The impact of small improvements in prediction accuracy can be deceptive. The mundane process of filling in missing information can make prediction machines seem magical. In the 1990s, it was difficult to build large enough data sets. But by 2004, computer processing and storage had improved. With additional variables and customers, machine learning methods started to perform as well as humans. Machine learning models are good at determining which of many possible variables will work best and recognizing that some things don't matter and others do. It enables predictions based on unanticipated correlations. "Deep learning," relies on an approach called "back propagation." It's how natural brains do it, by learning through example.

Many problems have transformed from algorithmic problems ("what are the features of a cat?") to prediction problems ("does this image with a missing label have the same features as the cats I have seen before?"). The output of machine learning—prediction—is key to intelligence. Prediction is also the basis for human intelligence. Prediction accuracy improves by learning, and high accuracy enables machines to perform tasks once associated with human intelligence. It is the primary function of the neocortex, and the foundation of intelligence. The cortex is an organ of prediction.

Current AI algorithms cannot reason, and it is difficult to interrogate them to understand the source of their predictions. Just as cheap arithmetic enabled by computers proved powerful in ushering in dramatic changes to business and personal lives, similar transformations will occur due to cheap prediction. Advances of 21st-century computer science match previous advances in social and physical sciences from recognition that algorithms work best when structured probabilistically, based on data. Statistics emphasized being correct on average, whereas in machine learning the goal was operational effectiveness. This gave scientists freedom to experiment and drove rapid improvements that take advantage of the rich data and fast computers that appeared over the last decade. Machine learning has less need to specify in advance what goes into a model and can accommodate much more complex models with many more interactions between variables.

Data plays 3 roles. First, input data is fed to the algorithm to produce a prediction. Second, *training data* generates the algorithm. Then, *training data* trains the AI to get good enough to predict in the wild. Finally *feedback data* improves the algorithm's performance with experience. Most people use Google for both rare and common searches. Being even a little better in search can lead to a big difference in market share and revenue. Data might be most valuable if you have more and better data than your competitor.

It is critical to balance the cost of data acquisition with the benefit of enhanced prediction accuracy. Statistical and economic reasons shape whether having more data generates more value. From a statistical perspective, data has diminishing returns. When firms used an objective and verifiable test along with normal interviews, there was a 15% bump in job tenure of hires relative to hiring decisions based on interviews alone. There are *known knowns:* things we know we know. We also know there are *known unknowns*--we know there are some things we do not know. But there are also *unknown unknowns*—the things we don't know we don't know. It is the latter category that tend to be difficult. Known knowns are when we have rich data, known unknowns are when there is too little data, unknown unknowns are those events that are not captured by past experience or the data but are nonetheless possible--so prediction is difficult, although we may not realize it. Unknown knowns are when an association that appears to be strong in the past is the result of some unknown or unobserved factor that changes over time and makes predictions unreliable.

Humans are sometimes extremely good at prediction with little data. We can recognize a face after seeing it only once or twice. From a very young age, we can guess the trajectory of a ball. We are also good at analogy, taking new situations and identifying other circumstances similar enough to be useful in a new environment. "One-shot learning" prediction machines are not yet adequate. Perhaps the biggest weakness of prediction machines is that they sometimes provide wrong answers that they are confident are right. If the machine does not understand the decision process that generated the data, its predictions can fail.

"Counterfactual": what would have happened if you took a different action. You will never have data on the action not taken. Google's search results come from a secret algorithm largely determined by prediction machines that predict which links someone is likely to click. But prediction machines can be gamed. Google uses human judgment to re-optimize the machine in the face of such spam. Sometimes, the combination of humans and machines generates the best predictions, each complementing the other's weaknesses---a classic division of labor, both in mechanical and mental processes.

The human and the machine are good at different aspects of prediction. A human pathologist was usually right when saying there was cancer. The AI was much more accurate when saying the cancer wasn't there. Combining human and machine prediction overcame their weaknesses to dramatically reduced the error rate. Machine prediction can enhance the productivity of human prediction via 2 broad pathways: an initial prediction that humans can use to combine with their own assessments. It then provides a second opinion after the fact, or a path for monitoring. One excellent place to examine interactions is prediction of the creditworthiness of loan applicants. A credit scoring system score helped, though the improvement was largest with the score provided to the humans in advance of their assessment.

Prediction machines can scale in a way that humans cannot. Machines struggle to make predictions in unusual cases for which there isn't much historical data. Many human-machine collaborations take the form of "prediction by exception." The human is, in many respects, the prediction machine's supervisor, engaging the human's attention only when needed. Identifying and redacting confidential information from legal documents is a laborious procedure arising in many legal situations. The user of a "Chisel redactor" may set the redactor to be aggressive or light. Each redaction is a recommendation rather than a final decision.

Humans, including experts, make poor predictions under certain conditions. As prediction machines improve, businesses must adjust the division of labor between humans and machines in response.

