Career Transitions and Trajectories: A Case Study in Computing

Tara Safavi

University of Michigan, Ann Arbor
tsafavi@umich.edu

Maryam Davoodi

Purdue University
mdavoodi@purdue.edu

Danai Koutra

University of Michigan, Ann Arbor
dkoutra@umich.edu

ABSTRACT

From artificial intelligence to network security to hardware design, it is well-known that computing research drives many important technological and societal advancements. However, less is known about the long-term career paths of the people behind these innovations. What do their careers reveal about the evolution of computing research? Which institutions were and are the most important in this field, and for what reasons? Can insights into computing career trajectories help predict employer retention?

In this paper we analyze several decades of post-PhD computing careers using a large new dataset rich with professional information, and propose a versatile career network model, R^3 , that captures temporal career dynamics. With R^3 we track important organizations in computing research history, analyze career movement between industry, academia, and government, and build a powerful predictive model for individual career transitions. Our study, the first of its kind, is a starting point for understanding computing research careers, and may inform employer recruitment and retention mechanisms at a time when the demand for specialized computational expertise far exceeds supply.

ACM Reference Format:

Tara Safavi, Maryam Davoodi, and Danai Koutra. 2018. Career Transitions and Trajectories: A Case Study in Computing. In *KDD '18: The 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, August 19–23, 2018, London, United Kingdom*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3219819.3219863>

1 INTRODUCTION

From the invention of the Unix operating system in the 1970s to the ongoing artificial intelligence revolution, the importance and impact of computing research can hardly be overstated. The world has taken notice accordingly: the news media regularly covers everything from frontiers in computer design [7] to the earnings of AI experts [24]. Naturally, questions regarding computing research careers are becoming relevant. What happens after a PhD in computer science? Which organizations are, or were, central in computing research? How do expertise and talent flow between organizations?

In this study, we answer these questions by analyzing a unique career trajectory dataset of computer science PhD graduates from the 1970s to the present. Our goal, broadly, is to understand the evolution of computing research as a profession on the levels of

individual **career transitions** (movement between distinct employers), **organizations** (employers), and three respective **sectors** (industry, academia, and government). To do so we propose R^3 , a versatile career network model that captures *resource* flow, employer *retention*, and *relative* organizational growth. Combining R^3 with the HITS link analysis algorithm [18], which has not (to the best of our knowledge) been used in career analysis before, we demonstrate R^3 's versatility with insights of varying granularity:

- **System-wide evolution.** We identify key organizations, from startups to universities to industry leaders, in computing research history. R^3 captures crucial factors beyond size and popularity that contribute to organizational “importance”, demonstrating that *some organizations are important precisely for their small sizes, low retention, or short existences*.
- **Cross-sector career movement.** We examine post-PhD career transitions across sectors. Beyond finding evidence that cross-sector collaboration is increasing, we use R^3 to reveal *significant asymmetry in the frequency, timing, and “prestige” of career moves between academia and industry*.
- **Individual retention prediction.** Finally, we predict career transitions by combining R^3 network dynamics and individual career trajectory information. We demonstrate R^3 's immediate utility in boosting prediction power with *interpretable features that can inform employer recruitment and retention mechanisms*.

This work is a starting point for large-scale studies of computing career trajectories. Such analyses are becoming crucial as demand for computing expertise grows and our world increasingly depends on research innovations in computer science.

Outline. This paper is organized as follows: we first discuss some of our extensive data standardization pipeline and describe our post-processed dataset (Sec. 2). We then motivate and detail our R^3 career network model (Sec. 3). With R^3 we analyze computing research careers at several levels of granularity (Sec. 4). Finally, we outline related areas of work and discuss future directions based on our study's results and limitations (Secs. 5 through 7).

2 DATA

Data collection. To obtain our data, we automatically crawled the public online information of around 10 thousand PhD graduates from the 1970s to 2015 in computer science and related subfields. We matched these graduates from the Proquest Digital Library of PhD dissertations to an online public professional (LinkedIn) profile. To guide automatic data collection, we obtained data for those with PhDs from the top 50 US computer science graduate programs as specified in the 2014 *US News & World Report* (USNWR)¹. We do not use the actual USNWR *rankings*, which have been criticized [2], anywhere in our study. Per person, we retained the PhD school,

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.

KDD '18, August 19–23, 2018, London, United Kingdom

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5552-0/18/08...\$15.00

<https://doi.org/10.1145/3219819.3219863>

¹<https://www.usnews.com/best-graduate-schools/top-science-schools/computer-science-rankings>

(a) Top post-PhD employers.

(b) Pareto cumulative distribution of PhD employment.

Figure 1: Most CS PhDs concentrate at one of a few institutions: post-PhD employment distribution follows a power law. **Figure 2: Fictitious but plausible samples from our dataset.** **Figure 3: First employment sector of PhDs per year.**

graduation year, and all available post-PhD employers, job start/end dates, and job titles. Figure 2 gives two illustrative samples.

While it is not possible to fully verify whether online professional profiles are up-to-date or truthful, we manually validated our data by inspecting individuals' listed employers. We discarded profiles with suspicious employers: for example, overly generic names like "college" or companies with no online records of existence. Moreover, because we collected the profiles of people whose dissertations were verified by ProQuest, it is unlikely that their profiles were fake or set up by fake accounts. After validation, we retained 17,358 unique employment records for 6,781 PhDs over five decades. We make two anonymized versions of the data available at <https://github.com/tsafavi/career-transitions-data>.

Name and sector standardization. Using several publicly available collections of organization names^{2,3}, we created a centralized list of academia, industry, and government organizations. We standardized each organization in our dataset against this list. For instance, all variants of Microsoft—i.e., Microsoft Bing, Microsoft Skype, and Microsoft Research—were grouped under one umbrella. In the case of ambiguity, we automatically mapped names to their most well-known instances in computing, like "CMU" as an accepted acronym for Carnegie Mellon University. Universities without specified campuses mapped to their flagships. For example, the University of Michigan listed without one of Ann Arbor, Flint, or Dearborn became the University of Michigan Ann Arbor.

To categorize employers into sectors (one of industry, academia, or government), we used our centralized organization list, keywords like "LLC" and "college", and rule-based automatic classification. For the 444 organization names that our system failed to categorize,

we provided 6 expert assessors with those organizations and a set of publicly available rules⁴. The inter-rater agreement on a subset of those employer names was 73.6% using Fleiss' kappa [12], which quantifies the degree of inter-rater agreement over that expected by chance. This "substantial" inter-rater agreement [22] demonstrates the relative simplicity and interpretability of our rule-based system.

