

Prediction Machines

Study Guide

Introduction and Chapters 1 - 19
(full study guide)

Meet the Authors



AJAY AGRAWAL

@professor_ajay

Ajay has a Ph.D. from the University of British Columbia and is the Munk Professor of Entrepreneurship at the Rotman School of Management. He is the founder of the Creative Destruction Lab and Next Canada.



JOSHUA GANS

@joshgans

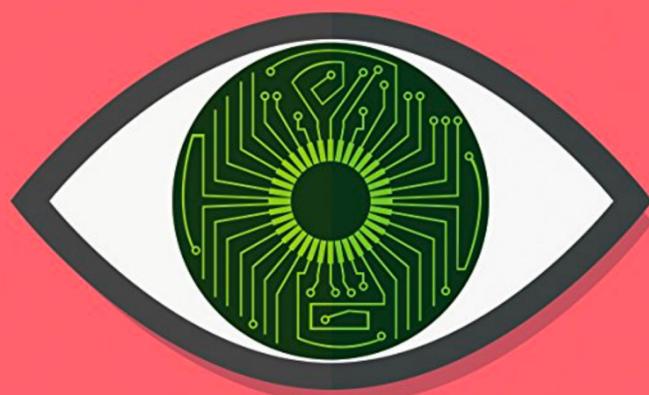
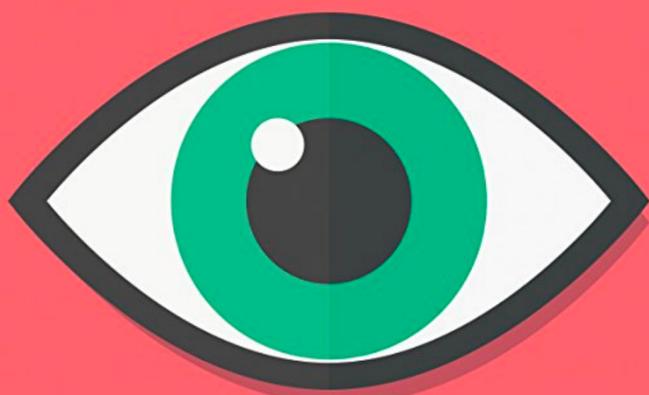
Joshua has a Ph.D. from Stanford and is Professor of Strategic Management and holds the Jeffrey S. Skoll Chair in Technical Innovation and Entrepreneurship at the Rotman School of Management.



AVI GOLDFARB

@avicgoldfarb

Avi has a PhD from Northwestern and is the Ellison Professor of Marketing at the Rotman School of Management. He is the author of numerous papers investigating digitisation in economics.



Summaries are taken from the end of each chapter. The authors were kind enough to provide summaries.

Chapter 1 - Introduction Machine Intelligence

Key Themes

- The current wave of advances in artificial intelligence doesn't actually bring us intelligence but instead a critical component of intelligence: prediction
- Prediction is a central input into decision making. Economics has a well-developed framework for understanding decision-making. The new and poorly understood implications of advances in prediction technology can be combined with the old and well-understood logic of decision theory from economics to deliver a series of insights to help navigate your organization's approach to AI.
- There is often no single right answer to the question of which is the best AI strategy or the best set of AI tools because AIs involve trade-offs: more speed, less accuracy; more autonomy; less control; more data; less privacy. We provide you with a method for identifying the trade-offs associated with each AI-related decision so that you can evaluate both sides of every trade in light of your organization's mission and objectives and then make the decision that is best for you.

Chapter 2 - Cheap Changes Everything

Key Themes

- Economics offers clear insights regarding the business implications of cheaper prediction. Prediction machines will be used for traditional prediction tasks (inventory and demand forecasting) and new problems (like navigation and translation). The drop in the cost of prediction will impact the value of other things, increasing the value of complements (data, judgement, and action) and diminishing the value of substitutes (human prediction).
- Organizations can exploit prediction machines by adopting AI tools to assist with executing their current strategy. When those tools become powerful, they may motivate changing the strategy itself. For instance, if Amazon can predict what shoppers want, then they may move from a shop-then-ship model to a ship-then-shop model- bringing goods to homes before they are ordered. Such a shift will transform the organization.
- As a result of the new strategies that organizations pursue to take advantage of AI, we will be faced with a new set of trade-offs related to how AI will impact society. Our choices will depend on our needs and preferences and will almost surely be different across different countries and cultures. We structured this book in five sections to reflect each layer of impact from AI, building from the foundations of prediction all the way up to the trade-offs for society: (1) Prediction, (2) Decision making, (3) tools, (4) Strategy and (5) Society.

