AI-Enabled Training in Manufacturing Workforce Development

Beverly Woolf, Aritra Ghosh, Andrew Lan, Shlomo Zilberstein, Tom Juravich, Andrew Cohen, Olivia Geho
University of Massachusetts Amherst
bev, arighosh, andrewlan, shlomo, juravich, alc, ogeho@umass.edu

Abstract

A highly productive workforce can evolve with the integration of digital devices, such as computer interfaces to operating machines, interconnected smart devices, and robots, in the workplace. However, this potential cannot be realized with the current state-of-the-art systems used to train workers. This problem is acute in manufacturing, where huge skills gaps are evident; most workers lack the necessary skills to operate or collaborate with autonomous systems. We propose to address this problem by using intelligent tutoring systems and worker data analysis. The worker data includes: i) fine-grained on-job performance data, ii) career path data containing the entire career paths of workers, and iii) job posting data over a long period of time indicating the required skills for each job. We will collect and analyze worker data and use it to drive new methods for training and reskilling workers. We detail ideas and tools to be developed by research in intelligent tutoring systems, data science, manufacturing, sociology, labor analysis, education, psychology, and economics. We also describe a convergent approach to developing effective, fair, and scalable software solutions and dynamic intelligent training.

Motivation

Workforce reskilling is of crucial importance to the US economy and society. As workers move to new jobs within the same occupation, they need to develop skills either on-the-fly or as part of a life-long learning process. Technologies, such as computer interfaces used to operate machines and connected devices in the internet of things, are already l leading to significant skills gaps between the current workforce and requirements of future jobs. Workforce reskilling is further complicated by the aging nature of the workforce where many workers lack the desire and commitment for upskilling. New competencies are required where they were not required in the past, e.g., basic computer interface skills, critical thinking, communication, and the ability to collaborate. A skill gap will result in an estimated 2.4 million jobs left unfilled between 2018 and 2028 (Deloitte & The Manufacturing Institute 2018). Current workers are critically aware of this skill gap, for example, 87% of polled manufacturing workers realized the importance of retraining and reskilling (Pew Research Center 2016).

Our research, funded by NSF[1], will train workers to use technology as part of addressing the future of work in the human-technology frontier. In the past, training workers typically involved teaching a fixed and static body of knowledge intended to be repeated by workers for many years in traditional workplaces. In contrast, artificial intelligence (AI) systems will fundamentally change the nature of work in many industries, which, in turn, will require the acquisition of new skills that may not be part of traditional training programs. And these are not just new static skills, but skills that will themselves require constant changes on the shop floor. For example, in manufacturing, a cobot might be programmed to thread a pipe one day and sort and pack for final packaging the next. Workstations will change over time as new robots and devices are incorporated. Therefore, an urgent need exists to seek training solutions that are flexible, effective and scalable for upskilling. This need raises key research questions about workforce training and upskilling: how can people be trained to work with and interact with autonomous systems to perform light reconfiguration, customization, or even collaborative tasks?

Partnerships

To address these issues, we formed partnerships with industry (Stanley Black & Decker Inc.), government (the City of Holyoke, Massachusetts and MassHire Holyoke Career Center), and academia (computer science, education, labor analysis, sociology, economics, and psychology). Stanley Black & Decker (SBD), is the number one tool provider in the world with 60,000 employees and 100 operating plants worldwide. SBD is rapidly adopting cutting-edge technologies and ensuring that a global workforce is prepared for these changes. One goal of SBD is to help build on Western New England's manufacturing legacy and establish it as a leading hub for the fourth industrial revolution. SBD's world-class Manufactory 4.0 facility, see Figure 1, is a global front-runner in advanced manufacturing, enabling the evolution of SBD in the new era

