

The Wrong Kind of AI?

Daron ACEMOGLU
MIT

Pascual RESTREPO
Boston University

Artificial Intelligence and the Future of Labor Demand



Dear friends,

We received this white paper on AI and its impact of job creation by TNIT member **Daron Acemoglu** of MIT and **Pascual Restrepo** of Boston University. We were so impressed and excited that we could not resist immediately sending it to you as an early holiday season gift. Daron and Pascual argue that AI can be the basis of two types of technological progress: I) automation, which displaces human activity and II) enhancement, which complements and enriches human work. They argue that there is scope for public policy to ensure that resources are allocated optimally between these two types of AI in order to ensure fulfillment of its potential for growth, employment and prosperity.

We hope you enjoy this exciting read and we wish you a very happy holiday season and new year,

Jacques Crémer,
TNIT Coordinator

Yassine Lefouili,
TSE Digital Center Director

Artificial Intelligence (AI) is one of the most promising technologies currently being developed and deployed. Broadly speaking, AI refers to the study and development of “intelligent (machine) agents”, which are machines, softwares or algorithms that act intelligently by recognizing and responding to their environment¹.

There is a lot of excitement, some amount of hype, and a fair bit of apprehension about what AI will mean for our future security, social lives and economy. But a critical question may have been largely overlooked: are we investing in the “right” type of AI for generating prosperity? We are not currently in a position to provide definitive answers to this question - nobody is.

But this does not mean that it should not be asked. In fact, this may be the right time to ask it, while we still have the ability to think, learn and take actions about these issues.

AI as a Technological Platform

Human (or natural) intelligence comprises several different types of mental activities. These include, among others, simple computation, data processing, pattern recognition, prediction, hand-eye coordination, various types of problem solving, judgment, creativity, and communication. Early AI, pioneered in the 1950s by researchers from computer science, psychology and economics, such as Marvin Minsky, Seymour Papert, John McCarthy, Herbert Simon and Allen Newell, aimed at developing machine intelligence capable of performing all of these different types of mental activities.² The goal was nothing short of creating truly intelligent machines.

Herbert Simon, for example, claimed in 1958 *“there are now in the world machines that think, that learn and that create. Moreover, their ability to do these things is going to increase rapidly until - in a visible future - the range of problems they can handle will be coextensive with the range to which the human mind has been applied.”*

These ambitious goals were soon dashed, however. AI came back into fashion starting in the 1990s, but with a different and in some ways more modest ambition: to replicate and then improve upon human intelligence in pattern recognition and prediction (pre-AI computers were already better than humans in computation and data processing). Many decision problems and activities we routinely engage in can be viewed as examples of pattern recognition and prediction. These include recognizing faces (from data we obtain visually), recognizing speech (from data we obtain by hearing speech), recognizing abstract patterns in data we are presented with, making decisions on the basis of past experience and current information, to name just a few. Though there are researchers working on “Artificial General Intelligence”, most of the work that goes under the name of AI and almost all commercial applications of AI are in these more modest domains and is thus sometimes (and appropriately) referred to as “Narrow AI” - even if the relevant applications are numerous and varied.

The big breakthroughs and the renewed excitement in AI are coming from advances in hardware and algorithms that enable the processing and analysis of vast amounts of unstructured data (for example, huge amounts of speech data, which are not labeled in a form that facilitates their processing and cannot be represented or processed in the usual structured ways, such as in simple, Excel-like databases. Central to this renaissance of AI have been methods of machine learning (which are the statistical techniques that enable computers and algorithms to learn, predict and perform tasks from large amounts of data without being explicitly programmed) and what is called “deep learning” (which involves the development of algorithms using multi-layered programs, such as neural nets, for improved machine learning, statistical inference and optimization).

