Future Defining Innovations: Trustworthy Autonomous Systems

Ashley J. Llorens, Erin N. Hahn, Marisa J. Hughes, Glen H. Fountain, Bruce A. Swett, and Cara E. LaPointe

ABSTRACT

Intelligent systems are already having a remarkable impact on society. Future advancements could have an even greater impact by empowering people through human–machine teaming, addressing challenges with vast geographic scales, and accelerating interstellar discovery. Creating intelligent systems that can be trusted to operate autonomously is a grand challenge for humanity. In this article, we explore potential futures for trustworthy autonomous systems, identify some of the significant challenges, and illustrate potential pathways by describing developments underway at the Johns Hopkins University Applied Physics Laboratory (APL).

INTRODUCTION

Popular narratives in books and movies portray a future in which machines act with levels of intelligence and autonomy that greatly exceed today's technology. Netflix's *I Am Mother* offers a future in which an attempt to create robots to protect humanity inadvertently results in robots attempting to redesign our entire species to make it better. In contrast, Disney's *WALL-E* presents a future in which robots play a key role in helping humanity survive and return to Earth.

As compelling as these stories are, we can take for granted neither the advance of the technologies needed to achieve these futures nor the potential outcomes that could result. Institutions like APL play a key role in helping to guide the advance of technology toward ensuring the security and prosperity of our nation—and human society more broadly—far into the future. This requires envisioning the outcomes we hope to achieve and guiding the evolution of the enabling technologies toward integrating into the complex ecosystems in which they will operate.

To accomplish complex goals, machines need intelligence—the ability to **perceive** and understand their environment, to **decide** on a course of action that best achieves the system's design goals, to **act** with a degree of autonomy to carry out those goals, and to do all this as part of a **team** with humans and other machines. Although we are focused in this article on the advancement of intelligent systems that can perform complex tasks autonomously, we recognize that there is no such thing as a fully autonomous system. Autonomy results from the delegation of a decision to an authorized entity to take action within specific boundaries.[1] To enable effective interaction and allocation of tasks among humans and machines, it is critical to develop an effective trust relationship.

Creating intelligent systems that can be trusted to operate autonomously in uncontrolled, open-world environments is a grand challenge for humanity. In this article, we explore the potential futures for trustworthy autonomous systems, identify some of the major challenges, and use current developments underway at APL to illustrate potential pathways. We examine the potential for trustworthy autonomous systems to empower people through enhanced human–machine teaming. Then we explore the potential impacts and challenges in promoting global health and prosperity and look toward accelerating discovery within and beyond or solar system. Lastly, we discuss the importance of developing an effective trust relationship with autonomous systems and offer concluding thoughts.

AUTONOMOUS SYSTEMS TOMORROW

Empowering People through Human–Machine Teaming

Humans and machines have different strengths and weaknesses. The human brain has evolved to perform a vast array of cognitive functions, enabling us not only to survive in a complex world but to set and achieve long-term goals for society. Yet, as impressive as human intelligence is, it is far from perfect. As behavioral economics teaches us, we can make irrational decisions that are not ultimately in our best interest.[2] Our attention spans can be short and our perception of the world is inherently biased. Our ability to evaluate a broad assortment of potential courses of action is limited. The bodies in which our brains reside have adapted for the particular conditions on the surface of our planet.

In contrast, machines can be constructed to operate in the vacuum of space or on the ocean floor but can't gracefully adapt if conditions change too far beyond what human designers may have imagined. Machines can vigilantly attend to complex data feeds for as long as they have power but don't always focus on the right things. Machines can consider millions of possible courses of action but can't evaluate the quality of their outcomes with the nuance of our best human decision-makers.

Effective teaming can elevate human capabilities by leveraging the unique strengths of humans and machines, but perfecting the human–machine partnership poses an enduring challenge (Figure 1). Making progress will require the development of machines capable of teaming on complex tasks in dynamic and uncertain environments by maintaining shared situational awareness with human teammates informed by an understanding of their intent, capabilities, and current state.

