

RESEARCH ARTICLE

Automatically assembling a full census of an academic field

Allison C. Morgan^{1*}, Samuel F. Way¹, Aaron Clauset^{1,2,3}

1 Department of Computer Science, University of Colorado, Boulder, CO, United States of America,

2 BioFrontiers Institute, University of Colorado, Boulder, CO, United States of America, **3** Santa Fe Institute, Santa Fe, NM, United States of America

* allison.morgan@colorado.edu

Abstract



OPEN ACCESS

Citation: Morgan AC, Way SF, Clauset A (2018) Automatically assembling a full census of an academic field. PLoS ONE 13(8): e0202223. <https://doi.org/10.1371/journal.pone.0202223>

Editor: Leonidas G Koniaris, Indiana University, UNITED STATES

Received: April 20, 2018

Accepted: July 30, 2018

Published: August 29, 2018

Copyright: © 2018 Morgan et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All data and analyses are available on Github at https://github.com/allisonmorgan/academic_census.

Funding: All authors were supported by NSF award SMA 1633791. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

The composition of the scientific workforce shapes the direction of scientific research, directly through the selection of questions to investigate, and indirectly through its influence on the training of future scientists. In most fields, however, complete census information is difficult to obtain, complicating efforts to study workforce dynamics and the effects of policy. This is particularly true in computer science, which lacks a single, all-encompassing directory or professional organization. A full census of computer science would serve many purposes, not the least of which is a better understanding of the trends and causes of unequal representation in computing. Previous academic census efforts have relied on narrow or biased samples, or on professional society membership rolls. A full census can be constructed directly from online departmental faculty directories, but doing so by hand is expensive and time-consuming. Here, we introduce a topical web crawler for automating the collection of faculty information from web-based department rosters, and demonstrate the resulting system on the 205 PhD-granting computer science departments in the U.S. and Canada. This method can quickly construct a complete census of the field, and achieve over 99% precision and recall. We conclude by comparing the resulting 2017 census to a hand-curated 2011 census to quantify turnover and retention in computer science, in general and for female faculty in particular, demonstrating the types of analysis made possible by automated census construction.

1 Introduction

Tenured and tenure-track university faculty play a special role in determining the speed and direction of scientific progress, both directly through their research and indirectly through their training of new researchers. Past studies establish that each of these efforts is strongly and positively influenced through various forms of faculty diversity, including ethnic, racial, and gender diversity. As an example, research shows that greater diversity within a community or group can lead to improved critical thinking [1] and more creative solutions to complex tasks [2, 3] by pairing together individuals with unique skillsets and perspectives that complement and often augment the abilities of their peers. Additionally, diversity has been shown to

produce more supportive social climates and effective learning environments [4], which can facilitate the mentoring of young scientists. Despite these positive effects, however, quantifying the impact of diversity in science remains exceedingly difficult, due in large part to a lack of comprehensive data about the scientific workforce.

Measuring the composition and dynamics of a scientific workforce, particularly in a rapidly expanding field like computer science, is a crucial first step toward understanding how scholarly research is conducted and how it might be enhanced. For many scientific fields, however, there is no central listing of all tenure-track faculty, making it difficult to define a rigorous sample frame for analysis. Further, rates of adoption of services like GoogleScholar and ResearchGate vary within, and across disciplines. For instance, gender representation in computing is an important issue with broad implications [5], but without a full census of computing faculty, the degree of inequality and its possible sources are difficult to establish [6]. Some disciplines, like political science, are organized around a single professional society, whose membership roll approximates a full census [7]. Most fields, on the other hand, including computer science, lack a single all-encompassing organization and membership information is instead distributed across many disjoint lists, such as web-based faculty directories for individual departments.

Because assembling such a full census is difficult, past studies have tended to avoid this task and have instead used samples of researchers [8–11], usually specific to a particular field [12–16], and often focused on the scientific elite [17, 18]. Although useful, such samples are not representative of the scientific workforce as a whole and thus have limited generalizability. One of the largest census efforts to date assembled, by hand, a nearly complete record of three academic fields: computer science, history, and business [19]. This data set has shed considerable light on dramatic inequalities in faculty training, placement, and scholarly productivity [6, 19, 20]. But, this data set is only a single snapshot of an evolving and expanding system and hence offers few insights into the changing composition and diversity trends within these academic fields.

In some fields, yearly data on faculty numbers and composition are available in aggregate. In computer science, the Computing Research Association (CRA) documents trends in the employment of PhD recipients through the annual Taulbee survey of computing departments in North America (cra.org/resources/taulbee-survey). Such surveys can provide valuable insight into trends and summary statistics on the scientific workforce but suffer from two key weaknesses. First, surveys are subject to variable response rates and the misinterpretation of questions or sample frames, which can inject bias into fine-grained analyses [21, 22]. Second, aggregate information provides only a high-level view of a field, which can make it difficult to investigate causality [23]. For example, differences in recruitment and retention strategies across departments will be washed out by averaging, thereby masking any insights into the efficacy of individual strategies and policies.

Here, we present a novel system, based on a topical web crawler, that can quickly and automatically assemble a full census of an academic field using digital data available on the public World Wide Web. This system is efficient and accurate, and it can be adapted to any academic discipline and used for continuous collection. The system is capable of collecting census data for an entire academic field in just a few hours using off-the-shelf computing hardware, a vast improvement over the roughly 1600 hours required to do this task by hand [19]. By assembling an accurate census of an entire field from online information alone, this system will facilitate new research on the composition of academic fields by providing access to complete faculty listings, without having to rely on surveys or professional societies. This system can also be used longitudinally to study how the workforce's composition changes over time, which is particularly valuable for evaluating the effectiveness of policies meant to broaden participation or

improve retention of faculty. Finally, applied to many academic fields in parallel, the system can elucidate scientists' movement between different disciplines and relate those labor flows to scientific advances. In short, many important research questions will benefit from the availability of accurate and frequently-recollected census data.

Our study is organized as follows. We begin by detailing the design and implementation of our web crawler framework. Next, we present the results of our work in two sections. The first demonstrates the validity and utility of the crawler by collecting census data for the field of computer science and comparing it to a hand-curated census, collected in 2011 [19]. The second provides an example of the type of research enabled by our system and uses the 2011 and 2017 censuses to investigate the “leaky pipeline” problem in faculty retention.

2 Background

Comprehensive data about academic faculty can be compiled from web-based sources, but is widely distributed and inconsistently structured across computer science departments. Here, we introduce a topical web crawler to retrieve and assemble these data into a comprehensive census. As a method for distributed information discovery, a topical web crawler navigates the Web, searching for relevant documents [24]. A crawler's search can be broad, such as the Never-Ending Language Learner, which continuously crawls the Web to learn new properties and relationships among persons, places or things [25]. Or, it can be narrowly focused, such as for building domain-specific Web portals [26]. Our crawler falls into the latter category in that its search space is restricted to academic webpages, with a goal of navigating to and extracting information from faculty directories.

Our search algorithm is an adaption of a “best-first search” [27–31] and can be described as follows: the crawler starts from a department's homepage, and scores each outgoing hyperlink to estimate the probability that it leads to the corresponding department's faculty directory. Then, the crawler visits the links in a greedy order, based on their computed score. If the visited page is not a directory, any additional links found on that page are scored and added to the existing priority queue. Once the topical crawler encounters the faculty directory, it follows the task of extracting the desired information from the page. Like the link structure leading to the page, faculty directories lack a common markup language [32] and are instead formatted in a variety of ways. Our method for extracting faculty information from directories therefore must thus be robust and adaptive. We describe our approach in following sections.