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[1]NSF's Convergence Accelerator Pilot, Team B-9615. See https://www.nsf.gov/od/oia/convergence-accelerator/index.jsp

of manufacturing. Our government partner, Holyoke, MA., a city with a high population of working-class immigrants, is one of the planned industrial cities that dotted the New England landscape. Between 1955-1970, its mills closed due to the changing economic landscape of early globalization and deindustrialization, which caused one of every two industrial jobs to vanish. Despite economic and social difficulties, the immigrant population grew significantly during the 1980s, aided by the 300% increase in the buying power of the Latino community from 1990 to 2016. Today, Latinos form the city of Holyoke's largest minority group, and the city has the largest Puerto Rican population per capita of any American city at 44.7%. As part of the diversification of its economic base, a number of precision manufactures have located and expanded in Holyoke, providing great opportunities for the upskilling of its workforce.

Figure 1: Stanley Black & Decker's Advanced Manufacturing Center of Excellence. The Manufactory 4.0 center is the home of the company's additive manufacturing accelerator program.

Our Approach

Our approach involves three branches: 1) Identify workers' needs; 2) Provide continuous skill development; and 3) Recommend a career path. Branch 1 is to conduct one-on-one interviews to identify the industry's needs, pain points, and opinions. Here we interviewed a variety of stakeholders (e.g., human resources executives, employment center leaders, industry executives), manufacturing workers, and unemployed people. Table 1 provides highlights and quotes from a subset of these interviews. The interviewees were keenly aware of the emerging importance of automation in manufacturing and both expressed concerns about the coming changes and recognized that the current level of training is inadequate.

Branch 2 is to develop software for continuous skill assessment (diagnose skills and provide training for upskilling) for manufacturing workers. Software tools will guide individual workers through the process of job selection and upskilling during their entire career. We will collect and analyze data from online job postings (Dave et al. 2018; Ramanath et al. 2018) and from data collected by corporate partners.

This software framework will perform worker profile DIagnosis, training program RECommendation, and intelligent Training platform development (DIRECT) for the purpose of continuous workforce development. DIRECT will help workers identify desirable future jobs and recommend training programs to prepare for those jobs. The software consists of three consecutive and intertwined components: (i) a skill level diagnosis and assessment component that uses cognitive models to assess worker skill levels from on-job data, (ii) a training experience development component that uses intelligent tutoring concepts to help workers acquire new skills and (iii) a skill gap identification component that uses labor market analysis to identify high-demand jobs and the skill gaps between a worker and their desirable job, see Figure 2.

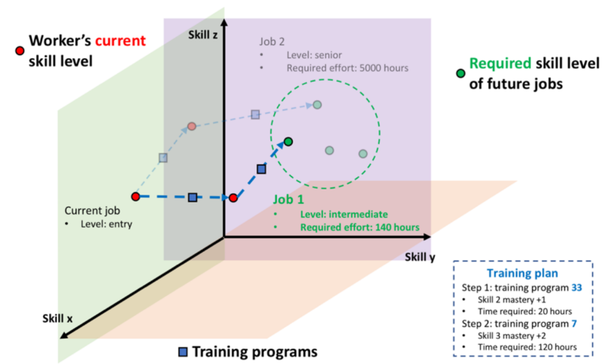

Figure 2: Continuous job and training program recommendation user interface. Workers are provided with visualizations of their current skill level (red dot, left), required skill level for alternative future jobs (four green dots), and several pathways of training programs to help them acquire the needed skills for a target future job (two dotted paths).

Branch 3 is to recommend career path choices, guide workers through the process of planning future career paths and provide a list of possible future jobs, including their seniority, compensation level, estimated upskilling effort needed to acquire the necessary skills, and the predicted likelihood that the worker will successfully get the job, given the worker's current skill levels. The future job and training program recommendation component will use predictive artificial intelligence algorithms to connect workers to future jobs and select training programs to acquire the necessary skills. It will support workers when selecting jobs they want to pursue using all the information provided and according to their personal interests. After a target job has been identified, the software will visualize i) the worker's current skill level, ii) the required skill levels of selected future jobs, and iii) personalized pathways consisting of a series of training programs for workers to acquire the skills needed for the jobs, together with the goal, estimated effect, and estimated effort required for each training program. DIRECT will support flexible worker inputs; if a user overrides its recommendation and manually constructs part of the training pathway, DIRECT will select training programs to complete the rest of the pathway.

