



# Communications and Digital Committee

## Corrected oral evidence: Digital regulation

Tuesday 23 November 2021

4.10 pm

[Watch the meeting](#)

Members present: Lord Gilbert of Panteg (The Chair); Baroness Buscombe; Viscount Coville of Culross; Lord Foster of Bath; Lord Griffiths of Burry Port; Lord Lipsey; Lord Stevenson of Balmacara; Baroness Stowell of Beeston; Lord Vaizey of Didcot; The Lord Bishop of Worcester.

Evidence Session No. 5

Heard in Public

Questions 43 - 48

### Witnesses

I: Mira Murati, Senior Vice-President of Research, Product & Partnerships, OpenAI; Tabitha Goldstaub, Co-founder, CogX, and Chair, AI Council.

### USE OF THE TRANSCRIPT

This is a corrected transcript of evidence taken in public and webcast on [www.parliamentlive.tv](http://www.parliamentlive.tv).

## Examination of witnesses

Mira Murati and Tabitha Goldstaub.

Q43 **The Chair:** Welcome to this virtual session. Mira Murati is senior vice-president for research at OpenAI, which is a research lab. We look forward to hearing a bit more about that. Tabitha Goldstaub is co-founder of CogX and chair of the UK Government's AI Council. Thank you both very much indeed for joining us and giving us evidence today. The session will be broadcast online, and a transcript will be taken. In the event that we have a short burst of bells, which are the Division Bells here in Parliament, we will pause momentarily but will not suspend the meeting.

I will start by asking you both to say a brief word of introduction, and then we will go to members of the committee and ask questions. In your introduction, could you just briefly summarise the most significant recent developments in AI—artificial intelligence—and, in soothsaying mode, look to the future and tell us how you see the next 10 or 15 years?

**Tabitha Goldstaub:** Thank you very much for the opportunity to join you all today. I wanted to start by making it clear that I am here to represent the views of the AI ecosystem. I am no expert in regulation, but as the chair of the AI Council I am able to share the current temperature from those at the forefront of this research. You ask about significant developments, and we are lucky that we have Mira here from OpenAI, who no doubt will talk about tech developments.

I thought it would be useful for me to focus on the recent shifts in the impact of AI. I am sure you have all heard the stats before, but it is important to remind ourselves that it is predicted that UK GDP will be 10.3% higher in 2030 as a result of AI. That is the equivalent of an additional £232 billion, making it one of the most commercial opportunities in today's economy, according to PWC and others who are predicting these changes.

Unlike some technology fields, the research and the development are happening hand in hand at breakneck speed. It is corporate research and AI, rather than just academia, that is pushing the envelope. Half of the top papers presented at the leading AI conferences are from corporate research centres—Google, Microsoft and Facebook—as well as the big universities.

This means that there is a move away from the open nature of research—only 15% per cent of papers are publishing their code—which obviously harms accountability and reproducibility in AI. But it also means that wonderful things are happening. Companies are being encouraged to use AI, with incredible results. Fifty per cent of companies surveyed by McKinsey confirmed that they have now adopted AI in at least one business function.

AI technologies are expanding in scale, scope and complexity, resulting in a diverse range of applications with relevance in all areas of our social and economic life. In the next five years we will continue to see this pace

increase. AI experts across the globe disagree hugely on what the next 10 to 15 years, as you ask, look like. Everybody still disagrees on how to solve intelligence, even what intelligent intelligence is, and as such it is almost impossible to predict which methods will be the next breakthroughs and in which areas they will be. That is why it is so important that you are looking at regulation now.

The good news is that we are not starting from scratch. There have also been recent developments in the Government's approach to AI. The AI Council published our AI road map earlier this year, which led to the commitments in national AI strategy being published two months ago by DCMS and BEIS. This included the Office for Artificial Intelligence committing to publishing a White Paper on a pro-innovation national position, on governing and regulating AI, and on commissioning the Alan Turing Institute's public policy programme to produce a study on common regulatory capacity for AI. I know that both are going well, which is why I am so thrilled that I get to be here today to answer some of the questions.

