SPAWAR
Systems Center
PACIFIC

TECHNICAL REPORT 3057
April 2016

Recommendations for Future Human-Autonomy Research Investments: Sensor Fusion, Machine Learning and Resilient Control Architectures

Jamie Lukos
Benjamin Migliori
Doug Lange
John Reeder
Leah Kelley
SSC Pacific

Daniel Donavanik
Army Research Lab-Sensors & Electron Devices Directorate (ARL-SEDD)

MaryAnne Fields
Army Research Lab-Vehicle Technology Directorate (ARL-VTD)

Jessie Chen
Army Research Lab-Human Research & Engineering Directorate (ARL-HRED)

SSC Pacific
San Diego, CA 92152-5001

SSC Pacific
San Diego, California 92152-5001

G. M. Bonitz, CAPT, USN
Commanding Officer

C. A. Keeney
Executive Director

ADMINISTRATIVE INFORMATION

The work described in this report was performed by the C2 Innovation Branch (Code 53624), the Advanced IO Concepts Branch (code 56150), the Unmanned Systems Advanced Development Branch, and the Anti-Terrorism Force Protection (ATFP) Surveillance Systems Branch (71740), Space and Naval Warfare Systems Center Pacific (SSC Pacific), San Diego, CA.

Released by
Jaime Lukos, Head
ATFP Surveillance Systems
Branch

Under authority of
A.D. Ramiez, Head
Advanced Systems & Applied
Sciences Division

SB

ACKNOWLEDGMENT

Special thanks to Lieutenant Commander Max Long, the operations officer of Littoral Combat Ship Crew 103, for an impressive tour of the LCS USS Freedom and keynote address on “Trust in Automation” at the Autonomy Research Pilot Initiative (ARPI) Technical Exchange Meeting.

EXECUTIVE SUMMARY

Performers of the Autonomy Research Pilot Initiative (ARPI) came together 17–19 February 2016, at the Liberty Station Naval Training Center in San Diego to discuss current issues impeding the success of human–autonomy teaming and strategies moving forward to push though these barriers.

The meeting included three small group breakout sessions, each focused on a different control system topic, followed by interactive outbriefs of each group to the wider audience. The goals of the breakout sessions were to determine the logical next steps to advance the field of autonomy, how the ARPI program aided in creating the foundations for these future endeavors and on what issues long-term research investments should focus.

Questions provided to the groups included:

- Is there a coherent understanding of the research in the field or competing theories?
- How it is currently applied?
- How important is it for the functioning of future autonomy? Why?
- What did you learn from ARPI work?
- How does enabling this technology/methodology help the warfighter to complete his/her mission?
- If you could extend ARPI another 3 years, what would you focus on?
- If you had to incorporate this into a functional prototype in 5 years, what would you focus on?
- If you had $2M to design a 5 year research program, what would you do?-
- What should be the “after-ARPI” plan? How should we continue to progress the areas of research that showed potential?
- Where should long-term S&T investments be placed to ensure technical superiority in this field in the future?
- What skillsets are needed to pursue this line of research?
- What are the bottleneck issues preventing this technology/methodology from being applied to current systems?
- What ARPI objectives you attempted remain unsolved?
- Where will the solutions ultimately come from?
- How long will it take?

This report is a pointed summary of recommended near-term and long-term research objectives going forward as well as any other commentary regarding the trajectory of S&T in this field.

This report will be delivered to ASD(R&E) for their consideration and posted online on the DoDTechipedia ARPI page: https://www.dodtechipedia.mil/dodc/x/SQKwAQ.

After action items include a follow up meeting in September 2016 in conjunction with the Final ARPI Review to reassess these issues and evaluate ways to continue cross-service collaborative research in this critical domain.

CONTENTS

1. TOPIC AREAS

Given the interdisciplinary and intraoperative nature of the challenges underlying human–autonomy integration, many topics could have been discussed and for much greater periods of time than available. We chose to focus on three topics that were relevant to the research objectives of all seven Autonomy Research Pilot Initiative (ARPI) projects.

