CIVIL LAW RULES ON ROBOTICS

Introduction

1. This is the response of the General Council of the Bar of England and Wales ("The Bar Council") to the spring 2017 consultation conducted by the Legal Affairs Committee of the European Parliament (JURI) on EU Civil Law rules for Robotics[1]. The Bar Council represents over 15,000 barristers in England and Wales. It promotes the Bar's high quality specialist advocacy and advisory services; fair access to justice for all; the highest standards of ethics, equality and diversity across the profession; and the development of business opportunities for barristers at home and abroad.

2. A strong and independent Bar exists to serve the public and is crucial to the administration of justice. As specialist, independent advocates, barristers enable people to uphold their legal rights and duties, often acting on behalf of the most vulnerable members of society. The Bar makes a vital contribution to the efficient operation of criminal and civil courts. It provides a pool of talented men and women from increasingly diverse backgrounds from which a significant proportion of the judiciary is drawn, on whose independence the Rule of Law and our democratic way of life depend. The Bar Council is the Approved Regulator for the Bar of England and Wales. It discharges its regulatory functions through the independent Bar Standards Board.

3. This paper draws on the experience and knowledge of practitioners who practise in the following areas: aviation; insurance; personal injury; product liability; road traffic accidents; European and conflict of laws. In the time available, the Bar has chosen to focus its response on Part 2 of the specialised questionnaire, in particular on the concrete policy options on the themes of ethics and (tortious) liability. Following the logic of the Bar's analysis of these topics, we are presenting our views below in the reverse order to that in the questionnaire. We also make a couple of points on data protection.

4. We note, however, that the Bar has significant expertise and experience across other relevant fields such as contract (contractual liability of machine-to-machine contracts is one of the areas we consider needs most urgent attention), employment, and intellectual property law, to name but a few. In addition to contributing further on the topics examined below therefore, we would be delighted to provide input across these other areas, as the EU institutions' work on this file develops. We will also be examining the potential impact of AI

[1] https://opinio.secure.europarl.europa.eu/opinio/s?s=secured_robotics&lang=en

on the legal system itself, which data we will share going forward. We thus take this opportunity to indicate our willingness to participate in Working Group(s) and other fora (drawing together industry experts, academics and legal practitioners) that the EP and the other EU institutions have, or may, set up to discuss and agree these issues going forward.

5. The Bar Council is registered on the EU Transparency Registry, and we are content for reference to be made to our remarks in the context of this consultation.

Liability (Section 2 of the specialised questionnaire)

6. The Bar considers that within the body of European law, particularly in the transport and insurance field, there are present the underlying principles which permit EU law to build a robust and sufficient regulatory framework without the need for the invention of a sui generis legal structure. In this regard, we reject the suggestion that it is necessary or possible to personalise AI/robotics or to ascribe legal personality to robots, at least during the transitional phase (the next 10 to 15 years) with which this paper is concerned.

7. By applying principles to be found in the fields of motor vehicle accidents and aviation accidents, it is possible to separate principles of insurance protection from issues of liability.

8. For the reasons set out below: we strongly support the provision of compulsory insurance to cover the owner/user/operator of AI; and the provision of a compulsory insurance fund in relation to accidents involving uninsured and untraced robots. If there is a compulsory system of insurance, including a direct right of action against the insurer by the innocent injured third party, then such a right should be available as a matter of strict liability, as against the insurer only, whenever there is an accident involving a robot. Of lesser importance is the need for harmonisation of rights of subrogation or apportionment of liability between the insurer and the person at fault for the accident (whether the manufacturer; owner; operator; repairer of the robot, or another party involved in the accident).

9. Given the suggested imposition of a strict liability regime, care will need to be taken to define what is a robot/AI which requires a system of mandatory insurance.

Establishing a system of Universal Insurance Coverage

10. An injured Third Party is best protected where that person has a direct right of action against an insurer in circumstances where there is a mandatory (public) duty to cover that risk (cf the Motor Insurance Directives System).

11. Mandatory public liability accident insurance should in principle be available and carried by the owner/operator/user of the AI robot in any case where the AI robot is involved in an accident which causes damage.

12. Where such persons are operating the AI robot as part of commerce or a public function, then those persons can be regulated, licensed and subject to mandatory insurance.

13. Where the activity is a regulated activity (for example driving a car), then the requirement to carry mandatory insurance can be applied to that activity.

14. Where the AI robot is being operated in a private context (e.g. teenagers flying drones), then it should be mandatory for the sale of the robot to a private individual to include mandatory life-time (that is for the life-time of the robot) insurance cover in relation to accident liability. That insurance should include cover for persons who receive the robot as a gift or acquire it by way of subsequent onward sale and therefore are an insured person but not the original purchaser of the robot.

15. Sale of the insurance cover should include a system for the registration of the identity of the insurer and the registration of the robot equipment (manufacturing number/identification number). (Cf the system for the identification of motor insurers; some countries also have systems for employer's liability insurance in relation to latent industrial diseases.)

16. In this regard, we consider that the creative use of technology – for example the use of transponders and the requirement to register use of a drone – may better facilitate the design of a system which allows for the easy identification of the owner/user/operator and the relevant insurer when things go wrong.

