

Forecasting Task-Shares and Characterizing Occupational Change across Industry Sectors

Subhro Das^{1*}, Sebastian Steffen², Prabhat Reddy¹, Erik Brynjolfsson², Martin Fleming¹

¹MIT-IBM Watson AI Lab, IBM Research

²Massachusetts Institute of Technology

Abstract

Artificial Intelligence (AI) has started to transform our economy and society, more specifically, AI has the potential to make both labor and machines more productive while displacing certain human tasks and simultaneously introducing new tasks into the economy. Using online job postings data, this paper proposes novel methodologies to characterize the dynamic evolution of occupational task-share demands across different industries in the U.S. labor market and estimates the implied US-\$ market values for skills. The paper develops a multi-variate and multi-step long short term memory (LSTM) network architecture to estimate 12-month and 24-month ahead forecasts of task-shares with 10% root mean-squared error. The industry-specific insights on occupation evolution and forecasts on task-shares will facilitate the policy-makers and strategy leaders' decision-making to transform the current workforce for the future.

1 Introduction

Two of the seventeen United Nations' Sustainable Development Goals agendas are on 'Decent Work and Economic Growth' and 'Industry, Innovation, and Infrastructure'. The large-scale development and adoption of Artificial Intelligence (AI) and other automation technologies is fueling massive progress in *Industry, Innovation, and Infrastructure*, while simultaneously transforming the occupations and wages of workers and thereby impacting the *Decent Work and Economic Growth* agenda. In particular, the tasks required to do certain occupations will change, such that workers will be required to reskill in order to be employed productively. It is therefore critical to document and predict which skills are and will be demanded by the labor market and how the implied values of skills change.

The paper develops a forecasting pipeline using long short-term memory (LSTM), an artificial recurrent neural network

(RNN) architecture, for 12-months and 24-months ahead prediction of occupational task-shares with bounded root mean-squared errors. This study, being the first of its kind, introduces AI based tools to model and forecast the evolution of occupations in the US labor market. The timely forecasts of shifts in occupational task-shares will provide workers and employers enough time to retrain themselves or their employees, respectively, so as to stay relevant to the demands of the industries or job market in general. For several occupations, in particular low-wage ones, AI is predicted to outperform humans within the next decade leading to significant risks of long-term unemployment [Grace *et al.*, 2018], [de Troya *et al.*, 2018]. Characterization of *Future of Work* will facilitate the people of the country to fight against such unemployment challenges faced by the economy. With the increasing use of new AI technologies in job interview assessments [Hemamou *et al.*, 2019] and [Shen *et al.*, 2018], the future job seekers should accordingly prepare themselves to appropriately demonstrate their capability to carry out the demanded tasks.

1.1 Related Work

The demand for skills required to do occupational tasks is in constant flux. General Purpose Technologies (GPTs), such as AI [Brynjolfsson *et al.*, 2018], have the ability to induce significant societal change [Bresnahan and Trajtenberg, 1995], including the employment and the tasks of occupations. The second half of the 20th century saw a tremendous rise in demand and supply of 'high skill' workers, i.e. college graduates until, by around 2000 the wage premium to skill went down again to 1915 levels [Beaudry *et al.*, 2015]. In hindsight, this outcome seems evident from an economic perspective, but the race between technology and education is far from over. If anything, it has changed in favor of a race between man and machine [Acemoglu and Restrepo, 2018], as the 'second machine age' [Brynjolfsson and McAfee, 2014] begins. Technologies are achieving super-human performance on more and more tasks [Webb, 2019], in particular in routine tasks [Acemoglu and Restrepo, 2019], and also complement many non-automatable tasks, i.e. via data-driven decision making [Brynjolfsson and McElheran, 2016].

The significant fall of the labor share of GDP [Autor *et al.*, 2017] may be a sign of the scale tipping too far in favor of machines. This could result in long-term technological unem-

*This work was supported by the MIT-IBM Watson AI Lab. Corresponding authors: Subhro Das (subhro.das@ibm.com), and, Martin Fleming (fleming1@us.ibm.com).




