iS
CHANNEL

Towards Interaction Machines

Antti Lyyra

PhD Candidate in Information Systems
Department of Management
London School of Economics and Political Science

KEYWORDS

Automation
Autonomy
Autonomous Artefacts
Robots and Autonomous Systems
Interactive Computation

ABSTRACT

Machines are an integral part of modern societies; they extend and expand the ability of the humankind to manipulate their environment by transcending the natural limitations intrinsic to humans. Advances in computationally enabled sensing, learning, action and control mechanisms and related techniques allow a wider variety of tasks and activities to be automated and passed from humans to machines. This paper aims to outline characteristics of this phenomenon by examining the foundations of machinery, automation and computation and consequently comparing the characteristics of automated and autonomous artefacts. As a result, the difference between the concepts is brought forward and implications to the design of autonomous machines and artefacts are discussed.

Introduction

Machines are ubiquitous in modern societies; individuals and organisations alike rely heavily on them in their daily routines. Given the wide spectrum of tasks, machines performing them also appear in various forms and operate on different functional foundations and principles.

Notwithstanding the differing appearances, they have some features that are common. To a certain extent most of them could be considered as extensions of human capabilities as well as embodiments of human knowledge that is built into machines to fulfil human needs and purposes. They are also artificial constructs that do not exist in the world without human involvement. In the context of automation, this involvement can be considered as a process of transforming life processes to mechanised operations that are automatically operated. These phenomena of mechanisation and automation have brought humankind to the machine age where a great deal of both material and immaterial outputs are produced by machines.

Modern machines, despite the great level of automation, require people to supervise and operate them because they are not very capable of adapting to the changes emerging from their environment. In order to loosen the coupling between machines and their operators, there are demands to equip machines and computer alike with capabilities to operate autonomously. The sustained efforts to build such artefacts have proved this to be challenging, although at the same time somewhat rewarding. This paper is set to compare and contrast the paradigms related to automation and autonomy with an aim to provide clarity on some foundational differences.

To begin the exploration, the second chapter outlines a brief history of machinery from early tools to the machine age along with some limitations of that machinery. After that, the third chapter discusses the concepts of agency, automation and autonomy. The fourth chapter reflects automation and autonomy against the backdrop of closed and open systems, after which computing techniques that enable autonomous behaviour in open systems are presented. The underlying characteristics of autonomous techniques are contrasted to that of Turing Model. Finally, the results are discussed in chapter five before concluding remarks.

Modern Machines

According to archaeologists, our ancestors started using simple stone tools in the Stone Age around

Corresponding Author
Email Address: A.K.Lyyra@lse.ac.uk (A. Lyyra)

3.5 million years ago (McPherron et al. 2010). Later, although the exact timing is not known, early primitive machines were contrived, a prime example being the one for making fire that consists of a fast spinning stick driven by bowstring (Paz et al. 2010). The foundations of modern machines were set around the 3rd century BC in Greece where Archimedes discovered the principle of mechanical advantage in the lever while studying levers, pulleys and screws (Wikipedia). Various machines and mechanical principles have become widely applied since their invention (see Nof, 2009).

The utilisation of machinery started at greater scale during the Industrial Revolution in the mid 18th century. At first, manufacturing facilities housed mechanically controlled machinery and production lines that were powered by steam engines (Paz et al. 2010). At the beginning of 20th century, the electrification of factories decoupled machines from the engines that powered them. With electricity also emerged electromechanical devices that enabled more sophisticated automation by providing means to operate and control machines and production lines automatically (Nof 2009).

In the first half of the 19th century, the first versions of mechanical machines capable of performing numerical calculations were introduced (see Grier 2005). Later, roughly a hundred years later, the first versions of digital reprogrammable computers started to emerge (Bissell 2009). Unlike their mechanical predecessors, these computers were able to process various types of information as long as it was presented in a correct format and reprogrammability made them pliable to various tasks*.

