


15th Berlin Debate on Science and Science Policy

Funding a revolution. How artificial intelligence might change the research system as we know it.

7 Nov 2020

SUMMARY



The following topics were identified as preliminary anchors for further reflection and debate:

- Specificity of AI: What are the challenges unique to the field?
- Fundamental vs. applied research in AI: Does the differentiation between fundamental research in publicly funded institutions and applications in the private sector still make sense? If not, how do we capture the mutual engagements of these two variants of knowledge generation?
- Brain drain: How can academia create attractive conditions to keep promising young AI scholars that have been trained in universities and not to lose them too early to the private sector?
- Collaboration or competition between public and private institutions active in the field of AI: How can innovative formats of exchanging talent between the public and the private sector look like?
- Training the next generation: What is the best way to do it? Where should it be done? How can we instil the idea that AI has to serve the public good?
- Context matters: Research being done cannot be detached from the context in which it is being done. How does this need to be reflected when wanting to create just and inclusive forms of AI?
- Diversity is key: The field of AI lacks diversity, e.g. in terms of gender, cultural background, or age. How can we contribute to changing this?
- Values and ethics: Who defines the values and norms to be considered when developing AI systems and how do we assure that they are being integrated embedded into technological developments?
- AI ethical guidelines and principles: How can they be made concrete and actionable, acknowledging the diversity of actors and contexts? How can accountability be ensured?

Organised in an unusual year, the year of the coronavirus pandemic, the 15th Berlin Debate took place in an unusual format. For the first time since its launch in 2001, it was held virtually. Since the initial gathering, the goal of the Berlin Debate has been the same: to connect stakeholders from economics, science, politics and civil society to discuss highly topical science and science policy issues. This goal remained the same in 2020: to foster connections between experts and practitioners in science and research and to provide a safe and exclusive intellectual space where views and perspectives can be exchanged freely and where convergent and divergent positions can be understood and discussed.

The 15th Berlin Debate gathered an outstanding and highly complementary group of experts and practitioners from academia, the private sector, policy making and research funding to discuss the growth of Artificial Intelligence (AI) and the relevance of this field as a game changer for others. While there is no shortage of debates on the benefits and related risks of applications of AI in a broad range of contexts, one aspect, however, attracts only little attention, namely the fact that a large part of AI research is done by private companies – and that this has the potential to fundamentally change the relation to academic (i.e. publicly funded) research. With this year's Berlin Debate, the Robert Bosch Foundation invited its guests to jointly and openly explore the consequences of this growing trend and to discuss the role that global AI research and its private and public funding sources and drivers are playing, as well as its implications on issues such as governance, ownership, education, or ethics.

Science has always had multiple funding sources, with industry funding playing a pivotal and indispensable role – in particular in turning scientific findings into viable products and solutions. Private industry funds the majority of R&D conducted around the world. The rapid growth of AI in the last decade, however, is fundamentally changing the balance between industry and public research funding. Therefore, AI has a disruptive potential not only for our economies, political systems and societies, but also for the functioning of the entire research system.

Building on these considerations, the 15th Berlin debate with its title “Funding a revolution. How Artificial Intelligence might change the research system as we know it” posed two fundamental and highly relevant interrelated questions; one addressing money-related interests and push- and pull-factors in the development of AI, the other speaking to the changes in knowledge generation due to AI – as was underlined by the Debate's moderator Ulrike Felt in her opening remarks. In particular, the Debate sought to address the following questions: How will AI change the research system as we know it, with its introduction in ever more domains? How will it change our knowledge ecologies, i.e. the different kinds of knowledge that we are generating and working with to understand the world, to make sense of it and to act in it? What should a fruitful cooperation between private and public research efforts look like? How can the research system and society as a whole benefit from the investments and expertise of private players? How can ethics and democracy be built into AI research?

