The Labor Market Impact of Artificial Intelligence

Ajay Agrawal, Joshua S. Gans and Avi Goldfarb

October 2018

1. Introduction

Perhaps the most popular topic with respect to the potential impact of artificial intelligence (AI) concerns what this technology means for jobs. As AI develops and mimics increasing levels of cognitive functions, the scope of jobs that might be impacted is great. This has motivated investigations into the nature of cognitive skills required for a wide array of occupations in order to identify those most likely to be impacted (e.g. Brynjolfsson, Mitchell, and Rock 2018; Frey and Osborne 2017).

Beyond these attempts at skills accounting, the literature on the broader, aggregate implications of AI on jobs emphasizes the ability of more productive capital to substitute for labor at the task level. Building on Acemoglu and Autor (2011), Acemoglu and Restrepo (2018a, b, c, d) provide what we anticipate will become the canonical model for understanding the impact of automation on labor. The central assumption is that, for an individual task, new technologies allow machines to substitute for workers if the machines are cheaper (in quality-adjusted pricing terms). This displacement induces a round of follow-on effects that can lead to lower or higher wages depending upon on what happens in other tasks that are complementary to the newly automated task, particularly with respect to the elasticity of substitution between tasks, and the ability to add more productive (or complex) labor-performed tasks to the production process (e.g. Acemoglu and Restrepo 2018a; Aghion, Jones, and Jones 2018; Brynjolfsson, Rock, and Syverson 2018; Guerreiro, Rebelo, and Teles 2018).

In our view, estimating the impact of AI on labor market outcomes requires an understanding of the particular tasks that AI will directly effect. Our goal in this article is to specify the characteristics of the technological change brought about by AI, and then to demonstrate how understanding these details provides useful insight into the labor market consequences.

Our starting point is the recognition that the current surge in AI development is driven by advances in a particular subfield of AI called machine learning. As we emphasize elsewhere (Agrawal, Gans, and Goldfarb, 2018a), this does not represent an increase in general intelligence of the kind that could substitute machine for human cognition, but rather one particular aspect of intelligence: *prediction*. We define prediction in the statistical sense of taking existing data to fill in missing information. Deep learning pioneer Geoffrey Hinton emphasized this view saying, "[t]ake any old problem where you have to predict something and you have a lot of data, and deep learning is probably going to make it work better than the existing techniques." (Hinton 2016)

Prediction is useful because it is an input into decision-making. Prediction specifies the confidence of a probability associated with an outcome under conditions of uncertainty. Thus, prediction is an input into decision-making under uncertainty. Decisions require predictions about the current and future states of the world. However, prediction is not the only element of a decision. Effective decision-making also requires data, the ability to take an action based on a decision, and an understanding of the objective function. The decision can be seen as a separate task from the prediction, though a decision is a perfect complement to prediction in the sense that a prediction has no value without a decision.

When AI is characterized specifically as a prediction technology, the labor market impact becomes clearer. Decision-making under uncertainty is essential to most occupations. Teachers decide how to educate students, managers decide who to recruit and reward, and janitors decide how to deal with a given mess. This wide breadth of application means that recent developments in AI represent a General Purpose Technology (GPT) as characterized by Bresnahan and Trajtenberg (1995).

Our goal in this paper is to provide a microeconomic characterization of the substitutes and complements of AI, as applied to Acemoglu and Restrepo's task-based framework. That framework examines the use and productivity of labor and capital at the task level. Specifically, a task, z's, output, $Y(z)$ is given by:

$$Y(z) = A_L \gamma_L(z) l(z) + A_K \gamma_K(z) k(z)$$

where $\{l(z), k(z)\}$ are the quantities of labor and capital employed on task z, $\{\gamma_L(z), \gamma_K(z)\}$ are task specific productivity terms, and $\{A_L, A_K\}$ are factor-augmenting technology terms common across tasks. This production function involves strict substitution between labor and capital at the task level. If W is the wage for labor and R is the rental rate of capital, then a task will be automated (that is, done by capital rather than labor) if:

$$\frac{A_L \gamma_L(z)}{A_K \gamma_K(z)} < \frac{W}{R}$$

In their framework, an automation technology is one that causes $\gamma_L(z)/\gamma_K(z)$ to fall for some task by a sufficient amount that capital substitutes for labor on that task. However, we argue that although this conception captures the notion of automation, the impact of AI on one task can simultaneously involve a distinct productivity impact on other tasks.

As we discuss above and in our prior work (Agrawal, Gans, and Goldfarb 2018a, b), the recent developments in AI are best seen as prediction performed by machines. For this reason, if a task primarily involves prediction (p), then the direct effect of employing AI is to substitute machines for humans in that task. However, every prediction task is associated with a decision task (d). That decision and associated actions can be characterized as a separate task, distinct from the prediction task, because they may be performed by different factors.

Importantly, there is always a task (d) that is a strict complement to the prediction task (p). Prediction reduces uncertainty and hence enables state-contingent decisions. First, this potentially improves the productivity of those decisions regardless of whether they are handled by a human or machine. Second, a reduction in uncertainty may have differential impacts on $\gamma_L(d)$ and $\gamma_K(d)$ for that task. As we will argue, there are cases where $\gamma_L(d)/\gamma_K(d)$ is likely to fall or rise as a result of AI being employed in the complementary prediction task (which, by definition, is represented as an increase in $A_K\gamma_K(p)$). In other words, while AI may involve some automation, it has direct flow-on effects to strictly complementary decision tasks, that have ambiguous impacts on the relative productivity of capital and labor in those tasks. Thus, it may not be appropriate to equate AI with past automation in order to predict its labor market impact either directly or in equilibrium.

We categorize the direct impact of AI on labor into four types. First, AI directly substitutes capital for labor in prediction tasks. Some of these tasks are already characterized as prediction tasks, such as demand forecasting. Others are not considered prediction tasks traditionally but are transformed into prediction-oriented tasks in response to the drop in the cost of prediction that results from advances in AI technology. For example, many human resources tasks that were previously considered to require emotional intelligence are now being recognized as prediction tasks. Recruiting, for example, is the task of predicting which subset of applicants will perform best in the job. Similarly, promotion is predicting which subset of existing employees will perform best in the next level up position. Likewise, retention involves predicting which star employees are most likely to leave. These predictions become inputs into a decision of whether to hire, promote, or retain. Whether a task was historically a prediction task or is newly-recognized as a prediction task, the effect of AI is the same – direct substitution of labor by capital. As we emphasize below, this may increase or decrease overall demand for labor within a workflow, depending on the effects on complementary tasks.

Automated prediction can increase or decrease the relative returns to capital versus labor in complementary decision tasks. When the relative return to capital increases, automated prediction leads to the complete automation of a complementary decision task. This is the second type of direct impact of AI on labor. An example is driving. A key task in driving is predicting what will happen next. For example, how will pedestrians, cyclists, and other drivers behave among all the static objects along the road? Autonomous driving programs outfit their vehicles with sensors to collect data about their surroundings and then train them to predict what will happen next and then to predict how a human driver would react to these predictions. Sometimes, the AI is able to make better predictions than a human could because it has access to higher fidelity input data, such as feeds from cameras, RADAR, and LIDAR around the car. For example, an AI detected a crash two cars ahead via RADAR - before it was humanly possible to predict.[1]

Once the prediction task is automated, it increases the returns to automating some of the complementary tasks, such as those associated with controlling the vehicle. For example, human reaction times are slower than those for machines. The value of a machine predicting a potential accident a few seconds or even a fraction of a second before a human would predict the accident is higher when the response time is faster. Thus, automating the prediction task increases the returns to also automating the decision tasks associated with vehicle control.

