Foresight Brief

The *Foresight Brief* is a new publication from the European Trade Union Institute (ETUI) focused on strategic thinking about the future challenges for the world of work. It is produced by the ETUI's Foresight Unit, whose work concentrates on two priority areas: climate change adaptation and new technologies. The *Foresight Brief* is also available in French under the title *Notes de prospective*.

A law on robotics and artificial intelligence in the EU?

Aída Ponce Del Castillo
Senior researcher at the European Trade Union Institute

Automation, robots and artificial intelligence are already deeply embedded within our society. They will continue to have an increasing impact on the way we live and will soon become commonplace in the workplace, working alongside humans with increasing levels of autonomy and self-reliance. This raises many concerns: in particular, what happens if these advanced technologies go wrong? Could a robot be taken to court? What should be done if a mistake is made by an automated car, a police-robot or the care-robot that had seemed so failsafe when it was purchased by your local hospital to replace nurses? Some technologists claim that robots are becoming increasingly safe, but we all know that the unexpected can happen. Who is ultimately liable: the robot, its owner, the manufacturer? This *Foresight Brief* introduces the European Parliament Resolution on Civil Law Rules on Robotics. It examines the regulatory aspects of existing and future technologies, drawing attention to several key issues, such as the visibility, accountability and liability of all stakeholders. We hope that this brief will contribute to the discussions on our evolving interaction with robots, AI and technology, both now and in the future.

#02 – September 2017

Editors of the Foresight Brief series:
Christophe Degryse
Philippe Pochet
Aída Ponce Del Castillo

Editor of this issue:
Christophe Degryse,
cdegryse@etui.org

More information on
www.etui.org > Foresight Unit

ISSN 2507-1548

foresight unit

etui.

Introduction

Regulating new technologies is a challenging exercise, as it implies dealing with uncertainty and fast-moving, often invisible technologies (Graeme *et al.* 2012). Understanding where technologies are heading, what they can achieve and, more importantly, how they are designed and used is very difficult.

It is equally difficult to see where we are heading as human societies. Predicting how technologies and societies will interact and influence each other is, understandably, even more complicated.

In the case of technologies such as robotics and artificial intelligence, understanding how they are going to affect our societies involves discussing persistently latent societal and political issues with a view to determining precisely where these technological developments belong in our regulatory system.

The European Parliament Resolution on the regulation of robotics and artificial intelligence comes at an opportune time, as these technologies affect and shape the relationship that we have with our jobs and our lives, rendering them increasingly fast, automated and digitalised. Whether in manufacturing or services, the level of automation can range from executing simple tasks to making complex and speedy decisions. As such, it is a driver of social, employment and environmental change. However, as with any disruptive technology, its impacts on society are significant and include changes in work organisation, working conditions, and the quality and other aspects of employment such as the outsourcing of responsibility (Broersen 2014; Degryse 2017).

At present, no specific legal provision on robotics exists (with the possible exception of the Machinery Directive 2006/42/EC and the Product Liability Directive 85/374/EEC). A minimum set of rules is therefore necessary in order to establish a transparent and democratic level playing field, and some pieces of EU legislation would need to be reviewed accordingly (Kritikos 2016).

