Query-Specific Knowledge Summarization with Entity Evolutionary Networks

Carl Yang*, Lingrui Gan*, Zongyi Wang, Jiaming Shen, Jinfeng Xiao, Jiawei Han

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

{jiyang3, lgan6, zwang195, js2, jxiao13, hanj}@illinois.edu

ABSTRACT

Given a query, unlike traditional IR that finds relevant documents or entities, in this work, we focus on retrieving both entities and their connections for insightful knowledge summarization. For example, given a query “computer vision” on a CS literature corpus, rather than returning a list of relevant entities like “cnn”, “imagenet” and “svm”, we are interested in the connections among them, and furthermore, the evolution patterns of such connections along particular ordinal dimensions such as time. Particularly, we hope to provide structural knowledge relevant to the query, such as “svm” is related to “imagenet” but not “cnn”. Moreover, we aim to model the changing trends of the connections, such as “cnn” becomes highly related to “imagenet” after 2010, which enables the tracking of knowledge evolutions. In this work, to facilitate such a novel insightful search system, we propose SETEVOLVE, which is a unified framework based on nonparanormal graphical models for evolutionary network construction from large text corpora. Systematic experiments on synthetic data and insightful case studies on real-world corpora demonstrate the utility of SETEVOLVE.

KEYWORDS

knowledge summaries, network construction, evolution analysis

ACM Reference Format:

Carl Yang*, Lingrui Gan*, Zongyi Wang, Jiaming Shen, Jinfeng Xiao, Jiawei Han. 2019. Query-Specific Knowledge Summarization with Entity Evolutionary Networks. In *The 28th ACM International Conference on Information and Knowledge Management (CIKM '19)*, November 3–7, 2019, Beijing, China. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3357384.3358068>

1 INTRODUCTION

Given a query, for knowledge summarization, it is often crucial to understand its relevant entities and their inherent connections. More interestingly, the entities and connections may evolve over time or other ordinal dimensions, indicating dynamic behaviors and changing trends. For example, in biological scientific literature, the study of a disease might focus on particular genes for a period

of time, and then shift to some others due to a technology breakthrough. Capturing such entity connections and their changing trends can enable various tasks including the analysis of concept evolution, forecast of future events and detection of outliers.

Related work. Traditional information retrieval (IR) aims at returning a ranked list of documents according to their relevance to the query [1]. To understand the results and distill knowledge, users then need to pick out and read some of the documents, which requires tedious information processing and often leads to inaccurate conclusions. To deal with this, recent works on entity search [7, 14] aim to search for entities instead of documents, but they only return lists of isolated entities, thus incapable of providing insights about entity connections. Existing works on graph keyword search [8] and natural language processing [2, 15] have been using graph structures for query result or knowledge summarization, but they do not consider the evolution of entity connections.

Present work. In this work, we advocate for the novel task of *entity evolutionary network construction* for query-specific knowledge summarization, which aims to return a set of query-relevant entities together with their evolutionary connections modeled by a series of networks automatically constructed from large-scale text corpora. Mathematically, we model entities as variables in a complex dynamic system, and estimate the connections among them based on their discrete occurrence within the documents.

Regarding techniques, recent existing works on network structure estimation have studied the inference of time-varying networks [5, 16]. However, in our novel problem setting, there are two unique challenges: 1) Identifying the query-relevant set of entities from text corpora and 2) Constructing the evolutionary entity connections based on discrete entity observations. To orderly deal with them, we develop SETEVOLVE, a unified framework based on the principled nonparanormal graphical models [4, 10].

For the first challenge, we assume a satisfactory method has been given for the retrieval of a list of documents based on query relevance (e.g., BM25 [1]) and develop a general post-hoc method based on document rank list cutoff for query-relevant entity set identification. The cutoff is efficiently computed towards entity-document support recovery for accurate network construction, with theoretical and empirical analysis on the existence of optimality.

For the second challenge, we formulate the problem as an evolutionary graphical lasso, by modeling entities as variables in an evolving Markov random field, and detecting their inherent connections by estimating the underlying inverse covariance matrix. Moreover, we leverage the robust nonparanormal transformation to deal with the ordinal discrete entity observations in documents. Through theoretical analysis, we show that our model can capture the true conditional connections among entities, while its computational efficiency remains the same to standard graphical lasso.