3 Problem formalization

Measuring the composition of computer science is complicated by differences in department structures and definitions of tenure-track faculty. Organizational structures for faculty can evolve and change over time, e.g. reflecting the dynamic and expanding nature of computing research. In our census of Computer Science, we make several ontological choices about how to define a Computer Science (CS) department in practice. These choices influence the output of the crawler and hence represent important factors in interpreting the outcomes of any downstream use of that data. For instance, many departments emphasizing computing scholarship were formed decades ago as Electrical Engineering and Computer Sciences (EECS) or Computer Science and Engineering (CSE) departments, reflecting the early intellectual closeness of Electrical Engineering (EE) and CS research and teaching. Today, many of these departments have formally or informally split into separate units, but not all have. More recently, there also exist Information Science departments, or more broadly, Colleges or Schools of Computing. The scope of these academic units varies by university, and over time,

reflecting the great diversity of computing scholarship and organizational approaches to formalizing that diversity across academic units.

For our purposes, these complexities induce ambiguities in how we define a field and therefore which specific departments we pass as input to the crawler. If we exclude an EECS or CSE department in the crawl, then we miss all the CS faculty of that university. But, if we do include it, we may overestimate the number of CS faculty because the EE faculty are also included. Similarly, if a department splits or merges over the course of its life, for whatever reason, our estimate of its size may change dramatically as a result of processes not directly related to hiring or attrition. This fact is not a failing of the data collection approach we describe here, and instead simply reflects the complicated ways the faculty that comprise a field can be organized into academic units across different universities. These ontological issues are common in other fields, as well, including the biological sciences, medicine, business, and the social sciences. In our subsequent analysis, we address these specific issues and discuss how they influence our results and the overall use of the crawler to perform an automated census.

We define the field of computer science to be the 205 North American, PhD-granting institutions from the CRA's Forsythe List (archive.cra.org/reports/forsythe.html). Here, the input to our system is a list of department homepages corresponding to each of these institutions, however, we note that searches can proceed from any listing of department or university homepages. Faculty employment information for these universities is contained in web-based faculty directories maintained by each department, yet assembling data from all institutions into a combined census is a difficult and expensive task. Crowdsourcing census construction, is complicated by the fact that domain expertise is required to distinguish tenured or tenure-track (TTT) faculty from other faculty or staff positions, and for distinguishing CS from EE or computer engineering (CE) faculty (see below). For example, the title of "associate professor" generally implies a full-time, tenured position, while an "adjunct associate professor" is neither full-time, nor tenured. To make such distinctions, workers must receive specific training when collecting data by hand, increasing both the cost and duration of a survey. The 2011 census [19] took a trained pool of workers about 1600 hours and cost \$16,000. Hence, to generate regular census snapshots, for multiple disciplines, would be prohibitively expensive and require a dedicated, trained workforce. Our topical web-crawler provides a cheap, accurate, and scalable alternative.

The crawler simplifies the overall task by finding and parsing departmental directories in four steps: first, (i) efficiently navigate to a department's directory, then (ii) identify the HTML structure separating entries within the directory, then (iii) extract every faculty record by identifying names, titles, webpages, and email addresses, and finally (iv) filter this list to include only TTT faculty members (Fig 1). In steps (i) and (iv), our approach favors higher recall by preferring false positive errors, since false negatives imply either the missed opportunity to scrape a directory or the omission of a TTT faculty member. In this setting, false positives can be corrected via downstream analyses, typically at the cost of extra computation with the parsing of a candidate's resume or the manual verification of very specific information using services like Amazon's Mechanical Turk (mturk.com) or CrowdFlower (crowdflower.com). In the following sections, we discuss each of the outlined tasks in the order of their completion.

Navigate to the directory

Our crawler's navigation strategy has two primary components: (i) navigate efficiently from a department's homepage to their directory, and (ii) identify whether a page appears to be a directory. First, in order to navigate to the desired faculty directories, our crawler must decide which hyperlinks to follow. Starting from a department's homepage, it adds all outgoing

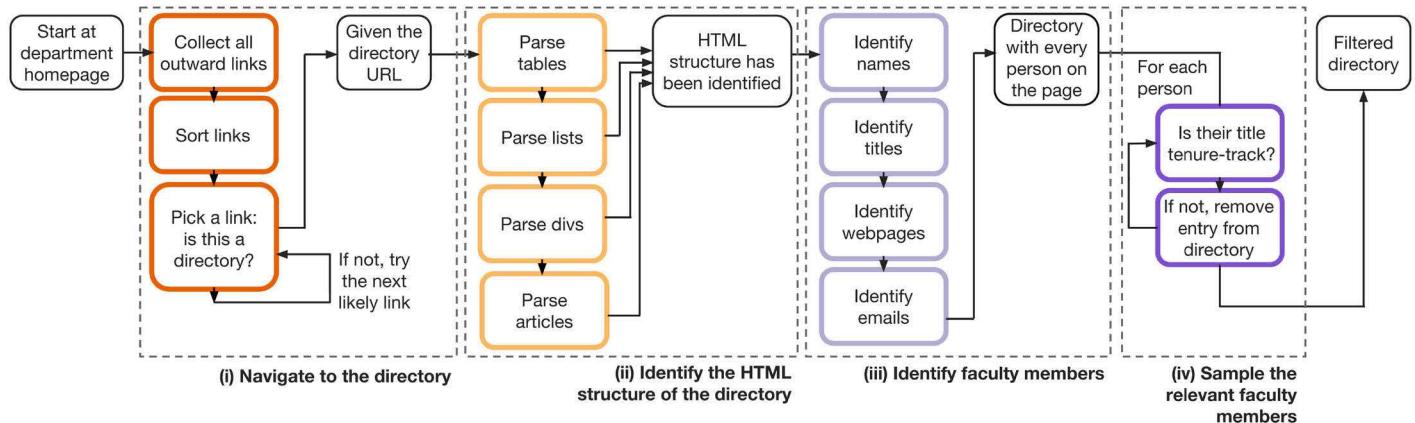


Fig 1. General schematic of our solution to the academic census problem. Starting from a department's homepage, our web crawler builds a census of its faculty in the following steps: (i) navigate to the department's faculty directory page, (ii) identify the logical structure of the directory, (iii) parse the directory to resolve potential faculty members, and finally (iv) sample and return a list of the relevant faculty members.

<https://doi.org/10.1371/journal.pone.0202223.g001>

hyperlinks to a max-priority queue, with priorities set equal to the number of keywords found within each URL and its surrounding text. This keyword list has 10 words ([S1 Appendix](#)), including “faculty,” “directory,” and “people,” which were manually extracted from common features of departments’ directory URLs. The crawler then visits pages in order of their priority, keeping track of any URLs that have already been visited, and adding newly discovered URLs to the queue as it goes, until it reaches a directory page ([Fig 2](#)).