A critical observation is that even if we focus on its narrow version, AI should be thought of as a technological platform - there are many ways in which AI technology can be developed as a commercial or production technology, with widely varying implications. To some degree, this is true of all clusters of technologies, but it is more emphatically so for AI. This can be seen, for instance, by comparing AI to a related but distinct new technology, robotics. Robotics is distinguished from other digital technologies by its defining focus on interacting with the physical world (moving around, transforming, rearranging or joining objects), even though it often makes use of AI and other digital technologies for processing data

as well. Industrial robots are already widespread in many manufacturing industries and in some retail and wholesale establishments. But their economic use is quite specific, and centers on automation, that is, the substitution of machines for certain tasks previously performed by labor.³

The Implications of New Technologies on Work and Labor

How new technologies impact the nature of production and work, and consequently, the employment and wages of different types of workers, is one of the central questions of the discipline of economics. The standard approach, both in popular discussions and academic writings, presumes that any advance that increases productivity (value added per worker) also tends to raise the demand for labor, and thus employment and wages. It is recognized, of course, that technological progress might benefit workers with different skills unequally and productivity improvements in one sector may lead to job loss in that sector. But even when there are sectoral job losses, the standard narrative goes, other sectors will expand and contribute to employment and wage growth for all workers.

This view is critically underpinned by the way in which the economic effects of new technology is usually conceptualized - as enabling labor to become more productive in pretty much all of the activities and tasks that it performs. Yet, this description of technologies not only lacks descriptive realism (which technology makes labor uniformly more productive in everything?), but may end up painting an excessively rosy picture of the implications of new technologies. Indeed, in such a world Luddites’ concerns about the disruptive and job displacing implications of technology would be misplaced, and they would have smashed all of those machines in vain.

The reality of technological change is rather different. Many new technologies - those we will call automation technologies - do not increase labor’s productivity, but are explicitly aimed at replacing it by substituting cheaper capital (machines) in a range of tasks performed by humans⁴. As a result, automation technologies always reduce the labor’s share in value added (because they increase productivity by more than wages and employment). They may also reduce overall labor demand because they are displacing workers from the tasks they were previously performing. The reason why they need not necessarily reduce labor demand, even if labor demand fails to keep up with productivity, is that some of the productivity gains may translate into greater demand for labor still employed in non-automated tasks as well as in other sectors producing complementary goods and services.

This discussion clarifies an important point: in an age of rapid automation, labor would be particularly badly affected if new technologies are not raising productivity sufficiently - if these new technologies are not great but just “so-so” (just good enough to be adopted but not so much more productive than the labor they are replacing). With so-so technologies, labor demand declines because of the displacement that automation creates, but does not rebound due to the lack of powerful productivity gains.

Is this far-fetched? Not really. We have previously studied the implications of one of the most important automation technologies, industrial robots⁵. Industrial robots are not technologies aimed at increasing labor’s productivity but are designed to automate tasks that were previously performed by production workers on the factory floor. The evidence is fairly clear that industries where more industrial robots are introduced experience declines in labor demand (especially for production workers) and sizable falls in their labor share. More importantly, local labor markets more exposed to the introduction of industrial robots, such as Detroit MI or Defiance OH, have significantly lower employment and wage growth. All of this is despite the fact that industry-level data also suggest significant productivity gains from the introduction of robots.

Automation is not a recent phenomenon. Many important breakthroughs in the history of technology have been centered around automation. Most notably, the spectacular advances in the early stages of the Industrial Revolution in Britain were aimed at automating weaving and spinning, and the focus on automation then shifted to the factory floors of other industries⁶. The mechanization of agriculture and the interchangeable parts system of American manufacturing are other prominent examples of automation.

But if automation tends to reduce the labor share and has mixed effects on labor demand, why did productivity growth go hand-in-hand with commensurate wage growth over the last two centuries? The answer is that at the same time as automation technologies are being introduced, other technological advances enable us to create new tasks in which labor has a competitive advantage. This generates new activities for labor - tasks in which it can now be reinstated -

¹ / See Russell and Norvig (2009) and Neapolitan and Jiang (2018).

² / See Nilsson (2009) for the history of AI.