The third type of direct impact of AI on labor also applies to the impact of automated prediction on complementary decision tasks. This time, however, automating the prediction task enhances the relative returns to labor over capital. This arises when decisions are complex and/or involve trade-offs in payoffs that cannot easily be codified. In this situation, automated prediction may have no impact on the productivity of capital performing a complementary task but increase the productivity of labor. Consider ODS Medical. This company developed a way of transforming brain surgery for cancer patients. Previously, a surgeon would remove a tumor and surrounding tissue based on previous imaging (say, an MRI). However, to be certain all cancerous tissue is removed, the surgeon would end up removing more brain matter than would be necessary if they had better information. That information, however, could only be gained once the tissue is exposed.

Enter ODS. Their device, which resembles a connected pen-like camera, uses AI to predict whether an area of brain tissue has cancer cells or not. Thus, while the operation is taking place,

---

[1] https://www.youtube.com/watch?v=FadR7ETT_1k (accessed on October 3, 2018)

the surgeon can obtain an immediate recommendation as to whether a particular area should be removed. By predicting with over 90 percent accuracy whether a cell is cancerous, the device enables the surgeon to reduce both type I (removing non-cancerous tissue) and type II errors (leaving cancerous tissue). As a result, the efficiency of the brain surgeon is enhanced as prediction-contingent decisions are enabled during surgery. The effect of this is to augment labor as performed by brain surgeons. Put simply, given a prediction, human decision-makers are able to make more nuanced choices – effectively moving towards optimizing where they were previously satisficing.

The fourth and final type of direct impact of AI on labor involves the creation of new decision tasks. This happens when automated prediction sufficiently reduces uncertainty as to enable new decision tasks that did not exist before. The new tasks can be performed by capital or labor, depending on the relative costs of each. In other words, improvements in prediction enable new decisions (Agrawal, Gans & Goldfarb 2018b). In the absence of any prediction, actions are taken in a state or signal independent manner. For instance, without a weather prediction, one might choose to always carry or not carry an umbrella depending on personal preferences regarding being dry rather than wet and umbrella carrying costs. A weather prediction allows an agent to decide based on the outcome of that prediction. Specifically, if the likelihood of rain is high, then the agent is more likely to decide to carry an umbrella. In more complex environments, a prediction allows a decision that might otherwise not be considered.

This relates to the reinstatement force in Acemoglu and Restrepo (2018d) where a freeing up of labor as a result of automation increases the returns to technologies that use labor for new tasks. Examining the scope for new decisions in the context of our characterization of AI provides some guidance on where those new tasks might appear. In particular, cheap prediction means that some tasks that are not economically viable when uncertainty is high become viable as prediction technology reduces the level of uncertainty. One context where this effect will likely be important is innovation. Uncertainty limits innovation, and so better prediction might enable new types of innovation to occur. As highlighted by Cockburn, Henderson, and Stern (2018), recent advances in machine learning can be seen as the invention of a method of inventing, yielding new opportunities as Griliches (1957) first highlighted in the context of hybrid corn.

However, this path is not only a labor reinstatement effect. Labor is not the only factor that can benefit from state-contingent decision-making. Capital, and in particular computer software, can be augmented and hence introduced to new tasks if it can be programmed to react to environmental changes in a state-contingent manner (Agrawal, Gans & Goldfarb, 2018b, 2019). Thus, while AI might lead to automation and a substitution of capital for labor in existing tasks, it is also possible that it may lead to capital being preferred for new tasks. Unlike labor, however, the price effects do not necessarily reinforce such investments. Furthermore, we emphasize that, at this point, the hypothesis that AI will lead to new tasks—for humans or machines—is mostly speculative. It is difficult to identify new tasks that have already arisen because of recent advances in prediction technology.

To summarize, the overall impact of AI on the labor market in the short and long-term depends upon the balance of substitution, complementarity, and demand expansion. In the context of our four direct effects, the net long-run effect of AI on labor demand is determined by:

Net Reduction to Aggregate Labor Demand
= Labor Time Saved by Automating Prediction Tasks + Labor Time Saved by Automating Decision Tasks – Labor Time Increased by Increasing the Demand for Decision Tasks Performed by Labor – Labor Time Created by New Decisions

We are not able to assess the net effect of these types of labor demand changes. Instead, our goal is to enhance the interpretation of Acemoglu and Restrepo’s model by increasing the specificity with regards to the types of tasks that can be automated and the resultant effect on other tasks by carefully considering the particular characteristics of machine learning. We provide anecdotal evidence consistent with our interpretation by describing some early applications of commercial AI.

Using data and cases, we provide examples to illustrate each type of effect of AI on labor. Before doing so, we establish some building blocks. In the next section, we explain why recent advances in AI are best seen as prediction technology. We describe how prediction technology fits into a broader task-based model of technology, automation, and labor markets. Examining 199 AI startups, we show the degree to which machine learning leads to the automation of jobs, providing support for a task-based model in which only certain elements of a job are done by machine. We then provide examples of the four effects through which advances in prediction technology affect labor in a task-based framework: automation of prediction tasks, automation of decision tasks, augmenting labor on decision tasks, and the generation of new tasks due to decreased uncertainty. We conclude by emphasizing that we cannot assess the net effect of AI on labor, even in the short run. Instead, most applications have multiple forces that impact the effect of AI on jobs, both increasing and decreasing the demand for labor. The net effect is an empirical question and will vary across applications and industries.

2 Prediction, Tasks, and Automation

Before we move to an in-depth discussion of the impact of AI, it is important to define our use of the term. The majority of recent achievements in AI are the result of advances in machine learning, a branch of computational statistics. Much of the material in standard machine learning textbooks (e.g. Alpaydin 2010 and Hastie, Tibshirani, and Friedman 2009) is familiar to economists: regression, maximum likelihood estimation, clustering, non-parametric regression, etc. Other tools have recently begun to enter the econometrician's toolkit: regression trees, neural networks, and reinforcement learning. Over the past decade or so, advances in computer speed, data collection, data storage, and algorithms have led to substantial improvements such that commercial applications are proceeding rapidly.

Much of the public attention on AI focuses on the technology's potential to automate processes. A process (or task sequence) is automatic if it is performed without human assistance. Automation occurs when a process previously performed by a human becomes automatic. Thus, by definition, automation leads to the direct substitution of machines for humans in the conduct of certain tasks. If AI facilitates automation, then there is a direct substitution effect. Any increase in the role for human work has to come through changes in processes and opportunities that arise because of automation.

A key element to anticipating the degree to which AI will lead to automation involves understanding how the technology will be utilized at the task level which in turn requires specifying the scope of particular tasks. In the context of machine learning, there are two natural ways to define tasks. First, making a prediction is a common task in many workflows. Thus, advances in prediction technology raise the possibility of substituting machines for labor in those tasks. Second, making a decision, which includes deciding to take an action, is a task. An interesting and important link between these two types of tasks is that better prediction (task type 1) could increase the value of human labor in decision-making (task type 2).

Classifying commercial applications as labor-replacing – a motivation for distinguishing between prediction versus decision tasks

To develop a sense for how AI is being used in commercial applications, we examined 199 AI startups that were accepted into the Creative Destruction Lab (CDL) program we founded that helps science-based startups scale. These startups participated in the CDL programs in Toronto and Montreal in 2016 and 2017. The data consists of program applications submitted by each of these startups in response to a standard series of questions. Of these 199 companies, 168 have products targeted towards enterprise customers, 15 towards end consumers, and the remainder towards both.

Two research assistants classified whether each startup described their value proposition as labor-replacing and whether each would automate a business process or provide a tool that would enhance productivity but not fully automate a process. The two main questions used in the coding were “Describe the value proposition of your company” and “Describe how a customer would use your product.”

