



# Department of Economics Discussion Paper Series

## Re-Thinking the Capabilities of Machines in Economics

Daniel Suskind

Number 825

Originally Published: May 2017

Revised: April 2018

# Re-Thinking the Capabilities of Machines in Economics

Daniel Susskind\*

This version: 12 April 2018

## Abstract

The literature exploring the effects of technological change on the labour market often relies on a very particular understanding of the capabilities of machines – known as the ‘ALM hypothesis’. However, this hypothesis has often led this literature to underestimate these capabilities. Tasks that were believed to be out of reach of automation can now be automated. I set out two explanations for this underestimation – one that is explored in the recent literature and maintains the ALM hypothesis, and a new explanation that challenges it. I propose a new hypothesis that contains the ALM hypothesis as a special case.

Keywords: Technological Change; Computerization; Automation; Job Tasks; Wages.

JEL: J20; J21; J23; J24; J30; J31; O31; O33.

## 1 Introduction

Accurately forecasting the future capabilities of machines is very difficult. The traditional ‘task-based’ literature that explores the effect of technological change on the labour market has at times underestimated them. For

---

\*Fellow, Balliol College, Oxford University. This is a revised version of Chapter II of my doctoral dissertation, ‘Technology and Employment’. Thanks to David Autor, Brian Bell, Martin Ellison, Frank Levy, Alan Manning, Margaret Stevens, Larry Summers, Richard Susskind, and David Vines for helpful input. Contact: daniel.susskind@balliol.ox.ac.uk.

instance, Autor et al. (2003) noted that the task of driving a car could not be readily automated, but a type of driverless car appeared two years later;<sup>1</sup> Autor et al. (2013) noted that order-taking and table-waiting could not be readily automated, but later that year the US restaurants Chili's and Applebee's announced they were installing 100,000 tablets to allow customers to order and pay without a human waiter;<sup>2</sup> Autor (2015) noted that the task of identifying a species of bird based on a fleeting glimpse could not be readily automated, but later that year an app was released to do that as well.<sup>3</sup>

These examples suggest that the traditional literature's conception of how machines operate and the capabilities that this implies – known as the 'ALM hypothesis' – may be incorrect.<sup>4</sup> These tasks were believed to be out of reach of automation because they were 'non-routine' rather than 'routine'. But in practice this has proven not to be the case.<sup>5</sup> The initial contribution of this paper is to set out two distinct explanations for why this might be so. The first explanation is that proposed in the recent literature (for instance, Autor 2014 and 2015, and Remus and Levy 2016). It maintains the ALM hypothesis, and argues that tasks which we thought were 'non-routine' are, in fact, more 'routine' than we realised. Here, the role of technology is to move the boundary between 'non-routine' and 'routine' tasks. The second is a new explanation. It challenges the ALM hypothesis, and argues that the 'non-routine' and 'routine' distinction itself is no longer a compelling way to think about the constraints on automation. Here, the role of technology is instead to allow machines to perform tasks in such a way that the 'non-routine' and 'routine' distinction no longer applies. This, I argue, is an increasingly compelling explanation. A further contribution of this paper is to propose a new hypothesis about the capabilities of machines, based on this second explanation, that nests

---

<sup>1</sup>The Society of Automotive Engineers defines five levels of vehicle 'autonomy'. These early cars were at a low level. Since 2005, further progress has been made. I thank Frank Levy for this point.

<sup>2</sup>See Pudzer (2016).

<sup>3</sup>See <http://merlin.allaboutbirds.org/photo-id/>.

<sup>4</sup>'ALM', after 'Autor, Levy and Murnane' – the authors of Autor et al. (2003).

<sup>5</sup>It is important to note that the traditional literature did recognise that the ALM hypothesis may not hold indefinitely. Autor et al. (2003), for instance, refers to "*present* technologies" and says that the ALM hypothesis is only "*at present*" a binding constraint (emphasis added).

the ALM hypothesis as a special case.