AI is an ‘enabling’ technology, which means, put simply, that “it makes things possible” – in different fields and in new ways. It is often difficult to see the different ways in which AI is already present today and will be even more so in the near future. Thus, we need to inquire not only how AI might change the research system, e.g. through new kinds of collaboration, novel funding streams, and unorthodox ways of working, but also how it changes what we know, how we know, and what values are prevalent. These wider concerns were reflected in the 15th Berlin Debate.

Science and money in AI: Understanding the nature of the challenge(s)”

Setting the scene for the discussion, the first session “Science and money in AI: Understanding the nature of the challenge(s)” addressed the issue of financial flows. David B. Resnik

from the National Institute of Environmental Health Science shared his insights into the relationship between science and funding sources. In almost all countries around the world, he argued, private funding for research significantly outweighs governmental funding. Private funding streams are, therefore, of central importance for research activities. This has significant implications on how and on what kind of research is being done. Dr Resnik underlined that a high percentage of scientists have financial interests related to their research, as do academic institutions. He stressed the need to acknowledge the economic realities of science. *“We have to come to terms with the economic realities of science. Money is a basic unit for exchanging goods and services. We live in a capitalist society, and there is no way around that. And science follows money. So you really can’t take the money out of science. You can’t even take private money out of science without damaging science, industry, academia, and restricting freedom of speech and economic activity. Even if you tried to have a completely socialised science, there would still be ethical and political problems related to the influence of government money on science.”*

According to Dr Resnik, relevant lessons on how to deal with these tensions and to manage the influence of money can be learned from other fields of research, such as biomedical research, e.g. the implementation of disclosure policies, requirements to share data, research protocols and materials, the need to declare certain kinds of conflicts of interests or education and training in responsible research, to name just a few.

Stefan Heumann from Stiftung Neue Verantwortung then shared his thoughts on the global AI race. AI is the key technology of the digital economy in the 21st century, creating opportunities but also conflicts and competition. Mr Heumann outlined three main areas of competition that shape the field, the first being geopolitical competition, mainly expressed in an increasingly hostile environment between the US and China that other actors have to navigate. The second area consists in a growing competition between the public and private sector and the impact of AI on fundamental rights. *“The struggle between governments and AI industries will shape how the technology is used, and how much it will serve public and private interests.”* The third area identified by Mr Heumann is growing competition regarding the openness of and access to AI technologies. AI is no longer a mainly academic field, since capabilities have shifted from academia to private industry. There is much more demand for AI experts than supply. Industry is increasingly winning this competition, resulting in significant brain drain from academia to industry. *“We see a huge concentration of capabilities, data pipelines, computing resources and human expertise in very few, very large companies. The cutting-edge research has become increasingly privatised.”* These are issues that public policy and science policy will have to grapple with.

The ensuing discussion focussed on the implications of the connection between science and money. While acknowledging that the relation between science and money is an important topic in the field of AI research, some participants argued that the tension between the two is also prevalent in other domains – it is not unique to science, and it is not unique to AI. A focus on biases or the foregrounding of specific values in AI developments was seen as potentially more important to focus on. The discussion on the specific challenges that AI rises and that differentiate it from other fields was a common thread throughout the debate.

Deliberating on the character of the global competition for AI talents, participants acknowledged the existence of a power play to win the war of talents, with the speed in which industrial research leads to products as its driving force. Jack Stilgoe from University College London pointed out that what we need to focus on is not the speed of the AI race but rather the trajectory of research that is taken. *“The trouble is not how AI research gets diverted from a sort of purity – rather the question should be: where is AI going? The big question to me seems to be who benefits: How do we stop AI being a gigantic multiplier for inequality.”*

Another important question debated in this first session related to the commonly accepted separation between fundamental research on AI being done in universities and publicly funded research labs on the one hand and applied research being done by industry on the other. Does the differentiation between fundamental research in public institutions and applications in the private sector still work or even make sense? In fact, researchers in companies produce a very significant share of AI publications, hence also seem to drive fundamental research.

