

Catalogue of questions for the Committee on the Digital Agenda's public discussion on 'The effects of robotics on economics, labour and society', 22 June 2016: Professor Ryan Calo

1. The advancing development and increasing application of robotics, automation and artificial intelligence are changing the world of work and society in general. In what fields do you see new and/or particular opportunities and options for growth in Germany and Europe's industries and economies?

I am not familiar with the specifics of the German economy. My understanding, however, is that robotics and artificial intelligence present particular opportunities in agriculture, transportation, healthcare, and eldercare and will continue to transform the manufacturing sector. I also believe personal and service robotics will grow into a major industry in the coming decade.

What is required, from your point of view, if Germany and Europe are to be able to stand at the forefront of global competition in this segment of industry?

I believe Germany should continue to invest in basic research through public-private partnerships. Robotics in particular still has far to go. For example, for drone delivery to become a reality, drones will need better autonomous navigation, more efficient batteries, and more reliable mechanisms to sense and avoid obstacles. Also, as robots leave the factory and enter everyday life, human-robot interaction will become an increasingly important field.

I also believe the German government itself should accrue expertise in cyber-physical systems; without such expertise, no government will be in a position to generate a wise law and policy infrastructure for robotics and artificial intelligence.

What regulatory issues will legislators face?

Robots and artificial intelligence raise countless potential issues for regulators. Here are a few:

- *Safety standards. Regulators will have to know when a drone or driverless car is safe enough to be permitted to transport people or goods. The prospect of emergent learning behaviour will make this task particularly challenging.*
- *Certification. Robots are beginning to develop the capability to perform tasks otherwise reserved for certified professionals. For example, several teams are working on autonomous surgical robots. Where and by what standard should such devices be certified?*
- *Liability. In many instances, liability for robots will be straightforward. But there will be circumstances where ordinary legal principles break down. For example, German courts or lawmakers will have to decide whether an [open robotic platform](#) should be liable for third party applications the user chooses to run on the robot. Also, the prospect of [emergent behaviour](#) renders concepts such as foreseeability (in tort) and mens rea (in criminal law) less useful—at least in American law.*

- *Intellectual property. Patent law must attempt to strike a balance between fostering collaboration in robotics and incentivizing innovation, as with any technology. But artificial intelligence presents the further prospect that a system will either generate or infringe upon intellectual property.*
- *Privacy. Robots and artificial intelligence generate [novel privacy issues](#). For example, artificial intelligence may prove capable of making predictions and guesses about individuals or groups that fall outside traditional notions of gathering personal data. Moreover, because we are hardwired to react to anthropomorphic technology [as though it were a person](#), citizens may have fewer opportunities for solitude as interfaces become more social.*
- *Taxation. One plausible scenario says that robots and artificial intelligence will displace jobs in the short run and create jobs in the long run. If so, regulators may need to plan for a revenue gap in terms of income taxes.*

Apart from legal questions, for example concerning responsibility and liability, where will ethical questions, in particular, also arise with regard to the use of artificial intelligence or as a result of the aggregation of information and algorithms?

Robots and artificial intelligence raise just as many ethical questions as legal ones. We might ask, for instance, what sorts of activities we can ethically outsource to machines. Does Germany want to be a society that relegates the use of force, the education of children, or eldercare to robots? There are also serious challenges around the use of artificial intelligence to make material decisions about citizens in terms of minimizing bias and providing for transparency and accountability—issues already recognized to an extent by the EU Data Directive.

2. What expectations do you have of this development – do you expect one or several industrial revolutions – and if so, of what kind(s)? Might it even be possible to talk of a cultural revolution in this connection? Can comparisons be drawn with previous industrial revolutions?

I lack the expertise to address this question.

3. How do you view proposals for the appointment of an (independent) expert or ethics commission to answer social, legal and ethical questions, and potentially put forward concrete recommendations for action?

[I have argued](#) for the establishment of a Federal Robotics Commission in the United States because of my perception that U.S. regulators lack sufficient expertise in robotics and artificial intelligence to generate an adequate law and policy framework. Independently, I have argued that companies should establish [ethical review boards](#) for the research they conduct on consumers—including through the use of machine learning.

