A University Center of Excellence supported by the Air Force Research Laboratory (AFRL).
https://madlab.ml.wisc.edu/

Newsletter – Q1 2019

Table of Contents

Welcome Message from Professor Robert Nowak, University of Wisconsin-Madison and Lee Seversky, Air Force Research Laboratory

On behalf of the MADLab AFRL University Center of Excellence (UCoE) team for Efficient and Robust Machine Learning (ERML) and the Air Force Research Laboratory center's team, we would like to share the center's first quarterly newsletter. The goal of the newsletter is to provide a venue for showcasing research accomplishments, team members and projects, as well as share information about upcoming events with the academic and Air Force machine learning communities.

What is the ERML UCoE? The Efficiency and Robust Machine Learning University Center of Excellence is focused on developing the next generation of machine learning methods to address the operational Air Force learning challenges of efficiency and robustness. The ERML UCoE was awarded to the University of Wisconsin-Madison in 2018 for up to 5 years. The center is a joint effort between the Air Force Research Laboratory, Information Directorate (AFRL/RI) and AFOSR. It is focused on advancing the state-of-the-art in efficient and robust machine learning methods as well as fostering a collaborative research environment between the university and AFRL government scientists and engineers.

Air Force Machine Learning Opportunities and Challenges: Machine Learning (ML) continues to emerge as a critical field of Artificial Intelligence (AI) and has a critical role to play in shaping the future Air Force. From shifting from today's high data processing and analysis burden from the Airman to machines, to increasing the overall quality and speed of decision making - advancements in Machine Learning hold great potential for the Air Force. However, for ML systems to be deployed in Air Force operational settings they must perform

reliably in complex operational situations as they learn, make decisions, and act. Further, AF systems operate in degraded and uncertain environments. As such, current ML techniques that do not explicitly consider these challenges have little hope of achieving the high performance necessary for trusted intelligence systems to be deployed in operational environments.

The research goals of this UCoE are to study ML techniques under the lens of these operational learning settings and to develop novel, principled techniques that overcome challenges presented by operational constraints. Specifically this UCoE is focused on the challenges of efficient and robust machine learning which are tackled through four research thrusts: Data efficiency, Computational efficiency, Operational robustness, and Adversarial robustness.

The second focus of the UCoE is to establish a close collaboration between the MADLab university team and AFRL government scientists and engineers. The goal of this UCoE is to foster a collaborative environment to support joint research projects, university & government exchanges, and special events such as hack-a-thons, seminars, and challenge problems to maximally engage both university and AFRL participation. This is especially critical for this fast-paced and high demand area of machine learning where often real-world problems and data drive discovery of new methods and the application of machine learning techniques to problems and data identify new challenges. The UCoE seeks to bring the two communities together to jointly advance the fundamental and applied research to Air Force application.

Meet the Team

MADLAB Team

University of Wisconsin-Madison (UW)
Toyota Technological Institute at Chicago (TTIC)
University of Chicago (UC)

Center Director: Robert Nowak (UW)*

Thrust Leads

Data-Efficiency, Robert Nowak (UW) & Greg Shakhnarovich (TTIC)
Computational Efficiency, Mikko Lipasti (UW) & Dimitris Papailiopoulos (UW)
Adversarial Robustness, Jerry Zhu (UW) & Yingyu Liang (UW)
Operational Robustness, Rebecca Willett (UC) and Karen Livescu (TTIC)

For more information please see the MADLab Website

AFRL Center Team

Executive Committee

Dr. Lee Seversky, Lee.Seversky@us.af.mil, AFRL, Information Directorate*
Dr. Erik Blasch, erik.blasch.1@us.af.mil, AFRL, Office of Scientific Research
Dr. Qing Wu, qing.wu.2@us.af.mil, AFRL, Information Directore

Thrust Leads

Data-Efficiency, Dr. Walter Bennette, walter.bennette.1@us.af.mil & Dr. Lee Seversky, Lee.Seversky@us.af.mil, AFRL, Information Directorate
Computational Efficiency, Clare Thiem, clare.thiem@us.af.mil, AFRL, Information Directorate
Adversarial Robustness, Ryan Luley, ryan.luley@us.af.mil, AFRL, Information Directorate
Operational Robustness, Dr. Ashley Prater-Bennette, ashley.prater-bennette@us.af.mil, AFRL, Information Directorate

*Team leads

Research Vignettes

Research Vignettes highlight research projects from the prior quarter.

