Imperial College London

Machine Learning for Defence Signal Processing and Communications

SSPD 2023

Athanasios (Thanos) Gkelias and Kin K. Leung

Imperial College

September 13, 2023

EPSRC Pioneering research and skills

[dstl]

Talk Outline

- Generative Adversarial Networks (GAN) for Limited EM Signals
- Federated Learning (FL) with Resource Constraints
- Use Transfer Learning to Adapt to New Operating Environments

GAN-Based Detection of Adversarial EM Signal Waveforms

- **Conventional techniques**
  - RF fingerprinting: channel-fingerprint, device-fingerprint
  - Transient-based, steady state-based
- **Challenge**: unauthorized transmissions often originate from unidentified devices with unknown EM fingerprints.
  - **No samples** - they appear for the first time
  - **Samples of insignificant size** to be efficiently modelled

Identification through traditional supervised learning or signal processing techniques is extremely difficult

A. Gkelias and K. K. Leung, "GAN-Based Detection of Adversarial EM Signal Waveforms," MILCOM 2022

Anomaly Detection

Use only available data to learn how to detect irregularities or unobserved patterns/features in new sets of data

- Technique known as "**anomaly**" or "**novelty**" **detection**
  - "**normal**": already known waveforms
  - "**anomaly**": previous unseen waveform, with features different to the former ones
- First, train the system on "**normal**" – consider as friendly waveforms
- Then, identify unknown waveforms as "**anomalies**" – potentially adversarial

Generative Adversarial Networks (GANs)

- GANs **only exploits the discriminative benefit** of the network
  - i.e., minimize distance between real and generated sample distributions
- Susceptible to **training instabilities** and **mode collapse** – **difficult to train**

$$\min_G \max_D V(D,G) = \mathbb{E}_{\boldsymbol{x}\sim p_{data}(\boldsymbol{x})}[\log D(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{z}\sim p_{\boldsymbol{z}}(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z})))].$$