Description. (1) *Education.* The school with the largest graduate representation in our dataset, at 6.7%, is the University of Illinois Urbana-Champaign (UIUC). Carnegie Mellon, MIT, UC Berkeley, and Stanford follow closely. As expected, PhD production volume increased substantially in the last 15 years (Fig. 3), during which interest in computing became widespread and compute resources and data availability skyrocketed. The rate of computer science PhD graduation between 2003 and 2008 grew on average 18.33% per year until a peak volume in 2008, then dropped briefly around the beginning of the Great Recession. Of these graduates, 11.7% had postdoctoral experience, with an average of 1.13 postdocs per person. The volume of PhD graduates beginning postdocs spiked between 2009 and 2012, again with the onset of the recession.

While representation bias in our dataset is possible, as collecting data on *all* graduating CS PhDs is difficult, our data are corroborated by the Computing Research Association's Taulbee Survey⁵. This survey collects, among other figures, data from US and Canada higher education institutions on graduating PhDs in computer science and related fields. The trend of our PhD production data strongly correlates (Pearson $r = 0.75$) with the Taulbee PhD production numbers publicly available from 2002 to 2015.

(2) *Employment.* The sector distribution of organizations in our dataset is 83.5% industry, 14% academia, and 2.5% government.

²<https://www.crunchbase.com>

³<http://www.nasdaq.com/screening/company-list.aspx>

⁴<http://bit.ly/2ErexBh>

⁵<https://cra.org/resources/taulbee-survey/>

Google is by far the most popular employer (Fig. 1a), with nearly 15% of the entire dataset having worked there at least once since its inception in the late 1990s. The most popular destination in academia, at around 1% of all PhDs in the dataset, is Carnegie Mellon. Like many other well-documented phenomena [6, 10], computer science PhD employment among organizations appears to follow a power-law distribution (Fig. 1b), demonstrating that most computing PhD talent has concentrated in a few companies and universities.

Although a PhD is often considered a gateway to academia, a majority of computer science PhDs in our dataset immediately work in industry. On average, 57% go to industry, 39% go to academia, and 4% go to government per year (Fig. 3). However, while industry jobs are more popular, academic jobs have higher longevity. The mean retention rate for industry employers in our dataset is 4.65 years; for academia, 5.84 years; for government, 4.91 years, with significant differences between academia and the others ($p \ll .00001$ academia/industry, and $p = 0.002$ academia/government, two-sided t -test). While this may be related in part to academic tenure policies, we do consider postdoctoral positions at academic institutions, which are intended to be short, and positions beyond tenured professorships as part of academia here.

3 R^3 TRANSITION NETWORK MODEL

To analyze the evolution of computing research with our unique dataset, we need an employer “desirability” or “importance” measure for computing PhDs. Such a measure quantifies hierarchies between organizations and helps us anchor our analysis around key representative institutions of the profession. For this two components are necessary: (1) a *network representation* that captures the dynamics of career paths; and (2) an *organizational ranking method* that captures both employee influx and outflux.

(1) Network representation. Among the various ways to model trajectory or sequence data [3, 19, 27, 33], a natural first choice is the transition network, which is a directed graph that here captures the post-PhD career transitions between employers (states). This representation is often called an **aggregate flow network** or “talent flow graph” or “job transition/hop network” [17, 26, 32]. In this representation $\mathcal{G}_f(\mathcal{V}, \mathcal{E})$, each node $v \in \mathcal{V}$ is an industry, academia, or government organization. Each directed edge $(u, v, t) \in \mathcal{E}$ is a set of employee transitions from organization $u \in \mathcal{V}$ to organization $v \in \mathcal{V}$ in year t . The weight of edge (u, v, t) , which we denote as $W_f^t(u \rightarrow v)$, captures the total number of employees making a career transition from u to v during year t .

The aggregate flow network $\mathcal{G}_f$ is simple and interpretable. However, it can obscure important insights, which we demonstrate in Sec. 4, for several reasons. For one, our data show that most PhD-trained talent in computing concentrates in very few organizations (Fig. 1b). Ranking organizations by aggregate flow *heavily* favors these organizations, which are mostly large companies, whereas organizational size is but one determinant of importance in the real world. Furthermore, capturing only aggregate transition volume cannot answer important temporal questions encoded in career sequence data. Which organizations have higher turnover than normal? Which are growing quickly relative to their size? Which are desirable for fresh graduates versus senior engineers, distinguished researchers, and program directors? To answer these questions, we

Figure 4: A hypothetical transition network comprising a stable company STABLE-LLC, a university UNI, a declining company DECLINE-LLC, and a fast-growing STARTUP. Each node is labeled with the number of employed PhDs before transitions. The edge weights denote the number of PhDs moving to or from each node.

propose the R^3 **transition network** model. Each R in R^3 transforms $\mathcal{G}_f$ ’s edge weights to capture a specific career dynamic, which we define as *resource transfer* (R_{SRC}), employee *retention* (R_{TN}), and *relative organizational growth* (R_{GR}).

(2) Ranking. PageRank and HITS are two of the most well-known node centrality measures on directed graphs [13]. The former, which has been used in career mining [17, 26] and many other settings [14, 20, 21, 29], ranks nodes by the quantity and “importance” of their in-links, and outputs one set of scores. By contrast, HITS [18] outputs two sets of scores, one for **hubs** and one for **authorities**. Hubness measures each node’s “indexing power” by the number and strength of its outgoing links to authority nodes. Authority measures each node’s “relevance” by the number and strength of its incoming links from hub nodes.

While the R^3 model proposed in this section can work with PageRank, HITS, and any other link analysis algorithm on directed weighted graphs, we design it with HITS in mind. In career analysis, both the in- and out-links of nodes, which respectively capture organizational influx and outflux, characterize employer roles and rankings in the flow graph. For this reason we posit that **identifying both hubs and authorities best captures the natural meaning of career transitions**, and the full dynamics of transitions from an organizational perspective. Intuitively, authority organizations attract talent and expertise from hub organizations. For brevity, we do not cover the theory of HITS (see [18] for details).