35 of the 184 startups with enterprise customers describe the value proposition of their business as labor replacing. An additional 3 describe how their technology will enable enterprise customers to fully automate processes, even though they do not explicitly mention that the technology is labor-replacing. Furthermore, 3 of the 15 aimed at end-consumers describe the value proposition as related to automating processes. For example, one is focused on automating language instruction. Thus, it is labor-replacing in the sense that it claims to enable someone to learn a language without the need to converse with a human or to have a human instructor.

Of course, these companies are not a representative sample of AI applications in business, or even a representative sample of AI startups. They are startups that were accepted into a particular program based in Canada. Recognizing the non-representativeness of this sample, we interpret our examination of these 199 startups to suggest that there are a large number of AI applications that do not involve automation of broad tasks. Instead, many applications automate a very specific prediction task.

The process of coding which innovations were labor-replacing and which automated business processes yielded an additional insight. It was relatively straightforward to code whether a given company’s technology would automate a business process. It involved little more than reading the company’s description of how a customer would use the product and assessing whether that meant a process previously done by a human would now be done by a machine. However, it was much harder to assess whether a technology was labor-replacing or labor-enhancing. In the end, we coded whether the description highlighted some labor-replacing process but did not code whether there were also labor-enhancing effects.

In doing this exercise, it became clear that most labor-replacing technologies were also labor-enhancing. The overall effect depends on the elasticity of demand for the end product and the degree of complementarity between the labor-saving task and the other tasks that humans perform in the workflow. Even for companies whose products seem to directly replace labor, we found their descriptions highlighted some labor-enhancing aspects. For example, A&K Robotics builds robots that autonomously clean office floors. It was coded as labor replacing and automating a business process. Nevertheless, in describing their value proposition, the company emphasizes the simultaneously labor-replacing and labor-augmenting points that one employee can clean more square footage.

Furthermore, the data highlight the link between prediction tasks and decision tasks. Sometimes automating the prediction increases the returns to automating the decision task as well. These are cases of full automation. Other times, automating the prediction increases demand for labor in the complementary decision task. We provide a number of examples below. This distinction motivates our focus on specifying the impact of AI on particular tasks and distinguishing between prediction and decision tasks.

3 Automating Prediction Tasks

Many jobs involve prediction as a key task in the workflow. Taxi drivers predict what other drivers will do and the time required to travel different routes. Financial analysts predict the returns to different assets. Translators predict the set and order of words in one language that match the meaning of sentences and paragraphs in another language. Professional athletes predict the trajectory of a ball.

In other words, predictions are important in most if not all occupations. In some situations, such as anticipating the path a ball might take in order to catch it, predictions become intuitive and thus require little time to generate. In others, generating predictions takes significant time and effort. When an AI generates a prediction, the human time is saved.

In this section, we describe examples that highlight substitution in the prediction task. Importantly, each example in this section also relates to another section: either the automation of the prediction task increase the relative returns to capital in the decision task and lead to automation of the decision, or it will increase the relative returns to labor in the decision task and increase demand for labor in the decision.

Prediction in legal services

Legal work involves a significant amount of prediction. A number of AI applications substitute capital for labor by automating prediction tasks in legal work, while still leaving the decision tasks to the human lawyer. We describe three examples.

Consider the process of redacting (or blacking out) sensitive information so that documents can be disclosed to a wider audience. There are penalties for over-redacting and not disclosing relevant, non-confidential information. At the same time, there are costs to disclosing confidential information, thus requiring care in identifying sensitive words, phrases and paragraphs. Chisel is a company that uses AI to automate the redaction process. The Chisel AI scans documents and then predicts the elements that are confidential and should be redacted. A lawyer then reviews those elements and confirms the redaction. In doing this, Chisel errs on the side of identifying potentially confidential items with the lawyer doing the job of deciding whether or not to follow the prediction. Thus, the AI does not replace the human making the redaction decision, but instead it reduces the time the human spends predicting what might require redaction, which then serves as an input into the human decision.

In another example, the Kira Systems AI scans contracts and summarizes relevant content. This may involve predicting which party in a lease agreement is liable for what actions or expenses. Or it may involve scanning all of the contracts signed by a firm to predict which ones would be impacted if that firm were involved in a merger or acquisition. Like Chisel, it is still up to human lawyers to make the decisions (as regulation requires) but Kira's technology predicts the relevance of clauses and information in a fraction of the time it would take a lawyer or paralegal.

Finally, AIs are being used to predict outcomes based on a wide corpus of legal judgments. Blue J Legal's AI scans tax law and decisions to provide firms with predictions of their tax liability. Consider income classification. Judges take case facts and use legal reasoning to come to a judgment. In tax law, there is often ambiguity on how income should be classified. One example is in the classification of trading in securities as income from business or as capital gains.[2] At one extreme, if someone trades multiple times per day and holds securities for a short time period, then the profits are likely to be classified as business income. In contrast, if trades are rare and assets are held for decades, then profits are likely to be classified by the courts as capital gains. Of course, most legal cases sit between these extremes. Currently, when a lawyer takes on a case, they collect the facts of the case from the client and conduct research on similar cases. They think carefully about what critiques are likely to arise from the other side in terms of matching the facts of the case to past judicial decisions. Blue J Legal uses machine learning to predict the outcome of new fact scenarios in tax and employment law cases. In addition to a prediction, the software provides a "case finder" that identifies the most relevant cases that help generate the prediction.

The end result of this process is not certainty. In the securities trading example above, the AI predicts the likelihood of particular case facts being classified as business income or capital gains. As Blue J Legal founder Benjamin Alarie describes it, judges take the input of facts found at trial and output a judgment, using legal reasoning as a mapping function from inputs to outputs. In contrast, the Blue J Legal AI utilizes test-facts that are assumed and entered into the system by the user, rather than the case facts found by trial, to generate a prediction of what a judge would decide with those facts. Instead of legal reasoning, the mapping function is a prediction generated by machine learning that is based on training data of past cases.[3] Blue J Legal claims 90 percent accuracy. Given the uncertainty, a lawyer still makes the ultimate decision.

These examples show how enhanced prediction, delivered by a machine, can be used to improve the productivity of decisions. In legal work, as lawyers still make the ultimate decision, it is a complement to the lawyer's skills. That said, it is saving labor with respect to the time it would take a human to generate the predictions. Thus, it is hard to predict how the effect of these AIs will show up in the aggregate labor statistics for legal work. AIs substitute for lawyers' prediction tasks, but may create opportunities at the decision-task level because better prediction might affect prices and quantities in a way that increases demand for decisions overall.

[2] https://www.youtube.com/watch?v=-pgbQejGVPc&submissionGuid=c9852990-2f9e-46e6-8cd2-b7507bebcfe9
[3] https://www.youtube.com/watch?v=-pgbQejGVPc&submissionGuid=c9852990-2f9e-46e6-8cd2-b7507bebcfe9

Prediction in driving

Prediction technology is changing the way we drive. While the headline grabbing technology is the potential for mass adoption of fully autonomous vehicles, prediction technology is already changing driving in a number of ways that do not replace human drivers with machines.

Vehicle manufacturers use AI to warn drivers about imminent risks. In other words, the AI predicts hazards and provides the driver warnings like “there is probably a car in your blind spot” or “there is likely a pedestrian behind your car” in the form of a beep or blinking light. The machine provides the prediction but the driver is still responsible for the decision of whether to stop, turn, or proceed.