A promising approach to creating human-aware machines is to enable them to measure human intent or state directly through physiological and neurological sensors. For the Defense Advanced Research Projects Agency (DARPA) Revolutionizing Prosthetics program, APL developed algorithms to derive robotic control signals for the Modular Prosthetic Limb (MPL) via neural activity measured through electrode arrays implanted in the motor cortex.[3,4] Machine learning algorithms translate these signals into motor intent and then into actions executed by the MPL. Through this coupling, the MPL has been controlled as an intuitive extension of the human body, performing manipulation tasks informed by an understanding of signals in the human brain and body as part of the overall operational context.

In many applications, human intent and state must be inferred even when they cannot be directly measured. Under DARPA's Air Combat Evolution (ACE) program, APL is helping to create an intelligent tactical autopilot that can assist human pilots in a dogfight by performing tactical maneuvers learned autonomously in simulation.[5] Effective human–machine teaming in this kind of combat situation requires shared understanding of the task and the operating environment. The pilot must have a sufficient mental model of the autopilot capabilities to ensure that they are evoked at the appropriate time. Conversely, the autopilot must have some understanding of operator and commander's intent, as well as the limitations and capabilities of the pilot. Overcoming these challenges can enable the human–machine team to achieve superhuman performance in many real-world applications, as illustrated by the culmination of DARPA's AlphaDogfight Trials during which an artificial intelligence (AI) scored a 5-0 victory over a skilled F-16 pilot in a simulated dogfight.[6]

Figure 1. Human–machine teaming. Effective teaming can elevate human capabilities by leveraging the unique strengths of humans and machines, but perfecting this partnership poses an enduring challenge.

Thought-control of complex systems and human–machine teaming in tactical combat offer compelling examples of emerging capabilities that rely on effective symbiosis of human intelligence and AI. Thoughtful advancement of enabling technologies for human–machine teaming stands to enhance a broad range of applications in national security, space exploration, and health. Achieving capabilities that are both effective and trustworthy will require addressing the challenge of ensuring that complex systems perform according to human intent even as they become capable of greater autonomy.

PROMOTING GLOBAL HEALTH AND PROSPERITY

The systems that support human populations, such as supply chains for food and health care, power and energy grids, water distribution, and evolving land use are increasingly complex and global in scale. Emerging trends are poised to stress our human–Earth systems in novel ways. Populations are increasing in some nations even as they are beginning to decrease in others.[7] The climate is growing warmer and more variable, causing natural disasters with increasing intensity.[8] The rapid spread of the novel coronavirus has demonstrated the vulnerabilities of our interconnected systems to biological threats, as well as the inherent challenges in addressing them effectively.

Intelligent autonomous systems offer the potential to help address global challenges through a combination of large-scale analytics and coordinated action (Figure 2). While effective combinations of these two capabilities have yet to be demonstrated at scale, the viability of near-real-time analytics has been demonstrated through their use in recent responses to natural disasters and COVID-19.

Researchers from APL recently partnered with the Joint Artificial Intelligence Center (JAIC) of the US Department of Defense to aid in humanitarian assistance and disaster relief efforts after Hurricane Dorian struck in the fall of 2019.[9] The APL team created deep-learning algorithms designed to find flooding, blocked roads, and damaged buildings by processing high-resolution overhead imagery. Given the high-resolution imagery, the current automated analysis capability can quickly generate data products used to provide situational awareness during a disaster, which allows decision-makers to prioritize flyovers and other first-responder operations.

The widespread use of the Johns Hopkins University COVID-19 Resource Center[10] is another powerful illustration of how large-scale analytics can help inform solutions that require coordinated behaviors across cities, states, and nations. The promotion of social behaviors such as social distancing to slow coronavirus transmission has been bolstered by the near-real-time feedback from analytics. The situational awareness visualizations and maps in the COVID-19 Resource Center are supported by a semiautonomous data collection infrastructure that continuously harvests global public health data on the outbreak. Independent asynchronous collection agents feed a data fusion pipeline that includes embedded analytics for anomaly detection along with interfaces that support validation and error correction by human operators.

