Received: 15 May 2022 | Revised: 9 December 2022 | Accepted: 25 January 2023

DOI: 10.1111/jems.12524

SPECIAL ISSUE ARTICLE

Journal of Economics & Management Strategy WILEY

What innovation paths for AI to become a GPT?

Timothy Bresnahan

Department of Economics, Stanford University, Stanford, California, USA

Correspondence
Timothy Bresnahan, Department of Economics, Stanford University, 755 Serra St. Stanford, Stanford, CA, USA.
Email: tbres@stanford.edu

Abstract

Early commercial applications of artificial intelligence technologies (AITs) were narrow but extremely profitable. Comparable uses of those technologies throughout the economy would lead to a growth boom. Firms which emulated the early applications successfully would make tremendous strategic gains. This is a situation familiar from earlier rounds of information and communication technology. However, for AITs to become a general-purpose technology across many commercial applications sectors will require some new innovations. This paper examines the innovation paths that could lead to that desirable outcome, the ones that have stalled, the ones in the process now, and the ones that might occur in the future. Strikingly, early AIT use, both commercial and with technical customers, occurred where Digital Transformation was not required for it to succeed. The innovation paths all require Digital Transformation as key steps.

1 | INTRODUCTION

Artificial intelligence technologies (AITs) may become general-purpose technologies (GPTs) even though the important early applications are, while deep, largely confined to a narrow range.[1] This paper takes up the question of how the narrow range of early applications could become general and trigger new growth.

The question is important for both strategy and economics. Early (largely 2015–2017) commercial applications of modern AITs were very valuable; these applications run the core production processes of firms with sales in the hundreds of billions of dollars. Other firms considered emulating these leaders. Success would mean remarkable gains for the emulators. Broad success would accelerate enterprise Digital Transformation. From an economy-wide perspective, if applications of AITs like those early exemplary ones were to diffuse to a broad fraction of the US economy, it would create dramatic productivity growth. These "importance" observations lead me to a specific version of the question. What paths could lead to a broad, general use of AITs as significant technologies including in commercial applications?

GPT loops—positive feedback loops through GPT improvements and Application Sector (AS) coinvention—can create substantial value. The process is not automatic. Inventions in both GPT and AS start the loop, expand its scope, and build momentum. While that invention is ignored in recipe approaches to the creation of GPT loops,[2] it is central in the world and in this paper. The task of this paper will be to examine each of the plausible inventive paths, including some that have already been attempted, and to state with particularity what invention would be required for them to succeed. This part of the exercise can be made without speculation. As a last step in the discussion of each potential path, I compare it to what seems to be the correct historical analog to comment on its likely scope and timing. This last step involves forecasting the future, of course, but in a disciplined way.

Stanford University and NBER. I have greatly benefited from comments from the participants in the conference on The Business Revolution of Digital Transformation at USC.

The initial conditions for AITs in the midteens seemed very propitious; substantial breakthroughs in AITs themselves and some very impressive applications. Further, there was no shortage of discussion in business communities of the prospect for much wider use of these technologies. While technical progress in AITs and within the narrow ranges of applications continued, efforts to start a positive feedback loop have been far slower.

A striking feature of all the early applications of AITs is that they have succeeded where Digital Transformation was not required. The two big blockages to Digital Transformation, poor visibility (Bresnahan, 2012) and Organizational Adjustment Costs (OACs) (Bresnahan & Greenstein, 1996) have largely been absent. In contrast, as we shall see one by one, the various paths to a commercial-applications GPT loop require invention that creates Digital Transformation (and that has visibility problems and high OAC). That does not bode well for a rapid transition to a thrumming GPT feedback loop.

2 | A GREAT START, DEEP BUT NARROW

"Artificial intelligence" (AI) advanced from the '50s to the aughts as an academic enquiry with little practical impact. A switch to statistical approaches led to dramatic improvements in longstanding research areas, such as machine vision (MV) and natural language processing (NLP), and to the very rapid growth and improvement of machine learning (ML) and specific statistical techniques used to create a wide range of prediction engines.[3] It became practical to talk about "AITs," that is, ready-for-application engineering realities.

The supply side of these technologies advanced impressively. There was a dramatic improvement in AITs themselves.[4] Ease-of-programming for applications based on AITs advanced, as all major enterprise-serving development platforms, from Windows to AWS to multiple Clouds to Salesforce, added AITs.[5] Engineering supply grew very rapidly, as universities added new curricula and both young and midcareer engineers flocked to them.[6] There was an explosion of entrepreneurship in AIT areas.[7] In short, a planned business system using cutting-edge AITs would be *programmed* based on market-available products, frameworks, and programmers. The tricky, expensive bit would be, as usual with commercial information and communication technology (ICT) use, deciding specifically what the system should do.

Like some but not all new ICTs, "AI" came to enterprise computing with terrific buzz. Why? First, the technologies actually worked. Second, demonstration projects such as driverless vehicles got a great deal of press. Third, there were some early very valuable applications. Fourth, the (false) idea that there had been breakthroughs in "general AI" that would allow the automation of white-collar work captured the public imagination. The buzz, in short, swept up not only technologists but also businesspeople.[8] The buzz had a very strong element of technological determinism, which as we shall see led to trouble.

