The Rise of Robots and the Law of Humans*

Horst Eidenmüller†

I. Artificial Intelligence (AI) and the law

When lawyers enter the discussion, the fun part is usually over. Engineers and computer scientists enjoy a similar reputation. In this article, I consider robots and the law. The prospects for entertainment may therefore be limited. However, the interaction of law and Artificial Intelligence (AI) poses exciting and important questions, and the answers to these questions will undoubtedly shape the future of mankind in the decades to come.

AI is now rapidly changing how we live and work. As routine tasks (both manual and cognitive) become increasingly automated, it is anticipated that robots ('embodied AI'[1]) will take approximately 1/3 of jobs in traditional professions by 2025.[2] The law will shape the future of AI. It will determine the permissible uses of AI, the costs of new products and technologies, among other things. Further, the initial regulatory decisions will be crucial. They may create path dependencies and make it hard to change regulatory course later.

Regulating AI is going to be challenging and difficult. After all, the law is – and always has been – made by humans and for humans. Just think of fundamental concepts such as 'personhood' and 'legal personality'. Historically, these concepts related to humans, i.e. natural persons. AI will thus strain the legal system: how shall we deal with robots? Shall we

* This article is based on my inaugural lecture as the Freshfields Professor of Commercial Law at the University of Oxford on 9 June 2016. The lecture format has been retained.

† Prof. Dr. Horst Eidenmüller, LL.M. (Cantab.), Freshfields Professor of Commercial Law at the University of Oxford and Professorial Fellow of St. Hugh's College, Oxford.

[1] See *Calo*, Robotics and the Lessons of Cyberlaw, California Law Review 103 (2015), pp. 513, 532 et seq.

[2] See *Hernaes*, Artificial Intelligence, Legal Responsibility and Civil Rights, 22 August 2015, https://techcrunch.com/2015/08/22/artificial-intelligence-legal-responsibility-and-civil-rights/ (last visited on 10 March 2017). For a thorough treatment of the problem including sensible policy options, see *Ford*, The Rise of Robots: Technology and the Threat of Mass Unemployment (London: Oneworld, 2015).

Electronic copy available at: https://ssrn.com/abstract=2941001

accord them legal personality, give them the right to acquire and hold property and to conclude contracts, etc.?[3]

In this article, I attempt to answer these and other fundamental questions raised by the rise of robots and the emergence of 'robot law'. The main theses developed in this article are the following: (i) robot regulation must be robot- and context-specific. This requires a profound understanding of the micro- and macro-effects of 'robot behaviour' in specific areas. (ii) (Refined) existing legal categories are capable of being sensibly applied to and regulating robots. (iii) Robot law is shaped by the 'deep normative structure' of a society. (iv) If that structure is utilitarian, smart robots should, in the not too distant future, be treated like humans. That means that they should be accorded legal personality, have the power to acquire and hold property and to conclude contracts. (v) The case against treating robots like humans rests on epistemological and ontological arguments. These relate to whether machines can *think* (they cannot) and what it *means* to be human.

I will develop these theses primarily in the context of self-driving cars – robots on the road with a huge potential to revolutionize our daily lives and commerce.[4] However, in order to illustrate the massive potential influence that robots will have on the fabric of our societies, I begin with a broader range of examples.

II. Varieties of robots and robot features

1. Robot applications

Self-driving cars are currently among the most discussed robot developments.[5] Indeed, most car manufacturers have experimented with self-driving cars, and these cars are already being

[3] For a summary of the issues, see, for example, *Kersten*, Menschen und Maschinen, Juristenzeitung (JZ) 2015, pp. 1, 6-8.

[4] See *Eidenmüller*, Whose Fault? Firms, Products and Liability in the Age of Artificial Intelligence, BMW Welcomes: Artificial Intelligence, Munich 21 April 2016, https://www.youtube.com/watch?v=4b6qElmM_Ng&t=14s (last visited on 7 March 2017).

[5] See, for example, *Lipson/Kurman*, Driverless: Intelligent Cars and the Road Ahead (Cambridge, MA: MIT Press, 2016).

Electronic copy available at: https://ssrn.com/abstract=2941001

tested on roads worldwide.[6] Google appears to have the lead in this development.[7] A key feature of its car is a rotating rooftop camera. It consists of an array of 64 laser beams that create 3D images of objects, allowing the car to orient itself. The car's driving behaviour is controlled by complex software.