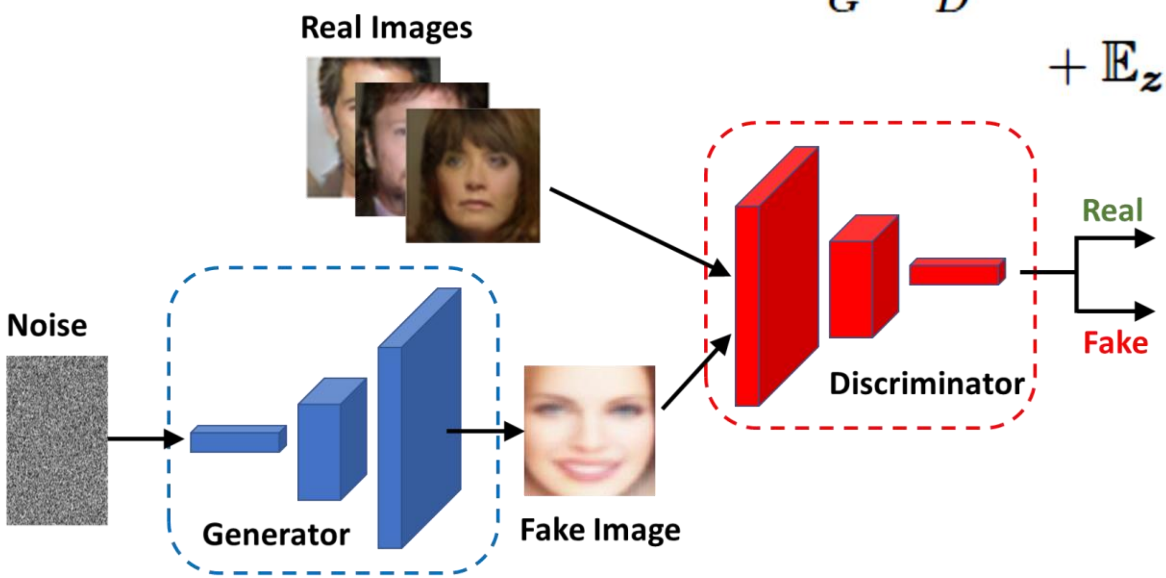

Imperial College London

Dual AE enhanced GAN Architecture

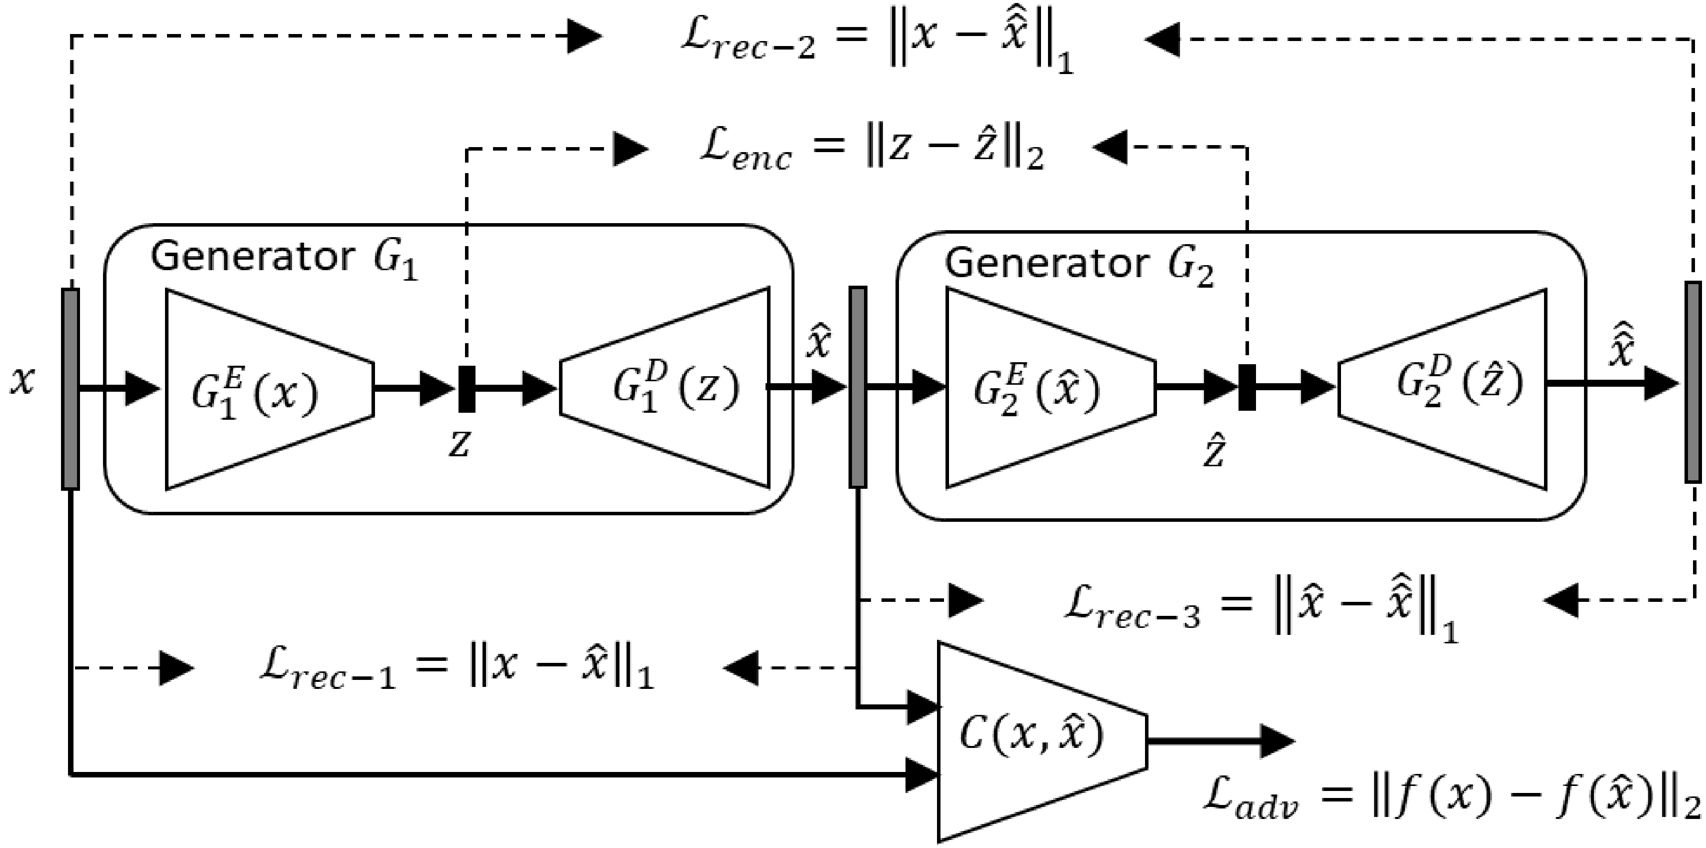

Imperial College London

Dual AE enhanced GAN Architecture

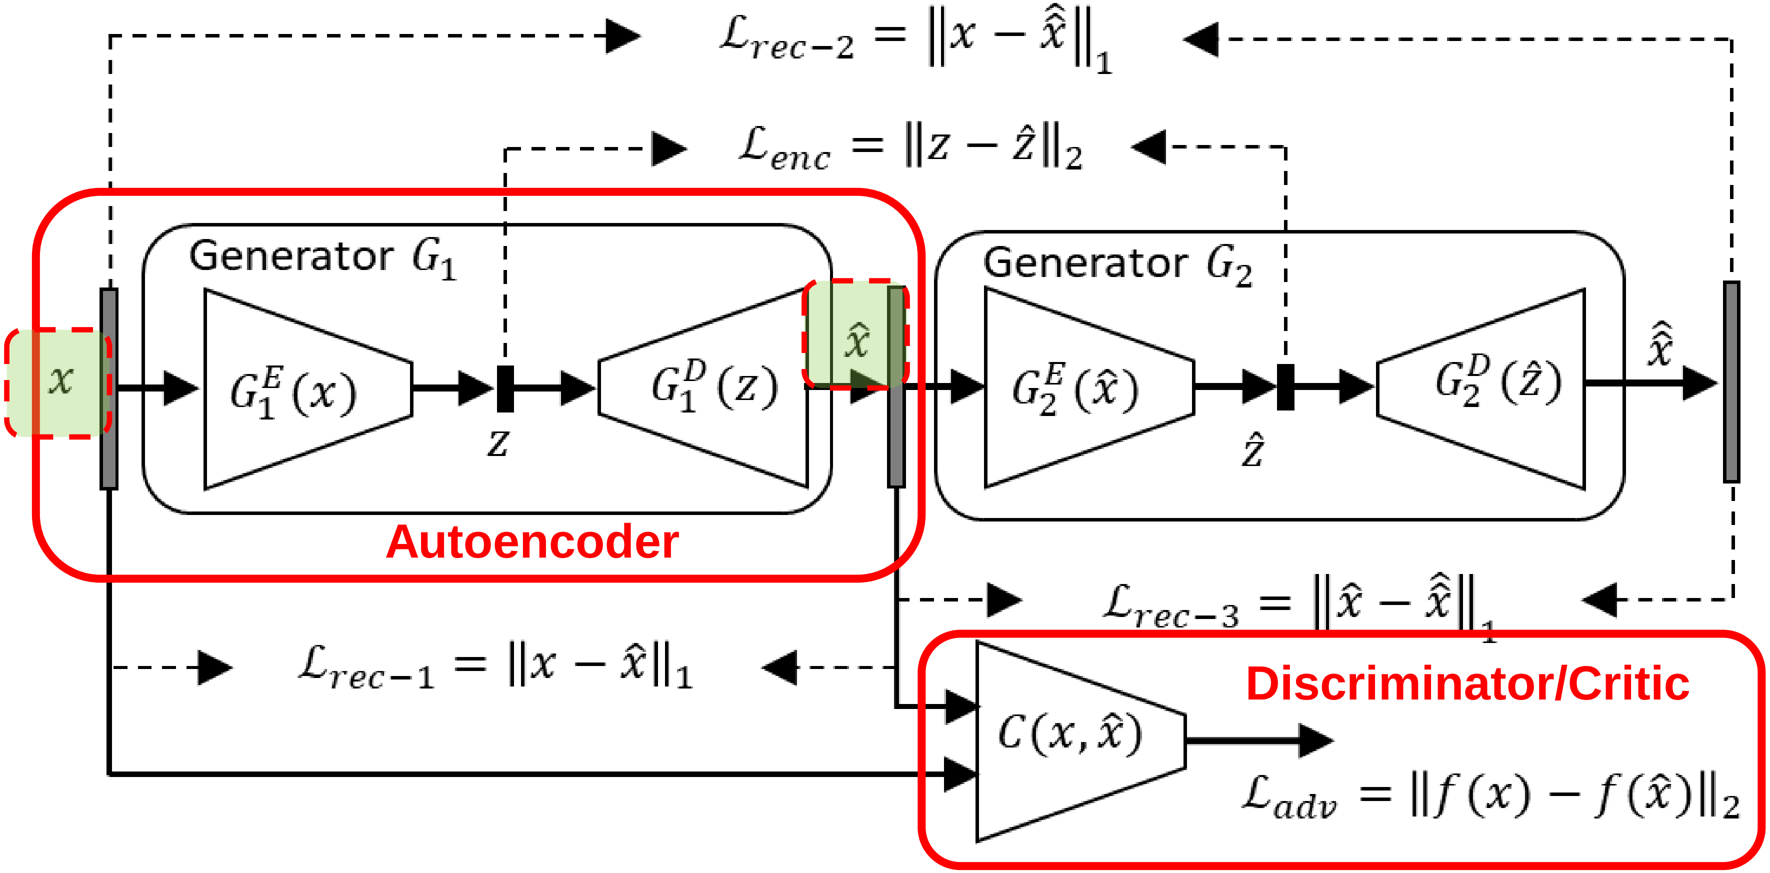

Adversarial Loss: $\mathcal{L}_{adv} = -\mathbb{E}_{\boldsymbol{x} \sim p_x}[C(G_1(\boldsymbol{x}))]$

