THE IMPENDING TECHNOLOGICAL UPHEAVAL
APPROACHES TO POLICY IN TRANSFORMATIVE TIMES

MAY 2018

Bottom Line Up Front

Emerging technologies such as Artificial Intelligence, Machine Learning, Autonomous Systems, and Multi-agent Swarms present both large opportunity and risk across the government. Policymakers must set the appropriate conditions to leverage opportunities and minimize the associated risks.

In earlier eras, those with access to vast resources had a distinct competitive advantage because of their ability to produce expensive, high-impact systems. Today, thanks to the plummeting costs of computing power, custom manufacturing, and access to specialized knowledge, those with fewer resources have the opportunity to gain advantage by virtue of their ability to innovate using cheap componentry, setting the stage for upheaval in both the political and industrial landscapes.

In order for the U.S. to keep pace and ideally lead in this transforming and uncertain environment, a two-layered strategy to system development and integration must be adopted.

- **Active Design:** At the project scale, Active Design processes must be leveraged that emphasize direct engagement between developers and end-users with regular and rapid iterations that generate fieldable functionalities on a short timescale.
- **Evolutionary Innovation:** Across projects, an environment for Evolutionary Innovation must be established that allows novel and unanticipated solutions that are effective to emerge and propagate, and minimizes risk of misallocating large-scale funding to projects that do not prove to be effective or timely by providing off-ramps.
- **Setting Policy:** Policymakers can create the environment for Active Design and Evolutionary Innovation to meet known requirements and explore what is possible by incentivizing the funding of small businesses with lean funding vehicles and decentralizing funding decisions.

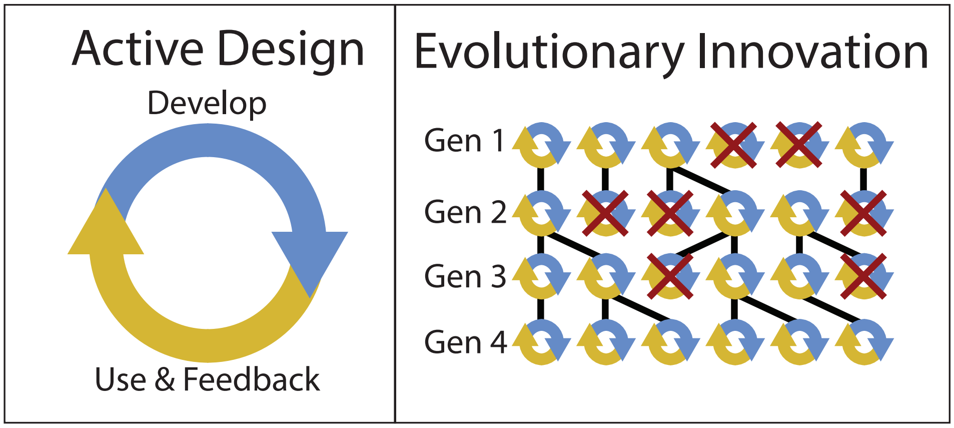

A two-layered strategy. Active Design and Evolutionary Innovation to thrive in uncertainty.

Discussion

The unceasing development of novel technologies that expand the possible ways we interact, cooperate, and compete is impacting all aspects of life, importantly in politics and security. Emerging technologies, to include Artificial Intelligence (AI), Machine Learning (ML), Virtual and Augmented Reality (VR, AR), Autonomous systems, and Swarm systems, require a radical shift in government policy in order to develop and leverage them to their greatest effect. Policy should both (1) promote the development and integration of new capability sets enabled by novel and often unanticipated tech, and (2) develop appropriate responses to adversaries that leverage them against our interests (*see shaded boxes for a few of the most visible emerging technologies and some of their features that have the potential to disrupt and transform*).

