Measuring Employment Demand Using Internet Search Data

Stevie Chancellor*
Georgia Tech
Atlanta, GA USA
schancellor3@gatech.edu

Scott Counts
Microsoft Research
Redmond, WA USA
counts@microsoft.com

ABSTRACT
We are in a transitional economic period emphasizing automation of physical jobs and the shift towards intellectual labor. How can we measure and understand human behaviors of job search, and how communities are adapting to these changes? We use internet search data to estimate employment demand in the United States. Starting with 225 million raw job search queries in 2015 and 2016 from a popular search engine, we classify queries into one of 15 fields of employment with accuracy and F-1 of 97%, and use the resulting query volumes to estimate per-sector employment demand. We validate against Bureau of Labor Statistics measures, and then demonstrate benefits for communities, showing significant differences in the types of jobs searched for across socio-economic dimensions like poverty and education level. We discuss implications for macroeconomic measurement, as well as how community leaders, policy makers, and the field of HCI benefit from this information.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords
internet search; employment; job search; big data

INTRODUCTION
The United States is in a transitional economic period facilitated by technological advances in both machinery and computing. The reverberations of these shifts are felt across many industries. While the output of American manufacturing has grown since 2000, the number of actively employed individuals continues to decline [22]. Technology jobs, from software development to customer service representatives, face continued threats of outsourcing [46]. Self-driving vehicles endanger four million trucking and transportation jobs [5].

To understand these emerging economic trends, researchers traditionally look to macroeconomic measures of employment, like unemployment and payroll data, provided by the Bureau of Labor Statistics (BLS) or the Census Bureau. While these measures are well-established and indispensable to understanding our economy, they have limitations in their application to understanding communities affected by economic downturn [73, 77]. First, many of these measures have a moderate to severe lag. Local unemployment data provided by the BLS has a modest two month lag[1], whereas local industry output and sector employment provided by the Census lag by two years[2]. For municipalities struggling with these changes, large delays can dramatically impact their ability to react.

Second, traditional macroeconomic measures focus on company production and output. By design, these measures translate creation of new jobs or goods/services produced into numbers that help measure the state of economy. However, these measures do not directly address the behavioral side of finding new jobs or how trends of job search change across geographic, demographic, and socioeconomic dimensions. Psychologists and sociologists have studied job search through interviews and surveys [79, 82] but cannot scale to study behaviors in a naturalistic, observational way across many individuals.

From this, we ask: what types of jobs do people look for? How can we better measure and understand demand for jobs by people and communities? How does this demand relate to demographic and socioeconomic measures like income and education level? Studying job search at scale could lead to design of better policy interventions to cope with shifting economic trends as well as designing technological interventions to facilitate gainful employment [51]. What is needed, then, is an at-scale, observational measure of job search that adequately reflects human behavior towards finding new employment.

In this paper, we conduct a study of *employment demand* at scale using 225 million search queries from Bing in 2015 and 2016. Employment demand is the measure of job seeking of prospective employees, both in the number of people looking for work and in the types of work they seek. We use employment demand as the counterpart to demand for labor demand, or "employment supply," which is the need for companies to hire workers to complete specific tasks.

Raw internet search is ideal to measure employment demand because 80% of job seekers use search engines to find new employment [71], and search can honestly represent the needs of individuals [26]. Aggregated search data has been used to study macroeconomic trends [16, 35], but only through heavily normalized and abstracted data analysis where search is a tool

*This work was conducted while author was at Microsoft.

[1] https://www.bls.gov/lau/
[2] https://www.census.gov/programs-surveys/cbp/data/tables.html

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

DOI: https://doi.org/10.1145/3173574.3173696

to predict another phenomenon. A measure of employment demand from internet search data, therefore, may reflect both the scale of existing government macroeconomic measures and the aspirations/perspectives of job seekers.

To accomplish this task, we first created a powerful job classification system to measure employment demand. Using a multi-class Support Vector Machine classifier, we map 225 million job search queries to fields of employment. Taking this measure of employment demand, we first explore its relation to existing economic measures like unemployment and monthly non-farm payrolls. We then compare it to demographic characteristics of U.S. counties such as population size, and socioeconomic measures like poverty and education.

In addition to improving our understanding of job seeking strategies, we demonstrate how a measure of employment demand can support meaningful policy and technical strategies that facilitate gainful employment. We show how employment demand may offer insights to municipalities struggling with shifting economic trends. Finally, we conclude with a discussion of the theoretical implications of our work, how employment demand can help communities cope with changing economies, and how HCI can benefit from our study.

RELATED WORK

Searching for Employment

To find new employment, individuals primarily use two techniques: leveraging social capital and interpersonal networks, and using online technologies like search.

Social Capital: Broadly defined, social capital is the resources, imagined and physically realized, gained by building and maintaining relationships with others [27]. The seminal explorations of social capital were done by sociologists Mark Granovetter [37, 38] and Ronald Burt [14], who found strong and weak ties play different roles in finding new employment. They both found that individuals leverage their social capital with personal connections to find employment. This line of research examining the use of social capital for job search has been extended into HCI. Burke and Kraut examined the use of Facebook by the unemployed in connecting with strong and weak ties to alleviate stress and find new jobs [13]; further work found that "strong ties" may actually be more valuable on social network sites [31]. Other angles have been studied, both for how people look for jobs and make connections [76] as well as the abuse of social capital through deception [41].

Internet Job Search: Nearly 80% of people use the internet to help get a new job [71], and web search is one primary method to finding new employment online [50]. In other fields, research has explored how the unemployed search for jobs online [61], how socioeconomic status impacts internet job search [39], and how online job search disrupts traditional job search strategies [49]. For those looking for jobs while employed, researchers studied predictors of new employment [44] as well as mental health outcomes of changing jobs [80].

In HCI, research has focused on understanding inequality of access and use of technology for job search. Dillahunt in particular has led numerous studies on the impacts of disruptive job technologies on economically and socially disadvantaged communities [23]. Dillahunt and Malone used a participatory design approach to understand how the sharing economy promises to but does not deliver on assisting those in disadvantaged communities find jobs [24]. Jen et al. examined whether modern online technologies like LinkedIn and TaskRabbit help economically distressed individuals [45]. In other areas, Thebalt-Spieker et al. investigated the inequality and biases of Uber and TaskRabbit in low socioeconomic status areas [75]. Using survey data, Hargittai and Litt explore privacy-preserving behaviors that different demographics adopt online while looking for a job [42].

Our work builds on this literature by offering a quantitative bridge between work in HCI and the study of how individuals use online search to find jobs.

Prediction of Economic Trends and Internet Search

In the last decade, "big data" analyses have been used to understand behavioral trends of individuals [3]. Economists too have realized the power of large-scale data analysis to augment traditional macroeconomic data and aid computationally powerful methods for economic problems [77]. Prominent examples include using mobile phone data to understand poverty in developing countries [9] and predicting the stock market using social media and other data sources [10, 32, 48, 58, 59].

In HCI, search is one such data source for large-scale explorations of behavior. Internet search data is a powerful assessment of peoples' thoughts and curiosities and is employed across fields to understand user behavior [26]. For complex search tasks, research has explored how search can best support information acquisition [43], how search can be exploratory rather than merely answering questions [81], and how search compares to human responses for similar queries [63]. In HCI and computational social science, search data is most often applied to understand health and wellness. Ginsburg et al. famously used raw search data to predict influenza outbreaks [33]. Search data has been harnessed to understand behaviors of those seeking health information online [21], exploring pregnancy and new motherhood [30], and learning about cancer diagnoses [66, 67].

In economics, researchers have capitalized on aggregated search data to predict macroeconomic measures. Nearly all of this work relies on Google Insights (formerly Google Trends) to gather heavily aggregated, normalized weekly trends of keyword searches. Choi and Varian's seminal work analyzed Google Trends to predict economic measures like unemployment claims and car sales [16]. Other researchers used this data to predict unemployment rates, both in the United States [20, 28] and in other countries [2, 68]. Another project studied the relationship between job search and depression and anxiety [74]. Search data has also been used to understand financial markets more broadly, mostly to predict changing prices and volumes in stocks [11, 70] and trading volumes [69]. Closest to our work is Baker, which gathered data on the keyword "jobs" to explore unemployment insurance search volumes [4].