*Both authors contribute equally.

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

CIKM '19, November 3–7, 2019, Beijing, China

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6976-3/19/11...\$15.00

<https://doi.org/10.1145/3357384.3358068>

2 SETEVOLVE

2.1 Overview

Given a query Q on a text corpus C , we aim to provide a *clear* and *interpretable* summarization S over the retrieved knowledge $\mathcal{K}$, which is extracted from C based on the relevance to Q . In principle, S should help users easily understand the key *concepts* within $\mathcal{K}$, as well as their *interactions* and *evolutions*.

To achieve this goal, we get motivated by recent success on text summarization with concept maps [2], and propose to represent S as a series of concept networks $\mathcal{S} = \{\mathcal{N}_1, \dots, \mathcal{N}_T\}$, by stressing the evolutionary nature of concept interactions. In each network $\mathcal{N}_t \in \mathcal{S}$, t denotes a window in an arbitrary ordinal dimension of interest, such as time, product price, user age, *etc.* Without loss of generality, we will focus on time as the *evolving dimension* in the following. We further decompose $\mathcal{N}_t$ into $\mathcal{N}_t = \{\mathcal{V}_t, \mathcal{E}_t\}$, where $\mathcal{V}_t$ is the set of relevant concepts, and $\mathcal{E}_t$ is the set of concept interactions, both in the time window denoted by t .

In this work, we propose SETEVOLVE, a unified framework based on nonparanomal graphical models [2, 4, 5, 10, 18], for the novel problem of constructing query-specific knowledge summarization from free text corpora. In particular, empowered by recent available structured data of knowledge bases and well-developed techniques of entity discovery [6, 12], we leverage *entities* to represent concepts of interest, and derive a theoretically sound method to identify the stable set of query-relevant entities. Moreover, to model concept interactions in a principled way, we consider the corresponding entities as interconnected variables in a complex system and leverage graphical models to infer their connection patterns based on their occurrence in the text corpus as variable observations. Finally, to capture knowledge evolution, we consider observations associated with time (or other evolving dimension of interest), and jointly estimate a series of networks for characterizing changing trends.

2.2 Entity Set Identification

To construct the knowledge summarization S , we first need to identify a set of entities to represent the key concepts that users care about. Entity set search without the consideration of network evolution is well studied [13]. Differently from them, in this work, we develop a general post-hoc entity set identification method based on the theoretical guarantees of nonparanomal graphical models for document-entity support recovery, as shown in Theorem 2.2. Particularly, we assume that based on a list of documents $\mathcal{D} = [d_1, d_2, \dots]$ ranked by relevance, the *optimal* query-relevant entity set $\mathcal{V}^*$ will appear consistently within the top ranked documents. In other words, as we bring in more documents, we get more complete entity sets and supportive information for constructing the entity network, but too many documents will bring in less relevant entities and redundant information. Therefore, we aim to find an optimal cutoff on $\mathcal{D}$ to extract $\mathcal{V}^*$ and further facilitate accurate evolutionary network construction.

Without loss of generality, we use the classic BM25 [1] to retrieve documents $\mathcal{D}$, while various other advanced methods like [17] can be trivially plugged in to better meet users' particular information need regarding query relevance. To extract entities from $\mathcal{D}$, we assume the availability of non-evolving knowledge bases and utilize an entity linking tool called CMNS [6] to convert $\mathcal{D}$ into an entity

set list $\mathcal{V}_D = [\mathcal{V}_1, \mathcal{V}_2, \dots]$, where $\mathcal{V}_i \in \mathcal{V}$ is the set of entities in $d_i \in \mathcal{D}$. To quantify the optimality, we propose to track the following metric sequentially as documents are taken in one-by-one and use it to determine the optimal entity set $\mathcal{V}^*$:

$$\gamma(n) = \frac{n}{\log |\mathcal{V}_1 \cup \mathcal{V}_2 \dots \cup \mathcal{V}_n|}. \quad (1)$$

where $|\mathcal{V}|$ is the size of an entity set $\mathcal{V}$, and n is the number of documents taken in from the rank list $\mathcal{D}$.