17. Just as for motor insurance, minimum levels of insurance cover should be required for private users and there should be limited defined circumstances when insurers can avoid cover. The minimum levels of cover for robots should be at least the same levels of cover required for motor vehicles (without prejudice to Member States being able to provide higher levels of cover).

18. Also, just as for motor vehicles, there should be a Robot Insurers Fund (RIF) for Untraced or Uninsured Robots involved in an accident which has caused loss and damage. (Cf the Motor Insurers Directives System). Such a system should apply to personal injury and property damage claims.

Civil Liability and Protection of the Injured Third Party

19. Given the complexities of AI in relation to establishing what went wrong and why, any burden of proof should not rest with an injured Third Party to establish causation and fault in relation to an insured liability.

20. It is sufficient to ensure adequate protection of the injured Third Party if there is no-fault liability arising out of an accident (cf aviation liability), an accident being defined as an event which causes damage. We prefer the concept of an accident, rather than providing that liability should be demonstrated where the robot/AI "caused" loss and damage. Causation would inevitably require the Claimant to prove underlying facts as to how the accident occurred, which could involve complex issues concerning the design and construction, maintenance and operation of the AI, including access to software and data, outwith the knowledge and understanding of the injured Claimant. It might also require the Claimant to

sue both the AI insurer and human third parties where the accident was potentially caused by two or more parties.

21. In the aviation sector, establishment of strict liability with capped limits (absent any fault) did not stunt the growth of the aviation industry, and we see no reason not to adopt a similar model, with there being an uncapped liability where damage is done either (i) with intent or recklessly and with knowledge that such damage would probably result or (ii) negligently (with the burden being on the defendant to prove an absence of fault) for claims over the minimum liability limit, and/or the adoption of a regime with more generous minimum liability limits.

22. Accordingly, where there is an accident, an injured Victim is entitled to bring a claim by way of direct right of action against any insurer of any AI involved in that accident for compensation in relation to damage caused by the accident. Any issue of fault and apportionment of liability is a matter for determination at the secondary stage.

Specific Liability Regimes

23. So far as drones are concerned, they should in any event be classified as a form of aircraft, subject to aircraft liability rules.

24. In relation to driverless cars, either there is dual liability regime (one liability rule for autonomous vehicles or dual use vehicles in autonomous mode, and one rule for non-autonomous vehicles), or a single regime of liability for all vehicles involved in a road traffic accident. If there is strict liability for driverless cars and normal liability for other cars in cases where two cars are involved, one driverless and one not, the authors foresee, as a matter of insurance market economics, that this may lead to distortions of liability or distortions of insurance risk coverage, such that insurers of AI driverless cars carry a disproportionate burden. The alternative would be the mandatory application of the no-fault accident rule to all vehicles, driverless and not (per the Loi Badinter in France).

25. In relation to healthcare robots which interact with the person and/or are implanted into the person; it is likely that a specific liability regime will apply both to define what is damage/an accident and in relation to insurable risk. It is suggested that in any case where there is a risk which must be insured, there is either a care provider or a provider of medical services in respect of which there will be contractual or quasi-contractual duties of care. Any duty to insure can be applied to the provider of services and/or the provider of equipment.

26. In respect of healthcare robots, the concept of what constitutes damage may need careful consideration whether to include circumstances where there is damage to the robot requiring replacement (including surgical intervention) and events which cause distress and upset because something has gone wrong short of actual damage (also perhaps requiring medical intervention). Such an eventuality may more easily be regulated by the imposition of positive duties on the care provider or supplier. It is suggested that such regulation should be sector specific.

Capped liability

27. It is suggested that there are two scenarios which can give rise to substantial payouts: damage to multiple individuals and significant property damage (eg a robot causing a fire).

28. If, as currently, there are caps on insurance liability in the maritime and aviation sectors, and a minimum level of indemnity cover in the motor insurance field (all sectors where in the future there will be increased use of AI/robots), then it would appear inconsistent to permit those caps to continue whilst having an uncapped liability in relation to other AI/robots. Some level of cap is desirable in order to ensure that the development of this field is not stifled.

Secondary allocation of liability

29. The Bar considers that where the insurer carries the burden of compensating the loss and damage of the injured Third Party, there is then no need to harmonise the rights and entitlements of the insurer to pursue, by way of recovery action or subrogated claim, a claim for damages against the responsible person. Such harmonisation would inevitably have to deal with very difficult issues including any or all of the following: the establishment of rules of allocation of risk: whether on a fault basis; whether by the application of burdens of proof or presumptions; or whether on the basis of apportionment of risk.

30. The current conflict of laws rules provide for multiple liability and rights of subrogation. We see no need for the articulation of new rules of liability in relation to cross-border protection.

Ethics (Section 1 of the specialised questionnaire)

31. The field of AI/robotics ethics is vast and emergent; this response is limited to observations on questions posed in the consultation questionnaire. Although some of the questions are purely ethical, many bear on the law. The questions relate to risks (question 1), a common European definition of 'robot' (questions 2 and 3), a registration system for advanced robots (questions 4 and 5), a framework for ethical development of technology (questions 6 to 10) and provide an opportunity to make any other observations (question 11).

32. Of these, questions 2 to 5 are most obviously linked to the legal schema proposed in the consultation document. At paragraph 9 of this Paper, we have said: 'Given the suggested imposition of a strict liability regime, care will need to be taken to define what is a robot/AI which requires a system of mandatory insurance.' Paragraphs 10 to 18 above set out detailed recommendations.