In the remainder of this section we define R^3 and demonstrate its effects on a simple but plausible example (Fig. 4). As job transition frequencies have been shown to follow a yearly cyclic pattern [32], from here on we assume yearly time units.

3.1 R_{SRC} : Modeling resources

Our first feature, **resources** (R_{SRC}), captures **the level of cumulative employee expertise** in inter-organization transitions. The intuition of R_{SRC} is that a longer career leads to more advanced individual expertise and organizational value. For example, one

might rise from a software engineer to a directorate role, or else from assistant to full professorship, with time. Our goal with R_{SRC} in terms of HITS is to capture organizational hubs and authorities for experienced people, who are skilled *resources*.

To quantify each employee's expertise level, we use a variant of the logistic skill-gain model from economics and organizational theory [35]. In more detail, as shown in Fig. 6, we model the expertise level of a PhD p making a career transition in year t as a sigmoid function of her career length up to that year, $\ell(p, t)$:

$$R_{SRC}(p, t) = \left(1 + \exp\left[-\frac{\ell(p, t) - \bar{\ell}(t)}{\alpha}\right] \right)^{-1}.$$

In the above formulation, α controls the curve's steepness and $\bar{\ell}(t)$, the sigmoid midpoint, is the system-wide average career length at year t (10 years in Fig. 6). R_{SRC} thus scores each transitioning PhD based on her experience relative to her peers. The transitions of those who entered the system earlier are deemed more valuable for the source and target organizations, although in our examples and analyses we set $\alpha = \bar{\ell}(t)/2$ (the least steep curve in Fig. 6, orange) to avoid over-penalizing those with fewer years of experience.

We transform each directed edge of the aggregate flow network $\mathcal{G}_f$ to concentrate flow in the graph around the movement of experienced people:

$$\begin{aligned} W_{R^3}^t(u \rightarrow v) &= W_f^t(u \rightarrow v) \cdot \overline{R_{SRC}(u, v, t)} \\ &= \sum_{\text{PhD } p: u \rightarrow v | t} R_{SRC}(p, t), \end{aligned} \quad (1)$$

where $\overline{R_{SRC}(u, v, t)}$ is the average R_{SRC} score of employees p moving from the source node u to the target node v during year t , each denoted as PhD $p : u \rightarrow v | t$ above. Note that our logistic model does not account for skill loss over time. While "productivity decline" in academia has been studied for research publication rates over time, among other phenomena, this narrative has been recently questioned [31]. As such we do not include it in our model.

Example. Assume that the few people transitioning to and from UNI (green, Figs. 4 and 5) are distinguished professors and industry leaders with 20 years of experience. Also, assume all others among STABLE-LLC, DECLINE-LLC, and STARTUP have 5 years of experience. Given a system-wide average career length of 10 years, transforming $\mathcal{G}_f$ with R_{SRC} results in UNI's authority score increasing from 0.12 to 0.31, reflecting its centrality as an employer of highly skilled people.

Figure 5: R_{SRC} effects.

3.2 R_{TN} : Modeling retention

Retention (R_{TN}) captures how well organizations retain talent, which has been shown to be crucial in career transition graphs [17]. Indeed, inter-employer transitions alone are comparatively sparse, with only 22% of our dataset transitioning on average per year. Our motivation for R_{TN} is that low retention may signify a variety of real-world meanings in organizations, from undesirability to a

Figure 6: Computing R_{SRC} for UNI's transitions (Fig. 4), assuming a system-wide average career length of 10 years.

short existence (i.e., startups that fail or get acquired quickly). In the context of HITS, we use R_{TN} to identify low-retention hub organizations that serve as "stepping stones" to other authorities.

Since employers with higher retention are better able to develop their employees' job-specific skills, we model retention by capturing organizational "expertise" on a sigmoid curve. To first account for significant differences in sector retention rates (Sec. 2), we stratify organizations by sector. We then model the retention of an organization v at year t as a sigmoid function, comparing v 's average PhD retention rate at year t , $\ell(v, t)$, to its sector's current average PhD retention rate $\ell(\sigma(v), t)$, where $\sigma(v)$ is v 's sector:

$$R_{TN}(v, t) = \left(1 + \exp\left[-\frac{\ell(v, t) - \ell(\sigma(v), t)}{\beta}\right] \right)^{-1}.$$

The idea here is that employers with higher-than-average PhD retention in their sector receive a higher R_{TN} score, and vice versa, although as before we smooth the curve by setting $\beta = \ell(\sigma(v), t)/2$. Our goal with R_{TN} is to capture hubs, so we transform *outgoing* edges for each node v in $\mathcal{G}_f$ with $1 - R_{TN}(v, t)$:

$$W_{R^3}^t(v \rightarrow u) = W_f^t(v \rightarrow u) \cdot (1 - R_{TN}(v, t)). \quad (2)$$

Taking the converse increases outflux from low-retention employers and decreases outflux from high-retention employers.

Example. Returning to Fig. 4, assume that STABLE-LLC's average retention matches the industry average; DECLINE-LLC's is around one-half the industry average due to its decline; and UNI's is twice the academia average due to its prestige and tenure policies. By transforming the flow network $\mathcal{G}_f$ with $1 - R_{TN}$, DECLINE-LLC's hub score increases to 0.91 and UNI's hub score drops to 0.01, magnifying the respective retention abilities of these institutions.

Figure 7: R_{TN} effects.

3.3 R_{GR} : Modeling relative growth

Our last feature, **relative growth (R_{GR})**, quantifies growth relative to organization size. The goal of R_{GR} is to boost the authority of small, fast-growing organizations. Its HITS interpretation is that employers with high R_{GR} , like buzzworthy startups or fast-growing university computer science departments, should gain authority even with relatively low influx.

Figure 9: Relative growth of DECLINE-LLC and STARTUP.