Figure 1: Average monthly task-shares of task cluster families in the Finance, Manufacturing and Retail industry sectors.

ployment [Acemoglu and Restrepo, 2019] as routine-biased technological change progress and further increases inequality and wage polarization [Autor and Dorn, 2013], [Goos *et al.*, 2014]. Shifts in the relative demand for tasks have been shown to explain a significant part of the increases in earnings inequality [Atalay *et al.*, 2019].

However, whether the latest wave of technological change is different remains to be seen as it also brings about new tasks and occupations [Acemoglu and Restrepo, 2019]. While these new tasks include complex, cognitive, creative, and problem-solving tasks [Bartel *et al.*, 2007], which will not be automatable in the near future, they also include poorly-paid routine tasks, i.e. for labeling training datasets [Gray and Suri, 2019] or in the gig and service economy. Even STEM workers’ earning profiles, while starting high, stayed relatively flat [Deming and Noray, 2018], due to the fast rate of change in skill demand [Wu *et al.*, 2019] and skill obsolescence. This begs the question: which skills will stay in demand and what would be valuable reskilling investments?

To answer these important questions, this paper leverages a novel panel dataset of task demands within occupations and industries, which we derive from the near-universe of annotated US job postings between 2010 and 2017.

2 Task-Share Dynamics

In the pursuit to study the evolution of occupations in the United States labor market using online job postings data from Burning Glass Technologies (BGT)¹, we adapt a framework similar to [Fleming *et al.*, 2019] and [Das *et al.*, 2020]. This dataset consists of nearly 170 million annotated online job postings from employers all across the United States during the last decade (2010-2017). From each posting, the job titles are extracted and mapped to the Standard Occupational Classification (SOC) codes² and the North American Industry Classification System (NAICS) codes³. The SOC define a taxonomy of occupations are unique 6-digit codes assigned to 867 distinct occupations in the US. As per the US Department of Labor’s Bureau of Labor Statistics (BLS), all jobs are mapped to one of these detailed occupations (SOC codes) according to their occupational definition. Similarly, the NAICS codes are unique 6-digit codes that defines 317 unique industry groups. The first two-digits of the NAICS codes represent the industry sector, namely Finance, Manufacturing, Healthcare etc.

¹ <https://www.burning-glass.com>

² <https://www.bls.gov/soc>

³ <https://www.naics.com/history-naics-code>

Once each job posting is tagged to an occupation (SOC code) and an industry sector (NAICS code), the tasks, mentioned in the posting, were extracted and mapped to BGT’s taxonomy of around 17,000 tasks. The employers demand for these tasks to be performed by the workers in that occupation. The tasks are hierarchically clustered into 572 task clusters and 28 task cluster families. For example, ‘Book-keeping’ is a task in the ‘General Accounting’ task cluster which falls under the ‘Finance’ task cluster family. Note that a job posting might have tasks that falls under the ‘Finance’ task cluster family, but it may not be in the ‘Finance’ industry sector – a hospital which is under the ‘Healthcare’ industry sector has a job opening for financial bookkeeping.

2.1 Task-Shares across Industry Sectors

For all the job postings within each occupation and industry sector pairs, the number of mentions of each tasks are counted and the task demand frequency are computed on a monthly basis. This provided us the frequency counts of each tasks for every occupation-industry pairs for all the months from Jan-2010 to Dec-2017. The task frequency is normalized using the total number of task mentions, the number of employees in each occupation and the total number of employees across all occupations (obtained from the Annual BLS report), to derive the monthly task-shares of each task for an occupation-industry combination. The monthly task-shares are then aggregated for all tasks within a task cluster family and for all occupations within an industry sector to compute the task-shares of each task cluster family within an industry sector.