Bresnahan (2018) examined the business logic of early AIT-using applications. These early applications are narrow, but some are extremely valuable. There were applications running central production processes at Amazon, Google, Facebook, Netflix, and others.[9] There were new user interfaces (UIs), involving advances in machine speech and hearing and in NLP. There were some applications of MV, both in productive use and in the most famous demonstration projects, driverless cars. There were applications of prediction—modern AITs are about prediction—in applications where prediction is normal and familiar, notably in security, predicting stockouts, and science and engineering. This list alone suggests that there might be multiple potential paths for AITs to become very widely and deeply used; the early applications were diverse and drew on diverse technologies.

Across all those diverse areas of early application, the value propositions of the new technical developments and the process of invention had at least two commonalities:

- Substitution of AIT capital for human labor was irrelevant to the invention of AIT applications. AIT-using systems typically either used AIT to replace other ICT capital or to undertake tasks never undertaken before.
- The narrow range is characterized by low OACs and very good inventive visibility. To translate to the language of this Special Issue, the narrow range is characterized by the absence of any need for Digital Transformation.

3 | IMITATION OF THE MOST SUCCESSFUL APPLICATIONS STALLS

Propitiously for the hope that AITs would transform much of white-collar work, some very valuable early (2015 or so) applications were in commercial application, selling products at Amazon and Netflix and advertisements at Google and Facebook. Hundreds of billions of dollars in revenues from systems run by AITs are an impressive achievement in the application of ICT.[10] They are remarkable both for their depth and for the speed with which they arrived. They were also narrow, meaning as yet no GPT.

The speed, depth, and narrowness of these applications have a common origin. The firms making these fabulous AIT applications had no need to undertake a Digital Transformation. They were already successful with digital production. More specifically:

- The tasks performed by AITs were already algorithmic before.
- The production architecture was already modular, with the tasks that were performed by AITs separate from other functions. Intertask communication did not need radical change.
- Data were already used intensively.
- The human side of the organizations was largely staffed with smart, flexible people.

The existing digital production systems lowered application invention costs at the internet giants (IGs). The first point is visibility. Ex ante these inventors needed only see, for example, that predictive AITs could replace an existing algorithm which predicted the probability a specific user will click on a specific ad. With modularized production, the visibility enquiry starts and ends there, for the inventor need not see all the system-wide implications of the change to AITs.[11]

The second point is OACs. These firms did not have a complex human organization running production ex ante, they had an ICT system doing it. So the AITs could be implemented without changing any human jobs in production. Communication about what the AIT systems did could pass through software interfaces.

In a more typical firm's organization, both visibility and OAC would be problems. Adding new technology, even AITs, to a production or marketing locus typically changes what the rest of the organization needs to know and do. System-wide changes have problems of visibility, and new communications paths between the changed locus and the rest of the organization create OAC. As a result, AIT application would have higher costs of invention at firms not as far down the path of Digital Transformation.

Early experiments by the Internet Giants revealed another limitation in the range of applications of these AIT-based product/consumer matching systems. These systems worked well in a low-stakes context and poorly when the stakes were higher. In their contexts, low stakes typically meant the consumer could decide, without fuss, which product not to consume or which ad not to click; high stakes meant that merely being offered the decision might offend the consumer.[12] Other versions of high stakes arise in the numerous industries not served by the IGs. This, too, could raise the costs or slow the speed of launching a GPT loop.

I now turn to enterprise computing adoption that might have ignited a GPT loop.

3.1 | AIT-based business systems for other kinds of businesses

One path to creation of a GPT loop is simple. The announcement of the technical capabilities of the GPT causes widespread applications invention.[13] This path is rare for commercial ICT GPTs (Bresnahan, 2012). The problem lies in the invention in applications sectors, which can be blocked or long delayed by difficult problems of visibility and OAC. The buzz around "AI" led to an extraordinary volume of AIT-using experiments very quickly, so we can be quite sure this simple path to creation of a GPT loop stalled.

A number of sources document the small fraction of firms which have an AIT-based system in use, including Census.[14] Within those sources, the Gartner surveys also attempt to measure whether firms have an experimental or exploratory project. From 2016 to 2019, so after the early successes at the Internet Giants but before now, those surveys reported a high level of experimentation—over 40% of enterprises several years in a row, peaking at almost half. They also reported, and continue to report, a low rate of AIT-based systems in operation.[15] This is a higher experimental failure rate even than the usual high rate for radical ICT projects.

The same surveys report, and note the oddity of, the high importance of senior business people (not technical people) in pressing for these exploratory projects (e.g., Bretheneux, 2018). That likely reflects the underlying fact that the support from senior management was based on buzz, not on a careful assessment of the overlap between business needs and technical capabilities. More specifically, it was based on the false hope that AITs were uniquely powerful technologies that could replace workers one for one: "get me an AI" to replace selected workers. A large-scale effort to quickly spread AITs to important roles in enterprise systems stalled out, leaving the range of applications narrow.

