Volume Title: The Economics of Artificial Intelligence: An Agenda

Volume Authors/Editors: Ajay Agrawal, Joshua Gans, and Avi Goldfarb, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-61333-8 (cloth); 978-0-226-61347-5 (electronic)

Volume URL: http://www.nber.org/books/agra-1

Conference Date: September 13–14, 2017

Publication Date: May 2019

Chapter Title: Artificial Intelligence as the Next GPT: A Political-Economy Perspective

Chapter Author(s): Manuel Trajtenberg

Chapter URL: http://www.nber.org/chapters/c14025

Chapter pages in book: (p. 175 – 186)

6

Artificial Intelligence as the Next GPT
A Political-Economy Perspective

Manuel Trajtenberg

6.1 Introduction

Artificial intelligence (AI) and related technologies are being heralded as "the next big thing," one that promises to revolutionize many areas of economic activity and thus to have a profound impact on economic growth. However, the rise of AI coincides with a recent wave of pessimism in terms of productivity growth, expressed forcefully by prominent economists such as Larry Summers (2016), and more thoroughly by Robert Gordon (2016).

Side by side with the gloom, the new "technology enthusiasts" envision a not-too-distant future in which AI will displace most (*all*?) human occupations while unleashing tremendous gains in productivity. This view poses once again disturbing questions about the future of employment, the distributional consequences of mass displacement, and so forth.

Nobody holds the crystal ball, hence rather than arguing about the inscrutable future, it is at least as important to inquire into what we can learn from history regarding episodes like this, that is, the appearance of a major new technology that is posed to have profound economic implications. Of course, the future is never a replay of the past, but it may provide a useful benchmark against which to assess the unfolding of the new technology.

Mokyr (2017) sounds a cautionary note in that regard: ever since the dawn of the Industrial Revolution in the late eighteenth century, both the pessimists and the enthusiasts have almost invariably been proven wrong.

Manuel Trajtenberg is a professor in the Eitan Berglas School of Economics at Tel Aviv University and a research associate of the National Bureau of Economic Research.

This is a follow-up to my discussion at this conference of Joel Mokyr's paper "The Past and the Future of Innovation: some lessons from Economic History." For acknowledgments, sources of research support, and disclosure of the author's material financial relationships, if any, please see http://www.nber.org/chapters/c14025.ack.

Moreover, Mokyr dismantles with solid historical and present day evidence Gordon's claim that technological advance is bound to slow down in a deterministic fashion (in particular, the claim that "all low-hanging scientific and technological fruit has already been picked").

However, nothing can be taken for granted—as Mokyr skillfully describes, institutions (including government policies) may play a key role enabling or retarding innovation. This is precisely the focus here: given that AI is poised to emerge as a powerful technological force, I discuss ways to mitigate the almost unavoidable ensuing disruption, and enhance AI's vast benign potential. This is particularly important in present times, in view of political-economic considerations that were mostly absent in previous historical episodes associated with the arrival of new general purpose technologies.

6.2 Is This Time Different? The Political Economy of Technological Disruptions

The presumption here, well argued in other papers in this conference,[1] is that AI has the potential of becoming a general purpose technology (GPT) in the foreseeable future, [2] thus bringing about a wave of complementary innovations in a wide and ever-expanding range of applications sectors. Such sweeping transformative processes always result in widespread economic disruption, with concomitant winners and losers.

The "winners" are primarily those associated with the emerging GPT sector itself, and those that are at the forefront of the deployment of the GTP in the main applications sectors. They tend to be young, entrepreneurial, and equipped both with the technical knowledge and the skills that are made relevant by the new GPT. The labor force composition of Silicon Valley offers a grand view of who are the winners in the present information and communication technologies (ICT)/internet era. There are further winners in those sectors that are ancillary to the core GTP circle, be it in services that directly benefit from the growth of the GPT (e.g., the venture capital (VC) industry, patent lawyers, designers, etc.), or in others that just ride on the localized boom (e.g., upscale restaurants and entertainment, gyms, tourism, etc.).

