

CAPER: Enhancing Career Trajectory Prediction using Temporal Knowledge Graph and Ternary Relationship

Yeon-Chang Lee
Ulsan National Institute of Science
and Technology (UNIST)
Ulsan, Korea
yeonchang@unist.ac.kr

JaeHyun Lee
Hanyang University
Seoul, Korea
dldja4@hanyang.ac.kr

Michiharu Yamashita
The Pennsylvania State University
University Park, PA, USA
michiharu@psu.edu

Dongwon Lee
The Pennsylvania State University
University Park, PA, USA
dongwon@psu.edu

Sang-Wook Kim*
Hanyang University
Seoul, Korea
wook@hanyang.ac.kr

Abstract

The problem of *career trajectory prediction* (CTP) aims to predict one's future employer or job position. While several CTP methods have been developed for this problem, we posit that none of these methods (1) jointly considers the mutual ternary dependency between three key units (*i.e.*, user, position, and company) of a career and (2) captures the characteristic shifts of key units in career over time, leading to an inaccurate understanding of the job movement patterns in the labor market. To address the above challenges, we propose a novel solution, named as CAPER, that solves the challenges via sophisticated temporal knowledge graph (TKG) modeling. It enables the utilization of a graph-structured knowledge base with rich expressiveness, effectively preserving the changes in job movement patterns. Furthermore, we devise an extrapolated career reasoning task on TKG for a realistic evaluation. The experiments on a real-world career trajectory dataset demonstrate that CAPER consistently and significantly outperforms four baselines, two recent TKG reasoning methods, and five state-of-the-art CTP methods in predicting one's future companies and positions—*i.e.*, on average, yielding 6.80% and 34.58% more accurate predictions, respectively. The codebase of CAPER is available at <https://github.com/Bigdasgit/CAPER>.

CCS Concepts

• Information systems → Data mining.

Keywords

career trajectory prediction; job market; temporal knowledge graph

ACM Reference Format:

Yeon-Chang Lee, JaeHyun Lee, Michiharu Yamashita, Dongwon Lee, and Sang-Wook Kim. 2025. CAPER: Enhancing Career Trajectory Prediction using Temporal Knowledge Graph and Ternary Relationship. In *Proceedings of*

*Corresponding author.



This work is licensed under a Creative Commons Attribution-NonCommercial International 4.0 License.

KDD '25, Toronto, ON, Canada

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1245-6/25/08

<https://doi.org/10.1145/3690624.3709329>

the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1 (KDD '25), August 3–7, 2025, Toronto, ON, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3690624.3709329>

1 Introduction

Background. With the advent of online professional network platforms (*e.g.*, LinkedIn and Indeed), there has been a surge of research efforts in understanding and improving diverse tasks in the labor market [4, 5, 22, 24, 38, 75, 76], such as talent recruitment [1, 42, 44], talent demand/supply prediction [13, 79, 82], and career trajectory prediction [27, 36, 60, 77]. With the fast-evolving job market and its immediate benefits to both employees and employers, in particular, the problem of **Career Trajectory Prediction (CTP)** (*i.e.*, predict one's future employer or job position) has increasingly attracted attention in recent years [27, 36, 60, 77], especially when the life-long career movement has become a new norm (*e.g.*, according to the U.S. Bureau of Labor Statistics, people in the U.S. hold an average of 8.6 jobs from ages 18 through 34¹). Effective solutions to the CTP problem can bring significant benefits to various groups, including governments, companies, and job seekers. Governments can devise better policies through an in-depth understanding of the labor market, while companies can refine their recruitment and retention strategies. Also, job seekers can design their own career paths more effectively by identifying promising future career opportunities.

Motivated by such wide applicability, several CTP methods have been developed recently [27, 36, 60, 77]. In a nutshell, they were designed based on the assumption that users tend to have common patterns in career movements. Specifically, they divide each user's career trajectory into his/her transition sequences of companies and positions. Then, they capture the frequent transition patterns over users with respect to companies and positions via graph recurrent neural networks (GRNN) models [49] (*e.g.*, long short-term memory (LSTM), gated recurrent unit (GRU), and Transformer [59]). Finally, these methods predict a user's next career movement, including a future company and a position, based on the learned GRNN model.

Challenges. However, in this paper, we explore two challenges that have been overlooked in existing studies: (**C1: Career Modeling**) a mutual ternary dependency exists among *three* key units of a

¹<https://www.bls.gov/opub/ted/2022/people-born-in-early-1980s-held-an-average-of-8-6-jobs-from-ages-18-through-34.htm>




Figure 1: An extrapolated career reasoning task on the input temporal knowledge graph consisting of career trajectories.

“career”—i.e., (*user, position, company, timestamp*); (C2: **Temporal Dynamics**) the characteristics per unit may change over time.

Regarding (C1), note that each career represents an intense interaction between a user, a position, and a company, which are highly related with each other. For instance, a current career (Bob, back-end engineer, Amazon, 2023) can be regarded as an output of jointly considering Bob’s past experience (e.g., computer science), Amazon’s industrial characteristics (e.g., technology company), and the back-end engineer’s required role (e.g., distributed programming). Meanwhile, only one method [77] attempted to directly model mutual dependency between *two units*, i.e., company and position. However, omitting users in modeling careers is not a good idea since users often play a key role in career decisions [12].

Regarding (C2), it is important to consider the *timestamp*, in which each career occurred, to address the dynamic nature of the labor market. Since each user’s experience evolves over time, his/her different careers at different timestamps should be regarded as the consequence of the change in his/her characteristic. Suppose that Bob started his career at Meta as a back-end developer intern in 2018 and then moved to Amazon as a back-end engineer in 2019. This suggests that Bob’s internship experience at Meta influenced (partially) his move to Amazon. Furthermore, each company’s characteristics can also alter over time due to various societal/economic issues (e.g., changes in industry trends). For instance, Meta started as a social-media-service business in 2004, but since 2021, its main focus has (reportedly) shifted to the metaverse. However, existing methods [27, 36, 60, 77] use the temporal dimension to only identify each user’s career sequence, ignoring the characteristic shifts of key units in career over time.

Proposed Ideas. To address the above challenges, we propose a novel solution, named as **CAPER** (CAreer trajectory Prediction approach based on tEmporal knowledge gRaph), which jointly learns the mutual dependency for the ternary relationships and the characteristic shifts over time. It consists of three key modules:

- **Modeling Career Trajectories (for (C1) and (C2)):** We model the career trajectories as a form of a knowledge graph (KG), preserving the ternary relationships on career information. To add the temporal dimension, we span the KG as a sequence of temporal KG (TKG) snapshots with timestamps, each snapshot of which contains the careers that co-occur at a timestamp.
- **Learning Mutual Dependency (for (C1)):** Given each KG snapshot, we obtain users, positions, and companies as low-dimensional vectors (i.e., embeddings) by aggregating their career information in the input KG snapshot. Here, we represent the same user and company at different timestamps as different temporal embeddings, while representing the position as a unified

embedding regardless of timestamps. We assume that, unlike the user and company units, the position unit has static properties, which will be discussed in Sections 4 and 5.

- **Learning Temporal Dynamics (for (C2)):** Based on temporal embeddings, we capture the changes in the characteristics of users and companies between previous and current timestamps.

Finally, we learn the embeddings of users, companies, and positions so that the likelihood of each career (i.e., ground truth) at each timestamp is maximized. We show experimentally that (1) our modules successfully solve the challenges of existing methods and (2) CAPER substantially outperforms the state-of-the-art CTP methods, including NEMO [27], HCPNN [36], AHEAD [77], TACTP [60], and NAOMI [69]. For a realistic evaluation, we devise an additional extrapolation reasoning task (see Figure 1) that aims to forecast the next careers that are likely to occur at the non-existent future timestamps in the input TKG, which will be elaborated in Section 4.

Contributions. Our contributions are as follows:

- **New Perspective:** We model the career trajectories of users as a TKG and then formulate the CTP problem as the **extrapolated reasoning task on TKG**. We are the first who leverages the rich expressiveness of a graph-structured knowledge base for ternary relationships to address the CTP problem.
- **Effective CTP Approach:** CAPER jointly learns the mutual dependency for the ternary relationships and the characteristic shifts over time via **TKG modeling**, thereby successfully addressing the challenges of existing methods.
- **Comprehensive Validation:** Through extensive experiments using a real-world career trajectory dataset, we show that CAPER consistently and significantly outperforms the state-of-the-art CTP methods, yielding **6.80% and 34.58% more-accurate company and position prediction on average**, respectively, compared to the best competitors.

2 Related Work

CTP Methods. Career trajectory prediction (CTP) is an essential task of predicting one’s next position and company [41, 55]. Li et al. [27] proposed a way to encode different types of entities (e.g., employees, skills, and companies) by integrating the profile context as well as career path dynamics. Meng et al. [36] leveraged a hierarchical neural network with the attention mechanism to jointly characterize internal and external career transition. Wang et al. [60] devised a unified time-aware framework that predicts not only the next company and position but also one’s transition timing. Zhang et al. [77] utilized the mutual dependency for pairwise relationships of position and company via heterogeneous graph modeling. Lastly, [69] predicts a future career pathway via the Transformer mechanism [59] and neural collaborative reasoning by jointly utilizing graph and BERT [47] embeddings. In contrast to these previous works, our work proposes a TKG-enhanced CTP approach that can jointly learn the mutual dependency for ternary relationships and the characteristic shifts over time. The main differences between the existing CTP methods and CAPER will be elaborated in Section 4.4.

Meanwhile, several attempts [2, 3, 11, 42, 73] have been made to model the matching between a job posting and a resume, which can be applied to the job or talent recommendation problems. However,

Table 1: Statistics of the career trajectory dataset

	All	Training	Test
Users	333,399	333,399	183,575
Companies	9,000	9,000	8,510
Positions	425	425	421
Careers	8,218,162	6,453,546	1,764,616
Avg. # Careers per User	24.64	19.35	9.61
Period	1968 ~ 2020	1968 ~ 2015	2016 ~ 2020
Granularity	year	year	year

as discussed in [41], we believe that these problems have different goals from that of the CTP problem, which requires an understanding of the underlying job movement patterns in the labor market.

TKG Embedding Methods. Temporal knowledge graph embedding (TKGE) methods aim to represent entities and relations in a given TKG into low-dimensional embeddings, which preserve the structural properties in the TKG. According to [40, 53], recent TKGE methods were designed for the purpose of addressing an interpolation task [9, 17, 23, 25, 37, 64, 80] or an extrapolation task [6, 14, 31, 32, 40, 53, 81]. Our extrapolated career reasoning task is inspired by the extrapolation task in the TKGE research. However, they are inappropriate for the career trajectory dataset, failing to fully leverage the intrinsic characteristics of the job market.