The training programs can be either on-job training that will neither interrupt the daily workflow nor reduce productivity or it might be traditional training that requires workers to commit significant effort after work. The on-job training

On Automation in Manufacturing
“Robots are doing the time consuming yet easy and repetitive tasks. . . . [This] increases the amount of things that [workers] can do. So in manufacturing . . . [floor employee’s] time is filled up to the max. No matter how many machines you get it’s not going to make your job easier ’cause they’re just going to add more things for you to do.”
“It saddens me in a way, because [robots] are replacing humans. But then again, employers can’t find the humans to do the work. Is this what the future is going to be like?”
“Manufacturers can’t find workers, so they feel that increased automation makes sense . . . although the potential ‘saddens them’. A lot of young folks aren’t interested in manufacturing, partially due to the stereotypes . . . pushed by grandparents who used to work in manufacturing and remember it being dirty and dangerous.”
On skills needed in the workplace
“The skill set is different. It’s more scientific, it is more mathematic. It requires a different level of problem solving, even though there are still jobs that are basically to hit this button every 20 minutes, and then hope nothing bad happens.”
“Knowledge and willingness to embrace computers, measurements skills and ability to read engineering documents.”
“Being comfortable with computers and technology . . . nomenclature, knowing the terms, and knowing what different things mean. Being able to take good measurements and being able to read measurements and being able to read maybe an engineering document.”
“Management rarely ever has experience [as a line worker] on the floor.”
On introducing new technology
“Some rollouts . . . [include] adding a touch screen or a tablet to each machine. And the machine operators will now need to interface with that as a part of their job. Before they weren’t doing any of that at all, or were doing it by paper at the end of their shift.”
“And hopefully, you are able to train operators potentially to be . . . in charge of coordinating cobots and marking the points that they need to pick up.”
“Machinery is not connected and [companies] are rolling out new software and hardware. I am hopeful about training existing staff to work with cobots. The company is trying out implementation of technology in a wide variety of facilities and I [hope they can] roll it out companywide.”
“CNC operators are basically glorified button pressers . . . You have to have a good ear to recognize what are good sounds, what are bad sounds.”
“It takes more than three months to learn these [CNC] codes and to operate these machines . . . No, the training was definitely not enough to place me in front of a $600,000 machine, and run these codes.”
“Automation is a really big threat to people working in manufacturing.”
On floor workers in manufacturing
Floor workers often work 12 hour-days with no overtime pay.
Office workers have emails, laptops and access to a variety of computer-based training that this does not apply to plant employees. “So, how do you, even communicate with [floor workers]? And, how do you deliver them training when some of them . . . maybe have a flip phone, they don’t even have a smartphone?” Pilot projects have introduced tablets and other technology for line workers.
“Getting jobs in manufacturing is very unsafe due to increased automation.”
On training floor workers
“We are still training people for a previous economy. We need to build technical and computer-based technology more into education at all levels.”
“School didn’t prepare me for anything physical. K-12 prepares you for more academia.”
“Former prisoners entering the workforce typically have a 9th grade education.”

Table 1: Highlights and quotes from interviews with people in the manufacturing industry. Interviewees included manufacturing workers, executives, unemployed people, and executives at a career center.

option has the potential to avoid conflict with family responsibilities and to provide manageable upskilling experiences. We will develop smart intelligent tutors that are brief and adaptable, and will track human learning while workers interface with autonomous systems for agile manufacturing. Such software will teach basic principles for providing workers with supervised practice in applying agile principles and interacting with robots. Figure 3 shows an intelligent tutoring system we previously developed for manufacturing (Poli, Grosse, and Woolf 1999). We will evaluate and update this software with industry and government partners to ensure that DIRECT is effective and practical in this new context.