The "losers" are mostly those employed in sectors that structurally cannot benefit from the unfolding GPT ("laggards"), and those in industries where the adoption of the new GPT renders many existing competencies and skills obsolete, thus bringing about massive layoffs. They tend to be middle-aged, have lower than average educational levels, and reside in areas that do not have much diversified sources of employment.

As economists, we tend to view the big sweep of economic growth since

1. See Cockburn, Henderson, and Stern (chapter 5, this volume).
2. See Bresnahan and Trajtenberg (1995).

the Industrial Revolution as the very embodiment of the "Idea of Progress" (as conceived in the Enlightenment), and hence the rate of growth of gross domestic product (GDP) as an unequivocal uptick in the welfare of society as a whole. Sure, we do acknowledge that there are distributional consequences, and sure, ever since Pareto we know that we are not allowed to "sum-up utilities" (and thus the "minuses" of losers do not cancel out with the "pluses" of winners). But those half-hearted qualifications become just lip service—the truth is that we rarely dwell into the balance of winners and losers, and in particular we do not pay much attention to the later. Paraphrasing the well-known dictum of Isaac Newton, we may say that

> We enjoy today higher standards of living because we are standing on the broken backs of those that paved the way for technological progress, but did not live long enough to benefit from it.

Partly in response to these inequities, the post-World War II era saw the creation of the welfare state, including unemployment insurance, transfers to the disadvantaged, some form of health insurance, retraining programs, and so forth. These "safety nets" were supposed to provide a reasonable palliative to "losers," but the truth is that we still do not have effective mechanisms to prevent or ameliorate the costs of major technologically induced transformations.[3] Moreover, existing safety nets will quite likely fail to cope with the juxtaposition of two new and powerful phenomena: (a) much larger flows of GPT-displaced workers *and* (b) a new "great demographic transition." Let us examine each in turn.

Regarding the extent of displacement: technological change always causes disruption, as brilliantly articulated by Schumpeter's notion of "creative destruction." Furthermore, there are inflection points as a new GPT starts working its way through the economy, when in relative short notice very many sectors, competencies, and skills became laggards and obsolete.

However, as clearly envisioned in this conference, AI in its various incarnations seems to go much further, in that it has the potential to replace a very wide swath of human occupations. Many argued forcefully that there are ***no*** occupations that cannot be eventually replaced by AI, and that the *vast majority* of present occupations will indeed vanish within a generation.

The consensual view seems to be that a large proportion of employment as we know it today will give way to smart machines, and therefore that x percent of workers will be displaced, whereby $\boldsymbol{x}$ is thought to be *significantly larger than in previous GPTs.* At the same time, the extent to which new, presently unforeseeable occupations may arise (denote them y percent) seems to be constrained by the very nature of AI: presumably AI will be able to

3. Typically, these safety nets function reasonably well when dealing with the consequences of not-too-pronounced business cycles or with small, *temporarily* deprived groups of the population. Not so when there are major structural transformations or when the underlying conditions that led to welfare dependency become permanent.

perform most of the new tasks, and hence they will not constitute a good enough counterbalance to the disappearing jobs, as has been the case in the past. The prevailing view is then that the net displacement of employment ($x - y$) will turn out to be significantly larger for AI than in previous episodes of technological disruption, posing a serious challenge to traditional economic policies.

The second part of the challenge entails a steep drop in birth rates together with the extension of life expectancy (which has been steadily growing for well over a century). These powerful demographic forces have resulted in aging populations, with the concomitant increase in the dependency ratio and the looming threat on the long-term viability of the pension system. Notice that life expectancy is now increasing well past the retirement age, so that a typical person in her fifties contemplates a further stretch of twenty-five to thirty years of life. Thus, the prospect of being permanently laid off at that stage in life has dire consequences for the displaced individual as much as for society as a whole.

The joint effect of a large influx of displaced workers at the seemingly unique inflection point posed by AI, together with their longer life expectancy, may thus create a formidable challenge that even the most advanced welfare state will be hard pressed to cope with. Put differently, we cannot afford to have many more, and longer-lived, unemployed or underemployed people. This is what is at stake with the advent of AI.

