Internet of Battlefield Things

COLLABORATIVE RESEARCH ALLIANCE

Resiliency to Adversarial Deception & Battlefield Conditions for Decision Loop Analytics

Lance Kaplan, Research Area Lead
DEVCOM, Army Research Laboratory

Paulo Tabuada, Research Area Lead
University of California, Los Angeles

March 18, 2022

OVERVIEW

Resiliency for Decision Loop Analytics

GOAL

- Detection of changes in the environment and/or adversarial deception
- Adaptation of IoBT resources to obtain accurate situational awareness in the face of highly dynamic environments and deception
- Safe actuation of resources to shape the battlespace despite uncertainty due to dynamic environments and deception

NEW SCIENCE

- Resilient Filtering
  - Cleanse manipulated data before ingesting it
- Resilient Inferencing
  - Detect that adversarial inputs have been ingested
  - Architectures insensitive to adversarial data manipulation
- Safe Reinforcement Learning
  - Make safe resource allocation decisions despite changes in the environment created by the adversary

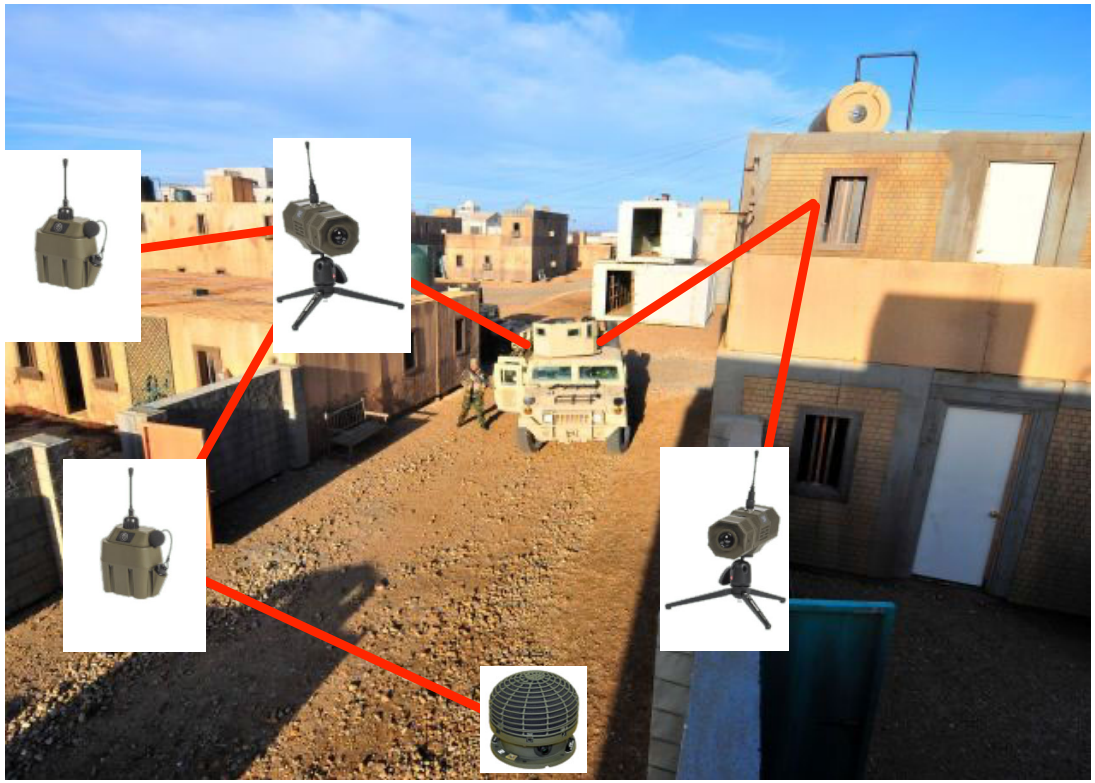

Resiliency for Decision Loop Analytics

GOAL

- Detection of changes in the environment and/or adversarial deception
- Adaptation of IoBT resources to obtain accurate situational awareness in the face of highly dynamic environments and deception
- Safe actuation of resources to shape the battlespace despite uncertainty due to dynamic environments and deception

NEW SCIENCE

- Resilient Filtering
  - Cleanse manipulated data before ingesting it
- Resilient Inferencing
  - Detect that adversarial inputs have been ingested
  - Architectures insensitive to adversarial data manipulation
- Safe Reinforcement Learning
  - Make safe resource allocation decisions despite changes in the environment created by the adversary

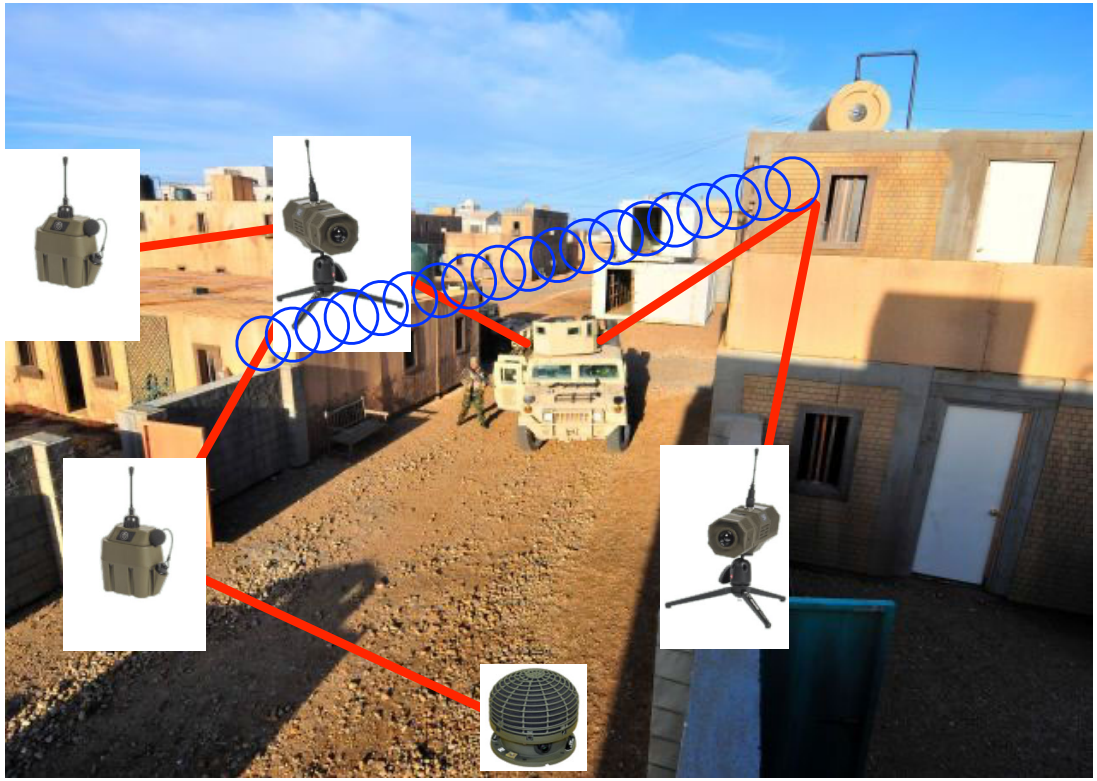

RESEARCH CONTRIBUTIONS

IoBT REIGN

Resiliency for Decision Loop Analytics

Executing distributed, heterogeneous, and scalable mission-essential tasks and responding to dynamically changing battlefield and threat conditions within contested/constrained environments

Detection and Adaptation to Contextual Shifts

- Detection of out-of-distribution and adversarial inputs to NN models
- Detection and mitigation of sensor attacks in target tracking
- Detection and mitigation of data attacks in decentralized learning

Intrinsically Resilient Systems

- New NN architectures for improved resiliency and interpretability
- Robust inference in the presence of distribution shifts

Safe Resource Allocation

- Characterization of tradeoffs in accuracy, latency and resiliency due to noisy data, adversarial manipulation and lack of labeled training data
- Reinforcement learning algorithms that produce safe policies in nonstationary environments

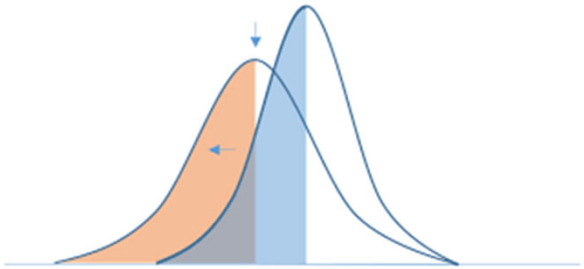

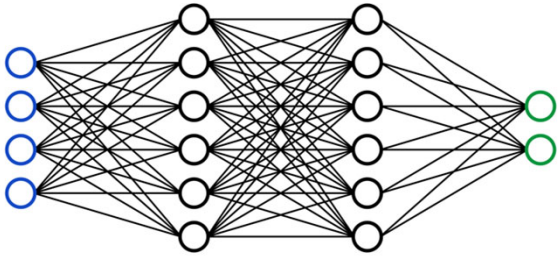

RESEARCH CONTRIBUTIONS

IoBT REIGN

Resiliency for Decision Loop Analytics

Executing distributed, heterogeneous, and scalable mission-essential tasks and responding to dynamically changing battlefield and threat conditions within contested/constrained environments

