



Special Notice N00014-19-S-SN08
Special Program Announcement for 2019 Office of Naval Research
Basic and Applied Research Opportunity:
“Science of Artificial Intelligence –
Basic and Applied Research for the Naval Domain”

I. INTRODUCTION

This announcement describes a research thrust entitled “Science of Artificial Intelligence – Basic and Applied Research for the Naval Domain” to be launched under Fiscal Year (FY) 19 Long Range Broad Agency Announcement (BAA) for the Navy and Marine Corps Science and Technology, N00014-19-S-B001 (and any amendments), which can be found at <https://www.onr.navy.mil/en/work-with-us/funding-opportunities/announcements> This Special Notice is not a solicitation. Anyone interested in submitting a white paper or full proposal for any of the research opportunities described below shall follow the submission requirements and procedures set forth in the FY19 Long Range BAA and any supplemental guidance provided in this Special Notice.

The research opportunities described in this announcement fall under the Technology Areas listed at: <https://www.onr.navy.mil/our-research/technology-areas>

These Technology Areas include:

Machine Learning, Reasoning and Intelligence:

<https://www.onr.navy.mil/Science-Technology/Departments/Code-31/All-Programs/311-Mathematics-Computers-Research/machine-learning-reasoning-intelligence>

Computational Methods for Decision Making:

<https://www.onr.navy.mil/Science-Technology/Departments/Code-31/All-Programs/311-Mathematics-Computers-Research/computational-methods-automated-image-understanding>

Command Decision Making:

<https://www.onr.navy.mil/Science-Technology/Departments/Code-34/All-Programs/human-bioengineered-systems-341/command-decision-making>

Computational Neuroscience:

<https://www.onr.navy.mil/Science-Technology/Departments/Code-34/All-Programs/human-bioengineered-systems-341/computational-neuroscience>

The purpose of this announcement is to (1) focus the attention of the scientific and technical community on specific areas of interest related to the advancement of artificial intelligence, (2) encourage dialogue amongst those interested in this area with the Office of Naval Research (ONR), and (3) provide a timetable for the submission of white papers and proposals.

II. TOPIC DESCRIPTION

ONR is interested in receiving white papers and proposals in support of advancing artificial intelligence for future naval applications. Work under this program will consist of basic and applied research, and therefore projects would be funded under Budget Activities 1 & 2 (as defined in the DoD Financial Management Regulation Vol. 2B, Ch. 5). The overall S&T efforts will be conducted at the Technology Readiness Level (TRL) 1-5 stage.

Topic 1

Title: AI for Predictive Maintenance (AI Applied Research)

Background:

While AI approaches have proven effective for a limited number of mechanical systems, major gaps exist in effectively diagnosing faults in complex machinery. Examples of critical military systems include helicopter drivetrains, shipboard reduction gearboxes, electrical power generation and distribution systems and land system power and drivetrains.

There is increasing interest in using unmanned surface vessels for long duration voyages. This will mandate more sophisticated predictive maintenance systems with longer term, more effective diagnostic and prognosis algorithms. Approaches to predictive maintenance are needed to manage the ship's machinery and electrical systems when there is no one on board to address it. For instance, algorithms will need to automatically detect a fault and then reconfigure the system to achieve a minimally sufficient performance of the maritime vessel, to allow it to complete its mission or task despite potential system degradations.

Objective:

New and powerful AI techniques developed in recent years show promise to address issues mentioned above; however, development of AI-based condition indicators (CIs) for predictive maintenance are lagging behind many other AI applications. Specifically:

- Existing parametric sensor data from platform health and usage monitoring systems come (predominately) from normal operations and contains very little data from actual faulted components.
- Lack of data for both labeled fault modes and verified system faults limits the development of supervised learning AI techniques.
- Unsupervised learning AI techniques can address the challenges associated with the lack of labeled data, but still need verified condition assessments to validate outputs.
- AI based CIs need to be validated against equipment condition through appropriate teardown engineering inspections and documentation of fault types, severity and prognostic failure data.
- Traditional health and usage monitoring systems for complex electrical and mechanical systems rely on high sample rate vibration (for mechanical systems), voltage and current (electrical), and pressure (hydraulic) data for feature extraction and classification. AI techniques have traditionally not focused on wideband time series data implied by these traditional monitoring techniques.
- Supplemental platform usage and maintenance data can provide valuable insights into the nature of faults and need for maintenance predictions. Unfortunately, usage and maintenance data are gathered asynchronously and are captured at rates (orders of magnitude) less frequently than parametric sensor data. Compounding the data-rate problem is the fact that both usage and maintenance historical data are often incomplete and inaccurate.
- AI-based CIs do not exist for new and powerful fluid analysis technologies that determine fault type, severity and rate of progression through debris characterization using real time optical imaging. Moreover, hybrid machinery prognostics models do not exist that effectively integrate inputs from vibration, fluid analysis and platform usage. Relationships between fluid condition, fluid debris and machinery vibration remain unknown and unproven.