The major limitation of these "nowcasting" approaches is Google Insights heavily aggregates and normalizes weekly search volumes. In this way, the data is obfuscated, and researchers cannot conduct fine-grain analyses to understand nuances in trends of employment – searches for generic keywords like "jobs" are vague about the types of jobs people are actually seeking. Also, "nowcasting" research uses search data as a variable to predict another object of interest, in this case unemployment. Instead, we use the data as-is, and believe that the search data itself indicates exciting trends of job seeking behaviors. Our work finds 225 million raw search queries about this new unemployment demand.

DATA

In this section, we explain how we gathered and filtered our dataset of job search queries and the ethics review of our study.

Dataset Acquisition and Filtering for Job Search Queries

Our dataset contains a sample of English-language queries from Bing for 24 months between January 2015 and December 2016. These are from both mobile and desktop devices.

To identify job search queries, we filtered for the appearance of four keywords: "job," "jobs," "career," and "careers." Other terms, like "employment" or "positions" were considered but generated more false positives than actual results (*e.g.*, "positions for trumpet scales"). We removed queries with co-occurring salacious words, URLs, and the most common keywords for false positives ("steve," "nose," and "stats").

FIPS County Identification

To tie search query analysis to communities, we placed each query in its corresponding U.S. county using the Federal Information Processing Standards (FIPS) county code system, similar to the approach used by Culotta in [19]. FIPS county codes are five digit identifiers that uniquely map to counties, cities, and distinguishing territories in the U.S.

We developed a bounding system to map latitude/longitude coordinates of queries to FIPS counties provided by the Census Bureau's TIGER/Line system[3], where every county draws property division lines. Derived from a reverse-IP look up, latitude/longitude pairs for searches were provided to us in our search sample and used to place each search into their county by these bounding lines so there is no overlap in geographic position. Reverse-IP lookup for this dataset has a median error distance of about .25km, meaning that there may be some softening of county borders, similar to Culotta's study [19]. Those that fall outside the U.S or that had invalid latitude/longitude pairs were discarded, bringing the final total to 225 million queries. All but 8 of the 3142 counties had search results in our dataset, giving us coverage of 3134 counties in the U.S.

Ethics Review and Data Protection

This study was found in line with the Common Rule for exemption by the Microsoft Research Ethics Advisory Board under protocol 7. Our data was gathered historically; there was no interactions with users by changing search results. All data was anonymized and aggregated to county level. No session information is used in our dataset. Our use and storage of this data is in agreement with Bing's End User License Agreement and Privacy Policy.

[3] https://www.census.gov/geo/maps-data/data/tiger-line.html

EMPLOYMENT DEMAND CLASSIFICATION SYSTEM

To understand employment demand, we need a system that categorizes jobs into fields of employment. One potential method is appropriating schema of industries provided by national organizations like the Bureau of Labor Statistics (BLS), where companies are mapped to broad sectors of private industry output. These categorizations measure the size and productivity of the labor force, economic growth, etc. For example, the North American Industry Classification System (NAICS) has 22 categories, like Information and Manufacturing.

However, these schema do not always match the models that individuals use to search for jobs. We observed that most job searches were for classes of jobs, like "software engineer jobs in san francisco" or "hr career in little rock". A software developer could work at multiple companies, like Amazon or Facebook. However, these companies are mapped to very different industries of output by NAICS – Amazon as Retail and Facebook as Information – and their profits are therefore mapped to entirely different industries. Said another way, these classification systems reflect corporate production and output and not demands for jobs.

One workaround is gathering a list of job titles from other sources and string-matching queries to categories of fields of employment, where we could quickly label search queries into their respective categories for analysis. However, string matching alone cannot model the sheer diversity of job titles in search data. As a result, generating all possible permutations of job titles and categories is impossible.

To overcome this, we bolster hand-validated job titles with a classifier that unwraps contextual meaning with word embeddings. Based on neural networks, word embeddings quantify contextual similarities between words using vector representations of these relationships [62]. In HCI, word embeddings have expanded model feature space to understand language use in online communities [15, 83]. We use word embeddings to infer contextual similarity between words similar to but not identical to our job titles, thereby expanding our keyword set.

In this section, we outline our category creation strategy, then our implementation of a multi-class SVM, and finally our classification results, described in Figure 1.

Job Title Matching and Category Creation

We began by building lists of high-level job categories loosely taken from the BLS's non-farm payroll categories [64] as well as our own observations of the data. We started with 18 potential categories as well as a Generic Job Query category to capture job queries that were nonspecific (*e.g.* "jobs in austin TX"). Two researchers then generated an initial set of 500 job titles taken from BLS and Census data and searches on sites like Glassdoor.com and indeed.com.

We then gathered random samples of 250 queries, and two researchers hand-annotated which of the 18 categories they matched to, noting job titles or variants missing from the dataset. Each addition, modification, or deletion to the job titles lists were approved by both researchers to represent the most precise list of job titles. These lists were updated iteratively until relative saturation was reached (less than 10 new

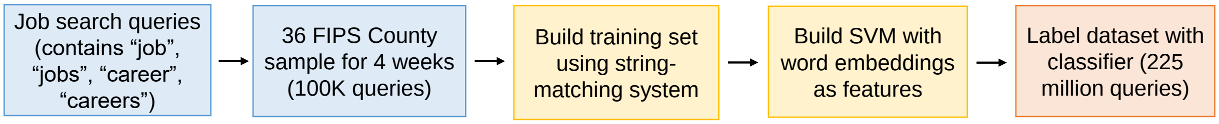

Figure 1. Flowchart of our classification pipeline.

	Raw #	%	Examples from Our Dataset
Generic Job Queries	184,500,000	82	"careers in nz", "fedex careers", "most commonly asked interview questions for a job"
Job-Specific Queries	40,500,000	18	-
Total	225,000,000	100	-
Architecture/Engineering	648,000	1.6	"entry level biomedical engineering jobs", "auto-cad jobs in central IN", "engineering careers firearms"
Art	1,539,000	3.8	"freelance writing jobs for beginners", "voice acting careers", "winterthur museum curator job"
Business	3,118,500	7.7	"vp of operations jobs", "marketing and product preference jobs", "hr career springfield ma"
Construction	1,134,000	2.8	"construction laborer jobs in reno nv", "welder jobs in wisconsin", "construction inspection jobs 06415"
Education	6,520,500	16.1	"community college professor jobs", "atlanta nanny jobs", "washoe county school district careers"
Finance	4,090,500	10.1	"financial banking jobs in vt", "mortgage lender careers", "medical insurance specialist job"
Food	810,000	2.2	"bartender jobs in minneapolis", "foodservice jobs", "craigslist dishwasher jobs"
Healthcare	12,555,000	31.0	"surgical tech jobs", "mental health jobs westernmass", "clinic job rn jax"
Leisure/Hospitality	1,822,500	4.5	"hollywood casino jobs", "fitness jobs in new hampshire", "laundry jobs in hotel"
Manufacturing	1,377,000	3.4	"machine operator jobs in columbia sc", "jobs in shipfitting in jacksonville", "machinist jobs in nj"
Retail	1,255,500	3.1	"electric boat jobs", "clothing store job applications online", "retail career at outlets near me"
Science	850,500	2.1	"psychology research associate jobs", "jobs in r&d in dc", "boston scientific careers"
Technology	1,701,000	4.2	"computer jobs in the army", "software architect career", "sql dba jobs near me"
Transportation	3,078,000	7.6	"chicago airport runway jobs", "cdl jobs in boise id", "railroad jobs in kansas"

Table 1. Output and examples of our 15 Job categories (14 specific counties + Generic Job Query). Totals are rounded.

titles per 250 searches), which took eight rounds of searching and hand-annotation.

Finally, we gathered a validation dataset of 10,000 searches and ran a string-matching system on our queries to label their respective job category. Any category with less than 2% representation or that had strong overlap with other categories were combined into other categories or eliminated. For example, “Community Outreach and Non-Profit Work” was dissolved because of its frequent overlap with many other fields (*e.g.* “nonprofit lawyer positions” overlapped with Business).

This produced 14 categories as well our 15th category, the Generic Job Query, and over 1200 job titles. Our 15 categories and example queries are available in Table 1.