## 2 The Traditional Approach

### 2.1 Context

Autor et al. (2003) provides the foundations for the new literature exploring the consequence of technological change on the labour market. Autor et al. (2003) is motivated by a disappointment with the dominant ‘skills-biased technical change’ (SBTC) thesis.<sup>6</sup> The SBTC thesis has empirical support.<sup>7</sup> But the SBTC thesis does not explain *why* technology has this ‘skills-biased’ property. As Autor et al. (2003) notes:

“It fails to answer the question of what it is that computers do – or what it is that people do with computers – that causes educated workers to be relatively more in demand.” (pg. 1280)

The purpose of Autor et al. (2003) is to answer this question and to provide a deeper explanation for the SBTC thesis. Autor et al. (2003) introduces two innovations to do this. The first is the ‘task-based’ approach. This breaks production into two stages – ‘*factor-based production functions for tasks*’ that describe how *factors* with particular *skills* and *capabilities* combine to perform different types of *task*, and ‘*task-based production functions for goods*’ that describe how different types of *tasks* are combined to produce different types of *good*. When these two distinct sets of functions are taken together, the resulting aggregate function describes how particular *factors* combine to produce different types of *goods*.<sup>8</sup> The second innovation is a theory about the capabilities of “computer capital”.

### 2.2 The ALM Hypothesis

The ‘ALM hypothesis’ is the theory about capabilities of computer capital developed in Autor et al. (2003). Based on “an intuitive set of observa-

---

<sup>6</sup>For an overview of the SBTC thesis, see Acemoglu (2002).

<sup>7</sup>See Autor et al. (2003) for the “wealth of quantitative and case-study evidence” to support it.

<sup>8</sup>Goos, Manning and Salomons (2014) describe this as a “two-stage setup”.

tions” from economists and others, Autor et al. (2003) make two assumptions:

“(1) that computer capital substitutes for workers in carrying out a limited and well-defined set of cognitive and manual activities, those that can be accomplished by following explicit rules (what we term “routine tasks”); and (2) that computer capital complements workers in carrying out problem-solving and complex communication activities (“nonroutine tasks”).” (pg. 1280)

These two assumptions are critical. Together, this distinction between ‘routine’ and ‘non-routine’ tasks, and these two assumptions with respect to that distinction, are the ALM hypothesis. A *weak* interpretation of the ALM hypothesis is that it implies machines are more capable at performing ‘routine’ tasks than ‘non-routine’ tasks. A *strong* interpretation of the ALM hypothesis is that it implies machines are capable of performing ‘routine’ tasks but cannot perform ‘non-routine’ tasks. In the literature both interpretations are used, often within the same paper. But the models that rely on the ALM hypothesis mostly use the strong interpretation (see, for instance, Autor et al. 2003, Autor and Dorn 2013, and Goos, Manning and Salomons 2014). This is the interpretation I will focus on.

To see why the ALM hypothesis no longer holds, two questions must be answered. First, what is the basis for the ‘routine’ vs. ‘non-routine’ distinction? Secondly, what is the basis for the assumption that machines can only perform ‘routine’ tasks?

#### *‘Routine’ vs. ‘Non-Routine’ Tasks*

In Autor et al. (2003), the distinction between ‘routine’ and ‘non-routine’ tasks is based on the work of Michael Polanyi, a philosopher – in particular Polanyi (1966). The authors argue that ‘non-routine’ tasks are “tasks fitting Polanyi’s description”, where Polanyi is quoted as follows:

“*We can know more than we can tell* [p. 4] . . . The skill of a driver cannot be replaced by a thorough schooling in the theory of the motorcar; the knowledge I have of my own body differs altogether from the knowledge of its physiology; and the rules of rhyming and prosody do not tell me what a poem told me,

*without any knowledge of its rules* [p. 20].” [Emphasis added]  
(pg. 1283)

Autor et al. (2003) therefore implies that ‘non-routine’ tasks are those that are performed by human beings using knowledge about which “we can know more than we can tell”, that human beings can perform without “any knowledge” of the rules they follow.<sup>9</sup> Put simply, ‘non-routine’ tasks are those that rely on what Polanyi called ‘tacit’ knowledge – knowledge that people struggle to articulate when called upon to do so. For example, a world-class sportsman may struggle to articulate the heuristics that allow him to throw a ball unrivalled distances. But this is still ‘knowledge’ – it is just not ‘explicit’ knowledge that can readily be articulated, but is ‘tacit’. Autor (2014), a more recent re-statement of the task-based approach, again defines ‘non-routine’ tasks with reference to Polanyi’s concept of ‘tacit’ knowledge, arguing that ‘non-routine’ tasks are those we only “tacitly understand how to perform”, that involve processes we “do not explicitly understand”.

