

Tell Me How to Survey: Literature Review Made Simple with Automatic Reading Path Generation

Jiayuan Ding^{*†}
Michigan State University
dingjia5@msu.edu

Tong Xiang^{*†}
Virginia Tech
xtong@vt.edu

Zijing Ou^{*}
Sun Yat-sen University
zijingou.mail@gmail.com

Wangyang Zuo
Zhejiang University of Technology
zuowangyang@foxmail.com

Ruihui Zhao
Tencent Jarvis Lab
zacharyzhao@tencent.com

Chenhua Lin
University of Sheffield
c.lin@sheffield.ac.uk

Yefeng Zheng
Tencent Jarvis Lab
yefengzheng@tencent.com

Bang Liu[✉]
Université de Montréal & Mila
bang.liu@umontreal.ca

Abstract—Recent years have witnessed the dramatic growth of paper volumes with plenty of new research papers published every day, especially in the area of computer science. How to glean papers worth reading from the massive literature to do a quick survey or keep up with the latest advancement about a specific research topic has become a challenging task. Existing academic search engines return relevant papers by individually calculating the relevance between each paper and query. However, such systems usually omit the prerequisite chains of a research topic and cannot form a meaningful reading path. In this paper, we introduce a new task named Reading Path Generation (RPG) which aims at automatically producing a path of papers to read for a given query. To serve as a research benchmark, we further propose SurveyBank, a dataset consisting of large quantities of survey papers in the field of computer science as well as their citation relationships. Furthermore, we propose a graph-optimization-based approach for reading path generation which takes the relationship between papers into account. Extensive evaluations demonstrate that our approach outperforms other baselines. A real-time Reading Path Generation (RePaGer) system has been also implemented with our designed model. Our source code and SurveyBank dataset can be found here ¹.

Index Terms—Reading Path Generation, Academic Search Engine, Automatic Dataset Creation

I. INTRODUCTION

With the explosion of the quantity of scientific papers, it becomes more challenging for researchers to glean papers worth reading from a large amount of literature for a quick review or keeping up with the latest advancement. Given a research topic that a user is not familiar with, he/she may try to read some survey papers about the topic and refer to the reference list in the surveys to get more details. However, a high-quality and up-to-date survey is not always available for any given research topic. A user may also utilize academic search engines to find out related papers to read. Existing academic search engines usually retrieve a list of articles by calculating semantic similarity between user’s query and

each individual article separately, ignoring the prerequisite relationship among searching results. Therefore, it will be highly valuable if we can retrieve a list of core papers with reading order for a given query to show users a reading path. In this way, a user can read the generated reading path and quickly get familiar with a new research field.

In recent years, research about scientific paper analysis has attracted a lot of attention. However, most existing works focus on citation recommendation [1]–[5], which aims at recommending citations for a given text. Another category of research works focus on citation classification [6]–[11], aiming to reason author’s intention for citing a paper. There are also research works about generating a reading list to users. [12] proposed ThemedPageRank to create reading lists by re-ranking topic-related papers. [13] built up a reading list for a given query set by exploring collaborative and content-based filtering techniques. But these methods ignore the prerequisite relationship among the returned papers. [14] designed a concept graph-based approach to generate structured reading lists, which takes into account the prerequisite concepts of the topic. However, this method strongly relies on the quality of existing concept graph.

Previous research about reading list generation mainly solves the problem of “what to read” by producing a list of papers for a given research topic. However, to serve users unfamiliar with a research field, two critical problems are omitted by existing research: “how to read” and “how to understand”. For the first problem, we need to organize different research works in a structured manner, so that a user can read the papers according to the paths in the structure, and get a sense of how the research about a specific topic evolves during history. Note that simply organizing papers by publication time is not enough, as the correlation between different research is overlooked. For the second problem, we need to include some papers in our generated path which may be not directly related to the queried research topic, but are quite helpful or essential to understand the highly related papers in our generated list.

Motivated by above analysis, in this paper, we introduce a new task *Reading Path Generation (RPG)*, which aims to

^{*} Equal contributions.

[†] The work was done during internship at Mila.

[✉] Corresponding author. Bang Liu is also affiliated with Canada CIFAR AI Chair.

¹<https://github.com/JiayuanDing100/Reading-Path-Generation>

automatically create a reading path for a given query. Our task is designed to help researchers who are new to a field to do a quick survey by solving the two problems mentioned above. Specifically, we are aiming to retrieve a list of papers which includes the research works highly related to the given query, as well as prerequisite research works that are essential to understand the highly related works. Furthermore, we organize these papers to form a reading path, which takes both the chronological order and the correlation between different research works into account. We hope that by reading the retrieved papers according to the reading path, a user can easily understand each research work, quickly complete a survey and form a landscape of the research history of any topic.

Due to the lack of benchmark datasets for evaluating Reading Path Generation, we further propose the *SurveyBank* dataset which is specifically designed for our newly proposed task. We create SurveyBank without any human labeling based on two key insights. The first one is that a lot of existing high-quality survey papers are already available, and each of them is written by professional researchers, encompassing highly valuable information about how to form a survey about a research topic. Based on this observation, we utilize the key phrases extracted from survey titles as the inputs of RPG task, and reading paths inferred from surveys’ reference lists as true labels. We initially collect 41,194 surveys in the field of computer science from Google Scholar and S2ORC [15]. After that, 9,321 high quality surveys are selected to form SurveyBank. It covers 10 topics of computer science like Artificial Intelligence, Computer Network and Database. The second key insight is that the citation relationship between different papers reflects the coherence between different papers. Therefore, we construct a large-scale scientific citation graph with 6 million papers from S2ORC [15] in the field of computer science by citation relationship, and this citation graph covers reference lists of each survey. To the best of our knowledge, SurveyBank is currently the largest survey dataset.

Last but not least, we also propose and implement a Real-time Reading Path Generation System (RePaGer) which takes key phrases as inputs and generates a reading path of papers for users to read. In RePaGer, we start by getting back top- K articles called initial seed papers from Google Scholar. After that, we re-allocate seeds by first capturing neighbors of initial seed papers as seed candidates and then selecting articles with high co-occurrence. We infer that these new selected seeds are more likely to be prerequisite concepts of a topic since multiple articles directly relevant to the topic cite them in their papers. Finally we propose Node-Edge Weighted Steiner Tree (NWST), a graph optimization-based approach, to identify a reading path with prerequisite chains for a given query. Our model not only considers the importance of each paper in a whole citation network but also takes into account the relationships among them.

We run extensive evaluations based on the SurveyBank dataset and compare our proposed approach for Reading Path Generation with existing academic search engines and baseline methods. The evaluation results demonstrate that our approach

outperforms other baselines and can generate insightful reading path for different research topics.

II. TASK DEFINITION

In this section, we compare the articles from Google Scholar and a corresponding survey’s reference list to motivate our work. We then formalize our proposed new task.

A. Background and Motivation

Given a research topic, people usually rely on academic search engines such as Google Scholar to retrieve a list of papers to read. Our key observation is that the paper list returned by search engines has dramatic difference with the reference papers in a survey for the same topic. This is explained by the fact that the returned ranked list from such search engines only considers relevance with the queried topic but not includes prerequisite chains of the queried topic. However, a reference list in a survey paper is comprehensive to cover more concepts including prerequisite chains or sub-topics of the queried subject.

Experimental Setup Figure 1 illustrates an example of comparing the searching results from Google Scholar and the reference list from a survey paper [16]. The query in Google Scholar is the key phrases extracted from a survey title. For searching setup, the time range is set from anytime to the year when this survey is published so that articles published after the survey will not be returned. We also remove the survey paper itself if found in the result of Google Scholar. To avoid contingency, we select a subset of SurveyBank (details in Section III) with high scores s for further statistics, where the score is calculated as follow:

$$s = \frac{citation}{2020 - year + 1}$$

where *citation* refers to the number of citations of a survey paper and *year* is when the survey is published.

Evaluation Because the importance of each paper in a survey’s reference list is not the same, as on the right side in Fig. 1, we create three labels for each survey based on how many times a referred paper is mentioned (or cited) in the survey: a full reference list (#occurrences ≥ 1), a reference list where each paper is cited at least twice (#occurrences ≥ 2) and a reference list where each paper is cited at least three times (#occurrences ≥ 3). For top 30 and top 50 searching results from Google Scholar, we respectively count the ratio of overlapping papers between scholar response of 0^{th} order, 1^{st} order, 2^{nd} order and the three reference lists.

