

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Artificial intelligence and machine learning for future army applications

Fossaceca, John, Young, Stuart

John M. Fossaceca, Stuart H. Young, "Artificial intelligence and machine learning for future army applications," Proc. SPIE 10635, Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR IX, 1063507 (4 May 2018); doi: 10.1117/12.2307753

SPIE.

Event: SPIE Defense + Security, 2018, Orlando, Florida, United States

Artificial Intelligence and Machine Learning for Future Army Applications

John M. Fossaceca, Stuart H. Young
U.S. Army Research Laboratory, Adelphi, MD USA
john.m.fossaceca.civ@mail.mil, stuart.h.young.civ@mail.mil

ABSTRACT

Based on current trends in artificial intelligence (AI) and machine learning (ML), we provide an overview of novel algorithms intended to address Army-specific needs for increased operational tempo and autonomy for ground robots in unexplored, dynamic, cluttered, contested, and sparse data environments. This paper discusses some of the motivating factors behind US Army Research in AI and ML and provides a survey of a subset of the US Army Research Laboratory's (ARL) Computational and Information Sciences Directorate's (CISD) recent research in online, nonparametric learning that quickly adapts to variable underlying distributions in sparse exemplar environments, as well as a technique for unsupervised semantic scene labeling that continuously learns and adapts semantic models discovered within a data stream. We also look at a newly developed algorithm that leverages human input to help intelligent agents learn more rapidly and a novel research study working to discover foundational knowledge that is required for humans and robots to communicate via natural language. Finally we discuss a method for finding chains of reasoning with incomplete information using semantic vectors. The specific research exemplars provide approaches for overcoming the specific shortcomings of commercial AI and ML methods as well as the brittleness of current commercial techniques such that these methods can be enhanced and adapted so as to be applicable to army relevant scenarios.

Keywords: Artificial Intelligence, Machine Learning, Online Machine Learning, Unsupervised Learning, Human in the loop, Reinforcement Learning, Deep Learning, Artificial Reasoning, Natural Language, Embedded Semantic Reasoning

1. INTRODUCTION

The battlefield of 2050 will surely be different than it is today. Not only will inevitable advancements in technology such as processor speed, memory size, and power efficiency—in addition to increases in overall lethality—shape the way warfare is carried out, but the multi-domain battlefield of the future will profoundly affect our overall tactics and strategies. The US Army has anticipated this situation for some time as indicated in the Training and Doctrine Command's (TRADOC) "Army Operating Concept" (AOC) – "Win in a Complex World". In fact, the AOC recognizes that rapid advances in technology and application to military uses as a forcing function for the US military for rapidly integrating advanced technologies in to the force and this integration requires collaboration between federal science labs, industry, academia, and international partners. One of the "key technology focus areas" called out by the AOC is "autonomy-enabled systems" that will act as "force multipliers" with the ability to learn on the fly with built-in artificial intelligence. Such technology will enhance soldier situational awareness and reduce cognitive load, while providing "persistent surveillance and reconnaissance" across all domains, even though inaccessible to humans both physical and in cyberspace. The US military and its allies will increasingly depend on AI and ML technologies in order to maintain overmatch and ensure dominance in the multi-domain battlefield.

What exactly do we mean by AI? One definition used at ARL is "the ability of computer-enabled agents (e.g., robots) to perceive the world, reason and learn about it, and propose actions that meet the agent's goals." [1] With such a definition one can imagine that robots of the future will be able to assist soldiers by working alongside them in the battlefield with the ability to understand the true intention of commands and able to reason about the current situation, understand the context, and communicate in a fairly natural way with their human teammates. The ability to reason about the physical world and perform autonomously is challenging for robots today under the most ideal conditions, let alone a dynamic battlefield. The subfield of AI called machine learning refers to "improving machine knowledge and performance via interactions with the environment, data, people, etc." [1]. An alternative definition of ML is given by [2] and is attributed to Tom Mitchell as "a branch of artificial intelligence", and is "a scientific discipline that is concerned with the design and development of algorithms that allow computers to evolve behaviors based on empirical data". Wikipedia states that ML

is “a field of computer science that gives computer systems the ability to ‘learn’ (i.e., progressively improve performance on a specific task) with data, without being explicitly programmed” [3].

2. THE US ARMY’S UNIQUE CHALLENGES

AI and ML have undergone a revival in the past decade with the emergence of several advances, including deep learning with convolutional neural networks (CNNs) on massive labeled datasets[4] and significant progress towards the realization self-driving cars with some estimates predicting that by 2040 self-driving vehicles will comprise 95% of all automobiles [5]. Nevertheless, the commercial pipeline of AI and ML techniques have not yet found their way into military applications. There are good reasons for this, as these modern techniques currently suffer from some deficiencies that make them unsuitable for use in the cluttered, contested, complex, and fast-moving situations in which the US Army operates [1]. The so-called “fog of war” will only increase as the battlefield splinters into multiple domains, ranging from conventional to urban warfare, from physical space to cyber space, and from clearly demarcated battle zones to a seemingly inscrutable social terrain. ARL is working towards resolving some of these issues with both commercial and organically developed AI/ML techniques in order to bridge the gaps between state-of-the-art methods and the military requirements for operationalizing this technology for the battlefield of the future. This paper describes these gaps in some detail and then provides an overview of some of ARL’s initiatives aimed at addressing those issues. Finally, we discuss in some specific research approaches that ARL’s CISD Information Sciences Division (ISD) is exploring in order to tackle several of the fundamental hurdles posed by and ways to transform commercial AI/ML methods such that they can be successfully applied to military operational scenarios.

The US Army is cognizant of the challenges of the future battlefield and has published a number of documents outlining approaches and requirements for expanding the US military’s capabilities in the coming decades. For example, in TRADOC’s “Win in a Complex World” [6] 8 specific technology focus areas are called out, including “autonomy-enabled systems” that will be used as force multipliers to increase soldiers’ situational awareness and mobility, and allow them keep up with battlefield operational tempo. Furthermore, AI-enabled autonomous systems will be capable of learning in the battlefield and aid in “force protection” while “providing increased capabilities to maintain overmatch” [6].

ARL has strategically aligned its research portfolio to support the needs of the future Army, including the establishment of 9 essential research areas (ERAs) that the laboratory has determined are required to support the Army of 2050 [7]. These ERAs are:

- Human Agent Teaming
- Artificial Intelligence and Machine Learning
- Cyber and Electromagnetic Technologies for Complex Environments
- Distributed and Cooperative Engagement in Contested Environments
- Tactical Unit Energy Independence
- Manipulating Physics of Failure for Robust Performance of Materials
- Science for Manufacturing at the Point of Need
- Accelerated Learning for a Ready and Responsive Force
- Discovery

In addition, the Chief of Staff of the Army (CSA) has outlined 6 nearer-term modernization priorities [8]:

- Long Range Precision Fires - (LRPF) – advanced systems with improved range, munitions, and target acquisition capability
- Next Generation Combat Vehicle (NGCV) – highly mobile, lethal and maneuverable including manned, unmanned, and optionally manned variants
- Future Vertical Lift Platforms (FVL) – ruggedized and highly maneuverable military helicopters including manned, unmanned, and optionally manned variants

- Army Network Command, Control and Communication and Intelligence (Network C3I) - sufficiently mobile and expeditionary - that can be used to fight cohesively in any environment where the electromagnetic spectrum is denied or degraded
- Air and Missile Defense (AMD) – protects forces from modern and advanced air and missile delivered fires, including drones
- Soldier Lethality (SL) – includes next-generation individual and squad combat weapons improved body armor, sensors, radios, loadbearing exoskeletons, advanced training and simulation

Clearly, AI/ML approaches can and will be applied to many of these areas. However, in this paper we focus our attention on aspects of the Human Agent Teaming and AI/ML ERAs, specifically some of ARL CISD’s Information Sciences Division (ISD) research that addresses these 2 areas. ISD’s fundamental research programs also support the CSA’s modernization priorities most directly in the areas of the NGCV but will also impact the other CSA areas.