Chapter 3 - Prediction Machine Magic

Key Themes

- Prediction is the process of filling in the missing information. Prediction takes the information you have, often called “data,” and uses it to generate information you don’t have. In addition to generating information about the future, prediction can generate information about the present and the past. This happens when prediction classifies credit card transactions as fraudulent, a tumour in an image as malignant, or whether a person holding an iPhone is the owner.
- The impact of small improvements in prediction accuracy can be deceptive. For example, an improvement from 85 percent to 90 percent accuracy seems more than twice as large as from 98 percent to 99.9 percent (an increase of 5 percentage points compared to 2). However, the former improvement means that mistakes fall by a factor of twenty. In some settings, mistakes falling by a factor of twenty are transformational.
- The seemingly mundane process of filling in missing information can make prediction machines seem magical. This has already happened as machines see (object recognition), navigate (driverless cars), and translate.

Chapter 4- Why It's Called Intelligence

Key Themes

- Machine learning science has different goals from statistics. Whereas statistics emphasized being correct on average, machine learning did not require that. Instead, the goal was operational effectiveness. Predictions could have biases so long as they were better (something that was possible with powerful computers). This gave scientists the freedom to experiment and drove rapid improvements that take advantage of the rich data and fast computers that appeared over the last decade.
- Traditional statistical methods require the articulation of hypotheses or at least of human intuition for model specification. Machine learning has less need to specify in advance what goes into the model and can accommodate the equivalent of much more complex models with many more interactions between variables.
- Recent advances in machine learning are often referred to as advances in artificial intelligence because: (1) systems predicated on this technique learn and improve over time; (2) these systems produce significantly more accurate predictions than other approaches under certain conditions, and some experts argue that prediction is central to intelligence; and (3) the enhanced prediction accuracy of these systems enable them to perform tasks, such as translation and navigation, that were previously considered the exclusive domain of human intelligence. We remain agnostic on the link between prediction and intelligence. None of our conclusions rely on taking a position on whether advances in prediction represent advances in intelligence. We focus on the consequences of a drop in the cost of prediction, not a drop in the cost of intelligence.

Chapter 5- Data Is the New Oil

Key Themes

- Prediction machines utilize three types of data: (1) training data for training the AI, (2) input data for predicting, and (3) feedback data for improving prediction accuracy.
- Data collection is costly; it's an investment. The cost of data collection depends on how much data you need and how intrusive the collection process is. It is critical to balance the cost of data acquisition with the benefit of enhanced prediction accuracy. Determining the best approach requires estimating the ROI of each type of data: how much will it cost to acquire, and how valuable will the associated increase in prediction accuracy be?
- Statistical and economic reasons shape whether having more data generates more value. From a statistical perspective, data has diminishing returns. Each additional unit of data improves your prediction less than the prior data; the tenth observation improves prediction by more than the one thousandth. In terms of economics, the relationship is ambiguous. Adding more data to a large existing stock of data may be greater than adding it to a small stock - for example, if the additional data allows the performance of the prediction machine to cross a threshold from unusable to usable or from below a regulatory performance threshold to above, or from worse than a competitor to better. Thus, organizations need to understand the relationship between adding more data, enhancing prediction accuracy, and increasing value creation.

