Continuous-Time Relationship Prediction in Dynamic Heterogeneous Information Networks

SINA SAJADMANESH and SOGOL BAZARGANI, Sharif University of Technology

JIAWEI ZHANG, Florida State University

HAMID R. RABIEE, Sharif University of Technology

Online social networks, World Wide Web, media, and technological networks, and other types of so-called *information networks* are ubiquitous nowadays. These information networks are inherently *heterogeneous* and *dynamic*. They are heterogeneous as they consist of multi-typed objects and relations, and they are dynamic as they are constantly evolving over time. One of the challenging issues in such heterogeneous and dynamic environments is to forecast those relationships in the network that will appear in the future. In this article, we try to solve the problem of continuous-time relationship prediction in dynamic and heterogeneous information networks. This implies predicting the time it takes for a relationship to appear in the future, given its features that have been extracted by considering both heterogeneity and temporal dynamics of the underlying network. To this end, we first introduce a feature extraction framework that combines the power of meta-path-based modeling and recurrent neural networks to effectively extract features suitable for relationship prediction regarding heterogeneity and dynamicity of the networks. Next, we propose a supervised non-parametric approach, called *Non-Parametric Generalized Linear Model* (NP-GLM), which infers the hidden underlying probability distribution of the relationship building time given its features. We then present a learning algorithm to train NP-GLM and an inference method to answer time-related queries. Extensive experiments conducted on synthetic data and three real-world datasets, namely Delicious, MovieLens, and DBLP, demonstrate the effectiveness of NP-GLM in solving continuous-time relationship prediction problem vis-à-vis competitive baselines.

CCS Concepts: • **Information systems** → **Data mining; Social recommendation**; • **Computing methodologies** → *Machine learning*;

Additional Key Words and Phrases: Link prediction, social network analysis, heterogeneous network, non-parametric modeling, recurrent neural network, autoencoder

ACM Reference format:

Sina Sajadmanesh, Sogol Bazargani, Jiawei Zhang, and Hamid R. Rabiee. 2019. Continuous-Time Relationship Prediction in Dynamic Heterogeneous Information Networks. *ACM Trans. Knowl. Discov. Data* 13, 4, Article 44 (July 2019), 31 pages.

<https://doi.org/10.1145/3333028>

This work is partially supported by NSF through Grant IIS-1763365.

Authors' addresses: S. Sajadmanesh, S. Bazargani, and H. R. Rabiee, Department of Computer Engineering, Sharif University of Technology, Azadi Ave, Tehran 1458889694, Iran; emails: sajadmanesh@ce.sharif.edu, sogolb2@illinois.edu, rabiee@sharif.edu; J. Zhang, IFM Lab, Department of Computer Science, Florida State University, 1017 Academic Way, Tallahassee, Florida 32304; email: jiawei@ifmlab.org.

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.

© 2019 Association for Computing Machinery.

1556-4681/2019/07-ART44 \$15.00

<https://doi.org/10.1145/3333028>

1 INTRODUCTION

Link prediction is the problem of prognosticating a certain relationship, like interaction or collaboration, between two entities in a networked system that are not connected already [23]. Due to the popularity and ubiquity of networked systems in the real world, such as social, economic, or biological networks, this problem has attracted a considerable attention in recent years and has found its applications in various interdisciplinary domains, such as viral marketing, bioinformatics, recommender systems, and social network analysis [43]. For example, suggesting new friends in an online social network [21] or predicting drug-target interactions in a biological network [7] are two quite different problems, but can both cast as the prediction task of friendship links and drug-target links, respectively.

The problem of link prediction has a long literature and is studied extensively in the last decade. Initial works on link prediction problem mostly concentrated on homogeneous networks, which are composed of single type of nodes connected by links of the same type [21, 22, 40]. However, many of today's networks, such as online social networks or bibliographic networks, are inherently *heterogeneous*, in which multiple types of nodes are interconnected using multiple types of links [31, 37]. For example, a bibliographic network may contain author, paper, venue, and so on as different node types; and write, publish, cite, and so on as diverse link types that bind nodes with different types to each other. In these heterogeneous networks, the concept of a link can be generalized to a relationship, which can be constructed by combining different links with different types. For instance, the author-cite-paper relationship can be defined in a bibliographic network as a combination of author-write-paper and paper-cite-paper links. Analogously, one can generalize the link prediction to *relationship prediction* in heterogeneous networks that tries to predict complex relationships instead of links [34].

While most of the studies on the link/relationship prediction in heterogeneous networks utilize a static snapshot of the underlying network, many of these networks are *dynamic* in nature, which means that new nodes and linkages are continually added to the network, and some existing nodes and links may be removed from the network over time. For example, in online social networks, such as Facebook, new users are joining in the network every day, and new friendship links are being added to the network gradually. This dynamic characteristic causes the structure of the network to change and evolve over time, and taking these changes into account can significantly boost the quality of link prediction task [28].

In recent years, newer studies have shifted from traditional link prediction on static and homogeneous networks toward newer domains, considering heterogeneity and dynamicity of networks [10, 12, 14, 25, 29]. However, most of these works merely focus on one of these aspects, disregarding the other. Although there are quite a few studies that address both the challenges of heterogeneity and dynamicity [2, 30], to the best of our knowledge, all of them have ultimately formulated the link prediction problem as a binary classification task, i.e., predicting *whether* a link will appear in the network in the future. However, in dynamic networks, new links are continually appearing over time. So a much more interesting problem, which we call it *continuous-time link prediction* in this article, is to predict *when* a link will emerge or appear between two nodes in the network. Examples of this problem include predicting the time at which two individuals become friends in a social network or the time when two authors collaborate on writing a paper in a bibliographic network [34]. Inferring the link formation time in advance can be very useful in many concrete applications in different disciplines, such as sociology, economics, biology, and epidemiology, where the interactions between entities can be modeled via timed links. For example, in the biological context, predicting the marker proteins interaction time in a gene regulatory network will lead to predicting tumor progression and prognosis [38]. As another example in online social networks, if the recommender system could predict the relationship building time between two people, then it

can issue a friendship suggestion close to that time since it will have a relatively higher chance to be accepted. Good continuous-time link prediction results will lead to denser connections among users, and can greatly improve users' engagement that is the ultimate goal of online social networks [20].

In this article, we aim to solve the problem of continuous-time relationship prediction, in which we forecast the relationship building time between two nodes in a dynamic and heterogeneous environment. This problem is very challenging from the technical perspective, and cannot be solved trivially for three main reasons. First, the formulation of continuous-time relationship prediction is quite different from the conventional link prediction due to the involvement of temporal dynamics of the network and the necessity of considering network evolution time-line. Second, we only know the building time of those relationships that are already present at the network and for the rest of them that are yet to happen, which are excessive in number versus the existing ones, we lack such information. Finally, as opposed to the works concerning the binary link prediction, there are very rare works in the literature on continuous-time link prediction that attempt to answer the "when" question. To the best of our knowledge, the only work that has studied the continuous-time relationship prediction problem so far is proposed by Sun et al. [34]. They infer a probability distribution over time for each pair of nodes given their features and answer time-related queries about the relationship building time between the two nodes using the inferred distribution. However, the drawback of their method, not to mention neglecting the temporal dynamics of the network, is that it mainly relies on the assumption that relationship building times are coming from a certain probability distribution that must be fixed beforehand. This assumption though simplifying is very restrictive, because in real applications this distribution is unknown, and considering any specific one as *a priori* could be far from reality or limit the solution generality.

In order to address the above challenges, we propose a supervised non-parametric method to solve the problem of continuous-time relationship prediction. To this end, we first formally define the continuous-time relationship prediction problem and formulate the approach to solve it generally. Then, we introduce our novel feature extraction framework, which leverages meta-path-based modeling and recurrent neural networks to deal with heterogeneity and dynamicity of information networks. Next, we present *Non-Parametric Generalized Linear Model* (NP-GLM) that models the distribution of relationship building time given the extracted features. The strength of this non-parametric model is that it is capable of learning the underlying distribution of the relationship building time, as well as the contribution of each extracted feature in the network. Afterward, we propose an inference algorithm to answer queries, like the most probable time by which a relationship will appear between two nodes or the probability of relationship creation between them during a specific period. Finally, we conduct comprehensive experiments over a synthetic dataset to verify the correctness of NP-GLM's learning algorithm, and on three real-world dataset—DBLP, Delicious, and MovieLens—to demonstrate the effectiveness and generality of the proposed method in predicting the relationship building time versus the relevant baselines. As a summary, we can enumerate our major contributions as follows:

- (1) The proposed feature extraction framework can tackle heterogeneity of the data as well as capturing the temporal dynamics of the network by incorporating meta-path-based features into a recurrent neural network based autoencoder.
- (2) Our non-parametric model takes a unique approach toward learning the underlying distribution of relationship building time without imposing any significant assumptions on the problem.
- (3) Extensive evaluations over both synthetic and real-world datasets are performed to investigate the effectiveness of the proposed method.

- (4) To the best of our knowledge, this article is the first one that studies the continuous-time relationship prediction problem in both dynamic and heterogeneous network configurations.

The rest of this article is organized as follows. In Section 2, we provide introductory backgrounds on the concept and formally define the problem of continuous-time relationship prediction. Then, in Section 3, we introduce our novel feature extraction framework. Next, we go through the details of the proposed NP-GLM method in Section 4, explaining its learning method and how it answers inference queries. Experiments on synthetic data and real-world datasets are described in Sections 5 and 6, respectively. Section 7 discusses the related works, and finally in Section 8, we conclude the article.

2 PROBLEM FORMULATION

In this section, we introduce some important concepts and definitions used throughout the article and formally define the problem of continuous-time relationship prediction.

2.1 Heterogeneous Information Networks

An information network is *heterogeneous* if it contains multiple kinds of nodes and links. Formally, it is defined as a directed graph $G = (V, E)$, where $V = \bigcup_i V_i$ is the set of nodes comprising the union of all the node sets V_i of type i . Similarly, $E = \bigcup_j E_j$ is the set of links constituted by the union of all the link sets E_j of type j . Now, we bring the definition of the *network schema* [35], which is used to describe a heterogeneous information network at a meta-level:

Definition 2.1 (Network Schema). The schema of a heterogeneous network G is a graph $\mathcal{S}_G = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of different node types and $\mathcal{E}$ is the set of different link types in G .

In this article, we focus on the following three different heterogeneous and dynamic networks: (1) DBLP bibliographic network¹; (2) Delicious bookmarking network²; and (3) MovieLens recommendation network.³ The schema of these networks is depicted in Figure 1. As an example, in the bibliographic network, $\mathcal{V} = \{Author, Paper, Venue, Term\}$ is the set of different node types, and $\mathcal{E} = \{\text{write, publish, mention, cite}\}$ is the set of different link types.

Analogous to homogeneous networks where an adjacency matrix is used to represent whether pairs of nodes are linked to each other or not, in heterogeneous networks, we define *Heterogeneous Adjacency Matrices* to represent the connectivity of nodes of different types:

Definition 2.2 (Heterogeneous Adjacency Matrix). Given a heterogeneous network G with schema $\mathcal{S}_G = (\mathcal{V}, \mathcal{E})$, for each link type $\varepsilon \in \mathcal{E}$ denoting the relation between node types $v_i, v_j \in \mathcal{V}$, the heterogeneous adjacency matrix M_ε is a binary $|V_{v_i}| \times |V_{v_j}|$ matrix representing whether nodes of type v_i are in relation with nodes of type v_j with link type ε or not.

For instance, in the bibliographic network, the heterogeneous adjacency matrix M_{write} is a binary matrix where each row is associated with an author and each column is associated with a paper, and $M_{\text{write}}(i, j)$ indicates if the author i has written the paper j .