3. NOTIONAL US ARMY REQUIREMENTS FOR AI & ML

When examining the nature of US Army mission sets, it is clear that deployed systems must be reliable, rugged, and robust. As always, the Army must adapt to a wide range of missions sets. The US Army, along with several experts, have predicted in studies such as [9], [10] and [11] that the Army will operate in “Dense Urban Environments,” megacities that are congested with people, complex infrastructure, and overlaid with a cyber-social terrain that will require a multi-domain approach and involve artificial intelligence that can rapidly and dynamically respond to ever-changing threats, reducing the cognitive load on the soldier while improving overall situational awareness. In fact, according to United States Army Capabilities Integration Center (ARCIC), disruptive technologies such as robotics, autonomy, artificial intelligence, and cyber warfare will be key for operations in megacities and artificial intelligence is singled out as possibly “the most disruptive technology of our time”. [11]

ARL’s basic research program has identified 3 major research challenge areas for the AI & ML ERA (depicted in Figure 1) that help to focus research efforts on studies that will lead to necessary and expanded capabilities for the soldier:

- Learning in Complex Data Environments
- Resource Constrained Processing at the Point of Need
- Generalizable and Predictable AI



Figure 1. Envisioning the Future Battlefield. In the future the multi-domain battlefield will be extremely complex and consist of highly distributed heterogeneous assets including intelligent agents that incorporate many AI & Technologies that improve situational awareness, assist soldiers with tasks, offload cognitive tasks and help the force make better decisions quickly [12].

3.1 Learning in Complex Data Environments

The US Army's operational environment is truly complex and often unpredictable. New AI & ML methods are needed that effectively handle datasets that have few samples, noisy data and operate in environments that contain a lot of "clutter" and "congestion" beyond just physical entities but also in terms cyberspace and social information. Not only are these data hard to collect ahead of time and still more difficult to process, but the data consists of a wide variety of data types from audio, video, text, network traffic, social media, news, wireless transmissions, and many other heterogeneous forms of information. Some of the events and information flows in this environment (e.g. megacity) may be natural and expected, while others may be deceptive and adversarial—intentionally created to thwart the Army's progress. Contested and deceptive environments will be especially challenging when dealing with sophisticated and near-peer adversaries, spawning the need for AI and ML methods that are not only robust and resilient but can also infer, quantify, and reason about the uncertainty of various situations [12].

3.2 Resource Constrained Processing at the Point of Need

Unlike the commercial AI and ML application pipeline which largely depends upon large and stable training data as well as fairly static operational environments with access to significant computing resources, the operational environment of the military is one of learning in complex, dynamic and under largely unpredictable circumstances. Access to powerful cloud computing facilities will likely be very limited or completely unavailable in military contexts. Power is always at a premium in the battlefield, and physical loading must be optimized to ensure that the weight and size of components can be supported by the mission. In the future battlefield consisting of many interconnected devices, this internetwork of "battlefield things" will require algorithms that run effectively in a distributed manner over a potentially large set of less powerful CPUs working with unreliable connectivity to solve complex learning problems and to provide useful, timely analytics and services to the warfighter [12,13].

3.3 Predictable and Generalizable AI

As mentioned earlier, the quantification of uncertainty is a requirement for AI & ML methods that will be employed in the future battlefield. This implies that these techniques need to be predictable so that failures are not catastrophic and performance limits are understood. Today's commercial AI and ML methods tend to be specific and narrow, and not general. Research is needed to create methods that generalize and that can transfer learning from one model to another, even in another domain. This will require methods that have built-in quantitative models that not only produce an "optimal" decision or assessment based a set of observations, but need to have some notion of the accuracy of the recommendation or assessment and explain its reasoning process, providing justification for the course of action to the soldier. It may even be possible to have provisions for human-in-the-loop training that will allow algorithms to be programmed to meet the performance requirements of specific missions [12, 14].

To summarize, the overarching goal for Army research in AI and ML is to develop artificially intelligent agents that are both heterogeneous and distributed; and that rapidly learn, adapt, reason, and act in contested, austere, and congested environments [12].

4. GAPS IN THE COMMERCIAL AI & ML PIPELINE

In order to address the aforementioned goal, we turn to the AI & ML commercial pipeline in order to see what techniques might be applicable to Army situations. The commercial AI & ML Pipeline has made some dramatic advances in the past decade, as evidenced by breakthroughs such as AlexNet's ability to classify images [4] followed by many successful applications of ML [15]. ARL has explored the many challenges that AI and ML algorithms face in a military context, and has identified a number of gaps that are described in the rest of this section.

4.1 Lack Enormous, Relevant Datasets with Labeled Examples

When examining the state of the art of AI & ML research, it is apparent that large amounts of example data are required in order to train many of the most successful algorithms [16]. In military contexts such massive datasets are rarely available. Even when there is a large amount of collected data, often it is not labeled and, therefore, not suitable for training algorithms that depend on the paradigm of supervised learning. With supervised learning, the learning algorithm must be presented with examples that each belong to a known class. The algorithm then trains on the labeled examples. Unfortunately, many algorithms such as convolutional neural networks with deep learning do not perform well when presented with sparse, incomplete datasets rendering them virtually useless for supporting military operational scenarios.

Even under the most ideal conditions, such as with self-driving cars and extremely large training sets, there have been situations where a human driver would have likely avoided an accident where the self-driving car failed. Ideally, new breakthroughs in AI and ML research could help to address this issue.

4.2 Batch-Based Processing Algorithms that Work Poorly with Sparse Data

Another important limitation of many commercial AI & ML algorithms is their dependence on “batch processing” of training datasets. With batch processing all training data is required at once in order for the learning algorithm to build a useful model that can be deployed into an operational setting. These “learned models”, rather than those handcrafted from domain expertise and first principles, have gained increasing relevance for a variety of engineering and pattern recognition tasks, from speech recognition to autonomous control of complex robotic platforms. The dominant approach to learning methods is based on neural networks, whose training requires optimization over millions of parameters and data points in advance of their use. This training procedure further requires cloud infrastructure to store massive amounts of data. When neither extensive data nor cloud services are available, learning in situ is required, for which training deep networks is intractable. The question then becomes how to design learning methods that can learn stably online while performing well after only tens to hundreds of data points, and strike at the benchmarks of deep learners as the number of observations becomes very large.

4.3 Dependence on Supervised Learning Algorithms

With supervised learning, exemplar training data is required that is completely representative of all operational scenarios that the trained agent will face. As discussed earlier, collection of this magnitude and quality of training data is generally not realistic for Army-relevant operations. Nevertheless, many of the recent successful machine learning algorithms are based on supervised learning [17]. In supervised learning systems, the learning agent is provided with the correct answer or “ground truth” label for each example input in the dataset—that is, mapping a set of inputs, X , to a set of outputs, Y . In the example provided earlier, AlexNet uses a deep convolutional neural network to classify 1.2 million high-resolution images into 1000 different classes achieving better than previous state of the art error rates [4]. Ideally, military applications of machine learning should not depend solely on supervised methods and if possible newer techniques that use semi-supervised or unsupervised approaches should be utilized. However, the current state of the art unsupervised learning approaches are not sufficiently robust to handle militarily relevant situations. In order to operationalize AI and ML, research is needed to develop new semi-supervised and unsupervised learning algorithms that are less susceptible to noise and can effectively and accurately generalize with either sparsely labeled or totally unlabeled data sets.