3 Preliminaries

3.1 Data Description

In this study, we obtained a real-world career trajectory dataset from a global career platform, FutureFit AI². From this dataset, we randomly extracted resumes (i.e., users) with at least five valid career transitions in the U.S. from 1968 to 2020. In total, we obtained 333K+ valid resumes. Each resume includes a user-related information (e.g., user ID, demographics, and education history), career-transition information (e.g., companies and positions), and timestamps (i.e., starting and ending date of each work experience). Any private elements in resumes were pre-anonymized. Table 1 summarizes statistics about our final dataset after pre-processing. The pre-processing procedures, additional statistics, and a sharing statement are presented in the Appendix.

3.2 The Problem Formulation: CTP

Let $\mathcal{U} = \{u_1, u_2, \dots, u_i\}$, $\mathcal{P} = \{p_1, p_2, \dots, p_j\}$, $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$, and $\mathcal{T} = \{t_1, t_2, \dots, t_l\}$ denote sets of i users, j positions, k companies, and l timestamps, respectively. In our dataset, the granularity of a timestamp is the year. Also, we denote a set of career trajectories across all users as $\mathcal{D} = \{\mathcal{D}_{u_1}, \mathcal{D}_{u_2}, \dots, \mathcal{D}_{u_i}\}$. Each user u_x 's trajectory is denoted as $\mathcal{D}_{u_x} = \{(u_x, p_{(x,1)}, c_{(x,1)}, t_{(x,1)}), (u_x, p_{(x,2)}, c_{(x,2)}, t_{(x,2)}), \dots, (u_x, p_{(x,s_x)}, c_{(x,s_x)}, t_{(x,s_x)})\}$, where $(u_x, p_{(x,n)}, c_{(x,n)}, t_{(x,n)})$ represents the n -th career of u_x , i.e., u_x worked in a company $c_{(x,n)}$ with a position $p_{(x,n)}$ at a timestamp $t_{(x,n)}$, and s_x indicates the total length of the working interval for u_x (i.e., the difference between the first and the last timestamps in \mathcal{D}_{u_x}). Then, our problem can be defined as follows [69]:

²<https://www.futurefit.ai>

Table 2: Notations used in this paper

Notation	Description
$\mathcal{U}, \mathcal{P}, \mathcal{C}, \mathcal{T}$	Sets of users, positions, companies, and timestamps
u_x, p_y, c_z, t_m	User (i.e., subject entity), position (i.e., relation), company (i.e., object entity), and timestamp
$\mathcal{D}, \mathcal{D}_{u_x}$	Set of career trajectories and u_x 's career trajectory
$p_{(x,n)}, c_{(x,n)}$	n -th position and company of u_x in \mathcal{D}_{u_x}
$\mathcal{G}, \mathcal{G}_{t_m}$	TKG and KG snapshot w.r.t t_m
$\mathbf{u}_x(t_m), \mathbf{c}_z(t_m)$	Temporal embeddings of u_x and c_z w.r.t t_m
$\tilde{\mathbf{u}}_x(t_m), \tilde{\mathbf{c}}_z(t_m)$	Evolution embeddings of u_x and c_z w.r.t t_m
\mathbf{p}_y	Embedding of p_y
L	Number of GCN Layers
$N_{u_x}(\mathcal{G}_{t_m})$	Set of pairs of a company $c_{z'}$ and a position $p_{y'}$ included in the careers, which include u_x , at t_m
$N_{c_z}(\mathcal{G}_{t_m})$	Set of pairs of a user $u_{x'}$ and a position $p_{y'}$ included in the careers, which include c_z , at t_m

PROBLEM 1 (CAREER TRAJECTORY PREDICTION (CTP)). Given a query user u_x and her/his career trajectory \mathcal{D}_{u_x} , predict the next M careers of u_x , including M companies $\{c_{(x,s_x+1)}, c_{(x,s_x+2)}, \dots, c_{(x,s_x+M)}\}$ and M positions $\{p_{(x,s_x+1)}, p_{(x,s_x+2)}, \dots, p_{(x,s_x+M)}\}$.

Next, we formulate the CTP problem as an extrapolated career reasoning task on a temporal knowledge graph (TKG). Let $\mathcal{G} = \{\mathcal{G}_{t_1}, \mathcal{G}_{t_2}, \dots, \mathcal{G}_{t_l}\}$ denotes a TKG [15, 57] that consists of a sequence of knowledge graph (KG) snapshots with timestamps based on career trajectories. Each KG snapshot $\mathcal{G}_{t_m} = \{(u_x, p_y, c_z, t_m)\}$ at t_m is a bidirected multi-relational graph and stores the careers that co-occur at t_m . For a career (u_x, p_y, c_z, t_m) , in a KG snapshot \mathcal{G}_{t_m} , the user u_x (i.e., subject entity) and the company c_z (i.e., object entity) are represented as nodes, and the position p_y (i.e., relation) is represented as a label of an (undirected) edge between u_x and c_z (see Figure 1). Such a TKG modeling to career trajectories will be more elaborated in Section 4.2. Then, the CTP problem can be restated to the extrapolated career reasoning task as follows:

PROBLEM 2 (EXTRAPOLATED CAREER REASONING). Given a TKG $\mathcal{G} = \{\mathcal{G}_{t_1}, \mathcal{G}_{t_2}, \dots, \mathcal{G}_{t_l}\}$ and a query user u_x , predict the unknown (next) careers of u_x that are likely to occur at future M timestamps $\{t_{(l+1)}, t_{(l+2)}, \dots, t_{(l+M)}\}$ in the input TKG \mathcal{G} .

Table 2 summarizes a list of notations used in this paper.

4 The Proposed Approach: CAPER

4.1 Overview

In this section, we present the overall procedure of CAPER. As shown in Figure 2, CAPER performs the following three key modules: (M1) modeling career trajectories, (M2) learning mutual dependency, and (M3) learning temporal dynamics.

Given a set \mathcal{D} of career trajectories across all users, in (M1), CAPER models them as a sequence of KG snapshots with timestamps, $\mathcal{G} = \{\mathcal{G}_{t_1}, \mathcal{G}_{t_2}, \dots, \mathcal{G}_{t_l}\}$. Then, CAPER represents the entities (i.e., users and companies) and relations (i.e., positions) in \mathcal{G} as d -dimensional embeddings. Specifically, it constructs the same user u_x and company c_z at different timestamps t_m as different temporal embeddings $\mathbf{u}_x(t_m)$ and $\mathbf{c}_z(t_m)$, respectively, while constructing the same position p_y as a single unified embedding \mathbf{p}_y regardless

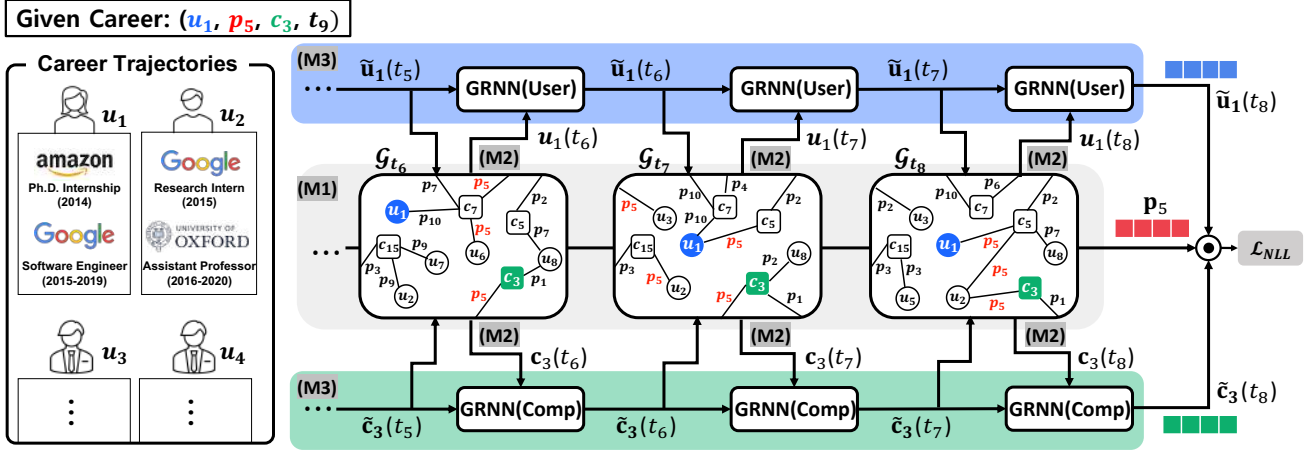


Figure 2: Overview of CAPER, which consists of three key modules: (M1) modeling career trajectories, (M2) learning mutual dependency, and (M3) learning temporal dynamics. Given a career (u_1, p_5, c_3, t_9) , CAPER infers the corresponding career’s likelihood using the temporal and evolution embeddings of u_1 and c_3 before t_9 and the embedding of p_5 .

of timestamps. The intuition behind this design choice is to differentiate the dynamic natures of users and companies from the static nature of positions, which will be elaborated in Section 4.2.

Next, CAPER repeatedly performs (M2) and (M3) whenever each KG snapshot in \mathcal{G} is given sequentially. Given a KG snapshot \mathcal{G}_{t_m} , in (M2), CAPER updates the temporal embedding $u_x(t_m)$ of u_x with respect to t_m so that the embedding can preserve the mutual dependencies on his/her career; in the same manner, it updates $c_z(t_m)$ of c_z with respect to t_m . Using $u_x(t_m)$, in (M3), CAPER obtains the evolution embedding $\tilde{u}_x(t_m)$ of u_x with respect to t_m , which captures the characteristic shift of u_x between the previous and current timestamps t_{m-1} and t_m ; in the same manner, it obtains $\tilde{c}_z(t_m)$ of c_z with respect to t_m using $c_z(t_m)$. Then, $\tilde{u}_x(t_m)$ and $\tilde{c}_z(t_m)$ with respect to t_m are used to initialize the temporal embeddings $u_x(t_{m+1})$ and $c_z(t_{m+1})$ with respect to t_{m+1} .

Lastly, given a query career (u_x, p_y, c_z, t_{m+1}) in $\mathcal{G}_{t_{m+1}}$, CAPER learns the temporal and evolution embeddings of u_x and c_z before t_{m+1} and the embedding of p_y via the negative log-likelihood (NLL) loss so that the corresponding career’s likelihood can be maximized.

4.2 Key Modules

(M1) Modeling Career Trajectories. This module aims to model the given career trajectories as a graph structure that can simultaneously represent the mutual dependency for the ternary relationships on a user, a position, and a company (i.e., (C1)) and the changes in characteristics of users and companies over time (i.e., (C2)).