Advances in decision aids and robotic platforms may offer the possibility to leverage analytics like these to drive not only human action, as in these two examples, but also autonomous action by intelligent systems.

Unmanned ground, air, and marine platforms have advanced on a parallel path alongside the use of large-scale analytics. For example, it is increasingly commonplace to see the use of unmanned aerial vehicles by tactical teams, such as fire and rescue, to provide tactical situational awareness.[11] The Center for Robot-Assisted Search and Rescue provides robotic support for disaster operations, such as debris clearing and search reconnaissance.[12] At present, robotic systems like these are transported, operated, and maintained by teams of roboticists. Further advancements in the capabilities and resilience of intelligent systems may accelerate and enhance first response by automating the preparation of complex operating environments before the first human arrives and then providing persistent, self-sufficient support.

Figure 2. Systems that support human populations. Intelligent autonomous systems offer the potential to help address global challenges through a combination of large-scale analytics and coordinated action.

Successfully combining large-scale analytics with autonomous action could unlock new capabilities, such as swarms of intelligent systems that autonomously aid with reforestation by identifying fertile grounds from overhead imagery, planting seeds, monitoring growth, and adjusting local conditions within predetermined bounds. Distributed agents could dynamically optimize the physical placement, orientation, and maintenance of solar cells while also managing the power grids that they feed. Fleets of autonomous buoys fueled by microplastics could patrol the ocean to monitor marine life, track down pollutants, and control carbon-sequestering algae blooms. Analogous to the more tactical examples of human–machine teaming, large-scale applications of intelligent autonomous systems will require the appropriate calibration of trust to ensure that the design goals of these systems remain aligned with beneficial outcomes and that the limitations and risks are sufficiently understood and adequately accounted for.

ACCELERATING INTERSTELLAR DISCOVERY

Beyond assisting humans in tackling global challenges, autonomous systems are poised to help us deepen our understanding of our solar system and the universe beyond (Figure 3). As of September 4, 2020, Voyager 1 is the farthest-ranging spacecraft and is now at ~150 Astronomical Units (AU) from Earth—that is, 150 times the distance between the Earth and the Sun.[13] NASA has recently commissioned a study of the trade space available to a pragmatic near-term Interstellar Probe, now being conducted at APL. The aspirational goals for the Interstellar Probe include designing a mission to operate for 50+ years, with the capability to communicate data back to Earth at up to 93 billion miles away from the Sun (1,000 AU), all with technology that would be launch-ready no later than 2030.

At this distance from Earth, space exploration systems must also be able to fix themselves if something goes wrong. Advances in onboard autonomous fault management beyond those made over the past decade may be helpful as the opportunities for direct human intervention become increasingly limited given the growth of the round-trip light time, and hence significant delay in communications with Earth. For example, the APL-built MESSENGER spacecraft had a sunshade to keep the spacecraft from getting too hot while orbiting Mercury. If an attitude error occurred that exposed the spacecraft to the Sun, MESSENGER would have overheated in 30 minutes. A similar issue experienced by the currently operating Parker Solar Probe spacecraft would require a recovery time of 10 seconds.

The NASA-funded Dragonfly mission, now in development at APL, is scheduled to launch in 2026 and arrive at Saturn's moon, Titan, in 2034.[14] Dragonfly will require onboard systems to fly and land the rotorcraft on Titan, and subsequent flights on the surface will have to occur without Earth in the loop. This will be accomplished via a suite of sensors, actuators, and flight software applications responsible for autonomously executing flight profiles, assessing landing site safety, managing power, and addressing faults during flight. The system design leverages heritage sensors, elements of the Parker Solar Probe autonomy system, and algorithms developed under the Autonomous Landing Hazard Avoidance Technology and Robotic Lunar Lander programs.[15]

In the case of the Interstellar Probe concept, operating at 1,000 AU means a round-trip light time of 11.5 days and severely limited telemetry rates. Current data compression techniques will likely allow acceptable data downlink; however, to fully utilize the instruments Interstellar Probe might fly, more significant compression on the order of 100:1 may be required. In addition to its improved fault tolerance, this mission concept may be enhanced by algorithms capable of analyzing instrument data onboard the spacecraft and making autonomous decisions as to which data to transmit back to Earth—including those that contain anomalous or unexpected results.