4.4 Inability to Operate Autonomously and Reason about the Environment

Intelligent agents need a lot of human supervision and cannot operate in the environment autonomously. Most robotics currently used in the field are teleoperated by soldiers and even if programmed to operate in an autonomous manner, most robots are unpredictable and unable to adapt to unexpected environmental conditions. Current intelligent systems cannot reason about or learn from the environment in which they operate. As conditions change, intelligent agents need to adjust to the changing conditions. For humans, this comes naturally. For example, when a person walks outside on a sunny, dry day and suddenly it begins to rain, they will sense that the ground surface seems slippery and automatically adjust the way they walk. Likewise, on a cold winter day if icy patches appear on a sidewalk, people will deliberately change the way they take their steps, differently than they would in the rain, perhaps taking each step tentatively until they get an intuitive sense of the surface traction, automatically anticipating the possibility of slipping and falling. This innate ability is partially pre-wired and partially learned from the time the person is a small child just learning to walk. No one has to tell a human what to do in this instance, they learn through experience. Likewise, if the driver of a car senses that the road feels slippery, they will adjust their speed accordingly. An inexperienced driver may slam on the brakes the first time they encounter a slick road and need to stop, but will quickly learn that this may cause loss of control and will adjust future behavior accordingly. In fact, once a driver has gained experience, they may see glaze on the road surface and automatically change the way they drive in anticipation of potentially dangerous road conditions. Unfortunately, this ability to learn from and reason about the environment is lacking in today’s robots. In fact, most robots are tethered to humans via a wired or wireless link. If they get stuck, humans must intervene. In reality, most robots are a remote controlled extension of the human who must focus a significant amount of their attention on the task of controlling the robot. Intelligent agents have difficulty learning from a human operator either by example or by direct feedback. Instantiating truly autonomous capabilities into robots remains a challenging problem and active area of research.

4.5 Lack of Common Sense

In order for robots to operate effectively as teammates for soldiers, they will not only have to reason about the environment but they also need to have an innate “common sense,” such that a soldier can ask a robot to complete a task and robot will understand how to properly carry out the order. For example, a soldier may ask a robot to explore an area and report back if they find anything “of interest”. This task would be straightforward for humans, who will identify things such as weapons, people in the area, vehicles, buildings, and any of a number of “interesting” facts that cannot be anticipated in advance. Basic research is needed in order for intelligent agents to be able to understand the “intent” of the commander’s orders; furthermore, the intelligent agent should be able to ask questions in order to clarify instructions that it has been given. Part of the reason this problem exists is that intelligent agents don’t have a “shared understanding” of the situation with the human commander. Research breakthroughs are required that will help an intelligent agent to share a common “frame of reference” about the environment and the operational context with the human. For example, the robot should know if the current order should be executed in a normal or urgent manner.

Often, AI and ML algorithms can be used to make decisions about or classify a situation but are not able to explain what particular attributes drove the algorithm to come to that decision. This ability to “explain” is lacking in most AI and ML methods.

Not only do intelligent agents such as robots lack common sense, they rarely have the ability to explain their decisions or even communicate in a straightforward way with their commanders. If these systems could use natural language communication and have the ability to query the commander for clarification this would provide for smoother more seamless operation in the field and even lead to the possibility of an intelligent agent to “explain” their decision process and let the human commander know how a decision was reached or the agent could seek clarifying information.

Furthermore, AI and ML methods tend to be “brittle” and experience severe failures that one would not expect from a human being. For example, there are many papers that demonstrate the apparent ease of tricking CNNs into thinking an image belongs to a completely different unrelated class. For example, in one paper an image that is obviously a “school bus” with a little bit added, but imperceptible pixel noise, is classified as an “ostrich” [18]. For a military context, this would be a catastrophic failure and a mistake that is extremely unlikely for a human to make. The fact that a few pixels in the image need to be manipulated in order to fool the CNN is troubling and represents a vulnerability that adversaries could exploit. This issue is exacerbated due to the inability of the AI/ML to explain how it came to a final classification decision thereby making it virtually impossible to understand why the image was incorrectly classified.

4.6 Summary of Gaps

Thus far we have posited 5 AI and ML gaps for military relevant operations. The rest of the paper reviews some of the basic research that is being undertaken by ISD in order to address these challenges. To review, the gaps discussed were:

1. There is a lack of large scale, labeled military relevant datasets upon which commercial AI and ML methods rely.
2. AI and ML methods often are required to run in a “batch-based” learning mode which need to process an entire data set at the same time to produce a useful model.
3. With many commercial AI and ML methods there is an over reliance on supervised learning techniques which requires the training data set to have complete and accurate labels.
4. AI and ML systems struggle to operate autonomously and reason about the environments in which they operate. Furthermore, it is difficult for intelligent agents to learn from their environment.
5. Intelligent agents do not effectively interact and communicate with humans because they lack shared understanding of the world and “common sense”. AI and ML systems tend to be non-transparent, “black boxes” that lack the ability to “explain” their decisions while intelligent agents (robots) require constant human attention and control. Many AI and ML methods have brittle and unpredictable performance in the presence of noise and adversaries.

5. BASIC RESEARCH TO ADDRESS COMMERCIAL AI & ML SHORTCOMINGS

ISD’s fundamental research program spans several areas executed by 3 distinct branches, Battlefield Information Processing (BIP), Multilingual Computing and Analytics (MCA), and Asset Control and Behavior (ACB) [19]. The BIP branch of the ISD conducts basic and applied research in supporting technologies for information superiority, global

connectivity, and enhanced situational awareness for the warfighter at the tactical edge of the conflict. Research concentrates on the adaptation and development of technologies to mitigate the unreliability of the current tactical network through the introduction of middleware, which acts as a mediator between the software applications and the network services. In addition, the BIP branch develops innovative methods to exploit the quantitative and qualitative nature of the expanded data available within the tactical environment. The MCAB branch conducts basic and applied research in supporting technologies that automatically process and interpret human language input. Research concentrates on the recognition of speech and document images, in both English and foreign languages, and military-specific machine translation for low density foreign languages. In addition, the MCA branch develops social network analysis (SNA) techniques to map and measure the relationships, and flows between people, groups, organizations, computers, and other information or knowledge processing eEntities encountered within the tactical environment. The ACB branch conducts basic and applied research in supporting technologies such as AI and ML for the control of networked assets (i.e., robots) and their corresponding behaviors. ACB's research concentrates in the areas of autonomous mobility for small unmanned systems and tactical behaviors for groups of heterogeneous assets in accomplishing a coordinated task. This research enables dismounted operations of small scale robots in urban and complex terrains.

In the rest of this section we focus on specific research programs that address the 5 gap areas identified previously in section 4. In response to the 5 gap areas, we have identified several research thrusts to address each gap in ISD's ongoing research portfolio.

5.1 Collection of Relevant Datasets (when possible)

Although the US military collects and possesses a lot of data, it is very challenging to label enough data to train AI and ML algorithms thoroughly such that systems will operate effectively in military scenarios. Perhaps the most famous example of data collection in the commercial world is ImageNet, with over 14 million images that have been labeled by humans [4, 9]. Although data collection is difficult and some question whether it is even possible to collect enough data to feed state-of-the-art algorithms such that they would perform well under militarily operational conditions, the DoD is working to do just that. In April 2017, Robert Work, former Deputy Secretary of Defense, announced the formation of Project Maven and established a an Algorithmic Warfare Cross-Functional Team (AWCFT) that is working towards fielding technology to augment and automate the processing, exploitation, and dissemination of data for tactical Unmanned Aerial System (UAS) and Mid-altitude Full-Motion Video (FMV). [20,21] In addition to efforts to automate the labeling of data, the AWCFT enlisted ARL's ISD and other ARL researchers to advise on automatic labeling techniques and to help consolidate existing algorithm-based technology initiatives related to mission areas of the Defense Intelligence Enterprise, including initiatives that develop, employ, or field artificial intelligence, automation, machine learning, deep learning, and computer vision algorithms. The overarching objective of Project Maven is to "turn the enormous volume of data available to DOD into actionable intelligence and insights at speed."

ARL also maintains internal efforts to collect data for research and to train AI and ML algorithms through various research programs. One such facility is the Intelligent Systems Center (ISC), which is part of ARL's Open Campus initiative with the goal of will facilitating innovation and encouraging cross-disciplinary research with a focus on long-term basic and applied research. The Center will leverage the strength of its current research program by focusing on systems that interact with the physical world [22]. The ISC is a highly collaborative environment with cross-discipline opportunities spanning many areas touching AI and ML including traditional robotics (intelligence, perception, and mobility/manipulation), adaptive control, autonomous networking, cognitive architectures, natural language, semantics and reasoning with facilities for test, evaluation, validation and verification. The ISC has many Robotic experimental platforms (ground and air) as well as an intelligent algorithm software repository, unique military datasets and simulation tools as well as unique facilities for experimentation for ground- and air-based platforms.