Classification

Features. Our features come from word embeddings provided by the pre-trained Google News word2vec model [62] because it provides contextual data for employment trends in the U.S. We use the python library gensim[4] for our word embeddings.

To build our feature vector, we convert each search query into a list of lowercase words, remove stop words and punctuation, as well as our job search terms (“job”,“jobs”,“career”, “careers”). Then, we take each word, look up its embedding, and then average it using the coordinate-wise mean to produce a composite, 300-dimensional feature vector.

Training Data. When creating our classifier, we noticed random sampling of the search stream overbiased the classifier and job categories to trends in large metropolitan areas. To counteract this, our training/testing data come from 4 weeks of data from 36 hand-selected counties from the 2015 search stream. We hand-selected these counties to provided geographic, urban/rural, socioeconomic, and population diversity in the U.S. We sampled 4 random weeks in the dataset, one for each quarter of the year, avoiding weeks with major U.S. holidays. This generated about 100,000 searches, which were labeled by running string matching on these queries from our job titles lists. We sampled 80% (n=80,000) using stratified sampling for training, and holdout 20% (n=20,000) for testing.

Classifier Details. Our algorithm uses a one-vs-all multi-class Support Vector Machine Classification trained using 10-fold cross validation (RBF kernel, gamma=0.01, c=1000 selected using parameter sweeping). We experimented with other algorithms, such as Random Forests, Logistic Regression, and Naive Bayes, but found that the SVM dramatically outperformed other classifiers. We used the implementation of SVM in the python library scikit-learn.

Results

We split our 20% heldout dataset into two test sets. Test Set A is the whole 20% heldout dataset that illustrates the classifier’s performance in practice. Given the absence of a known baseline for this classification task, we assume the baseline is the natural prevalence of the largest class size (Generic Job Query) in our training dataset, or 77%. To verify the performance of the classifier in identifying relevant job categories and not just Generic Job Queries, Test Set B is the heldout dataset with “Generic Job Query” removed (n= 4000). For this baseline, the natural prevalence of the largest class size (Healthcare) is 30%.

In Table 2, we offer the results of our classifier and show accuracy as well as weighted precision/recall/F1. Weighted

[4] https://radimrehurek.com/gensim/

classification metrics use a weighted average, balancing each class's metric by its overall size in the dataset [72]. We note that the macro statistics, or the unweighted averages across classes, are within 5% of the weighted for all statistics, and each of our 15 classes performs above baseline (77%) for accuracy and F1.

	Acc	Prec	Rec	F1	AUC
(A) 20% Heldout (n=20,000)	0.967	0.967	0.967	0.968	0.968
(B) No Generics (n=4000)	0.907	0.921	0.888	0.904	0.909

Table 2. Results of our classification system.

Our results show that our employment demand classifier predicts the correct category of job with over 90% accuracy and F1 in both datasets, outperforming our baseline by least 13%. In Test Set A, our classifier achieves .967 F1 with comparable precision and recall. For Test Set B, we see a slight drop in performance (.967 to .907), although this performance substantially beats its baseline of 30%. The area under the curve for Test Sets A and B is 0.968 and 0.909, respectively.

Error Analysis. Finally, we conduct a brief error analysis on the classifier's performance of 250 hand-labeled queries to both identify novel matches not in our job title list as well as find errors. Two researchers labeled 250 job search queries to one of the 15 categories (ref. Table 1). The raters disagreed on three labels and resolved those differences through discussion.

First, we looked at what queries the classifier correctly identified that had no matches to the string matching lists. "x-ray jobs in minnesota" was identified as a Healthcare query despite x-ray never appearing in our Healthcare job titles. "client service rep careers" was another query that was correctly mapped into Business. "rep" is actually a word in Retail and "client service" in Business – we were pleased that the classifier learned that a "client service rep" is closer to a Business job.

We also examined errors the classifier made. For example, "assembly jobs" could be a Manufacturing job associated with building on a factory line, a Construction jobs where people assemble various structures, or a Technology career in an assembly programming language. The classifier identified this as Manufacturing, whereas our two raters initially disagreed on its category and resolved it as a Generic Job Query because of ambiguity. Without knowing more about the searcher's intentions, it would be difficult for any classification system to disambiguate this query.

Final Dataset. Of 225 million queries, the classifier labeled 184.5 million as Generic Job Queries, or 81% of our dataset. 18% of our queries were labeled as Non-Generic or Specific Job Types, for a total of 40.5 million. The most predominant searches are for Healthcare followed by Education, Finance, Business, and Transportation jobs, in decreasing order. The numbers and percentages of each type of job search query can be found in Table 1.

FINDINGS AND INSIGHTS

We now have 225 million searches organized by job category and location – this is our measure of employment demand. What can we understand about job search with this measure, and how might it improve our understandings of individuals and communities in a changing employment landscape?

We conduct two groups of analyses. First, we analyze macroeconomic measures closely related to changes in employment: labor force and unemployment, and non-farm payrolls. Second, we illustrate how employment demand relates to individuals and communities. We focus on four comparisons done at the U.S. county level: relationship of employment demand to population size and diversity, and to socioeconomic measures of poverty and education.

Relationship to Existing Employment Measures

Unemployment and Labor Force. Measuring the labor force and unemployment rate are major functions of the BLS. Formally defined, the labor force is the percent of the population who is working or looking for work; the unemployment rate is the percent of the labor force not currently working. As of August 2017, the national unemployment rate was 4.4% and the national labor force participation rate is about 63% [65]. However, both the labor force participation and unemployment rates vary widely across counties. Labor force participation typically ranges from a third to two thirds of a county's population, and unemployment from 1% to 25%.

Using Pearson's correlation coefficient and controlling for population and search volume, the correlation between the total number of job searches per population in a county and the unemployment rate is -0.11, or a weak negative relationship. However, univariate correlation of the size of the labor force and employment demand per population is 0.47 across all U.S. counties. Additionally, the Pearson's correlation between job searches and percent of a county that is employed is 0.47.

Our results suggest that when comparing across counties, search-based employment demand does not measure unemployment directly. Research and surveys have shown that those who are employed do consistently search for jobs [8, 40, 34]. The number of unemployed individuals is eclipsed by the much larger group looking for jobs including the unemployed and the employed. Our findings corroborate these findings and suggests that employment demand captures significant activity of the entire labor force, not just unemployed individuals.

Non-farm Payrolls. Arguably the most timely and important government indicator of employment is the monthly non-farm payroll (NFP) release from the BLS [64]. Released on the first or second Friday of the month, NFP shows the number of employees added to payrolls in the month prior, excluding the farming industry. NFP releases can have significant effects on stock markets and the value of the U.S. dollar, as strong NFP numbers indicate growing businesses and newly employed people with more money to spend on goods and services [6]. We expect NFP would correlate negatively with employment demand; that is, when NFP results increase, the search for new jobs goes down as more individuals are employed.

To understand the relationship between NFP and employment demand, we tabulate the month-over-month change values, or delta values, for both changes in NFP and aggregated, national employment demand. We run a univariate correlation between the resulting employment demand time series and NFP using

Signal	NFP Correlation
Total searches	-.51
Total searchers	-.45
Arch/Eng	-.54
Art	-.65
Business	-.59
Construction	-.57
Education	-.39
Finance	.22
Food	-.65
Generic	-.48
Healthcare	-.61
Leisure/Hospitality	-.55
Manufacturing	-.33
Retail	-.42
Science	-.67
Technology	-.65
Transportation	-.59

Table 3. Univariate correlations of month over month delta values in employment demand and non-farm payrolls.

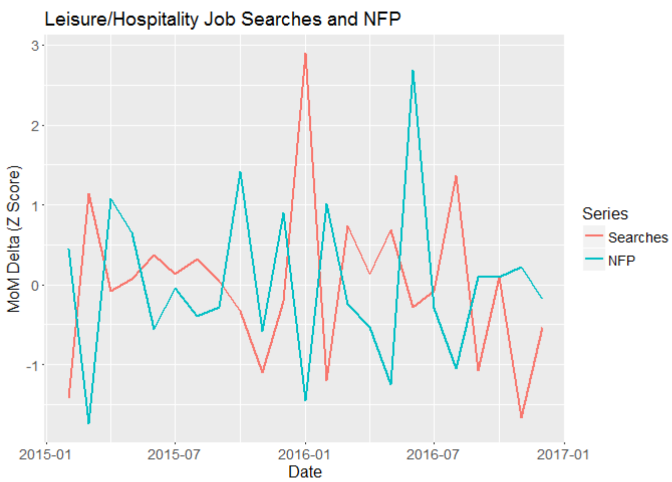

Figure 2. Searches for Leisure/Hospitality jobs versus non-farm payrolls. Each time series has been normalized to z-scores for y-axis alignment.