For example, is it necessary to impose restrictions on the deployment of algorithms or artificial intelligence for forecasting and concrete decisions, e.g. in the context of what is known as ‘predictive policing’? In your opinion, how will it be possible to ensure the non-discriminatory neutrality and transparency of automated decision-making processes? According to what principles will algorithms – self-learning algorithms, in particular – define decisions as correct or incorrect?

Decisions about people using algorithms could, in theory, have greater accuracy and less bias than decisions made entirely by people. The issue is that bias is very difficult to eliminate, in part because the designers of algorithms do not understand their own biases, and in part because training data can itself be biased. The German government could encourage the development of artificial intelligence that is transparent and accountable, i.e., that presents not just recommendations but reasons. The

trouble is that many of today's most promising machine learning techniques—neural nets, for instance—do not admit of scrutiny in this way. Thus, testing for disparate impact on the vulnerable and providing for opportunities for citizens to challenge algorithmically determined results becomes crucial.

4. It has been discussed over and over again in the course of digitisation whether jobs will fall victim to automation and robotics or whether new jobs can be created. What studies are you aware of, and how would you assess them? In what key technologies will robots replace human labour? Where do you see opportunities and options for the creation of new jobs in Germany and Europe, especially as a – direct or indirect – result of robotics and automation? Can the experience gained in other countries be applied seamlessly to Germany? What do you see as the advantages of the developments that are taking place – not only from economic points of view, but also as far as the design of modern workplaces, more flexible working time models and/or a better work/life balance are concerned?

I lack the expertise to address this question but recommend the book [The Second Machine Age](#) by Erik Brynjolfsson and Andrew McAfee at MIT.

5. What sort of role will human beings have in the future world of work? Will there be a blurring of the boundaries between the human and the machine?

For the foreseeable future, I hope and expect that artificial intelligence will mostly help humans make decisions rather than replace human decision-makers entirely. This arrangement is advantageous insofar as human-machine teams tend to outperform humans or machines alone. The challenge becomes training people who work with artificially intelligent systems to understand them as tools and not oracles. Robots will also help people perform physical tasks more safely and efficiently.

If so, in what fields will this happen and/or where is it already happening today?

The blurring of lines between person and instrument has been [happening for decades](#) in any number of fields. Among other issues, American courts have had to decide whether robots represent something “animate” for purposes of import tariffs, whether robots can “perform” as that term is understood in the context of a state tax on performance halls, and whether a salvage team “possesses” a shipwreck it visits with an unmanned submarine at maritime law. Moreover, many coteremporary systems—from airplanes to cars—blend human and machine control for safety reasons.

Where do you see ethical dilemmas? What solutions can you propose to overcome these dilemmas, where necessary?

An important ethical dilemma involves accountability. This problem runs both ways: We can imagine human decision-makers or actors “blaming” the machine for any mistakes. Conversely, we can imagine systems designed to keep a human being in the loop for the sole purpose of absorbing blame or liability should something go wrong—what M.C. Elish has [labelled](#) a “moral crumple zone.”

6. What role is Germany playing in the development of robotics and artificial intelligence? Is Germany competitive in international comparisons? What measures do you feel are necessary in order to ensure continuing success? What is required, from your point of view, if German and European start-ups in the field of robotics and automation are to be able to stand at the forefront of global competition?

One way the German government can foster innovation in robotics is to resolve the question of whether the manufacturer of an [open robotic platform](#) is liable when a user runs third party applications on that platform. As with personal computers or smart phones, personal robots will be

more useful if they are sufficiently open to third-party contributions of software and hardware. Just as you don't want Apple to have to write every app for the iPhone, you don't want a robot company to have to write every app for its home robot. But open robots also present the potential for inestimable legal liability, which may lead entrepreneurs and investors to abandon open robots in favour of products with more limited functionality so as to avoid liability for what users do with robots.

7. How do you view the creation of common standards at the European and international levels? From your perspective, is there a need for far-reaching amendments to legislation, e.g. in the field of labour law, and will such legislative processes potentially be accompanied by the lowering of existing standards of protection? What regulatory challenges will have to be coped with to ensure secure, efficient flows of data traffic for networked, automated systems?

I know too little about European law to answer this question well. I will say that, from my perspective, Europe has in place a robust data protection framework that will be relatively well-suited to address the privacy challenges posed by robots and artificial intelligence.

8. The development of artificial intelligence, in particular, is playing a major role in the discussion about robotics. What is your attitude to artificial intelligence?

