

# Replies to Senate Queries

Emeritus Professor Gary Marcus

13 June 2023

**Senator Richard Blumenthal**  
**Chair, Subcommittee on Privacy, Technology, and the Law**  
**Written Question for Gary Marcus**  
**Professor Emeritus, NYU**  
**May 23, 2023**

1. Training data is crucial to foundational models like GPT-4, where content such as news, art, music, and research papers are used to create and refine AI systems, largely material aggregated from the internet. This content represents the labor, livelihoods, and careers of artists, experts, journalists, and scientists.

**How should we make sure AI systems respect, acknowledge, and compensate the labor of individuals whose work is used to train AI models?**

This is, like all copyright law, ultimately a moral choice: how do we wish to structure our society; should artists, writers, and so forth be compensated for our work, such that we have a flourishing society?

Existing copyright laws were written to protect artists, writers, and so forth, but in a time before machines could digest human-produced works at rapid pace. I believe that we should update those laws —particularly because current generative AI techniques are black boxes, in which human-produced works can be assimilated with minimal modification, verging on plagiarism, *without appropriate credit being assigned*.

As a society, we should strive to move to a future generation of software in which systems that synthesize works at large scale from multiple underlying sources would be required to identify their sources and compensate authors for those works. Since this is not yet possible, we should (a) encourage the development of such software, and (b) strike a compromise in the meantime.

My recommendations are as follows

- Work that is under copyright should be used only on an **opt-in** basis.
- Companies should be encouraged to negotiate with artists on a good-faith basis to compensate artists whose work is assimilated by their systems. If you train on it, you buy a license to it.
- In order to encourage the development of more transparent generative software, a sundown period should be set —perhaps three years-- after which point, pure black box systems that cannot ascertain the sources for a given work of art (illustration, writing etc) and compensate their creators appropriately shall no longer have legal access to copyrighted materials.
- Generative AI systems should be transparent about their data sources, such that there is a way for government inspectors to ask whether any given artwork is in a given system's training set.
- The companies producing those generative AI systems should bear the inspection costs, e..g, for random or periodic audits.
- There should be enforcement mechanisms to assure compliance.

# Replies to Senate Queries

Emeritus Professor Gary Marcus

13 June 2023

**Senator Dick Durbin**  
**Chair, Senate Judiciary Committee**  
**Written Questions for Gary Marcus**  
**Professor Emeritus, New York University**  
**May 23, 2023**

1. In February, the Judiciary Committee held a hearing on kids' online safety. During that hearing, witness Emma Lembke testified regarding the toll social media took on her as she grew up. She explained, "As my screen time steadily increased, my mental and physical health suffered."

This is an experience shared by too many kids today. From 2015 to 2021, the time kids spent each day on social media rose to nearly three hours—an increase of almost 60 percent in six years. Over a similar time period, CDC data showed a massive spike in negative mental health outcomes for our kids—particularly teen girls. By 2021, 42 percent of teens reported persistent feelings of sadness or hopelessness, and nearly one in three teenage girls said they had seriously considered suicide.

I am concerned that artificial intelligence could exacerbate this problem.

- a. **In light of our experience with social media, is there a way to safely deploy artificial intelligence so it does not make the current mental health crisis our kids are experiencing even worse?**
  - b. **If so, what specific protections are necessary to minimize the potential harms artificial intelligence may pose to kids?**
2. **What specific guardrails and/or regulations do you support that would allow society to benefit from advances in artificial intelligence while minimizing potential risks?**
3. During the hearing, Mr. Altman testified that "a new framework" is necessary for imposing liability for harms caused by artificial intelligence—separate from Section 230 of the Communications Decency Act.
  - a. **Do you agree with Mr. Altman that a new framework is necessary for imposing liability for harms caused by artificial intelligence?**
  - b. **If so, what features do you consider most important for such a liability framework?**

1. There are many forms of AI that don't pose any particular threats to teenagers (eg GPS routing systems, traditional search engines), but generative AI chatbots can influence people's beliefs, and

# Replies to Senate Queries

Emeritus Professor Gary Marcus

13 June 2023

so there is some risk; being candid, I don't immediately see a way to guarantee their safety. In my Senate testimony I referred to a recent article that describes how these systems can influence our beliefs, and a described a recent system that was made available to millions of children, that told a person posing as a thirteen-year-old, how to lie to her parents about a trip with a 31-year-old man. These are clear warning flags. I don't see how current technology can fully avoid such things, unless one simply restricts access for people under a certain age, which is worth considering. *Future technology that could reason about human values might offer a more trustworthy solution, but is a long way away*. Government funding for AI that could reason about human values would be a good thing, as I have long argued. Generative AI systems are, in my expert opinion, not a good substrate for this for reasoning about human values.

