ITU Webinar

Distributed* Machine Learning and Wireless Networks: A Closer Union

Walid Saad

Electrical and Computer Engineering Department,
Network sciEnce, Wireless, and Security (NEWS) Group

Virginia Tech

Email: walids@vt.edu

Group: http://www.netsciwis.com

Personal: http://resume.walid-saad.com

*** A small part of the talk may violate the distributed angle**

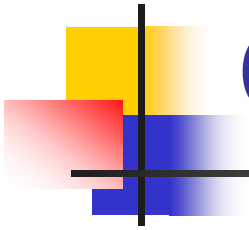

Outline

- Brief introduction
- Towards generalizable machine learning in wireless
  - Experienced deep reinforcement learning
  - Distributed generative adversarial networks
  - Multi-agent meta-reinforcement learning
- Edge learning as a use case of wireless systems
  - Overview and key results
- Conclusions

AI: Today vs Tomorrow

- **Artificial intelligence** is enabled by **machine learning** which enables "machines" to use data and connectivity for intelligent and autonomous decision making

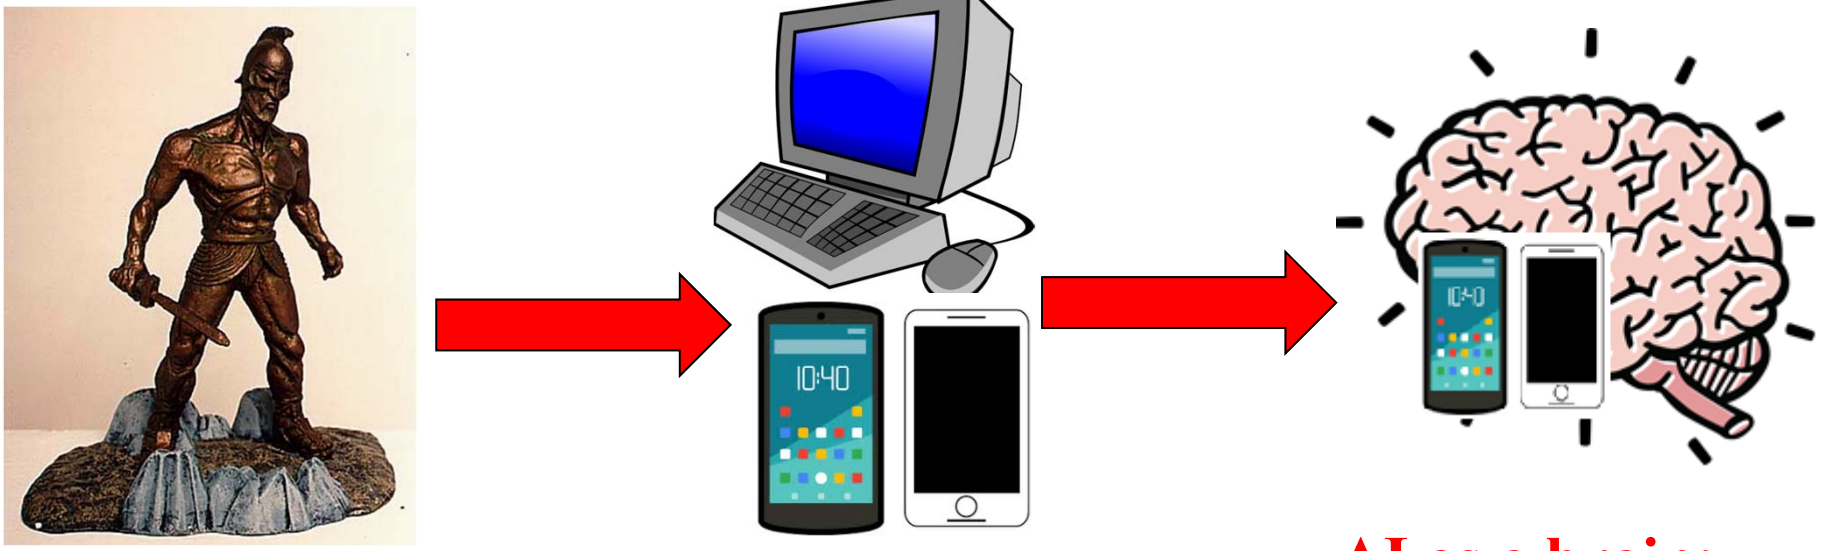

AI as fiction:
Talos – the "bronze" man of Greek myths

AI to compute:
Computers – can do arithmetic, math, etc

AI as a brain:
Intelligence – can mimic human brain

AI: Today vs. Tomorrow

Today's AI	Tomorrow's AI
• Mostly **centralized, single agent** at a controller	• **Distributed, multi-agent** systems at the edge
• Relies on **big data**	• Learns with **small data**
• Model-based **or** model-free	• **Explainable**
• Training-**dependent**	• Nearly **training-free**
• Learns **specific** "tasks", unreliable	• **Reliable**, learns "skills" and **generalizes** to unseen tasks and environments

- AI systems will rely on wireless systems in two ways
  - Learning (at the edge) to communicate?
  - Communication for learning (joint design)?

New Concept: Experienced Deep Reinforcement Learning for Reliable Communications

A. T. Kasgari, W. Saad, M. Mozaffari, and H. V. Poor, "Experienced Deep Reinforcement Learning with Generative Adversarial Networks (GANs) for Model-Free Ultra Reliable Low Latency Communications", *IEEE Transactions on Communications*, vol. 69, no. 2, pp. 884 - 899, February 2021.

Reliable AI for Reliable Communications

- Ultra-reliable low latency communications (URLLC) will be a staple of 6G, and will have to go "extreme"
- URLLC has been around for a while but prior art…
  - Focused on *IoT sensors* (uplink) – **autonomous vehicles/drones are different (downlink? large packets?)!**
  - Assumes *known models* for traffic (M/M/1 etc.)– **latency has many components, hard to model!**
  - Considers slow deep reinforcement learning (DRL) – **learning in URLLC must handle extreme, rare conditions!**
- Fundamental question: Can we design **reliable AI** that can work well under extreme conditions to achieve URLLC or, more broadly, reliable links?

System Model

- Consider the downlink of a single-cell wireless network whose base station is sending latency-sensitive control message to autonomous vehicles
- We consider a downlink OFDMA system with resource blocks that must be allocated
- The downlink rate from the BS to a user *i* will be

$$r_i(t) = \sum_{j=1}^{K} \rho_{ij}(t) B \log_2 \left(1 + \frac{p_{ij}(t) h_{ij}(t)}{\sigma^2}\right)$$

RB allocation indicator

Bandwidth

Power allocated over RB *j*

Channel gain

Problem Formulation

- Reliability is defined as the probability of end-to-end packet delay exceeding a threshold
- We can map this to the following constraint:

$$r_i(t) > \phi(\lambda_i(t), \beta(t), \gamma_i, D_i^{\max}) > \lambda_i(t)\beta_i(t)$$

Unknown relationship

Arrival rate

Packet size

- We do not make any assumptions for a delay model
  - Delay is intrinsically hard to model, most models are often unrealistic and have some hidden drawbacks
  - Delay has many components, hard to model their combination precisely

Problem Formulation

- Our goal is to solve the following problem

$$\min_{p_{ij},\rho_{ij}} \quad \lim_{t\to\infty} \frac{1}{t} \sum_{\tau=1}^{t} \sum_{i=1}^{N} \sum_{j=1}^{K} p_{ij}(\tau),$$

$$\text{s.t} \quad \Pr\{D_i > D_i^{\max}\} < 1 - \gamma_i^*, \quad \forall i \in \mathcal{N},$$

$$r_i(t) > \lambda_i(t)\beta_i(t), \quad \forall i \in \mathcal{N}, \quad \forall t$$

$$p_{ij}(t) \geq 0, \quad \rho_{ij}(t) \in \{0, 1\},$$

$$\forall i \in \mathcal{N}, \quad \forall j \in \mathcal{K}, \quad \forall t,$$

$$\sum_i \rho_{ij}(t) = 1, \quad \forall j \in \mathcal{K}, \quad \forall t.$$

Reliability constraint

Rate constraint

Feasibility constraints

- Explicit rate guarantees imposed
- Challenging to solve because of our model-free assumption

Handling Model-Free

- In reality, a wireless network can empirically measure the delay

$$\gamma_i(t) = 1 - \Pr\{D_i > D_i^{\max}\} \approx 1 - \frac{\mu_i'(t)}{\mu_i(t)}$$

Ratio of number of packets with delay excess and total number of packets

- Network can "learn" the delay once it connects with a user
- How to learn? Reinforcement learning is natural but…
  - ….classical solutions cannot handle the large state space
- Solution: Deep reinforcement learning
  - Deep RL used because it is appropriate to handle our large state space not because it is "fashionable"

Deep-RL for Model-Free URLLC

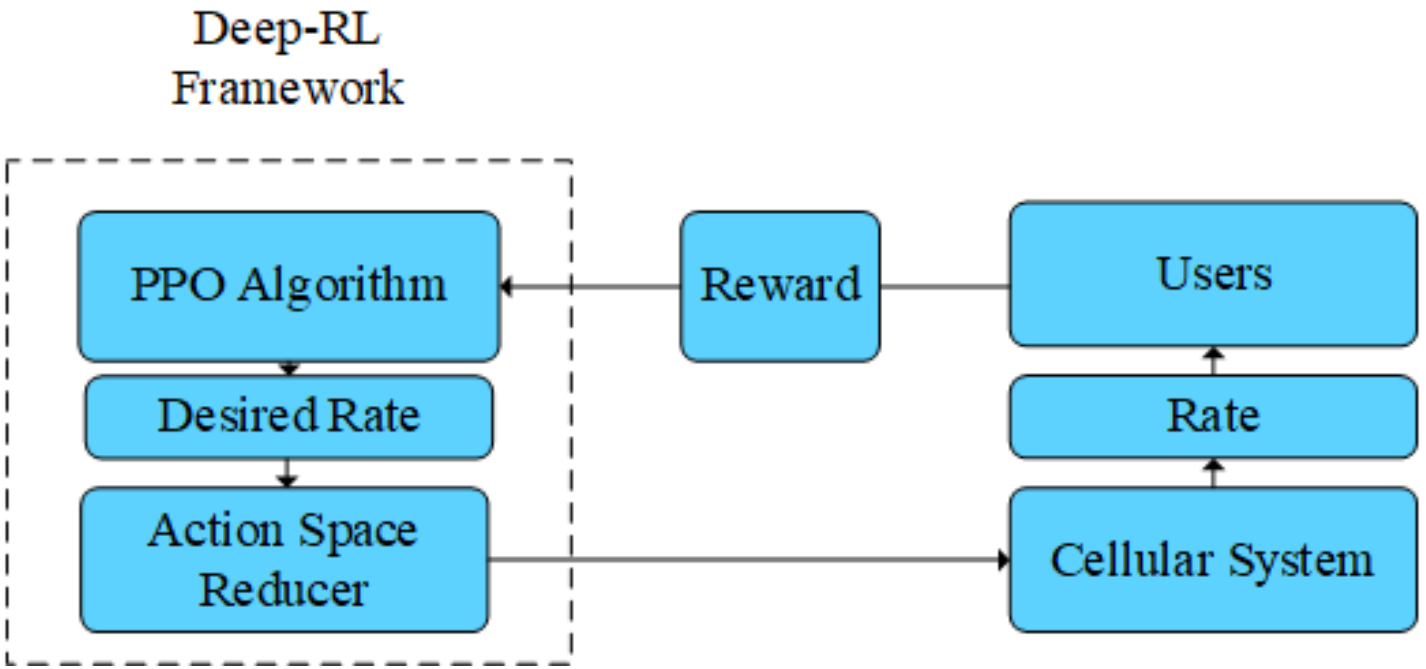

- **State space:** number of packets transmitted, packet size, and channel gains
- **PPO:** Proximal policy optimization determines target rates
- **Action space reducer:** Deep-RL made tractable