Adversarial Machine Learning

Research project conducted by Jerry (Xiaojin) Zhu, Professor of Computer Science at University of Wisconsin-Madison

Zhu is the Sheldon & Marianne Lubar Professor in Computer Science at the UW-Madison. His research focuses on machine learning, in particular optimal teaching, active learning, and semi-supervised learning. He is a recipient of a National Science Foundation CAREER Award in 2010 and several best paper awards, and was the co-chair for AISTATS 2017. Zhu leads the thrust on adversarial ML.

Professor Zhu along with his research colleagues uncovered a significant security threat while studying adversarial attacks on stochastic multi-armed bandit algorithms. The study used multi-armed bandit algorithms which gets its

name from the “one-armed bandit,” a gambler who needs to decide which slot machines to play, how many times and in which order. In this way, multi-armed bandit algorithms balance exploitation with exploration with the goal of testing the arms just enough times to determine which of k arms has the highest average payoff without wasting resources. Stochastic multi-armed bandits are used regularly for purposes such as recommending news articles, improving search results, displaying advertisements and even allocating medical treatment. However, Professor Zhu and colleagues revealed a type of previously unknown security threats against multi-armed bandit algorithms. In these type of threats, an adversary can control the behaviors of the algorithm -- e.g. what article the computer shows the user, or what medical treatment the computer recommends to a doctor -- without hacking into the code of the computer algorithm. Instead, an adversary can do so simply by subtly modifying the reward signal received by the multi-armed bandit algorithm. Their calculations showed that an attacker needs to shape the reward very little in order to manipulate the multi-armed bandit algorithm, which poses a significant security threat.

For more information please see the recently published article, Adversarial Attacks on Stochastic Bandits by Kwang-Sung Jun, Lihong Li, Yuzhe Ma, and Jerry Zhu.

Deep Learning for Inverse Problems

by Rebecca Willett, Professor of Statistics at the University of Chicago



Willett is a Professor Statistics at the University of Chicago. Her research focuses on developing machine learning and signal processing theory and methodology that exploit underlying low-dimensional models, including sparse and low-rank representations of data. Willett has received numerous awards, including the Air Force Office of Scientific Research Young Investigator Program award. Willett leads the thrust on operational robustness.

Left to right: Rebecca Willett, Greg Ongie-Postdoc, Davis Gilton- Ph.D. Student

In many image recovery problems, we wish to estimate an image from noisy, indirect measurements. Examples include image deblurring, inpainting, compressed sensing reconstruction, MRI, radar, and more. Traditionally, images were estimated in an optimization framework with a pre-specified regularizer, such as Tikhonov regularization or terms that promote sparsity in a known basis. More recent efforts focus on leveraging training images to learn a regularizer represented by a deep neural network. Our team has developed a new approach for learning a neural network in which the network architecture is constrained by the physical model that describes how the image data is collected. This approach yields compelling empirical results in a variety of application settings and outperforms the current state-of-the-art approaches. Our theoretical work has led to novel insights into why our approach is so effective, and we see that our learned estimator is consistent with the optimal estimator in certain settings.

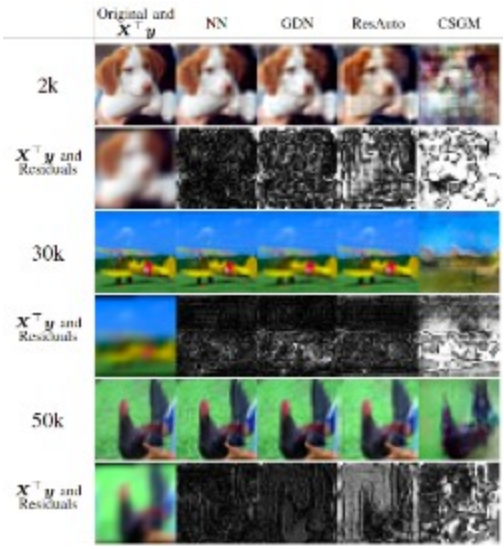

Figure 1: A qualitative comparison of the reconstructions produced for the deblurring problem at different training set sizes, along with the associated residual images. NN = proposed Neumann network, GDN = unrolled gradient descent network, ResAuto = residual autoencoder, CSGM = Compressed Sensing using Generative Models.