The monthly task-shares (averaged over an year) of the 28 task cluster families in finance, manufacturing and retail industry sectors is illustrated in Figure 1. The figure demonstrates that certain task cluster families were in high demand in specific industries for both 2010 and 2017, whereas others were not in such demand in that industry sector. However, the key thing to note here are the changes in task-shares of different task cluster families from 2010 to 2017. The task-shares of both *Design* and *Supply Chain & Logistics* task cluster families decreased within the manufacturing industry sector, in contrast both of them experience an increased demand in the retail sector. There was a significant growth in the task-shares of the *Information Technology* task cluster family in both finance and retail industries, however not-so-much growth in the manufacturing industry. The domain *Industry Knowledge* task cluster family saw a steady decline in its task-shares across all the three industries. Our methodology and the derived data can be leveraged to extract similar insights on the occupational evolution of other industry sectors in the US labor market.

3 Returns to Tasks

Estimating the returns to schooling, skills, as well as tasks goes back at least 70 years in the Economic literature to Jacob Mincer’s famous ‘Earnings Regressions’ [Mincer, 1958]. In his seminal work he estimates:

$$\ln w(s, x) = \alpha_0 + \rho_s s + \beta_0 x + \beta_1 x^2 + \epsilon,$$




Figure 2: Market Values (in Real \$) of Skills implied by Industry FE Regressions.

where $w(s, x)$ is the wage at schooling level s and work experience x ⁴. However, given the rise of college-educated workers, estimates of ρ , the returns to an additional year of schooling, are not a useful proxy for skill anymore. Furthermore, estimating returns from workers, i.e. the skill supply side, suffers from omitted variables bias (OVB) due to grit ([Duckworth and Gross, 2014]) and idiosyncratic preferences for work hours, location, or culture among others. Other worker-side proxy variables such as ability test scores [Hanushek *et al.*, 2015] or wage percentiles [Autor and Dorn, 2013] suffer from the same issue.

Instead, we use annotated job postings data, i.e. data from the skill demand side, to estimate the implicit market values that the labor market assigns to each skill. Skills in this case refers to the 28 skill cluster families from our data providers’ skill taxonomy. In particular, for each year we estimate:

$$\ln w_{ijt} = \theta_i + \theta_j + \sum_{k=1}^K r_{ijk t} s_{ijk t},$$

where w_{ijt} is the wage of occupation i in industry j at time t , the $r_{ijk t}$ are the implied returns to each skill k in occupation i and industry j at time t , the $s_{ijk t}$ are the skill shares of each skill k in occupation i and industry j at time t , and θ_i and θ_j are the occupation and industry-specific fixed effects, respectively. In particular, $r_{ijk t}$ can be interpreted as the dollar increase in wages associated with a 1% point change in skill k , all else equal. The implied returns to each skill k over time can be seen in figure 2.

We can see that the skills with the highest implied market value mainly fall into the cognitive and technical domains: Economics & Policy, Analysis, Design, Engineering, and IT skills are reach implied values of over \$1,000, ahead of Education & Training as well as Administration, Public Safety & National Security-related skills. We plan on extending our analysis into the future by rerunning our model with the predicted skill share values from our LSTM model described in section 4.1.

⁴See [Heckman and Carneiro, 2003] and [Firpo *et al.*, 2011] for excellent overviews.




Figure 3: The network architecture of multi-variate LSTM.

4 Task-Share Forecasting

Most state-of-the-art research on *Future of Work* studies the qualitative aspect of how the occupations are changing in the US Labor market, and more generally its global counterpart. Based on these qualitative trends, experts qualitatively predict the future state of occupations five to ten years from now. In contrast, this study quantitatively characterized the occupations as a bundle of tasks and derived their corresponding monthly task-shares in a time series format. This design of the study create a framework to learn the dynamics of the task-shares of different task cluster families and then use the trained models to predict future task-shares. Occupations being bundles of different tasks, such task-share forecasts will shed a quantitative light on the future state of occupations.