3.2 | AITs as automation technology

Many scholars and public commentators frame AITs as automation technologies. This is an old mistake, based in the first instance on the tendency of definitions of "AI" to refer, imprecisely and metaphorically, to a system that can do a cognitive task "normally done by humans" or "previously done by humans."[16] The "driverless vehicle" discussion has revitalized this metaphor of automation. But it is a metaphor, not a precise engineering or business statement. In any case, the metaphor has become an axiom: "At the heart of our framework is the idea that automation and thus AI and robotics replace workers in tasks that they previously performed" (Acemoglu & Restrepo, 2018, p. 198).

The metaphor of automation implies a theory of the innovation path to a GPT loop. It implies an AIT diffusion path—across jobs. As AITs become more capable, the popular theory goes, they can perform more and more jobs. Another advantage is that many of the jobs that seemed plausible to be automated early in this process are common across industries, so generality would arise quickly. Finally, automating jobs straight up would avoid the pesky problems of visibility and OAC. A single worker is replaced by "an AI" and all other workers and systems go on unchanged.

But the automation metaphor is false. AITs have not been substituted for human labor to any significant degree. They have been substituted for other capital (as at the IGs) or have performed tasks previously not done at all. The desire to find "an AI" is an anthropomorphic error; there are not AITs which are like workers yet. The advice literature for enterprise computing people moved, not long after the automation buzz began to spread, to warning computer departments against these unrealistic expectations.

Will this change as AITs improve? Of course it might. People invent extraordinary things.

As a forecast path toward a GPT loop, automation theories need to deal with three related problems. First, many jobs are complex combinations of very different tasks. Is the truckdriver part of the transportation system or part of the payload or both? (Consider a delivery truck where the driver chases accounts receivable or a plumber's truck.) If the driver is part of the transportation system, is that confined to driving? (Consider a UPS delivery truck, where the driver does the last 50 yards on foot.) These examples, by the way, are drawn from the most favorable kinds of jobs for avoiding OAC, ones at the edges of organizations.

Second, many organizations are highly nonmodular and have complex links from one job to another. For example, there may be communications systems across workers which call for very different sets of skills than the tasks assigned to each job.

Third, and overlapping a good deal with the first two, in many firms the current structure of jobs and of organizations is not well understood by anyone, so that visibility of the development project of replacing a particular job with an AIT-based system will be low.

These three problems stand out in the empirical literature on enterprise systems development (Bresnahan et al., 2002; Bresnahan & Greenstein, 1996). This makes the job-automation path, even with advances in AITs, seem unpromising. Changing enterprise production systems to more digital form likely will continue to entail the difficult work of Digital Transformation, and thus need to overcome, rather than evade, problems of visibility and of OACs.

3.3 | Consumer-product matching systems

Another mechanism for starting a GPT loop would have a successful AS coinvention *itself* diffuse across sectors. A famous historical example occurred after computers were used in life insurance for business data processing; similar applications later appeared in the financial functions of large firms in nonfinancial industries. Another famous historical example is transactions processing systems, which were first invented by IBM to serve electrical utilities (Campbell-Kelly, 2003; sec. on Industry apps and on CICS). The application of AIT prediction engines in product/consumer matching at the Internet Giants suggested the imitative invention of matching engines for other contexts. Matching experiments flourished inside corporations; they flourished as products from startups, and they flourished as features of enterprise software.

These widespread attempts to make customer/product matching engines like those at the IGs ran into several problems. First, few firms have a consumer/product matching problem as central to their operations as the IGs. Second, stakes are far higher in other markets, even other mass markets (consider healthcare), lowering the value of imitative systems. Third, few firms are far enough down the path of Digital Transformation to avoid the adjustment costs of inventing and implementing such a system.

Other matching engines, even outside mass-customization industries, form a second diffusion path. These are numerous if not profound. Consider personalized versions of Frequently Asked Questions, prescreening the resumes of potential hires, suggesting a response to a text, and many others. These are formally like the IG systems, as a prediction engine suggests a short list to a human for choice, but they are not economically like the IG systems. These systems are useful, in that, for example, it is somewhat easier to find the command you are looking for in Windows, Salesforce, and other software, and rushed hiring managers can prescreen resumes more quickly. However, they are typically not close to running a major function, not creating new ICT-based systems that move companies closer to the efficiency and growth of the Internet Giants. That awaits significant coinvention and Digital Transformation.

The contrast is stark. One Internet Giant executive, thinking of his own highly modular production process, told me "AI is like oxygen around here, we try it in everything. And then I meet these salesmen who could build us a customized FAQ." The heterogeneity is important to understand the transition to generality. Firms that have not yet made a Digital Transformation are using AITs along these diffusion paths, but using them in ways that create comparatively incremental value. Firms already through the Digital Transformation, like the IGs, thus far gain much more from the AIT opportunity.

3.4 | Information for improvements in AITs

These more routine, less revolutionary, efforts to diffuse AITs through enterprise computing have led to a number of distinct improvements in AITs. Technologists have learned from the mismatches between enterprise applications and existing AITs. (The learning spread from technologists in enterprise to academic Computer Science and then to Business scholars.) The learning leads to more valuable technologies, but not to revolutionary ones. This is a central element of any GPT loop, and there are a number of important examples of it with AITs.