For each visited page, the crawler must decide whether it is a directory to parse. To avoid parsing every likely page for faculty members, the crawler uses a random forest classifier to

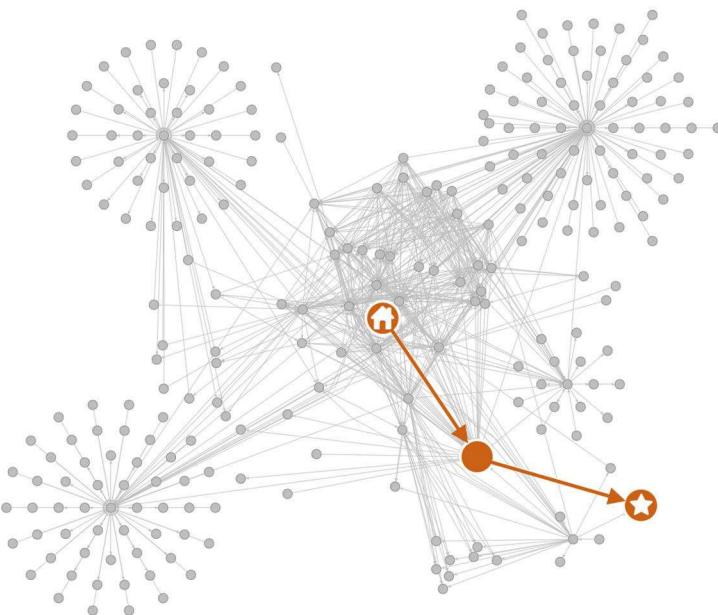


Fig 2. Example hyperlink network surrounding a department homepage. The network of all reachable webpages within two hops from the Department of Computer Science at University of California, Davis homepage (home icon). Shown in orange is the shortest possible path—and the one our crawler takes—to reach the targeted faculty directory (star icon).

<https://doi.org/10.1371/journal.pone.0202223.g002>

decide whether a page is likely to be worth fully parsing for faculty listing information. Each page is characterized by counting motifs commonly found on faculty directories, such as names, phone numbers, email addresses, and job titles. Since faculty directories typically contain little other text, a page's feature set includes counts of these motifs as a fraction of all words ([S2 Appendix](#)). Across these, the four aforementioned motifs are the most important for the classifier's accuracy. A false negative, overlooking a faculty directory, is an unrecoverable error, and induces a group of correlated false negatives for faculty in the census. We prefer a directory classifier that has no false negatives at the expense of more false positives, so any pages that yield a likelihood greater than zero are passed to the next stage (see below). Additionally, parsing a non-directory page is relatively inexpensive in terms of computational time and, since no faculty will likely be extracted from such a page, these pages are easy to subsequently identify as false positives.

Identify the HTML structure of the directory

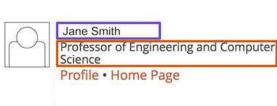
Once the crawler discovers a directory, it must extract information from a variety of HTML formatting conventions ([Fig 3](#)). In practice, despite enormous variation in the visual styling of these pages, there is a short list of common HTML tags that separate faculty members from each other: divs, tables, lists and articles. These four structural tags are used to format repeated faculty entries within a directory. Our crawler attempts to segment a directory according to each of these tags, separately, and ultimately selects the segmentation resolving the largest number of faculty records. Following this procedure for each of the 205 CS departments, we found that 100 directories were formatted with divs, 80 with tables, 24 with lists, and 1 with articles.

Finally, as part of this step, the crawler detects whether the faculty directory is distributed across multiple pages by searching for div or list tags containing common “pagination” or “pager” classes. If detected, the crawler collects the list of links and applies a parser to each. If no faculty members are collected from the page, the crawler logs the output and moves to the next highest priority URL in the queue.

Identify faculty members

After identifying the HTML structure separating faculty members from each other, the next step is to extract faculty information from the page. Each directory consists of repeated HTML

Regular Faculty



```

<ul>
  <li>
    <ul>
      <li>Jane Smith</li>
      <li>Professor of Engineering and Computer Science</li>
    </ul>
  </li> ...
</ul>

```

Filter by Faculty, Staff or Students

Faculty

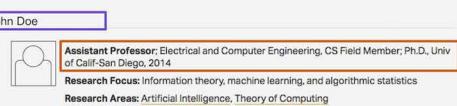
Show All | Show people whose last name begins with: A | B | C | D | E | F | G | H | I | J | K | L | M | N | O | P | Q | R | S | T | V | W | X | Y | Z

NAME	TITLE	OFFICE	EMAIL	PHONE
Clark, James	Associate Professor	9231 Gates and Crellin Laboratories	umul@cs.cmu.edu	412-268-8791

```

<table>
  <tbody>
    <tr>
      <td>Clark, James</td>
      <td>Associate Professor</td>
    </tr> ...
  </tbody>
</table>

```



```

<div>
  <h6>John Doe</h6>
  <strong>Assistant Professor</strong> ...
</div> ...

```

Fig 3. Faculty directories are formatted in a wide variety of styles, but using common structural elements. Three real examples of directories formatted using lists (left), tables (center), and divs (right). Highlighted are the pieces of information extracted by our crawler from these pages: faculty names (purple) and titles (orange).

<https://doi.org/10.1371/journal.pone.0202223.g003>

elements, and all faculty directories contain similar information: first and last names, titles, email addresses, and often faculty homepage links (Fig 3). This repetition in HTML and content allows the crawler to distinguish individual faculty records, and extract the target information. Our approach to identifying each faculty attribute is based on a set of keyword-matching heuristics, each based on a whitelist of known relevant strings.

To detect and extract names, we constructed a whitelist of first and last names from the 2011 computer science faculty census [19] (S3 Appendix). If a string contained a single substring that can be found in this set of names, the crawler classifies that string as a name. This set contained 6,798 entries. As directories were scraped, they were manually inspected and any previously unseen names were added to our list. This procedure added 260 new names (4%) to the whitelist. A similar, more exhaustive list of names could be constructed from other publicly available data, e.g., family names from the U.S. Census ([census.gov/topics/population/genealogy/data.html](https://www.census.gov/topics/population/genealogy/data.html)) or author names in bibliographic databases like DBLP (dblp.uni-trier.de).

The crawler then extracts appointment titles and email addresses from the text between names. For titles, we employ a whitelist comprising the set of all conventional titles for TTT and non-TTT faculty using partial string matches (S4 Appendix). This list is intentionally large, such that we avoid misidentifying a faculty member's title. If the crawler cannot find a title relevant to a name, it omits that entry from the directory. Typically email addresses can be identified using simple regular expressions. In some cases, emails are obfuscated on a directory page; however, in most cases circumventing such efforts is trivial. The most common obfuscation method is to remove any shared suffix ("@colorado.edu"). In these cases (4.9% of all CS departments), the domain can be trivially inferred from the web domain in the directory URL. Faculty email addresses could not be identified in this way for only 21.5% of departments.

Lastly, as they are often available, the crawler also searches for faculty webpage URLs included as links surrounding faculty member names. Although they are not utilized in this work, these URLs could be used as input for subsequent collection of faculty curriculum vitas, a direction we leave for future work.

At the end of this stage, the crawler has derived an exhaustive list of every person on the directory. This list will contain true positives, the records of TTT faculty, as well as false positives, which are any other individuals. This set of records is a superset of the in-sample faculty we seek. The next stage is to remove these false positives.