1.1 SENSOR FUSION

Situational awareness in complex, ever-changing environments requires a perpetual collection of information from disparate sources at varying time scales. Developing ways to ingest, process, interpret and act in a real-time fashion is a non-trivial task, particularly with the ever-growing sophistication of fielded sensors. Traditionally, there has been a chasm between researchers in the field of sensor fusion between those that focus on the ingest and processing side and those that focus on the interpretation and action side. However, the meshing of these disciplines is where the future of autonomous intelligence will lie. Researchers from these divergent, yet dependent fields were brought together under the ARPI to discuss methodologies of transforming large amounts of raw data into meaningful information.

1.2 MACHINE LEARNING

The resurgence of machine learning methods in the past decade has brought some intriguing success stories of machine intelligence capabilities and seems very promising to enable better human–autonomy integration. However, in some cases, excessive hype and general lack of understanding of the underlying processes due to black box nature of machine learning has brought with it a lack of trust from many end users of autonomous systems. There must exist a critical balance between the transparency of machine actions interpretable at a human level with autonomous execution a given task to allow a human to attend other priority objectives. ARPI brought together a diverse group of machine learning experts with those that study trust issues in autonomy to investigate this problem together.

1.3 RESILIENT CONTROL ARCHITECTURES

The foundation of success for man–machine teaming relies on the infrastructure of the network controlling the information flow and task management between human and autonomous agents. Modern warfare continues to increase in complexity while the time allotted to make critical decisions in the ever-quickening cyber age dwindles. As the warfare environment is one of deception and unforeseen surprises in often unpredictable territories, robust and resilient control architectures are necessary to seamlessly switch between human, autonomy, and automation roles to empower and strengthen each agent at any given time. Control systems experts and behavioral neuroscientists have worked together in the ARPI project to begin to push the understanding of the actions and needs of diverse teams of hybrid agents to develop dynamic heterogeneous networks.

2. SENSOR FUSION

2.1 WORKING DEFINITION

Sensor fusion is the exploitation of available and potentially sparse information to resolve situational ambiguity. Traditionally, sensor fusion techniques have been used in various ways to combine disparate sources of data, such as determining the three-dimensional (3-D) trajectory of an object based on the x, y, and z coordinates. From this fusion, one can interpret and potentially predict the movement of an agent in 3-D space more certainty than when using information from less than three sensors.

2.2 IMPORTANCE AND CRITICAL FEATURES

The problem space that we envision human is complex. Using autonomous teams to navigate space requires a more holistic understanding of the agents themselves, their environments, and their relationships within a given context. Building a robust understanding of the world requires multiple streams of information and importantly, effective sensor fusion techniques to interpret environment and agent state changes. Critical features to allow for successful sensor fusion states estimates include dynamic real-time gathering and analysis of data, combining multiple types of data whose timescales are not consistent, and continuously updating the interpretation of data based on situational context.

2.3 CURRENT RESEARCH METHODOLOGIES AND TECHNIQUES

Current research in multisensor fusion techniques largely takes place close to the *signal* level, and prior to transformation from subsymbolic to symbolic space. The latter transformation is necessary for the transition to higher levels of *information fusion.* Well-understood techniques such as variant Kalman filtering techniques (Extended, Stochastic Cloning, Sigma Point, etc.), and graph-based techniques generally assume that the sensors have fixed transformations between their frames, i.e. they are mounted on the same rigid body platform. These techniques require customized derivation of how the *state affects the measurement* (inverse view) for the addition of new sensor types.

From ARPI, we have learned that even small-scale integration sensors with highly variable time dynamics (e.g., brain activity from a human agent with environmental variables) may be problematic with existing frameworks. The assumption of *continuous* sensor fusion does not address scenarios which call for *event-driven* sensor fusion. Additionally, the assumption of rigid-body transforms may preclude the use of many of these techniques on fusion of *distributed* sensors. Nevertheless, current techniques can be used to address major problems in current and future autonomy, namely:

- Multiple sensors (visual, inertial, biometric) to resolve ambiguity and noise
- Varying time-dynamics among multiple sensors
- Resilience of sensor-fusion platforms that require contingency management capabilities (i.e. the need to "hot-plug" sensors for operational reasons).

