Poster Session 1: Data Efficient and Computationally Efficient Machine Learning

Title: Memory Efficiency of Matricies
PoC: Chien-Cu Chen

Title: Diastolic Arrays: Efficient Neural Network Inference Acceleration
PoC: Michael Mishkin and Mikko Lipasti
Abstract: The vast majority of neural network operations are the multiplication and accumulation associated with dot-product calculation. Neural network acceleration based on diastolic arrays facilitates energy-efficient neural network inference acceleration based on systolic arrays with shallow pipelines of complex cells each containing multiple multiplier units and an adder tree to perform partial reductions. These pipelines contain fewer flip-flops than conventional systolic array implementations of matrix-multiplication, which leads to substantial energy savings. Further performance improvements are achieved due to lower latency propagation through shallower pipelines, but this latency reduction is easily overshadowed by bandwidth limitations. Performance is further improved by operating multiple smaller diastolic array tiles in parallel in order to improve array utilization. The increased power consumption from tiling is offset by diastolic array power savings resulting in an optimal energy-delay product when combined.

Title: Student Acoustically grounded word embeddings for improved acoustics-to-word speech recognition
PoC: Shane Settle

Title: Student Multi-view representation learning for sequences
PoC: Qingming Tang

Title: Semantic query-by-example speech search using visual grounding
PoC: Herman Kamper, Aristotelis Anastassiou, Karen Livescu
Abstract: Learning in the presence of an additional modality can be a very successful form of weakly supervised learning. Here we investigate how to learn speech models from untranscribed speech by leveraging accompanying images at training time. We apply the idea to query-by-example (QbE) search, the task of retrieving utterances relevant to a given spoken query. We are particularly interested in semantic QbE, where the task is not only to retrieve utterances containing exact instances of the query, but also utterances whose meaning is relevant to the query. We follow a segmental QbE approach where variable-duration speech segments (queries, search utterances) are mapped to fixed-dimensional embedding vectors. Using an embedding function trained on visually grounded speech data, this approach outperforms a purely acoustic QbE system in terms of both exact and semantic retrieval.

Title: Context aware zero shot recognition
PoC: Ruotian Luo
Abstract: We present a novel problem setting in zero-shot learning, zero-shot object recognition and detection in the context. Contrary to the traditional zero-shot learning methods, which simply infers unseen categories by transferring knowledge from the objects belonging to semantically similar seen categories, we aim to understand the identity of the novel objects in an image surrounded by the known

objects using the inter-object relation prior. Specifically, we leverage the visual context and the geometric relationships between all pairs of objects in a single image, and capture the information useful to infer unseen categories. We integrate our context-aware zero-shot learning framework into the traditional zero-shot learning techniques seamlessly using a Conditional Random Field (CRF). The proposed algorithm is evaluated on both zero-shot region classification and zero-shot detection tasks. The results on Visual Genome (VG) dataset show that our model significantly boosts performance with the additional visual context compared to traditional methods.

Title: Deep Learning for Organometallic Compound Property Prediction
PoC: Nathaniel Krakeur
Abstract: Currently there is no practical and cost-effective way to explore the properties of a significant fraction of the possible molecular systems through chemical synthesis. Machine learning allows chemists to predict properties of molecules in silico to efficiently guide design of optimal materials. We use deep learning methods from graph methods and train the model on a machine readable chemical notation (SMILES string) representation of molecules. We are presently developing a database and model to accelerate property exploration of organic molecules, starting with flash point. We are assessing the ability of SMILES based machine learning models to predict this property across a wide-range of different molecule types. We are also exploring the ability to predict properties of the specific class of organosilicon molecules, which are of interest as electrolyte solvents possessing attractive properties for Li-ion batteries.

Data Efficient Machine Learning Thrust

Leads: Robert Nowak UW-Madison, Lee Seversky AFRL/RI, Walter Bennette AFRL/RI, Greg Shakhnarovich TTIC

Title: Adaptive Online Learning, Monitoring, and Sampling for Big Data Streams
PoC: Xiochen Xian
Abstract: With the rapid advancement of sensor technology, a huge amount of streaming data is generated in various applications, which poses new and unique challenges for the online learning, monitoring, and sampling of such data. We are interested in big data analytics for offline learning and modeling, as well as online monitoring general heterogeneous big data streams in the context of limited resources, where only a subset of observations are available at each acquisition time. In particular, such methods integrates rank-based CUSUM, data augmentation schemes, and innovative ideas that incorporates partial observations, which can effectively detect a wide range of possible mean shifts for non-normal data streams. Two theoretical properties on the sampling layout of the proposed algorithm are investigated when the process is in control and out of control. Both simulations and case studies are conducted under different scenarios to illustrate and evaluate the performance of the proposed method.

Title: Learning Nearest Neighbor Graphs from Noisy Distance Samples
PoC: Blake Mason
Abstract: We consider the problem of learning the nearest neighbor graph of a dataset of n items. The metric is unknown, but we can query an oracle to obtain a noisy estimate of the distance between any

pair of items. This framework applies to problem domains where one wants to learn people's preferences from responses commonly modeled as noisy distance judgments. In this paper, we propose an active algorithm to find the graph with high probability and analyze its query complexity. In contrast to existing work that forces Euclidean structure, our method is valid for general metrics, assuming only symmetry and the triangle inequality. Furthermore, we demonstrate efficiency of our method empirically and theoretically, needing only $O(n \log(n)\Delta-2)$ queries in favorable settings, where $\Delta-2$ accounts for the effect of noise. Using crowd-sourced data collected for a subset of the UT Zappos50K dataset, we apply our algorithm to learn which shoes people believe are most similar and show that it beats both an active baseline and ordinal embedding.