There is yet another significant development that magnifies the challenge, and that is the *democratization of expectations*. The growth in income per capita involves not only a rise in material standards, but in other no less important dimensions of well-being, including reduced uncertainty and a concomitant heightened sense of control over our own lives, which entails also the expectation of having a voice in processes that affect us (Hirschman 1970). Not by coincidence, economic growth and expanding democracy have more often than not gone hand in hand within, as well as across, countries.

The Luddites of the early nineteenth century surely had their voice heard, as did their like-minded emulators over the following decades. However, they could hardly expect to make a dent on their fate: democracy was still highly limited and living standards still very low for the vast majority, so that most people were just consumed by the need to provide for their basic needs.

Much has changed since, and nowadays virtually every individual in advanced western countries has come to expect to be entitled, at least in principle, to full participation in every realm of society: the political, the economic, and the cultural. The expectation is not just to vote in periodic elections, but to have an influence via "participatory democracy"; not just to hold a job, but to partake in the benefits of economic growth—this is what constitutes "the democratization of expectations."

We claim that in such context it has become much harder to have some (many?) bear the costs of technological disruption (the losers), while others

reap the benefits (the winners). Moreover, the losers have become much more skeptical of the vague promise that *eventually* the benefits will "trickle down" to them as well. With good reason: experience shows that the losers typically remain on the downside, even if the welfare state somehow softens their human costs. In advanced, democratic societies, people have become more impatient, more demanding of government, more intolerant of false promises, as well as of collective failures. Again, this should be surely considered a highly *positive* by-product of the rise in living standards.

The sharp split between winners and losers, if left to its own, may have serious consequences far beyond the costs for the individuals involved: when it coincides with the political divide, it may threaten the very fabric of democracy, as we have seen recently both in America and in Europe. Thus, if AI bursts into the scene and triggers mass displacement of workers, and demography plays out its fateful hand, the economy will be faced with a formidable dual challenge that may require a serious reassessment of policy options:

- Governments may have to assume a wider responsibility for navigating effective transitions from old to new GPTs, and not just for alleviating some of the costs. As said above, the democratization of expectations will not allow just for cosmetic adjustments—the political economy of it will eventually force real change.
- In so doing, governments may have to consider courses of action aimed *inter alia* at reducing significantly the number of those that fall in between the cracks during such transitions: actual and potential losers are bound to become much less tolerant of their fate. This should be done ***not*** by attempting to slow down the pace of technical change (that would be silly and ineffectual), but on the contrary, by making sure that many more can be brought to partake in it.

6.3 From Threat to Promise: Strategies for the AI-GPT Era

In order to meet the above-mentioned challenges, governments will have to design innovative strategies in the following key areas:

1. education: search for ways to provide for the changing nature of skills required for the AI era;
2. personal services: these are the fastest-growing occupations, but as defined at present cannot benefit from AI; and
3. direction of technical change: strive to human-enhancing innovations, not human replacing.

6.3.1 Education: The Upcoming Revolution

As already mentioned, the expectation is that AI will become the dominant GPT of the coming era, spreading throughout the economy, and

displacing in the process a great many occupations. At the same time, the remaining occupations and new ones that may spring up as complementary to AI will require a new set of skills that are not quite those currently provided by the education system, at all levels.

This is not new: the first and second industrial revolutions in the course of the nineteenth century required, and were accompanied by, corresponding revolutions in education. The need to rely on a more skilled, educated workforce, as well as a more disciplined one, fed educational reforms first in Prussia (already in the late eighteenth century), then in the United Kingdom and in the United States, that led gradually to the institutionalization of free and universal education, with highly structured, government-set curriculums.

From the late nineteenth century to this day, this "factory model" of education spread widely, expanding *quantitatively* in all dimensions: more hours spent at school, more subjects covered, and more years of study. Thus, for example, the average years of schooling in the UK adult population was less than 1 in 1870, whereas at present it stands at over 13. Universal education now starts at age three to four in many countries, high school became compulsory in the second half of the twentieth century, and in the past three decades some form of tertiary education has become commonplace.

It is now widely accepted that this "factory model" needs to be revised and perhaps totally revamped in view of twin pervasive developments: first, the internet revolution, which in this context means the availability of information/knowledge on any subject, at all times and virtually at no cost; second, the rapidly changing requirements for meaningful employment.