2.4 NEAR-TERM ACTIONABLE ITEMS

2.4.1 Integration of Sensor Fusion with Knowledge Representation

In the next 3 years, a push should be made to derive new mathematical frameworks for integrating *sensor fusion* (bottom up) with *knowledge representation* (top-down). This action would mitigate some of the time-dynamical issues by executing sensor-fusion based algorithms on aspects of the knowledge representation that the system is trying to build. It would also allow for *selective processing* of resources; i.e., the system would not necessarily need to update its state vector for *every* sensor measurement. Focus on *event-driven* rather than *continuous* processing is also a recommended shift based on experimentation performed under the ARPI.

2.4.2 Time-Varying Data in Denied Environments

We believe there should be focused research on further integration of sensors with very different characteristics and time dynamics in intermittent environments. Since our ability to gather may regularly be compromised due to changes in bandwidth, need for stealth, or enemy threat research should focus on teams of sensors without rigid transformations among them to remain operationally robust in denied environments.

2.4.3 Event-Cloud for Information Fusion

As mentioned above, a gap exists between sensor fusion and knowledge representation that limits our ability to gain critical information from the multitude of sensors we currently employ. To mitigate this problem, we recommend building a *dynamic graph* of events, where an event is roughly semantically equivalent to a sensor measurement transformed into symbolic space. The dynamic weighted relations of the graph act as the main knowledge representation for top-down fusion.

2.5 LONG-TERM RESEARCH INVESTMENTS

The above concept of an event-cloud knowledge representation to assist with *top down sensor fusion* indicates additional research into enabling technologies, as discussed in the following subsections.

2.5.1 Hardware Technology

Neuromorphic architectures such as IBM™ TrueNorth may provide a natural logical fit for accommodating multiscale sensors in an event-driven framework. The major concept at work is the natural processing of information at different layers of a feed-forward hierarchical network at different time scales, such as *Hierarchical Temporal Memory*.

2.5.2 Software Technology

Novel graph-based algorithms enable *world modeling* at the level required for a given agent/scenario to enable the type of dynamic, event-based processing described above:

- Current graph-based algorithms such as *union find* deal with dynamic insertion of new information but not removal
- A rule-based system for *information decay*, to remove outdated pieces of information, would likely be far too rigid
- Problems with the brittleness of ontologies would likely manifest immediately and would require robust investigation
- Representation of individual items in the graph must be sufficiently reusable
- Computationally, the graph must be processed in a tractable manner.

2.6 PERVASIVE ISSUES

Big data bottleneck issues have been a pervasive issue for a long time, but may largely be resolved if the above research investments are made. However, the big picture is that we are driving towards something higher level than sensor fusion: a *knowledge representation* to enable prosecution of the mission as well as *top-down* sensor fusion in the traditional sense. We currently lack graph/ontological algorithms as well as hardware that is sufficiently flexible to accommodate this vision in a complex, ever-changing environment.

For better understanding of augmentation human–autonomy integration, a better understanding of the dynamic complex behaviors of humans would be very valuable. However, derivation of a unified set of measurement equations from a human to enable a diverse array of sensors as biometrics (brain activity features, heart rate, polygraph sensors, etc.) into a single-system state function to be incorporated into a system remains frontier research.

The accumulation of knowledge representation over time and space results in a *world model*. A world model enables reasoning about connections between objects or events that are observed at different times or locations. Included in a world model may be short-term information that is collected by a robot as it conducts its current mission or long-term information that is collected over several missions. Short-term information has been used by current robotic systems recognize "loop closures" in navigation tasks. Long-term information may include a knowledge base built up from both human knowledge and the experiences of other autonomous agents. Such a knowledge store can support lifelong learning enabling the autonomous systems to apply prior observations and experiences to current tasks. However, as the previous section points out, care must be taken to manage the knowledge—eliminating superfluous data and combining/categorizing objects, activities and events. Complicating the knowledge management is task relevancy—information considered superfluous to one task may be vital to another. Cognitive approaches could quantify the association between specific information and specific tasks and should be further investigated. Additionally, *machine reflection* approaches need to be developed to relate similar activities and events, resolve conflicts, and update world models effectively.