**The Chair:** Thank you. We have many more questions for you. Mira, welcome, and thank you for standing in at short notice. We really appreciate it.

**Mira Murati:** Thank you very much for inviting me to speak with you today to introduce OpenAI and to make clear to the committee the perspective that we are bringing. I will tell you briefly about OpenAI and our work. We are a small AI research and deployment company based in San Francisco, and my focus in the company is very much on the development of the technology as well as its deployment. These are the matters that I think about every day, and many of my thoughts today will come from that lens.

Over the past couple of years, we have announced several major advances in AI. In May 2020, we announced GPT-3—GPT stands for generative pretrained transformer. This is an AI system with astounding linguistic competence that is sometimes indistinguishable from what humans can write. It is very interesting, because it is like a chameleon; you can give it any task, and it will adapt to doing that as long as it is in natural language.

When you think about where we were two or three years ago, with text-based systems, you had to have one system that could do, say, customer service for an insurance company, and another AI system that could do code or translation from one language to another. You had to have all these different systems for any specific parts, whereas with GPT-3 you have just one system that can do it all, and it shows the first signs of this general competence in a specific domain.

In January 2021, we announced DALL-E, which is a version of GPT-3 but, instead of text data, we are now introducing text and image pairs. We see that this system is capable of much more. It can do highly realistic and complex images from only a simple description as input. In August 2021, we went from text to text and code, so we built a system that could code in multiple programming languages.

We developed each of these systems as part of our effort towards building general-purpose artificial intelligence—what we refer to as AGI—and we envisioned this to be highly autonomous systems that outperform humans at most economically valuable tasks. Our mission is to ensure that once we reach AGI we develop and deploy it in ways that benefit all of humanity.

To ensure this, we have also thought about the set-up of our organisation, and we have organised OpenAI as a capped-profit company governed by a non-profit. This means that employees and investors are due limited return and the residual economics flow back to the non-profit. We have commercial efforts that help us to invest in technology and talent while maintaining the checks and balances that are needed to actualise our mission.

We have a dedicated safety team that works to ensure that we are developing and deploying the technologies in a responsible way. This spans from developing novel safety methods and research for aligning our systems with what humans value to developing mechanisms such as content filters for systems that are already in deployment and ensuring that we remove harmful or toxic content.

Our approach to the development and deployment of these technologies is very much controlled and iterative. We want to test this technology in the real world and react quickly, learn from how developers are using it, and deploy it when we are confident and comfortable that it will be used in the right way at scale. We ensure that all our partners abide by strict use-case guidelines and content restrictions in order to work with us. We also make continual updates to the technology when it is in the field. But we know that, even with all the controls and guidelines that we can put in place, we cannot achieve our mission without regulatory involvement, because we cannot possibly expect every lab or company that is in the AI space to share our commitment to safety. We believe that regulation is essential to build and sustain public trust in the technology and to ensure that it is deployed in a safe, fair and transparent way.

**Q44** **Baroness Buscombe:** Mira, where you have just finished is actually the key to all this, is it not? It is how you build public trust in the system. Listening to you both, I feel a huge sense of optimism, but also that I should listen more to those who, like you, can help to explain AI. It is complex, and it is a real challenge for regulators—I am looking at a brief here from our clerk; this has not come from me—to understand how we, for example, audit black-box algorithms whereby it is possible to guess at how the algorithm is working only by observing its effects rather than being able to trace how it operates.

To what extent will AI pose new challenges for regulation? What kind of regulation can we apply that will even begin to keep pace with this fabulous progress—let us call it that—in innovation, but also with the dangers that can come with that? That, of course, is the balance that we think about: how we keep the trust of the public as this progress develops.

