Artificial Intelligence Technologies and Aggregate Growth Prospects

By Timothy Bresnahan*

Version IV, May 2019

Abstract

Artificial Intelligence Technologies are already in large scale use in two kinds of applications. We examine those applications, consumer/product matching engines and user interfaces, as well as the early diffusion of AITs and some categories where laboratory AIT and application are closely linked. In the leading applications, AITs are core production technologies with hundreds of billions of dollars in revenue. These production applications of AITs involve no substitution of machine for human work at the task level. The tasks undertaken by AIT in the leading applications were previously done by capital, not by human workers. Whatever your view of AITs, it should be updated to deal with this basic fact about their initial highly valuable applications. There is capital deepening, not factor substitution, at the task level.

The absence of task-level substitution is unsurprising to scholars of ICT-based production. The transition to ICT-based production has largely proceeded at the system level, not the task level. A modern example is the AIT-using Amazon store and mall, which as a system competes with brick and mortar stores and malls. That competition involves many changes which are not cost minimization, i.e., not the substitution capital for labor with fixed outputs. Instead, system-level substitution is driven as much or more by output characteristics – convenience, breadth of choice, effective recommendations – as by cost. In forecasting the long run factor demand implications of AITs, there is no good reason to see them as different from other ICT – working at the system level to slowly change over to a more capital intensive, less labor intensive, more human capital intensive form of production, but not a form of production whose primary new feature is different factor use.

Looking not only at the leading applications but at the early diffusion, AIT-based systems appear to be on the same track as most other highly innovative ICT-based systems of the last 25 years. They have spread out over a range of activities, to be sure, but have their largest impact in a narrow range of industries and functions, notably in consumer-oriented mass market production, distribution and marketing. The invention of consumer-oriented services like the WWW and consumer oriented devices like smartphones and tablets represents the diffusion of ICT to consumption. AITs, however, so far are participate in a very different 21st century trend, capital deepening in a narrow range of consumer-oriented markets and functions with divergence across firms in the rate of technical progress. Examining the strengths and weaknesses of AIT-based production illuminates the narrowness and depth of ICT application more generally.

---

* Department of Economics, Stanford University. I am grateful to many colleagues and students for teaching me about ICT technology and applications. I want to acknowledge particular debts to Shane Greenstein and Ken Hendricks for penetrating comments on an earlier draft and Bill Janeway and David Yoffie for deep discussions. The errors are all mine – if you spot one, let me know! tbres@stanford.edu

6.1) Introduction

This chapter examines the commercial application of Artificial Intelligence Technologies (AITs), seeking to address questions about these technologies specifically and about 21st century technical progress. I focus on the highly valuable *applications* of AITs today, in production systems at the Internet Giants, in new User Interfaces, and elsewhere. My empirical conclusion about these applications is that the lazy idea of "Artificial Intelligence" -- i.e. 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 AITs' value in use. 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 ICT-based production, 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 is different than 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 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 detailed sections below, we examine the complementarity or AITs with specific aspects of existing capital-intensive production processes that explain the tendency toward capital deepening. Capital-capital complementarities and scale economies at the firm level are an important element. We also examine the aspects of new capital-intensive production processes, with and without AITs, that have largely limited their 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 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

---

[1] Many recent economics 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) and Agrawal, Gans, and Goldfarb (2018), the latter a terrific overview.

[2] GPT analytics emphasize the innovational complementarities between the GPT itself and inventions of applications (see Bresnahan and Trajtenberg (1995), Rosenberg and Trajtenberg (2004), Helpman and Trajtenberg (1998).) Shane Greenstein and I (1996) emphasize the role of difficult-to-invent applications 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).

are like the other big 21st 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, and 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 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 explains most writing about AITs. This metaphor underlies 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 actually exist, like all software technologies, are designed into systems. In this paper, we look at those systems. Spoiler alert: Do not hope for a lot of sci-fi in these applications.

---

[3] Artificial Intelligence is "The 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." Oxford Living Dictionary. 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.

In this paper, we look at the applications. It is not too early, 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 user interface (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, we will examine other growing AIT applications such as driver assist, the rebirth of expert systems, and improved decision support.

We 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 AI Technology, 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 AI Technology and of ICT more generally.

That leads us to a discussion of where ICT application has been going in the 21st century, and to an attempt to understand it. The easy to understand thing: there has been a remarkable series of waves, beginning with the web browser, of new technologies that serve consumers directly and which enable mass-market and mass-marketing commercial applications. Another easy to understand thing: 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: why the very impressive technical progress in those areas has had limited, i.e. some but only limited, impact on applications in the rest of the economy, where ICTs went first. We'll finish on this growth question.