Deep-RL for Model-Free URLLC

- The reward function used by deep-RL:

$$R(\boldsymbol{a}_t, \boldsymbol{s}_t) = -\sum_{i \in \mathcal{N}} w_i(t)(1 - \gamma_i(\boldsymbol{a}_t, \boldsymbol{s}_t)) - \alpha P(\boldsymbol{a}_t)$$

$$w_i(t+1) = \max\{w_i(t) + \gamma_i^* - \gamma_i(t), 0\}$$

Time-varying weight that control the reliability

- **Theorem 1:** By maximizing this reward, after convergence of the deep-RL algorithm, the reliability of each user is guaranteed, such that:

$$\gamma_i(t) \geq \gamma^* \, \forall i \in \mathcal{N}$$

  - Implicitly ensures rate requirements as well

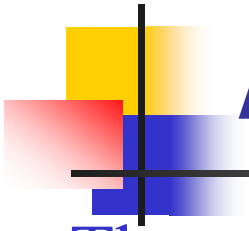

Action Space Reduction

- The action space for the deep-RL is too large
  - Non-wireless prior work: Small action space (e.g. Atari)
  - Wireless prior work on deep-RL does not handle the large action space, but maintains complexity
- Two-step solution
  - Use the PPO algorithm optimize rate, rather than RB/power
  - Map PPO outcomes to original actions (action space reducer)
  - Action space reducer: a re-formulated optimization problem
- But, is deep-RL reliable and suitable for URLLC?
  - No! Can be **slow** to converge and **unreliable** extreme cases
  - Solution? Use generative adversarial networks (GANs)!

What is GAN?

- A **generative** model seeks to create data that is not seen before, but fits some input data distribution

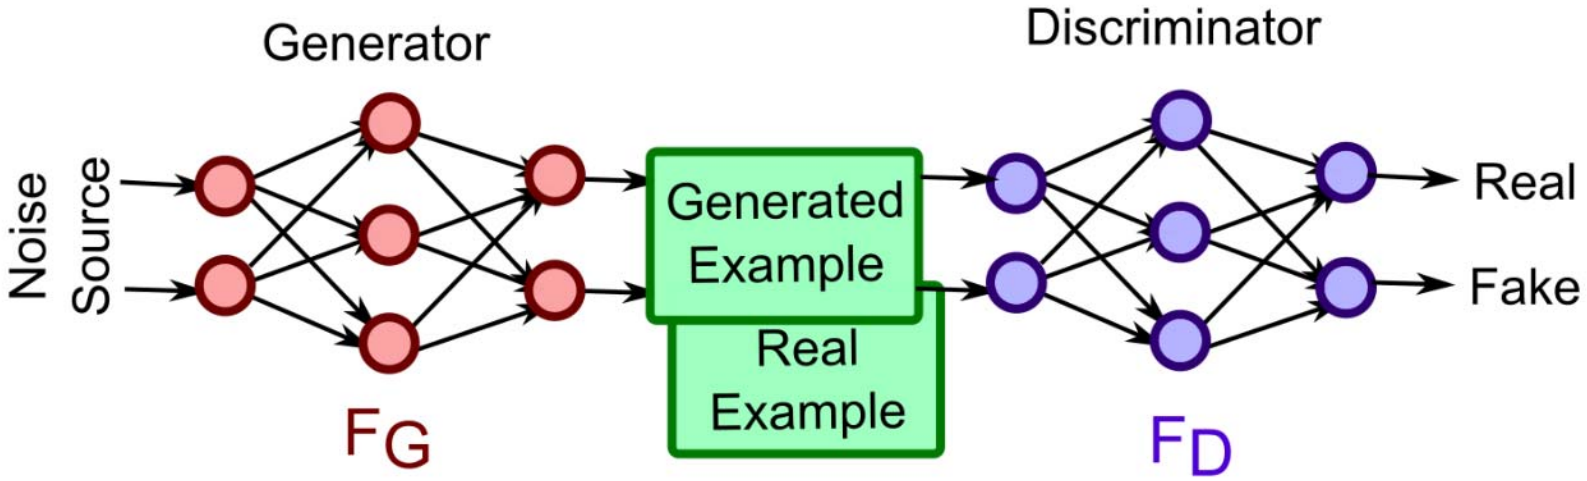

Figure source: http://hunterheidenreich.com/blog/what-is-a-gan/

- Generator: Tries to generate fake data
- Discriminator: Figure out whether data is fake or real
  - Adversarial interactions between the two (game theory)

Experienced Deep RL

- Use GAN to create a "virtual environment" for training

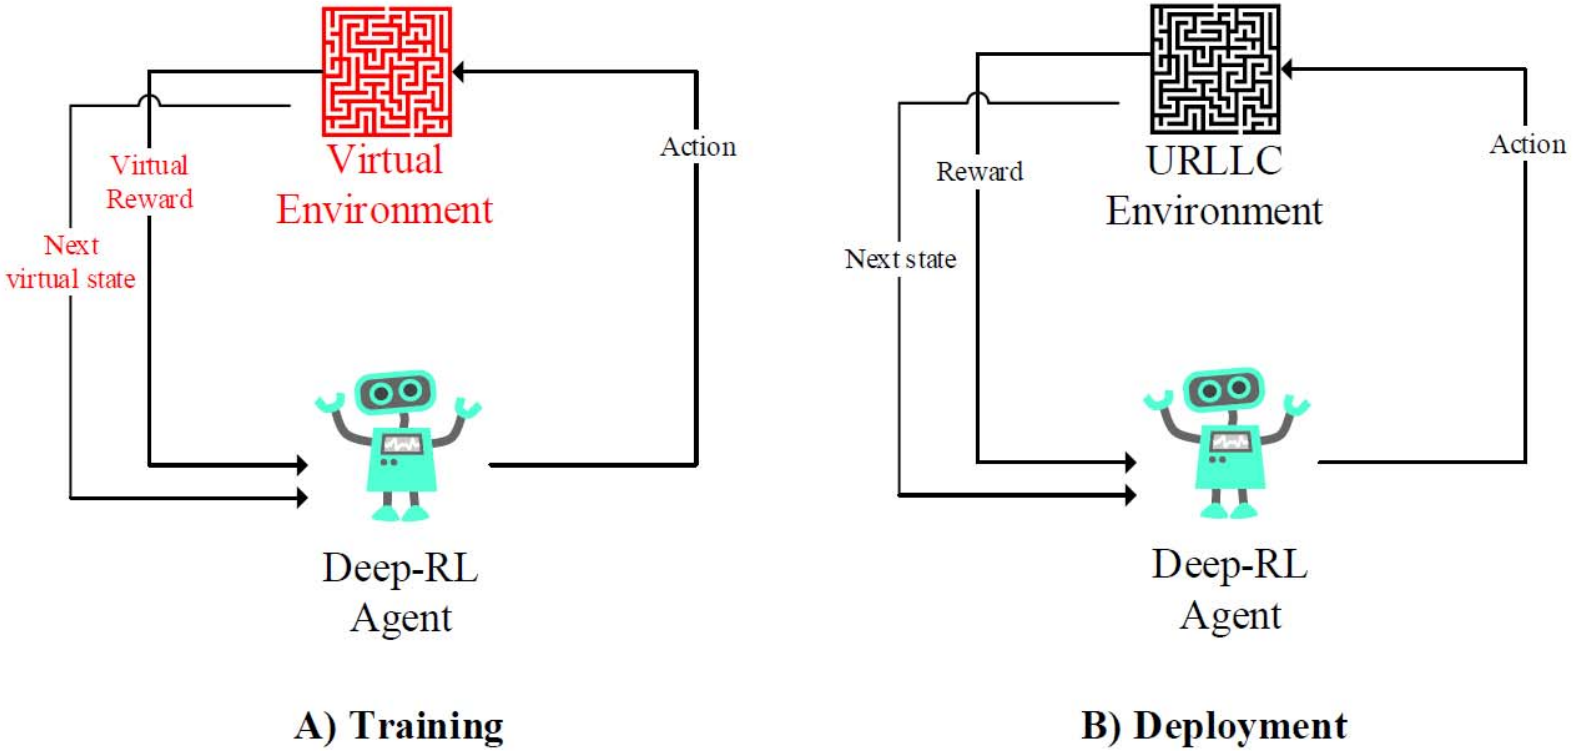

- Virtual environment is created by GAN using a mix of (limited) real data and synthetic (simulated) data

Experienced Deep RL

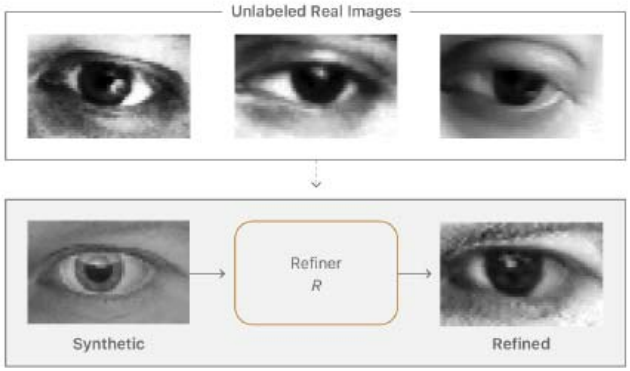

a)

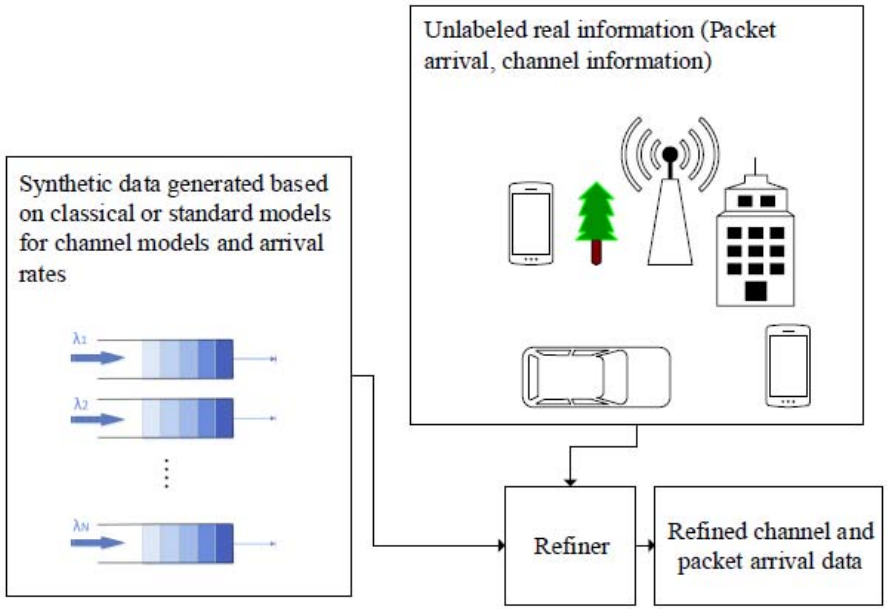

b)