**Mira Murati:** There are a few things that regulators can do to keep up with the advancements of these very powerful systems. One is that we are making the systems available through APIs—application programming interfaces. Systems like GPT-3 or Codex are available through APIs.

There are several dozen developers currently building different tools using our API in the UK, and many are using it in ways that we could not have predicted. A company called Elomia, for example, has created an AI companion to help people who are struggling with anxiety or cannot see a therapist. We have teams that monitor how people use the API, and we have policies that will limit harmful behaviour by users and customers, but we are worried that similar tools will proliferate via APIs that are less rigorously managed. We predict that they will have a comparable impact to social media platforms', and perhaps even more so, given the capabilities of these models in the coming decades. Very little attention is being paid to how these systems are being used right now, so this is the time when we can still take action to understand the risks and opportunities of these increasingly powerful AI systems.

More work needs to be done to understand the capabilities and benefits of any of the systems before they become widely available, as well as understanding the trade-offs. We could study the implications of human-like text and image generation, including the potential for misuse for disinformation purposes, as well as what the implications for economic productivity could be. UK agencies could study these issues directly or fund work by academic researchers on them. Studying this AI API is one thing that that would have high impact.

Another thing would be exploring third-party auditing or evaluation of these AI systems. Regulators will not be able to write rules that capture all the risks of these rapidly evolving AI systems, but they can audit them to make sure that they work in the ways developers say they are using them. Given the complexity and the increased capabilities of the systems, we are currently relying on developers self-reporting on whether the systems that they are building are safe, secure and fair, but regulators need to be able to verify these claims. I think that third-party auditing could strike the right balance between that and developers wanting to protect commercial secrets or privacy of data. Third-party auditors could be given privileged and secure access to private information, and they can be tasked with assessing whether safety claims made by AI companies are accurate.

**Baroness Buscombe:** What you are really talking about is ongoing practical input to pre-empt and prevent to the best of our collective ability to prevent harm. Tabitha, as chair of the UK AI Council, is this at the core of what the AI Council is looking at now? This may not necessarily be just regulation, but let us focus on regulation and how you are tackling it and will embrace it.

**Tabitha Goldstaub:** This is one of the most important areas because, as you said, it is about the pace of change. At the moment, private companies are able to deploy AI systems with potential for substantial

harm or misuse in almost unregulated markets. We need to ensure that Governments are able to scrutinise these systems effectively. As we have just heard from Mira, there are some really good ideas out there about how to do that, but ultimately we will need a really special set of skills to do that.

The Alan Turing Institute's new research in this space, commissioned by the Office for Artificial Intelligence although not yet published, is in the process of finding that there really is no common capacity across regulators to be able to do an investigation that has all the things that Mira is talking about—the cognitive skills, the practical abilities, the technical expertise—which are very difficult to get your head around.

It is not just the speed and the skills; there are also three unique areas where AI poses a very different challenge to regulation than we have seen before. There is a sort of grey area of harms. They are raised as technology challenges, but ultimately they are not really technology challenges; they are extensions of non-AI issues, because AI is rapidly developing. They are counterparts of AI systems, the second being the use of AI.

**Baroness Buscombe:** Can you give me an example of that?

**Tabitha Goldstaub:** A good example is if you can see challenges where you have systems where we say, "That was the technology's fault", but really it was the system and the power play around it that made the decisions. Recently, an insurance company called Lemonade decided to use the way you were writing online to see whether you were falsifying a claim. Now, is that an AI problem or is that just an insurance problem? There is a real tension when it comes to how you make these decisions. Also, it might not just be a sector problem, in this case insurance; it is also a business problem, because that could be used in finance. Technically it could be used in any other area—job interviews, for example. That, I think, poses a new challenge that we have not seen before.

**Baroness Buscombe:** In other words, we are spanning everything—in which case, the role that you have is critical to that. Do you feel you have the sufficient tools at the council at the moment to explore a lot of this?