3.4.1 | Switch from automation to augmentation

After "AI is Automation" grew unpopular with enterprise customers, technologists switched, on a wide range of fronts, to using AITs to augment human work.[17] This is quite normal for ICT markets. The stated value propositions of new technologies ex ante are often partly wrong, particularly in low-visibility environments like enterprise computing, but ex post market selection can still be a powerful force. It was, here.

One example of augmentation is both old and new—expert systems. It was strikingly easy to convert supposed decision-making systems, originally assessed as substitutes for human work, to decision-support systems. This was true in earlier waves of decision-support systems, and it was true this time. The applications improved in some cases where decision support was longstanding—underwriting, diagnosis—and moved into some new areas as well.[18] The path of this improvement will be limited by many forces, including the need to change the rest of the organization to deal with a different decision process.

Another important switch, and a rapid one, was from "driverless vehicle" automation technologies to driver support augmentation technologies. In this case, the switch was rapid for two reasons. First, some technologies invented as part of driverless vehicles have been repurposed as driver-assist features. Second, these features were those long discussed in the literature on driver-assist, so they had excellent inventive visibility. These features sell well, with many new US vehicles having some of them (and nearly all new EU vehicles having some of them because of regulation).[19]

This represents a move of AITs into broad, existing trends in enterprise software and consumer categories. There is an existing, broad, "digital tools" trend in enterprise computing. For years, digital tools have been provided to support a wide variety of workers. Bringing AITs to this trend will increase the pace of expansion of that list of tools and improve their quality. In this regard, while the deployment of AITs is not transformative, it is useful. It respects the incrementalism of enterprise systems development through progress at the single-worker, not organizational, level.

The switch to augmentation follows a longstanding pattern. When there is low ex ante visibility, as is the norm for commercial ICT applications, the predicted value proposition of new technology is often based more on a fond hope than a market reality. Automation of white-collar work and the driverless delivery truck fall well within the "fond hope" category. Once invented, demanders can respond to the actual features of the initial inventions, and demand can shape the direction of technical progress ex post. Then value creation becomes a market reality.

Another response to market events was efforts to give AITs more positive features from a consumer or user perspective. These changes, too, may help the diffusion of AITs toward generality, if not overcome the important bottlenecks. One is making AITs seem more trustworthy and, relatedly, to make their logic easier to explain. In low-stakes environments, AITs can draw on the enormous power of predictive statistics; in such environments an inexplicable prediction does little damage to trust. In higher-stakes environments, an inexplicable prediction may harm trust. The diffusion of AITs to higher stakes environments may thus benefit from an increase in perceived trustworthiness. Similarly, any statistical procedure, including training AITs, performs worse in smaller-sample-size environments. This means AITs are at a minimum noisier in minority populations. Solving this problem by using data from both majority and minority populations can be worse, as it can make the prediction depend on minority status or be highly correlated with minority status. AIT researchers are working on overcoming these issues technically (Zhang et al., 2022; Chap. 3).

Advances in trustworthiness and avoiding bias may quickly benefit existing AIT adopters like the Internet Giants. These firms all have varying problems of being perceived as trustworthy and unbiased by customers or complementors.

Whether these advances also increase the pace of diffusion toward important generality of use in enterprise is an entirely different question. These advances remove some substantive problems with AITs, but do not remove the problems of OAC and visibility. This leaves organizations that have not yet used AITs for core business functions with the problems of discovering what functions can be based on AITs and the problems of redesigning the organization to accommodate that change. Once again, success in the widespread diffusion of AITs calls for a dramatic increase in the rate of Digital Transformation, an increase that would be valuable with or without AITs.

3.5 | "Generative" AITs

Another area is the use of AITs to write prose, draw pictures, and other creative activities. Early uses are various forms of play, though now these technologies may come to aid workers who write or draw as part of their work. That would be valuable, if (at the beginning) limited in scope. The value would likely arise in the same way as did AIT-based improvements in programmer tools, increasing the effectiveness of a given class of professionals. One can also imagine problems of trustworthiness arising if, for example, AITs write college term papers.

However, two aspects of such "generative" AITs suggest there may be more opportunities here.[20]

First, the "generative" AITs are based on moving the boundary between the general technology and the application through what are called "foundation models." This can exploit social increasing returns to scale and lower the development cost of applications. Models such as BERT, DALL-E, and GPT-3 are "trained" on enormous data sets and can be adapted with a modest effort to specific applications.[21] The foundation models can, sometimes, be adapted to an application with very low effort. This lowers the cost of applications development.

Second, some technologists conjecture that this ease of development will open the door to new and better human–machine interaction. A foundation model can accept informal or incomplete direction, and it will often work well. See, for example, Weisz et al. (2022) for the introduction to a conference on this point. AIT-based UIs already have already been valuable, perhaps this low-cost programming will expand their scope.

4 | UI IMPROVEMENTS DIFFUSING AWAY FROM IGs

Another diffusion path is centered around UI improvements based on AITs, particularly NLP and machine speech and hearing. The beginnings of these diffusion paths were, once again, in the "tech" firms.[22] Examples include Apple's Siri, a voice UI which is much discussed, and Google's switch in its search product to better guesses about what the user is looking for—the switch "from strings to things." This has set off two different diffusion paths, both useful, and neither, as yet, transformative.