Extending the literature in ecology and stock analysis on growth rate modeling [16, 23], we model an organization v 's relative growth at year t as the difference between the logarithms of v 's PhD influx and outflux at year t . We normalize this difference by the number of PhDs working at v during year t before in- or out-transitions:

$$R_{GR}(v, t) = \frac{\log(\# \text{ PhDs joining } v) - \log(\# \text{ PhDs leaving } v)}{\log(\# \text{ PhDs at } v) + 1} = \frac{\log(\sum_u W_f^t(u \rightarrow v) + 1) - \log(\sum_w W_f^t(v \rightarrow w) + 1)}{\log(W_f^t(v \rightarrow v) + 1) + 1},$$

where smoothing is used to address noise and correct for zeroes. R_{GR} is oriented toward fast-growing authorities, so we apply this value to *incoming* edges of nodes in $\mathcal{G}_f$. Since we ultimately want a value between 0 and 1 to retain the same edge-weighting scale as R_{SRC} and R_{TN} , we transform edges with a normalized exponential function of R_{GR} :

$$W_{R^3}^t(u \rightarrow v) = \frac{W_f^t(u \rightarrow v) \cdot \exp[\gamma \cdot R_{GR}(v, t)]}{\max_{v,t}(\exp[\gamma \cdot R_{GR}(v, t)])}. \quad (3)$$

In the above formulation, γ controls the steepness of the exponential growth curve (Fig. 8). For our examples and analyses, we weight the curve with $\gamma = 1.5$, chosen by cross-validation to slightly boost influx toward fast-growing organizations without over-valuing them. The denominator in Eq. (3), which ensures that each transformation is between 0 and 1, normalizes by the maximum exponential growth observed during year t .

Example. Returning to our example scenario, we transform $\mathcal{G}_f$ with the normalized exponential growth function of R_{GR} . STARTUP's incoming edge weights are magnified (Fig. 8) and its authority increases from 0.22 to 0.53, overtaking STABLE-LLC. The latter's hub score also increases from 0.13 to 0.43 as several of its employees transition to fast-growing authority STARTUP.

Figure 8: R_{GR} effects.

3.4 Unifying R_{SRC} , R_{TN} , and R_{GR}

Thus far we have demonstrated each R 's separate effects. However, ultimately we apply all three R 's on the same graph in succession to obtain a *single* set of hubs and authorities rather than different rankings per R . In our running example, the unified R^3 scores, which reflect the different temporal career dynamics captured by R^3 , are shown in the table of Fig. 10. Under the unified

Figure 10: HITS scores from Fig. 4 placed in the three separate dimensions of R^3 (points), and unified (table). The hubs and authorities reflect the nuances of the system's dynamics better than in Fig. 4.

model, the largest employer, STABLE-LLC, no longer dominates the authorities. Instead, authority is distributed evenly among the three organizations STABLE-LLC, UNI, and STARTUP, reflecting the unique importance of each organization in the system. While DECLINE-LLC still dominates the hubs, STABLE-LLC's hub score also increases because it is a "stepping stone" to the fast-growing authority STARTUP. Overall, these changes allow for more meaningful analysis of the system's dynamics because they reflect different kinds of real-world organizational importance.

4 R^3 -DRIVEN ANALYSIS

With R^3 defined, we analyze our unique career trajectory dataset and demonstrate R^3 's nuance and versatility. To do so we: (1) explore the diverse kinds of institutional "importance" that R^3 captures in computing research; (2) use R^3 to characterize significant asymmetry in cross-sector career transitions; and (3) build a strong predictive model with R^3 to inform organizational retention.

4.1 System-wide evolution

How do the R 's in R^3 affect organization rankings on the career transition network $\mathcal{G}_f$, and what insights can we gain? Here we use our model to understand different ways organizations became "important" throughout computing research history.

Methodology. In five-year intervals from 1980 to 2015, we obtain HITS hub and authority rankings for all nodes on the $\mathcal{G}_f$ and R^3 career transition networks. We regress the R^3 rankings against the $\mathcal{G}_f$ rankings and identify the nodes with the highest standard error. These nodes adhere least to the hypothesis that R^3 transformations lead to identical HITS scores as $\mathcal{G}_f$. They thus capture the most R^3 -specific information, their updated rankings a result of the dynamics captured by R^3 . Figure 11 displays several "important" historical moments that only became visible after applying the R^3 model. In these moments, the shown organizations' rankings changed significantly with the R^3 transformations. For brevity, here we only cover a small selection of organizations, each of which was outside the top-50 hubs or authorities using HITS on $\mathcal{G}_f$. After applying R^3 , each of these organizations moved up at least 10 ranks and subsequently ranked as a top-50 hub or authority.

Figure 11: R^3 hubs and authorities (bolded) with their respective historical context. The shown organizations’ important historical moments were only revealed after using HITS on the R^3 transition network as opposed to the aggregate flow network. Each of these organizations became a top-50 hub or authority after applying R^3 to G_f .

Results. (1) *Hubs.* The R^3 network exposes a variety of organizations, some well-known and some relatively obscure, that were important hubs for different reasons. The hub timeline in Fig. 11 first shows the supercomputer company Kendall Square Research (KSR). Founded in the late 1980s and bankrupt by 1994, KSR’s short existence and subsequent low retention (R_{TN}) contributes to its increased hub ranking, moving from the 96th-ranked hub on G_f to the 38th on R^3 . A few years later, R^3 designates a host of small companies—Valence Research Inc, Softway Systems, VXTreme Inc, and Vermeer Technologies—as top-20 hubs. The common theme here is that they all existed for a short time (in consequence, low R_{TN}) before acquisition by Microsoft, their employees thus transitioning to one of the top authorities of the era.

In the mid to late 2000s, we observe an increased out-flow of experienced researchers (R_{SRC}) from several government and industry organizations. Adobe, VMware, Disney Research, and NASA become top-20 hubs, which may be in part related to the global economic crisis. According to Wikipedia, in 2008 Adobe laid off 8% of its workforce and VMware fired its CEO after disappointing financial performance. NASA’s funding cuts in the 2000s are also well-documented [11]. The most recent discovered hubs shown are, like the startups acquired by Microsoft earlier, small short-lived companies that were led by CS PhDs and bought by large companies. Figure 11 shows Katango and Timeful, two recent acquisitions by Google that moved from rank 100+ to top-30 R^3 hubs. Table 1 highlights ITA Software, also acquired recently by Google.

(2) *Authorities.* The HITS authority rankings on R^3 also lead to interesting discoveries not captured by G_f . In the late 80s, R^3 designates Johns Hopkins University as the first-ranked authority, which is unusual because universities naturally have lower in-flow. We found that computer science at Johns Hopkins officially became a department in 1986, whereas before computing had been part of statistics and operations research. This naturally led to an increased growth rate (R_{GR}) and an influx of academics (R_{SRC}) to a school that, prior to 1987, had no representation in our dataset.