Detection and Adaptation to Contextual Shifts

- Detection of out-of-distribution and adversarial inputs to NN models
- Detection and mitigation of sensor attacks in target tracking
- Detection and mitigation of data attacks in decentralized learning

Intrinsically Resilient Systems

- New NN architectures for improved resiliency and interpretability
- Robust inference in the presence of distribution shifts

Safe Resource Allocation

- Characterization of tradeoffs in accuracy, latency and resiliency due to noisy data, adversarial manipulation and lack of labeled training data
- Reinforcement learning algorithms that produce safe policies in nonstationary environments

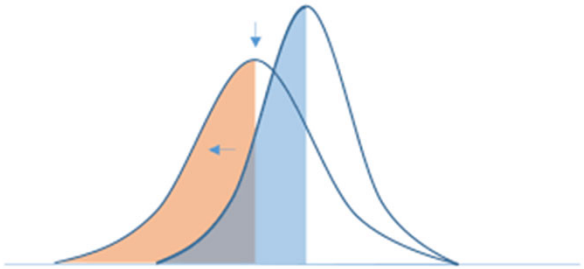

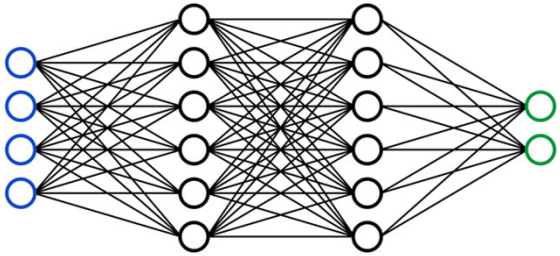

DETECTION OF OUT-OF-DISTRIBUTION AND ADVERSARIAL INPUTS

IoBT REIGN

DEVCOM ARMY RESEARCH LABORATORY

Design of ML models with formal robustness
and resiliency properties notwithstanding reduced training data and time

Army Relevance and Value Proposition

IoBTs need to withstand complex battlefield environments as well as malicious attacks by purposeful and intelligent adversaries

Prior State-of-the-Art

- Existing machine learning techniques are vulnerable to different types of attacks making them too brittle for the battlefield

Technical Approach

- Use a theory of mind approach to monitor ML models in IoBTs to detect “surprise” and provide confidence on predictions
- Combine information-theoretic approaches for training with inference-time monitors to yield formal guarantees of robustness and resiliency

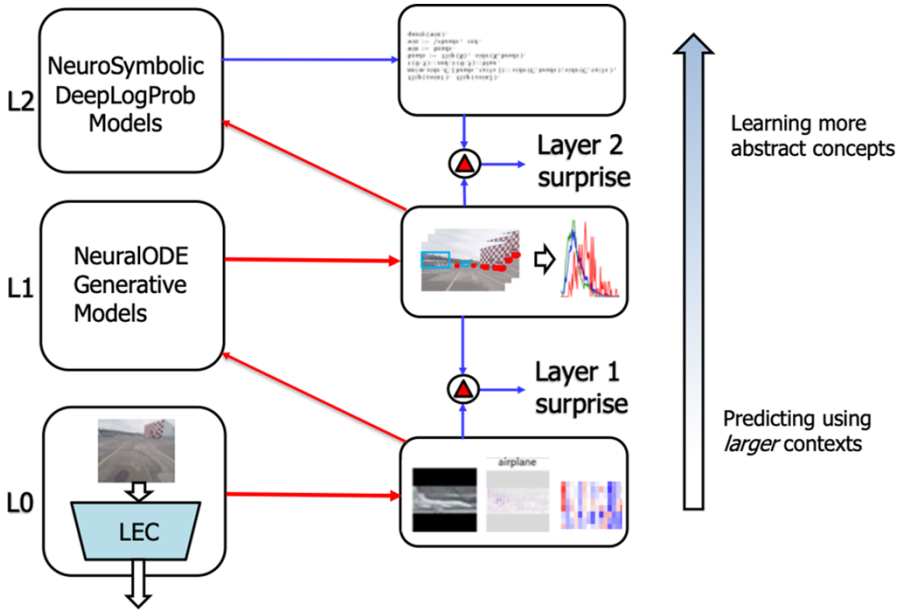

Contribution: Theory of mind, information theory, and control theory offer new pathways to develop robust and resilient IoBTs

DETECTION OF OUT-OF-DISTRIBUTION AND ADVERSARIAL INPUTS

Army Relevance: Build trust and resiliency in inferences made by deep neural network models

Contribution: Ability to discard predictions made by out-of-distribution or adversarial inputs.

- Use **surprise** or **novelty** to determine if a prediction can be trusted.
  - Sample the neighborhood of an input to determine if the output is consistent (no surprise).
  - Use attribution (e.g., integrated gradients) to make the sampling process efficient.
- Higher level knowledge can be used to detect surprise at more abstract levels, e.g., generative models or NeuroSymbolic models.

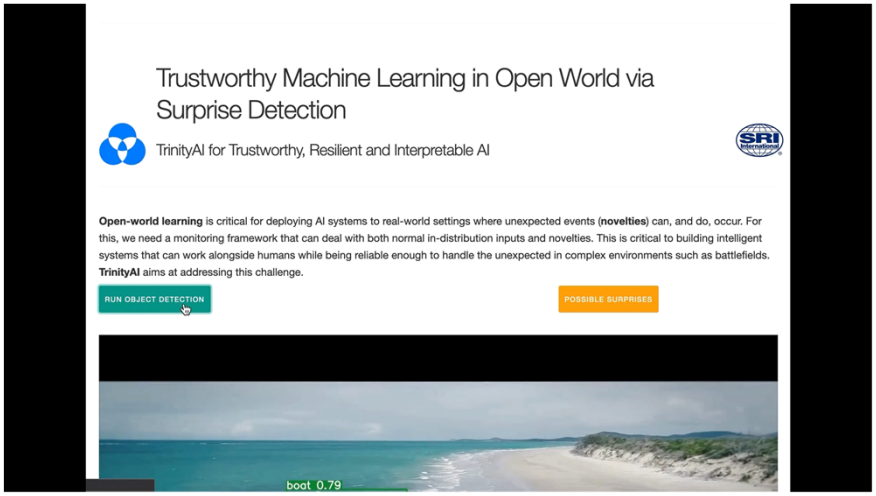

Illustrating low confidence of predictions from out-of-distribution inputs.

Full demo video available:
Trustworthy ML with Surprise Detection

Key Publications:

Attribution-Based Confidence Metric For Deep Neural Networks
S. Jha, S. Raj, S.L. Fernandes, S.K. Jha, S. Jha, B. Jalaian, G. Verma, A. Swami, *NeurIPS* 2019.

Runtime Monitoring of Deep Neural Networks Using Top-Down Context Models Inspired by Predictive Processing and Dual Process Theory
A. Roy, A. Cobb, N. D. Bastian, B. Jalaian, S. Jha, *Designing Artificial Intelligence for Open Worlds, AAAI Spring Symposium*, 2022.