Annual supply of industrial robots 2014-2015 and forecast for 2016-2019

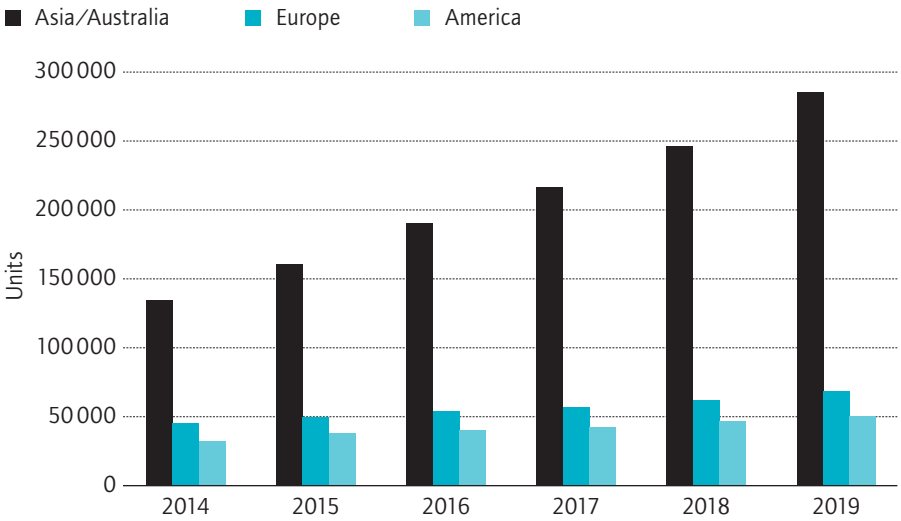

Source: IFR/World Robotics

This paper discusses the European Parliament's Resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (European Parliament 2017). It provides a brief summary of the content of the Resolution and looks at its basic principles and *raison d'être*. It also touches on the issue of defining robots and their liability. In so doing, it suggests a twofold shift in the rationale of Parliament's recommendations. Using a prospective approach and taking into consideration the views of scholars who are specialised in analysing robotics and artificial intelligence, this paper proposes that Parliament's recommendations could go further by addressing a much broader spectrum of artificial agents and artificial intelligence, instead of focusing on specific categories of robots.

It then looks at the responsibility, visibility and liability of those who have decision-making powers over the design, development and deployment of robots and artificial intelligence, including designers and developers.

The grounds for Parliament's initiative and robot politics

The European Parliament took the initiative to draft a report tabling a motion for a resolution of a legislative character and including recommendations related to Civil Law Rules on Robotics made by the Committee on Legal Affairs. The Resolution is not in itself a legislative initiative but rather a set of recommendations destined for the European Commission, asking the executive body to draw up a legal framework for the civilian use of robots. After an intense debate, Parliament voted in favour of the text on 16 February 2017. In its follow-up to the Resolution, the European Commission (EC/2015/2103 (INL)) expressed its intention to explore different solutions with a view to tackling liability and establishing a comprehensive registry, while speculating as to whether a definition was necessary for regulatory purposes. Even more interesting is that the Commission highlighted the importance of maintaining societal buy-in.

In its Resolution, the Parliament calls for legal solutions to civil liability caused by robots.

Although it does not state as such in the body of the document, it is clear that the Resolution is based on the results of the FP7 project RoboLaw – 'Regulating Emerging Robotic Technologies in Europe: Robotics facing Law and Ethics', funded by the European Commission and conducted between 2012 and 2014 (Palmerini *et al.* 2014); it also draws on the study 'European Civil Law Rules in Robotics' commissioned by the European Parliament's Legal Affairs Committee (Nevejans 2016).

The first draft of the Resolution followed the same lines as the project. The Resolution, adopted following the debate in the European Parliament, is set out in several sections on the development of robotics and artificial intelligence solely for civil use, namely: ethical principles such as transparency, and the well-established bioethical principles of beneficence, non-maleficence, autonomy and justice; the creation of a European agency to provide technical, ethical and regulatory expertise and guidelines on best practice; intellectual property rights and the flow of data; standardisation, safety and security; autonomous vehicles; care robots; medical robots; human repair

and enhancement; drones; education and employment forecast; liability and international aspects thereof; as well as a code for research ethics committees.

The main assumptions of the Resolution are the human right to privacy, respect for human frailty, maximal, reasonable transparency in the programming of robotic systems, and the need for predictability of robotic behaviour.

The Resolution provides an indication of what should be included in a definition of 'smart autonomous robots' and recommends establishing an ethical framework for the design, production and use of robots. The ultimate goal is to shape the technological revolution by incorporating the outlined principles into European legislation and codes of conduct.

It calls for legal solutions to civil liability for damage caused by robots, such as the establishment of a compensation fund, adoption of strict and proportional liability standards, establishment of an obligatory insurance scheme supplemented by a fund to ensure that reparation can be made for damage, allocation of registration numbers and creation of a specific legal status for robots in the long run.

In the field of employment, the Resolution calls on the Commission to develop digital abilities and monitor long-term job trends, and highlights the need for more flexible training, the importance of creative, social and digital skills, and robotics' potential for creating new risks.