Title: Fast convergence for mini-batch stochastic gradient descent with adaptive batch sizes
PoC: Scott Sievert
Agenda: Abstract: Stochastic gradient descent (SGD) approximates the objective function's gradient with a constant and typically small number of examples a.k.a. the batch size of mini-batch SGD. Small batch sizes can present a significant amount of noise near the optimum. This work presents a method to grow the batch adaptively with model quality that requires no more computation than standard SGD. With this method, convergence is significantly improved for strongly convex, convex and a certain class of non-convex functions in terms of the number of model updates. This method is easier to tune because the hyper-parameters have no time dependence and is more amenable to distributed systems. Simulations and experiments are performed to confirm and extend theoretical results.

Title: Verification Guided Tree Search for Deep Reinforcement Learning
PoC: Daniel Melcer
Abstract: The field of reinforcement learning has been revolutionized in recent years, due in part to the mass adoption of deep convolutional neural networks and the resurgence of powerful methods to refine decision-making policies. However, the problem of sparse reward signals remains pervasive in many non-trivial domains. While various reward-shaping mechanisms and imitation learning approaches have been proposed to mitigate this problem, a mathematically rigorous structure of the underlying objective is rarely exploited. In this paper, we resolve this by representing objectives using temporal logic and utilizing the automata that encode said objectives in order to define novel reward shaping functions that mitigate the sparse rewards problem within Monte Carlo Tree Search (MCTS) methods. We further demonstrate that such verification-guided reward shaping can be utilized to facilitate transfer learning between different environments when the objective is the same.

Title: Low algebraic dimension matrix completion
PoC: Greg Ongie

Title: No Clash Teaching Dimension of Smoothly Parameterized hypothesis spaces
PoC: Rahul Parhi
Abstract: Machine teaching as gained interest in recent years with several applications including robotics, trustworthy AI, and pedagogy. The difference between machine learning and machine teaching is that a teacher chooses a training set so that a learner can learn a specific hypothesis. There are several existing models of teaching and associated notions of ``teaching dimension'' to capture the complexity

of teaching a particular hypothesis. In this work we consider the setting where unfair collusion between the teacher and learner is not allowed. This is a suitable assumption to impose on teaching since collusion would trivialize the task of teaching. In this work we prove that for smoothly parameterized hypothesis spaces, the no-clash teaching dimension is upper bounded by the dimension of the parameter manifold.

Title: MaxGap Bandit: Adaptive Algorithms for Approximate Ranking
PoC: Ardhendu Tripathy
Abstract: This paper studies the problem of adaptively sampling from K distributions (arms) in order to identify the largest gap between any two adjacent means. We call this the MaxGap-bandit problem. This problem arises naturally in approximate ranking, noisy sorting, outlier detection, and top-arm identification in bandits. The key novelty of the MaxGap bandit problem is that it aims to adaptively determine the natural partitioning of the distributions into a subset with larger means and a subset with smaller means, where the split is determined by the largest gap rather than a pre-specified rank or threshold. Estimating an arm's gap requires sampling its neighboring arms in addition to itself, and this dependence results in a novel hardness parameter that characterizes the sample complexity of the problem. We propose elimination and UCB-style algorithms and show that they are minimax optimal. Our experiments show that the UCB-style algorithms require 6 ~ 8x fewer samples than non-adaptive sampling to achieve the same error.

Title: Sample Complexity of Species Tree Estimation From a Linear Combination of Internode Distances
Poc: Harrison Rosenberg & Sebastian Roch
Abstract: We consider the problem of estimating the species tree from large numbers of unrooted gene tree topologies in the presence of incomplete lineage sorting, a phenomenon which enforces heterogeneity among the set of gene trees. This phenomenon is modeled with the Multi-Species Coa- lescent Process. We make progress towards deriving a sample complexity bound for species tree reconstruction methods based on sample averages of graph distances. The sample complexity de- pends polylogarithmically on N , the number of species in the tree.

Title: Classifying Satellite Light Curves
AFRL PoC: Walter Bennette
Affiliation: AFRL/RIEA
Email: walter.bennette.1@us.af.mil
Abstract: The Air Force is interested in Resident Space Objects to help maintain Space Situational Awareness. The sheer number of objects tracked and volume of data associated with characterizing Resident Space Objects makes this a labor intensive task. Therefore, we explore automating the characterization of Resident Space Objects -- leveraging photometric data from non-resolved objects -- to reduce the labor costs of human analysts. Previous studies have only shown these abilities for

simulated data. In contrast, we develop and analyze techniques using real datasets. Specifically, this work discusses the curation of a real dataset of Resident Space Objects and investigates different classification strategies.

Title: Uncovering High-Confidence Mistakes for Classification Models
AFRL PoC: Walter Bennette
Affiliation: AFRL/RIEA
Email: walter.bennette.1@us.af.mil
Abstract: Assessing the predictive accuracy of black box classifiers is challenging in the absence of labeled test datasets. In these scenarios we may need to rely on a human oracle to evaluate individual predictions; presenting the challenge to create query algorithms to guide the search for points that provide the most information about the classifier's predictive characteristics. Previous works have focused on developing utility models and query algorithms for discovering unknown unknowns -- misclassifications with a predictive confidence above some arbitrary threshold. However, if misclassifications occur at the rate reflected by the confidence values, then these search methods reveal nothing more than a proper assessment of predictive certainty. We are unable to properly mitigate the risks associated with model deficiency when the model's confidence in prediction exceeds the actual model accuracy. We propose a utility model and corresponding greedy query algorithm that instead searches for overconfident unknown unknowns. Through robust empirical experiments we demonstrate that the greedy query algorithm with the facility locations utility model consistently results in oracle queries with superior performance in discovering overconfident unknown unknowns than previous methods.