In the 90s, EMC (later bought by Dell) becomes a top-20 authority around the time of its rapid growth to billions of dollars in revenue. This is reflected by its sudden large representation in our dataset in 1997 (R_{GR}). R^3 next captures the R_{SRC} flow to Avaya via its spinoff from the telecommunications giant Lucent in 2000, as well

as Google’s rise. Indeed, R^3 “discovers” Google before its mid-00s representation increase in our dataset: its authority ranking moves from 45th on G_f to 4th on R^3 in the late 90s. In the following decade, R^3 rankings differ even further from G_f as experienced PhDs (R_{SRC}) moved from established large hubs like Microsoft, Google, and IBM to small, fast-growing “unicorn” startups (R_{GR}). Twitter, Dropbox, Snapchat, Square, and Uber all move from rank 100+ to the top-50 in or after the late 00s. Others of the same vein include Baidu, Netflix, Light, Magic Leap, Databricks, and the nonprofit Allen Institute for Artificial Intelligence (AI2).

(3) *Top-10 rankings.* Beyond observing the organizations for which rankings changed the most, we report the top hubs and authorities for two 10-year segments in Table 1. The bolded names appear only in the R^3 top 10. We note the mix of organizations across sectors and sizes—some that persisted or rose in rankings, some that dropped or disappeared altogether, some that moved from top-ranked authorities to top-ranked hubs—throughout the years. R^3 is thus capable of capturing various meanings of organizational “importance” despite differences in sector, size, and hiring volume. That said, while R^3 shows that different organizations can gain importance in their own rights, these rankings are of course subjective. This is especially true in an opportunity-rich field where “importance” is highly dependent on individual goals and interests.

Table 1: Top employer hub and authority rankings for two time windows. The names in bold are those that only appear in their respective top 10s in the R^3 transition network.

	1995–2005		2005–2015	
	Hub	Auth.	Hub	Auth.
1	IBM	IBM	Microsoft	Google
2	Microsoft	Intel	Yahoo	Microsoft
3	Intel	Google	Intel	IBM
4	Bell Labs	Microsoft	Google	Intel
5	Compaq	Siemens	Siemens	Facebook
6	Sun	EA	ITA Software	Yahoo
7	HP	CMU	Amazon	Snaphchat
8	NASA	MIT	Apple	Qualcomm
9	Disney	Ask.com	IBM	MIT
10	Docomo	HP	HP	Amazon

Figure 12: Transitions *from* industry are usually made to higher-ranked organizations.

4.2 Cross-sector career movement

We next study career movement across sectors. In investigating what R^3 reveals about the “prestige” associated with cross-sector transitions, we discover insights into *how* people transition between employment sectors and *when* they make these transitions.

Methodology. Given a year t and an organization v , we obtain v ’s R^3 HITS rankings for the five-year interval preceding t . We do this to obtain v ’s most relevant rankings. As we showed in the previous section, institutional importance in computing changes quickly, and rankings from more than a few years ago may not be relevant. Moreover, due to relatively few transitions per year, a yearly granularity is not appropriate for our analysis.

Note in this section that we consider several thousand organizations in ranking, unlike most well-known ranking systems (i.e., university rankings) that only consider a few hundred institutions. In accordance with the size and complexity of the computing research professional system, R^3 ranking differences between organizations naturally vary more than ranking differences between organizations in smaller-scale studies.

Results. A common narrative in computing research careers is that it is easier to transition to industry from other sectors—in particular, from academia—than vice-versa due to differences in factors like salary and work-life balance. Without making any claims about the causes of this phenomenon, we do find that our data uphold the narrative. Out of all cross-sector career transitions in our dataset (16.3% of all transitions), nearly two-thirds are to industry.

That said, PhDs transitioning *from* industry appear to gravitate toward more “prestigious” institutions than their current industry employers. Without transforming $\mathcal{G}_f$ with R^3 , a PhD transitioning from industry to academia or government moves up on average 34 ranks in HITS authority rankings, whereas a PhD transitioning to industry from academia or government moves *down* on average 47 ranks (we exclude postdocs from these analyses). Indeed, around 15% of all PhDs leaving industry in our dataset go to the highly-ranked Stanford, UC Berkeley, MIT, and Carnegie Mellon. Moreover, PhDs transitioning *to* industry often transition from top-ranked schools to startups, which are naturally ranked lower. Even those moving up in rankings do not usually move up significantly because of their school’s already-high rankings.

The R^3 transformations magnify these differences. A career transition from industry in the R^3 network results in an average authority ranking gain of 127 places. This is partially because the resource factor (R_{SRC}) captures the imbalance of experience level in these transitions. The average career lengths of PhDs who transition

Figure 13: Hard/soft transition examples.

Figure 14: Side projects are becoming more common: percentage of yearly transitions that were “soft”, 1995 to 2015 (fit line slope $m=0.180$; p -value=0.037).

to and from industry—2.6 years versus 8.23 years respectively—are significantly different ($p \ll .00001$, two-sided t -test). This suggests that PhDs in industry tend to establish themselves first before leaving it, whereas those in academia and government more often move to industry at earlier stages in their careers.

Notably, about 1 in 10 PhDs making a career transition from industry go to consulting or advising, positions often taken up in parallel with other engagements. Indeed, inter-organization movement is fluid, made complex by the diverse and myriad opportunities in computing for collaboration. To capture this, we further categorize career transitions into **hard transitions**, made when an employee *leaves* one organization before joining another, and **soft transitions**, made when an employee joins an organization *without* leaving her previous employer (Fig. 13). Soft transitions make up around 21% of all transitions in our dataset. Such transitions have been increasing slowly, with some evidence of a linear upward trend since the mid-90s (Fig. 14). This suggests that multiple venues of professional engagement via side projects, collaborations, and startups are becoming more common.

With these definitions, we revisit cross-sector career transitions in more detail. Soft transitions *from* industry are quite common. Over one-third of all transitions made from industry are soft, moving *up* on average 64 authority ranks on $\mathcal{G}_f$ and 174 ranks on R^3 . This suggests that PhDs tend to cross over to prestigious academia and government institutions without fully leaving industry. By contrast, fewer than 1 in 10 transitions *to* industry are soft, moving *down* on average 133 authority ranks on $\mathcal{G}_f$ and 248 ranks on R^3 . Nearly half of these rare soft transitions to industry are startup-related, involving a professor taking on a chief role in a budding company. Two real-world examples are the professor-led startups Timeful and Katango (Fig. 11), which became top hubs in the 2010s.

