \mathbf{h}^{q-1}(\mathbf{u}) + \mathbf{b}^q). \quad (11)$$

$Q$  is the number of layers in the value function,  $\mathbf{W}^q$  and  $\mathbf{b}^q$  are the parameters of the  $q$ -th layer. The output of the last layer  $\mathbf{h}^Q$  is a  $\kappa$ -dimensional vector for the actor network  $\mu$  (a single number for the critic network  $\nu$ ), and  $\mathbf{h}^0(\mathbf{u}) = \mathbf{u}$ .

Now we prove the first property, *i.e.*,  $\forall c' \in \mathcal{N}_\xi(c)$ ,  $\|\mathbf{f}(S, c) - \mathbf{f}(S, c')\|_2^2 \leq \eta\xi$ .

*Proof.* Since  $c' \in \mathcal{N}_\xi(c)$ , we have

$$\|(\mathbf{S} + \mathbf{u}_c) - (\mathbf{S} + \mathbf{u}_{c'})\|_2^2 = \|\mathbf{u}_{c'} - \mathbf{u}_c\|_2^2 \leq \xi. \quad (12)$$

According to Eq. 10 and 11,  $\mathbf{g}(\cdot)$  is *smooth*, *i.e.*,  $\mathbf{g} \in C^\infty$  (this conclusion is actually non-trivial, but its proof is beyond the scope of this work). Based on the smoothness of  $\mathbf{g}$ , we have

$$\begin{aligned} \forall \xi, \mathbf{S}, \mathbf{u}_c, \mathbf{u}_{c'}, \|\mathbf{g}(\mathbf{S} + \mathbf{u}_c) - \mathbf{g}(\mathbf{S} + \mathbf{u}_{c'})\|_2^2 \\ \leq \|\nabla \mathbf{g}\|_2^2 \|\mathbf{u}_{c'} - \mathbf{u}_c\|_2^2, \end{aligned} \quad (13)$$

where

$$\begin{aligned} \|\nabla \mathbf{g}\|_2^2 &= \sup_{\mathbf{u} \in \Omega} \|\nabla \mathbf{g}(\mathbf{u})\|_2^2, \\ \Omega &= \{\mathbf{S} + \lambda \mathbf{u}_c + (1 - \lambda) \mathbf{u}_{c'}\}_{\lambda \in [0, 1]}. \end{aligned} \quad (14)$$

Therefore, we have

$$\begin{aligned} \forall \xi, S, c, c' \in \mathcal{N}_\xi(c), \exists \eta < \infty, s.t. \\ |\mathbf{f}(S, c) - \mathbf{f}(S, c')| \leq \eta\xi. \quad \square \end{aligned}$$

Now we prove the second property, *i.e.*,  $\forall c' \in \mathcal{N}_\xi(c)$ ,  $\|\nabla \mathbf{f}(S, c) - \nabla \mathbf{f}(S, c')\|_2^2 \leq \zeta\xi$ .

*Proof.* Similarly as in Eq. 13, based on the smoothness of  $\mathbf{g}$  (therefore, the smoothness of  $\nabla \mathbf{g}$ ), we have

$$\begin{aligned} \forall \xi, \mathbf{S}, \mathbf{u}_c, \mathbf{u}_{c'}, \|\nabla \mathbf{g}(\mathbf{S} + \mathbf{u}_c) - \nabla \mathbf{g}(\mathbf{S} + \mathbf{u}_{c'})\|_2^2 \\ \leq \|\mathbf{H}\|_2^2 \|\mathbf{u}_{c'} - \mathbf{u}_c\|_2^2, \end{aligned} \quad (15)$$

where  $\mathbf{H}$  is the Hessian matrix of  $\mathbf{g}$ , and

$$\begin{aligned} \|\mathbf{H}\|_2^2 &= \sup_{\mathbf{u} \in \Omega} \|\mathbf{H}(\mathbf{u})\|_2^2, \\ \Omega &= \{\mathbf{S} + \lambda \mathbf{u}_c + (1 - \lambda) \mathbf{u}_{c'}\}_{\lambda \in [0, 1]}. \end{aligned} \quad (16)$$

Therefore, similarly to  $\mathbf{f}$ , we have

$$\begin{aligned} \forall \xi, S, c, c' \in \mathcal{N}_\xi(c), \exists \zeta < \infty, s.t. \\ \|\nabla \mathbf{f}(S, c) - \nabla \mathbf{f}(S, c')\|_2^2 \leq \zeta\xi. \quad \square \end{aligned}$$

## APPENDIX B: Details of Cell Embedding

To enable efficient Gaussian exploration in a continuous space to fully leverage the data cube structure, we propose and design the novel process of *cell embedding* upon existing data cube technology. Particularly, we aim to compute a distributional representation for each cell, which captures their semantics *w.r.t.* the corresponding multi-dimensional labels. As we show in Section III, since such continuous cell representations can facilitate the utility estimation of close-by cells efficiently by avoiding the need of exhaustive search, they are critical to the success of an efficient reinforcement learning algorithm for efficient cube exploration and network construction.