To address (C1), we first construct the career trajectories as a KG, where the users u_x and companies c_z are represented as nodes, and the positions p_y are represented as labels of the edges between u_x and c_z . Suppose that there are the following career trajectories of two users u_1 and u_2 (see Figure 3-(a)): $\mathcal{D}_{u_1} = \{(u_1, \text{Ph.D. internship, Amazon, 2014}), (u_1, \text{software engineer, Google, 2015}), \dots, (u_1, \text{senior develop engineer, Meta, 2020})\}$ and $\mathcal{D}_{u_2} = \{(u_2, \text{research intern, Google, 2015}), (u_2, \text{assistant professor, University of Oxford, 2016}), \dots, (u_2, \text{assistant professor, University of Oxford, 2020})\}$. Figure 3-(b) shows a KG

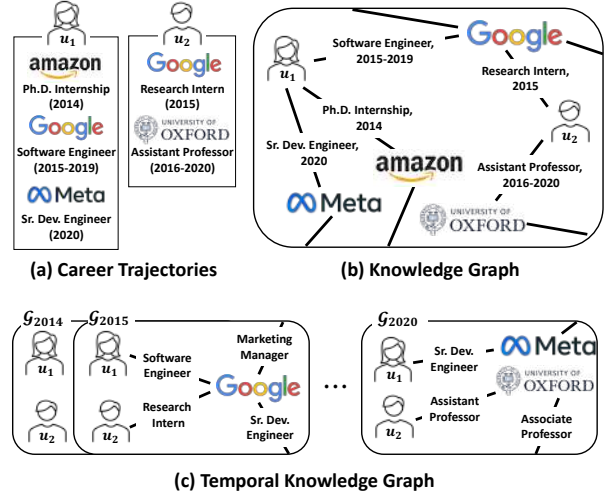


Figure 3: A toy example of modeling career trajectories.

constructed based on the career trajectories. By doing so, we can successfully preserve the intense interaction on three highly-correlated units, i.e., a user, a company, and a position, which form each career. The built KG only contains the actual careers of the users.

To address (C2), we then span the input KG as a sequence of KG snapshots with timestamps, each of which contains the careers that co-occur at a specific timestamp. Figure 3-(c) shows KG snapshots with timestamps constructed by the aforementioned career trajectories of two users. For example, \mathcal{G}_{2015} includes both careers $(u_1, \text{software engineer, Google, 2015})$ and $(u_2, \text{research intern, Google, 2015})$ that co-occur at 2015. By doing so, we can separately preserve the mutual dependency on careers that co-occur at a specific timestamp, capturing characteristic shifts over time. To the best of our knowledge, none of the existing methods consider the fact that the characteristics of some units can alter over time.

(M2) Learning Mutual Dependency. This module aims to reflect the mutual dependency on careers that co-occur at t_m into temporal embeddings of users and companies with respect to t_m . As mentioned above, the intuition behind this design choice is that each career reveals a high correlation between a user, a position, and a company in the corresponding career.

Therefore, we design a timestamp-aware graph convolutional networks (GCN) that learns the mutual dependency on careers at a timestamp. CAPER obtains each user u_x 's (resp. each company c_z 's) temporal embedding $\mathbf{u}_x(t_m)$ (resp. $\mathbf{c}_z(t_m)$) with respect to t_m by aggregating career information of u_x (resp. c_z) at t_m :

$$\mathbf{u}_x(t_m)^l = \mathbf{u}_x(t_m)^{l-1} + \sum_{(c_{z'}, p_{y'}) \in \mathcal{N}_{u_x}(\mathcal{G}_{t_m})} \frac{1}{c} (\mathbf{c}_{z'}(t_m)^{l-1} \circ \mathbf{p}_{y'}), \quad (1)$$

$$\mathbf{c}_z(t_m)^l = \mathbf{c}_z(t_m)^{l-1} + \sum_{(u_{x'}, p_{y'}) \in \mathcal{N}_{c_z}(\mathcal{G}_{t_m})} \frac{1}{c} (\mathbf{u}_{x'}(t_m)^{l-1} \circ \mathbf{p}_{y'}), \quad (2)$$

where $\mathbf{u}_x(t_m)^0$, $\mathbf{c}_z(t_m)^0$, $\mathbf{u}_{x'}(t_m)^0$, $\mathbf{c}_{z'}(t_m)^0$, and $\mathbf{p}_{y'}$ are randomly initialized without using any of their handcrafted features. In addition, \circ and $l \in \{1, \dots, L\}$ represent the element-wise product and the l -th GCN layer, respectively. Also, $\mathcal{N}_{u_x}(\mathcal{G}_{t_m})$ represents a set of company-position pairs $(c_{z'}, p_{y'})$ included in u_x 's careers at t_m . Similarly, $\mathcal{N}_{c_z}(\mathcal{G}_{t_m})$ represents a set of user-position pairs $(u_{x'}, p_{y'})$ included in c_z 's careers at t_m . In addition, c is a normalization constant, which is normally defined by the degree of nodes in GCNs [15, 18, 20, 21, 50]. Finally, CAPER considers $\mathbf{u}_x(t_m)^L$ (resp. $\mathbf{c}_z(t_m)^L$) obtained from the L -th GCN layer as u_x 's (resp. c_z 's) final temporal embedding $\mathbf{u}_x(t_m)$ (resp. $\mathbf{c}_z(t_m)$) with respect to t_m .

Here, we note that, unlike companies, users have few careers at each timestamp. On our real-world dataset, we confirmed that in each KG snapshot \mathcal{G}_{t_m} , each company c_z is included in 29.78 careers on average, while each user only has 2.37 careers on average.³ In this case, accordingly, there remains a question of whether each user's temporal embedding with respect to each timestamp is informative. Basically, we admit that each user's temporal embedding with respect to the timestamp, where he/she had the first career, may not be that informative. However, we believe that like a user's experience evolves while increasing his/her working interval over time, it is a natural way that the user's temporal embeddings also evolve over time gradually and thus become more informative.

Furthermore, recall that each position is modeled as a label of an edge between two entities of a user and a company in the input TKG. Therefore, following the existing TKG-based GCN mechanism (e.g., [25, 32]), we only use each position embedding as the aggregated information when obtaining the temporal embeddings of users and companies included in the corresponding position's careers.

(M3) Learning Temporal Dynamics. This module aims to capture the characteristic shifts of users and companies between the previous timestamp t_{m-1} and the current timestamp t_m . The intuition behind this design choice is that each user's experience evolves over time and each company's characteristics can also change over

time [16, 19, 74]. On the other hand, we assume that most positions have static properties so that their characteristics do not change significantly even over time.⁴ That's why we do not consider characteristic shifts of positions in (M3). We experimentally show in Section 5 that not considering the characteristic shifts of positions is more effective than considering them.

We design a GRNN that learns the evolution patterns of users and companies over time. Specifically, CAPER obtains each user u_x 's (resp. each company c_z 's) evolution embedding $\tilde{\mathbf{u}}_x(t_m)$ (resp. $\tilde{\mathbf{c}}_z(t_m)$) with respect to t_m as follows:

$$\tilde{\mathbf{u}}_x(t_m) = \text{GRNN}(\tilde{\mathbf{u}}_x(t_{m-1}), \mathbf{u}_x(t_m)), \quad (3)$$

$$\tilde{\mathbf{c}}_z(t_m) = \text{GRNN}(\tilde{\mathbf{c}}_z(t_{m-1}), \mathbf{c}_z(t_m)), \quad (4)$$

where GRNN indicates a recurrent function; CAPER can employ any functions, such as recurrent neural network (RNN), gated recurrent unit (GRU), and long short-term memory (LSTM). Also, $\tilde{\mathbf{u}}_x(t_{m-1})$ and $\tilde{\mathbf{c}}_z(t_{m-1})$ represents the evolution embedding of u_x and c_z , respectively, with respect to the previous timestamp t_{m-1} . Here, $\tilde{\mathbf{u}}_x(t_{(u_x,0)})$ with respect to $t_{(u_x,0)}$ (i.e., a timestamp right before u_x 's first timestamp) is randomly initialized. In the same manner, $\tilde{\mathbf{c}}_z(t_{(c_z,0)})$ with respect to $t_{(c_z,0)}$ is randomly initialized as well. See the Appendix for a detailed implementation of GRNN.

4.3 Training and Inference

Training. Given each career (u_x, p_y, c_z, t_{m+1}) in $\mathcal{G}_{t_{m+1}}$, CAPER obtains evolution embedding $\tilde{\mathbf{u}}_x(t_m)$ of u_x with respect to t_m , evolution embedding $\tilde{\mathbf{c}}_z(t_m)$ of c_z with respect to t_m , and embedding \mathbf{p}_y of p_y . Then, the likelihood between u_x and c_z and that between u_x and p_y are defined by the following two score functions:

$$\begin{aligned} \phi(u_x, c_z, t_{m+1}) &= \tilde{\mathbf{u}}_x(t_m) \cdot \tilde{\mathbf{c}}_z(t_m), \\ \phi(u_x, p_y, t_{m+1}) &= \tilde{\mathbf{u}}_x(t_m) \cdot \mathbf{p}_y, \end{aligned} \quad (5)$$

where " \cdot " represents the dot product.

CAPER learns the embeddings of users, positions, and companies by optimizing the following negative log-likelihood (NLL) loss:

$$\mathcal{L} = - \sum_{m=1}^{|T|} \sum_{(u_x, p_y, c_z, t_m) \in \mathcal{G}_{t_m}} \log p(c_z | u_x, t_m) + \log p(p_y | u_x, t_m), \quad (6)$$

$$\begin{aligned} p(c_z | u_x, t_m) &= \frac{\exp(\phi(u_x, c_z, t_m))}{\sum_{c_{z'}} \exp(\phi(u_x, c_{z'}, t_m))}, \\ p(p_y | u_x, t_m) &= \frac{\exp(\phi(u_x, p_y, t_m))}{\sum_{p_{y'}} \exp(\phi(u_x, p_{y'}, t_m))}, \end{aligned} \quad (7)$$

where $p(c_z | u_x, t_m)$ and $p(p_y | u_x, t_m)$ represent the likelihood between u_x and c_z and that between u_x and p_y with respect to t_m , respectively. Also, $c_{z'}$ and $p_{y'}$ represent the companies and positions, respectively, which have never been included in u_x 's careers before t_m . Intuitively, the goal of the NLL loss is to achieve a high likelihood for each existent career and a low likelihood for each non-existent career. By doing so, CAPER can accurately preserve the mutual dependency of units and their characteristic shifts over time in the embedding space.

³As shown in Figure 2, when a user u_1 's total length of the working interval is 3 and he/she has four careers for 3 years (i.e., $\mathcal{D}_{u_1} = \{(u_1, p_{10}, c_7, t_6), (u_1, p_{10}, c_7, t_7), (u_1, p_5, c_5, t_7), (u_1, p_5, c_5, t_8)\}$), the average number of u_1 's careers per timestamp becomes 1.33 (spec., $(1 \text{ (for } t_6) + 2 \text{ (for } t_7) + 1 \text{ (for } t_8)) / 3$).