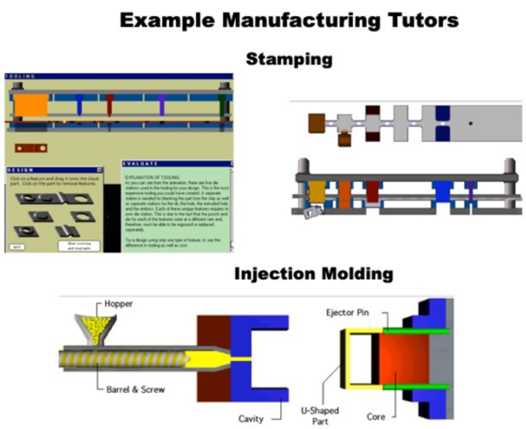

Figure 3: Example intelligent tutor for design for manufacturing. The Stamping Tutor (top) shows the tooling required to stamp the flat strip shown in the design window (top, left). The worker designs the slots and tabs (top right) to achieve the proper stamping template. The Injection Molding Tutor (bottom) demonstrates how pellets are fed from the hopper into the barrel to create the shape ejected.

Fairness and Transparency

The application of predictive algorithms in real-world settings faces ethical challenges. Without regularization and monitoring, algorithms trained on large data will inevitably develop biases against minorities and underrepresented workers if these people are not well represented in the database. For example, LinkedIn disproportionately advertised high-paying jobs to men (The Seattle Times 2016) and Facebook displayed considerable racial prejudice in censorship (USA Today 2019). This challenge is especially important for career planning and workforce development, since any bias in job and training recommendations can potentially impact employment outcomes and the well-being of millions.

For continuous training, we will develop algorithms that avoid potential AI-based discrimination and will create methodologies to detect and mitigate bias through the life cycle of AI applications. AI algorithms should not discriminate against protected classes of workers on the basis of gender, race, age, income level, disability status, and the like. We intend to develop algorithms that recognize that not every worker enters the workplace with the same resources. For example, a female worker may need to commit more effort to childcare than does a male worker; therefore, the algorithms need to take this aspect into account when making recommendations. Consequently, more sophisticated algorithmic and computational approaches are necessary to avoid making biased recommendations.

We note that equity is not the same as equality; equality means providing the same opportunities to all, while equity means providing specific resources and support as needed by an individual. For example, disadvantaged workers in a subgroup might need special training to bring them up to the same opportunity level as their peers. Equity means that personal and social factors are considered, as far as possible, in provided advice so that workers receive opportunities and experiences appropriate for them.

We will invite industry and government partners, human resources personnel and union representatives to identify workers in special classes to protect. We will conduct a series of pilot studies to qualitatively evaluate and adjust our algorithm to provide equity. We will also use a series of existing fairness metrics, e.g., individual fairness that requires that users with similar feature values be treated similarly (Dwork et al. 2012), outcome fairness that requires parity in the predicted probability of each outcome across user groups (drawn using sensitive attributes) (Zafar et al. 2017), conditional outcome fairness that requires parity in the predicted probability of each outcome given actual comes and regardless of sensitive attributes (Hardt et al. 2016), and counterfactual fairness that requires that the predicted outcome for each user remain mostly unchanged if the sensitive attribute changes (Russell et al. 2017).

Then, we will use these defined metrics of fairness to evaluate existing biases in workplaces. We will use data on worker career paths and on the recommendations they receive to identify whether biases exist against certain worker subgroups in past recommendations. We will start with aggregated statistics that are already collected by our partners, including statistics of workers with low-income, those who have disabilities, and those who are Black, Hispanic, and those who did not complete high school. This analysis will enable us to pinpoint underserved worker groups and to visualize trends in worker behavior and achievement. Finally, we will develop methods to minimize inequality in our recommendation algorithms.