Table 2: Prediction performance metric averages and standard deviations per value of n , highest two per category bolded.

	Features	$n = 1$	$n = 2$	$n = 3$	$n = 4$	$n = 5$
AUC	IND	0.625 $\pm$ 0.00	0.637 $\pm$ 0.00	0.654 $\pm$ 0.01	0.644 $\pm$ 0.01	0.656 $\pm$ 0.02
	IND + $\mathcal{G}_f$	0.639 $\pm$ 0.01	0.660 $\pm$ 0.02	0.666 $\pm$ 0.02	0.658 $\pm$ 0.03	0.663 $\pm$ 0.03
	IND + $\mathbf{R}^3$	0.656 $\pm$ 0.01	0.675 $\pm$ 0.02	0.677 $\pm$ 0.02	0.665 $\pm$ 0.02	0.670 $\pm$ 0.03
	ALL	0.649 $\pm$ 0.01	0.668 $\pm$ 0.02	0.674 $\pm$ 0.02	0.664 $\pm$ 0.02	0.669 $\pm$ 0.03
F1	IND	0.357 $\pm$ 0.05	0.459 $\pm$ 0.01	0.536 $\pm$ 0.01	0.574 $\pm$ 0.01	0.601 $\pm$ 0.04
	IND + $\mathcal{G}_f$	0.396 $\pm$ 0.00	0.473 $\pm$ 0.01	0.542 $\pm$ 0.00	0.577 $\pm$ 0.01	0.601 $\pm$ 0.04
	IND + $\mathbf{R}^3$	0.404 $\pm$ 0.01	0.488 $\pm$ 0.01	0.549 $\pm$ 0.01	0.576 $\pm$ 0.01	0.595 $\pm$ 0.03
	ALL	0.398 $\pm$ 0.01	0.488 $\pm$ 0.01	0.550 $\pm$ 0.00	0.578 $\pm$ 0.01	0.610 $\pm$ 0.03

Figure 15: Feature importances across all features (ALL) for employee transition prediction with $n = 1$ and $n = 5$.

Overall we observe not-insignificant cross-sector movement, including increasingly popular soft transitions, which highlights the field’s various opportunities for collaborations and suggests an increasing connectedness between sectors. We also observe significant asymmetry in the frequency, timing, and rankings of moves to and from industry, magnified by $\mathbf{R}^3$ ’s emphasis on experience. While our data alone cannot point to the *cause* of this asymmetry, a possible explanation is that because soft transitions are less risky or permanent, they allow for bigger leaps, both for those entering highly-ranked academic institutions and those leaving academia for high-risk, high-reward ventures like startups.

4.3 Individual retention prediction

Finally, we use *individual* career trajectories, which appear as ordered sequences of organizations ($v_1^t, \dots, v_n^t$) in the transition network, along with the $\mathbf{R}^3$ network dynamics previously discussed, to predict future employee transitions. Here we show that $\mathbf{R}^3$ boosts the performance of a model with strong predictive power, adding important and interpretable features that can inform organizations seeking to recruit or retain computing PhDs.

Methodology. Given a year t , can we predict which PhDs p will make a career transition within the next n years? We group our prediction features (40 total) into 3 categories:

- **IND:** individual features about p ’s career trajectory. These include: number of years since graduation; career length in years; number of employers total; average number of years spent with an employer; total years at current employer; number of jobs held in/outside of industry; number of inter-sector transitions; the sectors of p ’s first and most current employers; number of hard transitions; number of soft transitions; number of postdocs

done; and whether p ’s most current job is senior-level, founder or CEO, professor, researcher, engineer, and/or visiting.

- **$\mathcal{G}_f$:** features from the aggregate flow career transition network. These include: the HITS rankings and scores of p ’s current employer v , taken from transitions in the 5-year interval up to and including t ; and the HITS rankings and scores of v , taken from transitions in the 5-year interval up to and including the year that p started working at v .
- **$\mathbf{R}^3$:** features from our proposed career transition model. These include: the resources (R_{SRC}) score of p at year t ; the retention (R_{TN}) score of p ’s current employer v at year t ; v ’s relative growth (R_{GR}) score during year t ; and the HITS rankings and scores as with the $\mathcal{G}_f$ features, but computed on the $\mathbf{R}^3$ network.

For those who currently hold more than one job (i.e., soft transitioned to a second job while holding a first), we consider the job started earlier as the “most current”, since usually soft transitions and side projects occur after full-time employment has already begun. As we are interested in predicting immediate career transitions in the modern-day professional system, we vary $n \in [1, 5]$, predicting whether each person will transition between 1 and 5 years “in the future”, and predict for years $t \in (2000, 2010)$.

For our prediction tasks, the average ratio of positive labels per value of n is 23.5%, 32%, 38%, 43%, and 46.5% for $n \in [1, 5]$ respectively. For all tasks, we train a gradient boosting tree classifier from the open-source XGBoost library, performing a grid search over the learning rate (eta), training data subsampling (subsample), and label balance (scale_pos_weight) parameters to handle the label imbalance for smaller values of n . We report performance metrics averaged over 10-fold cross-validation in Table 2 for the **IND** features alone, the **IND** and $\mathcal{G}_f$ features together, the **IND** and $\mathbf{R}^3$ features together, and all groups of features (**ALL**).

Results. The results make it clear that system-wide network dynamics substantially boost prediction performance, justifying our initial choice of using HITS on a career transition network. Moreover, $\mathbf{R}^3$ adds extra power to the model. The top feature groups by performance are $\mathbf{IND} + \mathbf{R}^3$ and $\mathbf{ALL}$, which do about as well each other. While AUC remains relatively stable, the higher class imbalance makes the prediction tasks more difficult for smaller values of n , highlighting $\mathbf{R}^3$'s strength in improving F1.

Using XGBoost's built-in feature importance tools, we also found that $\mathbf{R}^3$ features were consistently considered "important" for prediction, and especially so for lower values of n compared to the other feature groups. The top 10 most important features across all feature groups for $n = 1$ and $n = 5$ years are given in Fig. 15. We immediately observe that all of the most important features beyond $\mathbf{IND}$ are related to an organization's hub scores and ranks. Most interestingly, we find that for smaller values of n , the $\mathbf{R}^3$ hub ranks and scores of p 's employer *at the time they started working there* are most important, which suggests that employer outflux rate, retention, and volume are about as predictive of short-term retention as individual-level information.

