

Response to questions from Senator Tillis for Mr. Ben Brooks, Head of Public Policy

Artificial Intelligence and Intellectual Property – Copyright

Judiciary Subcommittee on Intellectual Property

August 2023

1. Given generative AI is developing all over the world and countries are responding to it in different ways, are there policies or regulations being adopted elsewhere that you recommend that the U.S. consider or avoid?

The United States can demonstrate global leadership by developing a measured response to AI that realizes the full potential of these technologies while addressing emerging risks. As Congress considers the future of AI oversight, we encourage policymakers to vigorously promote open innovation in AI. These technologies will be the backbone of our digital economy, and it is essential that the public can scrutinize their development. Specifically, we urge policymakers to:

- **Foster innovation in open models.** Open models promote transparency, competition, and security in AI. By inspecting open models, researchers and authorities can ‘look under the hood’ to verify performance, identify risks, and develop new mitigations. By customizing open models, creators and developers can build new AI tools or launch new AI ventures without spending millions on research and computing. By building on open models, organizations can establish their own secure and independent AI capabilities without relying on a handful of firms for critical AI infrastructure. Future policy should account for the role of open models within the AI ecosystem, and support the development and deployment of open models as a public resource.
- **Adopt a risk-based approach to oversight.** Performance requirements for AI systems should be proportional to risk. Requirements may be more stringent where the likelihood or consequence of harm is greater, without imposing those requirements on lower risk systems or lower risk applications. For example, an AI system deployed in a higher-stakes domain such as healthcare, finance, education, or public administration may be subject to more rigorous requirements governing reliability, interpretability, and assurance than an AI system deployed in a lower-stakes domain. In addition, oversight frameworks should be adaptable. They should have the flexibility to respond to new technology and new research. Oversight frameworks should avoid prescribing a single means of compliance that may be ineffective or outdated: safety will depend on layers of mitigations across the AI supply chain.
- **Support diversity in the AI ecosystem.** The AI ecosystem is diverse, from large firms building proprietary technology through to everyday developers freely sharing their technology. Future policy should nurture this diversity. In particular, policymakers should

carefully consider the impact of regulation on grassroots innovation. For example, ‘one size fits all’ requirements that assume every AI developer is a well-funded corporation may have a chilling effect on collaborative research by independent developers, researchers, or entrepreneurs. Developers may be reluctant to contribute to AI research or share AI resources, leaving the community dependent on a small number of firms for access to AI technology.

These principles can help to ensure that AI is deployed in a way that meets public expectations for safety while promoting transparency and competition in AI. Other jurisdictions recognize the importance of these principles. For example, the United Kingdom has committed to a sector-specific regulatory approach that focuses on the risks presented by AI in specific domains.¹ The European Union has recognized the importance of open innovation, and the diversity of AI supply chains.² However, it is yet to be seen whether the final version of the draft AI Act will adopt a risk-based approach to AI models, and we continue to advocate for amendments to the final Act to protect grassroots developers.³ In addition, a range of jurisdictions including Singapore, Japan, and the European Union have reformed their copyright laws to create safe harbors for AI training that achieve similar effects to fair use.⁴ These reforms aim to establish a culture of ‘open learning’ that promotes access to large and diverse datasets, helping to make AI safer, more effective, and less biased.

2. A recent survey on how consumers view AI found that most consumers – nearly 80% – believe the use of AI should be explicitly disclosed. Do you agree? Why or why not?

Users should know when they are interacting with AI applications that generate content or make significant decisions.⁵ To that end, policymakers may consider disclosure obligations for application developers who provide certain AI services to a user. Additionally, policymakers may consider privacy obligations that require users to affirmatively consent prior to the collection of their data for AI training, since this data may be used to target the user with personalized content.

3. What are the benefits and disadvantages of requiring an AI company to keep records of everything that is ingested and to make those records publicly available? Under what circumstances, if any, should an AI company NOT be required to make its records of everything that is ingested by the AI publicly available? Under what circumstances, if

¹ British Government, *A Pro-innovation Approach to AI Regulation*, March 2023 at [3.2].

² European Parliament, draft AI Act, May 2023, Recitals 12a, 12b, 12c. See also draft Commission and Council text.