Chapter 6 - The New Division of Labour

Key Themes

- Humans, including professional experts, make poor predictions under certain conditions. Humans often overweigh salient information and do not account for statistical properties. Many scientific studies document these shortcomings across a wide variety of professions. The phenomenon was illustrated in the feature film Moneyball.
- Machines and humans have distinct strengths and weaknesses in the context of prediction. As prediction machines improve, businesses must adjust their division of labour between humans and machines in response. Prediction machines are better than humans at factoring in complex interactions among different indicators, especially in settings with rich data. As the number of dimensions for such interactions grows, the ability of humans to form accurate predictions diminishes, especially relative to machines. However, humans are often better than machines when understanding the data generation process confers a prediction advantage, especially in settings with thin data. We describe a taxonomy of prediction settings (i.e., known knowns, known unknowns, unknown knowns, and unknown unknowns) that is useful for anticipating the appropriate division of labour.
- Prediction machines scale. The unit cost per prediction falls as the frequency increases. Human prediction does not scale the same way. However, humans have cognitive models of how the world works, and this can make predictions based on small amounts of data. Thus, we anticipate a rise in human prediction by exception whereby machines generate most predictions because they are predicated on routine, regular data, but when rare events occur, the machine recognizes that it is not able to produce a prediction with confidence, and so calls for human assistance. The human provides prediction by exception.

Chapter 7 - Unpacking Decisions

Key Themes

- Prediction machines are so valuable because (1) they can often produce better, faster, and cheaper predictions than humans can; (2) prediction is a key ingredient in decision-making under uncertainty; and (3) decision-making is ubiquitous throughout our economic and social lives. However, a prediction is not a decision - it is only a component of a decision. The other components are judgement, action, outcome and three types of data (input, training, and feedback).
- By breaking down a decision into its components, we can understand the impact of prediction machines on the value of humans and other assets. The value of substitutes to prediction machines, namely human prediction, will decline. However, the value of complements, such as the human skills associated with data collection, judgement, and actions, will become more valuable. In the case of the London cabbies who each invested three years to learn “The Knowledge”- how to predict the fastest route from one location to another at a particular time of day - none became worse at their job because of prediction machines. Rather, many other drivers became a lot better at choosing the best route by using prediction machines. The cabbies’ prediction skills were no longer a scarce commodity. Drivers who weren’t cabbies had driving skills and human sensors (eyes and ears) that were effectively enhanced by prediction machines, enabling them to compete.
- Judgement involves determining the relative payoff associated with each possible outcome of a decision, including those associated with “correct” decisions as well as those associated with mistakes. Judgement requires specifying the objective you’re actually pursuing and is a necessary step in decision-making. As prediction machines make predictions increasingly better, faster, and cheaper, the value of human judgement will increase because we’ll need more of it. We may be more willing to exert effort and apply judgement to decisions where we previously has chosen not to decide (by accepting the default).

Chapter 8 - The Value of Judgement

Key Themes

- Prediction machines increase the returns to judgement because, by lowering the cost of prediction, they increase the value of understanding the rewards associated with actions. However, judgement is costly. Figuring out the relative payoffs for different actions in different situations takes time, effort, and experimentation.
- Many decisions occur under conditions of uncertainty. We decide to bring an umbrella because we think it might rain, but we could be wrong. We decide to authorize a transaction because we think it is legitimate, but we could be wrong. Under conditions of uncertainty, we need to determine the payoff for acting on wrong decisions, not just right ones. So, uncertainty increases the cost of judging the payoffs for a given decision.
- If there are a manageable number of action-situation combinations associated with a decision, then we can transfer the judgement from ourselves to the prediction machine (this is “reward function engineering”) so that the machines can make the decision itself once it generates the prediction. This enables automating the decision. Often, however, there are too many action-situation combinations, such that it is too costly to code up all the payoffs associated with each combination, especially the very rare ones. In these cases, it is more efficient for a human to apply judgement after the prediction machine predicts.

Chapter 9 -Predicting Judgement

Key Themes

- Machines may learn to predict human judgement. An example is driving. It is impractical for humans to code their judgement about how to handle every possible situation. However, we train autonomous driving systems by showing them many examples and rewarding them for predicting human judgement: What would a human do in this situation?
- There are limits to the ability of machines to predict human judgement. The limits relate to lack of data. There is some data that humans have that machines do not, such as individual preferences. Such data has value, and companies currently pay to access it through discounts on using loyalty cards and free online services like Google and Facebook.
- Machines are bad at prediction for rare events. Managers make decisions on mergers, innovation, and partnerships without data on similar past events for their firms. Humans use analogies and models to make decisions in such unusual situations. Machines cannot predict judgement when a situation has not occurred many times in the past.