Finally, the relevance of training promising young AI scholars and keeping them in academia resonated a lot with participants. As Jean-Pierre Bourguignon from the European Research Council pointed out, it is difficult for universities to retain promising young AI talent because of competing offers from industry where they find good financial conditions but also access to data, resources that the academic environment cannot compete with. A bit surprisingly, some companies do offer some free time to conduct independent research. Helga Nowotny from the ERA Council Forum Austria (and former President of the European Research Council, ERC) underlined that in a field as popular as AI, universities have an important role to play in installing the conviction that AI has to serve the public good. There is a clear need to reinvent the ways AI and other technologies are being taught and engaged with throughout the training of researchers.

AI and the research system: disruption ante portas?

The following session on the relationship between AI and the research system was introduced by Nuria Oliver from the Data Pop Alliance who presented the European Lab for Learning and Intelligent Systems ELLIS. ELLIS' mission is to create a diverse European network that promotes research excellence and advances breakthroughs in AI, as well as a pan-European PhD programme to educate the next generation of AI researchers. ELLIS also aims to boost economic growth in Europe by leveraging AI technologies.

Dr Oliver pointed out that ELLIS had been created to address a number of the issues that were discussed in the first session, notably the importance to train, nurture and keep promising talents in academia. In particular, she emphasised the importance of training and educating the next generation of AI research talents and the central role of transmitting and instilling values.

Building on Dr Oliver's Europe-focussed presentation, Moustapha Cisse from the Google AI Centre Accra provided a picture of the AI research landscape in Africa. Dr Cisse pointed out that the challenges faced in the Global South, and in particularly in Africa differ significantly from the situation in Europe or the US. No country on the African continent spends 1% or more of its GDP on research. For Dr Cisse, this is very problematic, in particular given the development of critical technologies such as AI. The main challenge faced by African countries is not brain drain. It is also not the competition between academia and industry but rather the challenge of educating AI talent in Africa. *"If we solve this problem, if we successfully tackle that challenge, then we can make sure that the benefits of AI arrive globally, and not only in the traditional centres of expertise."* In order for this to happen, Dr Cisse argued, we need to build centres of expertise outside of Europe and North America. He also emphasized the important role played by non-governmental funding institutions in creating educational programmes and underlined the need for research to happen in both academia and industry.

The following debate focussed very quickly on the relevance of context. Building on Dr Cisse's remarks, Bernhard Schölkopf from the Max Planck Institute for Intelligent Systems pointed out that *"research done cannot be detached from its context in which it is being done. It is very difficult to predict what will come out of it but for sure the influence of the society in which it is being done is going to be significant for the kind of AI that is going to be devel-*

oped.” Context also matters with regard to data and access to data, an area of central importance to the field of AI. It is also mirrored in the importance of diversity, as was highlighted by Dr Oliver. Currently AI as a field lacks diversity, given the fact that central players are being concentrated in few regions in the world. Thus, it is of central importance for the teams developing AI technologies to strive for diversity and to implement diversity strategies.

Referring to Dr Cisse’s comments on the need for the public and private sectors to work together to nurture AI talent, Subra Suresh from Nanyang Technological University Singapore provided an insightful example of fruitful university-industry relations and elaborated on how to foster the retention of faculty through innovative formats of collaboration and exchange with the private sector. This was supported by Carlo Rizzuto from the Central European Research Infrastructure Consortium CERIC-ERIC who argued that sending some of academia’s best talents in a planned way to industry and policy-making will also prove to be very beneficial for the ones who remain in research.

Highlighting dissemination as a key part of research systems, Magdalena Skipper from Nature added another relevant topic to the discussion, asking to what extent AI can disrupt the way we disseminate research and primary research findings. To what extent can AI impact how we evaluate research output, both in terms of funding but also of publications? Dr Skipper also suggested to collectively work towards standards that would allow us to surface rather than censor outputs from both academia and industry. *“I think we should rather know about the research that’s going on than not know it.”*

AI and (research) governance: How to build democracy and ethics into AI?