**Tabitha Goldstaub:** I have to be honest. Luckily, the AI Council does not have to resolve these problems. As an independent council, our job is to support the Office for Artificial Intelligence to make good decisions on what the next steps are to be. We are an independent group of 23 members who meet once a quarter. The people who really need to resolve this are the regulators themselves, and they need the support of groups like the Alan Turing Institute and others to come together to think, "How are we going to do this collectively?" The AI Council will be one of those parties, but we are not responsible for navigating this.

Q45 **Baroness Buscombe:** Who is teaching the regulators the questions they need to ask?

**Tabitha Goldstaub:** That is probably the best question. In the last six months we have found that the regulators are doing some incredible work themselves. Let me look up the exact names, otherwise I will get this wrong. Regulators have come together to create a digital regulation co-operation forum, and they are also formalising the existing regulators and AI working group and the Digital Markets Taskforce. We have seen from the Alan Turing research into where the capabilities are that there are pockets, but ultimately we are far from being able to have common capacity across all the regulators to be able to make sure that there is fair distribution of the skills.

So I think the answer, sadly, is no, the regulators do not have the right skill sets, and we will need to bring that together. This is what the Alan Turing Institute's research is looking to do. We should have this published in the new year. It was commissioned by the Office for Artificial Intelligence, so ultimately that should provide us with at least a path to how we will upscale the individual regulators themselves—and not just as individuals; I think the superpower will be once they can come together to see how the AI challenges move across the sectors. The idea that is floating around is thinking about it as a common capacity, a hub for this expertise to come together—not just as regulators, academics, industry, the UK, but all of that together in one place.

Q46 **Lord Foster of Bath:** This has been a fascinating discussion, but it is just worth reflecting on what you, and the AI Council, have said. You said that, “The UK will only feel the full benefits of AI if all parts of society have full confidence in the science and the technologies and in the governance and regulation that enables them”.

Following on from what Baroness Buscombe and you have talked about, which is the fact that it is not the AI Council that has to do it, you said, “Actually, it is the regulators who are responsible”. But, in truth, if the public are to have confidence in those regulators, they have to have confidence in the regulations themselves. I would be interested in your reflections on that and in particular on the work that has been done by quite a number of very eminent AI companies on self-regulation and whether that is a sufficient starter for 10 at the present time.

Just to make life a bit more confused, I have been looking at the Government's plan for digital regulation, in which they are very clear that they do not want heavy-handed regulation, because they fear that it will stifle innovation and so on. Just develop your thinking on that.

**Tabitha Goldstaub:** Thank you. You have definitely chosen my favourite line from the AI road map and saved me reminding you all of it. That confidence is the most important thing that we have to find. People talk about making sure that we have these trustworthy systems in order to have confidence. There are two parts to this. One, as you say, is that we are starting to see that it is not just about the regulators; it is about the regulation. We have seen, since 2019, the OECD, then the G20 Ministers, then the European Commission all coming out with their principles—or, as they are calling them in some cases, advice—at the same time as the

big corporates are coming out with their principles and their advice. I think we all realise that that is not enough.

The 2021 European Commission proposals for a harmonised regulation on AI in Europe were really where the rubber hit the road. Some say that that is not the right sentiment, but it excited me because it made me feel like a group of people had come together and said, "Okay, standards and self-regulation are not enough. We need to move to the EU regulation". Now we are obviously starting to see the challenge with the fact that Europe as a whole is behind China and the US in the development of regulations.

What is exciting, I think, is that the UK has a chance to play a leading role in actually being able to speak with one voice, and I am starting to see other countries really look to the UK for leadership on what the actual regulations should be. The European Commission came up with some good and some not so good ideas. Ultimately, whatever they decide will have impact across the whole world in the same way that GDPR has. We have to be a part of, and think about, how the UK is going to regulate. That is why the Office for Artificial Intelligence is set to publish this White Paper in January/February, and I think the AI Council's job is—and I implore you all as your job—to keep the pace and the speed up on this.