Vehicle maintenance scheduling is a prediction problem. Rust (1987) developed an empirical model of Harold Zurcher, the superintendent of maintenance at the Madison Wisconsin Metropolitan Bus Company. Using statistical predictions of Zurcher’s decisions, the model could be used to substitute for his predictions about when buses would break down. Today, this practice is increasingly common. Advances in sensors and prediction algorithms have led to many new products that predict when a vehicle will break down, and thus inform the decision of whether to bring a vehicle in for maintenance.

Finally, prediction is changing commercial driving by providing effective predictions of the most efficient route between two locations at a given time. Perhaps the most dramatic example is the case of London taxi cabs. For decades, earning a taxi license in London required learning ‘The Knowledge.’ This involved learning the location of every address in London as well as the shortest route between two addresses. To pass the resulting test took two to four years of study with the help of specialist training schools. The end result was that London cab drivers were very knowledgeable about their city.

Today, AI -powered best-route prediction apps like Waze deliver ‘The Knowledge’ to any driver with a smartphone. This enables ride sharing services such as Uber to compete with London taxis. This competition would not have been possible without the enabling prediction technology. Thus, although the skill of London cabbies did not diminish, their competitive advantage was seriously eroded by AIs that provided labor-augmenting support for those who had not invested in learning The Knowledge. The end result on employment is ambiguous. While it is surely negative for London cabbies, overall it may be positive as more drivers enter the market (assuming they have permission from regulators).

Predictions in email responses

People in almost all professions respond to email. While we might not think of composing an email response as a prediction problem, it can be formulated as one. For example, Google developed Smart Reply for its email service, Gmail, that uses artificial intelligence to scan incoming emails and predict possible responses. Smart Reply doesn't automate sending the email response, but rather predicts possible responses and provides the user with three suggestions. The user then makes the final decision about whether to send one of the predicted responses or to compose and send their own response. In 2018, within weeks of Google rolling out Smart Reply as a default setting for all its active 1.4 billion active Gmail accounts, 10% of all Gmail responses sent were generated by the AI.[4] For most workers, this represents the automation of a prediction, not a decision. It saves the user the time of composing a response in cases where one of the three predicted replies is sufficient. However, the user must still decide whether to send a predicted response or to compose one directly.

Unlike the previous examples of automating prediction, for some jobs this automates a decision task. For example, even though the tool works in precisely the same fashion for executive assistants as for other occupations, it automates *their* decisions. For the executive assistant, the decision was to draft a possible response that someone else would decide on regarding whether to send or not. In the O*NET data, this task is described as "Prepare responses to correspondence containing routine inquiries." O*NET lists eight other occupations that involve a similar task: Correspondence Clerks, Tellers, Receptionists and Information Clerks, License Clerks, Legal Secretaries, Insurance Policy Processing Clerks, Medical Secretaries, and Loan Interviewers and Clerks. This raises the possibility that the same AI implementation might automate a prediction in some occupational settings, where the worker must still apply judgment about the benefits and costs of a particular decision before deciding or taking an action but automate the full decision in other occupations where the central task is the prediction because the judgment is provided by someone else in a different occupation. Such automated decisions are the subject of the next section.

[4] https://www.wsj.com/articles/very-interesting-awesome-love-it-gmail-users-confront-chipper-smart-reply-1537282569?mod=searchresults&page=1&pos=10 (accessed on September 30, 2018)

In this section, we provided examples of prediction tasks traditionally performed by humans that may increasingly be done by machines. In these cases, machine prediction displaces human prediction. This is common to *all* examples considered here – AIs automate prediction at the task level. However, we can say little about the overall effect on jobs. The humans performing prediction tasks that are replaced by machines may find new opportunities in other tasks. Alternatively, other humans may identify new opportunities if prediction is removed as a bottleneck to their participation in the labor market, as in the London taxicab example. We offer these examples to provide some specific context for what it means for humans to be replaced by machines in prediction tasks in the context of AI.

4 Automating Decision Tasks

We now describe the conditions under which automating the prediction task increases the relative returns to automating the decision task, compared to performing the decision task with human labor. In other words, both the prediction and the decision are automated. This represents the full automation of tasks a la Acemoglu-Restrepo. For this to happen, it must be possible to specify the desired action to be taken for each realization of uncertainty (that is, for each realization of a prediction). Thus, the more complex the external environment, the more difficult it is to code a machine to take the decision. For this reason, the most common type of machine-based decisions are those that are binary – say, to reject or accept a credit application or to recommend or reject a candidate for a job interview. However, as AI improves, it can provide better predictions in more complex environments where the action space is correspondingly larger. In other words, better prediction will lead to more automation of decisions in situations where a human is not required to be reactive within a decision context.

We have started to see this type of automation in environments where machine learning techniques are applied to mimic human decision-making. For example, a machine fitted with sensors is trained by observing the choices made by a human operator. With sufficient observations, the machine can learn what action a human would take given different sensory inputs. For instance, the autonomous operation of vehicles on public roads has been advanced by humans driving millions of kilometers in appropriately equipped vehicles that are able to collect both the perception data regarding environmental conditions on the road (input) and the action data regarding the decisions made by human drivers behind the wheel (output) in response to the perception data. Even in this case, however, there is still a measure of human supervision required depending on the probability of ‘edge’ cases arising for which the machine has not been appropriately trained.

Autonomous driving is an area for considerable research precisely because so much time is spent by so many people driving. In other words, the labor-saving time from automation is potentially large. However, this is also an area where the challenge of completely removing human operation is difficult because the cost of failure can be so high.

Cleaning illustrates a more pedestrian attempt at automation. Tasks such as household floor cleaning have been transformed by iRobot with its Roomba and now a wealth of competitor products. In commercial applications, however, where the cleaning demands are higher and the scale of the operations larger, automated cleaning is more of a challenge. A&K Robotics is using AI to tackle this problem. Their approach is to take existing, human operated cleaning devices, and to retrofit them with sensors and a motor to operate autonomously. In other words, they are basically adding prediction to existing machines. If successful, the human operators will no longer be necessary. AI enables prediction of the correct path for the cleaning robot to take, and of the identity of unexpected surprises that appear in that path, such as humans. Given these predictions, it is possible to pre-specify what the cleaning robot should do, and so the decision task can also be automated. The company emphasizes how this can increase workplace productivity, reduce workplace injuries, and reduce costs.

AI has also enabled the automation of driving vehicles in warehouses. A key task is to move items from the storage part of the warehouse to the packing and shipping part. Much of this automation occurs without machine learning, by simply using dedicated tracks for delivery vehicles. Without machine learning, however, humans play an important role in predicting hazards and guiding vehicles. Recent applications of AI enable swarms of robots to predict optimal routes and avoid collisions, eliminating the need for human controllers to decide on route planning. Under these conditions, warehouse vehicles can be fully automated.

Similarly, vehicle automation is growing in the mining industry, in particular for remote operations. For example, in Australia’s Pilbara region, the iron ore mining sites are over 1000 miles from the nearest major city. Given the remoteness of the region, and the extremely hot temperatures, it is expensive to bring human workers on site. Truck drivers are therefore unusually expensive. Mining giant Rio Tinto initially solved this problem by driving the trucks remotely from the offices in Perth, but in 2016 the company went a step further. It deployed dozens of self-driving trucks, saving operating costs. AI made this automation of the steering decision task possible by predicting hazards in the trucks' way and by coordinating the trucks with each other. As with robot cleaning and robots in warehouses, better prediction was the last step in removing humans from the decisions involved in driving in these relatively controlled settings. Unlike autonomous vehicles on public roadways, in these environments with far fewer edge cases, cheap prediction has already led to widespread automation of the decisions.