Research Focus Areas:

New AI-based techniques are needed to bring together parametric sensor data with usage data, maintenance data and life-cycle support information to enable determination of both asset health and the “fight” left in a platform. New learning techniques are required to determine system health from multiple data sources with widely different time scales, and incomplete and/or inaccurate data. These new learning techniques must also be capable of learning optimal time horizons for predicting platform remaining useful life.

Finally, solutions are needed to improve the collection of data from real fleet operational platforms to support AI-based algorithm development for complex mechanical and electrical systems.

For shipboard systems, approaches are sought in which the state of the vessel's machinery and electrical systems can be recognized through machine learning algorithms, based on sensor outputs at various locations in the system. In a larger ship, this will be a very large number of sensors. "State" in this context means either a normal operating condition or a fault condition. The machine learning algorithm should be able to identify the fault based on patterns in the sensor readings, and reconfigure the system to compensate for the fault. Offerors are expected to obtain or develop machinery/electrical sensor datasets to support their work. The diesel engine and drivetrain used on the Sea Hunter autonomous USV is of special interest. Datasets gleaned from simulations or land-based testbeds will be considered.

Research Topic Chiefs:

Dr. Thomas McKenna, ONR 34, 703-696-4503, tom.mckenna@navy.mil

Dr. Robert Brizzolara, ONR 331, 703-696-2597, robert.brizzolara@navy.mil

Topic 2

Title: Rapid Learning of Task Procedures (AI Applied Research)

Background:

- Procedural knowledge abounds in various natural language sources—from more structured documents such as scientific articles, user guides, Standard Operating Procedures (SOPs) documents, operating manuals, news articles, and web pages to less structured ones such as blog posts, wiki pages, and chat forums. While machine-learning-based natural language processing techniques have gotten very good at certain tasks such as intent recognition, sentiment classification, entity extraction, etc., extracting and validating procedural knowledge from text that can then be used to define best practices, SOPs, and/or developing procedural automation remains an open problem. The domains of command and control tasks, deriving operational procedures and best practices, and mission planning, as well as deriving maintenance (and predictive maintenance) procedures, are of particular interest for this topic.

Objectives:

- Existing knowledge capture techniques are labor intensive and involve multiple individuals working through multiple iterations of knowledge capture. Automated techniques for capturing knowledge are brittle, and generally inadequate for capturing and generalizing procedural learning. These approaches limit the ability for automated learning of workflow that could serve as real-time decision support/task assistance.
- Knowledge capture needs also differ depending on the intended use for the captured information for example, capturing procedures for explanation vs automation. Techniques are needed that automate the capture of both procedures as well as deriving rationale and explanations of task context in order to create proactive decision aids that deliver timely decision aids for mission planning and execution.
- Of particular interest is the capturing of procedures from written documents. Tasks are described with different degrees of structure depending on the knowledge source and domain. These procedures are typically described with significant implicit knowledge that humans can infer from previous learning, experience or context. For example, handwritten scripts to process procedural text from sources with fixed formats only work for those sources, and do not apply to less-structured text (e.g., scientific articles, blog posts, chat forums). Moreover, learned models are often tied to specific domains, and there is little ability to generalize and reuse procedures that are derived ad hoc. Techniques for identifying and deriving implied knowledge for AI/ Machine Learning (ML) algorithms will be needed in dynamically inferring processes from written documents as well as observed behaviors.

- Text based knowledge extraction to date has focused on building knowledge ontologies and knowledge graphs. What is needed is a greater focus on deriving task procedures, best practices from data to codify as knowledge.
- Procedures are complex entities and tasks can be described in different ways. Many sequential procedures are described using complex relationships expressed over multiple sentences. Tasks can have arguments of various types, tasks may be conditional, and there may be relationships of various types between tasks (subtask, ordering, conjunctive, disjunctive) implicitly or explicitly described. Techniques to model and validate procedural knowledge are needed.
- Other challenges arise due to: different domains exhibiting domain-specific jargon and classes of procedures, the need to detect and characterize incomplete, missing, implied, incorrect knowledge, and the information source may be in the form of question and answer sessions (spoken, text, chat). Techniques to detect, inference implicit knowledge are needed for these applications.
- Note, this research is not intended to focus on natural language processing per se (it is presumed those technologies already exist and will be exploited). The desired research will focus on automating learning that has been codified from various natural language text-based sources.