Common European Definition (questions 2 and 3)

33. The consultation asks whether a common European definition is desirable and what such a definition should include.

34. As to whether a common definition is desirable, the Bar has concerns about (i) *terminology*, (ii) *feasibility*, and (iii) *utility*, of a definition.

35. Our concerns about *terminology* are as follows.

a. First, it is difficult to know what entity we are asked to define, because of the way that the consultation questionnaire is drafted. Question 2, asking for a definition, refers both to 'smart' robots and 'autonomous' robots. Question 1 referred to 'autonomous' robots. Other questions refer, in turn, to simply 'robots' or 'advanced robots'. It is difficult to propose a common definition of something which is itself described in varying ways.

b. Even if the previous concern was merely semantic, pedantic, or a function of translation, other concerns about terminology remain. It is unclear what relationship any of these terms bear to the question of 'artificial intelligence', and to the overall draft framework. For example, are smart/autonomous/advanced robots the same thing as 'artificial intelligence'? Or are they merely embodied instantiations of artificial intelligence? Is it then necessary to define artificial intelligence? Does the definition of smart/autonomous/advanced robots define the scope of the legislation, define the jurisdiction of the European Agency on Robotics/AI, and undergird all the regulation (including of liability)?

c. In terms of the latter question, the Commission is called upon to define 'smart autonomous robots *and their subcategories*'.[2] In the Bar's view, it would be preferable simply to define the subcategories and provide sector-specific regulation, e.g. of autonomous vehicles, RPAS (drones), etc. This can be done with greater precision, avoiding problems of overreach and underreach in definitional scope.

36. Problems of overreach and underreach lead onto the second concern, that of *feasibility*. The concerns about feasibility are as follows.

a) Even within the academic literature, the term 'robot' is notoriously difficult to define (see Lin et al 2011).[3] The law has even greater need for accuracy because real consequences flow from law.
b) The European RoboLaw project was only able to provide a 'tentative' definition, based on the interaction of five criteria.[4] It preferred to focus on specific examples.
c) The RoboLaw project raised a further concern: that of individuation. The authors of that project found it impossible 'to focus on a single technology since the majority of technologies do not work in isolation but rather as components of

[2] Paragraph 1 of the Motion ('General Principles').

[3] Patrick Lin, Keith Abney, and George Bekey 'Robot ethics: Mapping the issues for a mechanized world' (2011) 1755-6 Artificial Intelligence 942-949 <http://www.sciencedirect.com/science/article/pii/S0004370211000178>

[4] E. Palmerini, A. Bertolini, F. Battaglia, B.-J. Koopsc, A. Carnevale, P. Salvini, 'RoboLaw: Towards a European framework for robotics regulation' Robotics and Autonomous Systems 86 (2016) 78–85.

technological systems.'[5] Schafer asks: 'How is a component system within a larger system to be identified? How do we count robots, and how do we identify individual specimens? If a company owns one hundred cars, each with identical software, all communicating constantly with each other and a central server, is this one (distributed) robot, or one hundred? If the latter, why would an autonomous car, which will have several hardware and software components that constantly talk to each other and a central processor, not also count as several robots?'[6] It is only useful to define *a single* smart/autonomous/advanced robot if there is a genuine utility to doing so. This leads to the next concern.

37. A generalised definition of a smart/autonomous/advanced robot is only worth creating if it will serve a useful function. This raises the question of what purpose the definition is intended to serve. The Bar's concerns about *utility* are as follows.

a) The impression given by the Motion[7] is that the definition of smart/autonomous/advanced robots sets the scope of the legislation, defines the jurisdiction of the European Agency on Robotics/AI, and undergirds all the regulation (including of liability). Unless this is a mistaken impression, then it seems unduly laborious to construct such a definition. In the Bar's view, this burden is too great for a single definition to bear.
b) In the general law, it is rarely thought that one's moral/legal status by itself settles all questions, including questions of liability. For example, the question of whether a principal is liable for the actions of its agent do not turn on whether either or both are natural persons, corporate persons, etc. Detailed provisions of law are required to settle such questions. Indeed, a similar view was concluded in the Agentlink workshop discussions.[8]
c) Creating a new category of legal entity will require a lot of additional labour. It would be inefficient and imprudent to carry out that labour at this stage. We have not yet created general artificial intelligence. Such artificial intelligence as we have remains domain-specific. Without general artificial intelligence, there is no ontological basis for ascribing personhood or legal status to electronic persons. The only basis for doing so would be a functional basis. The functions that the status of person confers can, in the case of machines, be met through bespoke obligations and rights that are specific to current requirements. Thus the Bar considers that regulation should be domain-specific at least as long as the technology remains domain-specific.
d) Creating a new legal status will require the resolution of vexed questions,[9] including: *What are robots' rights* inter se, *when they interact with one another? Will a human person's*

[5] E. Palmerini, A. Bertolini, F. Battaglia, B.-J. Koopsc, A. Carnevale, P. Salvini, 'RoboLaw: Towards a European framework for robotics regulation' Robotics and Autonomous Systems 86 (2016) 78–85.