In this section, we describe the conditions under which automation of the prediction task enables automation of the decision task. Again, we can say little about the overall effect on jobs. However, with full automation of both the prediction and the decision, we identify the types of jobs that are unlikely to expand, such as driving vehicles in controlled environments like factories and warehouses. While this might make other tasks more efficient, it identifies the types of jobs where the core bottleneck to automation is prediction and thus are candidates for elimination.

5 Augmenting Labor on Decision Tasks

There is a tendency in current discussions to equate AI with full automation. As we have just shown, there are situations where this happens. Generally, better prediction improves decision-making regardless of who is making the decision – a human or machine. In this section, we discuss examples where the automation of prediction can enhance the productivity of labor, relative to capital, in the decision task. In other words, the examples show how AI can improve decision-making by humans and consequently the productivity of labor – specifically, allowing them to make state-contingent decisions that reduce errors, enhancing payoffs.

Bail decisions

One of the decisions that judges make concern whether or not to grant bail.[5] Kleinberg et al (2017) study the predictions that inform this decision. "Soon after arrest, a judge *decides* where defendants will await trial, at home or in jail. By law, this decision should be based solely on a *prediction*: What will the defendant do if released? Will they flee or commit a new crime?...

[5] Bail as defined by the Oxford English Dictionary: The temporary release of an accused person awaiting trial, sometimes on condition that a sum of money is lodged to guarantee their appearance in court.

Currently the predictions on which these decisions are based are, in most jurisdictions, formed by some judge processing available case information in their head.” (emphasis added by us)

Although AIs enhance prediction accuracy, they do not automate the decision. This point is somewhat obfuscated in Kleinberg et al (2017). On the one hand they write “By law, this decision should be based solely on a *prediction…”* On the other hand, later on in the same paragraph, they note that judges must apply judgement to predictions when making a decision. Essentially, they must weigh the relative cost of type I and II errors. “A judge must trade off these risks [flee or commit a new crime] against the cost of incarceration. This is a consequential decision for defendants since jail spells typically last several months (or longer); recent research documents large costs of detention even over the long term. It is also costly to society: at any point in time the US has over 750,000 people in jail, disproportionately drawn from disadvantaged and minority populations.”

Thus, in this application, AIs automate the prediction, but not the decision—because judges weigh the relative costs of errors and because the US legal system requires human judges to decide. So, these AIs will not replace judges, but they could enhance their productivity. How much would this affect jobs? The amount of time spent making bail decisions varies across judges. Overall, according to O*NET’s characterization of the occupation “Judges, Magistrate Judges, and Magistrates” only two of the 19 tasks are associated with making bail decisions: 1) “Impose restrictions upon parties in civil cases until trials can be held”, and 2) “Conduct preliminary hearings to decide issues such as whether there is reasonable and probable cause to hold defendants in felony cases.” So, on the one hand, the effect of this type of AI may not be significant for judges as they only support a small fraction of overall tasks. On the other hand, prediction accuracy improvements could be consequential, especially given that the police arrest over 10 million people per year in the United States. Based on AIs trained on a large historical dataset to predict decisions and outcomes, the authors report simulations that show enhanced prediction quality could enable crime reductions up to 24.7% with no change in jailing rates, or jailing rate reductions up to 41.9% with no increase in crime rates. In other words, if judicial output were measured in a quality-adjusted way, output and hence labor productivity could rise significantly.

Emergency medicine

Prediction technology has the potential to make medicine more efficient and effective through the personalization of treatment. In particular, machine learning can identify those patients for which a given treatment will be most effective.

There is a widespread sense that much of healthcare spending is wasteful. For example, in the United States in particular, many prominent policymakers, economists, and medical researchers have argued that doctors test too much. [6]

One area where over-testing appears widespread is testing for heart attack among patients who arrive at the emergency department of hospitals. The decision of who to test is predicated on a prediction. In this context, all patients that need treatment require a test. Tests performed on patients who subsequently are assessed as not needing treatment are considered a waste.

Most arguments for over-testing emphasize averages: the average return to testing is less than the average cost. Mullainathan and Obermeyer (2018) emphasize that prediction tools enable doctors to use the theoretically relevant object: marginal benefits and marginal costs. Using machine learning, Mullainathan and Obermeyer demonstrate that a large number of very high-risk patients go untested, while many low risk patients are tested. In other words, there is not only a problem of over-testing, but there is also a problem of under-testing. Using better prediction models when patients arrive at the emergency department could substantially increase health outcomes for the same spending (or substantially reduce spending for equal health outcomes).

Automation of the prediction of who to test does not change the workflow. The process of testing and the associated treatment remain as is. The machine prediction saves little time for the doctors, nurses, and staff working in the emergency department, who would otherwise make the testing decision quickly, but it improves efficiency and outcomes, both reducing the number of unnecessary tests and increasing necessary ones, thus saving lives. In other words, automating the prediction task improves the productivity of emergency medicine in the context of heart attacks without much substitution for human work. Furthermore, it increases demand for labor in complementary tasks such as surgery.

---

[6] This discussion draws heavily on Mullainathan and Obermeyer (2018).

Drug discovery

Machine prediction is also improving the productivity of scientists and researchers. For instance, Atomwise uses AI to enhance the drug discovery process. Traditionally, identifying molecules that could most efficiently bind with proteins for a given therapeutic target was largely based on educated guesses and, given the number of potential combinations, it was highly inefficient. That meant that downstream experiments to test whether a molecule could be of use in a treatment were compromised by the noise of poor-quality candidate molecules.

Atomwise automated the task of predicting which molecules have most potential for exploration. Their software classifies foundational building blocks of organic chemistry and predicts the outcomes of real-world physical experiments. Automating this process, enabling lower cost and higher accuracy, increases the returns to the downstream lab testing procedure that is conducted by people. As a consequence, the demand for labor to conduct such testing is likely to increase. Furthermore, higher yield due to better prediction of which chemicals might work increases the number of humans needed in the downstream tasks of bringing these chemicals to market.

This example brings attention to a broader point. AI is already having an impact in innovation. As Cockburn, Henderson, and Stern (2018) emphasize, machine learning tools are already widespread in the sciences. In some cases, like Atomwise, this leads to automated prediction, and the increased use of already-existing complementary tasks. In other cases, it is possible that the increased innovation will lead to the invention of new tasks, a subject we return to below. For example in the context of Atomwise, new tasks may arise if it means that drug discovery is more efficient and drugs can be better-targeted to narrower populations.

This drug discovery example also highlights another point. Automation of a prediction task can effect decision tasks that are significantly downstream from the prediction task, performed by different people in downstream occupations. In other words, better prediction can affect a number of different decisions.

Language translation

Another example of machine prediction affecting a wide variety of downstream decisions is the effect of machine language translation on trade. Machine learning has been used to automate translation by framing the task as a prediction problem. The machine predicts how a human translator would translate a string of characters from one language into another. In one of the first attempts to estimate the economic impact of a commercial deployment of AI, Brynjolfsson, Hui, and Liu (2018) measure the effect of an improvement in the quality of translation by an AI on the volume of trade conducted on the online platform eBay. The authors find that moving to AI-based translation resulted in a 17.5% increase in the volume of trade. So, an AI-driven improvement in the prediction task results in a significant increase in downstream trade activity, much of which we can assume is performed by human labor. Of course, increased trade has many forms of impact on overall economic activity and so we cannot draw any conclusions on the overall impact of this implementation of AI on labor in equilibrium.

In this section, we highlighted how automation of the prediction task can enhance labor in the decision task. The prediction tasks in these examples do not require much human labor, and so the net effect could be labor-enhancing with better and sometimes more decision tasks done by labor.