³ European Parliament, draft AI Act, May 2023, Articles 2, 3, and 28b.

⁴ See the respective text and data mining exceptions. See also the *Regulation of Technologies Review* conducted by the UK Government Chief Scientific Adviser in March 2023 at 9: “If the government’s aim is to promote an innovative AI industry in the UK, it should enable mining of available data, text, and images (the input) and utilise [sic] existing protections of copyright and IP law on the output of AI.”

⁵ Stability AI, ‘Advocating for Open Models in AI Oversight’, May 2023. See our submissions to the National Telecommunications and Information Administration and the Office of Science and Technology Policy.

any, should an AI company be required to make its records of everything that is ingested by the AI publicly available?

Datasets should be assessed for safety, bias, and suitability before they are used to train an AI model. Independent auditing is one way to help promote scrutiny of these datasets. By inspecting datasets, researchers, developers, and users can anticipate potential risks or limitations in a model. For example, by inspecting image datasets, researchers can determine whether an image model fairly represents a particular community or demographic. By inspecting language datasets, researchers can determine whether a language model accounts for cultural, political, or language diversity. Developers can work to correct these limitations and adjust undesirable behavior, either by training a new model or by ‘fine-tuning’ an existing model with additional data.

Open datasets may not be feasible in all cases. For example, models may be customized or ‘fine-tuned’ with proprietary data, and a developer may not be able or willing to share that data publicly. The content and structure of a dataset affects the performance of an AI model, and we acknowledge that developers may choose to keep their datasets closed. Nonetheless, model developers should ensure that datasets are screened for unsafe content, bias, and opt-out requests from creators. In sensitive contexts – such as healthcare, finance, or public administration – additional dataset requirements may be determined by regulators.

4. Do you think that generative AI prompts provided by users are copyrightable? And if so, under what circumstances could they be copyrightable? Do you think that whether the prompt used is copyrightable or not should impact the copyrightability of the resulting AI output generated as a result of the provided prompt?

Whether a prompt is copyrightable is a fact-sensitive inquiry. Certain prompts may consist of unique and highly detailed instructions, with original expressive content, created by a user after significant experimentation. These prompts may satisfy the criteria for copyrightability in certain circumstances. In our view, the copyrightability and registrability of an AI output is a separate inquiry (see response to Question 12 below).

5. What does the impact of generative AI have on the creative industry? Specifically, what are your thoughts regarding the concern that the proliferation of generation AI will take over jobs?

We believe AI is a tool that can help creators express themselves, but it is not a substitute for creators. Instead, AI can help to accelerate the creative process. AI tools can help existing creators boost their productivity, experiment with new concepts, and perform complex tasks as part of a wider workflow. In addition, AI can lower barriers to entry for people who do not have the resources or training to realize their creative potential. In this way, AI can help open up

economic opportunities for new and existing creators, much like earlier tools such as computers, software, smartphones, and social media.

However, there is a risk that the economic benefits of AI accrue disproportionately to AI firms. That's why Stability AI is committed to putting AI in the hands of everyday creators and developers through open models. Open models enable everyday creators to participate in this new industrial revolution as builders – not just consumers – of technology. Creators can adapt, integrate, and experiment with open models to develop AI applications that best support their work (see examples from our written testimony). They can develop these AI tools without spending millions on research or computing power. In this way, we expect the economic benefits of open models will be shared by a broader community of creators, developers, and entrepreneurs, not just Silicon Valley.

In the long term, we are sensitive to the implications of AI for routine tasks and the effects of AI on certain kinds of work. To that end, we encourage policymakers to carefully scrutinize the impact of AI across the national economy. Over time, it may be necessary to reflect on how we learn, train, and work in light of these tools.

6. If a generative AI system is found to infringe a copyrighted work, who should be liable for the infringement – the AI company, the user providing the prompts to the AI tool, or both?

Stability AI is unable to comment on matters subject to litigation.

7. In your opinion – currently or in the foreseeable future – can AI generated material ever replace the quality of human created work?