The optimal set $\mathcal{V}^*$ extracted from $\mathcal{D}^*$ is determined based on two criteria: 1) the convergence of $\gamma'(n)$, which is used to measure the completeness of the entity set; and 2) $\gamma(n)$ is larger than a prefixed threshold (e.g., 10), which provides the theoretical guarantees in Theorem 2.2 for document-entity support recovery.

Figure 1: Entity-document support consistency (holder).

Figure 1 shows $\gamma'(n)$ computed on the PubMed and Yelp datasets with different queries. As we can see, the derivation of $\gamma(n)$ often converges to 0 rapidly, which corroborates our assumption of entity consistency. Particularly, as we take in more documents, the total number of entities converges, so we can cutoff the document list at $\mathcal{D}^*$ once the aforementioned two criteria are met, and use $\mathcal{D}^*$ to identify $\mathcal{V}^*$ and construct the evolutionary networks over $\mathcal{V}^*$.

2.3 Evolutionary Network Construction

After fixing $\mathcal{V}^*$, we consider the retrieval of connections among $\mathcal{V}^*$. A straightforward way to construct links in an entity network is to compute and threshold the document-level entity co-occurrence as suggested in [11]. However, it is not clear how to weigh the links and set the thresholds. More importantly, because the thresholding method [11] only models the pairwise *marginal* dependence between entities v_i and v_j , and does not consider the interactions between $\{v_i, v_j\}$ and all the other entities $v_{-\{i,j\}}$, the links generated may be messy and less insightful, as shown in [13].

To find the essential entity connections in a principled way, we propose to uncover the underlying connection patterns among entities as a graphical model selection problem [3, 4]. Assume we have n observations $V^1, \dots, V^n$ of a set of entities $\mathcal{V} = \{v_1, \dots, v_m\}$. In a standard GGM, we assume the n observations identically and independently distributed (*iid*) follow a multivariate Gaussian distribution $V^1, \dots, V^n \sim \mathcal{N}(0, \Theta^{-1})$, and our target is to detect the support of the precision matrix Θ . It is well-known that $\Theta_{ij} = 0$ in standard GGM is equivalent to the conditional independence of v_i and v_j , given all other variables, i.e., $v_i \perp v_j \mid v_{-\{i,j\}}$. If the Gaussianity and *iid* assumptions of GGM hold, insightful conditional dependence network links can be generated systematically.

However, the GGM assumptions do not exactly fit the context we consider, i.e., entity evolutionary network construction, because

1) GGM assumes observations are Gaussian, but entity occurrences in documents are discrete; 2) The entity network is evolving over time, so the *iid* assumption does not hold.

Motivated by previous works in graphical models that address the non-Gaussian data [10] and time dependent data [9, 19], we propose an *evolutionary nonparanormal graphical model*, which detects the conditional dependence structure even when data is both *discrete* and *evolving*. Without loss of generality, in the following, we use *time* as the evolving dimension to describe our model.

Denote the m -dimensional discrete observation at time t as $Y^{(t)} = (Y_1^{(t)}, \dots, Y_m^{(t)})$. In the proposed model, we assume $Y^{(t)}$ satisfies the following relationship:

$$f(Y^{(t)}) \sim \mathcal{N}(0, \Theta(t)^{-1}), \quad (2)$$

where $f(\cdot)$ is an m -dimensional Gaussian copula transformation function defined in Eq. (3) and $\Theta(t)$ is assumed to evolve smoothly from t to $t + 1$. A definition of the smoothness assumption of the evolution pattern we use is as Assumption S in [9], where the smoothness is quantified through the boundedness of the first and second derivative of the changes in $\Theta(t)_{i,j}$ over time t . No further assumption on the parameter (distribution) form of the evolution pattern is required for the model.

As shown in the following proposition, the model (2) can capture the conditional dependence links even when data is *evolving* and *discrete*. The proposition can be shown through standard matrix calculation and we refer to Lemma 2 in [10] for the proof.