- GAN-based refiner
  - Proposed by Apple for computer vision
- Inputs
  - Unlabeled real data
  - Synthetic model data
- Output
  - Refined (and larger) dataset that includes new network conditions (extreme events) that can train your deep RL

Experienced Deep RL

- We train our deep RL using the GAN-refined data
  - We now have an experienced agent that has been exposed to extreme (rare) network conditions/events
  - The experienced agent will be able to better cope with extreme events as well as to converge faster in a URLLC system by eliminating transient period
- The refiner (which is a neural network) is trained as follows:

$$\theta_R^* = \arg\min_{\theta_R} \max_{\theta_D} f(\theta_R, \theta_D) = \arg\min_{\theta_R} f(\theta_R, \theta_D^*(\theta_R)),$$

$$\text{s.t.} \quad \mathbb{E}_{z \sim g_{\text{sim}}}[\|F(z;\theta_R) - z\|] < \epsilon_r,$$

Refiner weights

Discriminator weights

Refined, real-like data

Control
Real vs. synthetic

Experienced Deep RL

- Theorem 2: The refiner cannot be trained (i.e., problem is infeasible) if:

$$\epsilon_r < \epsilon_r^t \qquad \epsilon_r^t = \sqrt{\|\mu_R\|^2 + \|\mu_z\|^2 - 2\mu_R^T \mu_z}.$$

  - Threshold is function of the expected values of synthetic data and refiner output
  - We can control how our data is being generated
  - There is also an upper bound but hard to characterize mathematically
- Using our GAN and these results, we can create a training environment for ANY deep RL agent

Simulation Results

- We use a real dataset with specific packet sizes and inter-arrival times (with some modification)

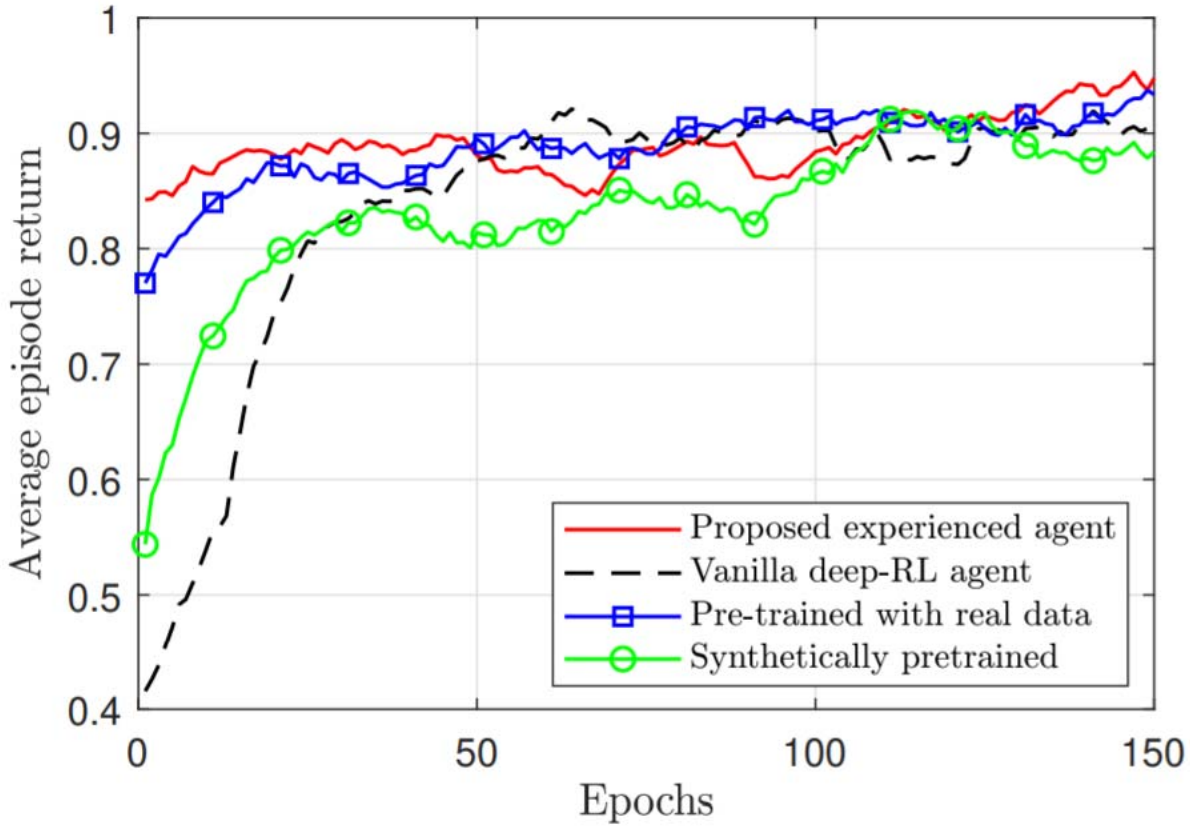

Simulation Results

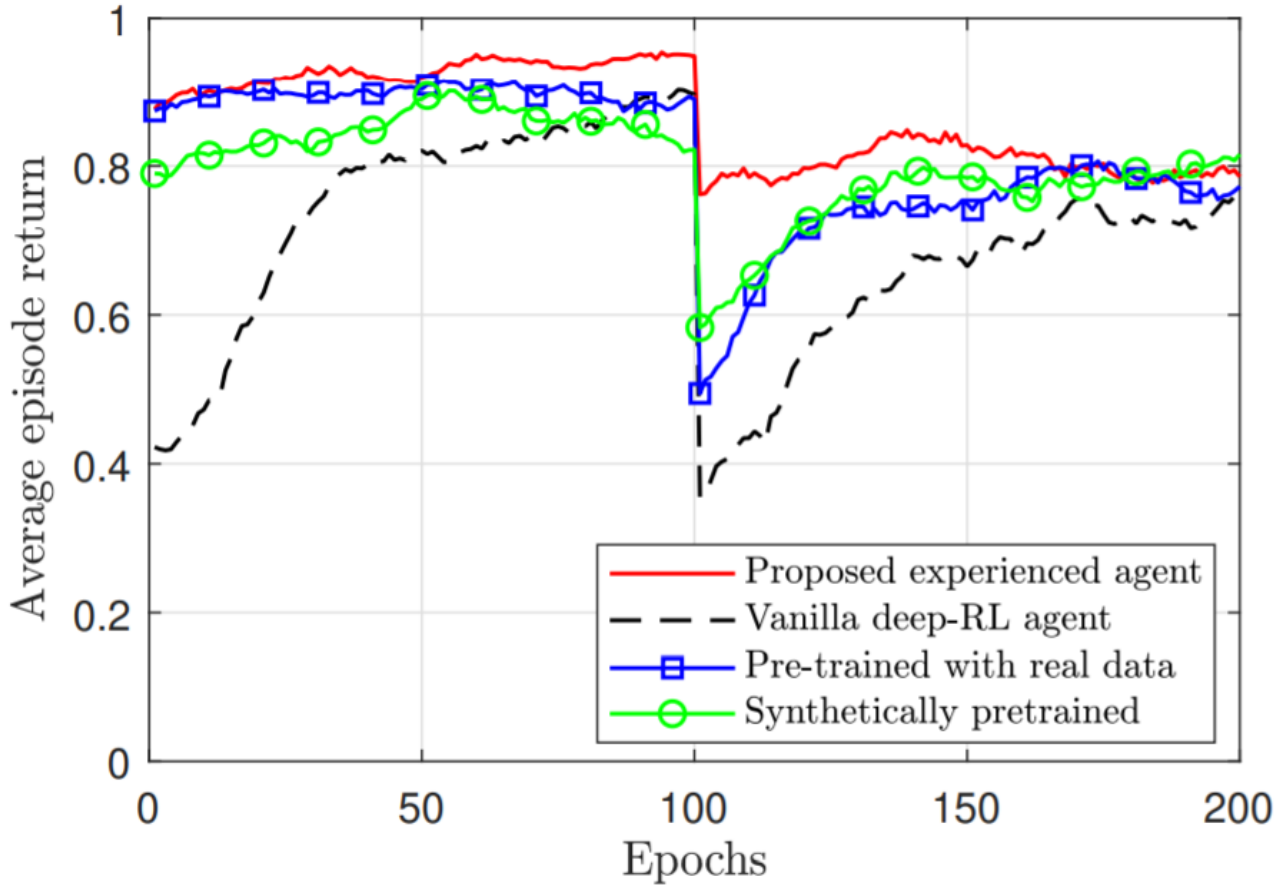

- Experience allows a very smooth handling of extreme events compared to vanilla deep-RL

Simulation Results

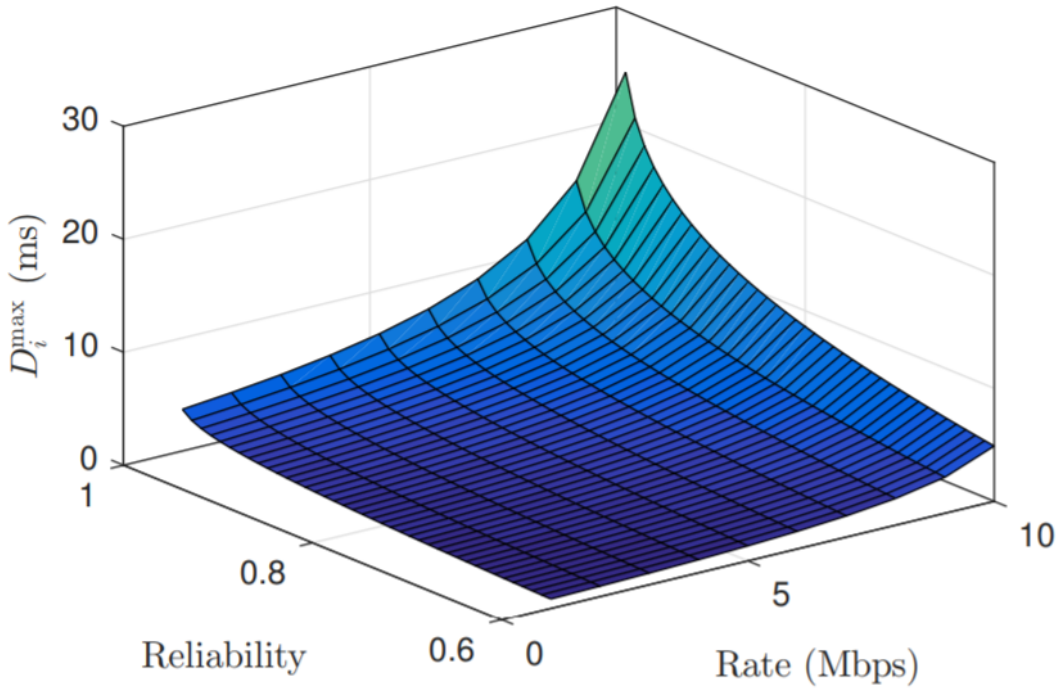

- Rate-reliability-latency tradeoff
  - 99% reliability, 4.2 ms latency, but rate of 1 Mbps
  - To gain 1% reliability, 47% lower delay but 7-times lower rate
  - Higher rate, higher power needed to have higher rates

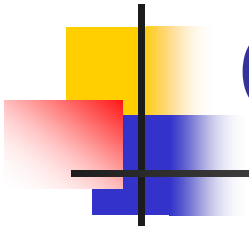

Open Questions

- Explore disentangled representations (e.g., identify "pieces" of the data that can potentially be synthesized)
- Extensions to multi-agent scenario is a very interesting aspect (how to look at scale in that case)
  - Can we have a multi-agent GAN? (Discussed next)
  - Can we have a multi-agent, generalizable RL? (discussed next)
  - Can we work with "distributions" not averages? (see Q. Zhang, W. Saad, et al.)
- We used a deep Q network
  - Can we design new deep RL architectures with a more interesting backbone ANN? (see M. Chen, W. Saad, et al.)

