SAND2017-5327R

SANDIA REPORT

SAND2017-XXXX
Unclassified Unlimited Release
Printed XXXX 2017

Autonomy and Complexity at Sandia Executive Summary of Academic Alliance Workshop on Autonomy and Complex Systems

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2017-xxxxC
Unlimited Release
Month 2017

Autonomy and Complexity at Sandia
Executive Summary of Academic Alliance Workshop on Autonomy and Complex Systems

Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185

Sandia has identified autonomy as a strategic initiative and an important area for providing national leadership. A key question is, "*How might autonomy change how we think about the national security challenges we address and the kinds of solutions we deliver?"* Three workshops at Sandia early in 2017 brought together internal stakeholders and potential academic partners in autonomy to address this question. The first focused on programmatic applications and needs. The second explored existing internal capabilities and research and development needs. This report summarizes the outcome of the third workshop, held March 3, 2017 in Albuquerque, NM, which engaged Academic Alliance partners in autonomy efforts at Sandia by discussing research needs and synergistic areas of interest within the complex systems and system modeling domains, and identifying opportunities for partnering on laboratory directed and other joint research opportunities.

Table of Contents

Workshop Goals....6
Autonomy – What Is It and How Do We Think about It at Sandia?....8
Defense Systems and Assessments....8
Broadening National Security Perspectives for Autonomy....10
Autonomy and Complexity – What Are the State of the Art, Opportunities, Challenges, and Risks?....13
Perspectives from Purdue University (Professor Dan DeLaurentis)....13
Purdue University Research Overview....13
State of the Art, Opportunities, Challenges, and Risks....13
Perspectives from University of New Mexico (Professor Meeko Oishi)....14
UNM Research Overview....14
State of the Art, Opportunities, Challenges, and Risks....14
Perspectives from University of Illinois Urbana-Champaign (UIUC)....14
UIUC Research Overview (Professor Geir Dullerud)....14
State of the Art, Opportunities, Challenges, and Risks (Professor Naira Hovakimyan)....14
Perspectives of Georgia Tech University (Professor Fumin Zhang)....15
Georgia Tech Research Overview....15
State of the Art, Opportunities, Challenges, and Risks....15
Perspectives of University of Texas at Austin (Professor Ufuk Topcu)....16
University of Texas (UT) Research Overview....16
State of the Art, Opportunities, Challenges, and Risks....16
Discussion....17
Autonomy and Complexity – What Are the Research Partnership Opportunities?....19
Common Core Discussion Questions for Breakout Groups....19
Human Machine Teaming (HMT) Breakout Group....19
HMT Specific Discussion Questions....19
Discussion....19
Summary....22
Robust, Reliable, and Trusted Systems (RRTS) Breakout Group....22
RRTS Specific Discussion Questions....22
Discussion....23
Summary....23
Distributed Control and Cooperative Systems (DCCS) Breakout Group....24

DCCS Specific Discussion Questions 24
Discussion 24
Summary 25
Common Themes, Follow-Up Questions, Key Takeaways 25
Common Themes 25
Follow-up Questions 26
Key Takeaways 26
Next Steps 27
Appendix A: Workshop Participants 28
Appendix B. Workshop Agenda 30
Appendix C. Breakout Group Participants 31

LIST OF FIGURES

Figure 1 Autonomy, Complexity, and Artificial Intelligence 7
Figure 2 Systems Framework for Discussing Autonomy at Sandia 8
Figure 3 Levels of Autonomy 9
Figure 4 Dimensions of Autonomy Mission Applications and Operational Needs 11
Figure 5 Autonomous System Types and Operating Environments 12
Figure 6 USAF Framework for Insertion of Autonomy into Aircraft Systems (from Autonomous Horizons report) 13

WORKSHOP GOALS

Recent advances in artificial intelligence (AI), spurred on by the availability of new hardware that is enabling deeper and cheaper machine learning, has raised interest in autonomous systems from academia, to government, to industry. The precision and accuracy of self-driving cars, Google Translate, drone technology, and radiological classification systems have all benefitted from these developments. These and many other new enabling technologies for autonomous capabilities present both opportunities and challenges for national security.

Sandia has accordingly identified autonomy as a key priority for national leadership. We are evaluating the opportunities and challenges for adopting autonomy into mission areas as well as implications for strategic engagement.[1] We wish to remove barriers to progress even as we facilitate dialogue and awareness of opportunities, while leveraging capabilities to solve common challenges across areas of application. For example, a common concern across all potential applications is, how do we build trust in autonomous systems to know that they are performing in the way that they are supposed to, know when they are not, and know that correct actions and controls are in place to respond when they are not?[2] How do we design trusted systems that must respond at machine speed (e.g., in the cyber dimension) yet incorporate humans-in-the-loop?

Realizing the potential of autonomy for national security requires perspectives other than technical – to include policy, ethical, and legal dimensions – to ensure adequate measures of safety, security, and reliability that build confidence among all stakeholders in these systems. Solving the toughest of these national science and technology challenges and capitalizing effectively on the opportunities presented by autonomy requires partnerships between our national laboratories and universities. The Academic Alliance initiative links faculty, students, and researchers at key universities with Sandia scientists and engineers to develop collaborative solutions to mission-critical challenges.[3] ***The goal of the workshop with Academic Alliance partners was to identify key complex systems and systems modeling research needs as applied to autonomy and define potential future research directions and collaboration opportunities.***

The research & development (R&D) landscape for autonomy is complex, involving, at a minimum, the natural and physical sciences (e.g., physics, mathematics, engineering), computer and information sciences, neurosciences, and social sciences. All workshop participants have a long history of R&D in many of the areas that underpin autonomous systems, such as pattern recognition, machine learning, computing and information sciences, and robotics, as well as a substantial body of R&D in complex systems science. Attendees explored the use of complex

[1] Prior workshops at Sandia have explored the following questions: In what contexts should Sandia seek to employ autonomy system technologies? In which aspects of autonomy should Sandia provide leadership? How well is Sandia positioned today and what new capabilities will we need to acquire? Who should we partner with? What will commitment to autonomy require?

[2] Sandia's experience in providing such assurance for the nuclear weapons deterrence capability of the United States has led to the "Always-Never" principle of surety that is likely to have relevance for most autonomous systems.

[3] The Alliance schools are Georgia Institute of Technology, Purdue University, University of Illinois at Urbana-Champaign, University of New Mexico, and the University of Texas at Austin. The topic of autonomy is one that has the potential to provide strategic collaboration opportunities.

systems approaches across autonomy-related topics such as human-machine teaming, AI and machine learning, and distributed autonomy and cooperative systems. In addition, the workshop covered the use of complex systems to facilitate the process of designing systems with autonomy, incorporating autonomy into existing systems, and assessing the resilience of an autonomous system (Figure 1). Participants were asked to discuss how the application of complex systems science can help to better understand the opportunities and risks in these research areas, and to identify the technical gaps that complex systems research can address.

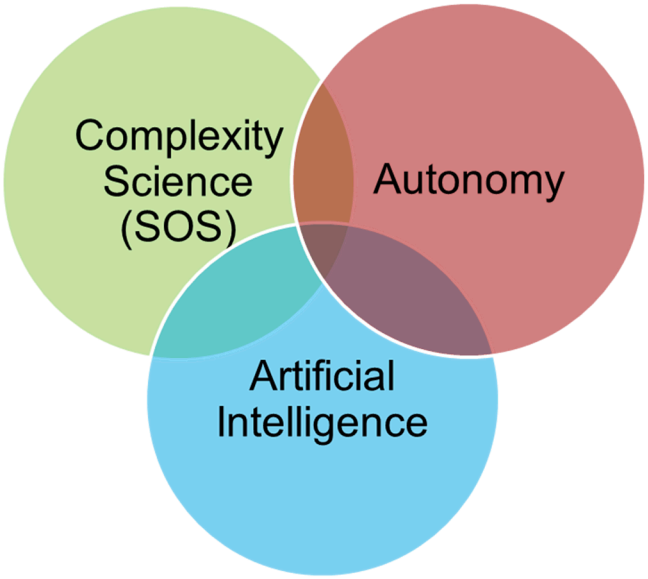

Figure 1 Autonomy, Complexity, and Artificial Intelligence

The format of the workshop allowed for introductory presentations from Sandia, followed by a panel discussion and breakout sessions in the afternoon. Participants were invited to bring their research interests and ideas and were asked to give a brief talk on those ideas during the breakout sessions. Reports from the breakout session summarized areas for potential collaboration.