Computers were superior in performing calculations and processed information significantly faster than their human counterparts and quickly started forming structures that could be called information systems. Those systems record, manipulate and display information and transfer it over distances (Kallinikos 2001), making various types of data and information widely available and accessible for people or other systems. These systems could be considered as neural networks of modern societies (Arthur 2011).

Aforementioned technological inventions have helped transcending some physical and mental limitations intrinsic to humans. If tools used in the Stone Age served as a medium to extend the reach of human intelligence beyond the physical limits (Lovejoy 1981), the industrial revolution and its aftermath compares with growing the muscle of humankind (Weizenbaum 1984). In similar fashion, the modern computer technology could be reflected as an extension of cognitive and communicative capabilities.

According to W. Brian Arthur (Arthur 2007) technologies can be defined as a means to fulfil human purpose regardless of what the purpose is or how clearly it is defined:

> *As a means to fulfil a purpose, a technology may be a method or process or device: a particular speech recognition algorithm, say, or a filtration process in chemical engineering, or a type of diesel engine.*

The purposes that need fulfilling are human constructs, and while the human needs may be abundant, the technologies and techniques available to satisfy them typically are not.

This paper concentrates on the gap between the needs and enabling technologies in the field of automation. The great majority of aforementioned machines and computers introduced trail predefined procedures; should something unexpected happen, they quickly render themselves unable to operate. Therefore, while recognising that some of them may run reasonable long on their own, they ultimately need to be subjected to human supervision or be operated by humans.

By loosening the coupling between the machines and their operators, running a machine would not be contingent upon the availability of human operators or supervisors to the extent as it is now. Thereby, in order to reduce human involvement, techniques that would allow machines with a greater degree of autonomy are under development. This artificial simulation of human cognitive abilities, if embodied in machines, would extend the reach of human intelligence

A great majority of research on automation and autonomy have been carried out under the labels of electrical engineering, computer science, robotics and artificial intelligence (Siciliano & Khatib 2008; Brooks 1986; Winfield 2012). While approaches and problem areas vary, they rely on digitally enabled interaction in order to provide as a means to machines and computers with autonomous capabilities.

To conclude, humans have a long history of pushing boundaries; they build various types of

* Given the reprogrammable nature of the digital computer, it is capable of performing various information processing tasks, as long as the one requiring attention can be formalised (programmed) by specifying inputs (data) and desired outputs through a suitable sequence of instructions for data manipulation (algorithm) using the digital binary (0/1) numbers understandable by a computer (bits). Digitisation results as loose coupling between the type of information and its processor upon the assumption that all digitised information adopts the same form (bits). While digitised data format is flexible, the Turing model as computing paradigm is less so as it assumes the computer as a single processor that takes a input and performs calculations defined in a given algorithm until all steps have been completed. (Tilson et al. 2010; Kallinikos et al. 2013; Yoo et al. 2010; Weizenbaum 1984)

artificial machines as well as manipulate the natural environment they occupy. While great benefits have been received through mechanisation and automation, there is a will to provide machines with autonomy to relax the coupling between machines and their human operators. To further discuss this theme, in the following chapter, we examine the characteristics of automation and autonomy.

Automation and autonomy

Given that the creation of autonomous artefacts is grounded in robotics and artificial intelligence, we start this chapter by looking into the definition of a robot and the nature of agency. After that the origins and meaning of the terms autonomy and automation are discussed.

Autonomous artefacts

There are many ways to define and classify robots depending on their structure (see Siciliano & Khatib 2008) or areas of application (Haidegger et al. 2013). However, the definition provided by the Institute of Electrical and Electronic Engineer's (IEEE) Ontologies for Robotics and Automation (ORA) working group is presented as it emphasises the core aspects of robotics at a more abstract level (Prestes et al. 2013):

> *For our general purposes, robots are agentive devices in a broad sense, purposed to act in the physical world in order to accomplish one or more tasks. In some cases, the actions of a robot might be subordinated to actions of other agents, such as software agents (bots) or humans. A robot is composed of suitable mechanical and electronic parts. Robots might form social groups, where they interact to achieve a common goal. A robot (or a group of robots) can form robotic systems together with special environments geared to facilitate their work.*