AI is a tool, not a replacement for creators. AI can help to accelerate the creative process, but AI depends on a human in the loop to provide creative direction. AI may deliver comparable performance to humans in the technical execution of specific creative tasks (e.g. extending, filling in, or editing a photorealistic image). However, audiences value a range of attributes in art beyond technical execution. They take into account the composition as a whole, the context of the work, the creator, and subjective interpretations of meaning. By itself, AI cannot impart these attributes to content. We expect the community will continue to value human works – indeed, we may value them at a premium. Smartphones didn't destroy photography, and word processors didn't diminish literature, despite radically transforming the economics of creation. Instead, they created new fields of artistic expression, giving rise to new demand for services, new markets for content, and new creators.

8. A balance needs to be struck in terms of how to encourage innovation, how to be responsible, and how to ensure that there is clarity for all using this technology. How do

you propose we do this in the copyright space in a way that allows the U.S. to stay competitive and remain the global leader?

We believe that training AI models is an acceptable, transformative, and socially-beneficial use of existing content that is protected by the fair use doctrine and furthers the objectives of copyright law, including the Constitutional goal of ‘promot[ing] the progress of science and useful arts’.⁶ These models learn the unprotectable and non-expressive ideas, facts, and structures within a visual or textual system. ‘Open learning’ of these facts about our world is essential to recent developments in AI, and it is doubtful that these groundbreaking technologies would be possible without it. The US has established global leadership in AI due, in part, to a robust, adaptable, and principles-based fair use doctrine that balances creative rights with open innovation. We encourage policymakers to preserve this culture of ‘open learning’ in any future reform.

However, there are specific harms that may warrant future intervention. For example, the improper use of likeness can be problematic if it wrongfully implies a person’s endorsement of, affiliation with, or promotion of a work or idea. The improper use of personal likeness should be governed by clear rules that specify impermissible use. These may include right of publicity laws (for performance or voice likeness) and deepfake laws that prohibit specific kinds of misleading conduct in political, commercial, or personal settings (including non-consensual intimate imagery). Policymakers may consider strengthening these laws. Further, policymakers can help to fortify the wider information ecosystem by encouraging the adoption of content provenance standards by AI applications and advanced content moderation systems by social media or streaming platforms.

9. In the copyright context, what differentiates the technology of generative AI from other machine-aided creativity, such as photography, video cameras, electronic music, and the like, all of which allow the public to develop and advance knowledge?

AI models review existing content to learn the relationship between words, ideas, and fundamental visual or textual features. They can apply this knowledge to help a user generate new and unseen content as part of a creative workflow. In this sense, AI is comparable to other assistive technologies, from cameras to software. Like other tools, AI may accelerate the creative process; increase the volume of creative outputs; and change the economics of creation. Like other tools, AI may be used properly or improperly by a user. However, these characteristics do not relevantly distinguish AI from other technologies that alter the ease, speed, volume, or cost of performing a creative task.

10. What steps can and should the creative community take today to ensure that their work is more easily attributed to them, regardless of whether their work is used for training an AI model? For example, indicating authorship and contact information via the metadata of the author’s digital content.

⁶ U.S. Constitution, Article I, Section 8, Clause 8.

Machine-readable metadata can help to communicate essential information about a particular work. For example, Stability AI is exploring machine-readable metadata that indicates whether a creator would like to be removed from datasets used for AI training. This metadata would be additional to existing protocols, such as robots.txt, that indicate whether a web page can be ‘crawled’ for purposes such as search engine indexing. Today, machine-readable opt-outs are required in order for a rightsholder to exercise certain opt-out rights under the European Union text and data mining exception.⁷

11. Are existing laws and regulations sufficient to deal with the issues relating to transparency and record keeping by AI companies?

Existing frameworks are sufficient to address emerging challenges in AI, including statutory authorities that apply across sectors (e.g. the *Federal Trade Commission Act*), domain-specific regulatory authorities (e.g. FDA, DOT, SEC, and DOL), and existing product liability principles. The risk of an AI system is determined by how that AI system is deployed for a specific application in a specific environment. Transparency, auditability, and documentation requirements will vary accordingly, and domain-specific rules will best account for these variations.