Observations From Fig. 1, we can observe that only one paper shown in the top 5 result belongs to the survey’s reference paper list (i.e., Paper 1). From Fig. 2, we can also notice that the overlapping ratio between initial retrieved articles from Google Scholar and the survey’s reference paper list is pretty low no matter it is TOP 30 or TOP 50. However, if we further explore the papers cited by the top K result of Google Scholar (1^{st} order neighbors), more survey’s references will show up (i.e., Paper 5, 7, and 8). We can repeat this operation by further exploring the papers cited by the 1^{st} order neighbor papers

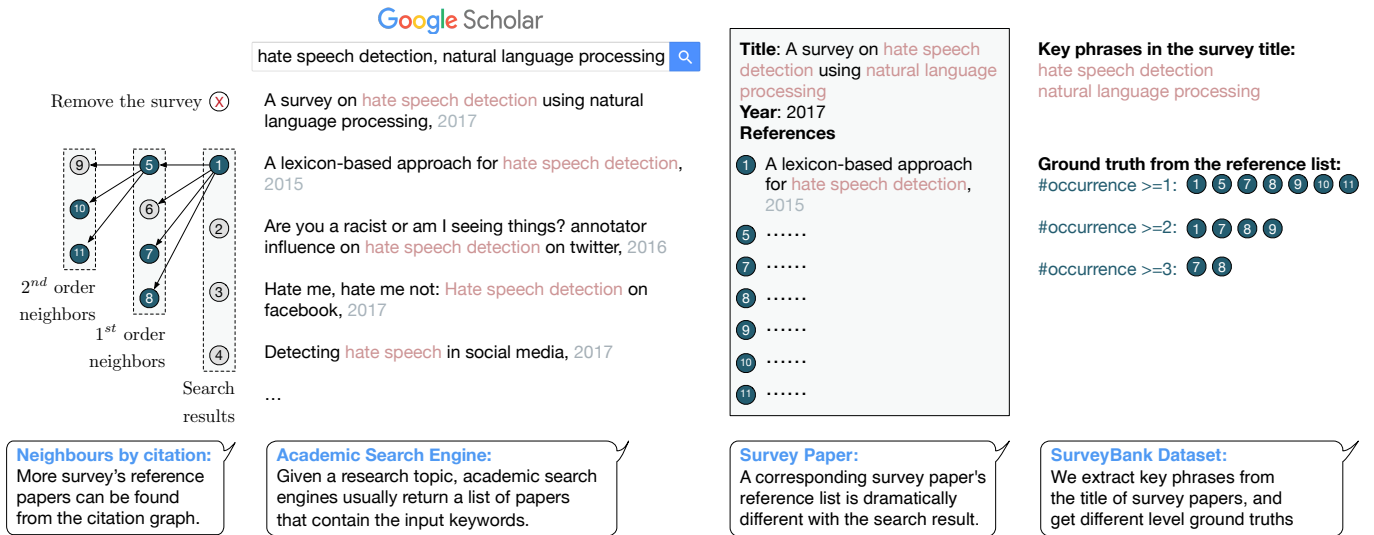


Fig. 1: A real-world example of the comparison between Google Scholar and a survey paper. In this figure, each circle represents a paper, and dark ones are overlapping papers between the searching result of Google Scholar and a survey’s reference list. An arrow between two papers indicates the reference relationship. For example, Paper 1 \rightarrow Paper 5 means Paper 1 cites Paper 5. The right side in this figure illustrates the evaluation benchmark for our Reading Path Generation task. The number of occurrences at the right side indicates how many times a reference paper was cited in a survey.

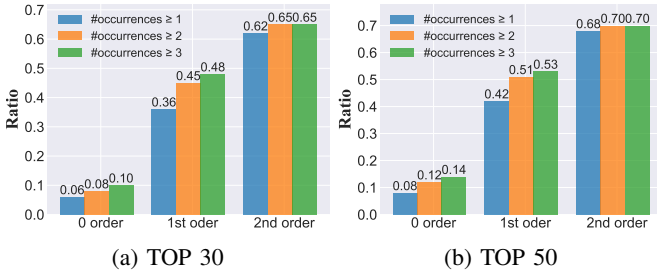


Fig. 2: TOP30/TOP50 refers to the number of initial seed papers returned from Google Scholar

(i.e., the 2nd order neighbors), then more papers belonging to the survey’s references would be found (Paper 10 and 11). It is also proved from Fig. 2 that if we further explore 1st order and 2nd order neighbors of initial TOP 30/50 seed papers, the overlapping ratio will greatly increase. To conclude, we can summarize these into two main observations below:

- **Observation I:** For the same topic, the articles directly returned from Google Scholar have a huge gap with the reference list of a survey.
- **Observation II:** Although there is a gap between the two, most of the missing papers can be retrieved through the neighbors (1st order and 2nd order citations) of the initial seed nodes (i.e., Google Scholar results).

Understandings In this subsection, we provide the understanding and explanation for aforementioned observations.

- **Understanding of Observation I:** The results from Google Scholar are only directly related to the topic of the query, while a survey is comprehensive to cover more concepts including prerequisite chains or sub-topics of the subject.

- **Understanding of Observation II:** In each scientific paper, pre-requisite work or sub-topics of the object can be mentioned beforehand in related work section for introductory purpose. It is reasonable that pre-requisite work or sub-topic papers can be found in the neighbors of initial seed papers.

The observations and understandings above motivate us to design a new task called Reading Path Generation, which not only returns papers highly relevant with the queried topic but also retrieve prerequisite papers to form a path or paths to help better understand the topic or those highly relevant papers. The observations also motivate the design of our graph-optimization-based approach to solve this problem in Section IV.

B. The Task of Reading Path Generation

Motivated by above observations, we propose the task of Reading Path Generation. Here we first introduce some necessary definitions, and then formally formulate our task.

Evaluation Benchmark. The evaluation benchmark for our task consists of a set of surveys $\mathcal{D} = \{d_1, d_2, \dots, d_{|\mathcal{D}|}\}$. As shown in Fig. 1, each survey d is accompanied by a set of key phrases $\mathcal{K} = \{k_1, k_2, \dots, k_{|\mathcal{K}|}\}$ extracted from its title, as well as a set of true labels \mathcal{V} inferred from its reference list. For each survey, since the importance of each paper in a survey’s reference list is not the same, we design multiple true labels in \mathcal{V} by counting how many times a reference paper was cited in a survey: $\mathcal{V} = \{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3\}$, where \mathcal{L}_i represents a list of papers where each was cited at least i times in the survey.

Reading Path Generation Task. Given input key phrases \mathcal{K} , our goal is to learn a mapping function $\mathcal{F}: \mathcal{K} \rightarrow \mathcal{O}$ to generate a reading path which covers the reference papers \mathcal{V}

as much as possible, where \mathcal{O} is a list of generated papers with reading order. The reading order in our task refers to the direct/indirect citation relationships between papers.

C. Task Evaluation

How to define a reasonable reading path for a list of papers is an open question, and it is not accurate to evaluate whether the specific reading path is correct or not with metric evaluation. Because for a given list of papers, there may be many optionally reasonable reading paths to be organized. In this paper, we present that the citation relationship is one of the intuitive and reasonable ways to generate the structured reading path for the given list of papers. Of course, it is encouraged to incorporate other useful information to generate a more ingenious reading path, we leave it for future work. Therefore, we divide the evaluation on the generated reading path into two steps:

- i) **Overlapping metric evaluation on the reading list;**
- ii) **Human evaluation on the generated reading path.**

Why not metric evaluation like MAP on the ranked list instead of overlapping metric evaluation on the reading list? It is not possible to define a reading path as a ranked list $\{P_1, P_2, \dots, P_n\}$ where the paper at the front of the list has higher importance to the queried topic. Because a sequence of papers in the reading path only indicate the reading order instead of the importance of each paper, and it is hard to differentiate the importance of each paper in the reading path as well. Considering these, we evaluate the flattened reading path with overlapping metrics.

Can the citation relationship between papers determine unique reading order? Yes, it can. It's hard to define the ground-truth reading order given a reading list. Different experts may have different opinions. We propose that the citation relationship combined with publish time is one of intuitive and reasonable ways to generate the structured reading path for the given list of papers. Once the reading list is determined, the reading direction between two papers can be easily and uniquely obtained from our constructed citation graph based on citation relationship and published time. For the case of multiple citation paths between two papers, we will assign all paths to them. It can be explained by the fact that there can be multiple sub-topics related with the queried topic.

III. THE SURVEYBANK DATASET

Here we introduce the *SurveyBank* dataset as the benchmark for the evaluation of RPG task. We first present the overview of the dataset and the construction process, and then highlight the statistical properties of SurveyBank. Moreover, we briefly introduce the building procedure of the citation graph.