This addresses the first question. Following Autor et al. (2003) and Autor (2014), ‘non-routine’ tasks are those that require tacit knowledge to perform. The second question follows – why does the fact that a task is ‘non-routine’, and so requires tacit knowledge, imply that a machine is poorly suited to perform it?

### *‘Computer Capital’ and Routine Tasks*

Autor et al. (2003) also introduces the argument that machines can only perform ‘routine’ tasks. This claim is grounded on a particular understanding of how machines must operate. In Autor et al. (2003) this is that a machine can only perform a task if it can “follow explicit programmed rules”. This implies that a machine requires a task to be “exhaustively specified with programmed instructions” and without these “explicit programmed instructions” a machine cannot perform the task. The same point is made in Autor (2014) but with an additional claim:

---

<sup>9</sup>In Autor, Levy, and Murnane (2002) there is an informal discussion of the ideas that would become the ALM hypothesis. The argument in this earlier paper is that “what computers actually do” is “the execution of procedural or “rules-based” logic”. This is similar to the ALM hypothesis.

“For a computer to accomplish a task, a programmer must first fully understand the sequence of steps required to perform that task, and then must write a program that, in effect, causes the machine to precisely simulate these steps.” (pg. 6)

In Autor (2015) this is described as the “traditional” approach to programming. But the claim in Autor (2014; 2015) is not simply the original one made in Autor et al. (2003), that a machine must be set an explicit set of programmed rules; it is stronger – that those explicit rules must originate with, and precisely reflect, the thinking process of a human being. It is from this stronger claim that the assumption that computer capital can only perform ‘routine’ tasks follows. Autor (2015) puts this formally:

“But the scope for this kind of substitution is bounded because there are many tasks that people understand tacitly and accomplish effortlessly but for which neither computer programmers nor anyone else can enunciate the explicit “rules” or procedures.” (pg. 11)

And Autor (2014) more prosaically:

“At a practical level, Polanyi’s paradox means that many familiar tasks, ranging from the quotidian to the sublime, cannot currently be computerized because we don’t know the rules.” (pg. 8)

This provides an answer to the second question, about the basis for the claim that machines can only perform ‘routine’ tasks. If human beings cannot articulate their thought process in performing a task, if they do not ‘understand’ how they perform the task, then they cannot articulate a set of rules for a machine to follow, and the machine will not be able to perform the task. Those tasks for which human beings cannot articulate their thought processes are the ‘non-routine’ ones – and so, accordingly to the ALM hypothesis, ‘non-routine’ tasks cannot be performed by machines. Conversely, those tasks that human beings can articulate their thought processes for are ‘routine’ tasks, and so those are the tasks that machines can perform.

## 2.3 Machines Can Perform ‘Non-Routine’ Tasks

The ALM hypothesis has a clear implication – machines cannot perform ‘non-routine’ tasks. Yet in practice this no longer holds. Increasingly, machines perform a wide range of ‘non-routine’ tasks. This is an important challenge to the ALM hypothesis. Driverless cars are a good case. Polanyi (1966) considered driving a car as a definitive case of a task that requires tacit knowledge:

“The skill of a driver cannot be replaced by a thorough schooling in the theory of the motorcar.” (pg. 20)

Given the ALM hypothesis draws on Polanyi (1966), it is understandable that Autor et al. (2003) also use the task of driving a car as a definitive case of a ‘non-routine’ task:

“Navigating a car through city traffic or deciphering the scrawled handwriting on a personal check – minor undertakings for most adults – are not routine tasks by our definition. The reason is that these tasks require visual and motor processing capabilities that cannot at present be described in terms of a set of programmable rules.” (pg. 1283)

In the task-based literature that followed Autor et al. (2003), the task of driving a car is also used as a definitive case of a ‘non-routine’ task that cannot be automated. Consider Levy and Murnane (2004):

“The bakery truck driver is processing a constant stream of information from his environment .... To program this behaviour we could begin with a video camera and other sensors to capture the sensory input. But executing a left turn across oncoming traffic involves so many factors that it is hard to imagine discovering the set of rules that can replicate the driver’s behaviour.” (pg. 20)

In 2004, this position – that the task of driving a car necessarily required a human being – was consistent with the available evidence. Yet this no longer holds. Almost all major car companies are now developing driverless cars.<sup>10</sup>

---

<sup>10</sup>See Gibbs (2016).