[6] Schafer, B 2016, 'Closing Pandora's box? The EU proposal on the regulation of robots' Pandora's Box – The Journal of the Justice and the Law Society of the University of Queeensland, vol 19, pp. 55-68.

[7] e.g. Motion Explanatory Statement, page 21.

[8] Emily Weitzenboeck, 'Electronic agents and the formation of contracts' (2001) 9 *International Journal of Law Information Technology* 204; Giovanni Sartor, 'Cognitive automata and the law: electronic contracting and the intentionality of software agents' (2009) 17 *Artificial intelligence and law* 253.

[9] For a sample, see: Mark Coeckelbergh, 'Robot rights? Towards a social-relational justification of moral consideration' (2010) 12(3) Ethics and Information Technology 209-221; Luciano Floridi and JA Sanders, 'On

right to employment prevail over that of an electronic person? etc. It is not yet necessary to settle these questions.

e) As above, the RoboLaw project raised a further concern about the utility of a definition: that of individuation.

38. In summary, there are concerns about terminology, feasibility and utility. Even if the first two concerns were overcome, it would still be premature to create a definition. Such a step (particularly if it equates to creating the status of 'electronic person') will not be necessary until the advent of general artificial intelligence. Given that paragraph 25 of the Motion asks for a legal instrument dealing with the next 10-15 years, **the Bar's recommendation is that the Commission concentrate its efforts on domain-specific artificial intelligence in both embodied and unembodied forms**. A short discussion on unembodied forms of artificial intelligence is found at the conclusion to this section.

39. Question 3 asks what a definition of smart/autonomous/advanced robots should include. Even if a definition is adopted, for whatever purpose, the criteria proposed in question 3 are unduly restrictive or otherwise unsuitable. For example:

a. As to criterion 3.1 and 3.4, the terminology is again unclear. As Lin et al (2011) say: 'An intuitive definition could be that a robot is merely a computer with sensors and actuators that allow it to interact with the external world; however, any computer that is connected to a printer or can eject a CD might qualify as a robot under that definition, yet few roboticists would defend that implication.'[10] As with the proposed framework, they adopt a criterion of autonomy. They acknowledge, however, that 'this definition merely postpones our task and invites another question: What does it mean for machines to have autonomy?'

b. As to criterion 3.2, the terminology is vague. The term 'self-learning' is inappropriate because machine learning is not self-directed. For example, a machine can be programmed to ignore certain datasets. The terms 'incremental or acquired' are better, because the learning may cease after a period, or the algorithm may simply make predictions following training on a static dataset. Thus the Bar recommends that this definition is replaced by something like: 'an algorithm or system that has incremental or acquired

---

the Morality of Artificial Agents' (2004) 14 Minds and Machines 349; Andreas Matthias, 'The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata' (2004) 6 Ethics and Information Technology 175; Dante Marino and Guglielmo Tamburrini, 'Learning Robots and Human Responsibility' (2006) 6 International Review of Information Ethics 46; David J Calverley, 'Imagining a Non-Biological Machine as a Legal Person' (2008) 22 AI & Society 523; Bertram F Malle et al, 'Sacrifice one for the good of many?: People apply different moral norms to human and robot agents' (2015) Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction. ACM, 2015; Burkhard Schafer et al, 'A fourth law of robotics? Copyright and the law and ethics of machine co-production' (2015) 23(3) Artificial Intelligence and Law 217-240; Woodrow Neal Hartzog, 'Et Tu, Android? Regulating Dangerous and Dishonest Robots' (2016) 5(3) Journal of Human-Robot Interaction 70-81; Toni M Massaro and Helen Norton, 'Siri-Ously? Free Speech Rights and Artificial Intelligence' (2016) 110(5) Northwestern University Law Review 1169-1194.

[10] Patrick Lin, Keith Abney, and George Bekey 'Robot ethics: Mapping the issues for a mechanized world' (2011) 1755-6 Artificial Intelligence 942-949 <http://www.sciencedirect.com/science/article/pii/S0004370211000178>

acuity at a domain-specific task(s)'. This describes machine learning. Machine learning is the most likely path to general artificial intelligence in any event.[11] It would be anachronistic to regulate robotics but not machine learning. Regulating domain-specific intelligence will lead most naturally into regulation of general artificial intelligence, if and when we come to it. The Bar therefore recommends that the remit of the Motion be extended to regulate unembodied machine learning.

c. As to criterion 3.3, again the terminology is vague. It is not clear whether the wording is intentionally different to that at page 13 of the Motion, where the proposal is 'the *form* of the robot's physical support.' Most things in the physical universe have a physical basis. Even a remotely-operating server could well count as a 'physical support'. That is to say, most code will need some sort of hardware to run it. The motivation behind this criterion may be to exclude 'pure' code or 'pure' algorithms, but we argue that this is mistaken. As stated above, the definition should be tailored to include rather than exclude machine learning. Thus, we recommend a definition like 'an algorithm or system that has incremental or acquired acuity at a domain-specific task(s)'.