A. Overview of SurveyBank

Our goal of building SurveyBank is to provide a clean and scalable benchmark for RPG. Some previous works such as S2ORC [15] have been focusing on building a large-scale linked scientific corpus. These datasets cannot be directly transformed to RPG task since they i) mainly focus on citation

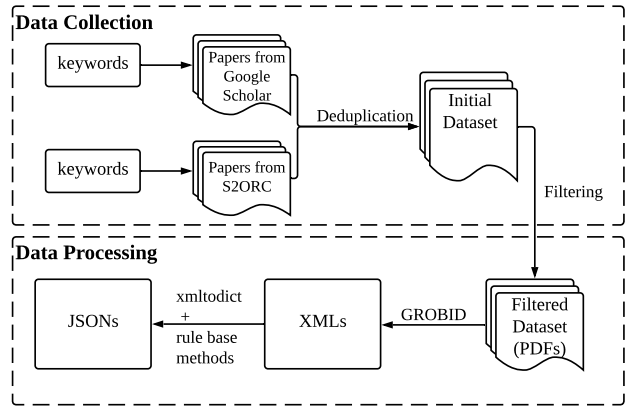


Fig. 3: Overview of the dataset construction process.

graph, and ii) suffer from low coverage of full text, disarray of paper types as well as topic domains. Other works [14] were focusing on tasks which are similar to RPG, but they require the assistance of domain experts which potentially lead to low scalability and additional human efforts. To mitigate the shortcomings of existing systems mentioned above and to better serve the task of RPG, we propose SurveyBank, a dataset constructed from a collection of survey papers without human labeling, as our benchmark. Each paper in it has its corresponding parsed full context containing hierarchical section information, key phrases extracted from the title and ground truth labels from the references.

Our dataset also contains the input key phrases and ground truth labels to better serve as a benchmark for RPG. As we have mentioned in Section II, RPG takes one or more keywords/key phrases as its input. Since only a few surveys contain keywords as an inherent part of their papers, we choose to infer keywords/key phrases from the titles since titles can indicate the topics of the paper in many cases [12]. For ground truth labels (the list of papers for generating desired reading path), we start with calculating the number of occurrences for each paper cited in the survey. The intuition behind this is simple: the more a paper is cited within the survey, the more likely this paper should be considered as significant. As described in Section II, we use generated lists to denote different significance levels for cited papers in the survey.

B. Dataset Construction

The overall process for constructing SurveyBank is shown in Fig. 3. Initially, we collect survey papers based on keywords that indicate whether the paper is a survey or not, as well as topics to restrict the domain of the surveys collected. Our papers come from two sources: Google Scholar and S2ORC [15]. To make sure that the collected surveys are in the domain of computer science, we aggregate topic concept phrases from LectureBank [17] and TutorialBank [18]. LectureBank provides 208 topic keywords across 5 different sub-domains in computer science, including Natural Language Processing (NLP), Machine Learning (ML), Artificial Intelligence (AI),

Deep Learning (DL), and Information Retrieval (IR), verified by experts. On the other hand, extra 306 topic keywords are collected from TutorialBank. These topic keywords are collected according to several criteria and also performed by experts, where the goal is to make the topic keywords more conceivable. We mix these topic keywords together and apply deduplication process to them, leaving us 441 unique topic keywords. We utilize them together with additional keywords such as “survey” as inputs to collect survey papers from Google Scholar. For S2ORC, we retrieve the papers containing survey-indicating keywords such as “survey” in their titles from the computer science subset. Finally, we obtain an initial collection with 41,194 papers in total. We further check paper titles in order to make sure there is no duplication among them and discard samples without proper full text.

We process the surveys’ full text in PDF format following a similar procedure as S2ORC [15]. Specifically, GROBID [19] is utilized to extract metadata, body text, and bibliography entries from PDFs into the XML format. In addition, we extract the keywords/key phrases from the title using the TopicRank algorithm [20] implemented by pke [21] and provide fine-grained structure for body text for each survey. The atomic units here are paragraphs instead of sections; for each section or subsection with a proper headline, we assign a notation to it to denote its hierarchical position within the survey. We hope that other researchers can make further progress in RPG task by leveraging the paper structure information. A survey is excluded from SurveyBank if it meets one of the following criteria: i) the corresponding PDF full text cannot be processed by the Python library PyPDF2 [22]; and ii) the survey is more than 100 pages or less than 2 pages (A paper with more than 100 pages is more likely to be a thesis or reports). These parsed XMLs are then transformed into the JSON format; we apply a rule-based methods to the data processing procedure in addition to the Python library xmldict [23] in order to mitigate certain inherent errors caused by GROBID and xmldict. Our SurveyBank dataset is publicly available and can be found [here](#).

C. Statistical Property

We calculate the statistics of SurveyBank, as shown in Fig. 4. Overall, SurveyBank contains 9,321 survey papers with high quality filtered from an initial collection with 41,194 survey papers. Each paper has approximately 58 references on average. One criterion for evaluating the impact of academic papers is how many times it has been cited overtimes. Intuitively, the more times one survey paper is cited, the more significant role the paper may be playing in its research domain throughout years of development. To reflect this criterion, we also calculate the citation numbers for each survey paper within SurveyBank. Overall, 17.8% of surveys have no opportunity of being cited by other papers, while 5.3% of the surveys were cited for more than 500 times over the years. A detailed reference quantity distribution can be found in Fig. 4a. In the case of publish date, more than 87.8% of survey papers in SurveyBank are published within the recent

TABLE I: Topic distribution of the survey papers in SurveyBank.

Domain	#Papers
Artificial Intelligence	1,151 (12.3%)
Interdisciplinary, Emerging Subjects	440 (4.7%)
Computer Network	424 (4.5%)
Computer Graphics and Multimedia	280 (3.0%)
Database, Data Mining, Information Retrieval	270 (2.9%)
Software Engineering, System Software, Programming Language	205 (2.2%)
Computer Architecture, Parallel and Distributed Computing, Storage System	196 (2.1%)
Network and Information Security	162 (1.7%)
Computer Science Theory	123 (1.3%)
Human-Computer Interaction and Pervasive Computing	86 (0.9%)
Uncertain Topics	5,984 (64.2%)
Total	9,321

20 years. The distribution of publication years and citation numbers for SurveyBank are shown in Fig. 4b. Eventually, we show the distribution of the number of papers cited by the survey papers in SurveyBank in Fig. 4c.

Finally, to explore what kind of papers we have collected, we look into the domain distribution of papers in SurveyBank by checking where they were published according to a venue collection released by CCF [24]. The distribution is listed in Table I. Here the category *Uncertain Topics* contains the papers that either do not contain a valid publication venue inside its full text or their publication venues are not included in the venue collection provided by CCF.

D. Citation Graph

We initially collect around 80 million scientific papers covering all kinds of subjects such as biology, chemistry and computer science from S2ORC. Then we pick all papers in the domain of computer science by the provided attribute “domain” to get around 6 million computer science papers. These papers would be linked by citation relationship and publish time to form citation graph. We also visualize a part of connected citation graph consisting of 10,000 papers randomly sampled from our whole citation graph as shown in Fig. 5.

IV. OUR APPROACH OF REPAGER SYSTEM

In this section, we first introduce the overall architecture of our RePaGer system and then present a detailed description of the Node-Edge Weighted Stiner Tree (NEWST) model for path selection over a citation graph.

A. The Architecture of RePaGer System

Our RePaGer system for Reading Path Generation takes key phrases as inputs and produces a reading path for users to read. Fig. 6 shows the overall architecture of RePaGer. There are five main steps, which would be detailed in the following:

- 1) **Initial seed nodes obtainment.** Our RePaGer system first leverages SerAPI [25] tool to get top K articles from Google Scholar by searching a user’s query.
- 2) **Weighted citation graph construction.** We construct a large-scale citation graph with 6 million papers from

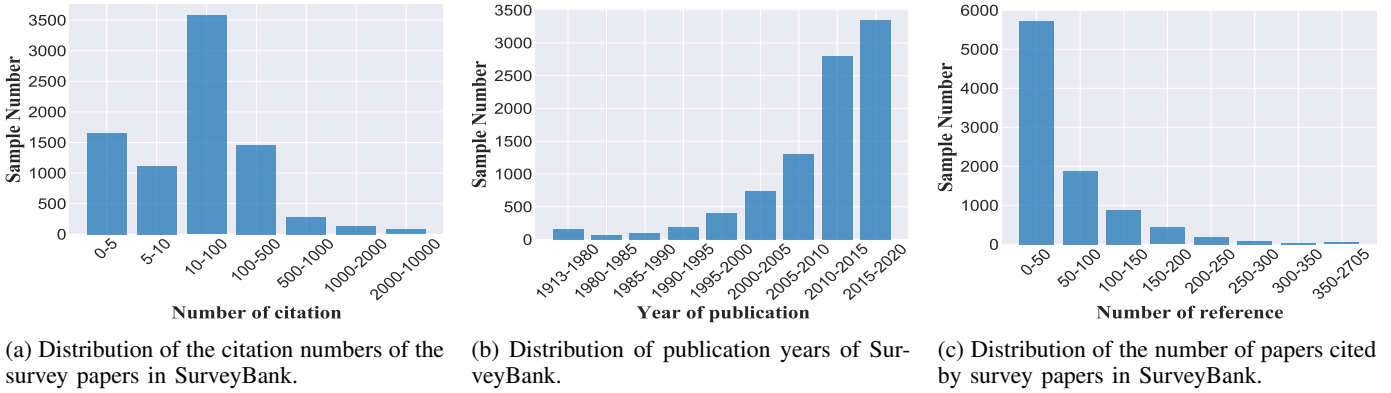


Fig. 4: Statistics of SurveyBank.