Unlike the technological upgrades of the recent past, which primarily increased the intensity or effectiveness of existing capabilities (e.g. more accurate and powerful rifles; faster, more nimble fighter jets; faster money transfers; higher definition video transmission), these technologies offer new kinds of capabilities that have the potential to transform the nature of interactions and engagements across society and societies. This in turn necessitates that we adopt strategies that can tolerate, or even ***thrive in***, uncertain conditions.

Further, unlike in previous eras where those with the most resources were able to dominate by constructing big-ticket, high-impact systems, in the coming era cheap componentry and easy access to specialized knowledge will de-emphasize those with amassed resources and reward those who are able to rapidly innovate, reconfigure, and evolve technological solutions.

Artificial Intelligence and Machine Learning

There are ongoing debates over what separates artificial intelligence (AI) from other forms of automated technology, but the resolution of this academic detail is of minor importance relative to the major impact that the continued development of AI systems are likely to have. Generally speaking, AI systems can sense, perceive, decide, and/or act in increasingly complex and nuanced ways, with minimal or no human intervention or oversight. The effects of the augmentation and replacement of human beings in decision-making processes can not be understated—ripples will be felt across all domains of human society.

Machine learning (ML) is part of the overarching vision of AI, and an important aspect of many AI systems is their ability to generate, update, and refine their activity based on acquired or provided information. ML *per se* is an umbrella term that covers a wide array of techniques that may be used in the context of AI, or separately, for instance in data analysis. ML is distinct from other software routines in that data is leveraged to discover configurations or parameters that 'tune' the internal model of the system, as opposed to those aspects being directly imposed by a developer.

The current policy environment will not effectively catalyze innovation or integration of these emerging and continuously-evolving technologies. Simply, it promotes practices that focus too much on decomposing well understood problems and applying standard solutions; a lengthy, burdensome process that results in decades-long development and production processes. What is needed instead is policy that creates an environment for active exploration of future possibilities in little-understood territory.

If the U.S. is to keep pace with the coming wave of transformational technology and navigate the sea of flux and uncertainty, policymakers must create the conditions for capabilities development to be ***active*** and ***evolutionary*** at all phases of system life cycle: research and development, acquisition, adoption and integration, deployment, C2, and oversight. If this is done appropriately, it will enable the U.S. to rapidly innovate and give the best opportunity for high-impact, unanticipated advances to expand and improve on existing practices.

UPHEAVAL

One of the hallmarks of this unfolding technological revolution, and the one that primes both public and private sectors for upheaval, is the cost structure associated with intelligent systems. Unlike previous eras, where high-impact systems were associated with high material, engineering, and manufacturing costs (e.g. aircraft carrier, state of the art fighter jet), the components and even development cycles of intelligent and autonomous systems are cheap and getting cheaper. In this new era, the winners will not be those who invest the most money into a small number of high-impact systems, but those who are able to innovate and iterate rapidly via investment in a large number of relatively cheap systems.

This situation may not merely level the playing field between nations and organizations with heterogeneous access to resources, it may in fact lend advantage to those that do not have a history of manufacturing big-ticket items. Those without such history will be poised to innovate rapidly through trial and error and low bureaucratic overhead.

The solution for the U.S. to keep up, or better, stay ahead, is first to not fall into the trap of trying to develop these systems in the same manner as traditionally engineered, big-ticket items. Second, adopt a two-layered strategy to developing, acquiring, integrating, deploying, and controlling systems leveraging these emerging technologies that have transformative potential and an associated high degree of uncertainty.

RUNAWAY FUNDING OF BIG-TICKET ITEMS

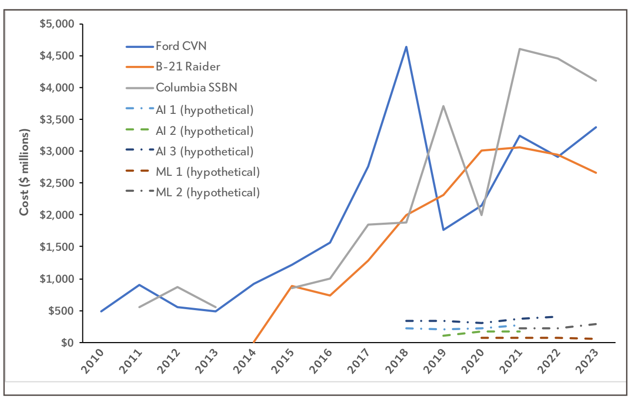

Funding over time. Large item acquisition can lead to runaway costs. There is an alternative approach.