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

6.2) Product/Consumer Matching Applications

We 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 invention of new and better ways to match product to consumer. I look at them together because they have a common role for 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 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 21st 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

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

recommendation engines has been a central part of one of the most valuable technical advances of the 21st century. A good place to start to understand 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 she was going to buy next.[5] 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 we shall return.

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 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; for the new products of outside sellers they rely on the seller for information about category, etc. For products that have been for sale for a while, Amazon knows who bought them and who did not.

For many years, algorithms made recommendations that, for example, suggest further products based on what a consumer has chosen to search for and what she has bought in the past.

[5] See Brandt (2011), p. 67.

More recently, machine learning and AITs drive the algorithms Amazon uses in these areas.[6] 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 (find what you want) and from an Amazon and merchant perspective (get matched to customers), consistent with 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 pre-existing recommendation system that was part of its well-functioning and *modularized* online store/mall. Making that recommendation system more targeted to the individual customer, 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.[7] Having a scalable production

---

[6] The language I use here draws heavily on Bezos (2017), which 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, what an Amazon app or a Kindle show a consumer, are managed by AITs.

[7] Modularity helps scalability of a system with multiple moving parts in a number of ways. Modularity permits adding processing power and storage where needed as loads change. Similarly, modularity permits adding new data streams (e.g. adding information from Gmail) or analytical elements (e.g. thwarting ill-behaved SEO at Google) as systems change. Dynamically, modularity permits improving the system architecture to 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 sub-chapter of Chapter 3 on "Dimensionless, scalable design rules".

system has, in turn, let the firm gain scale economies. To understand this, we need to examine the Economics concept of "economies of scale" next to the Computer Science concept of "scalability".

"Scale economies" mean that marginal cost is lower than average cost, and typically we mean this over a wide range. "Scalability" means that a system has been designed so that its workload can be increased without changes to its architecture (yes, there is (Simon 1962) again). Again, this is typically understood to cover a wide range. Scalability means more than scale economies – when workloads are uncertain, for example, a scalable production process has 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, participated in this modularized online store/mall architecture. Changing to AIT recommendation systems preserved the modularity and did not break the scalability. The firm already had good estimates of the costs and benefits of algorithmic recommendations, and could re-use 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 of sales.[8] The resulting low

---

Hennessy and Patterson (2017) make a similar observation about software when they write "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."

[8] 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

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, lead 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. Marginal cost would rise dramatically if human activity were required as part of each transaction. Fixed costs 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.

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 only be economic for large firms. Whether "large" means as many product recommendations as occur at Amazon or significantly smaller is a topic to which we return.

6.2.1.1) Stakes

Finally, the product recommendations made by Amazon are just that – recommendations. The consumer ultimately decides what to buy.[9] 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 has no broader negative consequences.

---

the firm applies AITs. One example is inventory prediction in the warehouses, already a statistical prediction problem before application of AITs.

[9] Agrawal, Gans, and Goldfarb (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.

The implications for the economics of adding AIT to the recommendation algorithm are straightforward. The fixed costs of switching to an AIT based production can be spread over the large volume of sales. AIT is based on statistical prediction, so AIT based matching systems have rates of:

true positives	suggestions made that lead to sales
false positives	suggestions made that do not lead to a sale
false negatives	suggestions that would have led to a sale but were not made
true negatives	unmade suggestions that would not have led to a sale.

The benefits to Amazon of even a small increase in the overall rate of true positives are substantial at Amazon's scale. So too the benefits of even a small reduction in true negatives. A sale results; much of the costs of being Amazon are 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. What about the error cost of a false positive? These are small, a false positive is just a recommendation not taken up by the consumer. The problem has a *low stakes lost function.*

The basis in enormous data, the use of AT Technologies to achieve large scale at low MC, the modular system before AI 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, they provide much of an explanation of the firm'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 we have just seen, though the difference between a product recommendation and an advertisement means 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. [10] I will review it only briefly, with emphasis on the parts that draw on AIT to match.

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.[11] When a consumer searches, two kinds of results are returned. The "organic" search results are Google's guess at what the consumer was looking for. There are also advertisements, i.e., information that is looking for 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

---

[10] See, e.g. Varian (2007) and Athey and Ellison (2011).

[11] 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.