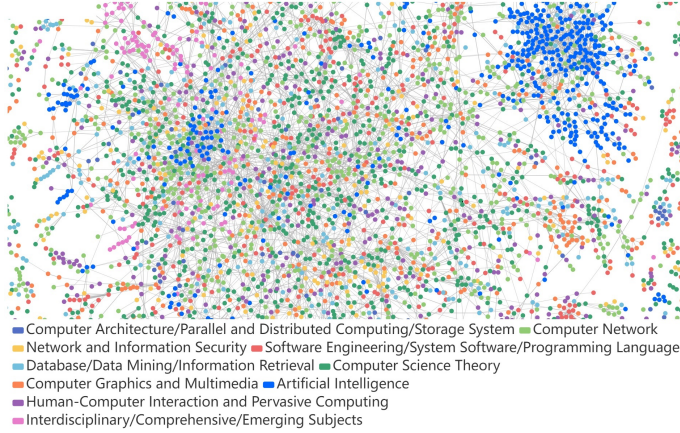


Fig. 5: A citation graph which consists 10,000 papers sampled from our whole citation graph. Different colors indicate different subjects in the domain of computer science.

S2ORC [15] in the field of computer science by citation relationship. We further apply our proposed methods on original citation graph to calculate the weight of each paper. The weights of nodes and edges will be introduced in Section IV-B.

- Sub-citation graph construction.** After acquiring the initial seed papers from Google Scholar and constructing the weighted citation graph, we capture the first-order and second-order neighbors of initial seeds as candidates to form sub-citation graph with weighted edges and nodes based on the observation that most of the papers relevant to the query can be found among the first-order and second-order of initial seeds as shown in Fig. 2.
- Seed nodes reallocation.** Query’s prerequisite chains hidden in citation network cannot be found in initial retrieved results from Google Scholar. In order to capture more preconditions of the topic and generate a logical reading path with prerequisite chains, we reallocate the initial seed papers by selecting articles with high co-occurrence as new seeds in weighted citation graph. For example, as in Fig. 6, paper 13 co-occurs in the reference list of paper 9 and paper 10, and then paper 13 would be considered as a new seed node because of high co-

occurrence. We infer that these new selected papers are more likely to be prerequisite articles for understanding a topic because multiple articles directly relevant to this topic mention them in their papers.

- Reading path generation by NEWST.** The NEWST is a graph optimization-based approach to address the Reading Path Generation problem. Given a weighted citation graph and a subset of vertices in this graph as seed papers (the result of step 4), the NEWST spans through the weighted citation graph, and finds an optimal tree interconnecting all papers in this subset at the minimum cost. Our algorithm not only considers how important each individual paper is in the whole graph, but also takes into account the edge weights between papers. We will describe more details about the NEWST in the following section.

B. The NEWST MODEL

Scientific paper citation networks usually consist of massive papers, and the prerequisite concepts of a given query are hidden in them. In this regard, we propose Node-Edge Weighted Steiner Tree (NEWST) model to identify a reading path for a given query. Note that our proposed system and algorithm is not only suitable for Reading Path Generation in the filed of computer science or other domains, but also easy to transfer to solve other weighted graph related problems to discover useful information hidden in the graph. We also give a heuristic solution for NEWST. The classical node-weighted Steiner tree problem [26]–[28] is defined as: given a graph $G = (V, E)$ with node weight $w : V \rightarrow R^+$ and a subset $S \subseteq V$, find a spanning tree which contains all the nodes in S with minimum overall cost.

Let $G = (V, E, S, w, c)$ be a connected and undirected network / graph, where V is the set of vertices, E is the set of edges, w is a function which maps each vertex in V to a positive vertex weight, and c is a function which maps each edge from E to a positive edge cost. Let $S \subseteq V$ be a subset of V called compulsory terminals in classical node-weighted Steiner tree problem. The compulsory terminals in our model are reallocated seed papers we retrieve from Step 3 Seed Papers Reallocation. The purpose of this task is to find

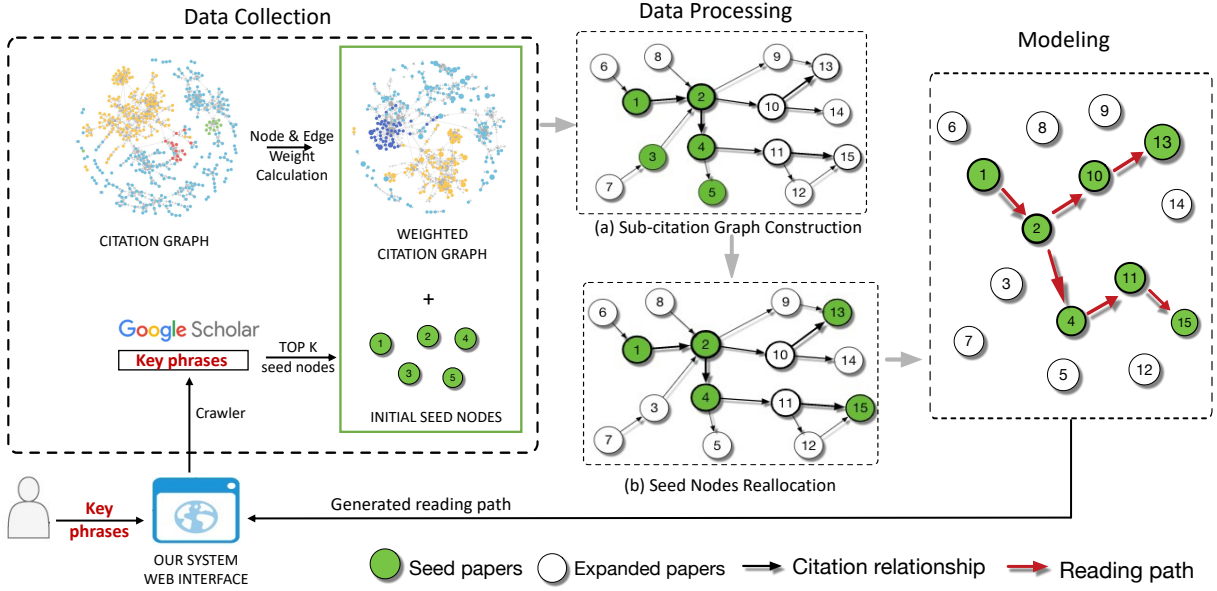


Fig. 6: Overall framework of our RePaGer system

an optimal Steiner tree T that spans S with minimal total cost on its edges and nodes, and T is defined below:

$$T = (V_T, E_T), S \subseteq V_T \subseteq V, E_T \subseteq E$$

which minimizes the objective function $cost$ of our whole model

$$\text{minimize } cost(T) = \sum_{e \in E_T} c(e) + \sum_{v \in V_T} w(v), \quad (1)$$

where e is an edge and v is a vertex in the optimal Steiner tree T .

The edge cost function c in Eq. (1) is defined as:

$$c(i, j) = \frac{\alpha}{con(i, j)^\beta}, \quad (2)$$

where i and j are indexes of two different papers, α and β are positive constant values, and $con(i, j)$ is a score that measures the relevance between paper i and paper j . In our model, con is calculated by how many times paper j cited in paper i or inversely.

The vertex cost function w in Eq. (1) is defined as:

$$w(i) = \frac{\gamma}{a * pgscore(i) + b * venue(i)}, \quad (3)$$

where i is a paper's index, γ is a positive constant value, $pgscore(i)$ denotes the PageRank [29] score of paper i in our scientific citation network, and $venue(i)$ is the paper i 's venue score that we calculate according to venue rankings from AMiner and CCF. This comprehensive venue rankings span 10 topics with a total of around 700 top journals and conferences in the field of computer science. This collection of journals and conferences are manually divided into three levels of rankings by experts in this area in CFF while AMiner also gives an influence score automatically for each of them based on the citations of the best papers in each journal or conference. We calculate the average of the two as a final

Algorithm 1 Heuristic Solution for NEWST.

Input: A connected undirected node-edge weighted graph $G = (V, E, S, w, c)$, with compulsory terminals $S \subseteq V$, vertex cost function w and edge cost function c ;

Output: An optimal Steiner tree $T = (V_T, E_T)$ containing all nodes in S , $S \subseteq V_T \subseteq V, E_T \subseteq E$;

- 1: Construct the complete undirected distance graph $G_1 = (V_1, E_1, S, w, c)$ from G ;
- 2: Find the minimum spanning tree, T_1 , of G_1 . (If there are several minimal spanning trees, pick an arbitrary one.);
- 3: Construct the subgraph, G_s , of G by replacing each edge in T_1 by its corresponding shortest path in G . (If there are several minimal spanning trees, pick an arbitrary one.);
- 4: Find the minimal spanning tree, T , of G_s . (If there are several shortest paths, pick an arbitrary one.) Then T is our optimal Steiner tree containing all nodes in S ;

venue score of each paper in our model; a and b are positive constant values. Note that we can revise the definition of the cost function of edges and weight function of nodes to incorporate more valuable information for generating a better reading path. For example, instead of only exploit the title and citation relationship in our model, we can further utilize the semantic information of the main text. We leave these to our future work.