Research Focus Areas:

- Research and develop a generalized approach to deriving representations for procedural knowledge based upon a range of text-based materials.
 - Demonstrate Domain Transfer—can (learn to) extract procedures from text across several domains
- The resulting algorithmic approaches should demonstrate the ability to operate on a range of text sources, from formal procedural manuals, to less structured and incomplete task representations. Algorithms should:
 - Be able to make inferences about missing or implied procedural steps
 - Suggest missing/implied steps and procedures
 - Characterize assumed, implied and missing knowledge
- Ability to extract complex tasks (hierarchical tasks, conditional tasks) and task attributes (arguments, preconditions, effects)
- Evaluation: Using a corpus of procedural documents provided by the Government, the new algorithms will be evaluated based on metrics to include:
 - Procedural **Completeness**. Number of procedural steps correctly called out relative to SME's conducting task analyses of the same document, as well as compared to other performers
 - Procedure Mapping **Speed**. Time required to derive procedures
 - **Autonomy**. Human Intervention Requirements. Number of interventions required for Subject Matter Expert (SME) assistance with analyses

Research Topic Chief:

Dr. Jeffrey Morrison, ONR 341, jeffrey.g.morrison@navy.mil 703-696-4875

Topic 3

Title: Scalable Verification and Validation Tools for Artificial Intelligence in the Naval Domain (AI Fundamental and Applied Research)

Background:

- One of the greatest challenges to rapid adoption and fielding of Artificial Intelligence (AI) methods in the naval domain is the lack of Verification and Validation (V&V) tools that are automated, and:
 - Can scale to naval sized problems under realistic assumptions;
 - Can be tailored to different domains that may have different needs in terms of safety, time criticality, mission reliability and constraints; and
 - Address the degree of human oversight that is possible or practical.
- V&V is particularly important for systems that either directly make consequential decisions or that can influence consequential decisions through their interface with people or more traditional algorithms/machines.
- New V&V challenges raised by AI include:
 1. Systems that can substantially change, due to different forms of machine learning, knowledge acquisition, and knowledge management,
 2. Complex, sequential forms of reasoning and decision-making,
 3. Systems that may have unexpected or emergent properties due to interactions between more than one type of AI method, and
 4. Systems that may interact with humans in new and complex ways that are not captured well by the types of human models currently used in verification processes including in handoff, communication, and interplay.
- In all of these cases, the algorithmic decision space can be either non-deterministic or intractably complex. As a result, this space cannot be exhaustively searched, examined, simulated or tested using many traditional V&V methods. On the other hand, formal methods often lack the scalability or expressability to extend to these types of systems.

Objective:

Developing methods and tools to enable the generation of Artificial Intelligence requirements that are mathematically expressible, analyzable, and automatically traceable to different levels or abstractions of system design.

Further, there is a need for protections against requirements that can lead to problems such as so called “reward hacking” in which Artificial Intelligence systems learn, optimize, or otherwise find ways to solve problems that are mathematically correct relative to cost functions, goals,

constraints, or requirements, but lead to pathologically poor behavior or results relative to what the designer or user intended.

Formal Verification and Validation methods that support the new challenges of Artificial Intelligence systems and for which the reasoning can be effectively automated as opposed to requiring a good deal of hand-tailored theorem proving by experts.

Developing tools, processes, and frameworks for rigorously balancing verification and validation across the life cycle of an Artificial Intelligence driven system including at progressive steps during the design process, during certification processes prior to fielding, at run time in operations, and over the lifetime of systems that may continuously learn and acquire new knowledge.

Developing methods and tools for Verification and Validation that can be aggregated, reused, and rapidly repaired when a priori assumptions become invalid as systems are upgraded or deployed and utilized by warfighters under conditions different from the original design assumptions

Research Areas:

Development and evaluation of new automated V&V methods, frameworks, and processes to one or more example naval systems (including experimental/research ones) or surrogate systems that have representative complexity and challenges to the rapid adoption of Artificial Intelligence in naval systems. Ideally, this should also show that the methods are versatile enough to support a reasonable breadth of AI applications to naval systems as opposed to be tailored only for a very specific point design.

New V&V methods and processes that will help enable a structured argument, supported by evidence, to justify that a system that includes a variety of AI components is correct relative to models/requirements, and will support the needs of the end user through a combination of offline tools, lifetime and real-time monitoring, prediction, and repair.