PROPOSITION 2.1. *If $\Theta(t)_{ij} = 0$, $Y_i^{(t)}$ and $Y_j^{(t)}$ are conditionally independent, i.e., $Y_i^{(t)} \perp\!\!\!\perp Y_j^{(t)} \mid Y_{-\{i,j\}}^{(t)}$.*

Motivated by [10], we utilize a Gaussian copula transformation function $f(\cdot)$ to handle the non-Gaussianity in the data, and the function $f(\cdot)$ is defined as follows:

$$f_j(x) = u_j + \sigma_j \Phi^{-1}(F_j(x)). \quad (3)$$

In (3), u_j and σ_j are the empirical mean and empirical standard deviation of Y_j respectively, $\Phi^{-1}(\cdot)$ is the quantile function of the standard norm distribution, and $F_j(x)$ is the Winsorized estimator of the empirical distribution suggested in [10] and defined as:

$$F_j(x) = \begin{cases} \delta_n, & \text{if } \hat{F}_j(X) < \delta_n; \\ \hat{F}_j(X), & \text{if } \delta_n \leq \hat{F}_j(X) \leq 1 - \delta_n; \\ 1 - \delta_n, & \text{if } \hat{F}_j(X) > 1 - \delta_n, \end{cases} \quad (4)$$

where $\delta_n = \frac{1}{4n^{1/4}\sqrt{\pi \log n}}$ and $\hat{F}_j(X)$ is the empirical cumulative distribution function.

This transformation allows us to estimate the evolving pattern of $\Theta(t)$ through a kernel based method [9, 19], where our target is to minimize the following kernel-based objective function:

$$\hat{\Theta}(t) = \arg \min_{\Theta(t) > 0} \left\{ -\log |\Theta| + \text{tr}(\Theta S_f(t)) + \lambda |\Theta| \right\}. \quad (5)$$

$S_f(t) = \sum_{t'} w_{t'}^t f(Y^{(t')}) f(Y^{(t')})'$ is a kernel estimator of the sample covariance matrix at time t , with weights $w_{t'}^t$ defined as:

$$w_{t'}^t = \frac{K_h(|t - t'|)}{\sum_{t'} K_h(|t - t'|)}, \quad (6)$$

where $K_h(\cdot)$ is a kernel function such as the box kernel $K_h(x) = \frac{1}{2} \times \mathbb{I}\{\frac{x}{h} \in [-1, 1]\}$.

2.4 Theoretical Analysis

In the following theorem, we show the accuracy of SETEVOLVE in detecting the true entity links.

THEOREM 2.2. *With the same assumptions of [9] on the true evolving graphs, define the maximum graph node degree as d . Suppose $S_f(t)$ is estimated using a kernel with bandwidth $h = O(n^{-1/6})$. If the number of documents n satisfies $n > C(d^2 \log m)^3$ for a sufficiently large constant C and we choose tuning parameter $\lambda = O(n^{-1/6} \log n \sqrt{\log m})$, then our estimation procedure can detect the links correctly with probability converging to 1.*

PROOF. $\|S_f(t) - (\Theta^0(t))^{-1}\|_\infty \leq \|S_f(t) - ES_f(t)\|_\infty + \|ES_f(t) - (\Theta^0(t))^{-1}\|_\infty$, where the first term is $O_p(\sqrt{\frac{\log m \log^2 n}{hn^{1/2}}})$ by Theorem 4 in [10] and the second is $O_p(h)$ by Lemma 11 in [9]. If the bandwidth h is chosen as $h = O(n^{-1/6})$, we have $\|S_f(t) - (\Theta^0(t))^{-1}\|_\infty = O_p(\sqrt{\frac{\log m \log^2 n}{n^{1/3}}})$, i.e., $\|S_f(t) - (\Theta^0(t))^{-1}\|_\infty < \sqrt{\frac{\log m \log^2 n}{n^{1/3}}}$.