⁴One can still argue that the characteristics of some positions may change significantly over time, e.g., due to advances in technology, the role of data scientists is expanding to a wide variety of fields. However, it is not trivial to identify such individual positions. Thus, we leave it as future work.

Multi-Step Inference. Given a TKG $\mathcal{G} = \{\mathcal{G}_{t_1}, \mathcal{G}_{t_2}, \dots, \mathcal{G}_{t_l}\}$, a query user u_x , and the number of future timestamps to predict M , CAPER sequentially predicts the top- k companies c_z and top- k positions p_y of u_x that are most likely to occur per future timestamp, i.e., $t_{(l+1)}, t_{(l+2)}, \dots, t_{(l+M)}$, via Eq. (5). To this end, we obtain the temporal and evolution embeddings of u_x and all companies c_z with respect to each future timestamp, and the embeddings of all positions p_y by performing GCN and GRNN encoders repeatedly.

CAPER requires no *additional* training during this multi-step inference process. Instead, since we do not know the career trajectories after t_l , the only extra task is to infer KG snapshots with respect to future timestamps that can be used for the GCN encoder. To obtain the temporal embeddings with respect to $t_{(l+1)}$, we use an already-built KG snapshot \mathcal{G}_{t_l} with respect to t_l , which contains the *actual careers* of the users. In contrast, to obtain the temporal embeddings after $t_{(l+1)}$, we sequentially construct KG snapshots, i.e., $\mathcal{G}_{t_{(l+1)}}, \mathcal{G}_{t_{(l+2)}}, \dots, \mathcal{G}_{t_{(l+(M-1))}}$, based on each user’s *inferred careers*, which consist of two pairs, i.e., (1st ranked company, 1st ranked position) and (2nd ranked company, 1st ranked position), between his/her top-2 companies and his/her top-1 position predicted at the previous timestamp. This design choice is based on the following observations: (1) the number of companies is much larger than the number of positions, as shown in Table 1, and (2) each user has an average of 2.37 careers in each KG snapshot.

As in other CTP work [36, 60, 69, 77], this paper does not assume the *inductive setting* typically used to address the prediction problem for new users, companies, and positions. However, as CAPER is designed based on KG-based GCN encoders, it can simply leverage the inductive ability [58, 65] inherent in GCNs by utilizing embeddings of existing neighbor nodes for new nodes. We leave the sophisticated design for an inductive setting as future work.

4.4 Discussion

4.4.1 Novelty. We further highlight the novelty of CAPER in contrast to state-of-the-art CTP methods [36, 60, 69, 77].

Mutual Dependency. NAOMI [69] preserves the structural dependency between positions or companies, while AHEAD [77] preserves the mutual dependency for pairwise relationships between each position and each company, disregarding ternary relationships in both cases. Notably, AHEAD [77] represents the career trajectories as a heterogeneous graph, where nodes indicate companies and positions, and edges indicate the career transitions between two companies or positions and whether a position belongs to a company [77]. However, we claim that omitting users in modeling careers is not a good idea since users often play a key role in career decisions [12]. Conversely, both HCPNN [36] and TACTP [60] model a user’s company sequence by leveraging job sequences and/or features (e.g., number of social connections and self-introduction), and differ significantly from CAPER in that: (1) they fail to accurately represent the position and user embeddings, as they only model users’ company sequences; (2) they are unable to benefit from the graph structure, which includes high-order relationships (i.e., more than 2 hops) among users, companies, and positions. Thus, we argue that none of the existing methods jointly model explicit and implicit correlation among the three units of a user, a position, and a company, which is critically important.

Characteristic Shifts. Existing methods of [36, 60, 69, 77] employed RNN encoders to learn the career sequences of users. This design aims to capture *common patterns in career movements* among users, rather than to capture the evolution in both of users and companies over time. Under such a design, each company is represented as a *single* learnable embedding [36, 60, 69, 77], while each user is represented as a *single static* (i.e., non-learnable) feature [36, 60]. Accordingly, they cannot capture the characteristics of each user/company at each timestamp, being unable to model their evolution. Thus, we argue that none of the existing methods consider the characteristic shifts of users and companies over time.

4.4.2 Complexity. CAPER employs GCN and GRNN to learn mutual dependency and temporal dynamics, respectively. As a result, the time complexity of CAPER is expressed as $l \cdot O(\text{GCN}) + O(\text{GRNN})$, where l indicates the number of timestamps (i.e., the number of KG snapshots). Specifically, the time complexity of GCN is defined as $O(L \cdot (n + j + d(d + 1)))$ [63], where L, n, j , and d represent the number of GCN layers, the number of nodes (i.e., users and companies), the number of edges (i.e., positions), and the dimensionality of embeddings, respectively. In addition, the time complexity of GRNN is given by $O(l \cdot (d_h(d_h + 1) + d_i))$ [48], where d_i and d_h indicate the input and hidden size of RNN, respectively.

Here, one might have concerns about the computational overhead of CAPER. Considering the architecture that repeats GCN and GRNN l times, the value of l can impact the performance of CAPER. However, we observed two key points: (1) the average inference time per user is only 0.89 msec with CAPER, suggesting that it provides real-time predictions, and (2) CAPER shows linear scalability as the number of users increases. Moreover, we can mitigate the computational overhead of CAPER on large datasets without significant accuracy loss by reducing the value of l for training each temporal embedding, as validated in the [online appendix](#).

5 Evaluation

To evaluate effectiveness of CAPER, we designed our experiments, aiming at answering the following key evaluation questions (EQs):

- (EQ1) Does CAPER outperform its competitors for career trajectory prediction?
- (EQ2) Is temporal knowledge graph modeling effective for career trajectory prediction?
- (EQ3) Is learning mutual dependency effective for career trajectory prediction?
- (EQ4) Is learning characteristic shifts over time effective for career trajectory prediction?

We also demonstrated that the accuracies of CAPER do not fluctuate a lot depending on the embedding dimensionality, and that CAPER works well in a different time granularity (see [online appendix](#)).

5.1 Experimental Settings

Dataset. Table 1 provides some statistics on our dataset. We used a real-world career trajectory dataset whose details were explained in Section 3.1. For the extrapolated career reasoning task (*additionally designed by us*) discussed in Section 4.1, we used the most recent five years of careers (i.e., $M=5$) for the test set and the remaining 47 years for the training set. It is worth noting that finding publicly

available datasets is very challenging in a CTP domain due to the proprietary nature of the datasets. Privacy concerns surrounding resume datasets further hinder their public release. Therefore, most prior literature, including NEMO [27], HCPNN [36], AHEAD [77], TACTP [60], and NAOMI [69], used a single private dataset of their own, none of which was shared publicly. For this reason, we had to rely only on our own dataset (obtained from a global career platform, FutureFit AI) for the evaluation.

For the benefit of the research community, our dataset is available upon request; interested researchers should contact FutureFit AI or the Penn State team to access the original dataset. Additionally, in [72], we introduced OpenResume, a publicly available, anonymized, and structured resume dataset, specifically curated for downstream tasks in the job domain.

Competitors. To evaluate the effectiveness of CAPER, we compare CAPER with four baselines, two recent TKG reasoning methods, and five state-of-the-art CTP methods – (1) baselines: Popular (*i.e.*, predicts the most-frequent companies and positions), RNN, GRU, and Transformer [59]; (2) TKG reasoning methods: DaeMon [10] and CENET [68]; (3) CTP methods: NEMO [27], HCPNN [36], AHEAD [77], TACTP [60], and NAOMI [69].

As we have carefully surveyed the relevant literature, including [41], it is important to note that, to the best of our knowledge, our competitors here cover all the relevant baselines in the CTP problem (refer to the ‘Career Mobility Prediction’ category in [41]). For evaluation, we used the source codes provided by the authors. Since none of the competitors was designed for multi-step inference, we call those algorithms repeatedly to predict the next companies and positions. We carefully tuned the hyperparameters of competitors and CAPER. For the full reproducibility, we provide complete implementation details in the Appendix.

Evaluation Tasks. Following [60, 77], we employ company and position prediction tasks. The goal of both tasks is to forecast each user’s next careers (*i.e.*, company or position) that are likely to occur at *non-existent future M timestamps* in the training set. Specifically, we consider only those users (*i.e.*, having their careers after 2015) included in the test set as test users. Given each test user’s career trajectory during the training period, we predict the next companies and positions that the user is likely to work in the test timestamps. Since we are already addressing the problem of long-term career prediction, we do not additionally perform the duration prediction task done in [36, 60, 77], which only predicts the single next career.

Following [27, 36, 60, 77], we employ the two popular metrics: mean reciprocal rank (MRR) and Accuracy (Acc) @ k ($k=1, 5, 10$). Using these metrics, we measure how accurately each method predicts each test user’s company or position per test timestamp (*i.e.*, from 2016 to 2020). We train each method over 3 random weight initializations and report the averaged accuracy.

Inadequateness of Random-Sampling-Based Evaluation. We note that some CTP methods [60, 77] evaluate their CTP solutions as follows: they randomly sample 10% of users and use their career trajectories to construct a test set. However, we believe this evaluation to be problematic.

Consider Figure 4 that depicts the career trajectories of three users. Suppose that the career trajectories of two users, A and B,



Figure 4: An example to show the limitations of the random-sampling-based evaluation employed in existing studies.

are used for the training set, while another user C’s career trajectory is used for the test set. In this case, where test user C has a career trajectory from 1999 to 2010, a CTP method should predict careers that occur at existent past timestamps in the training set. Specifically, since test user C has his/her last career at 2010 for some reason (*e.g.*, retirement), he/she doesn’t need predictions for the next career anymore. We believe such an evaluation is not realistic in terms of real-world applications (*e.g.*, future career prediction). We found that among all 333K users on our dataset (see Table 1), the proportion of users whose careers do not exist at the test timestamps (like user C) is 55.06% (*i.e.*, 149K users), which indicates they can be prone to be inappropriately sampled as test users under the random-sampling-based evaluation. To address this limitation, we consider only users (*e.g.*, A and B in Figure 4) who have careers at the test timestamps (*i.e.*, from 2016 to 2020) as test users, and their careers (*e.g.*, Google, Meta, EPFL, and Carnegie Mellon University in Figure 4) as the ground truth for evaluation.

5.2 Results

For EQ2, EQ3, and EQ4, we only show the results of CAPER when the value of M is 1 (*i.e.*, career prediction for the immediate next year). However, we confirmed that the performance for $M=[2-5]$ is consistent with that for $M=1$ (see [online appendix](#)).