DETECTION OF OUT-OF-DISTRIBUTION AND ADVERSARIAL INPUTS

Experiments and Results:

- Confidence of predictions in the presence of new object classes
- Confidence of predictions in the presence of Out-Of-Distribution (OOD) inputs

Novel object detection: Datasets: Tiny imagenet with 200 object classes. 20 classes are available during training and rest 180 classes considered as the novel objects

Novel object recognition

TinyImageNet	OpenMax (CVPR16)	G-OpenMax (BMVC17)	OSRCI (ECCV18)	C2AE (CVPR19)	CROSR (CVPR19)	Gen-dis (CVPR20)	Ours
AUROC	57.6	58.0	58.6	58.1	58.9	64.7	**73.26**

Closed set recognition

TinyImageNet	Gen-dis (CVPR20) Resnet-18	Gen-dis (CVPR20) WideResnet-28-10	Ours
DTACC	49.2	55.9	**74.74**

OOD detection: Datasets: MNIST, KMNIST, F-MNIST, CIFAR-10, CIFAR100, STL10, SVHN, LSUN, ImageNet

In-dist (model)	OOD dataset	TNR (TPR=95%)	AUROC	DTACC
MNIST (LeNet5)	KMNIST	67.72 / 80.52 / **91.82**	92.98 / 96.53 / **98.3**	85.99 / 90.82 / **94.01**
	F-MNIST	58.47 / 63.33 / **74.49**	90.76 / 94.11 / **95.55**	83.21 / 87.76 / **90.98**
CIFAR10 (ResNet34)	STL10	10.63 / 13.9 / **17.4**	61.56 / 66.47 / **67.52**	59.22 / 62.75 / **63.7**
	SVHN	72.85 / 53.16 / **88.2**	93.85 / 93.85/ **97.69**	85.4 / 89.173 / **92.14**
	Imagenet	46.54 / 68.41 / **74.53**	90.45 / 95.02 / **95.73**	83.06 / 88.63 / **89.73**
	LSUN	45.16 / 77.53 / **81.23**	89.63 / 96.51 / **96.87**	81.83 / 90.64 / **91.19**
	SCIFAR100	37 / 38.39 / **61.11**	86.13 / 88.86 / **94.74**	78.5 / 82.51 / **90.53**
CIFAR10 (ResNet50)	STL10	12.19 / 10.33 / **16**	60.29 / 61.95 / **66.39**	58.57 / 59.36 / **62.28**
	SVHN	86.61 / 34.49/ **91.06**	84.41 / **98.19** / 91.98	91.25 / 76.72 / **93.2**
	Imagenet	73.23 / 29.48 / **75.96**	94.91 / 84.3 / **95.79**	88.23 / 77.19 / **89.26**
	LSUN	80.72 / 32.18 / **81.38**	96.51 / 87.09 / **96.93**	90.59 / 80.07 / **91.79**
	SCIFAR100	47.44 / 21.06 / **48.33**	86.16 / 77.42/ **92.98**	78.69 / 71.43 / **88.27**
SVHN (DenseNet)	STL10	45.91 / 81.66 / **87.76**	77.6 / 96.97 / **97.63**	72.62 / 92.29 / **93.35**
	CIFAR10	37.23 / 80.82 / **86.42**	73.14 / 96.8 / **97.37**	68.92 / 92.27 / **92.86**
	Imagenet	62.76 / 85.44 / **93.44**	85.41 / 97.29 / **98.38**	79.94 / 93.39 / **94.53**
	LSUN	62.91 / 76.87 / **89.73**	86.06 / 96.37 / **97.73**	80.04 / 92.43 / **93.55**
	SCIFAR100	48.17 / 86.06 / **96.72**	78.94 / 97.43 / **98.24**	73.72 / 93.02 / **96.26**

Summary:

Highest performing confidence predictor for deep neural networks

DETECTION OF OUT-OF-DISTRIBUTION AND ADVERSARIAL INPUTS

Experiments and Results:

- Confidence of predictions in the presence of new object classes
- Confidence of predictions in the presence of Out-Of-Distribution (OOD) inputs

Novel object detection: Datasets: Tiny imagenet with 200 object classes. 20 classes are available during training and rest 180 classes considered as the novel objects

Novel object recognition

TinyImageNet	OpenMax (CVPR16)	G-OpenMax (BMVC17)	OSRCI (ECCV18)	C2AE (CVPR19)	CROSR (CVPR19)	Gen-dis (CVPR20)	Ours
AUROC	57.6	58.0	58.6	58.1	58.9	64.7	**73.26**

Closed set recognition

TinyImageNet	Gen-dis (CVPR20) Resnet-18	Gen-dis (CVPR20) WideResnet-28-10	Ours
DTACC	49.2	55.9	**74.74**

OOD detection: Datasets: MNIST, KMNIST, F-MNIST, CIFAR-10, CIFAR100, STL10, SVHN, LSUN, ImageNet

In-dist (model)	OOD dataset	TNR (TPR=95%)	AUROC	DTACC
MNIST (LeNet5)	KMNIST	67.72 / 80.52 / **91.82**	92.98 / 96.53 / **98.3**	85.99 / 90.82 / **94.01**
	F-MNIST	58.47 / 63.33 / **74.49**	90.76 / 94.11 / **95.55**	83.21 / 87.76 / **90.98**
CIFAR10 (ResNet34)	STL10	10.63 / 13.9 / **17.4**	61.56 / 66.47 / **67.52**	59.22 / 62.75 / **63.7**
	SVHN	72.85 / 53.16 / **88.2**	93.85 / 93.85/ **97.69**	85.4 / 89.173 / **92.14**
	Imagenet	46.54 / 68.41 / **74.53**	90.45 / 95.02 / **95.73**	83.06 / 88.63 / **89.73**
	LSUN	45.16 / 77.53 / **81.23**	89.63 / 96.51 / **96.87**	81.83 / 90.64 / **91.19**
	SCIFAR100	37 / 38.39 / **61.11**	86.13 / 88.86 / **94.74**	78.5 / 82.51 / **90.53**
CIFAR10 (ResNet50)	STL10	12.19 / 10.33 / **16**	60.29 / 61.95 / **66.39**	58.57 / 59.36 / **62.28**
	SVHN	86.61 / 34.49/ **91.06**	84.41 / **98.19** / 91.98	91.25 / 76.72 / **93.2**
	Imagenet	73.23 / 29.48 / **75.96**	94.91 / 84.3 / **95.79**	88.23 / 77.19 / **89.26**
	LSUN	80.72 / 32.18 / **81.38**	96.51 / 87.09 / **96.93**	90.59 / 80.07 / **91.79**
	SCIFAR100	47.44 / 21.06 / **48.33**	86.16 / 77.42/ **92.98**	78.69 / 71.43 / **88.27**
SVHN (DenseNet)	STL10	45.91 / 81.66 / **87.76**	77.6 / 96.97 / **97.63**	72.62 / 92.29 / **93.35**
	CIFAR10	37.23 / 80.82 / **86.42**	73.14 / 96.8 / **97.37**	68.92 / 92.27 / **92.86**
	Imagenet	62.76 / 85.44 / **93.44**	85.41 / 97.29 / **98.38**	79.94 / 93.39 / **94.53**
	LSUN	62.91 / 76.87 / **89.73**	86.06 / 96.37 / **97.73**	80.04 / 92.43 / **93.55**
	SCIFAR100	48.17 / 86.06 / **96.72**	78.94 / 97.43 / **98.24**	73.72 / 93.02 / **96.26**

Summary:

Highest performing confidence predictor for deep neural networks

DETECTION OF OUT-OF-DISTRIBUTION AND ADVERSARIAL INPUTS

Experiments and Results:

- Confidence of predictions in the presence of new object classes
- Confidence of predictions in the presence of Out-Of-Distribution (OOD) inputs

Novel object detection: Datasets: Tiny imagenet with 200 object classes. 20 classes are available during training and rest 180 classes considered as the novel objects

Novel object recognition

TinyImageNet	OpenMax (CVPR16)	G-OpenMax (BMVC17)	OSRCI (ECCV18)	C2AE (CVPR19)	CROSR (CVPR19)	Gen-dis (CVPR20)	Ours
AUROC	57.6	58.0	58.6	58.1	58.9	64.7	**73.26**

Closed set recognition

TinyImageNet	Gen-dis (CVPR20) Resnet-18	Gen-dis (CVPR20) WideResnet-28-10	Ours
DTACC	49.2	55.9	**74.74**

OOD detection: Datasets: MNIST, KMNIST, F-MNIST, CIFAR-10, CIFAR100, STL10, SVHN, LSUN, ImageNet

