MARK R. WARNER
VIRGINIA

United States Senate
WASHINGTON, DC 20510–4606

COMMITTEES:
FINANCE
BANKING, HOUSING, AND URBAN AFFAIRS
BUDGET
INTELLIGENCE
RULES AND ADMINISTRATION

August 8, 2023

Mr. Sundar Pichai
Chief Executive Officer
Alphabet Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043

Dear Mr. Pichai,

I write to express my concern regarding reports that Google began providing Med-PaLM 2 to hospitals to test early this year.[1] While artificial intelligence (AI) undoubtedly holds tremendous potential to improve patient care and health outcomes, I worry that premature deployment of unproven technology could lead to the erosion of trust in our medical professionals and institutions, the exacerbation of existing racial disparities in health outcomes, and an increased risk of diagnostic and care-delivery errors.

Over the past year, large technology companies, including Google, have been rushing to develop and deploy AI models and capture market share as the technology has received increased attention following OpenAI's launch of ChatGPT.[2] Numerous media outlets have reported that companies like Google and Microsoft have been willing to take bigger risks and release more nascent technology in an effort to gain a first mover advantage.[3] In 2019, I raised concerns that Google was skirting health privacy laws through secretive partnerships with leading hospital systems, under which it trained diagnostic models on sensitive health data without patients' knowledge or consent.[4] This race to establish market share is readily apparent and especially concerning in the health care industry, given the life-and-death consequences of mistakes in the clinical setting, declines of trust in health care institutions in recent years, and the sensitivity of health information. One need look no further than AI pioneer Joseph Weizenbaum's experiments involving chatbots in psychotherapy to see how users can put premature faith in even basic AI solutions.

According to Google, Med-PaLM 2 can answer medical questions, summarize documents, and organize health data. While AI models have previously been used in medical settings, the use of

[1] Kruppa, Miles, and Nidhi Subbaraman. "In Battle with Microsoft, Google Bets on Medical Ai Program to Crack Healthcare Industry." *Wall Street Journal*, July 8, 2023. https://www.wsj.com/articles/in-battle-with-microsoft-google-bets-on-medical-ai-program-to-crack-healthcare-industry-bb7c2db8.
[2] Roose, Kevin. "How ChatGPT Kicked Off an A.I. Arms Race." *The New York Times*, February 3, 2023. https://www.nytimes.com/2023/02/03/technology/chatgpt-openai-artificial-intelligence.html.
[3] Grant , Nico, and Karen Weise. "In A.I. Race, Microsoft and Google Choose Speed Over Caution." *The New York Times*, April 7, 2023. https://www.nytimes.com/2023/04/07/technology/ai-chatbots-google-microsoft.html.
[4] Needleman, Sarah E., and Rob Copeland. "U.S. Lawmakers Scold Google's 'Project Nightingale' Over Health-Data Privacy." *Wall Street Journal*, November 12, 2019. https://www.wsj.com/articles/senators-urge-scrutiny-of-health-data-deals-including-google-project-11573597883.

generative AI tools presents complex new questions and risks. According to the Wall Street Journal, a senior research director at Google who worked on Med-PaLM 2 said, “I don’t feel that this kind of technology is yet at a place where I would want it in my family’s healthcare journey.”[5] Indeed, Google’s own research, released in May, showed that Med-PaLM 2’s answers contained more inaccurate or irrelevant information than answers provided by physicians.[6] It is clear more work is needed to improve this technology as well as to ensure the health care community develops appropriate standards governing the deployment and use of AI.

Given these serious concerns and the fact that VHC Health, based in Arlington, Virginia, is a member of the Mayo Clinic Care Network, I request that you provide answers to the following questions.

1. Researchers have found large language models to display a phenomenon described as “sycophany,” wherein the model generates responses that confirm or cater to a user’s (tacit or explicit) preferred answers[7], which could produce risks of misdiagnosis in the medical context. Have you tested Med-PaLM 2 for this failure mode?
2. Large language models frequently demonstrate the tendency to memorize contents of their training data, which can risk patient privacy in the context of models trained on sensitive health information. How has Google evaluated Med-PaLM 2 for this risk and what steps has Google taken to mitigate inadvertent privacy leaks of sensitive health information?
3. What documentation did Google provide hospitals, such as Mayo Clinic, about Med-PaLM 2? Did it share model or system cards, datasheets, data-statements, and/or test and evaluation results?
4. Google’s own research acknowledges that its clinical models reflect scientific knowledge only as of the time the model is trained, necessitating “continual learning.” What is the frequency with which Google fully or partially re-trains Med-PaLM 2? Does Google ensure that licensees use only the most up-to-date model version?
5. Google has not publicly provided documentation on Med-PaLM 2, including refraining from disclosing the contents of the model’s training data. Does Med-PaLM 2’s training corpus include protected health information?
6. Does Google ensure that patients are informed when Med-PaLM 2, or other AI models offered or licensed by, are used in their care by health care licensees? If so, how is the disclosure presented? Is it part of a longer disclosure or more clearly presented?
7. Do patients have the option to opt-out of having AI used to facilitate their care? If so, how is this option communicated to patients?

[5] Kruppa, Miles, and Nidhi Subbaraman. “In Battle With Microsoft, Google Bets on Medical AI Program to Crack Healthcare Industry.” *Wall Street Journal*, July 8, 2023. https://www.wsj.com/articles/in-battle-with-microsoft-google-bets-on-medical-ai-program-to-crack-healthcare-industry-bb7c2db8.

[6] Singhal, Karan, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, et al. “Towards Expert-Level Medical Question Answering with Large Language Models.” arXiv, May 16, 2023. https://arxiv.org/pdf/2305.09617.pdf.

[7] Perez, Ethan, Sam Ringer, Kamile Lukosiute, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, et al. “Discovering Language Model Behaviors with Model-Written Evaluations.” In Findings of the Association for Computational Linguistics: ACL 2023, 13387–434. Toronto, Canada: Association for Computational Linguistics, 2023. https://aclanthology.org/2023.findings-acl.847.

8. Does Google retain prompt information from health care licensees, including protected health information contained therein? Please list each purpose Google has for retaining that information.
9. What license terms exist in any product license to use Med-PaLM 2 to protect patients, ensure ethical guardrails, and prevent misuse or inappropriate use of Med-PaLM 2? How does Google ensure compliance with those terms in the post-deployment context?
10. How many hospitals is Med-PaLM 2 currently being used at? Please provide a list of all hospitals and health care systems Google has licensed or otherwise shared Med-Palm 2 with.
11. Does Google use protected health information from hospitals using Med-PaLM 2 to retrain or finetune Med-PaLM 2 or any other models? If so, does Google require that hospitals inform patients that their protected health information may be used in this manner?
12. In Google’s own research publication announcing Med-PaLM 2, researchers cautioned about the need to adopt “guardrails to mitigate against over-reliance on the output of a medical assistant.”[8] What guardrails has Google adopted to mitigate over-reliance on the output of Med-PaLM 2 as well as when it particularly should and should not be used? What guardrails has Google incorporated through product license terms to prevent over-reliance on the output?

Sincerely,

Mark R Warner

Mark R. Warner
U.S. Senator

[8] Singhal, Karan, Shekoofeh Azizi, Tao Tu, S. Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, et al. “Large Language Models Encode Clinical Knowledge.” Nature 620, no. 7972 (August 2023): 172–80. https://doi.org/10.1038/s41586-023-06291-2.