Another important application of robots is in medicine. For instance, prototypes of nanotech medical robots with a size of a 1-10/1,000,000 of a millimetre have been developed.[8] These nanotech robots will travel through a patient's blood and into tumours where they will deliver a therapy that turns off an important cancer gene.

Robots are also beginning to enter the finance and financial consulting industry. 'Robo financial advisers' might shake up the brokerage business with low-cost, automated investing.[9] For example, UBS recently announced that advisers in its American wealth management division will use a robot to cater to wealthy clients.[10] This technology, which will be used by the company's 7,000 advisers, has been developed by a San Francisco start-up, SigFigWealth Management, which is one of a growing group of robo-advisers.[11]

Finally, AI is also going to fundamentally change the legal profession.[12] Indeed, AI systems already assist in the (automated) resolution of disputes,[13] and 'robo-lawyers' are entering the stage. In 2016, for example, the world's first artificially intelligent lawyer was hired by the US law firm, BakerHostedler, which licensed 'ROSS-Intelligence' for use in its

---

[6] See '33 Corporations Working On Autonomous Vehicles', https://www.cbinsights.com/blog/autonomous-driverless-vehicles-corporations-list/ (last visited on 20 March 2017).

[7] See 'Google self-driving car project', https://plus.google.com/+SelfDrivingCar (last visited on 10 March 2017).

[8] See 'Nanotech Robots Travel Through Blood to Turn Off Tumor Cells', http://inhabitat.com/nanotech-robots-travel-through-blood-to-turn-off-tumor-cells/ (last visited on 10 March 2017).

[9] See *Bernard*, Should a robot oversee your retirement money?, New York Times of 3 May 2016, p. 15.

[10] See 'UBS to offer SigFig's robo-platform to its financial advisers', http://www.investmentnews.com/article/20160516/FREE/160519939/ubs-to-offer-sigfigs-robo-platform-to-its-financial-advisers (last visited on 10 March 2017).

[11] See https://www.sigfig.com/site/#/home/am (last visited 10 March 2017).

[12] See *Remus/Levy*, Can Robots Be Lawyers? Computers, Lawyers, and the Practice of Law, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2701092 (last visited on 11 March 2017); *Susskind/Susskind*, The Future of the Professions: How Technology will Transform the Work of Human Experts (Oxford: Oxford University Press, 2015), pp. 66-71; *Kaplan*, Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence (New Haven and London: Yale University Press, 2015), pp. 145-149.

[13] See *Fatima/Kraus/Wooldridge*, Principles of Automated Negotiation (Cambridge: Cambridge University Press, 2014). The leading private service provider in this area is Modria, see http://modria.com/ (last visited on 11 March 2017).

bankruptcy restructuring and creditor rights department.[14] 'Robo-lawyers' will be deployed especially with respect to document searches and classification in discovery.[15] In England, predictive coding (classification of documents for discovery) was recently backed by the High Court in *Brown vs BCA Trading* on 17 May 2016.[16] 'Robo-lawyers' will also be involved in the (online) drafting of legal documents[17], and 'smart contracts' based on blockchain technology are around the corner.[18]

AI might also come in the form of 'robo-judges'. Based on a data set of 150,000 US felony cases, *Kleinberg et al*. found that a release rule (pending resolution of the cases) based on machine learning predictions would enable us to reduce the jail population by 25% without any increase in the crime rate, or let us reduce crime rates by 20% without changing the jail population.[19] Taken together, robo-lawyering and judging seem to bring significant positive developments in making legal advice more affordable, judging more accurate, and improving access to justice for many.

2. Robot features

Reflecting on the examples discussed above, certain 'robot features' emerge that are important when thinking about regulating robots. As already mentioned, robots are a form of embodied AI. They consist of a sensor or other input mechanism, a controlling algorithm, and the capacity to give feedback to the outside world.[20] These three features together constitute

[14] See 'The world's first artificially intelligent lawyer was just hired at a law firm', http://www.businessinsider.com/the-worlds-first-artificially-intelligent-lawyer-gets-hired-2016-5?IR=T (last visited on 10 March 2017).

[15] See *Sobowale*, How artificial intelligence is transforming the legal profession, 1 April 2016, http://www.abajournal.com/magazine/article/how_artificial_intelligence_is_transforming_the_legal_profession (last visited on 20 March 2017),

[16] Brown v BCA Trading Ltd [2016] EWHC 1464 (Ch) (17 May 2016).