³ / See Ayres and Miller (1983), Groover et al. (1986), Graetz and Michaels (2015) and Acemoglu and Restrepo (2018).

⁴ / This approach is developed in Zeira (1998), Autor, Levy and Murnane (2003), Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018a,b).

⁵ / Acemoglu and Restrepo (2018b).

⁶ / See Mantoux (1927) and Mokyr (1989).

and robustly contributes to productivity growth as new tasks improve the division of labor and increase the productive complexity of the production process⁷. Indeed, occupations in which there are new tasks have been at the forefront of employment growth in the US economy over the last three decades, and historical evidence documents the crucial role of new occupations and industries in labor demand growth during the 19th and early 20th centuries.

This perspective then suggests a different reinterpretation of the history of technology and a different way of thinking about the future of work - as a race between automation and new, labor-intensive tasks. Labor demand has not increased steadily over the last two centuries because of technologies that have made labor more productive in everything. Rather, many new technologies have sought to eliminate labor from tasks in which it previously specialized. All the same, labor has benefited from this process both because, even as they were automating previously labor-intensive tasks, these technologies have increased productivity, and more importantly because at the same time other technologies have enabled new labor-intensive tasks to be introduced. These new tasks have done more than just reinstating labor as a central input into the production process; they have also likely played a vital role in productivity growth.

Viewed from this perspective, if employment and wage growth has been disappointing over the last two decades, this is in part because productivity growth has been weak, and also in large part because the introduction of new tasks has been sluggish. The future of work will be much brighter if we can mobilize more of the technologies that increase labor demand and ensure healthier productivity growth.

Varieties of AI

This perspective provides a new way of thinking about the economic opportunities and challenges posed by AI. Most AI researchers and economists studying its consequences view it as a way of automating yet more tasks. No doubt, AI has this automating capability, and most of its applications to date have been of this mold - e.g., image recognition, speech recognition, translation, accounting, recommendation systems, product placement in advertising, and customer support. But there are two problems in accepting this as the primary way that these technologies can be and indeed ought to be used.

First, if all we do is continue down the path of automation, with no counterbalancing innovations to generate new tasks, the implications for labor are depressing. It will not be the end of work anytime soon⁸, but the trend towards lower labor share and anemic growth of labor demand will continue - with potentially disastrous consequences for income inequality and social cohesion.

Second, as we go deeper and deeper into automation, based on AI, robotics and other technologies, we are moving into areas in which human labor is quite good (for example think of image and speech recognition or hand-eye coordination), and machine productivity, at least to start with, is not always impressive, to say the least. Automation technologies aimed at substituting machines for humans in these tasks are thus likely to be of the so-so kind. As a result, we cannot even count on powerful productivity gains to increase our living standards and contribute to labor demand.

But it doesn't have to be this way. Since AI is not just a narrow set of technologies with specific, pre-determined applications and functionalities but a technological platform, it can be deployed for much more than just for automation - it can be used for restructuring the production process in a way that creates many new, high-productivity tasks for labor. If this type of "reinstating AI" is a possibility, there would be potentially very large societal gains both in terms of improved productivity and greater labor demand (which will not only create more inclusive growth but also avoid the social problems created by joblessness and wage declines).

Consider a few examples of the ways in which AI applications can create new tasks for labor. Notably, such new tasks can go well beyond those that are sometimes emphasized as "enablers" of AI (human tasks involved in training and monitoring new machines as they automate what the rest of us do).