Pearson's correlation. Table 3 shows the results of these correlations across the dataset, as well as by the total number of searches and searchers nationally and each of the individual job categories. In all but one case, employment demand negatively correlates with NFP at medium to large effect sizes (-0.33 to -0.67).

To explore this more, Figure 2 shows the relationship between searches for Leisure/Hospitality jobs and NFP – we scale both NFP and employment demand using a z-score to normalize the data and align it on the y-axis. This figure shows that in nearly every month, there is an inverse relationship between NFP and employment demand that persists all year.

Our results supported our expectation that NFP is negatively correlated with employment demand. As individuals gain jobs in the NFP results, we can see employment demand drops, and visa versa, in Figure 2. We argue that this shows the sensitivity of employment demand to variation in job search is comparable to the rate of new hires measured by NFP.

Finally, we provide an initial analysis of the predictive power of employment demand over the known offline baseline of NFP using time series analysis. We caveat that there is only 24 months of data, and there is no offline equivalent to employment demand, making perfect comparisons difficult. An autoregressive baseline explains less than 1% of variance in month-over-month NFP. We fit a simple linear model where employment demand predicts NFP in two categories, Business and Finance. With a 2 week lead, we explain 35% of adjusted variance in month-to month NFP. Even at 6 weeks lead, we explain 10% of variance over the autoregressive baseline.

Relationship to Communities

In this section, we address the links of employment demand to communities. We analyze the differences in high and low population counties, employment diversity, and socioeconomic measures of poverty and education.

Population. We hypothesize that there are differences in employment demand between counties with different populations. Specifically, we predict that people in larger population centers will seek more Business, Finance, and Technology jobs, as well as more niche job areas like Art and Science. In contrast, job seekers in lower population areas look for jobs in Manufacturing, Transportation, and Construction. Prior work has discussed the urban-rural differences in employment [71].

To test this, we segment U.S. counties into two groups, the highest quartile of counties by population (> 67282) and the lowest quartile (< 10960). To control for effects of population on search, we compare percentage across categories rather than raw counts of queries. Using a Welch's t-test, we test whether there are differences in searches between these two groups. Results for this test are in Table 4.

We note here that throughout the remaining analyses we use quartile splits to compare across demographic variables. We chose quartiles because they reflected meaningful separations in the data. Two counties near but on opposite sides of the median population are not very different, whereas those in the top and bottom quartiles align with commonly understood notions of small and large populations. We ran the same comparisons using median splits, with little substantive differences in outcomes.

Category	Lowest 25%	Highest 25%	t-stat	Sig
Arch/Eng	0.0157	0.0169	-1.2708	
Art	0.0357	0.0351	0.5227	
Business	0.0584	0.0721	-9.8869	***
Construction	0.0456	0.0285	10.715	***
Education	0.201	0.168	8.701	***
Finance	0.0846	0.0944	-4.400	***
Food	0.0162	0.0204	-5.0576	***
Healthcare	0.284	0.316	-7.481	***
Leisure/Hospitality	0.0374	0.0423	-3.008	*
Manufacturing	0.0420	0.0376	3.414	**
Retail	0.0246	0.0303	-4.9853	***
Science	0.0225	0.0216	0.932	
Technology	0.0346	0.0411	-5.390	***
Transportation	0.0874	0.0751	5.443	***

Table 4. Results for t-test between the highest population quartile (x > 67282) and lowest (x < 10960) (n=787 for both groups). Significance is represented with * = p < 0.05, ** 0.01, * 0.001, after Bonferroni correction($\frac{p}{14}$) for multiple testing.**

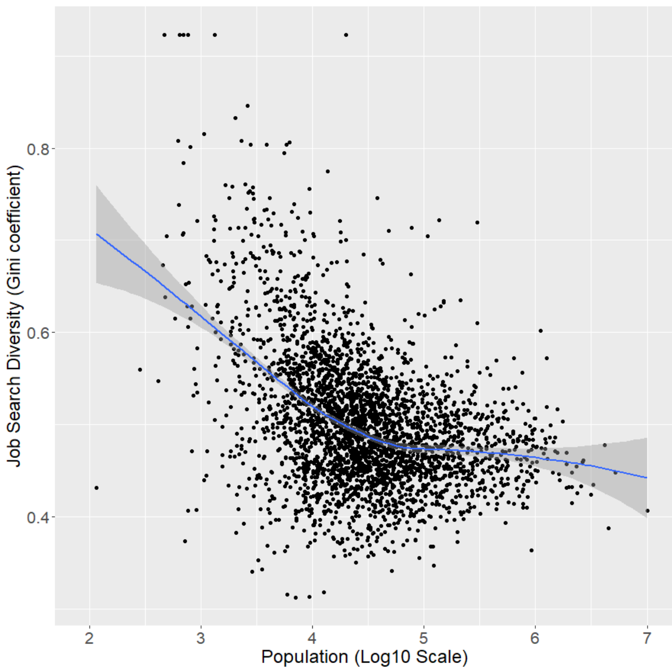

Figure 3. Diversity of employment demand versus population size over U.S. counties (2016). Loess fit line shown to illustrate the inflection point in job search diversity at population size of about 50,000.

We find that job categories more strongly associated with urban centers, like Business and Finance, show a significantly higher percentage of job searches in larger counties. Food and Leisure/Hospitality jobs also were significantly larger in urban centers, reflecting a higher concentration of hotels, restaurants, and sporting venues. Conversely, we also saw that Manufacturing and Construction jobs were a significantly higher percentage of searches in low population counties. Surprisingly, niche job categories like Architecture/Engineering, Art, and Science showed no significant differences.

Our hypotheses about population were mostly supported. We saw that jobs typically associated with urban centers were searched for in larger population counties (Finance, Business, Retail, etc), and that, as expected, Manufacturing and Construction jobs were more sought after in rural counties. Notably, we saw that Education jobs were a higher percentage of search in small counties than in large. Education jobs occur in all counties at high rates, and we suspect that Education appears more frequently because of its necessity as a service compared to other job types. We also see the most uncommon job types (Architecture/Engineering, Art, and Science) showed no significant difference between high and low population counties. While we expected these categories to be more common in more densely populated areas because of specialization, it appears that these jobs simply are rare in general despite population density.

Diversity of Job Searches. Beyond differences in population, we expect that smaller counties will see less diversity in job search. Smaller counties have fewer people and fewer businesses; therefore, residents of these counties are likely to explore a smaller range of employment options tied more closely to essential services.

To measure diversity of employment demand, we use the Gini coefficient. Gini measures the degree of inequality across a set of values, most commonly to understand income inequality [17]. Ranging from 0 to 1, higher Gini indicate more inequality; lower Gini indicates more equality.

We calculate a Gini coefficient for each county across searches in our 14 job categories, controlling again for search volumes by using percentages, then show its relationship to population (Figure 3). There is a clear and strong negative correlation (r = -.51) between population size and job diversity: as population increases, the Gini for a county drops. However this correlation nearly disappears when county size reaches 50,000.

Our analysis mostly supported our hypothesis. Generally, there is a negative correlation between employment demand and Gini in U.S. counties. It is important to emphasize that lower job search diversity does not imply a poorly performing economy. Some degree of inequality in job offerings is expected because it is a byproduct of needing relatively more essential services, like doctors, grocery stores, and schools. Research has suggested that low job diversity might be one factor that influences instability [57], but it alone is not evidence of a poorly performing economy. However, we did not expect Gini to stabilize so drastically in counties larger than 50,000 people. This may represent the convergence of talent at cities and counties naturally specializing in certain fields, like San Francisco and surrounding counties being a major Technology hub.