See Appendix A for workshop participants, Appendix B for agenda, and Appendix C for breakout group participants.

AUTONOMY – WHAT IS IT AND HOW DO WE THINK ABOUT IT AT SANDIA?

Sandia's framework for discussing autonomy employs a systems approach (Figure 2), considering *system utilization* to meet mission needs (to include applications and operational considerations); the *system-level characteristics* required to meet those needs (e.g., how systems are embodied, where they operate, how they are organized, and what level of autonomy is most appropriate); the *technical building blocks* required (e.g., what functions and what technologies); and the organizations and partners required throughout the *lifecycle of the system development* (e.g., providers of enabling science and technology; exploration of system concepts; system design, testing, evaluation, and certification; maintenance and forensic analysis). A specific example of system utilization for defense systems, and the requisite characteristics, building blocks, and development considerations was presented by Sandia, followed by more generalized considerations for utilizing autonomy in other mission areas.

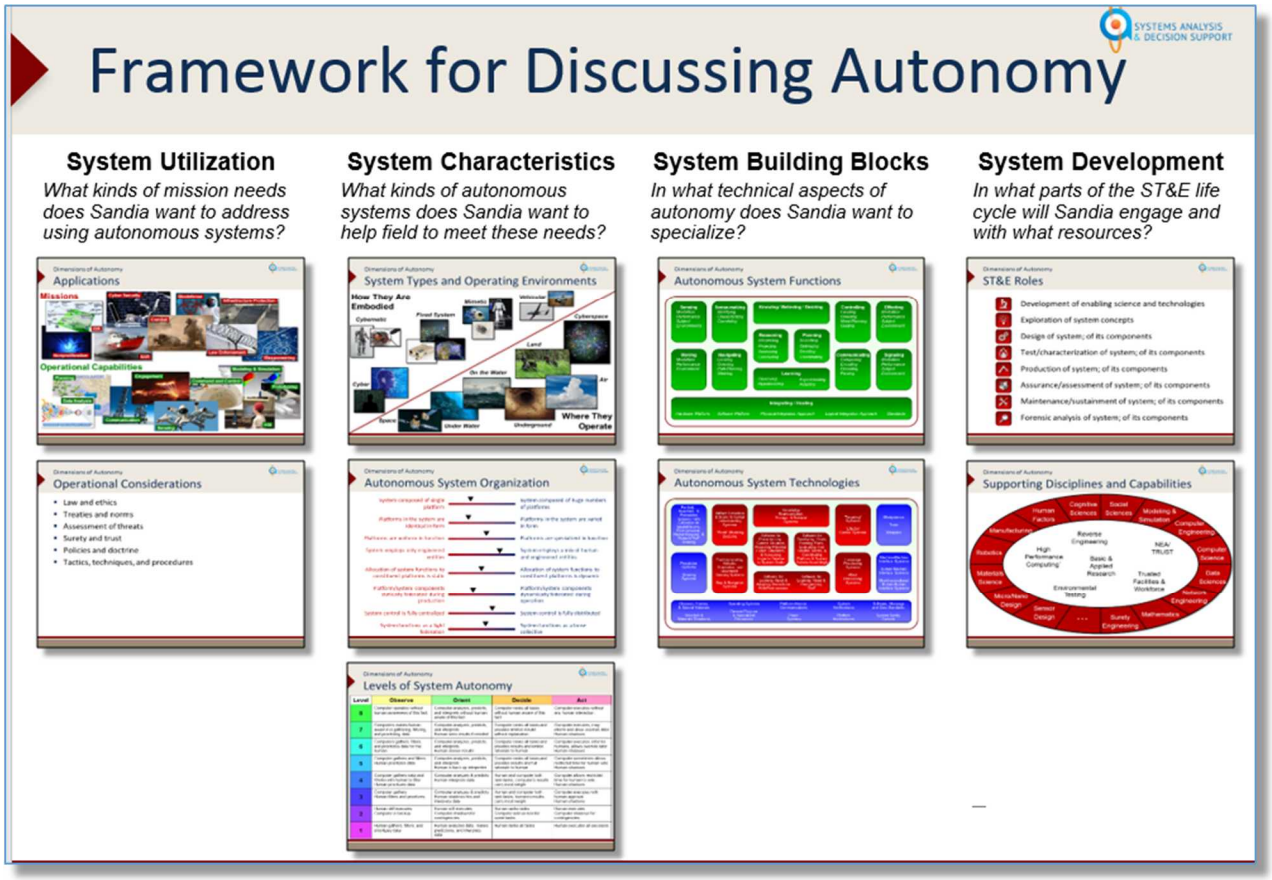

Figure 2 Systems Framework for Discussing Autonomy at Sandia

Defense Systems and Assessments

Among national security stakeholders in autonomy, the Department of Defense (DoD) is a leader in supporting the development of applications for national security. The working definition of

autonomy used in the workshop derived from that adopted by the Defense Science Board Summer Study 2016:

> *Autonomy is a capability (or a set of capabilities) that enables a **particular action** of a system to be…within programmed boundaries, "self-governing."*
>
> *- Defense Science Board July 2012 Task Force Report: The Role of Autonomy in DoD Systems*

In general, the DoD considers autonomy to be capabilities of machines to perform tasks that usually require human intelligence. The operational definitions used by the DoD Defense Science Board Summer Study 2016 locate autonomy along a continuum involving the degree of human and machines in the overall process of sense (observe) – think (orient, decide) – act – team (Figure 3). The operational value of autonomy for DoD is highest in unstructured, adversarial environments with high consequence actions especially when needing fast decision speeds (e.g., cyber operations, missile defense), dealing with high volumes of heterogeneous data leading to complex response actions (e.g., multi-mission operations), facing high danger levels to personnel (e.g., contaminated areas), and/or needing requirements for endurance and persistence (e.g., unmanned surveillance).

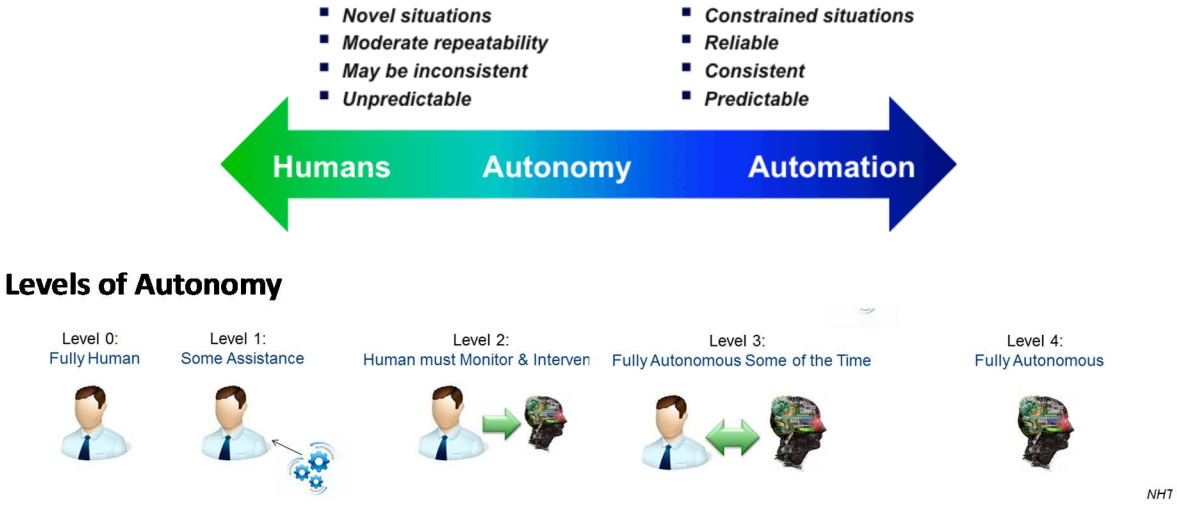

Figure 3 Levels of Autonomy

DoD frequently requires different performance characteristics in these applications than their commercial counterparts (e.g., extremely precise object identification, overcoming denial and deception, and onboard real-time data processing and communications in absence of reliable networks). Sandia supports the DoD in developing applications of autonomy in five of their six operational domains – air, land, space, cyber, and electromagnetic spectrum.[4] Primary areas of support are in the development of onboard sensing systems – including training and data analytics for those systems; using synthetic aperture radar (SAR) for navigation and terminal

[4] DoD also has an active program for autonomy at sea. However, Sandia is not heavily engaged in supporting DoD in this domain.

sensing in lieu of global positioning system (GPS) for Sandia advanced flight systems; and cyber defense and multi-domain cooperative systems.