While the definition provided is holistic and well rounded, it is worth noting that Prestes et al. (2013) consider robots as agentive devices with varying degree of autonomy that act in the physical world leaving out of the definition immaterial artefacts that operate only in the cyberspace. Although autonomous immaterial artefacts do not act in the physical world as such, they may still have concrete effects on the real world. To provide an example, trading robots may be used to monitor selected stocks and commodities in exchanges and to place sell and buy orders in the hope of gaining profits. Although trading robots, in other words, software programs running on digital computers connected to electronic market places (Lockwood et al. 2012), operate only in the cyberspace, the acts they perform have very concrete monetary consequences in the real life when the trades are cleared.

Thereby, in this paper, no borderline is drawn between the artefacts based on their manifestations or how they interact with the real world. Instead, the aim is to keep focus on the agentive and autonomous nature when explaining principles related to the phenomenon, hence we refer to them interchangeably as autonomous artefacts or autonomous machines.

Agency

The concept of agency is often present when robots and other autonomous artefacts are discussed. To avoid confusion, it is important to note that the term agent has different connotations depending on the context. In films an agent can be a spy working for a state, in business it can be a salesman contracted to act on the behalf of a company, and in philosophical discussion agent can be considered as a conscious, reflexive, intentional and rational agent symbolising an independent human soul (Rammert 2008).

However, when speaking of machines, agentive behaviour should be in principle considered in the light of representing someone else because machines as such do not have any inherent reflexivity or intentionality in them; should desires, beliefs and goals be embodied in machines, they would have been designed and implemented by humans, although some reservations must be left for the techniques and systems that are built on the idea of evolution, transformation, learning or emerging behaviour (Hayles 2005).

John Searle (1995), when presenting his argument regarding institutional and non-institutional agentive functions that are assigned to various artefacts, described the agentive function as follows:

> *…use to which we [conscious agents] intentionally put these objects. (p. 20)*

As an example Searle provided a stone that can be assigned with a function of paperweight. This definition reflects the human origins of the agency that is assigned to artefacts, and while it may leave the definition of the nature of agency debatable in terms of autonomy and how it may be perceived, it does provide the language and flexibility to discuss the evolving nature of agency in the context of post-modern machines. As an example, if a general-purpose machine has a capability and can be instructed to perform different tasks and therefore assigned to many different uses, the agentive function may change over time.

In this paper, the pragmatic and functional notion of agentive functions as presented by Searle is embraced and further exploration continue with the terms automation and autonomy.

Automation

The terms and automation and autonomy are quite often conflated because a great degree of automation may lead to an appearance of autonomy. However, this notion is somewhat misleading and therefore the difference between autonomy and automation is discussed here.

According to Richard D. Patton and Peter C. Patton automation is a combination of two words, automatic and operation (R. D. Patton & P. C. Patton 2009). The Oxford English Dictionary defines automatic as *"working by itself with little or no direct human control"*. The word has it roots in the Greek word *autómatos*, which means *"acting of itself"*, self-dictating and self-moving. The Springer Handbook of Automation automatic describes it as follows (Nof 2009):

> *A key mechanism of automatic control is feedback, which is the regulation of a process according to its own output, so that the output meets the conditions of a predetermined, set objective. (p. 23)*

From automatic we return back to combination of the words automatic and operation, to the definition of automation that Patton & Patton (2009) present in their chapter in the Handbook of Automation:

> *Automation is fundamentally about taking some processes that itself was created by a life process and making it more mechanical… it can be executed without any volitional or further expenditure of life process energy. (p. 305)*

By automation through mechanisation Patton & Patton mean the ability to perform linear and step-wise algorithmically defined processes with clear inputs and clear outputs.

This is similar to the definition, which Herbert A. Simon (1996) has provided on artificialness: Simon made a distinction between the natural sciences and the science of artificial, noting that natural science aims to find patterns hidden in apparent chaos whereas engineers and other designers of artificial systems aim to create new functional compositions out of patterns discovered in nature (pp. 1-10). Therefore, even if mechanistic automation may rely on phenomena that can be found in nature, it is not natural; it is not inevitable in the world and would not exist without human involvement.