6 The Case of Radiology: Are AIs Labor-Reducing?

In this section, we describe the ambiguous and nuanced effects on labor in a context that is often provided as an obvious example of AI leading to automation: Radiology. At an AI conference in 2016, deep learning pioneer Geoffrey Hinton publicly asserted "We should stop training radiologists now," comparing the profession to Wile E. Coyote from the Road Runner cartoon who has run off the cliff but hasn't looked down yet.[7]

His remark was motivated by the progress of AI tools that are increasingly applied to identify abnormalities in medical images. IBM and GE have already commercialized AI tools in radiology, identifying breast, lung, and other cancers from medical images. Smaller companies and startups have also commercial similar products. Zebra Medical Vision received FDA approval to predict whether coronary heart disease is present in a CT scan. Zebra also develops tools to predict the presence of various medical issues, including bone, liver, and lung disease.[8]

[7] https://www.youtube.com/watch?v=2HMPRXstSvQ Accessed August 10, 2018.
[8] www.zebra-med.com Accessed August 10, 2018.

A common practice is to embed AI-based image recognition into the software the radiologists use to read scans.[9] The software highlights areas predicted to be abnormalities. Radiologists examine the highlighted image when interpreting and reporting on the results. Some evidence suggests that AI is best-used this way, to augment the diagnosis decisions of humans rather than replace them altogether (Wang et al 2016). In these cases, the work flow of radiologists remains unchanged. The human remains in the loop for each scan, but the readings become faster and more accurate. If the number of scans stays fixed, then the demand for radiologists declines. On the other hand, if readings are faster and more accurate, then the number of scans could increase enough to counteract the increased number of scans read per radiologist. In this scenario, a human is in the loop for every read. The AI delivers a prediction that is used as an input into a decision task that is performed by a human. Thus, under current practice, AI is used for prediction but the decision task remains with the human. In this way, if we focus on current practice, the discussion of AI in radiology belongs in the previous section where AI automates the prediction task, but not the decision.

Recent research suggests, however, that machine prediction can meet or even surpass human diagnostic accuracy in detecting some types of disease.[10] If so, then full automation of the task of interpreting imaging results becomes possible. While the current level of technology suggests a human should remain in the loop, it is likely that over time AI will lead to full automation of the image interpretation task. This will lead to a substantial change in the radiology workflow and potentially impact on the labor component significantly. First, if the "interpret imaging results" task is done by machine, then that step is no longer part of the radiologist's work flow. To the extent that this task takes up a significant fraction of the overall time, then automating this task will reduce the demand for radiologists.

Nevertheless, many tasks in the workflow of diagnostic radiologists would remain: choosing the exam, directing the technologists, reporting on the results, and deciding on an action given the probabilities reported by the machine. Many radiologists serve as the "doctor's doctor", communicating the meaning of images to other patient-facing doctors.[11] The interpretation of scans is often probabilistic and radiologists have expertise in interpreting probabilities to help the patient-

[9]Interview with Toronto-based radiologist on February 20, 2018.

[10]https://www.ncbi.nlm.nih.gov/pubmed/28275919.

[11] https://pubs.rsna.org/doi/full/10.1148/radiol.2512090177

facing doctor recommend a course of action. Thus, reporting on the results may require a human intermediary between the machine prediction and the doctor who requested a test. For example, a human is needed to determine payoffs in order to recommend a course of action. What is the cost of conducting a biopsy if no disease is present? What is the cost of failing to conduct a biopsy if disease is present? In other words, what is the probability of disease threshold over which a biopsy (or some other further action) should be conducted? How does that vary based on patient characteristics, whether fully codifiable (such as age and medical history) or not (such as the doctor's sense of the patient's personality and preferences). In fact, as the prediction task becomes better, faster, and cheaper, the demand for these related, complementary tasks may increase. In other words, it is plausible that automating the image prediction task, while reducing the demand for labor to perform that task, may increase the overall demand for labor due to an increased demand for labor to perform the complementary tasks.

In Table 1, we list the 29 different tasks that comprise the radiologists' workflow according to the occupational classification database O*NET. Only two of these tasks are directly affected by an image recognition AI: Task #3: "Perform or interpret the outcomes of diagnostic imaging procedures including magnetic resonance imaging (MRI), computer tomography (CT), positron emission tomography (PET), nuclear cardiology treadmill studies, mammography, or ultrasound." and Task #25 "Interpret images using computer-aided detection or diagnosis systems." On the one hand, this seems like a small fraction of the overall workflow - two twenty-ninths. On the other hand, this raises the question of whether interpreting the outcome of diagnostic imaging is the main task—the defining area of expertise—of radiologists. Perhaps another type of professional could do the other tasks, like obtaining patients' histories, preparing interpretive reports, transmitting images, and communicating results to families?

Twenty-nine Tasks Associated with the Occupation Radiologists (source: O*NET)
1. Obtain patients' histories from electronic records, patient interviews, dictated reports, or by communicating with referring clinicians.
2. Prepare comprehensive interpretive reports of findings.
3. Perform or interpret the outcomes of diagnostic imaging procedures including magnetic resonance imaging (MRI), computer tomography (CT), positron emission tomography (PET), nuclear cardiology treadmill studies, mammography, or ultrasound.

4. Review or transmit images and information using picture archiving or communications systems.

5. Communicate examination results or diagnostic information to referring physicians, patients, or families.

6. Evaluate medical information to determine patients' risk factors, such as allergies to contrast agents, or to make decisions regarding the appropriateness of procedures.

7. Provide counseling to radiologic patients to explain the processes, risks, benefits, or alternative treatments.

8. Instruct radiologic staff in desired techniques, positions, or projections.

9. Confer with medical professionals regarding image-based diagnoses.

10. Coordinate radiological services with other medical activities.

11. Document the performance, interpretation, or outcomes of all procedures performed.

12. Establish or enforce standards for protection of patients or personnel.

13. Develop or monitor procedures to ensure adequate quality control of images.

14. Recognize or treat complications during and after procedures, including blood pressure problems, pain, oversedation, or bleeding.

15. Administer radiopaque substances by injection, orally, or as enemas to render internal structures and organs visible on x-ray films or fluoroscopic screens.

16. Participate in continuing education activities to maintain and develop expertise.

17. Participate in quality improvement activities including discussions of areas where risk of error is high.

18. Supervise and teach residents or medical students.

19. Implement protocols in areas such as drugs, resuscitation, emergencies, power failures, or infection control.

20. Schedule examinations and assign radiologic personnel.

21. Provide advice on types or quantities of radiology equipment needed to maintain facilities.

22. Participate in research projects involving radiology.

23. Perform interventional procedures such as image-guided biopsy, percutaneous transluminal angioplasty, transhepatic biliary drainage, or nephrostomy catheter placement.

24. Administer or maintain conscious sedation during and after procedures.

25. Interpret images using computer-aided detection or diagnosis systems.

26.	Serve as an offsite teleradiologist for facilities that do not have on-site radiologists.
27.	Develop treatment plans for radiology patients.
28.	Treat malignant internal or external growths by exposure to radiation from radiographs (x-rays), high energy sources, or natural or synthetic radioisotopes.
29.	Conduct physical examinations to inform decisions about appropriate procedures.

Table 1: Tasks in a radiologist's workflow, according to O*NET[12]

Overall, the 29 tasks reveal that even if image interpretation becomes fully automated, plenty of tasks for humans remain. The key open question is whether those tasks are best conducted by a radiologist. Instead, at least some of these tasks might be better performed by medical practitioners with different expertise. Judgment on the best course of action for a patient might be best done by a primary care physician or perhaps a social worker. The supervision of radiology technologists might be better done by more experienced radiology technologists.