We introduce a straightforward way of computing the cell embedding, based on the recent success of word embedding techniques, and theoretically show the effectiveness of it in capturing the semantics of cells. Following is a summarization of our cell embedding approach:

- **Step 1:** For each cell  $c \in \mathcal{C}$ , we firstly decompose it into the corresponding set of labels in the multiple cube dimensions, *i.e.*,  $c = \{l_c^1, l_c^2, \dots, l_c^P\}$ , where each  $l_c^p \in \mathcal{L}^p$  is the label of  $c$  in the  $p$ -th dimension. The embedding of  $c$  is then computed as  $\mathbf{u}_c = [\mathbf{u}_{l_c^1}, \mathbf{u}_{l_c^2}, \dots, \mathbf{u}_{l_c^P}]$ , where  $\mathbf{u}_{l_c^p}$  is the label embedding of label  $l_c^p$  and  $[\dots]$  is the vector concatenation operator.
- **Step 2:** For each label  $l \in \mathcal{L}$  ( $\mathcal{L} = \cup_{p=1}^P \mathcal{L}^p$ ), we firstly split it into single words, *i.e.*,  $l = \{w_l^1, w_l^2, \dots\}$ . The embedding of  $l$  is then computed as  $\mathbf{u}_l = \sum_i \mathbf{u}_{w_l^i}$ .
- **Step 3:** For each word  $w \in \mathcal{W}$ , where  $\mathcal{W}$  is an assumably complete vocabulary, look up the embedding of  $w$ , *i.e.*,  $\mathbf{u}_w$ , from a pre-trained word embedding table. For generality, we use the popular GloVe table<sup>6</sup>. Additionally, we assign the zero vector to all stop words and unmatched words.

Let us still take DBLP data as an example. Consider a cell  $c$ , *e.g.*,  $c = \langle 200X, \text{SIGMOD}, \text{data mining} \rangle$ , which has three dimensions, *i.e.*, *decade*, *venue* and *topic*. For the *decade* dimension, the label 200X is actually an aggregation of 10 words (2000-2009). Therefore, we can directly look up the 10 words from the embedding table and add them up as the label representation of 200X. For the textual dimensions like *topic*, we firstly split the phrase *data mining* into two words, *i.e.*, *data* and *mining*, then look up their embeddings in the table, and finally add them up. The *venue* label SIGMOD is firstly mapped back to its full name *Sig on management of data* and then processed in the same way as *topic*. After getting the embeddings of all three labels, the cell embedding is computed as a  $\kappa$ -dimensional vector concatenation of the corresponding label embeddings.

Although this approach of generating cell embedding is straightforward, we theoretically show that in this way, we can actually capture the semantic proximities among cells. Specifically, we prove the following theorem.

**Theorem VI.1.** *Given our particular cell embedding approach,  $\forall c_i, c_j \in \mathcal{C}$ ,  $\|\mathbf{u}_{c_i} - \mathbf{u}_{c_j}\|_2^2 \leq \epsilon S(c_i, c_j)$ , where  $\epsilon$  is a constant and  $S(c_i, c_j)$  is the semantic gap between  $c_i, c_j$ .*

*Proof.* This proof is based on the assumption that the pre-trained word embedding  $\mathcal{U}_W$  has captured word semantics in the vocabulary  $\mathcal{W}$ , so that  $\forall w_i, w_j \in \mathcal{W}$ ,  $\|\mathbf{u}_{w_i} - \mathbf{u}_{w_j}\|_2^2 < \epsilon S(w_i, w_j)$ , where  $S(w_i, w_j)$  is the semantic gap between  $w_i$  and  $w_j$ .

We further define the semantic gaps among labels. Specifically,  $\forall l_i, l_j \in \mathcal{L}^p$ , where  $p \in \mathcal{P}$  is any cube dimension, suppose  $l_i = \{w_i^1, w_i^2, \dots, w_i^{K_i}\}$ , where  $K_i$  is the number of words in  $l_i$ ;  $l_j = \{w_j^1, w_j^2, \dots, w_j^{K_j}\}$ , where  $K_j$  is the number of words in  $l_j$ , the semantic gap  $S(l_i, l_j)$  can then be defined as  $S(l_i, l_j) = \sum_{k_i=1}^{K_i} \sum_{k_j=1}^{K_j} S(w_i^{k_i}, w_j^{k_j})$ .