Research Topic Chief:

Marc Steinberg, Code 351, marc.steinberg@navy.mil 703-696-5115

Topic 4

Title: Brain-Inspired Deep Learning with Spiking Neurons (AI Fundamental Research)

Background:

Deep Neural Networks have achieved strong performance in tasks such as image classification. However, they require significant time and computing resources to train. While the steady advance of Moore's Law has enabled their success, integrated circuit technology is reaching fundamental physical limits due to thermal and stability issues. At the same time, there is a need in the Department of Defense (DoD) for advanced compute capabilities on platforms severely limited in Size, Weight, and Power. Another need is Machine Learning (ML) systems that can learn without access to large amounts of labeled training data.

The brain is a highly decentralized event-driven computing device, processing multiple asynchronous streams of sensorimotor data in real-time. Neurons communicate through spikes - brief impulses transmitted to other neurons through synapses. Experimental evidence shows that not only the rate of spike firing, but also the precise timing of spikes can be important for processing information

Deep-learning models solve problems by assuming static units that produce analog output, which describes the time-averaged firing-rate response of a neuron. These rate-based artificial neural networks (ANNs) are easily differentiated, and therefore can be efficiently trained using stochastic gradient descent learning. The recent success of deep learning demonstrates the computational potential of trainable, hierarchical distributed architectures.

Simulations of deep learning networks is highly computation intensive, which consumes power and limits the efficiency of mobile devices. Neuromorphic hardware based on spike communication between chips is a thousand times more energy efficient and more compact than digital chips. Because spikes are discontinuous, which precludes computing gradients; it has not been possible to use stochastic gradient descent to solve complex problems. Research in spike-based computation has been impeded by the lack of efficient supervised learning algorithm for spiking networks.

Objective:

The goal of this project is to develop new learning algorithms for spiking neurons that will allow deep learning spiking networks to be built that can solve complex real-world problems. The consequences would be far reaching in terms of both the practical applications and the theoretical insights into how to compute with spikes.

Nature is an existence proof that this is a solvable problem.

Research Focus Areas:

- Develop new learning algorithms that can be used to train recurrent networks of spiking neurons.
- Apply the new learning algorithms to deep feedforward networks.
- Generalize the static neuron models to dynamical neuron models such as the Hodgkin-Huxley model, the Morris-Lecar model and the FitzHugh-Nagumo model.
- Learn gated network models of spiking neurons such as long short-term memory models.
- Show that it is possible to achieve a level of performance that is as good as traditional rate-coded deep learning networks.
- Implement a deep spiking network in hardware
- Build hardware for a deep spiking network that can learn on a chip.
- Novel adaptive weight devices for rapid learning
- Demonstrate the performance on a challenging classification problem with time series data (RF, sonar, vibration) and/or imagery.

Research Topic Chief:

Dr. Thomas McKenna, ONR Code 341, 703-696-4503, tom.mckenna@navy.mil

Topic 5

Title: Brain-based computation (AI Fundamental Research)

Background:

Computational neuroscience has served as an inspiration for artificial neural networks, including recent deep learning architectures and algorithms. However, only a limited set of neural principles have been fully explored in deep learning architectures. The Office of Naval Research (ONR) seeks novel approaches that draw upon a more complete understanding of brain circuits and mechanism that can extend the state-of-the-art in deep learning and recurrent networks. These two areas of interest are:

1. Brain-inspired memory processes

Memory is central to cognition and there are a number of promising directions for enhancement of deep learning or for alternative architectures. Working memory serves to maintain a set of items in an active store of information for a finite duration where it is available to for decision making, and interaction with long term memory. Working memory and related cognitive control has been regarded as analogous to LSTM (long short-term memory) although the latter is implemented in a manner that is not consistent with neural mechanisms. Memory consolidation, and the abstraction of significant events, memory replay and memory retrieval has been recently intensively studied in terms of cortical – hippocampal transactions. Neocortex and hippocampal areas have distinctive circuit architectures and translating cortical-hippocampal interaction into new deep learning architectures lead to promising new learning capabilities. Episodic memory is an important element of human daily experience and is also linked to the segmentation of events. Even though episodic memory has “one shot learning” like properties, we are still able to segment out and reuse elements such as the actors, actions, locations, relations and temporal sequences within episodes and reuse them later to recount events or mentally simulate new events. This is a rich area for translation into machine learning technology.

Objective:

The goal of this topic is to characterize the computational principles, at circuit and system levels, of brain memory systems sufficient to build more powerful brain-inspired machine learning systems with the power of human associative memory, the ability to link new memories to existing memories in a constructive process, and the ability to extract actors, actions, locations and temporal sequence from segmented events and later recombine these elements in analysis and simulation of new events.

Research Focus Areas:

- Computational models of working memory, based on cognitive neuroscience research, that elaborate its role in attention, cognitive control, perception and decision-making.