d. As to criterion 3.5, this is very problematic and should be rejected.
   i. There is no consensus about what 'alive in the biological sense' means, even in the scientific world (see, e.g., the debate over viruses[12]). Further, the criterion 'not alive' is contestable philosophically too (e.g. is an early foetus 'alive'?); It will not serve any useful function except to transfer the debate into the question of what counts as 'alive' or 'biological';
   ii. Some robots have already been developed using only chemical and biological components.[13] These would presumably be unintentionally excluded by this criterion;
   iii. Other robots have been developed which are powered by chemical reactions in the human digestive system, dubbed 'ingestible electronics'.[14] Still other robots might be developed to combine biological components (e.g. enzymes) with mechanical ones, even outside the human body. Their proper classification is left ambivalent by this criterion;
   iv. There are human persons who may wish to attain the status of 'electronic person' (e.g. artist Neil Harbisson[15] and other frontrunners.[16] Further, those with biomedical implants, such as paraplegics, and those who wear advanced prosthetics, orthoses or exoskeletons, might reasonably wish this to be reflected in

[11] Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press 2014) page 61.
[12] Steven A Benner, 'Q&A: life, synthetic biology, and risk', BMC Biology 8 (2010) 77.

[13] <http://spectrum.ieee.org/the-human-os/biomedical/devices/celllike-robot-programmable-with-dna>

[14] <http://spectrum.ieee.org/the-human-os/biomedical/devices/powering-ingestible-electronics-with-gut-fluids>
[15] <https://www.theguardian.com/artanddesign/2014/may/06/neil-harbisson-worlds-first-cyborg-artist>

[16] <https://www.theguardian.com/artanddesign/architecture-design-blog/2015/aug/14/body-hackers-the-people-who-turn-themselves-into-cyborgs>

their identity or status).[17] This option should not be foreclosed without being given due consideration.

e. The problem may well be that 'robot' is a term that owes more to popular thought than to technical specification. The focus should rather be on the underlying technology whether this is embodied in a 'robot' or unembodied code in a computer.

Registration System for Advanced Robots (questions 4 and 5)

40. In view of the recommendations in our Paper above, the answer to question 5 should be that a registration system is 'voluntary for all advanced robots but compulsory for specific categories of robots'. Such robots will include, but are not restricted to, those over whom there is a scheme of mandatory insurance. Registration will be necessary to administer such a scheme.

41. It may be thought desirable that robots in high-risk scenarios, such as healthcare, should also be registered (regardless of whether mandatory insurance exists). Thought should be given as to how voluntary registration can be incentivised or encouraged. This could be done by a manufacturer, perhaps, if a customer applies for a warranty. No doubt further incentives or points of intervention are possible. If paragraph 22 of the Motion's General Principles is to be workable, a registration system will be required. And the more accurate the picture gained by the Commission is, the more successful monitoring and regulation will be.

The other questions (questions 1, 6 to 10)

42. Some brief observations on the remaining questions:

a. Question 1 concerns risks. It is difficult to say anything useful at this level of abstraction, unless the 'free text' option is chosen. In these early stages of development and adoption, human safety is a paramount consideration. However, as time goes on, the risk to human safety is more likely to come not from robots but from humans performing a task that is better left to automation. Conversely, threats to non-discrimination,[18] self-determination and privacy[19] will increase rather than decrease with time. Question 1 is a multiple-choice

[17] Likewise, an animatronic baby seal was granted a birth certificate in Japan, listing its inventor as its father: Jennifer Robertson, 'Human rights vs. robot rights: Forecasts from Japan' (2014) 46 Critical Asian Studies 571.

[18] For detailed consideration, see all of the annual conference papers of the Fairness Accountability and Transparency in Machine Learning organisation, which are available at <http://www.fatml.org> See further: <http://www.techrepublic.com/article/why-microsofts-tay-ai-bot-went-wrong/> ; Friederike Eyssel and Frank Hegel, '(S)he's got the look: gender stereotyping of robots' *Journal of Applied Social Psychology*, 2012, 42, 9, pp. 2213–2230; <http://www.techrepublic.com/article/big-data-can-reveal-inaccurate-stereotypes-on-twitter-according-to-upenn-study/>

[19] Calo points out: 'Robotics and artificial intelligence implicate privacy in at least three ways. First, they increase our capacity for surveillance. Second, they introduce new points of access to historically private spaces such as the home. Finally, they trigger hardwired social responses that can threaten several of the values privacy protects. Responding to the privacy implications of robotics and artificial intelligence is likely to require a combination of design, law, and education.' See: Ryan Calo, 'Peeping HALs: Making Sense of Artificial Intelligence and Privacy' (2008) 2 Eur J Legal Stud 168 <http://www.ejls.eu/6/83UK.pdf> and Ryan M Calo, 'Peeping Hals' (2011) 1755-6 Artificial Intelligence 940-941

question, so it is difficult to delve into the issues it raises in any serious way. However, some links to further reading can be found in the footnotes below.

b. Questions 6 to 10 concern the development of an ethical framework for technology. They involve consideration of who the framework should apply to, and the weighting of factors to be included in the framework. The ethical framework proposed in the Annex to the Motion (Detailed Recommendations) will be voluntary, not binding. If such a framework is to be capable of guidance, rather than being a list of vague desiderata, it will need to be far more fine-grained and context-specific. Rather than a single ethical framework, several frameworks are needed which are sensitive to different industries and different applications of AI/robotics. A European Agency for Robotics will be better-placed than the European Parliament or Commission to draft such guidance and monitor compliance.