Title: ESCAPE Multi-Modal Data Set for ML Research
PoC: Peter Zulch
Affiliation: AFRL/RIGC
Email: peter.zulch@us.af.mil
Abstract: Over the last decade there has been a technological explosion of advanced, digital, solid state, software controlled, and low size, weight, power and cost (SWaPC) sensors and payloads. The sensor advancement allowed engineering practitioners the ability to easily perform a variety of remote sensing operations and consider a wide range of modalities measuring the environment simultaneously. However, there has not been a common multi-modal data set to compare data fusion methods. This poster describes a multi-mode data set performed by the Air Force Research Laboratory, Information Directorate, to enable multi-modal signature data-fusion research. The Experiments, Scenarios, Concept of Operations, and Prototype Engineering (ESCAPE) collection brings together electro-optical, infrared, distributed passive radio-frequency, radar, acoustic and seismic data in a common scenario for the application of advanced fusion methods for aerospace systems. The poster details hardware, scenarios, and data collection specifics. Scenarios involved disparate moving emitting ground vehicles, challenging vehicle path patterns, and differing vehicle noise profiles. The purpose of the data collection, and the resulting data sets, is to engage the data fusion community in advanced upstream heterogeneous data analytics, design, and understanding.

Title: Low Shot Learning at DSTL

PoC: Todd Robinson
Affiliation: DSTL
Email: trobinson2@mail.dstl.gov.uk

Abstract: The Defence Science and Technology Laboratory (Dstl) ensures that innovative science and technology contribute to the Defence and Security of the UK. At Dstl, we are interested in applying low-shot learning across Defence and are developing novel methods which can be trained effectively with little data. This poster presents the work done by Dstl in the area of low shot learning, in particular in the domains of image classification and image segmentation.

Title: Calibration Models and System Development for Compressive Sensing with Micromirror Arrays

PoC: Rebecca Profeta
Affiliation: AFRL/RYAT
Email: rebecca.profeta@us.af.mil

Abstract: Compressive sensing (CS) is an active research field focused on finding solutions to sparse linear inverse problems, i.e. estimating a signal using fewer linear measurements than there are unknowns. In this research, we present the development of a hardware CS imaging system using a Digital Micromirror Device (DMD) providing spatial light modulation via an array of micromirrors. Additionally, we develop a number of new DMD-specific calibration models intended to capture the physical attributes of micromirrors and the end-to-end data collection system. The resultant CS reconstructions demonstrate a substantial reduction in image estimation error while reducing the number of required measurements by fifty percent, relative to current baseline calibration methods. The research shown uses a Bayesian approach to reconstruct the compressed measurement. The authors wish to extend this work and compare these baseline results to a deep learning based approach.

Title: Characterizing Inter-Layer Functional Mappings of Deep Learning Models

PoC: Donald Waagen
Affiliation: AFRL/RWWI
Email: donald.waagen@us.af.mil

Abstract: Deep learning architectures have demonstrated state-of-the-art performance for object classification and have become ubiquitous in commercial products. These methods are often applied without understanding (a) the difficulty of a classification task given the input data, and (b) how a specific deep learning architecture transforms that data. To answer (a) and (b), we illustrate the utility of a multivariate nonparametric estimator of class separation, the Henze-Penrose (HP) statistic, in the original as well as layer-induced representations. Given an N-class problem, our contribution defines the C(N,2) combinations of HP statistics as a sample from a distribution of class-pair separations. This allows us to characterize the distributional change to class separation induced at each layer of the model. Fisher permutation tests are used to detect statistically significant changes within a model. By comparing the HP statistic distributions between layers, one can statistically characterize: layer adaptation during training, the contribution of each layer to the classification task, and the presence or absence of consistency between training and validation data. This is demonstrated for a simple deep neural network using CIFAR10 with random labels, CIFAR10, and MNIST datasets.

Title: Comparing Classifiers that Exploit Random Subspaces
PoC: Jamie Gantert
Affiliation: AFRL/RWWI
Email: jamie.gantert.1@us.af.mil
Abstract: Many current classification models, such as Random Kitchen Sinks and Extreme Learning Machines (ELM), minimize the need for expert-defined features by transforming the measurement spaces into a set of ”features” via random functions or projections. Alternatively, Random Forests exploit random subspaces by limiting tree partitions (i.e. nodes of the tree) to be selected from randomly generated subsets of features. For a synthetic aperture RADAR classification task, and given two orthonormal measurement representations (spatial and multi-scale Haar wavelet), this work compares and contrasts ELM and Random Forest classifier performance as a function of (a) input measurement representation, (b) classifier complexity, and (c) measurement domain mismatch. For the ELM classifier, we also compare two random projection encodings.

Title: USAF Relevant Remote Sensing Data Sets with Labels
PoC: Todd V. Rovito
Affiliation: AFRL/RYAT
Email: todd.rovito@us.af.mil
Abstract:
AI/ML research needs data sets in order to make advancements. Academia has created well known data sets such as ImageNet which has generated community interest in solving large scale recognition problems. But ImageNet is not relevant to the USAF because these data sets include a large number of pixels per example and don’t contain atmospheric turbulence or motion blur from moving platform. We present three public released data sets that are captured from a moving aircraft, include labels, and offer multiple sensor types. The data sets are:

- Columbus Large Image Format (CLIF) https://github.com/AFRL-RY/data-clif-2007
- UNIfiedCOincidentOptical and Radar for recognitioN(UNICORN) https://github.com/AFRLRY/data-unicorn-2008
- The Minor Area Motion Imagery (MAMI) https://www.sdms.afrl.af.mil/index.php?collection=mami2013

Computational Efficiency Thrust

Leads: Mikko Lipasti UW-Madison, Clare Thiem AFRL/RI, Dimitris Papailiopoulos UW-Madison

Title: Bitstream computing
PoC: Kyle Daruwalla

Title: Neuromorphic Computing for Computational Efficiency

PoC: Clare Thiem

Affiliation: AFRL/RITB

Email: clare.thiem@us.af.mil

Abstract: The Air Force Research Laboratory's Information Directorate is pursuing a variety of neuromorphic computing hardware technologies to achieve computational efficiency for future Air Force systems. Data can be collected faster than traditional computing approaches can process it. The goal of the research activity is to obtain brain-inspired, extremely low size, weight, and power (SWaP), intelligent computing at the edge for dynamic and contest mission environments. The poster will provide an overview of the research activity including motivation for the research, trade-offs and technologies being considered along with possible future capabilities that might be realized.