- Development of brain-based networks that emulate memory consolidation of “significant” memories, including linkage to existing memories that can be retrieved in the appropriate context (re-consolidation),
- Analysis and development of networks and models of episodic memories, in which elements of events, such as actors, actions, objects, locations, and temporal sequence can be extracted and later recombined for the analysis of events and for the simulation of events, including counterfactual reasoning, planning and prediction
- Collaborations between AI researchers and neuroscientists is encouraged

Research Topic Chief:

Dr. Thomas McKenna, ONR 341, 703-696-4503, tom.mckenna@navy.mil

2. Brain-based visual object recognition

State-of-the-art computer vision (CV) models based on feedforward convolutional neural networks (CNNs) have been highly effective on benchmark image datasets. However, CNNs often exhibit performance shortfalls with images they have not been trained on—for example, images of objects at untrained poses, scales, or illuminations. In contrast, biological visual systems exhibit these shortfalls little if at all. Recent findings suggest one likely cause for the performance differences between biological vision and CNN-based CV systems is the massive bi-directionality of biological vision. The goal of this topic is to deepen understanding of the computational role of recurrence in the visual brain and to exploit this understanding to develop and evaluate a deep recurrent neural network - object recognition system that approaches biological vision in its ability to robustly recognize objects in challenging images. The result carries the potential to yield a significant advance in the state-of-the-art in computer vision, providing a principled basis for the integration of bottom-up and top-down projections.

Objective:

The goal of this topic is to deepen understanding of the computational role of recurrence in the visual brain and to exploit this understanding to develop and evaluate a deep recurrent neural network object recognition system that approaches biological vision in its ability to robustly recognize objects in challenging images. The result carries the potential to yield a significant advance in the state-of-the-art in computer vision, providing a principled basis for the integration of bottom-up and top-down projections.

Research Focus Areas:

- Recurrent projections appear throughout the ventral (or “what”) stream of primate visual system and at several scales. They appear within processing stages, as in the inferotemporal (IT) and extrastriate cortices, and across stages, extending from the lateral

geniculate nucleus through IT cortex and beyond, including perirhinal, and prefrontal cortices. Exploit recently developed, high-resolution neural pathway silencing techniques to determine the relative contributions to recognition performance of these recurrent projections.

- Object recognition is made difficult for feedforward CNNs by introducing any of several forms of transformations between training and test images, including occlusion, changes in pose, blurring, and clutter. Exploiting silencing and other techniques, determine what roles, if any, the various recurrence projections play in facilitating biological performance in the face of these transformations.
- Determine mechanistically the computational role, that is, the computational problem solved, by recurrence in biological vision and its implications for visual reasoning and deliberation.
- Exploit resolution of the foregoing issues to create, demonstrate, and systematically assess the merit of a deep recurrent artificial neural network-based object recognition system

Research Topic Chiefs:

Dr. Harold Hawkins, ONR 341, 703-696-4323,
harold.hawkins@navy.milharold.hawkins@navy.mil

Dr. Behzad Kamgar-Parsi, ONR 311, 703-696-5754,
behzad.kamgarpars@navy.milbehzad.kamgarpars@navy.mil

Topic 6

Title: Explainable AI Systems (AI fundamental and applied Applied Research)

While the explainability of machine learning systems is always challenging, the challenges may be greater when the developed product makes thousands of decisions at machine speed as in the case of deep reinforced learning. While these methods have relevance and benefit to naval mission planning and execution, the explainability shortfalls need to be addressed. Research that addresses the following capability needs is of specific interest but other related approaches may be responsive.

1. Explainability of end-to-end AI: While explainability of deep learning is always challenging, progress is possible. In computer vision, explainability progress can happen by sharing the meaning of shape activations within a neural network. For example, an object classification of car is explained by lower level recognitions of wheels and bumpers. Translating this visibility to settings using deep reinforced learning that involves algorithm decisions across logistics, surveillance, maneuver and fires space is more challenging. New methods for explainability is needed for end-to-end AI implementations across a diverse decision space.
2. Develop capabilities or methods to interpret the actions of a system based on deep reinforced learning: In end-to-end AI systems with algorithms taught to play games such as StarCraft, what fails to be revealed is the motivation/intent behind observed decisions. Novel methods to develop AI that can reason about and report on the intent of an end-to-end AI system is desired.
3. Methods to compare hierarchical implementations of AI services versus end-to-end implementations: It is generally easier to build explainability into hierarchical AI systems as opposed to end-to-end since a classification of a car can be supported by observations of a wheel or a bumper. However, past research suggests that end-to-end systems can achieve a higher level of performance. Novel methods are needed to predict the potential gain of end-to-end versus hierarchical. Additional research is also needed to assess the value of the explainability of hierarchical AI systems for a specific application.