Heuristic Solution for NEWST. Since the classical Steiner tree problem in graphs is NP-hard, and our *NEWST* is a more general version of the Steiner tree, therefore the *NEWST* problem is also NP-hard, which means that there may not be an algorithm that can solve the problem with a large number of instances in polynomial time. One common solution to deal with large network inefficiency is a heuristic algorithm, which finds sub-optimal solutions in large networks in a short time. In *NEWST*, we leverage a heuristic algorithm [30] to solve this

problem and briefly summarize it in Algorithm 1. The general idea is that the minimum of NEWST can be approximated by computing the minimum spanning tree of the subgraph of the metric closure of G induced by compulsory nodes S , where the metric closure of G is the complete distance graph in which each edge is weighted by the shortest path distance between the nodes in G . A shortest path from paper P_i to P_j is a path from P_i to P_j whose distance, including node costs and edge weights, is minimal among all the possible paths from P_i to P_j . The minimum spanning tree of a graph is a subset in this graph where the total distance on its edges and nodes is minimal among all spanning trees. As described in Algorithm 1, the worst case time complexity of Step 1-4 is $O(|S||V|^2)$, $O(|S|^2)$, $O(|V|)$ and $O(|V|^2)$, respectively. Hence, our heuristic algorithm has a worst case time complexity of $O(|S||V|^2)$ and guarantees to output a tree that spans S with total cost on its nodes and edges no more than $2(1 - \frac{1}{l})$ times of the global optimal value, where l is the number of leaves in the optimal Steiner tree T .

V. USER INTERFACE OF REPAGER SYSTEM

The user interface of **Reading Path Generation (RePaGer)** system is a web application with input panel same as normal academic search engines like Google Scholar. An user can input an arbitrary string for an interested research topic, which can be a single key phrase or even a combination of multiple key phrases. As shown in Fig. 7, after an user inputs a query, the web will present five components consisting of input panel (a), navigation bar (b), generated reading path panel (c), detailed information of the clicked paper (d) and node and edge weight annotation (d), which would be detailed in the following.

- **Input panel.** This allows users for the subsequent query for a new research topic without coming back main UI interface.
- **Navigation bar.** Users can leverage the navigation bar to browse papers in the generated path in a flattened way. Each paper is provided with meta information including a title, authors and published year.
- **Generated reading path panel.** In the generated path, the direction between two papers indicates a reading order. The color of each paper and edge in the path denotes the importance of this paper in the whole path and the relevance between two papers respectively.
- **Detailed information of the clicked paper.** When a paper either in b or c is clicked, more detailed information about this paper would show up in this component.
- **Node and edge weight annotation.** Side weight legend annotates the importance between two papers, which means how important a cited paper to its citing paper. From left to right of the legend indicates the importance from low to high. For the node weight legend, it annotates the importance of the paper in a whole reading path.

The RePaGer system also provides two ways for users to interact with the generated reading path. One way is to start with navigation bar to browse papers and then click on

interested papers. After that, the color of the clicked paper in component c would become red, and users can easily get relevant information of this paper in a whole reading path. As shown in Fig. 7, the paper of “Attention is all you need” in component b is clicked and then the corresponding paper in component c gets red. Another way is to directly click one interested paper in component c, and then corresponding detailed information of the clicked paper should be presented in component d.

VI. EXPERIMENTS

A. Experiment Setup

Experimental Details We evaluate the NEWST model on the proposed SurveyBank Dataset. The key phrases extracted from the title of each survey paper are taken as the model input and the true label is reference papers of the corresponding survey. It is worth noting that we remove the corresponding survey paper of input key phrases to avoid data leakage if it appears in the response from Google Scholar. Since the limitation of the publication time for each survey, the ground truth papers cannot cover the newest literature. Therefore, all papers considered during the metric computation are published before the time when the survey was issued. Following the setting of [28], the parameters $\{\alpha, \beta, \gamma, a, b\}$ used in the NEWST algorithm are set as $\{3, 2, 5, 0.7, 0.3\}$. Moreover, we set the number of initial seed papers from Google Scholar to be 30 for all experiments.

Baselines We compare our method with the following models:

- **Google Scholar:** A paper retrieval engine developed by Google which can return a list of papers according to the input phrases. For each query phrase, we select the top- K papers in the retrieval results as the generated results.
- **Microsoft Academic:** A paper retrieval engine developed by Microsoft. Correspondingly, the top- K retrieved papers of the query phrase are chosen as the final results.
- **AMiner:** Another academic search engine released by Tsinghua University. Similarly, we take the top- K retrieval results in our experiments.
- **PageRank [29]:** A well-known Web page ranking algorithm which orders pages based on their citation relationship. Similar to the NEWST, we first expand initial seed nodes returned from Google Scholar to their neighbors as candidates, and then the PageRank algorithm is applied to reorder initial seeds and expanded candidates together.
- **SciBERT [31]:** An awesome variant of the BERT model [32] trained on scientific literature from the corpus of Semantic Scholar [33]. We train a matching model using SciBERT to score the matching degree of queries with paper titles and abstracts. During the inference phase, we also expand the seed nodes returned from Google Scholar and then re-rank them via our trained matching model.

Task Evaluation As described in Section II, we combine metric evaluation and human evaluation to estimate the quality of generated reading paths. On one hand, we first downgrade the evaluation of reading paths to the evaluation of

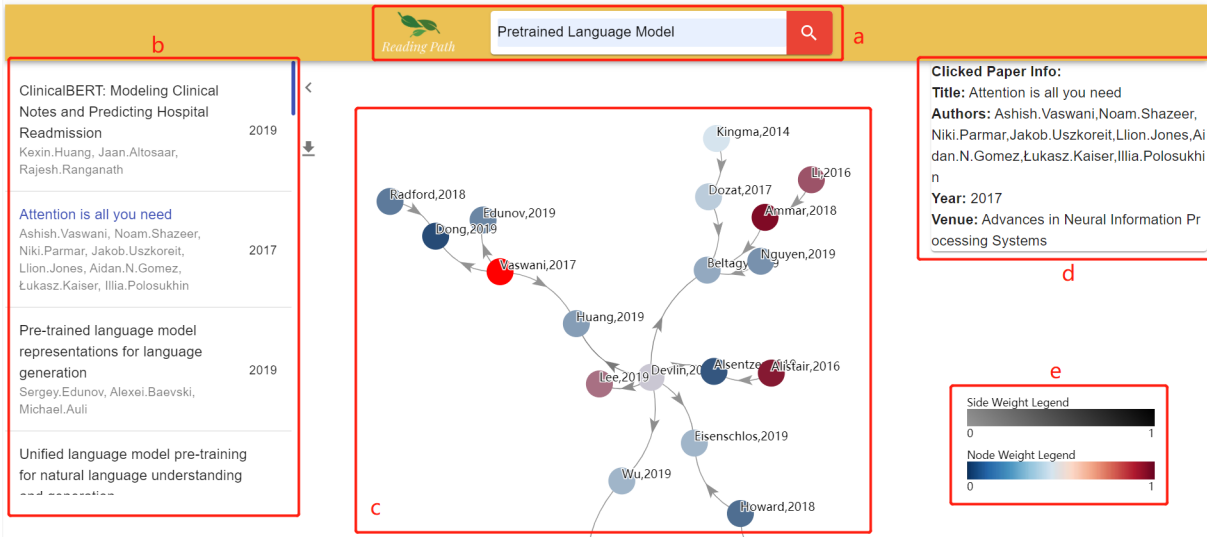


Fig. 7: User interface of RePaGer system. Panel (a), (b), (c), (d) and (e) indicate “input panel”, “navigation bar”, “generated reading path panel”, “detailed information of the clicked paper” and “node and edge weight annotation” respectively. The content shown in the figure is a reading path generated for the query of “Pretrained Language Model”.