c. Some excellent literature on robo-ethics can be found, but it is beyond the scope of this brief response to discuss the various proposals for ethical frameworks in depth. Please note the following, however:

   i. As to paragraph L of the recitals, regarding Asimov's laws, Riek and Howard point out: 'Asimov's laws of robotics… have so dominated discussion about the ethics of human-robot interaction as to eclipse the day-to-day ethical challenges facing [people in] research, development, and marketing.'[20] As Schafer notes: 'Asimov did not advocate them as a solution, if anything, his stories show how difficult it can be to reduce normative decision making to simple rule following.'[21] See also Weng et al (2009),[22] and Leenes & Lucivero (2014).[23]

   ii. To give an example of how fine-grained and context-specific guidance is needed, consider the following clause from the draft Licence for Designers: 'You should ensure that robots are identifiable as robots when interacting with humans'. This clause does not take into account the human tendency to anthropomorphise things

<http://acelscdncom/S0004370211000166/1s20S0004370211000166mainpdf?_tid=941bfee8f51c11e6a94400000aacb35e&acdnat=1487341557_60c64264c89df09002cbaefe62c2896c> .
See also Min Kyung Lee et al, Understanding users' perception of privacy in human-robot interaction." (2011) Proceedings of the 6th international conference on Human-robot interaction; Anrdrea Bertolini, 'Wearable Robots: A Legal Analysis' in J González-Vargas et al (eds) *Wearable Robotics: Challenges and Trends* (Proceedings of the 2nd International Symposium on Wearable Robotics 2017); Dag Sverre Syrdal et al, 'He knows when you are sleeping-privacy and the personal robot companion' (2007) Proc. workshop human implications of human–robot interaction, Association for the Advancement of Artificial Intelligence (AAAI'07) <http://www.aaai.org/Papers/Workshops/2007/WS-07- 07/WS07-07-006.pdf>

[20] Laurel D Riek and Don Howard 'A code of ethics for the human-robot interaction profession' (2014) <https://www3ndedu/~dhoward1/a-code-of-ethics-for-thehuman-robot-interaction-profession-riek-howardpdf>

[21] Schafer, B 2016, 'Closing Pandora's box? The EU proposal on the regulation of robots' Pandora's Box – The Journal of the Justice and the Law Society of the University of Queeensland, vol 19, pp. 55-68.

[22] Yueh-Hsuan Weng, Chien-Hsun Chen, and Chuen-Tsai Sun 'Toward the human–robot co-existence society: on safety intelligence for next generation robots' (2009) 14 International Journal of Social Robotics 267.

[23] Ronald Leenes and Federica Lucivero, 'Laws on Robots, Laws by Robots, Laws in Robots: Regulating Robot Behaviour by Design' (2014) 6(2) Law, Innovation & Technology 193-220.

that we interact with; see Duffy (2003).[24] We do this to our pets, and some therapeutic robots are effective precisely *because* they interact with humans in a relatable way. Thus PARO, the animatronic baby seal, is an effective companion only because the user can forget he is 'not real'.[25] Further, the clause does not take into account hybrid human-machine entities. This is the so-called 'Wizard of Oz' phenomenon, where remote human operators are behind at least some robotic behaviour.[26] This is not to say that no guidance is needed,[27] but rather that the various considerations cannot be dispensed with in a single sentence.

iii. Significant ethical concerns remain over automated weapons. The inference from the principle of non-maleficence in the Motion is that the development of such weapons is prohibited, at least ethically, in the European Union. In this regard, paragraph 34 of the Motion regarding the Dual Use Regulation is well-advised. There is not the scope to consider automated weapons and military robots further here.

iv. Common European definition: Returning to the discussion about the desirability of a common European definition, it may be more appropriate in the context of the Ethical Guidelines than in the context of law. If its purpose is merely to define the scope of the Ethical Guidelines, and if such Guidelines are purely voluntary, then less precision is needed than if the definition were to be legal. However, the Bar considers that the Ethical Guidelines should also be directed at embodied as well as unembodied machine learning. Therefore, a broader definition than the one proposed in the consultation is needed.

[24] Brian R Duffy, 'Anthropomorphism and the social robot' (2003) 42(3) Robotics and autonomous systems 177-190. <https://facultyutrgvedu/richardfowler/csci6174/papers/Duffy_RAS03pdf>
On the tendency to attribute stereotypically 'female' qualities to a long-haired robot and stereotypically 'male' qualities to a short-haired robot, see Friederike Eyssel and Frank Hegel, '(S)he's got the look: gender stereotyping of robots' *Journal of Applied Social Psychology*, 2012, 42, 9, pp. 2213–2230.

[25] PARO was granted a birth certificate in Japan, listing its inventor as its father: Jennifer Robertson, 'Human rights vs. robot rights: Forecasts from Japan' (2014) 46 Critical Asian Studies 571. Again, this demonstrates the complexities associated with anthropomorphism and the creation of a status of 'electronic person'.