The Resolution begins by referring to the oft-cited 'Three Laws of Robotics' devised by science-fiction author Isaac Asimov in *Runaround* (1942). The Three Laws, which establish that robots should serve their human master, provide as follows: a robot may not injure a human being or, through inaction, allow a human being to come to harm; a robot must obey orders given it by human beings except where such orders would conflict with the First Law; finally, a robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

The Three Laws formulated by Asimov have been used and applied for the development of computer software and other digital technologies (Feitelson 2007). Although not presented as core legal principles, they are an important part of the document's rationale.

One of the most interesting aspects of the Resolution is that this is not the first time that Parliament has taken an initiative to address the impact of technology on our society. An earlier example would be Parliament's high level of involvement in regulating nanotechnologies and nanomaterials (European Parliament 2009, 2014) which pushed the European Commission to consider adopting a regulatory framework.

Why the concept of 'artificial agents' makes more sense than 'robots'

We interact with robots and other AI systems constantly in our daily life, in all sorts of settings. Rapidly and constantly changing technologies become more and more intertwined with humans in their private and professional life. One obvious and simple example of this is how we use mobile devices (containing personal data) and how, for some individuals, being connected in an 'onlife' (Floridi 2014) world has become almost an extension of the self.

There are countless examples of how AI has become embedded in the workplace. In the financial sector, high-speed algorithms make decisions on buying and selling shares and have somehow replaced the need for human brokers. In journalism, algorithms can identify, calculate and select news items to be published by any given media. In the healthcare sector, robots can fetch objects, measure a patient's vital signs, guide elderly people in care homes and even demonstrate 'social behaviour'. In the security sector, robots merge with other technology systems such as the 'Internet of Things' (IoT), algorithms and platforms, with limited human interaction. In manufacturing, collaborative robots or worker-robot systems operate alongside humans, and their respective working spheres or envelopes[1] overlap.

The robots that deserve more attention are those with deep-learning capabilities.

The robots that deserve more attention are those with deep-learning capabilities: robots that can 'learn' from their environment, sense their surroundings, identify patterns, change the way a problem is originally framed, adjust their behaviour in response to their environment and function semi-autonomously (at this point in time, there are no fully functional long-term autonomous robots). Their interaction with humans can develop in many directions and across many layers as they play an active role in the decision-making process (robots/surgeons, algorithms/stock exchange, etc.).

Defining and categorising robots

The word 'robot' comes from the Czech word 'robota', meaning 'serf labour', 'drudgery' or 'hard work'; however, there is no consensus on the precise meaning of the term, and its numerous negative connotations make defining it very difficult. Even categorising robots can be a never-ending exercise.

Focusing on what roboticist Alan Winfield highlights as the three key defining qualities of robots can offer a solution. Winfield talks of 'an artificial device that can *sense* its environment and *purposefully act* on or in that environment; an *embodied* artificial intelligence; or a machine that can *autonomously* carry out useful work' (Winfield 2012).

It should be noted that, in its Resolution, the European Parliament does not propose a straightforward definition or classification of 'smart' autonomous robots. It does provide, in the Annex to the Resolution, certain criteria to be taken into consideration and calls on the Commission to propose a definition of cyber physical systems, autonomous systems, smart autonomous robots and their subcategories, which in itself can be an interminable process and raises a number of concerns.

Indeed, we can compare this process to the one used to define nanotechnologies and nanomaterials in 2010 (Nanowerk 2013). A divergence of scientific opinion led to the definitions becoming political compromises, and the negotiations on their revision are still ongoing. We can perhaps learn a useful lesson here: trying to achieve a strict definition and categorisation for regulatory purposes is not always the best option, and the European Parliament

1. *Envelope* refers to a robot's working space or range of movement.

should perhaps be careful not to become bogged down in discussions aimed at defining that which is scarcely definable.

Artificial agents

In our view, rather than trying to define robots, which can be divided into millions of sub-types depending on their functionalities, the Commission should adopt a wider approach, encompassing algorithms and AI, and take into consideration machines that have the capacity to learn, evolve and eventually become semi- or, maybe one day, fully autonomous.