For onboard sensing, key mission drivers include the need to reduce transmission of data and manpower requirements while maintaining persistence and endurance. Technical solutions involve integrating SAR and automatic target recognition into autonomous surveillance and strike operations, drawing on new capabilities in sensor cuing, generative models, and neural-inspired remote sensing. Key needs for autonomy in advanced flight systems are reduced reliance on GPS for navigation, development of terminal sensing capabilities, exploitation of swarm behavior in offensive and defensive systems, and simulation of potential adversary systems. Within the cyber domain, there is a need to reduce reliance on humans with autonomous systems that can conduct analysis and response to attacks with machine speed, especially across distributed sensor grids. For multi-domain cooperative systems, the DoD requires robust collection management, in which autonomy facilitates rapid re-direction of sensors in real time, with distributed control across platforms that coordinate to achieve better theater coverage across different domains, while meeting the challenges of diverse data owners and the need to avoid detection.

From the perspective of defense systems, key questions for autonomy are

- Robust, reliable and trusted systems:
    - How do we know the machine is doing what it is supposed to be doing?
    - How do we ensure credibility of data – either in training or in operations?
    - How do we abstract from limited data models to get reliable real-world performance?
    - How do we quantify and communicate uncertainty on the fly in autonomous systems?
- Distributed control and cooperative systems
    - How can the system architecture that is optimized for command and control also be resilient and secure?
    - How can goals be imparted to isolated agents such that their joint actions achieve a common autonomous goal?
- Human-machine teaming
    - When do machines work best on their own and when should they be paired with humans?
    - How should trust be calibrated depending on the level of autonomy versus human intervention in a system?
    - What variables do you use to assess trust?
    - How do humans and machines communicate mutual trust levels and metrics?

Broadening National Security Perspectives for Autonomy

Sandia's mission areas in global security, energy and climate, and nuclear weapons have also explored how autonomy shapes the future landscape of national security needs and the operational considerations that must be addressed. A common theme across these mission areas is that, in addition to the applications of autonomy in warfighting scenarios for national security,

autonomy holds significant promise to improve physical security of sites and materials of importance to national security. These applications may be embodied in many different types of systems – fixed, mimetic, cyber, vehicular, or cybernetic, and may operate in many different types of physical domains – air, underground, underwater, etc. (Figure 4 and Figure 5).

Figure 4 Dimensions of Autonomy Mission Applications and Operational Needs

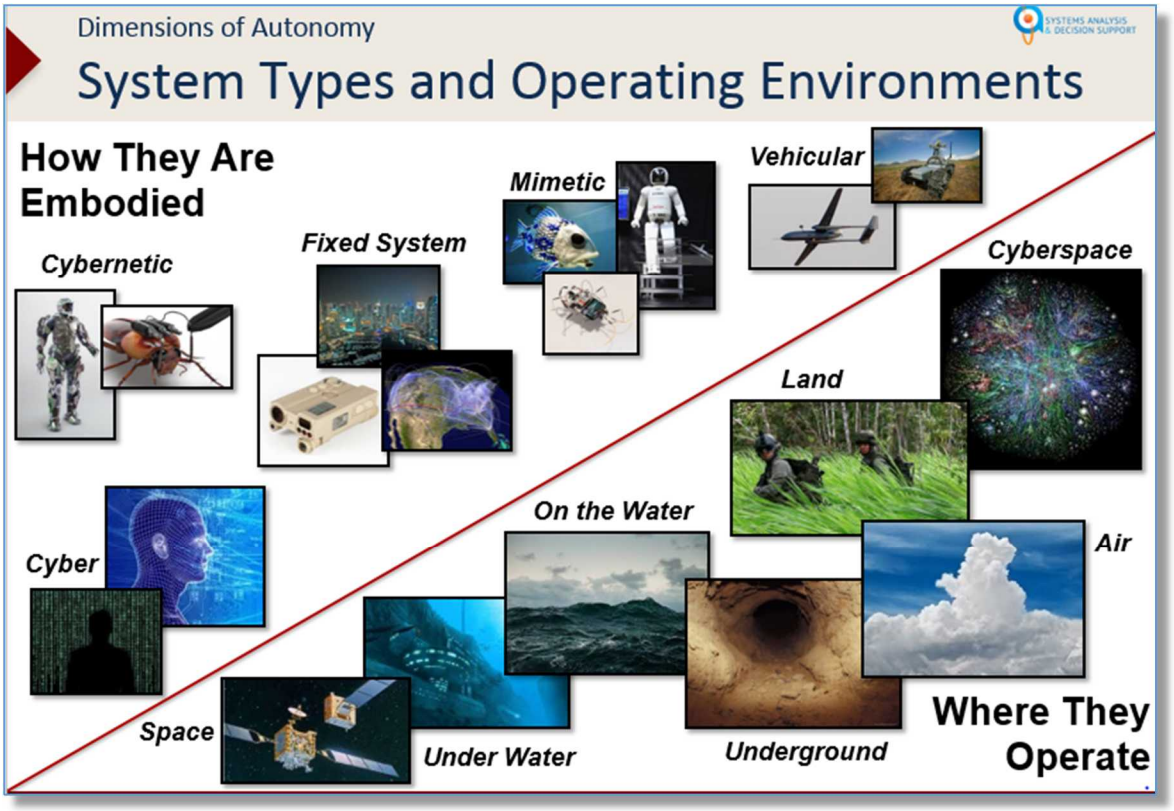

Figure 5 Autonomous System Types and Operating Environments

Within each of these areas, operational considerations – law and ethics, treaties and norms, assessments of new threats and vulnerabilities, surety and trust, policies and doctrine, new techniques and procedures – must be addressed from a systems perspective for a full understanding of potential national security impact. As we develop these applications, we must be prepared to counter their development by potential adversaries and anticipate unintended consequences, such as the impact of widespread adoption of autonomous systems on the economy and accompanying increases in (and perceptions of) social injustice that fuel political stabilities. Management of the transition to wide-scale adoption of autonomy will be a primary consideration for national security in the future.

AUTONOMY AND COMPLEXITY – WHAT ARE THE STATE OF THE ART, OPPORTUNITIES, CHALLENGES, AND RISKS?

Academic Alliance partners presented high-level overviews of work in autonomy and complexity at each of their institutions, followed by a panel discussion of the state of the art, opportunities, challenges, and risks in autonomy, considered through the lens of complex systems research today, tomorrow, and in the future. Group reflection, questions, and discussion followed the presentations.

Perspectives from Purdue University (Professor Dan DeLaurentis)

Purdue University Research Overview

Recognizing the inherent interdisciplinary nature of autonomy, Purdue University has organized teams from across the university for R&D in autonomy around the capability needs in the framework presented in the US Air Force report, *Autonomous Horizons*[i] (Figure 6). Each research team targets applications identified by consumers, particularly in the US defense community.

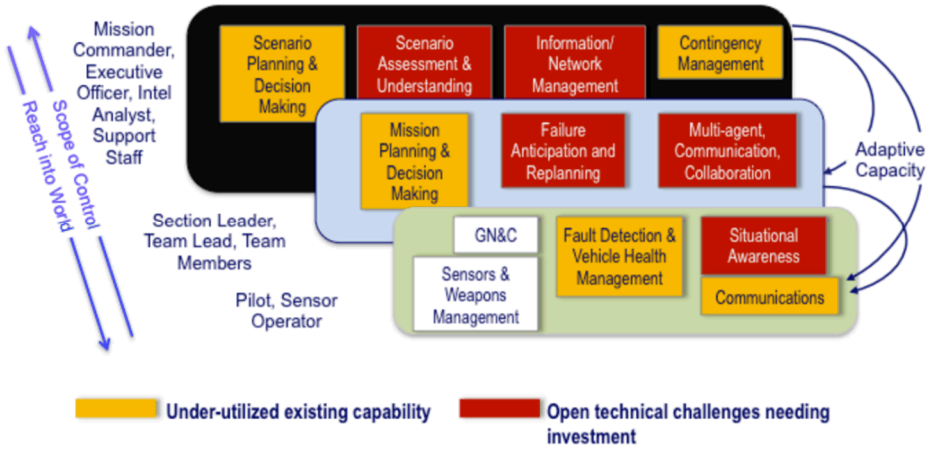

Figure 6 USAF Framework for Insertion of Autonomy into Aircraft Systems (from Autonomous Horizons report)

An example of intersection of the autonomy research at Purdue with complexity is the work of Professor Hwang on reachability analysis for control systems.[ii]

State of the Art, Opportunities, Challenges, and Risks

We need system-of-systems (SOS) thinking that goes beyond data fusion and machine learning. How much do you expand the boundary of the system to think about all the aspects that will have an impact on meeting the ultimate system objective function? For example, the tactical mission success of unmanned aerial vehicle (UAV) strikes by the United States in Pakistan may have unintended consequences in the long term with respect to strategic goals of the US mission in Pakistan. To achieve this SOS thinking, we need to learn others' language and constructs in this interdisciplinary world, including constraints and bounds.