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.[12]

Google's system, loosely called "quality score," has elements that are calculated in real time for each advertisement in each auction.[13] As the relevant economic theory makes clear, a central part of Google's profit maximization problem is to predict the probability that *this* consumer views *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. It is in this advertising ranking for a particular consumer search that the AIT comes in.

Predicting the probability of a click on an ad is a near-ideal for use of predictive AIT with machine learning. The situation is complex. The probability is specific to a given advertisement, advertiser, searcher, the search terms, the device the searcher is using and the time of day, and interactions among all those things, e.g. relevance of the advertisement to the search. Google has big data on all those things. 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.

---

[12] The interaction between this ranking and the incentive-compatible elements 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 only get clicked on very rarely -- but who make a large profit if clicked on -- will bid their way to the top. Great for them; bad for Google; likely bad for searching consumer.

[13] The elements are click through rate, the "relevance" of the ad to the users' search, and the "landing page experience" if the user clicks on the ad and goes to the advertisers' website. "Quality score" refers both to a number and to the system that generates the number, and which can change the number in real time without telling anyone. 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 lessons the effectiveness of the gaming.

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. Here, "prediction" means exactly what it meant when you took basic statistics. The "deep learning" part of all this 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 <u>why</u> a group of data elements are good predictors of a match, only <u>that</u> 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 (here, probability of clicking on the ad and/or revenue). Deep learning algorithms work well with many, complex data elements that might be used to predict <u>if</u> they have large sample sized to work with. 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 here 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.[14] The economics of this early highly valuable application of AIT

---

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

also follow the 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.[15] All that the AIT for matching needs to do to cover its (high) fixed cost of invention is increase the probability of a successful match a small amount.

6.2.3) Facebook

Facebook is also ad-supported, and also uses AIT to match particular advertisements to particular consumers.[16] The business logic of that subsystem is like those we have 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.[17]

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 it uses to decide what information, including ads, to show to users, in contrast to Amazon's and Google's observations

---

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

[16] Mark Zuckerberg on Facebook 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".

[17] 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, machine translation. Facebook has a number of efforts to detect problems such as suicidal users, posts not from real users, etc., an area we shall revisit below.

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 explicit an explicit product search by the user. This was long algorithmic and is now based on AIT. Facebook now uses AIT in the long-algorithmic function of deciding what non-advertising information to show, i.e., in populating users' "news feed."

The algorithm used to 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 we 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.[18] 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 as above apply. Facebook has scale, needs scalability (low MC through automation) has very "big data," had, before widespread use of AITs, of necessity a modularized production system. Matching an ad to a user is a difficult prediction for an algorithm but the vast amount of

---

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

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.

It is when we go beyond commercial advertisements and the narrow limits of reading posts from ones' own friends that we can learn a few of the important limitations to the use of AI Matching Technologies which 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 posers, 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, apologize, and repair.[19] Typically these experiments cannot be contained at a sample of users, for users interact and overlap. This makes Facebook a great place to look for 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 place where Facebook has experienced limits to the use of AIT is flagging "inappropriate content."[20] Why is this? Partly, it is because some "inappropriate content" involves higher stakes when a user is shown either a post or an advertisement they do not like. Mr. Guy Rosen of Facebook made the table I post just below, and it is clear that a number of the

---

[19] 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.

[20] Another area us algorithms for suggesting "people you may know" which sometimes suggest people you really don't want to know or, perhaps, to 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.

categories are ones where the mere fact that the content was shown to the user is seen as a significant negative by the user.[21]

This problem has been made worse by the complexity of the "social graph" and by the creation of communities to 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, e.g. in Reddit, to plan coordinated assaults. These communities create problems by 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.

Why all this human effort? Guy Rosen, Facebook VP of Product Management explains[22]:

> *"...we 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."*

Mr. Rosen cites three problems for AITs. The first two are about statistical power to discriminate between problematic and regular 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. (One thinks of a new example of hate speech, early

---

[21] We saw this with the reaction to the Cambridge Analytica scandal. People, and governments, felt that Facebook had crossed important limits. My point here 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 else. Instead, what is relevant to the enquiry about value creation from use of AITs is that the scandals reveal difficulties in algorithmically policing political posts because the reader can be offended by seeing one.

[22] In a blog post at https://newsroom.fb.com/news/2018/05/enforcement-numbers/ The Facebook people discussing this, 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 CS people.

in its ugly life.) We can easily understand these first two problems for AITs with analysis of a simple table based on Mr. Rosen's post:

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"

Success in AIT fighting spam (the first line) 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, has come to a high 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 the AIT's rate of winning the human/machine race. This is easy to explain as high-loss-function hate speech poses severe problems for AITs. Human complaints about hate speech are a rapid process in the 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 difficult 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 throw people power at hate speech.