While the level of AI and autonomy required for Interstellar Probe may seem advanced, space systems developed over the last 60 years show a clear trend in the direction of the required technological advancements.

Figure 3. Systems that advance our understanding of space. Autonomous systems are poised to help us deepen our understanding of our solar system and the universe beyond.

Starting in 1996 with the Near Earth Asteroid Rendezvous (NEAR) mission, these system elements have been implemented with autonomy rules, macros, storage variables, and computed telemetry that can be uploaded directly to the spacecraft (personal communication, T. Adrian Hill, APL). Space exploration systems have become more and more sophisticated over the past two decades—autonomous decisions and actions made by future systems will increasingly influence the direction of our scientific investigation and understanding.

TRUSTING AUTONOMOUS SYSTEMS

In the previous sections, we explored the potential value to society if we are successful in advancing intelligent autonomous systems. Yet, realization of this potential requires trust—from the individual level to the societal level. Developing trust between humans and autonomous systems is a complex undertaking that includes both ensuring that the technology itself is developed with trustworthy characteristics and that society is employing the technology in a trustworthy way. Underlying this trust relationship, there is potential danger in over-trusting systems (i.e., delegating tasks to systems that are not capable of performing them reliably) and potential inefficiencies through under-trusting systems (i.e., failing to delegate tasks to systems when there is a clear benefit over the best alternative).[16] Achieving an optimal trust relationship between humans and machines will require targeted advancements across the full technology cycle,[17,18] as well as the associated legal, ethical, and policy frameworks.

Understanding the trustworthiness of autonomous technologies can be particularly challenging for military systems where operators must have some idea of the environment and mission states under which a system is capable of performing a particular task.[19] Distinguishing one state from another for a complex system is not straightforward. Whereas people often understand state at the semantic level (e.g., "day" versus "night"), machines experience environment states as statistical distributions of input variables and sensor measurements interpreted according to complex code and parameters, some of which may have been programmed explicitly while others may have been generated through machine learning.

Given these complexities, systematically equipping operators with a mental model of how an intelligent system will perform as a dynamic function of environment state is yet another enduring challenge. Making progress will require new perspectives, methods, and technologies over the entire life cycle of the system from research, design, and development through test and evaluation, operation, and maintenance. Current research underway at APL seeks to advance testing and evaluation capabilities for intelligent autonomous systems, such as onboard watchdogs,[20] advanced "simulation to reality" capabilities,[21,22] and scalable formal methods techniques.[23] (See the short article by Kouskoulas et al., in this issue, for a glimpse into some of APL's work in formal methods).

Recognizing the critical role of trust in realizing the full potential of intelligent systems, Johns Hopkins University has recently established the Johns Hopkins Institute for Assured Autonomy (IAA) to drive toward a future where autonomous systems are trustworthy contributors to society. By working across the entire spectrum of foundational research to application, IAA focuses on the complex intersections among technology, ecosystem, and policy and governance to advance the foundational underpinnings for calibrating trust in autonomy. See the article by LaPointe et al., in this issue, for more on the IAA.

CONCLUSION

Intelligent systems are already having a tremendous impact on society. In this article, we explored applications where future advancements could have an even greater impact by empowering individuals through human–machine teaming, addressing challenges that span large geographic scales, and accelerating interstellar discovery.