The third and last session focused on how AI relates to, responds to and is shaped by our visions of democracy and ethics. As could be seen from the previous sessions, this question is particularly challenging since visions of democracy and ethics can be very diverse. While it is easy to agree that democracy and ethics are important elements to consider in the development and application of AI, we might have different understandings of what these notions entail. Effy Vayena from ETH Zurich launched the debate with an introduction to the global landscape of AI ethics guidelines. Drawing on a research project that investigates the broad field of these guidelines, she pointed out that the majority of guidelines come from the Global North and are issued mainly by private sector and public sector institutions. They share a focus on AI as an enabling new technology and lay out a number of broad principles. While many of these principles can be found in the various guidelines, they do not always convey the same meaning. Many guidelines also include checklists aimed at the development and deployment of AI. However, as Professor Vayena pointed out, context and practice orientation appear to be glaringly missing in the development of these documents, hence failing to be instructive in terms of what needs to be done and implemented in practice. She also noted a drive to ‘soft’ law governing AI.

Chinmayi Arun from Yale Law School then talked about AI and its risks for the Global South. To her, what distinguishes AI from other technologies is the degree to which it permeates our daily lives and the ways we behave, having a tremendous impact on how humans navigate the world around them and how they organise societies. To her, the ever-increasing relevance of data points to a new generation, a new form of capitalism. In order to regulate companies that work with this data, it is important for regulators to understand their work and to closely engage with them. However, in particular in the Global South, this is rarely the case. Citizens run the risk of being perceived as data sets – which is why Prof. Arun highlighted the need to understand local contexts and what human flourishing and well-being actually mean.

Participants were highly intrigued by the issue, nature and proliferation of AI ethics guidelines. Again, the question of the specificity of AI was posed since many of the guidelines and the principles they were putting forward are of a very general nature and can be applied to

many different contexts. What is relevant, what needs to be prioritised for AI? And how can these very soft guidelines be translated into concrete and contextualised actions for the people building and deploying AI? Also, the guidelines and principles reflect very different conceptions of ethics, fairness and other values that need to be put into context and spelled out in more detail. Participants argued that it was time to stop the proliferation of AI guidelines and rather move to a culture of reflection and finding a way in which institutions can set parameters and be more specific in terms of context and what it actually means to be transparent and to protect privacy.

How can we ensure that what we build into our AI systems is ethically sound and in line with our democratic values, without becoming all too paternalistic? Raising this question, Dr Cisse noted that the way we build ethics and fairness into machine learning models so far is strictly normative. It is then translated into mathematical terms and used as constraints in some optimisation problem. Even more problematic, according to Dr Cisse, is the fact that there is barely any research on positively enabling (as opposed to pointing to what needs to be avoided) approaches to building ethics into AI: *“In the absence of matters and ways of building learning models and machines that truly take into account that local contextual definition of what is considered to be ethical, how do we go about the law itself, the regulation?”* Jack Stilgoe advised not to concentrate mainly on the technical specificities of AI, but rather to consider the lessons we can learn from previous technologies. Looking at health technologies, for instance, it becomes clear that it takes much more than ethical principles: *“It takes institutions of technology assessment, it takes holding hype to account, it takes real regulatory organisation in order to do these things.”*

Concluding remarks

In her concluding remarks, Ulrike Felt pointed out that the debate was most definitively only a beginning. Throughout the discussions, it became evident that AI redefines the relationships between different kinds of actors within the research system but also beyond, requiring us to question and reconsider many of our established categories, such as basic and applied, industrial actors versus public actors, and how they relate to one another.

“We have opened up sensitive areas and pointed to some of the issues we need to work on, we need to think about and we need to connect to, be it as funding agencies, as institutional leaders, as people who closely relate to industrial actors or as industrial actors looking to collaborate with academia, or as people who try to foster the engagement with different actors. It’s about continued learning from each other.”