Dual AE enhanced GAN Architecture

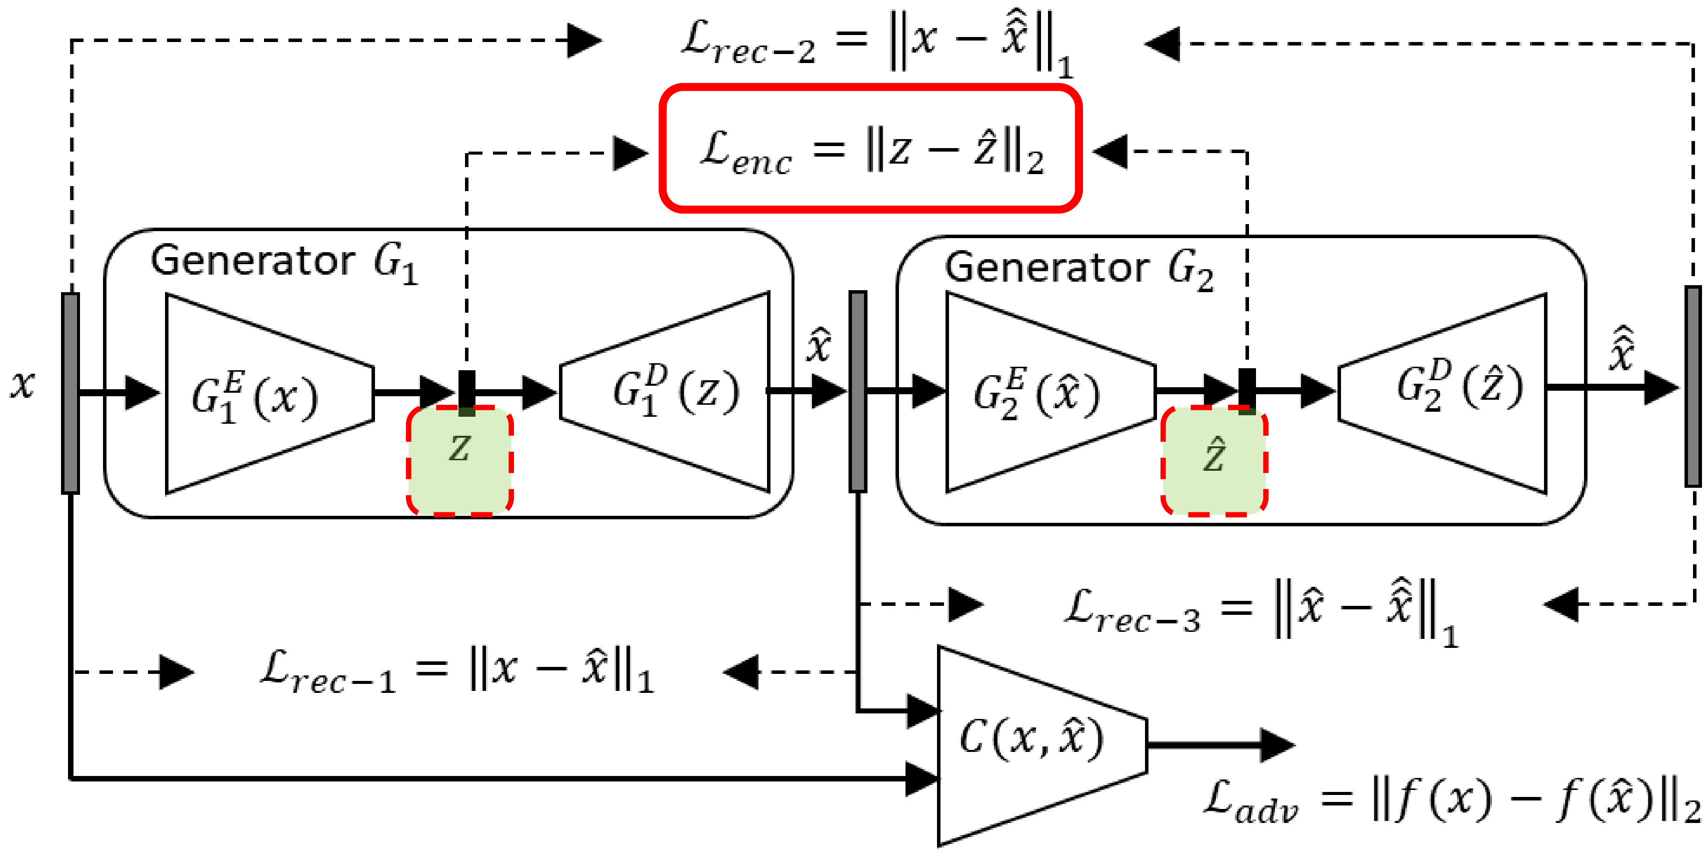

Encoder Features Loss: $\mathcal{L}_{enc} = \mathbb{E}_{\boldsymbol{x}\sim p_{\boldsymbol{x}}} \parallel G_1^E(\boldsymbol{x}) - G_2^E(G_1(\boldsymbol{x})) \parallel_2$