These futures are by no means guaranteed and will require focused investments to thoughtfully advance the ability of machines to perceive and understand their environment, make complex decisions that align with human design goals, and team with humans in carrying out complex tasks. To create systems that are worthy of trust, we must adapt the entire technology life cycle to align with societal goals and institutional best practices. To appropriately calibrate our trust, we must integrate an understanding of technological capabilities and limitations into our legal, ethical, and normative frameworks. Making meaningful progress along all these fronts will enable us to continue turning what may seem like science fiction today into trustworthy capabilities tomorrow.

REFERENCES

[1] Defense Science Board, "Report of the Defense Science Board Summer study on autonomy," Department of Defense, Washington, DC, Jun. 2016, https://dsb.cto.mil/reports/2010s/DSBSS15.pdf.

[2] B. Appelbaum, "Nobel in economics is awarded to Richard Thaler," *New York Times*, Oct. 9, 2017, https://www.nytimes.com/2017/10/09/business/nobel-economics-richard-thaler.html.

[3] G. Hotson, D. McMullen, M. Fifer, M. Johannes, K. Katyal, et al., "Individual finger control of the Modular Prosthetic Limb using high-density electrocorticography in a human subject," *J. Neural Eng.*, vol. 13, no. 2, pp. 1–25, 2016, https://doi.org/10.1088/1741-2560/13/2/026017.

[4] A. D. Ravitz, M. P. McLoughlin, J. D. Beaty, F. V. Tenore, M. S. Johannes, et al., "Revolutionizing Prosthetics—Phase 3," *Johns Hopkins APL Tech. Dig.*, vol. 31, no. 4, pp. 366–376, 2013, https://www.jhuapl.edu/Content/techdigest/pdf/V31-N04/31-04-Ravitz-RP3.pdf.