- **EDUCATION:** Education is one of the areas in which there has been the least penetration of AI. That may be partly due to the fact that automation is not an attractive or even feasible option for most of the core tasks in education. But using AI to create new tasks would be a different way of deploying this new technology in a fruitful and productive way. Consider, for example, classroom teaching. This has not changed for over 200 years. A teacher teaches to the whole class, even if occasionally he or she or an aide may also engage in one-on-one instruction or provide help for some subset of students. There is evidence, however, suggesting that many students have different

"learning styles", and what works for one student will not work for another, and even what works for one student in one subject will not work for him or her in another subject.⁹

At the moment, individualized teaching, targeted and adaptively adjusted for each student or for small subsets of students, is impossible, and not just because the resources in terms of teacher time, not to mention teacher skill, are not there in schools; it is also because nobody has the information (and cannot easily acquire and process the information) to determine the optimal learning style of a student in a specific subject or for a specific topic within a subject. This is something that can change with AI. It is quite feasible to design AI software and other interfaces to collect and process data about the specific reactions, difficulties and successes of students to different subject areas, especially when taught or expositied in different styles, and then make recommendations for improved individualize teaching. The potential improvements in terms of educational productivity could be quite large (we just don't know). Societal benefits could exceed these direct benefits as AI-powered teaching methods may do better in terms of providing students with skills that will be more valued in future labor markets (rather than the more backward-looking curricula and teaching emphasis currently prevailing in schools). The development and deployment of such technologies would also increase the demand for human labor in teaching - there would be need for more teachers to do the individualized teaching, even with help from AI software and other technologies.

- **HEALTHCARE:** The situation in healthcare is very similar. Though there has been more effort devoted to introducing digital technologies into healthcare, the focus has not been in creating new tasks (in fact, some of the uses of AI, for example in radiology, are very much in the mold of automation). New, adaptable AI applications that collect information can significantly empower nurses, technicians and other healthcare providers to offer a wider range of services and more real-time health advice, diagnosis and treatment. The benefits in terms of greater labor demand and productivity are very similar to the education case.

- **AUGMENTED REALITY:** The third area in which the use of AI can significantly change the production process in a way that may be favorable to labor is in the use of augmented reality (including virtual reality) technologies in manufacturing. The most advanced manufacturing technologies of the last three decades have focused on automation. But companies such as Amazon and Tesla have also discovered that automating all factory-floor and manual tasks is not economically rational, because some tasks are still more productively performed by human labor. One difficulty facing companies introducing industrial robots, however, is that these new technologies do not necessarily work very well together with humans for at least two reasons.

First, most of robotics technology is cordoned off from workers because of safety concerns. Second, human work may not mesh very well with the degree of precision required and achieved by robotics technology. Augmented reality technologies - which use interactive interfaces in order to increase the ability of humans to perceive, monitor and control objects - have the promise of enabling workers to work alongside machines and improve their ability to perform high precision production tasks. This will not just help workers keep some of the tasks that we might have otherwise been automated; it could also create new tasks in which humans, augmented by digital technology and sensors, can be employed and contribute to productivity¹⁰.

Why the Wrong Kind of AI?

If there are potentially productive and profitable uses of AI beyond simple automation, can we then count on market forces and the innovation strategies of existing companies to take us there? Should there be any reason to worry about productive, reinstating applications of AI not being exploited and resources just pouring into the wrong kind of AI?

Economists tend to have great trust in the ability of the market mechanism to allocate resources in the right, most efficient way. But most experts also recognize that the magic of the market ceases to shine as brightly when it comes to innovation. There are several reasons for this in general, as well as some specific reasons that are specifically important in the context of AI.

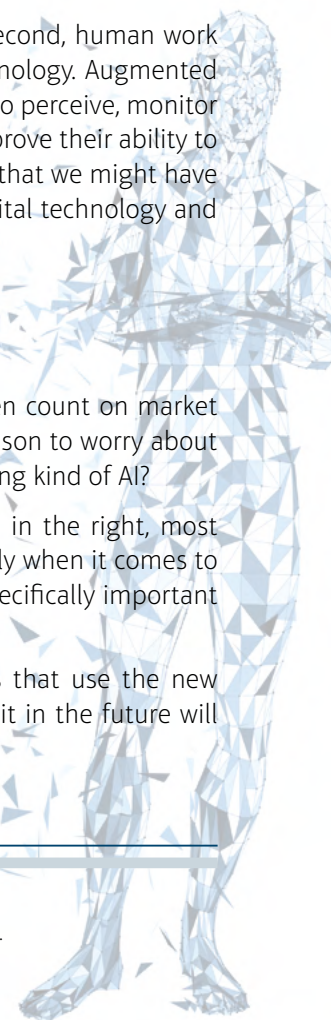
- Innovation activities generally create externalities - not just the innovator, but the workers that use the new technology, firms that deploy it, and most importantly other firms and researchers building on it in the future will benefit from it. Markets do not generally do a good job in the presence of such externalities.