By contrast, for higher values of n the $\mathbf{IND}$ features become more informative. The most important feature is how many years individuals have worked at their current employers. This is intuitive given the fast-paced nature of the computing profession and the fact that those who have worked at an employer longer are more likely preparing to transition. This is especially true in industry, where the employment length mean and variance are shorter. Indeed, sector-related features become important for higher values of n (`num_jobs_not_in_industry`). Whether or not a person is currently in a visiting position also becomes important, since visiting positions are often limited in duration.

We conclude by noting that while $\mathcal{G}_f$ features rarely appear in the top 10 most important features, $\mathbf{R}^3$ features are consistently important, and R_{GR} is one of the top features for $n = 1$. These results show the immediate utility of $\mathbf{R}^3$ for retention prediction, which becomes important from an organizational perspective as the demand for computational expertise increases.

5 RELATED WORK

Career path mining. Mining professional career paths has attracted recent interest. One of the first studies to mine career trajectories proposes a similarity measure between professional profiles using temporal sequence alignment on user career paths [34]. More recently, Xu et al. [32] detect "talent circles" in job transition networks to find qualified candidates for jobs, and Kapur et al. [17] apply PageRank on career transition networks as an intermediary step for ranking and recommending universities. Unlike [32] and [17], our goals are not recommendation-oriented. Moreover, while we are not the first to design a weighted HITS scheme [8, 30], we are the first, to the best of our knowledge, to propose one for career trajectory mining.

Academic career trajectories. Most work in academic career trajectory analysis concerns career movement *within* academia. For example, Clauset et al. [5] find that academic prestige correlates with higher productivity and better faculty placement, and Deville et al. [9] find that transitions between academic institutions are

influenced by career stage and geographical proximity. That said, a few studies on the distribution of PhD graduates between academia and industry exist. For example, Sauermann and Roach [28] find in a survey that most students' career preferences shift from academia toward industry over the course of a PhD, and Balsmeier and Pellens [1] consider how and why scientists *leave* academia. By contrast, we are interested in *all* directions of cross-sector career movement (Sec. 4), not just transitions from academia. Moreover, our study is long-term, following individuals beyond their PhDs, and observational, as we do not use individually-reported preferences to explain the causes of the phenomena we observe.

Publication trajectories in computing. Several recent works study computing career "trajectories" in terms of *publishing* productivity or *citation* counts. For example, Way et al. [31] study faculty "productivity trajectories" in computer science, providing evidence that publishing trends in the field do not follow the conventional "early peak and gradual decline" narrative. Most recently, Chakraborty and Nandi [4] study scientific "success trajectories" in computer science and physics by analyzing paper citation counts. Similar to [31], they question established notions of scientific success, finding multiple distinct trajectories of successful scientific papers beyond the "early rise" trajectories of immediately impactful papers. These works come from the larger body of research devoted to study of bibliographic data [15, 25], which we do not consider here (see discussion).

6 DISCUSSION

Our data-driven study of long-term computer science PhD employment dynamics is the first of its kind. Naturally, many directions for future work remain. One such direction is increasing the scope of our study in terms of data. Many important people in computing research obtained PhDs outside the US. Some did not obtain a PhD at all. Although we only considered PhD graduates from a subset of schools to ensure the accuracy of our data, an ideal dataset would include those who made contributions to computing research regardless of degree or background.

A related direction is that of merging bibliographic data with existing career trajectory data. This task is challenging for large datasets due to the difficulty of entity resolution across databases, which in our case amounts to matching online professional profiles with Google Scholar or DBLP profiles. However, such data would address questions never before answered: How do publishing rates compare across sectors? Do "impactful" authors concentrate in "important" institutions? Is a person's publishing history predictive of their future career transitions?

Future studies could also perform further data validation. One inherent limitation of our study is that little standardized data on post-PhD careers exist, and the data that *do* exist are hard to verify. These concerns are not unique to our study, but they are important. A future larger-scale study could consider multiple levels of automatic and manual data validation using online CVs, resumes, surveys, and/or news articles as available.

A final interesting direction is comparison of different sub-groups in computing research, for example of continents or countries, those with or without a postdoc, and historically underrepresented groups in computing. Examining such group-specific differences could lead

to actionable organization- and individual-level insights, although such analyses would require self-reporting of gender, race, country of origin, etc. Again, this direction reduces to the problem of gathering reliable large-scale data. We hope this will become easier in the future as the topic we study gains traction.

7 CONCLUSION

In this work we examine the career transitions and trajectories of computer science PhDs on the individual, organizational, and sector levels. We propose R^3 , a versatile model for temporal career network dynamics. Using the HiTS link analysis algorithm in conjunction with our R^3 model, we:

- Provide new insights into the meaning of institutional “importance” in computing research careers;
- Reveal a significant asymmetry, from several perspectives, between post-PhD career moves to and from industry; and
- Demonstrate R^3 ’s immediate utility in supporting prediction of individual career transitions.

While our study is the first of its kind, we conclude by emphasizing that these analyses may be repeated in the future when more data, made possible by increasing PhD graduation rates, are available. As computing research continues to grow in importance and worldwide presence, this is certain to happen.

ACKNOWLEDGEMENTS

The authors thank Rada Mihalcea, Jiongsheng Cai, Shuo Chen, Chengkai Hu, Cole Hudson, Natalia Jenuwine, Jennings Jin, Yinning Wong, and Jiabin Zhu for their early involvement in this project their valuable feedback and contributions.