(EQ1) Comparison with 11 Competitors. We conducted comparative experiments on our career trajectory dataset. to demonstrate the superiority of CAPER over the four baselines, two TKG reasoning methods, and five state-of-the-art CTP methods. In Table 3, the values in **boldface** and underlined indicate the best and 2nd best accuracies in each column, respectively.

From Table 3, we can see that CAPER consistently and significantly outperforms all competitors in almost all cases. Overall, CAPER yields 6.80% and 34.58% more-accurate company and position predictions on average, respectively, compared to the best competitor (*i.e.*, TACTP). Specifically, for company prediction, CAPER yields up to 20.29% and 6.91% higher MRR than the best competitor when the values of M are 1 and 2 (*i.e.*, short-term predictions, which are the most practical scenario in our setting), respectively. For long-term predictions (*i.e.*, $M=3,4,5$), CAPER also provides accuracy slightly better than or comparable to the best competitor. Additionally, for position prediction, CAPER yields up to 61.55%, 34.64%, 23.65%, 25.69%, and 32.06% higher MRR than the best competitor when the values of M are 1, 2, 3, 4, and 5, respectively. We

Table 3: Accuracies of four baselines, two recent TKG reasoning methods, and five state-of-the-art CTP methods, and CAPER. We conducted t -tests at a 95% confidence level, showing statistically significant differences between CAPER and competitors, with p -values below 0.05 (except for $M=4$ on $\text{Acc}@10$ in Table 3-(a)).

(a) Company Prediction																				
Methods	MRR					Acc@1					Acc@5					Acc@10				
	M=1	M=2	M=3	M=4	M=5	M=1	M=2	M=3	M=4	M=5	M=1	M=2	M=3	M=4	M=5	M=1	M=2	M=3	M=4	M=5
Popular	1.51	1.47	1.45	1.49	1.5	0.42	0.38	0.33	0.34	0.34	0.90	0.85	0.80	0.84	0.85	3.12	3.06	3.05	3.16	3.18
RNN	37.32	24.17	19.15	17.32	16.42	32.29	20.52	16.15	14.59	13.91	41.93	26.56	20.88	18.87	17.77	46.42	30.82	24.56	22.20	20.82
GRU	40.90	27.05	21.76	19.84	18.99	36.11	23.57	18.77	17.07	16.43	45.30	29.26	23.58	21.48	20.50	49.55	33.38	27.12	24.79	23.53
Transformer	38.88	32.30	28.77	26.42	25.07	27.96	24.92	21.46	17.47	15.17	45.46	34.57	30.18	29.63	28.94	50.23	38.62	34.73	34.18	34.10
DaeMon	53.70	42.51	38.24	35.16	30.25	44.32	35.24	25.26	21.20	18.57	59.83	40.25	35.22	30.88	26.80	<u>64.27</u>	50.20	36.90	32.48	29.51
CENET	<u>60.03</u>	47.52	42.74	39.30	33.81	56.39	44.83	32.13	26.97	23.62	<u>62.96</u>	42.35	37.06	32.49	28.20	63.09	49.27	36.20	31.88	28.96
NEMO	41.99	31.19	26.24	24.00	22.46	37.17	27.31	22.77	20.69	19.26	47.54	34.92	29.43	26.97	25.29	50.57	38.08	32.32	29.82	28.08
HCPNN	44.55	35.02	30.97	29.57	29.06	40.93	32.51	28.60	27.24	26.80	48.18	37.15	32.93	31.47	30.88	50.76	39.32	35.07	33.61	32.95
AHEAD	44.72	33.71	29.91	28.66	28.15	40.60	30.93	27.29	26.04	25.55	49.39	36.04	32.02	30.75	30.17	52.04	38.50	34.44	33.24	32.70
NAOMI	44.22	43.58	41.75	42.13	43.20	38.53	38.10	36.33	36.37	37.37	48.64	48.62	44.99	46.14	47.91	50.17	48.29	46.33	46.98	49.10
TACTP	59.08	<u>54.38</u>	<u>49.83</u>	<u>48.70</u>	<u>48.18</u>	<u>57.57</u>	<u>52.95</u>	<u>49.12</u>	<u>47.24</u>	<u>46.89</u>	60.23	<u>55.25</u>	<u>51.12</u>	<u>49.91</u>	<u>48.68</u>	61.10	<u>56.58</u>	<u>51.99</u>	51.70	<u>49.16</u>
CAPER	71.07	58.14	51.35	49.70	49.43	69.34	57.09	49.44	48.24	48.58	73.27	59.05	51.75	50.75	50.09	74.10	59.93	52.87	<u>51.57</u>	50.88

(b) Position Prediction																				
Methods	MRR					Acc@1					Acc@5					Acc@10				
	M=1	M=2	M=3	M=4	M=5	M=1	M=2	M=3	M=4	M=5	M=1	M=2	M=3	M=4	M=5	M=1	M=2	M=3	M=4	M=5
Popular	10.66	10.64	10.54	10.45	10.43	3.51	3.29	2.99	2.87	2.89	8.82	8.94	8.88	8.73	8.73	25.12	25.61	25.99	26.06	25.92
RNN	27.41	18.90	15.79	14.31	13.47	19.29	11.04	8.06	6.35	5.74	32.61	23.69	20.24	19.07	17.44	45.43	34.45	31.35	30.41	28.83
GRU	30.05	20.66	18.32	17.57	17.72	22.05	12.62	10.40	9.41	9.53	34.79	25.49	22.99	22.71	22.60	47.92	36.57	34.18	33.87	33.50
Transformer	28.83	26.64	22.05	21.92	20.75	20.13	18.89	13.31	12.76	11.63	34.01	29.81	25.32	22.36	21.28	48.13	44.07	40.08	36.47	33.93
DaeMon	40.85	33.84	21.54	19.60	17.54	31.50	18.88	13.54	12.74	11.44	51.40	44.63	35.12	34.76	31.58	59.38	50.50	42.67	41.00	40.51
CENET	<u>51.08</u>	<u>42.31</u>	26.93	24.51	21.93	<u>45.46</u>	27.25	19.54	18.38	16.51	<u>54.89</u>	<u>47.66</u>	37.50	37.12	33.72	<u>60.44</u>	51.40	43.43	41.73	41.23
NEMO	34.32	23.22	20.11	18.43	17.95	25.49	13.82	11.08	9.24	8.97	42.96	32.00	28.23	26.89	25.57	53.45	42.20	38.82	37.60	36.42
HCPNN	11.39	11.20	11.10	11.08	10.95	4.98	4.19	4.15	4.14	4.08	15.84	14.75	14.62	14.60	14.40	26.04	24.87	24.59	24.52	24.24
AHEAD	33.59	25.17	22.66	21.26	20.98	24.58	15.62	13.06	11.27	11.07	42.07	34.09	31.54	30.91	30.07	53.42	44.32	42.33	41.88	41.14
NAOMI	30.81	30.33	30.24	30.13	29.48	23.61	23.34	23.38	23.30	22.99	40.64	40.68	36.99	38.14	39.91	48.17	46.29	44.33	44.98	47.10
TACTP	37.09	36.14	<u>35.00</u>	<u>34.33</u>	<u>33.34</u>	28.73	<u>27.71</u>	<u>26.84</u>	<u>25.99</u>	<u>24.82</u>	43.53	42.78	<u>41.00</u>	<u>40.08</u>	<u>39.69</u>	52.91	<u>51.97</u>	<u>50.89</u>	<u>50.64</u>	<u>48.51</u>
CAPER	59.92	48.66	43.28	43.15	44.03	50.87	40.62	35.74	35.62	36.63	69.69	56.56	50.27	50.42	50.77	76.35	62.47	56.61	56.53	56.58

attribute the superiority of CAPER to its sophisticated career modeling and capturing of characteristic shifts of entities over time. That is, it validates our claims for the significance of considering mutual dependency and characteristic shifts in the CTP problem, as discussed in Section 4.4.

(EQ2) Effectiveness of TKG Modeling. In (M1) of Section 4.2, we modeled the given career trajectories as a TKG to represent the mutual dependency between key units and their characteristic shifts over time. For (EQ2), we compare CAPER with its two variants: CAPER_{KG} models the career trajectories as a KG without consideration of temporal dimension; and CAPER_{HetG} models the career trajectories as a heterogenous graph, where nodes indicate companies and positions, and edges indicate the career transitions between two companies or two positions and whether a position belongs to a company [77]. As shown in Table 4-(i), we see that CAPER consistently outperforms its two variants. For company (resp. position) prediction, CAPER yields up to 10.35% and 15.48% (resp.

28.94% and 30.97%) higher MRR than CAPER_{KG} and CAPER_{HetG}, respectively. For position prediction, CAPER yields up to 28.94% and 30.97% higher MRR than CAPER_{KG} and CAPER_{HetG}, respectively. It shows that not only preserving the mutual dependency for the ternary relationships on users, positions, and companies (*i.e.*, CAPER and CAPER_{KG} vs. CAPER_{HetG}) but also adding a temporal dimension to KG (*i.e.*, CAPER vs. CAPER_{KG}) is helpful in better modeling career trajectories.

(EQ3) Effectiveness of Learning Mutual Dependency. In (M2) of Section 4.2, we designed a GCN encoder to learn mutual dependencies for the ternary relationships of a user, a position, and a company. For (EQ3), we compare CAPER with its variant CAPER_{w/o GCN} (*i.e.*, not employing a GCN encoder and forwarding randomly initialized temporal embeddings into the GRNN encoder without aggregating career information). Table 4-(ii) shows that CAPER dramatically outperforms CAPER_{w/o GCN}. The results show

Table 4: Ablation studies for CAPER when the value of M is 1. The goals of (EQ2), (EQ3), and (EQ4) are to demonstrate the effectiveness of (M1), (M2), and (M3) in CAPER, respectively. The results show that not only each design choice in CAPER is effective, but also CAPER integrating all of them is the most effective.

(a) Company Prediction												
Metrics	(i) EQ2		(ii) EQ3	(iii) EQ4-1				(iv) EQ4-2				CAPER
	KG	HetG	w/o GCN	All	w/o \tilde{u}_x	w/o \tilde{c}_x	w/o \tilde{p}_x	w/o GRNN	RNN	GRU	Transformer	
MRR	64.40	61.54	7.44	69.94	55.36	69.17	71.07	56.01	69.99	69.83	66.87	71.07
Acc@1	60.64	58.11	4.65	68.67	51.50	67.76	69.34	53.48	68.20	68.08	63.71	69.34
Acc@5	66.28	65.21	9.41	70.51	59.54	69.78	73.27	56.97	70.80	70.70	70.28	73.27
Acc@10	71.35	67.86	12.49	72.10	62.55	71.61	74.10	60.69	73.12	72.87	72.53	74.10

(b) Position Prediction												
Metrics	(i) EQ2		(ii) EQ3	(iii) EQ4-1				(iv) EQ4-2				CAPER
	KG	HetG	w/o GCN	All	w/o \tilde{u}_x	w/o \tilde{c}_x	w/o \tilde{p}_x	w/o GRNN	RNN	GRU	Transformer	
MRR	46.47	45.75	3.85	16.70	40.40	17.74	59.92	30.37	58.95	58.01	52.74	59.92
Acc@1	28.82	29.07	0.06	7.94	27.74	8.66	49.87	21.21	49.77	48.00	40.79	49.87
Acc@5	60.41	66.95	2.65	22.96	55.15	24.71	69.69	32.91	64.71	64.50	66.40	69.69
Acc@10	75.73	73.5	9.83	34.74	63.39	36.59	76.35	48.23	74.09	74.67	72.93	76.35

that our convolutional layer benefits from learning mutual dependency for ternary relationships.