Narrow vs General Intelligence

A crucial distinction lies in the difference between narrow and general intelligence. Narrow intelligence solves problems where (1) desired end states are well known, (2) rules of the operating environment and choices or behavioral options are known in full. An example of narrow intelligence is a system that plays chess. While chess is a challenging game with a very large number of possibilities, (1) it is well- known what winning looks like, and (2) what kinds of moves, options, or behaviors are available (and are not available) at any moment are well-known. Thus, while it may require a considerable amount of intelligence to win a chess game against a master, if that is the only thing a system can do, it is rather narrow. Still, the ‘progress’ that narrowly intelligent systems have made is impressive. Google’s AlphaGo system, for instance, became the first machine to defeat a world champion in the game of Go, a game whose potential complexity dwarfs that even of chess.

Beyond narrow intelligence, the crown jewel of AI is in the synthesis of systems with general intelligence. That is, intelligence that, like our own human intelligence, can problem solve in contexts where (1) desired end states are not necessarily known, nor final, and (2) the rules of the environment and options for action to achieve proximate and distal goals are not explicit. In other words, general intelligence not only solves problems, and not only solves problems across domains, but also continuously figures out what problems to solve to generally achieve progress that is not always (often not) well-defined.

By most accounts, artificial general intelligence (AGI) has not yet been achieved, and its achievement is expected by many to be a transformational breakthrough whose consequences are unknown. By some estimations this uncertainty carries with it considerable, even *existential*, risk.

A TWO-LAYERED STRATEGY FOR INNOVATING UNDER UNCERTAINTY: POLICIES FOR ACTIVE AND EVOLUTIONARY DESIGN

As noted above, a large degree of uncertainty exists around the capabilities and impacts that will result from the development of a suite of emerging technologies. At least some of the innovations to come will not only augment or enhance current capabilities, but generate novel capabilities that require rethinking the approaches and political structures that support leveraging them in responsible and effective ways. At the same time many attempts at development will not move the needle much. Yet it will be very difficult to predict which will be which without testing them out.

Uncertainty scares bureaucracy, thus, policymakers must set the conditions for a two-layered strategy of active and evolutionary design and innovation to minimize risk. This approach can lead to success and superiority in this uncertain environment where actions, interventions, and outcomes can only be considered probabilistically. Short of this, more nimble adversaries with less bureaucratic baggage will seize the day.

ACTIVE DESIGN

The first layer is to engage in active design across the life cycle of a system that takes advantage of novel technological capabilities; when developing, testing, fielding, and deploying. An active approach leverages rapid iterative cycles, engaging all developers, customers, and other stakeholders on a regular basis to refine capabilities, adjust based on experience and feedback, and redeploy – repeating this iterative process for the duration of system life cycle. This process blurs the traditional phases of R&D, procurement, and deployment, and similarly in a policy environment, policy establishment, direction, and execution, as it is understood refinement and development is an ongoing process that continues after initial deployment. Moreover, it emphasizes the exploration of the possible, rather than attempting to fulfill predefined requirements, a process with well documented limitations in the face of complexity and uncertainty. Shifting focus from responding to requirements lists as articulated by a minority of stakeholders to capability exploration is necessary if we are to discover high-impact and effective uses of emerging technology in a timely manner—and represents a disruptive move to the historical practice.

The recognition of the need for iterative refinement is in line with visible shifts from DoD who have

Autonomous Systems

Autonomous systems make many or all decisions in the absence of direction or oversight. Autonomy comes in degrees; semi-autonomous systems may make some decisions or take some actions without oversight, but seek or require it for others. Intelligence enables autonomy. A system that is able to process complex information, make decisions, and take action that is useful relieves the need for oversight. The more general the intelligence, the wider the possible autonomy—one of the reasons AGI is considered risky by some.