In particular, the advent of AI as the new GPT, with its expected pervasive impact on employment, may call for a new education revolution, very much like the industrial revolutions of the nineteenth century. The key to it appears to be the shift away from imparting knowledge per se, to developing skills relevant for an AI-based economy. Likewise, such educational revolution will in all likelihood aim toward "personalized education," departing from the quest for uniformity that has characterized education systems ever since Prussian reforms.

What are likely to be the top skills required for employment in the upcoming AI era? There is a great deal of heated discussion in this area, but some agreement is emerging around a core set of skills, such as those listed in table 6.1.

There is a great deal of similarity between these three lists of skills, and in fact they can be classified into the following (nonexhaustive) main types:

- **Type I: analytical, creative, adaptive**
  - critical and creative thinking
  - analytical and research
  - sense-making

Table 6.1 Skills sought for employment (from websites)

UNICEF 10 life skills	MyStartJob.com	Top10onlinecolleges.org
1. Problem-solving	1. Communication skills	1. Sense-making
2. Critical thinking	2. Analytical and research	2. Social intelligence
3. Effective communication	3. Flexibility-adaptability	3. Novel adaptive thinking
4. Decision-making	4. Interpersonal abilities	4. Cross-cultural competency
5. Creative thinking	5. Decision-making	5. Computational thinking
6. Interpersonal relationships	6. Plan, organize, prioritize	6. New media literacy
7. Self-awareness	7. Wear multiple hats	7. Transdisciplinary
8. Empathy	8. Leadership/management	8. Design mind-set
9. Coping with stress	9. Attention to detail	9. Manage cognitive load
10. Coping with emotions	10. Self-confidence	10. Virtual collaboration

  - novel adaptive thinking
  - design mind-set
- **Type II: interpersonal, communication**
  - effective communication
  - interpersonal relationships/abilities
  - social intelligence
  - virtual collaboration
- **Type III: emotional, self-confidence**
  - self-awareness
  - empathy
  - coping with stress
  - manage cognitive load
  - coping with emotions

The important point to notice is that most of these skills are neither imparted in the current K–12 system, nor in academia. The whole system is still geared primarily toward the transmission of knowledge, highly structured and uniform, and not toward skills, let alone those skills. Pupils of all ages are now very aware of the fact that school-like information is available at the tip of their fingers, they are less receptive to frontal classes, their attention span is much shorter, and the sort of stimuli that makes them tick is different. This is also true at the tertiary level, and in addition, we are witnessing there the rise of the massive open online courses (MOOCS) and of other such online-based teaching tools.

In view of these trends, educational strategies may need to undergo equally significant changes away from the "factory system," and the fact that the incipient GPT may render many existing occupations obsolete, provides it with renewed urgency. These are some of the issues to tackle:

- Invert the pyramid: it is now widely recognized that critical skills, hard and soft, cognitive and social, are acquired very early on. Furthermore,

failure to do so at the earliest stages may be hard (even impossible) to remedy later on (see, e.g., Heckman et al. 2014). Thus, we may have to consider investing much more in early childhood education, from birth to age six.
- Find ways to incorporate the development of skills (of the three types sketched above) as an integral part of teaching in every discipline and at all stages, including in academia.
- Effective educational methods are hard to come by, thus it is important to engage in bottom-up experimentation in pedagogy, school design, and social skills development in the context of flexible, creative, teaching environments.
- Reconsider the prevalent norm of uniform (typically government-mandated) curriculums and educational models, vis à vis diversity and open-innovation communities built around educational institutions.
- Foster research on the effectiveness of new educational models, their adequacy to shifting needs, the extent to which they promote equal opportunity, and so forth. This type of research will be crucial given the move away from "top-down" models and the emphasis on widespread experimentation.

6.3.2 Upgrading Personal Services

A Bureau of Labor Statistics (BLS) study[4] projects that virtually all of the employment gains in the decade to 2024 will be in services, and within the service sector particularly in health care and social assistance (see table 6.2).

Many of these occupations as performed today require little training and minimal educational attainment. Not surprisingly, most confer low wages, low status, and are supported by very little complementary technology. As the projections suggest, those occupations are at present not seriously threatened by AI—on the contrary, they will grow significantly. Thus, the overall prospects look rather gloomy when not only employment is considered by also wages: major upscale occupations are projected to remain stagnant or decline, whereas low-scale occupations are expected to grow.