A preprint is online at https://arxiv.org/abs/1901.03707.

More specifically, we consider solving linear inverse problems in imaging in which an image, β, is observed via noisy linear projections as y = Xβ + noise. This general model is used throughout computational imaging, including in deblurring, tomographic reconstruction, magnetic resonance imaging, radar imaging, image inpainting, compressed sensing, interferometry, and many others. The task of estimating β from y is referred to as image reconstruction. Classical reconstruction methods assume some prior knowledge about β, such as smoothness, sparsity in some dictionary or basis, or geometric properties, and attempt to estimate β that is both a good fit to the observation y and which conforms to this prior knowledge. In general, a regularization function r(β) measures the lack of conformity of β to this prior knowledge and β is selected so that r(β) is as small as possible while still providing a good fit to the data.

However, recent work in computer vision using deep neural networks has leveraged large collections of “training” images to yield unprecedented image recognition performance and an emerging body of research is exploring whether this training data can also be used to improve the quality of image reconstruction. In other words, can training data be used to learn how to regularize inverse problems? We propose a novel neural network architecture

based on the Neumann series expansion that we call a Neumann network. Neumann networks, which directly incorporate the forward operator X into the network architecture, can learn to reconstruct images with fewer training samples than competing approaches, making them more amenable to applications where training datasets may be smaller.

Figure 2: Proposed Neumann network architecture. Inspired by the Neumann series expansion for computing the inverse of an operator, a Neumann network maps a linear function of the measurements to a reconstruction by successive application of a nonlinear operator while summing the intermediate outputs of each block. Here R is a trained neural network, and the scale parameter η is also trained. Unlike other networks based on unrolling of iterative optimization algorithms, the series structure of Neumann networks lead naturally to skip connections (highlighted in red) that route the output of each dashed block to directly to the output layer.

Figure 2

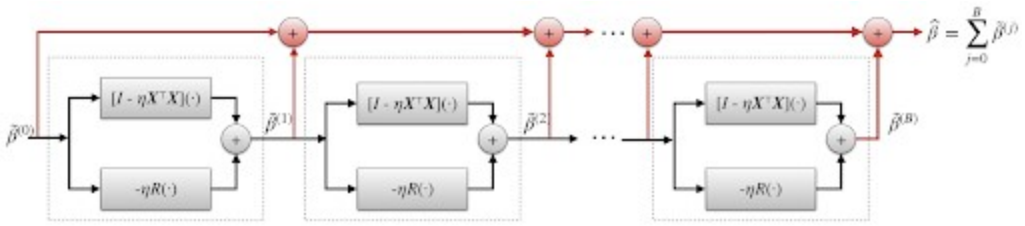

Neumann networks naturally yield a block-wise structure with “skip connections” between blocks that appear to make training the network easier by yielding an optimization landscape that's easier to navigate. Furthermore, when the images of interest lie on a union of subspaces, then there exists a set of weights such that the resulting Neumann network estimator corresponds to an oracle minimum variance unbiased estimator (MVUE). After training the Neumann network on simulated data drawn from a union of subspaces, the trained network is consistent with the oracle MVUE. A simple preconditioning step combined with the Neumann network further improves empirical performance in a manner consistent with our theoretical analysis. The empirical performance of the Neumann network on superresolution, deblurring, compressed sensing, and inpainting problems exceeds that of competing methods.