2. There are many guardrails and regulations I would suggest.

- Research funding for AI systems that could reason about human values, rather than simply regurgitating statistics with little comprehension.
- Building a national AI agency that is agile enough to scan current and developing technologies both for opportunities and risks, coordinating the many existing agencies but also finding gaps in existing legal coverage.
- Working towards a global AI agency to coordinate international policy.
- Creating an FDA-like regulatory regime for AI that evaluates large-scale deployment, balancing risks and benefits.
- Creating a system for auditing systems post-deployment, with government support to mandate compliance by companies that create AI systems, focusing on issues such as bias, misinformation, etc medical misinformation, etc
- Requiring the licensing and watermarking of AI systems, such that liability can be traced.
- Mandate far greater transparency into how current systems work and especially on what they are trained on.
- Empower independent scientists to evaluate systems for potential risks and ways to remediate those risks.

3. Yes, I agree with Mr Altman that we need a new liability framework. Fundamentally, our liability laws did not anticipate the current situation, in which AI systems can produce harm at large scale, and with little transparency into their mechanics or use.

I would focus on

- Mandating the **licensing and watermarking** of all deployed generative AI systems, such that harms can be tracked.
- Developing a framework for **making companies responsible for harms that they cause indirectly**, e.g., if their tools are used in the widespread, automatized production of harmful misinformation, or in the production of weaponized armies of bots committing cybercrimes.
- **Empowering a national agency** to act quickly as new risks become realized.

# Replies to Senate Queries

Emeritus Professor Gary Marcus

13 June 2023

**Sen. John Kennedy**

## **Questions for Gary Marcus, Professor Emeritus, New York University**

1. What are the most important factors for Congress to consider when crafting legislation to regulate artificial intelligence?
1. In your judgment, what is the most appropriate and measured approach for reviewing the safety of artificial intelligence models prior to their introduction to the public?
2. If you were to plan a monitoring agency, please describe the structure which you would implement.

1. The fundamental challenge is the rapid speed at which AI is moving, relative to the traditionally slower speed at which government works, coupled with the reality that current systems are opaque (we don't really know how they function, or how to predict with precisions what they do) .There are many guardrails and regulations I would suggest, such as the following.
  - Research funding for AI systems that could reason about human values, rather than simply regurgitating statistics with little comprehension.
  - Building a national AI agency that is agile enough to scan current and developing technologies both for opportunities and risks, coordinating the many existing agencies but also finding gaps in existing legal coverage.
  - Working towards a global AI agency to coordinate international policy.
  - Creating an FDA-like regulatory regime for AI that evaluates large-scale deployment, balancing risks and benefits.
  - Creating a system for auditing systems post-deployment, with government support to mandate compliance by companies that create AI systems, focusing on issues such as bias, misinformation, etc medical misinformation, etc
  - Requiring the licensing and watermarking of AI systems, such that liability can be traced.
  - Mandate far greater transparency into how current systems work and especially on what they are trained on.
  - Empower independent scientists to evaluate systems for potential risks and ways to remediate those risks.
2. The most appropriate and measured approach for large-scale deployments of new AI systems is perhaps something analogous to clinical trials for new medicines. Fundamentally there should be parallels to academic peer review. *When releasing systems for large-scale deployment, companies should be obligated put forward a safety case*, Outside experts must evaluate those systems, and must be able to request additional data, detail etc. Small-scale deployments for research should

# Replies to Senate Queries

Emeritus Professor Gary Marcus

13 June 2023

have minimal requirements (e.g. simple registration), but larger-scale deployments should be subject to significantly greater scrutiny, with examiners able to request more detail, further safeguards, etc. (Although the current emphasis has been primarily on generative AI and chatbots, parallel issues arise e.g., in rollouts of large-scale deployments of driverless cars, large fleets of robots that might interact with the public, etc.)

3. I am not an expert in the exact structure of regulatory agencies, but certainly see parallels to the [FDA](#), and one might see its [org chart](#) as a departure point for discussion, with a Commissioner, legal and operational support, etc, and different Directors perhaps focusing on different forms of AI (eg generative AI vs driverless cars) and/or different potential risks (cybercrime, misuse for market-manipulation, state-sponsored disinformation, and so on).