6

# Artificial Intelligence Technologies and Aggregate Growth Prospects

Timothy Bresnahan

## 6.1 Introduction

This chapter examines the commercial application of Artificial Intelligence Technologies (AITs), seeking to address questions about these technologies specifically and about twenty-first-century technical progress and its current and potential impact on economic growth. I focus on the highly valuable *applications* of AITs today, in production systems at the Internet Giants, in new user interfaces (UIs), and elsewhere. My empirical conclusion about these applications is that the lazy idea of AI – that is, of computer systems that are able to perform productive tasks previously done by humans – is irrelevant to understanding how these technologies create value. Here "irrelevant" does not mean that substitution of machine for human tasks is less important than other determinants of the value in use of AITs. It means irrelevant: task-level substitution of machine for human plays no role in these highly valuable systems.

The absence of task-level substitution is unsurprising to scholars of production based on information and communications technology (ICT), and it does not mean that there has been no factor substitution at all. The transition to ICT-based production has largely proceeded at the production system level, not the task level. Consider examples from the largest area of AIT use so far: consumer–product matching and targeting. The production systems by which Google and Facebook present targeted advertisements to individual consumers differ from those of the older advertising business. To be sure, some of the differences are in factor utilization – the new advertising industry production systems run on ICT capital. But other differences are equally important, such as in the distinction between targeted and mass-media advertising. System-level



https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

substitution, generally, is driven as much or more by output characteristics such as ad targeting as by cost minimization.

That leads to analysis of the characteristics of the AIT-using systems and the structure of incentives and opportunities to invent new AIT-based production processes. In their economically important initial applications, and in the early stages of diffusion, AIT-using systems are largely capital deepening in already capital-intensive production processes and services. As with several other recent important new ICTs, the largest applications are in mass-market marketing and distribution, focusing on consumers. Media markets, advertising markets, and the marketing functions of consumer products and services companies appear likely to get the deepest investments in AITs. In the sections that follow, I examine the complementarity of AITs with specific aspects of existing capital-intensive production processes that explain the tendency toward capital deepening. AIT capital–other capital complementarities and scale economies at the firm level are an important element. I also examine the aspects of new capital-intensive production processes, with and without AITs, that have limited their range of application to mass-market environments with low-stakes transactions.

Many observers hope that AITs will become General Purpose Technologies (GPTs).[1] That appears to be half right in the early going, but it leads us away from AITs' visible role in growth. The half that is clearly right is positive feedback loops running through improvements in AITs and their applications.[2] Positive feedback loops are associated with social scale economies and thus, potentially, with growth (Romer 1986). But thus far there is little indication that the diffusion of AIT-based systems will contribute most of its value through *broadening* the range of applications of new capital to a range of industries and functions. While we can anticipate widespread use of AITs, thus far the economically important applications lie in capital *deepening* in a narrow range of industries and functions. In that important regard, AITs are like the other big twenty-first

[1] Many recent economic papers on AI pose the question of whether all of it will become a GPT. See Brynjolfsson, Rock, and Syverson (2017), Taddy (2018), and Cockburn, Henderson, and Stern (2018). Most importantly, see Trajtenberg (2018), as well as Agrawal, Gans, and Goldfarb (2018), which provides an excellent overview of these issues.

[2] GPT analytics emphasize the innovational complementarities between the GPT itself and inventions of applications (Bresnahan and Trajtenberg, 1995; Rosenberg and Trajtenberg, 2004; Helpman and Trajtenberg, 1998). Bresnahan and Greenstein (1996) emphasize the role of difficult-to-invent applications in slowing the diffusion of, and easy-to-invent applications accelerating the diffusion of, ICT GPTs. Rosenberg (1997) writes about the role of post-invention uncertainty (often about the most important applications).

-century waves of ICT, such as Web 2.0 and mobile. Earlier, ICT spread out over more and more economic activity for many decades, from a few functions in large firms, to many functions, to system access by individual workers, to extensive consumer applications. Recently, that nature of the positive feedback loop driving ICT invention and ICT-application invention has moved from broadening to deepening. In our era, rapid ICT technical progress leads to some universal benefits but, importantly, also leads to ICT-capital deepening in particular firms, industries, and functions.

For many years, there has been a research area – General Artificial Intelligence – with the imprecise goal of designing computer systems that can do tasks previously requiring human intelligence.[3] Taking this research goal as a metaphor, looking at laboratory phenomena and demonstration projects and adding technological determinism is the basis of most writing about AITs. This metaphor underlies the focus on task-level substitution in the literature. General AI research goals have not been met. Instead, the statistical turn in AI research of a generation ago dramatically accelerated progress in a number of separate but related laboratory technologies. These technologies are based on the idea of statistical prediction – if in very different domains – ranging from "seeing" pictures to forecasting what book a consumer might read next. The actual AITs that exist – like all software technologies – are designed into systems. In this chapter, I look at those systems. Spoiler alert: Do not hope for a lot of sci-fi in them.

It is not too early to look at the systems that embody AITs, and the analysis is not based on speculation. AITs are now central elements of working commercial systems generating revenues in the hundreds of billions of dollars. AITs are broadly used in UI subsystems. Both the earliest production uses and the UIs have begun to diffuse away from those first applications, enough to at least examine the early diffusion path. There is another category of AIT applications, smaller at this stage, where AIT laboratory phenomena are very close to already-algorithmic production steps. Finally, I will examine other growing AIT applications

[3] The Oxford Living Dictionary defines Artificial Intelligence as "[t]he theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages." Technologists recognize the imprecision of this definition, which stems both from "normally" and from "such as." What is a "task normally requiring human intelligence"? Computers have been doing tasks previously done using human intelligence for over seventy years, starting with arithmetic.

such as driver assist, the rebirth of expert systems, and improved decision support.

I will have two research goals in mind. First, what has been the business and economic logic of the AIT applications? Once we learn that AIT application does not emphasize the task-level substitution of machine thinking, seeing, etc. for human labor, we come to the question of what the important forces determining AIT value creation might be. Second, does the commercialization of AIT, at this early stage, exhibit continuity or discontinuity with prior rounds of technological commercialization of ICT? Either continuity or discontinuity (along a particular dimension) will offer valuable clues to the future direction of technical progress in the application of AIT and of ICT more generally.

That leads to a discussion of where ICT application has been going in the twenty-first century, and to an attempt to understand it. An easy-to-understand thing is that there has been a remarkable series of waves – beginning with the web browser – of new technologies that serve consumers directly, enable mass markets, and support the creation of mass-marketing commercial applications. Another easy-to-understand thing is that many of those waves, including Web 2.0, mobile, cloud, and now AITs, have led to substantial capital deepening in those same areas (consumption, mass marketing, etc.). A harder-to-understand thing is why the very impressive technical progress in those areas has had limited – that is, some, but only limited – impact on applications in the rest of the economy, where ICTs went first. The conclusion, in Section 6.5, will focus on this growth question.

In forecasting the long-run factor demand implications of AITs, there is as yet no evidence that they are different from other ICTs – working at the system level to slowly change to a more capital-intensive, less labor-intensive, more human-capital-intensive form of production, but not a form of production in which the main new feature is different factor use.

## 6.2 Product–Consumer Matching Applications

I begin with AIT-based product–consumer matching engines at Amazon, Google, Facebook, Netflix, and other consumer-oriented Internet Giants. Based on machine learning using these firms' considerable "big data" assets, these applications have created substantial economic value for their inventors. These are not demonstration projects or experiments. They are production systems generating revenues in the hundreds of billions of dollars.

These systems are impressive business and engineering accomplishments, involving not only use of new AITs but also the invention of new and better ways to match products to consumers. I look at them together because they have a common role in their use of AIT. These engines match a specific potential buyer to a specific potential seller. At Amazon, this matching yields product recommendations for a particular customer; at Google, it ranks advertisements that a specific searcher might see.

This use of AITs is no small thing. Improvements in targeting buyer–seller matches amount to a marketing revolution in the twenty-first century.[4] The private returns to the invention driving that revolution – returns captured thus far mostly by the Internet Giants – have been enormous, significantly increasing capital's share of output. In short, the use of AITs in product recommendation engines has been a central part of one of the most valuable technical advances of the twenty-first century – and is thus a good place to start understanding how AITs create value.

### 6.2.1 Amazon

Amazon, both in its own store and now in its online mall, recommends products to consumers. Excellence at recommendation has been a goal of the company from the outset. One of the firm's earliest employees attributed to founder Jeff Bezos the idea that the firm's web page should display for each consumer one book, the one they are going to buy next (Brandt, 2011). Product recommendation at Amazon has been algorithmic for many years; today, the firm uses recommendation algorithms built with AIT and machine learning. Amazon also has other AIT applications and products, to which I shall return later in this section.

Amazon's recommendation system responds to users' input with lists of potential products to examine or to buy. The data used for these systems have grown over time. Amazon has long known a great deal about what products individual customers have searched for or bought and where the customer is in the search process. The span of that information has increased as Amazon's range of products has grown from books, to many products, to hosting a mall. More recently, a number of Amazon services and products, such as Amazon Prime, Kindle, and Alexa, have increased the amount of data associated with individual customers. Amazon also has much product information, including both sellers'

[4] See Goldfarb (2014) on the targeting revolution in marketing.

descriptions of category, etc. and data on which consumers did or did not consider or choose the product.

For many years, algorithms have made recommendations that, for example, suggest additional products based on what a consumer has chosen to search for and what they have bought in the past. More recently, machine learning and AITs drive the algorithms that Amazon uses in these areas.[5] Amazon has sufficient data to do a good job of predicting what a particular consumer will look at, buy, etc., based on machine-learning-based product–consumer matching systems. One outcome is improved recommendations from a consumer perspective (who find items they want) and from an Amazon and merchant perspective (who get matched to customers interested in their products), consistent with the strategic goals for the company.