As we mentioned in the Introduction section about heterogeneous networks, the concept of a link can be generalized to a relationship. In this case, a relationship could be either a single link or a composite relation constituted by the concatenation of multiple links that together have a particular semantic meaning. For example, the co-authorship relation in the bibliographic network

¹<http://dblp.uni-trier.de/>.

²<http://delicious.com/>.

³<https://movielens.org/>.

Fig. 1. Schema of three different heterogeneous networks. Underlined characters are used as abbreviations for corresponding node types.

with the schema shown in Figure 1(a) can be defined as the combination of two *Author-write-Paper* links, making *Author-write-Paper-write-Author* relation. When dealing with link or relationship prediction in heterogeneous networks, we must exactly specify what kind of link or relationship we are going to predict. This specific relation to be predicted is called the *Target Relation* [34]. For example, in DBLP bibliographic network we aim to predict if and when an author will cite a paper from another author. Thus, the target relation, in this case, would be *Author-write-Paper-cite-Paper-write-Author*.

2.2 Dynamic Information Networks

An information network is *dynamic* when its nodes and linkage structure can change over time. That is, in a dynamic information network, all nodes and links are associated with a birth and death time. More formally, a dynamic network at the timestamp τ is defined as $G^\tau = (V^\tau, E^\tau)$ where V^τ and E^τ are, respectively, the set of nodes and the set of links existing in the network at the timestamp τ .

In this article, we consider the case that an information network is both dynamic and heterogeneous. This means that all network entities are associated with a type, and can possibly have birth and death times, regardless of their types. The bibliographic network is an example of both dynamic and heterogeneous one. Whenever a new paper is published, a new *Paper* node will be added to the network, alongside with the corresponding new *Author*, *Term*, and *Venue* nodes (if they do not exist yet). New links will be formed among these newly added nodes to indicate the *write*, *publish*, and *mention* relationships. Some linkages might also form between the existing nodes and the new ones, like new *cite* links connecting the new paper with the existing papers in its reference list.

In order to formally describe the state of a heterogeneous and dynamic network at any timestamp τ , we define the *time-aware heterogeneous adjacency matrix* in the following.

Definition 2.3 (Time-Aware Heterogeneous Adjacency Matrix). Given a dynamic heterogeneous network G^τ with schema $\mathcal{S}_G = (\mathcal{V}, \mathcal{E})$, for each link type $\varepsilon \in \mathcal{E}$ denoting the relation between node types $v_i, v_j \in \mathcal{V}$, the time-aware heterogeneous adjacency matrix M_ε^τ is a binary matrix representing if nodes of type v_i are in relation with nodes of type v_j with link type ε at the timestamp τ . More formally, for $a \in v_i$ and $b \in v_j$ we have

$$M_\varepsilon^\tau(a, b) = \begin{cases} 1, & \text{if } (a, b) \in \varepsilon \text{ and } bt(a, b) < \tau \leq dt(a, b), \\ 0, & \text{otherwise} \end{cases},$$

where $bt(a, b)$ and $dt(a, b)$ denote the birth and the death time of the link (a, b) , respectively.

2.3 Continuous-Time Relationship Prediction

Suppose that we are given a dynamic and heterogeneous information network as G^τ lastly observed at the timestamp τ , together with its network schema $\mathcal{S}_G$. Now, given the target relation R , the aim of continuous-time relationship prediction is to forecast the building time $t \geq \tau$ of the target relation R between any node pair (a, b) in G^τ .

In order to solve this problem given a pair of nodes like (a, b) , we try to train a supervised model that can predict a point estimate on the time it takes for the relationship of type R to be formed between them. The input to such a model will be a feature vector $\mathbf{x}$ corresponding to the node pair (a, b) . The model will then output with a continuous variable t that indicates when the relationship of type R will be built between a and b . To train such a model, we need to assemble a dataset comprising the feature vectors of all the node pairs between which the relation R have already been formed. The process of selecting sample node pairs, extracting their feature vector, and training the supervised model are explained in the subsequent sections.

3 FEATURE EXTRACTION FRAMEWORK

In this section, we present our feature extraction framework that is designed to have three major characteristics: First, it effectively considers different type of nodes and links available in a heterogeneous information network and regards their impact on the building time of the target relationship. Second, it takes the temporal dynamics of the network into account and leverages the network evolution history instead of simply aggregating it into a single snapshot. Finally, the extracted features are suitable for not only the link prediction problem but also the generalized *relationship prediction*. We will incorporate these features in the proposed non-parametric model in Section 4 to solve the continuous-time relationship prediction problem.

3.1 Data Preparation For Feature Extraction

To solve the problem of continuous-time relationship prediction in dynamic networks, we need to pay attention to the temporal history of the network data from two different points of view. First, we have to mind the evolutionary history of the network for feature extraction, so that the extracted features reflect the changes made in the network over time. Second, we have to specify the exact relationship building time for each pair of nodes that have formed the target relationship. This is because our goal is to train a supervised model to predict a continuous variable, which in this case is the building time of the target relationship. Hence, for each sample pair of nodes, we need a feature vector $\mathbf{x}$, associated with a target variable t that indicates the building time of the target relationship between them.

Suppose that we have observed a dynamic network G^τ recorded in the interval $t_0 < \tau \leq t_1$. According to Figure 2, we split this interval into two parts: the first part for extracting the feature $\mathbf{x}$, and the second for determining the target variable t . We refer to the first interval as *Feature Extraction Window* whose length is denoted by Φ , and the second as *Observation Window*, whose

Fig. 2. The evolutionary timeline of the network data.

length is denoted by Ω . Now, based on the existence of the target relationship in the observation window, all the node pairs in the network will fall within either one of the following three different groups:

- (1) Node pairs that form the target relationship before the beginning of the observation window (in the feature extraction window).
- (2) Node pairs that form the target relationship in the observation window for the first time (not existing before in the feature extraction window).
- (3) Node pairs that do not form the target relationship (neither in the feature extraction window nor in the observation window).

The node pairs in the second and third categories constitute our data samples, and will be used in the learning procedure to train the supervised model. For such pairs, we extract their feature vector $\mathbf{x}$ using the history available in the feature extraction window. For each node pair in the second category, we see that the target relationship between them has been created at a time like $t_r \in (t_0 + \Phi, t_1]$. So we set $t = t_r - (t_0 + \Phi)$ as the time it takes for the relationship to form since the beginning of the observation window. For these samples, we also set an auxiliary variable $y = 1$ that indicates that we have *observed* their exact building time. On the other hand, For node pairs in the third category, we have not seen their exact building time, but we know that it should be definitely after t_1 . For such samples, that we call *censored* samples, we set $t = t_1 - (t_0 + \Phi)$ that is equal to the length of the observation window Ω , and set $y = 0$ to indicate that the recorded time is, in fact, a lower bound on the true relationship building time. These type of samples are also of interest because their features will give us some information about their time falling after t_1 . As a result, each data sample is associated with a triple $(\mathbf{x}, y, t)$ representing its feature vector, observation status, and the time it takes for the target relationship to be formed, respectively.

3.2 Dynamic Feature Extraction

In this part, we describe how to utilize the temporal history of the network in the feature extraction window in order to extract features for continuous-time relationship prediction problem. We first begin with the meta-path-based feature set for heterogeneous information networks, and then incorporate these features into a *recurrent neural network based autoencoder* to exploit the temporal dynamics of the network as well. Hereby, we begin by defining the concept of meta-path [35]:

Definition 3.1 (Meta-Path). In a heterogeneous information network, a meta-path is a directed path following the graph of the network schema to describe the general relations that can be derived from the network. Formally speaking, given a network schema $\mathcal{S}_G = (\mathcal{V}, \mathcal{E})$, the sequence $v_1 \xrightarrow{\varepsilon_1} v_2 \xrightarrow{\varepsilon_2} \dots v_{k-1} \xrightarrow{\varepsilon_{k-1}} v_k$ is a meta-path defined on $\mathcal{S}_G$, where $v_i \in \mathcal{V}$ and $\varepsilon_i \in \mathcal{E}$.

Meta-paths are commonly used in heterogeneous information networks to describe multi-typed relations that have concrete semantic meanings. For example, in the bibliographic network whose schema is shown in Figure 1(a), we can define the co-authorship relation by the following

meta-path:

$$Author \xrightarrow{write} Paper \xleftarrow{write} Author,$$

or simply by $A \rightarrow P \leftarrow A$. Another example is the author citation relation, which in this article is used as the target relation for DBLP network. It can be specified as

$$Author \xrightarrow{write} Paper \xrightarrow{cite} Paper \xleftarrow{write} Author,$$

abbreviated as $A \rightarrow P \rightarrow P \leftarrow A$.

We can extend the concept of the heterogeneous adjacency matrix, which is used to indicate relationships between nodes of different types, to *meta-path adjacency matrix*, which we will use to indicate the number of path instances between two nodes of (possibly) different types, as explained below.

Definition 3.2 (Meta-Path Adjacency Matrix). Given a heterogeneous network G with schema $S_G = (\mathcal{V}, \mathcal{E})$, and the meta-path $v_1 \xrightarrow{\varepsilon_1} v_2 \xrightarrow{\varepsilon_2} \dots v_{k-1} \xrightarrow{\varepsilon_{k-1}} v_k$ defined over S_G denoting the relation between node types $v_i, v_j \in \mathcal{V}$, the meta-path adjacency matrix M_Ψ is defined as follows:

$$M_\Psi = \prod_{i=1}^{k-1} M_{\varepsilon_i},$$

which indicates the number of path instances between any node pair $u \in v_1$ and $v \in v_k$ following the meta-path Ψ . The time-aware counterpart of meta-path adjacency matrix is defined analogously by using the time-aware heterogeneous adjacency matrix.

Among the possible meta-paths that can be defined on a network schema, there are some that capture the similarity between two nodes. For example, the co-authorship meta-path $A \rightarrow P \leftarrow A$ in a bibliographic network creates a sense of similarity between two *Author* nodes. These type of meta-paths, called *similarity meta-paths*, are widely used to define topological features for link prediction problem in heterogeneous networks [29, 33, 46]. Table 1 presents a number of similarity meta-paths that can be defined on DBLP, Delicious, and MovieLens networks to capture the heterogeneous similarity between different node types.