[17] On 'computational law' see, for example, *Kaplan*, Artificial Intelligence: What Everyone Needs to Know (Oxford: Oxford University Press, 2016), pp. 95-97.

[18] See, for example, *Eyers*, Lawyers prepare for 'driverless M&A' as smart contract era dawns, 19 June 2016, http://www.afr.com/technology/lawyers-prepare-for-driverless-ma-as-smart-contract-era-dawns-20160616-gpknyz (last visited 11 March 2017).

[19] See *Kleinberg/Lakkaraju/Leskovec/Ludwig/Mullainathan*, Human Decisions and Machine Predictions, NBR Working Paper No. 23180, February 2017, http://www.nber.org/papers/w23180 (last visited on 11 March 2017).

[20] See *Franklin*, Introduction, in: Calo/Froomkin/Kerr (eds.), Robot Law (Cheltenham, UK and Northampton, MA: Edward Elgar, 2016), pp. x, xi.

the so-called 'Sense-Think-Act Paradigm'.[21] The sensor or other input mechanism may draw from multiple sources as is the case, for example, with respect to 'networked cars'.[22] As described, they use cameras, traffic data from GPS, and geographical data taken from the internet.

Smart robots have machine learning capabilities, i.e. they not only use huge amounts of existing data, but also use data from experiences and other new information to adapt their behavior.[23] Therefore, to some extent, these robots are unpredictable by design.[24] An interesting question is whether robots, at some point in time, might reach a super-human intelligence level.[25] This is often referred to as 'singularity'.[26] Robots might learn how they are controlled by humans, taking over control themselves. For the time being, though, the debate about 'singularity' is more academic (and fictional) than real.

Amongst the various forms of robots, an import subclass is 'social robots'.[27] These are specifically designed for interactions with humans. Just think of toys for children that have human features.[28] Research has identified 'social valence' of 'social robots': humans treat anthropomorphic robots like humans.[29] This has potentially important regulatory consequences.

---

[21] See, for example, *Hexmoor*, Essential Principles for Autonomous Robotics (San Rafael, CA: Morgan & Claypool, 2013), p. 25.

[22] See, for example, *Lim/Balakrishnan/Eriksson/Gifford/Madden/Rus*, Intelligent Transportation with Networked Cars, https://groups.csail.mit.edu/drl/wiki/images/0/0f/LimMobisysDemo08.pdf (last visited on 26 March 2017).

[23] See *Brynjolfsson/McAffe*, The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies (New York, NY and London: W. W. Norton, 2014), pp. 89-96; *Kaplan* (note 17 *supra*), pp. 27 et seq. ("So four trends – improvements in computing speed and memory, the transition from physically to electronically stored data, easier access ..., and low-cost high-resolution digital sensors – were prime drivers in the refocusing of efforts from symbolic reasoning to machine learning", ibid. at p. 39).

[24] See *Millar/Kerr*, Delegation, relinquishment and responsibility: The prospect of expert robots, in: Calo/Froomkin/Kerr (note 20 *supra*), pp. 102, 107.

[25] See generally *Bostrom*, Superintelligence: Paths, Dangers, Strategies (Oxford: Oxford University Press, 2014).

[26] See *Bostrom* (note 25 *supra*), pp. 1-3; *Ford* (note 2 *supra*), pp. 225-245.

[27] See KPMG, Social Robots: 2016's new bread of social robots is ready to enter your world, 2016, https://assets.kpmg.com/content/dam/kpmg/pdf/2016/06/social-robots.pdf (last visited on 26 March 2017).

[28] On such toys see, for example, 'Netzagentur ruft Eltern auf, Puppe "Cayla" zu zerstören', 16 February 2017, http://www.sueddeutsche.de/digital/ueberwachung-im-kinderzimmer-netzagentur-ruft-eltern-auf-puppe-cayla-zu-zerstoeren-1.3383009 (last visited on 26 March 2017).

[29] See *Darling*, Extending legal protection to social robots: The effects of anthropomorphism, empathy, and violent behaviour towards robotic objects, in: Calo/Froomkin/Kerr (note 20 *supra*), pp. 213, 216 et seq.

III. Regulating self-driving cars

1. The potential of self-driving cars

The first thesis developed in this article is that robot regulation must be 'robot-context specific'. It requires a profound understanding of the micro- and macro-effects of 'robot behavior' in specific areas. I will illustrate this thesis through the example of accident liability for fully autonomous cars – a good test case for 'robot law'.