In-dist (model)	OOD dataset	TNR (TPR=95%)	AUROC	DTACC
MNIST (LeNet5)	KMNIST	67.72 / 80.52 / **91.82**	92.98 / 96.53 / **98.3**	85.99 / 90.82 / **94.01**
	F-MNIST	58.47 / 63.33 / **74.49**	90.76 / 94.11 / **95.55**	83.21 / 87.76 / **90.98**
CIFAR10 (ResNet34)	STL10	10.63 / 13.9 / **17.4**	61.56 / 66.47 / **67.52**	59.22 / 62.75 / **63.7**
	SVHN	72.85 / 53.16 / **88.2**	93.85 / 93.85 / **97.69**	85.4 / 89.173 / **92.14**
	Imagenet	46.54 / 68.41 / **74.53**	90.45 / 95.02 / **95.73**	83.06 / 88.63 / **89.73**
	LSUN	45.16 / 77.53 / **81.23**	89.63 / 96.51 / **96.87**	81.83 / 90.64 / **91.19**
	SCIFAR100	37 / 38.39 / **61.11**	86.13 / 88.86 / **94.74**	78.5 / 82.51 / **90.53**
CIFAR10 (ResNet50)	STL10	12.19 / 10.33 / **16**	60.29 / 61.95 / **66.39**	58.57 / 59.36 / **62.28**
	SVHN	86.61 / 34.49 / **91.06**	84.41 / **98.19** / 91.98	91.25 / 76.72 / **93.2**
	Imagenet	73.23 / 29.48 / **75.96**	94.91 / 84.3 / **95.79**	88.23 / 77.19 / **89.26**
	LSUN	80.72 / 32.18 / **81.38**	96.51 / 87.09 / **96.93**	90.59 / 80.07 / **91.79**
	SCIFAR100	47.44 / 21.06 / **48.33**	86.16 / 77.42 / **92.98**	78.69 / 71.43 / **88.27**
SVHN (DenseNet)	STL10	45.91 / 81.66 / **87.76**	77.6 / 96.97 / **97.63**	72.62 / 92.29 / **93.35**
	CIFAR10	37.23 / 80.82 / **86.42**	73.14 / 96.8 / **97.37**	68.92 / 92.27 / **92.86**
	Imagenet	62.76 / 85.44 / **93.44**	85.41 / 97.29 / **98.38**	79.94 / 93.39 / **94.53**
	LSUN	62.91 / 76.87 / **89.73**	86.06 / 96.37 / **97.73**	80.04 / 92.43 / **93.55**
	SCIFAR100	48.17 / 86.06 / **96.72**	78.94 / 97.43 / **98.24**	73.72 / 93.02 / **96.26**

Summary:

Highest performing confidence predictor for deep neural networks

DETECTION AND MITIGATION OF SENSOR ATTACKS

Approaches to detect attacked sensors so their malicious data can be discarded when tracking targets

Army Relevance and Value Proposition

Adversarial spoofing of sensors can be detected and mitigated to provide resiliency to near peer attacks

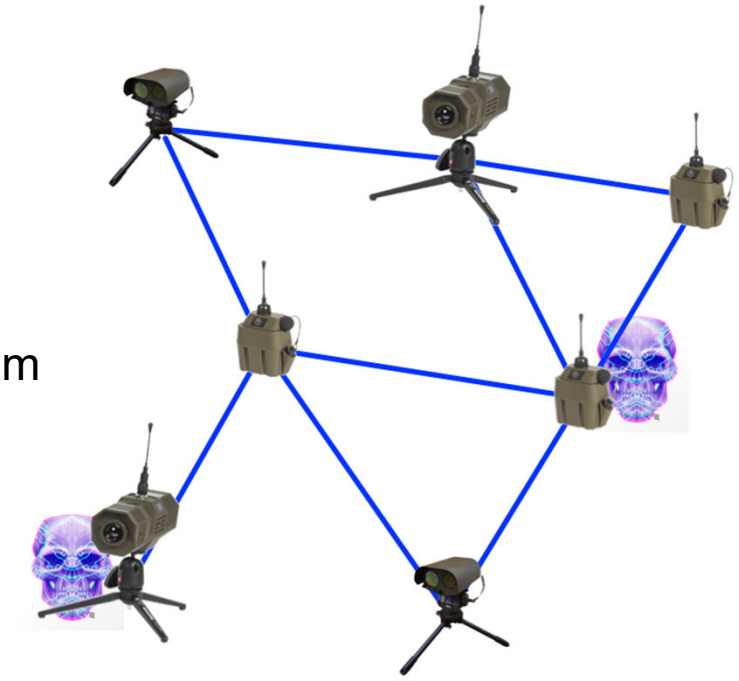

Prior State-of-the-Art

- Sensor spoofing can lead to errors in target tracking
- Finding attacked sensors is a hard combinatorial problem

Technical Approach

- Use a model of the target dynamics to detect attacks
- Identify classes of problems that can be solved in polynomial time
- Deploy efficient centralized and decentralized algorithms

Contribution: Incorporation of target dynamics improves attack detection and mitigation

DETECTION AND MITIGATION OF SENSOR ATTACKS

Army Relevance: Resiliency to sensor attacks will limit and delay adversarial action

Contribution: Detect attacked sensors and track targets despite the attacks

- A model for the target dynamics is used to determine if sensor data is compatible with physics
- The dynamics are decomposed into small components to simplify the checking of consistency with sensor data
- Attacked sensors are identified and associated data not used for tracking
- Decentralized implementation involves a consensus step that compresses sensor data while keeping enough redundancy for attack detection

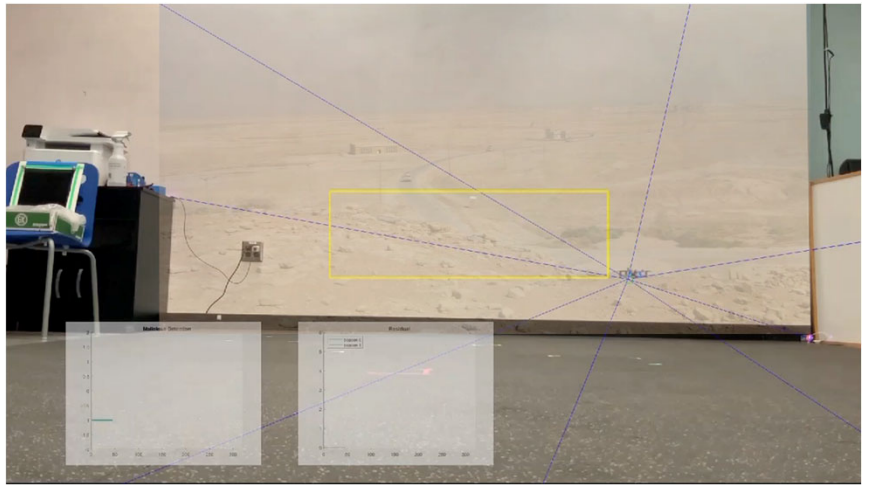

Illustrating detection of sensor attacks

Full demo video available: Resilient Localization

Key Publications:

Secure State-Reconstruction Over Networks Subject to Attacks
Y. Mao, S. Diggavi, C. Fragouli, P. Tabuada, *IEEE Control Systems Letters*, 5(1), 2021.

On the Computational Complexity of the Secure State-Reconstruction Problem
Y. Mao, A. Mitra, S. Sundaram, P. Tabuada, *Automatica*, 136, 2022.

DETECTION AND MITIGATION OF SENSOR ATTACKS

Experiments and Results:

- Identified classes of problems for which this problem is solvable in polynomial time (it is NP-hard otherwise)
- Attack detection in milliseconds for large numbers of sensors (more than 100)
- Polynomial growth of computation time with number of sensors, number of attacked sensors, and dimension of the model for the target dynamics