Chapter 10- Taming Complexity

Key Themes

- Enhanced prediction enables decision makers, whether human or machine, to handle more “ifs” and more “thens”. That leads to better outcomes. For example, in the case of navigation, illustrated in this chapter with the mail robot, prediction machines liberate autonomous vehicles from their previous limitation of operating only in controlled environments. These settings are characterized by their limited number of “ifs” (or states). Prediction machines allow autonomous vehicles to operate in uncontrolled environments, like on a city street, because rather than having to code all the potential “ifs” in advance, the machine can instead learn to predict what a human controller would do in any particular situation. Similarly, the example of airport lounges illustrates how enhanced prediction facilitates more “thens” (eg. “then leave at time X or Y or Z,” depending on the prediction of how long it will take to get to the airport at a particular time on a particular day), rather than always leaving early “just in case” and then spending extra time waiting in the airport lounge.
- In the absence of good prediction, we do a lot of “satisficing”, making decisions that are “good enough” given the information available. Always leaving early for the airport and often waiting once you arrive because you’re early is an example of satisficing. That solution is not optimal, but it’s good enough given the information available. The mail robot and the airport lounge are both inventions designed in response to satisficing. Prediction machines will reduce the need to satisfice and thus reduce the returns to investing in solutions like mail robot systems and airport lounges.

We are so used to satisficing in our businesses and in our social lives that 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, this, more complex decisions in more complex environments.

Chapter 10- Taming Complexity

Key Themes

It's not intuitive for most people to think of airport lounges as a solution to poor prediction and that they will be less valuable in an era of powerful prediction machines. Another example is the use of biopsies, which largely exist in response to weaknesses in prediction from medical images. As the confidence in prediction machines go up, the impact from medical imaging AIs may be much greater on the jobs associated with conducting biopsies because, like airport lounges, this costly and invasive procedure was invented in response to poor prediction. Airport lounges and biopsies are both risk management solutions. Prediction machines will provide new and better methods for managing risk.

Chapter 11- Fully Automated Decision Making

Key Themes

- The introduction of AI to a task does not necessarily imply full automation of that task. Prediction is only one component. In many cases, humans are still required to apply judgement and take an action. However, sometimes judgement can be hard-coded or if enough examples are available, machines can learn to predict judgement. In addition, machines may perform the action. When machines perform all elements of the task, then the task is fully automated, and humans are completely removed from the loop.
- The tasks most likely to be fully automated first are the ones for which full automation delivers the highest returns. These include tasks where: (1) the other elements are already automated except for prediction (e.g., mining); (2) the returns to speed of action in response to prediction are high (e.g., driverless cars); (3) the returns to reduced waiting time for predictions are high (e.g., space exploration).
- An important distinction between autonomous vehicles operating on a city street versus those in a mine site is that the former generates significant externalities while the latter does not. Autonomous vehicles operating on a city street may cause an accident that incurs costs borne by individuals external to the decision maker. In contrast, accidents caused by autonomous vehicles operating on a mine site only incur costs affecting assets or people associated with the mine. Governments regulate activities that generate externalities. Thus, regulation is a potential barrier to full automation for applications that generate significant externalities. We anticipate a significant wave of policy development concerning the assignment of liability driven by an increasing demand for many new areas of automation.

