July 27, 2023

NSF RFI: 88 FR 26345
Organization: The Center for Security and Emerging Technology (CSET)
Respondent type: Academic institution / Think tank
Primary respondent: Catherine Aiken PhD, Director of Data Science and Research

The Center for Security and Emerging Technology (CSET) at Georgetown University offers the following comments in response to NSF's Request for Information on *Developing a Roadmap for the Directorate for Technology, Innovation, and Partnerships*. A policy research organization within Georgetown University, CSET provides decision-makers with data-driven analysis on the security implications of emerging technologies, focusing on artificial intelligence, advanced computing, and biotechnology. We appreciate the opportunity to offer these comments.

Our response pertains to five of the topics from the enumerated list provided in the RFI:

1. Prioritization
2. Suitability
3. Workforce
4. Additions
5. Other topics

Summary Comments

- The proliferation of technology lists poses a challenge to strategic investment and development. We recommend efforts to (i) consolidate and align existing lists and (ii) consider alternative approaches to identifying critical areas of science and technology.
- AI/ML is a technology well-suited for critical use-inspired and translational research, which can and should be prioritized.
- AI/ML is a technology that already has notable workforce needs, and there are ways TIP can foster workforce development and upskilling in AI/ML.
- Semiconductors and biotechnology are two other areas with use-inspired research areas and workforce needs that should be prioritized.

1. Prioritization

At CSET, we consider a range of metrics when trying to assess technological competitiveness, some of which are listed in Table 1. Which metrics we use, or what additional metrics we explore, depends on the analytic question at hand. For example, examining applications of artificial intelligence to enable novel military capabilities requires different metrics than analyzing how large language models can spur disinformation. Table 1 provides a starting point for evidence that exists to help NSF determine priorities among the listed technologies. We illustrate use of these metrics with some examples from CSET research, which are linked in Table 1. See also our best practices for studying tech competition through research output.

Table 1. Technological Competitiveness Concepts and Metrics

Concept	Metric(s)	Example Uses for Prioritization
Research production	Publication output	What research areas are growing? Who is contributing to these areas? (e.g., Comparing the United States' and China's Leading Roles in Science, AI Research Funding Portfolios and Extreme Growth)
Research focus	Publication content	What topics is research focused on? Who is contributing? (e.g., Trends in AI Research for the Visual Surveillance of Populations, Research Almanac, The Inigo Montoya Problem for Trustworthy AI)
Research translation/transfer	Citations, publication content	Where is research in one field being applied in another field / in the market? (e.g., Exploring trends in AI and genetics)
Research collaboration	Co-authorship	What countries are collaborating in a research area? (e.g., Research Impact, Research Output,

		and the Role of International Collaboration, Headline or Trend Line)
Research impact	Citations, media coverage, software implementation	Who is producing impactful research? (e.g., Comparing U.S. and Chinese Contributions to High-Impact AI Research, ETO Map of Science)
New technology or capability	Patent filings, company financial filings, publications, media coverage	What novel applications or technologies are resulting from research? What are recent advances or breakthroughs? (e.g., Small Data's Big AI Potential, Trends in Robotics Patents, Country Activity Tracker)
Industry interest	Affiliated publication output, Patent filings	What research is happening in industry? How are companies translating research to products? (e.g., Comparing Corporate and University Publication Activity in AI/ML)
Military interest	Affiliated publication output, Research funding	What research is being done or funded by military institutions? What tech is the military interested in? (e.g., Revisiting China's Security Forces' AI Research Output)
Talent and skills development	Job postings, Employee/researcher profiles, education statistics, degree conferrals	What countries have a sufficient tech workforce? Where is top talent working? (e.g., U.S. AI Workforce, Voices of Innovation)
Research & Development	Grants, Budgets, Labs	What research and tech are organizations investing in? (e.g., Mapping U.S. Multinationals' Global

		AI R&D Activity, U.S. Military Investments in Autonomy and AI)
Software Development	Open-source software	How is research being implemented in software? (e.g., GitHub Data: Capturing Open Source Software)
National Policies	Policy documents	How are countries developing and deploying technologies? (e.g., Agile Alliances)
Supply Chains and Manufacturing Capacity	Company financial filings, trade publications, media coverage	What countries have an advantage in a tech supply chain? Where are critical tech suppliers located? (e.g., Supply Chain Explorer, Betting the House)

Having these metrics "on the shelf" is important, but must be combined with input from subject matter experts. Additional evidence could be systematically collected from subject matter experts, whether through surveys, interviews, or literature reviews.

Beyond metrics to consider for prioritization, we offer the following recommendations for research areas well-suited for use-inspired research:

(1) Artificial intelligence, machine learning, autonomy, and related advances.