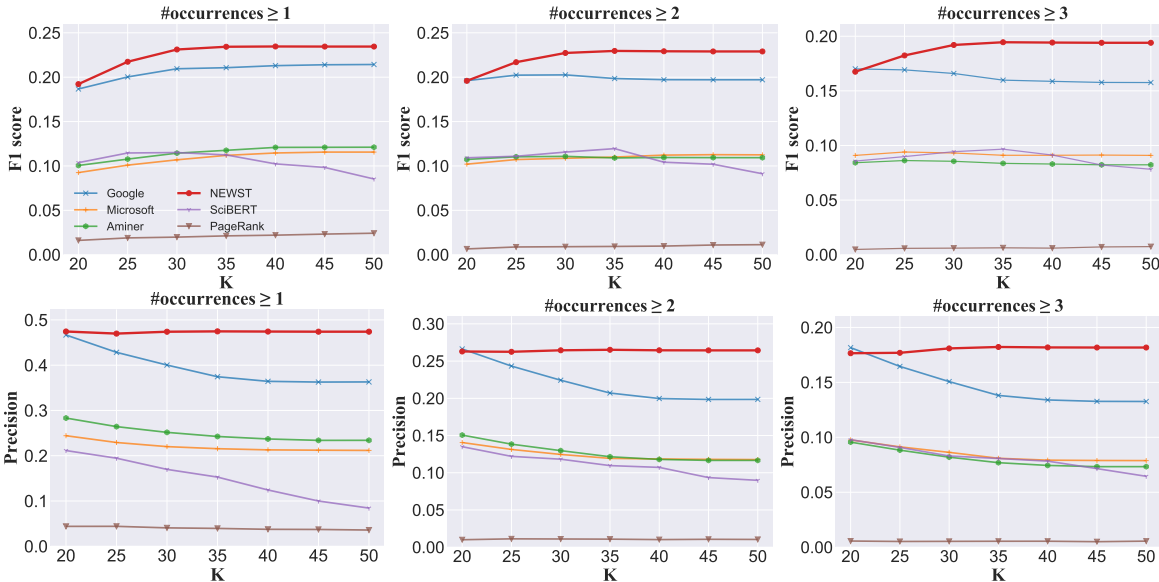


Fig. 8: The performance comparison of NEWST, Google Scholar, Microsoft Academic, AMiner, PageRank and SciBERT on F1 score and precision metrics for top- K recommended papers.

reading lists. The reason is that **once the reading list is determined, the reading direction between two papers can be easily and uniquely obtained from our constructed citation graph based on citation relationship and published time**. Therefore, the quality of Reading Path Generation solely depends on the generated reading lists. As our model returns a list of candidate papers for each input query phrase, we adopt $P@K$ (precision) and $F1@K$ (F1 score) to evaluate the performance of our model, where K denotes the number of generated papers. The results averaged over all query phrases are reported as the final performance. On the other hand, we further perform human evaluation on the generated reading paths to further evaluate their quality.

B. Overlapping Metric Evaluation

Overall Performance To test the performance of our model, we compare it with several outstanding methods on the task of generating a list of top- K papers, where K ranges from 20 to 50 as each survey at least cites 20 papers. The results are shown in Fig. 8 in which we aim to generate the ground truth papers that appear in the survey paper with different times. From these experimental results, it can be seen that for almost all evaluation scenarios, our proposed NEWST model outperforms all other methods, especially when the number of generated papers becomes large. In particular, when the number is large than 25, NEWST consistently outperforms the academic search engine (i.e., Google Scholar, Microsoft

Academic and AMiner) by a substantial margin. This is caused by the fact that existing search engines solely return the paper whose title contains query phrases.

It is worth noting that by exploring the citation relationship among literature, NEWST performs significantly better, compared with SciBERT based matching model, which plainly employs semantic analysis, highlighting the importance of using the Steiner tree algorithm to discover high-order reference chain. Although introducing citation information may attribute to the quality of generating reading list, solely considering this type of information is inclined to deteriorate the performance. This phenomenon can be clearly observed in Fig. 8, in which the PageRank method achieves the worst results, as it always returns the papers whose citation number is the largest, which is fallacious since in the real world scenarios, the most relevant papers may not be the most cited papers. The proposed NEWST model relieves this issue by introducing the node-edge weighted Steiner tree algorithm and, narrows down the candidate papers via seed nodes selection and attains significant performance gains. Another interesting observation is that NEWST preserves superior performance no matter K is large or small, especially in terms of precision. For large K , the performance of all baseline methods generally tends to drop significantly, while our model achieves even improved precision, indicating that our method is more robust and more suitable for learners who like to acquaint more relevant knowledge.

Comparison of Different Seed Paper Numbers To investigate the impacts of the seed nodes selection, we conduct a parameter sensitive study with different numbers of seed nodes and the performance results are demonstrated in Table II. It can be observed that the performance of our model is still robust even though the node number varies in a large interval. Specifically, increasing the seed node number, the performance rises steadily. The reason for this expected phenomenon lies in the fact that the more seed nodes to be used, the more ground truth papers can be included after expansion, highlighting the importance of considering multi-order citation relationship among the papers returned from Google Scholar, which solely takes into account of the keywords in user queries. However, overloading the quantity decreases the performance of precision, which might be caused by introducing ineffectual noise paper when using too many seed nodes.

Effect of Seed Papers Reallocation We experiment with four variants for our model to investigate the effect of seed papers reallocation: i) NEWST: with high co-occurrence papers as compulsory nodes for the NEWST; ii) NEWST-W: with initial top-30 seed papers as compulsory nodes; iii) NEWST-U: with the union of above two; iv) NEWST-I: with the intersection of above two. Table III presents the results of all variant methods. Comparing NEWST with NEWST-W, it can be observed that the injection of seed papers reallocation has improved the robustness of generating reading list, resulting in better performance. Moreover, NEWST-I, which uses the intersection of initial seed papers and co-occurrence papers as compulsory nodes, achieves comparable results with

TABLE II: Impact of the number of seed nodes on the performance of NEWST model.

#seed nodes	10	15	20	25	30	40	50
F1 score	0.1924	0.2094	0.2197	0.2251	0.2343	0.2385	0.2387
Precision	0.4279	0.4454	0.4565	0.4668	0.4743	0.4755	0.4658

NEWST, further confirming the superiority of considering co-occurrence times during reallocating the compulsory nodes for the NEWST algorithm. However, by using the union of them, NEWST-U achieves even better results in F1 score but worse in precision. This may be caused by that the union of papers increases the number of compulsory nodes, which results in including more ground truth papers such that rising up the recall rate, but also introducing noise to reduce the precision.

Impact of the Node-Edge Weight To further evaluate the impacts of node and edge weights, we propose the following variant methods: i) NEWST-C: method that uses reallocated papers as final results (i.e., ignoring step 4); ii) NEWST-N: excluding node weights in Algorithm 1 and iii) NEWST-E: excluding edge weights. As shown in Table III, by introducing node and edge weights, NEWST-N, NEWST-E, and NEWST consistently outperform NEWS-C in terms of F1 score, which means that the proposed Steiner tree algorithm can effectively involve the ground truth papers in the generation of reading path. Compared with NEWST-C, although the exploitation of NEWST algorithm may damage precision, our model still achieves comparable performance. More importantly, due to the inability of path generation, NEWST-C cannot take account of the prerequisite relationship. However, the proposed NEWST model not only tells users “what to read”, but also “how to read”.

C. Human Evaluation

To further examine the quality of our generated reading path, we conduct a human evaluation of Google Scholar and our RePaGer System. We conducted a human evaluation to examine the quality of the generative reading path. The following is a detailed description of the evaluation procedure.

Datasets We used the proposed SurveyBank dataset in human evaluation. Specifically, we randomly selected 20 query samples from the domain of Artificial Intelligence and Information Retrieval, respectively. As shown in Fig. 7, for each query sample, the reading path and reading list retrieval results were generated by RePaGer system and Google Scholar, respectively. These retrieval results and query sentences were used as evaluation materials to compare the usability of the two systems.

Participants We recruited 16 participants to evaluate our system. All of them are graduate students of computer science, mainly majoring in artificial intelligence (AI) or information retrieval (IR). According to their expertise in these two fields, we divided them equally into two groups, and then assigned them to the corresponding evaluation tasks (i.e., 8 persons for 20 queries in AI, 8 persons for 20 queries in IR).

Procedure At the beginning, we briefly introduced the SurveyBank dataset and RePaGer system. Then we take the

TABLE III: The performance of variant NEWST models validated in ablation study.

Methods	F1 score	Precision	Methods	F1 score	Precision
NEWST	0.2343	0.4743	NEWST	0.2343	0.4743
NEWST-W	0.2269	0.4533	NEWST-C	0.2064	0.5080
NEWST-I	0.2345	0.4738	NEWST-N	0.2195	0.4978
NEWST-U	0.2462	0.4188	NEWST-E	0.2203	0.4710

TABLE IV: Running time of our model under different retrieval cases.

	#nodes	#edges	Time (seconds)
Case 1	851	1313	9.28
Case 2	1276	2208	27.99
Case 3	1495	2235	35.31
Avg. (test set)	1754	2938	60.42

query phrase “pretrained language model” as an example, to demonstrate how to conduct evaluation. The retrieval results of RePaGer are shown in Fig. 7, respectively. Each participant were asked to explore the recommended papers and think whether these papers are conducive to help them to quickly grasp the subject matters of this field. After careful consideration, the participants filled out a self-evaluation questionnaire and began evaluating the selected 20 query samples.