Ensuring Verifiable Long-Term Behavior in Autonomous Agents

by Alvaro Velasquez, AFRL, Information Directorate, alvaro.velasquez.1@us.af.mil

Alvaro Velasquez received his PhD in computer science in 2018 from the University of Central Florida, where he was a recipient of the National Science Foundation Graduate Research Fellowship. He is now a research scientist at the Air Force Research Laboratory working under the Autonomous Command and Control (AC2) competency. Alvaro has made contributions in the areas of logic and circuit design, coding theory, networking, and artificial intelligence. His current research interests lie in the area of verifiable decision-making and the broader field of trustworthy autonomy.

The application of autonomous decision-making systems has become widespread in recent years. From medical diagnoses [1] to robotic navigation [2], the adoption of these methods has attained ubiquitous appeal. While such systems boast impressive performance, there are generally no guarantees on their behavior. This behavior is typically governed by a decision-making policy $\pi : S \times A \mapsto [0, 1]$ whose goal is to maximize some environmental reward signal $R : S \times A \mapsto R$ within the Markov Decision Process (MDP) $M = (S, A, T, R)$ that models the agent-environment dynamics, where S is a set of states, A is a set of actions that an agent can take, and $T : S \times A \times S \mapsto [0, 1]$ is the transition probability function. This myopic goal can lead to failures in decision-making and presents a significant challenge to the adoption of traditional reinforcement learning and control methods in safety-critical domains. We seek to mitigate this by developing a framework of decision-making algorithms which maximize their expected reward while satisfying a set of desirable behavioral specifications. In particular, we have been successful in imposing guarantees on the long-term behavior of automated decision-makers by solving a constrained variant of the average-reward maximization problem. See Figure 1 for an example.

Figure 1: Example of an agent within an environment whose dynamics can be

modeled by an MDP M = (S, A, T, R), where the state-transition probability is given by T(sj |si , a) = 1 for the appropriate action, e.g. T(s2|s1, →) = 1 and T (∗|s1, ←) = 0. (right) A policy that solves the control problem satisfying the following steadystate behavioral specifications: (i) Stay in blue telecommunication links s3, s4, s7, s8 at least 70% of the time; (ii) Visit the yellow rendezvous point between 1% and 10% of the time; (iii) Avoid the red dangerous states s9, s10, s12.

Figure 1

$$\pi(a|s) = \begin{array}{c} \begin{array}{cccc} \leftarrow & \uparrow & \rightarrow & \downarrow \end{array} \\ \begin{pmatrix} 0.0 & 0.0 & 0.53543 & 0.46457 \\ 0.25545 & 0.0 & 0.54298 & 0.20157 \\ 0.12793 & 0.0 & 0.58195 & 0.29012 \\ 0.66285 & 0.0 & 0.0 & 0.33715 \\ 0.0 & 0.53893 & 0.46107 & 0.0 \\ 0.25369 & 0.25807 & 0.48824 & 0.0 \\ 0.09746 & 0.176 & 0.63112 & 0.09542 \\ 0.7107 & 0.2893 & 0.0 & 0.0 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.0 & 0.56166 & 0.0 & 0.43834 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.0 & 0.0 & 1.0 & 0.0 \\ 0.2939 & 0.0 & 0.7061 & 0.0 \\ 0.45054 & 0.26172 & 0.28774 & 0.0 \\ 1.0 & 0.0 & 0.0 & 0.0 \end{pmatrix} \end{array} \begin{array}{c} s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \\ s_6 \\ s_7 \\ s_8 \\ s_9 \\ s_{10} \\ s_{11} \\ s_{12} \\ s_{13} \\ s_{14} \\ s_{15} \\ s_{16} \end{array}$$

In Figure 1, a robotic agent has been placed in an environment whose underlying dynamics are given by the MDP M = (S, A, T, R), where S = {s1, . . . , s16} and A = {←, ↑, →, ↓} denote the state and action spaces, respectively. The robot starts in the top-left corner of the map and observes a positive reward R(s9, ↓), R(s14, ←) ∈ R + associated with recharging its energy source in state s13. The mission of the agent is to attempt to establish communication with a satellite via the blue telecommunication links in states s3, s4, s7, s8. Thus, we want to make sure the agent spends at least 70% of its time in these states. Successfully establishing communication will cause a spaceship to arrive at the yellow rendezvous point s16 and take the agent home, thereby accomplishing the mission. The agent does not know if or when the spaceship will arrive, so it must check back periodically to state s16, but not so often that it interrupts the task of establishing communication or recharging batteries. We therefore want to visit s16 at least 1% of the time, but no more than 10% of the time. During this mission of establishing communication, recharging batteries, and visiting the rendezvous point, the agent must avoid red states s9, s10, s12 which contain radioactive material. Note that we are interested in the continuing task setting where there are no terminal states and we must instead reason about the agent's long-term behavior through an infinite interaction with its environment. This is in contrast to the episodic case, where the underlying MDP contains terminal states which naturally end a simulation of the agent-