[5]D. Javorsek, "Air Combat Evolution (ACE)," Defense Advanced Research Projects Agency, https://www.darpa.mil/program/air-combat-evolution. Accessed Sep. 4, 2020.
[6]J. Surowiec, "AI bests human fighter pilot in AlphaDogfight trial at Johns Hopkins APL," press release, Laurel, MD, APL, Aug. 28, 2020, https://www.jhuapl.edu/PressRelease/200828-AI-bests-human-fighter-pilot-in-AlphaDogfight-trial-at-APL.
[7]A. Cilluffo and N. G. Ruiz, "World's population is projected to nearly stop growing by the end of the century," *Fact Tank*, Pew Research Center, Jun. 17, 2019, https://pewrsr.ch/2WJzNHf.
[8]C. Harvey, "Scientists can now blame individual natural disasters on climate change," *Scientific American*, Jan. 2, 2018, https://www.scientificamerican.com/article/scientists-can-now-blame-individual-natural-disasters-on-climate-change/.
[9]K. Elkharbibi, "APL shaping an intelligent approach to disaster response and relief," press release, Laurel, MD, APL, Sep. 26, 2019, https://www.jhuapl.edu/PressRelease/190926.
[10]Johns Hopkins University and Medicine, Coronavirus Resource Center, https://coronavirus.jhu.edu/.
[11]M. LaBerge, "Drone flights to begin over Bear Creek Greenway," NBC, Jul. 29, 2020, https://kobi5.com/news/drone-flights-to-begin-over-bear-creek-greenway-133560/.
[12]E. Pittman, "How robots are changing disaster response and recovery," *Government Technology*, Jan. 29, 2015, https://www.govtech.com/products/GT-How-Robots-Are-Changing-Disaster-Response-and-Recovery.html.
[13]"Voyager mission status," NASA Jet Propulsion Laboratory, https://voyager.jpl.nasa.gov/mission/status/. Accessed Sep. 4, 2020.
[14]M. Buckley, "Destination: Titan," press release, Laurel, MD, APL, Jun. 27, 2019, https://www.jhuapl.edu/PressRelease/190627b.
[15]T. G. McGee, D. Adams, K. Hibbard, E. Turtle, R. Lorenz, et al., "Guidance, navigation, and control for exploration of Titan with the Dragonfly rotorcraft lander," in *Proc. AIAA Guid., Nav., and Contr. Conf.*, Kissimmee, FL, Jan. 8–12, 2018, pp. 1–7, https://doi.org/10.2514/6.2018-1330.
[16]M. R. Endsley, "Autonomous horizons: System autonomy in the Air Force—A path to the future, Volume 1: Human autonomy teaming," US Department of the Air Force, AF/ST TR 15-01, Jun. 2016, https://www.hsdl.org/?abstract&did=768107.
[17]J. M. Mueller, "The ABCs of Assured Autonomy," in *Proc. 2019 IEEE Int. Symp. Technol. Soc. (ISTAS)*, Medford, MA, 2019, pp. 1–5, https://doi.org/10.1109/ISTAS48451.2019.8938010.
[18]M. Brundage, S. Avin, J. Wang, H. Belfield, G. Krueger, et al., "Toward trustworthy AI development: Mechanisms for supporting verifiable claims," 2020, https://arxiv.org/abs/2004.07213.
[19]B. A. Swett, E. N. Hahn, and A. J. Llorens, "Designing robots for the battlefield: state of the art," presented at the *Pontifical Acad. Sci. (PAS) and the Pontifical Acad. Social Sci. (PASS) Conf. Rob., AI, and Humanity: Science, Ethics, and Policy*, Vatican City (Casina Pio IV), May 16–17, 2019.
[20]D. Scheidt and W. Van Besien, "Safe Testing and Execution of Autonomy in Complex, Interactive Environments," presented at the *Safe and Secure Syst. Softw. Symp. (S5)*, Jul. 12, 2016, http://www.mys5.org/Proceedings/2016/Day_1/2016-S5-Day1_1405_Scheidt_no_video.pdf.
[21]G. E. Mullins, P. G. Stankiewicz, R. C. Hawthorne, J. D. Appler, M. H. Biggins, et al., "Delivering test and evaluation tools for autonomous unmanned vehicles to the fleet," *Johns Hopkins APL Tech. Dig.*, vol. 33, no. 4, pp. 279–288, 2017, https://www.jhuapl.edu/Content/techdigest/pdf/V33-N04/33-04-Hawthorne.pdf.
[22]G. E. Mullins, A. G. Dress, P. G. Stankiewicz, J. D. Appler, and S. K. Gupta, "Accelerated testing and evaluation of autonomous vehicles via imitation learning," in *Proc. 2018 IEEE Int. Conf. Robot. Automat. (ICRA)*, Brisbane, QLD, 2018, pp. 5636–5642, https://doi.org/10.1109/ICRA.2018.8460965.
[23]I. Papusha, R. Wu, J. Brulé, Y. Kouskoulas, D. Genin, and A. Schmidt, "Incorrect by construction: Fine tuning neural networks for guaranteed performance on finite sets of examples," in *Proc. 3rd Workshop on Formal Methods for ML-Enabled Autonomous Systems (FoMLAS 2020)*, Aug. 2020, arXiv:2008.01204 [cs.LG], https://arxiv.org/abs/2008.01204.

Ashley J. Llorens, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Ashley J. Llorens was chief of APL's Intelligent Systems Center, where he directed research and development in machine learning, robotics, autonomous systems, and neuroscience toward addressing the nation's most critical challenges. He received a BS in computer engineering from the University of Illinois and an MS in electrical engineering, with a concentration in digital signal processing, from the University of Illinois at Urbana-Champaign. During his 15 years at APL, Ashley led technical teams in applying machine learning to advance autonomous sensing capabilities for the US Armed Services and served as an advisor to the Defense Science Board. He now works for Microsoft.