3. MACHINE LEARNING

3.1 WORKING DEFINITION

Machine learning is traditionally viewed as a method of allowing computer systems to learn particular features of data without the need of a human explicitly programming it to do so. Algorithms are created that can be trained to understand aspects of a dataset, e.g., the presence of a cat in an image, and apply the knowledge to other unseen data to make predictions about the learned features on the new data.

3.2 IMPORTANCE AND CRITICAL FEATURES

As systems become more intelligent and capable of complex decision making, it becomes increasingly important for the human to understand the reasoning process behind the system's output to provide input when necessary. As machine learning algorithms increasing act as the intelligence underlying our systems, it is critical that they maintain some level of human interpretability. However, transparency for machine learning-based systems poses a challenge, as explanations may not be easily generated. Specific features that trigger responses in machine learning algorithms do not always help the user to understand the machine intelligence because these features usually do not provide meaningful information to the end user. A common example is a layer of a neural network that focuses on detecting edges in an image. The output of this layer looks like a set of Gabor filters and is not helpful for developing an understanding the capabilities or weaknesses of the system. However, without proper transparency—at least high-level explanations—it may be difficult for the human to provide input (based on information that only the human has) to the system. Furthermore, predictability may be an issue for machine learning systems as they continue to evolve. Without proper transparency, the human may find it difficult to calibrate his/her trust in the system.

Thus, a critical feature moving forward in machine learning is the appropriate amount of transparency into the system that allows the human users to understand system responses.

3.3 CURRENT RESEARCH METHODOLOGIES AND TECHNIQUES

There are many different methodologies that can be described as machine learning. More traditional methodologies such as k-means clustering, regression analyses, decision trees, Bayesian networks and neural networks are used in a plethora of ways to tackle data classification and interpretation issues. Though many of these methods have been investigated for decades, much attention has been given to machine learning more recently because of the proliferation of easy to access data and computing processing power enabling high-speed, high-throughput data algorithms. With these catalysts, the field of deep learning and use of multilayer neural networks has increased exponentially. These methods have proven capable of accurate classification on benchmarked datasets, particularly in the field of computer vision, far beyond traditional methodologies. These techniques have allowed us to explore data in ways never before imaged, such as real-time object detection, classification, segmentation, and representation of objects in our environment that will soon allow driverless cars to populate our roadways. Deep learning is not only tackling problems that have been nagging computer scientists and machine intelligence researchers for decades, but is beginning to solve problems we have never even considered before.

3.4 NEAR-TERM ACTIONABLE ITEMS

3.4.1 Methods for Debugging/Interpreting Learning Algorithms

Currently, established models for “debugging“ neural networks are inadequate, particularly deep learning networks. Networks in crucial roles should be able to either expose inner working to allow for a man-on-the-loop validated outcome, or be able to provide feedback on what and how a decision is made. Given the nature of neural networks, we recommended a focused effort on machine learning metacognition with the development of a tool that allows identification of user responses or nonsensical results.

3.4.2 Controlled Experiments with Systems Trained with/without Man in the Loop

Currently, machine learning can only solve very specific problems. Broader testing outside the few concept scenarios for each program are needed to make machine learning a tool to use for generalizable real-world scenarios. One possible way to inspire understanding and confidence in a machine learning system that is currently being explored in the ARPI program is to train the users along with the system. Giving users exposure to the learning process, and allowing them to watch the performance increase, may result in higher overall system performances as well as greater trust in the system, which would help to improve both the human and autonomy aspects of the system.

3.4.3 Fielded Testing with Real-Time Autonomous Agents

For many projects within ARPI, the autonomy of the individual elements in a system was simulated, with the assumption that such autonomous agents and their enabling technologies required would be developed by the time a system matured. If our intended audience is the warfighters, we need to address availability of the machine learning systems, the impact of the systems on the functioning of the autonomous agents, and the impact on overall training objectives on the human user. Currently, the impact of implementing machine learning techniques on the human–autonomy system as a whole has yet to be tested robustly. This includes the need to consider the amount of time/data required to train a machine-learning system for a given mission.