The concept of similarity meta-paths can be extended to define heterogeneous features suitable for relationship prediction problem, where we have a target relation. Here, we follow the same approach as in [34], which suggests the following three meta-path-based building blocks to describe features for relationship prediction problem, given a target relation between two nodes of type A and B :

- (1) $A \overset{\text{similarity}}{\rightsquigarrow} A \overset{\text{target}}{\rightarrow} B$
- (2) $A \overset{\text{target}}{\rightarrow} B \overset{\text{similarity}}{\rightsquigarrow} B$
- (3) $A \overset{\text{relation}}{\rightsquigarrow} C \overset{\text{relation}}{\rightsquigarrow} B$

where $\rightsquigarrow$ denotes a meta-path, with labels *similarity* and *target* denoting a similarity meta-path and the target relation, respectively. The *relation* label denotes an arbitrary meta-path relating two nodes of possibly different types. The first block tells that there are some nodes of type A similar to a single node of the same type that has made the target relationship with a node of type B . Therefore, those similar nodes may also form the target relation with the type B node. An analogous intuition is behind the second block. For the third, it says that some nodes of type A are in relation with some type C nodes, which are themselves in relation with some nodes of type B . Hence, it is likely that type A nodes form some relationships, such as the target relationship,

Table 1. Similarity Meta-Paths in Different Networks

Network	Meta-path	Semantic meaning
DBLP	$A \rightarrow P \leftarrow A$	Authors co-write a paper
	$A \rightarrow P \leftarrow A \rightarrow P \leftarrow A$	Authors have common co-author
	$A \rightarrow P \leftarrow V \rightarrow P \leftarrow A$	Authors publish in the same venue
	$A \rightarrow P \rightarrow T \leftarrow P \leftarrow A$	Authors use the same term
	$A \rightarrow P \rightarrow P \leftarrow P \leftarrow A$	Authors cite the same paper
	$A \rightarrow P \leftarrow P \rightarrow P \leftarrow A$	Authors are cited by the same paper
Delicious	$U \leftrightarrow U \leftrightarrow U$	Users have common contact
	$U \rightarrow B \leftarrow U$	Users post the same bookmark
	$U \rightarrow B \rightarrow T \leftarrow B \leftarrow U$	Users post bookmarks with the same tag
MovieLens	$M \rightarrow A \leftarrow M$	Movies share an actor
	$M \rightarrow C \leftarrow M$	Movies belong to the same country
	$M \rightarrow D \leftarrow M$	Movies have the same director
	$M \rightarrow G \leftarrow M$	Movies have the same genre
	$M \rightarrow T \leftarrow M$	Movies have the same tag
	$U \rightarrow M \leftarrow U$	Users rate common movie
	$U \rightarrow M \rightarrow A \leftarrow M \leftarrow U$	Users rate movies sharing an actor
	$U \rightarrow M \rightarrow C \leftarrow M \leftarrow U$	Users rate movies from the same country
	$U \rightarrow M \rightarrow D \leftarrow M \leftarrow U$	Users rate movies of the same director
	$U \rightarrow M \rightarrow G \leftarrow M \leftarrow U$	Users rate movies with the same genre
	$U \rightarrow M \rightarrow T \leftarrow M \leftarrow U$	Users rate movies with the same tag

with type B nodes. We refer to the meta-paths that are created using these three blocks as *feature meta-paths*.

As an example in DBLP bibliographic network, for the target relation, we use $A \rightarrow P \rightarrow P \leftarrow A$ as a meta-path denoting the author citation relation. In Addition, Paper-cite-Author ($P \rightarrow P \rightarrow A$) and Author-cite-Paper ($A \rightarrow P \rightarrow P$) are also used as the arbitrary relations, and the similarity meta-paths for DBLP network from Table 1 are used to define the features for author citation relationship prediction.

After specifying feature meta-paths, we need a method to quantify them as numeric features. Due to the dynamicity of the network, different links are emerging and vanishing from the network over time. Therefore, the quantifying method must handle this dynamicity. Here, we formally define *Time-Aware Meta-Path-Based Features*:

Definition 3.3 (Time-Aware Meta-Path-Based Feature). Suppose that we are given a dynamic heterogeneous network G^T along with its network schema $S_G = (\mathcal{V}, \mathcal{E})$, and a target Relation $A \rightsquigarrow B$. For a given pair of nodes $a \in A$ and $b \in B$, and a feature meta-path $\Psi = A \xrightarrow{\varepsilon_1} v_1 \xrightarrow{\varepsilon_2} \dots v_{n-1} \xrightarrow{\varepsilon_n} B$ defined on S_G , the time-aware meta-path-based feature at the timestamp τ is the number of path instances between a and b following Ψ :

$$f_{\Psi}^{\tau}(a, b) = M_{\Psi}^{\tau}[a, b].$$

This way, for any pair of nodes, we can quantify the number of path instances of a particular meta-path at any specific timestamp τ . Although this quantification requires matrix multiplication, it can be done efficiently due to the following reasons:

- (1) The heterogeneous adjacency matrices are highly sparse, thus for calculating meta-path adjacency matrices, we can considerably reduce the time complexity of each single matrix multiplication by using fast sparse matrix multiplication algorithms [18].
- (2) The process of calculating the meta-path adjacency matrices is highly parallelizable, as the corresponding meta-paths decouples into simpler similarity meta-paths, which themselves decouple further into link types. Therefore, we can calculate the adjacency matrix of different similarity meta-paths in parallel, and then multiply them together to obtain the feature meta-path adjacency matrices.
- (3) Due to the similarity meta-paths sharing common sub-paths, computation time for the similarity meta-paths can also be saved using dynamic programming to avoid recalculating previously computed products. For example, for the DBLP dataset, if the target relation is $A \rightarrow P \rightarrow P \leftarrow A$, then by using the similarity meta-paths shown in Table 1, the path $A \rightarrow P \rightarrow P$ will appear in all the following feature meta-paths:

$$A \rightarrow P \rightarrow P \leftarrow P \leftarrow A$$

$$A \rightarrow P \rightarrow P \rightarrow P \leftarrow A$$

$$A \rightarrow P \rightarrow P \leftarrow A.$$

Therefore, we can calculate $M_{A \rightarrow P \rightarrow P}$ once and then reuse it in the calculation of the adjacency matrices of the above meta-paths.

- (4) Finally, the symmetry of the similarity meta-paths further reduces the number of products, because we can calculate the matrix corresponding to half of the path, and then multiply the resulting matrix by its transpose. For instance, the adjacency matrix of the similarity meta-path $A \rightarrow P \leftarrow V \rightarrow P \leftarrow A$ can be calculated as $X \cdot X^T$, where $X = M_{\text{write}} \cdot M_{\text{publish}}$, reducing the number of multiplications from three to two.

So far we proposed a method to calculate the time-aware meta-path-based features, which is the number of path instances of a particular meta-path at the timestamp τ . If we set this timestamp to the end of the feature extraction window, it is as though we are aggregating the whole network into a single snapshot observed at time $t_0 + \Phi$. In order to avoid such an aggregation, we divide the feature extraction window into a sequence of k contiguous intervals of a constant size Δ , as shown in Figure 2. By doing so, we intend to extract time-aware features in each sub-window that results in a multivariate time series containing the information about the temporal evolution of the topological features between any pair of nodes. With this in mind, we define *Dynamic Meta-Path-based Time Series* as follows:

Definition 3.4 (Dynamic Meta-Path-Based Time Series). Suppose that we are given a dynamic heterogeneous network G^τ observed in a feature extraction window of size Φ ($t_0 < \tau \leq t_0 + \Phi$), along with its network schema $\mathcal{S}_G = (\mathcal{V}, \mathcal{E})$ and a target relation $A \rightsquigarrow B$. Also, suppose that the feature extraction window is divided into k fragments of size Δ . For a given pair of nodes $a \in A$ and $b \in B$ in $G^{t_0+\Phi}$, and a meta-path Ψ defined on $\mathcal{S}_G$, the dynamic meta-path-based time series of (a, b) is calculated as follows:

$$x_\Psi^i(a, b) = f_\Psi^{t_0+i\Delta}(a, b) - f_\Psi^{t_0+(i-1)\Delta}(a, b) \quad i = 1 \dots k.$$

For each feature meta-path designed using the triple building blocks described before, we get a unique time series. For each time step, we put the corresponding values from all the time series into a vector. Consequently, we get a multivariate time series where each time step is vector-valued. For example, if we have d feature meta-paths Ψ_1 to Ψ_d , then each time step of the resulting time

Fig. 3. The architecture of the LSTM Autoencoder used for dynamic feature extraction. The first k steps depicts the manner of the working of the encoder LSTM, while the second k steps describes the decoder LSTM. The output of the k th stage is used as the feature vector x , which is fed into the decoder k times to produce the input sequence in the reversed order.

series for any node pair (a, b) will become

$$\mathbf{x}_{a,b}^i = \left[x_{\Psi_1}^i(a, b), \dots, x_{\Psi_d}^i(a, b) \right]^T, \quad i = 1 \dots k.$$

We refer to this vector-valued time series as *Multivariate Meta-Path-Based Time Series*. Such multivariate time series reflect how topological features change between two nodes across different snapshots of the network. Based on the level of the network dynamicity, it can capture increasing/decreasing trends or even periodic/re-occurring patterns.

Now it is time to convert the multivariate meta-path-based time series into a single feature vector so that we can use it as the input to our non-parametric model that will be discussed in the next section. A trivial solution would be to stack all the vector-valued time steps of the multivariate time series into a single vector. However, this approach will result in a very high-dimensional vector as the number of time steps increases and can lead to difficulties in the learning procedure due to the curse of dimensionality. This is in contrast with our expectation that more time steps would bring more information about the history of the network and should result in a better prediction model. To overcome this problem, we combine the power of recurrent neural networks, especially Long Short Term Memory (LSTM) units [17], which have proven to be very successful in handling time series and sequential data, with Autoencoders [4], which are widely used to learn alternative representations of the data such that the learned representation can reconstruct the original input. Our goal is to transform the multivariate meta-path-based time series into a compact vector representation such that the resulting vector holds as much information from the original multivariate time series as possible.

Inspired by the work of Dai and Le on semi-supervised sequence learning [9], we design an autoencoder that learns how to take a multivariate time series as input and compress it into a latent vector representation. The architecture of such autoencoder is illustrated in Figure 3. The autoencoder consists of two components as follows: (1) the encoder, which takes the input data and transforms it into a latent representation; and (2) the decoder, which takes the encoded representation and transforms it back to the input space. The autoencoder is trained in such a way that it can reconstruct the original input data.

As the purpose of using the autoencoder in this article is to compress multivariate time series, instead of using simple feed-forward neural networks, both encoder and decoder are built using LSTMs. The input to the encoder LSTM is a multivariate time series of length k . The encoder accepts the vector-valued time steps of the input multivariate time series sequentially. After receiving the k th time step, the output of the encoder LSTM will be the compressed feature vector that we will use as the input to the NP-GLM method. In order to train the encoder to learn how to compress the input time series, it is matched with a decoder LSTM. The decoder LSTM receives k copies of the compressed feature vector one after another, and with a proper loss function (such as mean squared error) it is forced to reconstruct the original multivariate time series in reverse

order. Reversing the output sequence will make the optimization of the model easier since it causes the decoder to revert back the changes made by the encoder to the input sequence.

The benefits of using the LSTM autoencoder is three-fold as follows: (1) since the autoencoder can reconstruct the original time series, which reflects the temporal dynamics of the network, we get minimum information loss in the compressed feature vector; (2) as we can set the dimensionality of the compressed feature vector to any desired value, we can evade the curse of dimensionality; and (3) due to the inherent dynamicity of recurrent neural networks and LSTMs, when we receive $(k + 1)$ th snapshot of the network, we can easily fine-tune the previous autoencoder that was learned with k snapshots to consider the new snapshot as well, instead of repeating the whole learning procedure from scratch.

To conclude this section, we quickly review the whole procedure of processing the network data, training the autoencoder, and assembling a training dataset for the supervised model to predict the building time of a particular target relation:

- (1) The network evolution timeline is split into the feature extraction window and the observation window.
- (2) Those node pairs that have either formed the target relationship in the observation window (observed samples) or have not formed the target relationship at all (censored samples) are selected as sample node pairs.
- (3) By extracting feature meta-paths based on the target relation and similarity meta-paths, a multivariate time series can be obtained for each sample node pair. Thus, if we have N sample node pairs, we will have a dataset of N multivariate time series.
- (4) The LSTM autoencoder is trained using the dataset of N multivariate time series to learn how to compress time series into feature vectors.
- (5) For each sample node pair, the corresponding multivariate time series is compressed into a feature vector $\mathbf{x}$ using the learned encoder LSTM.
- (6) For each observed node pair, the feature vector $\mathbf{x}$ is labeled with $y = 1$ and associates with the variable t denoting the time it takes for the node pair to form the target relationship. For censored node pairs, y is set to zero and t becomes equal to the size of the observation window.
- (7) Finally, we will have a dataset of the form $\{\mathbf{x}, y, t\}_i, i = 1 \dots N$ that will be used to train the supervised model.