Decisions are at the core of most occupations. Schoolteachers. Managers. Truck drivers. Police officers. Doctors. Parents confront decisions often under conditions of uncertainty. Decisions require applying judgment to a prediction and then acting. Prediction machines are valuable because (1) they can often produce better, faster, and cheaper predictions; (2) prediction is a key ingredient in decision making under uncertainty; and (3) decision making is common throughout our economic and social lives. Judgment determines the relative payoff associated with each possible outcome of a decision,

Better predictions raise the value of judgment. More precise predictions on fraud reduce the likelihood of mislabeling legitimate transactions as fraudulent. "Reward function engineering" determines the rewards to various actions. Most of us already do some of that but for humans, not machines. Parents teach their children values. Mentors teach new workers how the system operates. Managers give objectives to their staff and then tweak them to get better performance. Prediction machines increase the returns to judgment because, by lowering the cost of prediction, they increase the value of understanding the rewards associated with actions. If there are a manageable number of action-situation combinations associated with a decision, we can transfer judgment to a prediction machine

We can automate decisions by rewarding the prediction machine for predicting what a human would do. Humans have 3 types of data that machines don't. 1) Human senses are powerful. 2) Humans are the ultimate arbiters of our own preferences. 3) Privacy concerns restrict the data available to machines. Machines may learn to predict human judgment within data limits, but they are bad at predicting rare events.

You can train prediction machine with examples. By enabling more complex decisions, better predictions can lower risk. Enhanced prediction enables decision makers to handle more "ifs" and "thens." That leads to better outcomes. In the absence of good prediction, we do a lot of "satisficing," making decisions that are "good enough" given the information available. It will take practice to imagine the vast array of transformations possible as a result of prediction machines that can handle more "ifs" and "thens" and more complex decisions in more complex environments.

US carmakers agreed with the Department of Transportation to make automatic emergency braking standard on vehicles by 2022. Automation arises when a machine undertakes an entire task, not just prediction. We must determine the returns to machines performing other elements (data collection, judgment, actions) to decide whether a task should be fully automated. Mining giant Rio Tinto determined the real cost to mining companies is not people but downtime. In 2016 it implemented 73 self-driving trucks that could operate autonomously, saving 15% in operating costs. The mine runs its trucks 24 hours a day, without drivers. The trucks do not need a front and back, meaning they do not need to turn around, No human driver needs to watch over the truck's safety on-site or even remotely. And there are fewer humans around to create safety risks.

When speed is needed, the benefit of ceding control to a prediction machine is high. Automation occurs when the return to machines handling all functions is greater than the returns to including humans in the process. It is much easier to send robots into space than humans. The need for human intervention limits the autonomy of prediction machines even when they might operate on their own. What distinguishes the "within factory" environment for a robot from the "open road" is the possibility of what economists call "externalities"—costs felt by others, rather than the key decision makers. Playing with children, caring for the elderly, and many other actions that involve social interaction may also be inherently better from a human. The introduction of AI to a task does not imply full automation of that task. Prediction is only one component. Tasks most likely to be automated first are ones which deliver the highest returns.

AI is a general-purpose technology. It has the potential to affect every decision, because prediction is key to decision making. AI requires rethinking processes. AI tools are point solutions. Each generates a specific prediction, and most are designed to perform a specific task. Companies will break their work flows down into tasks, estimate the ROI for building or buying an AI to perform each task, rank-order the AIs in terms of ROI, and then start from the top.

One thing that separates us from the high primates is that we're tool builders. AI tools predict intention of speech (Amazon's Echo), command context, what you want to buy (Amazon's recommendations), which links will connect you to the information you want to find (Google search), when to apply the brakes to avoid danger (Tesla's Autopilot). Each delivers a predictive component to make a decision. Judgment takes 2 forms in Radiology: targeting the disease and understanding potential side effects. Given data on protein characteristics, the startup Atomwise is involved in discovering promising pharmaceutical drug prospects. It can predict which molecules have the highest binding affinity and whether molecules that have never been produced are likely to have high binding affinity. The way to decompose the selection task: ACTION: What are you trying to do? For Atomwise, it is to test molecules to help cure or prevent disease. PREDICTION: What do you need to know to make the decision? Atomwise predicts binding affinities of potential molecules and proteins. JUDGMENT: How do you value different outcomes and errors? Atomwise and its customers set the criterion regarding the relative importance of targeting the disease and the relative costs of potential side effects. OUTCOME: What are your metrics for task success? For Atomwise, it's the results of the test. Ultimately, did the test lead to a new drug? This AI tool supports a prediction task in its customers' drug discovery workflow. Its value is in reducing the cost and increasing the likelihood of success for drug development.

Who will be a best or high-value MBA student? It is strategic in that it is global rather than local and is looking for impact rather than, say, maximizing student income or creating wealth. The value of accepting a particular student is a matter of judgment. Tasks must be decomposed to see where prediction machines can be inserted. The AI canvas is an aid to help with the decomposition process. Prediction is at the center of the AI canvas.