3.4.4 Collection of Empirical Warfighter Data

Similar to the need for testing machine learning algorithms on autonomous agents, we also have the need to collect data on our warfighters. Often a requirement for machine learning to work well is a large initial data test to train neural network models or other machine learning algorithms. If we expect our autonomous systems to someday take the place of human squad members to perform tasks, we must first understand how traditional squads act in theater. However, obtaining and accessing *real* warfighter data, beyond laboratory or simulated scenarios, is currently a difficult challenge for federal researchers. We suggest determining methods to allow researchers to obtain warfighter data to allow for the effective incorporation of squad dynamics into hybrid agent training activities.

3.5 LONG-TERM RESEARCH INVESTMENTS

3.5.1 System Generalization

System generalization (whether to one-shot learning, out-of-training tests, distortions, or chop-the-last-layer-and-retrain networks) will be a critical component of successful machine-learning systems deployed on autonomous systems. However, a significant amount of work is still required to determine how layered networks can be formed to allow for all-purpose networks in conjunction with specialized networks to provide robustness across a multitude of problem spaces.

3.5.2 Cyber in Machine Learning

An important aspect that deserves attention in this domain is cybersecurity. As with all things in the digital age, cybersecurity needs to be considered a forethought rather than an afterthought. If we ask our machines to learn and make decisions based on information not directly provided by trusted human agents, we must develop mechanisms to make sure these decisions are not manifested from data coming from enemy sources. Conversely, we should also consider if non-transparency be positively leveraged in some ways to enable less detectability by our enemies.

3.5.3 Trust in Machine Learning

Graceful, predictable, and understandable degradation are the types of behavior that inspire "forgiveness" of a user. Current machine learning systems generally lack these behaviors. Without these, systems that are "brittle" and fail catastrophically are not surprisingly the type that will not be trusted in a critical situation. A recommended starting point to address this issue is by applying the concept of *situational awareness based agent transparency*, which provides a framework of information requirements from the system to the human in order for the human to obtain proper situation awareness of the system's actions and plans in its mission environment: system's current actions and plans (Level 1), the reasoning process behind (Level 2), and projection of future outcomes and uncertainty (Level 3).

3.6 PERVASIVE ISSUES

Today's bureaucratic mechanisms for pushing technology to the warfighter put evolving algorithms in a precarious situation. Satisfying a requirement versus optimizing a system is a problem that arises when discussing machine learning systems for military applications vs civilian applications. Sometimes a system that is satisfactory, and correct most of the time, will be better than a system that is inscrutable but with higher accuracy.

4. RESILIENT CONTROL ARCHITECTURES

4.1 WORKING DEFINITION

Resilient control architectures, labelled the “next generation” control systems, take into account many different aspects of the task, environment, and agents to allow for robust and adaptive actions in complex scenarios. The discussion at the ARPI meeting was focused on a particular effort currently under investigation by the Navy, but one that can be applied to many critical DoD problem spaces. Specifically, the problem scenario discussed was the command and control of a heterogeneous multirobot human collaboration in a disconnected, intermittent, high latency (DIL) (24 days to 24 hours) communications environment. Such domains include underwater environments and areas where communication provides great risk of detection, including in areas behind enemy lines.

4.2 IMPORTANCE AND CRITICAL FEATURES

When dealing with a variable, heterogeneous communications environment, determination of the role players is important. In a conventional military setting, there is a central commander who issues orders to subordinates. The subordinates may or may not have some degree of decision-making authority. For example, a platoon may be assigned a mission objective by a central commander, but they may determine the exact sequence of actions they take to accomplish the mission. When humans team with unmanned systems, the level of autonomy given to each team has to be determined. For each situation encountered by a human–machine team, the functions handled by humans and by autonomous agents may change. The situation may dictate which humans command, and which are commanded, which level of autonomy each artificial intelligence (AI) agent may be permitted, and which functions of an unmanned system may be handled by automated scripts.