With the notion that automation refers to artificial automatic operations that are self-dictating mechanised forms of real life, we move on to autonomy.

Autonomy

To start, in the Oxford English Dictionary the term *autonomy* is defined as *"the right or condition of self-government"*, and it has its roots in the Greek word *autónomos* that has the meaning of *"having its own laws"*. According to Froese et al. (Froese et al. 2007; Froese & Ziemke 2009) autonomy could be further defined in terms of external behaviour and internal autonomy, something they refer to as behavioural and constitutive autonomies. The former is generally required for having stable and flexible communication with environment whereas the latter is related to self-production, mutation and evolution as well as natural intentionality.

If we examine some of the most complex machines that humans have built such as digital interaction systems that control self-driving cars (Thrun et al. 2006), deep down we find nothing but formalised linear step-wise instructions (algorithms) and data that are represented by strings of binary numbers in order for necessary calculations to be performed. These rules that govern the self-dictating are realised and inscribed in machines by their designers. Thereby, it appears that such machines are automatic instead of autonomous in the sense that they do not posses the right or condition of self-government in the sense as a free person does, namely constitutive autonomy.

Joseph Weizenbaum (1984) describes a distinction between automatic and autonomous behaviour as follows:

> *Most automatic machines have to be set to their task and subsequently steered and or regulated by sensors or human drivers. An autonomous machine is one that, once started, runs itself on the basis of internalized model of some aspects of the real world. (p. 24)*

With *automatic machines have to be set to their task* Wiezenbaum presumably means that their behaviour including inputs and outputs must be formalised - this is what Patton & Patton refer to as mechanisation of a life process. When referring to autonomous machines, Weizenbaum indicates that they run themselves based on the some aspects of the real world.

The early attempts to build artefacts with autonomous behaviour were founded on the idea of predefined aspects of the real world and mechanised, computational models of the decision-making, rationalistic reasoning and cognition. This approach, currently known as GOFAI, Good Old Fashioned Artificial Intelligence (Haugeland 1985), assumes that objects existing in systems are presented in the form of meaningful symbolic knowledge presentations and logical step-by-step deductions used in problem solving were grounded on these knowledge representations.

These early attempts were heavily criticised due to their narrow view to real-life problem solving, human behaviour and decision-making. These shortcomings were debated by prominent philosophers, computer scientists, roboticists and anthropologists (H. Dreyfus & S. E. Dreyfus 1986; Winograd 2006; Suchman 1987; Brooks 1986; Winograd & Flores 1986).

However, through failed attempts to create autonomous behaviour it started to become increasingly evident that there are a multitude of modalities that guide human behaviour in different situations. As an example, Lucy Suchman (1987) argues that actions are always situated:

> *…insofar as actions are always situated in particular social and physical circumstances, the situation is crucial to action's implementation. (p. 178)*

Moreover, psychologist Daniel Kahneman (2011) argues the human brain consists of two systems, fast (1) and slow (2): the system 1 being fast, automatic, reactive and subconscious whereas the system 2 is slow, logical, calculating and conscious. In addition to human behaviour, other natural forms of interaction and communication have provided inspiration for technologists creating autonomous artefacts.

Thereby, it became evident that the early attempts to build autonomous artefacts did not recognise the multitude of modalities inherent to human communication or the role of context in action. In this light it seems that Weizenbaum's notion on the some aspects of the real world do not only refer to the aspects that can be captured, formalised and embodied into a machine, but also to the aspects that cannot be subjected to such mechanisation.

Open systems and interaction machines

In order to reflect the influence of surrounding environment, we continue with Patton & Patton (2009). They write that mechanical means non-context sensitive and discreet and also highlight that machine theory is the opposite of general systems theory. By *general systems* they mean open systems or in other words, systems that can locally overcome entropy and are self-organizing. Moreover, open nonlinear context-sensitive systems are fundamentally different from the computational algorithms inscribed into machines in the sense that everything else in the systems affects the behaviour, not only the previous step in an algorithm (p. 306).