Imperial College London

Dual AE enhanced GAN Architecture

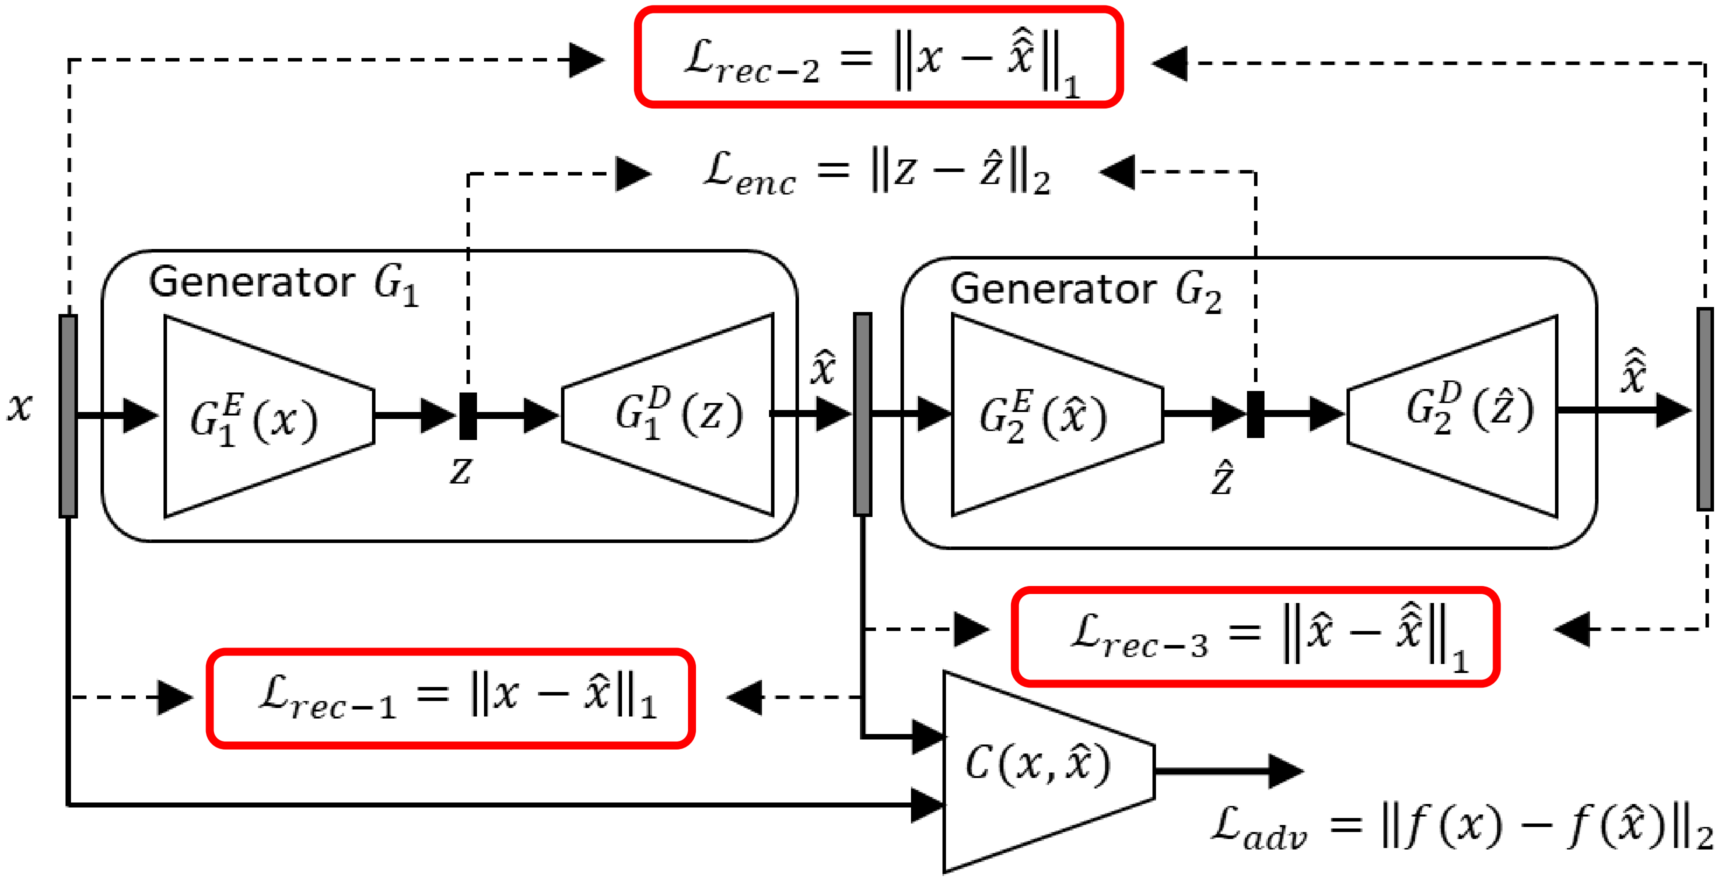

Reconstruction Loss : $\mathcal{L}_{rec} = v_1\mathcal{L}_{rec-1} + v_2\mathcal{L}_{rec-2} + v_3\mathcal{L}_{rec-3}$

Anomaly Score

- During training, minimize:

$$\mathcal{L}_{gen} = w_{adv}\mathcal{L}_{adv} + w_{rec}\mathcal{L}_{rec} + w_{enc}\mathcal{L}_{enc}$$

- Anomaly Score

$$A_{score}(\dot{\boldsymbol{x}}) = \| G_1^E(\dot{\boldsymbol{x}}) - G_2^E(G_1(\dot{\boldsymbol{x}})) \|_1$$

- Anomaly Score as probability

$$\mathcal{S} = \{s_i : A_{score}(\dot{x}_i)\} \qquad s_i' = \frac{s_i - \min(\mathcal{S})}{\max(\mathcal{S}) - \min(\mathcal{S})}$$

Wasserstein GAN-GP

- Training instabilities (Discriminator is optimised faster than the Generator and mode collapse)
- Use **Wasserstein GAN-GP** to overcome these issues:
  - **Wasserstein-1 distance** as discriminator's objective function
  - Output is now a **scalar score** rather than a probability
  - "**gradient-penalty**" for weight regularization to enforce 1-Lipschitz continuity

$$\mathcal{L}_{critic} = \mathbb{E}_{\boldsymbol{x}\sim p_x}[C(G_1(\boldsymbol{x}))] - \mathbb{E}_{\boldsymbol{x}\sim p_x}[C(\boldsymbol{x})] + \lambda_{gp}\mathbb{E}_{\tilde{\boldsymbol{x}}\sim p_{\tilde{x}}}[(\| \nabla_{\tilde{\boldsymbol{x}}} C(\tilde{\boldsymbol{x}}) \|_2 - 1)^2]$$

Experimental Set-up and EM Dataset

- Synthetic (MATLAB simulated) dataset
  - **Radar**: Rectangular, LFM, Barker Code
  - **Communications**: BPSK, QPSK, PAM4, GFSK, CPFSK
  - **Channel-impairments**: AWGN, Rician multipath fading, Doppler shift

- 3 different evaluation scenarios
  - Only Radar waveforms
  - Radar and Communications waveforms
  - Only LFM Radar waveforms

Imperial College London

Performance Results

TABLE I
AUROC VALUES FOR RADAR ONLY EM WAVEFORMS

SNR	Pulsed Radar Waveforms		
	LFM	**Rectangular**	**Barker**
-18dB	0.981	0.832	0.999
-12dB	0.986	0.911	0.999
-6dB	0.998	0.915	0.999
0dB	0.999	0.968	0.999
6dB	0.999	0.976	0.999

