

CYBER SECURITY & INFO TECHNOLOGIES

The USSC Channels the Strategic Computing Program

JANUARY 13, 2025

Mark Twain is often quoted saying (perhaps apocryphally) that “history does not repeat itself, but it rhymes.” Whoever spoke these words first, the U.S.-China Economic and Security Review Commission (USSC) is determined to prove them right.

The USSC released its 2024 [annual report](#) in November with 32 recommendations, 10 of which it singles out for “particular significance” (p. 27). The very first recommendation reads as follows:

1. Congress establish and fund a Manhattan Project-like program dedicated to racing to and acquiring an Artificial General Intelligence (AGI) capability. AGI is generally defined as systems that are as good as or better than human capabilities across all cognitive domains and would usurp the sharpest human minds at every task. Among the specific actions the Commission recommends for Congress:
 - Provide broad multilayer contracting authority to the executive branch and associated funding for leading artificial intelligence, cloud, and data center companies and others to advance the stated policy at a pace and scale consistent with the goal of U.S. AGI leadership; and
 - Direct the U.S. secretary of defense to provide a Defense Priorities and Allocations System “DX Rating” to items in the artificial intelligence ecosystem to ensure this project receives national priority (p. 27).

The USSC is recommending that the US Congress establish a program that authorizes the Executive Branch to inaugurate public-private partnerships in the service of AGI advancement backed by sufficient funding. The report does not recommend [specific investment strategies](#). It also urges Congress to direct the US Secretary of Defense to give AI items within the defense

contract ecosystem a DX Rating—this would make the programs housing the AI items “[the highest national priority](#).”

There is value in providing a sufficiently broad mandate to conduct research and development with the interdisciplinarity required of AGI, and a Manhattan Project-like program for AGI points in this direction.

However, the historical analogy between the race to build the atomic bomb and the race to build AGI oddly neglects a more appropriate comparison: the U.S. Defense Advanced Research Projects Agency’s (DARPA) 1983-1993 Strategic Computing Program. Strategic Computing was a major, multi-layered initiative that pursued, in part, the construction of a “generic” (application-independent) system during the peak of the twentieth century’s AI enthusiasm.

Strategic Computing failed to achieve its highest objectives; it was too ambitious. It serves as a cautionary tale for US AI policy today.

The [technical report](#) detailing Strategic Computing’s plans laid out the challenge it sought to address: the increasing computerization of defense operations brought relative advantages in capabilities, but these computers, “having inflexible program logic, are limited in their ability to adapt to unanticipated enemy behavior in the field...We are now challenged to produce adaptive, intelligent systems having capabilities far greater than current computers.” The *opportunity* lies in advances in “artificial intelligence, computer science, and microelectronics” and the “mechanization of the practical knowledge and the reasoning methods of human experts in many fields” through Expert Systems (p. i) (Expert Systems were the peak of “[good-old-fashioned AI](#),” or Symbolic AI).

These technologies would support at least three new defense [applications](#): *autonomous vehicles* capable of sensing, planning, reasoning, and communicating for the Army; a *battle management system* to analyze uncertain data, produce and evaluate options, and explain rationales for the Navy; and a *pilot’s associate* to aid pilots in the face of overwhelming data for the Air Force (p. 17, Figure 4.2).

The kicker was Strategic Computing’s higher and interdependent goal to develop “generic software systems that will be substantially independent of particular applications” (p. 30). An Expert System with a sufficiently large underlying knowledge base, and an architecture

permitting operations over incomplete and faulty data, “will be substantially generic in nature so that it will significantly advance expert systems capabilities and support a wide-range of applications for both the Government and industry” (p. 41).

[Emma Salisbury](#) wrote (in 2020) that Strategic Computing fell short of its highest goals. Progress, where made, was disappointing relative to the Program’s ambition. But *why* it failed is notable: Strategic Computing’s architects—including Information Processing Techniques Office Director Robert Kahn and DARPA Director Robert Cooper—promised the U.S. Congress a research agenda that seeks advances in *separate* research areas, on *fixed* deadlines, and structured in such a way that if one research track were to fall short, the *entire* program would suffer.

To be sure, the USSC is not calling for a plan of this interdependency, instead abstaining from specific investment strategy recommendations and deferring to Congress. There is wisdom in this deference.

But it was Congress that authorized Strategic Computing forty years ago. And back then, it was Japan, with its [Fifth Generation Computer Systems](#) program, that embodied America’s chief [technological rival](#). A historical rhyme, if you will.