"Fake news," while not cataloged in Mr. Rosen's list, also is getting a large jolt of people power at Facebook.[23] 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.[24] The loss function has high stakes.

Mr. Rosen flags a third problem, (3) "we're up against sophisticated adversaries who continually change tactics to circumvent our controls, which means we must continuously build and adapt our efforts." This is an old market regulation problem, long familiar to economists (though we have thought about it more in terms of public policy than business policy.) AITs based on machine learning are about statistical prediction – not about 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 predition in much the same way it applies, familiarly, to changes in public policy. The policy change might cause changes in behavior that break the prediction. Facebook discovered this the hard way. When Facebook introduces an AIT-based set of "controls" designed to block certain bad behaviors,

---

[23] https://www.cnet.com/news/can-facebook-mark-zuckerberg-new-hires-take-on-troll-farms-and-data-privacy-after-cambridge-analytica/

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

their “sophisticated adversaries” then “change tactics” –breaking the statistical prediction model. This is a limitation of the range of AIT’s usefulness. The Facebook problems are using AITs in policies to avoid high-stakes loss function errors where the problematic behavior will change in response to policy changes. All of hate speech, trolling or made-up news stories of high emotional range fall outside this boundary. Because the problematic actors change behavior in response to a policy change, there will always be periods in 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. This is one of the current boundaries of application – and consistent with the initial large-scale production uses of AIT recommendation systems being 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 show would this user like to consider next? The number of choices Netflix can present is limited not only by customers’ attention, but by the clumsiness of the TV screen and TV remote. Netflix’ business model also rewards it heavily for matching customers to content they like; it is particularly profitable to match users to content that falls outside the list of currently most popular movies. Finally, Netflix has impressive big data on users’ past choices.

[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.

Netflix long used an algorithmic 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 publically about "big data" at this firm than almost anywhere else. 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 are liked by some viewers and disliked by others. (The example they discussed publically 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," i.e. 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 UIs have been introduced – voice search for TV programming at cable companies is an example. A number of entertainment-deliver 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, but the voice UI permits the user merely to say the search. These examples bring together the virtues of

---

[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.

the Netflix example we saw above (limited range of vocabulary) and the virtues of the UI applications on cellphones (high value of UI improvements in difficult environment for typing.)

One could list more Internet firms and related prediction applications, but I hope the point is made. Matching, at large scale, something like a product or an advertisement to a specific consumer in specific circumstances, hoping the consumer will choose it, using an engine that predicts whether the choice will be made and/or what the value to the firm will be if the choice is made.[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 "artificial intelligence" – not as sci-fi as an anthropomorphized robot.

Exactly. The AITs here 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, the AITs used here are like earlier ICTs. They combine technical progress, tools for invention of applications, and technical progress embedded in capital.

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. The zero is not a feature of AITs. But to understand the scope, direction, speed, etc. of these technologies, there is no role for "tasks usually done by humans" at all. TLS is irrelevant.

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

[27] One things also of LinkedIn – jobs you may be interested in, and of Waze, Google maps, best route, local vendors, or of Google image search for similar images, or of TaskRabbit matching rabbit to task.

brick and mortar supply chain including 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 has multiple determinants, of which static cost saving, i.e., 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., in 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 TLS.

6.2.6) Matching Engines like Earlier ICT Waves

The AIT-based matching engines we have seen 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 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, particularly since the great turn to service 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 21st century. Facebook, as we have seen, 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 in just a firm.

6.3) UI Improvements Based on AITs

A second important area of AIT application is in user interface (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"[29] 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, 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 the users. It all adds up to a remarkable technology deployment that is another very widespread and valuable use of AITs.

Voice UIs tend to make casual observers think of the Turing test. If the user "can't tell" that she is talking to a machine, but rather thinks she is talking to a human, then the machine has

---

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

[29] D'Onfro (2016) reports an interview with Jonathan Jarvis after he left Google, reporting on why the firm decided not to give the UI a personality. "We always wanted to make it feel like you were the agent, and it was more like a superpower that you had and a tool that you used."

achieved “artificial intelligence.”[30] This makes an elementary mistake about what a user interface 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. Here is a clarification from a technologist, Al Lindsay the manager of Amazon’s Alexa service (Oremus (2018)).