Assuming conditions in Theorem 2.2 hold, it implies Proposition 3 and 4 in [9] hold with $w(n, m) = n^{-1/6} \log n \sqrt{\log m}$ and $\lambda = Cn^{-1/6} \log n \sqrt{\log m}$. We then apply Proposition 2 in [9] to conclude our proof on link detection accuracy. $\square$

In terms of the computation efficiency, in SETEVOLVE, we first compute $S_f(t)$ using equation (6), and then compute $\Theta(t)$ for each t using the state-of-the-art Graphical Lasso algorithm [3]. The computation complexity is thus exactly same as the state-of-the-art evolutionary network inference algorithms [5, 16].

3 SYNTHETIC EXPERIMENTS

Since there is no ground-truth for knowledge summarization with evolutionary networks, we follow the common practice in recent works on network inference [5, 16] to demonstrate the effectiveness of SETEVOLVE through comprehensive synthetic experiments.

Data Generation. Following Hallac et al. [5], we generate two evolutionary networks of sizes 20 and 100 and then randomly generate the ground-truth covariance and sample discrete observations over 100 timestamps, where the global and local evolutions occur at time $t = 50$. At each t , we generate 10 independent samples from the true distribution, with maximum variable value set to 10.

Compared Algorithms. We compare SETEVOLVE to two baselines: 1) *Static* [10], which applies the nonparanormal transformation for discrete observations, but aggregate observations from all timestamps to infer a single static network without evolutions, and 2) *TVGL* [5], which leverages graphical lasso for evolutionary network construction but does not consider discrete observations.

Performance Evaluations. Table 1 shows the performance of compared methods. Following [5, 16], we compute macro-F1 scores for link reconstruction over evolutionary networks. As we can see, all algorithms perform better on smaller networks, partly because the observations are denser and the essential correlations are easier to capture. Moreover, all algorithms perform better with local evolutions, because the overall changes are smaller, while *TVGL* and SETEVOLVE perform better on global evolutions compared with *Static*. Finally, SETEVOLVE consistently outperforms both *Static* and *TVGL* on networks of different sizes and evolution types, which

indicates its effectiveness in detecting correlations among evolutionary discrete variables. The scores of TVGL are lower than those in the original work due to the discreteness of variables. Besides accuracy, SETEVOLVE can construct networks with fewest links, generating clear views of networks for deriving the most essential entity connections.

Networks	Static		TVGL		SetEvolve	
	F1	E	F1	E	F1	E
20-global	0.3698	97	0.4058	104	0.4592	58
20-local	0.4567	90	0.5148	85	0.6260	35
100-global	0.3239	1573	0.3372	2020	0.3670	687
100-local	0.3684	1548	0.3683	1895	0.4045	637

Table 1: Link detection performance.

We evaluate the robustness of SETEVOLVE by randomly adding Poisson noises to synthetic observations and show results in Figure 2. We also find that runtimes of TVGL and SETEVOLVE are roughly the same, which indicates the good scalability of SETEVOLVE.

Figure 2: Performance with noisy observations.

4 CASE STUDIES

We show the practical utility of SETEVOLVE on three datasets: *PubMed* with 14 GB texts of 3.3M articles¹, *Yelp* with 5.4 GB texts of 6M reviews², and *Semantic* with 19 GB texts of 20M papers³.

Figure 3, 4 and 5 show the constructed networks on each corpus. The results are insightful. Figure 3, for example, shows that the long-used OGTT test for Diabetes was replaced by a new standard HbA1 test in the 2000's, and Type 2 Diabetes (DM2) draws more research attention with newly found proteins like Osteocalcin in recent years. In Figure 4, we notice that Chinese restaurants get poorly rated often due to the quality of classic dishes (e.g., orange chicken, sour soup), while the higher ratings focus on authentication and customer service. In Figure 5, we observe that papers about MOOC and transfer learning have accumulated fewer citations, partly because these fields are quite new, whereas papers about question answering, knowledge base and language model are usually cited more. We will release the full implementation of our framework and interested readers can pose queries to find more insightful knowledge.

Figure 3: Querying Diabetes on PubMed.

¹<https://www.ncbi.nlm.nih.gov/pubmed/>

²<https://www.yelp.com/dataset/challenge>

³<https://labs.semanticscholar.org/corpus/corpus/archive>

Figure 4: Querying Chinese food on Yelp.