Now we first prove that the label embedding  $\mathcal{U}_L$  captures the label semantics, so that  $\forall l_i, l_j \in \mathcal{L}^p$ ,  $\|\mathbf{u}_{l_i} - \mathbf{u}_{l_j}\|_2^2 \leq \epsilon S(l_i, l_j)$ . For the label embeddings  $\mathcal{U}_L$ , ( $\forall \mathbf{u}_l \in \mathcal{U}_L$ ,  $\mathbf{u}_l = \sum_{k=1}^K \mathbf{u}_{w_l^k}$ ),

$$\begin{aligned}
& \sum_{k_i=1}^{K_i} \sum_{k_j=1}^{K_j} \|\mathbf{u}_{w_i^{k_i}} - \mathbf{u}_{w_j^{k_j}}\|_2^2 - \|\mathbf{u}_{l_i} - \mathbf{u}_{l_j}\|_2^2 \\
&= (K_j - 1) \sum_{k_i=1}^{K_i} \mathbf{u}_{w_i^{k_i}}^T \mathbf{u}_{w_i^{k_i}} + (K_i - 1) \sum_{k_j=1}^{K_j} \mathbf{u}_{w_j^{k_j}}^T \mathbf{u}_{w_j^{k_j}} \\
&\quad - 2 \sum_{k_i^1 \neq k_i^2; k_i^1, k_i^2=1}^{K_i} \mathbf{u}_{w_i^{k_i^1}}^T \mathbf{u}_{w_i^{k_i^2}} - 2 \sum_{k_j^1 \neq k_j^2; k_j^1, k_j^2=1}^{K_j} \mathbf{u}_{w_j^{k_j^1}}^T \mathbf{u}_{w_j^{k_j^2}} \\
&= (K_j - K_i) \left( \sum_{k_i=1}^{K_i} \mathbf{u}_{w_i^{k_i}}^T \mathbf{u}_{w_i^{k_i}} - \sum_{k_j=1}^{K_j} \mathbf{u}_{w_j^{k_j}}^T \mathbf{u}_{w_j^{k_j}} \right) \\
&\quad + \sum_{k_i^1 \neq k_i^2; k_i^1, k_i^2=1}^{K_i} \|\mathbf{u}_{w_i^{k_i^1}} - \mathbf{u}_{w_i^{k_i^2}}\|_2^2 \\
&\quad + \sum_{k_j^1 \neq k_j^2; k_j^1, k_j^2=1}^{K_j} \|\mathbf{u}_{w_j^{k_j^1}} - \mathbf{u}_{w_j^{k_j^2}}\|_2^2 \geq 0. \tag{17}
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
\|\mathbf{u}_{l_i} - \mathbf{u}_{l_j}\|_2^2 &\leq \sum_{k_i=1}^{K_i} \sum_{k_j=1}^{K_j} \|\mathbf{u}_{w_i^{k_i}} - \mathbf{u}_{w_j^{k_j}}\|_2^2 \\
&\leq \epsilon \sum_{k_i=1}^{K_i} \sum_{k_j=1}^{K_j} S(w_i^{k_i}, w_j^{k_j}) = \epsilon S(l_i, l_j) \tag{18}
\end{aligned}$$

Moreover,  $\forall c_i, c_j \in \mathcal{C}$ , suppose  $c_i = \{l_i^1, l_i^2, \dots, l_i^P\}$ , and  $c_j = \{l_j^1, l_j^2, \dots, l_j^P\}$ , where  $P$  is the number of cube dimensions, we then define the semantic gap  $S(c_i, c_j)$  as

$$S(c_i, c_j) = \sum_{p=1}^P S(l_i^p, l_j^p). \tag{19}$$

Now we prove that the cell embedding  $\mathcal{U}_C$  captures the cell semantics, so that  $\forall c_i, c_j \in \mathcal{C}$ ,  $\|\mathbf{u}_{c_i} - \mathbf{u}_{c_j}\|_2^2 \leq \epsilon S(c_i, c_j)$ .