Existing regulatory agencies are equipped with the legal tools to respond to these AI deployments, from AI decision making (e.g. automated vehicles) to AI content (e.g. healthcare interactions with a patient or business interactions with a consumer). However, regulatory agencies may not be adequately resourced to respond to these challenges. Prior to further legislative intervention, we encourage Congress to conduct a rapid gap analysis to determine (i) where regulatory competence falls short in any sector or domain and (ii) where agency funding and personnel may need to be supplemented.

12. Have you reviewed the U.S. Copyright Office’s Registration Guidance for “Works Containing Material Generated by Artificial Intelligence” and, if so, what are your views on the guidance? Do you think that the Copyright Office got it right? Are there aspects of the guidance that could stand to be clarified or revised?

In principle, we acknowledge a threshold of authorship below which a work with negligible human input may not qualify for registration. However, we are concerned that the U.S. Copyright Office’s Registration Guidance and recent decisions (such as *Kashtanova*) do not account for the many ways in which human input may rise above the threshold. For example, a creator (i) may fine-tune their own model on their own content to better evoke their established style, (ii) provide detailed creative direction via prompts, inputs, and settings that narrowly define the range of possible outcomes, or (iii) use the model’s suggestion as a starting point, and refine the initial image via editing.

⁷ European Union, Directive 2019/790 (Copyright Directive), 2019, Article 4.

Sufficient creative control demonstrated in one or more of these ways should qualify a work for registration. A user with clear expressive intent, and who has demonstrated that they directed the model, should be able to register their work. We welcome further clarification on this issue. Uncertain or discretionary guidance means that professional creators may be unfairly disadvantaged by their use of AI tools within a wider creative workflow.

13. Both the U.S. Patent and Trademark Office and the U.S. Copyright Office have engaged in extensive outreach regarding AI. Have you participated in this outreach and, if so, how did you find it? What more can and should these offices do?

Stability AI participated in the U.S. Copyright Office listening session on AI and visual arts in May 2023. These sessions were comprehensive and professional, featuring a diverse range of viewpoints. Notably, these sessions were structured around different modalities to reflect important nuances in the copyright treatment of different media. While we object to aspects of the recent guidance on AI (see response to Question 12 above), we welcome further engagement with the U.S. Copyright Office in relation to AI, including registrability.

14. Your CEO has publicly suggested that, in the future, Stability AI and other AI models will be fully licensed. While hearing this is promising, what more can be done now – across the AI community – to address the current concerns amongst the copyright community?

There are practical steps that AI developers can take today to address emerging concerns. These include respecting opt-outs in training data; watermarking and tagging AI-assisted content to help users and platforms distinguish it from human-generated content; and mitigating defective behavior such as overfitting. These are detailed in our written testimony. In addition, we are actively exploring new ways to incentivize creators in the digital economy. These are early days, but we look forward to sharing more information with the Subcommittee as these initiatives unfold.

15. Some AI developers have said that the ingestion of copyrighted works is transformative and qualifies as fair use. What impact does the Supreme Court's recent decision in *Andy Warhol Foundation v. Goldsmith* have on that position?

For the reasons given in our testimony, we believe training is an acceptable, transformative, and socially-beneficial use of content that is supported by the fair use doctrine.

16. Scraping the Internet for data – text, images, audio, video, etc. – for use in training AI models has all the current focus. However, once this has been done the focus may shift to sources of data that are not as readily accessible, such as private user data. Do you foresee companies using cloud-based file storage systems – such as Microsoft OneDrive, Google Drive, Dropbox etc. – as a potential source of data to be scraped? What are your thoughts on this?

A ‘base’ AI model is typically trained on broad datasets, including content aggregated from the Internet. A base model will demonstrate some level of competence in basic tasks. However, to make AI models useful, they will need to be customized or ‘fine-tuned’ for specific tasks using high-quality data. That may include proprietary or confidential data. This data should be accessed and handled appropriately, and we urge greater transparency in how private data is used in the development of AI models. Further, user data should be treated with care, since it may be used to train AI models that target a particular individual with personalized content. We have previously advocated for robust privacy obligations that require affirmative user consent before an AI application is able to collect or retrieve personal data for AI training.