Chapter 12 - Deconstructing Work Flows

Key Themes

- AI tools are point solutions. Each generates a specific prediction, and most are designed to perform a specific task. Many AI startups are predicated on building a single AI tool.
- Large corporations are comprised of workflows that turn inputs into outputs. Work Flows are made up of tasks (e.g., a Goldman Sachs IPO is a workflow comprised of 146 distinct tasks). In deciding how to implement AI, companies will break their workflows 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 of the list and begin working downward. Sometimes a company can simply from an AI tool into their workflow and realize an immediate benefit due to increasing the productivity of that task. Often, however, it's not that easy. Deriving a real benefit from implementing an AI tool requires rethinking or “reengineering” the entire workflow. As a result, similar to the personal computer revolution, it will take time to see productivity gains from AI in many mainstream businesses.
- To illustrate the potential effect of an AI on a workflow, we describe a fictitious AI that predicts the ranking of any MBA application. To derive the full benefit from this prediction machine, the school would have to redesign its workflow. It would need to eliminate the task of manually ranking applications and expand the task of marketing the program, as the AI would increase the returns to a greater applicant pool (better predictions about who will succeed and lower cost of evaluating applications). The school would modify the task of offering incentives like scholarships and financial aid due to increased certainty about who will succeed. Finally, the school would adjust other elements of the workflow to take advantage of being able to provide instantaneous school admission decisions.

Chapter 13- Decomposing Decisions

Key Themes

- Tasks need to be decomposed in order to see where prediction machines can be inserted. This allows you to estimate the benefit of the enhanced prediction and the cost of generating that prediction. Once you have generated reasonable estimates, rank-order the AIs from highest to lowest ROI by starting at the top and working your way down, implementing AI tools as long as the expected ROI makes sense.
- The AI canvas is an aid to help with the decomposition process. Fill out the AI canvas for every decision or task. This introduces discipline and structure into the process. It forces you to be clear about all three data types required: training, input and feedback. It also forces you to articulate precisely what you need to predict, the judgement required to assess the relative value of different actions and outcomes, the action possibilities, and the outcome possibilities.
- At the center of the AI canvas is prediction. You need to identify the core prediction at the heart of the task, and this can require AI insight. The effort to answer this question often initiates an existential discussion among the leadership team: “What is our real objective, anyhow?” Prediction requires a specificity not often found in mission statements. For a business school, for example, it is easy to say that they are focused on recruiting the “best” students, but in order to specify the prediction, we need to specify what “best” means- highest salary offer upon graduation? Most likely to assume a CEO role within five years? Most diverse? Most likely to donate back to the school after graduation? Even seemingly straightforward objectives, like profit maximization, are not as simple as they first appear. Should we predict the action to take that will maximize profit this week, this quarter, this year, or this decade? Companies often find themselves having to go back to basics to realign on their objectives and sharpen their mission statement as a first step in their work on their AI strategy.

Chapter 13- Decomposing Decisions

Key Themes

- Tasks need to be decomposed in order to see where prediction machines can be inserted. This allows you to estimate the benefit of the enhanced prediction and the cost of generating that prediction. Once you have generated reasonable estimates, rank-order the AIs from highest to lowest ROI by starting at the top and working your way down, implementing AI tools as long as the expected ROI makes sense.
- The AI canvas is an aid to help with the decomposition process. Fill out the AI canvas for every decision or task. This introduces discipline and structure into the process. It forces you to be clear about all three data types required: training, input and feedback. It also forces you to articulate precisely what you need to predict, the judgement required to assess the relative value of different actions and outcomes, the action possibilities, and the outcome possibilities.
- At the center of the AI canvas is prediction. You need to identify the core prediction at the heart of the task, and this can require AI insight. The effort to answer this question often initiates an existential discussion among the leadership team: “What is our real objective, anyhow?” Prediction requires a specificity not often found in mission statements. For a business school, for example, it is easy to say that they are focused on recruiting the “best” students, but in order to specify the prediction, we need to specify what “best” means- highest salary offer upon graduation? Most likely to assume a CEO role within five years? Most diverse? Most likely to donate back to the school after graduation? Even seemingly straightforward objectives, like profit maximization, are not as simple as they first appear. Should we predict the action to take that will maximize profit this week, this quarter, this year, or this decade? Companies often find themselves having to go back to basics to realign on their objectives and sharpen their mission statement as a first step in their work on their AI strategy.