Self-evaluation questionnaire In the questionnaire, we asked the participants to self-rate their preference of the systems from three folds: i) prerequisite: whether the retrieval results contain prerequisite relationship, i.e., the system should not only tells the reader “what to read”, but also “how to read”; ii) relevance: whether the retrieval results are consistent with the query phrase; iii) completeness: whether the retrieval results contain comprehensive knowledge of the query domain. Further, the participants are asked to rank them in three folds: i) prefer the PaReGer; ii) prefer the Google Scholar; or iii) prefer the two systems equally. The ratings averaged over all participants are reported as final performance.

Evaluation results The results of human evaluation are shown in Table V. Thanks to the incorporation of citation relationship, the generated reading path of NEWST contains more *prerequisite* chains, indicating that our model can generate a more reasonable structured reading list. The superiority of our model can be further observed by that the *consistency* of the two models is basically the same. However, the comparison on the *completeness* shows that NEWST can provide readers more comprehensive knowledge, which further confirms that our model can help users to comprehend this topic profoundly.

D. Empirical Study of Computational Efficiency

To show the computation efficiency of our proposed method, we also report the running time among several retrieval cases on a machine with 2.40GHz Intel(R) Xeon(R) 14 CPU cores and 132 GB memory. The results are demonstrated in Table IV, where #nodes and #edges represent the number of nodes and edges in the citation graph constructed in Step 2 of our method, respectively. It can be seen that because of the extra work to generate the reading path from the citation graph, the

TABLE V: Human evaluation on the domain of artificial intelligence (AI) and data mining (DM). A and B represent the system of Google Scholar and NEWST, respectively.

Domain	Criterion	Prefer A (%)	Same (%)	Prefer B (%)
AI	Prerequisite	0	7.90	93.10
	Relevance	37.93	37.93	24.14
	Completeness	20.69	41.38	37.93
DM	Prerequisite	0	23.33	76.67
	Relevance	26.67	43.33	30.00
	Completeness	20.00	36.67	43.33

proposed model is a little bit more time-consuming than the general research engines, but delivers the final results around 1 minute on average. For the readers looking for literature on their own, our model effectively tailors the learning order for them, greatly saving their learning time.

E. Reading Path Visualization

Figure 9 illustrates the reading path of query “Pretrained Language Model”, where each arrow indicates a reading order between two papers. Due to the space limitation, the reading path demonstrated in Fig. 9 is a subtree of the result returned from our RePaGer system. As we can see, by introducing the prerequisite structure into reading path, learners can easily grasp the technological evaluation of a new field and learn the main concept quickly. More importantly, our system not only teaches users “how to read”, but also tells them “how to understand”. This desirable virtue can be observed in the 4-th, 6-th and 12-th nodes of the reading path. These prerequisite concepts of Attention, BERT and Contextualized Word Embedding are less directly relevant to the topic of “Pretrained Language Model” but essential to help understand technology evolution for the topic profoundly. However, these pre-requisite concepts can not be found in the retrieved TOP list from Google Scholar.

VII. RELATED WORKS

With the explosion of the quantity of scientific papers, generating a reading list from a large amount of literature according to the given query phrases has been shown to be essential. Existing research can be roughly divided into two categories: i) reading list generation that ignore the prerequisite relationship among the generated literature, i.e., reading order; ii) generating structured reading path which can effectively indicate the readers to learn step by step.

Among the reading order irrelevant methods, the mainstream methods mainly focus on the task of retrieving the most relevant documents for a given query [34]–[36]. Specifically, [37] introduced the problem of pedagogical paper recommendation which takes account of learner’s interest, knowledge, goals, etc, and resolves it by leveraging collaborative filtering techniques [38]. To generate introductory reading lists, [13] explored several collaborative and content-based filtering methods for rating the importance of a paper within the literature. [39] proposed another method combined with the gradient decent algorithm to construct customized authority

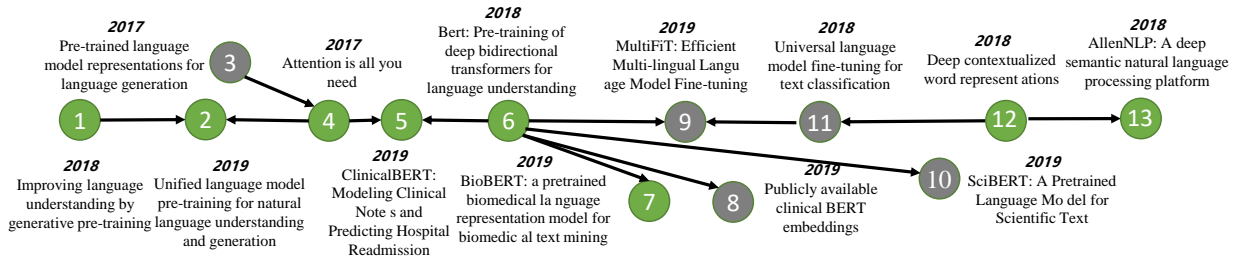


Fig. 9: The reading path generated by our NEWST model in terms of query phrase “Pretrained Language Model”, where the reading order is determined by arrows. Each green circle denotes a paper not showing up in the TOP 30 papers from Google Scholar while the gray one denotes the paper appearing in the TOP 30 papers from Google Scholar.

lists. Recently, [12] proposed to produce reading lists in an unstructured manner in which several relevant documents of the given query phrases are first chosen by the LDA topic model [40], and then a ThemedPageRank approach is further developed to re-rank these topic-related papers. However, these methods neglect the prerequisite chain [17] of the literature, downgrading the problem to item-based recommendation, which abominably hinders the application of the models.

In the pedagogical scenario, learners are likely to spotlight the suitable literature and prioritize them based on the prerequisite relationship. To this end, the reading order relevant methods aim to incorporate prerequisite chain more than solely generating reading lists. Regretfully, due to the frustration of extracting prerequisite relation among documents, there are few efforts invested in this field. Recently, [14] proposed to address this issue by adopting concept graph [18], [41], which is used to model the concept dependency among knowledge domains and relative documents. However, the concept dependency mining algorithm used in the construction of concept graph is somewhat crude and inaccurate. Besides, the concept graph is not able to cover different domains. Beyond the document recommendation, some data inherently appears with the prerequisite relationships, like lecture data created in the MOOCs [42], [43], as different lectures credited by the same student are time-variant. This notorious problem can be automatically sidestepped by leveraging citation relationship.

The problem coming with the generation of reading lists is how to evaluate its quality. In [14], 33 experts were recruited to take part in the evaluation. It is no doubt that such a direct method is impractical, time and cost consuming. The reference lists of papers to some extent reflect the professional knowledge in their scientific field. Therefore, some prior works [13], [44], [45] have been proposed to evaluate the coverage rate of the reference list given a query paper. Inspired by that, we publish a public SurveyBank dataset to employ an offline analysis. Different from the aforementioned evaluation methods, in this paper, we opportunely exploit the key phrases extracted from the survey title as evaluation query, and its reference list as the ground truth expected to be generated, which is more reasonable and intuitive.

It is worth to note that the proposed reading path generation problem is somewhat relevant to the exploratory data analysis (EDA) problem [46]–[48] and the composite item formation

and exploration (CIFE) problem [49], [50]. EDA aims at building an interactive system for users to allow easy-to-use data exploration and visualization. CIFE studies a problem of effective construction and exploration of large sets of packages associated with a central item. At first glance, the targets of EDA and CIFE are to some extent similar to RPG, which focuses on generating a paper reading path for someone who lacks preliminary knowledge of a queried research field. However, the methods used in EDA and CIFE cannot be directly applied to RPG. More specifically, recent EDA models generally exploit reinforcement learning to automatically generate the data analysis process [46]–[48]. These methods cannot be applied in RPG since RPG is not an interactive process. Furthermore, to apply CIFE in the context of RPG, we need to build up a huge structure of concepts in the domain, which would be another totally different research topic.

VIII. CONCLUSION AND FUTURE WORK

In this paper, we proposed an effective and efficient academic paper retrieval system, which cannot only tell users “what to read”, but also guide them about “how to read” and helps with “how to understand”. To this end, we formulated a new task named Reading Path Generation and proposed the first corresponding dataset SurveyBank to facilitate future research. To generate a useful learning path, a reasonable and scalable model was further proposed, in which we first expanded the initial seed papers returned from Google Scholar to construct a weighted citation network, and then reallocated the seed papers by selecting the co-citing papers. Finally, the reading path was generated by using a novel node-edge weighted Steiner tree algorithm. The effectiveness of our framework was demonstrated with extensive experiments. Since how to define a reasonable reading path for a given topic is an open question, and we propose it based on citation relationship, we are looking forward to more research works on this in the future.

IX. ACKNOWLEDGMENT

This work was in part supported by the Canada CIFAR AI Chair Program, the NSERC Discovery Grant RGPIN-2021-03115, and the UK Engineering and Physical Sciences Research Council under Grant EP/P011829/1.