Title: Assessing Cognitive Workload via a Passive Brain-Computer Interface with Ensemble Learning

PoC: Justin R. Estepp

Affiliation: AFRL/711 HPW

Email: justin.estepp@us.af.mil

Abstract: A passive brain-computer interface (pBCI) allows implicit inference to be made about an operator's cognitive processes through direct measurement of the brain's activity. Typically, machine learning methods are used to decode the relationship between signals measured from the brain and a desired objective function (i.e., the desired output of the pBCI system). One objective function that can be determined in a pBCI is an estimate of an operator's cognitive workload. Knowledge of cognitive workload, in real-time, would be advantageous to human-machine teaming applications where the machine teammate can actively (and, perhaps, proactively) adjust its behavior to the moment-to-moment cognitive demand placed upon its human partner. Previous work has shown that a learning set consisting of data sampled from multiple days leads to more generalizable workload models a result that is hypothesized to be a mitigation against non-stationarity in the brain's response over time. However, there is currently no known, objective method to determine how these physiological data should be ideally sampled to achieve generalization. We look to ensemble learning methods, such as Adaptive Boosting (AdaBoost), to obtain a lower model prediction error and variance on an independent test set and leverage theoretical performance of the ensemble as compared to the observed performance to make inference about generalization in different learning set sampling strategies. In this analysis, theoretical performance is established by comparing error and variance of the ensemble to the best-performing individual learner within the ensemble. The results presented here indicate that these ensemble learning approaches show convergence toward theoretical performance as the sampling of data included in the training set spans a greater period of time (up to 4 unique days of data collection), but they did not reach convergence to theoretical performance in either accuracy or variance with the data collected. By comparing observed ensemble performance to theoretical, we may be able to make inference on model generalization that cannot be achieved by simple calculation of prediction error of a base learner. This, in addition to reduced bias and variance expected of the ensemble approach, will lead to better methods for training machine learning decoders in pBCI systems. Future work in this area will utilize several new datasets that are being expanded to 10 days of data collection that may of interest to the ERML CoE community.

Title: Complete & Orthogonal Replication of Hyperdimensional Memory via Elementary Cellular Automata
PoC: Nathan McDonald
Affiliation: AFRL/RITB
Email: nathan.mcdonald.5@us.af.mil
Abstract: Hyper-dimensional computing (HDC)/ Vector Symbolic Architectures (VSA) implements associative learning using very large binary vectors. For example, this approach has been used to model online learning and transfer learning in a foraging honeybee. In real-world systems, this learned memory will be copied onto multiple agents, e.g. swarm systems; however, simply copying the memory vectors across multiple agents makes all agents vulnerable to the same attack by a malicious entity. The challenge is to replicate the parent agent's item memory and compositional memory such that all learned associations are preserved yet the clone's memory vectors are maximally uncorrelated with the parent's memory vectors. This work evaluated all 256 elementary cellular automata (ECA) rules for this task and identified 8 rules that satisfied these replication requirements. To the best of the authors' knowledge, this is the first report of complete and orthogonal replication of HDC memory using ECA.

Title: Machine Learning on a Neuromorphic Computer with TrueNorth Processors (1 of 2)
PoC: Clare Thiem
Affiliation: AFRL/RITB
Email: clare.thiem@us.af.mil
Abstract: The Air Force Research Laboratory is pursuing Machine Learning from several different perspectives including the use of neuromorphic computers. These posters provide insight as to why computing hardware such as TrueNorth Processors are being examined by Air Force researchers and how they might exploit the technology to help the Air Force perform its mission.

Title: Machine Learning on a Neuromorphic Computer with TrueNorth Processors (2 of 2)
PoC: Clare Thiem
Affiliation: AFRL/RITB
Email: clare.thiem@us.af.mil
Abstract: The Air Force Research Laboratory is pursuing Machine Learning from several different perspectives including the use of neuromorphic computers. These posters provide insight as to why computing hardware such as TrueNorth Processors are being examined by Air Force researchers and how they might exploit the technology to help the Air Force perform its mission.

Title: Inverse Design of All-Dielectric Metasurface Reflectors
PoC: Meghan N. Weber
Affiliation: AFRL/RXAP
Email: webermn@protonmail.com
Abstract: Optical metasurfaces are able to achieve performance that is unattainable in bulk materials due to the engineered nanoscale geometry. These features allow for the ability to exhibit near perfect, low loss reflection via light scattering. To design the next generation of functional materials, scientists

and engineers must create new methods to evaluate the near-infinite number of possible patterns. Here we present our work using machine learning to accelerate the process of metasurface design. Our goal is to design a silicon pillar on silicon dioxide substrate metasurface to meet target specifications. To validate the process, we aim to achieve a reflection peak at 1550 nm. This process uses artificial neural networks (ANNs) to solve the inverse design problem. Utilizing a dataset of devices simulated with rigorous coupled wave analysis (RCWA), we create an ANN-accelerated simulator, achieving a speed up of 10^6 over the relatively quick RCWA simulation method. We couple this simulator with another ANN, known as the predictor, to directly design optical metasurfaces with targeted performance. Together, the simulator/predictor ANNs provide a framework to rapidly evaluate and design potential optical metasurface devices beyond what is currently possible using simulation and optimization methods.