UIs can serve as valuable access technologies. By an access technology I mean a complement to existing systems that permits those systems to be more used—more easily, more conveniently, more safely, or by more users.

Since early AIT-based UIs came from the IGs, they focused on improving consumer access to IG services. The voice UIs, for example, were particularly valuable to consumers in environments like kitchens or automobiles where traditional UIs might be inconvenient or dangerous. The same benefit has driven voice UIs into work environments where convenience or safety is important, such health care providers where voice UIs for communications systems can save much handwashing.[23] This is a lowering of the cost of systems access and an increase in its convenience.

Is this a breakthrough that will lead to voice UIs quickly becoming the key to a new GPT loop? It is invention, so many things might happen. However, past inventions of access technologies have made a large rapid impact when conditions of visibility were very good, specifically when businesspeople had already recognized the loci where existing systems needed better or broader access. Those initial conditions were present for both of the big successes for access technologies, the widespread use of the Internet and the widespread use of GUI-based PCs. In the Internet example, businesspeople had long seen that widespread access was the key to expanding electronic commerce, communications, and content dissemination to mass markets (Bresnahan, 2012). In the PC/GUI example, businesspeople had long seen the limitations of nontechnical workers accessing enterprise systems through "green screen" terminals.[24] While some opportunities for improved access are visible today, there is little business discussion of widespread contexts where a voice UI would be a dramatic improvement. In the near term, this diffusion path seems more likely to be as important an access technology as touch screens rather than as important as the widespread use of the Internet.

New NLP-based interfaces have also improved enterprise systems software after their initial use in consumer applications. Finding the desired command is far easier in many software packages as a result. Finding the desired turn of phrase in an email or text is easier, as well.[25] These interfaces are spreading into more valuable territory, making enterprise software more of a "do what I meant" product. Here, too, the AIT-based advances are complementary to existing systems, and the economic impact of the uses is to increase existing systems' value.

This is a different move forward toward a GPT loop. Language processing was an existing area of AIT research; early applications are where something much like that research product is valuable. Different ones may follow, as did business data processing uses for "computers" (literally making computations) after a lag.

4.1 | Two smaller trends

Many observers put robots and AIT-based "bots" in the same category of automation technologies. This is merely a category error with regard to the AIT-based systems. Some robots use MV, which could contribute to advances in MV by providing a steady, if not very large, source of demand.

Few jobs have seen as steady or rapid technical progress over as long a period of time as that of the programmer. AIT technologists tend to say that ML *already* means that "the computer is programming itself." In the narrow sense in which this is true, prediction machines are easier to write than in the days of many algorithmic switches. Another narrow sense in which AITs can make the work of programmers easier is automated coding and debugging. Testing and debugging is more than half of programmers' work, so there is room for productivity gains. If this proves correct, it will continue the longstanding tendency of new ICTs to succeed much more quickly in domains where the demander is technical.

5 | PREDICTION WAS FAMILIAR

A number of areas of AIT use had high visibility ex ante because prediction was an already familiar element and, as noted above, modern AITs are based on statistical prediction. These include applications in science, in security, and in production environments where there was already digital prediction of stock outs, capacity crunches, and the like (in an establishment and across a supply chain).

AITs can help in R&D activities. Scientists and engineers are comfortable with uncertainty and prediction. As computers did many decades ago, predictive AITs are lowering the costs of R&D and enabling frontier projects. This form of usage has spread quickly, and clearly predictive AITs are already in general use in these technical environments (Cockburn et al., 2018).

Prediction engines can help guide production to avoid problems like running out of capacity or inventory, accepting a stolen credit card, and so forth. MV is also used to identify individuals in some security applications. Thus far, these applications appear in low-OAC and high visibility environments. The difficulty with rapid diffusion along this path arises when prediction needs to find a new, hard-to-foresee use (visibility) or requires reorganization (OAC).

Agrawal et al. (2018) (AGG) suggested a possible path to a GPT. Prediction could be the core technology output that might draw complements, continue to improve dramatically, and then break out of applications that are literally and visibly prediction into other applications where the application never entailed prediction before. As they point out, such transformations have occurred in the past; computer arithmetic grew cheaper, attracted complements, and diffused to new uses.[26]

Before concluding that such a path will be rapid or even occur, one should consider the recombinatory steps that made arithmetic the core of a computing GPT in the historical example. The key transition there, we know ex post, was the diffusion of computing from science to business data processing, which eventually became enterprise computing. Then as now, that was a much larger and more general form of end use than science and engineering.

The recombinatory inventions that brought electronic computing (in the sense of arithmetic) to business data processing and that then made it very important in some BCD applications were not inevitable. The first step, which brought electronic computation to BCD, was the creation of a computer which was compatible with *existing* BCD data storage, retrieval, and input/output devices, and which permitted incremental improvements to *existing* BCD applications that were already doing *some arithmetic*.[27] OAC were low because there was no change to the organization, and initial applications' visibility was very localized—what if a card tabulator could do more arithmetic than it was already doing? With the volume of prediction application experiments undertaken using AITs so far, it seems unlikely that an opportunity like this has been missed.