AI will also affect radiology in a variety of other ways, separate from predicting abnormalities in scans. For example, radiologists often dictate their reports. Past practice was that the recorded reports were then sent to a (human) transcription service. This is Task #4: "Review or transmit images and information using picture archiving or communications systems." Many radiology departments already use AI-based transcription services to automate the transcription task. The human can be completely replaced because the task of predicting the text spoken by the radiologist is a stand-alone task. The words in the recording simply need to be converted to text. While this may reduce costs and reduce wait times for radiologists and patients, the direct effect is the elimination of transcription-related jobs.[13]

Thus, in radiology, the overall effects on jobs are hard to identify. There is automation of the prediction related to reading scans. Currently, this does not change the workflow of radiologists, though it might if the AI gets better. Radiologists perform many other non-prediction tasks, and so AI is unlikely to automate these tasks; however, it is not clear that radiologists will be the humans who perform these tasks if reading scans becomes automated. Finally, humans

[12] https://www.onetonline.org/link/summary/29-1069.10.
[13]Interview with Toronto-based radiologist on February 20, 2018; https://www.jacr.org/article/S1546-1440%2817%2931671-X/fulltext.

working in radiology who are not radiologists and do not work on scans—such as those providing transcription services—may have their jobs automated completely.

7 New tasks through new decisions

AI may allow for new decisions to be made where previously it was impossible or too costly to do so. As Herbert Simon emphasized (Simon 1972), when rationality is bounded – in terms of being able to adequately distinguish between important outcomes in a complex environment, for example – rather than make optimal, state-contingent choices, economic agents will instead resort to rules. Those rules may be followed by individuals, be part of operational procedures in companies or be embedded in machines that take a single action but otherwise do not react to environmental conditions.

State-contingent choices can be consequential for companies. For example, Google deployed AI developed by its DeepMind unit to optimize the use of air conditioners in its data centers. The end result was a 40 percent reduction in energy used in a highly energy intensive operation.[14] State-contingent choices enabled new automated decisions on energy usage. In other words, new tasks were generated by better prediction. In this application, the new tasks were performed by machines.

New tasks may also be performed by humans. As mentioned above, AI is already having an impact is in scientific research. Uncertainty is pervasive in many aspects of research, and so prediction technology is likely to have a large impact on the production of science. Cockburn, Henderson, and Stern (2018) show that machine learning is used by scientists in a wide variety of fields. Figure 1 shows the number of publications in computer science and applications journals by AI field. The blue, dark red, and dark green lines show publications in three different AI subfields of computer science: machine learning, robotics, and symbolic logic. The results show a slow and steady increase in publications in all three, with the largest increase in machine learning. The most striking result in the figure, however, is the yellow line. It shows the increase in publications outside of computer science that mention machine learning. In other words, it demonstrates that, since 2012, the biggest change in AI publications did not occur in computer

[14] https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/

science. It occurred in other fields of science that use machine learning. The same is not true for robotics and symbolic logic (light blue and pink lines).

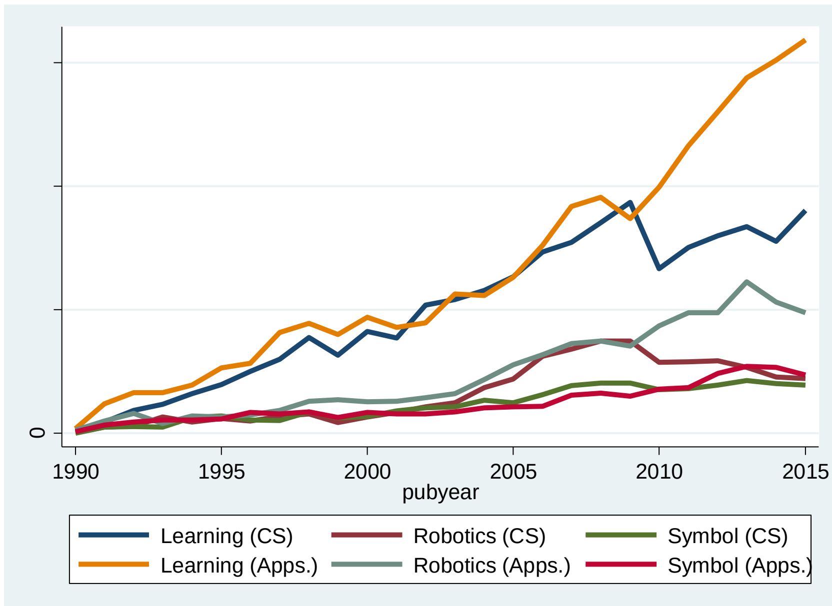

Figure 1: Publications in Computer Science versus Application Journals, by AI Field

Source: Cockburn, Henderson, and Stern (2018). Figure 4

This result is apparent in our own field of economics. We are increasingly using machine learning to improve our statistical models and advance our knowledge of economics (Athey 2018). The same forces are at play in many other scientific fields, such as chemistry, public health, and neuroscience.

In this way, the recent advances in machine learning can be seen as an invention in the method of inventing, as highlighted in Griliches (1957) for hybrid corn. Cockburn, Henderson, and Stern (2018) describe this insight: “The challenge presented by advances in AI is that they appear to be research tools that not only have the potential to change the method of innovation itself, but also have implications across a wide range of fields.” Agrawal, McHale, and Oettl (2018) similarly describe AI as a General Purpose Technology and explain how that may influence the knowledge production function. In other words, in addition to increasing demand for existing tasks (as in the Atomwise example above), AI is likely to create innovations that lead to new industries—and new types of jobs with new tasks in those industries.

Of course, at this early stage in the technology’s diffusion, the challenge in writing about the new tasks that AI will bring is that these tasks do not yet exist. Examples are scarce. In determining what new tasks and industries may arise, the discussion is necessarily more speculative. Nevertheless, we provide two examples of academics and companies applying AI in new areas that suggest the potential for new industries, new jobs, and new tasks.

University of Toronto professor Alan Aspuru-Guzik's research group is developing a "self-driving chemistry lab" that enables the discovery of chemicals and materials at a fraction of the price of a current lab. Using advances in robotics and machine learning, the lab could be deployed in thousands of locations around the world, without the need for a local workforce with deep expertise in chemistry. This would enable industries in rural areas and developing countries to have access to a wide variety of materials. The inventors emphasize that this tool could "provide the scientific community with an easy-to-use package to facilitate novel discovery at a faster pace" (Roch et al. 2018, p. 1) and "democratize autonomous discovery" (p. 12). One can imagine many new tasks associated with the arrival of autonomous on-site discovery.

The space industry is an example of how machine learning, as an input into the scientific process, could generate a new industry at commercial scale. For example, the latency of radio signals that arises when deploying machines in outer space or in highly volatile settings such as a rocket launch means that some processes are more efficient when automation is possible (no need to wait for instructions from mission control). As another example, the company Seer Tracking has built an AI to predict the trajectory of space debris. This creates a variety of new opportunities in space. First, it could create a set of (human and machine) tasks focused on moving space assets out of the way of the debris. Second, it could enable more commercial opportunities in space by reducing the risk that a space asset will be destroyed by debris. In other words, the uncertainty associated with space debris may mean that some decision tasks are never undertaken. Resolving this uncertainty could enable new commercial opportunities in space. This is, of course, still speculative. We highlight this as an informed example of a new industry that could arise, where startup founders and investors are already allocating scarce resources.