In addition to deciding the roles humans and machines will play, the question of how multiple agents cooperate seamlessly to accomplish the mission and adapt to change arises. In a DIL environment, the functions that must be considered are as follows:

- Navigation
- Communication
- Mission execution
- Mission planning and replanning
- Responding to real-time changes in the mission parameters or environment
- Distributed data processing
- Vehicle or robot status
- Infrastructure status and predicting its next state
- Mission optimization

4.3 CURRENT RESEARCH METHODOLOGIES AND TECHNIQUES

Research in planning algorithms has been conducted by the Artificial Intelligence and Control Theory communities for the last several decades. Much of the early work focused on centralized, domain-independent solvers where a single AI agent is responsible for generating a plan for all actors. Centralized planning methods include tree-based search methods, constraint satisfaction methods, satisfiability solvers, model checking, planners that exploit hand-coded control rules, and hierarchical task networks. Unlike centralized planning, decentralized planning involves multiple agents, which may or may not communicate with one another. An extreme example of decentralized planning occurs in swarm application, where each swarm member decides its own actions, which may or may not involve reasoning over what other swarm members may do, and may or may not consider the environment during decision making. A domain-dependent planning algorithm is one that takes advantage of the unique structure of a particular problem and is optimized to solve only that problem. Finally, research in domain modeling and planning over mixed continuous-discrete environments is ongoing.

Research in planning algorithms and AI reasoning has focused on addressing the issue of uncertainty. Hidden Markov Models (HMMs), Markov Decision Processes (MDPs), and Partially Observable Markov Decision Processes (POMDPs) have all been used to model uncertainty in problem domains, and methods of solving planning problems described as such. Another current approach develops inference reasoning algorithms to solve processes modeled as Bayesian Non-Parametric Models, which are adaptive models that can refine their complexity as new data are observed. Such methods may be applicable to planning in the DIL environment.

4.4 NEAR-TERM ACTIONABLE ITEMS

4.4.1 Adaptive Planning Algorithms

The planning algorithms currently used in practice today do not sufficiently address the DIL environment, and are not robust to change. Research and development of adaptive planning algorithms, as well as development of multimodal communication and system integration is needed to field teams of unmanned systems in a DIL environment. In an adaptive planning algorithm, the set of UxVs, communications outposts in the field, and central command will need to be dynamically partitioned, based on who can communicate with whom. Communication must be both flexible and intermittent in a DIL environment. Plans will need to be based on both the world/state models of the members in the set, as well as assumptions of what the other members of the other sets will do, especially in decentralized planning algorithms.

An important question that near-term research must answer is: How frequently should replanning take place? Clearly, planning that is not frequent enough can result in lack of pertinent updated information that may be vital to understanding dynamic task, environment, and agent conditions. Too frequent replanning can result in churn, a state of oscillation between plans resulting in no action. Both may result in inappropriate or suboptimal task execution, and possible mission failure. Certainly, the frequency of replanning will depend on the changes within the operational environment, communication ability, and urgency of tasks.

4.4.2 Multimodal Communication

Multimodal communication also needs consideration when deciding the type of planning to use. In an underwater environment, acoustic, optical, and magnetic methods of communicating exist. Acoustic methods are generally effective across distances on the order of kilometers, but have very low data rates (e.g., on the order of bits per second). Optical methods can facilitate higher transmission rates, but require systems in extremely close proximity. Magnetic communication is still an area of active research. Additionally, an underwater vehicle may be equipped with a radio-frequency (RF) antenna and the ability to communicate with satellites upon surfacing, although such communications are more susceptible to detection. AI planning can be used on board each fielded agent to determine which method of communication is most suitable, depending on the situational needs.

4.4.3 System Integration

Along with research in adaptive planning algorithms and planning for multimodal communication, physical systems that demonstrate their use must be developed and tested, which is important to consider early in the development phase, as some of its hardware requirements may be dictated by the choice of AI planning methods used. For example, an algorithm may require a minimum amount of on-board computational power to plan at sufficient speed. Specific sensors may also be required to provide sufficient data. It is also important to consider how physical design may evolve as sensors, control architectures and algorithms mature.

4.5 LONG-TERM RESEARCH INVESTMENTS

4.5.1 Shared World Views and Dynamic Updating

Longer term research investments include shared world views and dynamic updating. In a multi-agent system, each actor will have its own world view. Collaborative planning can be difficult when world views are vastly different. So, how can the agents effectively share their world views for planning purposes, especially in low-communication environments? Current research has demonstrated sensor fusion for situational awareness between unmanned vehicles. Such techniques may be applicable to shared world modeling. When communication between agents exists, world models can be updated to resolve conflicts and synchronize observations.