> *“... when 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 only make economic sense with easier access. In this sense UIs are GPTs -- but this is not the sense observers have in mind when they hope “AI becomes a GPT.”

It is easiest to understand the impact of the new AIT-base 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 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 WWW 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 web sites were enabled by the access improvements of the PC/browser combination.

---

[30] Turing (1950). 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.

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 touch-screen 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 the technological and business reality. I again quote Mr Lindsay of Amazon (from the same 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*

Exactly. Mr Lindsay works for Amazon, so he hopes what it improves is shopping, but the general point is that UI improvements "remove friction" to increase access to applications.

The rate of technical progress in the AITs called natural language processing (NLP) increased when they got a statistical-prediction basis. They are like the matching engines we examined in the previous section but are distinct technologies. The value creation we are seeing with 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.[31] Google also exposes an 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, e.g. Siri "suggests" opening a particular app, but foolish, as this is precisely the same as the Google website "suggesting" a particular ad. Using devices in bad-for-typing and bad-for-reading environments like cars and kitchens raises the value of voice UI, 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 when there were already widely distributed devices that could benefit from them, and those same devices were telephones – terrific locations to gather "big data" on voice. Scale and competition played a role with Apple, Google, Amazon and others in the mix at very large scale.

The improvements to smartphone voice UIs are made more economic by the very large scale deployment of those devices, their use by UI-sensitive demanders (consumers) and the

[31] 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. Made an oddly worded version of a common query? It guesses the underlying query. At the public announcement of RankBrain, Google was quoted as saying it was the third most frequently used algorithm. See Clark (2015) These NLP technologies, characterized in the Google-sphere as the switch "from strings to things," are valuable both to the searcher and to advertisers. They help searchers find what they are looking for, and they help communicate to advertisers what the searcher might truly be looking for.

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 between the Internet Giants plays 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.[32] 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 an home-device company with large scale, and drawing competitive responses from other internet giants.[33]

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.[34] Echo devices running Alexa have only the voice interface, not 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 an extension of the 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 unsurprising that software clients for both Kindle

---

[32] 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 this 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 vs Amazon Prime music, e.g.) swirl in the vast Apple rumor mill.

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

[34] 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.

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.

Despite their difference, I put Kindle and Alexa in parallel here 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 the smartphone voice UIs, 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.[35] Siri and Google assistant & c. 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 – capital deepening in marketing technologies 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

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

forces have not yet pulled AITs, as they did not pull earlier rounds of impressive technical progress like the smartphone, far beyond that range of consumer and mass market applications.

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 (new application built on platform) or in a broader sense (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.[36] 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 voice UIs 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 are 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.

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

No task level substitution to here! These UIs substitute capital for capital to increase convenience and access. New software technology, based on NLP (AIT) replaces old as new user interfaces replace earlier user interfaces in some uses. They also expand the use of devices to new activities. Voice user interfaces partially replace touch-screen UIs, WIMP 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 clear 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 on consumer time could become a related sense of saving on 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, broadening access is.

6.4) Broader Diffusion – Other Marketing Applications

The early successful applications have created interest among firms other than the Internet Giants. The interest extends to 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,"[37] so there are many experiments.

[37] 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 roll over the CFO because of the buzz around a technical area.

Programmer toolkits are available for many AITs.[38] 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 very little application of the AITs outside the Internet Giants at this (Spring 2018) stage. However, over the last 2 years, large firms' approaches to AIT applications has crossed from speculation and investigation to experimentation and development.[39] Surveys of firms about their applications intentions now mention specific use cases. There are applications plans, applications experiments, and ideas.[40]

The largest category of experiments is *marketing* applications. Using the AITs underlying product/consumer matching systems and NLP, "chatbots" and the like are applied in marketing interactions with customers; not just initial customer acquisition (advertising) but 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).[41] Surveys of firms about their AIT application plans

---

[38] Through AWS, Amazon is an important supplier of services for Web, Cloud, etc., applications, 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.

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

[40] A Gartner survey of chief information officers is representative (Gartner (2018)): "Only 4% 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."

[41] The areas include "customer engagement," "customer satisfaction," "customer support" "customer experience," and others. See, e.g., Murray (2018), which looks at consumer package goods manufacturers and reports that CX is the top priority for marketing technology spending. Of the technology areas that might be used to improve ICT-based marketing, "Artificial Intelligence is still emergent as a purchase driver," meaning just under

typically show this as the largest area of planned applications growth. This is the main direction of diffusion in the present.[42]

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, pre-existing 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 improving ICT systems and advanced them along their existing path.