Poster Session 2: Operational and Adversarial Robustness

Title: Processing light curve data to hopefully be able to classify satellite bus types.
PoC: Curt Lockhart

Title: Intersection of optimization and control
PoC: Laurent Lessard

Title: DIODE: A Dense Indoor and Outdoor DEpth Dataset
PoC: Greg Shakhnarovich
Abstract: We introduce DIODE, a dataset that contains thousands of diverse high resolution color images with accurate, dense, wide-range depth measurements. DIODE (Dense Indoor/Outdoor DEpth) is the first public dataset to include RGBD images of indoor and outdoor scenes obtained with one sensor suite. This is in contrast to existing datasets that focus on just one domain/scene type and employ different sensors, making generalization across domains difficult.

Title: Neumann Networks for Inverse Problems in Imaging
PoC: Davis Gilton

Title: Adversarial Machine Learning
PoC: Yingyu Liang

Title: Standardized Machine Learning Software to Enable Rapid Iteration of Materials Science Research
PoC: Benjamin Affenbach

Title: Estimating Network Structure from Incomplete Event Data
PoC: Benjamin Mark
Abstract: Multivariate Bernoulli autoregressive (BAR) processes model time series of events in which the likelihood of current events is determined by the times and locations of past events. These processes can be used to model nonlinear dynamical systems corresponding to criminal activity, responses of patients to different medical treatment plans, opinion dynamics across social networks, epidemic

spread, and more. Past work examines this problem under the assumption that the event data is complete, but in many cases only a fraction of events are observed. Incomplete observations pose a significant challenge in this setting because the unobserved events still govern the underlying dynamical system. We develop an approach to estimating the parameters of a BAR process in the presence of unobserved events via an unbiased estimator of the complete data log-likelihood function. We propose a computationally efficient estimation algorithm which approximates this estimator via Taylor series truncation and establish theoretical results for both the statistical error and optimization error of our algorithm. We further justify our approach by testing our method on both simulated data and a real data set consisting of crimes recorded by the city of Chicago.

Title: Fast Adversarial Training with Stale Perturbation Updates
PoC: Yang Guo
Abstract: Many state-of-the-art machine learning models have been repeatedly shown to be vulnerable to test-time adversarial attacks. While various defense methods have been proposed, one of the most popular and successful is adversarial training using projected gradient descent (PGD). In this framework, adversarial perturbations are computed at training time using PGD and are used to retrain the model so as to reduce vulnerability. Unfortunately, this method suffers from a high training cost [1]. The time of computing PGD attacks typically dominates the time of computing mini-batch stochastic gradient updates, and is proportional to the total number of samples.

Operational Robustness Thrust

Leads: Rebecca Willett UC, Ashley Prater-Bennette AFRL/RI, Karen Livescu TTIC

Title: Autonomous Swarms for Information-aware Mission Operation with Verification
PoC: Jeffrey Hudack
Affiliation: AFRL/RIEA
Email: jeffrey.hudack@us.af.mil
Abstract: This Autonomous Swarms for Information-aware Mission Operation with Verification (ASIMOV) program at AFRL's Information Directorate is demonstrating new methods for designing software payloads that can be deployed on operational small UAS (Unmanned Air Systems) platforms. The use of sUAS for autonomous operations introduces new technical challenges for object classification using on-board sensors, computation, communications and software. With reduced size comes more restrictive payload weights, limiting the number and diversity of sensors that can be located on a single platform, requiring them to share and leverage multi-sensor data and models over limited communications channels.

Title: Deep Learning based Super Resolution of Aerial and Satellite Imagery
PoC: Asif Mehmood
Affiliation: AFRL/RYAT
Email: asif.mehmood.1@us.af.mil
Abstract: Super-resolution is the process of creating high-resolution (HR) images from low-resolution (LR) images. SISR is challenging because high-frequency image content typically cannot be recovered

from the low-resolution image and absence of high-frequency information thus limits the quality of the high-resolution image. Furthermore, SISR is an ill-posed problem because a low-resolution image can yield several possible high-resolution images. To address this issue numerous techniques being proposed but recently deep learning based methods are being employed. Deep learning based methods have shown great success in numerous computer vision tasks achieved by the convolutional neural networks (CNNs). Therefore, it is worthwhile to make use of CNN to address this challenging problem. This paper presents a deep learning based super resolution (DLSR) approach to find a HR image from its LR counterpart by learning the mapping between them. This mapping is possible because LR and HR images have similar image contents and differ primarily in high-frequency details. In addition, DLSR utilizes residual learning strategy- the network learns to estimate a residual image. DLSR is applied to both aerial and satellite imagery and resulting estimates are compared against the traditional methods using Peak Signal to Noise Ratio (PSNR), Structure Similarity Index Metric (SSIM) and Naturalness Image Quality Evaluator (NIQE), which is also called perceptual quality index. These results show that DLSR outperform the traditional approaches.

Title: Data-Informed Design of Algorithms with Provable Guarantees
PoC: Zola Donovan
Affiliation: AFRL/RISA
Email: zola.donovan@us.af.mil
Abstract: In general, different algorithms used to solve the same computationally challenging problem have incomparable performance. Depending on the performance metric, one algorithm may do better on some instances, but not on others. To address this issue, we seek to use theory (beyond worst-case analysis) to decide which is best? One approach has been to cast the problem of selecting the best algorithm for a poorly understood application domain as one of learning the optimal algorithm with respect to an unknown instance distribution. Our goal is to leverage the power of machines, gain further insights, and learn near-optimal algorithms with provable guarantees to solve computationally challenging problems specific to highly dynamic environments.