Poverty. Next, we examine the relationship between employment demand and poverty. We hypothesize that counties with higher percentages of poverty might favor lower paying jobs like Retail and Transportation versus higher paying jobs like Finance and Business.

To explore this relationship, we again segment U.S. counties into quartiles based on percentage of population below the poverty line from the Census Small Area Income and Poverty Estimate (SAIPE) in 2015. For simplicity's sake, we refer to the top quartile (19.7-47.4%) as the "Poorest 25%" and the bottom quartile (3.4-11.5%) as the "Wealthiest 25%". To complement the data available, we use 2015 search data and control for effects in population and search volume. We again run a Welch's t-test on these two groups.

Category	Poorest 25%	Wealthiest 25%	t-stat	Sig
Arch/Eng	0.0180	0.0173	0.739	
Art	0.0352	0.0350	0.166	
Business	0.0614	0.0648	-2.856	
Construction	0.0394	0.0392	0.109	
Education	0.179	0.181	-0.608	
Finance	0.077	0.0999	-9.852	***
Food	0.0162	0.0183	-2.749	
Healthcare	0.316	0.297	4.820	***
Leisure/Hospitality	0.0362	0.0419	-3.410	**
Manufacturing	0.0450	0.0389	4.637	***
Retail	0.0289	0.0262	2.892	
Science	0.0218	0.0222	-0.501	
Technology	0.0373	0.0390	-1.571	
Transportation	0.0874	0.0796	3.717	**

Table 5. Percent of jobs sought by category in the poorest 25% and wealthiest 25% of counties in the US, as measured by percent of population below the poverty level. * = p < 0.05, ** 0.01, * 0.001, after Bonferroni correction($\frac{p}{14}$) for multiple testing.**

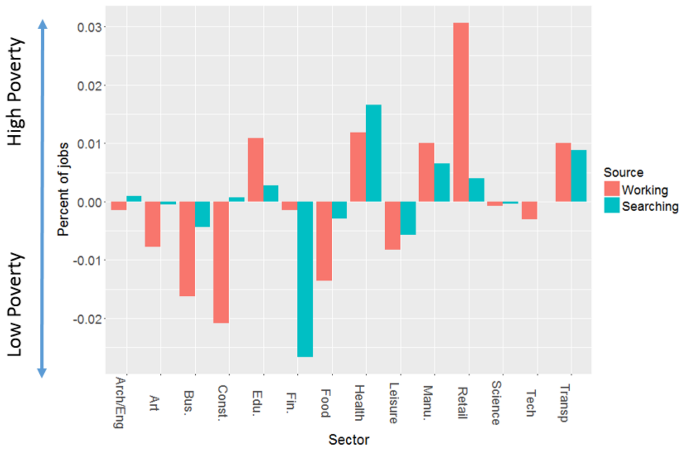

Figure 4. People working at types of businesses and types of jobs searched for: high versus low poverty counties. Bars with positive values indicate skew toward high poverty counties.

Table 5 shows our results. People in the Poorest 25% search for significantly more Healthcare, Manufacturing, and Transportation jobs compared to counties in the bottom quartile. Conversely, people in wealthier counties seek more Finance and Leisure/Hospitality jobs. We see no significant difference between search patterns in Education and Technology jobs.

Recall from the Introduction that jobs at risk from transitions to the digital economy were in sectors like Manufacturing and Transportation as well as Technology [5, 22, 46]. These jobs were significantly more searched for in the Poorest 25% of counties, excepting the lack of difference for Technology jobs, suggesting that people in poorer counties search for jobs that are at risk of going away or at least undergoing significant change. One explanation is that people in poorer areas simply are searching for the types of jobs available where they live.

To test this possibility, we compared what jobs people search for to what jobs are available in the Wealthiest and Poorest 25% of counties. We do this by comparing the percent of searches for each job category to the percent of jobs per county as reported by BLS. We used annual totals for 2015 from the BLS Quarterly Census of Employment and Wages, broken out into employment sectors. With the exception of aligning Technology job searches to the Information sector, all other employment demand sectors had an exact counterpart in the BLS data.

Figure 4 shows the results of this analysis. It shows broad agreement between the types of jobs people work and search for: those in higher poverty counties skew both in the jobs they are working and searching for toward sectors like Retail and Transportation. This suggests a "vicious cycle" in that people in some communities are stuck working lower paying jobs and those are the same jobs they search for when looking for work. Further, higher paying jobs such as those in Finance, skew strongly toward lower poverty, or "Wealthy" counties.

However, when one compares the relative skew of jobs working versus being searched for, the employment demand measures shows an aspiration in higher poverty counties away from these lower paying jobs. For instance, while higher poverty counties skew strongly toward working Retail jobs, the skew in employment demand for those jobs is modest. Business jobs working skews strongly toward lower poverty counties, but Business jobs searched for skews toward those counties only modestly. The skews for Architecture/Engineering and Art jobs are small but actually flip direction with a skew toward low poverty counties for jobs working and toward high poverty counties for jobs sought.

Education. Finally, we suspected that education levels might also relate to employment demand. Per common intuition about education and work, we expect that jobs with higher education needs (Business, Education, Healthcare, Science, Technology) would be searched for more in counties where people have higher levels of education. We investigated this by examining the relationship between education, measured by the percentage of individuals with at least a Bachelor's degree, and employment demand. Education data was obtained from the United States Department of Agriculture's Economic Research Service for 2015. Again, we selected for the highest quartile of education (24.5-78.8%) to compare against the lowest quartile (0.1-14.3%). We again use a Welch's t-test to test our hypothesis.

Category	Lowest 25%	Highest 25%	t-stat	Sig
Arch/Eng	0.0162	0.0165	-0.283	
Art	0.0346	0.0345	0.0545	
Business	0.0585	0.0672	-7.116	***
Construction	0.0398	0.0281	10.142	***
Education	0.171	0.164	2.165	
Finance	0.0769	0.0907	-6.926	***
Food	0.0151	0.0201	-6.254	***
Healthcare	0.307	0.282	5.474	***
Leisure/Hospitality	0.0327	0.0439	-7.161	***
Manufacturing	0.0516	0.0309	20.134	***
Retail	0.0318	0.0258	5.713	***
Science	0.0191	0.0227	-5.0278	***
Technology	0.0370	0.0387	-1.458	
Transportation	0.0902	0.0639	13.726	***

Table 6. Results for a t-test between the highest 25% counties with bachelor's degrees (x > 24.5) and lowest 25% of counties in the US (x < 14.3). (n=803 for each class).

Results for this comparison are in Table 6. We observe that in counties in the highest quartile of educational attainment, individuals have significantly more demand for jobs in Business, Finance, Leisure/Hospitality, and Science. For counties with lower prevalence of Bachelor's degrees, individuals look more for jobs in Construction, Food, Healthcare, Manufacturing, Retail, and Transportation.

We find that intuitions about education levels and the type of work generally are supported by our analysis. Interestingly however, we see no significant difference in searches for Technology jobs, as well as Architecture/Engineering and Art. In fact, searches for Healthcare jobs, generally requiring some kind of degree, are more prevalent in counties with fewer Bachelor's degrees.

DISCUSSION

In this paper, we provide one of the first attempts at a large-scale measure of employment demand – over 225 million queries and two years, we capture the behavioral trends of individuals looking for jobs. Further, we employ a hybrid classification approach using 1200 human-generated job titles with word embeddings to label each job search query with one of 15 categories of employment. We show that these measures correlate with traditional economic measures of labor demand, and that they illuminate outcomes relevant to communities, like socioeconomic measures of poverty and education.

A near real-time measure of employment demand such as this combats many challenges of timeliness limited by traditional macroeconomic measures of job search. Measures from BLS and Census can lag by months and in some cases years, delaying action that could be taken using these data. In contrast, internet search data is quickly gathered, aggregated, and analyzed, even as soon as the data is produced. In our example of NFP (ref. Findings and Insights), monthly projections for NFP from employment demand can be calculated immediately at month's end, beating out the BLS release by one or two weeks. For measures other than NFP, we believe that timeliness provides considerable value. For instance, data can be aggregated in the stream to understand emergent trends in employment demand, which could be used to make timely assessments of the economy months or even years ahead of national organizations.