Autonomous systems may be physical (i.e. robotic) or virtual (i.e. reside in digital ecosystems), and potentially exotic mixtures of the two; something we currently lack even the language, nevermind the policy, to grapple with. The weaponization and harmful side effects of autonomous systems are inevitabilities, presenting many ethical, legal, and practical challenges. Policymaking will be crucial in determining how we respond to harmful effects as we discover them, including determining liability and mitigating risk while balancing opportunity. These issues are fast approaching as just last month the first pedestrian death was recorded as a result of the behavior of an autonomous vehicle developed by Uber.

indicated the desire to move beyond traditional requirements-based engineering approaches. The National Defense Strategy Summary released in January 2018 states:

> *"A rapid, iterative approach to capability development will reduce costs, technological obsolescence, and acquisition risk. … Prototyping and experimentation should be used prior to defining requirements … Platform electronics and software must be designed for routine replacement instead of static configurations that last more than a decade."*

Active design is iterative and experimental, recognizing that the best solutions to complex problems come from prototyping, fielding, and refining; not from attempting to exhaustively list all desired 'requirements' in a sterile vacuum or one-way processes that demand unattainable foresight.

Swarms

Most of the popular and much of the academic focus in AI is on individual intelligence. That is, a central decision maker that processes information and produces output(s). For a fuller view of the potential impacts of AI, collective or swarm intelligence and behavior must be considered as well.

Collective intelligence is exhibited when groups of agents operate in parallel and interact with the environment and one another (either directly or indirectly through the environment, i.e. 'stigmergy'). It is possible for groups of relatively simple individual agents to solve complex problems and display complex collective behaviors.

The natural world gives us many examples where collectives of somewhat simple agents organize themselves and interact to solve complex problems none of them can individually. Ants and bees locate and exploit food sources and calculate optimal waste sites, termites build structures whose temperature is internally regulated for biological efficiency, starlings form complex dynamic aerial patterns to confuse and evade predators. These so-called swarms are not centrally controlled, like traditional C2 systems, but rather embody a distributed control architecture, where each (somewhat) intelligent and/or autonomous agent makes decisions based on local information.

The decentralized and potentially low-cost nature of artificial swarms make them both an attractive area of exploration for future capabilities, as well as a major challenge for current concepts of C2, both in projection and defense. As an illustration, a swarm of 100, or even 1000, cheap autonomous drones operating as a swarm could be able to outperform a F-35 in many if not most operational settings for a fraction of the cost and dramatically increased replaceability.

EVOLUTIONARY INNOVATION

Active design can help find effective solutions to complex problems where traditional, linear/sequential approaches fail. However, to more fully explore the possibilities that emerging technologies engender, and to hedge the risk of failure of individual approaches/systems, a large number of active design efforts should be undertaken in parallel. These parallel efforts may address the same or similar problem sets, or different ones, but what is important is their approach is in some way different.

The ensemble of parallel efforts will produce an array of outcomes, out of which harmful or those with the least promise can be scrapped, and novel approaches or variants on existing ones can fill the available resource space. In other words, evolution can take place across the ecosystem of active policy projects.

As in evolutionary systems in biology, evolutionary innovation has the potential to discover surprising solutions to complex problems simply by (1) generating and fostering variety, (2) exposing systems to stress, (3) culling systems that fail under stress, and (4) allowing those that perform reasonably well to spread, adapt, mutate, interact, and hybridize.

This approach casts a wide net across the space of possible innovators, giving the greatest chances for discovering unanticipated, high-impact innovations. The novelty of emerging tech is such that well-known systems developers and manufacturers do not have any clear innovational advantage over small, lean, and agile development teams. The temptation must be resisted to "pick winners" based on legacy contracts for big-ticket items; it simply cannot be predicted who the major advances will come from. This strategy would have the added benefits of economic stimulation across a large number of small companies, encouraging an entrepreneurial environment, distributing experience and expertise rather than siloed, and enhancing the health of public-private relationships and cross-fertilization.