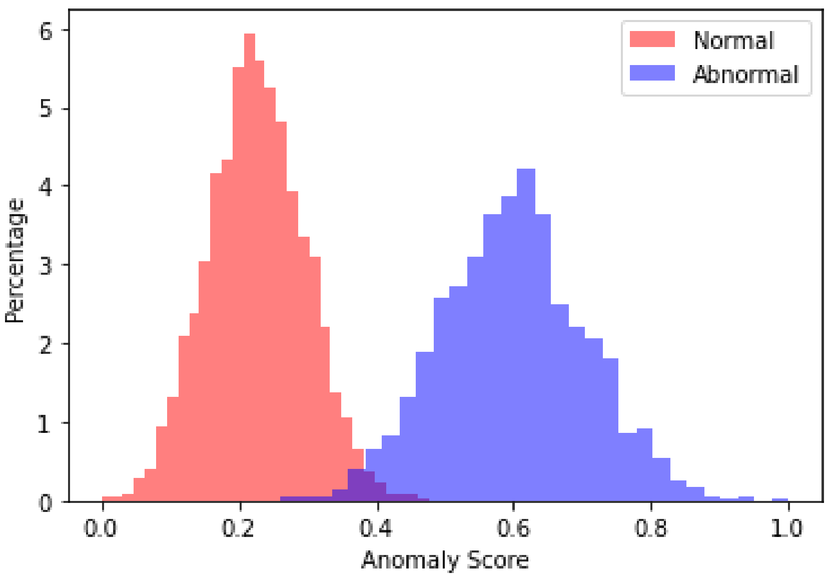

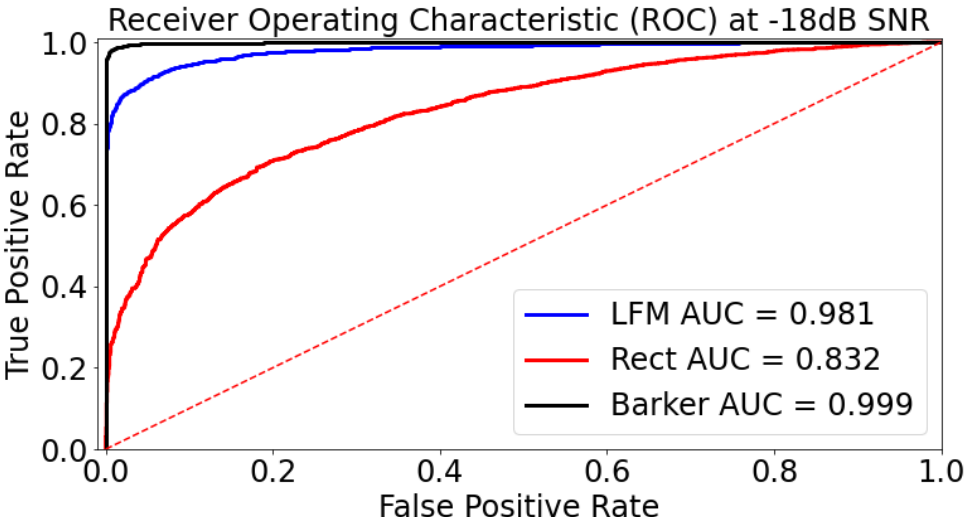

Challenge for Federated Learning (FL)

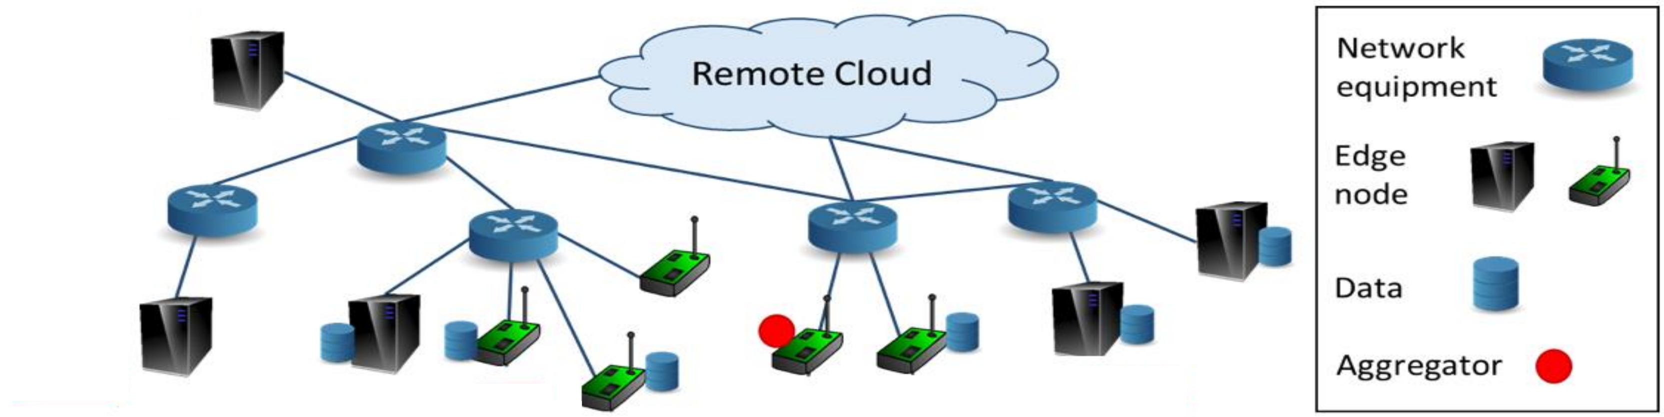

Often need to train machine learning (ML) models using data collected at different locations

Not to share data from multiple locations for data privacy or lack of communication bandwidth

➡ Federated Learning (FL) is a possible solution

Edge computing or other systems often have limited resources (e.g., bandwidth, processing power, response time)