Enterprise computing, once it was electronic, did mostly integer arithmetic. (The scientists, mostly floating-point calculations.) Arithmetic computation became important in this domain some decades later, as advanced software systems, plus experience with databases, permitted the invention of analytical applications. More and more applications inventions have occurred since then, some using complex arithmetic. (AITs are just the latest example.) This was a long, slow, indirect path, and one that made arithmetic coequal with data storage and transmission as hardware functionalities, not itself the center of a GPT loop.

Clearly a recombinatory path to a GPT is possible. It happened with computer arithmetic, if not quickly. Perhaps there will be a prediction revolution that will be much faster than the arithmetic revolution. The close complements for AI-based prediction are data, data storage, data management, and analysis technologies and their intelligent use in companies. The data *technologies* are out in the marketplace and have been deeply linked in all the important development platforms to AITs. The *intelligent use* of data remains rare and valuable in enterprise. So it is difficult to predict a “prediction revolution” faster than the arithmetic revolution, other than by noting that someone might invent something. That critical invention is unlikely to be technical, as the bottleneck lies on the use side. Once again, it is widespread Digital Transformation, including in firms where it will be difficult, that would be the key to this path.

6 | REFRAME GPT DISCUSSION

What do we mean when we say “AITs could become GPTs?”

> **Problem 1.** What exactly do we mean when we say “GPT” when there is a technology stack? By “stack” I mean that applications use a number of different general-purpose components, mostly complementary to one another, in varying proportions.

Conceptually, there are clearly two different answers to this question. One answer is that the whole stack is the GPT and a new layer has a particular kind of improvement to it. Another answer is that each general-purpose component that has a positive feedback loop is a GPT. Each of these answers can be more useful depending on what is being analyzed.

Prediction based on AITs is so tightly complementary with other ICT, especially data-related technologies, that there may not be any analytical purpose where it is better to think of it as a separate GPT. More generally, even including the novel uses of NLP in UIs, most applications of AITs so far exhibit powerful capital–capital complementarities.[28] AITs increase the depth and flexibility of the ICT stack.

Thus far, it appears that AITs have been easiest to adopt in commercial applications when the applications were already data-based and algorithmic. As often happens in the commercial application of ICT, this might mean that the long-run benefits arrive via an indirect path.[29] The most profitable path at the firm level might involve future investment in data technologies followed much later by AIT deployment.

Problem 2 (Technological determinism). It is an error to think that being a GPT is an attribute only of a technology, not of applications sectors. When applications sectors are themselves largely engineering, as in the examples of factory electrification (Du Boff, 1966) or machine tools (Rosenberg, 1976) or computer *programming* (Bresnahan, 2019). However, as we have seen in this paper, the bottlenecks to launching a GPT loop for many past ICT technologies are in inventing key applications. And, again as we have seen through this paper, many of the more promising paths to widespread productive use of AITs call for the invention of new organizational structures, new relationships with customers, and other difficult tasks. Those paths will be slow.

Both the problems raised in this section point to a core error in much academic and policy-arena analysis of technical progress, especially ICT-using technical progress. Improvements in the underlying GPT are necessary but not sufficient for advance. Brilliant AIT inventions or large public investments in AITs are likely to be less efficacious than brilliant application inventions in moving the whole loop forward.

However, throughout this paper we have seen that an AIT GPT loop might arise by paths similar to a number of earlier GPTs. Those paths are quite distinct from one another, so there are many ways for these new technologies to create powerful feedback loops. Those paths also varied dramatically in how long they took to be realized, and the quicker ones (discussed in Section 3.1, above, now seem less likely.

This paper has not addressed some of the widely discussed potential problems with AITs: improvements in security systems may become improvements in autocratic state surveillance; biases arising from the backward-looking "training" may limit adoption; AIT invention may migrate from very open-science as it is today toward a single provider (as did computers and PCs but not the Internet) with the attendant innovation problems. These are serious topics that may become very important if AIT application widens.

7 | CONCLUSION

AIT use in enterprise systems is narrow but deep. Other AIT use is in specialized environments, technical environments, or ones where the prediction was already central. Academics, who see AITs in use in a large number of scientific and engineering disciplines, tend to already see "AI everywhere." That is, everywhere except where large gains to firms and to economic growth might arise; the large gains occur in just a few firms in just a few industries.

This paper has examined the main paths that might lead to AITs becoming widely and deeply used in commercial environments beyond the early areas of applications. All of those paths involve not only improvements to AITs themselves, but also widespread Digital Transformation. Technical progress has not, however, lowered the invention costs which make Digital Transformation difficult and scarce. The important bottlenecks remain in coinvention.

ORCID

Timothy Bresnahan http://orcid.org/0000-0003-2854-7035

ENDNOTES

[1] See, for example, Brynjolfsson et al. (2017), Zolas et al. (2021), and Bresnahan (2018). These papers draw on very different data sources, finding the narrowness in all. The third paper also has cites to some prominent versions of the GPT conjecture, and the first examines the GPT possibility at length.