4.5.2 World View Interpretation

Interpreting what is sensed in the environment, especially in opaque, uncertain, and often novel environments remains a challenge. For example, threat detection in a DIL environment can be perceived as an AI classification problem in a partially observable world. An agent may have policies and actions based on goals, costs, rewards, and/or confidence levels, but defining what is classified as a “threat” in relation to the current state of the world remains a challenging endeavor.

4.6 PERVASIVE ISSUES

It remains a challenge to develop an optimal resilient control architecture that can be generally applied to human–machine teaming applications. The specificity of the agents, the operating domain and the mission, among other factors, can all heavily influence the configuration of the networked infrastructure needed for dynamical control.

5. DISCUSSION AND CONCLUSIONS

5.1 GENERAL RECOMMENDATIONS

5.1.1 Research in Autonomy

There is no doubt that human–autonomy teaming is critical to the future of warfare, which has been explicitly stated by many of our DoD leaders in the past year, most notably the Third Offset Strategy outlined by Deputy Secretary of Defense Robert Work in January 2015. ARPI was a unique opportunity for government researchers to accelerate this field in a direction that will benefit the DoD as a whole. Within only the first 2 years of the program, *more than 90 research articles and conference proceedings* were generated. ARPI was a great accelerant for DoD laboratories, but also uncovered just as many new important avenues for research as it did in answering questions. The knowledge gained at the DoD laboratories has the potential to continue to grow and mature internally if given the proper chance. Additionally, enabling unfettered and mission-oriented research across the DoD laboratories is a critical retention tool for young, talented researchers, which is especially important as the tech industry competition for jobs continues to grow.

5.1.2 Cross-service Collaboration

The modern technological challenges the DOD is addressing require large, diverse teams to achieve success. The ability to work between services was an extremely valuable experience for the technical performers, but additionally had a large return on investment to the DOD from the knowledge and resource sharing between federal researchers, which will manifest for many years to come. From the technical performers' standpoint, the ARPI model was successful in establishing unique environments for investigating high-risk/high-reward research objectives within the walls of the government laboratories.

However, enabling a unified vision with disparate researchers at disparate sites led to communication and collaboration challenges, which is a solved problem outside the DoD, with many modern technologies and tools facilitating virtual collaborations. For example, Github™, Dropbox™, Slack™, Google™ documents, etc., allow for team-level data input/output, with the ability to both store and operate on big data within the structure. This advancement would significantly accelerate collaborative data analysis compared to the current methods of personnel flying with bags of hard drives to do data transfer. Enabling these tools within unclassified the DoD S&T community would not result in a cybersecurity risk and would have enormous positive impact for communication and collaboration.

5.1.3 Knowledge Maintenance and Continuity

With the time, effort, and money invested in the ARPI Program, a need exists for a sustainment plan to keep relationships alive and the knowledge base intact. The goals and culture of ARPI can continue to thrive as an ongoing network of like-minded researchers. All the services greatly benefited from interservice collaboration, and if ASD(R&E) chooses to support other RDT&E efforts, it is recommended that leadership at each service works to sustain these collaborations beyond the initial years. Additionally, the success of ARPI will continue to surface as research progresses and applications transition. To appreciate the full value of this initiative, opportunities to regularly come together for technical exchange meetings to share achievements and progress as well as host a virtual shared space for discussions, updates, planning, and data sharing would be very worthwhile and furthermore help the DoD to track return on investments.

6. ABOUT ARPI

The Autonomy Research Pilot Initiative (ARPI) seeks to promote the development of innovative, cross-cutting science and technology for autonomous systems able to meet future DOD system and mission requirements. The focus is on those projects with the potential to radically advance capabilities 5,10 or more years in the future in important warfare areas. Envisioned technology will allow military systems to complete complex military missions in dynamic environments with the right balance of warfighter involvement. The ARPI is a pilot test of an OSD-sponsored innovation program, directed by ASD (R&E), and executed by the Services, with support from the DOD Priority Steering Council (PSC) for Autonomy.