➡ We propose an approach to optimizing FL subject to resource constraints

Imperial College London

Federated Learning: Distributed Gradient Descent

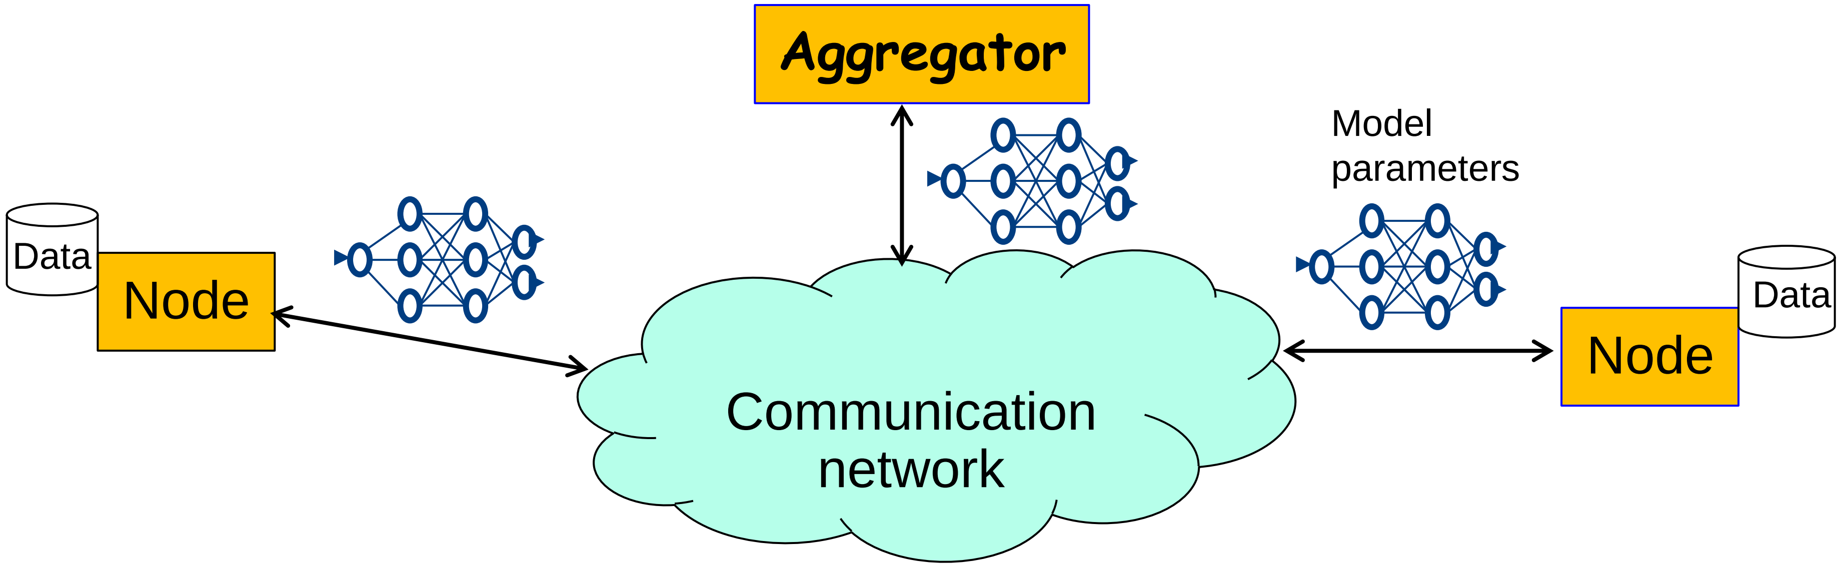

Question: How many local updates between two global aggregations subject to available resources?

- Distributed and centralized gradient descent are NOT equivalent: Divergence of local gradient by local updates depends on data distribution at nodes
- Infrequent aggregation saves communication cost, but affects learning

➡ Minimize learning error by finding the optimal number of local updates between two global aggregations given available resources

Imperial College London

Approximate Solution for Optimal Training

Derive and use upper bound as an approximation to the loss function

Original problem:

$$\min_{\tau,T} \quad F(\mathbf{w}(T))$$
$$\text{s.t.} \quad T\left(c+\frac{b}{\tau}\right) \le R$$

Approximation:

$$\min_{\tau,T} \quad \frac{1}{T\left(\omega\eta\left(1-\frac{\beta\eta}{2}\right)-\frac{\rho h(\tau)}{\tau\varepsilon^2}\right)}$$
$$\text{s.t.} \quad T \le \frac{R\tau}{c\tau+b}$$

$$\boxed{\tau^* = \arg\max_{\tau} G(\tau) \qquad T^* = \frac{R\tau^*}{c\tau^*+b}}$$

where $G(\tau) \triangleq \frac{\tau}{\tau+a}\left(\eta\left(1-\frac{\beta\eta}{2}\right)-\frac{\rho h(\tau)}{\tau\varepsilon^2\omega}\right)$

$a = b/c$

τ: number of local updates between two aggregations

- $G(\tau)$ has a unique maximum (strictly concave)
- τ^* is found using binary search

S. Wang, T. Tuor, T. Salonidis, K.K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive Federated Learning in Resource Constrained Edge Computing Systems," IEEE JSAC, 2019

Imperial College London

Experiment Results: Loss Functions

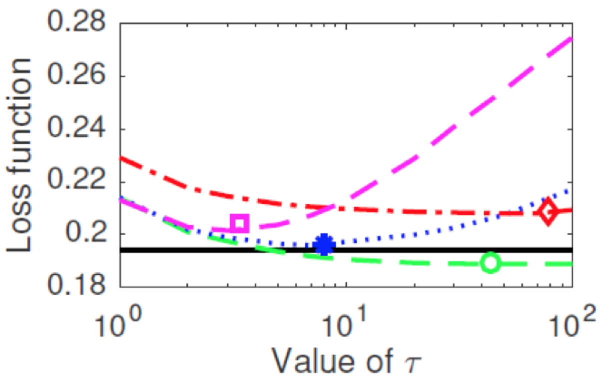

(a) SVM (DGD)

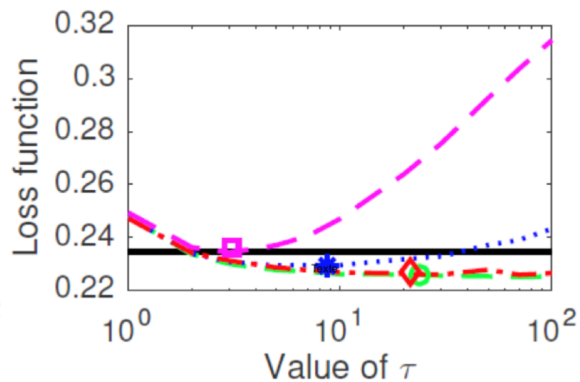

(b) SVM (SGD)

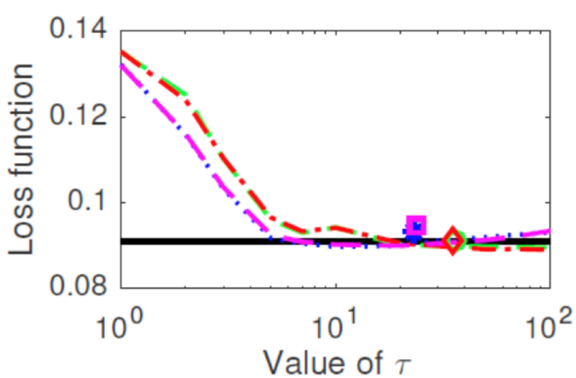

(c) K-means (SGD)