We also need to recognise the importance of using the correct terminology in policy-making and in the drafting of legislation. The words that are used to describe these machines will have an impact on the kind of policy and legislation that will result. As these technologies are so closely integrated into all areas of human life, we have to proceed with caution. Should we refer to robots, artificial agents or autonomous artificial agents, and what are the likely consequences of that distinction?

The Commission should take into consideration machines that have the capacity to learn, evolve and eventually become semi- or fully autonomous.

The term 'artificial agent', understood as a spectrum concept, makes sense because it covers a wide diversity of 'soft' and 'hard' agents: decision-making algorithms, automated machines, digital agents, hybrid multi-agents, Internet bots, robots, nano-robots, drones, etc. These agents have the capacity to operate and learn through experience and interaction, without the direct intervention of humans or other agents. They are not (yet) fully autonomous; they are *artificial* because they are produced or constructed by humans, and they are *agents* because they take action (Chopra and White 2004; Chopra 2010; Floridi 2014, 2015).

The use of such a term could be a first step in the development of a sound legal framework that aims at achieving transparency and accountability through some form of auditing, and could help us avoid falling into a regulatory black hole in an effort to be too specific.

Attributing electronic personhood and liability – robots as electronic persons?

The European Parliament recommends creating a specific status for robots as '**electronic persons**', with specific rights and obligations, and applying it to cases where robots make decisions or interact with third parties (paragraph 59(f) of the section on Liability), a status currently unknown in our legal system.

The fact that some robots perform functions, such as managing operations, delegating tasks, resolving complex issues and making decisions in real time, raises legal questions. According to Parliament, attributing electronic personhood to them would resolve these legal issues.

According to current legal theory, granting legal personality to artificial agents is complex. It is not a case of simply equating robots to corporations. This opens the personhood debate, which has always been a source of controversy (as evidenced by the status of slaves or women in the past, or other beings and corporations more recently) (Chopra and White 2011).

What would be the implications of this approach in the case of robots?

Legal personhood means having rights and duties and the capacity to take civil action and be held accountable for their actions (even to sue and be held accountable for criminal acts). A *legal person* is also able to express moral values and become active politically. Once personhood is attributed to robots or to any other autonomous artificial agents, they become *subjects* (as opposed to things or objects) and enter the universe of *legal persons* (Chopra 2010; Chopra and White 2011).

> How *autonomous* could artificial agents be? Could they be aware of their intentions? Could they know that they are responsible and hence liable?

A discussion should be held about the nature of those rights and duties and the ability of autonomous artificial agents to exercise them. To illustrate this complexity, Samir Chopra reminds us that not all persons have the same rights and duties. Some of them depend on age, such as the right to vote, marry or do military service. Exercising these rights and duties require autonomy and the capacity to self-govern which, for autonomous artificial agents, could mean having the capacity to work without any human supervision.

In this connection, several questions emerge with regard to the European Parliament's Resolution: exactly how *autonomous* could artificial agents be? Could they be *aware* of their intentions? Could they *know* that

Number of multipurpose industrial robots (all types) per 10 000 employees in the manufacturing industry (ISIC rev.4: C) 2014