Research Topic Chief:

Martin Kruger, ONR 341, 703-696-5349, martin.kruger1@navy.mil

Topic 7

Title: Mission-focused AI (AI fundamental and applied Research)

Future watchfloor and planning cell execution will be conducted at machine speed. To enable this, novel AI methods are needed to help human staff and watchstanders plan, monitor and make adjustments based on the large streaming datasets that feed the common tactical and common intelligence pictures. Research that addresses the following capability needs is of specific interest but other related approaches may be responsive.

1. Multi-class mission state classifiers: Naval watchfloors monitor the track of a common tactical and intelligence picture (CTP/CIP) to constantly assess how an exercise or conflict is progressing across multiple measures of effectiveness. The development of methods enabling machine learning to assess a track of a CTP/CIP across multiple classes is desired. Researchers may use games such as StarCraft to generate data or may use unclassified datasets generated from mission simulators to support research.
2. Apply transfer learning across mission contexts: In transfer learning applied to AI methods, the goal is to accelerate the learning needed to answer one question by jump starting learning by using the lower layers of a neural network trained to answer another question. In a military domain, the challenge is to show that (for example) a decision tool to coordinate long range targeting developed in one mission context (opponent, location) can be used to rapidly train a service in a second mission context.
3. Effective mission state embeddings for deep reinforced learning: Complex mission simulation or game contexts can have a very large representation of the state of a pixel or gridded location. For example, if an opponent has 20 different capabilities the state of each pixel or grid can have 2 to the 20 different states. Learning policies relevant to this number of different states for each location of interest in a larger area of interest can be computationally very expensive. Novel methods are needed to represent complex states in much lower dimensional states that do not degrade the quality of learned policies.
4. Methods to classify decision tracks into coherent courses of action: Each sequence of decisions an AI algorithm makes in a game or mission simulation setting can be thought of as a course of action. In the current mission planning process, staffs typically consider three enemy and own force courses of actions. Deep reinforced learning methods are, however, capable of comparing an almost infinite number of decision tracks as these algorithms learn policy rules for different states. This method could enable staffs to have the time to examine a much larger number of courses of action. Approaches are needed, however, to cluster the very large decision space examined by the machine into a reasonable number of courses of action. Clusters of decision tracks

formed, need to be a good sample of the overall decision space examined but not so numerous that staffs would not have the time to reason over them.

5. Methods to discover the biases of an enemy force commander given deviations between observed decisions and a learned optimum decision sequence to exploit those biases: Decision policies learned by reinforced learning approaches result in decision recommendations leading to the optimal value of measures of effectiveness. A human decision maker will not, however, always follow the decision recommendations from an algorithm. Research is needed to develop methods to introduce human bias and variance into a machine generated decision track. This will enable trainers to be built that emulate the expected actions of a specific threat.
6. Relative to a complex decision sequence required by a large force conflict, develop methods to estimate maximum human level performance, maximum possible performance and predicted variances with changes in force composition or mission objectives: As part of the training process for deep reinforced learning, the AI developer needs to estimate highest possible system performance as well as highest achievable human level performance. For tasks such as computer vision this is normally straightforward. For contexts such as gaming, it is less clear how to measure optimal machine or human performance since it can vary with context. Methods are desired to define the measures needed in designing and executing an AI algorithm training program.
7. AI enabled training environments that test warfighter skill against a thinking opponent whose abilities can be tuned: Methods exist to enable users of games like chess to specify the expertise level of the AI opponent. These approaches restrict the AI's ability to consider future states. Approaches to translate this ability to complex mission simulations/games that are enabled by deep reinforced learning are needed to create trainers or games to vary their skill level.

Research Topic Chief:

Martin Kruger, ONR 341, martin.kruger1@navy.mil 703-696-5349

Topic 8

Title: Predictive Adaptations to Support Human Performance and Injury Prevention (Applied Research)

Background:

The tempo of military operations require personnel, in particular those serving in combat related specialties, to endure heavy physical demands ranging from short-term, high intensity workloads to sustained long-duration operations. Under these highly variable workloads, military personnel must have the appropriate physical abilities for successful operations. These circumstances also exist within the preparatory training and education pipelines for military personnel; completing physically demanding military courses such as the Navy Basic Underwater Demolition / Special Warfare (BUD/S) or Marine Corps Infantry Officer Course requires a certain level of physical fitness. To prepare for these courses and the physical rigors of operations entails physical conditioning and training.