Title: An Approach to a Goal-Oriented Conversational Agent
PoC: Krista Anken
Affiliation: AFRL/RIEA
Email: krista.anken@us.af.mil
Abstract: Recent advances in goal-oriented dialogue systems have made it possible for companies to begin to augment traditional business practices through the use of task specific conversation chat-bots. The Air Force Research Laboratory is perusing R&D in this area by paring open-source approaches to intent detection, dialogue state tracking and end-to-end question and answering with structured data-sources to create a baseline information retrieval tool called Roger.

Title: Specification-Guided Tree Search for Reinforcement Learning
PoC: Alvaro Velasquez
Affiliation: AFRL/RISC
Email: alvaro.velasquez.1@us.af.mil

Abstract: The field of reinforcement learning has been revolutionized in recent years, due in part to the mass adoption of deep convolutional neural networks and the resurgence of powerful methods to refine decision-making policies. However, the problem of sparse reward signals remains pervasive in many non-trivial domains. While various reward-shaping mechanisms and imitation learning approaches have been proposed to mitigate this problem, a mathematically rigorous structure of the underlying objective is rarely exploited. In this paper, we resolve this by defining objectives using temporal logic and utilize the automaton that defines said objectives in order to define novel reward shaping functions that mitigate the sparse rewards problem within Monte Carlo Tree Search (MCTS) methods. We further demonstrate that such specification-guided reward shaping can be utilized to facilitate transfer learning between different environments when the objective is the same.

Title: StreamlinedML (SML)
PoC: Amy Wong
Affiliation: AFRL/RISC
Email: amy.wong.2@us.af.mil
Abstract: To foster open machine learning (ML) exploration and enable the rapid, low-cost deployment of state-of-the-art (SOTA) ML capabilities for mission-tailored applications, we are developing an end-to-end ML ecosystem. The StreamlinedML Ecosystem provides tools and services for accessing and developing SOTA algorithms using a wide variety of datasets. We have also released a Model Integration Software Toolkit (MISTK) for agile development of ML techniques that will seamlessly integrate into the ecosystem. MISTK enables developers to develop and test their approaches without recreating the entire environment. This allows them to concentrate on the problem that their ML techniques address, instead of the support infrastructure and standards. MISTK is publicly available at https://github.com/mistkml.

Title: Robust Pattern-Preserving Tensor Decomposition for Spatiotemporal Data Analysis and Dimensionality Reduction
PoC: Charlotte Ellison
Affiliation: U.S. Army Corps of Engineers, Geospatial Research Laboratory
Email: Charlotte.L.Ellison@erdc.dren.mil
Abstract: The amount of spatiotemporal data has sky-rocketed in recent years, causing growth in many new Machine Learning areas of research. The knowledge gained by merging different modes of data (such as moving object trajectories, social networks, or event data) and extracting latent patterns from them could provide key insights to work towards the Army's goal of improved emergent pattern identification in the GEOINT-HUMINT spectrum of operations. Hence, this research hypothesizes that merging spatial trajectories with multi-modal datasets in a tensor representation (i.e., high-order array with dimensions or modes ≥ 3) and extending its factor decomposition with a within-mode coherence model will enable the discovery of hidden links between trajectories and other data classes in a robust and generalizable way. This has the added benefit of reducing the dimensionality of the data while maintaining these structures. Specifically, latent patterns such as similarity in the paths of users through time were identified by imposing temporal coherence. A novel decomposition technique based on dual tensor-matrix CP decomposition with alternating least squares updating was applied to a tensor. The

addition of within-mode similarity constraints enforced temporal coherence by altering the update function corresponding to the time mode. The method was tested on the BerlinMOD dataset, a set of simulated users' car trajectories moving through Berlin. This was accomplished by performing the structured decomposition on the simulated data, data subjected to a simple time shift, and randomly reordered data. The novel decomposition correctly identified the distinct temporal pattern shared by the original and time-shift data, while a traditional CP decomposition failed to differentiate these data from the randomly reordered data. Even when noise was introduced by coding location errors in half of the dataset, the novel structured tensor decomposition method correctly recovered the temporal pattern of the data. Next, when additional contextual information concerning the users was included, the novel tensor decomposition revealed both latent temporal patterns in trajectory data and information about group membership. Given the generalizable nature of the structured tensor decomposition method, it could readily be applied to other data sources or different modal coherence models and contextual information. Results from the tensor decomposition could also effectively reduce the dimensionality of higher order data, while preserving salient *a priori* information, prior to applying other machine learning techniques such as clustering or classifying the data.

Title: Composite Tensors with Sparse Tucker Representations
PoC: Ashley Prater-Bennette
Affiliation: AFRL/RISC
Email: ashley.prater-bennette@us.af.mil
Abstract: A common tool to process and interpret multimodal data is to represent the data in a sparse Tucker format, decomposed as a sparse core tensor with dictionary matrices for each modal dimension. In real-world applications one may be presented with a composition of several tensors, each with its own sparse Tucker representation and collection of dictionaries. The Tucker model and associated recovery algorithms struggle to accurately separate composite tensors in this situation, either having difficulty with the overcomplete dictionaries or not fully taking advantage of the special structure of the decomposition. To address these deficiencies, we introduce an overcomplete sparse Tucker model and an iterative algorithm to separate a composite sparse Tucker tensor. The method, which is based on soft-thresholding shrinkage techniques, is demonstrated to effectively separate overcomplete tensors and recover the sparse components tensors on real-world datasets, and to do so more accurately than other Tucker methods.

Title: CNN Classifier Performance w.r.t. Image Contrast
PoC: Christopher Menart
Affiliation: AFRL/RYAT
Email: christopher.menart@us.af.mil
Abstract: Deep neural networks demonstrate high performance at classifying high-dimensional signals, but often fail to generalize to data that is different from the data they were trained on. In this paper, we investigate the resilience of convolutional neural networks (CNNs) to unforeseen operating conditions. Specifically, we empirically evaluate the ability of CNN models to generalize across changes in image contrast. Multiple models are trained on electro- optical (EO) or near-infrared (IR) data, and are evaluated in environments with degraded contrast compared to training. Experiments are replicated

across varying architectures, including state-of-the-art classification models such as Resnet-152, and across both synthetic and measured datasets. In comparison to models trained and evaluated on identically-distributed data, these models can generalize well when contrast invariance is built up through data augmentation. Future work will investigate CNN ability to generalize to other changes in operating conditions.