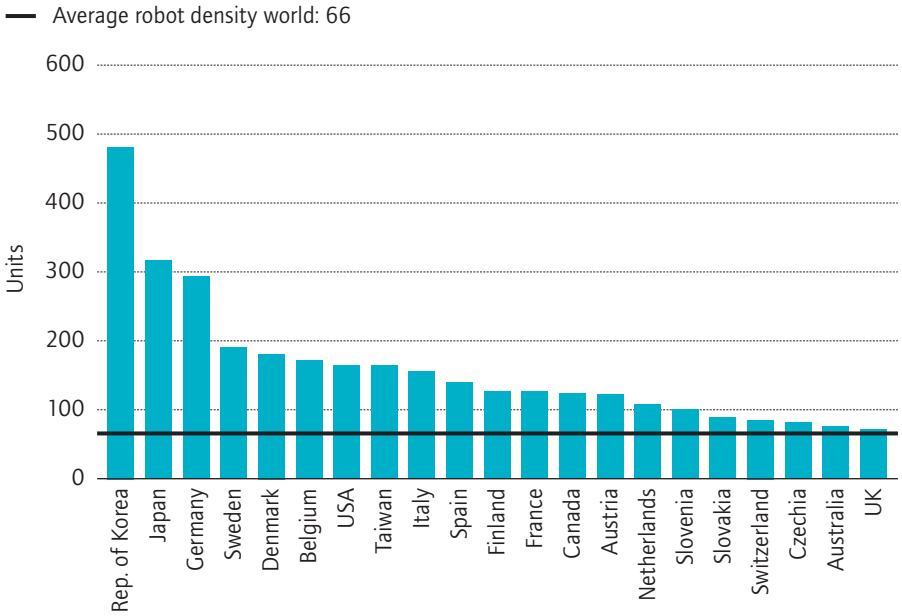

Source: IFR/World Robotics

they are responsible and hence liable? If so, they would have to assume moral responsibility and face the economic consequences of risks or accidents.

The Resolution identifies a solution based on insurance schemes and compensation funds (paragraphs 56 to 59 of the section on Liability).

This whole question of responsibility is a difficult one, and granting 'personhood' to autonomous artificial agents would require a fundamental shift in legal thinking. If that point is ever reached, humans and robots would indeed become players in the same legal framework.

What driver for European regulation?

In the case of complex technologies whose possibilities and outreach remain unknown, **potential risks** are equally unknown. Robots and artificial intelligence could replace humans and transform jobs, but at what cost? Some scientists, such as Stephen Hawking and Bill Joy, are concerned about the potential for abuse and have predicted the destruction of humanity (Hern 2016). Others see it from a different perspective and believe that, in order to surpass human beings, machines in their hard or soft version would need more than intelligence and deep learning: they would need to master abstraction, have the ability to understand the world visually and language semantically (Knight 2016), and be creative and even compassionate. Others, like Luciano Floridi, Professor of Ethics and Philosophy of Information at the University of Oxford, believe that there is little to worry about, for there are many things that machines cannot do, such as think, know and be conscious (Floridi 2016).

> Regulation should ensure that we have visibility and traceability over who is responsible for what. The aim is to promote a better interplay between makers and workers/users.

At all events, the purpose of regulation is to limit the possible misuse or 'stupid' use of power, and in so doing minimise the risks and potential adverse effects of new technologies. The development of exposure scenarios to identify levels of risk is essential. This applies to artificial agents and is why a proper regulatory framework needs to be envisaged, as opposed to soft measures such as codes of conduct, certification or standards.

Codes of conduct can help to promote ethical behaviour and introduce certain values and guidelines for professionals to follow, but they are not instruments for governance. When nanotechnology regulation was at the top of the agenda, the European Commission produced a code of conduct for responsible research in nanosciences and nanotechnologies (European Commission 2009). This code is aimed at all stakeholders in the field, but it has two main disadvantages: it is voluntary and its application is not monitored by any EU or national body.

Certification for designers is not the solution here either, as it is a privately driven action where certifiers, i.e. private companies, set their own criteria for what they want to certify. Their main objective is to compel designers to work in pursuit of specific interests, not to protect society, and they often fail to take societal views into account.

The same rationale applies to international technical standards. International standardisation organisations function on the basis of paid membership, which allows some actors to contribute to the development and implementation of technical standards. Again, the concerns here are that this process does not include the participation of societal stakeholders and that it cannot replace legislation owing to its voluntary nature (Hauert *et al.* 2015).

None of these three options leads to effective governance or solves the issue of accountability. Regulation of robots and artificial intelligence is crucial, but not because of a fear that they might take decision-making powers away from humans. This paper advocates that the driving force behind European regulation should, in simple terms, be to ensure that we have visibility and traceability over who is responsible for what. The aim is to promote a **better interplay between makers and workers/users**, whose relationship is very intense and closely intertwined. Artificial agents and humans should co-exist, but the former should remain useful to humans, while respecting human values and the rules of democratic society.