environment dynamics. To this end, we leverage the steady-state probability distribution (Pr∞ π (s1), . . . , Pr∞ π (s|S|)) over the state space, where Pr∞ π (s) denotes the long-term probability of being in state s under policy π. The preceding behavioral specifications can then be imposed as constraints on these probabilities in order to shape the policy so that the resulting agent maximizes its expected reward while satisfying behavioral specifications on its long-term behavior. The control policy π in Figure 1 (right) solves this example and yields Pr∞ π P (s16) u 0.0179, s∈{s3,s4,s7,s8} Pr∞ π (s) u 0.70967, and P s∈{s9,s10,s12} Pr∞ π (s) = 0. Thus, π satisfies all of the desired specifications.

One challenge that arises in finding a policy which satisfies steady-state specifications is the need for recurrence in the Markov Chain induced by the policy [3]. That is, all states must be reachable from each other in the agent-environment dynamics resulting from the given policy [4]. When this condition does not hold, the equations used to determine steady-state probabilities can admit solutions which are not indicative of the true behavior of the system. This problem is often mitigated in the literature by assuming that the underlying MDP is recurrent in that the Markov chain induced by any policy is itself recurrent [5]. In our work, we make no such assumptions on the underlying MDP. The method proposed finds a recurrent Markov chain in a possibly non-recurrent MDP, if one exists. We accomplish this by combining techniques from linear programming and network flow theory. The foregoing can be seen as a generalization of the steady-state control problem [5], which has been shown to be in PSPACE. We have improved upon this result by showing that this problem is actually in P. Furthermore, we proved that when a deterministic policy π : S 7→ A is required, the problem becomes NP-hard. This result is established via a reduction from the Hamiltonian cycle problem.

References

[1] Florin C Ghesu, Bogdan Georgescu, Tommaso Mansi, Dominik Neumann, Joachim Hornegger, and Dorin Comaniciu. An artificial agent for anatomical landmark detection in medical images. In International Conference on Medical Image Computing and Computer-Assisted Intervention, pages 229–237. Springer, 2016.

[2] Gregory Kahn, Adam Villaflor, Bosen Ding, Pieter Abbeel, and Sergey Levine. Self-supervised deep reinforcement learning with generalized computation graphs for robot navigation. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 1–8. IEEE, 2018.

[3] Sheldon M Ross. Introduction to probability models. Academic press, 2014.

[4] Erhan Cinlar. Introduction to stochastic processes. Courier Corporation, 2013.

[5] Sundararaman Akshay, Nathalie Bertrand, Serge Haddad, and Loic Helouet. The steady-state control problem for markov decision processes. In International Conference on Quantitative Evaluation of Systems, pages 290 –304. Springer, 2013.

Looking Ahead

Go here for information on monthly MADLab/AFRL WebEx's featuring a short research discussion

The next WebEx will be Feb 6 from 11:30-12 Central time. Professor Jerry Zhu will give a presentation on adversarial ML topics.

Recent and Upcoming SILOS http://silo.ece.wisc.edu/web/content/seminars
SILO is a weekly seminar series which hosts a catered lunch every Wednesday at the Wisconsin Institute for Discovery for graduate students from across campus. Researchers from Computer Science, Engineering, Mathematics and Statistics make up the core of SILO, but those from other fields are strongly encouraged to participate.

Save the date: MADLab/AFRL Workshop June 5th 8:30-5
The purpose of the workshop is to facilitate engagement with AFRL government scientists and to foster a collaborative research environment.

Registration information coming soon, see the website for more information.