Title: Image2RF: Deep Signal Strength Prediction in an Urban Environment Leveraging Visual Features

PoC: Dylan Elliot

Affiliation: AFRL/RISC

Email: dylan.elliot@us.af.mil

Abstract: In this paper, we introduce a novel approach for predicting the received signal strength indicator (RSSI) between pairs of radio transmitters in an urban environment by exploiting both the rich information content and high availability of overhead imagery combined with the descriptive power of neural networks. We introduce a deep approach for learning the complex physical electromagnetic propagation properties from visual data alone. Specifically, given the locations for a pair of radios (a transmitter and a receiver), as well as a satellite image containing the locations of these radios, we train a multiple-input deep learning model that combines positional with image features to estimate the true RSSI. We demonstrate the effectiveness of our approach on a popular RSSI measurement dataset and show that our approach outperforms several well-known theoretical/ empirical path loss prediction models. Additionally, we show through an ablation study the advantage of using combined positional and image input on the overall model performance.

Title: Upstream Unsupervised Sensor Fusion Using ML

PoC: Peter Zulch

Affiliation: AFRL/RIGC

Email: peter.zulch@us.af.mil

Abstract: This poster presents a processing pipeline for fusing 'raw' and / or feature-level multi-sensor data – upstream fusion – and initial results from this pipeline using Full Motion Video (FMV) and Passive Radio Freqeuency (P-RF) signal data to determine which tracked object, among several, hosts an emitter of interest. Correctly making this determination requires fusing data across these modalities. Our approach performs better than standard fusion approaches that make detection / characterization decisions for each modality individually and then try to fuse those decisions – downstream (or post-decision) fusion. Our approach (1) fully exploits the inter-modality dependencies and phenomenologies inherent in different sensing modes, (2) automatically discovers compressive hierarchical representations that integrate structural and statistical characteristics to enhance target / event discriminability, and (3) completely obviates the need to specify features, manifolds, or model scope a priori. This approach comprises a unique synthesis of Deep Learning (DL), topological analysis over probability measure (TAPM), and hierarchical Bayesian non-parametric (HBNP) recognition models. Deep Generative Networks (DGNs – a deep generative statistical form of DL) create probability measures that provide a basis for calculating homologies (topological summaries over the probability measures). The statistics of the resulting persistence diagrams are inputs to HBNP methods that learn to discriminate between target types and distinguish emitting targets from non-emitting targets, for example. This approach overcomes the inadequacy of pre-defined features as a means for creating efficient, discriminating, low-dimensional representations from high-dimensional multi-modality sensor data collected under difficult, dynamic sensing conditions. Machine learning makes adaptivity a central feature of our approach. Adaptivity is critical because it enables flexible processing that automatically accommodates a broad range of challenges that non-adaptive, standard fusion approaches would typically require manual intervention to begin to address.

Title: Clustering Partial State Measurements & Space Tracking
PoC: Christopher T. Diggans
Affiliation: AFRL/RISC
Email: christopher.diggans@us.af.mil
Abstract: Many measurement devices do not fully characterize the state of the phenomenon they are meant to record; some record quantities that are indirectly related to the state of interest. In such cases, it is often desirable to partition a large set of observations by their sources in order to allow data driven state estimation. Standard clustering approaches fail when pairwise comparisons based on distance in the measurement space yield no usable information for defining a similarity kernel. A framework is described for data-aware clustering where an algorithmic similarity measure incorporates the data set for context, and expert knowledge and heuristics are used to determine a likelihood of pairwise identification. In this way, a normalized affinity matrix is constructed and by either basic spectral clustering or through a novel use of diffusion mapping, the data set is partitioned by neighborhoods in the state space of interest. A use case in the domain of space tracking is provided to illustrate how to apply this framework to a domain with expert knowledge and heuristics.

Title: Demonstrating a Tactical Server Concept Leveraging Big Data, Deep analytics, Machine Learning (ML), and Artificial Intelligence (AI) Algorithms
PoC: Ying Zhao
Affiliation: Naval Postgraduate School
Email: yzhao@nps.edu
Abstract: In the recent years, there has been tremendous advancement in commercial applications recently using the big data, deep analytics including machine learning (ML), and artificial intelligence (AI) methods. These methods provide emerging technologies with applications to address the unique challenges in the Common Tactical Air Pictures (CTAP) and Combat Identification (CID). Decision superiority at the tactical edge is imperative including meeting the ongoing needs and requirements for the CTAP and CID.
We have leveraged these advancements and demonstrated the role of big data in CTAP and CID. This year, we built on the research results towards an integrated demonstration (the "tactical server" concept) by showing the feasibility of using the advanced algorithms contributing decision superiority in CTAP and CID. The goal is to demonstrate a logical, incremental introduction of the technologies towards improved engagement in doctrine, tracking and identification, meanwhile addressing more challenges and needs.
We will present and show the initial progress of the tactical server concept by leveraging the state-of-the-art big data, deep analytics, and ML/AI algorithms. Specifically, we will show the progress towards

an integrated demonstration of the tactical server at the GENSER level in order to participate in an annual Navy exercise, in effort to demonstrate the reality of the improvement of the fidelity and latency of identifying neutral and other interesting airborne objects upon the traditional methods such as the Collaborative Engagement Capability (CEC) composite ID. The work has been realized by building virtual airways (patterns of life) using (through features) of the big and open source data such as Automatic dependent surveillance—broadcast (ADS–B) data. We will report and demonstrate the resulted algorithms, models, and evolving and new concepts from the simulation or replay/reconstruction of a real-life CEC event for testing in support of realistic functions of CTAP and CID, combat systems, and kill chain leveraging ML/AI. Future work includes class and type recognition with more data sources such as ICAO, FAA, ELINT (e.g. RF) and discovery of patterns of life using unsupervised learning and planning algorithms.