Beyond the large volume of data (not all of it high quality – this is "big data") and those strategic goals, there are a number of features of the Amazon mall that make it particularly suitable for an application of AIT. Amazon had a preexisting recommendation system that was part of its well-functioning and *modularized* online store/mall. Making that recommendation system more targeted to the individual customer by using the AIT of machine learning would not require changing other elements extensively because of the modularization. This benefit of a modular production process has been known to economists since Simon (1962), and a deep literature has turned it into practical management doctrine (e.g., Baldwin and Clark, 2000).

The Amazon production process was already highly modularized in part because it was already algorithmic and software developers see many benefits of modular systems. Among those benefits is that modularization at Amazon and the other Internet Giants permits them to gain *scalability* in the face of growing and changing loads.[6] Having a scalable production

[5] The language I use draws heavily on Bezos (2017), who also notes that related AITs based on machine learning do other matching functions, such as product and deal recommendations, merchandising placements, etc. In short, a lot of what Amazon.com shows a consumer and what an Amazon app or a Kindle show a consumer is managed by AITs.

[6] Modularity helps the scalability of a system with multiple moving parts in a number of ways. Modularity permits the addition of processing power and storage as needed when loads change. Similarly, modularity permits the addition of new data streams (e.g., adding information from Gmail) or analytical elements (e.g., thwarting ill-behaved search engine optimization at Google) as systems change. Dynamically, modularity permits improvement of the system architecture such that it can do new things while not undercutting the scalability of the existing components. A business discussion of this topic can be found in Baldwin and Clark (2000), notably in the subchapter of chapter 3 on "dimensionless, scalable design rules." Hennessey and Patterson (2012, p. 260) make a similar observation

system has, in turn, let the firm gain scale economies. To understand this, we need to examine how the economic concept of "economies of scale" interacts with the computer science concept of "scalability."

"Scale economies" mean that marginal cost is lower than average cost. "Scalability" means that a system has been designed so that its workload can be increased without changes to its architecture (again following Simon, 1962). Scalability is not just scale economies – when workloads are uncertain, for example, a scalable production process has the flexibility to be changed quickly.

Amazon is a large-scale firm in the sense that it and its online mall tenants engage in a large number of transactions with a large number of consumers. The firm's online store and now its mall have an architecture that permits scalability through modularization. A number of complex systems, most centrally the product recommendation systems, are elements of this modularized online store/mall architecture. Changing to AIT recommendation systems preserved modularity and did not affect scalability. Amazon already had good estimates of the costs and benefits of algorithmic recommendations, and could reuse those in an AIT algorithmic system.

One implication for the firm's economic costs is that, with an automated selling system including an automated recommendation system, the level of sales can be increased with approximately zero contribution of human work to the marginal cost (MC) of selling.[7] The resulting low MC, together with the fixed costs of designing the selling system, including using AITs in the recommendation system, and the need for a large body of data on multiple customers, led to considerable scale economies at the firm level. The day-to-day production process that leads to an advertisement shown or a product recommendation made is carried out by capital. MC would rise dramatically if human activity were required as part of each transaction. The fixed costs (FC) of these systems, on the other hand, are large and include much human work. The architecture of the system is designed, however, by extremely smart humans, not by machines.

about software: "Scalability is also not free in software. To build software applications that scale requires significantly more attention to load balance, locality, potential contention for shared resources, and the serial (or partly parallel) portions of the program."

[7] Amazon has vertically integrated into several complementary businesses, such as warehousing, where human workers do contribute to MC. The large number of Amazon businesses means that there are several places where the firm applies AITs. One example is inventory prediction in the warehouses, already a statistical prediction problem before application of AITs.

What about the shift to AIT from earlier algorithms? MC falls with the transition to AIT if AIT does a better job of recommending than did the prior algorithm. The human efforts to design and specify the AITs themselves contribute to the large FC of the Internet Giants. Invention that has a large FC to lower MC will be economic only for large firms. Whether "large" means as many product recommendations as occur at Amazon or a significantly smaller number is a topic to which I will return in Section 6.5.

Finally, the product recommendations made by Amazon are just that – recommendations. The consumer ultimately decides what to buy.[8] This has important implications for the loss function associated with bad match predictions. Choosing to recommend a product the consumer does not buy may be a lost revenue opportunity for Amazon, but, typically, it has no broader negative consequences.

The implications for the economics of adding AIT to the recommendation algorithm are straightforward. The FC of switching to an AIT-based production system can be spread over a large volume of sales. AIT is based on statistical prediction, so AIT-based matching systems have the four possible outcomes listed in Table 6.1.

A core feature of AIT product recommendation systems is that they are statistical. The profitabilityof the system thus depends on its effectiveness, on the rate of true positives, and on the error costs. Increasing the rate of true positives increases profits by increasing sales.[9] At Amazon's scale, modest increases in that rate represent considerable dollar profit. Profits increase whether the rate of true positives increases at the expense of either false negatives or true negatives – a profitable sale occurs either way. The second element of payoffs is error costs, which are small in this context. The error costs associated with false negatives are no more than the lost profits to Amazon and the lost purchase opportunity to the customer. Similarly, a false positive is just a sales recommendation that was not accepted by the customer. The role of the recommendation engine

[8] Agrawal et al. (2018), in a "thought experiment," consider the possibility that the Amazon prediction engine might become so good that it ships products without the consumer choosing them, and point out that Amazon has clearly done some technological development that might lead toward this. While such a change would fulfill one version of the firm's founder's early vision, it would require significantly more than a statistical improvement in the prediction engine. That last step of consumer product choice after recommendation makes the current system much more forgiving of errors in prediction.

[9] Much of Amazon's costs is the costs of getting the customer to the point of receiving a recommendation. The incremental costs of making an additional sale are largely limited to costs of goods sold and of fulfillment.

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

Table 6.1 *AIT-based matching system results*

| Results | Outcomes |
|---|---|
| True positives | Made suggestions that led to sales |
| False positives | Made suggestions that did not lead to sales |
| False negatives | Unmade suggestions that would have led to sales |
| True negatives | Unmade suggestions that would not have led to sales |

as advisor to the user means that its output is not the final word; the user can simply turn down the recommendation. This lowers the stakes for false positives. In Sections 6.22, 6.23, and 6.5, I shall examine other AIT applications where the costs of false positives are significant. Those costs are not, however, significant for Amazon, which has a *low stakes loss function.*

The availability of enormous data, the use of AI to achieve large scale at low MC, the modular system in place before AIT was deployed, the readily available payoff function, and the low-stakes loss function for match errors will reappear as systematic features of applications at the internet giants. Together with Amazon's strategic goals and position, and the firm's terrific technical capability, these features provide much of the explanation of Amazon's successful adoption of AIT for product-matching prediction.

### 6.2.2 Google

Google's largest revenue product is targeted advertising. The firm's online, mobile, and voice search products match particular searching consumers to particular advertisers. Google runs an auction to decide which advertisements, in which order, each searching consumer sees. Each consumer is more likely to click on some ads than others if they are shown. Part of the complex rules of the auction makes it easier for an ad to win if that particular consumer is more likely to click on it. To implement those rules, Google uses an AIT engine to predict specific searcher–ad click rates. This function is similar to the Amazon product recommendations engine described in Section 6.2.1, though the difference between a product recommendation and an advertisement means that the details are different.

Google uses AITs in a variety of ways: to attract users (e.g., translation), to communicate with users (e.g., Google Assistant), and, critically, to rank advertisements in a way that is targeted to each user. Google Search has been extensively studied in marketing and economics, as has Google's auction-based system for selling searchers' attention to advertisers (see, e.g., Varian,

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

2007; Athey and Ellison, 2011). I will review it only briefly, with emphasis on the parts that draw on AIT to match consumers and advertisers.

Consumers use Google to find information, including information about products and services they might buy. Google has significant big data about many of these consumers, based on both their searching activities and on other use of Google products, such as Gmail.[10] When a consumer searches, two kinds of result are returned. The "organic" search results are Google's guess at what the consumer was looking for. There are also advertisements, that is, information that is seeking the consumer's attention. A particular advertiser's ad will be displayed, or not, depending on the outcome of an advertising auction.

For each search, Google runs an auction to sell the searcher's attention. These auctions can scale to millions of searches per second because, on Google's side, they are automated. The auctions are granular; advertisers can choose to target detailed "AdWords." The AIT forms one important element of the auctions. Google is paid by advertisers when users click on their ads, not just for showing the ad. To maximize ad revenue, Google uses a system for ranking advertisements.[11]

Google's system, loosely called "quality score," has elements that are calculated in real time for each advertisement in each auction.[12] As the relevant economic theory makes clear, a central part of Google's profit maximization problem is predicting the probability that *this* consumer will view *this* ad if it is shown in a particular slot; that probability – not just the advertisers' bids in the auction – determines Google's expected ad revenue. AIT is used in determining this advertising ranking for a particular consumer search.

[10] Many Google products use AIT to offer the consumer a better service so that the firm gains more user data. Gmail, for example, has a smart reply function. Searched-to pages or entered text can be machine translated and recorded.

[11] The interaction between this ranking and the incentive-compatible elements of bidding in the auction are very well explored in the relevant economics literature. The core incentive idea is simple: if the best slot for an advertisement is filled by the highest-bidding advertiser, advertisers who get clicked on only very rarely – but who make a large profit if clicked on – will bid their way to the top. This is great for them, but bad for Google and likely bad for searching consumers.

[12] The elements are the click-through rate, the "relevance" of the ad to the user's search, and the "landing page experience" if the user clicks on the ad and goes to the advertiser's website. "Quality score" refers both to a number and to the system that generates the number, which can change the number in real time without informing either advertisers or consumers. Google has excellent reasons for imprecision in its public discussion of its search products, since websites and advertisers might otherwise game it. They game it anyway, but the imprecision lessens the effectiveness of the gaming.