⁷ / Acemoglu and Restrepo (2018a).

⁸ / See Dreyfus (1992) and Autor (2015).

⁹ / See Allport (1937), Cassidy (2004) Honey and Mumford (1986) and Ramirez and Casteneda (1974).

¹⁰ / See Ong and Nee (2013) and <https://www.ge.com/reports/game-augmented-reality-helping-factory-workers-become-productive/>.



► Particularly problematic for market allocations is when there are alternative, competing technological paradigms. When one paradigm is ahead of the other, it may be in the interest of both researchers and companies to follow that paradigm, even if the alternative one could be more productive. Moreover, in such a situation, once the wrong paradigm pulls ahead, it may be very difficult to reverse this trend and benefit from the possibilities offered by the alternative paradigm. To the extent that different types of approaches to AI constitute such alternative, competing paradigms, our trust in the market mechanism getting it right should be even lower.¹¹

► The way in which research effort is allocated towards areas of high priority often includes public-private partnership, for example, governments steering the research direction by grants and by creating demand for different types of new technologies and products. The US government has traditionally been heavily involved in many of the leading technologies, including the Internet, sensors, pharmaceuticals, biotech and nanotechnology¹².

But more recently, the US government has been more timid its support for research overall and in its determination to steer the direction of technological change. This shift is likely to be a result of a combination of factors, including the reduction in resources devoted to government support of innovation and the increasingly dominant role of the private sector in setting the agenda in high-tech areas (can government officials and researchers meaningfully influence the direction of inventive activity in Silicon Valley?). This shift will continue to make it more difficult for factors related to future promise (that are not immediately reflected in profitability) and other social objectives (such as employment creation) to influence the direction of technological change.

► Innovation is not a purely economic activity. Several noneconomic rewards affect what types of technologies attract the attention and imagination of researchers. It is possible that the ecosystem that has emerged around the most creative clusters in the United States, such as Silicon Valley, excessively reward automation, and pays insufficient attention to uses of frontier technologies for other purposes. This may be partly because of the values and interests of leading researchers (consider for example the ethos of companies like Tesla that have ceaselessly tried to automate, and likely ended up excessively automating their production processes), and partly because of the prevailing business model and vision of the companies, such as Google (Alphabet), Facebook, Amazon, Microsoft and Netflix which are the source of most of the resources going into AI, has focused on automation and removing the (fallible) human element from the production process.

This last consideration may have become even more critical recently as the vast resources of several leading companies are pouring into academia and impacting the teaching and research missions of leading universities, and as a result of this and of the attractive job market opportunities, the best minds in the current generation are gravitating towards computer science, AI and machine learning, but with a close to singular focus on automation. An ecosystem that is biased would become much more stifling for the direction of technological change when it becomes all-encompassing.

► The last two factors may have played a particularly important role in swinging the balance away from the creation of new tasks and towards greater focus on automation over the last two decades. With companies whose workforces are miniscule relative to their assets and whose business model relies on the production and marketing of digital goods and services setting the research agenda for the entire nation or even the world, and with the US government abdicating its role of shaping the direction of technological change, the market process may have gone down the path of greater and greater focus on automation while technologies with the capacity to increase the demand for labor fell by the wayside.