Sample the relevant faculty members

In addition to TTT faculty, department faculty rosters often list other kinds of individuals, including affiliated, courtesy, teaching or research faculty, various staff or non-faculty administrative positions, and sometimes trainees like postdocs or graduate students. In Section 5, we focus our analysis on TTT faculty for direct comparison to the 2011 census, and hence here we discuss selecting out TTT individuals. This filtering criteria reflects a choice; another filtering criteria could be applied here to produce a different kind of directory, e.g., all research faculty, contingent faculty (adjunct, adjoint, etc.), or teaching faculty. Another choice we made is the restriction to faculty whose primary affiliation is within CS, which excludes affiliated faculty and courtesy appointments.

To perform this filtering, we construct a blacklist of titles that signify non-TTT faculty and staff (such as "adjoint," "staff," "emeritus," and "lecturer"). This list contains 81 titles and was constructed by the manual evaluation of faculty records (S5 Appendix). Faculty records containing these restricted titles are removed from the output directory. Often universities publish online their definitions of non-TTT appointments (e.g., faculty.umd.edu/policies/ten_titles.html

or ap.washington.edu/ahr/academic-titles-ranks). A more sophisticated approach might collect these documents to build department specific filters.

Some CS faculty are housed in joint EECS departments, and so the crawler also checks whether a person is flagged as computer science faculty. For example, if a title contains the substrings “of,” “from,” or “in,” it checks whether that string contains a computing related word from a short custom built whitelist. However, in most cases, information about which field, CS or EE, a faculty member officially belongs to is not available on the directory. We address this issue manually in Sec. 5. Previous work has shown that faculty research interests can be distinguished using topic modeling on publication titles [6]. In the future, filtering faculty by research field in this stage could potentially be automated using publication data.

4 Results

The modular design of our system allows us to evaluate both how individual stages behave independently of each other and collectively. First, we evaluate each of the four stages separately, discussing errors and where future work could improve the system’s behavior. Then, we analyze their combination as a single system. All evaluations of the timing of our system have been made with any HTML already requested and stored locally, which controls for variability due to network latency and server liveness. Finally, we assess the generality of the system by deploying it on two additional fields, noting potential improvements for further expansion.

4.1 Navigate and classify

We evaluate the efficiency of our crawler’s navigation strategy by comparing its traversal to the shortest path from the homepage to the directory (Fig 4). A difference of zero means that our crawler makes as many HTTP requests as the shortest path. For more than half (56%) of departments, our navigation heuristic is optimal, and on average makes only 0.88 excess HTTP requests relative to optimality.

Next, to evaluate the performance of our directory classifier, we run a stratified five-fold cross-validation test. The positive training set consists of all 205 department directories, and the negative training set contains a uniformly random 50% sample (4206) of non-directory

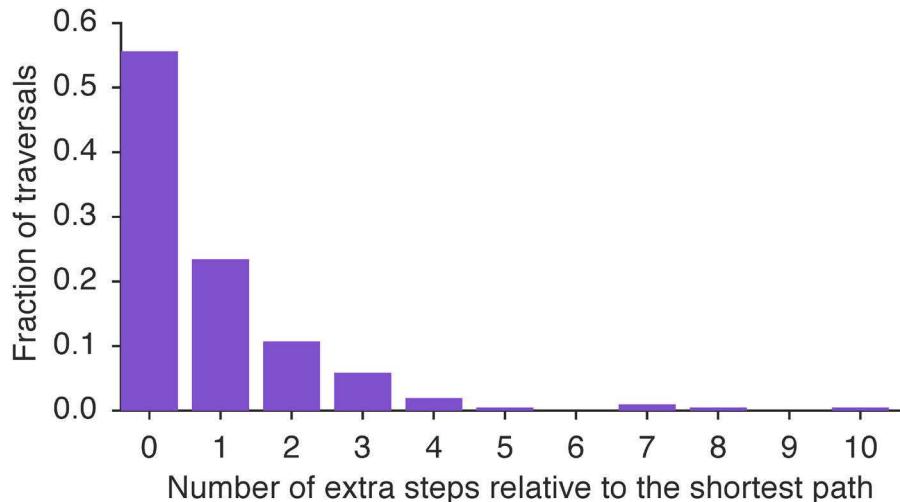


Fig 4. The distribution of extra steps taken in navigating to faculty directories. The number of steps taken by the crawler, subtracting the minimum path length from each department’s homepage to the corresponding faculty directory. In 79% of cases, our crawler commits only 1 extra step beyond the optimal path length.

<https://doi.org/10.1371/journal.pone.0202223.g004>

pages linked from the department homepages. As suggested above, the crawler was designed to avoid false negatives. In this case, a false negative would cause the crawler to not parse a directory and therefore induce a group of correlated false negatives in the census. To reduce this likelihood, the classifier returns a positive if the directory likelihood for a page is greater than zero. The resulting classifier has perfect recall across all five folds, at the expense of precision, as intended. The average accuracy—fraction of correct classifications (positive and negative)—is 0.82 due to the over-classification of non-directory webpages as faculty directory pages (standard deviation of 0.02), and the average area under the ROC curve is 0.99. The non-directory pages that are particularly difficult for the classifier to distinguish are primarily pages listing campus or administrative contact information. These pages often have similar features to directories (names, phone numbers and email addresses) and little other text. For similar reasons, pages that contain job postings or directories of affiliated or courtesy faculty are also commonly flagged as directories. False positives produced here are largely filtered out as non-TTT faculty in fourth and final stage, as described below.

Combining efficient navigation and directory classification yields a considerable improvement over a naive breadth-first search. A breadth-first crawl visits 62 pages, on average, to find the directory. The average time to parse a page is 24 CPU seconds. (CPU seconds are the amount of time the computer processor spent executing the program. This measure is shorter than the total elapsed time, but does not include time spent waiting to read from or write to files, or multitasking other programs.) Thus, the most naive implementation of a crawler would take about 1488 CPU seconds per department. In comparison, we find empirically that the navigation approach detailed here, without the directory classifier, takes 57 CPU seconds on average, while navigating intelligently and using the directory classifier takes 55 CPU seconds on average.

4.2 Parse and filter

We evaluate the performance of our four parsing methods and our ability to recover the correct attributes of a faculty record, by manually verifying their output on a subset of departments. This subset is composed of 69 departments, chosen uniformly at random but conditioned on having at least one representative for each of the four HTML structures ([S1 Table](#)). To each of these departments, we apply the correct parser directly to its faculty directory and inspect the results by hand.

To evaluate our parsing method's accuracy, precision and recall are measured by manually counting the number of TTT faculty. The 69 directories in our evaluation group list 1872 TTT faculty, of which 1868 are correctly identified, leaving 4 members missing due to either ill-formatted HTML or a missing title. The parsers also misclassify 12 individuals, calling them TTT computer science faculty when they are actually emeritus, affiliated faculty, or staff.

On this sample, the parser's recall is 99.97%, indicating that only a small fraction of true TTT faculty are missed. And, the system's precision is 99.36%, indicating that only a small fraction of non-TTT faculty are incorrectly included. The directory parsing stage is the most time intensive step of our system, taking on average 47 CPU seconds per department. As we will discuss in Sec. 4.3, this is a dramatic improvement over previous work.