Predicting the probability of a click on an ad is a near-ideal use of predictive AIT with machine learning. The situation is complex. The probability is specific to a given advertisement, advertiser, searcher, search terms, device the searcher is using, and time of day, as well as to the interactions among all those factors, for example, the relevance of the advertisement to the search. Google has big data on all these items. Finally, for profitability, the ad ranking does not need to predict very well; it just has to predict well enough for Google to achieve significant revenue from the ad auctions. The prediction system that ranks ads for the auction has relied heavily on predictive AITs for some time.

Prediction – in this case prediction of the revenue that will come to Google from showing a particular advertisement to a particular consumer, or to Amazon from recommending a particular product to a particular consumer – is one of the important parts of modern AIT. In this case, "prediction" means exactly what it means in basic statistics. The "deep learning" part of the system is automated estimation of the prediction model. Automation scales well. Google's computer can create prediction models that apply to a very large number of searchers, a very large number of searches, and a very large number of advertisements/advertisers/products without human intervention at the margin. Since Google has vast data, the prediction can be based on a great many data elements associated with the searcher, the search, and the ad. It is just prediction – the deep learning algorithm is not about *why* a group of data elements are good predictors of a match, only that they *are*.

When will deep learning work? It needs a lot of data, a lot of computer power, and a quantitative measure of a good prediction (in this case, the probability of clicking on the ad and/or revenue). Deep-learning algorithms work well in predicting based on many complex data elements <u>if</u> the sample sizes are large. All of these conditions for success by a deep-learning algorithm are satisfied for the problem of ranking different advertisements in the Google AdWords auction. Deep learning also needs to avoid bad product–consumer matches if they are high stakes.

The role of the AIT in this case lies in matching specific potential buyers (the searchers) to specific potential sellers and their products and services (the advertisers). In this regard, the purpose of the AIT application at the heart of Google Search is quite like the product recommendation at Amazon.[13] The economics of this early highly valuable application of AIT also follows the

[13] AITs are used elsewhere in Google Search, notably NLP technologies in RankBrain and in voice search.

same logic as at Amazon. Google has very large scale with enormous big data for the machine-learning part of the systems, has a search profit model that requires its systems to scale at low MC, had an already modular system before the AIT was deployed, has a readily available payoff function for the learning engine to maximize, and, finally, has a low-stakes loss function as the displayed advertisements are advisory to the consumer.[14] To be profitable, AIT matching needs only to cover its (high) FC of invention by increasing the probability of a successful match a small amount.

### 6.2.3 Facebook

Facebook is also ad-supported and it also uses AIT to match particular advertisements to particular consumers.[15] The business logic of that subsystem is like those just seen. It is a product–consumer matching problem, finding the advertisements that will draw a response from a particular consumer in particular circumstances. Facebook already had a scalable, modularized system based on very large volumes of data (including social connections data among billions of users) before it began to use AIT. Like Google, Facebook uses AITs to make consumer-attraction features outside its core production process.[16]

Facebook is also different in ways that will help us illuminate the economics of AITs. Facebook has big data on the "social graph" among users, information that it uses to decide what information, including ads, to show to users, in contrast to Amazon's and Google's observations of consumer searches for information, products they might buy and topics they want to know about. Facebook is used by consumers as a communications medium, so there is a flow of information across the social graph influenced both by users and by Facebook. As a result of these information differences, Facebook has a different set of algorithmic tasks. Facebook chooses which advertisements to show users without an explicit product search by the user. This was long algorithmic and is now based on

[14] Google blocks ads on certain kinds of searches because some searchers or advertisers might find them offensive. This avoids a high-stakes loss problem.

[15] Quoting Mark Zuckerberg on Facebook's July 2017 earnings call: "Now you can put a creative message out there, and AI can help you figure out who will be most interested."

[16] Facebook has both an AI research group and an Applied Machine Learning engineering function. Their inventions include automatic picture tagging, which uses photo recognition AITs and is deployed as a user-decision-support system, other social recommendations made to users for their potential action, and machine translation. Facebook also attempts to detect problems such as suicidal users and posts that are not from real users, etc., an area I shall revisit later in this section.

AIT. Facebook now uses AIT in the long-algorithmic function of deciding what nonadvertising information to show, that is, in populating users' "news feed."

The algorithm used by Facebook to prioritize items for the news feed uses "who" information as well as "what" information. The Facebook "social graph" forms the background to the news feed, so the "who" includes the poster of information as well as the reader of it. A tricky bit – to which I shall return – is the interesting difference between *I want to see this item on my news feed* and *You want me to see this item.* All of these factors were part of Facebook's move from its "EdgeRank" algorithm to one using significantly more AIT to fill the News Feed in 2013.[17] The feasibility conditions for using machine learning – tons of data and much information about what users like (literally "like," or read, or don't hide, etc.) to give the optimizer a quantitative goal – are well satisfied here.

In a closer parallel to Google, Facebook also uses AIT to target advertisements to specific consumers. For commercial ads, where there is a low-loss function, many of the same positive conditions noted in Sections 6.2.1 and 6.2.2 apply. Facebook has scale, needs scalability (low MC through automation), has very "big data," and had, because it needed to, an already modularized production system before widespread use of AITs. Matching an ad to a user is a difficult prediction for an algorithm, but the vast amount of low-quality "big data" can make that prediction more accurate. There is a low-stakes loss function within commercial advertisements for both false negatives and false positives.

Beyond commercial advertisements and the narrow limits of reading posts from one's own friends, a few of the important limitations to the use of AI matching technologies have been revealed at Facebook. The complexity of the entire system of Facebook – readers, posters, friends, friends of friends, likes, comments, the whole social graph – makes it difficult both to decide what algorithms should do and to set rules for posters, advertisers, and so on. Problems with a policy change can ripple through the complex system and then blow up. Facebook has systematically used business model experiments to learn what constitutes a mistake, to apologize, and to repair.[18] Typically, these experiments cannot be contained within a sample of users, for users interact and overlap. This makes

[17] See McGee (2013) who notes that Facebook had switched to an algorithm based on machine learning.

[18] A typical apology from Facebook's Mark Zuckerberg can be seen in McCarthy (2006). There have been over a dozen significant experiment/problem/redesign/apologize cycles.

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

Facebook a great place to examine the limits of recommendation systems based on predictive AIT, especially the limits associated with actions associated with a higher loss function for false positives because the outcomes matter too much to consumers.

One area where Facebook has experienced limits to the use of AIT is flagging "inappropriate content."[19] Partly, this is because some "inappropriate content" involves higher stakes associated with false positives. In many cases, a user is shown a post they dislike so intently that the mere fact that it was shown to the user is seen as a significant negative by the user.[20]

This problem has been made worse by the complexity of the "social graph" and by the creation of communities that insert and push content that others see as inappropriate, where in some cases "others" is nearly everyone. These communities might meet in Facebook groups, but they might also meet elsewhere, for example, in Reddit, to plan coordinated assaults. These communities create problems using a number of strategies, including manipulating humans to like the post, creating fake humans, and so on. The combination of high-stakes loss categories of content and the organized pushing of such content has led Facebook to retreat from its AIT-based inappropriate content system and to hire tens of thousands of human content editors. Guy Rosen, Facebook Vice President of Product Management, provided the rationale for all of this human effort in a recent blog post: "[W]e have a lot of work still to do to prevent abuse. It's partly that technology like artificial intelligence, while promising, is still years away from being effective for most bad content because context is so important."[21] Rosen cites three problems for AITs. The first two are about statistical power to discriminate between problematic and regular

[19] Another area of algorithm use is suggesting "people you may know," which sometimes suggests people you really don't want to know or, perhaps, be reminded about. This can be a high-stakes loss function for an erroneous false positive recommendation, but it seems to be a problem that can be contained.

[20] One example is the reaction to the Cambridge Analytica scandal, where many people and governments felt that Facebook had crossed important limits. My point has nothing to do with the merits of any of those arguments politically, in terms of privacy policy, the regulation of "troll farms," or anything similar. Instead, what is relevant to the inquiry about value creation from the use of AITs is that the scandals reveal difficulties in algorithmically policing political posts because the reader can be offended by seeing one.

[21] Rosen, "Facebook Publishes Enforcement Numbers for the First Time," https://newsroom.fb.com/news/2018/05/enforcement-numbers/. The Facebook representatives discussing this problem, good technologists all, tend to say the problem is that AIT has not yet advanced enough. This error lies somewhere between the anthropomorphic metaphor – the future of AIT is anything people can do – and the ordinary techno-centrism of those with computer science training.

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

Table 6.2 *Incidence of and algorithmic effectiveness against inappropriate content on Facebook*

| Form of problematic content | Incidence* | Algorithmic effectiveness** |
|---|---|---|
| Spam | 837.0 | nearly 100% |
| Adult nudity and sexual activity | 21.0 | 96% |
| Graphic violence | 3.5 | 86% |
| Hate speech | 2.5 | 38% |

* Number of items removed by Facebook in the first quarter of 2018, in millions.

** The percentage "identified by our technology before it was reported to Facebook."

content: (1) telling the difference between "someone . . . pushing hate" and someone (else) telling of their own experience to raise awareness of a problem, and (2) that Facebook lacks sufficient data (!) for machine-learning "training" for problems that are not frequent (e.g., a new example of hate speech, early in its ugly life). To understand these first two problems for AITs, consider Table 6.2, which is based on the Rosen post.

