PRINCETON UNIVERSITY

CENTER FOR INFORMATION TECHNOLOGY POLICY
CITP.PRINCETON.EDU

To: the offices of Senator Heinrich and Senator Rounds
ASAP Request for Information
Via email: ASAP@heinrich.senate.gov and ASAP@rounds.senate.gov

June 30, 2025

Response to ASAP Request for Information

Thank you for the opportunity to respond to the American Science Acceleration Project (ASAP) Request for Information. We are academic researchers at Princeton University and write to share our perspective on the relationship between AI adoption and scientific progress.

We commend the initiative's goals of accelerating American technical innovation and agree with several key aspects of ASAP, particularly the emphasis on developing metrics to measure scientific innovation (Question 4) and the recognition that we need to understand bottlenecks in our innovation ecosystem. We also appreciate the focus on keeping human scientists at the center of the discovery process, which aligns with our concerns about maintaining human expertise (Question 3).

However, based on emerging evidence from metascience research, we believe that further incentivizing AI adoption in science may be counterproductive. In this response, we argue that: (1) AI adoption in science is already proceeding rapidly, (2) increased production enabled by AI might not translate to increased progress, and (3) rapid AI adoption may actually be slowing scientific progress through many specific mechanisms, which we discuss below.

1. AI adoption in science is already rapid

Researchers already face overwhelming professional incentives to adopt AI tools, making additional government incentives unnecessary and potentially distortive. Recent empirical evidence demonstrates the strength of existing adoption drivers:

Career advancement incentives are powerful. A 2024 study examining AI adoption across scientific fields found that researchers who adopt AI publish 67% more papers, receive three times as many citations, and become team leaders 4 years earlier than their peers. These dramatic career benefits create intense pressure for individual researchers to adopt AI tools regardless of whether they advance scientific understanding.

Adoption rates reflect these incentives. AI adoption across 20 scientific fields has quadrupled in the last decade. This mirrors the trend in the AI field itself: conference submissions in AI-related fields have increased tenfold since 2015. This rapid uptake occurred without specific programs to incentivize adoption.

Market forces are already driving adoption. The private sector has heavily invested in making AI tools accessible to researchers. Major cloud providers offer free or subsidized compute for academic research. Open-source communities have democratized access to state-of-the-art models. Scientific publishers and funding agencies increasingly favor research that incorporates AI methods.

The bottleneck is not adoption speed but adoption quality. Our research on reproducibility in machine learning-based science found that the rush to adopt AI has led to widespread methodological errors. In our past work, we collected flaws in over 600 papers across 30 scientific fields. In many cases, the affected papers constitute the majority of the surveyed papers, raising the possibility that in many fields, the majority of AI-enabled research is flawed. The problem is not that researchers lack access to AI tools or incentives to use them, but that they lack the expertise and institutional support to use them rigorously.

Most proposals to incentivize AI adoption don't address the actual bottleneck. Given these realities, we submit that further government efforts to accelerate AI adoption would be pushing on a string. Researchers who can productively use AI are already doing so. Those who aren't using AI either work in domains where it's genuinely not applicable or lack the expertise to use it properly—neither situation would be improved by additional adoption incentives.

2. Increased production has not translated to scientific progress

A growing body of metascience research reveals a troubling paradox: while scientific production has soared, most available indicators suggest that genuine scientific progress has slowed. This finding challenges the assumption that accelerating research output through AI will accelerate scientific breakthroughs. As we detail in the next section, several mechanisms specific to AI adoption may be exacerbating this production-progress gap.

Production metrics show explosive growth. The number of scientific papers published annually has grown exponentially. The number of researchers has similarly expanded, as has research funding in real terms. AI adoption has intensified this trend by making it easier to generate papers, run analyses, and produce results.

Progress metrics tell a different story. Multiple independent analyses have found evidence of slowing scientific progress despite increased production:

- Park et al. analyzed millions of papers and patents using the "consolidation-disruption index" as a measure of progress. They found that disruptive discoveries have declined as a fraction of total scientific output.
- Analysis of Nobel Prizes shows that recent awards increasingly recognize discoveries from decades past rather than recent breakthroughs.
- Studies tracking the introduction of novel concepts and terminology in scientific literature have found stagnation since the early 2000s.
- Surveys of scientists across fields indicate they view recent advances as comparable in importance to those from earlier decades, despite vastly more resources.

This evidence suggests that accelerating the *production* of research will not accelerate innovation. Simply enabling researchers to work faster or produce more outputs does not address the institutional and epistemological factors that determine whether research leads to genuine progress. While ASAP correctly identifies the need for metrics to measure innovation (Question 4), we urge that these metrics focus on genuine progress rather than production speed.