Research shows that humans are especially averse to algorithms that make unjustified decisions, even when the algorithms outperform humans (Dietvorst, Simmons, and Massey 2015); once an algorithm makes an error, e.g., in a prediction or an analysis, humans are intolerant of it, although humans are more tolerant of humans making worse mistakes. Therefore, a need exists to ensure that recommendations are explainable and reliable so workers develop trust using them. To ensure transparency, within our autonomous systems and continuous training, we are developing explainable AI algorithms. Understanding how AI models arrive at specific decisions is a key principle of trusted AI. As shown in Figure 2, we will ensure that DIRECT outputs explainable recommendations that workers can understand and thus develop trust.

Preliminary Results

We have started to develop a career path recommender system (Branch 3) using a publicly available LinkedIn dataset. One objective is to correctly predict the job/company/skill that a user will select in the future. The dataset contains snapshots of public LinkedIn profiles for over a million users; each profile includes i) educational experiences, with each record containing the school, major, degree, and duration information, ii) professional experiences, with each record containing the job title, company, and duration information, iii) a set of skills the user possesses, and iv) miscellaneous information, such as industry and locality. We are studying the problem of forecasting the next career move of each user given their past experiences, a task first studied in (Li et al. 2017). Because the data contains only skills that are tagged to each user at the time the profile was accessed, not when they were tagged, we are also including an auxiliary task of skill prediction given the user's past experiences. We preprocess the data to filter out entities with low occurrence, resulting in 20K job titles, 60K companies, and 10K skills.

We employ a dynamic latent state model (Kalman 1960), which uses a latent skill vector to characterize each user at each time instant. Each entity (school, major, degree, job title, company) is parameterized by an embedding vector (Dave et al. 2018). We use a long short-term memory network (LSTM) (Hochreiter and Schmidhuber 1997) as our state transition model; this model changes the latent skill levels of each user over time as they engage with educational and professional activities. We use a fully connected neural network as our career move prediction model, which utilizes user skill levels and job title/ company embeddings to predict the probability of the corresponding career move. The skill tags are predicted in a similar fashion using the user's latent skill level at the time their profile is accessed.

In our experiments, we use an 80%-20% split of the dataset into training and test sets and evaluate the performance of our model using the Precision@N metric, which is the proportion of the top-n documents that are relevant. Precision@N quantifies how frequently the actual job title/company/skill is located in our top-N predictions across users in the test set. Table 2 shows our experimental results. For 20-40% of the career moves, the model correctly lists the actual job title/company/skill among its top 10 predictions, despite the presence of a large number of entities (more than 10K). As N gets larger, the precision numbers increase.

Precision@	10	100	1000
Job title	40%	59%	77%
Company	21%	45%	73%
Skill	35%	70%	92%

Table 2: Precision@N numbers for job title, company, and skill prediction tasks on the LinkedIn dataset. Precision at n is the proportion of the top-n documents that are relevant.

One advantage of using a latent state model is that we can use it to recommend career paths for future users. For a new user at a given point in time, we can use the model to estimate their current skill level given their past experiences and then search into the future to map out candidate career paths and future jobs for them. One ultimate goal is to open up future career paths for users and help them plan ahead; we will use this model to provide users with information on possible career moves that lead to a maximization of their personal objective (e.g., cumulative income in the next 10 years), under their personal constraints (locality, working time, etc.).

Discussion

As more technology is involved in the workplace of the future, workers need to continuously improve their skills. Robots, automation, and artificial intelligence will perform more tasks and create a massive disruption of jobs. A wide variety of education and skills-building programs will be needed to meet future demands of the workplace. We described research to interview workers to understand their needs, pain points, and training issues; develop software to diagnose a worker's skills and skill gap; create intelligent tutors to provide short effective training sequences; and develop software to provide career path recommendation. Open questions about training for the future of work include:

- Which content-based skills and capabilities need to be taught?
- Which life skills (e.g., emotional, social, behavioral), attributes and competencies need to be taught?
- Which skills can be taught effectively during short periods of time, while the worker is on the floor?
- Which skills will be most difficult to teach at scale?
- Will employers step up their own efforts to train and retrain workers?
- Will jobholders themselves engage in increased self-teaching efforts as they take advantage of proliferating online opportunities?