We explain our proposed non-parametric model in the next section that takes the learned representation as the feature vector $\mathbf{x}$ and attempts to predict the corresponding event time t .

4 SUPERVISED NON-PARAMETRIC MODEL

In this section, we introduce our proposed model, called NP-GLM, to solve the problem of continuous-time relationship prediction based on the extracted features. Since the relationship building time is treated as a continuous random variable, we attempt to model the probability distribution of this time, given the features of the target relationship. Thus, if we denote the target relationship building time by t and its features by $\mathbf{x}$, our aim is to model the probability density function $f_T(t | \mathbf{x})$. A conventional approach to modeling this function is to fix a parametric distribution for t (e.g., Exponential distribution) and then relate $\mathbf{x}$ to t using a Generalized Linear Model [34]. The major drawback of this approach is that we need to know the exact distribution of the relationship building time, or at least, we could guess the best one that fits. The alternative way that we follow is to learn the shape of $f_T(t | \mathbf{x})$ from the data using a non-parametric solution.

Table 2. Characteristics of Some Probability Distributions Used for Event-Time Modeling

Distribution	Density function $f_T(t)$	Survival function $S(t)$	Intensity function $\lambda(t)$	Cumulative intensity $\Lambda(t)$
Exponential	$\alpha \exp(-\alpha t)$	$\exp(-\alpha t)$	α	αt
Rayleigh	$\frac{t}{\sigma^2} \exp\left(-\frac{t^2}{2\sigma^2}\right)$	$\exp\left(-\frac{t^2}{2\sigma^2}\right)$	$\frac{t}{\sigma^2}$	$\frac{t^2}{2\sigma^2}$
Gompertz	$\alpha e^t \exp\{-\alpha(e^t - 1)\}$	$\exp\{-\alpha(e^t - 1)\}$	αe^t	αe^t
Weibull	$\frac{\alpha t^{\alpha-1}}{\beta^\alpha} \exp\left\{-\left(\frac{t}{\beta}\right)^\alpha\right\}$	$\exp\left\{-\left(\frac{t}{\beta}\right)^\alpha\right\}$	$\frac{\alpha t^{\alpha-1}}{\beta^\alpha}$	$\left(\frac{t}{\beta}\right)^\alpha$

In the rest of this section, we first bring the necessary theoretical backgrounds related to the concept, then we go through the details of the proposed model. In the end, we explain the learning and inference algorithms of NP-GLM.

4.1 Background

Here, we define some essential concepts that are necessary to study before we proceed to the proposed model. Generally, the formation of a relationship between two nodes in a network can simply be considered as an event with its occurring time as a random variable T coming from a density function $f_T(t)$. Regarding this, we can have the following definitions:

Definition 4.1 (Survival Function). Given the density $f_T(t)$, the survival function denoted by $S(t)$, is the probability that an event occurs after a certain value of t , which means

$$S(t) = P(T > t) = \int_t^\infty f_T(t) dt. \quad (1)$$

Definition 4.2 (Intensity Function). The intensity function (or failure rate function), denoted by $\lambda(t)$, is the instantaneous rate of event occurring at any time t given the fact that the event has not occurred yet

$$\lambda(t) = \lim_{\Delta t \rightarrow 0} \frac{P(t \leq T \leq t + \Delta t \mid T \geq t)}{\Delta t}. \quad (2)$$

Definition 4.3 (Cumulative Intensity Function). The cumulative intensity function, denoted by $\Lambda(t)$, is the area under the intensity function up to a point t :

$$\Lambda(t) = \int_0^t \lambda(t) dt. \quad (3)$$

The relations between density, survival, and intensity functions come directly from their definitions as follows:

$$\lambda(t) = \frac{f_T(t)}{S(t)} \quad (4)$$

$$S(t) = \exp(-\Lambda(t)) \quad (5)$$

$$f_T(t) = \lambda(t) \exp(-\Lambda(t)). \quad (6)$$

Table 2 shows the density, survival, intensity, and cumulative intensity functions of some widely-used distributions for event time modeling.

4.2 Model Description

Looking at Equation (6), we see that the density function can be specified uniquely with its intensity function. Since the intensity function often has a simpler form than the density itself, if we learn the shape of the intensity function, then we can infer the entire distribution eventually.

Therefore, we focus on learning the shape of the conditional intensity function $\lambda(t | \mathbf{x})$ from the data, and then accordingly infer the conditional density function $f_T(t | \mathbf{x})$ based on the learned intensity. In order to reduce the hypothesis space of the problem and avoid the curse of dimensionality, we assume that $\lambda(t | \mathbf{x})$, which is a function of both t and $\mathbf{x}$, can be factorized into two separate positive functions as the following:

$$\lambda(t | \mathbf{x}) = g(\mathbf{w}^T \mathbf{x})h(t), \quad (7)$$

where g is a function of $\mathbf{x}$ that captures the effect of features via a linear transformation using coefficient vector $\mathbf{w}$ independent of t , and h is a function of t that captures the effect of time independent of $\mathbf{x}$. This assumption, referred to as *proportional hazards condition* [5], holds in GLM formulations of many event-time modeling distributions, such as the ones shown in Table 2. Our goal is now to fix the function g and then learn both the coefficient vector $\mathbf{w}$ and the function h from the training data. In order to do so, we begin with the likelihood function of the data that can be written as follows:

$$\prod_{i=1}^N f_T(t_i | \mathbf{x}_i)^{y_i} P(T \geq t_i | \mathbf{x}_i)^{1-y_i}. \quad (8)$$

The likelihood consists of the product of two parts: The first part is the contribution of those samples for which we have observed their exact building time, in terms of their density function. The second part on the other hand, is the contribution of the censored samples, for which we use the probability of the building time being greater than the recorded one. By applying Equations (5) and (6), we can write the likelihood in terms of the intensity function:

$$\prod_{i=1}^N \left[\lambda(t_i | \mathbf{x}_i) \exp\{-\Lambda(t_i | \mathbf{x}_i)\} \right]^{y_i} \exp\{-\Lambda(t_i | \mathbf{x}_i)\}^{1-y_i}. \quad (9)$$

By merging the exponentials and applying Equations (3) and (7), the likelihood function becomes

$$\prod_{i=1}^N \left[g(\mathbf{w}^T \mathbf{x}_i) h(t_i) \right]^{y_i} \exp \left\{ -g(\mathbf{w}^T \mathbf{x}_i) \int_0^{t_i} h(t) dt \right\}. \quad (10)$$

Since we do not know the form of $h(t)$, we cannot directly calculate the integral appeared in the likelihood function. To deal with this problem, we treat $h(t)$ as a non-parametric function by approximating it with a piecewise constant function that changes just in t_i s. Therefore, the integral over $h(t)$, denoted by $H(t)$, becomes a series:

$$H(t_i) = \int_0^{t_i} h(t) dt \simeq \sum_{j=1}^i h(t_j)(t_j - t_{j-1}), \quad (11)$$

assuming samples are sorted by t in increasing order, without loss of generality. The function $H(t)$ defined above plays an important role in both learning and inference phases. In fact, both the learning and inference phases rely on $H(t)$ instead of $h(t)$, which we will see later in this article. Replacing the above series in the likelihood, taking the logarithm and negating, we end up with the following negative log-likelihood function, simply called the *loss function*, denoted by L :

$$L(\mathbf{w}, h) = \sum_{i=1}^N \left\{ g(\mathbf{w}^T \mathbf{x}_i) \sum_{j=1}^i h(t_j)(t_j - t_{j-1}) - y_i \left[\log g(\mathbf{w}^T \mathbf{x}_i) + \log h(t_i) \right] \right\}. \quad (12)$$

The loss function depends on both the vector $\mathbf{w}$ and the function $h(t)$. In the next part, we explain an iterative learning algorithm to learn both $\mathbf{w}$ and $h(t)$ collectively.

4.3 Learning Algorithm

Minimizing the loss function (Equation (12)) relies on the choice of the function g . There are no particular limits on the choice of g except that it must be a non-negative function. For example, both quadratic and exponential functions of $\mathbf{w}^T \mathbf{x}$ will do the trick. Here, we proceed with $g(\mathbf{w}^T \mathbf{x}) = \exp(\mathbf{w}^T \mathbf{x})$ since it makes the loss function convex with respect to $\mathbf{w}$. Subsequent equations can be derived for other choices of g analogously.

Setting the derivative of the loss function with respect to $h(t_k)$ to zero yields a closed form solution for $h(t_k)$:

$$h(t_k) = \frac{y_k}{(t_k - t_{k-1}) \sum_{i=k}^N \exp(\mathbf{w}^T \mathbf{x}_i)}. \quad (13)$$

By applying Equation (11), we get the following for $H(t_i)$:

$$H(t_i) = \sum_{j=1}^i \frac{y_j}{\sum_{k=j}^N \exp(\mathbf{w}^T \mathbf{x}_k)}, \quad (14)$$

which depends on the vector $\mathbf{w}$. On the other hand, we cannot obtain a closed form solution for $\mathbf{w}$ from the loss function. Therefore, we turn to use Gradient-based optimization methods to find the optimal value of $\mathbf{w}$. The loss function with respect to $\mathbf{w}$ is as follows:

$$L(\mathbf{w}) = \sum_{i=1}^N \left\{ \exp(\mathbf{w}^T \mathbf{x}_i) H(t_i) - y_i \mathbf{w}^T \mathbf{x}_i \right\} + \text{Const.}, \quad (15)$$

which depends on the function H . As the learning of both $\mathbf{w}$ and H depends on each other, they should be learned collectively. Here, we use an iterative algorithm to learn $\mathbf{w}$ and H alternatively. We begin with a random vector $\mathbf{w}^{(0)}$. Then, in each iteration τ , we first update $H^{(\tau)}$ via Equation (14) using $\mathbf{w}^{(\tau-1)}$. Next, we optimize Equation (15) using the values of $H^{(\tau)}(t_i)$ to obtain $\mathbf{w}^{(\tau)}$. We continue this routine until convergence. Since this procedure successively reduces the value of the loss function, and as the loss function (i.e., the negative log-likelihood) is bounded from below, the algorithm will ultimately converge to a stationary point. The pseudo code of the learning procedure is given in Algorithm 1.

4.4 Inference Queries

In this part, we explain how to answer the common inference queries based on the inferred distribution $f_T(t \mid \mathbf{x})$. Suppose that we have learned the vector $\mathbf{w}$ and the function H using the training samples $(\mathbf{x}_i, y_i, t_i)$, $i = 1 \dots N$ following Algorithm 1. Afterward, for a testing relationship R associated with a feature vector $\mathbf{x}_R$, the following queries can be answered:

4.4.1 Ranged Probability. What is the probability for the relationship R to be formed between time t_α and t_β ? This is equivalent to calculating $P(t_\alpha \leq T \leq t_\beta \mid \mathbf{x}_R)$, which by definition is

$$\begin{aligned} P(t_\alpha \leq T \leq t_\beta \mid \mathbf{x}_R) &= S(t_\alpha \mid \mathbf{x}_R) - S(t_\beta \mid \mathbf{x}_R) \\ &= \exp \left\{ -g(\mathbf{w}^T \mathbf{x}_R) H(t_\alpha) \right\} - \exp \left\{ -g(\mathbf{w}^T \mathbf{x}_R) H(t_\beta) \right\}. \end{aligned} \quad (16)$$

The problem here is to obtain the values of $H(t_\alpha)$ and $H(t_\beta)$, as t_α and t_β may not be among t_i s of the training samples, for which H is estimated. To calculate $H(t_\alpha)$, we find $k \in \{1, 2, \dots, N\}$ such that $t_k \leq t_\alpha < t_{k+1}$. Due to the piecewise constant assumption for the function h , we get

$$h(t_\alpha) = \frac{H(t_\alpha) - H(t_k)}{t_\alpha - t_k}. \quad (17)$$

ALGORITHM 1: The learning algorithm of NP-GLM

Input: $\mathbf{X}_{N \times d} = (\mathbf{x}_1, \dots, \mathbf{x}_N)^T$ as d -dimensional feature vectors, $\mathbf{y}_{N \times 1}$ as observation states, and $\mathbf{t}_{N \times 1}$ as recorded times.