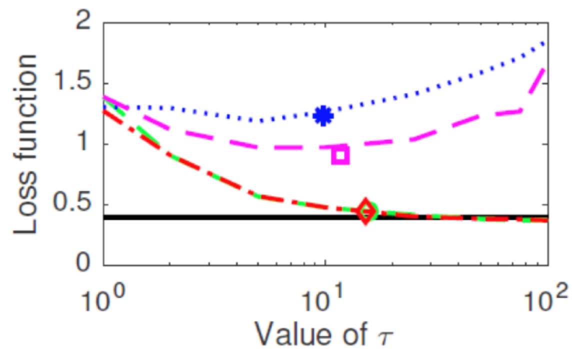

(d) CNN (SGD)

DGD: Deterministic gradient descent
SGD: Stochastic gradient descent

- Proposed approach (symbols in the left plots) performs close to the optimum for all cases and models
- Optimal value of is different for different cases and models
- In some cases, distributed approach can perform better than centralized approach for fixed available resources

Imperial College London

Model Pruning for Federated Learning (FL)

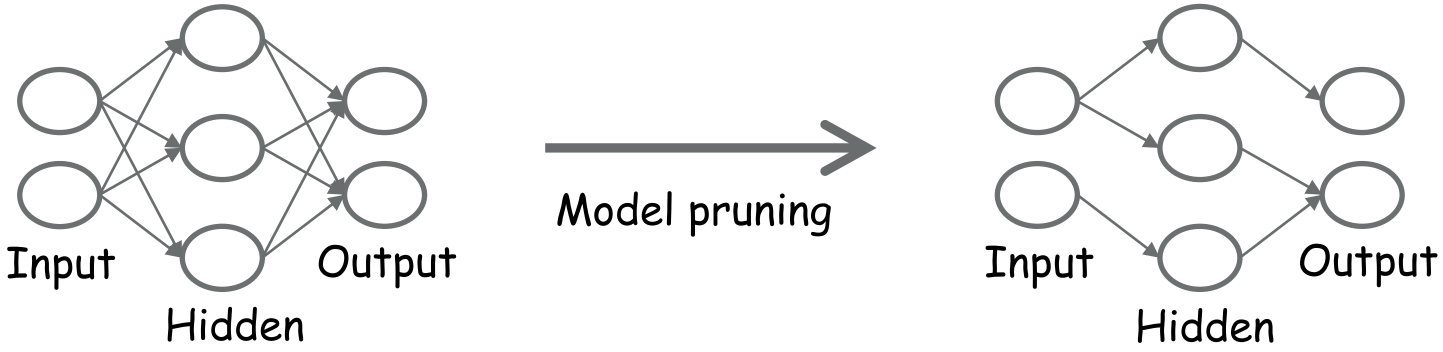

Key ideas of model pruning

- Removing unimportant parameters does not degrade performance
- Fewer model parameters reduce computation and communication

Adaptive selection of a subset of model parameters

- Among all subsets, select the subset of parameters to maxmize the ratio of decrease in loss function by pruned parameters to time needed to process the parameter subset
- Optimal greedy algorithm for identifying pruned parameters
- Established convergence bound on the loss function by the pruned parameters