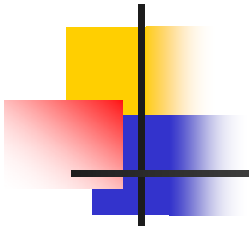

Distributed, Brainstorming Generative Adversarial Networks (BGANs): Framework and Applications

A. Ferdowsi and W. Saad, "Brainstorming Generative Adversarial Networks (BGANs): Towards Multi-Agent Generative Models with Distributed Private Datasets", arXiv:2002.00306.

Q. Zhang, A. Ferdowsi, W. Saad, and M. Bennis, "Distributed Conditional Generative Adversarial Networks (GANs) for Data-Driven Millimeter Wave Communications in UAV Networks", *IEEE Transactions on Wireless Communications*, to appear, 2022.

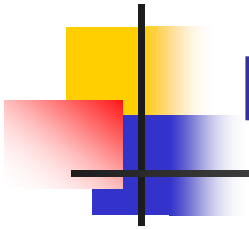

Distributed GANs?

- Existing GAN models (including variants such as InfoGAN, conditional GAN, etc.) are ***centralized***
- What if the data of interest is:
  - Distributed among multiple agents
  - Scarce (each agent has partial data)
  - Private (agents do not want to share their data)
- Can we learn the distribution of the total data **without** sharing the raw data between the agents and **without** relying on a central server?

Existing Distributed GAN Solutions

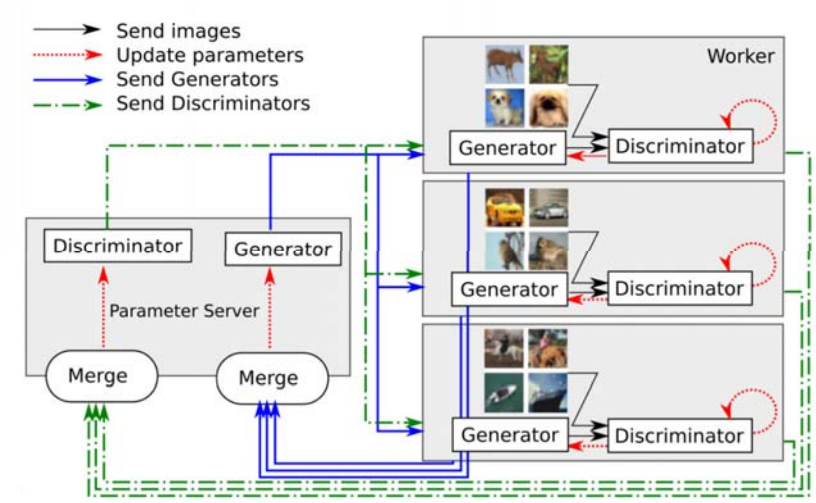

Federated learning (FLGAN)

- Not fully distributed

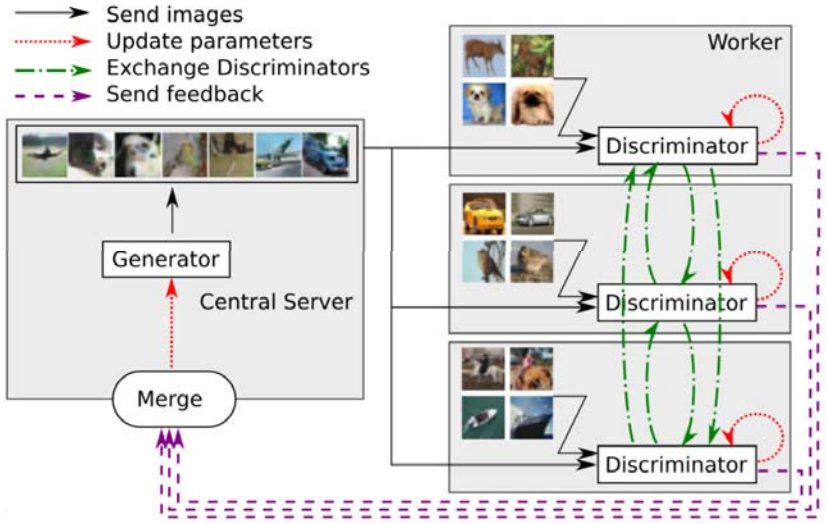

Multi-discriminator (MDGAN)

Forgiver-First Update (F2UGAN)

Drawbacks of the state-of-the art

- Not fully distributed (need a central controller)
- Expensive communication requirements particularly for MDGAN and F2UGAN
- Agents cannot have different neural network architectures and, thus, they must be homogeneous
- Agents do not own their generators
- Can we create a **fully-distributed** solution with multiple, **heterogeneous** agents?

- A. Ferdowsi and W. Saad, "Brainstorming Generative Adversarial Networks (BGANs): Towards Multi-Agent Generative Models with Distributed Private Datasets", arXiv:2002.00306.

Brainstorming GANs

- Architecture allows each agent to have their own generator and discriminator

- **Brainstorming:** Share the generated data points **(ideas)** with other agents at every training epoch

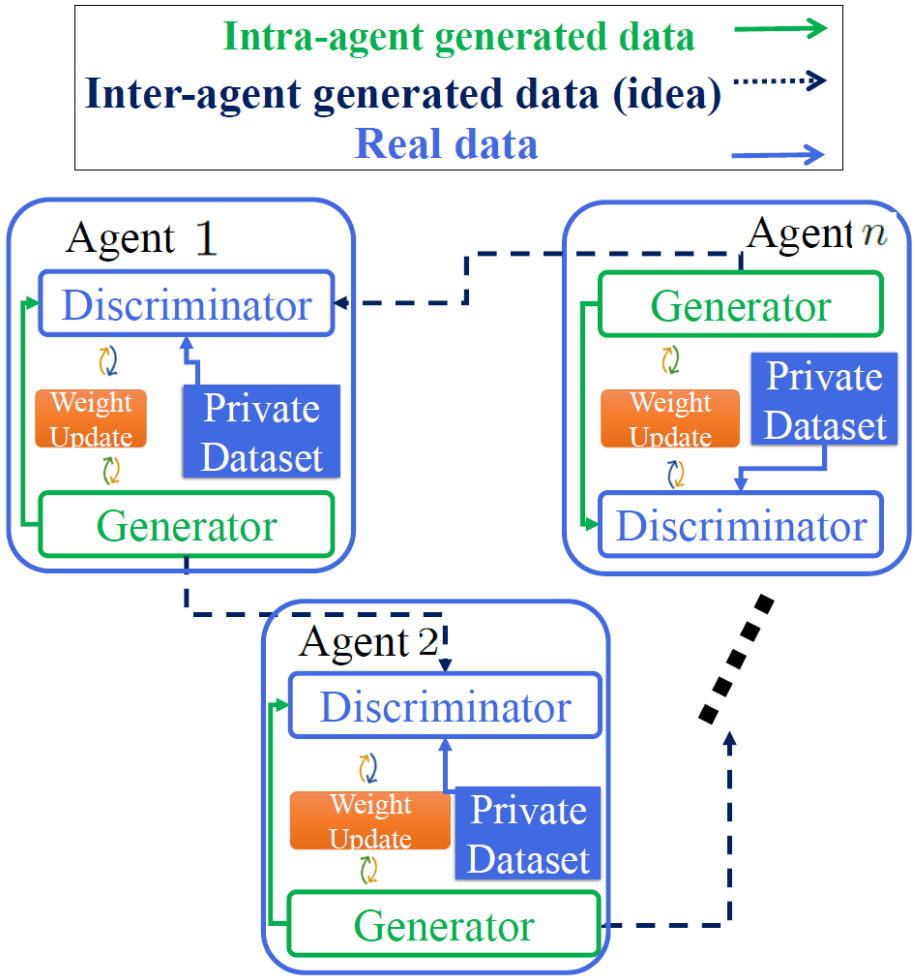

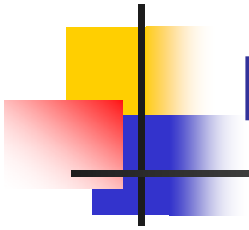

Benefits of BGANs

- Fully distributed, we do not require a central controller or aggregator
- Agents can have different neural network architectures
  - Capabilities-tailored neural networks
- Less communication overhead than most baselines (depends on the dimensions of the data points rather than the neural network parameters)
  - Will be shown to be more efficient in practical cases
- We theoretically show that this GAN architecture admits an equilibrium and is effective

Simulation Results

- If the agents own partial, non-overlapping data, can they figure out the entire distribution?

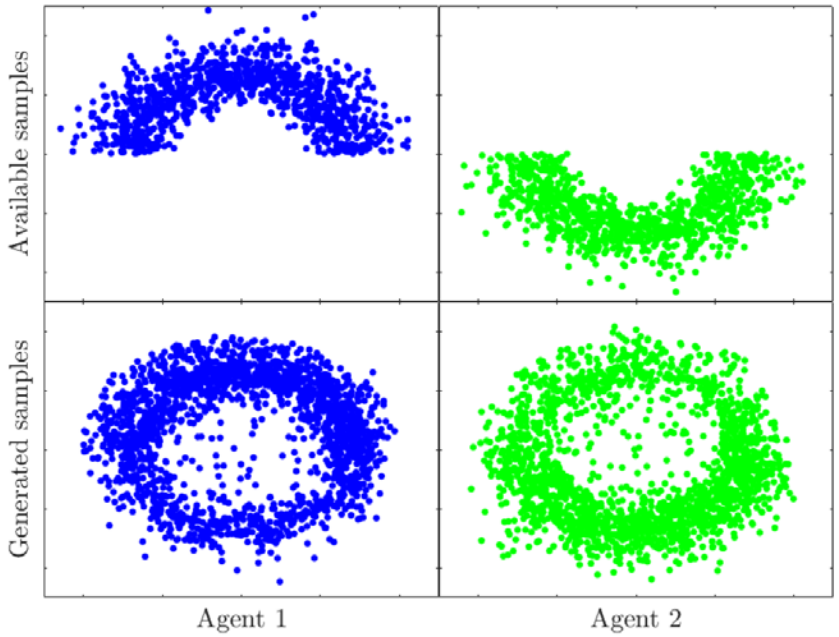

- Each agent owns part of the circle
- Each agent owns a single digit

Simulation Results

- If the agents own partial, non-overlapping data, can they figure out the entire distribution?