We also argue that honest and deliberate intentions captured by search log analysis is critical for measuring the process of finding employment. Changing jobs is stigmatized, and many job seekers suffer mental health impacts from their search, even if they already have employment [29, 78, 79, 82]. Additionally, job seekers may be embarrassed to admit to family and friends that they are looking for new employment [80]. Search engine data, with its unbiased and honest data source of people's inquiries [26], can capture these intentions not obscured by social stigma. In short, search data is a window into the strategies of many kinds of job seekers which avoids discussing the stigmatized action of finding a new job [36].

In sum, we provide one of the first large-scale studies of employment demand from the perspective of human behavior. Our work bolsters prior work in psychology and sociology that explores job seeking behaviors, like how individuals use the internet to cope with unemployment [61] or find jobs while still employed [44]. We envision our employment demand data used alongside both traditional macroeconomic data as well as alongside psychological methods to explore new questions about job search in the 21st century.

Community Implications

For communities challenged with changes to the national economy, timely and relevant economic analyses are crucial to understanding these trends and to developing appropriate policy. In this section, we describe several ways that real-time employment demand can bridge these gaps and provide communities better insights.

Configurable Data Sources. One benefit of highly granular employment demand data is that the data are highly configurable to multiple contexts for communities. For instance, employment demand can be used at local scale to explore the immediate impacts of a shopping mall closure on job search behaviors in Retail. Employment demand can also be used at a high level to understand long-terms shifts in job searching behaviors for a city struggling with the slow decline of Manufacturing or rise of job seekers for Healthcare positions. In this way, employment demand can be configured atomically for diverse communities with diverse economic and social needs.

Retraining and Reemployment. We see our measure as helping to inform retraining and reemployment opportunities driven by communities. Specifically, post-secondary education opportunities could be better tailored to the aspirations of the population rather than "one-size-fits-all" approaches to remediating economic hardship and job loss [56]. Local job placement agencies may uncover skills or goals in their local populations from employment demand to better target the acquisition of new opportunities for clients. Employment demand is not a replacement for cultural and community knowledge, an idea captured by Lazer et al's warning against "hubris" in search data analysis [53], where we assume that big data substitutes and replaces traditional analyses. To the contrary, we see our large-scale measure of employment demand as one of many tools for communities and social organizations to better respond to changing trends in economic situations.

Targeting the Underserved. We also think employment demand facilitates better understandings of communities underserved by traditional economic analyses. Because of limitations of traditional economics measures, economists themselves have argued that "big data" is one solution to their limitations [73, 77]. As shown by our correlations to labor force and employment percentages (ref. Findings and Insights), search data can help understand the job search trends of those not served by regular economic analyses, such as those casually looking for jobs but not considered part of the workforce, or those working in the "gig economy." Employment demand may also provide a naturalistic lens to study job searching behaviors of regions under particular economic duress, like the "Rust Belt" in the northern U.S. [18]. This provides insights into community job search strategies, offering these communities more agency.

Local governments and municipalities in particular often do not have the availability of data that states and large metropolitan areas have, nor access to the resources to hire analysts. A timely, localized measure of employment demand from search data would give these organizations tools to understand trends of job search within their communities. As one example, we saw that there was no significant difference in searches for Technology jobs between counties with the highest and lowest poverty rates (ref. Findings and Insights, Poverty). This was also true with other job types like Science, Architecture/Engineering, and Education. Given that we measure what individuals search for, this highlights that all communities are responding to economic changes fostered by digital technology advancements at least in their search patterns.

Because employment demand captures some notion of the type of work people want in the future, we can use it to assist social organizations and educational institutions to design better interventions for economies. For example, higher than average trends of searches for Technology jobs may lead community colleges to prioritize retraining and skills acquisition in Technology-related areas. Such targeted efforts may be of particular use in smaller population communities, where employment demand is less diverse (ref. Findings and Insights, Population Diversity). Training could either better match the job types searched for, or it could be used to expand on the types of jobs for which community members feel qualified. Lower income areas that are at risk of getting stuck in a downward employment spiral in which people continue to seek the same lower wage jobs they have been working (ref. Findings and Insights, Poverty) can be identified as areas where job awareness and retraining would be of particular benefit.

Gainful Employment, Public Policy, and HCI

The methods we use for understanding employment demand from search data can be used for other problems in large-scale social computing, behavioral research log analysis, and other quantitative research in HCI. For instance, the hand-labeled job categories combine with word embeddings is a very useful technique to scale human labels with computational language modeling. Additionally, the human-in-the-loop machine learning system we present is extendable to other areas of research in HCI. We hope that other researchers use our methodology to improve their own work going forward.

More importantly, our measure of employment demand highlights larger questions for how HCI can speak to broader goals of political and social importance. In recent years, there has been a growing push from within HCI to consider impacts to other fields, such as sustainable HCI [25] and implications of cannabis legalization and consumption [47]. In particular, HCI researchers like Lazar and Bederson have argued CHI, and by proxy HCI, should more deeply consider public policy implications of work in the field [7, 51, 52].

In relation to employment, prior work in CHI and CSCW has started addressing how the sharing economy impacts individuals and communities [23, 75], and how individuals leverage socio-technical systems to find new employment [13, 41]. We see these works as fundamental building blocks for understanding the technology behind job search and employment more broadly, and hope our work adds to this foundation.

Therefore, we urge HCI researchers to use their technological expertise and human-centric research methods to bear on conversations about gainful employment and public policy. We see this happening in two ways. First, given the field's interest in large-scale data analysis with search data and beyond, HCI can contribute analytical studies that promote deeper understandings of how individuals use technology to find jobs, and emergent socio-technical issues. Our work on employment demand aligns with these goals because of its power to speak to the behavioral aspect of achieving gainful employment. We also see HCI building on prior work [23] to build technology-driven interventions to improving employment outcomes for communities. HCI is well-equipped to design, test, and deploy technologies that can facilitate gainful employment and speak more broadly to new technology in employment. In sum, we strongly advocate that HCI consider research questions, like issues of gainful employment that we study, that actively engage with issues of growing societal and political importance.

Limitations and Future Work

Alkhatib et al. caution against "the hazards of predicting the future" [1], and we are aware of the methodological limitations of such analyses [12, 53, 54].

Employment demand is constrained to measuring what kind of jobs individuals search for using search engines. This means that we cannot understand job search strategies that people conduct through other venues, likely through their social contacts and interpersonal networks [13, 14, 38]. We notice this prominently with jobs in agriculture, mining, and forestry; these job searches were very rare in our dataset. Although online job search is known to happen across all demographic and urban/rural lines [71], rural job seekers in agriculture and forestry leverage their social ties more aggressively than their urban counterparts [55, 60].

However, we do not claim to model employment demand for all individuals in the U.S, nor expect that it would be feasible to do so. Such a metric would need to combine how individuals use their social networks to find jobs, multiple online sources for jobs, and offline jobs search strategies. We see our measure of employment demand as a major step to understanding how individuals and communities think about finding employment through their online search strategies.

Second, we currently have no means to externally validate this metric, as we are using observational analysis of search queries on a popular search engine. As reported, this measure does correlate fairly strongly with the "gold standard" BLS measure of employment growth. However, employment demand is intentionally not redundant with the BLS measure, and we hope that future work validates what we see as an important and timely way to understand job search.

In addition to new collaborations and analyses we present above, we see several areas for future work to improve our measure of employment demand. First, distinguishing the unemployed from the employed job seeker would be beneficial in understanding how the two strategies differ. This would require outside knowledge or prior individualized search history. Second, we could unpack the linguistic search trends of how people describe jobs in search or click-thru to their intentions of searching for a job. This information could be incorporated into the classifier about the type of job they may be looking for with a particular search query. Finally, using our location data, we could understand the "radius of gyration" that individuals have to find jobs in counties outside their residence. This could help communities understand search strategies for jobs outside of the immediate county.

CONCLUSION

In this paper, we proposed a new measure of job search trends called employment demand that captures the demands individuals have while searching for a job. We used 225 million search queries in 2015 and 2016 from Bing to map and explore trends in employment demand. We developed a classification system that labels these search queries into one of 15 categories of job search that performs with .97 F1. We validated our measure by showing expected correlations with changes in monthly non-farm payrolls, and across geographies by showing relationships to demographic characteristics of U.S. counties such as population size, existing labor force measures, and socio-economic measures. We believe that employment demand can be used to explore many community, policy, and economic trends, and encourage HCI to engage in research around gainful employment.