- AI "safety" research - CSET work classifying English-language research literature using an AI safety definition found that only 2% of AI research is focused on AI safety, and such research has a disproportionately large share of industry affiliated authors. U.S.-based authors are contributing to roughly 43% of research in this space. This finding suggests a gap in this area of research that the U.S. is poised to fill, and one where academic voices could be amplified and public-private collaboration could be fruitful.
- AI assurance and implementing AI safeguards - there is a gap in our understanding of the feasibility and assurance prospects of increasingly general-purpose systems in open-ended domains (including, but not limited to, large language or "foundation" models). This includes research in assured autonomy, ML robustness, and ML interpretability. Current

efforts focus on the assurability of small, specialized ML-based components trained to perform a specific function in a specific context. While there is a need for more early research in this area (see CSET's Foundational Research Grant call for proposals), TIP could support use-inspired research within a 3-year timeframe to address the challenge of planning for and implementing assurance techniques for general-purpose ML systems.

- Impacts on human behavior and outcomes - can include ways human behavior conditions the adoption/operation of AI/ML and how people interact with AI/ML (see here and here).
- Autonomous cyber defense capabilities - recent CSET work outlines the current limitations and potential related to this type of AI application (see here and here), as well as research demonstrating that China is actively working to develop these capabilities.

(2) High performance computing, semiconductors, and advanced computer hardware and software.

The global semiconductor industry (including academia, government, and the private sector) is pursuing three lines of effort as Moore's Law's comes to an end (see here, pg. 72):

- More Efficient Architectures and Packaging (Timeframe: Present Day - 10+ years): Systems on chip (SoC), Advanced Packaging, Photonics, Heterogeneous Integration, and Chiplets
- New Models of Computation (Timeframe: Present Day onwards): Domain specific, asynchronous, approximate, neuromorphic, analogous, and quantum computing are all being explored, with no prohibitive favorite yet.
- New Materials and Devices (10-20+ year time horizon): Spintronics, carbon nanotubes, graphene, and superconductors are all being explored as alternatives to CMOS-based integrated circuit manufacturing.

All of these technologies are seeing substantial investment by industry and, in some cases, by the U.S. government. Below are three examples, one from each line of effort, of use-inspired research topics that should be prioritized for investment in a 1- to 3-year time frame:

- **Advanced Packaging**: As the physical limits of transistor scaling are reached, the semiconductor industry is actively searching for ways to squeeze more performance out of chips. Advanced packaging (which encompasses concepts like chiplets and heterogeneous integration, among other things) is one area of particular interest due to its high technology readiness level, relatively low cost, and the promise of up to 20% performance improvements compared with conventional/traditional packaging techniques. CSET has published work describing the challenges and opportunities of advanced packaging (Re-Shoring Advanced Semiconductor Packaging).
- **AI for Microelectronics Design and Fabrication**: AI/ML is increasingly adopted by the semiconductor industry. Examples include chip design, chip fabrication, and printed circuit

board manufacturing. In each instance, the use of AI/ML generates cost savings for firms and performance improvements for end users. Increasing and emphasizing adoption of AI/ML in semiconductor front and back end manufacturing may confer asymmetric advantages for U.S. firms in the short to medium term.

- **Superconducting Electronics**: Unlike traditional semiconductor materials, superconductors have zero electrical resistance. This allows processors to move bits without any dissipation of energy, resulting in much lower energy consumption than standard electronics. For now, many technical hurdles remain before SCE achieves commercial viability, and the field as a whole remains relatively small, with around 100 papers published globally per year. But several countries, particularly the United States and Japan, have recognized the promise of superconductor electronics and have funded research in this space for many years. CSET has published work describing challenges and opportunities of superconducting electronics (Superconductor Electronics Research: National Competitiveness and Funding Activity).

(5) Natural and anthropogenic disaster prevention or mitigation.

In the context of infectious diseases like a pandemic, we find that the U.S. is not prepared for the next pandemic along the entire pipeline, from basic R&D to therapy approval and manufacturing. Increased attention on epidemiology, basic research, clinical research, and biomanufacturing is important.

(6) Advanced communications technology and immersive technology.

Industry is investing considerable resources into this, but prior CSET research has highlighted the potential pitfalls of failing to monitor and provide adequate government support for critical infrastructure, specifically telecommunications. There are use-inspired research areas that may not be a focus of industry and thus would benefit from government funding, one such case is use of this technology in operational helmets for professions that habitually wear helmets.

(7) Biotechnology, medical technology, genomics, and synthetic biology.