Perspectives from University of New Mexico (Professor Meeko Oishi)

UNM Research Overview

Professor Oishi presented an overview of her work in hybrid systems, while recommending that Sandia explore a broader representation of work in autonomy and complexity at UNM that includes her colleagues, Professors Stephanie Forrest (computer science) and Melanie Moses (computer science, biology) among many others. Dr. Oishi's work focuses on methods to improve verification of large complex systems, formally and explicitly incorporating humans-in-the-loop through probabilistic and stochastic reachability methods. The work relies on constrained optimization methods to create optimized human-machine interfaces.

State of the Art, Opportunities, Challenges, and Risks

The biggest challenges in autonomy are that the methods we have designed often assume that (i) there is no human in the loop, and (ii) things work perfectly all the time. We need to come to terms with the error of these assumptions. Human-centered design at a control level, as well as in an architectural sense, needs to be friendly to humans interacting with it for the system to be reliable and operate as expected. There are many hard problems in this area.

Perspectives from University of Illinois Urbana-Champaign (UIUC)

UIUC Research Overview (Professor Geir Dullerud)

University of Illinois is conducting research and development in autonomy at many different levels from local to global scales and across multiple physical domains. They are concerned with technology transition as well as basic research. For example, they have deployed control systems for aeronautic applications of NASA and Edwards Air Force Base. They work with the Navy using game theory to support large-scale operations using autonomy for intruder/defender exercises conducted in Chesapeake Bay. Other research applications include cooperative robots (both aerial- and ground-based using indoor GPS); human-machine learning and coexistence; cyber-secure autonomous systems; deep, multilayer neural nets for machine learning (with emphasis on determining required network sizes, causal inference and hidden time series). Intersections with complexity include graph theory and networks for machine learning, and statistical verification of hybrid systems (e.g., physically connected to software through discrete algorithms) through research in probabilistic model checking using complex specification of Markov processes. The latter generates possibility space of autonomous system components and rules for guaranteeing "correct" behavior.

State of the Art, Opportunities, Challenges, and Risks (Professor Naira Hovakimyan)

Autonomy is different depending on whether it is at the sensory/motor level, reactive level, or cognitive level. A significant amount of research has been done at each level, but much of it independently. The challenge is having a framework that brings together these interdisciplinary advances for human/machine/systems that think and act as desired. In developing that framework, we need common understanding and language for what we want "SMART" autonomy to do. For example, DoD wants to reduce pilot stress and burden to humans.

Ethics, laws, and social implications are critical to the processes and systems for autonomy. Some of the many questions are

- If a robot makes a mistake, who is responsible for it?
- When do autonomous systems have rights and responsibilities (e.g., self-driving cars on the road)? What should the legal status of autonomous systems be? How do those affect the design?
- What are the social and economic impacts of autonomy? Are some jobs too important for robots to take over? [iii] Are we providing the right training for the jobs that will be available?[iv] The way you think about a technology when under development is different than the way you will think about it after deployed, because it will have shaped the environment. For example, have we considered the lifecycle infrastructure needs and costs of unmanned aerial systems (equipment, services, maintenance)? What are the possible disruptions to economic/security/social/infrastructure SOS built on autonomous capabilities?
- We need to better integrate the intersection of psychology and engineering interests and research needs. We need a place to test these technologies in realistic human-machine-social settings. For example, are we ready (societally) for drones to be doing tasks for us controlled from personal iPhones? Some questions to consider are
    - How do we physically design robots with social considerations, such as materials, noise levels, efficiencies?
    - How do we protect robots from hackers, at what levels and under what scenarios? Who is responsible for that protection? What are the trade-offs between security and performance requirements and who decides?
      Virtual reality can be used to test some of these for level of human arousal in the presence of new technologies.

The bottom line is that we **need a multidisciplinary research approach cutting across ethics, social science, psychology, economics, and engineering to face problems that will arise when autonomy becomes part of daily lives.**

Perspectives of Georgia Tech University (Professor Fumin Zhang)

Georgia Tech Research Overview

In 2013, Georgia Tech formed the Institute for Robotics and Intelligent Machines (IRIM) – an interdisciplinary research program bringing together faculty and resources from the colleges of computing, engineering, and science. IRIM offers an interdisciplinary PhD in robotics with approximately 25 graduate students per year, as well as undergraduate summer programs and robotic training programs. IRIM currently has approximately $30M of federal research funding, supporting more than 30 labs and 70 researchers. They have some industry partners but would like more. Core research areas include mechanics, control systems, perception, artificial intelligence and autonomy, and human-machine interactions. These research areas focus on a variety of applications, with a key strength being around marine robotics for both the US Government (USG) and private customers. IRIM industry partnerships have resulted in startup tech companies as spin-offs from the academic program.

State of the Art, Opportunities, Challenges, and Risks

Converting data streams in autonomy, and having to transmit data across platforms are two of the big issues. Autonomy is a way to handle complexity, but we tend to make it more complex by

working in silos, and then must transmit data between parts of the system. That is, we are not taking an SOS approach in the development of autonomous systems. We end up with tremendous amounts of temporally and spatially correlated data that may be 40 layers deep. How do you make this data flow in real time and extract meaningful information? This is a cross-cutting problem across all domains of application of autonomy. A similar challenge is, how do you remove humans from the loop so you can compress data more meaningfully, quickly, and efficiently? Humans may be one nodule in "huge loop." How do we put these together?

Perspectives of University of Texas at Austin (Professor Ufuk Topcu)

University of Texas (UT) Research Overview

The relevant programs in autonomy and complexity research at UT are vast. Robotics is one of the interdisciplinary "Portfolio Programs" at UT, which allows students to obtain cross-disciplinary credits while completing graduate study in a discipline.[v] This allows students from across the university in different areas to explore applications and research in robotics for their fields, fostering collaborative interactions and networking across diverse disciplines, such as engineering and behavioral psychology. A noteworthy facility is the human-centered robotics laboratory. Professor Topcu's research is in dynamically changing environments with unforeseen conditions, such as imperfect perception. The work requires intersection of control theory, machine learning, and formal methods (software interacting with physics) integrated into infrastructures at scale in collaboration with humans. Sponsors are the Defense Advanced Research Projects Agency, the Office of Naval Research, the Air Force Research Laboratory, the National Science Foundation, the Air Force Office of Scientific Research, and the Jet Propulsion Laboratory.

State of the Art, Opportunities, Challenges, and Risks

Observations from interactions with DoD and NASA on autonomy are the following:

- DoD still needs to do low-level programming. Their aircraft can fly autonomously but cannot park themselves on the ground.
- Blindly pushing humans out of the loop results in limited acceptance and surprise, with limited gains in manpower efficiency (which is one of the primary operational pulls for DoD).
- Safety is critical but impossible without exhaustive testing (as evidenced by the accident report on the failure of Google self-driving car).
- As we are more successful in limited tasks, some of the safety, criticality and manpower efficiency, and overall sustainability problems will become greater. This begs the question, "How can we build **affordable** and **trusted** systems?" How can autonomous networked systems enable new business? This question is hard because the systems are heterogeneous, going from vehicle actuation to path planner to traffic planner to mission planner. Each planner uses different math and analysis frameworks, creating **communication challenges** that need to be resolved with a **SOS** solution.
- There will always be uncertainties and risk of failure. Protecting against, attributing, and correcting faults will be difficult.

- Autonomous systems are often distributed but integrated. The problems are at an integrated systems level, outside of any one expert's comfort zone, but there is something for "almost everyone" in solving the problem. How do you ensure that solutions from different subdomains work with the solution methods used on other subdomains? From a control perspective, the problem is not within the subdomains but between the subdomains. An example is the use of game theory versus control theory for decision making in different subdomains.
- We need adaption through learning without *a priori* knowledge about the environment.
- Joint control and learning under temporal logic specifications offer possible solution approaches to these challenges. See research in this area at UT.