REFERENCES

- [1] D. Zhou, S. Zhu, K. Yu, X. Song, B. L. Tseng, H. Zha, and C. L. Giles, "Learning multiple graphs for document recommendations," in *Proc. WWW*. ACM, 2008, pp. 141–150.
- [2] J. Tang and J. Zhang, "A discriminative approach to topic-based citation recommendation," *PAKDD*, pp. 572–579, 2009.
- [3] Q. He, J. Pei, D. Kifer, P. Mitra, and L. Giles, "Context-aware citation recommendation," in *Proc. WWW*. ACM, 2010, pp. 421–430.
- [4] X. Ren, J. Liu, X. Yu, U. Khandelwal, Q. Gu, L. Wang, and J. Han, "Cluscite: effective citation recommendation by information network-based clustering," in *Proc. SIGKDD*, 2014, pp. 821–830.
- [5] Y. Sun, H. Yin, and X. Ren, "Recommendation in context-rich environment: An information network analysis approach," in *Proc. WWW*, 2017, pp. 941–945.
- [6] D. Jurgens, S. Kumar, R. Hoover, D. McFarland, and D. Jurafsky, "Citation classification for behavioral analysis of a scientific field," *arXiv preprint arXiv:1609.00435*, 2016.
- [7] S. Teufel, A. Siddharthan, and D. Tidhar, "Automatic classification of citation function," in *Proc. EMNLP*, 2006, pp. 103–110.
- [8] C. Dong and U. Schäfer, "Ensemble-style self-training on citation classification," in *Proc. IJCNLP*, 2011, pp. 623–631.
- [9] H. Zhao, Z. Luo, C. Feng, and Y. Ye, "A context-based framework for resource citation classification in scientific literatures," in *Proc. SIGIR*, 2019, pp. 1041–1044.
- [10] A. Cohan, W. Ammar, M. van Zuylen, and F. Cady, "Structural scaffolds for citation intent classification in scientific publications," *arXiv preprint arXiv:1904.01608*, 2019.
- [11] D. Pride and P. Knoth, "An authoritative approach to citation classification," in *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries in 2020*, 2020, pp. 337–340.
- [12] J. G. Jardine, "Automatically generating reading lists," Computer Laboratory, University of Cambridge, Tech. Rep., 2014.
- [13] M. D. Ekstrand, P. Kannan, J. A. Stemper, J. T. Butler, J. A. Konstan, and J. T. Riedl, "Automatically building research reading lists," in *Proc. RecSys*, 2010, pp. 159–166.
- [14] J. Gordon, S. Aguilar, E. Sheng, and G. Burns, "Structured generation of technical reading lists," in *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, 2017, pp. 261–270.
- [15] K. Lo, L. L. Wang, M. Neumann, R. Kinney, and D. Weld, "S2ORC: The semantic scholar open research corpus," in *Proc. ACL*, Jul. 2020, pp. 4969–4983.
- [16] A. Schmidt and M. Wiegand, "A survey on hate speech detection using natural language processing," in *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, 2017, pp. 1–10.
- [17] I. Li, A. R. Fabbri, R. R. Tung, and D. R. Radev, "What should i learn first: Introducing lecturebank for nlp education and prerequisite chain learning," in *Proc. AAAI*, 2019, pp. 6674–6681.
- [18] A. Fabbri, I. Li, P. Trairatvorakul, Y. He, W. Ting, R. Tung, C. West-erfield, and D. Radev, "TutorialBank: A manually-collected corpus for prerequisite chains, survey extraction and resource recommendation," in *Proc. ACL*, Melbourne, Australia, Jul. 2018, pp. 611–620.
- [19] "Grobid," <https://github.com/kermitt2/grobid>, 2008–2020.
- [20] A. Bougouin, F. Boudin, and B. Daille, "TopicRank: Graph-based topic ranking for keyphrase extraction," in *Proc. IJCNLP*, Oct. 2013, pp. 543–551.
- [21] F. Boudin, "Pke: an open source python-based keyphrase extraction toolkit," in *Proc. COLING*, Osaka, Japan, December 2016, pp. 69–73.
- [22] "Pypdf2," <https://github.com/mstamy2/PyPDF2>.
- [23] "xmlltodict," <https://github.com/martinblech/xmlltodict>.
- [24] "Ccf," 1962–2021. [Online]. Available: <https://www.ccf.org.cn/>
- [25] "Serapi," 2018–2021. [Online]. Available: <https://serapi.com/>
- [26] A. Segev, "The node-weighted steiner tree problem," *Networks*, vol. 17, no. 1, pp. 1–17, 1987.
- [27] C. W. Duin and A. Volgenant, "Some generalizations of the steiner problem in graphs," *Networks*, vol. 17, no. 3, pp. 353–364, 1987.
- [28] Y. Sun, C. Ma, and S. Halgamuge, "The node-weighted steiner tree approach to identify elements of cancer-related signaling pathways," *BMC Bioinformatics*, vol. 18, no. 16, p. 551, 2017.
- [29] L. Page, S. Brin, R. Motwani, and T. Winograd, "The pagerank citation ranking: Bringing order to the web." Stanford InfoLab, Tech. Rep., 1999.
- [30] L. Kou, G. Markowsky, and L. Berman, "A fast algorithm for steiner trees," *Acta Informatica*, vol. 15, no. 2, pp. 141–145, 1981.
- [31] I. Beltagy, K. Lo, and A. Cohan, "Scibert: A pretrained language model for scientific text," *arXiv preprint arXiv:1903.10676*, 2019.
- [32] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [33] "Semantic scholar," 2015–2021. [Online]. Available: <https://www.semanticscholar.org/>
- [34] I. H. Witten, I. H. Witten, A. Moffat, T. C. Bell, T. C. Bell, and T. C. Bell, *Managing gigabytes: compressing and indexing documents and images*. Morgan Kaufmann, 1999.
- [35] T. I. Wang, K. H. Tsai, M.-C. Lee, and T. K. Chiu, "Personalized learning objects recommendation based on the semantic-aware discovery and the learner preference pattern," *ET&S*, pp. 84–105, 2007.
- [36] T. Bogers and A. van den Bosch, "Recommending scientific articles using citeulike," in *Proc. RecSys*, 2008, pp. 287–290.
- [37] Y. Tang, "The design and study of pedagogical paper recommendation," Ph.D. dissertation, University of Saskatchewan, 2008.
- [38] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Item-based collaborative filtering recommendation algorithms," in *Proc. WWW*, 2001, pp. 285–295.
- [39] H. Chang, D. Cohn, A. K. McCallum *et al.*, "Learning to create customized authority lists," in *ICML*, 2000, pp. 127–134.
- [40] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *Journal of Machine Learning Research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [41] J. Gordon, L. Zhu, A. Galstyan, P. Natarajan, and G. Burns, "Modeling concept dependencies in a scientific corpus," in *Proc. ACL*, 2016, pp. 866–875.
- [42] J. Yu, C. Wang, G. Luo, L. Hou, J. Li, J. Tang, and Z. Liu, "Course concept expansion in moocs with external knowledge and interactive game," *arXiv preprint arXiv:1909.07739*, 2019.
- [43] J. Yu, G. Luo, T. Xiao, Q. Zhong, Y. Wang, J. Luo, C. Wang, L. Hou, J. Li, Z. Liu *et al.*, "Moocube: A large-scale data repository for nlp applications in moocs," in *Proc. ACL*, 2020, pp. 3135–3142.
- [44] S. M. McNee, I. Albert, D. Cosley, P. Gopalkrishnan, S. K. Lam, A. M. Rashid, J. A. Konstan, and J. Riedl, "On the recommending of citations for research papers," in *Proc. CSCW*, 2002, pp. 116–125.
- [45] R. Torres, S. M. McNee, M. Abel, J. A. Konstan, and J. Riedl, "Enhancing digital libraries with techlens+," in *Proc. JCDL*, 2004, pp. 228–236.
- [46] O. Bar El, T. Milo, and A. Somech, "Automatically generating data exploration sessions using deep reinforcement learning," in *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, 2020, pp. 1527–1537.
- [47] H. P. Parette and J. B. Stoner, "Benefits of assistive technology user groups for early childhood education professionals," *Early Childhood Education Journal*, vol. 35, no. 4, pp. 313–319, 2008.
- [48] T. Milo and A. Somech, "Next-step suggestions for modern interactive data analysis platforms," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 576–585.
- [49] S. Basu Roy, S. Amer-Yahia, A. Chawla, G. Das, and C. Yu, "Constructing and exploring composite items," in *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, 2010, pp. 843–854.
- [50] S. B. Roy, G. Das, S. Amer-Yahia, and C. Yu, "Interactive itinerary planning," in *2011 IEEE 27th International Conference on Data Engineering*. IEEE, 2011, pp. 15–26.