For the cell embeddings  $\mathcal{U}_C$ , ( $\forall \mathbf{u}_c \in \mathcal{U}_C$ ,  $\mathbf{u}_c = [\mathbf{u}_{l_c^1}, \dots, \mathbf{u}_{l_c^P}]$ , where  $[\dots]$  denotes vector concatenation), the

<sup>6</sup><https://nlp.stanford.edu/projects/glove/>

proof is trivial because the norm of a concatenated vector equals to the sum of the norms of its components, *i.e.*,

$$\begin{aligned} & \|\mathbf{u}_{c_i} - \mathbf{u}_{c_j}\|_2^2 \\ &= \sum_{p=1}^P \|\mathbf{u}_{l_i^p} - \mathbf{u}_{l_j^p}\|_2^2 \\ &\leq \epsilon \sum_{p=1}^P S(l_i^p, l_j^p) = \epsilon S(c_i, c_j). \quad \square \end{aligned}$$

## REFERENCES

- [1] D. Chakrabarti, S. Funiak, J. Chang, and S. A. Macskassy, "Joint inference of multiple label types in large networks," in *ICML*, 2014.
- [2] J. J. McAuley and J. Leskovec, "Learning to discover social circles in ego networks," in *NIPS*, vol. 2012, 2012, pp. 548–56.
- [3] C. Yang, L. Zhong, L.-J. Li, and L. Jie, "Bi-directional joint inference for user links and attributes on large social graphs," in *WWW*, 2017, pp. 564–573.
- [4] C. Yang, L. Bai, C. Zhang, Q. Yuan, and J. Han, "Bridging collaborative filtering and semi-supervised learning: A neural approach for poi recommendation," in *KDD*, 2017, pp. 1245–1254.
- [5] M. Zhang and Y. Chen, "Weisfeiler-lehman neural machine for link prediction," in *KDD*, 2017, pp. 575–583.
- [6] T. Lyu, Y. Zhang, and Y. Zhang, "Enhancing the network embedding quality with structural similarity," in *CIKM*.
- [7] L. F. Ribeiro, P. H. Saverese, and D. R. Figueiredo, "struc2vec: Learning node representations from structural identity," in *KDD*, 2017.
- [8] J. Han, J. Pei, and M. Kamber, *Data mining: concepts and techniques*. Elsevier, 2011.
- [9] C. X. Lin, B. Ding, J. Han, F. Zhu, and B. Zhao, "Text cube: Computing ir measures for multidimensional text database analysis," in *ICDM*, 2008.
- [10] C. Yang, D. Teng, S. Liu, S. Basu, J. Zhang, J. Shen, C. Zhang, J. Shang, L. Kaplan, T. Harratty, and J. Han, "Cubenet: Multi-facet hierarchical heterogeneous network construction, analysis, and mining," in *KDD*, 2019.
- [11] C. Zhang, F. Tao, X. Chen, J. Shen, M. Jiang, B. Sadler, M. Vanni, and J. Han, "Taxogen: Unsupervised topic taxonomy construction by adaptive term embedding and clustering," in *KDD*, 2018, pp. 2701–2709.
- [12] F. Tao, C. Zhang, X. Chen, M. Jiang, T. Hanratty, L. Kaplan, and J. Han, "Doc2cube: Allocating documents to text cube without labeled data," in *ICDM*, 2018.
- [13] F. Tao, H. Zhuang, C. W. Yu, Q. Wang, T. Cassidy, L. R. Kaplan, C. R. Voss, and J. Han, "Multi-dimensional, phrase-based summarization in text cubes," *IDEB*, vol. 39, no. 3, pp. 74–84, 2016.
- [14] J. Shang, J. Liu, M. Jiang, X. Ren, C. R. Voss, and J. Han, "Automated phrase mining from massive text corpora," *TKDE*, 2018.
- [15] A. Gionis, M. Mathioudakis, and A. Ukkonen, "Bump hunting in the dark: Local discrepancy maximization on graphs," *TKDE*, pp. 529–542, 2017.
- [16] M. Sozio and A. Gionis, "The community-search problem and how to plan a successful cocktail party," in *KDD*. ACM, 2010, pp. 939–948.
- [17] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, p. 354, 2017.
- [18] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *NIPS*, 2000, pp. 1057–1063.
- [19] I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio, "Neural combinatorial optimization with reinforcement learning," in *ICLR*, 2017.
- [20] E. Khalil, H. Dai, Y. Zhang, B. Dilkina, and L. Song, "Learning combinatorial optimization algorithms over graphs," in *NIPS*, 2017, pp. 6351–6361.
- [21] O. Vinyals, S. Bengio, and M. Kudlur, "Order matters: Sequence to sequence for sets," in *ICLR*, 2015.
- [22] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [23] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.
- [25] Y. Sun, J. Han, X. Yan, P. S. Yu, and T. Wu, "Pathsim: Meta path-based top-k similarity search in heterogeneous information networks," *VLDB*, vol. 4, no. 11, pp. 992–1003, 2011.