(EQ4) Effectiveness of Learning Characteristic Shifts. In (M3) of Section 4.2, we represented the same unit (*spec.*, user and company) at different timestamps as different temporal embeddings. Based on the temporal embeddings, we then reflected the changes in the characteristics of users and companies in their evolution embeddings gradually. To verify the effectiveness of both ideas, we conduct experiments to answer the two sub-questions:

- **(EQ4-1)** Is it effective to consider characteristic shifts of all units?
- **(EQ4-2)** Is it effective to model the evolution over time of units based on GRNN?

For (EQ4-1), we compare the accuracy of the following 4 variants of CAPER: (1) CAPER_{all} represents all units at different timestamps as different temporal embeddings; (2) CAPER_{w/o \tilde{u}_x} represents each user as a single embedding regardless of timestamps; (3) CAPER_{w/o \tilde{c}_x} represents each company as a single embedding regardless of timestamps; (4) CAPER_{w/o \tilde{p}_x} represents each position as a single embedding regardless of timestamps. To implement CAPER_{all}, we add another GRNN encoder for positions in Figure 2. Here, **CAPER_{w/o \tilde{p}_x} is the same as the final version of CAPER.**

As shown in Table 4-(iii), CAPER_{w/o \tilde{p}_x} consistently outperforms other variants of CAPER. Notably, we see that for position prediction, using evolution embeddings of positions adversely affects the performance of CAPER. On the other hand, we note that jointly considering the characteristic shifts of users and companies is consistently helpful for solving CTP problem. We conjecture that the reason of this phenomenon is that, unlike users and companies, the characteristics of positions tend to have a more-static nature that does not change significantly over time. For instance, the nature of "software engineer" position remains more static between 2010 and 2020 (despite new positions getting created) while the skills of users or characteristics of companies change more dynamically.

For (EQ4-2), we compare CAPER (*i.e.*, CAPER_{LSTM}) with its variant CAPER_{w/o GRNN}, which does not employ any GRNN encoders and thus cannot construct evolution embeddings of users and companies. Table 4-(iv) illustrates the importance of learning

the evolution of users and companies over time by sequentially accumulating their past career information. Additionally, we found that employing RNN or GRU encoders yields comparable results. Conversely, we observed that the Transformer encoder exhibits limited effectiveness in terms of accuracy in the CTP domain. This suggests that the sequential order of each user's career is particularly important in this domain. From the results, we believe that our superior performance can be attributed to our novel modeling and utilization of CTP data, rather than specific techniques.

6 Conclusions

In this work, we found that existing CTP methods fail to jointly consider the mutual dependency for ternary relationships among a user, a position, and a company, and the characteristic shifts over time. To address the challenges, we proposed a novel CTP approach, CAPER, that consists of three key modules: (M1) modeling career trajectories, (M2) learning mutual dependency, and (M3) learning temporal dynamics. Furthermore, we formulated the CTP problem as the extrapolated career reasoning task on TKG, which aims to forecast the next careers that are likely to occur at non-existent future timestamps in the input TKG. Via experiments using a real-world career trajectory dataset, we demonstrated that CAPER consistently and significantly outperforms 4 baselines, 2 recent TKG reasoning methods, and 5 state-of-the-art CTP methods.

Acknowledgments

The work of Sang-Wook Kim was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2022-0-00352 and No.RS-2022-00155586, A High-Performance Big-Hypergraph Mining Platform for Real-World Downstream Tasks). Yeon-Chang Lee's work was supported by the Institute of Information & communications Technology Planning & Evaluation (IITP) grant, funded by the Korea government (MSIT) (No. RS-2020-II201336, Artificial Intelligence Graduate School Program (UNIST)). Dongwon Lee's was supported by NSF award #1934782. The dataset was provided by FutureFit AI under a Memorandum of Understanding.

References

- [1] Elias Abdollahnejad, Marilynn Kalman, and Behrouz H Far. 2021. A deep learning bert-based approach to person-job fit in talent recruitment. In *Proceedings of the International Conference on Computational Science and Computational Intelligence (CSCI)*. 98–104.
- [2] Shuqing Bian, Xu Chen, Wayne Xin Zhao, Kun Zhou, Yupeng Hou, Yang Song, Tao Zhang, and Ji-Rong Wen. 2020. Learning to Match Jobs with Resumes from Sparse Interaction Data using Multi-View Co-Teaching Network. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*. 65–74.
- [3] Shuqing Bian, Wayne Xin Zhao, Yang Song, Tao Zhang, and Ji-Rong Wen. 2019. Domain Adaptation for Person-Job Fit with Transferable Deep Global Match Network. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 4809–4819.
- [4] Michele A. Brandão and Mirella M. Moro. 2017. Social professional networks: A survey and taxonomy. *Computer Communications* 100 (2017), 20–31.
- [5] Xi Chen, Yiqun Liu, Liang Zhang, and Krishnaram Kenthapadi. 2018. How LinkedIn economic graph bonds information and product: applications in LinkedIn salary. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD)*. 120–129.
- [6] Zhongwu Chen, Chengjin Xu, Fenglong Su, Zhen Huang, and Yong Dou. 2023. Meta-Learning Based Knowledge Extrapolation for Temporal Knowledge Graph. In *Proceedings of the ACM Web Conference (WWW)*. 2433–2443.
- [7] Le Dai, Yu Yin, Chuan Qin, Tong Xu, Xiangnan He, Enhong Chen, and Hui Xiong. 2020. Enterprise Cooperation and Competition Analysis with a Sign-Oriented Preference Network. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*. 774–782.
- [8] Vachik S. Dave, Baichuan Zhang, Mohammad Al Hasan, Khalifeh AlJadda, and Mohammed Korayem. 2018. A Combined Representation Learning Approach for Better Job and Skill Recommendation. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*. 1997–2005.
- [9] Zifeng Ding, Yunpu Ma, Bailan He, and Volker Tresp. 2021. A Simple But Powerful Graph Encoder for Temporal Knowledge Graph Completion. *arXiv preprint arXiv:2112.07791* (2021).
- [10] Hao Dong, Zhiyuan Ning, Pengyang Wang, Ziyue Qiao, Pengfei Wang, Yuanchun Zhou, and Yanjie Fu. 2023. Adaptive Path-Memory Network for Temporal Knowledge Graph Reasoning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 2086–2094.
- [11] Chuyu Fang, Chuan Qin, Qi Zhang, Kaichun Yao, Jingshuai Zhang, Hengshu Zhu, Fuzhen Zhuang, and Hui Xiong. 2023. RecruitPro: A Pretrained Language Model with Skill-Aware Prompt Learning for Intelligent Recruitment. In *Proceedings of the ACM Conference on Knowledge Discovery and Data Mining (KDD)*. 3991–4002.
- [12] Itamar Gati, Nimrod Levin, and Shiri Landman-Tal. 2019. Decision-making models and career guidance. *International Handbook of Career Guidance* (2019), 115–145.
- [13] Zhuoning Guo, Hao Liu, Le Zhang, Qi Zhang, Hengshu Zhu, and Hui Xiong. 2022. Talent Demand-Supply Joint Prediction with Dynamic Heterogeneous Graph Enhanced Meta-Learning. In *Proceedings of the ACM Conference on Knowledge Discovery and Data Mining (KDD)*. 2957–2967.
- [14] Zhen Han et al. 2021. Explainable Subgraph Reasoning for Forecasting on Temporal Knowledge Graphs. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- [15] Shaoxiong Ji, Shirui Pan, Erik Cambria, Pekka Marttinen, and Philip S. Yu. 2022. A Survey on Knowledge Graphs: Representation, Acquisition, and Applications. *IEEE Trans. on Neural Networks Learning Systems (TNNLS)* 33, 2 (2022), 494–514.
- [16] Yiqiao Jin, Yeon-Chang Lee, Kartik Sharma, Meng Ye, Karan Sikka, Ajay Divakaran, and Srijan Kumar. 2023. Predicting Information Pathways Across Online Communities. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. 1044–1056.
- [17] Jaehun Jung, Jinhong Jung, and U Kang. 2021. Learning to Walk across Time for Interpretable Temporal Knowledge Graph Completion. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD)*.
- [18] Min-Jeong Kim, Yeon-Chang Lee, and Sang-Wook Kim. 2023. TrustSGCN: Learning Trustworthiness on Edge Signs for Effective Signed Graph Convolutional Networks. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. 2451–2455.
- [19] Min-Jeong Kim, Yeon-Chang Lee, and Sang-Wook Kim. 2024. PolarDSN: An Inductive Approach to Learning the Evolution of Network Polarization in Dynamic Signed Networks. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*. 1099–1109.
- [20] Taeri Kim, Yeon-Chang Lee, Kijung Shin, and Sang-Wook Kim. 2022. MARIO: Modality-Aware Attention and Modality-Preserving Decoders for Multimedia Recommendation. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*. 993–1002.
- [21] Taeyong Kong, Taeri Kim, Jinsung Jeon, Jeongwhan Choi, Yeon-Chang Lee, Noseong Park, and Sang-Wook Kim. 2022. Linear, or Non-Linear, That is the Question!. In *Proceedings of ACM International Conference on Web Search and Data Mining (WSDM)*. 517–525.
- [22] Yeon-Chang Lee, Jiwon Hong, and Sang-Wook Kim. 2016. Job recommendation in AskStory: experiences, methods, and evaluation. In *Proceedings of the 31st Annual ACM Symposium on Applied Computing (SAC)*. 780–786.
- [23] Yeon-Chang Lee, JaeHyun Lee, Dongwon Lee, and Sang-Wook Kim. 2024. Learning to compensate for lack of information: Extracting latent knowledge for effective temporal knowledge graph completion. *Inf. Sci.* 654 (2024), 119857.
- [24] Yujin Lee, Yeon-Chang Lee, Jiwon Hong, and Sang-Wook Kim. 2017. Exploiting job transition patterns for effective job recommendation. In *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 2414–2419.
- [25] Yeon-Chang Lee, JaeHyun Lee, Dongwon Lee, and Sang-Wook Kim. 2022. THOR: Self-Supervised Temporal Knowledge Graph Embedding via Three-Tower Graph Convolutional Networks. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*. 1035–1040.
- [26] Huayu Li, Yong Ge, Hengshu Zhu, Hui Xiong, and Hongke Zhao. 2017. Prospecting the Career Development of Talents: A Survival Analysis Perspective. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*. 917–925.
- [27] Liangyue Li, How Jing, Hanghang Tong, Jaewon Yang, Qi He, and Bee-Chung Chen. 2017. Nemo: Next career move prediction with contextual embedding. In *Proceedings of The ACM Web Conference (WWW)*. 505–513.
- [28] Lei Li, Guangzhi Zheng, Svetlana Peltsverger, and Chi Zhang. 2016. Career Trajectory Analysis of Information Technology Alumni: A LinkedIn Perspective. In *Proceedings of the Annual Conference on Information Technology Education*. 2–6.
- [29] Shan Li, Baoxu Shi, Jaewon Yang, Ji Yan, Shuai Wang, Fei Chen, and Qi He. 2020. Deep Job Understanding at LinkedIn. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. 2145–2148.
- [30] Yunqi Li, Michiharu Yamashita, Hanxiong Chen, Dongwon Lee, and Yongfeng Zhang. 2023. Fairness in Job Recommendation Under Quantity Constraints. In *AAAI-23 Workshop on AI for Web Advertising*.
- [31] Zixuan Li et al. 2021. Search from History and Reason for Future: Two-stage Reasoning on Temporal Knowledge Graphs. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL)*.
- [32] Zixuan Li et al. 2021. Temporal Knowledge Graph Reasoning Based on Evolutional Representation Learning. In *Proceedings of International ACM Conference on Research and Development in Information Retrieval (SIGIR)*.
- [33] Ye Liu, Luming Zhang, Liqiang Nie, Yan Yan, and David Rosenblum. 2016. Fortune Teller: Predicting Your Career Path. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 30.
- [34] Haiyan Luo, Shichuan Ma, Anand Joseph Bernard Selvaraj, and Yu Sun. 2019. Learning job representation using directed graph embedding. In *Proceedings of the International Workshop on Deep Learning Practice for High-Dimensional Sparse Data*. 1–5.
- [35] Sumit Maheshwari, Abhishek Sainani, and P Krishna Reddy. 2010. An Approach to Extract Special Skills to Improve the Performance of Resume Selection. In *International Workshop on Databases in Networked Information Systems*. 256–273.
- [36] Qingxin Meng, Hengshu Zhu, Keli Xiao, Le Zhang, and Hui Xiong. 2019. A hierarchical career-path-aware neural network for job mobility prediction. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD)*. 14–24.
- [37] Sebastian Montella, Lina Maria Rojas-Barahona, and Johannes Heinecke. 2021. Hyperbolic Temporal Knowledge Graph Embeddings with Relational and Time Curvatures. In *Proceedings of Annual Meeting of the Association for Computational Linguistics (ACL)*. 3296–3308.
- [38] Richard Jayadi Oentaryo, Ee-Peng Lim, Xavier Jayaraj Siddarth Ashok, Philips Kokoh Prasetyo, Koon Han Ong, and Zi Quan Lau. 2018. Talent Flow Analytics in Online Professional Network. *Data Science and Engineering* 3, 3 (2018), 199–220.
- [39] Ioannis Paparrizos, B. Barla Cambazoglu, and Aristides Gionis. 2011. Machine Learned Job Recommendation. In *Proceedings of the Fifth ACM Conference on Recommender Systems*. 325–328.
- [40] Namyoung Park, Fuchen Liu, Purvanshi Mehta, Dana Cristofor, Christos Faloutsos, and Yuxiao Dong. 2022. Evokg: Jointly modeling event time and network structure for reasoning over temporal knowledge graphs. In *Proceedings of ACM International Conference on Web Search and Data Mining (WSDM)*. 794–803.
- [41] Chuan Qin, Le Zhang, Rui Zha, Dazhong Shen, Qi Zhang, Ying Sun, Chen Zhu, Hengshu Zhu, and Hui Xiong. 2023. A Comprehensive Survey of Artificial Intelligence Techniques for Talent Analytics. *CoRR abs/2307.03195* (2023).
- [42] Chuan Qin, Hengshu Zhu, Tong Xu, Chen Zhu, Liang Jiang, Enhong Chen, and Hui Xiong. 2018. Enhancing Person-Job Fit for Talent Recruitment: An Ability-aware Neural Network Approach. In *Proceedings of International ACM Conference on Research and Development in Information Retrieval (SIGIR)*. 25–34.
- [43] Chuan Qin, Hengshu Zhu, Tong Xu, Chen Zhu, Liang Jiang, Enhong Chen, and Hui Xiong. 2018. Enhancing Person-Job Fit for Talent Recruitment: An Ability-Aware Neural Network Approach. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. 25–34.