[26] Laurel D Riek and Don Howard 'A code of ethics for the human-robot interaction profession' (2014) <https://www3ndedu/~dhoward1/a-code-of-ethics-for-thehuman-robot-interaction-profession-riek-howardpdf>

[27] Eric Leonardis and Ayse P Saygin 'Humanoid Robots and The Social Brain: Ethical Implications' (2015) 27 System 69-192 <http://wwwopenroboethicsorg/hri15/wpcontent/uploads/2015/02/Mf-Leonardis-Sayginpdf>
See also: Francesco Ferrari, Maria Paola Paladino, and Jolanda Jetten, 'Blurring Human–Machine Distinctions: Anthropomorphic Appearance in Social Robots as a Threat to Human Distinctiveness' (2016) 8(2) International Journal of Social Robotics 287.
<https://wwwresearchgatenet/profile/Jolanda_Jetten/publication/290973096_Blurring_Machine_Distinctions_Anthropomorphic_Appearance_in_Social_Robots_as_a_>

Data Protection

43. The level of data acquisition and storage necessary for AI manufacturers and users to be able to demonstrate the behaviour of AI – possibly over a period over years – is likely to pose significant challenges.

44. Drones and automated vehicles that scan their environment will necessarily acquire vast quantities of personal data, and a workable system of retention and data protection will be required.

45. Such data protection may well involve real time transmission of acquired data to storage facilities. This, in addition to the bandwidth required for the control and operation of AI may place significant pressure on usable bandwidth.

46. Adequate infrastructure is therefore a pre-requisite to the existence of a functioning regulatory system.

Conclusion and overall recommendation

47. The Bar Council congratulates JURI/EP for their foresight and consider that this developing area should be monitored and given detailed consideration. We applaud initiatives such as RoboLaw and the issuing of this consultation. However, the Motion is in one sense too ambitious, and in another sense not ambitious enough.

- It is too ambitious to the extent that it anticipates regulating anything approaching general artificial intelligence. The framework should concentrate on the next 10 – 15 years. A review can be undertaken then, or earlier if a technological breakthrough requires it.
- It is not ambitious enough to the extent that it makes provision for embodied artificial intelligence (robotics) but not unembodied artificial intelligence i.e. machine learning software and algorithms.
- 'Artificial intelligence' and 'robotics' are both enormous areas that require legislation. The Bar makes separate points as to both of these.

48. First, as to *robots*, context-specific laws are needed. It will become clear that they require regulation in different ways, from development to use. Indeed, how they are being used and who is using them also requires fine-grained analysis. Consider how it would be impossible to regulate the list below in one fell swoop:

- Autonomous weapons
- Medical devices – whether surgical aids, implants, wearables, etc. It may be appropriate for some to be regulated by the Medical Devices Directive while others require new and bespoke law.
- 'Sex bots'. It would leave gaps if, for example, there was there was regulation of hard-copy pornographic magazines and digital video, but not of sex robots. For example, would the development of life-like child-robots for sexual gratification be

permissible?[28] It is not possible to do justice to the discussion here, but see Richardson[29] (2016), and Sullins (2012),[30] among others.

- Care robots. The Paper above considers these. Further detail can be found in Simshaw et al (2016),[31] Sorell and Draper (2014),[32] and others.[33]
- Autonomous vehicles, considering the Society of Automotive Engineers' six levels of automation.[34] The Paper above considers these in depth.
- Drones (RPAS). The Paper above considers these in depth.
- Industrial robots. These are currently regulated by the Machinery Directive 2006/42/EC, and workplace safety legislation.

All these require regulation in different ways. Some may already be implicitly subject to regulation, e.g. by the General Product Safety Directive 2001/95/EC and the Consumer Protection Directive 1999/44/EC, or by sector-specific law (such as that governing road traffic or aviation). Interactions between these laws and new law should be examined so as to avoid unintended effect.

49. Secondly, as to *artificial intelligence*, the Bar recommends that the framework law is expanded in scope to embrace unembodied artificial intelligence i.e. machine learning algorithms and software. This is for the following reasons:

a) It would be anachronistic to regulate embodied AI but not unembodied AI. This would be to leave significant gaps in the law.
b) The European Parliament has already taken some steps toward regulating algorithms, such as within the General Data Protection Regulation. It would be better to undertake a complete review, rather than proceed by partial or piecemeal legislation.
c) Machine learning is the most likely path to general artificial intelligence in any event.[35] Regulating machine learning is the best way to anticipate developments and produce law incrementally.

---

[28] This might be thought to be a matter most suitable for national legislatures.

[29] Kathleen Richardson, 'The asymmetrical 'relationship': parallels between prostitution and the development of sex robots' ACM SIGCAS Computers and Society 453 (2016) 290-293 <http://ieeexplore.ieee.org.elib.tcd.ie/stamp/stamp.jsp?arnumber=7484884>

[30] John P Sullins, 'Robots, love, and sex: The ethics of building a love machine' (2012) 34 IEEE Transactions on Affective Computing 398-409 <http://ai2-s2-pdfss3amazonawscom/3b24/6eed37a61ce96ffa8baa08e30c502cac4d02pdf>

[31] Drew Simshaw et al., *Regulating Healthcare Robots: Maximizing Opportunities While Minimizing Risks*, 22 Richmond Journal of Law & Technology 3 (2016) <http://jolt.richmond.edu/v22i2/article3.pdf>

[32] Tom Sorell and Heather Draper 'Robot carers, ethics, and older people' (2014) 163 Ethics and Information Technology 163 183-195 <http://linkspringercom/article/101007/s10676-014-9344-7/fulltexthtml>

[33] e.g. A van Wynsberghe, 'Service robots, care ethics, and design' (2016) 18 Ethics Inf Technol 311–321.