This divergence between production and progress began before widespread AI adoption, driven by academic incentive systems that reward publication quantity. However, by making it dramatically easier to produce papers without necessarily improving our ability to make breakthroughs, AI could widen this gap.

3. How rapid AI adoption could further slow scientific progress

There are several specific mechanisms through which accelerated AI adoption could actively impede scientific advancement:

Proliferation of errors and erosion of research quality

The accessibility of AI tools has led to their adoption by researchers who might lack the expertise to use them properly, resulting in a reproducibility crisis in AI-based science. Our systematic review identified over 600 papers across 30 fields with fundamental errors in their use of machine learning, including using test data during training, which is like teaching to the test — it leads to an exaggerated accuracy estimate.

Others have found that AI tools are often used with inappropriate baseline comparisons, making it incorrectly seem like they outperform older methods. These errors are not just theoretical: they affect the potential real-world deployment of AI too. For example, Roberts et al. found that of 400+ papers using AI for COVID-19 diagnosis, *none* produced clinically useful tools due to methodological flaws.

These errors persist and compound over time. The scientific community's self-correction mechanisms have proven inadequate for catching AI-related errors, perhaps because scientists are still getting accustomed to software-driven research. When foundational research contains errors, subsequent work building on it wastes resources and leads research down unproductive paths. Meanwhile, quality control mechanisms are overwhelmed—peer reviewers often lack the specialized knowledge to evaluate AI methods outside their domain, researchers rarely review the code and data released alongside papers, and the sheer volume of papers has stressed review systems to the breaking point.

Systematic narrowing of research focus

AI adoption steers research toward questions amenable to current AI methods rather than the questions most important for scientific progress. Hao et al. analyzed six scientific fields and found that research that used AI shifted toward data-rich problems at the expense of research that generated new questions or theoretical advances. Fields showed decreased diversity in research topics and methods post-AI adoption.

This narrowing is reinforced by the "benchmark culture" in AI. Within AI research, the focus on achieving state-of-the-art performance on standard benchmarks has created a dynamic where marginal improvements receive more attention than novel approaches. As other fields adopt AI, they could import this culture, evaluating research by performance metrics rather than conceptual advances. The opportunity costs compound when funding and attention flow to AI-amenable research, while other crucial areas suffer. Young researchers, seeing the career benefits of AI adoption described in Section 1, could abandon research programs that might lead to breakthroughs but don't fit the AI paradigm.

Substituting prediction for understanding

AI excels at pattern recognition and prediction but often provides little insight into underlying mechanisms, potentially hindering theoretical progress. In drug discovery, AI models can predict molecular properties but offer limited insight into why molecules behave as they do. In social sciences, models may predict outcomes without illuminating causal mechanisms, limiting our ability to design effective interventions.

This creates a false sense of progress: high predictive accuracy can create an illusion of understanding. Researchers may stop seeking theoretical explanations when AI provides accurate predictions, similar to how celestial mechanics could predict planetary motion accurately for centuries using incorrect geocentric models, which depicted the Earth at the center of the Universe. The danger is that we optimize our research efforts around improving predictions within flawed frameworks rather than developing better fundamental understanding.

4. Recommendations

Based on this analysis, we recommend that ASAP:

1. **Redirect focus from adoption to quality.** Rather than incentivizing faster adoption, support the development of best practices, training programs, and quality control mechanisms for AI use in science.
2. **Invest in measuring genuine progress.** While we appreciate ASAP's emphasis on developing metrics (Question 4), we urge that these metrics should capture progress, not just publication count or the speed of research. This should include support for metascience research to understand which AI applications genuinely accelerate discovery versus merely accelerating production.
3. **Support research in AI-resistant domains.** Deliberately fund research in areas where current AI methods don't apply, ensuring we don't abandon important questions simply because they're not amenable to automation.
4. **Bolster scientific reproducibility and self-correction.** Fund systematic reviews and reproducibility studies to catch errors before they propagate through the literature. Support the development of AI-based tools for detecting and preventing common methodological flaws.
5. **Reform incentive structures.** We agree with the RFI's recognition (in the Process questions) that funding models and peer review need modernization. However, we emphasize these reforms should reward quality over quantity, genuine progress over incremental improvements, and careful methodology over rapid publication.

The evidence suggests that the bottleneck in scientific progress is not the speed of AI adoption but rather our ability to maintain research quality, pursue diverse research questions, and develop genuine understanding. We urge ASAP to focus on these deeper challenges rather than simply accelerating the adoption of AI in science.

We appreciate the opportunity to provide these comments and would welcome further discussion of these issues.

Respectfully submitted,

Sayash Kapoor
PhD Candidate, Department of Computer Science
Princeton University

Arvind Narayanan
Professor of Computer Science
Princeton University

Contact: sayashk@princeton.edu, arvindn@princeton.edu