Taking an expansive view of the term biotechnology, we recommend TIP focus on:

1. Interdisciplinary training. The next generation of biomedical technologies will incorporate AI, so having a solid, deep understanding of the biology and of AI/ML will be important and, at least in the near-term, require interdisciplinary training and collaboration, and;
2. Manufacturing. The advances coming out of biotechnology are less useful if the United States cannot manufacture them. A recent CSET report focused on biomanufacturing

finds gaps in vaccine biomanufacturing in the United States. Biomanufacturing relies on a specialized skill set, and there is growing evidence that the U.S. workforce may not be able to meet demand (see here).

(8) Data storage, data management, distributed ledger technologies, and cybersecurity, including biometrics.

There are areas for use-inspired research in data management, storage, and governance. Based on participation in two separate working groups on the topic, as well as forthcoming research drawing on new survey data, we are aware of specific gaps in organizations' knowledge of and confidence in managing data in a secure, ethical, and legal manner and adhering to best practices for data governance for AI adoption and deployment. Use-inspired research could, among other things, examine how researchers and organizations can implement existing frameworks and to what effect (see related work here).

2. Suitability

One technology well suited for use-inspired and translational research is AI/ML. Sometimes referred to as AI "convergence" (see here), AI is already impacting research and innovation in a range of sectors, fields, and topics. Early research that CSET is involved in finds that the use of AI is widespread throughout the sciences and most disciplines benefit from AI. Yet there is evidence of a misalignment between the teaching of AI skills and its impact on scientific research, where the supply of AI talent in scientific disciplines is not commensurate with demand for its use. This forthcoming research finds that the incorporation of AI in research poses growing knowledge demands on individual scientists, and women and under-represented minorities scientists benefit substantially less from AI advances. This poses several opportunities for TIP:

1. Invest in workforce and AI/ML upskilling, particularly for researchers,
2. Invest in curriculum redesign efforts to teach more AI skills,
3. Facilitate cross-department collaborations with AI experts, and;
4. Invest in female and under-represented groups to pursue study in AI-related fields

Additional evidence of the suitability of AI/ML for greater use-inspired and translation reach is early research which suggests AI inventions (measured via patents) and AI adoption increases firm productivity, output, and employment - but that this adoption is geographically concentrated. This poses an additional opportunity for TIP to encourage more dispersed AI adoption and invest in research and education on responsible AI deployment.

Investments should target talent, without neglecting the triad of algorithms, compute, and data. CSET survey research examined AI researchers' preferences for federal funding to support AI research and found most report a desire for research grants to further their work, while many respondents also desired greater compute and data resource provision. One way research funding could foster the talent required to do translational research in AI/ML is by encouraging, or requiring, interdisciplinary research teams, designs, and outputs (see here).

In addition to research grants, and provision of other research resources like compute and data, investment and funding approaches to foster talent development and upskilling will go a long way toward maturing AI/ML. We outline some specific recommendations on this below.

Two other listed technologies that are well-suited for investment in use-inspired are:

- High performance computing, semiconductors, and advanced computer hardware and software, specifically in the areas noted in the previous section - advanced packaging, AI for microelectronics design and fabrication, and superconducting electronics.
- Biotechnology, medical technology, genomics, and synthetic biology, specifically in terms of biomanufacturing, including vaccine biomanufacturing.

3. Workforce

While our research cannot speak to *which* listed technologies will have the greatest workforce needs in the next 1 to 5 years, it can speak to the need for greater investment in AI/ML (see Training Tomorrow's AI Workforce, U.S. AI Workforce) and semiconductors/hardware workforce development (see The Chipmakers, Reshoring Chipmaking Capacity Requires High-Skilled Foreign Talent). As noted in the previous section, our preliminary research also suggests a workforce gap and need for greater skills development in biomanufacturing.

One area where TIP could help advance AI/ML and a broader tech enabled workforce is the prioritization of K-12 teacher recruitment, retention, and professional development. When considering education at the K-12 level, foundational STEM training provides students with the general preparatory knowledge to pursue post-secondary and/or workforce pathways in any of the technologies listed above. Despite the importance of access to high-quality K-12 STEM education, there are ongoing shortages of K-12 STEM (especially in the computer science field, including cybersecurity). K-12 educational initiatives are only as strong as the educators prepared to support and enact them.

TIP could collaborate with other government and private organizations to provide state, local, and tribal governments with model curriculums that provide the knowledge, skills, and abilities needed to build pathways to prepare future workers and reskill current workers for entry into the key technology focus areas by:

- Supporting the development of an AI/ML skills framework, similar to the NICE framework for cybersecurity, and provide guidance for implementing the framework at the state, local, and tribal government levels. With no standardized way to map competencies/skills to AI-related roles/occupations, it is difficult for all stakeholders to speak the same language when talking about training and upskilling talent, and for employers to convey their needs.
- Ensuring regional workforce stakeholders (e.g., employers, state and local governments, nonprofits, and educators) collaborate when developing such programs.
- Empowering and investing in community colleges, which would expand reach to underserved communities.
- Adding new career pathways for programs in emerging technology fields, like the CyberCorps program, to expand beyond the federal workforce. For example, recipients could fulfill their workforce obligations through state, local, tribal employment.