Success in AIT fighting spam (the first row) is much like success at Google in ordering advertisements. Spam messages succeed by reaching an enormous number of readers of whom a tiny fraction click. This gives the spam detector AIT plenty of sample size – nearly a billion messages a quarter – to work with. Meanwhile, spam, while annoying, is not outrageous, so the loss function for false positives is low. The AIT spam detector usually wins the race against complaining readers.

The pornography and violence lines are different. Picture recognition, like other aspects of machine vision, is at an advanced state in AIT. On the other hand, people are likely to react strongly to being shown offensive images. The loss function for failure to recognize an offensive image is high, so the race between human complaints and AIT is won rather more frequently by the human complainers when there are new pornographic or violent images. The high-loss function for errors works against AIT.

Hate speech is even lower in frequency and in AIT's rate of winning the human v. machine race. This is easy to explain as high-loss-function hate speech poses severe problems for AITs. Human complaints about hate speech multiply rapidly in a social network if the speech arrives either at its (hated) target audience or at enough people appalled by it. This leaves every new hate-speech utterance with relatively low sample size before detection, and little information to assign a negative payoff in a machine-learning algorithm until after the hate speech has done significant harm.

Thus, human complaints typically win the race with machine learning for new hate-speech utterances. This is a structural problem of the application context, not something that improvements in AIT itself can easily remedy. Not surprisingly, after some experience with AITs, Facebook decided to utilize "people power" in combatting hate speech.

"Fake news," while not cataloged in Mr. Rosen's list, is also getting a large application of people power at Facebook.[22] This is wise on Facebook's part. Predicting what is "fake news" is a daunting statistical problem, as any particular piece of "fake news" is highly welcome by some readers – not always many – and highly disliked by others.[23] The loss function has high stakes.

Mr. Rosen flags a third problem: "[W]e're up against sophisticated adversaries who continually change tactics to circumvent our controls, which means we must continuously build and adapt our efforts."[24] This is an old market regulation problem, long familiar to economists (usually in the context of public policy rather than business policy). AITs based on machine learning are about statistical prediction rather than causation, selection, or other structural considerations. A change in policy can lead to changes in market behavior that break formerly reliable statistical predictions – the ones that formed the basis for the policy.

This applies to changes in business policies that rely on statistical prediction in much the same way that it applies – familiarly – to changes in public policy. The policy change might cause changes in behavior that invalidate the prediction. Facebook discovered this the hard way. When Facebook introduces an AIT-based set of "controls" designed to block certain bad behaviors, its "sophisticated adversaries" then "change tactics" – invalidating the statistical prediction model. This is a limitation on the usefulness of AITs. It is difficult to use any statistical prediction model, including those in AITs, to detect problems when the problematic behavior will change in response to policy changes. If the stakes are high and sophisticated adversaries are deliberately seeking to impose the false-positive losses, AIT will struggle. Hate speech, trolling, and fake news

[22] For example, see CNET, "Can Facebook's New Hires Take on Troll Farms and Data Privacy?," April 11, 2018, www.cnet.com/news/can-facebook-mark-zuckerberg-new-hires-take-on-troll-farms-and-data-privacy-after-cambridge-analytica/.

[23] Some technologists have suggested, implausibly, that advances in AIT will enable machines to determine the truth of all news. This is one of the few AIT overreaches not driven by the anthropomorphic metaphor.

[24] Rosen, "Facebook Publishes Enforcement Numbers for the First Time," https://newsroom.fb.com/news/2018/05/enforcement-numbers/.

stories of high emotional impact fall outside this boundary. Because the problematic actors change behavior in response to a policy change, there will always be periods during which the AIT seeking to predict the bad behavior will be catching up. Because of the high-stakes loss function for errors, waiting until there is enough sample size to train an AIT-based algorithm is more costly to the firm than replacing the technology with human workers.[25]

This illustrates the difficulty of applying AI prediction technology when the stakes and thus the loss function are not low, and it is thus one of the current boundaries of application of AITs. It is also consistent with the initial large-scale production uses of AIT recommendation systems having been developed in low-stakes loss function environments, like showing an ignorable advertisement.

### 6.2.4 Netflix and Others

The entertainment distributor Netflix has significantly fewer products available than Amazon. Over that restricted range, Netflix faces a similar product recommendation problem. What movies or shows would this user like to consider next? The number of choices Netflix can present is limited not only by customers' attention but also by the clumsiness of the TV screen and the TV remote. The Netflix business model also depends on matching customers to content they like; it is particularly profitable to match users to content that falls outside the list of the most popular shows and movies. Finally, Netflix has impressive big data on users' past choices.

Netflix long used a matching algorithm called CineMatch to suggest movies. A 2006 contest offered outsiders a reward for improving the algorithm and supported them with a Netflix dataset, so more is known publicly about "big data" at this firm than at almost any other. The contest winners did improve the algorithm. However, Netflix uses an internally developed machine-learning algorithm. For this application, it is fair to say that AI matching technologies are better than a wide variety of human-written algorithms.

Netflix has also faced, and solved, lower-stakes versions of problems like "fake news" and hate speech, after it discovered early on that some movies

[25] These last two paragraphs are dedicated to every econometrician or statistician who has suffered in one of those appalling machine learning seminars in which we were told that all of the traditional concerns of econometrics and statistics – other than prediction – are outdated.

are liked by some viewers and disliked by others. (The example it discussed publicly was *Napoleon Dynamite*.)

Netflix has other AITs in use.[26] More recently – as the firm has become a more important producer as well as distributor – it has made "trailers," that is, advertisements for shows and movies. Which trailer to show a customer is an advertising choice. AI matching technology makes consumer–trailer matching recommendations, leaving Netflix with a blend of product recommendation engines (like Amazon) and advertising choice engines (like Google or Facebook).

A few smaller-scale – and more specialized – voice user interfaces (Voice UIs) have been introduced – voice search for TV programming at cable companies is an example. A number of entertainment-delivery firms use Voice UI to permit users to choose content. Both Comcast and Dish Networks, for example, offer a "voice remote." Typing with a TV remote is tedious. In contrast, the Voice UI permits the user merely to say the name of the item for which they're searching. These examples bring together the virtues of the Netflix Voice UI (limited range of vocabulary) and the virtues of the UI applications on cell phones (high value of UI improvements in environments that are difficult for typing) discussed in Section 6.3.3.[27]

### 6.2.5 Technical Progress Based on AITs

Some readers will be disappointed. The "deep learning" in these systems does not resemble human learning, and the use of AIT in these systems is not what people were imagining when they heard about AI – not as sci-fi as an anthropomorphized robot. Instead, these AITs are software technologies, not sci-fi. They are tools that permit the design of new productive systems. They are embedded in the capital of those productive systems. In that regard, these AITs are like earlier ICTs. They combine technical progress, tools for invention of applications, and technical progress embedded in capital.

[26] For example, Netflix needs to predict the bandwidth-management version of inventory stockouts and allocate accordingly. This problem is made more difficult as some content is more bandwidth sensitive than others. Machine learning technology has proved useful in this function, which has long been a statistical prediction problem.

[27] There are many other Internet firms and related prediction applications, including LinkedIn (for jobs individuals may be interested in), and Waze, Google Maps, applications searching for a best route or local vendors, Google Image Search for similar images, and TaskRabbit matching workers and tasks.

Task-level substitution plays no role in these applications of AIT. These very valuable early applications are not ones in which labor was undertaking a task and was replaced by capital. Observers focus on task-level substitution not because it occurs but because the definition of general AI includes "tasks usually done by humans." Until general AI is commercialized, which is not likely in the foreseeable future, analysis should focus on the capabilities and applications of actual AITs. While there may be some task-level substitution in the future, it is unrelated to the value proposition of AITs.

System-level substitution is an entirely different matter. System-level substitution, such as using (the supply chain that includes) Amazon instead of (the parallel, partly overlapping bricks-and-mortar supply chain that includes) bookstores, is important in ICT-based production. System-level substitution has led to a great deal of substitution of capital for labor in the ICT era. The newly designed systems that are growing tend to be more capital intensive and more human-capital intensive than the old ones they replace. The pace, locus, and scope of that substitution have multiple determinants, of which static cost saving, that is, the degree to which new kinds of capital can be substituted for labor in particular production process tasks, has not systematically been the most important. Instead, it depends on the pace at which whole new production processes and business models are invented (Amazon's store and mall were invented and have been constantly improved). System-level substitution depends on the competition between old and new firms, and on the effectiveness of old firms at inventing competitive responses (e.g., Walmart inventing e-commerce services). In short, system-level substitution entails a wide variety of opportunities and barriers to invention, involves competition, and involves the development of complementary markets and services. It is the opposite of local and simple – and of task-level substitution.

### 6.2.6 Matching Engines Like Earlier ICT Waves

The AIT-based matching engines discussed in this section are like many recent waves of ICT technology. Their largest uses are in (mass) marketing, they are complements to existing ICT systems and assets, and, at a system level, they increase the growth of capital-intensive production processes. As with many other early rounds of ICT, they are deployed in already modularized production processes, or they call for difficult-to-invent modularization. The technologies are scale-using and increase the degree of scale economies. These are familiar features of new rounds of ICT,

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

particularly since the great turn toward consumption and mass markets that followed the widespread use of the Internet.

Another important sense in which the matching engines are like earlier rounds of ICT goes back much farther than the twenty-first century. Facebook, as discussed in Section 6.2.3, struggles to control the system-wide implications of new communications technology because it changes incentives. This is an old, old story inside large organizations moving to a more digital production process.[28] What is new is that lowered costs of access to ICT services through cheaper and easier-to-use devices have made the scale of the "organization" wider in society than just a firm.

## 6.3 UI Improvements Based on AITs

A second important area of AIT application is in UI improvements, which have contributed to the lowered cost of access to ICT services, especially for consumers.