Figure 5: Querying Knowledge summarization on Semantic.

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 HD-TRA11810026, and grant 1U54GM114838 awarded by NIGMS through funds provided by the trans-NIH BD2K initiative.

REFERENCES

- [1] Ricardo Baeza-Yates, Berthier de Araújo Neto Ribeiro, and others. 2011. *Modern information retrieval*. New York: ACM Press; Harlow, England: Addison-Wesley.
- [2] Tobias Falke and Iryna Gurevych. 2017. Bringing Structure into Summaries: Crowdsourcing a Benchmark Corpus of Concept Maps. In *ACL*.
- [3] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. 2008. Sparse inverse covariance estimation with the graphical lasso. In *Biostatistics*.
- [4] Lingrui Gan, Naveen N Narisetty, and Feng Liang. 2018. Bayesian Regularization for Graphical Models with Unequal Shrinkage. In *JASA*.
- [5] David Hallac, Youngsuk Park, Stephen Boyd, and Jure Leskovec. 2017. Network Inference via the Time-Varying Graphical Lasso. In *KDD*.
- [6] Faegheh Hasibi, Krisztian Balog, and Svein Erik Bratsberg. 2017. Entity linking in queries: Efficiency vs. effectiveness. In *ECIR*.
- [7] Shoaib Jameel, Zied Bouraoui, and Steven Schockaert. 2017. MEmBER: Max-Margin Based Embeddings for Entity Retrieval. In *SIGIR*.
- [8] Varun Kacholia, Shashank Pandit, Soumen Chakrabarti, S Sudarshan, Rushi Desai, and Hrishikesh Karambelkar. 2005. Bidirectional expansion for keyword search on graph databases. In *Vldb*.
- [9] Mladen Kolar and Eric Xing. 2011. On time varying undirected graphs. In *ICAI*.
- [10] Han Liu, John Lafferty, and Larry Wasserman. 2009. The nonparametric: Semiparametric estimation of high dimensional undirected graphs. In *JMLR*.
- [11] Hao Peng, Jianxin Li, Yu He, Yaopeng Liu, Mengjiao Bao, Lihong Wang, Yangqiu Song, and Qiang Yang. 2018. Large-Scale Hierarchical Text Classification with Recursively Regularized Deep Graph-CNN. In *WWW*.
- [12] Jiaming Shen, Zeqiu Wu, Dongming Lei, Jingbo Shang, Xiang Ren, and Jiawei Han. 2017. SetExpan: Corpus-Based Set Expansion via Context Feature Selection and Rank Ensemble. In *ECML/PKDD*.
- [13] Jiaming Shen, Jinfeng Xiao, Xinwei He, Jingbo Shang, Saurabh Sinha, and Jiawei Han. 2018. Entity set search of scientific literature. In *SIGIR*.
- [14] Jiaming Shen, Jinfeng Xiao, Yu Zhang, Carl Yang, Jingbo Shang, Saurabh Sinha, Peipei Ping, Zhiyong Lu, and Jiawei Han. 2018. SetSearch+: Entity-Set-Aware Search and Mining for Scientific Literature. In *KDD*.
- [15] Zhaowei Tan, Changfeng Liu, Yuning Mao, Yunqi Guo, Jiaming Shen, and Xinbing Wang. 2016. AceMap: A Novel Approach towards Displaying Relationship among Academic Literatures. In *WWW*.
- [16] Federico Tomasi, Veronica Tozzo, Saverio Salzo, and Alessandro Verri. 2018. Latent variable time-varying network inference. In *KDD*.
- [17] Xuanhui Wang, Cheng Li, Nadav Golbandi, Michael Bendersky, and Marc Najork. 2018. The lambdaloss framework for ranking metric optimization. In *CIKM*.
- [18] Carl Yang, Lin Zhong, Li-Jia Li, and Luo Jie. 2017. Bi-directional joint inference for user links and attributes on large social graphs. In *WWW*. 564–573.
- [19] Shuheng Zhou, John Lafferty, and Larry Wasserman. 2008. Time varying undirected graphs. In *COLT*.