Agent who owns airplane images

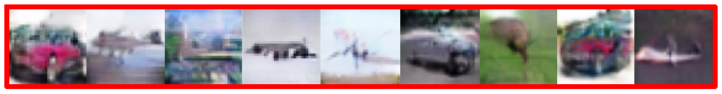

Agent who owns automobile images

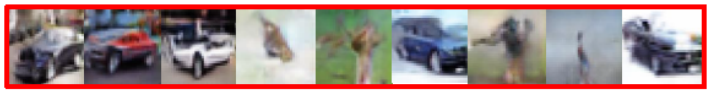

Agent who owns bird images

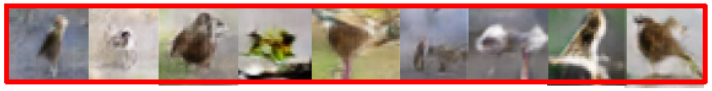

Agent who owns cat images

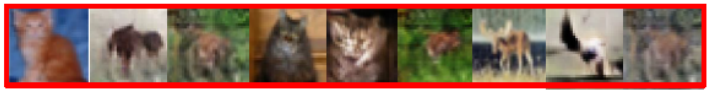

Agent who owns deer images

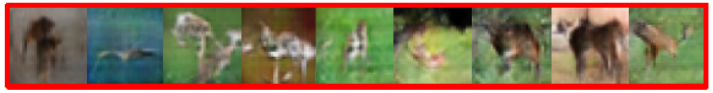

Agent who owns dog images

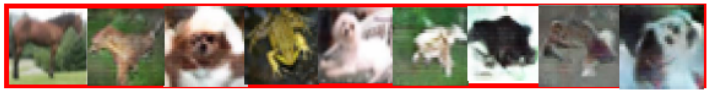

Agent who owns frog images

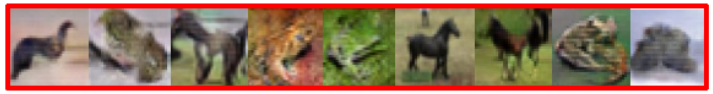

Agent who owns horse images

Agent who owns ship images

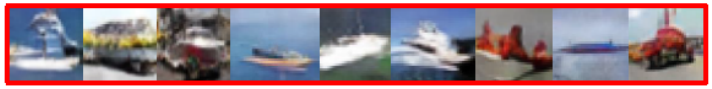

Agent who owns truck images

- CIFAR-10

Simulation Results

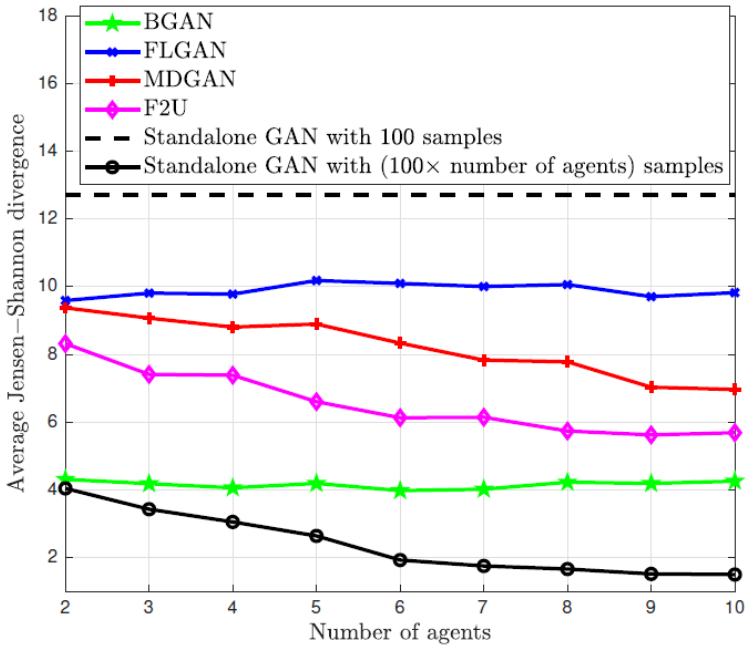

JSD comparison between BGAN, FLGAN, MDGAN and F2U on ring dataset

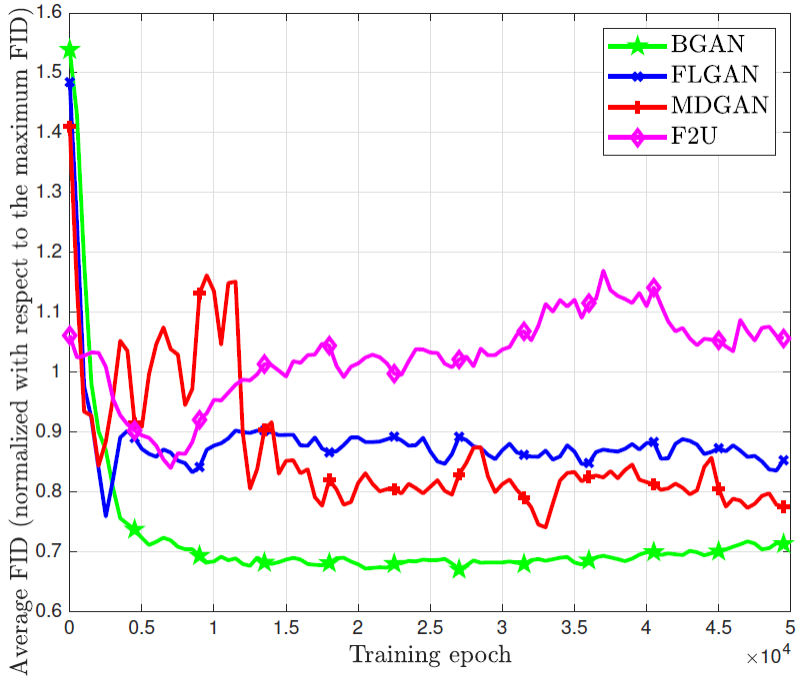

JSD comparison between BGAN, FLGAN, MDGAN, and F2U on the MNIST dataset

n	Number of agents
b	Batch size
$\lvert\boldsymbol{x}\rvert$	Data size
$\lvert\boldsymbol{\theta}_g\rvert$ $\lvert\boldsymbol{\theta}_d\rvert$	Neural network size

Architecture	Communication resources
BGAN	$\mathcal{O}\left(nb\lvert\boldsymbol{x}\rvert\right)$
MDGAN	$\mathcal{O}\left(n\left(b\lvert\boldsymbol{x}\rvert+\lvert\boldsymbol{\theta}_d\rvert\right)\right)$
FLGAN	$\mathcal{O}\left(n\left(\lvert\boldsymbol{\theta}_g\rvert+\lvert\boldsymbol{\theta}_d\rvert\right)\right)$
F2U	$\mathcal{O}\left(n\left(b\lvert\boldsymbol{x}\rvert+1\right)\right)$

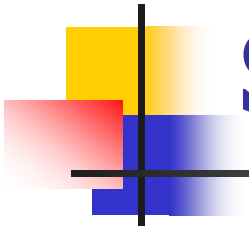

Summary

- We showed that distributed GAN models with partial datasets that are distributed across devices can be devised with no centralized control
- What can we do with this next?
  - Enhance security/privacy
  - Distributed BGAN discriminator for inference
  - More sophisticated network connectivity and graphs
  - Applications to security (intrusion detection, GC'19)
  - Applications of BGAN to wireless networks (from channel modeling to resource allocation and vehicular networking)
- Let's see an example application

BGAN-based UAV Channel Modeling

- Can we create a spatio-temporal map through UAV collaboration?
  - Assume stable air-to-air links
- Each UAV has its local dataset, a generator, and a discriminator
- Each UAV shares its generated channel samples with other UAVs, by forming a distributed learning network
- All generators collaboratively generate channel samples to fool all of the discriminators
- Reformulated with BGAN
- We then find the optimal architecture that enables a fast convergence time while heeding network resources

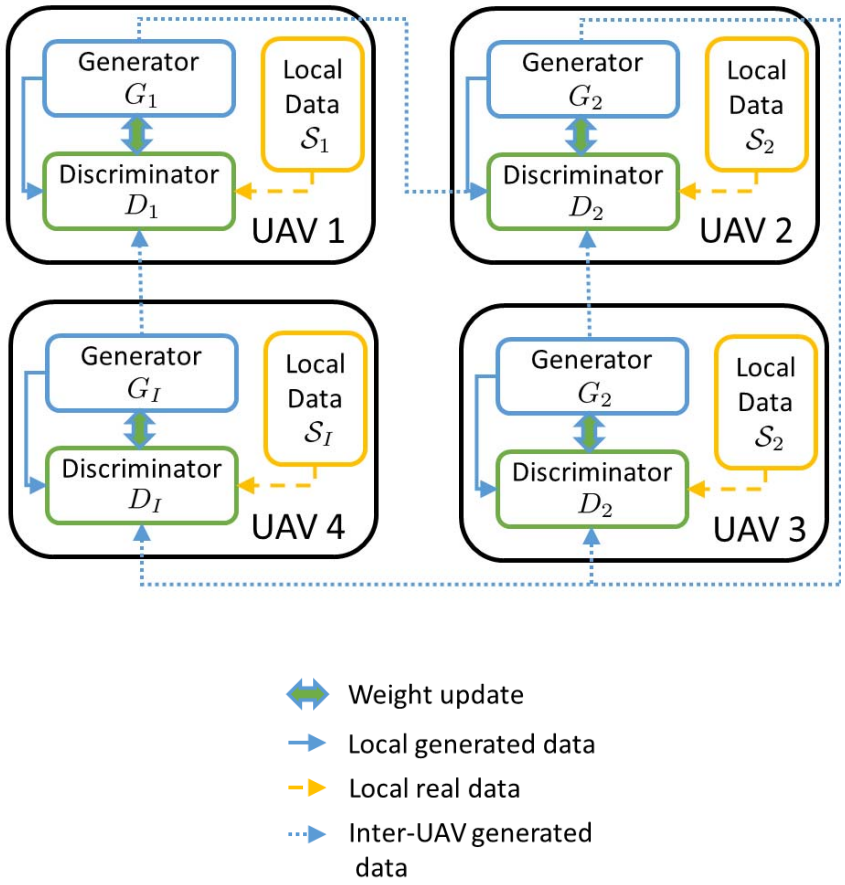

Simulation Results

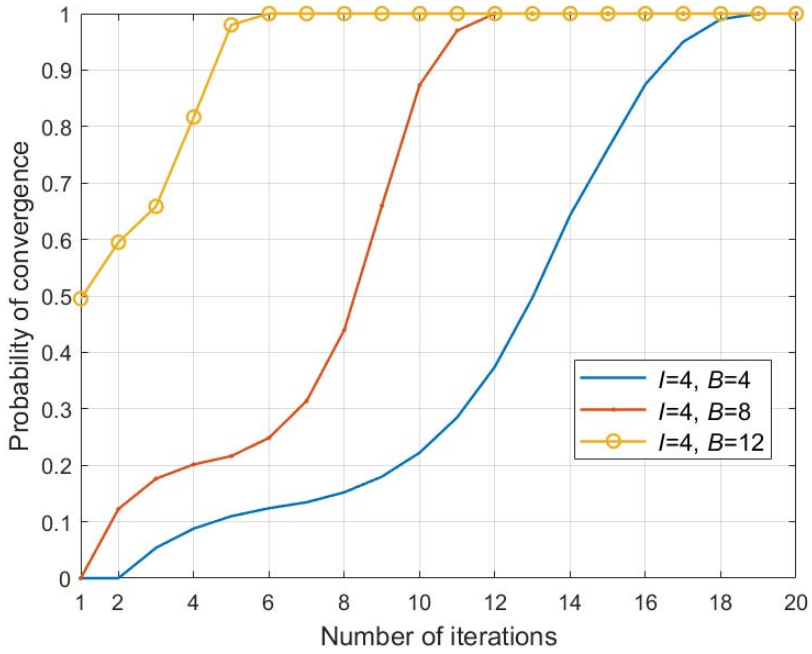

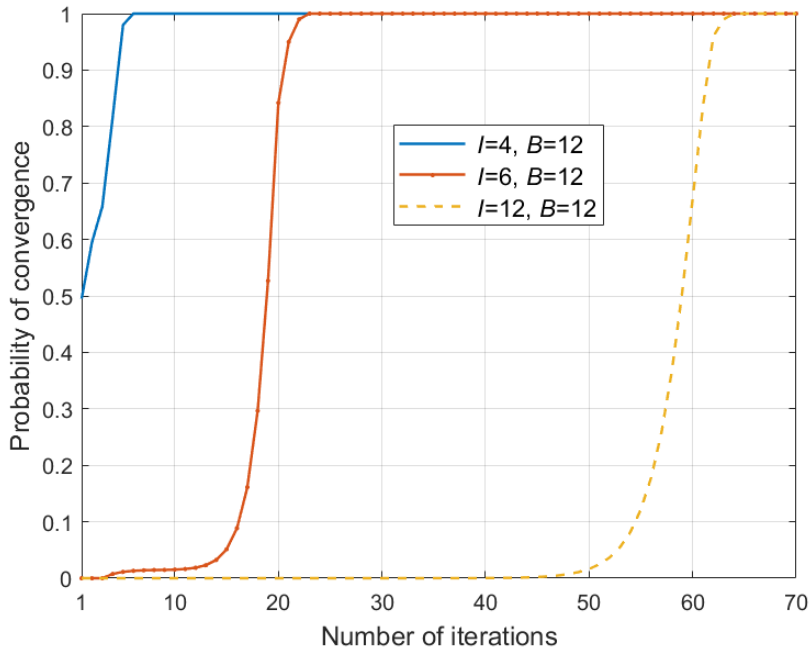

- Learning rate improves with more communication resources *B* for A2A transmissions, and decreases for larger networks (larger *I*).
- Q. Zhang, A. Ferdowsi, W. Saad, and M. Bennis, "Distributed Conditional Generative Adversarial Networks (GANs) for Data-Driven Millimeter Wave Communications in UAV Networks", arXiv:2102.0175.