In 20 to 25 years, approximately 75 % of cars on the road will be self-driving.[30] There are a lot of significant positive developments associated with this trend. First, we are going to see significantly fewer car accidents and, even more importantly, fewer casualties: currently, approximately 1.3 million lives are lost every year in car accidents worldwide.[31] Approximately 90% of these fatal accidents are caused by human fault, 9% by environmental conditions and less than 1% by technical defects.[32] Second, we are going to witness increased mobility of persons with disabilities, the elderly, etc. Third, the opportunity costs of driving will be significantly reduced. On average, we spend more than four years driving during our lifetime,[33] and we are not always driving cars that are supposed to give you 'Sheer Driving Pleasure' like a BMW.[34] We could use the time spent behind the wheel for more productive activities like reading, thinking, or just dreaming. Fourth, we are going to see fewer cars on the streets because of car-sharing. This is going to free up parking space, reduce travel time and also emissions.

---

[30] See *Kaplan* (note 12 *supra*), p. 195 ("... the expert consensus is that 75 percent of the vehicles on the road will be self-driving in twenty to twenty-five years."); *HopgoodGanim*, Autonomous Car: The regulatory and liability challenges, 8 April 2016, http://www.hopgoodganim.com.au/page/Publications/Autonomous_Car_The_regulatory_and_liability_challenges/ (last visited on 11 March 2017).

[31] See *Gurney*, Sue my car not me: products liability and accidents involving autonomous vehicles, Journal of Law, Technology & Policy 2013, pp. 247, 250.

[32] See Statistisches Bundesamt, Verkehr auf einen Blick, 2013, p. 36, https://www.destatis.de/DE/Publikationen/Thematisch/TransportVerkehr/Querschnitt/BroschuereVerkehrBlick0080006139004.pdf?__blob=publicationFile (last visited on 11 March 2017).

[33] See '30 Surprising Facts About How We Actually Spend Our Time', 7 January 2015, http://distractify.com/old-school/2015/01/07/astounding-facts-about-how-we-actually-spend-our-time-1197818577 (last visited on 20 March 2017).

[34] See http://www.bmw.com/com/en/insights/brandcampaign_pleasure/experience.html?cm=Facebook (last visited on 9 March 2017).

However, there are also negative developments. First, we are going to see more unemployment in certain professions. Taxi and bus drivers, for example, will clearly be affected. Second, there is a higher risk of 'high magnitude accidents' occurring, since cars will be networked and therefore liable to devastating cyber attacks, etc.[35] Finally, the odd 'small' accident will also occur, albeit rarely. Two such incidents happened in the first half of 2016. In one case, a Google car crashed into a bus in Mountain View, California, when trying to change lane.[36] The Google car suffered some damage, whereas the bus was scratched only very lightly. Much more serious damage occurred in another case, when a Tesla car in self-driving mode crashed into a lorry in Florida, killing the driver of the Tesla car.[37] Apparently, the sensor system in the Tesla failed to spot the approaching lorry.

What the foregoing demonstrates, therefore, is that effective robot regulation must be tailor-made to the macro- and micro-impacts of such technologies – it must account for the social changes caused by *specific* robots, while also accounting for the risks and opportunities arising day-to-day.

2. Accident liability for fully autonomous cars

I will now proceed to use the issue of accident liability for fully autonomous cars to develop my second thesis: that (refined) existing legal categories are capable of being applied to robots, to sensibly regulate their behavior. Specifically, the question I would like to pose and answer is: who shall be liable *in tort* if an accident occurs? I will not examine how existing legal systems would deal with the issue right now; [38] rather, I am interested in the best

---

[35] See *Lin*, The Ethics of Autonomous Cars, 8 October 2013, https://www.theatlantic.com/technology/archive/2013/10/the-ethics-of-autonomous-cars/280360/ (last visited on 11 March 2017).

[36] See https://www.youtube.com/watch?v=I9T6LkNm-5w (last visited on 9 March 2017).

[37] See https://www.theguardian.com/technology/2016/jun/30/tesla-autopilot-death-self-driving-car-elon-musk (last visited on 9 March 2017).