Conclusion and way forward

The regulatory aspects of robotics and artificial intelligence are starting to be discussed in the institutions of the European Union where various stakeholders are engaging in the debates and bringing their own values and interests to the table. This paper has sought to analyse the fundamental issues related to the difficulties inherent in defining robots and their categories, and highlights the concerns associated with 'electronic personhood'. It proposes that, before legislating issues such as human enhancement, the type of insurance that would be required for robots, the way in which a European agency for robotics should be organised and what its remit would be, etc., Parliament and the European Commission should consider fundamental questions related to responsibility and liability. As Broersen pointed out (2014), 'our tendency to delegate responsibilities to artificially intelligent systems will become a serious problem' for our society and for our legal systems globally. Therefore, with a view to regulating artificial agents, this paper identifies a number of key aspects that need to be addressed and taken into consideration before we become entrenched in debates that, more likely than not, will end in political compromise.

'Our tendency to delegate responsibilities to artificially intelligent systems will become a serious problem' for our society and for our legal systems globally.

Rules should be developed to clear the obfuscation that surrounds who does what among the makers, designers, data scientists, suppliers and companies responsible for creating artificial agents. Ensuring the **visibility** of actors and their responsibilities is a first step towards governing automation and artificial intelligence. In other words, we need to be able to ascertain who is responsible for what and ensure that these actors are **traceable** (Mittelstadt *et al.* 2016).

To that end, it is also necessary to be able to identify all the other **actors who interact with and use artificial agents**, such as workers, employers, consumers, patients, users and trainers.

Minimum regulatory standards need to be developed in order to attribute **responsibility and liability** in cases where the artificial agent has 'learning and teaching' features and is able to exercise unintended outcomes (Grodzinsky *et al.* 2008; Vanderelst and Winfield *et al.* 2016). These standards need to ensure that appropriate adjustments are made with a view to eliminating, as far as possible, any undesirable or non-programmed behaviours or consequences, and to address the issue of accountability should such behaviours or consequences arise.

The EU needs to implement requirements to exercise the '**right to explanation**' of models and decisions made by automated or artificially intelligent algorithmic systems.

The EU needs to implement binding requirements to exercise the '**right to explanation**' of models and decisions made by automated or artificially intelligent algorithmic systems, as already laid down in the General Data Protection Directive adopted in 2016 (Wachter *et al.* 2016). With specific regard to workers, regulation should guarantee their right to have the logic, functionality and consequences of automated decision-making systems explained to them, and it should identify when human involvement occurs and when a decision can be contested. Put simply, 'workers don't need to know the code, but they do need to know what the code seeks to achieve.'

We are creating a gigantic new space, full of machines and data, that is transforming our environment and our policies. It is crucial that regulation be enacted in response to this new world that we are building. It is also crucial that all key actors be made visible, accountable and ultimately liable, given that the ultimate decision about the purpose of a design, the relationships at play in the use of a machine, safety levels and the potential risks to society, health and the environment rests almost exclusively with them.

Given the integration of artificial agents into our society, with particular regard to their deployment in the workplace, we should look beyond the use of artificial agents purely for the purpose of increasing productivity and profitability. The consequences of technological development should be understood in reference to societal actors (Winner 1980). An effective regulatory framework is ultimately required in order to ensure that artificial agents co-exist harmoniously with humans and that they are specifically designed for, operate according to and are capable of adapting to human values and needs.

References

Asimov I. (1942) Runaround, in Asimov I., I, Robot, London, HarperVoyager, 2001.

Broersen J. (2014) Responsible intelligent systems. The REINS Project, KI - Künstliche Intelligenz, 28 (3), 209-214. DOI 10.1007/s13218-014-0305-4

Chopra S. (2010) Rights for autonomous artificial agents?, Communications of the ACM, 53 (8), 38-40.

Chopra S. and White L. F. (2011) A legal theory for autonomous artificial agents, Ann Arbor, The University of Michigan Press.

Chopra S. and White L.F. (2004) Artificial agents – personhood in law and philosophy, Proceedings of the 16th European Conference on Artificial Intelligence, Valencia, August 22-27, 2004, 635-639.

Degryse C. (2017) Shaping the world of work in the digital economy, Foresight Brief, #01- January 2017, Brussels, ETUI.