Chapter 14 - Job Redesign

Key Themes

- A job is a collection of tasks. When breaking down a work flow and employing AI tools, some tasks previously performed by humans may be automated, the ordering and emphasis of remaining tasks may change, and new tasks may be created. Thus, the collection of tasks that make up a job can change.
- The implementation of AI tools generates four implications for jobs:
 - AI tools may augment jobs, as in the example of spreadsheets and bookkeepers.
 - AI tools may contract jobs, as in fulfilment centres
 - AI tools may lead to the reconstitution of jobs, with some tasks added and others taken away, as with radiologists.
 - AI tools may shift the emphasis on the specific skills required for a particular job, as with school bus drivers.
-
- AI tools may shift the relative returns to certain skills and, thus, change the types of people who are best suited to particular jobs. In the case of bookkeepers, the arrival of the spreadsheet diminished the returns to being able to perform many calculations quickly on a calculator. At the same time, it increased the returns to being good at asking the right questions in order to fully take advantage of the technology's ability to efficiently run scenario analyses.

Chapter 15 - AI in the C-Suite

Key Themes

- C-Suite leadership must not fully delegate AI strategy to their IT department because powerful AI tools may go beyond enhancing the productivity of tasks performed in the service of executing against the organization's strategy and instead lead to changing the strategy itself. AI can lead to strategic change if three factors are present: (1) there is a core trade-off in the business model (e.g., shop-then-ship versus ship-then-shop); (2) the trade-off is influenced by uncertainty (e.g, higher sales from ship-then-shop are outweighed by higher costs from returned items due to uncertainty about what customers will buy); (3) an AI tool that reduces uncertainty tips the scales of the trade-off so that the optimal strategy changes from one side of the trade to the other (e.g., an AI that reduces uncertainty by predicting what a customer will buy tips the scale such that the returns from a ship-then-shop model outweigh those from the traditional model).
- Another reason C-suite leadership is required for AI strategy is that the implementation for AI strategy is that the implementation of AI tools in one part of the business may also affect other parts. In the Amazon thought experiment, a side effect of transitioning to a ship-then-shop model was vertical integration into the returned items collection business, perhaps with a fleet of trucks that did weekly pick-ups throughout the neighbourhood. In other words, powerful AI tools may result in significant redesign of work flows and the boundary of the firm.
- Prediction machines will increase the value of complements, including judgement, actions, and data. The increasing value of judgement may lead to changes in organizational hierarchy- there may be higher returns to putting different roles or different people in positions of power. In addition, prediction machines enable managers to move beyond optimizing individual components to optimizing higher-level goals, and this makes decisions closer to the objectives of the organization.

Chapter 15 - AI in the C-Suite

Key Themes

Owning the actions affected by prediction can be a source of competitive advantage that allows traditional businesses to capture some of the value from AI. However, in some cases, where powerful AI tools provide a significant competitive advantage, new entrants may vertically integrate into owning the action and leverage their AI as a basis for competition.

Chapter 16 - When AI Transforms your Business

Key Themes

- A key strategic choice is determining where your business ends and another business begins - deciding on the boundary of the firm (e.g., airline partnerships, outsourcing automotive part manufacturing). Uncertainty influences this choice. Because prediction machines reduce uncertainty, they can influence the boundary between your organization and others.
- By reducing uncertainty, prediction machines increase the ability to write contracts, and thus increase the ability to write contracts, and this increase the incentive for companies to contract out both capital equipment and labour that focuses on data, prediction, and action. However, prediction machines decrease the incentive for companies to contract out labour that focuses on judgement. Judgement quality is hard to specify in a contract and difficult to monitor. If judgement could be well specified, then it could be programmed, and we wouldn't need humans to provide it. Since judgement is likely to be the key role for human labour as AI diffuses, in-house employment will rise and contracting out labour will fall.
- AI will increase incentives to own data. Still, contracting out for data may be necessary when the predictions that the data provides are not strategically essential to your organization. In such cases, it may be best to purchase predictions directly rather than purchase data and then generate your own predictions.