Discussion

Discussion questions

- At Sandia, there is conundrum between focus on architectures and frameworks versus specific instantiation of application to put on a platform. Where is the national security need? Where is the payoff? Are they congruent?
- Greatest risk is in early adoption of autonomy. How do we handle that to get over the risk equation early on, especially in the national security environment?
- How do we address policy challenges? The United States has highest deaths by fire in the developed world. Insurance companies cover tails but we do not have regulations on middle levels of risk. Can our society afford the same approach with respect to autonomous systems and what are the implications for design?
- Can we assume a US-based jobs and skill base or will the largest job source for building autonomous systems be in developing countries? If so, can we assume that they will have the training capability to be sure that they do it right? They are in race with time to develop those capabilities, and will not have the benefit of industrial revolution development. What are the implications for safety, security, trust?
- What are appropriate areas for a Federally Funded Research and Development Center (FFRDC) such as Sandia to provide leadership on issues in autonomy important to the USG and the national interest? How can we be most effective?

Panel Responses

- Development of autonomous systems (by the USG) should strive for modular and "pluggable," with appropriate software and hardware that is interoperable. This has implications for safety and security, which would have to be designed and regulated to the component level at every step, as well as SOS level. This is an area where complexity research can be applied to understand at what level in a system problems will first present themselves.
- In the initial R&D stages, we do not tend to worry about these "big" system SOS level questions. However, after some success in deploying autonomous systems, these questions will be very important. Leadership is required early on to get ahead of these problems that will surely arise. There will be competition among institutions for leadership. Leadership in other emerging technology areas – such as robotic operating systems – has gone to those who "do the dirty work," building institutional capacity for research in all areas that were important and sharing that research broadly in open fora.

Success in early projects is also critical for establishing credibility as a leader on the national stage to include industry.

- The research endeavor in autonomy is itself a complex adaptive system. The community is in the process of discovering the right questions and objectives. We want to maximally glean knowledge as we go. Universities have the context and organizational structure to do generic, basic research – the "dirty work" – with a freedom that the national laboratories lack. This is a challenge that Sandia must address to be a leader.
- The transference of tasks and capabilities from humans to autonomous systems needs to be thought through at the social level. For example, we need to think about what humans need to feel safe in environments involving autonomous systems, and the role of the federal government and policy. What do those look like when there are mixed systems of humans and drones in the sky and on the roads?
- Commercially, all else being equal, whoever emerges as the technical leader in autonomy will set the rules of the game. See Google as an example. This is an area where the national labs need to assert themselves in the public interest, especially in areas where industry cannot test *a priori*. There is a current opportunity to survey people in neighborhoods in Austin where there is autonomous car testing versus other neighborhoods where there is not – who trusts these cars? Is there an impact due to familiarity? Who will sign off that these systems are safe? Sandia could have a role in certification that is useful "in the national interest."
- We do not know the impact of autonomous systems on economies yet. They could be like cell phone services in Africa – where they do not replace jobs but create opportunities for new ones. Understanding how cell phones impact economy and society of Africa might be an informative analogy for how autonomous systems may impact future societies/economies, especially in the developing world.

AUTONOMY AND COMPLEXITY – WHAT ARE THE RESEARCH PARTNERSHIP OPPORTUNITIES?

Three research areas for breakout group discussions on potential research partnership opportunities were identified in advance of the workshop by subject matter experts and program managers, considering results of the previously held workshops at Sandia and the expertise and research interests of the Academic Alliance partners. These topical areas were (i) Human-Machine Teaming, (ii) Distributed Control and Cooperative Systems, and (iii) Robust, Reliable, and Trusted systems. Participants in each breakout group were asked to address the common core questions listed below, as well as questions specific to each group that follow.

See Appendix C for participants in each group.

Common Core Discussion Questions for Breakout Groups

- What are the research gaps in autonomy for national security?
- What do complexity science and systems engineering bring for addressing those gaps?
- What are the relative strengths that academia, FFRDCs, such as the government owned national laboratories, and industry each bring to addressing research gaps?

Human Machine Teaming (HMT) Breakout Group

HMT Specific Discussion Questions

- What are the goals of the mission within the application space (i.e., physical security/ situation awareness at remote sites, collaborative systems in motion, collaborative control of large-scale critical systems)? Are the goals of the humans similar/different than those of the machines? If so, how?
- What machines are involved, where are they, what are they doing, when?
- Who are the people, where are they, what are they doing, when? How might that change dynamically? How does this inform systems-level design?
- How do we measure quality and effectiveness of system configuration, communication, situational awareness, learning? In what environments?
- How can science inform interface design? How can we measure interface effectiveness in ambiguous conditions (i.e., when there is no ground truth, when there is a large amount of uncertainty)?
- How can we design and maintain effective transitions from human to machine while maintaining transparency and explainability?

Discussion

Research Gaps and Challenges

- Bridging theoretical gaps during application: There are lots of theories about how people make decisions under certain circumstances. These paradigms do not always hold up in operation environments. How do we bridge the theoretical gap between the lab and the real world?

- New problems with no previous literature/data available: We are formulating new types of problems, applying new applications that have never been seen before. An example is in the personalized use of drones. There is a difference between real safety in terms of collision avoidance and perceived safety of the human. Variables such as size, distance, and velocity of the drone affect how safe the human feels. There are also research gaps in terms of collecting data in new systems such as virtual reality. Do we need more sensors or not? Sensors currently being used include head tilt, heart rate, and skin conductors.
- Missing guidelines and well-developed theory for human-centered design: There is a gap for people that are designing human-centered systems. There are no guidelines for designing new systems with which humans interface. In the context of the TSA, individual differences between people are the biggest explanation for differences in data. There is nothing so far that can predict the variability between individuals. In some cases, these differences can be a great thing to look at and help us. Other times they do not help at all. In cognitive psychology, converging theories are those that happen across multiple paradigms and are not idiosyncratic to a certain paradigm or context.
- No consolidation of data/expertise from different domains (exacerbated by duplication of research efforts): There are no resources or manpower to bring all the data together. A key need is to **identify appropriate state variables**. We do not want to find some ephemeral effect and mistake it for a robust one. Can we come up with common state variables such as situational awareness? Can we make interfaces adaptable to individuals based on how each individual reacts?
- Fear and lack of understanding of vulnerabilities introduced by autonomous systems and possible exploitation: Autonomous army vehicles carrying gear for a team of soldiers – Could these systems be hijacked to hurt us? Could they be disabled? Could someone else make hand motions to make the machine run you over?
- Impacts of the way in which humans operate in different contexts: The Army is interested in a system with a leader-follower kind of relationship. However, in combat, humans may not speak calmly to a system to issue clear voice commands. They could be shouting in a war zone and making lots of hand motions, using nonverbal body language. Alternatively, is there a situation where they cannot speak, because they are on a silent watch?
- Lack of trust in autonomous systems:
- Interface design challenges: It is important to minimize cost and testing of system. The military wants it implemented quickly. However, testing is important for assessing if system goals are met. Does the system do the job it is supposed to do? It is important to make something work in practical fashion before we optimize it to be the best it possibly can be.
- Risk quantification & analysis: What is the risk of putting a machine in a situation versus a human? Right now, we cannot strap enough people into virtual reality helmets to get enough data to be considered reliable. Could we use simulated data? What is the credibility of simulated data? Should risk taking be judged based on number of humans involved? How do we assess relative consequences, depending on number of humans involved and the context? Will humans and machines make the same risk judgements?
- Virtual environments to test systems: To quantify risk, we need a reliable and accurate way to simulate environment to test autonomous systems. Where is the "wind tunnel" in

this industry? Sometimes there are no methods in place to really test autonomous systems and simulate their environment. Do we need a "wind tunnel" for the systems or humans or both? What about virtual spaces where multiple people interact with each other and with the system?

- Repeatable/testable/verifiable/safe testing and protocols and environments: Regarding sample space, are we trying to detect an effect that is so miniscule that you need a sample space of a million people to detect it? Or do you just need a handful of people to get a robust baseline response? In the aerospace industry, they were invested in safety because otherwise the public would not even consider getting on a plane if they did not perceive it as safe. Public perception is important! A drone flying too close to someone will make them feel uneasy. Do soldiers feel safe with a robot following them? Do people feel safe when they find autonomous google cars on the road?
- Machine Learning: What is the right input data for machine learning? We need enough data, but we do not want to overwhelm the system. How do we deal with noise and complex environments likely to be encountered during learning in operational environments?
- Operator & human unpredictability in off-normal situations: In terms of grid, things can be tuned to day-to-day operations, but when there is a hurricane, these systems might need to adapt and change. We need to understand better how to "tune" for resilient situations and improve efficiency of unsupervised learning. Can we come up with paradigms of unsupervised learning to improve its efficiency and resiliency so that the lessons hold for different data sets and contexts, especially off normal or unanticipated situations?