While 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 to substitute 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 <u>not</u> an outbreak of cost reduction via TLS. Instead, the relevant demand-side forces are competition to improve consumer experiences and pre-existing 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 a VRU. Other parts of that knowledge comes from ineffective FAQs – text chatbots are a more targeted (AIT doing the targeting) FAQ page.

6.4.1) Diffusion of Matching Engines to Advise Employee, not Customer

Another area of potentially important application experimentation, the second largest in surveys, deploys AIT to advise an employee. 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

---

half of firms anticipate it will influence their "marketing technology" purchases going forward, while about ¾ anticipate that CX will influence those purchases. Other product areas (in similar reports not cited here) are similar.

[42] See, e.g., the IDC "use case" survey reported above. Gartner (Rollings 2018) is even more direct: Gartner: "Customer Experience Represents the Majority of AI Business Value Through 2020"

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.

Use of Alexa, Siri, or Google Assistant to undertake chores at work, much like the way they are undertaken 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 lossy outcomes. One early focus seems to be on simple chores, e.g. Alexa skills for sending an email, setting up a meeting, etc. The factor market implications appear to be for a modest increase in individual worker productivity.

Actual organizational productivity improvements following from this kind of use are harder to forecast enthusiastically. As did 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, lead a number of observers to forecast growth for AIT-based Cortana in this role (e.g. Finnegan (2018)). The software predicts what the user wants to do next and suggests that more prominently – low loss function, good use of matching technology, increases software user productivity.

A third area of experimentation is improvement of decision support 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 take it past 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.

All these fronts are areas of potential progress in AIT use that break out of the marketing function. They go to other loci where existing assets can be improved by AIT interfaces, at the boundary between ICT systems and human users. Their labor market implications are approximately the opposite of TLS: the human worker continues to do the 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 TLS technology – replace that driver with an "autonomous" vehicle. The very important inventions from this effort are rapidly moving into use as driver assist technologies, i.e., like a decision-support system or an advisory product choice system, not like TLS. Many observers made the – vapid – remark that a "driverless car" "proves" that "computers can do anything humans used to do." The vapor lies in the interpretation about proof and about TLS, not in the engineering. Commercial application of driverless trucks awaits large-scale modularization of truck drivers' jobs, which typically involve much more than driving. Market application of the "driverless" technologies as driver-assist safety features in high-end automobiles are growing rapidly. My bet is still open: Which of two events will occur faster?

1. Half of the driverless vehicles in actual use will be AIT based autonomous vehicles.[43]

[43] Factories have been using driverless vehicles for decades; the latest AI Technology based variants need technical progress to permit the driverless vehicles leave a controlled environment and go on the road.

2. AITs are deployed in over 100 million vehicles on the road to offer driver assist, not driver replacement.

I'll take event 2 at even money from up to 100 readers at up to $10 each. I'm an old guy, but event 2 is so close to occurring that I'll easily live to collect.[44]

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 involve 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 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.

6.5) Laboratory Results Close to Use Case

A third category of early applications is ones in which laboratory results are close to commercial use. These are not diffusion of production process inventions such as we have seen above. They are, instead, a series of parallel tracks. Most of these do not call for modularizing organizations or production processes.

---

[44] Today, lane assist, crash warnings, semi-automatic braking, and other AI Technology based improvements are largely marketed as high-end car safety features, but they are spreading rapidly into a wide range of vehicles.

Production scheduling, inventory management, shipment scheduling and related tasks 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 in AIT prediction models.

A more complex example can be seen in credit card fraud systems. These are statistical, but for years 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. Here 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 catching more fraud – need not increase the level of losses.

Science and engineering 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. See Cockburn, Henderson, and Stern (2018) who are particularly interested in such important commercial scientific areas as drug discovery and development. They offer an interesting analysis of re-organization 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 (of the eye), etc. We can expect some of this to (slowly, organizations change slowly) change many security systems outside computer systems access – for example, in air travel. Once again, we can see a path toward more general use with potentially higher 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, inter alia, are subject to security threats. Improvements in security systems permit ICT production systems to be delivered more conveniently, e.g. through "cloud" or mobile channels. Improvements in ICT security systems remove an externality, fraud or theft, which 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 vs older production systems, and those of ICT-based production of new goods and services.

Photo and voice recognition applications are diffusing to other kinds of non-security applications as well. Other applications are being developed. 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 spread out takes on some, but not a great deal, of adjustment costs. Likely the uses in which adjustment cost boundaries are being most tested and pushed are the security examples.