A number of voice-based "personal digital assistants" (PDAs) have been introduced, such as Alexa (Amazon), Siri (Apple), Cortana (Microsoft), and "Google Assistant." These form new UIs on mass-market general purpose consumer devices and offer the user a voice-based connection to many of the services running on or through those devices. Other new voice interfaces, such as voice search for video or audio entertainment on set-top boxes and new versions of telephone voice response units (VRUs), work on narrower domains. AIT-based improvements in text processing are making UI improvements on other dimensions, such as more-forgiving response to typed input from users. It all adds up to a remarkable technology deployment that is another very widespread and valuable use of AITs.

VUIs tend to make casual observers think of the Turing (1950) test. If the user "can't tell" that they are talking to a machine, but rather thinks that they are talking to a human, then the machine has achieved "artificial intelligence."[29] This confusion about the Turing test, however, makes an elementary mistake about what a UI is and does. A UI is an interface – it works between two things. In this case, it works between the user and a system or service, as indicated by Al Lindsay, the manager of Amazon's

28 See, e.g., Zuboff (1988) and Bresnahan and Greenstein (1996).

29 Many people recall the Turing test as "Could you tell?" – are you conversing with a human or a machine? The Turing test was set as a game. A computer and a human in isolated rooms, communicating only by text, vie to convince a (human!) judge that they are the human. Will the computer be better at appearing to be human than the human? Turing thus gave the machine a goal, anticipating much modern machine learning.

Alexa service (quoted by Oremus, 2018): "[W]hen I think about Alexa, I think about user-interface paradigm. I think about the voice interface only as a way to interact with technology, your platform, or a service that underlies it."

New and improved UIs can have very great economic importance. They can make user access to existing applications easier, and they can enable the invention of new applications that make economic sense only with easier access. In this sense, UIs are GPTs – but this is not the sense that observers have in mind when they hope that "AI will become a GPT."

It is easiest to understand the impact of the new AIT-based UI improvements by thinking about the series of mass-market UI improvements that preceded it. New UIs can make ICT systems available to more users by requiring less training in "using computers." The PC enabled computer use by a wide range of white-collar workers, especially after the deployment of graphical user interfaces (GUIs). It was economic to have more people, at more locations, having access to more computer systems with PCs as the UI devices than with the earlier "terminals." The World Wide Web and the web browser permitted access to applications and services from multiple locations. More and more enterprise and consumer-oriented systems could be used from more places, as many enterprise systems have a "web interface." And many consumer-oriented websites were enabled by the access improvements of the PC–browser combination.

The same analysis applies to a series of UI and access-device improvements in the consumer-ICT era. Smartphones and tablets are more portable than PCs, allowing access from more places. These new devices are "always on" and have touchscreen interfaces, permitting new kinds of access. These devices often connect to cell-phone networks (significantly more easily than PCs), permitting more access. Invention of complementary network technology, the "cloud," permits access to the same services from different devices and locations. These new UI and UI device technologies have increased the value of existing applications and enabled new ones, especially for consumers.This view of what UI improvements do reflects technological and business reality. Again quoting Al Lindsay of Amazon (from the same Oremus (2018) interview about Alexa): "I think about the voice interface as a natural evolution of those technology interfaces. . . . I think adding a voice capability to something like shopping just removes friction and it makes things easier for customers." Mr. Lindsay, of course, works for Amazon, so he hopes that UI improves shopping, but the general point is that UI improvements "remove friction" to increase access to applications.

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

The rate of technical progress in the AITs called natural language processing (NLP) increased when they became based on statistical prediction. They are similar to the matching engines examined in Section 6.2.6 but are distinct technologies. The value creation attributable to AITs is not associated with a broad, general scientific area called "Artificial Intelligence" but instead with specific technologies sharing and taking on directions of their own.

Google has invested in NLP technology, improving its Voice UI and its connection to underlying services. For example, Google "Voice Search" is an improvement in Google's core product: Search. Voice Search for YouTube videos has similar logic. Voice Control of Gmail has a different set of capacities, which are related to Gmail's increasing ability, based on a modularly separate AIT-based writing engine, to guess what message the author wants to write.[30] Google also exposes an application programming interface (API) for apps running on Android smartphones to take advantage of the Voice UI. Apple, the other important smartphone UI firm, also exposes an API to "Siri," so that non-Apple apps can use voice commands and use spoken output. Finally, both firms have (modularly separate) AIT services that try to predict how users will use their phone or tablet. It is easiest to describe that with the anthropomorphic metaphor, for example, Siri "suggests" opening a particular app, but it is also foolish, as this is precisely the same as the Google website "suggesting" a particular ad. Using devices in environments where it is difficult to type or read, such as in cars and kitchens, raises the value of VUI, of course.

These are statistical prediction technologies, so quality increases with the sample size of the voice data stream used to "train." Speech that can be linked to context and to the user's goals is particularly demanding of scale for training. After years as a moderately important technology, Voice UIs became important after cell phones were widely distributed. Cell phones are terrific locations to gather "big data" on voice and to use a VUI. Scale and competition played a role with Apple, Google, Amazon, and others in the mix at very large scale.

[30] Text UIs have also been improved using AITs. Since 2015, a Google algorithm called "RankBrain" running behind the search page tries to learn what the consumer might be interested in. Has a user typed a partial search? It suggests completions. Has a user made an oddly worded version of a common query? It guesses the underlying query. At the public announcement of RankBrain, Google stated that it was the third most frequently used algorithm (Clark, 2015). These NLP technologies, characterized in the Googlesphere as the switch "from strings to things," are valuable both to the searcher and to advertisers. They help searchers find what they are truly looking for and they help communicate what that is to advertisers.

The improvements to smartphone VUIs are made more economic by the very large-scale deployment of those devices, their use by UI-sensitive demanders (consumers), and the available voice data at large scale. So, this is – once again – a class of applications of AIT in which existing complementary assets, scale, and competition among the Internet Giants play a role.

### 6.3.1 Amazon: Alexa and Kindle

Alexa, a Voice UIT, and new hardware clients Echo were introduced together by Amazon. Alexa software also runs on other devices, including Android and iOS smartphones and tablets.[31] Many of the popular systems accessed through the Alexa Voice UI are home-control or media-demand applications. Alexa has been a hit, turning Amazon into a home-device company with large scale, and drawing competitive responses from other Internet Giants.[32]

Amazon has now made two successful attempts to create consumer clients, Kindle and Alexa. Compared to a smartphone or a tablet, each is more of a special purpose device. The Kindle functions primarily as an e-reader.[33] Echo devices running Alexa have only a voice interface and do not include a screen or keyboard. Both Kindle and Alexa embody new UI elements suitable to their goals. For both, Amazon's success represents not only impressive technical progress in the client devices and software themselves but also an extension of Amazon's large-scale store and rapidly growing digital media business to new distribution channels. One can read a book on Kindle; many of an Amazon Prime account's features, not least music, work through Alexa. Given this large-scale distribution strategy, it is not surprising that software clients for both Kindle and Alexa run on many non-Amazon devices. Kindle and Alexa are also complements to Amazon's product-matching AIT engines as well as to its other services, algorithmic or not.

[31] There are Kindle clients for PCs, for Android, and for iOS devices. There are Alexa clients for Windows PCs and Android phones. Alexa clients for Apple devices are a more complicated story at the time of writing, with Alexa running in the Amazon store app on iPhones but not yet as "Alexa," for example. Apple pleads the "app approval process" as usual, though suspicion of competitive motives (Apple iTunes versus Amazon Prime music, for example) swirls in the vast Apple rumor mill.

[32] Alexa involved fundamental advances in voice NLP, such as picking one voice out from ambient conversation or other "noise." It also involved building a new application system.

[33] Like most efforts to create an Android tablet to compete with iPad, versions of Kindle that run many Android apps have not had much market impact.

Despite their differences, I put Kindle and Alexa in the same group for two reasons. The first is to emphasize the importance of scale and marketing in the widespread deployment of new UIs. Amazon is a mass-market online store and media company, and had a powerful economic motivation to build a client presence in the mobile era. Amazon, particularly through its Amazon Prime volume discount program, was also well posed to encourage consumers to adopt Alexa. The scale of existing complements and the modularity of Amazon's existing systems encourage new UI invention.

### 6.3.2 Scale (of Complements Already Distributed) and Continuity of Invention

Like the technical improvements in smartphone VUIs, the technical improvements behind Alexa are impressive. Alexa can pick out individual voices in a crowded room with several people talking, making the UI more valuable in kitchens, living rooms and (soon) automobiles.[34] Siri, Google Assistant, and others can learn the voices of heterogeneous speakers. These are just some of the technical achievements in voice and text UIs in our era.

Despite this high rate of technical progress, the UI improvements do not, yet, materially alter the direction of technical change. They improve the ability of ICT-based systems to support media, retail, and related applications in large-scale, consumer-facing deployments. They have broadly the same economic implications as other important ICT advances since the widespread use of the Internet. They are capital-deepening marketing technologies deployed largely in mass markets. The existing complements and the large scale of existing applications have created a powerful economic incentive to use new UI inventions in this largely narrow capital-deepening direction. Demand forces have not yet pulled AITs far beyond that range of consumer and mass-market applications, just as they did not pull earlier rounds of impressive technical progress like the smartphone far beyond it.

### 6.3.3 What Will the UI Improvements Enable?

Platforms – in the narrow sense of that phrase, GPTs over which applications may be built – can enable applications in a narrow and immediate sense (a new application built on the platform) or in a broader sense (the

[34] Amazon acquired a startup working on those capabilities as part of its building of Alexa.

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

platform recombined with other elements in new systems). For the new UI elements in mobile devices, there are conjectures about both senses.