4.3 Deploying and evaluating the crawler

We now evaluate the performance of the entire system, applied to the full set of 205 computer science departments. Hence, the system now starts from each department homepage, navigates to its directory, parses all pages it classifies as being a potential directory, and finally writes out a directory of all TTT faculty. Running as a single-threaded process on an off-the-shelf laptop,

the overall time required to produce structured directories for all 205 computer science departments is roughly 3 CPU hours. The majority of this time is spent parsing directories, which could be potentially reduced using a more accurate directory classifier.

Compared to the 2011 manually collected census [19], which took 1600 hours of work by a team of 13 data collectors, our automated approach is substantially more efficient. In fact, the average time required to produce a single department's faculty directory is 55 CPU seconds. Launching 205 instances of our crawler, one for each department, in a modern cloud-computing environment would lower the running time to under a minute total. In such a setting, a full census of an academic field can be automatically assembled nearly 100,000 times faster than by hand.

For 509 professors (10%) of the 2017 census, our system could not obtain an unambiguous title from the departmental faculty listings. For instance, some directories include lists or tables of the names of faculty members with nothing more specific than headings like “full-time,” “tenure-track” or “professors.” We obtain these missing titles using crowd workers on Mechanical Turk. In a production-like environment, an automated system like ours would likely need to be complemented by a small amount of human labeling to correct such errors and missing information.

Our 2017 census of North American computer science departments contains 5237 faculty members: 2637 (50.4%) full professors, 1413 (26.9%) associate professors, and 1187 (22.0%) assistant professors.

4.4 Extending to other fields

To test the generality of our system on other academic fields, we have made a preliminary application of our system to 144 history departments and 112 business schools, both of which were also part of the 2011 manual census [19]. Our results suggest that relatively little customization is needed to adapt the system to other academic fields. Specifically, we visited the online directories for each of these academic units, selected the first person listed, and checked whether our 2017 automated census of these fields contained a record for that person. In 82% of history departments and 77% of business schools' directories, we correctly recalled these faculty members with no modifications to the system. Errors here were caused by particularly complicated (often multi-page, separated by sub-disciplines) directory formats or novel faculty names, both easily corrected, and not by novel faculty titles. The loss of accuracy due to faculty names is easily addressed by incorporating a more exhaustive list of names, e.g., all surnames recorded by the U.S. Census. Parsing novel directory structures will require modest additional software development to recognize and navigate these other forms of HTML pagination. Multi-page directories with search boxes or listed alphabetically by last name are uncommon in computer science, but more common in larger fields like business schools, and it should be straightforward to extend our crawler system to handle these more complicated formats. Extending our system to the directories of other countries would require a different whitelist for faculty names and non-English language versions of the other component white and black lists. This list could be seeded by a manually collected census of faculty, or government records. The extraction of titles would also need to be adapted to that country's system of academic ranks.

5 Retention in computer science

Having applied our system to the same 205 PhD-granting computer science departments as the 2011 manually-collected census [19], we can now compare this 6-year old snapshot with our 2017 automatically-generated census. This comparison illustrates the utility of a system for

automatically assembling an academic census and allows us to better characterize the kinds of errors it makes. We also use this comparison to quantify recent turnover and retention of computing faculty. We first perform this analysis for faculty as a whole, and then consider turnover and retention for female faculty specifically. This latter step allows us to provide new insights into a question of broad relevance in computing: Are women leaving the professoriate at greater rates than men?

In order to make our comparison fair and to improve the accuracy of our estimates of turnover and retention rates, a few additional post-processing steps were necessary. Of the 205 departments surveyed, 16 are Departments of Electrical Engineering and Computer Science (EECS) and 30 are Departments of Computer Science and Engineering (CSE), meaning that their faculty directories included both CS and EE faculty. The 2011 census manually separated and removed the EE faculty, and we repeat this process on the results of our system for the same 46 departments, using faculty research interests as the separating variables. We then performed approximate string matching based on the Levenshtein distance between the names of 2011 faculty and 2017 faculty. Faculty names were matched when the edit distance represented less than 5% of the name, and no better match could be made. The results of this operation divided the set of all faculty into three groups: (i) new faculty (1556 absent in 2011 and present in 2017), (ii) retained faculty (3461 present in 2011 and in 2017), and (iii) departed faculty (1776 present in 2011 and absent in 2017). We validated this matching procedure and the accuracy of the identified ranks of faculty by using crowdsourcing to obtain the current positions for uniformly random 10% samples of each of the assistant, associate, and full professor groups from the 2011 census (475 faculty total). Each current position was collected twice and 108 observed disagreements were then manually evaluated, producing a majority vote label aggregation. Additionally, we manually checked a uniformly random 10% (132) of the 2017 assistant, associate, and full professors who were new (not seen in 2011). The results of these efforts were tabulated in a confusion matrix representing the error rates for classifying faculty by their faculty rank and by their membership in the new, retained, and departed groups ([Table 1](#)).

This confusion matrix was then used to derive corrected counts for faculty by rank and membership, multiplying the distribution of transitions generated from our crawler by the MTurker's estimated transition rates. Aggregating these corrected counts across ranks yields 1393 new hires, 4608 retained faculty, and 582 departed faculty (not observed at any in-sample institution). Overall, we find that 88.8% of faculty observed in 2011 are also found in our 2017 census ([Fig 5](#)). Furthermore, the number of new hires (23.2%) is more than twice as large as the number of departed faculty (11.2%), reflecting the overall growth in computing over this time period.

Using our two censuses, we can also investigate differences in attrition within departments. An average attrition error rate across all departments was used to produce corrected departmental rates ([Table 2](#)). As noted previously, these differences in retention are not always directly related to changes in faculty attrition, and may instead reflect organizational differences. For example, Rice University shows relatively high retention with just one assistant professor leaving for industry, whereas Stanford University's lower retention reflects the 11 faculty who became emeritus during this period. Similarly, Oregon Health and Science University has a large fraction of tenure-track faculty that are primarily grant-funded unlike many other institutions, and seven of its faculty either moved to industry or non-tenure-track positions within academia during this period. The relatively high attrition rates for some departments presented here reflect differences not merely attributable to faculty retention. For example, our system classified 17 of Indiana University's Informatics faculty as non-TTT CS faculty, unlike the manually collected census in 2011. Furthermore, the relatively high attrition

Table 1. Estimated error rates for faculty rank transitions from 2011 to 2017. Estimated error rates (expressed as percentages; rows sum to 100%) for all possible transitions of the form $X \rightarrow Y$, where X is the rank of a faculty member in the 2011 manual census (where “New” indicates that they were unobserved in 2011) and Y is their rank in 2017 (where “Gone” indicates that they were unobserved at any institution in 2017). To construct this confusion matrix, we used crowdsourcing to determine Y for a 10% uniform random sample of the 2011 and 2017 faculty, and compared those titles (columns) to the output of our crawler (rows).