[38] For the US see *Gurney* (note 31 *supra*), pp. 257-271; *Karnow*, The application of traditional tort theory to embodied machine intelligence, in: Calo/Froomkin/Kerr (note 20 *supra*), pp. 51, 61-74; for Germany see *Schulte-Nölke*, Europäisierung des Haftungsrechts, in: E. Lorenz (Hrsg.), Karlsruher Forum 2015: Europäisierung des Haftungs- und des Versicherungsvertragsrechts (Karlsruhe: Verlag Versicherungswirtschaft, 2016), pp. 3, 31 et seq.; *Lutz*, Autonome Fahrzeuge als rechtliche Herausforderung, Neue Juristische Wochenschrift (NJW) 2015, pp. 119, 119-121.

'solution' if one were *to design* a liability system from scratch. Thus, I will look at various possibilities: (i) nobody is liable ('the loss is where it falls'), (ii) the car manufacturer is liable, (iii) the AI (device) producer is liable (if different from the car manufacturer), (iv) the car owner is liable (there are no 'drivers' with respect to fully autonomous cars), and, finally (v), the car itself is liable. It might appear to be quite fanciful to consider the last option. But things are evolving rapidly. In 2015, a robot that was part of an art installation in Switzerland bought ecstasy and a fake Hungarian passport, amongst other things, on the dark web.[39] The robot, not the artist or another human, was arrested by the St. Gallen police and freed after three months.

In view of the above five options, it appears that strict liability of the car manufacturer – i.e. regulatory option two – should be preferred, and for several reasons.[40]

To begin, holding nobody liable certainly would not be an acceptable solution. Expected (accident) costs would not be internalized by the producers. Hence, they would have the wrong incentives, producing too many cars at too low costs and prices. At the same time, we would be very reluctant to buy and/or use autonomous cars under a 'no liability regime' and might even refrain from such activity altogether.

Under a regime of strict liability of the car manufacturer, no fault would be required to establish liability, just a defective product and causation. The manufacturer seems to be best positioned to control risks and balance the benefits and costs of the technologies that are 'driving' fully autonomous cars: the manufacturer develops and uses the AI, and it has an intimate knowledge of the relevant technologies. This is apparent, for example, from the immediate and sophisticated reaction of Google to the car accident described above. Google was very quickly able to identify what had caused the problem.[41]

---

[39] See http://fusion.net/story/122192/robot-that-bought-ecstasy-and-a-fake-passport-online-released-from-swiss-prison/ (last visited on 9 March 2017).

[40] See *Eidenmüller* (note 4 *supra*); *Eidenmüller*, Wenn Maschinen töten, Süddeutsche Zeitung of 13 July 2016, p. 2.

[41] See 'Google self-driving car crashes into a bus (update: statement)', 29 February 2016, https://www.engadget.com/2016/02/29/google-self-driving-car-accident/ (last visited on 26 March 2017).

Further, in contrast to a strict liability regime, a fault-based liability regime of the car manufacturer would require the courts to determine the proper amount of care to be exercised when manufacturing and selling autonomous cars. This, as with all new and complex technology, is a very difficult exercise. Further, a fault-based liability regime does not control the 'activity level' of the producer, i.e. the number of cars produced and sold.[42] Hence, a strict liability regime seems to be the better option.

What about regulatory option three, i.e. liability of the AI (device) producer if this producer is different from the car manufacturer? I submit that it would be very difficult to disentangle the accident causes in complex technology products, as a general matter, and with respect to fully autonomous cars, in particular. Also, the car manufacturer 'controls' the system, including all component parts. Hence, the car manufacturer probably is the 'cheapest cost avoider'.[43] As far as tort liability vis-à-vis third parties is concerned, it should therefore be the only liability addressee. Of course, the car manufacturer could seek indemnity from the device producer based on their contractual relationship if a defective AI device ultimately caused an accident. Holding the car manufacturer strictly liable involves the risk of the manufacturer falling insolvent and therefore not being able to pay up. To cover this risk, manufacturers should be required by the law to purchase product liability insurance.

Indeed, the 'solution' to the liability problem developed above seems to be the one into which the market and private contracting practice is moving. Late in 2015, for example, Volvo announced that it would take responsibility for the actions of its self-driving cars.[44] With this announcement, Volvo sent a strong signal to the market: we have confidence in our own technology. That signal created a lot of pressure on competitors to follow suit. The 'solution' developed above also largely mirrors the legal status quo under the Product

[42] See, for example, *Jackson/Kaplow/Shavell/Viscusi/Cope*, Analytical Methods for Lawyers (New York, NY: Foundation Press, 2003), pp. 404-405.

[43] If parties were able to bargain for the applicable liability rule at zero (transaction) costs, they would contract for liability of the party that is best positioned to avert the expected accident costs at the lowest costs (Coase Theorem). See *Coase*, The Problem of Social Cost, Journal of Law and Economics 3 (1960), pp. 1 et seq.