4.1 Multi-variate LSTM for Multi-step Prediction

The task-shares dynamics of each task cluster family are trained using an Autoregressive Integrated Moving Average (ARIMA) model in [Das *et al.*, 2020]. The trained ARIMA models, one for each task cluster family, are then used to make one-month ahead predictions of the task-shares. The limitations of [Das *et al.*, 2020] paper are: (a) the linear ARIMA model doesn't capture the intricate non-linear dynamics of the task-shares evolution; (b) the prediction model does not take into consideration the coupling effects between different task cluster families; and, (c) one-month ahead prediction is limited in scope for practical utility. In this paper, we address all the three limitations by implementing a multi-variate Long Short-Term Memory (LSTM) [Hochreiter and Schmidhuber, 1997], an artificial recurrent neural network (RNN) architecture, trained on the task-shares time-series data to predict 12-month and 24-month ahead forecasts of the task-shares of different task cluster families. LSTM networks are well-suited to learn the inherent non-linear aspects of the task-shares dynamics, and the multi-variate nature of the LSTM incorporates the coupling effect between the task-shares of different task cluster families while learning its network parameters.

4.2 LSTM Architecture & Forecast RMSE

The task-share data contains 29 task cluster families, each of which is split into 3 series based on wage tercile (high, medium, or low wage). After dropping the two series with

Layer(type)	Output Shape	Param
LSTM	(None, 200)	228800
Repeat Vector	(None, 24, 200)	0
LSTM	(None, 24, 200)	320800
Dropout	(None, 24, 200)	0
Time Distributed	(None, 24, 85)	17085

Table 1: Network layers of Seq2Seq encoder-decoder LSTM with dropout. The number of total parameters and trainable parameters are both 566,685.

Model	RMSE	NRMSE
LSTM 12 month	0.00102	10.06%
LSTM 24 month	0.000849	9.51%

Table 2: RMSE of Persistence and Seq2Seq encoder-decoder LSTM on 12 and 24 month ahead predictions

insufficient data, there are 85 task cluster family-wage tercile series. To leverage the coupling effect, the task-share time-series dataset is modeled as 85 parallel time-series, each consisting of monthly task-shares from 2010 to 2017 (96 months), upon which a Seq2Seq encoder-decoder LSTM model based on [Cho *et al.*, 2014] is trained for 3000 epochs. The architecture of this model is shown in Figure 3. Using this LSTM model, we make 12 and 24 month ahead predictions of task-shares. We compare the root mean squared error (RMSE) and normalized root mean square error (NRMSE) of predictions of the LSTM model to a persistence model as shown in Table 2.

The NRMSE is around slightly lower for the 24 month ahead forecast, which seems to support the robustness of this model over longer periods of time, when the effects of outlying monthly fluctuations are reduced. Our trained LSTM model is capable of predicting shifts in occupational task-shares 2 years ahead of time with around 10% RMSE. It is a clear demonstration of the successful quantitative characterization of occupational changes, and, of the ability to predict the upcoming changes in labor demands.

5 Conclusions

The empirical findings of this paper brings new insights on the *Future of Work* by representing occupations in terms of task-shares of the task cluster families. Specifically, the task-share changes in different industry sectors and the task-share dynamics forecast indicate that jobs are changing but slowly, giving the workforce sufficient time in learning new technical skills to adapt with the occupational change. The increasing implied market value of some of the technical skills seems to imply that the technologies to which they are complementary, will continue to be adopted. To prepare for continued adoption and advancements in the technologies, the development of accurate and robust predictive models are crucial to be able to provide guidance to workers, employers, and new graduates on skills and tasks of the future.