There are also additional factors that may have specifically distorted choices over what types of AI applications to develop. The first one is that if employment creation has a social value beyond what is in the GDP statistics (for example, because employed people are happier and become better citizens, or because faster growth of labor demand improves income inequality), this social value will be ignored by the market mechanism. The second is related to the tax policies adopted in the United States and other Western nations, which subsidize capital and investment while taxing the use of workers. This would make the use of machines instead of labor more profitable, and these profits will encourage not just automation but also faster advancement of automation technologies. Finally, and complementing these factors, to the extent that firms take into account the cost of labor (wages), which tend to be higher than the social opportunity cost of labor - due to imperfections and noncompetitive elements in the labor market - they will have additional powerful reasons for adopting automation technologies beyond what is socially optimal. This once again implies greater incentives for developing automation technologies¹³.

► Additional factors might be blocking the way of reinstating applications of AI. Take the example of education mentioned above. It is not only that developing AI to create new labor-intensive tasks in education is not viewed as being the frontier or one of the “cool” areas, say compared to facial recognition applications. It is also that it is not an area that is on the radar screen of the biggest players in the AI space. Moreover, the complementary investments and resources to make this type of reinstating AI profitable may be missing completely. Educational applications of AI, for example, would make economic sense for developers only if there will be resources to hire more teachers to work with these new AI technologies (after all, that is the point of the new technology, to create new tasks in teaching, which will have to be filled with new teachers with the right skills). But without the resources - which would have to come mostly from governments and perhaps some from philanthropic private institutions - to do this, the demand for these AI technologies will remain dormant. In the case of healthcare, limited resources are not the problem (the share of national income devoted to health is continuing to grow). But highlighting other barriers to the use of new technologies to create new tasks, the way that hospitals, insurance companies and the whole medical profession (especially with the overwhelming and stifling control of the American Medical Association) is organized is likely to be in the way. If empowering, and increasing the productivity of, nurses and technicians is perceived to reduce the dominance and the demand for the services of doctors and challenge the current business model of hospitals, it will be strenuously resisted.

All in all, even though we currently lack any definitive evidence that research and corporate resources today are being directed systematically towards the wrong kind of AI, the way the market for innovation is organized there is no compelling reason to expect that the efficient balance between different uses of any technology, let alone AI, will be struck. Concerns that much of the energy and focus being directed to the wrong subareas become amplified when one sees the best minds in universities today pouring into data science, machine learning and AI fields with an almost singular focus on automating everything. If at this critical juncture the outcome ends up being insufficient attention devoted to inventing and creating demand for rather than just replacing labor, that would be the wrong kind of AI from the social and economic point of view, potentially with dismal consequences for growth, employment, inequality and our prosperity.

CONCLUSION

Artificial Intelligence is set to influence every aspect of our lives, not least the way production is organized in modern economies. But there should be no presumption that left to its own devices, the right types of AI technologies will be developed and implemented. Though many today worry about the security risks and other unforeseen (often non-economic) consequences of AI, we have argued that there are prima facie reasons for worrying about the wrong kind of AI from an economic point of view becoming all the rage and the basis of future technological development. The considerable promise of AI implies that we need care and serious thought to be devoted to its implications and to the question of how to best develop this promising technological platform - before it is too late.



¹¹ / See Nelson and Winter (1977), Dosi (1982) and Acemoglu (2012).

¹² / Mazzucato (2015).

¹³ / Acemoglu and Restrepo (2018a).



The Wrong Kind of AI?