Individuals respond/adapt to physical training differently,.. As such, the ability to sense and respond to the impact of training, or apriori predict activities and response presents an opportunity to maximize individual readiness. A recent survey of the physical training domain included these areas of highest priorities research: physical demands in operational and training environments, measurement of physical performance/fitness, physical training programs: strength and endurance, and musculoskeletal injuries: overuse injuries and mitigation programs. Technologies to assess and address physical training and fitness have been growing, and shown to reduce incidents of injury rates. While there has been some initial success with these technologies, there are also concerns about the accuracy and reliability of these technologies. Moreover, these technologies are often tuned to specific athletic domains, and do not generalize/transfer to military populations' tasks and activities.

Objective:

The goal of this project is to increase military personnel readiness (reduce military injuries, and enhance/increase overall fitness), by developing technologies and algorithms for assessing and predicting adaptation to physical training. Beyond identifying and characterizing adaptation to physical training, approaches must also include recommendations for future physical activities. Methods are also needed for detecting when models for classification and prediction are no longer reliable (e.g., due to changes in training methods, equipment, populations, etc.), communicating about reliability and other limitations, and recommending a course of action when tuning or re-testing of algorithms are necessary. These assessments must occur in an unobtrusive and rapid manner, and capable of scaling to support the military numbers. The results of this project extend beyond baseline fitness and impact both the practical applications within the military, and theoretical insights into physical adaption.

Research Focus Area:

Proposals for research should account for the Human Performance Optimization model and associated global data sets, along with a more centralized focus on areas that include validated metrics within the physical human performance domains. An integrated understanding of the stress-adaptation response influence upon autonomic nervous system control and the relationship to exercise-induced inflammation and overuse injuries are important to the general construct of this work. There is a strongly desire that commercial technology/hardware is utilized, and the focus is on developing software, algorithms, and knowledge to bring together physiological, biomechanical, subjective, or other validated metrics of performance with military relevant metrics to assess fitness and injury risk. New learning techniques are sought for assessment of military fitness and injury prediction from multiple data sources and timescales with missing or incorrect data. The focus should be on workload and allow for predictive modeling of adaptations of current individual status and future planning of performance-state within operational and training environments. In addition, new learning algorithms and measures are necessary to support stratification of responders to physical training regimens; and provide individual programmatic activities and training load recommendations. The initial focus for the military tasks are the close combat arms positions, which requires short-term high physical activities, and sustained long-duration operations. Proposals should describe the method by which they will share the products to academia, industry and government – e.g. datasets, software libraries, or knowledge repositories resulting from the research.

Research Topic Chief:

Dr. Peter Squire, ONR 341, peter.squire@navy.mil 703-696-0407

III. WHITE PAPER SUBMISSION

Although not required, white papers are strongly encouraged for all offerors seeking funding. White Papers will be evaluated by the Government to determine whether the technology advancement proposed appears to be of particular value to the Department of the Navy. Initial Government evaluations and feedback will be issued via e-mail notification from the Technical Point of Contact. The initial white paper appraisal is intended to give entities a sense of whether or not their concepts are likely to be funded.

A detailed Full Proposal (Technical and Cost Volumes) will be subsequently encouraged from those Offerors whose proposed technologies have been identified through the above referenced email as being of “particular value” to the Government. However, any such encouragement does not assure a subsequent award. Full Proposals may also be submitted by any Offeror whose White Paper was not identified as being of particular value to the Government or any Offeror who did not submit a White Paper.

For White Papers proposing efforts that are considered of particular value to the Navy, but either exceed available budgets or contain certain tasks or applications that are not desired by the Navy, ONR may suggest a Full Proposal with reduced effort to fit within expected available budgets or an effort that refocuses the tasks or application of the technology to maximize the benefit to the Navy.

White Papers shall not exceed 5 single-sided pages, exclusive of cover page, references, and resume(s) of principal investigator(s), and shall be in 12-point Times New Roman font with margins not less than one inch. White Papers shall be in Adobe PDF format (preferred) or in Microsoft Word format compatible with MS Office 2010.

The cover page shall be labeled “WHITE PAPER” and include the current Long Range BAA Solicitation Number, Special Notice Number N00014-19-S-SN08, proposed title, technical points of contact, telephone number, facsimile number (if available), and email address.

The 5-page body of the White Paper shall include the following information:

1. Principal Investigator(s);
2. Relevance of the proposed effort to the research areas described in Section II; relationship of the proposed work to current state-of-the-art.
3. Technical objective of the proposed effort;
4. Technical approach that will be pursued to meet the objective;
5. A summary of recent relevant technical breakthroughs; and
6. A funding plan showing requested funding per fiscal year.