Output: Learned parameters $\mathbf{w}_{d \times 1}$ and $\mathbf{H}_{N \times 1}$.

converged $\leftarrow$ *False*;
threshold $\leftarrow 10^{-4}$;
 $\tau \leftarrow 0$;
 $L^{(\tau)} = \infty$;
Initialize $\mathbf{w}^{(\tau)}$ with random values;
while *Not converged* **do**
 $\tau \leftarrow \tau + 1$;
 Use Equation (14) to obtain $\mathbf{H}^{(\tau)}$ using $\mathbf{w}^{(\tau-1)}$;
 Minimize Equation (15) to obtain $\mathbf{w}^{(\tau)}$ using $\mathbf{H}^{(\tau)}$;
 Use Equation (12) to obtain $L^{(\tau)}$ using $\mathbf{w}^{(\tau)}$ and $\mathbf{H}^{(\tau)}$;
 if $\|L^{(\tau)} - L^{(\tau-1)}\| < \textit{threshold}$ **then**
 converged $\leftarrow$ *True*;
 end
end
 $\mathbf{w} \leftarrow \mathbf{w}^{(\tau)}$;
 $\mathbf{H} \leftarrow \mathbf{H}^{(\tau)}$;

On the other hand, since h only changes in t_i s, we have

$$h(t_\alpha) = h(t_{k+1}) = \frac{H(t_{k+1}) - H(t_k)}{t_{k+1} - t_k}. \quad (18)$$

Combining Equations (17) and (18), we get

$$H(t_\alpha) = H(t_k) + (t_\alpha - t_k) \frac{H(t_{k+1}) - H(t_k)}{t_{k+1} - t_k}. \quad (19)$$

Following the similar approach, we can calculate $H(t_\beta)$, and then answer the query using Equation (16). The dominating operation here is to find the value of k . Since we have t_i s sorted beforehand, this operation can be done using a binary search with $O(\log N)$ time complexity.

4.4.2 Quantile. By how long the target relationship R will be formed with probability α ? This question is equivalent to find the time t_α such that $P(T \leq t_\alpha \mid \mathbf{x}_R) = \alpha$. By definition, we have

$$1 - P(T \leq t_\alpha \mid \mathbf{x}_R) = S(t_\alpha \mid \mathbf{x}_R) = \exp \left\{ -g(\mathbf{w}^T \mathbf{x}_R) H(t_\alpha) \right\} = 1 - \alpha.$$

Taking logarithm of both sides and rearranging, we get

$$H(t_\alpha) = -\frac{\log(1 - \alpha)}{g(\mathbf{w}^T \mathbf{x}_R)}. \quad (20)$$

To find t_α , we first find k such that $H(t_k) \leq H(t_\alpha) < H(t_{k+1})$. We eventually have $t_k \leq t_\alpha < t_{k+1}$ since H is a non-decreasing function due to non-negativity of the function h . Therefore, we again end up with Equation (19), by rearranging which we get:

$$t_\alpha = (t_{k+1} - t_k) \frac{H(t_\alpha) - H(t_k)}{H(t_{k+1}) - H(t_k)} + t_k. \quad (21)$$

ALGORITHM 2: Synthetic dataset generation algorithm

Input: The number of observed samples N_o , the number of censored samples N_c , the dimension of the feature vectors d , and the desired distribution $dist$

Output: Synthetically generated data $\mathbf{X}_{N \times d}$, $\mathbf{y}_{N \times 1}$, and $\mathbf{t}_{N \times 1}$.

$N \leftarrow N_o + N_c$;

Draw a weight vector $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d)$, where $\mathbf{I}_d$ is the d -dimensional identity matrix;

Draw scalar intercept $b \sim \mathcal{N}(0, 1)$;

for $i \leftarrow 1$ **to** N **do**

 Draw feature vector $\mathbf{x}_i \sim \mathcal{N}(0, \mathbf{I}_d)$;

 Set distribution parameter $\alpha_i \leftarrow \exp(\mathbf{w}^T \mathbf{x}_i + b)$;

if $dist == \text{Rayleigh}$ **then**

 | Draw $t_i \sim \alpha_i t \exp\{-0.5\alpha_i t^2\}$;

else if $dist == \text{Gompertz}$ **then**

 | Draw $t_i \sim \alpha_i e^t \exp\{-\alpha_i(e^t - 1)\}$;

end

Sort pairs $(\mathbf{x}_i, t_i)$ by t_i in ascending order;

for $i \leftarrow 1$ **to** N_o **do**

 | $y_i \leftarrow 1$;

end

for $i \leftarrow (N_o + 1)$ **to** N **do**

 | $y_i \leftarrow 0$;

end

By combining Equations (20) and (21), we can obtain the value of t_α , which is the answer to the quantile query. It is worth mentioning that if $\alpha = 0.5$, then t_α becomes the median of the distribution $f_T(t | \mathbf{x}_R)$. Here, again the dominant operation is to find the value of k , which due to the non-decreasing property of the function H can be found using a binary search with $O(\log N)$ time complexity.

4.4.3 Random Sampling. Generating random samples from the inferred distribution can easily be carried out using the Inverse-Transform sampling algorithm. To pick a random sample from the inferred distribution $f_T(t | \mathbf{x})$, we first generate uniform random variable $u \sim \text{Uniform}(0, 1)$. Then, we find k such that $S(t_{k+1} | \mathbf{x}) \leq u \leq S(t_k | \mathbf{x})$. We output t_{k+1} as the generated sample. Again, searching for the suitable value of k is the dominant operation that can be undertaken via binary search with $O(\log N)$ time complexity.

5 SYNTHETIC EVALUATIONS

We use synthetic data to verify the correctness of NP-GLM and its learning algorithm. Since NP-GLM is a non-parametric method, we generate synthetic data using various parametric models with previously known random parameters and evaluate how well NP-GLM can learn the parameters and the underlying distribution of the generated data.

5.1 Experiment Setup

We consider generalized linear models of two widely used distributions for event-time modeling, Rayleigh and Gompertz, as the ground truth models for generating synthetic data. Algorithm 2 is used to generate a total of N data samples with d -dimensional feature vectors, consisting N_o non-censored (observed) samples and remaining $N_c = N - N_o$ censored ones. For all synthetic

Fig. 4. Convergence of NP-GLM's average log-likelihood ($\log \mathcal{L}$) for different number of training samples (N). Censoring ratio has been set to 0.5.

Fig. 5. Convergence of NP-GLM's average log-likelihood ($\log \mathcal{L}$) for different censoring ratios with 1K samples.

experiments, we generate 10-dimensional feature vectors ($d = 10$). We repeat every experiment 100 times and report the average results.

5.2 Experiment Results

5.2.1 Convergence Analysis. Since NP-GLM's learning is done in an iterative manner, we first analyze whether this algorithm converges as the number of iterations increases. We recorded the log-likelihood of NP-GLM, averaged over the number of training samples N in each iteration. We repeated this experiment for $N \in \{1,000, 2,000, 3,000\}$ with a fixed censoring ratio of 0.5, which means half of the samples are censored. The result is depicted in Figure 4. We can see that the algorithm successfully converges with a rate depending on the underlying distribution. For the case of Rayleigh, it requires about 100 iterations to converge but for Gompertz, this reduces to about 30. Also, we see that using more training data leads to achieving more log-likelihood as expected.

In Figure 5, we fixed $N = 1,000$ and performed the same experiment this time using different censoring ratios. According to the figure, we see that by increasing the censoring ratio, the convergence rate increases. This is because NP-GLM infers the values of $H(t)$ for all t in the observation window. Therefore, as the censoring ratio increases, the observation window is decreased, so NP-GLM has to infer a fewer number of parameters, leading to a faster convergence. Note that as opposed to Figure 4, here a higher log-likelihood does not necessarily indicate a better fit, due to the likelihood marginalization we get by censored samples.

Fig. 6. NP-GLM's mean absolute error (MAE) versus the number of training samples (N) for different censoring ratios.

Fig. 7. NP-GLM's mean absolute error (MAE) versus the number of censored samples (N_c) for different number of observed samples (N_o).

5.2.2 Performance Analysis. Next, we evaluated how good NP-GLM can infer the parameters used to generate synthetic data. To this end, we varied the number of training samples N and measured the mean absolute error (MAE) between the learned weight vector $\hat{\mathbf{w}}$ and the ground truth. Figure 6 illustrates the result for different censoring ratios. It can be seen that as the number of training samples increases, the MAE gradually decreases. The other point to notice is that more censoring ratio results in a higher error due to the information loss we get by censoring.

In another experiment, we investigated whether censored samples are informative or not. For this purpose, we fixed the number of observed samples N_o and changed the number of censored samples from 0 to 200. We measured the MAE between the learned $\mathbf{w}$ and the ground truth for $N_o \in \{200, 300, 400\}$. The result is shown in Figure 7. It clearly demonstrates that adding more censored samples causes the MAE to dwindle up to an extent, after which we get no substantial improvement. This threshold is dependent on the underlying distribution. In this case, for Rayleigh and Gompertz it is about 80 and 120, respectively.

5.2.3 Running Time Analysis. Finally, we assess the running time of NP-GLM's learning algorithm against the size of the training data when it becomes relatively large. To this end, we varied the number of samples from 10K to 100M and measured the average running time of the learning algorithm of NP-GLM on a single machine whose specification is reported in Table 3. Figure 8 depicts the result in log-log scale for Rayleigh and Gompertz distributions under different censoring ratios selected from the set $\{0.05, 0.25, 0.50\}$. It can be seen from the figure that the running time scales linearly with the number of training samples since the number of parameters to be inferred

Table 3. PC Specification and Configuration

Operating system	Windows 10
CPU	Intel Core i7 1.8GHz
RAM	12GB DDR III
GPU	Nvidia GeForce GT 750
Disk type	SSD
Programming language	Python 3.6

Fig. 8. NP-GLM's average running time (T) measured in seconds versus the number of training samples (N) in log-log scale for different censoring ratios.

in NP-GLM as a non-parametric model depends on the size of the training data. The censoring ratio though negligible in scale can impact the running time of the algorithm, with more censoring ratio resulting in less running time. This is because higher censoring ratio reduces the observation window, which in turn reduces the number of parameters.

6 EXPERIMENTS ON REAL DATA

We apply NP-GLM with the proposed feature set on a number of real-world datasets to evaluate its effectiveness and compare its performance in predicting the relationship building time vis-à-vis state of the art models.

6.1 Datasets

6.1.1 DBLP. We use the DBLP bibliographic citation network, provided by [36], which has both attributes of dynamicity and heterogeneity. The network contains four types of objects: authors, papers, venues, and terms. The network schema of this dataset is depicted in Figure 1(a). Each paper is associated with a publication date, with a granularity of one year. Based on the publication venue of the papers, we limited the original DBLP dataset to those papers that are published in venues relative to the theoretical computer science. This resulted in having about 16k authors and 37k papers published from 1969 to 2016 in 38 venues.