Training Acceleration for Pruned FL

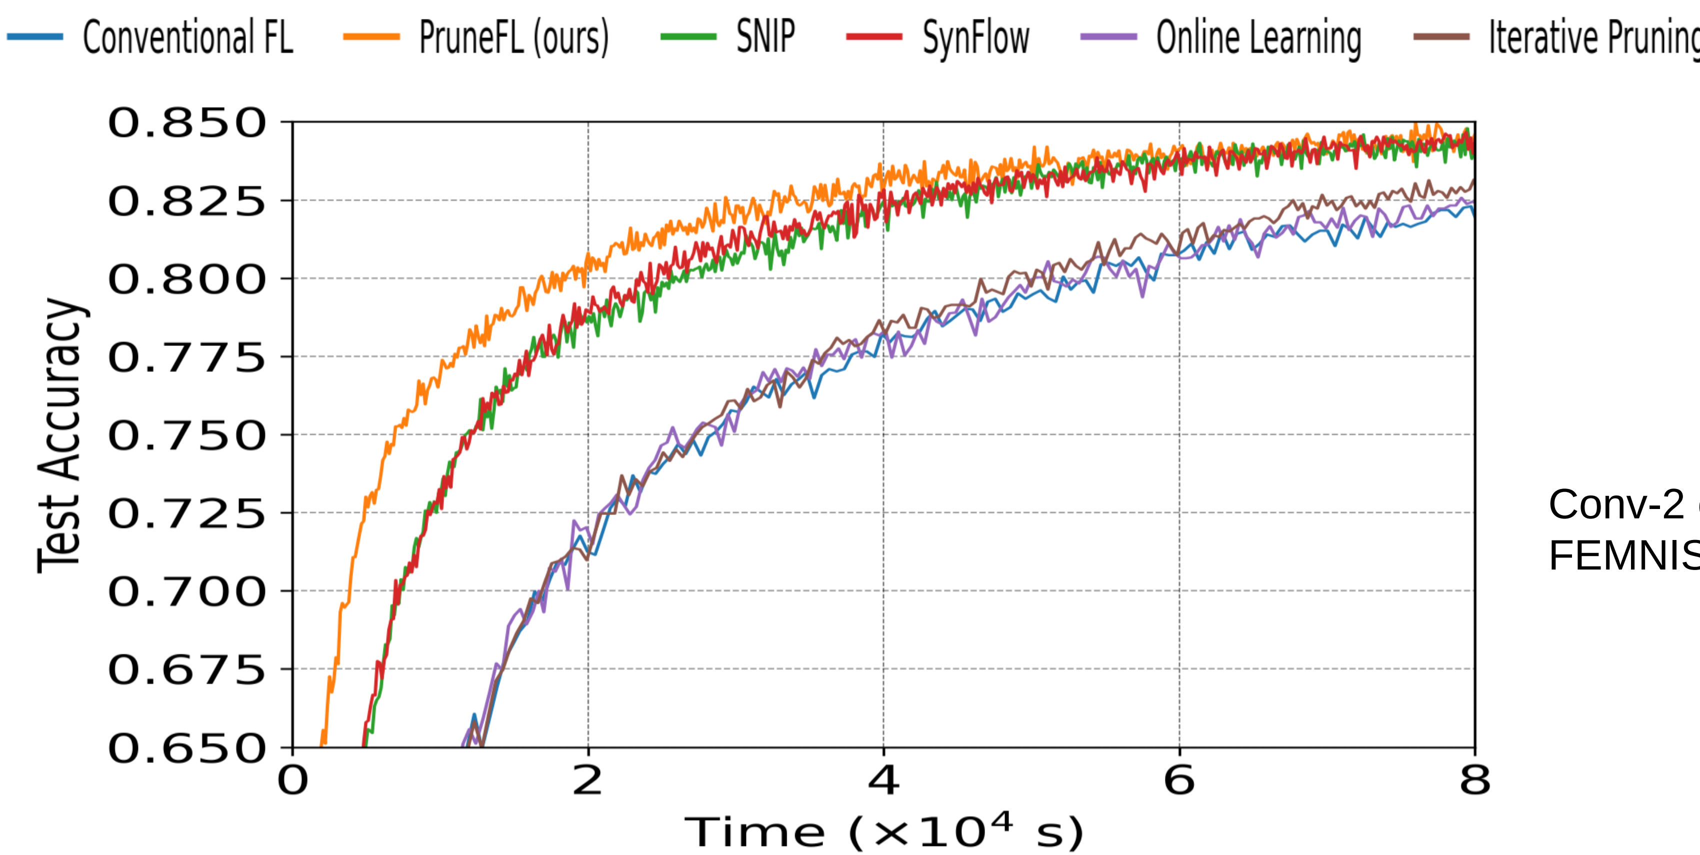

Conv-2 on FEMNIST

Observation: PruneFL accelerates training on various datasets

Y. Jiang, S. Wang, et. al., "Model Pruning Enables Efficient Federated Learning on Edge Devices," *IEEE Trans. on Neural Networks and Learning Systems*, early access, April 2022.

Change of EM Environments

Urban environment

Suburban environment

- ML Challenge: Change of EM Environments
  - Learned model is no longer valid in new environments
  - Learn from beginning in a new environment?
  - Can we re-use knowledge learned from one environment to a new one?
- Possible Technique to Adapt Learning in New Environments
  - Transfer learning (TL)

Imperial College London

RL + Transfer Learning for Environment Changes

- Issue: How can RL adapt to changes of operating environment?
- Joint Reinforcement and Transfer Learning (RL+TL)
  - Consider SDN fragmentation with 2 domains, focusing on data servers
  - Combine RL (e.g., sasRL) and TL based on generative adversary network (GAN) to synthesize data for learning in new environments
  - Combined RL+TL can significantly speed up RL when operating environment changes (e.g., SDN domain fragmentation and re-connection)

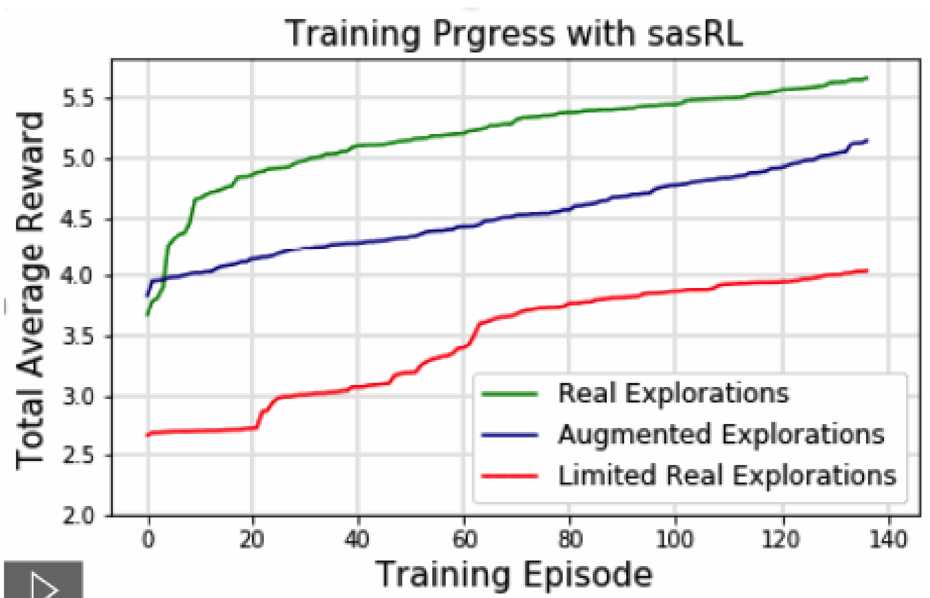

Results:
- The reward is inversely proportional to the service delay
- Real Explorations = 10,000 data samples
- Augmented (RL+TL) or Limited Explorations = 100 data samples (1% of Real Exploration sample size)

Summary on ML for Signal Processing & Communications

- Generative Adversarial Networks (GAN) for Limited EM Signals
  - Develop the GAN framework to generate training data for classifying EM signals (e.g., hostile signals)
  - Validated the proposed framework by simulated EM environments
- Federated Learning (FL) with Resource Constraints
  - Formulated and derived upper bound for the loss function to estimate optimal FL parameters using limited resources
  - Developed a technique to prune FL models
  - Validated new approaches by various datasets
- Transfer Learning (TL) to Adapt to New Operating Environments
  - Use reinforcement learning to illustrate transfer learning can speed up training following changes of operating environment
  - Experimentation indicates very significant speed up from TL

- ## Acknowledgment of Research Funding

  EPSRC/Dstl UDRC3 EM Theme Project, US/UK ITA Project

- ## Publications

A. Gkelias and K.K. Leung, “Generative Adversarial Networks for Waveform Classification in Congested EM Environments,” IEEE MILCOM 2022.

Y. Jiang, S. Wang, V. Valls, B.-J. Ko, W.-H. Lee, K.K. Leung, and L. Tassiulas, “Model Pruning Enables Efficient Federated Learning on Edge Devices,” *IEEE Transactions on Neural Networks and Learning Systems*, early access, April 2022.

Z. Zhang, A. Mudgerikar, A. Singla, K.K. Leung, E. Bertino, D. Verma, K. Chan, J. Melrose, and J. Tucker, “Reinforcement and Transfer Learning for Distributed Analytics in Fragmented Software Defined Coalitions,” SPIE, April 2021, Florida.

S. Wang, T. Tuor, T. Salonidis, K.K. Leung, C. Makaya, T. He, and K. Chan, “Adaptive Federated Learning in Resource Constrained Edge Computing Systems,” IEEE Journal on Selected Areas in Communications, 2019

S. Wang, T. Tuor, T. Salonidis, K.K. Leung, C. Makaya, T. He, and K. Chan, When Edge Meets Learning: Adaptive Control for Resource-Constrained Distributed Machine Learning, *IEEE INFOCOM*, Apr. 2019

T. Tuor, S.Wang, T. Salonidis, B.J. Ko, K.K. Leung, “Distributed Machine Learning at Resource-Limited Edge Node: Demo”, *IEEE INFOCOM* (Demo Session), Apr. 2018

THANK YOU

Q&A?