SETTING THE POLICY CONDITIONS

Policymakers must set the conditions for active design and evolutionary innovation across all sectors if the U.S. is to thrive throughout the coming technological wave. We offer several considerations that could help build such an environment that catalyzes innovation and minimizes monetary and other operational risks:

- **Incentivize contract awards for small and new companies.** Major players in this space will reveal themselves in time, many of which may not even exist yet. Moreover, a track-record in the production of systems from previous eras may not be a good predictor of performance in this emerging space.

- **Provide lean contracting vehicles.** Make development contracts brief by default, relatively small in scale, easy to initiate, renew, adjust, and terminate. The most important advances won't necessarily come from the largest injections of money. Lean contracting vehicles will facilitate the exploration of a large variety of projects with minimal risk due to unfavorable outcomes. This variety sets the environment up for an evolutionary selection process.

- **Decentralize (small) funding decisions.** Let the teams that will be directly impacted by

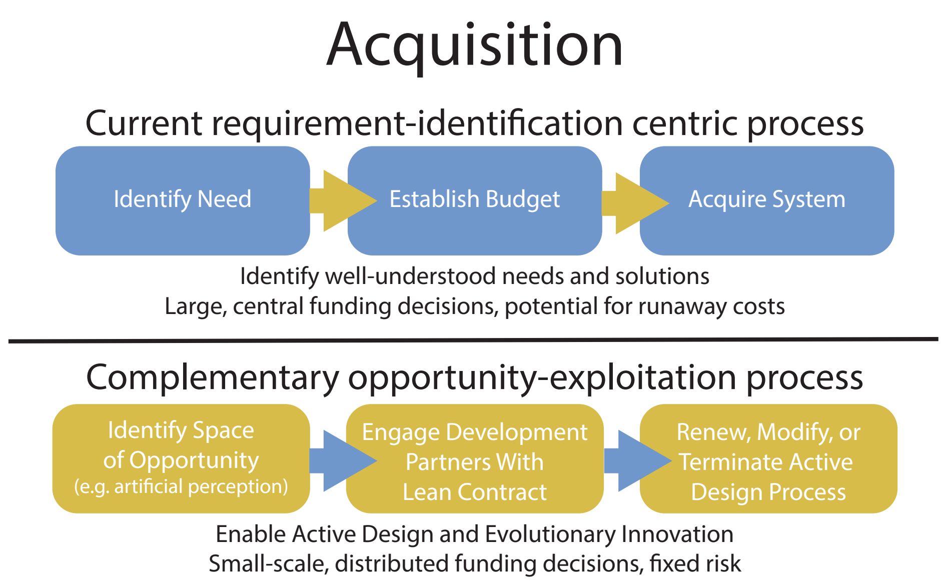

A complementary acquisition process. A process focused on exploring what is possible to achieve.

novel systems choose suitable partners and evaluate their performance. This will promote the engagement that active design demands.

CONCLUSION

Policymakers can set the conditions for success in leveraging emerging technologies for the benefit of society by enabling active design and evolutionary innovation.

Active design has the potential to move us from stilted, clunky, and ultimately ineffective system development towards organic processes that catalyze the exploration into possible futures, creating systems that achieve their aims.

Evolutionary innovation can move us from a process that was suited for an environment dominated by those with the most imposing high-impact systems, to a process that thrives in uncertainty. This environment of continuous innovation and disruption demand constant adaptation, wherein power is often derived from information and intelligence rather than physical force.

Neptune Advisory

202-827-0277
www.neptuneasc.com

415 8th Street SE
Washington, DC 20003

Scott Ellison scott@neptuneasc.com

Patrick McCarthy, CFA patrick@neptuneasc.com

Joe Norman jnorman@neptuneasc.com

David Schopler schop@neptuneasc.com