But driverless cars are only one example of the widespread automation of ‘non-routine’ tasks. Consider, for instance, a second task that Autor et al. (2003) suggests is ‘non-routine’ and so cannot readily be automated, that of “deciphering the scrawled handwriting on a personal check”. This task too can now be performed by machines.<sup>11</sup> Indeed, it can be performed in a variety of ways, and is a small part of a much larger field in AI research on ‘pattern recognition’.<sup>12</sup> So too for a third task noted by Autor et al. (2003), “medical diagnosis”.<sup>13</sup> In turn, in Autor and Dorn (2013) it is claimed that a further set of ‘non-routine’ tasks cannot readily be automated. Consider three of these – personal care, order-taking, and table-waiting. Traditionally, each of these may have required a human being. But whether this will continue is unclear. In Japan, for example, where the elderly population is the highest in the world and expected to rise further (26.7 percent 65 or older, Yoshida 2016), the burden of care is increasingly being shared between human beings and machines – commercial robots like ‘Paro’, a therapeutic toy seal, and ‘Pepper’, a personal robot, are used, and the market for nursing care robotics is expected to grow twenty-five-fold from 2015 to 2035 (Neumann 2016).<sup>14</sup> As for order-taking and table-waiting, there are now many restaurants and cafes where these processes are performed by machines.<sup>15</sup>

A further, important, recent development is that it now appears, in certain cases, that the opposite of the ALM hypothesis holds in practice – not only can certain ‘non-routine’ tasks be automated, but the ‘non-routine’ character of these tasks makes them more susceptible to automation. Consider Ng (2016), a leading AI researcher:

“If a typical person can do a mental task with less than one second of thought, we can probably automate it using AI either

---

<sup>11</sup>These systems are now ubiquitous. For example, for economists there is a smartphone app called ‘Mathpix’ which scans handwritten equations and solves them. This is available at <https://itunes.apple.com/gb/app/mathpix-solve-graph-math-using/id1075870730?mt=8>.

<sup>12</sup>Consider *Pattern Recognition*, published by Elsevier, a journal dedicated to this field.

<sup>13</sup>This is now a thriving field. See, for example, Moorefields (2016), describing the new partnership between Google DeepMind Health and Moorfields Eye Hospital in London.

<sup>14</sup><http://www.parorobots.com/> and <https://www.ald.softbankrobotics.com/en/cool-robots/pepper>.

<sup>15</sup>See Pudzer (2016).

now or in the future.”

In practice, it appears from Ng (2016), that tasks which require ‘tacit’ knowledge are at greater risk of automation. This is the opposite of what the ALM hypothesis supposes.

## 2.4 The First Explanation

The recent literature has noted that many ‘non-routine’ tasks can now be automated. The models in Acemoglu and Autor (2011) and Acemoglu and Restrepo (2017; 2018), for instance, allows for machines to displace labour in a wider range of tasks. In turn, Autor (2015) raises direct questions about the ALM hypothesis:

“Polanyi’s paradox – “we know more than we can tell” – presents a challenge for computerisation because, if people understand how to perform a task only tacitly and cannot “tell” a computer how to perform the task, then seemingly programmers cannot automate the task – *or so the thinking has gone.*” (pg. 24)  
[Emphasis added]

However, the explanation offered for *why* ‘non-routine’ tasks can now be automated still relies on the ALM hypothesis. The argument is that tasks which we thought were ‘non-routine’ are, in fact, more ‘routine’ than we realised. New technologies – advances in processing power, data retrieval and storage capabilities, and algorithm design – have made it possible to make ‘explicit’ more of the ‘tacit’ knowledge that human beings draw upon. That, the explanation goes, is why ‘non-routine’ tasks can now be automated. This argument is made, for instance, in Autor (2014):

“Contemporary computer science seeks to overcome Polanyi’s paradox [the inability to articulate rules for tasks that require ‘tacit’ knowledge] by building machines that learn from human examples, *thus inferring the rules that we tacitly apply but do not explicitly understand.*” (pg. 2) [Emphasis added]

in Autor (2015):

“[R]ather than teach machines rules that we do not understand, engineers develop *machines that attempt to infer tacit rules* from context, abundant data, and applied statistics.” (pg. 23)  
[Emphasis added]

and in Remus and Levy (2016) with respect to predicting the outcome of a legal dispute:

“Because the mental protocol [of a judge issuing a decision] is tacit – and not easily articulated – the judge may not experience his decisions as routine, but the *machine learning model makes the tacit protocol explicit* as a mathematical combination of characteristics taken from the case, which can then be used to predict future judicial decisions.” (pg. 13) [Emphasis added]

## 2.5 The Second Explanation

There is, however, a second explanation for why ‘non-routine’ tasks can be automated. The first explanation argues that new technologies allow us to uncover more of the tacit rules that human beings follow in performing ‘non-routine’ tasks; the second explanation argues that new technologies allow us to perform tasks with systems and machines that follow rules which do not need to reflect the rules that human beings follow at all, tacit or otherwise. This is a new and important conceptual point.

This second explanation is rooted in the observation that many of these new systems and machines are performing tasks in very different ways to human beings. The traditional literature recognises that this is possible. For instance, Autor et al. (2003) notes in a footnote that it is “a fallacy to assume that a computer must reproduce all of the functions of a human being to perform a task traditionally done by workers.” In the recent literature, Autor (2014) reflects this in a discussion of how driverless cars operate:

“It is sometimes said by computer scientists that the Google car does not drive on roads but rather on maps ... the Google car navigates through the road network primarily by comparing its real-time audio-visual sensor data (collected using LIDAR) against painstakingly hand-curated maps that specify the exact locations of all roads, signals, signage, obstacles, etc. ... while the Google car appears outwardly to be as adaptive and flexible as a human driver, it is in reality more akin to a train running on invisible tracks.” (pp. 33-4)

This system performs the task of driving a car in a different way to a human being. While both a human being and a machine rely upon sensor data

to drive the car – though their respective sensors vary – only the machine requires a “pre-specified” map.

The chess-playing system Deep Blue, developed by IBM, is the canonical case of a machine that performs a complex task in a different way to a human being. In 1997, Deep Blue defeated the then chess world-champion, Garry Kasparov, in a six game set. Kasparov (2010) noted afterwards that:

“Instead of a computer that thought and played chess like a human, with human creativity and intuition, they got one that played like a machine, systematically evaluating 200 million possible moves on the chess board per second and winning with brute number-crunching force.”

As a final example of a machine performing a task differently to a human being, consider the task of performing a medical diagnosis. Esteva et al. (2017) describes a system that can predict as accurately as dermatologists whether a picture of a skin discolouration is cancerous. To do so, it runs a pattern-recognition algorithm through a database of 129,450 past cases, searching for similarities between those cases and the particular discolouration in question. In a few seconds, it searches more possible cases than a human doctor could hope to review in their lifetime.

The fact that machines can increasingly perform tasks in different ways to human beings has very significant implications for the nature of the rules that these machines follow. This is not sufficiently recognised in the existing literature. In particular, when machines perform tasks in different ways to human beings *it is no longer necessary for the rules that these machines follow to resemble the rules that human beings follow*. In the cases of driving a car, playing chess, or making a medical diagnosis, there is no longer a need to understand, articulate, and replicate the thinking process of a human driver, chess-player, or doctor. It is wrong to imagine that the only rules that a machine could follow are those that human beings actually do follow. Put another way, the set of all possible rules a machine could follow is larger than the particular rules that human beings actually do follow.

The limitation with the ALM hypothesis follows. It assumes that the only way to automate a task is to understand, articulate, and replicate the thinking process of a human being when performing that task. The

only purpose of the ‘routine’ vs. ‘non-routine’ distinction is to distinguish between those types of tasks for which human beings can articulate their thought processes and those they cannot. And the basis for the assumption that machines can only perform ‘routine’ tasks is that those are the only tasks for which we can write a specific procedure, since those are the only tasks we can articulate how human beings perform. But when machines follow different rules to human beings, it is clear that ALM hypothesis no longer holds for those tasks. The inability of human beings to articulate their thinking processes is no longer a constraint.

The possibility that these systems and machines follow different rules suggests that first explanation for why certain ‘non-routine’ tasks can now be automated is incomplete. It may be true that it is now possible to automate certain ‘non-routine’ tasks because the ‘tacit’ rules followed by human beings can be made ‘explicit’. But if these machines follow rules that do not resemble those that human beings follow the first explanation is insufficient. The second explanation is necessary. Both arguments could explain why it is now possible to automate many more ‘non-routine’ tasks. But as machines become more powerful, there is less need to replicate the thinking-processes of human beings. This second explanation is becoming increasingly important.