The field of artificial intelligence is advancing rapidly, but narrowly. Many of the most promising techniques were developed in the 1970s. Suddenly, due to sufficiently low storage costs for data and access to greater and greater computing power, these same techniques yield astonishing results. For example, Skype can now provide real-time language translation and IBM's Watson can defeat our greatest trivia masters. Some systems can even learn to accomplish tasks in surprising ways. But artificial intelligence—and hence robots—remain fundamentally incapable of versatile problem solving, let alone forming intent or exercising judgment.

In future, will robots also be able to perform emotional tasks that enable them to replace human beings, e.g. in the care sector, road traffic or health services?

Robots can certainly detect and mimic emotions, which will make them more useful in a variety of settings. Note, however, that detecting and mimicking emotions does not in any way require robots to experience or understand them.

How will it be ensured that ethical criteria are complied with? Who will define these criteria, and can you give examples?

Robot ethics begins with designers. The people that make robots and design artificially intelligent systems should use existing methods such as value sensitive design (Batya Friedman) and responsible innovation (Jeroen van den Hoven) to attempt to anticipate and address stakeholder concerns. I also believe ethical review boards such as that of Deep Mind have a potential role—if conducted in a meaningful and transparent way.

This being said, I believe government can and should step in to police against abusive or manipulative behaviours. The fact that robots feel like [people opens the door](#) to questionable consumer practices.

The performance of such tasks would also involve value judgements and decision-making processes. Is it conceivable that automated systems and artificial intelligence will be able to judge situations and take decisions just as individually and situatively as human beings do? – Or even better?

There is a tension here. On the one hand, machines may be better able to make accurate and unbiased determinations in many circumstances. On the other hand, machines cannot display judgment, discretion, or compassion the way a person would. For the foreseeable future, artificially

intelligent systems should be thought of as tools and not oracles. This has several advantages. First, research suggests that human-machine teams perform better than humans or machines alone. Second, where people deviated from machine recommendations, we can reflect on why: Did they deviate because of bias or because of judgment or compassion?

9. Automation, robotics and artificial intelligence will increase rationalisation and boost productivity. For a long time, there have been proposals to introduce an 'automation dividend', 'robot levy' or 'robot tax', partly with a view to the future organisation of social security systems. Would such a levy be worthwhile or necessary in your opinion?

In 1984, in the United States, a petitioner formally requested that the Internal Revenue Service to impose a tax on companies that replaced workers with machines. The IRS declined, but there is much to recommend a modest tax. First, it could control activity levels by forcing firms to think carefully about whether the replacement of humans with robots makes fiscal sense. Second, the revenue generated could fund retraining for displaced workers or otherwise provide a cushion for the transition toward an economy with vastly greater automation. To be clear, however, I do not envision a world in which every human retires in favour of robots.

10. Apart from specific questions concerning liability, fundamental legal questions are increasingly being posed. In this respect, the central issue will be how decisions or actions that have been taken by robots or 'artificial intelligence' and caused damage can be contested – potentially by means of legal action. What fields of law are particularly affected by these issues, and what reforms would you propose?

My own study of robotics law and policy suggests three general categories:

- *Embodiment. Robots are software that can touch you. The law will have to strike a different balance now that bones instead of bits are at issue.*
- *Emergence. Robots are increasingly capable of useful but surprising behaviour. Many areas of law rely upon concepts such as foreseeability (tort), sufficient or exclusive control (tort, agency, maritime law), and intent (criminal law, copyright) that the prospect of emergent behaviour complicates.*
- *Social valence. To a greater degree, perhaps, than any previous technology, robots blur the line between person and instrument. We are hardwired to anthropomorphise them. Courts have already had to draw lines between robots and people—an issue that will only amplify with time, with implications for ethics, privacy, liability and other areas.*

11. Against the background of current trends in research and development, please sketch out briefly the direction in which, in your opinion, robotics and artificial intelligence will, or could, develop?

I believe robotics and artificial intelligence will continue to get better at domain specific tasks, and even to surprise us with new capabilities. I also believe we will see greater innovation as we open up robotic platforms to third party developers and to creative users. I doubt any field will remain untouched by artificial intelligence and robotics—these tools are simply too useful. But I do not expect, based on current trends, that the technology will approximate let alone surpass human intelligence in my lifetime.

Thank you for this opportunity to address the Committee. I am honoured and look forward to a great discussion.