R&D Partnership Opportunities

- Address data issues (UNM):
    - Labeled data sets that we can collaborate on and tell the system when it is successful in analyzing them.
    - Collaborate on keeping humans in the loops (human automation alignment methods) with autonomous systems: How will the human operators act when the humans lose situational awareness? How do we bring the humans back into the loop?
- Develop science-based design principles for human-centered automations (UNM)
- Adaptable control in resilient situations
    - How do we characterize what humans are going to attend to in relation to what the autonomous system is doing around them?
    - There is no research in truly operational environments, but there is lots of research in attention awareness.
- Bridge gaps from basic theory to operational environments
    - Example: dynamic distribution of tasks has implications with human-machine teams, machine-machine teams, and human-human teams.
- Uncertainty analysis and display at human-machine interface

    - Example: An autopilot alerts the operator and says, "Hey, I've never flown over this type of stuff so watch me a little closer."
- Human-in-the-loop testing
    - How do we get that human into the mix? Right now, we focus on all the hardware and everything. We need to remember that the human is part of this system of systems and that they are, in fact, the most complex part.
    - Creating an environment where we can explore the science more and collect data.
- Create testing and evaluation standards and facilities (What is our wind tunnel?)
    - We need to be careful about building this wind tunnel because certain virtual environments could cause stress and trauma to human testers.
- Identify and bring together the right communities to answer autonomy questions

Summary

We need formal specifications (e.g., metrics) regarding human-machine teaming for emerging systems. These must consider standards that allow interoperability, the different types of needs for getting information out of the system and integration of different human perspectives in the system. The human machine teaming must be adaptive, with guidance points for V&V to occur in real time. We need formal design and testing models that integrate human engagement and influences into the behavior of the system as a whole. This fits into constrained strategy class of modeling problems. Such models need to be able to bound real system performance in different operating environments in order to understand how the human machine teaming affects performance of the system as a whole. We need to make V&V methods for human machine teaming scalable, including virtual environments.

Robust, Reliable, and Trusted Systems (RRTS) Breakout Group

RRTS Specific Discussion Questions

- How do we measure "trust" status, including: health, integrity, operational space validity? In what environments?
- What are the effects of complexity and scale?
- Data issues: how do we establish and maintain necessary quality, quantity, precision, representativeness, as well as integrity in planning, training, and operations? How do we determine/assess situation-dependent "good enough" for decision making? How do we separate signal from noise in different environments?
- How does design and implementation of confidence levels, transparency, explainability change to account for higher consequence and risk, including ethics and policy issues (e.g., increasing lethality of system)?
- How do we determine in real time what sensor information adds value and where (upstream for sensor reconfiguration and targets versus downstream for analysis)?
- How do we use and aggregate confidence levels for single, multi-modal sensors/platforms for environments where distributed, multiple sensor systems are not possible?

Discussion

Research Gaps and Challenges

- We cannot observe everything. What kind of inferences can we draw, even if we do not learn the entire system? How do we maintain or communicate validation of trust of the system knowing that we are limited? It takes at least a partial understanding of how complex systems work. While partial observations have been studied for 60 years, the problem becomes more complicated when they involve hyper dynamics and hypotheses may not be falsifiable. You must quantify the uncertainty through modeling, simulation, and appropriate experimental design in some domains where there are no observations. What can be learned/deduced about robustness from interventions? How do we account for noise and covariance, and assure that the testing environment for the system translates to military or other application environments (e.g., transfer learning/knowledge across different situations using a common set of variables that do not change)?
- How is trust measured in different contexts? For human systems, can use "scorecards." How do we account for irrational nature of human trust? For machines and models, we can use verification of mathematics and logic. How do the two combine? Complex systems can never be truly verified. How do we use distribution measurements of trust? Confidence measurements? How do you know when a machine is learning or unlearning?
- How can we link control policy and trust? How do we account for different decision timescales in this linkage?

R&D Partnership Opportunities

- Start with less complex systems and focus on means for improving robustness, reliability, and trust in automated systems. Think about formal specifications for learning systems and standards that allow interoperability (UT).
- For hybrid human-machine systems, collaborate to provide different perspectives of human interaction – both the users and the designers. Develop a model that integrates these perspectives and helps to characterize a system's behavior (GT).
- Partner to gain more operational and experiential knowledge and confidence with designed systems (UIUC).
- Develop formal models of systems that include strategies of users and designers, then analyze strategies (including constraints and restrictions on strategies).
- Push component testing as far as possible and develop better understanding of scalability (Sandia).
- Collaboration mechanisms require teamwork, and might be facilitated by open call for and publications of ideas, use of student interns.

Summary

We need formal specifications, standards, and metrics for emerging systems that will facilitate interoperability. We need to address the R&D challenge of partial observations and how to get information required for establishing robust, reliable, and trusted systems and networks. We need to give priority to thinking about the user in the autonomous system, creating models to integrate the various human perspectives in the system for design, testing and evaluation. We

need adaptive or iterative, scalable verification and validation methods that learn as they go with a process that creates guidance points, and that are compatible with discrete connected algorithms. We need formal models that integrate human engagement for inferences into behavior-constrained strategy classes. We need capabilities to test and bound real system performance in different operating environments.

We need to work together on domain-specific problems where Sandia can provide value to the academic community as an interface with the user community for funding tied to national security applications. Collaborations can be facilitated by open sharing of research needs and ideas, use of student interns, and faculty sabbaticals.

Distributed Control and Cooperative Systems (DCCS) Breakout Group

DCCS Specific Discussion Questions

- How do we measure quality and effectiveness of system configuration, communication, situational awareness, learning? In what environments?
- How do we formalize explainability, confidence building, transparency, reliability, vulnerability of adaptive algorithms?
- How do we enforce and maintain transparency and "validation" of assumptions in design and training in real-time?
- How does machine learning/distributed control scale across time, physical space, complexity of system? How do we integrate machine learning into cooperative systems? Into control systems?
- Confidence in data based on compilation of ensemble of sensors versus individual sensor data – how/when can resolution requirements be decreased if data are going to be used in ensemble? How does that depend on distribution and level of cooperation between sensors?

Discussion

Research Gaps and Challenges

- Common language: Confusion around terms such as "control," 'adaptation," "automation," and "autonomy" among R&D community is a challenge and can lead to confusion, especially in distributed systems across multiple platforms and levels of cooperation.
- Communication in Automated and Autonomous systems: Learn from natural (animal and biological) systems such as fish, who have limited but effective communication capabilities based on judgment, is nonlinear (is amplified above certain thresholds), and makes effective use of local versus distributed communication systems and learn to eliminate what does not work. This is different than the way that engineers tend to think about communication systems. Can also use biological analogues, such as organisms and adaptive immune systems. How many levels of surveillance are necessary, when is it best to use swarming behavior? How do we accomplish data fusion?
- Appropriate level of complexity in control systems: When do we need simple simulations and when are multiple layers better?

R&D Partnership Opportunities

- Bio-inspired distributed control systems integrated with cognitive systems (GT)
- Intent-based, goal-driven optimal control through adaptable systems that can be bottom-up or top-down (GT)
- Dynamic machine learning – for example, algorithms that an individual UAV or machine can work by itself in a swarm to include new information, game theory (UIUC)
- Accounting for adversary machine learning, adversarial neural networks, and impacts on security (UIUC)
- Understanding modes of system failures and resilience to those failures
- Understanding what degree (level) of autonomy is desired and the appropriate control and types of communication systems at the human interface (including alternative communication paradigms, effective communication in rapidly changing environments)
- Design optimization for adaptable systems in changing environments, with changing goals as the environments change

Summary

Distributed control and cooperative systems in autonomy may have different goals. For example, the goal of a self-driving car is not to crash, while the goal of swarming fish is to achieve something together. These require different types of adaptive, distributed control and communication systems with flexible design optimization.