[2] By "recipe" approaches I mean two arguments every reader of this paper will have seen. One enumerates parts of the ex post definition of a GPT—a technology that is capable of continuous improvement, innovational complementarities with applications invention, and use throughout the economy. This approach must forecast all three, and typically fails to have a reliable basis for the second and third elements, which are distinctly ex post concepts, relying primarily on the first. That is, in the case of "AI," attractive but incomplete; AITs are powerful technologies capable of continuous improvement. The other recipe approach notes that everything that has grown very large in the past has had a period of rapid growth, typically following a period of slow growth, and falsely generalizes that slow growth predicts faster growth.

[3] Agrawal et al. (2018) offer good sources for the history.

[4] Measuring improvements in AITs has all the difficulties of measuring improvements in computing or software (Berndt & Griliches, 1993; Berndt et al., 1995). Zhang et al. (2022, Chap. 2) discuss the technical performance of AITs in eight categories, such as vision, general reinforcement learning, and so on, with different metrics such as the time and data requirements to "train," accuracy on several dimensions, and so on.

[5] See Baker et al. (2022) for developer tools in cloud-based platform services, the most common locus for ML and AIT services.

[6] See Zhang et al. (2022, Chap. 4.4 "AI Education").

[7] See Zhang et al. (2022, Chap. 4.2).

[8] Rollings (2018) documents the wave. Bretheneux (2018) has practical advice to IT personnel on "Managing Business Leadership Expectations" in light of all the "Artificial Intelligence Hype."

[9] These applications are discussed at some length in Bresnahan (2018) which also has citations.

[10] See Bresnahan (2018) for more detailed discussion of the business and technical logic of these systems.

[11] In remarks at the Conference, Joshua Gans interpreted this as a "conjecture" that AITs and production modularity are complements. No. It was an empirical observation about the process of invention of new production systems, not about substitution and complementarity in production processes once invented.

[12] See Bresnahan (2018) for discussion of the experiments that revealed the problem of stakes.

[13] This announcement need not be of the first version of the GPT, but of an improved version. Many GPT loops with technical users, including technical uses of ICT (not commercial ones) follow this path. See, for example, Bresnahan and Malerba (1999) for the distinction between commercial and technical-user ICT and its implications for diffusion and marketing of computers. AIT use where the customer is a technical worker has followed this path as well. A well-documented example is used by scientists: see Cockburn et al. (2018).

[14] See, for example, Vesset (2019), Austin et al. (2017), and Zolas et al. (2021).

[15] And a support structure to align the experimental systems building with business goals. The big commercial survey services for ICT like IDC and Gartner reported a stunningly wide "AI Readiness" distribution for US corporations.

[16] Engineers dislike this element of the definition. My Engineering colleague John McCarthy, who coined the term "AI," doubted the value of links to human intelligence (McCarthy, 1989).

[17] My colleagues in the (enormous) Stanford center on "Human Centered Artificial Intelligence" have dedicated their research to this change. See Zhang et al. (2022).

[18] See Zhang et al. (2022, p. 158) for data on where investment in AIT products is going. Medical expert systems are very important. See Agrawal et al. (2022) for a discussion of how this might spread to decisions beyond the traditional range of expert systems.

[19] The longstanding literature on driver-assistance technologies had been discussing most of these categories, like blind spot vision, safe following distance, lane control, and so on. Inventive visibility was therefore good. Considerable experience with computer control of vehicles, notably aircraft but some ground vehicle elements like antilock brakes, also lowered adjustment costs. The one technology that could be characterized as automation is self-parking, which is not a large fraction of driving. Of course, driverless vehicles will have commercial applications someday. Many technologies eventually loop back around to their initial photogenic applications.

[20] See, for example, Brown et al. (2020) for the first of these opportunities in language processing.

[21] "GPT-3" stands for third-generation "generative pre-trained transformer." The confusion with "GPT" arises because there are two fields of knowledge. The same explains a novel meaning of "generative," another name for "enabling."

[22] Readers with a strong technological determinism streak sometimes read this sentence to say the IGs were the original inventors of these UIs. That is an error.

[23] Scholars continue to focus on expert systems for diagnosis as the central AIT application in health care. So far, voice applications are dramatically more important in actual use. See, for example, Zolas et al. (2021, Table 16) which shows that the top three four-digit industries ranked by voice recognition are all in health care provision.

[24] See Chposky and Leonsis (1988, p. 10) for a discussion of IBM's consideration of the market opportunity.

[25] These nudge approaches are not without costs and AITs have no sense of humor or irony. A few paragraphs back, I was trying to write about the switch from automation to augmentation, and Word helpfully suggested "augmented reality." There is a paper to write: "What invention paths for AR to become a GPT?"

[26] AGG attribute this to me, but, as the next paragraph shows, I argued this is an arduous path.

[27] The relevant recombinatory computer was the IBM 650, which was compatible with existing IBM electromechanical (not electronic) equipment. Yates (2005) documents the way electronic computers from IBM and Remington Rand could be inserted into exiting data processing organizations without change in their structure.