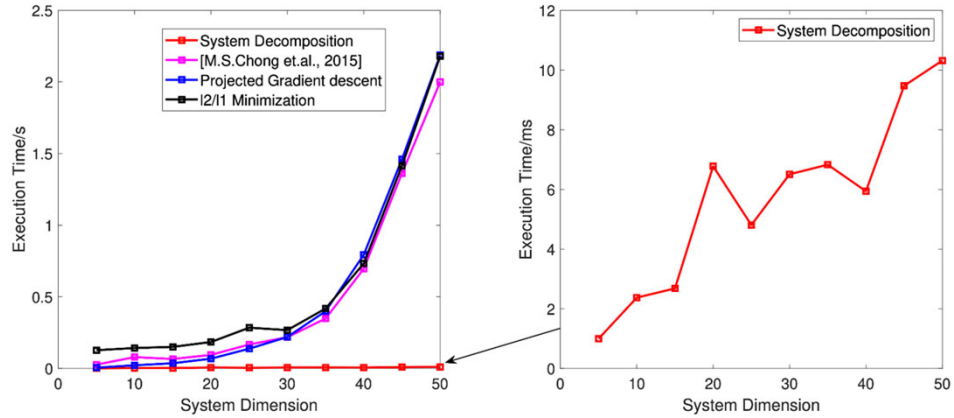

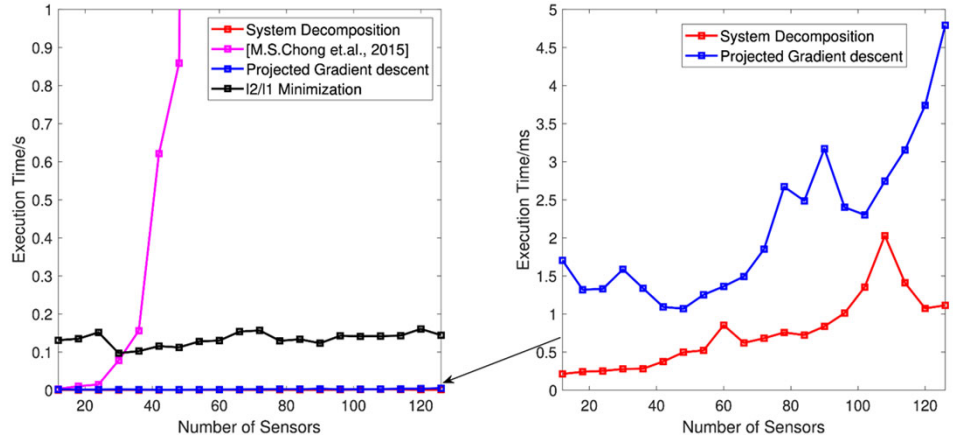

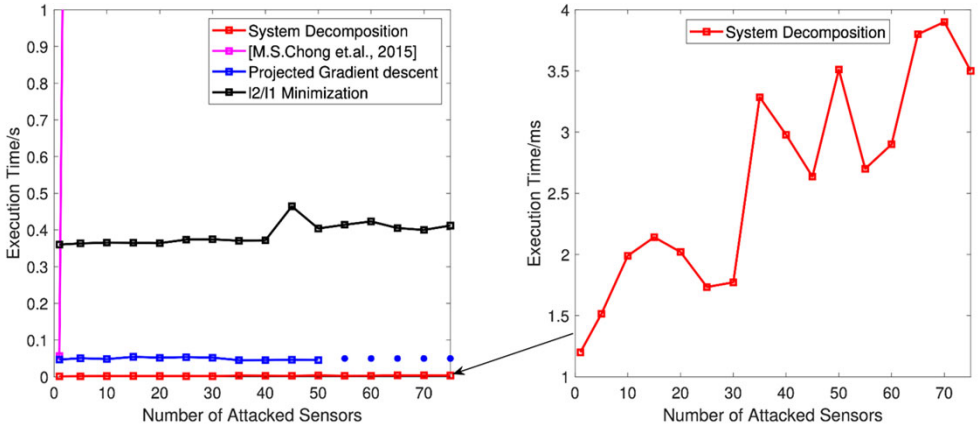

Summary:

10x faster sensor attack detection for target tracking using hundreds of sensors

DETECTION AND MITIGATION OF SENSOR ATTACKS

Experiments and Results:

- Identified classes of problems for which this problem is solvable in polynomial time (it is NP-hard otherwise)
- Attack detection in milliseconds for large numbers of sensors (more than 100)
- Polynomial growth of computation time with number of sensors, number of attacked sensors, and dimension of the model for the target dynamics

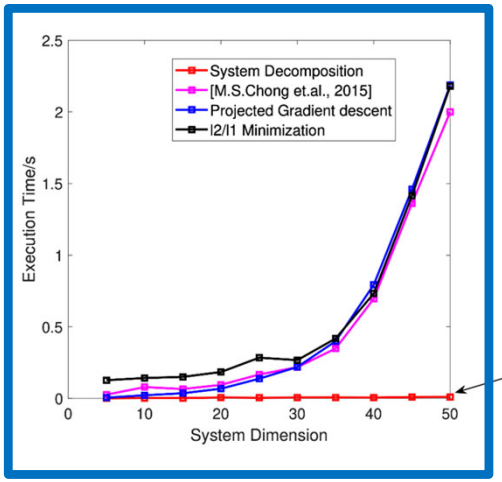

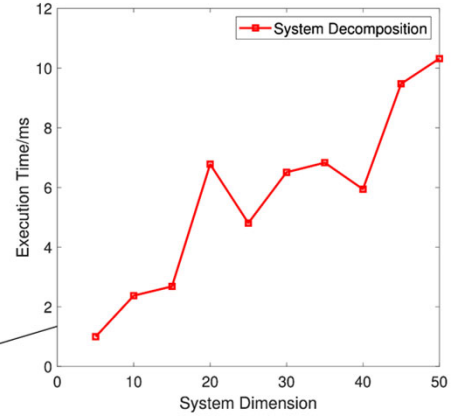

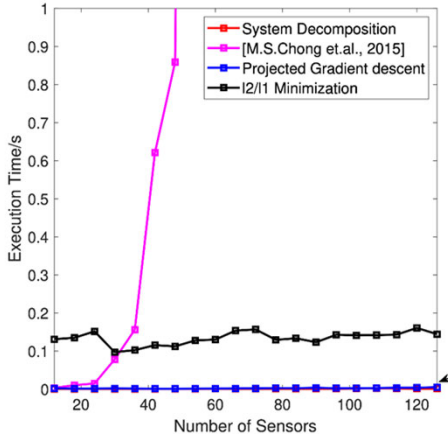

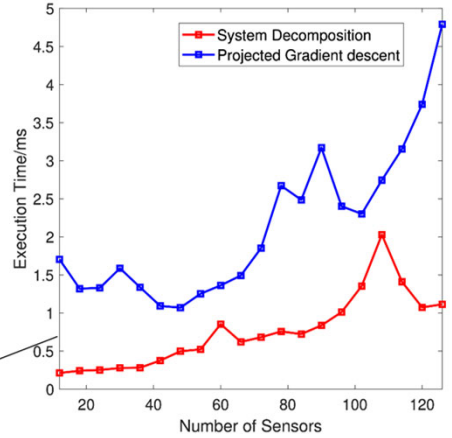

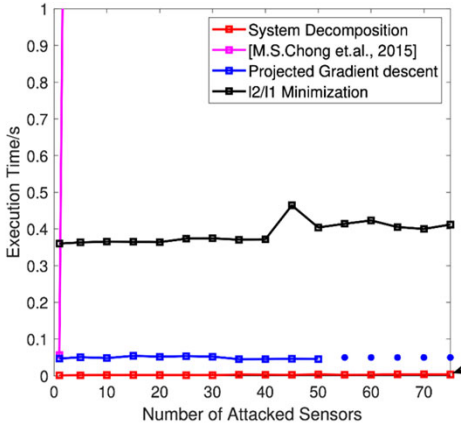

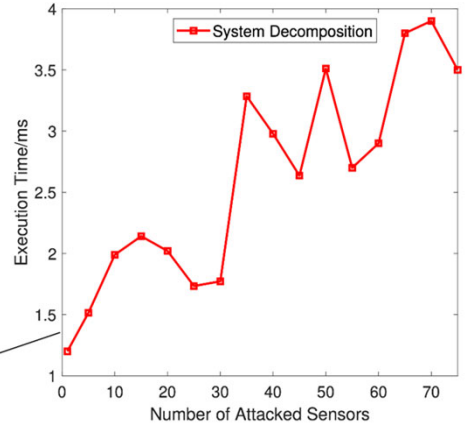

Summary:

10x faster sensor attack detection for target tracking using hundreds of sensors

DETECTION AND MITIGATION OF SENSOR ATTACKS

Experiments and Results:

- Identified classes of problems for which this problem is solvable in polynomial time (it is NP-hard otherwise)
- Attack detection in milliseconds for large numbers of sensors (more than 100)
- Polynomial growth of computation time with number of sensors, number of attacked sensors, and dimension of the model for the target dynamics