[34] <https://www.sae.org/misc/pdfs/automated_driving.pdf>

[35] Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press 2014) page 61.

d) We are not yet at the point when we have our own personal robot butler, but many of us have personal assistants (such as Siri or Amazon Home). These have a current and significant impact on our lives. They are the challenge facing us *now*, and a pan-European response is appropriate.

e) Not only is machine learning pervasive across modern Europe, it is increasingly used to make decisions that have significant impact on the lives of individuals. A broad range of organisations across the private and public sectors use algorithmic decision-making to supplement, or even replace, human decision-making on matters from consumer credit to welfare and criminal justice.[36] For example, decisions about whether to grant bail can be made with the aid of COMPAS, a commercial recidivism prediction tool.[37] Alarie et al (2016) argue that: 'Machine learning algorithms can help predict outcomes of court cases. This will allow regulators to provide faster, more consistent, and more reliable rulings. Second, machine-learning algorithms can become the law. Machine learning can be used to not only reflect the law, but also to refine and improve the law. Machine learning algorithms can predict consequences of human behaviour. This will facilitate the development of laws that are context-specific, tailored to every possible scenario. Such advances will fundamentally change the structure of law.'[38] They developed a machine learning programme to predict the outcome of tax cases in Canada, and foresee machine learning becoming an integral part of law itself. Other examples of first-generation legal systems have also been discussed in the literature.[39]

f) It is imperative that a public conversation is had at European level about the permissible role of machine learning, the limits on the use of machine learning, and minimal procedural fairness standards for such machine learning tools as are used in public life.

g) The impact of machine learning is not limited to public life, however. It is already in use in the provision of legal services (for example, to scan potentially discoverable documents), insurance and banking (for example, High Frequency Trading algorithms). It extends to even more 'private' activities than these. The algorithms on social media and search engine websites can tailor advertising, suppress dissenting speech and create 'filter bubbles' wherein 'fake news' flourishes.[40] Albanie et al warn:

---

[36] Bryce Goodman, 'Economic Models of (Algorithmic) Discrimination' (2016) NIPS Symposium, available at <http://www.mlandthelaw.org/>

[37] Nina Grgi´c-Hlaˇca et al, 'The Case for Process Fairness in Learning: Feature Selection for Fair Decision Making' (2016) NIPS Symposium, available at <http://www.mlandthelaw.org/>

[38] Benjamin Alarie, Anthony Niblett and Albert H Yoon, 'Regulation By Machine' (2016) NIPS Symposium, available at <http://www.mlandthelaw.org/>

[39] e.g. Taxman (discussed in Thorne McCarty, 'Reflections on Taxman: An Experiment in Artificial Intelligence and Legal Reasoning' (1997) 90 *Harvard Law Review* 837) or the Latent Damage System (discussed in Richard Susskind 'The latent damage system: a jurisprudential analysis' (Paper presented at Proceedings of the 2nd international conference on artificial intelligence and law (ICAIL 89), University of British Columbia, Vancouver, 1989) 23–32). See also W M Campbell et al, 'Predicting and Analyzing Factors in Patent Litigation' (2016) NIPS Symposium, available at <http://www.mlandthelaw.org/>

[40] Samuel Albanie, Hilary Shakespeare and Tom Gunter, 'Unknowable Manipulators: Regulation of Curation in Social Networks' (2016) NIPS Symposium, available at <http://www.mlandthelaw.org/>

'There now exists the potential for such an algorithm to engage in the manipulation of its users for several qualitative reasons: 1. Access to vast quantities of user data combined with ongoing breakthroughs in the field of machine learning are leading to powerful but uninterpretable strategies for decision making at scale. 2. The availability of an effective feedback mechanism for the short and long term user responses to curation strategies. 3. Techniques from reinforcement learning have allowed machines to learn automated and highly successful strategies at an abstract level, often resulting in non-intuitive yet nonetheless highly appropriate action selection.'[41]

h) Quite apart from the role of machine learning in public affairs, it is imperative that consumers and vulnerable end users are protected from manipulation from commercial or even antisocial actors.
i) The above sketch is brief and incomplete but, for all the reasons given, machine learning (unembodied artificial intelligence) is just as deserving as robotics (embodied artificial intelligence) of regulation.

50. The overall recommendation is not that the proposed project (*a lex robotica*) should be abandoned. Rather, this project should be pursued. Initially, legislation should be made with only the next 10-15 years in mind. The scope of the project should be expanded to include unembodied domain-specific artificial intelligence (i.e. machine learning). It would be beneficial for all these spheres of regulation to be considered alongside one another, and at the same time. Nonetheless, a single overarching Regulation is unlikely to be suitable or adequate. Detailed and context-specific laws and ethical frameworks are needed.[42]

Bar Council
April 2017

For further information please contact:

Evanna Fruithof Consultant Director
The General Council of the Bar of England and Wales, Brussels Office
Avenue des Nerviens 85, Brussels B-1040
Telephone: 00 32 486 561968
Email: evanna.fruithof@barcouncil.be

[41] Samuel Albanie, Hilary Shakespeare and Tom Gunter, 'Unknowable Manipulators: Regulation of Curation in Social Networks' (2016) NIPS Symposium, available at <http://www.mlandthelaw.org/>

[42] per Recital O of the Motion.