In this section, we provide examples of the type of new tasks that may arise as AI diffuses. AI reduces uncertainty. Uncertainty can render certain activities economically infeasible. Reduced uncertainty enables state-contingent choices and this can enable new opportunities and new tasks, either for capital or labor. DeepMind's energy efficiency predictions enables new tasks that are performed by machines. Cockburn, Henderson, and Stern point out that by reducing uncertainty in innovation, recent advances in AI can be seen as an invention of a method of invention thereby increasing the potential for accelerated innovation and new job opportunities – implying new tasks to be performed by labor.

Nevertheless, the discussion in this section is particularly speculative. While we observe real examples of the reduction of labor due to automating existing prediction tasks and we also see examples of increased demand for labor due to enhanced demand for certain existing tasks that are complements to prediction, we are still early in the commercial application of AI to see compelling examples of new complementary tasks for labor.

8 Conclusion

The task-based model presented by Acemoglu and Restrepo provides a useful framework for understanding the impact of AI on jobs. From a practical perspective, however, both the theoretical literature and the commonly used data (e.g., O*NET) define tasks in a seemingly arbitrary way. Our contribution to this line of inquiry is to provide specificity to the task-based model by drawing from our previous work that defines AI as a prediction technology and describes the role of prediction in decision-making. Specifically, we provide structure to the classification of tasks and in particular highlight the role of prediction tasks and decision tasks, where decision tasks are perfect complements to prediction tasks in the sense that prediction has no value without a decision. This structure describes how AI directly substitutes capital for labor in the case of prediction tasks and may indirectly effect decision tasks by increasing or decreasing the relative returns to labor versus capital for decision tasks. In addition, machine prediction may directly impact decision tasks by enabling new decisions due to reduced uncertainty.

We concur with Acemoglu and Restrepo in terms of their emphasis on the main effects from AI being displacement of labor in some *existing* tasks and reinstatement of labor in *new* tasks; however, we were surprised at how hard we found it to find examples of *new* tasks for labor. We suspect that new tasks will begin to emerge and that at the time of this writing the technology is simply too early and the diffusion too limited. Nevertheless, at this time the most likely source of gains to labor seem to come from increased demand for *existing* tasks that benefit from better, faster, and cheaper prediction. Some examples we cited include more productive brain surgery (due to better cancer cell predictions), more productive driving (due to better navigation predictions), more productive laboratory testing (due to better molecule binding efficacy predictions), and more productive stent implants (due to better heart attack predictions).

While the direct effect of deploying AI involves automation of prediction, we argue that AI does not fit easily into existing analyses of the effect of automation on labor markets. The reasons are threefold. First, prediction is always strictly complementary to other tasks – namely decision-related tasks. Those tasks can be existing or newly possible because of the better prediction. Second, better prediction improves decisions—whether taken by labor or capital—by enabling more nuanced decisions through the reduction of uncertainty. Finally, it is not yet possible to say whether this impact on decision tasks – whether existing or new – is likely to favor labor or capital. We have found important examples of both and there is no obvious reason for a particular bias to emerge. Thus, we caution on drawing broad inferences from the research on factory automation (e.g. Acemoglu and Restrepo 2017; Autor and Salomons 2018) in forecasting the near-term consequences of AI for labor markets.

References

Acemoglu, Daron, and David Autor. 2011. Skills, Tasks and Technologies: Implications for Employment and Earnings. Chapter 12. In *Handbook of Labor Economics. Eds.* David Card and Orley Ashenfelter, Elsevier, Volume 4, Part B, 1043-1171,

Acemoglu, Daron, and Pascual Restrepo. 2017. Robots and Jobs: Evidence from US Labor Markets. NBER Working Paper 23285.

Acemoglu, Daron, and Pascual Restrepo. 2018a. Artificial Intelligence, Automation and Work. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Acemoglu, Daron, and Pascual Restrepo. 2018b. Modeling Automation. *AEA: Papers and Proceedings*, May.

Acemoglu, Daron, and Pascual Restrepo. 2018c. The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, 108 (6): 1488-1542.

Acemoglu, Daron, and Pascual Restrepo. 2018d. Automation and New Tasks: The Implications of the Task Content of Technology for Labor Demand. *Journal of Economic Perspectives* (forthcoming).

Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2018a. *Prediction Machines: The Simple Economics of Artificial Intelligence*, Harvard Business Review Press.

Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2018b. Prediction Judgment and Complexity: A Theory of Decision Making and Artificial Intelligence. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2019. Dumb machines, smart machines, and humans: The impact of artificial intelligence on task allocation. Working paper, University of Toronto.

Agrawal, Ajay, John McHale, and Alexander Oettl. 2018. Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. 2015 The Skill Complementarity of Broadband Internet. *Quarterly Journal of Economics* 130(4), 1781–1824.

Alpaydin, Ethem. 2010. *Introduction to Machine Learning*. Second Edition. MIT Press, Cambridge MA.

Athey, Susan. 2018. The impact of machine learning on economics. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. Computing Inequality: Have Computers Changed the Labor Market? *Quarterly Journal of Economics*, 113(4), 1169–1213.

Autor, David, and Ana Salomons. 2018. Is automation labor-displacing? Productivity growth, employment, and the labor share. *Brookings Papers on Economic Activity*.

Bresnahan, T. and M. Trajtenberg. 1995. General Purpose Technologies 'Engines of Growth'? Journal of Econometrics. 65. 83-108.

Brynjolfsson, Erik, Xiang Hui, and Meng Liu. 2018. Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform. NBER Working Paper 24917.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. Machine Learning and Occupational Change. Working paper, MIT.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2018. Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Brynjolfsson, Erik and Lorin Hitt. 2003. Computing Productivity: Firm-Level Evidence. Review of Economics and Statistics. 85(4): 793-808.

Cockburn, Iain, Rebecca Henderson, and Scott Stern. 2018. The Impact of Artificial Intelligence on Innovation. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

David, Paul. 1991. Computer and Dynamo: The Modern Productivity Paradox in a Not-Too-Distant Mirror. In Technology and Productivity: The Challenge for Economic Policy, Paris. OECD Publishing. 315-47.

Frey, Carl Benedikt and Osborne, Michael A., (2017), The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, issue C, p. 254-280.

Fry, Hannah. 2018. Hello World: Being Human in the Age of Algorithms, Norton: New York.

Griliches, Z. 1957. Hybrid Corn: An Exploration in the Economics of Technological Change. Econometrica. 25(4). 501-522.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles. 2018. Should Robots Be Taxed? Working paper, Northwestern University.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data mining, Inference, and Prediciton*. Second Edition. Springer, New York.

Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2018. Human Decisions and Machine Predictions, *Quarterly Journal of Economics*, 133(1), 237–293

Mullainathan and Obermeyer. 2018. Using Machine Learning to Understand Human Decision-Making: Application to Health Care. Working paper, University of Chicago.

Roch, Loic, Florian Hase, Christoph Kreisbeck, Teresa Tomayo-Mendoza, Lars Yunker, Jason Hein, and Alan Aspuru-Guzik. 2018. ChemOS: An orchestration software to democratize autonomous discover. ChemRxiv. https://chemrxiv.org/articles/ChemOS_An_Orchestration_Software_to_Democratize_Autonomous_Discovery/5953606.

Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica, 55*(5), 999-1033.

Simon, Herbert. 1972. Theories of Bounded Rationality. In McGuire and Radner Eds. *Decision and Organization*. Chapter 8. North-Holland, Amsterdam.

Stevenson, Betsy. 2018. AI, Income, Employment, and Meaning. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Trajtenberg, Manuel. 2018. AI as the Next GPT: A Political Economy Perspective. In Agrawal, Gans, and Goldfarb Eds. *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Wang, D., A. Khosla, R. Gargeya, H. Irshad, and A. H. Beck. 2016. Deep Learning for Identifying Metastatic Breast Cancer. arXiv preprint arXiv:1606.05718.