REFERENCES

1. Ali Alkhatib, Michael S Bernstein, and Margaret Levi. 2017. Examining crowd work and gig work through the historical lens of piecework. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 4599–4616.
2. Nikolaos Askitas and Klaus F Zimmermann. 2009. Google econometrics and unemployment forecasting. *Applied Economics Quarterly* 55, 2 (2009), 107–120.
3. Sitaram Asur and Bernardo A Huberman. 2010. Predicting the future with social media. In *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, Vol. 1. IEEE, 492–499.
4. Scott R Baker and Andrey Fradkin. 2014. The Impact of Unemployment Insurance on Job Search: Evidence from Google Search Data. *Review of Economics and Statistics* (2014).
5. Anita Balakrishnan. 2017. Self-driving cars could cost America's professional drivers up to 25,000 jobs a month, Goldman Sachs says. (May 2017). `https://www.cnbc.com/2017/05/22/goldman-sachs-analysis-of-autonomous-vehicle-job-loss.html`
6. Pierluigi Balduzzi, Edwin J Elton, and T Clifton Green. 2001. Economic news and bond prices: Evidence from the US Treasury market. *Journal of financial and Quantitative analysis* 36, 4 (2001), 523–543.
7. Benjamin B Bederson, Jonathan Lazar, Jeff Johnson, Harry Hochheiser, and Clare-Marie Karat. 2006. Workshop on SIGCHI public policy. In *CHI'06 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1655–1657.
8. David M Blau and Philip K Robins. 1990. Job search outcomes for the employed and unemployed. *Journal of Political Economy* 98, 3 (1990), 637–655.
9. Joshua Evan Blumenstock. 2016. Fighting poverty with data. *Science* 353, 6301 (2016), 753–754.
10. Johan Bollen, Huina Mao, and Xiaojun Zeng. 2011. Twitter mood predicts the stock market. *Journal of computational science* 2, 1 (2011), 1–8.
11. Ilaria Bordino, Stefano Battiston, Guido Caldarelli, Matthieu Cristelli, Antti Ukkonen, and Ingmar Weber. 2012. Web search queries can predict stock market volumes. *PloS one* 7, 7 (2012), e40014.
12. Danah Boyd and Kate Crawford. 2012. Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, communication & society* 15, 5 (2012), 662–679.
13. Moira Burke and Robert Kraut. 2013. Using Facebook after losing a job: Differential benefits of strong and weak ties. In *Proceedings of the 2013 conference on Computer supported cooperative work*. ACM, 1419–1430.
14. Ronald S Burt. 2009. *Structural holes: The social structure of competition*. Harvard university press.
15. Stevie Chancellor, Yannis Kalantidis, Jessica A Pater, Munmun De Choudhury, and David A Shamma. 2017. Multimodal Classification of Moderated Online Pro-Eating Disorder Content. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 3213–3226.
16. Hyunyoung Choi and Hal Varian. 2012. Predicting the present with Google Trends. *Economic Record* 88, s1 (2012), 2–9.
17. Frank Cowell. 2011. *Measuring inequality*. Oxford University Press.
18. Sean Culey. 2017. Revitalizing The Rust Belt. `https://www.forbes.com/sites/realspin/2017/09/08/revitalizing-the-rust-belt`. (2017). [Online; accessed 18-September-2017].
19. Aron Culotta. 2014. Estimating county health statistics with twitter. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 1335–1344.
20. Francesco D'Amuri and Juri Marcucci. 2010. 'Google it!' Forecasting the US unemployment rate with a Google job search index. (2010).
21. Munmun De Choudhury, Meredith Ringel Morris, and Ryen W White. 2014. Seeking and sharing health information online: Comparing search engines and social media. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 1365–1376.
22. Drew Desilver. 2017. Most Americans unaware that as U.S. manufacturing jobs have disappeared, output has grown. (July 2017). `http://www.pewresearch.org/fact-tank/2017/07/25/most-americans-unaware-that-as-u-s-manufacturing-jobs\-have-disappeared-output-has-grown/`
23. Tawanna R Dillahunt, Nishan Bose, Suleman Diwan, and Asha Chen-Phang. 2016. Designing for disadvantaged job seekers: Insights from early investigations. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*. ACM, 905–910.
24. Tawanna R Dillahunt and Amelia R Malone. 2015. The promise of the sharing economy among disadvantaged communities. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 2285–2294.
25. Carl DiSalvo, Phoebe Sengers, and Hrönn Brynjarsdóttir. 2010. Mapping the landscape of sustainable HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1975–1984.
26. Susan Dumais, Robin Jeffries, Daniel M Russell, Diane Tang, and Jaime Teevan. 2014. Understanding user behavior through log data and analysis. In *Ways of Knowing in HCI*. Springer, 349–372.

27. Nicole B Ellison, Charles Steinfield, and Cliff Lampe. 2007. The benefits of Facebook "friends:" Social capital and college students' use of online social network sites. *Journal of Computer-Mediated Communication* 12, 4 (2007), 1143–1168.

28. Michael Ettredge, John Gerdes, and Gilbert Karuga. 2005. Using web-based search data to predict macroeconomic statistics. *Commun. ACM* 48, 11 (2005), 87–92.

29. Jane E Ferrie, Martin J Shipley, Michael Gideon Marmot, Stephen Stansfeld, and George Davey Smith. 1995. Health effects of anticipation of job change and non-employment: longitudinal data from the Whitehall II study. *Bmj* 311, 7015 (1995), 1264–1269.

30. Adam Fourney, Ryen W White, and Eric Horvitz. 2015. Exploring time-dependent concerns about pregnancy and childbirth from search logs. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 737–746.

31. Laura K Gee, Jason Jones, and Moira Burke. 2017. Social Networks and labor markets: How strong ties relate to job finding on Facebook's social network. *Journal of Labor Economics* 35, 2 (2017), 485–518.

32. Eric Gilbert and Karrie Karahalios. 2010. Widespread Worry and the Stock Market. In *ICWSM*. 59–65.

33. Jeremy Ginsberg, Matthew H Mohebbi, Rajan S Patel, Lynnette Brammer, Mark S Smolinski, and Larry Brilliant. 2009. Detecting influenza epidemics using search engine query data. *Nature* 457, 7232 (2009), 1012–1014.

34. Glassdoor. 2017. Statistical Reference for Recruiters: 50 HR and Recruiting Statistics for 2017. (2017). `http://resources.glassdoor.com/rs/899-LOT-464/images/50hr-recruiting-and-statistics-2017.pdf`

35. Sharad Goel, Jake M Hofman, Sébastien Lahaie, David M Pennock, and Duncan J Watts. 2010. Predicting consumer behavior with Web search. *Proceedings of the National academy of sciences* 107, 41 (2010), 17486–17490.

36. Erving Goffman. 1978. *The presentation of self in everyday life*. Harmondsworth.

37. Mark Granovetter. 1995. *Getting a job: A study of contacts and careers*. University of Chicago Press.

38. Mark S Granovetter. 1973. The strength of weak ties. *American journal of sociology* 78, 6 (1973), 1360–1380.

39. Anne E Green, Yuxin Li, David Owen, and Maria De Hoyos. 2012. Inequalities in use of the Internet for job search: similarities and contrasts by economic status in Great Britain. *Environment and Planning* 44, 10 (2012), 2344–2358.

40. Rodger W Griffeth, Peter W Hom, and Stefan Gaertner. 2000. A meta-analysis of antecedents and correlates of employee turnover: Update, moderator tests, and research implications for the next millennium. *Journal of management* 26, 3 (2000), 463–488.

41. Jamie Guillory and Jeffrey T Hancock. 2012. The effect of Linkedin on deception in resumes. *Cyberpsychology, Behavior, and Social Networking* 15, 3 (2012), 135–140.

42. Eszter Hargittai and Eden Litt. 2013. New strategies for employment? internet skills and online privacy practices during people's job search. *IEEE security & privacy* 11, 3 (2013), 38–45.