Adversarial Robustness Thrust

Leads: Jerry Zhu UW-Madison, Ryan Luley AFRl/RI, Yingyu Liang UW-Madison

Title: Does Data Augmentation Lead to Positive Margin

PoC: Shashank Rajput

Abstract: Data augmentation (DA) is commonly used during model training, as it significantly improves test error and model robustness. DA artificially expands the training set by applying random noise, rotations, crops, or even adversarial perturbations to the input data. Although DA is widely used, its capacity to provably improve robustness is not fully understood. In this work, we analyze the robustness that DA begets by quantifying the margin that DA enforces on empirical risk minimizers. We first focus on linear separators, and then a class of nonlinear models whose labeling is constant within small convex hulls of data points. We present lower bounds on the number of augmented data points required for non-zero margin, and show that commonly used DA techniques may only introduce significant margin after adding exponentially many points to the data set.

Title: Convergence and Margin of Adversarial Training on Separable Data

PoC: Zachary Charles

Abstract: Adversarial training is a technique for training robust machine learning models. To encourage robustness, it iteratively computes adversarial examples for the model, and then re-trains on these examples via some update rule. This work analyzes the performance of adversarial training on linearly separable data, and provides bounds on the number of iterations required for large margin. We show that when the update rule is given by an arbitrary empirical risk minimizer, adversarial training may require exponentially many iterations to obtain large margin. However, if gradient or stochastic gradient update rules are used, only polynomially many iterations are required to find a large-margin separator. By contrast, without the use of adversarial examples, gradient methods may require exponentially many iterations to achieve large margin. Our results are derived by showing that adversarial training with gradient updates minimizes a robust version of the empirical risk at a O(ln(t)^2/t) rate, despite non-smoothness. We corroborate our theory empirically.

Title: Adversarial images and human perception
PoC: Liam Marshall
Abstract: Adversarial attacks attempt to confound machine learning systems. By changing the input to a classifier slightly – for instance, tweaking image pixels – the output classification can be manipulated. Throughout the literature on attacking image classifiers, the visibility of attacks to a human observer (who might become suspicious, which we wish to avoid) is typically gauged (without adequate justification) by a p-norm of the difference between the original and modified image.
Via behavioral experiments, we show that human perception of image modifications is not well-described by any p-norm, nor several alternative measures.
This has significant impact on adversarial ML research; the robustness of attacks to human inspection relies on an accurate understanding of what humans will and will not see as "tampering".

Title: Adversarial Learning
PoC: Jerry Zhu
Affiliation: UW-Madison

Title: Critic After Convergence
PoC: Walter Bennette
Affiliation: AFRL/RIEA
Email: walter.bennette.1@us.af.mil
Abstract: A Generative Adversarial Network (GAN) is a neural network training regimen composed of a generative model and a critic model. Through a back and forth adversarial competition, the critic model guides the generative model to produce samples that mimic a chosen distribution. At convergence (or when the critic can no longer differentiate between real and generated samples), the desired output of a GAN is typically a generative model that is capable of producing realistic samples from that distribution. In this work we investigate the performance of the critic model after the GAN has converged. Specifically, we perform experiments on the MNIST dataset to see if a critic model can be updated outside of the adversarial regimen to again discriminate between real and generated digits.

Title: Using Machine Learning for Cyber Agility and Network Intrusion Detection
PoC: Nandi O. Leslie
Affiliation: Army Research Laboratory
Email: Nandi.O.Leslie.Ctr@mail.mil
Abstract: Mission-critical cyber-physical systems (CPS), such as Internet of Things (IoT), are increasingly necessary at the tactical edge and face mounting risks to cyberattacks adversaries ranging from non-state actors to peers. Using semi-supervised learning, we propose an anomaly-based network intrusion detection system (NIDS) to detect and classify anomalous and/or malicious traffic. With this proposed machine learning approach, we detect botnet traffic and distinguish it from the normal and background traffic in the IPv4 flow datasets. We evaluate the prediction performance results for the flow-based NIDS algorithms. We show an improvement in detection accuracy and reduction in error rates, when compared with signature-based NIDS and previous studies. In the future, we hope to combine our

network intrusion detection algorithms from machine learning with cyber agility approaches — these include assessing the impacts of bond and edge percolation to the size of the giant connected component and developing proactive recovery techniques — in an effort to create more resilient CPS to network breaches.

Title: Prototypical Networks are Robust to Adversarial Perturbations

PoC: Leslie N. Smith
Affiliation: Naval Research Laboratory
Email: leslie.smith@nrl.navy.mil

Abstract: The use of deep neural networks have made huge strides in numerous applications, such as computer vision, speech recognition, robotics, and security. Along with these remarkable improvements in performance, the potential for vulnerabilities has also increased, causing concern in safety critical applications. In particular, the introduction of adversarial examples, where imperceptible perturbations are added to the input that cause neural networks to give false outputs, are an active area of research. Currently there is an arms race between adversarial example attackers and defenders, which the attackers seems to be winning.

In this poster we demonstrate that prototypical networks are surprising robust to adversarial examples. In order to understand the factors for this robustness, we performed a series of experiments with various levels of a metric learning loss function, which is the basis of prototypical network's loss function. While metric learning loss functions are more robust than Softmax, our experiments clearly indicate that there are other factors that behind the robustness of prototypical networks that are yet to be determined.