Because it is not possible to model open systems, an artefact, in order to function as a part of self-organising open system, should be capable of orienting itself in such a system; it should be granted with capabilities to negotiate with and adapt to the surroundings it is located in and is a part of. In other words, artefacts should be modelled as complex adaptive systems similarly as the environment they operate in, such as road system and traffic, may be. John H. Holland (1992) describes such systems as follows:

> *Complex adaptive systems are evolving structures; these systems change and reorganize their component parts to adapt themselves to the problems posed by their surroundings.*

Embracing "*the right or condition of self-government*" as the general definition of autonomy and reflecting the notion of behavioural autonomy provide by Froese et al. (2007), in the context of artificial autonomous artefacts, autonomy could be considered as a behavioural model of an autonomous artefact that provides it with a local and situated capacity to act in an open and dynamic environment when it is performing an agentive function.

Because objects in open systems are in constant communication with their environment, several techniques have developed to simulate natural phenomena that allow interaction with and within open systems: examples include computing techniques such as neural networks for machine learning and speech, image and text recognition (Haykin 1994), embodied sense-react heuristics for direct interaction based behaviour modelling (Brooks 1986) as well as technologies and techniques for sensing, localisation and mapping, planning and actuation (see Siciliano & Khatib 2008) and communication (Mezei et al. 2010; Arumugam et al. 2010).

Some advanced compositions of these techniques have been brought together in form of autonomous vehicles (Thrun et al. 2006) and bipedal robots (Bekey 2005). While their behavioural autonomy is limited, they have been able to operate in loosely constrained systems somewhat successfully. These systems are built on foundations that are radically different compared to the Turing hypothesis that serves as a foundational concept for computation and states that any process that can be naturally called an effective procedure [algorithm] is be realised by a Turing machine (Vitanyi 2012).

The thesis defines a closed system where inputs, processing logic and outputs are clearly defined in symbolic format while simultaneously preventing undesired external impacts from entering the system. An ordinary personal computer, in its basic form provides a good analogy: a user instructs a machine, using a mouse and keyboard and validates the outcomes that are displayed on the screen. To outline the closed and artificial nature of this approach, Gordana Dodig-Crnkovic (2011) states:

> *The Turing Machine essentially presupposes a human as a part of a system—the human is the one who poses the questions, provides material resources and interprets the answers.*

To further illustrate the shifting paradigm of computational processes, Dodig-Crnkovic (2011) describes them as outlined below:

> *Computational processes are nowadays distributed, reactive, agent-based and concurrent. The main criterion of success of the computation is not its termination, but its response to the outside world, its speed, generality and flexibility; adaptability, and tolerance to noise, error, faults, and damage.*

However, the techniques to interact with surroundings are fundamentally, at their lowest level, automatically operated mechanisms that utilise a variety of feedback loops for controlling and steering the processes of adapting to the environment while pursuing for goals. These techniques could be referred to as sensing, thinking, acting and reacting. While the individual atomistic features and mechanisms can be composed and modelled as algorithms, together they may form an interaction machine that operate in a non-algorithmic manner. In Peter Wagner's (1997) words:

> *Interactiveness provides a natural and precise definition of the notion of open and closed systems. Open systems can be modeled by interaction machines, while closed systems are modeled by algorithms.*

Also, here we must note that such interaction machines may resemble Russian dolls by their nature and be compositions of different artefacts with varying degree of autonomy consequently forming open artificial systems. This is what Prestes et al. (2013) refer to when they argue that a robot can be a composition of robotic devices, a robot group a composition of robots and consequently a robot system may consist of robot groups.