To build successful AI products, organizations need to hire for technical AI occupations (e.g. machine learning engineers), in addition to product development and commercial AI occupations. Compared to technical AI occupations, a smaller share of product and commercial occupations are filled by individuals with STEM degrees, but these occupations are essential for diffusing the benefits of AI development across the U.S. economy. TIP is uniquely positioned to identify and retain these individuals, and to do so, the Directorate could:

- Work with the Department of Homeland Security to add non-traditional innovation-relevant degree programs to the Department's STEM Designated Degree Program List.
- Provide employment resources to recipients of the TIP Directorate's undergraduate scholarships, graduate fellowships and traineeships, or postdoctoral awards. These resources, which would benefit US residents and non-residents alike, should include employment counseling and connect awardees with potential employers—including but not limited to semi-finalists and finalists of the NSF Engines program.

We note that while expanding domestic workforce development is essential, such programs have long time horizons and might have minimal impact in the near-term. To meet immediate demand, the U.S. will also need to rely on high-skill non-resident individuals.

5. Additions

While we do not have recommendations of additional technologies to add to the current list, we recommend that TIP/NSF:

- Consolidate and align lists (see more below),
- Specify the technologies currently on the list (e.g., AI, biotechnology), and;
- Develop an alternative approach to identifying and tracking critical and emerging technologies, ideally one that starts with a set of criteria (see more below).

7. Other topics

We offer two recommendations as TIP moves forward developing a roadmap for investing in use-inspired and translational research.

1. Reconsider the use of a set list as the basis for goal-setting and prioritization for advancing U.S. tech competitiveness and addressing related societal, national, and geostrategic challenges. Instead focus on (i) strategic goal articulation, (ii) establishing criteria for key technologies, and (iii) aligning lists and monitoring efforts.
2. Build monitoring and analytic capacity to monitor S&T developments, specifically along the established criteria for key technologies.

To elaborate on the first recommendation, the proliferation of key technology lists and variation in relevant criteria presents numerous challenges. Forthcoming CSET research cites one count of the number of emerging technologies lists published by U.S. government, industry, academia, journalists, and non-profits in recent years at 78. This includes the National Science and Technology Council's Critical and Emerging Technologies, the Department of Defense's list of 14 critical technology areas, and the U.S. Department of Commerce's "emerging" and "foundational" technologies. There is a downside to agencies having their own criteria for what technology is key, critical, or important. A balance is required between defining organization or mission-specific criteria and alignment with related efforts within, and outside, government. The risks of different efforts being guided by different lists are non-trivial, especially when efforts are meant to work toward common goals at the national level. Additionally, the effort required to compile and maintain them over time is considerable and divergence across resulting lists introduces new challenges.

We recommend list alignment with articulated national goals. Without this, it is difficult to connect listed technologies to broader goals and determine what areas of research would contribute to progress toward goals. With goals in place, we can determine the criteria that must be met for a technology to be considered priority, key, or critical (accounting for mission-specific criteria). We can then have a more dynamic, fluid way to determine where investment is needed, without the need to annually update a static list of technologies as a starting point. This criteria-based approach may still result in a list, or small set of lists, of technologies, but one that is more data-informed, accessible for updating, and reflective of shared and specific goals.[1]

To elaborate on the second recommendation, we continue to encourage investment in national monitoring and analytic capabilities for S&T development and situational awareness (see here). This would require investing in data & analytic infrastructure that would enable TIP to answer questions like those in this RFI and monitor trends over time. It would also enable dynamic monitoring of established criteria - possibly leveraging the metrics listed in Table 1 - to more efficiently identify areas for investment and areas where prior investments are paying off. Monitoring trends to identify areas that meet established criteria or maintain a list of key technology areas requires robust analytic capabilities, and the existing gap in these capabilities within government is a limiting factor. In addition to building and supporting the infrastructure, talent, or data needed for such capabilities, TIP could help develop novel measures and analytic tools through use-inspired research.

[1] It could be argued that working from a set of criteria instead of an expert-curated static list enables a more data-driven approach to monitoring tech development. While this is true, we need to keep in mind that metrics are inherently limited, often to what can be readily counted or what has been used before. This is where the preliminary step of goal articulation is critical. We might still be able to identify in the data the “best” predictor(s) of development among a limited set, but we may be predicting or explaining an outcome that is not actually the one we aim to produce. For example, predicting technology competition, we might say leadership is observed as research output, but the best predictors in that case may not align with criteria to be “key” technology.