Before adding AIT elements to their UIs, Android and iOS enabled a great deal of "app" complementary innovation. Access by a mass market of consumers enabled consumer and entertainment applications such as games. Access to that same mass market of consumers enabled consumer product and services firms to create marketing and customer service apps.[35] There were, of course, some apps for tablets and smartphones that formed other parts of the production process. But consumption, sales, and marketing have been the center of it. As the UIs of smartphones and tablets improve with voice capabilities, the immediate direction of application change stays squarely within that area.

Alexa "skills" are examples of VUIs enabling applications in the narrow sense. Alexa hardware devices open up opportunities for consumers and for those who would like to reach consumers. Accordingly, there is a new range of applications, programmed into the "skills" APIs, that run on Alexa machines and on the networked system behind them. These are, once again, largely applications aimed at a consumer end-user and thus fall in the range of media, entertainment, sales, and marketing.

Task-level substitution also plays no role in these UIs, which substitute AIT capital for capital to increase convenience and access. New software technology – based on NLP (AIT) – replaces old as new UIs replace earlier UIs in some uses. They also expand the use of devices to new activities. Voice UIs partially replace touchscreen UIs, WIMP (windows, icons, menus, pointers) UIs, and so on. At a task level, this is substitution, not of machine for human but of machine for machine.

One sense in which there is substitution of capital for labor through UI improvements is increased convenience for the user. Some user time can be saved, either by the UI permitting a task to be done during less-expensive time or by the UI enabling an underlying system to respond more quickly to the user. "Alexa, play Fox News" (or CNN for other tribes) saves a walk across the kitchen, for example. This sense of saving consumer time could become a related sense of saving worker time, which is one of the directions of diffusion of AIT-based UIs under active consideration today. But, thus far, time saving is not the centerpiece of even consumer UIT; instead, the centerpiece is broadening access.

[35] The latter category was later than the former but surprisingly large (Bresnahan, Davis, and Yin (2015)).

### 6.4 Broader Diffusion: Other Marketing Applications

The early successful AIT applications have created interest among firms beyond the Internet Giants in a number of different AITs. These include the prediction technologies used in the matching and targeting applications, NLP technologies including voice and text, and perception technologies such as image recognition and matching. While there are no large areas of application, the early examples and the world of ideas make it easier to fund any project that can be labeled "Artificial Intelligence,"[36] so there are many experiments.

Programmer toolkits are available for many AITs.[37] This reduces the narrow programming costs of an AIT application but not the (usually) more difficult and expensive part of a novel application, its business specification. Each application must be invented, will still have its own costs (including error costs), will fit in existing systems modularly or not, etc.

There is little application of the AITs outside the Internet Giants as of spring 2018. However, over the last two years, large firms' approaches to AIT applications have crossed from speculation and investigation to experimentation and development.[38] Surveys of firms about their applications intentions now mention specific use cases; there are applications plans, applications experiments, and ideas.[39]

The largest category of experiments is *marketing* applications. Using the AITs' underlying product–consumer matching systems and NLP, "chatbots" and the like have been applied in marketing interactions with customers; the scope is not only initial customer acquisition (advertising) but

[36] This is one narrow sense in which the "AI Technology Boom" resembles the Internet Bubble (and its underlying boom). Then, as now, the CTO could dominate the CFO because of the widespread interest in a technical area.

[37] Through Amazon Web Services (AWS), Amazon is an important supplier of services for applications for Web, Cloud, etc., especially those with "big data" storage and programming needs. AWS now includes a large number of AITs. Microsoft is also an important supplier of tools and services for cloud computing, now including AITs. Other established firms, such as IBM, Google, and Oracle, are supporting their customers with new toolkits in this area, as are many startups. The big consulting houses are all seeking to establish "thought leadership" in AITs. A number of deep learning software technologies have been moved to open source.

[38] Compare the sources cited in Bresnahan and Yin (2017) to more recent sources, e.g., Chui et al. (2018), Walker, Andrews, and Cearley (2018), or Schubmehl (2017).

[39] A Gartner survey of chief information officers is representative (Gartner Research, 2018b, p. 20) "Only four percent of CIOs say their organization has deployed AI, but we expect a substantial increase in deployments as one-fifth say they are experimenting with AI, or have short-term plans for AI."

also answering customer queries, supporting/encouraging repeat purchases, retaining customers for the long run, etc. Surveys show that the area many firms see as a priority for ICT-based technical progress, generally, lies within the marketing function. I will summarize these priority areas as "customer experience" (CX).[40] Surveys of firms about their AIT application plans typically show this as the largest area of planned applications growth. This is the main direction of diffusion in the present.[41]

This early diffusion has many of the features of the applications at the Internet Giants. Scale is important, as is the presence of complementary capital assets, such as big data, preexisting systems that communicate with customers, and so on. The early stages of diffusion, thus far, are not revolutionizing the way in which ICT is deployed in production, but they are improving ICT systems, and advancing them along their existing path.

While the complementarity of new AIT capital with existing capital is more important, the sales and marketing applications to which AI matching technology and AI NLP technology are now starting to diffuse have some modest prospects for substituting machines for human work. We can expect both some expansion of the range of customer support activities and some replacement of customer support people. This is *not* an outbreak of cost reduction via task-level substitution. Instead, the relevant demand-side forces are competition to improve consumer experiences and preexisting knowledge of where automated CX might work. Some of that knowledge comes from old "VRU" efforts, which were often inconvenient for the customer. The voice chatbots are a better version of VRUs. Other parts of that knowledge come from ineffective FAQs – text chatbots are a more targeted (with AIT doing the targeting) FAQ page.

### 6.4.1 Diffusion of Matching Engines to Advise Employees Rather than Customers

Another area of potentially important application experimentation, the second largest according to surveys, deploys AIT to advise employees.

[40] The areas include "customer engagement," "customer satisfaction," "customer support," "customer experience," and others. For example, Murray (2018) looks at consumer package goods manufacturers and reports that CX is the top priority for marketing technology spending. He argues that of the technology areas that might be used to improve ICT-based marketing, AIT is still emergent as a purchase driver, as just under half of firms anticipate that it will influence their "marketing technology" purchases going forward, while about three-quarters anticipate that CX will influence those purchases. Other product areas (described in reports not cited here) are similar.

[41] See, e.g., Murray (2018), Rollings (2018), and Gartner Research (2018a).

This is a bit of a portmanteau category, as it includes at-work versions of "digital assistants" like Alexa and Siri, improved help functions in enterprise software like Cortana, as well as systems described as decision support. Some aspects of this category may be less certain and farther off in the future than others.

The use of Alexa, Siri, or Google Assistant to undertake chores at work, much like their use at home, is one direction of diffusion. It has very high visibility, changes at most the job of the worker at hand rather than the organization, and provides only output chosen by the worker, thus avoiding outcomes with losses. One early focus seems to be on simple chores, for example, Alexa skills for sending an email and setting up a meeting. The factor market implications appear to be a modest increase in individual worker productivity.

Actual organizational productivity improvements following from this kind of use are harder to forecast positively. As with many earlier rounds of ICT adoption, the use of smartphone-based email has led to accidental organizational changes – email from the boss at night. Voice UI on the email device is not the solution to this – and "AI-based screening of emails" confuses a technical problem with an organizational one.

Another early focus is in enterprise software ease-of-use. Cortana's role in Windows plus Microsoft's role in enterprise software have led a number of observers to forecast growth for AIT-based Cortana in this role (Finnegan, 2018). The software predicts what the user wants to do next and suggests that more prominently. This is a use of matching technology that increases the productivity of software users and may have a low loss function.

A third area of experimentation is improvement of decision support (DS) systems using AI matching technology. Indeed, some forecast that the category of decision support will be taken over by AITs entirely. That would be a substantial increase in the role of AIT in advising human decision-making. Adding AIT to existing DS systems may let them draw on big data to make better recommendations. This area is immature, so it is not obvious what recombination with what new inventions can or will advance it beyond low-error-cost applications. At least some AIT applications can take advantage of a modular boundary already drawn between the decider and business systems. Growth beyond those low-adjustment cost areas will be more difficult.

These are areas of potential expansion in AIT use beyond the marketing function to other areas where existing systems can be improved by AIT interfaces, at the boundary between ICT systems and human users. Their labor market implications are approximately the opposite of task-level

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

substitution: the human worker continues to do their job with better input from the AIT-based computer system.

Finally, some AIT demonstration projects have generated excitement and are spinning off useful applications. The most visible of these is the "driverless automobile" originally envisaged as a task-level substitution technology involving replacing a driver with an "autonomous" vehicle. The very important inventions from this effort are rapidly moving into use as driver-assist technologies rather than as task-level substitution. Many observers have made the – vapid – remark that a "driverless car" "proves" that "computers can do anything humans used to do." The vapidity lies in the laxity about proof and about task-level substitution, not in the engineering. Commercial application of driverless trucks awaits large-scale modularization of truck drivers' jobs, which typically involve much more than driving, especially for short-haul trucks. Market application of the "driverless" technologies as driver-assist safety features in high-end automobiles is growing rapidly.

### 6.4.2 Diffusion Implications

The early diffusion – today in experiments – of AITs away from the Internet Giants suggests a narrow range of very valuable applications of ICT-based production, many in the "customer experience" elements of mass-market selling efforts, most involving complementarity with existing capital assets of a particular form, and many involving scale and scalability. In short, the early stages of diffusion look much like the initial highly valuable applications. Other applications appear likely to increase individual worker productivity. In terms of scope, the early range of valuable diffusion looks like other recent waves of ICT, such as mobile devices and Web 2.0, with much capital deepening in the consumer-marketing-oriented industries. Scale economies at the firm level, at least for the high-value CX applications, appear important in the early diffusion path, just as in the initial applications.