More fundamentally, the ALM hypothesis draws on traditional reasoning about ‘artificial intelligence’ (AI). It relies on what is known as the first generation of AI research, conducted from the 1950s to 1980s.<sup>16</sup> Called ‘expert systems’ or ‘knowledge-based systems’, these required a human ‘domain specialist’ who was capable of articulating how they performed a particular task and then representing this in a set of rules for a machine to follow. Winston (1977), Hayes-Roth et al. (1983) and Russell et al. (1995) provide an overview of these approaches.<sup>17</sup> But the second generation of AI research, relying upon brute-force processing power, massive data storage capabilities, and advances in algorithm design, no longer require an articulation of those human rules. This does not imply at all that Polanyi’s distinction between ‘tacit’ and ‘explicit’ knowledge is wrong. But

---

<sup>16</sup>See Susskind and Susskind (2015).

<sup>17</sup>See also Weizenbaum (1976) and Dreyfus (1976; 1992). These first generation AI researchers also thought a Polanyi-type constraint, based on the inability to express tacit knowledge, was important.

it does imply that Polanyi's distinction is the wrong constraint when thinking about the capabilities of machines that operate in very different ways to human beings. Only when the way in which a machine performs a task must exactly reflect the way in which a human being performs that same task does a Polanyi-type constraint bind.

## 2.6 A New Hypothesis

These two explanations have different implications for the rules that machines follow in performing 'non-routine' tasks. The first explanation suggests that these machines are following the 'tacit' rules of human beings, which have been made 'explicit' – these 'non-routine' tasks were 'routine' after all. The second suggests that these machines are following rules that do not reflect the rules that human beings follow – the 'non-routine' and 'routine' distinction no longer matters.<sup>18</sup>

It follows that in thinking about automation, the appropriate criteria is not whether a task is 'routine' or 'non-routine' from the standpoint of a human being, as under the ALM hypothesis, but whether it has features that make it more or less *routinisable* from the standpoint of a machine. If a task is routinisable, a routine can be composed that allows a machine to perform it – but that routine may not necessarily reflect the way in which a human being performs the task. The concept of 'routinisability' nests the concept of 'routineness'. The latter is focused on whether a task has features that make it more or less feasible to articulate how a human being performs it in a set of rules for a machine to follow. The concept of 'routinisability' instead asks whether a task has features that make it more or less feasible to articulate a set of rules for a machine to follow. A new general hypothesis follows.

**HYPOTHESIS:** (1) *that machines substitute for workers in carrying out a set of tasks and activities that are 'routinisable'; and (2) that machines*

---

<sup>18</sup>One of features of the latest machines is that they are opaque – it is difficult to articulate the particular rules that they themselves follow. I thank David Autor for this point. But this makes the second explanation even more important – if it is not possible to verify that a particular set of rules resembles a human being's, then a broader explanation is needed to allow for the possibility that these rules may be different.

*complement workers in carrying out a set of tasks that are ‘unroutinisable’.*  
*The set of tasks and activities that are routinisable changes over time.*

This general hypothesis also nests the ALM hypothesis – in the special case where the rules that a machine follows must resemble those that human beings follow, the ‘routineness’ and the ‘routinisability’ of a task coincide. In that special case, the general hypothesis collapses to the ALM hypothesis. But in general, the routineness of a task is only one signal of a task’s routinisability. This is what recent literature exploring the new limits on automation have found (see, for instance, Brynjolfsson and Mitchell 2017). Most importantly, the set of tasks that were routinisable in the past is likely to be far smaller than the set of tasks that are routinisable in the future.

### 3 Conclusion

The traditional literature, based on the ALM hypothesis, supports an optimistic view about the threat of automation. This optimism relies on the claim that there exists a large set of types of tasks that cannot be automated and, in turn, that those “tasks that cannot be substituted by automation are generally complemented by it” (Autor 2015). But if, as argued in this paper, many ‘non-routine’ tasks can also be automated, then the set of types of tasks that offer a ‘refuge’ for labour will be far smaller than the traditional literature assumes. Intuitively this suggests that the ALM hypothesis, and the boundary it has imposed between what machines can and cannot do, may have created a misleading sense of optimism about the prospects for labour. This is an important argument for pessimism that requires further theoretical and empirical research.