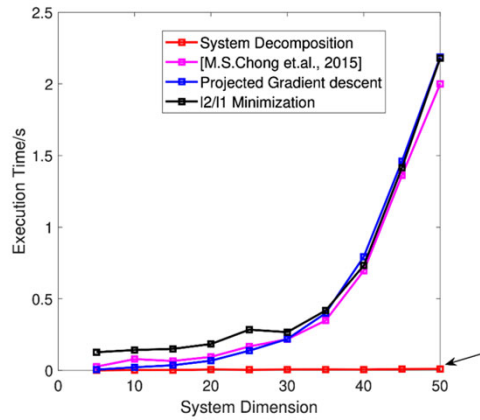

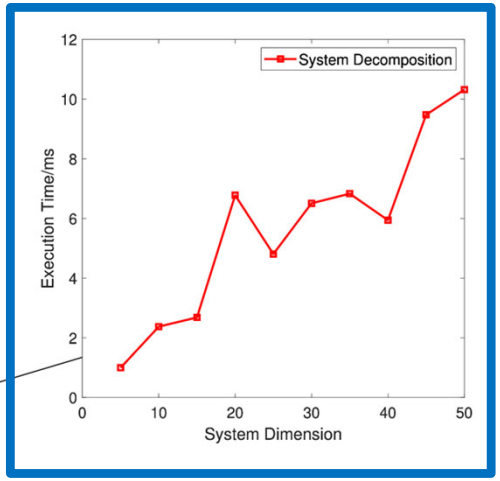

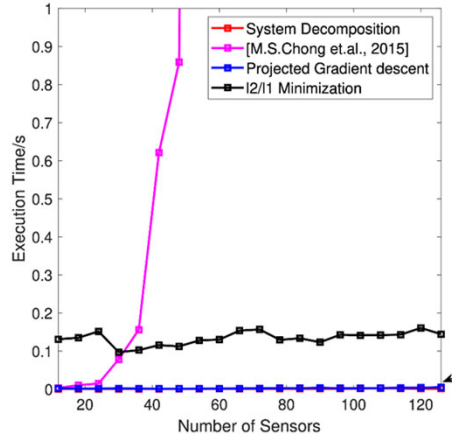

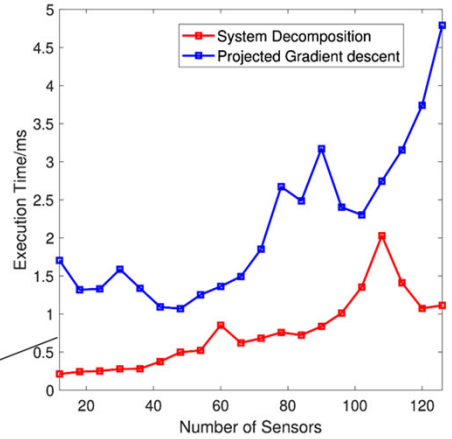

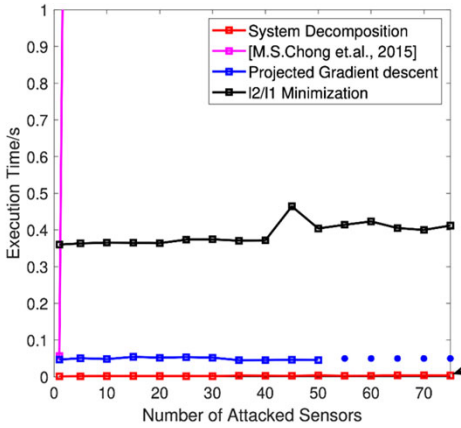

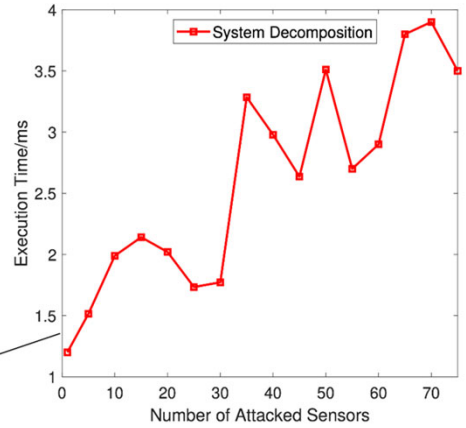

Summary:

10x faster sensor attack detection for target tracking using hundreds of sensors

DETECTION AND MITIGATION OF SENSOR ATTACKS

Experiments and Results:

- Identified classes of problems for which this problem is solvable in polynomial time (it is NP-hard otherwise)
- Attack detection in milliseconds for large numbers of sensors (more than 100)
- Polynomial growth of computation time with number of sensors, number of attacked sensors, and dimension of the model for the target dynamics

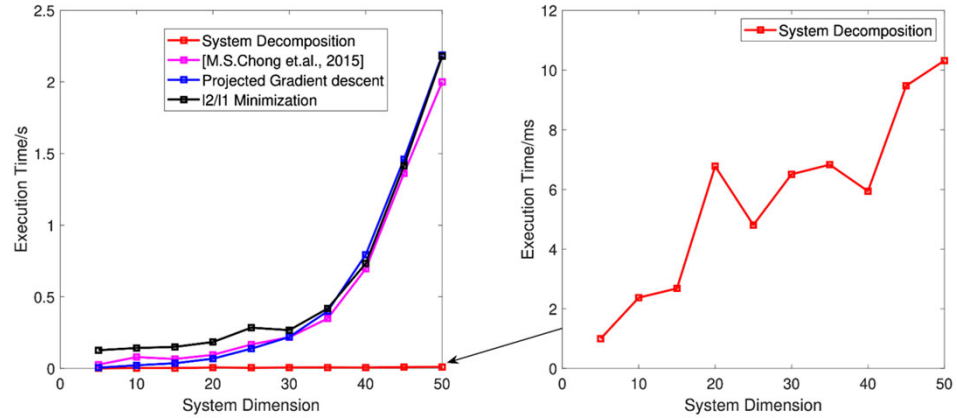

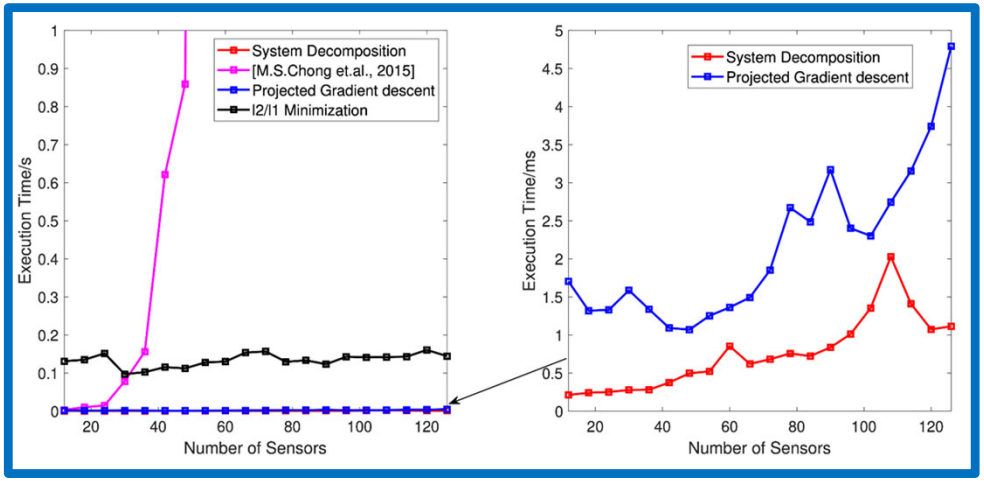

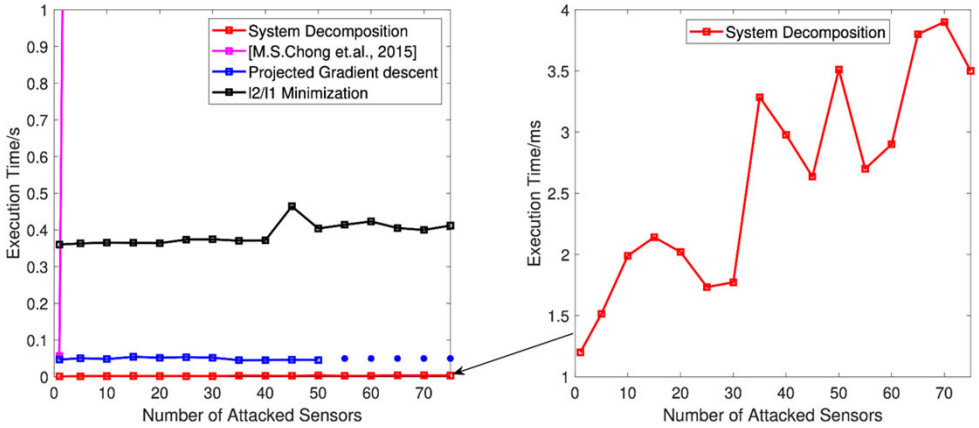

Summary:

10x faster sensor attack detection for target tracking using hundreds of sensors

DETECTION AND MITIGATION OF SENSOR ATTACKS

Experiments and Results:

- Identified classes of problems for which this problem is solvable in polynomial time (it is NP-hard otherwise)
- Attack detection in milliseconds for large numbers of sensors (more than 100)
- Polynomial growth of computation time with number of sensors, number of attacked sensors, and dimension of the model for the target dynamics