Generalizable Multi-Agent Reinforcement Learning with Meta-Learning

Y. Hu, M. Chen, W. Saad, H. V. Poor, and S. Cui, "Distributed Multi-agent Meta Learning for Trajectory Design in Wireless Drone Networks", *IEEE Journal on Selected Areas in Communications (JSAC), Special Issue on UAV Communications in 5G and Beyond Networks*, vol. 39, no. 10, pp. 3177 - 3192, Oct. 2021.

Towards Generalizable Multi-Agent RL

- Use case for wireless-aware drone trajectory planning

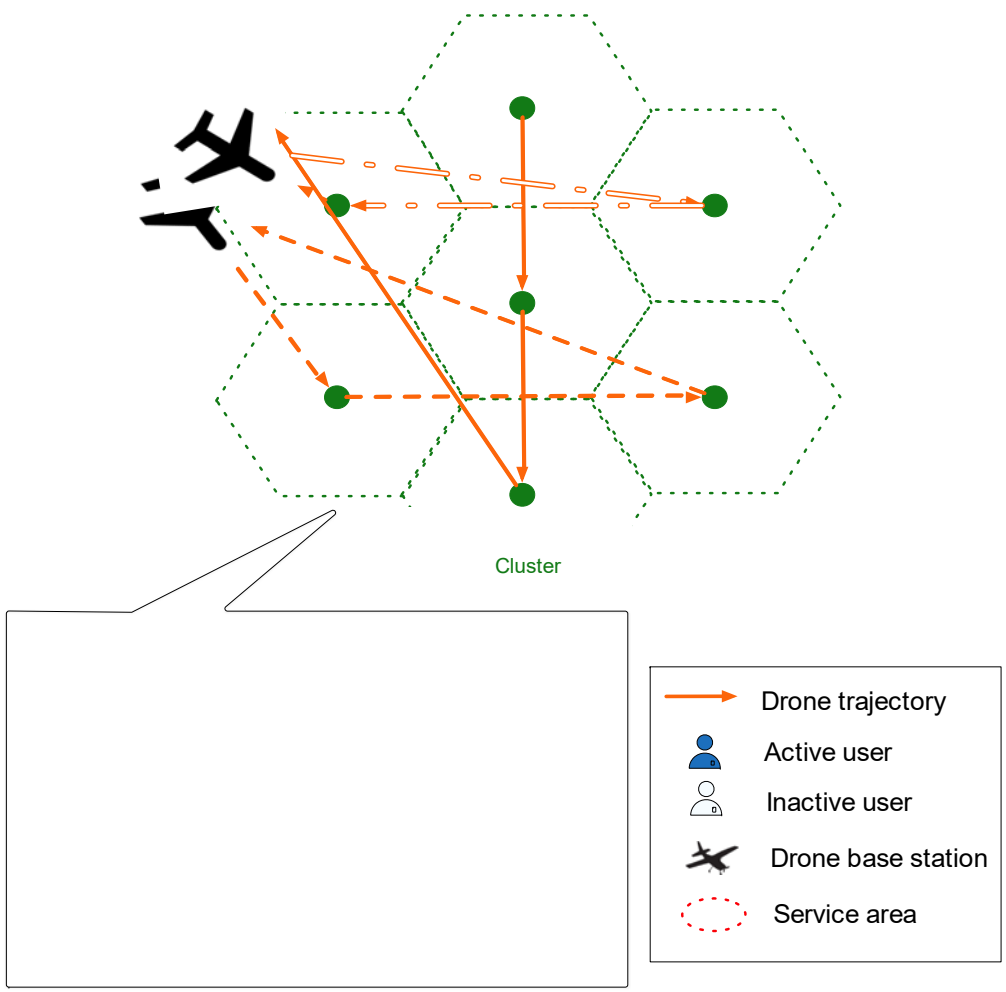

- **Key Goal:** Fly drone base stations in a way to maximize coverage (reliability) for the maximum number of users
- Problem can be formulated as a non-convex coverage optimization problem
- RL can be a solution but existing multi-agent RL have several drawbacks
  - Large overhead for coordination
  - Difficulty to generalize to unseen tasks/environments
  - High variance

Towards Generalizable Multi-Agent RL

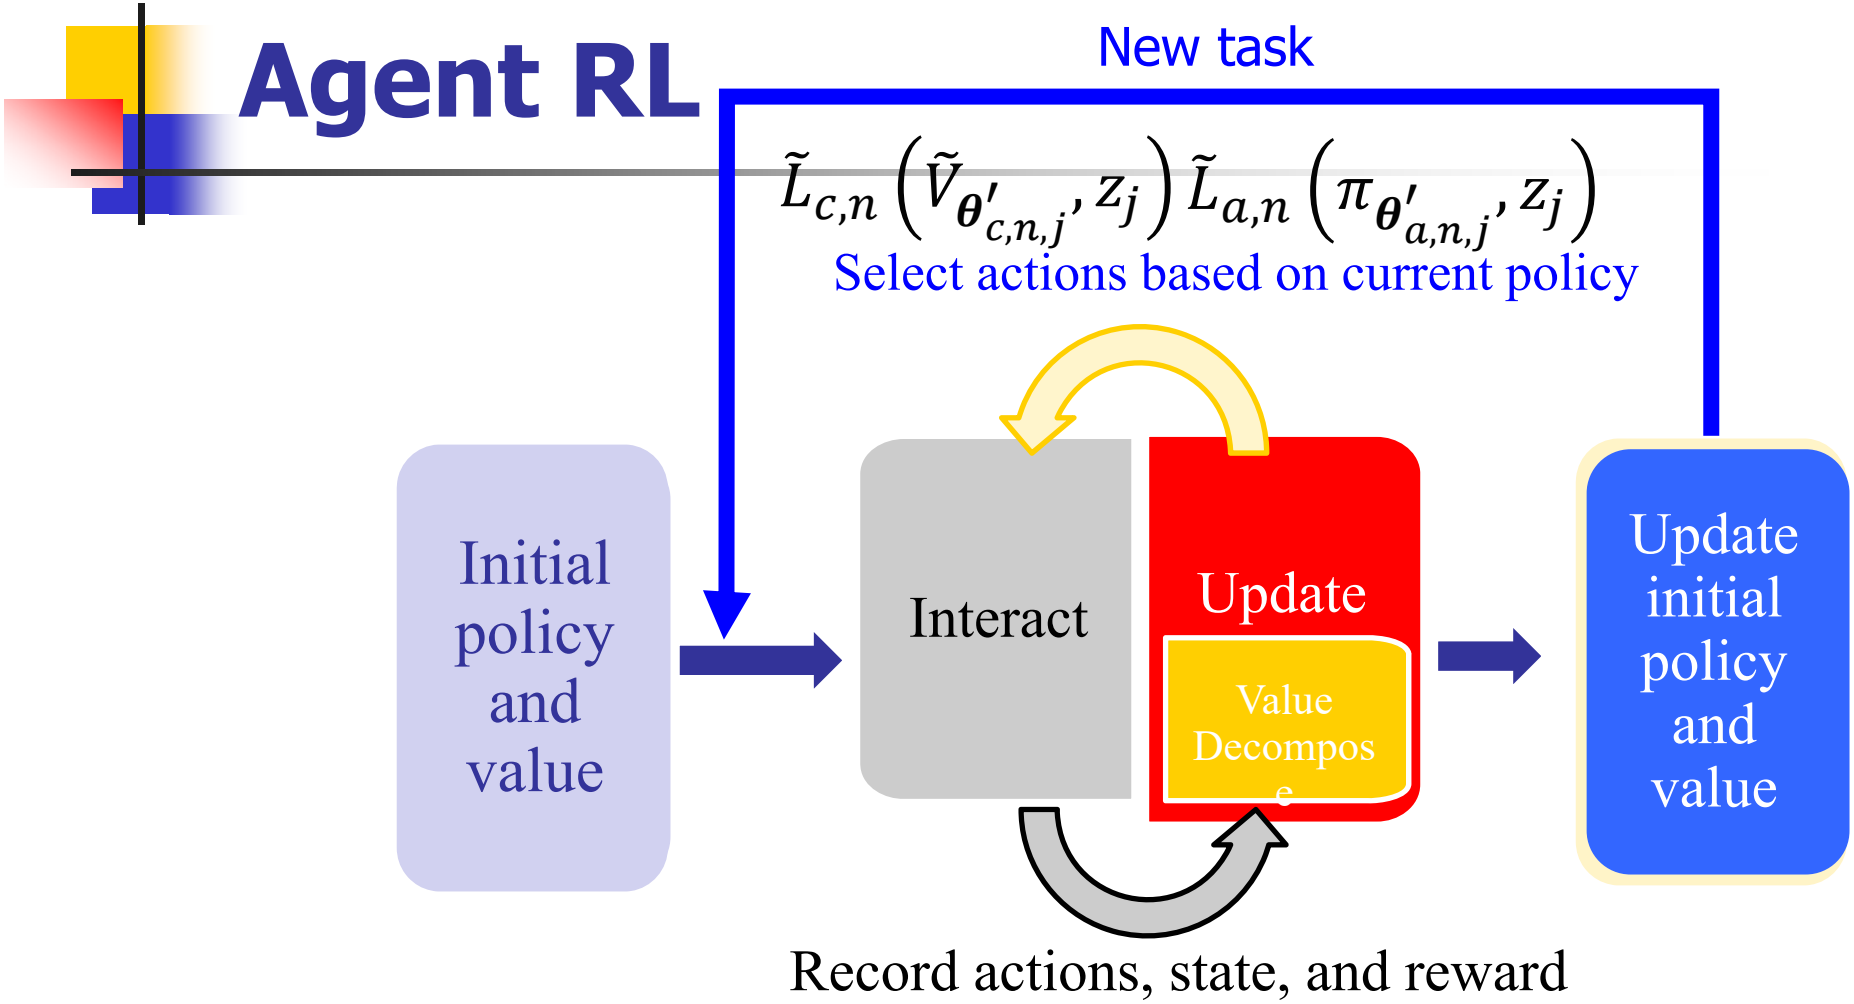

- Value decomposition => reduce overhead, use local observations
- Policy gradient to enhance variance
- Meta-learning to enable each agent to learn a "skill" and generalize, i.e., learn a group/distribution of tasks

