Knowledge Product: IoBT Resiliency (Fighting Network)

ARL POC: Lance Kaplan, *lance.m.kaplan.civ@army.mil*, (301) 394-0807
Consortium POC: Paulo Tabuada, *tabuada@ee.ucla.edu*, (310) 794-4266

OPPORTUNITY: Recent advances in machine learning compute uncertainty and defend against adversarial manipulation. These advances enable robust machine-learned systems with safety guarantees, but the state of the art mainly features centralized solutions operating in relatively benign environments. The IoBT has extended the field of robust learning to scalable, decentralized, and dynamic solutions. IoBT solutions adapt to dynamic (non-stationary) environments. They advance the Pareto frontier of policies that balance performance (obtaining an accurate and timely situational awareness) and resilience to various adversarial attacks that manipulate data.

GOAL: Realize a distributed automated system that can manage and control resources in an MDO effect/decision loop operating in a contested (e.g., intermittently connected), low bandwidth network environment with ever changing sensor phenomenology due to the fog of war and adversarial deception. The automated system needs to be able to control its decision cadence to quickly take confident actions when possible and slow down and recalibrate as the environment changes to ensure all actions are safe.

NEW SCIENCE: The IoBT CRA has advanced resiliency of battlefield analytics due to seven different sets of IoBT innovations: (1) identification of the Pareto frontier to characterize the tradeoffs in accuracy, latency and resilience due to noisy data, adversarial manipulation and lack of labeled training data, (2) new neural network architectures that tradeoff accuracy for resilience near the Pareto frontier, (3) novel out-of-distribution classifiers that can detect "surprises" due to either adversarial manipulation and/or distributional shifts, (4) new robust filtering concepts that naturally remove non-conforming data items/sources, (5) redundancy in sensing and network flows to ensure sufficient signal strength after conformal filtering (6) insertions of dynamic models of the exogenous environment into reinforcement learning to detect the need to adapt to new command and control policies, and (7) sequential determination of confidence bounds in reinforcement learning to provide guarantees for safer operation.

Stochastic queuing theory helps to establish the Pareto frontier for the possible tradeoffs to achieve resiliency. It has been demonstrated that neural stochastic differential equation architectures invoke attributional smoothness leading to resilience to adversarial manipulation. Furthermore, an ensemble of attributional measures, aleatoric and epistemic uncertainty measures, and top-down neuro-symbolic contextual models are able to detect subtle distribution shifts and adversarial manipulations. Network error correction and information fusion has led to more robust estimation for situational awareness by censuring the effects of outliers. Finally, the emerging theory of Seldonian reinforcement learning leverages confidence intervals to select safe state-action policies that are guaranteed with arbitrarily high probability to improve upon the current policy. The incorporation of dynamical modeling enables performance guarantees even in non-stationary environments.

SCIENTIFIC CHALLENGES AND RISK: The above contributions are noteworthy in overcoming such challenges as (1) distributed processing at scale, (2) reinforcement learning in non-stationary environments, (3) limited supervised data for on-line learning, and (4) theoretical understanding of the tradeoffs of machine learning systems. Systems resilient to adversarial manipulation tend to be less accurate, and incorporation of dynamical models to adapt to non-stationary environments stresses the training requirements for reinforcement learning methods. The risks are that the extreme environments and scalability requirements might not allow an admissible balance of accuracy, latency, and resilience.

RESULTS FROM EXPERIMENTATION: The above innovations were validated experimentally where possible and/or in simulation, where scale of the experiment warranted a simulation-based exploration. For example, the ability of machine learning models to detect surprise from adversarial examples and out-of-distribution inputs was validated on the problems of novel object detection and out of distribution detection. Tests on multiple data sets showed the proposed algorithms consistently outperformed existing alternatives. Similarly, the new neural network architectures developed to improve robustness and interpretability were tested on multiple data sets showing again improved performance over existing techniques. Results on safety of Reinforcement Learning (RL) in the presence of non-stationarity showed how the new safe RL algorithms always improve upon existing policies (despite the non-stationarity). In contrast, high-confidence safe RL algorithms that do not account for non-stationarity resulted in policies with worse performance.

TECHNOLOGY TRANSITION: Resiliency innovations, developed in the IoBT CRA, bootstrapped winning proposals from DARPA and IARPA. The **DARPA** award (on Assured Autonomy) demonstrates that attribution-based confidence algorithms developed in the CRA can be used to provide runtime assurance for autonomous systems. We also demonstrated our approach to detect novelty on data collected by Boeing and provided to DARPA for this program. This is a joint effort between SRI (an IoBT participant), MIT, and Caltech. The **IARPA** (TrojAI) award demonstrates that attributions can be used for static analysis of machine learning models and detecting whether the model has an embedded backdoor in it. This is a joint effort between SRI (an IoBT participant), Boston University, and Stonybrook. The knowledge gain in resilient IoBT networks will support and shape the ARL's Foundational Research Competencies in Network, Cyber and Computational Sciences (NC&CS) and Military Information Sciences (MIS). Specifically, (1) ensuring resilient networks for distributed analytics in MDO within NC&CS competency and (2) enabling delivery of timely and mission–aware information at speed and scale in all-domain operation within MIS competency. Resiliency innovations can also enrich the IoBT Digital Twin system being developed in an active transition to Boeing Inc. that encapsulates several IoBT innovations.