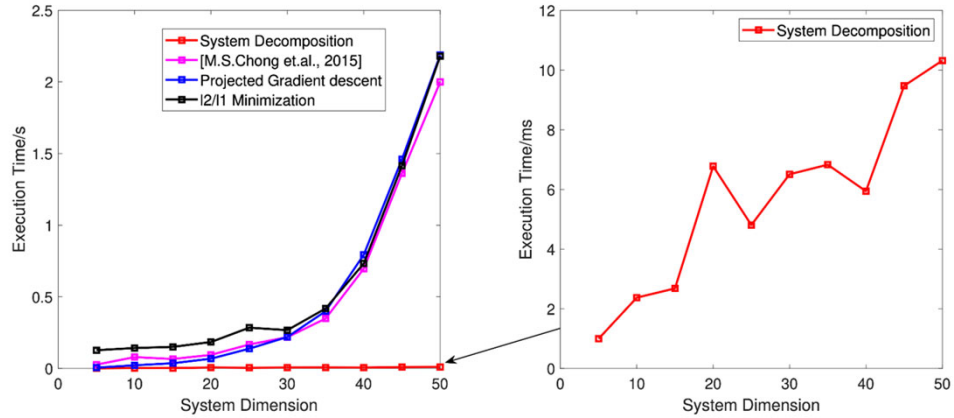

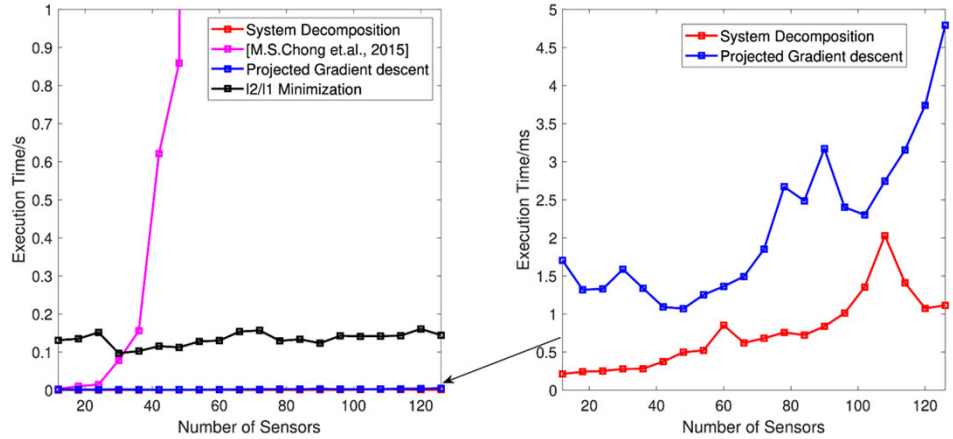

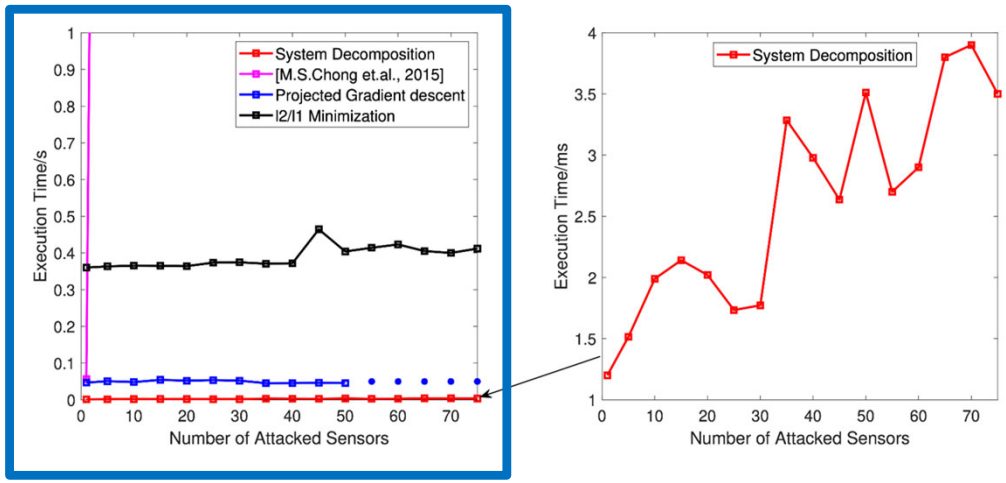

Summary:

10x faster sensor attack detection for target tracking using hundreds of sensors

UNIVERSAL OFF-POLICY EVALUATION (UnO)

Development of reinforcement learning algorithms with formally provable safety guarantees

Army Relevance and Value Proposition

Guarantee safety of actions derived from RL algorithms deployed to optimally control IoBT resources with limited training for the dynamic environments

Prior State-of-the-Art

- Reinforcement learning (RL) algorithms have seen recent successes in low-risk applications like video game playing.
- Standard RL algorithms are not safe for high-risk uses

Technical Approach

- Develop improved high-confidence bounds on the safety of newly proposed policies before they are used
- Importance sample to estimate the CDF with appropriate confidence interval
- Derive relevant statistics such as mean and conditional value at risk (CVaR) and corresponding confidence intervals
- Choose optimal policy that meets safety guarantees

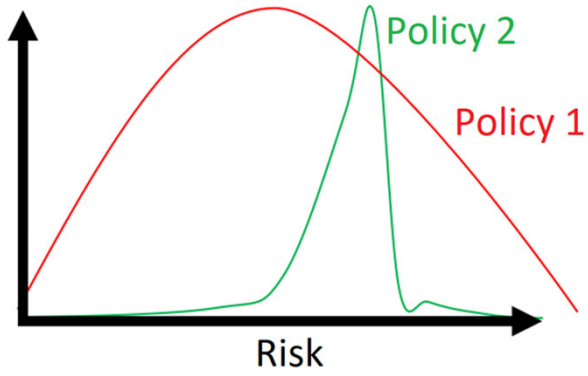

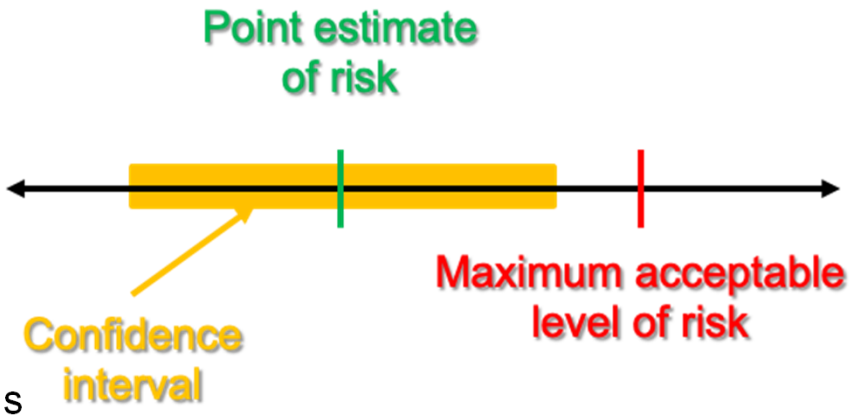

Contribution: Assured guarantees for aggressive decision making under uncertainty

UNIVERSAL OFF-POLICY EVALUATION (UnO)

Army Relevance: Provides performance guarantees for deployment to dynamic environments

Contribution: The first method for estimating and bounding *any* parameter of the performance distribution of a proposed new policy, using data collected with a different current policy. This enables verification of the safety of a potentially dangerous new policy before it is used

- Past methods use *importance sampling* to estimate the expected performance of the new policy
- The new method uses importance sampling to estimate and bound the entire cumulative distribution function (CDF) of performance of the new policy
- This results in confidence intervals for all parameters of the return distribution that hold simultaneously (i.e., without requiring further applications of the union bound)