Erin N. Hahn, National Security Analysis Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Erin N. Hahn is a senior national security analyst and group supervisor at APL. She earned a BA in philosophy from Salisbury University, an MS in public administration from Katholieke University, and a JD from the University of Maryland. Her work emphasizes legal, policy, and normative issues in the context of operational or technological challenges. She was formerly the associate director of the University of Maryland Center for Health and Homeland Security (CHHS) in Baltimore, Maryland, and prior to that, an attorney in private practice doing complex civil litigation. She serves on the International Panel on the Regulation of Autonomous Weapons (iPRAW), whose work informs the ongoing discussions of the United Nations Group of Government Experts for Lethal Autonomous Weapons Systems. She has authored or coauthored diverse works, including a book chapter in *Robot Ethics 2.0* (Oxford University Press, 2017). Her email address is erin.hahn@jhuapl.edu.

Marisa J. Hughes, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Marisa J. Hughes is assistant supervisor of the Reasoning Machines Section and the assistant program manager for environmental resilience in APL's Research and Exploratory Development Department. She has a BS in mathematics from Binghamton University and an MS and a PhD in mathematics from Cornell University. Marisa's research interests include methods to understand, mitigate, adapt, and reverse environmental threats including climate change, resource shortages, and threatening waste. She has experience in data science, data modeling, systems analysis, machine learning, and management, as well as in Earth systems, intelligence, missile defense, health, neuroscience, and submarine maintenance domains. Her email address is marisa.hughes@jhuapl.edu.

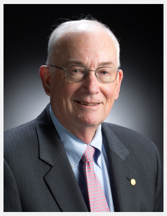

Glen H. Fountain, Space Exploration Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Glen H. Fountain is a program manager in APL's Space Exploration Sector. He received BS and MS degrees in electrical engineering from Kansas State University. During his early career at APL, he was part of the technical staff working to develop space missions for both the Department of Defense and NASA. Since 1974, he has held a number of line and project management positions. From 2004 through 2015, he was the project manager for the New Horizons mission to Pluto, which launched in January 2006 and made observations of the Pluto system in 2015. He has been recognized for his work in various ways, among them the AIAA Von Braun Award for outstanding space program management (2007), NASA Exceptional Public Service Medals (2008, 2016), and the SpaceX/AIAA Award for technical excellence (2015). His email address is glen.fountain@jhuapl.edu.

Bruce A. Swett, Northrop Grumman

Bruce A. Swett is the chief artificial intelligence architect in Northrop Grumman's Mission Systems Sector, where he is responsible for the design and implementation of integrated cloud computing and artificial intelligence (AI) capabilities across the enterprise. Bruce completed his PhD in neuroscience and cognitive sciences at the University of Maryland, College Park and his postdoctoral studies at the National Institute of Deafness and Communications Disorders at the National Institutes of Health. He serves as a subject matter expert and consultant in the areas of AI, brain–computer interfaces, and robotics—both nationally and internationally, having recently served on the Pontifical Academy of Science, advising Pope Francis on AI and robotic technologies. He has created intellectual property and patent applications on seven topics related to neutrally inspired AI. His experimental and computational research focused on using high-performance computing to understand how the brain learns and automates sequences, a topic that applies to novel forms of AI. He was with APL before joining Northrop Grumman.

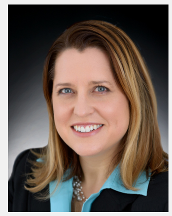

Cara E. LaPointe, Asymmetric Operations Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Cara E. LaPointe is a futurist who works at the intersection of technology, policy, leadership, and ethics. She earned a BS in ocean engineering from the US Naval Academy, an MPhil in international development studies from Oxford University, an MS in ocean systems management from the Massachusetts Institute of Technology (MIT), an Eng degree in naval engineering from MIT, and a PhD in mechanical and oceanographic engineering awarded jointly from MIT and the Woods Hole Oceanographic Institution. She is the co-director of the Johns Hopkins Institute for Assured Autonomy, focused on ensuring that autonomous systems throughout society are safe, secure, reliable and ethical. Cara has served as an advisor to numerous global emerging technology initiatives at the National Academy of Medicine, the United Nations, and the Organization for Economic Co-operation and Development. She is the author of *The Blockchain Ethical Design Framework*, a patented engineer, and a White House Fellow. Her email address is cara.lapointe@jhuapl.edu.