In this kind of open systems the overall functionality may emerge in a generative manner from the interaction of components - as Robin Milner (2006) puts it:

> *[I]n interactive systems everything can happen as soon as the interactions which trigger it have occurred.*

However, if interfaces between components of the system, are highly constrained, non-algorithmic systems may become algorithmic (Wegner 1997). The capabilities of components as well as their interconnections together define the capabilities and constraints of a given machine, in other words, its level of autonomy in a given context (Parasuraman et al. 2000). Reflecting the potential range of capabilities and assuming that only a subset of all interactional capabilities will be in use at any given one time, others remaining dormant, I refer to the whole set of possible interactions within a single artefact as *interactive affordances*, denoting the potential characteristics of *interactional performances*.

Although realisations of computing applications are evolving towards interaction machines, the Turing model is not perishing. Instead, it plays a central role as an atomic unit of interactive systems (Dodig-Crnkovic 2011), as it does in today's practical implementation of interactive systems such as self-driving cars or telecommunication networks. To better understand systems at the level of interactions, the focus of sense-making is expanding from single Turing machines and algorithms to computational processes, from computational prescriptions imposed on a computer to behavioural descriptions in terms of on-going interactions (see Goldin et al. 2006). According to Robin Milner (2006) computing has grown into *Informatics* that he calls as the science of interactive systems.

Discussion

In the previous chapters a brief history of machinery was described, before moving on to the emerging trend of granting machines with behavioural autonomy in open systems and discussing on how autonomy is different from automation along with the role of interactive computation. This chapter summarises the main ideas.

The examination of autonomous artefacts started from robotics. While Prestes et al. (2013) described robots as devices, it was suggested that immaterial autonomous artefacts operating in the cyberspace should also be taken into consideration when examining the emerging computing paradigm, given that the interactive nature of computation applies equally in both cyber and physical worlds. In robotics, physical features such as frames, sensors and actuators act as an interface to the real world (Dodig-Crnkovic 2011) whereas in the cyberspace that sensing and actuation are realised through electronic messaging interfaces. Therefore both material and immaterial realisations were referred interchangeably as *autonomous artefacts* or *autonomous machines*, indicating their human-made nature and autonomous behaviour.

After outlining briefly different meanings of agency, it was decided to follow John Searle's definition due to its pragmatic definition, suggesting that agency is manifested in the form of *agentive function* that is considered to be uses that we conscious agents intentionally put these objects. Here it is important to note that in the context of reprogrammable autonomous artefacts, the agentive function, the act, the goal to pursue may change over time.

In terms of automation, it was concluded that it refers to automatic operation and is a mechanised form of a life process, designed to work in closed systems without sensitivity to the context outside systems' boundaries. Furthermore, when autonomy was explored, it was considered that behaviourally autonomous machines, operating in open systems, are to proceed in a more non-deterministic manner by choosing an appropriate course of action from the spectrum of possible choices. In order for a machine to succeed in doing so, the machine should possess an ability to react to the changes emerging from the environment (open system) and to negotiate a solution that is in harmony with the local context and agentive function. These interactional capabilities could be referred as *interactional affordances* because these capabilities define the capacity to act, although may remain dormant as well. However, this autonomous behaviour is modelled through a series of interacting automated mechanisms that are designed and embodied into autonomous artefacts by the designers of artificial systems. Thereby, the point where autonomy ends and automation begins remains sometimes debatable.

Persistent efforts have been made to equip machines with autonomous capacities. The early attempts of creating artificial intelligence were not very successful given that system designers created knowledge representations and rational logics for solving selected problems problem. They were more suitable for cognitive reasoning within closed systems; in open systems problems they rendered themselves unintelligible because the knowledge embodied in them was not necessarily relevant *in situ* as its meaning and purpose was not grounded in the reality emerging from the environment.

After the failed attempts and realising that there is a multitude of modalities associated to human behavioural and interactional models in situated contexts, several computing technologies and techniques have been developed to provide a solution for particular communicational problem. As none of them is able to solve all communicational and interactional challenges at once, several technologies and techniques need to be combined in order to develop a desired behavioural model. This transforms the focus from a computable algorithm to behavioural models as compositions of interaction processes, which could be considered as a paradigm shift. Moreover, compositions of interacting artefacts can consequently form systems where overall behaviour may emerge from the interaction of artefacts. In these cases, the nature of perceived agency may vary depending on the point of view the observer has. A person building a machine may have a clear understanding of the inner workings and embodied logic, whereas, someone not familiar with the system could be tempted to speak *as if* the autonomous artefact had desires and beliefs because he or she is required to explicate the behaviour by interpreting the actions taken by an *interaction machine*.