[28] MV might have more duality than other AITs. There are some applications where MV is central, and those might form a positive feedback loop that is itself analytically important for some question. MV might also participate in other applications as a feature of a broader ICT GPT.

[29] Such indirect transitions have happened before. Consider the DBMS. The first steps were digital records for mundane purposes, followed by the DBMS emerging as a programmer tool. Once the tool and the data were in place "strategic" applications running on top of the DBMS were easier to invent and very valuable.

REFERENCES

Acemoglu, D., & Restrepo, P. (2018). The economics of artificial intelligence: An agenda. In A. Ajay, G. Joshua, & G. Avi (Eds.), *Artificial intelligence, automation, and work*. University of Chicago Press.

Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.

Agrawal, A., Gans, J., & Goldfarb, A. (2022). *Power and prediction: The disruptive economics of artificial intelligence*. Harvard Business Review Press.

Austin, T., Linden, A., & Rollings, M. (2017). *Hype hurts: Steering clear of dangerous AI myths*. Gartner.

Baker, V., Sicular, S., Bretheneux, E., Batchu, A., & Fang, M. (2022). *Magic quadrant for cloud AI developer services*. Gartner Group.

Berndt, E. R., & Griliches, Z. (1993). Price indexes for microcomputers: An exploratory study. In M. F. Foss, M. E. Manser, & A. Young (Eds.), *Price measurements and their uses*. NBER Studies in Income and Wealth, University of Chicago Press.

Berndt, E. R., Griliches, Z., & Rappaport, N. J. (1995). Econometric estimates of price indexes for personal computers in the 1990's. *Journal of Econometrics*, *68*, 243–268.

Bresnahan, T. (2012). Generality, recombination, and reuse. In J. Lerner & S. Stern (Eds.), *The rate and direction of inventive activity revisited*. University of Chicago Press.

Bresnahan, T. (2018). Artificial intelligence technologies and aggregate growth prospects. In Z. George & D. John (Eds.), *Prospects for economic growth in the United States*. Cambridge University Press.

Bresnahan, T. (2019). Technological change in ICT in light of ideas first learned about the machine tool industry. *Industrial and Corporate Change*, *28*, 331–349.

Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics*, *117*, 339–376.

Bresnahan, T., & Greenstein, S. (1996). Technical progress and co-invention in computing and in the uses of computers. *Brookings Papers on Economic Activity: Microeconomics*, *1996*, 1–83.

Bresnahan, T., & Malerba, F. (1999). Industrial dynamics and the evolution of firms' and nations' competitive capabilities in the world computer industry. In D. C. Mowery & R. R. Nelson (Eds.), *Sources of industrial leadership: Studies of seven industries*. Cambridge University Press.

Bretheneux, E. (2018). *Artificial intelligence hype: Managing business leadership expectations*. Gartner Group.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., & Amodei, D. (2020). *Language models are few-shot learners*. arXiv. https://doi.org/10.48550/arXiv.2005.14165

Brynjolfsson, E., Rock, D., & Syverson, C. (2017). *Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics NBER* [Working Paper No. 24001].

Campbell-Kelly, M. (2003). From airline reservations to Sonic the Hedgehog: A history of the software industry. In I. Bernard Cohaen & W. Aspray (Eds.), *History of computing*. MIT Press.

Chposky, J., & Leonsis, T. (1988). *Blue magic: The people, power, and politics behind the IBM personal computer*. Facts on File.

Cockburn, I. M., Henderson, R., & Stern, S. (2018). *The impact of artificial intelligence on innovation* [National Bureau of Economic Research Working Paper Series, No. 24449].

Du Boff, R. B. (1966). Electrification and capital productivity: A suggested approach. *The Review of Economics and Statistics*, *48*, 426–431.

McCarthy, J. (1989). *What is AI?/Basic questions*. http://jmc.stanford.edu/artificial-intelligence/what-is-ai/index.html (November 21, 2022).

Rollings, M. (2018). *Deliver artificial intelligence business value: A TrendInsight report*. Gartner.

Rosenberg, N. (1976). Technological change in the machine tool industry. In N. Rosenberg (Ed.), *Perspectives on technology*. Cambridge University Press.

Vesset, D. (2019). *Artificial intelligence: A slow-motion explosion*. IDC.

Weisz, J. D., Maher, M. L., Strobelt, H., Chilton, L. B., Bau, D., & Geyer, W. (2022). HAI-GEN: 3rd workshop on human-AI co-creation with generative models. In *27th International Conference on Intelligent User Interfaces*. Association for Computing Machinery.

Yates, J. A. (2005). *Structuring the information age: Life insurance and technology in the twentieth century*. Johns Hopkins University Press.

Zhang, D., Maslej, N., Brynjolfsson, E., & Etchemendy, J. (2022). *The AI index annual report*. Stanford Institute for Human-Centered AI.

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., & Dinlersoz, E. (2021). *Advanced technologies adoption and use by U.S. firms: Evidence from the annual business survey*. National Bureau of Economic Research.

How to cite this article: Bresnahan, T. (2023). What innovation paths for AI to become a GPT? *Journal of Economics & Management Strategy*, 1–12. https://doi.org/10.1111/jems.12524