- [44] Chuan Qin, Hengshu Zhu, Tong Xu, Chen Zhu, Chao Ma, Enhong Chen, and Hui Xiong. 2020. An enhanced neural network approach to person-job fit in talent recruitment. *ACM Transactions on Information Systems (TOIS)* 38, 2 (2020), 1–33.
- [45] Zhiwei Qin, Liangjie Hong, Rui Song, Hongtu Zhu, Mohammed Korayem, Haiyan Luo, and Michael I. Jordan. 2020. Decision Intelligence and Analytics for Online Marketplaces: Jobs, Ridesharing, Retail and Beyond. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*. 4898–4899.
- [46] Rohan Ramanath, Hakan Inan, Gungor Polatkan, Bo Hu, Qi Guo, Cagri Ozcaglar, Xianren Wu, Krishnamurthy Kothapadi, and Sahin Cem Geyik. 2018. Towards Deep and Representation Learning for Talent Search at LinkedIn. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*. 2253–2261.
- [47] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 3980–3990.
- [48] Michael Rotman and Lior Wolf. 2021. Shuffling Recurrent Neural Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 9428–9435.
- [49] Hojjat Salehinejad, Julianne Baarbe, Sharan Sankar, Joseph Barfett, Errol Colak, and Shahrokh Valaei. 2018. Recent Advances in Recurrent Neural Networks. *arXiv preprint arXiv:1801.01078* (2018).
- [50] Kartik Sharma, Yeon-Chang Lee, Sivagami Nambi, Aditya Salian, Shlok Shah, Sang-Wook Kim, and Srikanth Kumar. 2024. A Survey of Graph Neural Networks for Social Recommender Systems. *ACM Comput. Surv.* (2024), 265.
- [51] Baoxu Shi, Shan Li, Jaewon Yang, Mustafa Emre Kazdagli, and Qi He. 2020. Learning to Ask Screening Questions for Job Postings. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. 549–558.
- [52] Baoxu Shi, Jaewon Yang, Feng Guo, and Qi He. 2020. Saliency and Market-Aware Skill Extraction for Job Targeting. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*. 2871–2879.
- [53] Haohai Sun et al. 2021. TimeTraveler: Reinforcement Learning for Temporal Knowledge Graph Forecasting. In *Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [54] Ying Sun, Fuzhen Zhuang, Hengshu Zhu, Xin Song, Qing He, and Hui Xiong. 2019. The Impact of Person-Organization Fit on Talent Management: A Structure-Aware Convolutional Neural Network Approach. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*. 1625–1633.
- [55] Ibraiz Tarique and Randall S Schuler. 2010. Global talent management: Literature review, integrative framework, and suggestions for further research. *Journal of world business* 45, 2 (2010), 122–133.
- [56] Mingfei Teng, Hengshu Zhu, Chuanren Liu, Chen Zhu, and Hui Xiong. 2019. Exploiting the Contagious Effect for Employee Turnover Prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 33. 1166–1173.
- [57] Rakshit Trivedi, Hanjun Dai, Yichen Wang, and Le Song. 2017. Know-Evolve: Deep Temporal Reasoning for Dynamic Knowledge Graphs. In *Proceedings of the International Conference on Machine Learning (ICML)*, Vol. 70. 3462–3471.
- [58] Rakshit Trivedi, Mehrdad Farajtabar, Prasenjeet Biswal, and Hongyuan Zha. 2019. DyRep: Learning Representations over Dynamic Graphs. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems (NeurIPS)*. 5998–6008.
- [60] Chao Wang, Hengshu Zhu, Qiming Hao, Keli Xiao, and Hui Xiong. 2021. Variable interval time sequence modeling for career trajectory prediction: Deep collaborative perspective. In *Proceedings of The ACM Web Conference (WWW)*. 612–623.
- [61] Jingya Wang, Kareem Abdelfatah, Mohammed Korayem, and Janani Balaji. 2019. Deepcarotene-Job Title Classification with Multi-Stream Convolutional Neural Network. In *Proceedings of the IEEE International Conference on Big Data (BigData)*. 1953–1961.
- [62] Jian Wang, Yi Zhang, Christian Posse, and Anmol Bhasin. 2013. Is It Time for a Career Switch?. In *Proceedings of the International Conference on World Wide Web (WWW)*. 1377–1388.
- [63] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. 2019. A Comprehensive Survey on Graph Neural Networks. *CoRR abs/1901.00596* (2019).
- [64] Chengjin Xu et al. 2021. Temporal Knowledge Graph Completion using a Linear Temporal Regularizer and Multivector Embeddings. In *Proceedings of Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*. 2569–2578.
- [65] Da Xu, Chuanwei Ruan, Evren Körpeoğlu, Sushant Kumar, and Kannan Achan. 2020. Inductive representation learning on temporal graphs. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- [66] Huang Xu, Zhiwen Yu, Jingyuan Yang, Hui Xiong, and Hengshu Zhu. 2018. Dynamic Talent Flow Analysis with Deep Sequence Prediction Modeling. *IEEE Trans. Knowl. Data Eng.* 31, 10 (2018), 1926–1939.
- [67] Ye Xu, Zang Li, Abhishek Gupta, Ahmet Bugdayci, and Anmol Bhasin. 2014. Modeling Professional Similarity by Mining Professional Career Trajectories. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*. 1945–1954.
- [68] Yi Xu, Junjie Ou, Hui Xu, and Luoyi Fu. 2023. Temporal Knowledge Graph Reasoning with Historical Contrastive Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 4765–4773.
- [69] Michiharu Yamashita, Yunqi Li, Thanh Tran, Yongfeng Zhang, and Dongwon Lee. 2022. Looking further into the future: Career pathway prediction. In *Proceedings of the International Workshop on Computational Jobs Marketplace*.
- [70] Michiharu Yamashita, Jia Tracy Shen, Thanh Tran, Hamoon Ekhtiari, and Dongwon Lee. 2023. JAMES: Normalizing Job Titles with Multi-Aspect Graph Embeddings and Reasoning. In *2023 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE.
- [71] Michiharu Yamashita, Thanh Tran, and Dongwon Lee. 2024. Fake Resume Attacks: Data Poisoning on Online Job Platforms. In *Proceedings of the ACM on Web Conference (WWW)*. 1734–1745.
- [72] Michiharu Yamashita, Thanh Tran, and Dongwon Lee. 2024. OpenResume: Advancing Career Trajectory Modeling with Anonymized and Synthetic Resume Datasets. In *Proceedings of IEEE International Conference on Big Data (BigData)*.
- [73] Kaichun Yao, Jingshuai Zhang, Chuan Qin, Peng Wang, Hengshu Zhu, and Hui Xiong. 2022. Knowledge Enhanced Person-Job Fit for Talent Recruitment. In *Proceedings of the IEEE International Conference on Data Engineering (ICDE)*. 3467–3480.
- [74] Hyunsik Yoo, Yeon-Chang Lee, Kijung Shin, and Sang-Wook Kim. 2023. Disentangling Degree-related Biases and Interest for Out-of-Distribution Generalized Directed Network Embedding. In *Proceedings of the ACM Web Conference (WWW)*. 231–239.
- [75] Denghui Zhang, Junming Liu, Hengshu Zhu, Yanchi Liu, Lichen Wang, Pengyang Wang, and Hui Xiong. 2019. Job2vec: Job title benchmarking with collective multi-view representation learning. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*. 2763–2771.
- [76] Le Zhang, Tong Xu, Hengshu Zhu, Chuan Qin, Qingxin Meng, Hui Xiong, and Enhong Chen. 2020. Large-scale talent flow embedding for company competitive analysis. In *Proceedings of The ACM Web Conference (WWW)*. 2354–2364.
- [77] Le Zhang, Ding Zhou, Hengshu Zhu, Tong Xu, Rui Zha, Enhong Chen, and Hui Xiong. 2021. Attentive heterogeneous graph embedding for job mobility prediction. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD)*. 2192–2201.
- [78] Le Zhang, Hengshu Zhu, Tong Xu, Chen Zhu, Chuan Qin, Qingxin Meng, Hui Xiong, and Enhong Chen. 2019. Large-Scale Talent Flow Forecast with Dynamic Latent Factor Model. In *Proceedings of The Web Conference (WWW)*. 2312–2322.
- [79] Qi Zhang, Hengshu Zhu, Ying Sun, Hao Liu, Fuzhen Zhuang, and Hui Xiong. 2021. Talent Demand Forecasting with Attentive Neural Sequential Model. In *Proceedings of the ACM Conference on Knowledge Discovery and Data Mining (KDD)*. 3906–3916.
- [80] Sensen Zhang, Xun Liang, Zhiying Li, Junlan Feng, Xiangping Zheng, and Bo Wu. 2023. BiQCcap: A Biquaternion and Capsule Network-Based Embedding Model for Temporal Knowledge Graph Completion. In *Proceedings of International Conference on Database Systems for Advanced Applications (DASFAA)*. 673–688.
- [81] Cunchao Zhu, Muhao Chen, Changjun Fan, Guangquan Cheng, and Yan Zhang. 2021. Learning from history: Modeling temporal knowledge graphs with sequential copy-generation networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 35. 4732–4740.
- [82] Chen Zhu, Hengshu Zhu, Hui Xiong, Pengliang Ding, and Fang Xie. 2016. Recruitment Market Trend Analysis with Sequential Latent Variable Models. In *Proceedings of the ACM Conference on Knowledge Discovery and Data Mining (KDD)*. 383–392.
- [83] Yun Zhu, Faizan Javed, and Ozgur Ozturk. 2016. Semantic Similarity Strategies for Job Title Classification. *arXiv preprint arXiv:1609.06268* (2016).