In addition, ARL hosts the Automated Online Data Repository (AODR) and Sensor Information Testbed for Collaborative Research Environment (SITCORE), which are both used by ISC researchers. ARL and other government organizations, industry, and academia use AODR to store and disseminate multiple intelligence (Multi-INT) datasets collected at field exercises and demonstrations, facilitates research and advancement of analytical tools and algorithms supporting the Intelligence, Surveillance, and Reconnaissance (ISR) community. SITCORE is a virtual research development environment allowing collaboration from other locations, including DoD, industry, academia, and coalition partners. SITCORE combined with AODR provides a comprehensive end-to-end algorithm development, experimentation, demonstration, and validation platform [23].

5.2 Online, Anytime, Continuous Learning and Learning with Sparse Examples

Notwithstanding current Army efforts to collect and label large amounts of relevant data, there remain many situations where collection of adequate data required by many of the mainstream AI and ML algorithms will not be possible. For example, operations in enemy territories that are very dissimilar in nature to the archived environmental data in the Army's repositories could be hindered if modern AI and ML methods are utilized. Creating scenarios representative of the wide variety of operational situations that could be encountered by an intelligent system on some future mission is just not possible. Some early research is underway to simulate data in order to generate large and rich training sets that are derived from existing data. Although promising, it is not yet clear whether or not this approach will provide examples that will enable algorithms to perform well under operational conditions. From ISD's perspective, what is required are new algorithms that are capable of operating simultaneously in online learning and operating mode, receiving data one sample at a time and processing a small subset of the total observed set of samples without the need for the entire dataset to be present for training, as is the case for batch mode learning.

ISD has recently published research on a new algorithm called Parsimonious Online Learning with Kernels (POLK) that is an example of such an algorithm [24]. POLK proposes the use of nonparametric machine learning, which tends to be more accurate than parametric methods and allows for a more flexible hypothesis that can adapt to incoming data. An example of a parametric method, linear regression, is given in Equation 1.

$$\hat{y} = \beta_0 + \beta_1 * x + \epsilon \quad (1)$$

Note that with parametric models we must make specific assumptions about the data and the parameter selection will greatly impact how well the model performs. Some examples of parametric machine learning algorithms are [25]:

- Logistic regression
- Linear discriminant analysis
- Naive Bayes
- Simple neural networks

Parametric machine learning methods, albeit fairly simple, offer a constrained model representation and this limits their application to complex and noisy environments. On the other hand, nonparametric methods offer several advantages in that they make no a priori assumptions about the data and, therefore, are capable of providing a more flexible model that is more accurate in complex scenarios. However, models built from nonparametric methods tend to grow in size as they are exposed to more data, requiring a large amount of storage. Because of this, most ML practitioners do not apply nonparametric methods to large online learning applications. Typical examples of nonparametric ML methods are:

- k-Nearest neighbors
- Decision trees
- Support vector machines

With nonparametric methods, the model complexity grows with each additional data point such that the complexity is proportional to the iteration index and, therefore, does not readily lend itself to online learning; many techniques such as support vector machines rely on batch data.

POLK was specifically developed to address the complexity and memory requirements challenges of nonparametric approaches in order to operate in true streaming data scenarios, attaining state-of-the-art accuracy tradeoffs between model complexity, statistical accuracy, and stability, both for statistical inference (supervised learning) and learning-based control (reinforcement learning). This is accomplished by utilizing a stochastic gradient descent approach that uses a supervised "sparsification" scheme, which only keeps the data points that minimize the statistical loss. The POLK authors consider this to be a "greedy function subspace projection," where, for each gradient step, data samples are either retained or discarded depending on their ability to minimize the statistical loss, resulting in a technique better suited than traditional parametric techniques for dynamic environments with limited data.

For Army-relevant scenarios, POLK and similar approaches to machine learning offer unique advantages. First, online learning is important because underlying distributions of observed data can change as events unfold. Second, in a data-

impoverished environment, the ML algorithm may only observe a few examples of a phenomenon before needing to adapt to it.

POLK has a strong mathematical basis and is derived from prior online learning approaches such as Kivinen et. al’s Naïve Online Regularization Minimization Algorithm (NORMA), which uses a fixed number of data points and discards the oldest point [26]. Kivinen et. al. were the first to demonstrate how the application of classical stochastic gradient descent can lead to novel and practical algorithms for online learning using kernel methods. Subsequently, Wang et. al. introduced Budgeted Stochastic Gradient Descent (BGSD), which addresses the “curse of kernelization”, a phenomena that causes unbounded linear growth in model size and update time with data size [27]. With BGSD, any number of underlying learning algorithms can be used (such as NORMA) with a fixed number of data points, but the algorithm discards the data point with the smallest weight and “reprojects” weights onto the remaining data points. These memory reduction methods may cause learning instability due to the fact that the discarded data is not based on learning progress.

POLK provides an alternative method to discarding data that takes into account the value of the observed data and determines which data points are the most informative. This is shown pictorially in Figure 2.

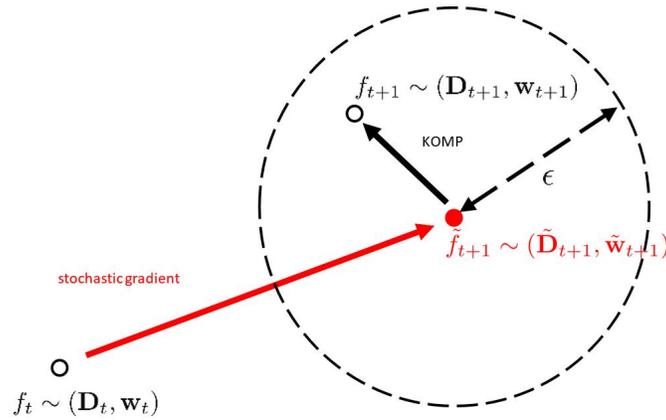


Figure 2. Visual depiction of POLK where epsilon is the size of the error neighborhood and f_t is the current data point, \tilde{f}_t is the estimate of the next SGD step and f_{t+1} is the POLK estimate. D_t is the kernel dictionary. The function f_t is parameterized by dictionary D_t and weight vector w_t . f_{t+1} is determined from output of DKOMP on \tilde{f}_{t+1} which is parameterized by $(\tilde{D}_{t+1}, \tilde{w}_{t+1})$ [18].

The stochastic projection is defined as follows (where P is the Project function) in Equations 2 and 3:

$$f_{t+1} = P_{H_{D_{t+1}}} [(1 - \eta_t \lambda) f_t - \eta_t \nabla_f \ell(f_t(x_t), f y_t)] \quad (2)$$

Where:

$$H_{D_{t+1}} = \text{span}\{k(\mathbf{d}_n, \cdot)\}_{n=1}^{M_{t+1}}. \quad (3)$$

Note that in in Equation 2, the term $P_{H_{D_{t+1}}}$ is the projection onto the most representative data points, while η is the learning rate and λ is the regularization parameter. Also note that $[(1 - \eta_t \lambda)]$ can be considered as a “forgetting factor”. $\eta_t \nabla_f \ell(f_t(x_t), f y_t)$ is the gradient step and ℓ represents the loss function. The estimated value for

$$\tilde{f}_{t+1} = (1 - \eta_t \lambda) f_t - \eta_t \nabla_f \ell(f_t(x_t), f y_t) \quad (4)$$

$$\tilde{D}_{t+1} = [D_t, \mathbf{x}_t] \quad (5)$$

$$\tilde{w}_{t+1} = [(1 - \eta_t \lambda) \mathbf{w}_t, -\eta_t \ell'(f(\mathbf{x}_t, \mathbf{y}_t))]. \quad (6)$$

Note that without using projection D_t is merely memorizing the most recent data point \mathbf{x}_t .