[44] See 'Volvo CEO: We will accept all liability when our cars are in autonomous mode', http://fortune.com/2015/10/07/volvo-liability-self-driving-cars/ (last visited on 9 March 2017).

Liability Directive 85/374/EEC in the European Union, with one exception: under the Directive, both the manufacturer of the car and the manufacturer of the component part would be jointly and severally liable (Articles 3(1), 5, 7(f)).[45]

One problematic aspect of a strict tort liability regime that holds the car manufacturer liable for accidents caused is the activity level of owners of autonomous cars. Thus, it is clear that the likelihood of accidents very much depends on the activity level: the more often a car is on the road, the more accidents will happen. It is likewise clear that the owner controls this activity level, depending on his or her personal preferences, business model, etc. Just think of a taxi company on the one hand and a private car owner on the other hand. Hence, controlling the activity level of car owners is thus an important element of an efficient liability regime.[46]

One can think of various potential 'solutions' to this problem. One would be co-liability of owners depending on their activity profile. Another could be tying the sale of the car to liability insurance with the premium determined by (i) the manufacturer, (ii) the type of car, and (iii) the owner / user profile. Such personalized insurance is available already today. It is often called 'black box insurance' because cars are fitted with a small 'black box' device, about the size of a smartphone, which records speed, distance travelled and the time of day or night that the car is on the road.[47] The device also assesses driving style by monitoring speeding and cornering. It further records the types of roads on which the car typically travels, and the times of day and night on which it is in operation, to build up a comprehensive driving profile.

[45] Under the Directive, the component manufacturer is not liable if the defect is attributable to the design of the product or to instructions by the manufacturer regarding the final product. Liability under the Directive is for damage caused by death or by personal injuries or damage to property other than the defective product itself, Article 9. The liability regime under the Directive does not affect contractual liability, Article 13.

[46] See, for example, *Shavell*, Liability for Accidents, in: Polinsky/Shavell (eds.), Handbook of Law and Economics, Volume I (Amsterdam: Elsevier, 2007), pp. 139, 146 et seq.

[47] See, for example, http://www.rac.co.uk/insurance/car-insurance/black-box-insurance (last visited on 9 March 2017); https://www.confused.com/car-insurance/black-box/telematics-explained (last visited on 10 March 2017).

III. Treating smart cars (machines) like humans?

So far, I have not discussed option five of the potential regulatory responses to accidents caused by fully autonomous cars, namely holding the car itself liable. However, this problem may be used to illustrate the third thesis of this article: that robot law is and will be shaped by the 'deep normative structure' of a particular society.[48] For example, how does a society's normative structure affect the response to the problem posed above, if this structure is utilitarian?[49]

As has already been mentioned, manufacturers cannot (fully) foresee the behaviour of smart cars because of machine learning. This is the starting point for thinking that we might, at some point in time, therefore, be forced to excuse the car manufacturer and instead hold *the car itself* liable. Conceptually, this would imply that we would acknowledge that a car has legal capacity. If it can be liable, we might conclude that it should also have the power to acquire property, conclude contracts, etc. Indeed, against a utilitarian background, there are a number of arguments that would seem to support moving in this direction.

First, smart cars *appear* (functionally) capable of purposive actions. They exhibit what can be called 'moral' or 'legal' agency, i.e. smart cars act similarly as humans would act in similar situations.[50] Second, we *treat* anthropomorphic robots like humans. Should we not then give them rights to send a signal against mistreatment of humans generally?[51] Protecting robots in this way would have a beneficial feedback effect on human interaction. Third, we accord corporations legal personality. There is a long debate amongst legal scholars whether

---

[48] See also *Eidenmüller* (note 4 *supra*); *Eidenmüller* (note 40 *supra*); *Eidenmüller*, Robots' Legal Personality, Oxford Business Law Blog 8 March 2017, https://www.law.ox.ac.uk/business-law-blog/blog/2017/03/robots%E2%80%99-legal-personality (last visited on 9 March 2017).