Caution in the interpretation of the USSC’s recommendation is therefore historically justified. It is feasible, given the advisory function of the body, to *reinterpret* its recommendation. This reinterpretation could be faithful to its spirit without committing to its letter by moving away from the language of AGI and total cognitive supremacy over humans *without* sacrificing specific actions like Executive branch contracting authority and a designation to assign a DX Rating to AI items in the defense contract ecosystem.

The immediate aim should be to [complement](#), rather than merely replicate or defer to, the techniques employed within private-sector AI research and development.

History’s rhyme provides guidance. A continuity between Strategic Computing and frontier AI research today is a failure to build systems of sufficient safety and reliability, particularly for deployment in mission-critical domains. The *unreliability* of Expert Systems in uncertain conditions was a criticism of Strategic Computing, in particular, given the “[notoriously unpredictable](#)” behaviors of humans in crisis conditions for which the system could not “be

fully tested in advance; nor can crisis conditions ever be fully simulated” (p. 14).

Indeed, some AI researchers and analysts are so concerned with raw capabilities—understood loosely as the sheer *ability* to reproduce human-like outputs—that they forget that safety guarantees like reliability of performance, back-ups and fail-safes, and auditability of outcomes are as much a part of the toolkit of an intelligent system as mathematical prowess.

The conflation of safety and intelligence may seem odd. But when humans deliberately construct systems that offer performance guarantees and regulate technological deployments to serve highly predictable and controllable ends (particularly in critical domains), they are exhibiting their own political and technical competencies. Humans rightly expect their machines to work. Even lacking the personal know-how to achieve this, individuals understand the value of controllable, explainable, human-serving devices.

Contrary to [sweeping](#) and [exorbitant](#) claims about the results of recent benchmark tests with OpenAI’s “o3” model ([tempering](#)—but [not dismissing](#)—a [previous](#) claim about the [ARC-AGI](#) benchmark), deep learning models often fail to offer *performance guarantees*; that is, they fail to provide the “[five nines \[99.999%\] of correctness](#)” in human-critical situations.

Repeated warnings about the [unreliability](#) of systems built atop [deep learning](#) models nevertheless fall on deaf ears. “Intelligence is capabilities,” the thinking goes. Yet, one such capability is the ability to say, “[I don’t know](#)” (when one genuinely does not know) and to *anticipate* one’s own shortcomings and limitations ahead of action, particularly when another’s interest is at stake. Generative models—from GPT-3.5 to o3—lack these capabilities. Congress should be wary of accounts of AI’s trajectory that depend on diminishing human performance to justify the unreliability of state-of-the-art models, in this way avoiding the fate of Strategic Computing in a new context.

The U.S. Congress, should it take up the USSC’s recommendation, should re-interpret it to avoid the pitfalls of Strategic Computing. Concretely, this will mean granting Executive branch contracting authority for the explicit support of the development of novel AI architectures that *robustly* support applications in the real-world, beyond the huffing-and-puffing of benchmark evaluations (which are directionally significant, though limited). This work will likely extend [beyond](#) machine learning, into areas including [Neuro-Symbolic AI](#).

Such research ventures could see the development of AI systems that are more *specialized* rather than the general models of commercial and USSC intrigue. Yet, without sacrificing general AI research, the successful development and deployment of specialized AI systems that offer performance guarantees and transparent operations is consistent with a [recommended revision](#) to US AI strategy vis-à-vis China. Capabilities are of diminished strategic interest if they cannot be managed, echoing the related “[capability-reliability gap](#).”

History does rhyme, and the USSC’s laser-like focus on AGI rhymes with Strategic Computing’s pursuit of a generic system and Congressional fear of Japanese technological rivalry. Yet, what appears obvious today without the benefit of history—that the U.S. government should go “all-out” to create AGI—is considerably less obvious with stock taken in the aftermath of the twentieth century’s peak AI moment. Ambition should be tempered, with U.S. government action grounded in neglected areas of AI research.

America can do better than AI systems that lack [principled constraints](#) on their outputs. Congress should grant the Executive branch the authority to invest accordingly.

Orion Policy Institute (OPI) is an independent, non-profit, tax-exempt think tank focusing on a broad range of issues at the local, national, and global levels. OPI does not take institutional policy positions. Accordingly, all views, positions, and conclusions represented herein should be understood to be solely those of the author(s) and do not necessarily reflect the views of OPI.