Projection is carried out by using a novel variation of Kernel Orthogonal Matching Pursuit, dubbed here as *Destructive Kernel Orthogonal Matching Pursuit* (DKOMP). Simply put, for each data point, DKOMP discovers the approximation error of the best function approximation if this point were removed and removes the point that incurs the least error. This process is repeated until the approximation error budget is reached as shown in Equation 7. See [18] for more details:

$$[f_{t+1}, \mathbf{D}_{t+1}, \mathbf{w}_{t+1}] = \text{DKOMP}(\tilde{f}_{t+1}, \tilde{\mathbf{D}}_{t+1}, \tilde{\mathbf{w}}_{t+1}). \quad (7)$$

Fundamentally, POLK leverages greedy compression to define a projection onto most representative observations.

Koppel et. al. proved that this approach has a guaranteed bounded model order. POLK has been shown to work well with synthetic and real datasets. For example, when run with the MNIST handwritten character dataset, the following results were obtained:

- libsvm (*batch*): 1.5% test error, ~15,555 model points
- POLK (*online*): 4.5% test error, ~450 model points

The end result is optimally compressed online kernel learning algorithms that offer an alternative to deep learning and overcomes their data hunger as well as cloud reliance. Thus, these methods will be vital to the development of intelligent systems that operate in communications-denied settings without much a priori information.

5.3 Semi-Supervised and Unsupervised Learning

In addition to developing techniques for online supervised learning, ISD is making strides in developing new AI and ML methods that operate in an unsupervised manner and label data collected by agents operating in an unknown environment. Although many AI and ML methods expect virtually perfectly labeled training data, there are many challenges with creating correctly labeled datasets including expense and reliability [28]. Data labeling is expensive because it is time-consuming and often requires experts to perform the labeling. Labeled datasets often contain mislabeling errors called “label noise”. Label noise comes from 4 main sources [28]:

- Inadequate or unclear description of labeling task or poor data quality.
- Experts often make errors during labeling.
- Classification can sometimes be subjective in nature so even experts may not always agree.
- Dataset encoding errors.

Labeling ambiguity issues also derive from changes in illumination, scale, perspective, color, and background clutter [29]. To address these issues ISD researchers Wigness et. al. have developed novel methods to hierarchically cluster image data in such a way that a wide “spectrum of visual concept granularities” is effectively captured in a way that greatly increases the efficiency of labeling data and reduces label noise [29]. This work was subsequently expanded to develop a comprehensive framework intended to more effectively label datasets by focusing on reducing human labeling effort using 4 metrics—discovery, efficiency, exploitation, and accuracy—by leveraging hierarchical clustering, a novel “interestingness” metric, and an iterative group-based labeling process based upon relationships between labeled and unlabeled data [30]. This new framework dubbed “Hierarchical Cluster Guided Labeling” (HCGL) requires no a priori knowledge of the label set and, unlike other methods, requires no re-clustering for refinement of groupings, which reduces overall process latency. In HCGL, discovery is the process of learning “concepts” or categories for images. In other words, the algorithm determines what labels belong on each image by setting up a taxonomy. For example an image could be clustered into a concept for “animals” while also belonging simultaneously to “dogs”. Ultimately, this improves learning performance during classification training. Labeling efficiency refers to the extent to which the human effort for labeling data can be reduced. This is determined by accounting for both the effort and time it takes humans to assist in labeling images. The idea behind exploitation is measuring the number of images that can be labeled with a single label while accuracy is the proportion of non-erroneous labels [30].

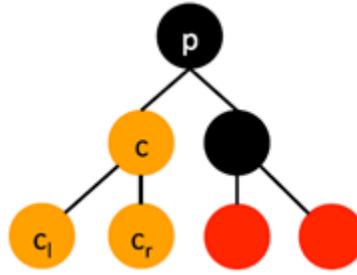


Figure 3. Illustration that depicts the relationships for group c in a local neighborhood of H , including its parent p and left and right children, c_l and c_r . [30]

HCGL utilizes several “selection heuristics,” including interestingness, exploration, and exploitation, in order to make group selections in the labeling process. Interestingness can be seen as the structural novelty within a local neighborhood of the hierarchy, while exploitation was discussed previously. Exploration is a measure that determines how similar a group is to previously labeled groups.

HCGL differs from other group-based labeling approaches in that the entire hierarchical clustering set of the training data is maintained and HCGL selects groups from this set, allowing for the encoding of many levels of “semantic granularity”. HCGL uses its novel hierarchical group selection criteria to quickly collect large amounts of labeled data, explore the feature space, and identify locations in the hierarchy that are more likely to contain groups of data that represent images from exactly one class. As compared with other group-based labeling techniques, HCGL produces higher performing visual classifiers with respect to labeling effort, which is very promising in for Army-relevant environments where time for labeling collected data is at a premium. The researchers also tested HCGL on a robot extending the algorithm so that the technique would achieve high quality visual perception requiring only a few hours of labeling effort from a human annotator [31].

While HCGL is a batch-based, semi-supervised learning algorithm, Wigness and Rogers developed a fully unsupervised, online semantic scene labeling algorithm called Unsupervised Semantic Scene Labeling (USSL) that continuously learns and adapts to incoming image data streams [32]. In this work, Wigness and Rogers address the shortcomings of batch-based processing, which prohibits online learning and discovery of novel concepts in incoming visual data streams. They also point out that even semi-supervised methods introduce processing delays as humans are still required to label some data. Unlike earlier segmentation preprocessing algorithms, USSL directly models scene semantics with regard to pixel connectivity in localized regions of the image.



Figure 4. Comparison of segmentation output for our technique which directly models semantics without regard to locality, and hierarchical graph-based (GBH) segmentation which emphasizes segmentation coherency using strict locality rules. USSL, ARL’s approach labels all four traffic cones consistently (column two), whereas GBH assigns unique labels to each cone (column three) [32].

Note that in Figure 4, USSL is capable of identifying the 4 cones as belonging to the same class, as denoted by the purple color, whereas the hierarchical graph-based segmentation method defines each cone as belonging to different classes. USSL uses a technique called “Agglomerative Clustering,” which is a “bottom-up” approach when each object is assigned to a single cluster and then is successively merged until a desired cluster structure is obtained based on some predefined criteria [33,34]. Therefore, USSL iteratively creates and adapts a set of semantic models from streaming image data without a priori knowledge of image classes or any parameters of the data and is, therefore, capable of modeling novel

images and concepts continuously. One of the key innovations in USSL is the use of overlapping sliding windows to create an ensemble of learners that are encoded in a graph structure and used to reconcile learning errors and generate label mappings and a global label set for the data stream, as depicted in Figure 5. According to the researchers, by “over-learning locally, noise introduced by unsupervised learning can be minimized, and remaining errors can be reconciled using an ensemble of local learners encoded in a graph-based structure.”

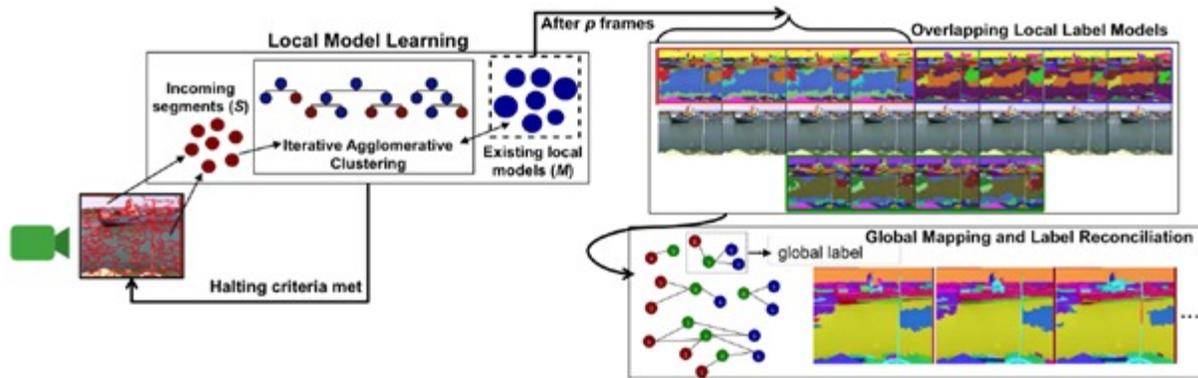


Figure 5. Overview of the USSL algorithm. The next image from the data stream is over-segmented, and segments are agglomeratively merged with existing models from previous frames in the stream. Overlapping local models are created for a window in the stream and these local label sets are mapped and reconciled using a graph encoding to generate a global label set. [32]

5.4 Autonomous and Semi-Autonomous Operation

As discussed previously, it is common in the field of AI and ML to train learning agents by large numbers of training examples but this is not feasible for many Army scenarios. Even with recent applications of AI and ML, such as self-driving vehicles, miles of miles of data has been collected but this may not be adequate. As Knox and Stone point out, “an autonomous driving agent should not learn to drive by crashing into road barriers and endangering the lives of pedestrians.” [35] Furthermore, these researchers point out the humans have domain knowledge that can be leveraged to increase the speed and learning while avoiding the need for complex or difficult to create training examples. Knox and Stone’s framework is called “Training an Agent Manually via Evaluative Reinforcement” (TAMER) where a human trainer to interactively “shapes” an agent’s policy using positive and negative reinforcement signals [35].