[49] An uncritical utilitarian perspective is adopted, for example, by *Schirmer*, Rechtsfähige Roboter?, Juristenzeitung (JZ) 2016, pp. 660, 663 et seq. ("[Roboter] sind insoweit rechtsfähig, wie es zweckmäßig ist, also immer dann, wenn entweder ihre Einordnung als selbständiger Akteur zu ihrem Schutz notwendig ist oder für den Rechtsverkehr einen Mehrwert bringt."). However, utilitarianism is of course a much disgraced legal philosophy. See, for example, *Eidenmüller*, Effizienz als Rechtsprinzip: Möglichkeiten und Grenzen der ökonomischen Analyse des Rechts (Tübingen: Mohr Siebeck, 4th ed. 2015), pp. 187 et seq.

[50] *Kaplan* (note 12 *supra*), p. 87.

[51] See *Darling* (note 29 *supra*), pp. 226-229.

corporate personality is based on a fiction (as *von Savigny* argued[52]) or whether there is something 'real' about corporations, whether there is something that could be called a 'real group-person' (as *von Gierke* argued[53]). Smart cars (robots) do seem no less 'real' than corporate persons. I should add that the word *persona* in Latin originally means mask or role. Roman law acknowledged various non-human right- and duty-bearing entities such as *municipia* or *collegia*.[54]

If we were to treat smart cars (machines) like humans and accord them legal personality, we could tie ownership of cars to liability insurance. Such treatment would also allow us to think about other innovative sanctions with respect to automobile accidents such as revoking the legal capacity of the car, detaining it for some time (like in the Swiss art installation case) or destroying it.

IV. The case against treating robots like humans

To be sure, most of us probably feel very uneasy when considering according smart cars (machines) legal personality. These apprehensions rest on epistemological[55] and ontological arguments, as I have suggested when introducing the theses developed in this article. Legal capacity or personhood, one can argue, is tied to humans because only humans understand the *meaning* of rights and obligations. Thinking involves active engagement with life, participation in culture, situation-specific 'know how' and common sense of the sort that can

[52] See *von Savigny*, System des heutigen Römischen Rechts, Zweyter Band (Berlin: Veit und Comp., 1840), p. 236 ("Die Rechtsfähigkeit wurde oben dargestellt als zusammenfallend mit dem Begriff des einzelnen Menschen (§ 60). Wir betrachten sie jetzt als ausgedehnt auf künstliche, durch bloße Fiction angenommene Subjecte. Ein solches Subject nennen wir eine juristische Person ...").

[53] See *von Gierke*, Rechtsgeschichte der deutschen Genossenschaft (Berlin: Weidmann, 1868), p. 1 ("Wie sich der Fortschritt der Weltgeschichte unwandelbar vollzieht, so erhebt sich in ununterbrochen aufsteigender Wölbung der erhabene Bau jener organischen Verbände, welche in immer größeren und umfassenderen Kreisen den Zusammenhang alles menschlichen Seins, die Einheit in seiner bunten Mannichfaltigkeit, zur äußeren Erscheinung und Wirksamkeit bringen.").

[54] See, for example, *Thomas*, Textbook of Roman Law (Amsterdam: North-Holland, 1976), pp. 387-388; *du Plessis*, Borkowski's Textbook on Roman Law (Oxford: Oxford University Press, 4th ed. 2010), pp. 85-87.

[55] For a general account of the epistemological issues raised by AI, see *Crane*, The Mechanical Mind: A Philosophical Introduction to Minds, Machines and Mental Representation (London and New York: Routledge, 3rd ed. 2016), pp. 77-90; *Kaplan* (note 17 *supra*), pp. 67 et seq.

never be formulated by rules.[56] Thinking is more than formal symbol manipulation (syntax). It involves sensitivity to the *meaning* (semantics) of these symbols, and robots don't have that sensitivity.[57] Robots can be programmed to conform to rules but they cannot follow rules.[58] Rule-following presupposes an understanding of the *meaning* of rules. Robots are not capable of such understanding. Robots are not active in the discipline of hermeneutics – and they never will be.

The second argument that I should like to submit, building the case against treating robots like humans, is an ontological one. The laws of a particular society in general, and the rights and obligations accorded to members of that society in particular, are an expression of the 'human condition'. Laws reflect what we believe is a precondition for an orderly intercourse between humans. But it also reflects what we believe lies at the heart of humanity, at the heart of what it means to be human. Just think of fundamental human rights in general and freedom of expression and speech in particular. But also think of such controversial issues as abortion or same-sex marriage. It simply and literally would be the dehumanizing of the world if we were to accord machines legal personality and the power to acquire property and conclude contracts, even though such machines may be smart – possibly even smarter than humans. So, treating robots like humans would dehumanize humans, and therefore we should refrain from adopting this policy.[59]

Is corporate personhood a good counter-argument, weakening the case against treating robots like humans? This paper submits it is not. The crucial difference between a corporation and a robot is that corporations always act 'through' humans. Humans sit on the boards whose actions are attributed to the corporation. It is true that also in firms, AI is becoming more and

---

[56] This point was succinctly made by *Dreyfus*, What Computers Still Can't Do: A Critique of Artificial Reason (Cambridge, MA: MIT Press, 1992).