Resume(s) of the principal investigator(s), not to exceed one (1) page per principal investigator, shall also be included after the 5-page body of the White Paper.

White Papers shall be submitted via email to tom.mckenna@navy.mil with “WHITE PAPER” in the subject line.

White Papers shall otherwise comply with requirements of the ONR Long Range BAA, N00014-19-S-B001, or current Long Range BAA.

To ensure full, timely consideration for funding, White Papers should be submitted **no later than 15 August 2019**. White Papers received after that date will be considered as time and availability of funding permit.

The planned date for completing the review of White Papers is **15 SEPTEMBER 2019**.

IV. FULL PROPOSAL SUBMISSION AND AWARD INFORMATION

Full proposals shall be submitted under the FY 19 Long Range BAA, N00014-19-S-B001, or current ONR Long Range BAA, by **15 October 2019**. Full proposals received after that date will be considered as time and availability of funding permit. Full proposals shall be submitted in accordance with the requirements of the current ONR Long Range BAA.

ONR anticipates that grants, contracts or OTAs will be issued for this effort.

Full proposals for contracts shall be submitted in accordance with the instructions of the current ONR Long Range BAA.

Full proposals for grants shall be submitted via Grants.gov. The following information must be completed as follows in the SF-424 to ensure that the application is directed to the correct individual for review: Block 4a, Federal Identifier: Enter N00014; Block 4b, Agency Routing Number: Enter the three (3) digit Program Office Code and the Program Officer's name: The relevant Program Officer will be the first Topic Chief for a topic. All attachments to the application should also include this identifier to ensure the proposal and its attachments are received by the appropriate Program Office.

Although ONR expects the above described program plan to be executed, ONR reserves the right to make changes.

Funding decisions should be made by **01 November 2019**. Selected projects will have an

estimated award date of **01 March 2020**.

V. SIGNIFICANT DATES AND TIMES

Event	Date	Time
White Paper Submission	15 August 2019	1400 Eastern Local Time
Notification of White Paper Valuation*	15 September 2019	
Full Proposal Submission	15 October 2019	1400 Eastern Local Time
Full Proposal Selections*	01 November 2019	
Awards*	01 March 2020	

Note: *These are approximate dates

VI. POINTS OF CONTACT

In addition to the points of contact listed in N00014-19-S-B001, the specific points of contract for this announcement are listed below.

Technical Points of Contact:

Topic 1:

Dr. Thomas McKenna, ONR 34, 703-696-4503, tom.mckenna@navy.mil

Dr. Robert Brizzolara, ONR 331, 703-696-2597, robert.brizzolara@navy.mil

Topic 2:

Dr. Jeffrey Morrison, ONR 341, jeffrey.g.morrison@navy.mil 703-696-4875

Topic 3:

Marc Steinberg, Code 351, marc.steinberg@navy.mil 703-696-5115

Topic 4:

Dr. Thomas McKenna, ONR Code 341, 703-696-4503, tom.mckenna@navy.mil

Topic 5:

Dr. Thomas McKenna, ONR 341, 703-696-4503, tom.mckenna@navy.mil

Dr. Harold Hawkins, ONR 341, 703-696-4323, harold.hawkins@navy.mil

Dr. Behzad Kamgar-Parsi, ONR 311, 703-696-5754, behzad.kamgarparsi@navy.mil

Topic 6

Martin Kruger, ONR 341, martin.kruger1@navy.mil 703-696-5349

Topic 7:
Martin Kruger, ONR 341, martin.kruger1@navy.mil 703-696-5349

Topic 8:
Dr. Peter Squire, ONR 341, peter.squire@navy.mil 703-696-0407

Primary Business Point of Contact:

Michelle Parrott

Contract Specialist, Code 252

michelle.parrott@navy.mil

Secondary Business Point of Contact:

Phillip Lee

Contracting Officer, Code 252

phillip.m.lee2@navy.mil

VII. SUBMISSION OF QUESTIONS

Any questions regarding this announcement must be provided to the Technical Point of Contact and Business Point of Contact listed in Section VII above. All questions shall be submitted in writing by electronic mail.

Answers to questions submitted in response to this Special Notice will be addressed in the form of an Amendment and will be posted to the following web pages:

- Federal Business Opportunities (FEDBIZOPPS) Webpage
<https://www.fbo.gov/>
- ONR Special Notice Webpage
<https://www.onr.navy.mil/en/Contracts-Grants/Funding-Opportunities/Special-Notices>

Questions regarding White Papers or Full Proposals should be submitted no later than two weeks before the dates recommended for receipt of White papers and/or Full Proposals. Questions after these dates may not be answered.