This material is based upon work supported by the National Science Foundation under Grant No. IIS 1743088, an Adobe Digital Experience research faculty award, a Google scholarship, the University of Michigan, and the GENI project. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or other funding parties. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

- [1] Benjamin Balsmeier and Maikel Pellens. 2014. Who makes, who breaks: Which scientists stay in academe? *Economics Letters* 122, 2 (2014), 229–232.
- [2] Michael N. Bastedo and Nicholas A. Bowman. 2010. U.S. News & World Report College Rankings: Modeling Institutional Effects on Organizational Reputation. *American Journal of Education* 116, 2 (2010), 163–183.
- [3] Ivan Brugere, Brian Gallagher, and Tanya Y. Berger-Wolf. 2018. Network Structure Inference, A Survey: Motivations, Methods, and Applications. *ACM Comput. Surv.* 51, 2, Article 24 (April 2018), 39 pages.
- [4] Tanmoy Chakraborty and Subrata Nandi. 2018. Universal trajectories of scientific success. *Knowledge and Information Systems* 54, 2 (2018), 487–509.
- [5] Aaron Clauset, Samuel Arbesman, and Daniel B. Larremore. 2015. Systematic inequality and hierarchy in faculty hiring networks. *Science Advances* 1, 1 (2015).
- [6] Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman. 2009. Power-Law Distributions in Empirical Data. *SIAM Rev.* 51, 4 (2009), 661–703.
- [7] Tim Cross. 2016. After Moore’s law. (2016). <http://www.economist.com/technology-quarterly/2016-03-12/after-moores-law>
- [8] T Deguchi, K Takahashi, H Takayasu, and M. Takayasu. 2014. Hubs and authorities in the world trade network using a weighted HiTS algorithm. *PLoS One* 9, 4 (2014).
- [9] Pierre Deville, Dashun Wang, Roberta Sinatra, Chaoming Song, Vincent D Blondel, and Albert-Laszlo Barabasi. 2014. Career on the Move: Geography, Stratification, and Scientific Impact. *Scientific Reports* 4 (2014).
- [10] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. 1999. On Power-law Relationships of the Internet Topology. *SIGCOMM Comput. Commun. Rev.* 29, 4 (Aug. 1999), 251–262.
- [11] Robert G Ferguson. 2013. NASA’s First A. The NASA History Series.
- [12] Joseph L. Fleiss. 1971. Measuring Nominal Scale Agreement Among Many Raters. 76 (11 1971), 378–.
- [13] Lise Getoor and Christopher P. Diehl. 2005. Link Mining: A Survey. *SIGKDD Explor. Newsl.* 7, 2 (Dec. 2005), 3–12.
- [14] David F. Gleich. 2015. PageRank Beyond the Web. *SIAM Rev.* 57, 3 (2015), 321–363.
- [15] Jorge E Hirsch. 2005. An index to quantify an individual’s scientific research output. *Proceedings of the National academy of Sciences of the United States of America* 102, 46 (2005), 16569.
- [16] William Hoffmann and Hendrik Poorter. 2002. Avoiding Bias in Calculations of Relative Growth Rate. 90 (2002), 37–42.
- [17] Navneet Kapur, Nikita Lytkin, Bee-Chung Chen, Deepak Agarwal, and Igor Perisic. 2016. Ranking Universities Based on Career Outcomes of Graduates. In *ACM KDD*. 137–144.
- [18] Jon M. Kleinberg. 1999. Authoritative Sources in a Hyperlinked Environment. *J. ACM* 46, 5 (Sept. 1999), 604–632.
- [19] Danaï Koutra, Paul N. Bennett, and Eric Horvitz. 2015. Events and Controversies: Influences of a Shocking News Event on Information Seeking. In *WWW. International World Wide Web Conferences Steering Committee*, 614–624.
- [20] Danaï Koutra, Abhilash Dighe, Smriti Bhagat, Udi Weinsberg, Stratis Ioannidis, Christos Faloutsos, and Jean Bolot. 2017. PNP: Fast Path Ensemble Method for Movie Design. In *KDD. ACM*, 1527–1536.
- [21] Danaï Koutra and Christos Faloutsos. 2017. Individual and Collective Graph Mining: Principles, Algorithms, and Applications. *Synthesis Lectures on Data Mining and Knowledge Discovery* 9, 2 (2017), 1–206.
- [22] J Landis and Gary G. Koch. 1977. The Measurement Of Observer Agreement For Categorical Data. 33 (04 1977), 159–74.
- [23] Gautier Marti, Frank Nielsen, Mikolaj Binkowski, and Philippe Donnat. 2017. A review of two decades of correlations, hierarchies, networks and clustering in financial markets. (03 2017).
- [24] Cade Metz. 2017. Tech Giants Are Paying Huge Salaries for Scarce A.I. Talent. (2017). <https://www.nytimes.com/2017/10/22/technology/artificial-intelligence-experts-salaries.html>
- [25] Mark EJ Newman. 2009. The first-mover advantage in scientific publication. *EPL (Europhysics Letters)* 86, 6 (2009), 68001.
- [26] Richard J Oentaryo, Xivier Jayaraj Ashok, Ee-Peng Lim, and Philis Kokoh Praseetyo. 2017. On Analyzing Job Hop Behavior and Talent Flow. In *ICDM Data Science for Human Capital Management Workshop*.
- [27] T. Safavi, C. Sripada, and D. Koutra. 2017. Scalable Hashing-Based Network Discovery. In *ICDM. IEEE*, 405–414.
- [28] Henry Sauermann and Michael Roach. 2012. Science PhD Career Preferences: Levels, Changes, and Advisor Encouragement. *PLOS ONE* 7, 5 (05 2012), 1–9.
- [29] Hanghang Tong, Christos Faloutsos, and Jia-Yu Pan. 2006. Fast Random Walk with Restart and Its Applications. In *ICDM. IEEE Computer Society*, 613–622.
- [30] Vincent S. Tseng, Jia-Ching Ying, Che-Wei Huang, Yimin Kao, and Kuan-Ta Chen. 2015. FrauDetector: A Graph-Mining-based Framework for Fraudulent Phone Call Detection. In *ACM KDD*. 2157–2166.
- [31] Samuel F. Way, Allison C. Morgan, Aaron Clauset, and Daniel B. Larremore. 2017. The misleading narrative of the canonical faculty productivity trajectory. *PNAS* 114, 44 (2017), E9216–E9223.
- [32] Huang Xu, Zhiwen Yu, Jingyuan Yang, Hui Xiong, and Hengshu Zhu. 2016. Talent Circle Detection in Job Transition Networks. In *ACM KDD*. 655–664.
- [33] Jian Xu, Thanuka L. Wickramaratne, and Nitesh V. Chawla. 2016. Representing higher-order dependencies in networks. *Science Advances* 2, 5 (2016).
- [34] Ye Xu, Zang Li, Abhishek Gupta, Ahmet Bugdayci, and Anmol Bhasin. 2014. Modeling Professional Similarity by mining Professional Career Trajectories. In *WWW*.
- [35] Louis E. Yelle. 1979. The Learning Curve: Historical Review and Comprehensive Survey. *Decision Sciences* 10, 2 (1979), 302–328.