ARL researchers Warnell, Waytowich, and Lawhern collaborated with one of TAMER’s creators—Stone—to extend TAMER to incorporate the use of high-dimensional state spaces and deep learning, resulting in Deep TAMER [36]. Whereas with TAMER, the agent receives signals from the non-expert human and the agent, in response, tries to estimate a hidden function that approximates the human’s feedback and adjusts its behavior to align with this estimate, Deep TAMER uses a deep neural network that is used to compute an estimate that minimizes the value of a cost function. In addition, TAMER learns once from each state-action pair, unlike Deep TAMER, which can learn multiple times for each state-action pair due to a built-in feedback replay buffer [36]. The feedback replay buffer, D increases the rate of learning by storing observed human feedback and sampling this running history at a fixed rate regardless of the rate of human feedback. That is, Stochastic Gradient Descent (SGD) updates to \hat{H} occur whenever the human provides new feedback and at a fixed rate using the data sampled from D . The Deep TAMER framework is depicted in Figure 6, and the initial research specifically considered the Atari game of BOWLING, which uses a pixel-level state space.

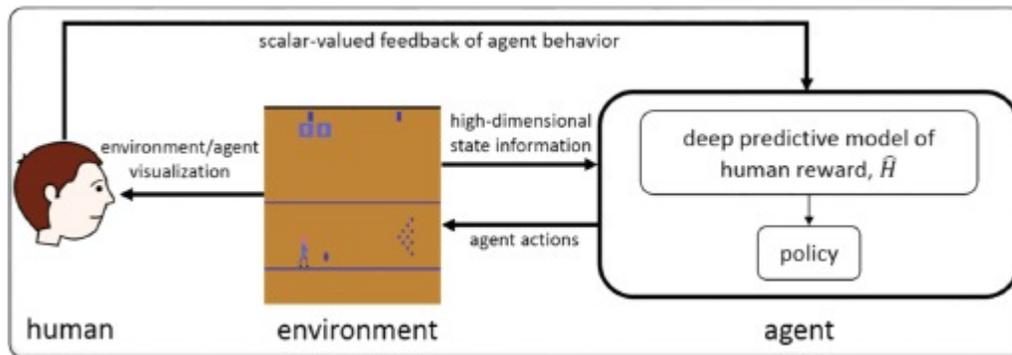


Figure 6. The Deep TAMER framework. A human observes an autonomous agent trying to perform a task in a high-dimensional environment and provides scalar-valued feedback as a means by which to shape agent behavior. Through the interaction with the human, the agent learns the parameters of a deep neural network, \hat{H} , that is used to predict the human's feedback. This prediction then drives the agent's behavior policy. [36]

In summary, Deep TAMER leverages the representational power of deep neural networks in order to learn complex tasks in just a short amount of time with a human trainer. ARL's researchers demonstrated that Deep TAMER is able to exceed the performance of average human trainers after only 7 minutes of training, and surpasses the performance of the expert human after 15 minutes. This was accomplished on the Atari game of BOWLING a task that has proven difficult for even state-of-the-art reinforcement learning methods [36].

5.5 Intent Awareness, Reasoning and Explainability

5.5.1 Natural Language Communication between Humans and Agents

In order to address the issues related to common sense, reasoning, and explainability, ARL has a research thrust that is looking at how humans can communicate with intelligent agents using natural language. The capabilities that will be derived from this research are useful in a number of ways. For example, if an intelligent agent is given an order such as "go through that door", it would be useful if it could respond with a query such as "Do you mean the door that is closest to me?" or "Do you mean the door on the left?" This grounding of the request is important and very natural when humans communicate with each other. ARL's "BOT Language" research program (BOTLang) aims to provide more natural ways for soldiers to interact and communicate with robots [37,38]. Current research involves using a "Wizard of Oz" approach that attempts to simulate capabilities of autonomous systems that we envision will exist in the future. This research not only informs researchers of human-system integration issues but also provides insights as to how autonomous systems should be designed.

BOTLang uses 2 wizards, the first is for robot navigation and the second is for dialog management. Experiments have the system explore an unknown environment with a human teammate, where the human gives verbal instructions to the robot from a remote location and the robot takes action based on the human requests, clarifying instructions through queries back to the human. The architecture is depicted in Figure 7.



Fig. 7. The participant provides instructions, unaware that the robot is composed of two wizards. The Dialog Manager (DM) wizard translates spoken language commands into a constrained action list for the Robotic Navigator (RN) wizard to execute. Map updates, image data, and the robot's dialogue responses are sent to the participant through a multimodal interface (upper right) [37,38].

The BOTLang architecture is very modular. The Robot Navigator (RN) can simply be thought of as autonomous navigation that leverages all available sensor data, while the Dialog Manager (DM) does all of the natural language processing and has the intelligence to translate verbal commands given by the participant into movement orders that will be understood by the RN. The BOTLang research has produced a corpus of data that is useful for training the AI and ML research algorithms that are being developed, and will eventually replace the RN and DM human “wizards”.

5.5.2 Common Sense Reasoning with Robustness

Often researchers will attempt to imbue intelligent agents with “common sense” by using a “knowledge base” (KB). A KB is a set of representations of objects in the world. An intelligent agent leverages a KB and an inference engine or inference algorithm such that when an agent is asked a question it can use its KB and inference engine to answer. [39]

ISD researcher Summers-Stay has noted that most common sense KBs are “brittle” and “incomplete” in that they are sensitive to how queries are posed and they are missing seemingly trivial knowledge facts [40]. In order to improve overall performance, research has been conducted that combines the deductive reasoning abilities of a KB inference engine with the ability to represent semantic similarity that is provided by “distributional semantic vector spaces”. This amounts to exploring the KB for semantically related terms to the query term. However, these systems sometimes do not result in an answer. To remedy this, Summers-Stay has developed a method for forming chains of “deductive logical reasoning,” which is akin to human “train of thought.”

Instead of attempting to find connective “paths between ideas” using a mathematical search technique to exhaustively search a knowledge graph, Summers-Stay uses a KB consisting of high dimensional vectors that are related to each via a predicate in form a triple $(e_n, \text{predicate}, e_m)$. Let us imagine that there are 2 words, g and p , and you need to determine if g implies p . The goal becomes finding a chain of entities that links g to p , which means finding a sum of “fact vectors” $(-e_n + e_m)$ that add up to $(-g + p)$. As it turns out, in cases where such a chain exists, this method returns a set of facts that constitute the chain, and even when the the chain does not exist the method will return an approximation of the correct links in the path leading to potentially new knowledge of undefined relations between closely related entities.

Figure 8 provides a visual and somewhat intuitive view of the method. One can envision words as “vectors”. The *word2vec* knowledge base that was utilized for this research has a 300-dimensional vector for each word. By leveraging

sparse approximation methods such as LASSO (least absolute shrinkage and selection operator) or Orthogonal Matching Pursuit (OMP), to obtain a sparse weight vector which is multiplied by the list of fact vectors in order to find a chain of entities linking g to p .