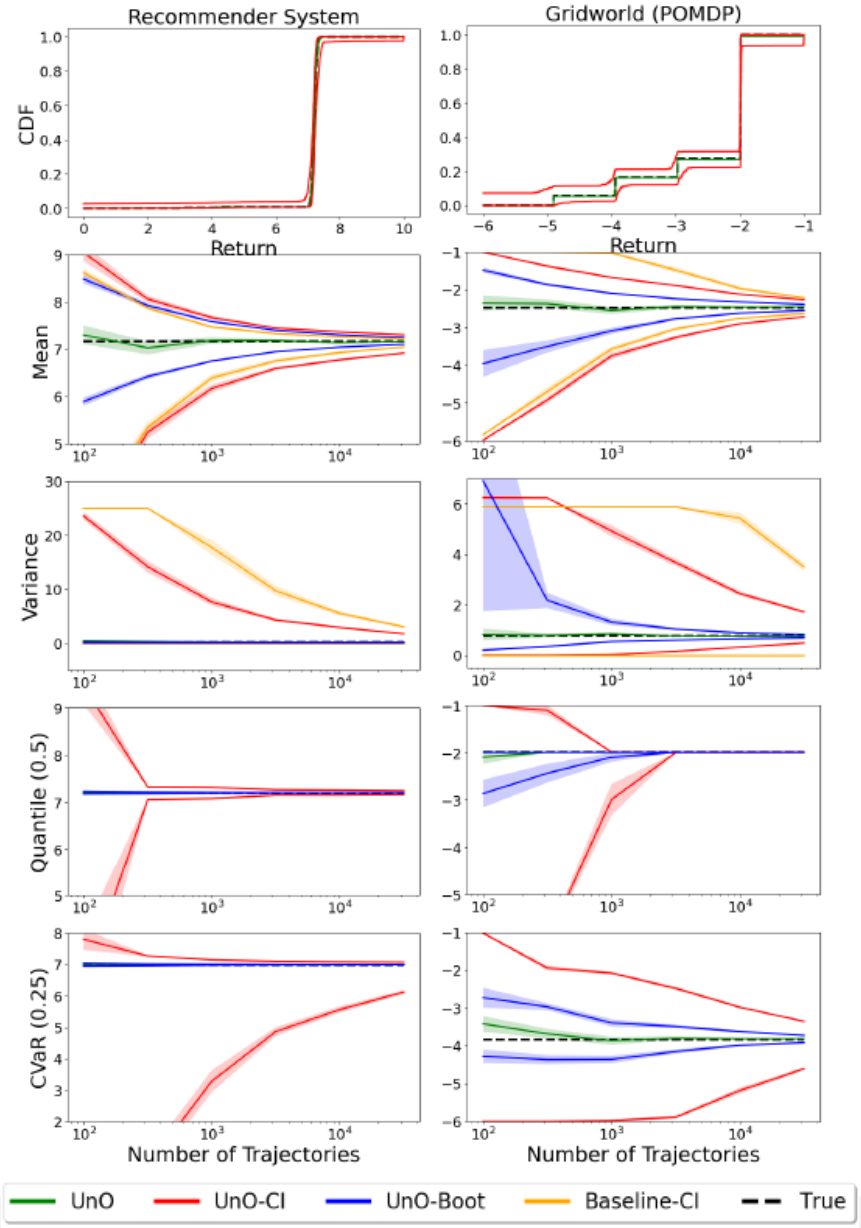

Key Publications:

Universal Off-Policy Evaluation

Y. Chandak, S. Niekum, B. Castro da Silva, E. Learned-Miller, E. Brunskill, P. S. Thomas.. *NeurIPS*, 2021.

UNIVERSAL OFF-POLICY EVALUATION (UnO)

Experiments and Results:

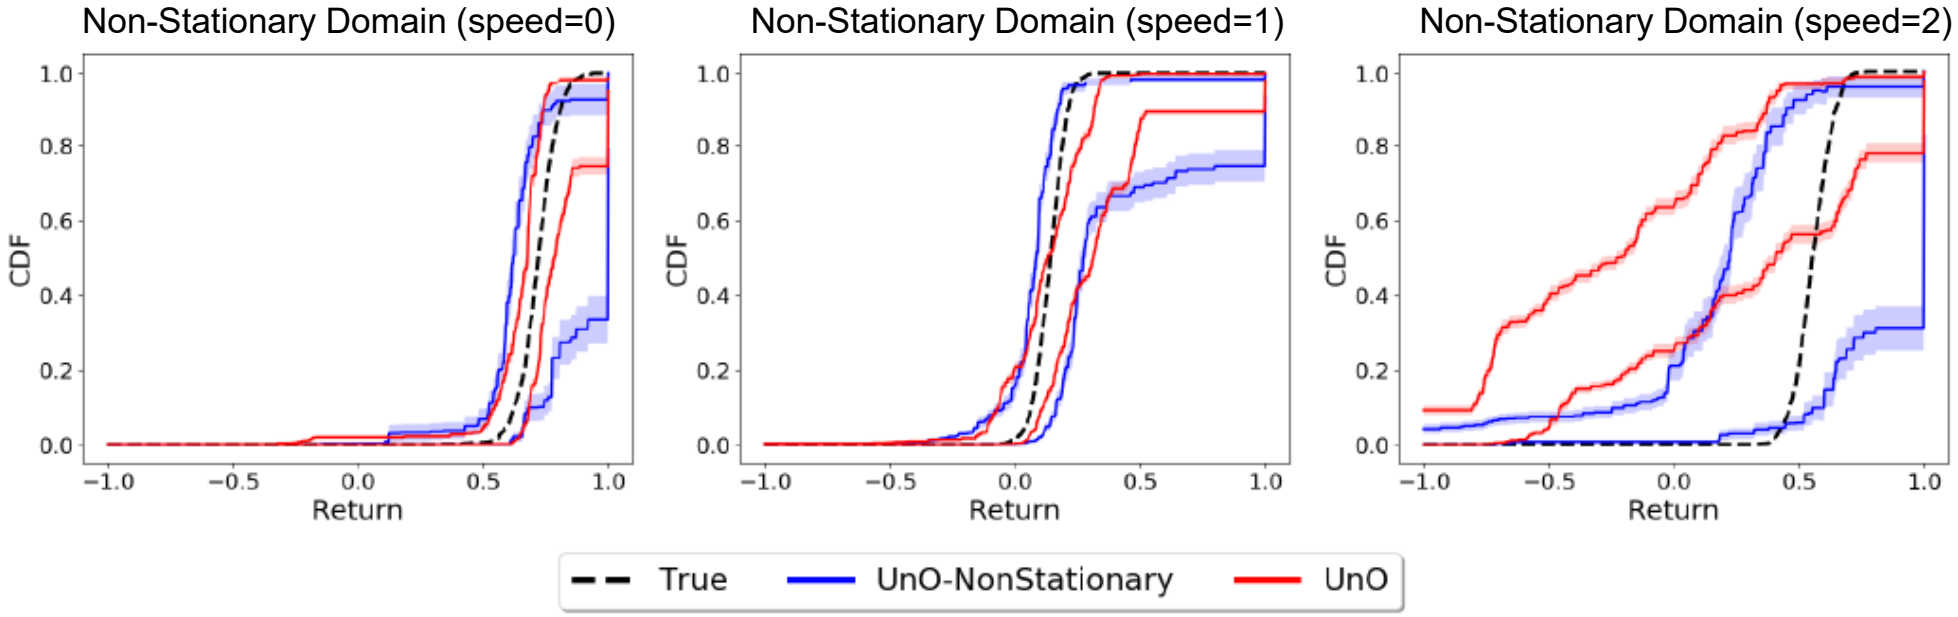

- Empirical results show confidence intervals for the variance of performance is state of the art
- Our method is the only method that provides confidence intervals for other parameters like conditional value at risk (CVaR)
- UnO can be extended to account for smooth nonstationarity of the MDP transition and reward functions

Summary:

Providing confidence intervals for the CDF and risk relevant statistics in non-stationary environments with smooth transitions

IOBT CRA IMPACT

Before IoBT CRA	With IoBT CRA
Neural network accuracy could not be predicted even for networks trained to be robust to attacks	Highest performing prediction of neural network output accuracy when fed adversarial or benign inputs
Sensor attack detection for target tracking too slow for real-time implementation	10x speed up in sensor attack detection with implementation on edge devices
Sensor attack detection not always possible in a decentralized manner	New decentralized algorithm works for all network topologies and target linear dynamics
Off-policy confidence intervals for *mean values* only	Confidence intervals for risk relevant statistics for off-policy evaluation

NEXT STEPS

IoBT Cornerstone Challenges: Robustness - Dynamics

- What is the best way to incorporate extant knowledge and supplement it with machine learning to enable adaptable, effective and safe processing, communication, and actuation of IoBT resources in **evolving systems**?
- How best to embrace uncertainty throughout online reinforcement learning that must adapt to latent changes in the environmental context?
- How do multiple agents optimize resources in **dynamic environments** changing at rates comparable to the data exchange rates?
- How to offer **resiliency**, **safety**, and sound **risk analysis** in **persistently transient systems**?
- Can theoretical bounds be discovered in **multi-objective optimization** contexts that characterize the plausibility of meeting requirements of decision latency, prediction accuracy, resilience to adversarial manipulation and robustness to exogenous environmental conditions across multiple domains?

Prospective Army Capabilities

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 overrides DDIL environment and offers protection through dispersion
- Information validation via multi-modal and multi-vantage sensing – leads to multiple corroborations of and stronger confidence in IoBT inferences

Questions?

Notable Awards/Honors:

- **10 year Most-influential Paper Award:** 42nd International Conference on Software Engineering (S. Jha)
- **Guest Editor:** Special Issue on AI of Transactions in Cyber-Physical Systems (S. Jha)
- **Guest Editor:** Special issue of Acta Informatica on Synthesis (P. Tabuada)
- **IFAC Fellow** (P. Tabuada)
- **Google Faculty Research** Award in Machine Learning (S. Diggavi)
- **2021 Wilbur Cross Medal, Yale University** (T. Başar)
- **Honorary Doctorate,** KTH Royal Institute of Technology, Stockholm, Sweden, 2019 (T. Başar)
- **Chair of HSCC Steering Committee** (P. Tabuada)
- **Safe AI work published in Science** (P. Thomas)

Notable Transitions:

- Some of the results developed in the IoBT CRA bootstrapped winning proposals from DARPA and IARPA

Demos & Posters: Visit https://abdelzaher.cs.illinois.edu/RMB22-Demos.html

- Trustworthy ML with Surprise Detection
- Risk Aware (Task) Placement
- Resilient Localization