The second thing is slightly different, but it is really important to mention. When we talk about ensuring the general public trust, we have to ensure that they actually have the education to do that. In the same road map we also said that having this confidence will really depend on having a data and AI-literate population. No one understands the rights given to them by GDPR, and they very rarely exercise those rights. In order for people to be able to trust, they have to understand the importance. So we stress hugely that it is about teaching the skills needed not just to grow talent, research and build AI, but to ensure that every child leaves school with the basics of AI and data literacy to be able to become engaged, informed, empowered consumers of the technology.

When I hear you talking, I cannot help but think that that is one of the missing pieces when we think about the next step. If we have government making regulations, and regulators regulating, we have to bring the general public into this process or they will feel no sense of comfort in this new world.

**Mira Murati:** To build on what Tabitha was saying, there is certainly a huge void in regulation in this space, and we can incentivise industry standards. It is critical that we develop some broadly agreed regulator-backed AI safety standards that can be adopted by AI companies as soon as possible. As this technology proliferates and diversifies, it will become increasingly more challenging to rally all, or even most, of the companies around a set of standards.

It would be easier to do it at the stage we are currently at. It will be difficult to come up with reasonable principles, but there are things that we could establish even today, such as industry standards on the

explainability of these systems, their robustness, their safety and reliability, and similar key issues that companies can be evaluated against. These are likely the most realistic option for ensuring safe development and deployment, because the safety of the deployment of these systems comes from the design stage, not just at the end when we are thinking about deploying them.

For example, when we released GPT-3 in May 2020 and launched it in our API in June 2020, we had a lot of restrictions in place because we were not sure that it could be used safely and reliably. It was only last week that we felt we had made enough progress on our safety protocols and systems for detecting bad behaviour and figured out how to catch it and how to mitigate it before opening it up for broad access. Sometimes this means that companies will need to move more slowly, which could mean a loss of competitive advantage. Regulator-backed industry standards can ensure that companies that are trying to ensure maximum safety for their users and society overall are not penalised. The large language models like the one we have built with GPT-3 could be a good place to start.

**Lord Foster of Bath:** I entirely accept that the education of the general public is crucially important. The truth for the vast majority of things we have been used to regulating in the past is that a fair proportion of the general public could understand the issues that we are dealing with. I suspect that in this case not necessarily the datasets but the technology that is used in AI will never be understood by the general public. Therefore, reliance and confidence in the regulatory system is absolutely critical.

The Department of Health and Social Care began to have real concern about something as basic as the datasets used in diagnostic AI technologies that were being developed that did not cover people from ethnic minority backgrounds, for example. So there is real concern even about the data, but there is certainly much more of a lack of understanding and potentially concern about the technology. I have not yet got from you a clear understanding of how, given that you, me and the committee want to have trust, we deal with that and give the public the trust in the developments that will be so big over the next few years.

**Mira Murati:** I think this goes to the standards that the regulators can put in place to incentivise a more coherent set of standards for the entire industry. One way to achieve that would be to set in place a standard on transparency and explainability of these AI systems. That does not mean that you have to peer into the depths of the neural networks and understand exactly what is going on, but the mitigations in place could be through the use of the system itself and understanding the strengths, capabilities, risks and hazards.

It goes back to the point I made earlier about auditing these systems rather than trying to keep up with the emergent capabilities. That could be one standard that could unify what the various companies in the industry are doing. Another could be about the reliability and safety of these systems, as with other technologies like automotives and

aeroplanes for which we have certain standards, and holding companies accountable for deploying the systems safely. Incentivising industry standards would go a long way to building public trust and ensuring safe use.

**Tabitha Goldstaub:** I agree. I often think about a story that Stuart Russell, a professor at Berkeley who is giving the Reith lecture soon, tells about how we do not understand how our car engines work but we trust that the cars will work. It is a really interesting challenge, because although I agree with him I also believe that we need the right leadership in order to make sure that we believe in those things.