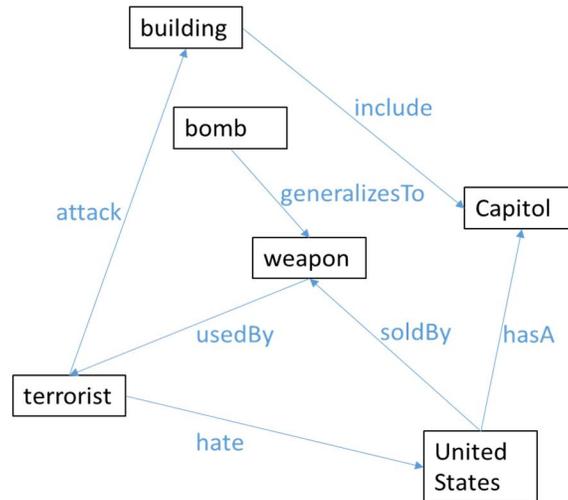


Fig. 8. Reasoning systems need to be able to give approximate answers when exact answers are impossible. Through sparse matrix techniques the intelligent agent is able to approximate reasoning with partial information and expand the knowledge base by exploiting distributional semantics at query time. In this simple example, terrorist is linked to Capitol [41].

An example from [40] is finding a path of relations between g : *Michael Jackson* and p : *Music*. The KB contains (Michael Jackson, *is a*, songwriter) and (musician, *composes*, music). Therefore, since songwriter and music are “nearby in semantic space,” therefore the connection between Michael Jackson and Music is made, demonstrating that this method can be used for analogical and associational connections. Another example is analogical and abductive reasoning where an analogy such as this is presented to the intelligent agent:

$$\text{bear:hiker::shark:X}$$

In this case the intelligent agent comes up with “snorkler”. How might this conclusion be made? Consider related contexts such as woods, sea, predator, and tourist. The vector for “bear” consists of a weighted average of all vectors related to “woods” and the average of all vectors related to “predators” plus noise. The analogy becomes:

$$\text{woods} + \text{predator} : \text{woods} + \text{tourist} + \text{noise}:: \text{sea} + \text{predator} : X$$

$$\text{while: snorkeler is: sea} + \text{tourist}$$

So the vector between bear and hiker is $-\text{predator} + \text{tourist} + \text{noise}$. Likewise the vector from shark to snorkeler is: $-\text{predator} + \text{tourist} + \text{noise}$. This is similar to the vector from shark to snorkeler. The algorithm looks for any vector that will make the sum have as low an error as possible and, therefore, is able to use analogical relations to complete a chain of argument and connect the concept shark to the concept snorkeler.

This method has also been extended to ontology merging, where simple questions can be answered by connecting relating concepts and even more complex questions containing multiple concepts.

This basic research is important not only because of its usefulness in helping intelligent agents to reason, but also because it can potentially explain “why” a certain conclusion was provided. Depending on how the system is constructed, it may be possible to ascertain how “certain” an intelligent agent is of its conclusion.

6. CONCLUSIONS AND FUTURE RESEARCH DIRECTIONS

This paper has reviewed some of the basic research being undertaken in ARL's Information Sciences Division to address the future needs of the US Army with regard to advanced AI and ML technologies, including relevant military datasets, new algorithms specifically targeted towards Army relevant scenarios such as dynamic, congested, and contested environments with little a priori data. The motivation for this research is derived from TRADOC, US Army priorities, gaps in the commercial AI and ML pipeline, and ARL's internal research focus areas. We briefly reviewed efforts by the US Military to collect and label relevant datasets with project Maven. However, with the realization that dataset collection and labeling is challenging, we discuss research that is making major inroads in the area of online, anytime, and continuous learning with the POLK algorithm with the ability to effectively operate with less data (i.e., sparse data environments). We also looked at ongoing research in unsupervised learning, which does not require labelled data with HGCL, and this was extended to an online learning mode with USSL. In order to achieve the goal of having fully autonomous intelligent agents in the future we reviewed successful efforts in reinforcement learning with a human in the loop called Deep TAMER to provide positive and negative signals to the agent in conjunction with a deep neural network. We also discussed strategies for creating training datasets via experiments for future systems under development for natural language communication between humans and robots with the BOTLang effort, which will eventually result in autonomous agents, such as robots working with soldiers, to execute tasks in the field. In addition to the ability for humans and robots to communicate, ARL researchers also engaged methods for automating reasoning with incomplete or noisy knowledge bases with sparse decomposition methods can be used to find chains of reasoning in a knowledge graph embedded in a distributional vector space, which can be extended to question-and-answer scenarios.

The surveyed research represents a small subset of ISD's and ARL's research portfolio but provides insight in the early successes and directions for AI and ML capabilities that will surely be derived from this work. Some challenges that remain and that are currently being explored are better methods for providing robust performance with noisy data, quantification of uncertainty, and dealing with adversarial attacks on data. Other ongoing efforts across ARL that have significant involvement of ISD researchers include the Robotics Collaborative Technology Alliance (RCTA) that has made significant inroads into perceiving and understanding dynamic and unknown environments, including the creation of a comprehensive model of the surrounding world [42]. The Distributed and Collaborative Intelligent Systems and Technology (DCIST) collaborative technology alliance is just getting underway and will explore AI and ML for operational effectiveness of intelligent system and soldier teams against dynamic threats in complex and contested environments, and provide technical and operational superiority through fast, intelligent, resilient and collaborative behaviors[43]. Finally, the Internet of Battlefield Things (IoBT) Collaborative Technology Alliance will develop a fundamental understanding of dynamically composable, adaptive, goal-driven IoBTs, enabling distributed analytics for intelligent command and control and battlefield services. IoBT will address unique Army challenges such as diverse and dynamic missions, rapid composition, and adaptability of battlefield devices with extreme heterogeneity that are forward-deployed in adversarial environments [44,45,46].

Future research will expand up the existing programs described here for more robust and autonomous learning approaches as well as a deeper focus on adversarial aspects of machine learning to answer fundamental questions about the resiliency of AI and ML methods in the face of real world disorder, confusion and noise, regardless of the source and whether natural or intentional.

ACKNOWLEDGEMENTS

The authors would like to thank ISD researchers Alec Koeppel, Maggie Wigness, John Rogers III, Douglas Summers-Stay, Garrett Warnell, Matthew Marge, and Ethan Stump for helpful conversations and explanations of their research.

REFERENCES

- [1] Kott, A., "The Artificial Becomes Real," Army Acquisition, Logistics and Technology Magazine, 18 January 18 2018, <http://asc.army.mil/web/2018/01/18/> (March 2018).
- [2] Hormozi, H., Hormozi E., and Nohooji, H. R., "The classification of the applicable machine learning methods in robot manipulators," Int J Mac Learn Comp 2.5, 560 (2012).