Towards Generalizable Multi-Agent RL

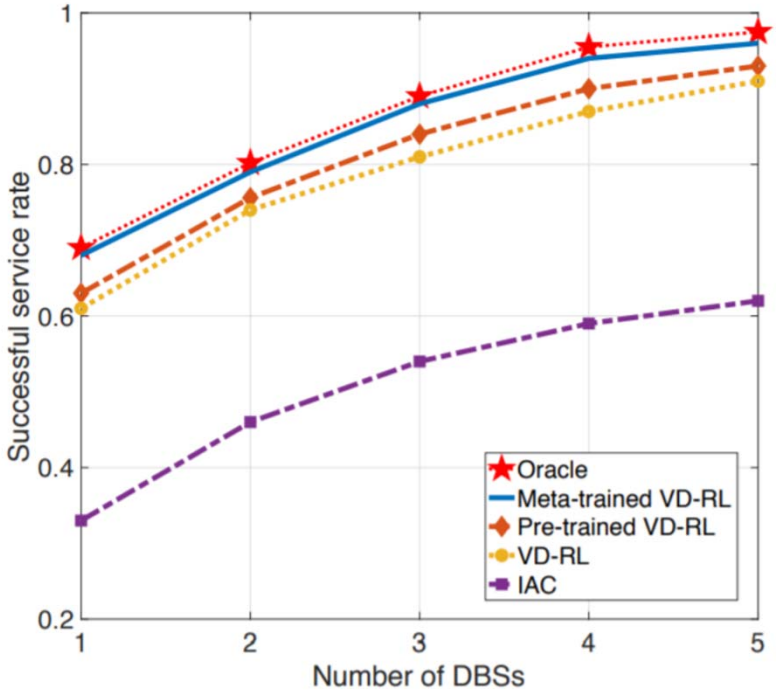

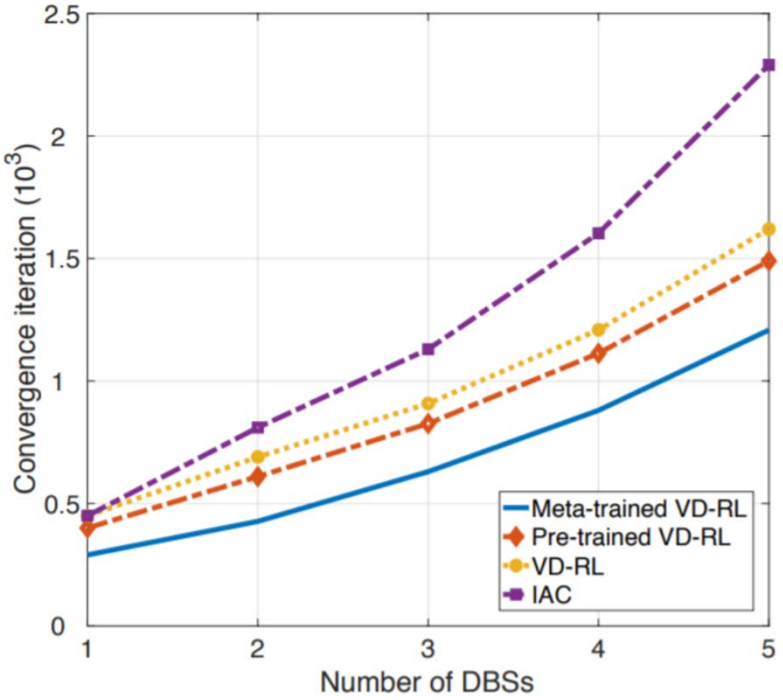

- Meta-RL approach is better in terms of both network performance, and convergence time
- Gap is significant compared to classical independent actor critic algorithm (IAC)

What's next in AI?

- Issues of reliability/generalization (particularly at high frequencies): Experienced/meta RL
- **Beyond** standard learning:
  - Learning is mostly training-based, can we get rid of training?
  - We are pursuing several ideas in this context:
    - Continual learning that can retain lifelong features
    - Generalize from small data (reasoning over data) through causal learning and/or other related concepts
  - Theoretical foundations for RL/explainable AI
- Rich field of application in 6G with specific needs
  - If we want AI-native networks, we must go towards generalizable learning frameworks

Joint Learning and Communications: Edge AI as a Wireless Use Case

M. Chen, Z. Yang, W. Saad, C. Yin, H. V. Poor, and S. Cui, "A Joint Learning and Communications Framework for Federated Learning over Wireless Networks", *IEEE Transactions on Wireless Communications*, vol. 20, no. 1, pp. 269 - 283, January 2021.

FL as a Wireless Use Case

- Centralized training

$$\min_{\boldsymbol{w}} \frac{1}{K} \sum_{i=1}^{U} \sum_{k=1}^{K_i} f\left(\boldsymbol{w}, \boldsymbol{x}_{ik}, y_{ik}\right), \quad \Rightarrow \quad \boldsymbol{w}_{t+1} = \boldsymbol{w}_t + \frac{1}{K} \sum_{i=1}^{U} \sum_{k=1}^{K_i} \nabla f\left(\boldsymbol{w}, \boldsymbol{x}_{ik}, y_{ik}\right),$$

- Distributed training

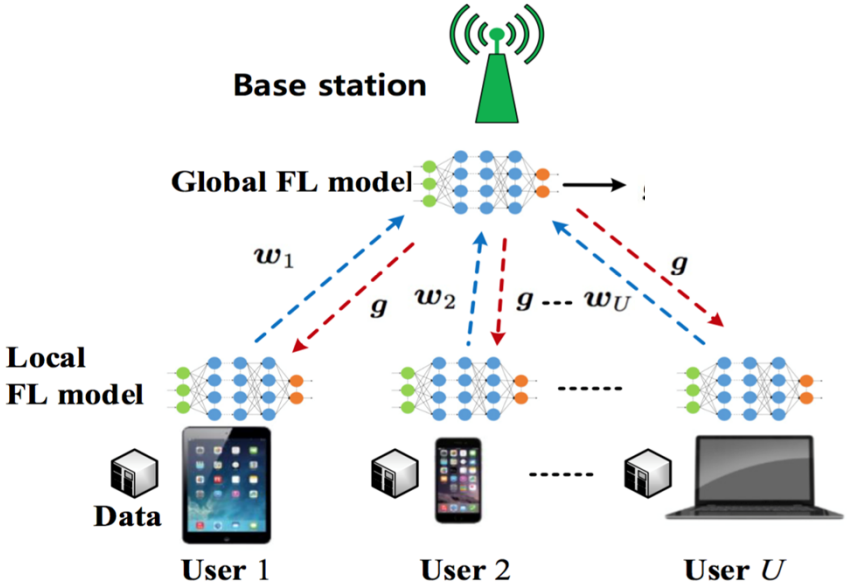

$$\boldsymbol{g}_t = \frac{\sum_{i=1}^{U} K_i \boldsymbol{w}_{i,t}}{K}$$

Without data transmission

$$\boldsymbol{w}_{i,t+1} = \boldsymbol{g}_t + \frac{1}{K_i} \sum_{k=1}^{K_i} \nabla f\left(\boldsymbol{g}_t, \boldsymbol{x}_{ik}, y_{ik}\right),$$

Joint FL-Wireless Design

- FL training process in a wireless network

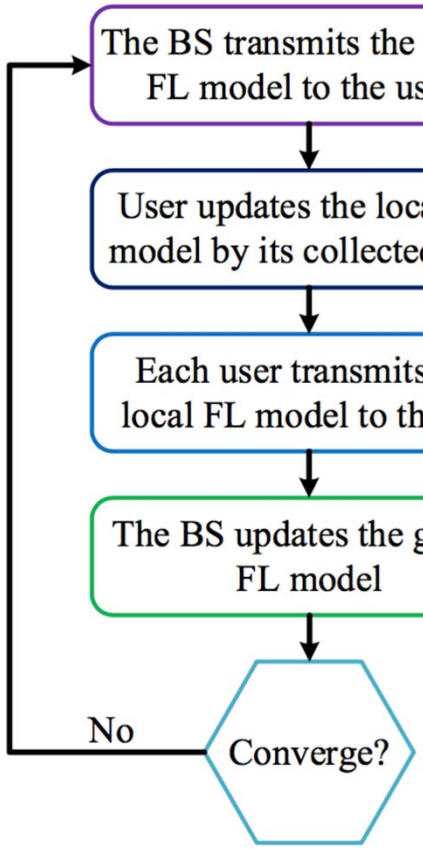

  - Since the global and local FL model parameters need wireless link exchanges, then ***wireless transmission errors will impact FL performance***
  - The base station (BS) must update the global FL model as it receives all of the local FL model transmitted from the users. Hence, ***transmission delay and energy consumption must be considered.***
  - ***Communication to support FL!***
  - M. Chen, Z. Yang, W. Saad, C. Yin, H. V. Poor, and S. Cui, "*A Joint Learning and Communications Framework for Federated Learning over Wireless Networks*", *IEEE Transactions on Wireless Communications*, 2021.

Global FL Model Analysis

- **Theorem 1:** An upper bound on the convergence point of FL over wireless networks can now be found:

$$\mathbb{E}[F(\boldsymbol{g}_{t+1}) - F(\boldsymbol{g}^*)] \leq \underbrace{\frac{2\zeta_1}{LK} \sum_{i=1}^{U} K_i \left(1 - a_i + a_i q_i \left(\boldsymbol{r}_i, P_i\right)\right) \frac{1 - A^t}{1 - A}}_{\text{Impact of wireless factors on FL convergence}} + A^t \mathbb{E}(F(\boldsymbol{g}_0) - F(\boldsymbol{g}^*)),$$

Total samples

Number of samples of user *i*

Global model at convergence

where $A = 1 - \frac{\mu}{L} + \frac{4\mu\zeta_2}{LK} \sum_{i=1}^{U} K_i \left(1 - a_i + a_i q_i \left(\boldsymbol{r}_i, P_i\right)\right)$

- This is a key characterization of FL performance over a wireless network
  - Convergence affected by PER and user association: wireless network must be reliable enough to support effective FL
  - We used full gradient descent but results extendable

Simulation Results

Identification accuracy as the number of users varies.