6.6) High Rate, Same Direction

AITs are 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 21st-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, wifi) [45] and in mass-market marketing and sales applications. I do not mean to imply that this is all of ICT application in our era, but rather that it is the subset which is creating large private returns to invention and to application, and to which technologists point when they claim we are in an era of rapid technical change. Over and over in this paper we have seen that the application of AITs follows these developments. What are these developments? Why did they occur? Why do AITs reinforce them rather than change their direction?

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

Complementary to the new devices, we have online mass market products and services, first based on the widely-used Internet and the WWW, later on "Web 2.0." These enabled many consumer-facing and other mass market services online. Another set of complements arose because the new consumer devices were largely mobile phones and tablets. This enabled mass market mobile web applications, then mobile apps. Complementarities among these technologies make them mutually reinforcing, as, for example, consumer-oriented web sites became mobile web sites 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,

---

[45] See Bresnahan and Yin (2010) for a discussion of the causes of ICT's turn to consumption and of the limited spread of consumer-oriented technologies beyond consumption and mass marketing applications.

particularly AIT matching services, are now being recombined with these existing assets into that web of mutual reinforcement. Again, there is substantial deepening of an existing trend in the early AIT era.

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 to us all. Perhaps less well known, but also economically important, is the trend in ICT generally and in the applications of ICT toward an emphasis on marketing technologies. Business people's original conceptualization of the value proposition of computers 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 steadily.[46] Their growth accelerated with the PC and the widely used Internet. The factor market implications of "strategic" applications typically arises, of course, through systems-level substitution not though TLS.

---

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

In short, much of the early application of AIT in highly valuable uses, and much of the planned (vs conjectured) path of diffusion, continues the direction of technical progress that has been most rapid and sustained over the last 25 years. Marketing, customer service, and their interaction in using customer service to create customer attention with is then monetized, whether in an advertising function or inside a standing buyer/seller relationship. AIT use re-starts a cycle of improvement which has been deepening in particular areas of economic activity since the mind 1990s.

This is not the only thing that has been happening in technical progress, not even in ICT and its applications. And, just as we should expect AITs to do, some of the previous rounds of new ICT in recent times have spread out beyond consumption and mass marketing. (Cloud is not only or even mostly client-facing; Internet improved many things; Mobile is somewhat better for non-marketing business after the Blackberry standard was replaced, etc.) With most of the invention of applications for AITs and (potentially) of new large complements in the future, it is possible that that "some spread" will go farther in the future. In the recent past, however, it mostly hasn't.

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. This technical and market opportunity was also obvious to technically oriented business people before the browser. Those business people saw ICT in organizations, ICT underlying non-mass-market e-commerce, non-mass-market online content and communications and saw the opportunity.[47] 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 25 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, to the extension of those web interfaces to mobile devices. The existing corporate applications – the reservation system was the example – in consumer products and services firms were complementary to advances like the Internet and mobile devices, and together the 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 we have seen, a central feature of the use of AITs.

All of these explain why ICT deepening has been a powerful force. But why not broadening? We 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 help explain the slow rate of cost-reducing (vs "strategic," including marketing, advances. It has always been difficult to modularize many bureaucratic production processes, and this has always been part of the slow rate of improvement of productivity through the application of ICTs in business. Production processes and the related

[47] And, in many cases, didn't just talk about the opportunity for mass market versions of all those services, but also launched large, mostly failed, experiments before the browser. See Bresnahan (2015).

markets and supply chains are difficult to modularize, so organizational adjustment costs have been large. Other applications of ICTs have required little organization change, e.g. workers seeing information on a cell phone they already could see on a PC, and thus have gone faster. Modern ICTs have raised the benefits of organization 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 longstanding picture.

Changing organizations and supply chains is a slow economic process, even when motivated by large opportunity. International factor cost differences and new technology each represent a large opportunity in our time. Exploiting those opportunities has been slow, since either replacing labor 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 the Technological Determinists. Spreading out of AITs beyond the applications seen thus far involves either breakthrough invention in applications (with modularization and organizational change) or addition of important new complements to the AITs, complements not known today, 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. So no modularization 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's 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.[48] The tendency of AIT to be scale using and to be complementary with existing capital assets at the firm level suggest that AIT use will continue these 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 invention are far from task-level substitution.

---

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

6.7) Bibliography

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018. 'Prediction machines : the simple economics of artificial intelligence'.

Athey, Susan, and Glenn Ellison. 2011. 'Position Auctions with Consumer Search', *The Quarterly Journal of Economics*, 126: 1213-70.