A Further Data Description

We pre-process and standardize job titles (i.e., position names) and company names. For example, the position of “software engineer” may be written as “software development engineer” or “SDE” in different resumes for various reasons (e.g., lack of industry standard taxonomy across companies or individual preference). However, in order to successfully solve the CTP problem, we need to identify and coalesce all differently-worded position names that actually refer to the same or very similar positions. For obtaining such normalized position names, we performed entity mapping based on the state-of-the-art job title mapping method [70], which cleans job titles



Figure I: Distributions of our dataset (i.e., x -axis indicates positions or companies in frequency order, while y -axis indicates the number of users in log scale).

Table I: Top ten frequent positions in the dataset

Normalized Position	Percentile (%)
policy administration professionals	2.49
systems analysts	2.39
social work and counselling professionals	2.30
advertising and marketing professionals	2.21
personnel and careers professionals	2.07
university and higher education teachers	1.97
sales and marketing managers	1.89
retail and wholesale trade managers	1.80
office supervisors	1.73
administrative and executive secretaries	1.56

into the ECSO-accepted occupations (i.e., the public job taxonomy). In the end, the position names in our dataset are mapped to 425 normalized position names. For cleaning company names, following the previous works [27, 77], we use the most frequently-occurred 9,000 companies while filtering out low frequent companies.

After standardizing the position names, the distribution of the normalized names is shown in Figure I-(a). On the other hand, Figure I-(b) shows the distribution of company names. In addition, Table I shows the top-10 most-frequent normalized position names.

Lastly, we conducted a detailed analysis by computing the average working duration between jobs for specific sectors. We filtered our dataset based on job titles, categorizing them into distinct domains such as Tech, Sales, University, and Medical. The criterion for filtering included specific keywords within job titles, such as “software” for Tech and “university” for University. Below, we provide examples from four domains that we examined:

- Tech: Mean duration of 42.8 months, Std of 22.9 months
- Sales: Mean duration of 43.1 months, Std of 22.1 months
- University: Mean duration of 43.8 months, Std of 27.5 months
- Medical: Mean duration of 43.4 months, Std of 23.8 months

The results revealed that while the mobility in the Tech sector is marginally higher, the differences across these domains are not statistically significant ($p < 0.01$). This suggests a relatively uniform mobility pattern across different professions. Also, it is important to note that career advancement and promotion opportunities exist within the same company or organization, which is also common in university-related jobs.

B Implementation Details

For hyperparameters of competitors, we used the best settings found via grid search in the following ranges: embedding dimensionality $\in \{50, 100, 150, 200, 250, 300\}$ (for DaeMon, CENET, NEMO, HCPNN, AHEAD, TACTP, and NAOMI) and the number of GCN layers $\in \{1, 2, 3\}$ (for AHEAD). Therefore, the optimal dimensionalities for embedding were 100 for DaeMon, NEMO, HCPNN, and AHEAD, 150 for TACTP, 200 for CENET, and 250 for NAOMI. Also, based on our experiments, the most optimal number of GCN layers for AHEAD is 1.

For CAPER, we set the embedding dimensionality, the number of GCN layers, and the max epoch for training as 150, 1, and 100, respectively. When we employ the LSTMs in (M3), GRNN can be formulated to obtain $\tilde{u}_x(t_m)$ of u_x with respect to t_m :

$$\begin{aligned}
 \mathbf{f}_{u_x}(t_m) &= \mathbf{W}_{f(u,1)}(\tilde{u}_x(t_{m-1})) + \mathbf{W}_{f(u,2)}(\mathbf{u}_x(t_m)), \\
 \mathbf{i}_{u_x}(t_m) &= \mathbf{W}_{i(u,1)}(\tilde{u}_x(t_{m-1})) + \mathbf{W}_{i(u,2)}(\mathbf{u}_x(t_m)), \\
 \mathbf{g}_{u_x}(t_m) &= \mathbf{W}_{g(u,1)}(\tilde{u}_x(t_{m-1})) + \mathbf{W}_{g(u,2)}(\mathbf{u}_x(t_m)), \\
 \mathbf{o}_{u_x}(t_m) &= \mathbf{W}_{o(u,1)}(\tilde{u}_x(t_{m-1})) + \mathbf{W}_{o(u,2)}(\mathbf{u}_x(t_m)), \\
 \mathbf{cell}_{u_x}(t_m) &= \mathbf{cell}_{u_x}(t_{m-1}) \circ \mathbf{f}_{u_x}(t_m) + \mathbf{i}_{u_x}(t_m) \circ \mathbf{g}_{u_x}(t_m), \\
 \tilde{u}_x(t_m) &= \mathbf{o}_{u_x}(t_m) \circ \tanh(\mathbf{cell}_{u_x}(t_m)),
 \end{aligned} \tag{8}$$

where $\mathbf{cell}_{u_x}(t_m)$ indicates a cell state for $\tilde{u}_x(t_m)$ of u_x with respect to t_m . Also, $\mathbf{f}_{u_x}(t_m)$, $\mathbf{i}_{u_x}(t_m)$, $\mathbf{g}_{u_x}(t_m)$, and $\mathbf{o}_{u_x}(t_m)$ represent the intermediate calculation processes to obtain $\tilde{u}_x(t_m)$. Lastly, $\mathbf{W}_{f(u,\cdot)}$, $\mathbf{W}_{i(u,\cdot)}$, $\mathbf{W}_{g(u,\cdot)}$, and $\mathbf{W}_{o(u,\cdot)}$ represent the learnable weight matrices; they are shared across all users. In the same manner, we can obtain $\tilde{c}_z(t_m)$ of c_z with respect to t_m .

C Data Sharing Statement

The rarity of accessible job/resume datasets underscores the significance of our commitment to reproducibility. It is noteworthy that none of the over 30 papers [7, 8, 26–30, 33–36, 38, 39, 43, 45, 46, 51, 52, 54, 56, 60–62, 66, 67, 69–71, 75, 77, 78, 83] that propose algorithms using job/resume datasets have opted to share their datasets due to the sensitive nature of these datasets. Privacy concerns surrounding resume datasets, which contain sensitive personal information, further hinder their public release. Consequently, the majority of CTP literature, including studies such as NEMO [27], HCPNN [36], AHEAD [77], TACTP [60], and NAOMI [69], has relied on a single private dataset, none of which have been publicly shared. The proprietary nature of the datasets in the CTP domain makes finding publicly available datasets very challenging.

As a result, we had to rely solely on our own dataset, obtained from a global career platform, for evaluation. However, with explicit consent from the company that donated our dataset, we pledge to make it available upon request, with all due legal protocols observed. Our dedication to openness extends beyond mere rhetoric—we have also shared our entire codebase. For the benefit of the research community, our dataset is accessible upon request; interested researchers should contact FutureFit AI or the Penn State team to access the original dataset. Additionally, in [72], we introduced OpenResume, an anonymized, structured, and publicly available resume dataset, specifically curated for downstream tasks in the job domain.