Stuart Russell talks about the different types of understanding. You need to be able to learn how to drive a car—you have driving tests in order to do that—but you do not have to understand how to open the bonnet. I do not think that will be enough in the AI world, and I can see from your head shaking that neither do you. I am starting to try to find the experts who are looking more deeply into things like differential privacy and explainability and the sorts of technologies that Mira was talking about, where we can actually say, “No, we’ll have to understand those things”, coupled with something that Rachel Caldicott, who I know was before you recently, said about the imagining that we can encourage the public to do, which is, “What kind of future do they want?” I do not know how to link those two things together, but that is the space that we have to bring closer in order for this to work. I just do not have the answer.

Q47 **Lord Bishop of Worcester:** Lord Foster has already asked the supplementary question I was going to ask, so we can move straight on. Thank you very much indeed for all that you have given. It has been fascinating and really helpful. You have spoken about the international context. I want to go into that a bit more, if I may.

There is great variation in regulation across the world, and I want us to think about the implications of this for the UK. Last week, Rachel Caldicott pointed out that in the UK we have inherited many communications technologies from the US and different approaches to the freedom of speech. The implication was that we are looked to as a leader, because there is some regulation on freedom of speech that would not be possible in the US, but tech giants might take to it if we went for it here.

On the one hand, I suppose, that encourages a slightly tighter regime than might be possible in the US. Sally Sfeir-Tait pointed out to us that the EU is seen to be a bit heavy handed in its regulation, and we would not want to go in that direction because it would stifle innovation. I wonder where both of you see the international situation. Could you expand on that a bit more and on how you see the UK in the midst of that? Mira, could you give us an international perspective first?

**Mira Murati:** The key here is for the UK, the US and the European Governments to work closely with industry to stay up to speed on the emerging issues. Each of these countries will have to have a view on AI. I think there will be a unification across specific matters that span

countries, and perhaps there will be specific views that countries take on issues of privacy or healthcare, for example, which could vary from country to country but not on most matters.

I think there will be some agreement on the general principles of a fair, transparent, beneficial use of AI, and on how to ensure that the systems are reliable and safe and are treating all people safely and fairly. The key here is to work very closely with industry, because the pace of innovation will vastly outstrip the regulators' ability to update the rules.

As I mentioned earlier, GPT-3, just a year away from GPT-2, became vastly more capable than its predecessor. We certainly need regulators in the mix from all over the world to incentivise safety and consider and mitigate any bad impacts on society. There should be the steady flow of information and working closely in partnership between industry and Governments to make sure that we are moving as quickly as possible.

At OpenAI, some of the things that we do with the US Government include providing previews for new technologies that we are building, and working with regulators to red-team some of our tools or make our researchers available to explain what kinds of things we are seeing, some of the problematic things that people might be doing with our tools, and what we are learning in real time. The Government could also boost funding for research on alignment and interpretability, which are broad issues that span borders. This is important and realistic and something that could be done right away.

**Lord Bishop of Worcester:** That seems quite a positive assessment. Tabitha, would you agree?

**Tabitha Goldstaub:** I am hopeful, but I am not as optimistic. Starting with the tensions going between the US, the EU and the Chinese ways of looking at this, there are definitely a lot of similarities where there are clear areas of what some people call the low-hanging fruit—areas that no one thinks is right, such as AI-created deep fakes. There is quite an obvious move towards that not being thought of as acceptable; that AI should always have to label itself as AI.

Then again, we are seeing vast variations on other things. There is the Chinese Government's regulation at the moment, which came out of cyberspace administration, on recommended systems and algorithms using AI. They are looking very, very clearly in their regulation at taking more control of their systems and increasing their control of the information flows in speech, for example. The US system is the opposite. It is having less and less control. The EU is looking much more at regulating the technology and the harms. So everyone has a different point of view.