6.1.2 Delicious. Another dynamic and heterogeneous dataset we use in our experiments is the Delicious bookmarking dataset from [6], with a network schema presented in Figure 1(b). It contains three types of objects, namely users, bookmarks, and tags, whose numbers are about 1.7k, 31k, and 22k, respectively. The dataset includes bookmarking timestamps from May 2006 to October 2010.

Table 4. Demographic Statistics of Real-World Datasets

Dataset	Time span	Entity	Count			
DBLP	From 1969 to 2016	Nodes	<i>Author</i>	15,929	<i>Venue</i>	38
			<i>Paper</i>	37,077	<i>Term</i>	12,028
		Links	write	100,797	publish	42,872
			cite	165,904	mention	284,156
Delicious	From May 2006 to Oct 2010	Nodes	<i>User</i>	1,714	<i>Tag</i>	21,956
			<i>Bookmark</i>	30,998		
		Links	contact	15,329	has-tag	437,594
			post	437,594		
MovieLens	From Sep 1997 to Jan 2009	Nodes	<i>User</i>	1,421	<i>Genre</i>	19
			<i>Movie</i>	5,660	<i>Tag</i>	5,561
			<i>Actor</i>	6,176	<i>Country</i>	63
			<i>Director</i>	2,401		
		Links	rate	855,599	has-genre	20,810
			play-in	231,743	has-tag	47,958
			direct	10,156	produced-in	10,198

6.1.3 MovieLens. The third heterogeneous dataset with dynamic characteristics has been extracted from MovieLens personalized movie recommendation website, provided by [15]. The dataset comprises seven types of objects, that are users, movies, tags, genres, actors, directors, and countries, as illustrated by the network schema in Figure 1(c). It contains about 1.4k users and 5.6k movies, with user-movie rating timestamps ranging from September 1997 to January 2009.

The demographic statistics of all datasets are presented in Table 4.

6.2 Experiment Settings

6.2.1 Comparison Methods. To challenge the performance of NP-GLM, we use a number of baselines introduced in the following:

- *Generalized Linear Model (GLM)*: This is the state-of-the-art method proposed in [34]. We use the GLM-based framework with Exponential and Weibull distributions, denoted as EXP-GLM and WBL-GLM used in [34].
- *Censored Regression Model (CRM)*: This model, also called type II Tobit model, is designed to estimate linear relationships between variables when there is censoring in the dependent variable. In other words, it is an extension to the ordinary least squares linear regression for censored data [39]. The structural equation in this model is

$$t^* = \mathbf{w}^T \mathbf{x} + \epsilon,$$

where ϵ is a normally distributed error term and t^* is a latent variable, which is observed within the observation window and censored otherwise. Accordingly, the observed t is defined as

$$t = \begin{cases} t^* & \text{if } y = 1 \\ \Omega & \text{if } y = 0 \end{cases}.$$

The coefficient vector $\mathbf{w}$ is learned using maximum likelihood estimation (more details in [3]).

- *Additive Regression Model (ARM)*: This model is another regression method suggested by Aalen for censored data [1]. Like NP-GLM, it specifies the intensity function, but instead of

a multiplicative linear model, the Aalen's model is additive:

$$\lambda(t \mid \mathbf{x}) = \sum_{i=0}^d w_i(t)x_i.$$

The learning algorithm infers $\int_0^t w_i(t)dt$ instead of estimating individual w_i s. For more details about the learning algorithm, the reader can refer to [19].

For all models, we consider the median of the distribution $f_T(t \mid \mathbf{x}_{test})$ as the predicted time for any test sample and then compare it to the ground truth time t_{test} .

To examine the effect of considering different feature extractors on the performance of the models, we use another dynamic feature extractor and a static one against the proposed LSTM Autoencoder:

- *Exponential Smoothing*: This dynamic feature extractor previously used in [14] is an exponentially weighted moving average over the features extracted in all the snapshots of the network, which is calculated as

$$\mathbf{f}^i = \begin{cases} \mathbf{x}^1, & \text{if } i = 1 \\ \alpha \mathbf{x}^i + (1 - \alpha) \mathbf{f}^{i-1}, & \text{otherwise,} \end{cases}$$

where $\mathbf{f}^i$ is the smoothed feature after i th snapshot, $\mathbf{x}^i$ is the i th step of the dynamic meta-path-based time series, and $\alpha \in (0, 1)$ is the smoothing factor. We then set $\mathbf{x} = \mathbf{f}^k$ as the final feature vector if we have k snapshots in total.

- *Single Snapshot*: This static feature extractor considers the whole network as a single snapshot, neglecting its temporal dynamics. This feature extractor is equivalent to the one proposed in [34].

6.2.2 Performance Measures. We assess different methods using a number of evaluation metrics that are described in the following:

- *Mean Absolute Error (MAE)*: This metric measures the expected absolute error between the predicted time values and the ground truth:

$$MAE(\mathbf{t}, \hat{\mathbf{t}}) = \frac{1}{N} \sum_{i=1}^N |t_i - \hat{t}_i|.$$

- *Mean Relative Error (MRE)*: This metric calculates the expected relative absolute error between the predicted time values and the ground truth:

$$MRE(\mathbf{t}, \hat{\mathbf{t}}) = \frac{1}{N} \sum_{i=1}^N \left| \frac{t_i - \hat{t}_i}{t_i} \right|.$$

- *Root Mean Squared Error (RMSE)*: This metric computes the root of the expected squared error between the predicted time values and the ground truth:

$$RMSE(\mathbf{t}, \hat{\mathbf{t}}) = \sqrt{\frac{1}{N} \sum_{i=1}^N (t_i - \hat{t}_i)^2}.$$

- *Mean Squared Logarithmic Error (MSLE)*: This measures the expected value of the squared logarithmic error between the predicted time values and the ground truth:

$$MSLE(\mathbf{t}, \hat{\mathbf{t}}) = \frac{1}{N} \sum_{i=1}^N (\log(1 + t_i) - \log(1 + \hat{t}_i))^2.$$

- Median Absolute Error (MDAE): It is the median of the absolute errors between the predicted time values and the ground truth:

$$MDAE(\mathbf{t}, \hat{\mathbf{t}}) = \text{median}(|t_1 - \hat{t}_1| \dots |t_N - \hat{t}_N|).$$

- Maximum Threshold Prediction Accuracy (ACC): This measures for what fraction of samples, a model have a lower absolute error than a given threshold:

$$ACC(\mathbf{t}, \hat{\mathbf{t}}) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}(|t_i - \hat{t}_i| < \text{threshold}).$$

- Concordance Index (CI): This metric is one of the most widely used performance measures for survival models that estimates how good a model performs at ranking predicted times [16]. It can be seen as the fraction of all the sample pairs whose predicted timestamps are correctly ordered among all samples that can be ordered, and is considered as the generalization of the Area Under Receiver Operating Characteristic Curve (AUC) when we are dealing with censored data [32].

6.2.3 Experiment Setup. For DBLP dataset, we confine the data samples to those authors who have published more than five papers in the feature extraction window of each experiment. Following the triple building blocks described for feature extraction in Section 3, and using the similarity meta-paths in Table 1, we start the feature extraction process with 19 feature meta-paths. In all experiments, the author citation relation ($A \rightarrow P \rightarrow P \leftarrow A$) is chosen as the target relation. For the case of the Delicious dataset, we select user-user relation ($U \leftrightarrow U$) as the target relation, and design six feature meta-paths via the similarity meta-paths in Table 1. Regarding the MovieLens dataset, we limit the actor list to the top three for each movie. To imply a notion of “like” relation between user and movie, we only consider ratings above four in the scale of five. For this dataset, the target relation is set to user rate movie ($U \rightarrow M$), based on which, we design 11 final meta-paths. For the sake of convenience, we convert the scale of time differences from timestamp to month in Delicious and MovieLens datasets.

Except for parameter settings analysis (Section 6.3.3) where we will analyze the effect of different parameters on the performance of different models, in the rest of the experiments in this section we set the length of the observation window Ω to 6 for all three datasets. For DBLP dataset, the number of snapshots k is set to 6, while for the other two datasets we set $k = 12$. We also fix the time difference between network snapshots Δ to 1 in all cases. These settings lead to having a feature extraction window of size $\Phi = 6$ years for DBLP, and $\Phi = 12$ months for Delicious and MovieLens. Accordingly, the number of labeled instances for DBLP, Delicious, and MovieLens are about 3.4K, 3.9K, and 7.8K, respectively. About half of the labeled samples are censored ones, which are picked uniformly at random among all the possible candidates.

We implemented the LSTM autoencoder using Keras deep learning library [8]. We used mean square error loss function, linear activation function, and Adadelata optimizer [45] with default parameters. For all datasets, we set the dimension of the encoded feature as twice as the input dimension and trained the autoencoder in 50 epochs. For exponential smoothing feature extractor, the smoothing factor α were tuned to maximize the performance on the training dataset. For Np-Glm, the data samples were ordered according to their corresponding time variables, as the model needs the samples sorted by their recorded time. We use five-fold cross-validation and report the average results for all the experiments in this section.

6.3 Experiment Results

In the rest of this section, we first assess how well different methods perform over various datasets and compare their performance based on different measures. Next, we discuss the efficiency of our proposed method by measuring and comparing its running time against the other baselines. Finally, we analyze the effect of different parameters and problem configurations on the performance of competitive methods.

6.3.1 Comparative Performance Analysis. In the first set of experiments, we evaluate the prediction power of different models combined with different feature extractors on DBLP, Delicious and MovieLens datasets. MAE, MRE, RMSE, MSLE, MDAE, and CI of all models using both dynamic and static feature sets has been shown in Table 5. We see that in all three networks, NP-GLM with the LSTM Autoencoder feature set is superior to the other methods under all performance measures. For instance, our model NP-GLM can obtain an MAE of 1.99 for DBLP dataset, which is 15% lower than the MAE obtained by its closest competitor, WBL-GLM. As of CI, NP-GLM achieves 0.62 on DBLP, which is 7% better than WBL-GLM. On Delicious dataset, NP-GLM improves MAE and CI by 11% and 23%, respectively, relative to WBL-GLM. Similarly, NP-GLM reduces MAE by 19% and increases CI by 25%. Comparable results hold for other performance measures as well. Accordingly, WBL-GLM, which has two degrees of freedom, has shown a better performance compared to other models. That is while NP-GLM, as a non-parametric model with highly tunable shape, outperforms all the other “less-flexible” models by learning the true distribution of the data.

Moreover, it is evident from Table 5 that using the dynamic features learned with the LSTM autoencoder has boosted the performance of all models over different datasets, and has outperformed the other feature extractors. Based on the results presented in Table 5, the alternative dynamic feature extractor, exponential smoothing, has performed better than the static single snapshot feature extractor, yet not better than the proposed LSTM Autoencoder. Comparing the LSTM Autoencoder with exponential smoothing feature extractor, over the DBLP dataset, the proposed feature extractor has achieved 7% less MAE and 17% more CI with NP-GLM. Over Delicious, NP-GLM with LSTM-based features reduces MAE by about 7% and improves CI by 7%. Finally, on the MovieLens dataset, combining LSTM Autoencoder and NP-GLM leads to an improvement of 8% and 7% under MAE and CI, respectively. The other models behave more or less similarly when they are combined with different feature extractors. This result clearly demonstrates that our feature extraction framework is performing well on capturing the temporal dynamics of the networks.