Chapter 17 - Your Learning Strategy

Key Themes

- Shifting to an AI-first strategy means downgrading the previous top priority. In other words, AI-first is not a buzzword- it represents a real tradeoff. An AI-first strategy places maximizing prediction accuracy as the central goal of the organization, even if that means compromising on other goals such as maximizing revenue, user numbers, or user experience.
- AI can lead to disruption because incumbent firms often have weaker economic incentives than startups to adopt the technology. AI-enabled products are often inferior at first because it takes time to train a prediction machine to perform as well as a hard-coded device that follows human instructions rather than learning on its own. However, once deployed, an AI can continue to learn and improve, leaving its unintelligent competitors' products behind. It is tempting for established companies to take a wait-and-see approach, standing on the sidelines and observing the progress in AI applied to their industry. That may work for some companies, but others will find it difficult to catch up once their competitors get ahead in the training and deployment of their AI tools.
- Another strategic decision concerns timing- when to release AI tools into the wild. AI tools are initially, trained in-house, away from customers. However, they learn faster when they are deployed to commercial use because they are exposed to real operating conditions and often to greater volumes of data. The benefit to deploying earlier is faster learning, and the cost is greater risk (risk to the brand or customer safety by exposing customers to immature AIs that are not properly trained). In some cases, the tradeoff is clear, such as with Google Inbox, where the benefits of faster learning outweigh the cost of poor performance. In other cases, such as autonomous driving, the trade-off is more ambiguous given the size of the prize for being early with a commercial product weighed against the high cost of an error if the product is released before it is ready.

Chapter 18 - Managing AI Risk

Key Themes

- AI carries many types of risk. We summarize six of the most salient types here:
 - Predictions from AIs can lead to discrimination. Even if such discrimination is inadvertent, it creates liability.
 - AIs are ineffective when data is sparse. This creates quality risk, particularly of the “unknown known” type, in which a prediction is provided with confidence but is false.
 - Incorrect input data can fool prediction machines, leaving their users vulnerable to attack by hackers.
 - Just as in biodiversity, the diversity of prediction machines involves a trade-off between individual - and system-level outcomes. Less diversity may benefit individual-level performance but increase the risk of massive failure.
 - Prediction machines can be interrogated, exposing you to intellectual property theft and to attackers who can identify weaknesses.
 - Feedback can be manipulated so that prediction machines learn destructive behaviour.

Chapter 19 - Beyond Business

Key Themes

- The rise of AI presents society with many choices. Each represents a tradeoff. At this stage, while the technology is still in its infancy, there are three particularly salient trade-offs at the societal level.
- The first trade-off is productivity versus distribution. Many have suggested that AI will make us poorer or worse off. That's not true. Economists agree that technological advance makes us better off and enhances productivity. AI will unambiguously enhance productivity. The problem isn't wealth creation; it's distribution. AI might exacerbate the income inequality problem for two reasons. First, by taking over certain tasks, AIs might increase competition among humans for the remaining tasks, lowering wages and further reducing the fraction of income earned by labour versus the fraction earned by the owners of capital. Second, prediction machines, like other computer-related technologies, may be skill-biased such that AI tools disproportionately enhance the productivity of highly skilled workers.
- The second trade-off is innovation versus competition. Like most software-related technologies, AI has scale economies. Furthermore, AI tools are often characterized by some degree of increasing returns: better prediction accuracy leads to more users, more users generate more data, and more data leads to better prediction accuracy. Businesses have greater incentives to build prediction machines if they have more control, but, along with scale economies, this may lead to monopolization. Faster innovation may benefit society from a short-term perspective but may not be optimal from a social or longer-term perspective.
- The third trade-off is performance versus privacy. AIs perform better with more data. In particular, they are better able to personalize their predictions if they have access to more personal data. The provision of personal data will often come at the expense of reduced privacy.

Chapter 19 - Beyond Business

Key Themes

Some jurisdictions, like Europe, have chosen to create an environment that provides their citizens with more privacy. That may benefit their citizens and may even create conditions for a more dynamic market for private information where individuals can more easily decide whether they wish to trade, sell or donate their private data. On the other hand, that may create friction in settings where opting in is costly and disadvantages European firms and citizens in markets where AIs with better access to data are more competitive.

For all three trade-offs, jurisdictions will have to weigh both sides of the trade and design policies that are most aligned with their overall strategy and the preferences of their citizenry.