None of the experts I speak to in the UK feels like anybody has cracked this, which is why it is so important for the UK to take this leadership role. What I am excited about from the UK's perspective is the way we are seeing this as a sectoral approach that in some cases has been working very well. It also has its challenges, and maybe I can talk about

that more in the next question, because this is not about the international.

If we look at our approach compared to other countries', I am hoping that our approach will look at both small and large companies. There are many wonderful things about GDPR, but one of the things that I think we struggled with was ensuring that small companies benefit from GDPR rather than it just being an absolute burden on them and an easy thing for big companies to be able to manage. That is an area where the UK can take a leadership approach in relation to the other countries, which we have not seen from their recommendations so far.

**Q48 The Chair:** Thank you both very much, Before I let our witnesses go, are there any other points that you think the committee should be thinking about in this area?

**Tabitha Goldstaub:** The only one is about this being the critical moment for the UK to decide whether we will continue with the sectoral approach or make tweaks to it. At the moment, there is the challenge of so many societal implications of AI, from the loss of agency, racism, bias, discrimination, privacy and climate harms to widening the digital divide. Every sector, regulation and regulator has gone very narrow in trying to resolve these things themselves because of the pace and the fact that this is obviously cross-sectoral. However, we are seeing some issues in that approach, because there is inconsistency across regulators. There are also overlaps and the fact that that is not the approach that we are starting to see internationally, as I touched on before.

Now, the Office for Artificial Intelligence and the Alan Turing Institute are looking into other ways to do this. There is: the removal of existing regulatory burdens, which is not proving to be very popular; retaining the current sector-based approach, which, as I just explained, is not popular; and introducing cross-sector principles, a set of rules, which is also proving unpopular.

That is why I wanted to raise this point, because I think that the option I gave right at the beginning of a common hub or pool of expertise, a place for collaboration to take place, is the only way we will be able to keep a sectoral approach but learn from and with each other, the wider ecosystem and so on, and really be the leader in this space. If we do not set this up quickly, the vacuum will be filled, because everybody needs to deal with this, whether it is regulators needing to deal with a specific problem or other countries deciding that they need to come up with guidance, principles or regulation. The time is now to create that common capacity.

**Mira Murati:** We are happy to provide early access to capabilities, just as we do with the US Government. If you need help with writing and understanding capabilities of the technology, we are happy to do that too.

I also wanted to emphasise the point about pace of progress and the extent to which these technologies are becoming increasingly more powerful. I talked earlier about where we are today and the capabilities

of the technology today. As we look into the future and where we are going, it is important to understand how we got here. Deep learning has been inspired by the brain. Nobody really knows how the brain works, even today, because biological neurons are so complex, but we can design artificial neurons that are inspired by the brain. We build the neural net out of the artificial neuron, and the key insight here is that we have a mathematically sound learning rule and we can train these neural nets. By contrast, the brain is super-mysterious. This was the first big discovery.

The second discovery has been that you put these large neural networks together with large data, large computers, and get us reliable and astounding AI progress. That is what OpenAI has been leaning on and we will continue to lean on this formula. We do that by picking an architecture, getting the data, deciding on the goal of the system, and then doing this over and over again. That is what we did with GPT-3. That is what we did with DALL-E and Codex.

You can see that if you continue on this trajectory, you can push it further and further and you can imagine very quickly getting to do things like having systems that write programmes just by telling the machine what you want it to do. You do not need any software engineering skills. Looking at it more broadly, if our goal as a society is to get lower-cost healthcare and accessible education to every human being who needs it while living on a planet that is not ruined by climate change, I think AI can help a lot with all of this. We can help to build systems that work with humans in a more scalable way to get preventive healthcare and equal access to good education and more.

**The Chair:** Thank you very much indeed for the evidence today and the offer to follow up on any further inquiries that the committee may have.