Is this a deterministic outcome? Not necessarily, and the case of nursing may be quite instructive. After World War II, nursing was one of the lowest-ranking occupations in the United States: in 1946 the average wage of a nurse was just one-third that of female workers in the garment industry.[5] In 1964 Congress passed the Nurse Training Act, which essentially redefined the occupation and turned it into a profession requiring an academic

4. See: Occupational Employment Projections to 2024, Monthly Labor Review, US Bureau of Labor Statistics, Dec. 2015. Also in https://www.bls.gov/opub/mlr/2015/article/occupational-employment-projections-to-2024.htm.

5. In 1946, the average registered nurse (RN) earned about one dollar an hour—or $175 a month.

Table 6.2 **US employment by major sector (millions)**

Sector	2014	2024*	Change*	Percentage growth*
Goods producing	19	19	~	~0
Services	121	130	+ 9.3	+ 7
Of which: health care and social assistance	18	22	+ 3.8	+ 20
Other	10	11	+ 0.5	+ 1
Total	151	160	+ 9.8	+ 6

*Forecast

degree, with an upgraded curriculum. Since then the nursing profession has risen in every dimension—salaries, status, academic requirements, range of responsibilities, and so forth. These days, the nursing profession spans a range of specializations, whereby the upper echelon commands annual wages as high as $100,000. Moreover, nurses now use advanced technologies, and these in turn contribute to upgrade the profession.

It could have been otherwise had it not been for the legislation of 1964, and so it is for other occupations in personal services. Thus, we need to consider proactive strategies for the *professionalization of personal services*, particularly in health care and education, setting standards and academic requirements.

Take for example early childhood education: in most countries there are virtually no such standards for caregivers of children age one to three, precisely the ages that are crucial for their development. Suppose now that they were required to have specialized academic degrees, with a curriculum that would include psychology, brain development, testing for learning disabilities, and so forth. Not only would the status and wages of these workers increase, but they would be much more likely to benefit from complementary advanced technologies.

The advent of AI would probably not threaten these growing occupations, and furthermore, if they were upgraded in the way just described, AI could bestow significant benefits to them as well. For that to happen smart interfaces between the practitioners of these occupations and the AI machines will have to be developed. Thus, imagine, for example, professional caregivers using AI to test very young children for learning disabilities, and then for treating them with specially tailored AI-based games.

To sum up: BLS projections indicate that the bulk of job creation in the decade to 2024 will be in personal services, particularly in personal care. As currently practiced, most of these occupations are at the low end of the scale and rather impervious to technological advances. However, there are viable options to upgrade these occupations, particularly by setting academic standards and advanced curriculums. If that were to happen, then

the changing composition of employment (i.e., more personal care, less of many others) would not adversely affect income distribution but perhaps to the contrary; furthermore, and more importantly here, AI may play a complementary role vis à vis these occupations, thus raising productivity in services and triggering a virtuous cycle.

6.3.3 The Direction of Technical Change: H-*Enhancing* or H-*Replacing*?

Although one of the seminal volumes in the economics of technological change is titled "The Rate and Direction of Inventive Activity," in fact the economic discipline has traditionally dealt much more with the "rate" than with the "direction." That may come as no surprise, since discussing direction requires getting into the guts of technology itself, and there is no reason to believe that economists have a comparative advantage in that regard.

Nevertheless, the extent and scope of technological advances that engulf us may require us to look more closely into the "black box" and try to understand, at the very least, what types of innovations we are facing and how they impact the economy. Furthermore, we would like to know whether there is room to affect the relative prevalence of the various types, in view of their differential economic effects.

Here is such an attempt: consider on the one hand innovations that mostly magnify, enhance, and extend sensory, motoric, analytical, and other human capabilities such as:

- In medicine: AI for diagnostics, for example, for reading and interpreting x-rays, CT scans and other imaging modalities; AI for robotic surgery (e.g., the da Vinci robot for prostate surgery); AI data mining of electronic medical records for follow-up evaluations of drug efficacy post-Food and Drug Administriation (FDA) approval, and so forth.
- In education: AI-based methods for "personalized teaching"; AI for online testing in MOOCS; (see also the above-mentioned applications for early childhood education), and so forth.