Computers made arithmetic cheap. One of the first killer apps was bookkeeping. This cenario—a job augmented as a machine takes over some, but not all, tasks—is likely to become common as a consequence of the implementation of AI tools. Many radiologists see themselves as the "doctor's doctor." A key part of their job is to communicate the meaning of images to primary care doctors. The challenging part is that interpretation of radiology images ("studies," in their language) is often probabilistic. Radiologists help them interpret the numbers so that the primary care doctors can work with patients to decide the best course of action. The role of the prediction machine is to increase a doctor's confidence in *not* conducting a biopsy.

Automating a particular task can emphasize other tasks that are important to a job that were previously underappreciated. A job is a collection of tasks. When breaking down a workflow and employing AI tools, some tasks previously performed by humans may be automated, and new tasks may be created.

Prediction and judgment are complements; as the use of prediction increases, the value of judgment rises. Who will capture the value that better prediction creates? Data has different grades. We highlight 3 types—training, input, and feedback data. Training data is used to build a prediction machine. Input data is used to power it to produce predictions. Feedback data is used to improve it. Only the 2 latter types are needed for future use. Spreadsheet value goes to users, the businesses that deployed them to make billions of better decisions. Prediction machines reduce uncertainty. As AI advances, we'll use prediction machines to reduce uncertainty more broadly. As the cost of AI falls, prediction machines will resolve a wider variety of strategic dilemmas. Prediction machines will increase the value of judgment, actions, and data.

Makers of luxury automobiles that manufactured their own parts improved faster from each model year to the next. They measured improvements at the customer end and could adapt more readily to customer feedback. But outsourced parts were better initially because the parts suppliers made better parts. More "ifs" means that a business can write contracts to specify what to do if something unusual happens. Suppose that AI allows airlines not only to forecast weather events but to generate predictions for how best to deal with weather-related interruptions. Prediction may give airlines and automakers the confidence to allow for better, more complex arrangements and products.

The introduction of ATMs produced a significant organization transformation. When performance measures change from objective (are you keeping the bank queues short?) to subjective (are you selling the right products?), human resource (HR) management becomes more complex. AI will shift HR management toward relational and away from transactional. AI will have an impact on labor different from its impact on capital, increasing the importance of human judgment. Prediction and judgment are complements, so better prediction increases demand for judgment. Data and prediction machines are complements. A key strategic choice is determining where your business ends and another business begins—the boundary of the firm.

AI requires learning, and it requires startups to be more willing to invest in this learning than their more established rivals. *Supervised learning*: when you already have good data. After the fact is called *reinforcement learning*. Apple plans to do AI in a way that respects privacy. It is making a big strategic bet that consumers will want control over their own data.

Automation could result in the deskilling of humans. A junior pilot may have almost 3000 hours in the air, but it was not quality experience. Most of the time, he has been flying the plane on autopilot. Experience is a scarce resource, some of which you need to allocate to humans to avoid deskilling. Shifting to an AI-first strategy means downgrading the previous top priority. AI tools are, initially, trained in house, away from customers. However, they learn faster when deployed into commercial use in real operating conditions and greater volumes of data.

Google's algorithms might amplify prejudices that already exist in society, a risk of AI. Ad-placement might result in what lawyers call "disparate impact." A person or an organization can be liable for discrimination, even if it is accidental. To figure out if AI is discriminating, you must look at the output. Humans discriminate more than machines. Such unknown knowns are a key weakness of prediction machines that require human judgment to overcome. Homogenous populations are at greater risk of disease and destruction. Such human monoculture can be individually beneficial but increase system-wide risk. Just as in biodiversity, the diversity of prediction machines involves a trade-off between individual and system level outcomes.

If a device is not connected to the cloud, a simultaneous attack becomes difficult. One risk is that when you have an AI (like Google's search engine), then a competitor can observe data being entered (query) and output being reported (such as a list of websites), and it has the raw materials to employ its own AI to engage in supervised learning and reconstruct the algorithm. Imitation can be easy for your competitors or detractors may deliberately try to train your prediction machine to make bad predictions. Also, AIs are ineffective when data is sparse.

Wisdom is breadth, as is broad framing. For the past few decades, labor's share of national income has been falling in favor of the share earned by capital. If AI is a new, efficient form of capital, then the capital share of the economy will likely continue to rise at the expense of labor. It disproportionately increases the wages of highly educated people and may decrease the wages of the less educated. If the skills needed to succeed with AI change more often, then the educated will benefit disproportionately. Technology-based monopolies are temporary due to a process that economist Joseph Schumpeter called "the gale of creative destruction."

Prediction is the process of filling in missing information. The current wave of advances in artificial intelligence doesn't bring us prediction, but organizations can exploit prediction machines by adopting AI tools to assist with executing their current strategy.

[Prediction is the basis for human intelligence. Prediction accuracy enables machines to perform tasks once associated with human intelligence. Humans discriminate more than machines.]