Save the date: Midwest Machine Learning Symposium (MMLS) June 6-7

The Midwest Machine Learning Symposium 2019 (MMLS'19) is being organized by Dimitris Papailiopoulos and Yingyu Liang at UW-Madison, together with colleagues Laura Balzano from U. Michigan, Abhishek Gupta from OSU, Steve Hanneke from TTIC, and Niao He from UIUC.

It aims to convene regional machine learning researchers for stimulating discussions and debates, to foster cross-institutional collaboration, and to showcase the collective talent of ML researchers at all career stages. The event combines individual research pursuits to form a broader perspective and helps to sustain a thriving ML community in the Midwest.

Publications

- Herman Kamper, Greg Shakhnarovich, and Karen Livescu, “Semantic speech retrieval with a visually grounded model of untranscribed speech,” IEEE/ACM Transactions on Audio, Speech, and Language Processing 27 (1): 89-98, January 2019

- Xiaojin Zhu. An optimal control view of adversarial machine learning. arXiv:1811.04422, 2018.

- Kwang-Sung Jun, Lihong Li, Yuzhe Ma, and Xiaojin Zhu. Adversarial attacks on stochastic bandits. In Advances in Neural Information Processing Systems, 2018.

- Hongyi Wang, Scott Sievert, Zachary Charles, Stephen Wright, Dimitris Papailiopoulos.
  "ATOMO: Communication-efficient Learning via Atomic Sparsification,” appeared in the 32nd Conference on Neural Information Processing Systems (NeurIPS), 2018.

- Lingjiao Chen, Hongyi Wang, Jinman Zhao, Paraschos Koutris, Dimitris Papailiopoulos.
  "The Effect of Network Width on the Performance of Large-batch Training,”
  Appeared in the 32nd Conference on Neural Information Processing Systems (NeurIPS), 2018.

- Urvashi Oswal and Robert Nowak. Scalable Sparse Subspace Clustering via Ordered Weighted l1 Regression. arXiv:1807.03746v1 [stat.ML] 10Jul2018

- Lingjiao Chen, Hongyi Wang, Zachary Charles, and Dimitris Papailiopoulos. DRACO: Byzantine-resilient Distributed Training via Redundant Gradients. In International Conference on Machine Learning, 2018.

- Zachary Charles, and Dimitris Papailiopoulos. Stability and Generalization of Learning Algorithms that Converge to Global Optima. In International Conference on Machine Learning, 2018.

- Hongyi Wang, Scott Sievert, Zachary Charles, Dimitris Papailiopoulos, and Stephen Wright. ATOMO: Communication-efficient Learning via Atomic Sparsification. arXiv preprint arXiv:1806.04090 (2018).

- Lingjiao Chen, Hongyi Wang, Jinman Zhao, Dimitris Papailiopoulos, and Paraschos Koutris. The Effect of Network Width on the Performance of Large-batch Training. arXiv preprint arXiv:1806.03791 (2018).

- Ayon Sen, Scott Alfeld, Xuezhou Zhang, Ara Vartanian, Yuzhe Ma, and Xiaojin Zhu. Training set camouflage. In Conference on Decision and Game Theory for Security (GameSec), 2018

- Yuzhe Ma, Kwang-Sung Jun, Lihong Li, and Xiaojin Zhu. Data poisoning attacks in contextual bandits. In Conference on Decision and Game Theory for Security (GameSec), 2018

- Jiefeng Chen, Xi Wu, Yingyu Liang, Somesh Jha. Improving Adversarial Robustness by Data-Specific Discretization.

- Yuanzhi Li, Yingyu Liang. Learning Overparameterized Neural Networks via Stochastic Gradient Descent on Structured Data. Neural Information Processing Systems, 2018.

- Huaizu Jiang, Erik Learned-Miller, Gustav Larsson, Michael Maire, Greg Shakhnarovich, Self-Supervised Relative Depth Learning for Urban Scene Understanding, ECCV 2018

- K. Krishna, S. Toshniwal, and K. Livescu, Hierarchical multitask learning for CTC-based speech recognition.

Subscribe to This Newsletter