In the next experiment, we investigated the performance of different models using the LSTM autoencoder feature extraction framework under maximum threshold prediction accuracy. To evaluate the prediction accuracy of a model, we record the fraction of test samples for which the difference between their true times and predicted ones are lower than a given threshold, called *tolerated error*. The results are plotted in Figure 9 where we varied the tolerated error in the range $\{0.5, 1.0, \dots, 3.0\}$. We can see from the figure that NP-GLM and WBL-GLM perform comparably, yet NP-GLM outperforms WBL-GLM in all cases. For example, on MovieLens dataset (Figure 9(c)), NP-GLM can predict the relationship building time of all the test samples with 100% accuracy by an error of 3 months, whereas for WBL-GLM, this is reduced to 90%. Similarly, on the Delicious dataset, NP-GLM with 3 months of tolerated error achieves around 80% accuracy, which is about 12% more than WBL-GLM.

6.3.2 Efficiency Analysis. In this part, we analyze and compare the running time of NP-GLM and WBL-GLM models utilizing different feature extractors, namely LSTM autoencoder, exponential smoothing, and single snapshot. All the algorithms were implemented in Python and were run on a Windows 10 PC with Intel Core i7 1.8 GHz CPU and 12GB of RAM. The full

Table 5. Comprehensive Performance Comparison of Different Methods

Dataset	Feature	Model	MAE	MRE	RMSE	MSLE	MDAE	CI
DBLP	LSTM Autoencoder (Dynamic)	NP-GLM	1.99	0.95	2.43	0.30	1.73	0.62
		WBL-GLM	2.33	1.10	2.85	0.36	2.08	0.58
		EXP-GLM	3.11	1.39	3.88	0.52	2.58	0.50
		CRM	3.08	1.06	3.32	2.04	2.98	0.37
		ARM	2.95	1.33	4.48	0.48	1.48	0.56
	Exp. Smoothing (Dynamic)	NP-GLM	2.15	1.07	2.54	0.32	1.98	0.53
		WBL-GLM	2.50	1.22	2.89	0.38	2.46	0.58
		EXP-GLM	3.20	1.49	3.73	0.51	3.06	0.45
		CRM	2.55	0.97	3.05	1.58	2.11	0.55
		ARM	6.75	2.83	7.86	1.17	6.39	0.60
	Single Snapshot (Static)	NP-GLM	2.76	1.35	3.07	0.44	2.88	0.50
		WBL-GLM	2.81	1.38	3.16	0.45	2.88	0.48
		EXP-GLM	3.28	1.57	3.70	0.53	3.30	0.14
		CRM	2.96	1.03	3.21	1.73	2.97	0.38
		ARM	3.89	1.76	5.45	0.66	2.11	0.46
Delicious	LSTM Autoencoder (Dynamic)	NP-GLM	2.10	1.20	2.55	0.35	2.05	0.70
		WBL-GLM	2.37	1.31	2.89	0.40	2.16	0.57
		EXP-GLM	3.21	1.58	3.84	0.54	2.89	0.55
		CRM	6.38	3.10	6.55	1.33	6.87	0.43
		ARM	5.20	2.56	6.23	0.86	4.99	0.52
	Exp. Smoothing (Dynamic)	NP-GLM	2.25	1.36	2.74	0.40	2.11	0.66
		WBL-GLM	2.61	1.64	3.20	0.47	2.17	0.56
		EXP-GLM	3.52	1.99	4.54	0.62	3.20	0.39
		CRM	3.28	3.69	3.84	2.07	2.88	0.43
		ARM	6.36	3.24	7.80	1.09	6.72	0.56
	Single Snapshot (Static)	NP-GLM	2.33	1.46	2.80	0.41	2.17	0.61
		WBL-GLM	2.65	1.62	3.23	0.47	2.26	0.43
		EXP-GLM	3.35	1.91	4.17	0.59	2.75	0.35
		CRM	3.06	2.05	3.47	1.53	2.84	0.38
		ARM	5.79	2.76	6.69	1.16	5.89	0.37
MovieLens	LSTM Autoencoder (Dynamic)	NP-GLM	2.48	3.08	3.04	0.55	2.14	0.70
		WBL-GLM	3.06	3.61	3.79	0.65	2.60	0.56
		EXP-GLM	3.79	2.70	4.60	0.78	3.48	0.45
		CRM	3.07	3.47	3.74	2.02	2.51	0.40
		ARM	5.53	5.63	7.41	1.12	3.80	0.53
	Exp. Smoothing (Dynamic)	NP-GLM	2.69	3.35	3.18	0.59	2.61	0.66
		WBL-GLM	3.09	3.62	3.59	0.66	2.95	0.52
		EXP-GLM	3.52	2.86	4.05	0.74	3.26	0.43
		CRM	3.18	3.37	3.68	1.90	2.56	0.48
		ARM	9.39	8.60	10.06	1.83	9.26	0.52
	Single Snapshot (Static)	NP-GLM	2.92	3.44	3.45	0.67	3.36	0.50
		WBL-GLM	2.99	3.52	3.51	0.69	3.37	0.49
		EXP-GLM	3.42	2.89	3.86	0.78	3.82	0.49
		CRM	3.14	3.48	3.63	2.20	3.55	0.35
		ARM	5.71	5.50	7.30	1.23	5.18	0.47

All bold values are statistically significant at p -value < 0.001 .

Fig. 9. Prediction accuracy of different methods versus the maximum tolerated absolute error on different datasets.

Table 6. Comparison of Computational Time Measured in Seconds

Dataset	Model	Feature extractor		
		Single snapshot	Exp. smoothing	LSTM autoencoder
DBLP	NP-GLM	35.92	98.35	93.59
	WBL-GLM	36.01	74.34	79.83
Delicious	NP-GLM	2.01	110.67	128.43
	WBL-GLM	1.89	97.44	123.53
MovieLens	NP-GLM	19.60	177.015	232.77
	WBL-GLM	19.70	154.86	213.7

specification of the host machine is reported in Table 3. We measured the running time of all the methods during a complete training and test procedure, including feature extraction, learning, and inference. For exponential smoothing feature extractor, we included the time required for tuning the smoothing factor α using a separate validation set, while for LSTM based framework the training time of the autoencoder is counted toward total running time. Table 6 presents the results over each of the DBLP, Delicious, and MovieLens datasets. Since a considerable amount of running time is spent on feature extraction, dynamic feature extraction frameworks require more time to process the network data as opposed to static single snapshot feature extractors. However, the proposed LSTM autoencoder performs comparably to exponential smoothing in terms of running time. Even though LSTM autoencoder is a bit slower than the other dynamic feature extractors, it demonstrates higher prediction performance compared to models utilizing exponential smoothing. For example, on MovieLens network with more than 20K nodes and 1 million links, NP-GLM with LSTM autoencoder requires less than 4 minutes to process the whole network, extract features and learn from about 6K labeled samples, and perform prediction for about 2K instances on a typical PC.

6.3.3 Parameter Setting Analysis. The performance of different models is influenced by two parameters, the number of snapshots k , and the time difference between snapshots Δ , as these parameters determine the length of the feature extraction window Φ . In this set of experiments, we investigate how these parameters affect the performance of our model NP-GLM and its closest competitor WBL-GLM over Delicious and MovieLens datasets using the proposed LSTM based feature extraction framework.

Firstly, the effect of increasing the number of snapshots on achieved MAE and CI by NP-GLM and WBL-GLM over Delicious and MovieLens datasets is illustrated in Figures 10 and 11, respectively.

Fig. 10. Effect of choosing different number of snapshots on performance of different methods using Delicious dataset.

Fig. 11. Effect of choosing different number of snapshots on performance of different methods using MovieLens dataset.

For both datasets, we set $\Delta = 1.5$ and $\Omega = 18$ and varied the number of snapshots in the range of 3–18. As we can see in both figures, increasing the number of snapshots results in lower prediction error and higher accuracy. This is due to the fact that as the number of snapshots grows, a longer history of the network is taken into account. Therefore, different models can benefit from more information about the temporal dynamics of the network given to them through the extracted feature vector.

Finally, the impact of choosing different values for Δ is analyzed on the performance of NP-GLM and WBL-GLM in terms of MAE and CI. The results for Delicious and MovieLens datasets are depicted in Figures 12 and 13, respectively. In this experiment, the number of snapshots and observation window length are accordingly set to 6 and 24. Different values of Δ are selected from the set $\{0.5, 1.0, \dots, 3.0\}$. As illustrated in both figures, by increasing Δ up to an extent, we witness that the performance of models improves gradually. That is because increasing the value of Δ leads to a wider feature extraction window. However, since the number of snapshots is constant, we see no performance improvement when the value of Δ becomes greater than a certain threshold. This is due to the fact that short-term temporal evolution of the network will be ignored when the value of Δ becomes too wide.

7 RELATED WORKS

The problem of link prediction has been studied extensively in recent years and many approaches have been proposed to solve this problem [41, 42]. Previous work on time-aware link prediction has mostly considered temporality in analyzing the long-term network trend over time [11].

Fig. 12. Effect of choosing different values for Δ on performance of different methods using Delicious dataset.

Fig. 13. Effect of choosing different values for Δ on performance of different methods using MovieLens dataset.

Authors in [28] have shown that temporal metrics are an extremely valuable new contribution to link prediction, and should be used in future applications. Dunlavy et al. focused on the problem of periodic temporal link prediction [13]. They concentrated on bipartite graphs that evolve over time and also considered a weighted matrix that contained multilayer data and tensor-based methods for predicting future links. Oyama et al. solved the problem of cross-temporal link prediction, in which the links among nodes in different time frames are inferred [26]. They mapped data objects in different time frames into a common low-dimensional latent feature space and identified the links on the basis of the distance between the data objects. Özcan et al. proposed a novel link prediction method for evolving networks based on NARX neural network [27]. They take the correlation between the quasi-local similarity measures and temporal evolutions of link occurrences information into account by using NARX for multivariate time series forecasting. Yu et al. developed a novel temporal matrix factorization model to explicitly represent the network as a function of time [44]. They provided results for link prediction as a specific example and showed that their model performs better than the state-of-the-art techniques.

The most relevant works to this study are available in [2, 14, 24, 30, 34]. The authors in [14] approach the problem of time series link prediction by extracting simple temporal features from the time series, such as mean, (weighted) moving average, and exponential smoothing besides some topological features like common neighbor and Adamic-Adar. But their method is designed for homogeneous networks and fail to consider the heterogeneity of modern networks. Aggarwal et al. [2] tackle the link prediction problem in both dynamic and heterogeneous information networks using a dynamic clustering approach alongside content-based and structural models. However, they aim to solve the conventional link prediction problem, not the continuous-time relationship

prediction studied in this article. In [30], the authors proposed a feature set, called TMLP, well suited for link prediction in dynamic and heterogeneous information networks. Although their proposed feature set copes with both dynamicity and heterogeneity of the network, it cannot be extended for the generalized problem of relationship prediction and is only designed for solving the simpler link prediction problem. Milani Fard et al. developed an approach called MetaDynaMix that utilizes a set of latent and topological features for predicting a target relationship between two nodes in a dynamic heterogeneous information network [24]. They combine meta-path-based topological features with inferred latent features that take temporal network evolutions into account, in order to capture both heterogeneity and dynamicity of the network.

Most of the aforementioned works answered the question of *whether* a link will appear in the network. To the best of our knowledge, the only work that has focused on the continuous-time relationship prediction problem is proposed by Sun et al. [34], in which a generalized linear model based framework is suggested to model the relationship building time. They consider the building time of links as independent random variables coming from a pre-specified distribution and model the expectation as a function of a linear predictor of the extracted topological features. A shortcoming of this model is that we need to exactly specify the underlying distribution of relationship building times. We came over this problem by learning the distribution from the data using a non-parametric solution. Furthermore, we considered the temporal dynamics of the network that has been entirely ignored in their work.