Education and training will need to adapt to prepare individuals for the changing labor market. Individuals need to engage in life-long learning not only to continue to be employable but also to achieve fulfilling and rewarding careers. Reskilling and upskilling strategies will be critical for companies to find the talents needed and to contribute to socially responsible approaches to the future of work.

References

[Dave et al. 2018] Dave, V. S.; Zhang, B.; Al Hasan, M.; AlJadda, K.; and Korayem, M. 2018. A combined representation learning approach for better job and skill recommendation. In *Proc. ACM International Conference on Information and Knowledge Management*, 1997–2005.

[Deloitte & The Manufacturing Institute 2018] Deloitte & The Manufacturing Institute. 2018. 2018 deloitte and the manufacturing institute skills gap and future of work study. online: http://www.themanufacturinginstitute.org/~/media/E323C4D8F75A470E8C96D7A07F0A14FB/DI_2018_Deloitte_MFI_skills_gap_FoW_study.pdf.

[Dietvorst, Simmons, and Massey 2015] Dietvorst, B. J.; Simmons, J. P.; and Massey, C. 2015. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144(1):114.

[Dwork et al. 2012] Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Proc. Innovations in Theoretical Computer Science Conference*, 214–226.

[Hardt et al. 2016] Hardt, M.; Price, E.; Srebro, N.; et al. 2016. Equality of opportunity in supervised learning. In *Proc. Conference on Advances in Neural Information Processing Systems*, 3315–3323.

[Hochreiter and Schmidhuber 1997] Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural Computation* 9(8):1735–1780.

[Kalman 1960] Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering* 82(1):35–45.

[Li et al. 2017] Li, L.; Jing, H.; Tong, H.; Yang, J.; He, Q.; and Chen, B.-C. 2017. Nemo: Next career move prediction with contextual embedding. In *Proc. International Conference on World Wide Web Companion*, 505–513.

[Pew Research Center 2016] Pew Research Center. 2016. The state of american jobs. online: https://www.pewsocialtrends.org/2016/10/06/the-state-of-american-jobs/.

[Poli, Grosse, and Woolf 1999] Poli, C.; Grosse, I.; and Woolf, B. 1999. Multimedia-based active tutors—A new approach to teaching design for manufacturing. In *Proc. ASME Design for Manufacturing Conference*, 1–9.

[Ramanath et al. 2018] Ramanath, R.; Inan, H.; Polatkan, G.; Hu, B.; Guo, Q.; Ozcaglar, C.; Wu, X.; Kenthapadi, K.; and Geyik, S. C. 2018. Towards deep and representation learning for talent search at LinkedIn. In *Proc. ACM International Conference on Information and Knowledge Management*, 2253–2261.

[Russell et al. 2017] Russell, C.; Kusner, M. J.; Loftus, J.; and Silva, R. 2017. When worlds collide: Integrating different counterfactual assumptions in fairness. In *Proc. Conference on Advances in Neural Information Processing Systems*, 6414–6423.

[The Seattle Times 2016] The Seattle Times. 2016. How linkedin?s search engine may reflect a gender bias. online: https://www.seattletimes.com/business/microsoft/how-linkedins-search-engine-may-reflect-a-bias/.

[USA Today 2019] USA Today. 2019. Facebook racism? Black users say racism convos blocked as hate speech. online: https://www.usatoday.com/story/news/2019/04/24/facebook-while-black-zucked-users-say-they-get-blocked-racism-discussion/2859593002/.

[Zafar et al. 2017] Zafar, M. B.; Valera, I.; Rogriguez, M. G.; and Gummadi, K. P. 2017. Fairness constraints: Mechanisms for fair classification. In *Proc. Conference on Artificial Intelligence and Statistics*, 962–970.