Crawler Transition	n	MTurk															
		New →				Asst →				Assoc →				Full →			
		Asst	Assoc	Full	Gone	Asst	Assoc	Full	Gone	Asst	Assoc	Full	Gone	Asst	Assoc	Full	Gone
New → Asst	89	94	0	0	5	1											
New → Assoc	17	6	53	0	17						24						
New → Full	26	0	4	46	27												23
Asst → Asst	14					50	50	0	0								
Asst → Assoc	35					0	91	9	0								
Asst → Full	6					17	50	33	0								
Asst → Gone	24					17	33	8	42								
Assoc → Asst	3									33	33	0	33				
Assoc → Assoc	65									0	100	0	0				
Assoc → Full	53									0	2	98	0				
Assoc → Gone	45									0	45	13	42				
Full → Asst	3													0	0	100	0
Full → Assoc	4													0	25	75	0
Full → Full	158													0	0	98	2
Full → Gone	65													0	0	69	31
Total	607																

<https://doi.org/10.1371/journal.pone.0202223.t001>

rates of Georgia Institute of Technology and California Institute of Technology are due to their particularly complex academic titles (e.g. titles which mention appointments in other departments, or coordinator or director roles). These errors highlight the difficulty in tuning a whitelist of TTT titles across universities in CS, and suggest that future research might consider a broader list of TTT titles to handle these complexities, and more human intervention when filtering TTT from non-TTT titles.

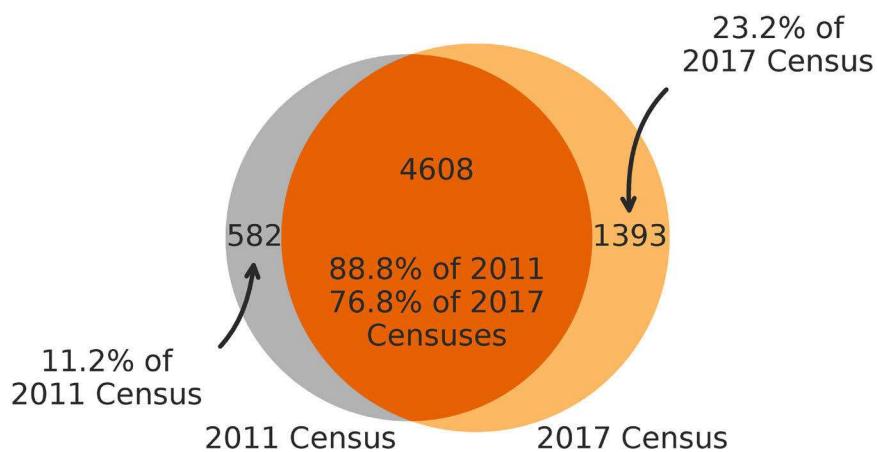


Fig 5. Faculty overlap in 2011 and 2017 censuses, adjusted for errors. The automatically collected 2017 census includes almost 90% of the faculty from the manually curated 2011 census. Non-overlapping faculty counts align with reported growth estimate from the CRA (see main text).

<https://doi.org/10.1371/journal.pone.0202223.g005>

Table 2. Departmental attrition rates, adjusted for errors. Differences in the numbers of tenure-track faculty between 2011 and 2017 for a diverse sample of departments. A complete listing of estimated department attrition rates is included in the supplement.

University	Attrition Rate	2011 Faculty	New Faculty
California Institute of Technology	0.835	16	8
Georgia Institute of Technology	0.641	89	6
Oregon Health & Science University	0.594	12	4
Colorado School of Mines	0.450	6	8
Indiana University	0.416	60	31
Stanford University	0.326	52	19
Harvard University	0.212	21	13
New York University	0.172	31	9
Tufts University	0.119	15	3
Rice University	0.050	18	9

<https://doi.org/10.1371/journal.pone.0202223.t002>

The CRA provides estimates of both department growth and losses based on information provided by a survey of the heads of departments. According to the CRA's 2011 and 2017 estimates of the number of employed tenure-track faculty from all US and Canadian CS departments (Table F1: cra.org/resources/taulbee-survey), there was an 11% growth in number of faculty. We find comparable net growth (16%) in the total number of computing faculty over 6 years. From 2012–2016, the CRA reports a total of 1206 computing faculty who left their existing positions, with 818 of these leaving academia entirely (Table F5: cra.org/resources/taulbee-survey). The size of the departed group is quite small compared to the CRA's own estimate of total faculty losses. This discrepancy likely stems from the fact that the CRA's data come from a social survey, while ours come directly from online directories and web searches. For instance, the CRA does not capture information about faculty who leave and then return to academia, while these faculty would appear in our data. A useful line of future work would involve a deeper comparison of the CRA's surveys with our faculty directory information.

Subdividing our three faculty groups (new hires, retained faculty, and departed faculty) according to each faculty member's rank (assistant, associate, or full) in 2011 and 2017, we can examine the flows of faculty into, through, and out of different career stages (Fig 6). Reflecting our finding of a substantial net growth in faculty, there is a relatively large inflow of new

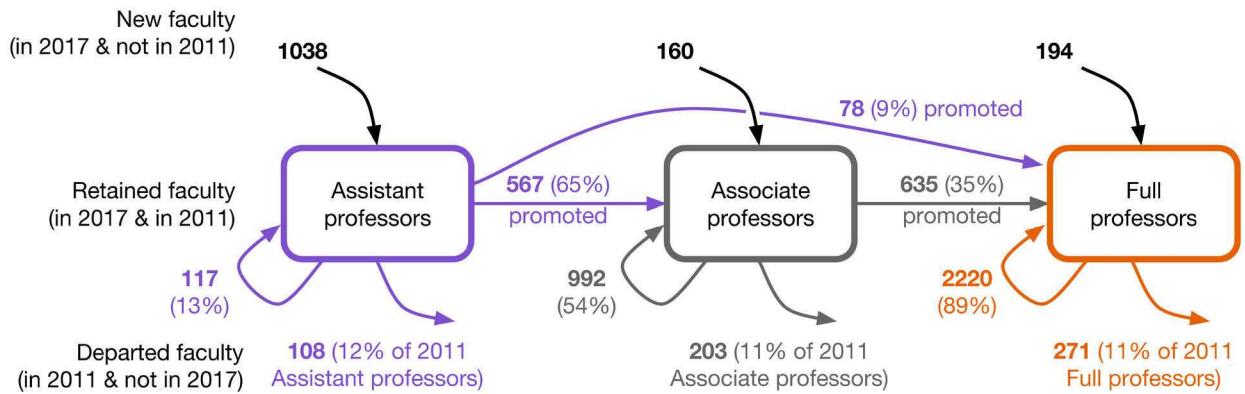


Fig 6. Faculty title transitions between 2011 and 2017 censuses. Flows of computer science faculty into, among, and out of the assistant, associate, and full professor ranks, comparing the 2011 manual census with the 2017 automated census. Counts are corrected for sampling errors in 2017 (see main text). Flows representing less than 1% of all faculty are omitted for clarity.

<https://doi.org/10.1371/journal.pone.0202223.g006>

Table 3. Faculty title transition probabilities differ slightly between men and women. Transition matrix showing the probability, based on corrected counts, that a female or male faculty member has one rank in 2011 and another in 2017. “Gone” indicates faculty not observed at any university in 2017, and this column gives the rank-level attrition rates of 2011 faculty. Total attrition rates are 0.155 for women and 0.143 for men.

		<i>n</i>	2017			
			Asst	Assoc	Full	Gone
2011	Women	Asst	274	0.293	0.462	0.150
		Assoc	193	0.192	0.434	0.067
		Full	274	0.000	0.000	0.892
	Men	Asst	1491	0.300	0.453	0.129
		Assoc	662	0.164	0.426	0.114
		Full	2116	0.000	0.000	0.886

<https://doi.org/10.1371/journal.pone.0202223.t003>

assistant professors, and large retention of associate and full professors. It is notable that the outflow rate of assistant professors is comparable to the outflow rates of associate and full professors. Naïvely, we might have expected the outflow to be larger at the assistant professor stage, reflecting the impact of negative tenure decisions.