8 CONCLUSION

In this article, we studied the problem of continuous-time relationship prediction in both dynamic and heterogeneous information networks. To effectively tackle this problem, we first introduced a novel feature extraction framework based on meta-path modeling and recurrent neural network autoencoders to systematically extract features that take both the temporal dynamics and heterogeneous characteristics of the network into account for solving the continuous-time relationship problem. We then proposed a supervised non-parametric model, called NP-GLM, which exploits the extracted features to predict the relationship building time in information networks. The strength of our model is that it does not impose any significant assumptions on the underlying distribution of the relationship building time given its features, but tries to infer it from the data via a non-parametric approach. Extensive experiments conducted on a synthetic dataset and real-world datasets from DBLP, Delicious, and MovieLens demonstrated the correctness of our method and its effectiveness in predicting the relationship building time.

For future work, we would like to design a unified architecture to combine feature extraction step with the learning algorithm in an integrated deep learning framework. Moreover, although the proposed method is able to scale to large information networks with thousands of nodes, it is not currently extensible to web-scale information networks where the number of nodes is in the scale of hundreds of millions. Learning temporal non-parametric models within an extremely huge dataset is a challenging problem and is an interesting and important future work. As calculating meta-path-based features are the primary computational bottleneck of our method, to make the learning process scalable, we set to investigate node embedding and approximation techniques.

REFERENCES

- [1] Odd O. Aalen. 1989. A linear regression model for the analysis of life times. *Statistics in Medicine* 8, 8 (1989), 907–925.
- [2] Cham Aggarwal, Yan Xie, and Philip S. Yu. 2012. On dynamic link inference in heterogeneous networks. In *Proceedings of the 2012 SIAM International Conference on Data Mining*. SIAM, 415–426.
- [3] Takeshi Amemiya. 1984. Tobit models: A survey. *Journal of Econometrics* 24, 1–2 (1984), 3–61.
- [4] Yoshua Bengio. 2009. Learning deep architectures for AI. *Foundations and Trends® in Machine Learning* 2, 1 (2009), 1–127.

- [5] Norman E. Breslow. 1975. Analysis of survival data under the proportional hazards model. *International Statistical Review/Revue Internationale de Statistique* 43, 1 (1975), 45–57.
- [6] Iván Cantador, Peter Brusilovsky, and Tsvi Kuflik. 2011. Proceedings of the 2nd workshop on information heterogeneity and fusion in recommender systems (HetRec 2011). In *Proceedings of the 5th ACM Conference on Recommender Systems (RecSys'11)*. ACM, New York, NY.
- [7] Xing Chen, Ming-Xi Liu, and Gui-Ying Yan. 2012. Drug–target interaction prediction by random walk on the heterogeneous network. *Molecular BioSystems* 8, 7 (2012), 1970–1978.
- [8] François Chollet et al. 2015. Keras. Retrieved from <https://keras.io>.
- [9] Andrew M. Dai and Quoc V. Le. 2015. Semi-supervised sequence learning. In *Proceedings of Advances in Neural Information Processing Systems*. 3079–3087.
- [10] Darcy Davis, Ryan Lichtenwalter, and Nitesh V. Chawla. 2011. Multi-relational link prediction in heterogeneous information networks. In *Proceedings of the 2011 International Conference on Advances in Social Networks Analysis and Mining (ASONAM'11)*. IEEE, 281–288.
- [11] Yugchhaya Dhote, Nishchol Mishra, and Sanjeev Sharma. 2013. Survey and analysis of temporal link prediction in online social networks. In *Proceedings of the 2013 International Conference on Advances in Computing, Communications and Informatics (ICACCI'13)*. IEEE, 1178–1183.
- [12] Yuxiao Dong, Jie Tang, Sen Wu, Jilei Tian, Nitesh V. Chawla, Jinghai Rao, and Huanhuan Cao. 2012. Link prediction and recommendation across heterogeneous social networks. In *Proceedings of the 2012 IEEE 12th International Conference on Data Mining (ICDM'12)*. IEEE, 181–190.
- [13] Daniel M. Dunlavy, Tamara G. Kolda, and Evrim Acar. 2011. Temporal link prediction using matrix and tensor factorizations. *ACM Transactions on Knowledge Discovery from Data* 5, 2 (2011), 10.
- [14] Alireza Hajibagheri, Gita Sukthankar, and Kiran Lakkaraju. 2016. Leveraging network dynamics for improved link prediction. In *Proceedings of the 9th International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior, SBP-BRIMS 2016, Washington, DC, June 28–July 1, 2016*. Springer, 142–151.
- [15] F. Maxwell Harper and Joseph A. Konstan. 2015. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems* 5, 4, Article 19 (December 2015), 19 pages.
- [16] Frank E. Harrell Jr, Robert M. Califf, David B. Pryor, Kerry L. Lee, and Robert A. Rosati. 1982. Evaluating the yield of medical tests. *JAMA* 247, 18 (1982), 2543–2546.
- [17] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [18] Ellis Horowitz, Sartaj Sahni, and Susan Anderson-Freed. 1983. *Fundamentals of Data Structures*. Vol. 20. Computer Science Press Rockville, MD.
- [19] David W. Hosmer Jr, Stanley Lemeshow, and Susanne May. 2011. *Applied Survival Analysis: Regression Modeling of Time-to-Event Data*. Vol. 618. John Wiley & Sons.
- [20] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. 2010. What is Twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web*. ACM, 591–600.
- [21] David Liben-Nowell and Jon Kleinberg. 2007. The link prediction problem for social networks. *Journal of the Association for Information Science and Technology* 58, 7 (2007), 1019–1031.
- [22] Ryan N. Lichtenwalter, Jake T. Lussier, and Nitesh V. Chawla. 2010. New perspectives and methods in link prediction. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 243–252.
- [23] Linyuan Lü and Tao Zhou. 2011. Link prediction in complex networks: A survey. *Physica A: Statistical Mechanics and its Applications* 390, 6 (2011), 1150–1170.
- [24] Amin Milani Fard, Ebrahim Bagheri, and Ke Wang. 2019. Relationship prediction in dynamic heterogeneous information networks. In *Proceedings of the European Conference on Information Retrieval (ECIR'19)*. Springer, 12 pages.
- [25] Behnaz Moradabadi and Mohammad Reza Meybodi. 2017. A novel time series link prediction method: Learning automata approach. *Physica A: Statistical Mechanics and Its Applications* 482 (2017), 422–432.
- [26] Satoshi Oyama, Kohei Hayashi, and Hisashi Kashima. 2011. Cross-temporal link prediction. In *Proceedings of the 2011 IEEE 11th International Conference on Data Mining (ICDM'11)*. IEEE, 1188–1193.
- [27] Alper Özcan and Şule Gündüz Öğüdücü. 2016. Temporal link prediction using time series of quasi-local node similarity measures. In *Proceedings of the 2016 15th IEEE International Conference on Machine Learning and Applications (ICMLA'16)*. IEEE, 381–386.
- [28] Anet Potgieter, Kurt A. April, Richard J. E. Cooke, and Isaac O. Osunmakinde. 2009. Temporality in link prediction: Understanding social complexity. *Emergence: Complexity and Organization* 11, 1 (2009), 69.
- [29] S. Sajadmanesh, H. R. Rabiee, and A. Khodadadi. 2016. Predicting anchor links between heterogeneous social networks. In *Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM'16)*. 158–163.

- [30] Niladri Sett, Saptarshi Basu, Sukumar Nandi, and Sanasam Ranbir Singh. 2018. Temporal link prediction in multi-relational network. *World Wide Web* 21, 2 (2018), 395–419.
- [31] Chuan Shi, Yitong Li, Jiawei Zhang, Yizhou Sun, and S. Yu Philip. 2017. A survey of heterogeneous information network analysis. *IEEE Transactions on Knowledge and Data Engineering* 29, 1 (2017), 17–37.
- [32] Harald Steck, Balaji Krishnapuram, Cary Dehing-Oberije, Philippe Lambin, and Vikas C. Raykar. 2008. On ranking in survival analysis: Bounds on the concordance index. In *Proceedings of the 20th International Conference on Neural Information Processing Systems*. 1209–1216.
- [33] Yizhou Sun, Rick Barber, Manish Gupta, Charu C. Aggarwal, and Jiawei Han. 2011. Co-author relationship prediction in heterogeneous bibliographic networks. In *Proceedings of the 2011 International Conference on Advances in Social Networks Analysis and Mining (ASONAM'11)*. IEEE, 121–128.
- [34] Yizhou Sun, Jiawei Han, Charu C. Aggarwal, and Nitesh V. Chawla. 2012. When will it happen? Relationship prediction in heterogeneous information networks. In *Proceedings of the 5th ACM International Conference on Web Search and Data Mining (WSDM'12)*. ACM, New York, NY, 663–672.
- [35] Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S. Yu, and Tianyi Wu. 2011. Pathsim: Meta path-based top-k similarity search in heterogeneous information networks. *Proceedings of the VLDB Endowment* 4, 11 (2011), 992–1003.
- [36] Jie Tang, Jing Zhang, Limin Yao, Juanzi Li, Li Zhang, and Zhong Su. 2008. ArnetMiner: Extraction and mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'08)*. ACM, New York, NY, 990–998.
- [37] Ben Taskar, Ming-Fai Wong, Pieter Abbeel, and Daphne Koller. 2004. Link prediction in relational data. In *Proceedings of the 16th International Conference on Neural Information Processing Systems*. 659–666.
- [38] Ian W. Taylor, Rune Linding, David Warde-Farley, Yongmei Liu, Catia Pesquita, Daniel Faria, Shelley Bull, Tony Pawson, Quaid Morris, and Jeffrey L. Wrana. 2009. Dynamic modularity in protein interaction networks predicts breast cancer outcome. *Nature Biotechnology* 27, 2 (2009), 199.
- [39] James Tobin. 1958. Estimation of relationships for limited dependent variables. *Econometrica* 26, 1 (1958), 24–36. Retrieved from <http://www.jstor.org/stable/1907382>
- [40] Chao Wang, Venu Satuluri, and Srinivasan Parthasarathy. 2007. Local probabilistic models for link prediction. In *Proceedings of the 7th IEEE International Conference on Data Mining (ICDM'07)*. IEEE, 322–331.
- [41] Peng Wang, BaoWen Xu, YuRong Wu, and XiaoYu Zhou. 2015. Link prediction in social networks: The state-of-the-art. *Science China Information Sciences* 58, 1 (2015), 1–38.
- [42] Tingli Wang and Guoqiong Liao. 2014. A review of link prediction in social networks. In *Proceedings of the 2014 International Conference on Management of e-Commerce and e-Government (ICMeCG'14)*. IEEE, 147–150.
- [43] Stanley Wasserman and Katherine Faust. 1994. *Social Network Analysis: Methods and Applications*. Vol. 8. Cambridge University Press.
- [44] Wenchao Yu, Charu C Aggarwal, and Wei Wang. 2017. Temporally factorized network modeling for evolutionary network analysis. In *Proceedings of the 10th ACM International Conference on Web Search and Data Mining*. ACM, 455–464.
- [45] Matthew D. Zeiler. 2012. ADADELTA: An adaptive learning rate method. arXiv:1212.5701.
- [46] Jiawei Zhang, Philip S. Yu, and Zhi-Hua Zhou. 2014. Meta-path based multi-network collective link prediction. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1286–1295.

Received May 2018; revised March 2019; accepted May 2019