Distributed control and communication at the human interface is a critical area for collaboration. Humans will be distributed in different locations and roles in autonomous systems, acting with incomplete information, limited or no communication and situational awareness and relying on machine learning to supplement gaps. We need to engineer resilient control systems drawing on biological inspiration where appropriate, depending on the level of autonomy and application environment (e.g., decision timescales, adversarial systems).

Common Themes, Follow-Up Questions, Key Takeaways

In the final session of the day, attendees reflected on what they had heard and shared observations on common themes, unanswered questions, and key takeaways.

Common Themes

Autonomous systems and the environments in which they are deployed are characterized by rapidly changing environments, multiple levels and scales of decision making, and both bottom-up and top-down control systems among distributed entities that are a fluid mix of humans and machines. Governing principles for these complex adaptive systems create constraints on traditional engineering approaches and insights on how to leverage knowledge from across other disciplines that include human cognitive and social systems, as well as biological systems. R&D success requires teaming across multidisciplinary fields and SOS thinking and evaluation approaches. Success at the national level requires leadership that reaches outside of engineered systems considerations and includes policy, ethical, and legal issues.

These themes, as well as the specific topics that emerged from the breakout groups, reflect the strategies articulated in the National Artificial Intelligence R&D Strategic Plan[vi], which are

- Make long-term investments in AI research
- Develop effective methods for human-AI collaboration
- Understand and address the ethical, legal, and societal implications of AI
- Ensure the safety and security of AI systems
- Develop shared public datasets and environment for AI training and testing
- Measure and evaluate AI technologies through standards and benchmarks
- Better understand the national AI R&D workforce needs.

Follow-up Questions

- How do we take these ideas and turn them into real actionable projects to take forward?
- How can we maintain these connections to support collaboration?
- Duplication of efforts across different fields/disciplines is bad! How do we connect them?
- How do we understand how much risk we are willing to take in designing systems and making assumptions?

Key Takeaways

- There are many people across all research disciplines that are interested in autonomy. Collaboration and communication forums with long-term vision are critical for a robust and effective R&D community in support of national security.
- There are many areas of common interest and potential collaboration between the Academic Alliance partners and Sandia. However, the R&D community is still formulating the right questions and roles for industry, academia, and government researchers.
- While SOS thinking and approaches are required, they are not necessarily supported at the level and modes of funding mechanisms provided for R&D.
- We're making machines do more and more and yet sometimes the human element that is integrated into that system is ignored. Not all problems are technical in nature.

NEXT STEPS

Campus Leads
Campus lead managers from Sandia are available as the first point of contact to facilitate communication between Academic Alliance partners and Sandia researchers for follow-on ideas and opportunities. These managers are

- UIUC – Russ Skocypec
- Georgia Tech – Rebecca Horton
- UT – Amanda Dodd
- Purdue University – Bill Hart
- UNM – Carol Adkins/Randy Shunk

LDRD Research
Sandia is currently going through the internal process to determine Laboratory Directed Research and Development (LDRD) priorities and allocations for FY17. The Academic Alliance effort at Sandia provides additional funding that can be sent to universities to partner with on the research proposals. The Office of the Chief Technology Officer (CTO) will survey the LDRD calls of possible autonomy related opportunities to share with universities.

Communications and Sharing Venues

- Sandia will be documenting results of workshops in reports on autonomy for sharing with internal audiences and external partners. We will create user-friendly "Ted talks" on findings to inspire creative, collaborative, interdisciplinary approaches to the challenges identified.
- Sandia will organize internal and external collaborative websites to identify research interests, teams, and publications.
- As an FFRDC, Sandia has a responsibility to help and advise the USG on issues important to national security. To that end, Sandia is in the formative planning stages for bringing together the external community in the September 2017 timeframe to engage on strategic issues of autonomy in the national interest. We are actively seeking input for key themes, stakeholders, and participants that represent thought leaders across diverse intellectual and policy frames of reference.

Creating Vision for Longevity and Impact
Academics need a long-term vision for research partnerships to have real impact. While 5 years is an important timeframe for achieving tenure, a typical timeframe from basic research idea to real impact on national security and social systems is 10 years. Strong institutional relationships between Academic Alliance partners and Sandia require continuity to the collaboration we start that transcends changes in people and funding agencies. This is especially important in disruptive technologies (e.g., autonomy). We should follow up on the idea to have a symposium at Sandia with the Vice-Presidents of research from the Academic Alliance schools around the technical and policy issues that have been identified in the workshop. This symposium should include a broad representation of relevant academic capabilities, including the policy and social science schools in addition to computer science, engineering, and human factors.

APPENDIX A: WORKSHOP PARTICIPANTS

Name	Organization	Email
Carol Adkins	06100	cladkin@sandia.gov
Phil Bennett	01463	pcbenne@sandia.gov
Diana Bull	00159	dlbull@sandia.gov
Ed Carroll	00158	ercarro@sandia.gov
Kerstan Cole	00431	kscole@sandia.gov
Ben Cook	01910	bkcook@sandia.gov
Richard Craft	00158	rlcraft@sandia.gov
Dan DeLaurentis	Purdue	ddelaure@purdue.edu
Lee DeVille	University of Illinois	rdeville@illinois.edu
Amanda Dodd	01914	ajbarra@sandia.gov
Geir Dullerud	University of Illinois	dullerud@illinois.edu
John Feddema	01460	jtfedde@sandia.gov
Pat Finley	06131	pdfinle@sandia.gov
Meghan Galiardi	SNL	mgaliar@sandia.gov
Jared Gearhart	06131	jlgearh@sandia.gov
Bill Hart	01913	wehart@sandia.gov
Nancy Hayden	00159	nkhayde@sandia.gov
Stephen Henry	06133	smhenry@sandia.gov
Matt Hoffman	06133	mjhoffm@sandia.gov
Rossitza Homan	06921	rhoman@sandia.gov
Marcey Hoover	06130	mlhoove@sandia.gov
Rebecca Horton	01900	rdhorto@sandia.gov
Naira Hovakimyan	University of Illinois	nhovakim@illinois.edu
Bobby Jeffers	06921	rfjeffe@sandia.gov
Dean Jones	06131	dajones@sandia.gov
Elizabeth Keller	00159	ejkisti@sandia.gov
Negar Kiyavash	University of Illinois	kiyavash@illinois.edu
Steve Kleban	06132	sdkleba@sandia.gov
Anne Lilje	06132	alilje@sandia.gov
Andrew Lucero	10619	aflucer@sandia.gov
Jerry McNeish	08954	jmcneis@sandia.gov
Alan Nanco	06114	asnanco@sandia.gov
Meeko Oishi	UNM	oishi@unm.edu
Sasha Outkin	06921	avoutki@sandia.gov
Alex Roesler	05440	awroesl@sandia.gov
Elizabeth Roll	00100	earoll@sandia.gov
John Rowe	5000	jcrowe@sandia.gov
Jon Salton	06533	jsalton@sandia.gov
Judi See	00151	jesee@sandia.gov
Russ Skocypec	00150	rdskocy@sandia.gov
Ann Speed	01462	aespeed@sandia.gov
Rayaduranm Srikant	University of Illinois	rsrikant@illinois.edu
Kevin Stamber	06132	alilje@sandia.gov

Name	Organization	Email
Laura Swiler	01441	lpswile@sandia.gov
Shawn Taylor	05629	setaylo@sandia.gov
Bruce Thompson	06133	bmthomp@sandia.gov
Ufuk Topcu	University of Texas	utopcu@utexas.edu
Tim Trucano	01400	tgtruca@sandia.gov
Stephen Verzi	06132	sjverzi@sandia.gov
Eric Vugrin	06613	edvugri@sandia.gov
Fumin Zhang	Georgia Tech	fumin@gatech.edu

APPENDIX B. WORKSHOP AGENDA

Academic Alliance Workshop on Autonomy and Complex Systems

Friday, March 3, 2017

Purpose	Engage Academic Alliance partners in autonomy efforts at Sandia by discussing research needs and synergistic areas of interest within the complex systems/systems modeling domain, and identify opportunities for partnering on laboratory directed and other joint research opportunities.
Desired Outcomes	A summary report that outlines common areas of research and ideas for joint proposals.
Classification	Presentations/Discussions – Official Use Only
Participants	Participants from previous internal Sandia workshops, Sandia Academic Alliance and CTO reps, Sandia Complex Systems Research Challenge and Systems Modeling Community, and Academic Alliance partners from the complex systems and systems modeling communities at the Academic Alliance universities.