The ARPI Invitation for Proposals was distributed in November 2012 to all DOD Federal S&T Enterprises with an emphasis on cross-service collaborations to focus on critical technology gaps that must be addressed in order for DOD to develop and field autonomous systems capable of carrying out complex military missions. Of the 47 submissions received, 7 ARPI projects were ultimately selected for FY13-FY16 funding, approximately a $25M effort. More information about ARPI and the individual projects can be found on the DoDTechipedia site: https://www.dodtechipedia.mil/dodc/x/SQKwAQ.

7. PRIMARY CONTRIBUTORS

This document was created by the following ARPI performers and collaborators:

Primary Contributors

Jamie Lukos	SSC Pacific – Technical Editor
Daniel Donavanik	ARL-SEDD – Sensor Fusion with Multi-scale Dynamics
MaryAnne Fields	ARL-VTD – Sensor Fusion with Multi-scale Dynamics
Ben Migliori	SSC Pacific – Machine Learning Transparency
Jessie Chen	ARL-HRED – Machine Learning Transparency
Leah Kelley	SSC Pacific – Resilient Control Architectures
Doug Lange	SSC Pacific – Resilient Control Architectures

Other Contributors

Ben Haynes	TARDEC – Sensor Fusion Outbrief Lead
John Reeder	SSC Pacific – Machine Learning Outbrief Lead
Eric Hendricks	SSC Pacific - Resilient Control Architectures Discussion Lead
Deji Coker	SSC Pacific
Will Curtis	AFRL/RW
Emily Doucette	AFRL/RW
Scott Douglass	AFRL/RH
Lorraine Duffy	SSC Pacific
Benjamin Files	ARL-West
Dan Gebhardt	SSC Pacific
Greg Hudas	TARDEC
Dai H Kim	ASD(R&E)
Zack Larrabee	ASD(R&E)
Raj Malhotra	AFRL/RY
Mike McCourt	AFRL/RW
Will Nothwang	ARL-SEDD
Victor Paul	TARDEC
Ryan Robinson	ARL-SEDD
Matt Yanagi	SSC Pacific

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-01-0188

1. REPORT DATE *(DD-MM-YYYY)*
April 2017

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*

4. TITLE AND SUBTITLE
Recommendations for Future Human-Autonomy Research Investments: Sensor Fusion, Machine Learning and Resilient Control Architectures

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHORS
Jaime Lukos
Benjamin Migliori
Doug Lange
John Reeder
Leah Kelley
SSC Pacific

Daniel Donavanik
ARL-SEDD

MaryAnne Fields
ARL-VTD

Jessie Chen
ARL-HRED

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
SSC Pacific
53560 Hull Street
San Diego, CA 92152–5001

8. PERFORMING ORGANIZATION REPORT NUMBER
TR 3057

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Assistant Secretary of Defense for Research and Engineering
Autonomy Research Pilot Initiative (ARPI)
3040 Defense Pentagon
Washington, DC 20301-3040

10. SPONSOR/MONITOR'S ACRONYM(S)
ASD(R&E)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Performers of the Autonomy Research Pilot Initiative (ARPI) came together 17–19 February 2016 at the Liberty Station Naval Training Center in San Diego to discuss current issues impeding the success of human–autonomy teaming and strategies moving forward to push though these barriers.

The meeting included three small group breakout sessions, each focused on a different control system topic, followed by interactive outbriefs of each group to the wider audience. The goals of the breakout sessions were to determine the logical next steps to advance the field of autonomy, how the ARPI program aided in creating the foundations for these future endeavors, and on what issues long-term research investments should focus.

15. SUBJECT TERMS
sensor fusion; machine learning; resilient control architecture

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
U	U	U	U	22	Jamie Lukos
					19B. TELEPHONE NUMBER *(Include area code)* (619)553-4178

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

INITIAL DISTRIBUTION

84300	Library	(1)
53603	D. Lange	(1)
53624	J. Reede	(1)
56150	B. Miglior	(1)
71720	L. Kelley	(1)
71740	J. Lukos	(1)

Defense Technical Information Center Fort Belvoir, VA 22060–6218	(1)

Approved for public release.

SSC Pacific
San Diego, CA 92152-5001