[57] *Searle*, Minds, brains, and programs, Behavioral and Brain Sciences 3 (1980), pp. 417-457.

[58] As should be clear from the text, this assessment is based on a conception of rule-following that is richer than the reductionist approach advocated by *Kripke*, Wittgenstein on Rules and Private Language (Cambridge, MA: Harvard University Press, 1982) (*Kripke* suggested a 'sceptical solution'. This takes the form of a social criterion of rule following: an individual follows a rule insofar as its behavior conforms to that of other members of their community.).

[59] On the threats for the 'human condition' by 'Big Data', see generally *White*, We, Robots: Staying Human in the Age of Big Data (Brooklyn, NY and London: Melville House, 2015).

more important.[60] At the same time, it is still the case that final decisions are taken by humans – at least for the time being and before 'singularity' is reached, i.e. a super-human intelligence level.

V. The (policy) road ahead

Ultimately, AI poses intricate new regulatory issues with a huge potential societal impact, which should be now clear from the foregoing discussion. As noted, the law will have to decide upon permissible forms of AI and upon a possible legal obligation to use (certain) AI devices in specific settings. Just think about 'medical experts' that are smarter than humans. The law will have to define rights and obligations of robots and/or their owners. The governance of firms will change fundamentally by being run more and more on AI.

Most importantly, however, the question of *access to AI* must be raised. There can be no doubt that such access is a significant source of power. Potent private actors might leverage themselves with smart technologies to shape transactions to their advantage. Less sophisticated parties may lose out. Are certain smart technologies public goods? Should they be accessible at low costs to all?

These and other issues raise deep philosophical problems, and robot law will be shaped by what I have called the 'deep normative structure' of a society. It very much matters whether a society is based on a utilitarian conception of 'the good' or whether it is rather based on a humanitarian or Kantian vision according to which not everything that is utility or wealth-maximizing is necessarily the better policy – quite to the contrary. What seems to be clear is that a utilitarian conception of 'the good' will tend to move a society into a direction in which robots eventually will take a prominent role – by virtue of the law.

This is one of the reasons why robot law will probably be characterized by much regulatory diversity and regulatory competition:[61] given significant differences in the 'deep

[60] See *Fenwick/Kaal/Vermeulen*, The 'Unmediated' and 'Tech-Driven' Corporate Governance of Today's Winning Companies, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2922176 (last visited on 26 March 2017).

normative structure' of different societies, it will be quite difficult for states to agree on common policies. Further, the incentives for states to attract investment in AI will also spur regulatory competition. It seems likely that 'utilitarian states' will enact more 'robot friendly' laws, putting pressure on other jurisdictions to follow suit.

Finally, as discussed above, AI will also fundamentally change law-making and the legal profession. This raises the intriguing question whether, at some point in time, smart (AI-based) law-making will assist us in regulating AI products and services. It is beyond doubt that smart technologies will be a great aid in enhancing the efficiency of law-making on a technical level. It is quite a different matter whether these technologies will be able to assist us to tackle complicated regulatory problems that require intricate value judgments. On this point, I am deeply sceptical. Again, machines cannot think, nor solve deep philosophical problems. Nobody expressed this more clearly than Goethe: 'Nur allein der Mensch vermag das Unmögliche: er unterscheidet, wählet und richtet …'.[62]

By the way, scientific research and writing can nowadays also be assisted by AI. There is software out there that writes scholarly papers.[63] Unfortunately, I did not find any that would have assisted me in the preparation of this article. It is all too human – hence its shortcomings.

---

[61] On regulatory competition in different fields of the law, see, for example, *Eidenmüller*, The Transnational Law Market, Regulatory Competition, and Transnational Corporations, Indiana Journal of Global Legal Studies 18 (2011), pp. 707 et seq.

[62] *Goethe*, Das Göttliche, 1783, Verse 37-40.

[63] See http://connectedresearchers.com/online-tools-for-researchers/ (last visited on 26 March 2017).