Finally, the 2011 manual census also includes information about each professor’s gender, allowing us to estimate gender differences in rates of retention, promotion, and attrition within the CS tenure-track pipeline (Table 3). These counts indicate that slightly more women than men were retained from the group of 2011 assistant professors (90.6% vs. 88.3%) and full professors (89.2% vs. 88.6%). At the same time, fewer women than men were retained from the groups of 2011 associate professors (69.2% vs. 70.4%). Future work should investigate the gender differences among new faculty, since the 2017 census does not contain information about gender.

Aggregating across ranks, attrition rates for women and men are similar (15.5% vs. 14.3%), but slightly higher for women. This modest difference is consistent with the “leaky pipeline,” a metaphor stemming from a large body of literature showing that women leave academia at slightly higher rates than men at all stages of an academic career [33–35], including computer science [36]. A key question, however, is whether these observed differences can be attributed to fluctuations. Our data cannot definitively answer this question. However, if we model the rates of retention and promotions across gender as independent random variables, then under a binomial test for each transition, the women’s attrition rate is not significantly different from the men’s ($p = 0.40$). (We also do not find any significant difference under a χ^2 test, $p = 0.42$.) That said, a standard hypothesis test may make unrealistic assumptions about independence in this setting, and so the lack of significance in comparing two somewhat arbitrarily dated snapshots should not be over interpreted. Longitudinal analysis of, for example, yearly censuses is surely necessary in order to correctly evaluate the true significance of the observed differences. An automated system like the one presented here should make that possible moving forward.

6 Conclusion

The ability to cheaply and quickly assemble a complete census of an academic field from web-based data will accelerate research on a wide variety of social and policy questions about the composition, dynamics, and diversity of the scientific workforce in general, and computing fields in particular. The past difficulty of performing such a census has limited such efforts, and researchers have instead used less reliable survey or sampling methods. The novel system we describe here, which uses a topical web crawler to automatically assemble an academic census from semi-structured web-based data, is both accurate and efficient. In a modern cloud

computing environment, this system could essentially run at scale in realtime, on as many fields as desired.

The modular design of our system enables independent incremental improvements to its overall performance, e.g., by developing better techniques for parsing the semi-structured information stored in departmental faculty listings or for selecting target individuals out of the full listings. That said, the high precision and recall of the system when applied to North American PhD-granting computer science departments suggests that it is already quite effective. We now focus on the limitations of our current system and outline specific recommendations for how future studies might enhance and extend our work to other disciplines.

First, the system's specification currently requires several hand-constructed whitelists or blacklists, or manual interventions in order to achieve high accuracy. An important direction for future work would be to automate these steps. For example, identifying which faculty members in a departmental listing are in-sample for a particular academic field can require manual investigation, as in the case of distinguishing EE versus CS faculty in our study. Any application to the biological or biomedical sciences would also require such separation, as the corresponding disciplines are mixed in complicated ways across many departments. This step could be automated to some degree by using topic models to cluster faculty interests based on their publications [6] or on their collaboration or citation networks [37]. Automating the discovery of distinguishing features would also drive the system's expansion to other languages, enabling new studies of the increasingly international scientific workforce.

Our system was unable to identify the faculty rank for about 10% of in-sample faculty, and we collected this information manually via crowd work. An easy way to improve the system's performance in this direction would be to perform deeper crawling for each identified faculty member, e.g., crawling their professional homepage, parsing their curriculum vitae, or performing targeted web queries. The information gained through this additional work would need to be evaluated carefully, however, as different sources of information will have different levels of authority or recency.

For a system like ours, some amount of manual work is essential in order to detect, characterize, and correct the census's errors. The detailed evaluation we performed in our comparison of our automated 2017 census of computer science with the manual 2011 census illustrates this point well, as the confusion matrix we constructed via crowd work allowed us to obtain more accurate estimates of counts of faculty at different ranks in 2017. Ideally, a more accurate automated system would make fewer such human-measurable errors, and constructing such a matrix serves to highlight where accuracy improvements could be made.

The large overlap between our system's 2017 census and the manual 2011 census demonstrates the utility of a cheap and efficient automated census system. We find close agreement between the CRA's official survey-based estimate of the net growth of computing faculty and our own automated estimate. Our analysis of the flows into, out of, and through faculty ranks overall, and for female faculty in particular, demonstrates that an automated census can provide detailed insights on important questions about the composition and dynamics of the scientific workforce. A thorough investigation of the patterns we observe, including the observation that slightly more female than male assistant professors from 2011 were retained as of 2017, while substantially fewer female full professors were retained (Table 3), would require a longitudinal study. Such a multi-year census effort should now be straightforward using the system described here.

As was evident in our analysis of the retention of female faculty from the 2011 census of computer science, a key future direction will be the development of longitudinal data, which would allow more detailed investigations of trends in hiring, promotion, retention, and attrition. The system presented here is fast and suitable for continuous collection of faculty

employment information over time. It could also be adapted to historical faculty listings stored in the Internet Archive (archive.org/web). We look forward to these and other developments, and the many scientific insights that will come from having an inexpensive and accurate method for automatically assembling a full census of an academic field.

Supporting information

S1 Table. Distribution of HTML formats. Counts of faculty directory HTML formats among PhD-granting CS departments in our sample.
(PDF)

S1 Appendix. Keyword list for navigating to a faculty directory.
(PDF)

S2 Appendix. Features of a faculty directory.
(PDF)

S3 Appendix. Collection of first and last names.
(PDF)

S4 Appendix. Keywords in TTT and non-TTT titles.
(PDF)

S5 Appendix. Keywords which should not be contained in TTT titles.
(PDF)

Acknowledgments

The authors thank Mirta Galesic and Daniel Larremore for helpful conversations. All authors were supported by NSF award SMA 1633791. Supplementary code and data can be found at https://github.com/allisonmorgan/academic_census.

Author Contributions

Conceptualization: Allison C. Morgan, Samuel F. Way, Aaron Clauset.

Data curation: Allison C. Morgan.

Formal analysis: Allison C. Morgan, Aaron Clauset.

Funding acquisition: Aaron Clauset.

Investigation: Allison C. Morgan, Samuel F. Way, Aaron Clauset.

Methodology: Allison C. Morgan, Samuel F. Way, Aaron Clauset.

Project administration: Aaron Clauset.

Software: Allison C. Morgan.

Supervision: Aaron Clauset.

Validation: Allison C. Morgan.

Visualization: Allison C. Morgan.

Writing – original draft: Allison C. Morgan, Samuel F. Way, Aaron Clauset.

Writing – review & editing: Allison C. Morgan, Samuel F. Way, Aaron Clauset.