### 6.4.3 Laboratory Results Close to Use

A third category of early applications is those in which laboratory results are close to commercial use. These do not reflect diffusion of production process inventions such as those described in Sections 6.2 and 6.3. They are, instead, a series of parallel tracks. Most of these do not call for modularizing organizations or production processes.

Production scheduling, inventory management, shipment scheduling, and related systems often relied on statistical prediction by algorithm before AIT. Inventory measurement has come to be automated at many factories, warehouses, and retail stores, and the problem of predicting inventory stockouts has long been statistical in many firms. Demand forecasting for capacity management at hotels, airlines, and so on – anything with capacity constraints and/or a queue – is a related area. AIT draws on machine learning to make a better statistical prediction of the same thing. Similarly, decades of automatic measurement for process control have led to algorithmic process control systems with elements of statistical prediction. In these areas, machine learning's value proposition – offering a better statistical prediction – is pushing on an open door: serving multiple rounds of control technologies, many with a statistical prediction element.

In finance, some asset market traders have prediction models of trades and price movements. These are also reliable adopters of new technologies that will let them get a slightly better or faster prediction. They need little organizational change to link to AIT prediction models.

A more complex example can be seen in credit card fraud systems. These are statistical, but for years they have faced a tradeoff between effective fraud detection and customer service. A phone call or text to a customer about a fraudulent transaction can stop the fraud. A false positive fraud warning, however, can annoy the customer. This is a problem of high stakes. Both statistical goals – predicting when a transaction is fraudulent and predicting when a message will annoy a customer – could be improved by AIT. This is one of those interesting examples where the system has been designed for decades to deal with the problem of a high-stakes-loss function for false positives (incorrect flagging of fraud). The addition of AIT that predicts better than an algorithm – better on both sides, not just in catching more fraud – need not decrease the level of losses.

Science and engineering applications, generally, have significantly smaller organizational adjustment costs than commercial applications. Further, many scientific and engineering disciplines have a strong statistical tradition. Unsurprisingly, AITs are coming into use as scientific toolkits. For example, Cockburn et al. (2018) are particularly interested in such important commercial scientific areas as drug discovery and development. They offer an interesting analysis of reorganization of the research process to take advantage of deep learning.

This might offer a different starting point to the diffusion of AIT as a GPT. Important ICTs, from the computer to the Internet, have started life as scientific and engineering tools. Leaps from "technical" to commercial domains typically require recombination and application invention and thus take decades.

AIT – including voice and image recognition – is being used for computer and network system security. This is an area in which the early applications were close to the laboratory, but considerable progress is being made in the field.

AIT-using security systems raise the return to better sensors more widely deployed. This is creating a great deal of experimentation with "biometrics," making an image of a person and matching it to a file of approved images to verify identity. The image could be a fingerprint, a photo of the iris, etc. These techniques may also diffuse to security systems that are more organizationally complex than computer systems access – for example, in air travel. Once again, we can see a path toward more general use with potentially greater and more widespread economic value creation with organizational changes.

Improvements in ICT security systems are complementary to the expansion of ICT-based production systems. ICT-based production systems, including marketing, entertainment, transacting, control, and the production of information goods, among others, are subject to security threats. Improvements in security systems permit ICT production systems to be delivered more conveniently, for example, through "cloud" or mobile channels. Improvements in ICT security systems remove an externality – fraud or theft – that holds back the broader use of ICT. Thus, security systems support the application of ICT-based production broadly and generally. The factor market implications of better security systems are those of system-level substitution of ICT-based production for older production systems as well as those of ICT-based production of new goods and services.

Photo and voice recognition applications are diffusing to other kinds of nonsecurity application as well. An early use was in flagging pornography at Google. Product search by photo makes product search generally more valuable. There are a number of examples in which pictures and sound recordings are being used in growing commercial systems. This diffusion involves some – but not a great deal of – adjustment costs. The adjustment cost boundaries that are being most tested and pushed are likely in the security examples.

## 6.5 Conclusion

AITs are a highly valuable group of technologies that represent a substantial increase in the *rate* of technical progress in ICT. They do not, however, represent a major change in the *direction* of technical progress in the applications of ICT.

These new waves of technology continue a twenty-first-century trend. Much of the profitable and inventive ICT in our century lies in consumer-oriented applications (retail, entertainment, mass-market product and services businesses, etc.) and devices (smartphones, tablets, etc.) creating value for new network forms (cell service, Wi-Fi)[42] and in mass-market marketing and sales applications. I do not mean to imply that this is all of the ICT applications in our era; rather, it is the subset that is creating large private returns to invention and to application, and to which technologists point when they claim that we are in an era of rapid technical change. The discussion in this chapter has repeatedly shown that the application of AITs follows these developments. What are these developments? Why did they occur? Why do AITs reinforce them rather than changing their direction?

The Internet era has seen much consumer-facing technical change. The most visible of these changes – and the longest awaited – are breakthroughs in consumer-oriented devices, cell phones, tablets, media players, e-readers, and now smart assistants. Now AIT in the form of NLP is reinforcing the trend toward devices that can be used more quickly by more people in more places.

Complementary to the new devices are online mass-market products and services, first based on the widely used Internet and the World Wide Web, and later on "Web 2.0." These enabled many mass-market services to be delivered online. Another set of complements arose because the new consumer devices were largely mobile phones and tablets. This enabled mass-market mobile web applications, and then mobile applications. Complementarities among these technologies make them mutually re-inforcing, as, for example, consumer-oriented websites became mobile websites and then mobile app-accessible services. Recombination, such as cloud technologies that link web, mobile, and e-reader access to the same content, communications services, and e-commerce, has sped the mutual reinforcement. AITs, particularly AIT matching services, are now being recombined with these existing assets into a web of mutual reinforcement. Again, there is substantial deepening of an existing trend in the early AIT era.

[42] See Bresnahan and Yin (2010) for a discussion of the causes of the emphasis of ICT on consumption and of the limited spread of consumer-oriented technologies beyond consumption and mass marketing applications.

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

That series of waves of ICT innovation, and the building web of mutual reinforcement, did not just contribute services directly to consumers. Firms who serve mass markets also took advantage of the new technologies. Many of the most valuable applications of all of the recent waves of ICT are as *marketing technologies* in consumer product and services firms, and to some degree in other mass markets. They have sales and service websites, mobile websites, and mobile apps. They advertise on new ICT-based media. A number of consumer-facing industries, notably the distribution and making of entertainment, retail, and retail finance, are going through dramatic structural changes as a result.

The consumer-facing products and services, and their relationship to mass marketing and mass markets, are familiar. Perhaps less well known, but also economically important, is the trend in ICT generally and in the applications of ICT more particularly toward an emphasis on marketing technologies. The original conceptualization of the value proposition of computers in business was "cost savings" through human work being replaced by computer work. The "cost savings" view of ICT application value creation has never disappeared, but it has been in steady decline as a portion of ICT applications for five decades. "Strategic" ICT applications – many in marketing and procurement – have been growing in its place.[43] Their growth accelerated with the PC and the widely used Internet. The factor market implications of "strategic" applications typically arise, of course, through systems-level substitution, not through task-level substitution.

In short, much of the early application of AIT in highly valuable uses, and much of the planned (versus conjectured) path of diffusion, continues the direction of technical progress that has been most rapid and sustained over the last twenty-five years: marketing, customer service, and their interaction in using customer service to create customer attention which is then monetized, whether in an advertising function or inside a standing buyer–seller relationship. AIT use restarts a cycle of improvement that has been deepening in particular areas of economic activity since the mid-1990s.

This is not the only thing that has been happening in technical progress, not even in ICT and its applications. Some of the previous rounds of new ICT in recent times have gone beyond consumption and mass marketing,

[43] The "cost savings" versus "strategic" language comes from Cortada (2004), reflecting an effort to quote or paraphrase businesspeople in many industries describing their goals and plans for technical progress.

including Cloud technologies, the Internet, and mobile telephony. AITs may join these earlier technologies in having a broader commercial impact in the future.

Any explanation of the great narrowing must begin with the underserved status of consumer computing and of mass markets before the widespread use of the Internet. It is now perfectly obvious that mass-market devices, communications services, online content, and e-commerce represent a substantial overlap of technological opportunity and market demand in our century. Similar applications, but not mass-market ones, were already present for e-commerce, online content, and electronic communication. Foresighted businesspeople saw the opportunity to extend these applications to mass markets.[44] As a result, much of the thinking about how to exploit the opportunity was in place before the widespread use of the Internet. The last twenty-five years or so have seen an explosion of mass-market, mass-marketing, and consumer-oriented ICT.

The extension of ICT production into mass marketing also drew on existing complements. Airline reservations systems, for example, were one of the oldest enterprise applications categories. They were complementary to direct passenger access through web interfaces and to the extension of those web interfaces to mobile devices. Existing corporate applications – the reservation system is the primary example – in consumer products and services firms were complementary to advances like the Internet and mobile devices, and together they enabled new rounds of rapid invention in consumer-facing marketing and customer services applications. That pattern of complementarity with existing capital supporting very rapid advance is, as has been discussed, a central feature of the use of AITs.

All of these factors explain why ICT deepening has been a powerful force. But why not broadening? I have identified two specific limitations of AITs – stakes and modularization.

Stakes help explain the subset of mass-market industries where AIT, like earlier rounds of ICT, has had big effects. It is more difficult to transform the marketing side of consumer-facing industries with high-stakes transactions. I am thinking here of health care, government, and many professional services. Modularity helps explain the slow rate of cost-reducing (versus "strategic," including marketing) advances. It has always been difficult to modularize many bureaucratic production processes, and this

[44] In many cases, these individuals did not just talk about the opportunity for mass market versions of all those services; they also launched large, mostly failed experiments before the browser (Bresnahan, 2012).

https://doi.org/10.1017/9781108856089.008 Published online by Cambridge University Press

has always been part of the slow rate of improvement of productivity through the application of ICTs in business. Production processes and the related markets and supply chains are difficult to modularize, so organizational adjustment costs have been large. Other applications of ICTs have required little organizational change, for example, workers seeing information on a cell phone they could already see on a PC, and thus have proceeded more rapidly. Modern ICTs have raised the benefits of organizational change, but the cross-section distribution of the costs of organizational change has been the key determinant of the diffusion of ICTs generally. AIT seems ill-posed to change this long-standing picture.