- [26] B. Perozzi, R. Al-Rfou, and S. Skiena, "Deepwalk: Online learning of social representations," in *KDD*. ACM, 2014, pp. 701–710.
- [27] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, "Line: Large-scale information network embedding," in *WWW*, 2015, pp. 1067–1077.
- [28] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *KDD*, 2016, pp. 855–864.
- [29] J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su, "Arnetminer: extraction and mining of academic social networks," in *KDD*, 2008.
- [30] L. Liu, L. Xu, Z. Wangy, and E. Chen, "Community detection based on structure and content," in *ICDM*, 2015, pp. 271–280.
- [31] R. Mazumder, T. Hastie, and R. Tibshirani, "Spectral regularization algorithms for learning large incomplete matrices," *JMLR*, vol. 11, no. Aug, pp. 2287–2322, 2010.
- [32] Y. Liu, T.-A. N. Pham, G. Cong, and Q. Yuan, "An experimental evaluation of point-of-interest recommendation in location-based social networks," *VLDB*, vol. 10, no. 10, pp. 1010–1021, 2017.
- [33] N. Barbieri, F. Bonchi, and G. Manco, "Topic-aware social influence propagation models," in *ICDM*. IEEE, 2012, pp. 81–90.
- [34] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing, "Mixed membership stochastic blockmodels," *JMLR*, pp. 1981–2014, 2008.
- [35] J. Yang, J. McAuley, and J. Leskovec, "Community detection in networks with node attributes," in *ICDM*, 2013, pp. 1151–1156.
- [36] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *NIPS*, 2013, pp. 3111–3119.
- [37] C. Yang, M. Liu, V. W. Zheng, and J. Han, "Node, motif and subgraph: Leveraging network functional blocks through structural convolution," in *ASONAM*, 2018, pp. 47–52.
- [38] Y. Shi, X. He, N. Zhang, C. Yang, and J. Han, "User-guided clustering in heterogeneous information networks via motif-based comprehensive transcription,"
- [39] C. Yang, H. Lu, and K. C. Chang, "Cone: Community oriented network embedding," in *arXiv preprint arXiv:1709.01554*, 2017.
- [40] C. Yang, D. H. Hoang, T. Mikolov, and J. Han, "Place deduplication with embeddings," in *WWW*, 2019, pp. 3420–3426.
- [41] C. Yang, M. Liu, F. He, X. Zhang, J. Peng, and J. Han, "Similarity modeling on heterogeneous networks via automatic path discovery," in *ECML-PKDD*, 2018, pp. 37–54.
- [42] C. Yang, C. Zhang, J. Han, X. Chen, and J. Ye, "Did you enjoy the ride: Understanding passenger experience via heterogeneous network embedding," in *Proc. of 2018, IEEE International Conference on Data Engineering*.
- [43] C. Yang, Y. Feng, P. Li, Y. Shi, and J. Han, "Meta-graph based hin spectral embedding: Methods, analyses, and insights," in *ICDM*, 2018, pp. 657–666.
- [44] C. Yang and K. Chang, "Relationship profiling over social networks: Reverse smoothness from similarity to closeness," in *ICDM*, 2019, pp. 342–350.
- [45] M. Gomez Rodriguez, J. Leskovec, and A. Krause, "Inferring networks of diffusion and influence," in *KDD*, 2010, pp. 1019–1028.
- [46] J. McAuley, R. Pandey, and J. Leskovec, "Inferring networks of substitutable and complementary products," in *KDD*, 2015.
- [47] D. Hallac, Y. Park, S. Boyd, and J. Leskovec, "Network inference via the time-varying graphical lasso," in *KDD*. ACM, 2017, pp. 205–213.
- [48] L. Zhou, Y. Yang, X. Ren, F. Wu, and Y. Zhuang, "Dynamic network embedding by modeling triadic closure process," in *AAAI*, 2018.
- [49] J. Li, H. Dani, X. Hu, J. Tang, Y. Chang, and H. Liu, "Attributed network embedding for learning in a dynamic environment," in *CIKM*, 2017, pp. 387–396.
- [50] J. Ma, P. Cui, and W. Zhu, "Depthlqp: Learning embeddings of out-of-sample nodes in dynamic networks," *AAAI*, 2018.
- [51] A. Vukotic, N. Watt, T. Abedrabbo, D. Fox, and J. Partner, *Neo4j in action*. Manning Publications Co., 2014.
- [52] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *NIPS*, 2012, pp. 1097–1105.
- [53] H. Dai, B. Dai, and L. Song, "Discriminative embeddings of latent variable models for structured data," in *ICML*, 2016, pp. 2702–2711.