References

1. Van Dyne L, Saavedra R. A naturalistic minority influence experiment: Effects on divergent thinking, conflict and originality in work-groups. *Br J Soc Psychol.* 1996; 35(1):151–167. <https://doi.org/10.1111/j.2044-8309.1996.tb01089.x>
2. McLeod PL, Lobel SA, Cox TH Jr. Ethnic diversity and creativity in small groups. *Small Group Res.* 1996; 27(2):248–264. <https://doi.org/10.1177/1046496496272003>
3. Page SE. *The difference: How the power of diversity creates better groups, firms, schools, and societies.* Princeton University Press; 2008.
4. Milem JF. The educational benefits of diversity: Evidence from multiple sectors. In: Chang MJ, Witt D, Jones J, Hakuta K, editors. *Compelling Interest: Examining the evidence on racial dynamics in higher education.* Stanford University Press; 2003. p. 126–169.
5. Hill C, Corbett C, St Rose A. *Why so few? Women in science, technology, engineering, and mathematics.* American Association of University Women; 2010.
6. Way SF, Larremore DB, Clauset A. Gender, productivity, and prestige in computer science faculty hiring networks. In: Proc. 25th Internat. Conf. on World Wide Web (WWW); 2016. p. 1169–1179.
7. Fowler JH, Grofman B, Masuoka N. Social networks in political science: Hiring and placement of PhDs, 1960–2002. *PS: Political Science & Politics.* 2007; 40(4):729–739.
8. Cole S. Age and scientific performance. *Am J Sociol.* 1979; p. 958–977. <https://doi.org/10.1086/226868>
9. Allison PD, Long JS. Departmental effects on scientific productivity. *Am Sociol Rev.* 1990; 55(4):469–478. <https://doi.org/10.2307/2095801>
10. Long JS, Fox MF. Scientific careers: Universalism and particularism. *Annu Rev Sociol.* 1995; 21:45–71. <https://doi.org/10.1146/annurev.so.21.080195.000401>
11. Xie Y, Shauman KA. Sex differences in research productivity: New evidence about an old puzzle. *Am Sociol Rev.* 1998; 63(6):847–870. <https://doi.org/10.2307/2657505>
12. Myers SA, Mucha PJ, Porter MA. Mathematical genealogy and department prestige. *Chaos.* 2011; 21(4):041104. <https://doi.org/10.1063/1.3668043> PMID: 22225334
13. Amir R, Knauff M. Ranking economics departments worldwide on the basis of PhD placement. *Rev Econ Stat.* 2008; 90(1):185–190. <https://doi.org/10.1162/rest.90.1.185>
14. Katz DM, Gubler JR, Zelner J, Bommarito MJ, Provins E, Ingall E. Reproduction of hierarchy? A social network analysis of the American law professoriate. *J Legal Educ.* 2011; 61(1):76–103.
15. Schmidt BM, Chingos MM. Ranking doctoral programs by placement: A new method. *PS: Political Science & Politics.* 2007; 40(3):523–529.
16. Hanneman RA. The prestige of Ph.D. granting departments of Sociology: A simple network approach. *Connections.* 2001; 24(1):68–77.
17. Zuckerman H. *Scientific elite: Nobel laureates in the United States.* Transaction Publishers; 1977.
18. Schlagberger EM, Bornmann L, Bauer J. At what institutions did Nobel laureates do their prize-winning work? An analysis of biographical information on Nobel laureates from 1994 to 2014. *Scientometrics.* 2016; 109(2):723–767. <https://doi.org/10.1007/s11192-016-2059-2> PMID: 27795592
19. Clauset A, Arbesman S, Larremore DB. Systematic inequality and hierarchy in faculty hiring networks. *Sci Adv.* 2015; 1(1):e1400005. <https://doi.org/10.1126/sciadv.1400005> PMID: 26601125
20. Way SF, Morgan AC, Clauset A, Larremore DB. The misleading narrative of the canonical faculty productivity trajectory. *Proc Natl Acad Sci USA.* 2017; 114(44):E9216–E9223. <https://doi.org/10.1073/pnas.1702121114> PMID: 29042510
21. Groves RM, Fowler FJ Jr, Couper MP, Lepkowski JM, Singer E, Tourangeau R. *Survey methodology.* vol. 561. John Wiley & Sons; 2011.
22. Weisberg HF. *The total survey error approach: A guide to the new science of survey research.* University of Chicago Press; 2009.
23. Imai K, Keele L, Tingley D, Yamamoto T. Unpacking the black box of causality: Learning about causal mechanisms from experimental and observational studies. *American Political Science Review.* 2011; 105(4):765–789. <https://doi.org/10.1017/S0003055411000414>
24. Menczer F. ARACHNID: Adaptive retrieval agents choosing heuristic neighborhoods for information discovery. In: Proc. 14th Internat. Conf. Machine Learning; 1997. p. 227–235.
25. Mitchell TM. Never-ending learning. In: Proc. 29th Conf. Artificial Intelligence (AAAI); 2015. p. 2302–2310.
26. McCallum AK, Nigam K, Rennie J, Seymore K. Automating the construction of internet portals with machine learning. *Information Retrieval.* 2000; 3(2):127–163. <https://doi.org/10.1023/A:1009953814988>

27. Cho J, Garcia-Molina H, Page L. Efficient crawling through URL ordering. *Computer Networks and ISDN Systems*. 1998; 30(1):161–172. [https://doi.org/10.1016/S0169-7552\(98\)00108-1](https://doi.org/10.1016/S0169-7552(98)00108-1)
28. Menczer F, Pant G, Srinivasan P. Topical web crawlers: Evaluating adaptive algorithms. *Transactions on Internet Technology*. 2004; 4(4):378–419. <https://doi.org/10.1145/1031114.1031117>
29. Pant G, Srinivasan P, Menczer F. Crawling the web. In: Levene M, Poulovassilis A, editors. *Web Dynamics*. Springer; 2004. p. 153–178.
30. Chakrabarti S, van den Berg M, Dom B. Focused crawling: A new approach to topic-specific web resource discovery. In: Proc. 8th Internat. Conf. on World Wide Web (WWW); 1999. p. 1623–1640.
31. Menczer F, Pant G, Srinivasan P, Ruiz ME. Evaluating topic-driven web crawlers. In: Proc. 24th Internat. ACM SIGIR Conf. on Research and Development in Information Retrieval; 2001. p. 241–249.
32. Butt AS, Haller A, Xie L. A taxonomy of semantic web data retrieval techniques. In: Proc. 8th Internat. Conf. Knowledge Capture; 2015. p. 9.
33. Xu YJ. Gender disparity in STEM disciplines: A study of faculty attrition and turnover intentions. *Research in Higher Education*. 2008; 49(7):607–624. <https://doi.org/10.1007/s11162-008-9097-4>
34. Kulis S, Sicotte D, Collins S. More than a pipeline problem: Labor supply constraints and gender stratification across academic science disciplines. *Research in Higher Education*. 2002; 43(6):657–691. <https://doi.org/10.1023/A:1020988531713>
35. Pell AN. Fixing the leaky pipeline: Women scientists in academia. *J Animal Sci*. 1996; 74(11):2843–2848. <https://doi.org/10.2527/1996.74112843x>
36. Jadidi M, Karimi F, Wagner C. Gender disparities in science? Dropout, productivity, collaborations and success of male and female computer scientists; 2017.
37. Rosvall M, Bergstrom CT. Maps of random walks on complex networks reveal community structure. *Proc Natl Acad Sci USA*. 2008; 105(4):1118–1123. <https://doi.org/10.1073/pnas.0706851105> PMID: 18216267