43. Ahmed Hassan Awadallah, Ryen W White, Patrick Pantel, Susan T Dumais, and Yi-Min Wang. 2014. Supporting complex search tasks. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*. ACM, 829–838.

44. Edwin AJ Van Hooft, Marise Ph Born, Toon W Taris, Henk Van Der Flier, and Roland WB Blonk. 2004. Predictors of job search behavior among employed and unemployed people. *Personnel Psychology* 57, 1 (2004), 25–59.

45. Benjamin Jen, Jashanjit Kaur, Jonathan De Heus, and Tawanna R Dillahunt. 2014. Analyzing employment technologies for economically distressed individuals. In *Proceedings of the extended abstracts of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 1945–1950.

46. Michael Kan. 2017. Laid-off IT workers worry US is losing tech jobs to outsourcing. (Mar 2017). `https://www.pcworld.com/article/3175682/techology-business/laid-off-it-workers-worry-us-is-losing-tech-jobs-to-\outsourcing.html`

47. Brian C Keegan, Patricia Cavazos-Rehg, Anh Ngoc Nguyen, Saiph Savage, Jofish Kaye, Munmun De Choudhury, and Michael J Paul. 2017. CHI-nnabis: Implications of Marijuana Legalization for and from Human-Computer Interaction. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 1312–1317.

48. Jaewoo Kim, Meeyoung Cha, and Jong Gun Lee. 2017. Nowcasting commodity prices using social media. *PeerJ Computer Science* 3 (2017), e126.

49. Kory Kroft and Devin G Pope. 2014. Does online search crowd out traditional search and improve matching efficiency? Evidence from Craigslist. *Journal of Labor Economics* 32, 2 (2014), 259–303.

50. Peter Kuhn and Hani Mansour. 2014. Is Internet job search still ineffective? *The Economic Journal* 124, 581 (2014), 1213–1233.

51. Jonathan Lazar. 2015. Public policy and HCI: making an impact in the future. *interactions* 22, 5 (2015), 69–71.

52. Jonathan Lazar, Julio Abascal, Simone Barbosa, Jeremy Barksdale, Batya Friedman, Jens Grossklags, Jan Gulliksen, Jeff Johnson, Tom McEwan, Loïc Martínez-Normand, and others. 2016. Human–computer interaction and international public policymaking: a framework for understanding and taking future actions. *Foundations and Trends® in Human–Computer Interaction* 9, 2 (2016), 69–149.

53. David Lazer, Ryan Kennedy, Gary King, and Alessandro Vespignani. 2014. The parable of Google Flu: traps in big data analysis. *Science* 343, 6176 (2014), 1203–1205.

54. David Lazer, Alex Sandy Pentland, Lada Adamic, Sinan Aral, Albert Laszlo Barabasi, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, and others. 2009. Life in the network: the coming age of computational social science. *Science (New York, NY)* 323, 5915 (2009), 721.

55. Colin Lindsay, Malcolm Greig, and Ronald W McQuaid. 2005. Alternative Job Search Strategies in Remote Rural and Peri-urban Labour Markets: The Role of Social Networks. *Sociologia Ruralis* 45, 1-2 (2005), 53–70.

56. Michael Luo. 2009. Job retraining may fall short of high hopes. *The New York Times* (2009).

57. Emil E Malizia and Shanzi Ke. 1993. The influence of economic diversity on unemployment and stability. *Journal of Regional Science* 33, 2 (1993), 221–235.

58. Huina Mao, Scott Counts, and Johan Bollen. 2011. Predicting financial markets: Comparing survey, news, twitter and search engine data. *arXiv preprint arXiv:1112.1051* (2011).

59. Huina Mao, Scott Counts, and Johan Bollen. 2015. *Quantifying the effects of online bullishness on international financial markets*. Technical Report. ECB Statistics Paper.

60. Ralph Matthews, Ravi Pendakur, and Nathan Young. 2009. Social capital, labour markets, and job-finding in urban and rural regions: Comparing paths to employment in prosperous cities and stressed rural communities in Canada. *The Sociological Review* 57, 2 (2009), 306–330.

61. Ruth G McFadyen and Jonathan P Thomas. 1997. Economic and psychological models of job search behavior of the unemployed. *Human Relations* 50, 12 (1997), 1461–1484.

62. Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.

63. Meredith Ringel Morris, Jaime Teevan, and Katrina Panovich. 2010. A Comparison of Information Seeking Using Search Engines and Social Networks. In *ICWSM*.

64. Bureau of Labor Statistics. 2010. Table B-1. Employees on nonfarm payrolls by industry sector and selected industry detail. `https://www.bls.gov/news.release/empsit.t17.htm`. (2010). [Online; accessed 17-September-2017].

65. Bureau of Labor Statistics. 2017. Labor Force Statistics from the Current Population Survey. `https://data.bls.gov/timeseries/LNS14000000`. (2017). [Online; accessed 14-December-2017].

66. Michael J Paul, Ryen W White, and Eric Horvitz. 2015. Diagnoses, decisions, and outcomes: Web search as decision support for cancer. In *Proceedings of the 24th International Conference on World Wide Web*. ACM, 831–841.

67. Michael J Paul, Ryen W White, and Eric Horvitz. 2016. Search and breast cancer: On episodic shifts of attention over life histories of an illness. *ACM Transactions on the Web (TWEB)* 10, 2 (2016), 13.

68. Jaroslav Pavlicek and Ladislav Kristoufek. 2015. Nowcasting unemployment rates with google searches: Evidence from the visegrad group countries. *PloS one* 10, 5 (2015), e0127084.

69. Tobias Preis, Helen Susannah Moat, and H Eugene Stanley. 2013. Quantifying trading behavior in financial markets using Google Trends. *Scientific reports* 3 (2013), srep01684.

70. Tobias Preis, Daniel Reith, and H Eugene Stanley. 2010. Complex dynamics of our economic life on different scales: insights from search engine query data. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 368, 1933 (2010), 5707–5719.

71. Aaron Smith. 2015. Searching for work in the digital era. *Pew Research Center: Internet, Science & Tech., Article published on* 19 (2015).

72. Marina Sokolova and Guy Lapalme. 2009. A systematic analysis of performance measures for classification tasks. *Information Processing & Management* 45, 4 (2009), 427–437.

73. Linnet Taylor, Ralph Schroeder, and Eric Meyer. 2014. Emerging practices and perspectives on Big Data analysis in economics: Bigger and better or more of the same? *Big Data & Society* 1, 2 (2014), 2053951714536877.

74. Nathan Tefft. 2011. Insights on unemployment, unemployment insurance, and mental health. *Journal of Health Economics* 30, 2 (2011), 258–264.

75. Jacob Thebault-Spieker, Loren Terveen, and Brent Hecht. 2017. Toward a Geographic Understanding of the Sharing Economy: Systemic Biases in UberX and TaskRabbit. *ACM Transactions on Computer-Human Interaction (TOCHI)* 24, 3 (2017), 21.

76. Sonja Utz. 2016. Is LinkedIn making you more successful? The informational benefits derived from public social media. *new media & society* 18, 11 (2016), 2685–2702.

77. Hal R Varian. 2014. Big data: New tricks for econometrics. *The Journal of Economic Perspectives* 28, 2 (2014), 3–27.

78. Tara Vishwanath. 1989. Job search, stigma effect, and escape rate from unemployment. *Journal of Labor Economics* 7, 4 (1989), 487–502.

79. Connie R Wanberg. 2012. The individual experience of unemployment. *Annual review of psychology* 63 (2012), 369–396.

80. Connie R Wanberg, Ruth Kanfer, and Joseph T Banas. 2000. Predictors and outcomes of networking intensity among unemployed job seekers. *Journal of Applied Psychology* 85, 4 (2000), 491.

81. Ryen W White and Resa A Roth. 2009. Exploratory search: Beyond the query-response paradigm. *Synthesis lectures on information concepts, retrieval, and services* 1, 1 (2009), 1–98.

82. Hans De Witte. 1999. Job insecurity and psychological well-being: Review of the literature and exploration of some unresolved issues. *European Journal of work and Organizational psychology* 8, 2 (1999), 155–177.

83. Diyi Yang, Zheng Yao, and Robert E Kraut. 2017. Self-Disclosure and Channel Difference in Online Health Support Groups.. In *ICWSM*. 704–707.