References

- [Acemoglu and Restrepo, 2018] Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 2018.
- [Acemoglu and Restrepo, 2019] Daron Acemoglu and Pascual Restrepo. Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 2019.
- [Atalay et al., 2019] Enghin Atalay, Phai Phongthientham, Sebastian Sotelo, and Daniel Tannenbaum. The evolution of work in the united states. *American Economic Journal: Applied Economics*, 2019.
- [Autor and Dorn, 2013] David H. Autor and David Dorn. The growth of low-skill service jobs and the polarization of the US Labor Market. *American Economic Review*, 2013.
- [Autor et al., 2017] David Autor, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. Concentrating on the Fall of the Labor Share. *American Economic Review: Papers & Proceedings*, 2017.
- [Bartel et al., 2007] Ann Bartel, Casey Ichniowski, and Kathryn Shaw. How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills. *The Quarterly Journal of Economics*, 2007.
- [Beaudry et al., 2015] Paul Beaudry, David A Green, and Benjamin M Sand. The Great Reversal in the Demand for Skill & Cognitive Tasks. *J. Labor Econ.*, 2015.
- [Bresnahan and Trajtenberg, 1995] Timothy F. Bresnahan and M. Trajtenberg. General purpose technologies 'Engines of growth'? *Journal of Econometrics*, 1995.
- [Brynjolfsson and McAfee, 2014] Erik Brynjolfsson and Andrew McAfee. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company, 2014.
- [Brynjolfsson and McElheran, 2016] Erik Brynjolfsson and Kristina McElheran. The Rapid Adoption of Data-Driven Decision-Making. *American Economic Review*, 2016.
- [Brynjolfsson et al., 2018] Erik Brynjolfsson, Tom Mitchell, and Daniel Rock. What Can Machines Learn and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, 2018.
- [Cho et al., 2014] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [Das et al., 2020] Subhro Das, Sebastian Steffen, Wyatt Clarke, Prabhat Reddy, Erik Brynjolfsson, and Martin Fleming. Learning Occupational Task-Shares Dynamics for the Future of Work. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 2020.
- [de Troya et al., 2018] Inigo de Troya, Ruqian Chen, Laura Moraes, Pranjal Bajaj, Jordan Kupersmith, Rayid Ghani, Nuno B Brás, and Leid Zejniliovic. Predicting, explaining, and understanding risk of long-term unemployment. In *NeurIPS Workshop on AI for Social Good*, 2018.
- [Deming and Noray, 2018] David J Deming and Kadeem Noray. STEM Careers and Technological Change. *NBER Working Paper*, 2018.
- [Duckworth and Gross, 2014] Angela Duckworth and James J. Gross. Self-Control and Grit. *Current Directions in Psychological Science*, 2014.
- [Firpo et al., 2011] Sergio Firpo, Nicole M Fortin, and Thomas Lemieux. Occupational Tasks and Changes in the Wage Structure. *IZA Discussion Paper Series*, 2011.
- [Fleming et al., 2019] Martin Fleming, Wyatt Clarke, Subhro Das, Phai Phongthientham, and Prabhat Reddy. The future of work: How new technologies are transforming tasks. In *MIT-IBM Watson AI Lab*, 2019.
- [Goos et al., 2014] Maarten Goos, Alan Manning, and Anna Salomons. Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review*, 2014.
- [Grace et al., 2018] Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans. When will AI exceed human performance? Evidence from AI experts. *Journal of Artificial Intelligence Research*, 2018.
- [Gray and Suri, 2019] Mary L Gray and Siddharth Suri. *Ghost work : how to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books, 2019.
- [Hanushek et al., 2015] Eric A Hanushek, Guido Schwerdt, and Simon Wiederhold. Returns to skills around the world: Evidence from PIAAC. *Eur. Econ. Rev.*, 2015.
- [Heckman and Carneiro, 2003] James Heckman and Pedro Carneiro. Human capital policy. Technical report, 2003.
- [Hemamou et al., 2019] Léo Hemamou, Ghazi Felhi, Vincent Vandembussche, Jean-Claude Martin, and Chloé Clavel. HireNet: a hierarchical attention model for the automatic analysis of asynchronous video job interviews. In *AAAI Conference on Artificial Intelligence*, 2019.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 1997.
- [Mincer, 1958] Jacob Mincer. Investment in human capital and personal income distribution. *J. Political Econ.*, 1958.
- [Shen et al., 2018] Dazhong Shen, Hengshu Zhu, Chen Zhu, Tong Xu, Chao Ma, and Hui Xiong. A joint learning approach to intelligent job interview assessment. In *International Joint Conference on Artificial Intelligence*, 2018.
- [Webb, 2019] Michael Webb. The Impact of Artificial Intelligence on the Labor Market. *NBER Working Paper*, 2019.
- [Wu et al., 2019] Xunxian Wu, Tong Xu, Hengshu Zhu, Le Zhang, Enhong Chen, and Hui Xiong. Trend-aware tensor factorization for job skill demand analysis. In *International Joint Conference on Artificial Intelligence*, 2019.