This change in the quality of computational processes from transaction to interaction processing and behavioural modelling at a system level provides engineers and computer scientist alike with challenging problems to tackle. The challenges revolve around how to transform and model continuous, analogue and open-ended world into machines as on-going series of discrete interactive computations. This requires new approaches and techniques such as multi agent systems and agent based modelling together with a solid theoretical foundation comparable to the Turing model that has served as a theoretical cornerstone of algorithmic computation (Wegner 1998; Goldin et al. 2006; Dodig-Crnkovic 2011).

Concluding remarks

Recent advances in computational techniques and approaches have made it possible to build autonomous artefacts that are able to perform tasks and activities in open environments, denoting the shift from algorithmic computation to interactive computational processes. Should these techniques be adopted at the speeds encountered with personal computers, the Internet and smart phones, it is quite possible that autonomous artefacts in various configurations will constitute a significant part of our digital infrastructures in the near future. For this reason, it would be important to expand information system research towards the fields of interactive computation and autonomous artefacts and study various technological, organisational and sociological implications they may arrive in the wake of the interaction machines.

References

Arthur, W.B., 2011. The second economy. *McKinsey Quarterly*, pp.90–99.

Arthur, W.B., 2007. The structure of invention. *Research Policy*, 36(2), pp.274–287.

Arumugam, R. et al., 2010. DAvinCi: A cloud computing framework for service robots. In Robotics and Automation (ICRA), 2010 IEEE International Conference on. IEEE, pp. 3084–3089.

Bekey, G.A., 2005. *Autonomous Robots*, London: MIT Press.

Bissell, C., 2009. A History of Automatic Control. In S. Y. Nof, ed. *Springer Handbook of Automation*. Heidelberg: Springer, pp. 53–69.

Brooks, R.A., 1986. A robust layered control system for a mobile robot. *Robotics and Automation, IEEE Journal of*, 2(1), pp.14–23.

Dodig-Crnkovic, G., 2011. Significance of Models of Computation, from Turing Model to Natural Computation. *Minds and Machines*, 21(2), pp.301–322.

Dreyfus, H. & Dreyfus, S.E., 1986. *Mind Over Machine*, Oxford: Basil Blackwell.

Froese, T. & Ziemke, T., 2009. Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence*, 173(3-4), pp.466–500.

Froese, T., Virgo, N. & Izquierdo, E., 2007. Autonomy: A Review and a Reappraisal. *Advances in Artificial Life*, 4648(Chapter 46), pp.455–464.

Goldin, D., Smolka, S.A. & Wegner, P., 2006. *Interactive Computation*, Springer.

Grier, D.A., 2005. *When Computers Were Human*, Woodstock: Princeton University Press.

Haidegger, T. et al., 2013. Applied ontologies and standards for service robots. *Robotics and Autonomous Systems*, 61(11), pp.1215–1223.

Haugeland, J., 1985. *Artificial Intelligence*, Cambridge, MA: MIT Press.

Haykin, S.S., 1994. *Neural networks*, Toronto: Maxwell Macmillan Canada.

Hayles, N.K., 2005. Computing the Human. *Theory, Culture & Society*, 22(1), pp.131–151.

Holland, J.H., 1992. Complex adaptive systems. *Daedalus*, 121(1), pp.17–30.

Kahneman, D., 2011. *Thinking, Fast and Slow*, Penguin UK.

Kallinikos, J., 2001. *The Age of Flexibility*, Lund: Academia Adacta.

Kallinikos, J., Aaltonen, A. & Marton, A., 2013. The ambivalent ontology of digital artifacts. *MIS Quarterly*, 37(2), pp.357–370.