One way to ensure that data is properly accessed and handled is by customizing AI models ‘in house’. For example, a regulated financial institution may need to customize an AI model to assist in analysis, decision making, or customer support. The financial institution will need to audit the performance of the model for reliability; train or fine-tune the model without exposing sensitive customer data to third-parties; and retain full control over the AI model without relying on a third-party provider.

Open models can help organizations to customize these AI models safely and securely. By building on open models, organizations can ‘look under the hood’ to identify potential risks or limitations. They can use the open model to train a custom model for specialized applications without exposing their proprietary data to a third-party AI firm. After customizing the model, they retain full control over their AI capabilities, including the distinctive settings or ‘parameters’ that define the model’s performance, without relying on a third-party AI firm. In this way, organizations can help to ensure the security of their proprietary or sensitive data.

17. Where does Stability AI get the training data that is used to train their AI systems?

Stable Diffusion is trained on a subset of two billion images from the LAION-5B dataset, a public dataset of five billion image URLs and text captions assembled by the nonprofit Large AI Open Network. We filter this dataset to remove unsafe images, among other criteria. By filtering training data before it ever reaches the AI model, we can help to prevent users from generating unsafe content. In addition, we filter this data for opt-out requests from creators, and we will incorporate these opt-out requests in future training.

18. What steps, if any, does Stability AI take to ensure that your training data does not include copyrighted material?

Stability AI has proactively solicited opt-out requests from creators, and will honor these over 160 million opt-out requests in upcoming training. Going forward, we are exploring new technical standards for machine-readable opt-outs, so that opt-out metadata follows the content wherever it goes. In addition, the datasets used for training our flagship model, Stable Diffusion, respected

industry-standard digital protocols like robots.txt, which indicate whether a website consents to automated data collection for ancillary purposes such as indexing or analysis.

19. Has Stability AI sought out licenses from copyright owners for copyrighted content that is included in your training data? Please explain why or why not.

We believe that training AI models is an acceptable, transformative, and socially-beneficial use of existing content that is protected by the fair use doctrine and furthers the objectives of copyright law, including the Constitutional goal of “promot[ing] the progress of science and useful arts”.⁸ Through training, these models develop an understanding of the relationship between words, concepts, and fundamental visual or textual features. The model doesn’t rely on any single work in the training data, but instead learns by observing recurring patterns over vast datasets – billions of image and caption pairs, and hundreds of billions or trillions of words. The model does not store the material in this training data. They do not “collage” or “stitch” together original works, nor do they operate as a “search engine” for existing content. Instead, AI models apply this generalizable knowledge to help a user create new and unseen content.

However, we are sensitive to emerging concerns. As detailed in our testimony, we are working to address these concerns through technology, standards, and good practices. In addition, we are actively exploring new ways to fairly incentivize creators in the digital economy, and we would be pleased to share more details with the Subcommittee as these initiatives unfold.

20. What considerations would impact your decision to use copyrighted material to train your model as opposed to non-copyrighted material?

We believe that training AI models is an acceptable, transformative, and socially-beneficial use of existing content that is protected by the fair use doctrine. Access to large and diverse datasets is essential to make AI safer, more effective, and less biased. However, as detailed in our testimony, we acknowledge emerging concerns among creators. We are working to address these concerns through technology, standards, and good practices.

21. Let’s assume that under *Andy Warhol Foundation v. Goldsmith* the use of copyrighted works for training AI is not considered transformative. Do you believe the use of these works would still qualify as fair use looking at the four factors? Which particular factors support your position?

We are unable to comment on matters subject to litigation. As detailed in our testimony, we believe training is an acceptable, transformative, and socially-beneficial use of content that is supported by the fair use doctrine.

⁸ U.S. Constitution, Article I, Section 8, Clause 8.

22. One concern about generative AI that has been raised by creators is that unauthorized copies of their works are being made during the process of collecting data and training a respective model. Could you please speak to how copies and how many copies of such data are made and when within the lifecycle of creating and executing an AI system – from start to end?

Image models are trained on vast datasets consisting of digital images and their associated text captions. During training, these images are reviewed by the AI model to understand the relationship between words, ideas, and fundamental visual features. These images may be reviewed a number of times as the model is optimized to improve performance and mitigate risks, but they are not retained in the trained model. The model can subsequently apply this knowledge to help a user generate new works.