IMPACT ON THE SCIENTIFIC COMMUNITY: The work produced multiple academic publications in premier peer-reviewed conferences and journals including IJCAI, IEEE Transactions on Automatic Control, Asilomar, MILCOM, AAAI, AISTATS, ICLR, ICML, NeurIPS, etc. The developed research was recognized by the relevant scientific communities in multiple ways including the Sandia National Labs Research Award, the guest editorship of the IEEE JSAIT Special Issue on Sequential, Active and Reinforcement Learning and of the Special Issue of the ACM Transactions on Cyber Physical Systems on Artificial Intelligence and Cyber-Physical Systems, the special issue of Acta Informatica on Synthesis, co-chairing the program committees of the 2019 9th ACM/IEEE International Conference on Cyber-Physical Systems, the 2019 IEEE International Symposium on Information Theory, the 2020 Nasa Formal Methods, the 2022 Design Automation for CPS and IoT, and receiving the the 10-year Most Influential Paper award at ICSE'2021.

POTENTIAL ARMY CAPABILITIES: A recent Focused Excursion organized by DEVCOM and FCC (a collaborative investigative process that supports concept and capability development with participation from FCC DoC, DoIS, DAC, and ARL) investigated IoBT implications on select learning demands of the Army Concepts Framework 2040. Among other hypotheses, the excursion suggested the hypothesis that "IoBT capabilities that provide robustness, resiliency, and adaptive services in the face of adversarial disruption and deception in conjunction with camouflaged sensing and communication, are essential to maintain commander confidence in data and reduce risk". Distributed processing will significantly reduce the negative effects of delayed/disconnected, intermittently connected, low-bandwidth communications and offer protection through dispersion. Information validation via multi-modal and multi-vantage sensing will lead to multiple corroborations of and stronger confidence in IoBT inferences.

EXPERIMENTAL DEMONSTRATIONS:

See videos on IoBT Edge Resiliency Innovations at: https://abdelzaher.cs.illinois.edu/RMB22-Demos.html

KEY PUBLICATIONS*:

1. T. Abdelzaher, N. Ayanian, T. Basar, S. Diggavi, J. Diesner, D. Ganesan, R. Govindan, S. Jha, T. Lepoint, B. Marlin, K. Nahrstedt, D. Nicol, R. Rajkumar, **S. Russell**, S. Seshia, F. Sha, P. Shenoy, M. Srivastava, G. Sukhatme, **A. Swami**, P. Tabuada, D. Towsley, N. Vaidya, V. Veeravalli, "Toward an Internet of Battlefield Things: A Resilience Perspective," *IEEE Computer*, 51(11), 2018.
2. S. Bagchi, T.F. Abdelzaher, R. Govindan, P. Shenoy, A. Atrey, P. Ghosh, and R. Xu. "New Frontiers in IoT: Networking, Systems, Reliability, and Security Challenges," *IEEE Internet of Things Journal* 7(12), 2020.
3. Y. Chandak, S. Niekum, B. Castro da Silva, E. Learned-Miller, E. Brunskill, and P. S. Thomas, "Universal Off-Policy Evaluation," *NeurIPS* 2021.
4. A. Cobb, **B. Jalaian**, **N. Bastian**, **S. Russell**, "Robust Decision-making in the Internet of Battlefield Things using Bayesian Neural Networks," In Proc. *Winter Simulation Conference*, December 2021.
5. S. Jha, S. Raj, S.L. Fernandes, S. Jha, **B. Jalaian**, **G. Verma**, and **A. Swami**, "Attribution-Based Confidence Metric for Deep Neural Networks," *NeurIPS* 2019.
6. S. Jha, **B. Jalaian**, A. Roy, **G. Verma**, "TRINITY: Trust, Resilience, and Interoperability of Machine Learning Models," *Game Theory and Machine Learning for Cyber Security* (Wiley IEEE Press), December 2021. https://github.com/SRI-CSL/Trinity
7. N. Karampatziakis, P. Mineiro, and A. Ramdas, "Off-policy Confidence Sequences," *ICML* 2021.
8. J. Kostas, Y. Chandak, S. M. Jordan, G. Theocharous, and P. Thomas. "High Confidence Generalization for Reinforcement Learning," *ICML* 2021.
9. B. Marlin, T. Abdelzaher, G. Ciocarlie, A. Cobb, M. Dennison, **B. Jalaian**, **L. Kaplan**, **T. Raber**, **A. Raglin**, **P. Sharma**, M. Srivastava, T. Trout, M. Vadera, and **M. Wigness**, "On Uncertainty and Robustness in Large-Scale Intelligent Data Fusion Systems, In Proc. *CogMI*, December 2020.
10. Y. Mao, S. Diggavi, C. Fragouli, and P. Tabuada, "Secure State-Reconstruction over Networks Subject to Attacks," *IEEE Control Systems Letters*, 5(1), 2021.
11. Y. Mao, A. Mitra, S. Sundaram, and P. Tabuada, "On the Computational Complexity of the Secure State-Reconstruction Problem," *Automatica*, 136, 2022.
12. A. Roy, A. Cobb, **N. Bastian**, **B. Jalaian**, and S. Jha, "Runtime Monitoring of Deep Neural Networks using Top-Down Context Models Inspired by Predictive Processing and Dual Process Theory," *AAAI Spring Symposium*, 2022.
13. P. Thomas, E. Brunskill, B. Castro da Silva, E. Learned-Miller, Y. Chandak, "Universal Off-Policy Evaluation," In Proc. *NeurIPS* 2021.
14. M. Vadera, **B. Jalaian**, and B. Marlin. "Generalized bayesian posterior expectation distillation for deep neural networks." In Conference on Uncertainty in Artificial Intelligence," In Proc. *PMLR* 2020.
15. **M. Wigness**, **T. Pham**, **S. Russell**, Tarek Abdelzaher, "Efficient and Resilient Edge Intelligence for the Internet of Battlefield Things," In Proc. *NATO STO IST Symposium on Artificial Intelligence, Machine Learning and Big Data*, October 2021.

*Note: Names in **blue** are government co-authors.