REFERENCES

- Acemoglu, Daron (2002) *“Technical Change, Inequality, and The Labor Market”*, Journal of Economic Literature, 40(1): 7–72.
- Acemoglu, Daron (2012) *“Diversity and Technological Progress”*, The Rate and Direction of Inventive Activity Revisited, University of Chicago Press, pp. 319–360.
- Acemoglu, Daron and David Autor (2011) *“Skills, Tasks and Technologies: Implications for Employment and Earnings”*, Handbook of Labor Economics, 4: 1043–1171.
- Acemoglu, Daron and Pascual Restrepo (2018a) *“The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment”*, American Economic Review, 108(6): 1488–1542.
- Acemoglu, Daron and Pascual Restrepo (2018b) *“Robots and Jobs: Evidence from US Labor Markets”*, NBER Working Paper No. 23285.
- Acemoglu, Daron and Pascual Restrepo (2018c) *“Modeling Automation”*, NBER Working Paper No. 24321.
- Acemoglu, Daron and Pascual Restrepo (2018d) *“Artificial Intelligence, Automation and Work”*, NBER Working Paper No. 24196.
- Allport, Gordon W. (1937) *“Personality: A Psychological Interpretation”*, New York: Holt and Co.
- Autor, David H. (2015) *“Why Are There Still So Many Jobs? The History and Future of Workplace Automation”*, Journal of Economic Perspectives, 29(3): 3–30.
- Autor, David H., Frank Levy and Richard J. Murnane (2003) *“The Skill Content of Recent Technological Change: An Empirical Exploration”*, The Quarterly Journal of Economics, 118(4): 1279–1333.
- Ayres, Robert U. and Steven M. Miller (1983) *“Robots: Applications and Social Implications”*, Ballinger Publishing Co.
- Cassidy, Simon (2004) *“Learning Styles: An Overview of Theories, Models, and Measures”*, Educational Psychology, 24(4): 419–444.
- Daugherty, Paul and H. James Wilson (2018) *“Human + Machine: Reimagining Work in the Age of AI”*, Harvard Business Review Press.
- Dreyfus, Hubert L. (1992) *“What Computers Still Can’t Do: A Critique of Artificial Reason”*, MIT Press.
- Dosi, Giovanni (1982) *“Technology Paradigms and Technological Trajectories”*, Research Policy, 11(3): 147–162.
- Graetz, Georg and Guy Michaels (2018) *“Robots at Work”*, The Review of Economics and Statistics, forthcoming.
- Groover, Mikell P., Mitchell Weiss, Roger N. Nagel, and Nicholas G. Odrey (1986) *“Industrial Robotics: Technology, Programming, and Applications”*, McGraw-Hill.
- Honey, Peter and Alan Mumford (1986) *“The Manual of Learning Styles”*, Peter Honey Associates.
- Kellner, Tomas (2018) *“Game On: Augmented Reality Is Helping Factory Workers Become More Productive”*, www.ge.com/reports/game-augmented-reality-helping-factory-workers-become-productive/, April 19, 2018.
- Mantoux, Paul (1927) *“The Industrial Revolution in the Eighteenth Century: An Outline of the Beginnings of the Modern Factory System in England”*, Harcourt, Brace & Co.
- Mazzucato, Mariana (2015) *“The Entrepreneurial State”*, Anthem Press.
- Mokyr, Joel (1989) *“Industrial Robotics: Technology, Programming, and Applications”*, Rowman & Littlefield Publishers.
- Neapolitan, Richard E. and Xia Jiang (2018) *“Artificial Intelligence: With an Introduction to Machine Learning”*, Chapman and Hall/CRC, Second Edition.
- Nelson, Richard and Sidney, Winter (1977) *“In search of useful theory of innovation”*, Research Policy, 6(1): 36–76.
- Nilsson, Nils J. (2009) *“The Quest for Artificial Intelligence: A History of Ideas and Achievements”*, Cambridge University Press.
- Ong, oh K. and Andrew Y.C., Nee (2013) *“Virtual and Augmented Reality Applications in Manufacturing”*, Springer Science and Business Media.
- Ramírez Manuel and Alfredo Castañeda, (1974) *“Cultural Democracy, Bicognitive Development, and Education”*, Academic Press.
- Russell, Stuart J. and Peter Norvig (2009) *“Artificial Intelligence: A Modern Approach”*, Prentice Hall, Third Edition.
- Zeira, Joseph (1998) *“Workers, Machines, and Economic Growth”*, Quarterly Journal of Economics, 113(4): 1091–1117.

We thank David Autor, Erik Brynjolfsson and Stu Feldman for useful comments.