We label these "human-enhancing innovations" (HEI)—in medicine they do not replace doctors, but rather augment their human-bound capabilities (think of the precision and consistency of robot surgery), thus making better doctors. Similarly for teachers, eventually for judges (ruling with the aid of AI-based analysis), and so forth.

On the other hand, consider "human-replacing innovations" (HRI), that is, technical advances that *replace* human intervention, and furthermore that often leave for humans mostly "dumb" jobs that are not worth yet replacing given the very low wages that they command (and often are indeed difficult to replicate by machines, the proverbial one being janitors).

Some HRIs lead to cutting-edge, virtually human-free factories (best exemplified by Tesla's new facilities to produce batteries for its e-cars) that

greatly improve productivity, even if reducing employment. Consider, however, the polar case of Walmart, the world's largest private employer (with over two million employees), having deployed advanced technologies along its whole chain of operations from logistics to retailing; it has turned a large proportion of its workers into "unthinking automatons," commanding very low wages with no prospect for improvement.

These then are two types of innovations (HEI and HRI) that have very different effects on key economic and social variables. It would seem that AI-based HEIs have the potential to unleash a new wave of *human creativity and productivity*, particularly in services (which to repeat are expected to be the fastest-growing occupations), whereas HRIs either decrease employment (e.g., Tesla), or create unworthy jobs.

Is it possible to design strategies to affect the direction of technical change in the sense of stimulating HEIs versus HRIs? It is hard to say, but it is certainly worthwhile investigating such possibility given the large impact that a change in direction may have on the economy. Incidentally, it would seem that in any case the traditional emphasis of economic policy on the "rate" of innovation, that is, on how much resources we devote to research and development (R&D), is misplaced—worldwide competition may be pushing us into too much investment in R&D, not too little (too many patents, too much replication, etc.). Some attention to the "direction" may bring much larger returns.

6.4 Concluding Remarks

The historical record suggests that dismal prophecies about the economic and social impact of great technological advances rarely come to pass. Thus, with AI poised to emerge as the new GPT, we should not necessarily envision a future whereby humans will be rendered obsolete and mass unemployment will be the "new normal." At the same time, as many occupations will indeed vanish, and many others will undergo significant changes, it is important to inquire into what sort of strategies may ameliorate the detrimental effects of AI and enhance the positive ones. This is all the more important given that in the twenty-first century the public at large has much less tolerance for bearing the costs of technical change and higher expectations for sharing into its benefits here and now.

Therefore, we need to anticipate the required institutional changes, experiment in the design of new policies (particularly in education and skills development) in the professionalization of service occupations, and in affecting the direction of technical advance. Furthermore, economists possess a vast methodological arsenal that may prove very useful for that purpose—we should not shy away from stepping into this area, since its importance for the economy cannot be overstated.

References

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

Gordon, Robert J. 2016. *The Rise and Fall of American Growth*. Princeton, NJ: Princeton University Press.

Heckman, James, Tim Kautz, Ron Diris, Bas ter Weel, and Lex Borghans. 2014. "Fostering and Measuring Skills: Improving Cognitive and Non-Cognitive Skills to Promote Lifetime Success." Report prepared for the Organisation of Economic Co-operation and Development, Paris. http://www.oecd.org/edu/ceri/Fostering-and-Measuring-Skills-Improving-Cognitive-and-NonCognitive-Skills-to-Promote-Lifetime-Success.pdf.

Hirschman, Albert. 1970. *Exit, Voice and Loyalty*. Cambridge, MA: Harvard University Press.

Mokyr, Joel. 2017. "The Past and the Future of Innovation: Some Lessons from Economic History." Paper presented at the conference on the Economics of Artificial Intelligence, Sept. 2017, Toronto.

Summers, Lawrence H. 2016. "The Age of Secular Stagnation: What It Is and What to Do About It." *Foreign Affairs*, Feb. 15. http://larrysummers.com/2016/02/17/the-age-of-secular-stagnation/.