Lockwood, J.W. et al., 2012. A Low-Latency Library in FPGA Hardware for High-Frequency Trading (HFT). In 2012 IEEE 20th Annual Symposium on High-Performance Interconnects (HOTI). IEEE, pp. 9–16.

Lovejoy, C.O., 1981. The origin of man. *Science, Technology, & Human Values*, 211(4480), pp.341–350.

McPherron, S.P. et al., 2010. Evidence for stone-tool-assisted consumption of animal tissues before 3.39 million years ago at Dikika, Ethiopia. *Nature*, 466(7308), pp.857–860.

Mezei, I., Malbasa, V. & Stojmenovic, I., 2010. Robot to Robot. *IEEE Robotics & Automation Magazine*, 17(4), pp.63–69.

Milner, R., 2006. Turing, Computing and Communication. In D. Goldin, S. A. Smolka, & P. Wegner, eds. Heidelberg: Springer Berlin Heidelberg, pp. 1–8. Available at: http://dx.doi.org/10.1007/3-540-34874-3_1.

Nof, S.Y., 2009. Automation: What It Means to Us Around the World. In S. Y. Nof, ed. *Springer Handbook of Automation*. Heidelberg: Springer, pp. 13–52. Available at: http://dx.doi.org/10.1007/978-3-540-78831-7_3.

Parasuraman, R., Sheridan, T.B. & Wickens, C.D., 2000. A model for types and levels of human interaction with automation. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 30(3), pp.286–297.

Patton, R.D. & Patton, P.C., 2009. What Can Be Automated? What Cannot Be Automated? In S. Y. Nof, ed. *Springer Handbook of Automation*. Heidelberg: Springer, pp. 305–313. Available at: http://dx.doi.org/10.1007/978-3-540-78831-7_18.

Paz, E.B. et al., 2010. *A brief illustrated history of machines and mechanisms*, Dordrecht: Springer.

Prestes, E. et al., 2013. Towards a core ontology for robotics and automation. *Robotics and Autonomous Systems*, 61(11), pp.1193–1204.

Rammert, W., 2008. Where the action is: Distributed agency between humans, machines, and programs. *Paradoxes of Interactivity*, pp.1–18.

Searle, J.R., 1995. *The Construction of Social Reality*, Simon and Schuster.

Siciliano, B. & Khatib, O., 2008. *Springer Handbook of Robotics*, Stanford, CA: Springer.

Simon, H.A., 1996. *The Sciences of the Artificial*, MIT Press.

Suchman, L.A., 1987. *Plans and Situated Actions*, Cambridge: Cambridge University Press.

Thrun, S. et al., 2006. Stanley: The robot that won the DARPA Grand Challenge. *Journal of Field Robotics*, 23(9), pp.661–692.

Tilson, D., Lyytinen, K. & Sørensen, C., 2010. Research Commentary—Digital Infrastructures: The Missing IS Research Agenda. *Information Systems Research*, 21(4), pp.748–759.

Vitanyi, P.M.B., 2012. Turing Machines and Understanding Computational Complexity. *arXiv.org*, 1201.1223v1, pp.1–9.

Wegner, P., 1998. Interactive Foundations of Computing. *Theoretical Computer Science*, 192(2), pp.315–351.

Wegner, P., 1997. Why interaction is more powerful than algorithms. *Communications of the ACM*, 40(5), pp.80–91.

Weizenbaum, J., 1984. *Computer Power and Human Reason from Judgement to Calculation*, London: Pelican Books.

Winfield, A., 2012. *Robotics: A Very Short Introduction*, Oxford University Press.

Winograd, T., 2006. Shifting viewpoints: Artificial intelligence and human–computer interaction. *Artificial Intelligence*, 170(18), pp.1256–1258.

Winograd, T. & Flores, F., 1986. *Understanding Computers and Cognition*, New Jersey: Intellect Books.

Yoo, Y., Henfridsson, O. & Lyytinen, K., 2010. Research Commentary--The New Organizing Logic of Digital Innovation: An Agenda for Information Systems Research. *Information Systems Research*, 21(4), pp.724–735.