Changing organizations and supply chains is a slow economic process, even when motivated by large opportunity. International factor cost differences and new technologies each represent a large opportunity in our time. Exploiting those opportunities has been slow, since replacing labor either with combinations of capital and human capital or with overseas labor, transportation, and communications calls for modularizing production processes and supply chains. There is progress on that front, but not the kind of instant breakthroughs suggested by technological determinists. Diffusion of AITs beyond the applications seen thus far involves either breakthrough invention in applications (with modularization and organizational change) or the addition of important new complements, unknown today, to the AITs, or both.

One hope for AITs as a GPT is that production processes and supply chains will not need to be modularized – workers will simply be replaced by machines without any change in job description. In this case, no modularization would be needed. This hope is driven by a metaphor, not by business and technological reality.

Finally, I note that there has been a steady increase in the relative success of the leading firms in a large number of industries relative to other firms in the industries. This is seen in a dramatic increase in measured firm effects, in a wide number of metrics, including wages, profits, growth, capital share, and firm size. These are typically associated with firm use of ICT. We have also seen an increase in industry concentration, positively correlated with the measured margins for the larger firms in the industry.[45] The tendency of AIT to be scale-using and to be complementary with existing capital assets at the firm level suggests that AIT use will continue these

[45] The resulting increase in concentration may or may not reflect a decline in competition, and the origin of firm effects in pro-consumer investments in ICT applications may or may not mean that the changes are efficient.

trends rather than change them. The capital deepening arises out of a cluster of characteristics that arise at the firm level: scale economies, complementarity of new AIT capital with existing capital assets, particularly with big data and with modularized production processes, and the ability to specify a quantitative goal for production in an environment with low stakes in the case of error. These important economic and technical elements of AIT-using systems inventions are far from task-level substitution.

## References


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

Athey, Susan, and Glenn Ellison, 2011. "Position Auctions with Consumer Search." *Quarterly Journal of Economics* 126 (3), 1213–70.

Baldwin, Carliss Y., and Kim B. Clark, 2000. *Design Rules*. MIT Press, Cambridge, MA.

Bezos, Jeff, 2017. "2016 Letter to Shareholders." https://blog.aboutamazon.com/company-news/2016-letter-to-shareholders.

Brandt, Richard L., 2011. *One Click: Jeff Bezos and the Rise of Amazon.com*. Penguin Group, New York, NY.

Bresnahan, Timothy F., 2012. "Generality, Recombination, and Reuse." In Josh Lerner and Scott Stern (eds.), *The Rate and Direction of Inventive Activity Revisited*, 611–56. University of Chicago Press for the NBER, Chicago, IL.

Bresnahan, Timothy F., Jason P. Davis, and Pai-Ling Yin, 2015. "Economic Value Creation in Mobile Applications." In Adam Jaffe and Benjamin Jones (eds.), *The Changing Frontier: Rethinking Science and Innovation Policy*, 233–86. National Bureau of Economic Research, Cambridge, MA.

Bresnahan, Timothy, and Shane Greenstein, 1996. "Technical Progress and Co-Invention in Computing and in the Uses of Computers." *Brookings Papers on Economic Activity: Microeconomics* 27 (3), 1–83.

Bresnahan, Timothy F., and Manuel Trajtenberg, 1995. "General Purpose Technologies: Engines of Growth?" *Journal of Econometrics* 65 (1), 83–108.

Bresnahan, Timothy, and Pai-Ling Yin, 2010. "Reallocating Innovative Resources around Growth Bottlenecks." *Industrial and Corporate Change* 19 (5), 1589–1627.

Bresnahan, Timothy, and Pai-Ling Yin, 2017. "Adoption of New Information and Communications Technologies in the Workplace Today." In Shane Greenstein, Josh Lerner, and Scott Stern (eds.), *Innovation Policy and the Economy* 17 (February), 95–124. National Bureau of Economic Research, Cambridge, MA.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, 2017. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." NBER Working Paper No. 24001. National Bureau of Economic Research, Cambridge, MA.

Chui, Michael, James Manyika, Mehdi Miremadi, Nicolaus Henke, Rita Chung, Pieter Nel, and Sankalp Malhotra, 2018. "Notes from the AI Frontier: Applications and Value of Deep Learning." Discussion Paper. McKinsey Global Institute, McKinsey & Company, New York, NY.

Clark, Jack, 2015. "Google Turning Its Lucrative Web Search Over to AI Machines." *Bloomberg Business*, October 26, www.bloomberg.com/news/articles/2015-10-26/google-turning-its-lucrative-web-search-over-to-ai-machines.

Cockburn, Iain M., Rebecca Henderson, and Scott Stern, 2018. "The Impact of Artificial Intelligence on Innovation." NBER Working Paper No. 24449. National Bureau of Economic Research, Cambridge, MA.

Cortada, James, 2004. *The Digital Hand*. Volume 1. Oxford University Press, Oxford, UK.

Finnegan, Matthew, 2018. "Report: AI Assistants and Chatbots Gain Traction in the Enterprise." *Computerworld*, April 4, www.computerworld.com/article/3268075/ai-assistants-and-chatbots-gain-traction-in-business.html.

Gartner Research, 2018a. "Gartner Says Global Artificial Intelligence Business Value to Reach $1.2 Trillion in 2018: Customer Experience Represents the Majority of AI Business Value Through 2020," April 25, www.gartner.com/en/newsroom/press-releases/2018-04-25-gartner-says-global-artificial-intelligence-business-value-to-reach-1-point-2-trillion-in-2018#:~:text=Customer%20Experience%20Represents%20the%20Majority,reach%20%243.9%20trillion%20in%202022.

Gartner Research, 2018b. "Market Guide for Conversational Platforms." Gartner Research, Stamford, CT.

Goldfarb, Avi, 2014. "What Is Different about Online Advertising?" *Review of Industrial Organization* 44 (2), 115–29.

Helpman, Elhanan, and Manuel Trajtenberg, 1998. "Diffusion of General Purpose Technologies." In Elhanan Helpman (ed.), *General Purpose Technologies and Economic Growth*, 85–120. MIT Press, Cambridge, MA.

Hennessey, John L., and David A. Patterson, 2012. *Computer Architecture: A Quantitative Approach*. Elsevier, Cambridge, MA.

McCarthy, Caroline, 2006. "Facebook's Zuckerberg: 'We Really Messed This One Up'." *CNET*, September 8, www.cnet.com/news/facebooks-zuckerberg-we-really-messed-this-one-up/.

McGee, Matt, 2013. "EdgeRank Is Dead: Facebook's News Feed Algorithm Now Has Close to 100 K Weight Factors." *MarketingLand*, August 16, https://marketingland.com/edgerank-is-dead-facebooks-news-feed-algorithm-now-has-close-to-100k-weight-factors-55908.

Murray, Gerry, 2018. "IDC's 'Next Tech' Survey of Marketing and MarTech in Consumer Packaged Goods Manufacturers." IDC, Framingham, MA, www.idc.com/getdoc.jsp?containerId=US43613518.

Oremus, Will, 2018. "When Will Alexa Know Everything? A Conversation with the Chief of Amazon's Voice Assistant on What Alexa Hears, What She Remembers, and How She's Learning." *Slate*, April 6, https://slate.com/technology/2018/04/alexa-chief-al-lindsay-isnt-worried-about-users-privacy-concerns.html.

Rollings, Mike, 2018. "Deliver Artificial Intelligence Business Value: A Gartner Trend Insight Report." Gartner Research, Stamford, CT.

Romer, Paul M, 1986. "Increasing Returns and Long-Run Growth." *Journal of Political Economy* 94 (5), 1002–37.

Rosenberg, Nathan, 1997. “Uncertainty and Technical Change.” In Ralph Landau, Timothy Taylor, and Gavin Wright (eds.), *The Mosaic of Economic Growth*, 334–56. Stanford University Press, Redwood City, CA.

Rosenberg, Nathan, and Manuel Trajtenberg, 2004. “A General-Purpose Technology at Work: The Corliss Steam Engine in the Late-Nineteenth-Century United States.” *Journal of Economic History* 64 (1), 61–99.

Schubmehl, David, 2017. “IDC Market Glance: Cognitive/AI Systems, Search, and Content Analytics.” IDC, Framingham, MA.

Simon, Herbert, 1962. “The Architecture of Complexity.” *Proceedings of the American Philosophical Society* 106 (6), 467–82.

Taddy, Matt, 2018. “The Technological Elements of Artificial Intelligence.” NBER Working Paper No. 24301. National Bureau of Economic Research, Cambridge, MA.

Trajtenberg, Manuel, 2018. “AI as the Next GPT: A Political-Economy Perspective.” NBER Working Paper No. 24245. National Bureau of Economic Research, Cambridge, MA.

Turing, Alan M., 1950. “Computing Machinery and Intelligence.” *Mind* 59 (236), 433–60.

Varian, Hal, 2007. “Position Auctions.” *International Journal of Industrial Organization* 25 (6), 1163–78.

Walker, Mike J., Whit Andrews, and David W. Cearley, 2018. “Top 10 Strategic Technology Trends for 2018: AI Foundation.” Gartner Research, Stamford, CT.

Zuboff, Shoshana, 1988. *In the Age of the Smart Machine: The Future of Work and Power*. Basic Books, New York, NY.