- [3] Wikipedia, "Machine Learning," Wikipedia. https://en.wikipedia.org/wiki/Machine_learning
- [4] Krizhevsky, A., Sutskever, I., and Hinton, G. E., "ImageNet classification with deep convolutional neural networks," *Advances in neural information processing systems*, (2012).
- [5] Munster, G., "Here's When Having a Self-Driving Car Will Be a Normal Thing," *Fortune Magazine*, 13 September 13 2017, <http://fortune.com/2017/09/13/gm-cruise-self-driving-driverless-autonomous-cars/>
- [6] US Army Training and Doctrine Command, "The US Army operating concept: Win in a complex world," US Army Training and Doctrine Command, Ft. Eustis, VA, TP: 525-3 (2014).
- [7] McNally, D., "Army Laboratory Shapes Future Research," 23 March 2017, <https://www.arl.army.mil/www/default.cfm?article=2977>
- [8] Milley, M. A., McCarthy, R. D., "Modernization Priorities for the United States Army," 3 October 2017, <https://admin.govexec.com/media/untitled.pdf>
- [9] Lawton, J., and Shields, L., "United States Army Mad Scientist: Megacities and Dense Urban Areas in 2025 and Beyond" Training and Doctrine Command (TRADOC) G-2, Fort Eustis, VA, (April 2016). <https://community.apan.org/wg/tradoc-g2/mad-scientist/m/mdua/170637/download>
- [10] Picatinny Arsenal Public Affairs, "Workgroup from across Army and beyond tackles problems of dense urban warfare," Picatinny Arsenal Public Affairs, (October 2017). https://www.army.mil/article/195475/workgroup_from_across_army_and_beyond_tackles_problems_of_dense_urban_warfare
- [11] ARCIC, "The Operational Environment and the Changing Character of Future Warfare," ARCIC, http://www.arcic.army.mil/App_Documents/The-Operational-Environment-and-the-Changing-Character-of-Future-Warfare.pdf (31 May 2017)
- [12] Pham, T., "ARO Workshop on Adversarial Machines Learning," Stanford University, 14 September 2017," <https://seclab.stanford.edu/AdvML2017/slides/pham.pdf>
- [13] ARL, "Then Internet of Battlefield Things," (IoBT) Program Announcement, <https://www.arl.army.mil/www/pages/3050/IOBT-Program-Announcement-AmendmentII.pdf> (7 February 2017)
- [14] Warnell, G., Waytowich, N., Lawhern, V., and Stone, P., "Deep TAMER: Interactive Agent Shaping in High-Dimensional State Spaces." (2017) arXiv preprint arXiv:1709.10163.
- [15] Ilachinski A., "AI, Robots, and Swarms: Issues, Questions, and Recommended Studies," CNA, DRM-2017-U-014796-Final, (2017), https://www.cna.org/cna_files/pdf/DRM-2017-U-014796-Final.pdf
- [16] Bean, R., "How Big Data is Empowering AI and Machine Learning at Scale," MIT Sloan Management Review, 8 May 2017, <https://sloanreview.mit.edu/article/how-big-data-is-empowering-ai-and-machine-learning-at-scale/>
- [17] Brynjolfsson, E., and McAfee, A., "The business of artificial intelligence," *Harvard Business Review* (2017).
- [18] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R., "Intriguing properties of neural networks," *ICLR 2014* (2014).
- [19] Computational and Information Sciences Directorate, US Army Research Laboratory, <https://www.arl.army.mil/www/default.cfm?page=30>
- [20] Work, R., "Establishment of an Algorithmic Warfare Cross-Functional Team (Project Maven)," https://www.govexec.com/media/gbc/docs/pdfs_edit/establishment_of_the_awcft_project_maven.pdf
- [21] Pellerin, C., "Project Maven Industry Day Pursues Artificial Intelligence for DoD Challenges," 27 October 2017, <https://www.defense.gov/News/Article/Article/1356172/project-maven-industry-day-pursues-artificial-intelligence-for-dod-challenges/>
- [22] ARL Open Campus, "Intelligent Systems Center (ISC) website," <https://www.arl.army.mil/opencampus/?q=centers/isc>
- [23] Ward, D. W., and Bennett K. W., "An overview of the US Army Research Laboratory's Sensor Information Testbed for Collaborative Research Environment (SITCORE) and Automated Online Data Repository (AODR) capabilities," *Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR VIII*, Vol. 10190. International Society for Optics and Photonics, (2017).
- [24] Koppel, A., et al, "Parsimonious online learning with kernels via sparse projections in function space," *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, (2017).
- [25] Brownlee, J., <https://machinelearningmastery.com/parametric-and-nonparametric-machine-learning-algorithms/>
- [26] Kivinen, J., Smola, A. J., and Williamson, R. C., "Online learning with kernels," *IEEE Trans. Sig. Process.* 52(8), 2165-2176 (2004).
- [27] Wang, Z., Crammer, K., and Vucetic, S., "Breaking the curse of kernelization: Budgeted stochastic gradient descent for large-scale svm training," *J. Mac. Learn. Res.* 13, 3103-3131 (October 2012).

- [28] Frénay, B., and Kabán, A. “A comprehensive introduction to label noise.” ESANN, April 2014.
- [29] Wigness, M., Draper, B. A., and Beveridge, R. J., “Efficient label collection for unlabeled image datasets.” Proc. IEEE Conf. Computer Vision and Pattern Recognition, 4594-4602 (2015).
- [30] Wigness, M., Draper, B. A., and Beveridge, J. R., “Efficient Label Collection for Image Datasets via Hierarchical Clustering,” *Int. J. Comp. Vision* 126(1), 59-85 (2018).
- [31] Wigness, M., Rogers, J. G., Navarro-Serment, L. E., Suppe, A., and Draper, B. A., “Reducing adaptation latency for multi-concept visual perception in outdoor environments.” *Int. Conf. Intelligent Robots and Systems (IROS)*, 2016 IEEE/RSJ, 2784-2791 (October 2016).
- [32] Wigness, M., and Rogers, J. G., “Unsupervised Semantic Scene Labeling for Streaming Data,” Proc. IEEE Conf. Computer Vision and Pattern Recognition 4612-4621 (July 2017).
- [33] Sasirekha, K., and Baby, P., “Agglomerative hierarchical clustering algorithm-a,” *Int. J. Sci. Res. Pubs.* 83 (2013).
- [34] Popat, S. K., and Emmanuel, M., “Review and comparative study of clustering techniques,” *Int. J. Comp. Sci. Info. Tech.* 5(1), 805-812 (2014).
- [35] Knox, W. B., and Stone, P., “Interactively shaping agents via human reinforcement: The TAMER framework,” Proc. 5th Int. Conf. Knowledge capture, ACM, 9-16 (September 2009).
- [36] Warnell, G., Waytowich, N., Lawhern, V., and Stone, P., “Deep TAMER: Interactive Agent Shaping in High-Dimensional State Spaces,” arXiv preprint arXiv:1709.10163, (2017).
- [37] Marge, M., Bonial, C., Byrne, B., Cassidy, T., Evans A. W., Hill S. G., and Voss, C., “Applying the Wizard-of-Oz Technique to Multimodal Human-Robot Dialogue,” Proc. IEEE RO-MAN, RO-MAN 2016, (2016).
- [38] Bonial, C., Marge, M., Artstein, R., Foots, A., Gervits, F., Hayes, C. J., Henry, C., Hill S. G., Leuski, A., Lukin, S. M., Moolchandani, P., Pollard K. A., Traum, D., and Voss, C. R., “Laying Down the Yellow Brick Road: Development of a Wizard-of-Oz Interface for Collecting Human-Robot Dialogue.” Proc. AAI Fall Symp. Series: Natural Communication for Human-Robot Collaboration, AAI FSS 2017, (2017).
- [39] Russell, S. J., and Norvig, P., [Artificial intelligence: a modern approach], 2nd Edition, Pearson Education, Inc., Upper Saddle River; New Jersey, (2003).
- [40] Summers-Stay, D., “Deductive and Analogical Reasoning on a Semantically Embedded Knowledge Graph,” *Int. Conf. Artif. Gen. Int.* Springer, Cham, 112-122 (August 2017).
- [41] Summers-Stay, Douglas. (2018, January). Deductive and Analogical Reasoning on a Semantically Embedded Knowledge Graph. Southwest Robotics Symposium, Phoenix, AZ.
- [42] RCTA, https://www.arl.army.mil/www/pages/392/New_Robotics_CTA.pdf
- [43] DCIST, <https://www.arl.army.mil/www/pages/3049/DCIST%20Opportunities%20Day%203-1-2017%20Piekarski.pdf>
- [44] IoBT, https://www.arl.army.mil/www/pages/3050/IOBT_CRA-Overview_20170327.pdf
- [45] “Team APG Theater: Internet of Battlefield Things,” MILCOM Panel, 24 October 2017, <http://events.afcea.org/milcom17/public/SessionDetails.aspx?FromPage=Sessions.aspx&SessionID=6327&SessionDateID=480>
- [46] Illinois leads \$25 million alliance to develop Internet of Battlefield Things,” <https://ischool.illinois.edu/articles/2017/10/illinois-leads-25-million-alliance-develop-internet-battlefield-things>