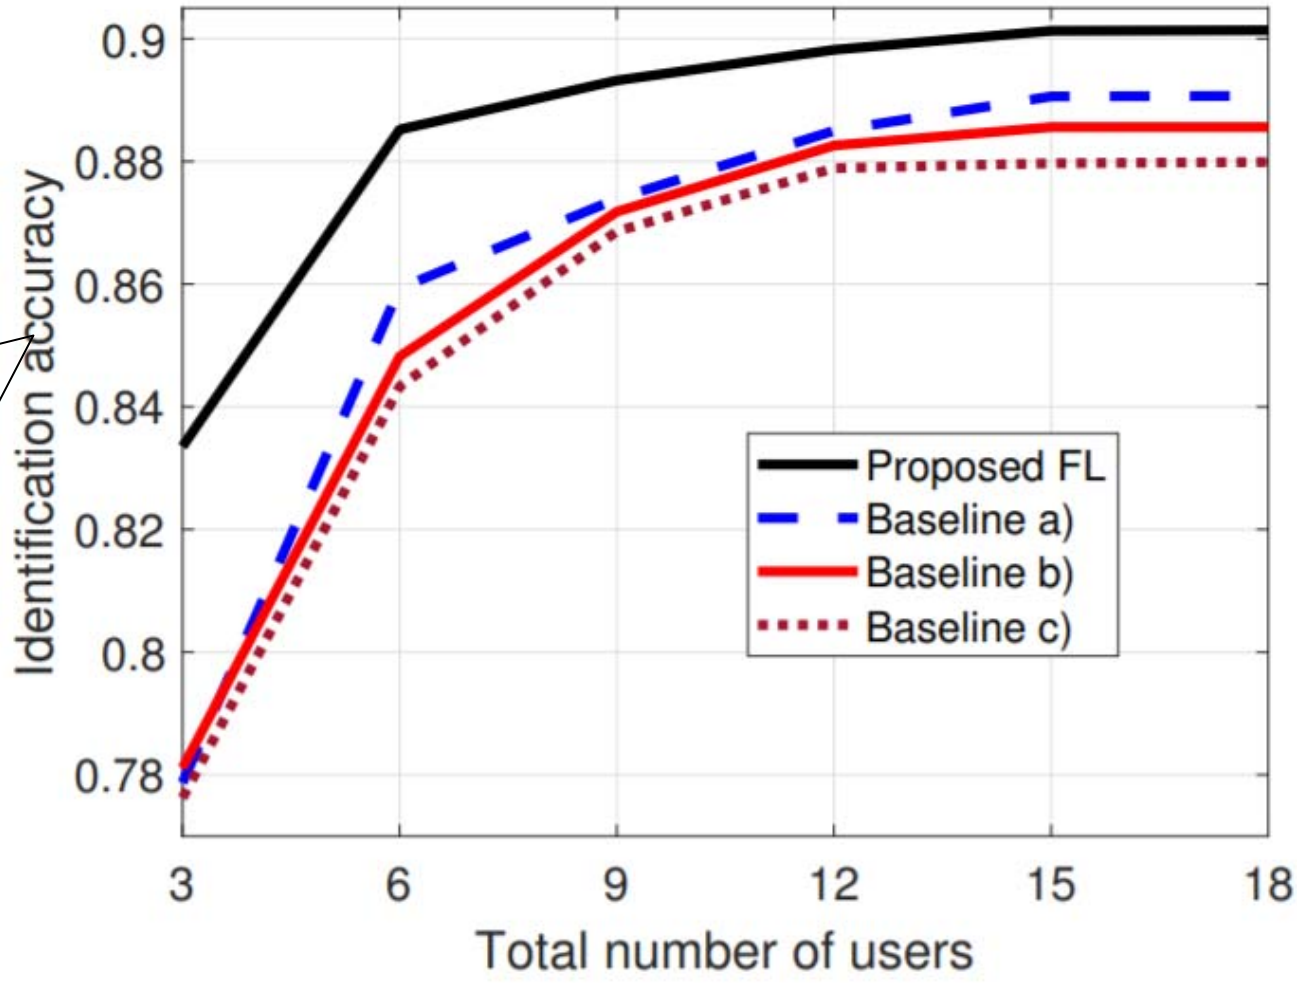

Energy vs. Precision vs. Accuracy

- Fundamental tradeoff: energy, precision, accuracy
  - Federated quantized neural networks

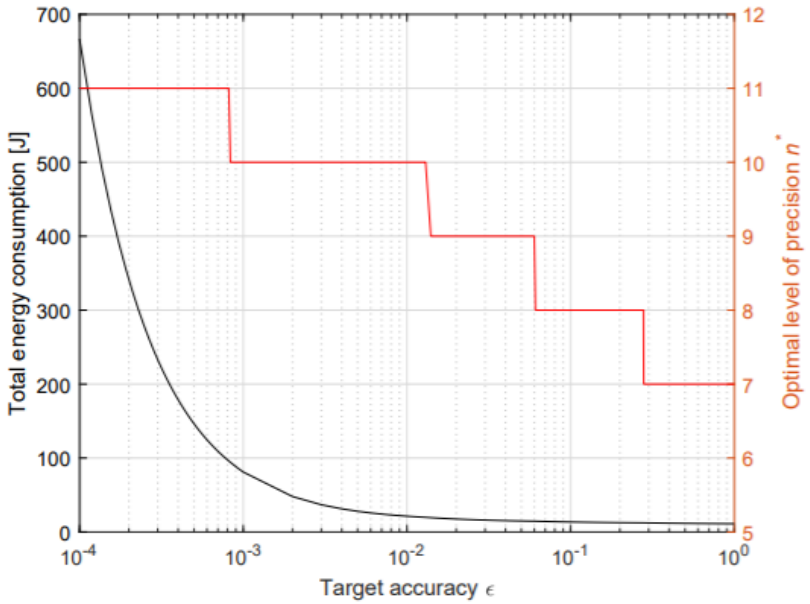

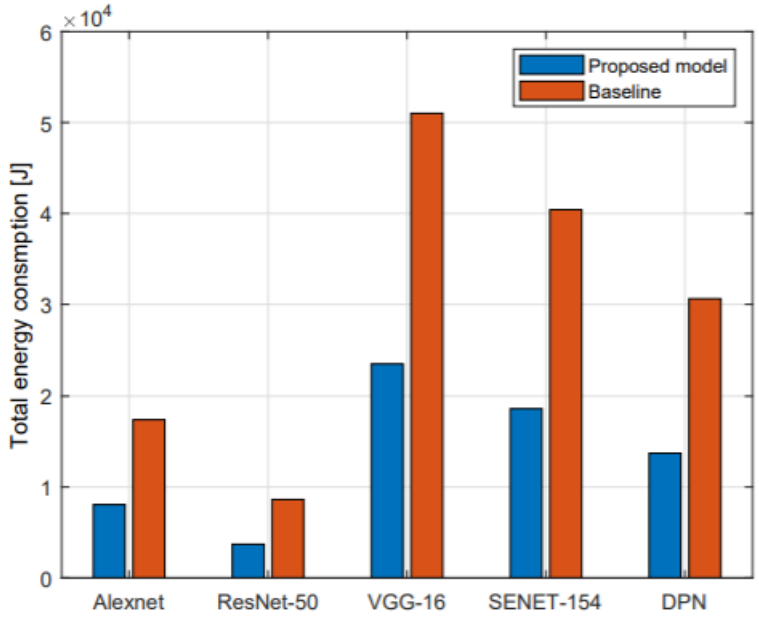

- M. Kim, W. Saad, M. Mozaffari, and M. Debbah, "On the Tradeoff between Energy, Precision, and Accuracy in Federated Quantized Neural Networks", in *Proc. of the IEEE International Conference on Communications (ICC), Green Communication Systems and Networks Symposium*, Seoul, South Korea, May 2022.

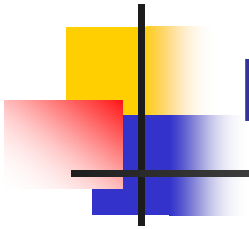

More on FL and Communications?

- Convergence time analysis (with M. Chen, IEEE TWC 2021, IEEE ICC 2020 – best paper award)
- Energy efficiency challenges (with Z. Yang, IEEE TWC 2021)
- Fully distributed FL with no central controller (with M. Chen, IEEE Communications Magazine, to appear 2021)
- Incentives for FL (with L. U. Khan, IEEE Communications Magazine, 2020)
- FL for vehicular networks (with S. Samarakoon, IEEE TCOM 2020)
- FL for virtual reality optimization (with M. Chen, IEEE TWC 2020)
- FL with drones and vehicles (with T. Zeng, IEEE ICC 2020, IEEE TWC submitted)

Other research areas

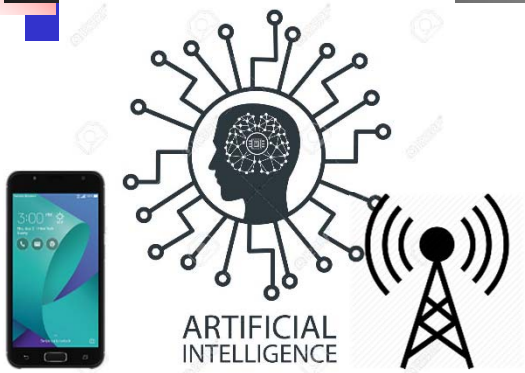

- 5G/6G/IoT systems
  - Reliable, low latency comm. with ML/GAN
  - Terahertz/RIS
  - AI-enabled XR/Twins

- AI-native semantic communications
  - Semantic information
  - Generalizable AI
  - End-to-end analysis

- Connected drones and autonomous vehicles
  - Distributed learning and control
  - Wireless connectivity
  - Sensing and comm.

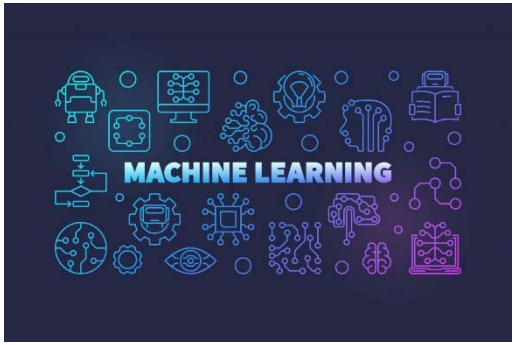

- Reliable, generalizable, distributed learning
  - Reliable machine learning
  - Meta-learning and training-free learning
  - Distributed and multi-agent learning
  - Continual learning

Other research areas

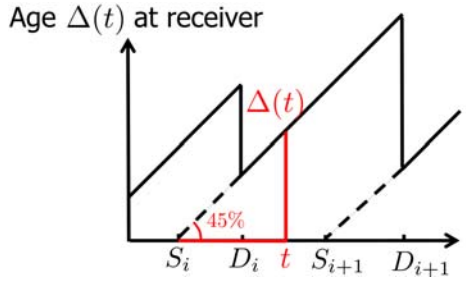

- Age of information
  - Performance analysis of Internet of Things systems with age of information considerations
  - Information management

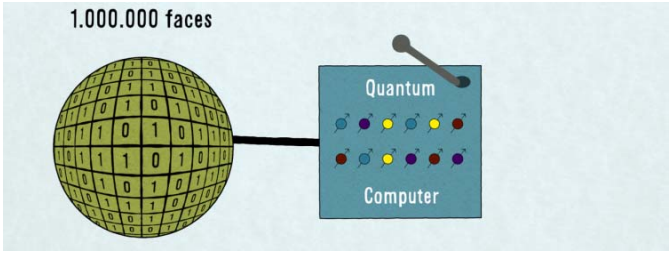

- Quantum networks
  - Communications with qubits
  - Quantum algorithms
  - Physics-informed networking

- Smart cities
  - Big data for smart city optimization
  - Air pollution
  - Security

- Game theory
  - Foundations
  - Applications to CPS, security, policy, wireless

- CPS security
- Blockchains
- Capsules

Conclusions

- Distributed learning is an exciting area, particularly when merged with wireless
- Distributed learning is not just federated learning
  - Distributed RL, multi-agent systems, GANs
- Generalizable machine learning is the way forward in wireless
  - Using a mixture of augmentation, multi-task, continual, and explainable AI
- BGAN is the first fully distributed GAN
- Abundant field of applications
  - From wireless design to semantic communications

Finally….

Thank You, Questions?