Sandia NM: CSRI/90

	Topic	Presenter/Moderator
7:30	Academic Alliance Visitors MUST Meet at Badge Office	Amanda Wilson
8:00	Continental Breakfast and Networking	
8:30	Welcome	Carol Adkins & Marcy Hoover
8:45	Introductions	University Reps
9:00	Executive Comments on Autonomy & National Security	Russ Skocypec
9:15	Autonomy at Sandia	Alex Roesler
10:15	Break/Transition to Panel	
10:30	Broadening Perspectives on Autonomy/Intro to Panel Session	Jon Salton
10:45	Panel Session "The Role of Complex Systems Research in Autonomy: Today, Tomorrow, and the Future"	Bill Hart - Moderator Dan DeLaurentis – Purdue Meeko Oishi – UNM Geir Dullerud – UIUC Ufuk Topcu – UT Fumin Zhang - GT
12:00	Lunch and Networking	Catered
1:00	Breakout instructions	Nancy Hayden
1:15	Breakouts • Human Machine Teaming - Room 147 • Distributed Control and Cooperative Systems – Room 148 • Robust, Reliable, and Trusted Systems – Room 137	 Rodriguez/Cole Schwartzwald/Vugrin Roll/Verzi
2:45	Break & Reconvene	
3:00	Breakout Sessions Readouts	Nancy Hayden
3:30	Common Research Themes Discussion	Nancy Hayden/Ben Cook
4:00	Adjourn	Marcy Hoover

APPENDIX C. BREAKOUT GROUP PARTICIPANTS

Robust, Reliable, and Trusted Systems:

Name	Organization	Email
Rebecca Horton	01900	rdhorto@sandia.gov
Jerry McNeish	08954	jmcneis@sandia.gov
Sasha Outkin	06921	avoutki@sandia.gov
John Rowe	5000	jcrowe@sandia.gov
Tim Trucano	01400	tgtruca@sandia.gov
Geir Dullerud	University of Illinois	dullerud@illinois.edu
Laura Swiler	01441	lpswile@sandia.gov
Ufuk Topcu	University of Texas	utopcu@utexas.edu
Elizabeth Roll	00100	earoll@sandia.gov
Stephen Verzi	06132	sjverzi@sandia.gov
Bruce Thompson	06133	bmthomp@sandia.gov
Negar Kiyavash	University of Illinois	kiyavash@illinois.edu

Distributed Control and Cooperative Systems:

Name	Organization	Email
Eric Vugrin	06613	edvugri@sandia.gov
Rayaduranm Srikant	University of Illinois	rsrikant@illinois.edu
Jon Salton	06533	jsalton@sandia.gov
Lee DeVille	University of Illinois	rdeville@illinois.edu
Matt Hoffman	06133	mjhoffm@sandia.gov
Pat Finley	06131	pdfinle@sandia.gov
Ed Carroll	00158	ercarro@sandia.gov
Andrew Lucero	10619	aflucer@sandia.gov
Dan DeLaurentis	Purdue	ddelaure@purdue.edu
Diana Bull	00159	dlbull@sandia.gov
Meghan Galiardi	SNL	mgaliar@sandia.gov
Jared Gearhart	06131	jlgearh@sandia.gov
Fumin Zhang	Georgia Tech	fumin@gatech.edu

Human Machine Teaming:

Name	Organization	Email
Stephen Henry	06133	smhenry@sandia.gov
Naira Hovakimyan	University of Illinois	nhovakim@illinois.edu
Bobby Jeffers	06921	rfjeffe@sandia.gov
Meeko Oishi	UNM	oishi@unm.edu
Phil Bennett	01463	pcbenne@sandia.gov
Bill Hart	01913	wehart@sandia.gov
Shawn Taylor	05629	setaylo@sandia.gov
Kerstan Cole	00431	kscole@sandia.gov
Alan Nanco	06114	asnanco@sandia.gov
Ann Speed	01462	aespeed@sandia.gov

Endnotes

[i] USAF Office of the Chief Scientist (2015) *AUTONOMOUS HORIZONS: System Autonomy in the Air Force – A Path to the Future*, AF/ST TR1501 http://www.af.mil/Portals/1/documents/SECAF/AutonomousHorizons.pdf?timestamp=1435068339702.

[ii] See https://engineering.purdue.edu/~ihwang/Publication.html.

[iii] See http://droneanalyst.com/ for research and discussions on the entire commercial unmanned aerial (drone) ecosystem, including end-users, technology vendors, service providers, and investors.

[iv] The National Academy of Sciences has initiated several studies exploring some of these questions. See for example, http://sites.nationalacademies.org/PGA/step/PGA_177034

[v] See https://robotics.utexas.edu/.

[vi] National Science and Technology Council, Networking and Information Technology Research and Development Subcommittee (2016), *"National Artificial Intelligence Research and Development Plan",* Office of Science and Technology Policy, Washington DC, October 2016.

Distribution

External Distribution

Electronic copies to:

Lee DeVille at University of Illinois (rdeville@illinois.edu)
Dan DeLaurentis at Purdue University (ddelaure@purdue.edu)
Geir Dullerud at University of Illinois (dullerud@illinois.edu)
Naira Hovakimyan at University of Illinois (nhovakim@illinois.edu)
Negar Kiyavash at University of Illinois (kiyavash@illinois.edu)
Meeko Oishi at University of New Mexico (oishi@unm.edu)
Rayaduranm Srikant at University of Illinois (rsrikant@illinois.edu)
Ufuk Topcu at University of Texas (utopcu@utexas.edu)
Fumin Zhang at Georgia Institute of Technology (fumin@gatech.edu)

1	MS 0116	Kent Meeks	(electronic copy)
1	MS 0127	Diana Bull	(electronic copy)
1	MS 0150	Russ Skocypec	(electronic copy)
1	MS 0151	Judi See	(electronic copy)
1	MS 0152	Kerstan Cole	(electronic copy)
1	MS 0158	Ed Carroll	(electronic copy)
1	MS 0158	Richard Craft	(electronic copy)
1	MS 0159	Nancy K. Hayden	(electronic copy)
1	MS 0159	Elizabeth Keller	(electronic copy)
1	MS 0159	Tom Nelson	(electronic copy)
1	MS 0351	Ben Cook	(electronic copy)
1	MS 0351	Bill Hart	(electronic copy)
1	MS 0351	Andrew Lucero	(electronic copy)
1	MS 0359	Karla Weaver	(electronic copy)
1	MS 0620	Rebecca Horton	(electronic copy)
1	MS 0672	Shawn Taylor	(electronic copy)
1	MS 0721	Carol Adkins	(electronic copy)
1	MS 0757	Eric Vugrin	(electronic copy)
1	MS 0980	John Rowe	(electronic copy)
1	MS 1002	Philip Heermann	(electronic copy)
1	MS 1003	Jon Salton	(electronic copy)
1	MS 1137	Meghan Galiardi	(electronic copy)
1	MS 1137	Bobby Jeffers	(electronic copy)
1	MS 1137	Sasha Outkin	(electronic copy)
1	MS 1137	Kevin Stamber	(electronic copy)
1	MS 1138	Rossitza Homan	(electronic copy)
1	MS 1138	Steve Kleban	(electronic copy)
1	MS 1138	Anne Lilje	(electronic copy)
1	MS 1188	Alan Nanco	(electronic copy)
1	MS 1138	Stephen Verzi	(electronic copy)
1	MS 1188	Pat Finley	(electronic copy)

1	MS 1188	Jared Gearhart	(electronic copy)
1	MS 1188	Stephen Henry	(electronic copy)
1	MS 1188	Matt Hoffman	(electronic copy)
1	MS 1188	Marcey Hoover	(electronic copy)
1	MS 1188	Dean Jones	(electronic copy)
1	MS 1188	Bruce Thompson	(electronic copy)
1	MS 1173	Alex Roesler	(electronic copy)
1	MS 1165	Elizabeth Roll	(electronic copy)
1	MS 1318	Laura Swiler	(electronic copy)
1	MS 1318	Tim Trucano	(electronic copy)
1	MS 1324	John Feddema	(electronic copy)
1	MS 1327	Phil Bennett	(electronic copy)
1	MS 1327	Ann Speed	(electronic copy)
1	MS 1327	John Wagner	(electronic copy)
1	MS 1371	Dianna Blair	(electronic copy)
1	MS 9004	Sheryl Hingorani	(electronic copy)
1	MS 9159	Jerry McNeish	(electronic copy)
1	MS 9957	Amanda Dodd	(electronic copy)
1	MS 0899	Technical Library	9536 (electronic copy)


Sandia National Laboratories