Baldwin, Carliss Y., and Kim B. Clark. 2000. *Design Rules* (MIT Press: Cambridge, Mass.).

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

Brandt, Richard L. 2011. *One click : Jeff Bezos and the rise of Amazon.com* (Portfolio/Penguin: New York).

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* (National Bureau of Economic Research, Inc).

Bresnahan, Timothy F., and M. Trajtenberg. 1995. 'General purpose technologies: Engines of growth'?', *Journal of Econometrics*, 65: 83.

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*, 1996: 1-83.

Bresnahan, Timothy, and Pai-Ling Yin. 2010. 'Reallocating innovative resources around growth bottlenecks', *Industrial and Corporate Change*, 19: 1589-627.

———. 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*.

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

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', *Mckinsey Global Institute*.

Clark, Jack. 2015. 'Google Turning Its Lucrative Web Search Over to AI Machines', *Bloomberg Business,* October 16, 2015.

Cockburn, Iain M., Rebecca Henderson, and Scott Stern. 2018. 'The Impact of Artificial Intelligence on Innovation', *National Bureau of Economic Research Working Paper Series*, No. 24449.

D'Onfro, Jillian. 2016. 'Why Google's smart assistant doesn't have a name like Siri, Alexa, or Cortana', *Business Insider,* May 19, 2016.

Finnegan, Matthew. 2018. 'Report: AI assistants and chatbots gain traction in the enterprise. Microsoft's Cortana AI assistant currently leads the charge in the workplace, thanks to Windows 10 integration.', *Computerworld,* APR 4, 2018.

Gartner. 2018. 'Market Guide for Conversational Platforms', *Gartner*.

Goldfarb, Avi. 2014. 'What is Different About Online Advertising?', *Review of Industrial Organization*, 44: 115.

Helpman, Elhanan , and Manuel Trajtenberg. 1998. 'Diffusion of General Purpose Technologies.' in Elhanan Helpman (ed.), *General Purpose Technologies and Economic Growth* (MIT Press: Cambridge, Mass.).

Hennessy, J.L., and D.A. Patterson. 2017. *Computer Architecture: A Quantitative Approach* (Elsevier Science).
McCarthy, Caroline 2006. 'Facebook's Zuckerberg: "We really messed this one up", *C/NET,* September 8, 2006.
McGee, Matt 2013. 'EdgeRank Is Dead: Facebook's News Feed Algorithm Now Has Close To 100K Weight Factors', *MarketingLand,* August 16, 2013.
Murray, Gerry 2018. 'IDC's 'Next Tech' Survey of Marketing and MarTech in Consumer Packaged Goods Manufacturers', *IDC*.
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 (Interview with Al Lindsay)', *Slate,* April 06, 2018
Rollings, Mike 2018. 'Deliver Artificial Intelligence Business Value: A Gartner TrendInsight Report', *Gartner*.
Romer, Paul M. 1986. 'Increasing Returns and Long-run Growth', *Journal of Political Economy*, 94: 1002-37.
Rosenberg, Nathan 1997. 'Uncertainty and Technical Change.' in Ralph Landau, Timothy Taylor and Gavin Wright (eds.), *The Mosaic of Economic Growth* (Stanford University Press).
Rosenberg, Nathan, and Manuel Trajtenberg. 2004. 'A General-Purpose Technology at Work: The Corliss Steam Engine in the Late-Nineteenth-Century United States', 64: 61.
Schubmehl, David. 2017. 'IDC Market Glance: Cognitive/AI Systems, Search, and Content Analytics', *IDC*.
Simon, Herbert. 1962. *The Architecture of Complexity*.
Taddy, Matt. 2018. 'The Technological Elements of Artificial Intelligence', *National Bureau of Economic Research Working Paper Series*, No. 24301.
Trajtenberg, Manuel. 2018. 'AI as the next GPT: a Political-Economy Perspective', *National Bureau of Economic Research Working Paper Series*, No. 24245.
Turing, Alan M. 1950. 'Computing Machinery and Intelligence', *Mind*, 59: 433-60.
Varian, Hal. 2007. 'Position auctions', *International Journal of Industrial Organization*, 25: 1163-78.
Walker, Mike J., Whit Andrews, and David W. Cearley. 2018. 'Top 10 Strategic Technology Trends for 2018: AI Foundation', *Gartner*.
Zuboff, Shoshana. 1988. *In the age of the smart machine : the future of work and power* (Basic Books: New York).