European Commission (2009) Commission recommendation on a code of conduct for responsible nanosciences and nanotechnologies research & Council conclusions on responsible nanosciences and nanotechnologies research, Luxembourg, Office for Official Publications of the European Communities. http://ec.europa.eu/research/science-society/document_library/pdf_06/nanocode-apr09_en.pdf

European Parliament (2017) European Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)).

European Parliament (2009) European Parliament resolution of 24 April 2009 on regulatory aspects of nanomaterials (2008/2208(INI)).

Feitelson D. G. (2007) Asimov's laws of robotics applied to software, IEEE Software, 24 (4), 112. DOI 10.1109/MS.2007.100

Floridi L. (2016) Should we be afraid of AI?, AEON. https://aeon.co/essays/true-ai-is-both-logically-possible-and-utterly-implausible

Floridi L. (2015) Toleration and the design of norms, Science and Engineering Ethics, 21 (5), 1095-1123. DOI 10.1007/s11948-014-9589-x

Floridi L. (2014) The fourth revolution. How the infosphere is reshaping human reality, Oxford, Oxford University Press.

Graeme L., Harmon S. HE and Arzuaga F. (2012) Foresighting futures: law, new technologies and the challenges of regulating for uncertainty, Law, Innovation and Technology, 4 (1), 1-33. DOI 10.5235/175799612800650626

Grodzinsky F., Miller K. W. and Wolf M.J. (2008) The ethics of designing artificial agents, Ethics and Information Technology, 10 (2-3), 115-121. DOI 10.1007/s10676-008-9163-9

Hauert C, Bütschi D., Graz J.-C., Audétat M. and Kaufman A. (2015) The international standardisation arena and the civil society participation stakes: results of the INTERNORM project, Policy Brief 14.2015, Brussels, ETUI. http://www.etui.org/Publications2/Policy-Briefs/European-Economic-Employment-and-Social-Policy/The-international-standardisation-arena-and-the-civil-society-participation-stakes-results-of-the-INTERNORM-project

Hern A. (2016) Stephen Hawking: AI will be 'either best or worst thing' for humanity, The Guardian, 19 October 2016. https://www.theguardian.com/science/2016/oct/19/stephen-hawking-ai-best-or-worst-thing-for-humanity-cambridge

Knight W. (2016) AI's unspoken problem, MIT Technology Review, 19 (5), 28-37.

Kritikos M. (2016) Legal and ethical reflections concerning robotics, STOA Policy Briefing, June 2016 – PE 563.501, Brussels, European Parliament Research Service. http://www.europarl.europa.eu/RegData/etudes/STUD/2016/563501/EPRS_STU(2016)563501(ANN)_EN.pdf

Mittelstadt B. *et al.* (2016) The ethics of algorithms: mapping the debate, Big Data & Society, December 1, 2016, 1-21. DOI: 10.1177/2053951716679679

Nanowerk (2013) Definition of the term 'nanomaterial'. http://www.nanowerk.com/spotlight/spotid=30804.php

Nevejans N. (2016) European civil law rules in robotics, Directorate General for Internal Policies. Policy Department C: Citizens' Rights and constitutional Affairs, Study PE 571.319, Brussels, European Parliament.

Palmerini E. *et al.* (2014) Guidelines on regulating robotics. RoboLaw Deliverable D6.2.
http://www.robolaw.eu/RoboLaw_files/documents/robolaw_d6.2_guidelinesregulatingrobotics_20140922.pdf

Vanderelst D. and Winfield A. (2016) An architecture for ethical robots, arXiv:1609.02931v1 [cs.RO] 9 September 2016. https://arxiv.org/pdf/1609.02931v1.pdf

Wachter S., Mittelstadt B. and Floridi L. (2016) Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation (December 28, 2016), International Data Privacy Law, forthcoming. Available at SSRN: https://ssrn.com/abstract=2903469

Winfield A. (2012) Robotics. A very short introduction, Oxford, Oxford University Press.

Winner L. (1980) Do artifacts have politics?, Daedalus,109 (1), 121-136. http://www.jstor.org/stable/20024652?origin=JSTOR-pdf&seq=1#page_scan_tab_contents

All links were checked on 06.09.2017.

Foresight Brief

The ETUI is financially supported by the European Union. The European Union is not responsible for any use made of the information contained in this publication.