### 4 Bibliography

- Acemoglu, Daron.** 2002. “Technical Change, Inequality, and the Labor Market”. *Journal of Economic Literature* 40 (1): 7-72.
- Acemoglu, Daron and David Autor.** 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings”. In: *Handbook of Labour Economics, Volume 4B* eds. David Card and Orley Ashenfelter, 1043-1171. Amsterdam: North-Holland.

- Acemoglu, Daron and Pascual Restrepo.** 2017. “The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares, and Employment.” NBER Working Paper No. 22252.
- Acemoglu, Daron and Pascual Restrepo.** 2018. “Artificial Intelligence, Automation, and Work.” NBER Working Paper No. 24196.
- Autor, David, Frank Levy and Richard Murnane.** 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration”. *The Quarterly Journal of Economics* 118 (4) 1279-1333.
- Autor, David.** 2013. “The “Task Approach” to Labor Markets: An Overview”, *Journal for Labour Market Research* 46 (3) 185 - 199.
- Autor, David and David Dorn.** 2013. “The Growth of Low-Skill Service Jobs and the Polarisation of the US Labor Market”. *The American Economic Review* 103 (5) 1553-1597.
- Autor, David.** 2014. “Polanyi’s Paradox and the Shape of Employment Growth”. NBER Working Paper No. 20485.
- Autor, David.** 2015. “Why Are There Still So Many Jobs? The History and Future of Workplace Automation”. *Journal of Economic Perspectives* 29 (3) 3-30.
- Erik Brynjolfsson and Tom Mitchell** 2017. “What can machine learning do? Workforce implications”. *Science* 358 (6730) 1530-1534.
- Dreyfus, Hubert.** 1976. *What Computers Can’t Do: The Limits of Artificial Intelligence* New York: Harper & Row.
- Dreyfus, Hubert.** 1992. *What Computers Still Can’t Do: A Critique of Artificial Reason* MA: MIT Press.
- Esteva, Andre, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan Swetter, Helen Blau and Sebastian Thrun.** 2017. “Dermatologist-level classification of skin cancer with deep neural networks.” *Nature* 542 155-188.
- Gibbs, Samuel.** 2016. “Self-driving cars: who’s building them and how do they work?” *The Guardian*, May 26.
- Goos, Maarten, Alan Manning and Anna Salomons.** 2014. “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring”. *The American Economic Review* 104 (8) 2509-2526.
- Hayes-Roth, Frederick, Donald Waterman and Douglas Lenat.** 1983. *Building Expert Systems*. MA: Addison-Wesley.
- Kasparov, Garry.** 2010. “The Chess Master and the Computer” *The New York Review of Books*, Feb 11.
- Levy, Frank and Richard Murnane** 2004. *The New Division of Labor*. India: New Age International Ltd.
- Mnih, Volodymyr, Koray Kavukcuoglu, David Silver, Andrei Rusu, Joel Veness, Marc Bellemare, Alex Graves, Martin Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis.** 2015. “Human-level control through deep reinforcement learning.” *Nature* 518 529-533.
- Neumann, Dana.** 2016. “Human Assistant Robotics in Japan” *EU Japan Centre for Industrial Cooperation*, March.
- Ng, Andrew.** 2016. “What AI Can and Can’t Do Right Now” *The Harvard Business Review*, November 9.
- Polanyi, Michael** 1966. *The Tacit Dimension*. London: University of Chicago Press.
- Puzder, Andy.** 2016. “Why Restaurant Automation Is on the Menu” *The Wall Street Journal*, March 24.
- Remus, Dana and Frank Levy.** 2016. “Can Robots be Lawyers? Computers, Lawyers, and the Practice of Law.” Available at SSRN: <http://ssrn.com/abstract=2701092>.

**Russell, Stuart and Peter Norvig** 1995. *Artificial Intelligence: A Modern Approach*. New Jersey: Pearson Education.

**Silver, David, Aja Huang, Chris Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel and Demis Hassabis.** 2016. "Mastering the game of Go with deep neural networks and tree search." *Nature* 529 484-489.

**Weizenbaum, Joseph** 1976. *Computer Power and Human Reason*. US: W.H. Freeman and Company.

**Winston, Patrick.** 1977. *Artificial Intelligence*. Reading: Addison Wesley.

**Yoshida, Reiji.** 2016. "Japan census report shows surge in elderly population, many living alone" *The Japan Times*, June 29.