

Federated Learning:

A Systems Engineering Approach to Smarter AI

Dr. Art Villanueva, ESEP

www.phronos.com

September 24, 2025



PHRONOS
PRACTICAL AI



Art Villanueva, DEng, ESEP



Dr. Art Villanueva lives in the intersection of **artificial intelligence, emergence, and systems engineering**. His current research focuses on complex adaptive systems (CAS) – from brain mechanisms to autonomous unmanned aerial vehicles (UAVs). He is the founder of Phronos, an AI services company dedicated to leveraging AI and SE across domains. He also serves as a Senior Staff Engineer at General Atomics Aeronautical Systems, where he applies his expertise to the next-generation of aerospace programs.

Dr. Villanueva previously served as Chief AI Technologist for Dell Technologies' Federal Strategic Programs, guiding Dell's AI posture across government sectors. His multidisciplinary experience also spans large defense and public transportation projects as well as clean-tech entrepreneurship in renewable energy.

An inventor with multiple U.S. utility patents and several peer-reviewed publications, Dr. Villanueva holds a Doctor of Engineering in Systems Engineering from Colorado State University, where he specialized in meta-algorithmics for natural language processing. He also holds master's degrees in Systems Engineering and Computer Science from UC San Diego, and a B.S. in Applied Mathematics from UCLA. He is an INCOSE-certified Expert Systems Engineering Professional (ESEP).



Disclaimer

The views and opinions expressed in this presentation are my own and do not represent those of General Atomics Aeronautical Systems, Inc. or any affiliated entities. This presentation is for informational purposes only and is unrelated to my professional role or work at General Atomics.


Agenda

- I. Previously...
- II. Intro and Motivation
- III. Types of Machine Learning
- IV. Federated Learning
- V. Types of Federated Learning
- VI. Examples of Federated Learning
- VII. Swarm Intelligence
- VIII. Technical Challenges
- IX. Conclusion
- X. Q-A-D: Question, Answer, Discussion

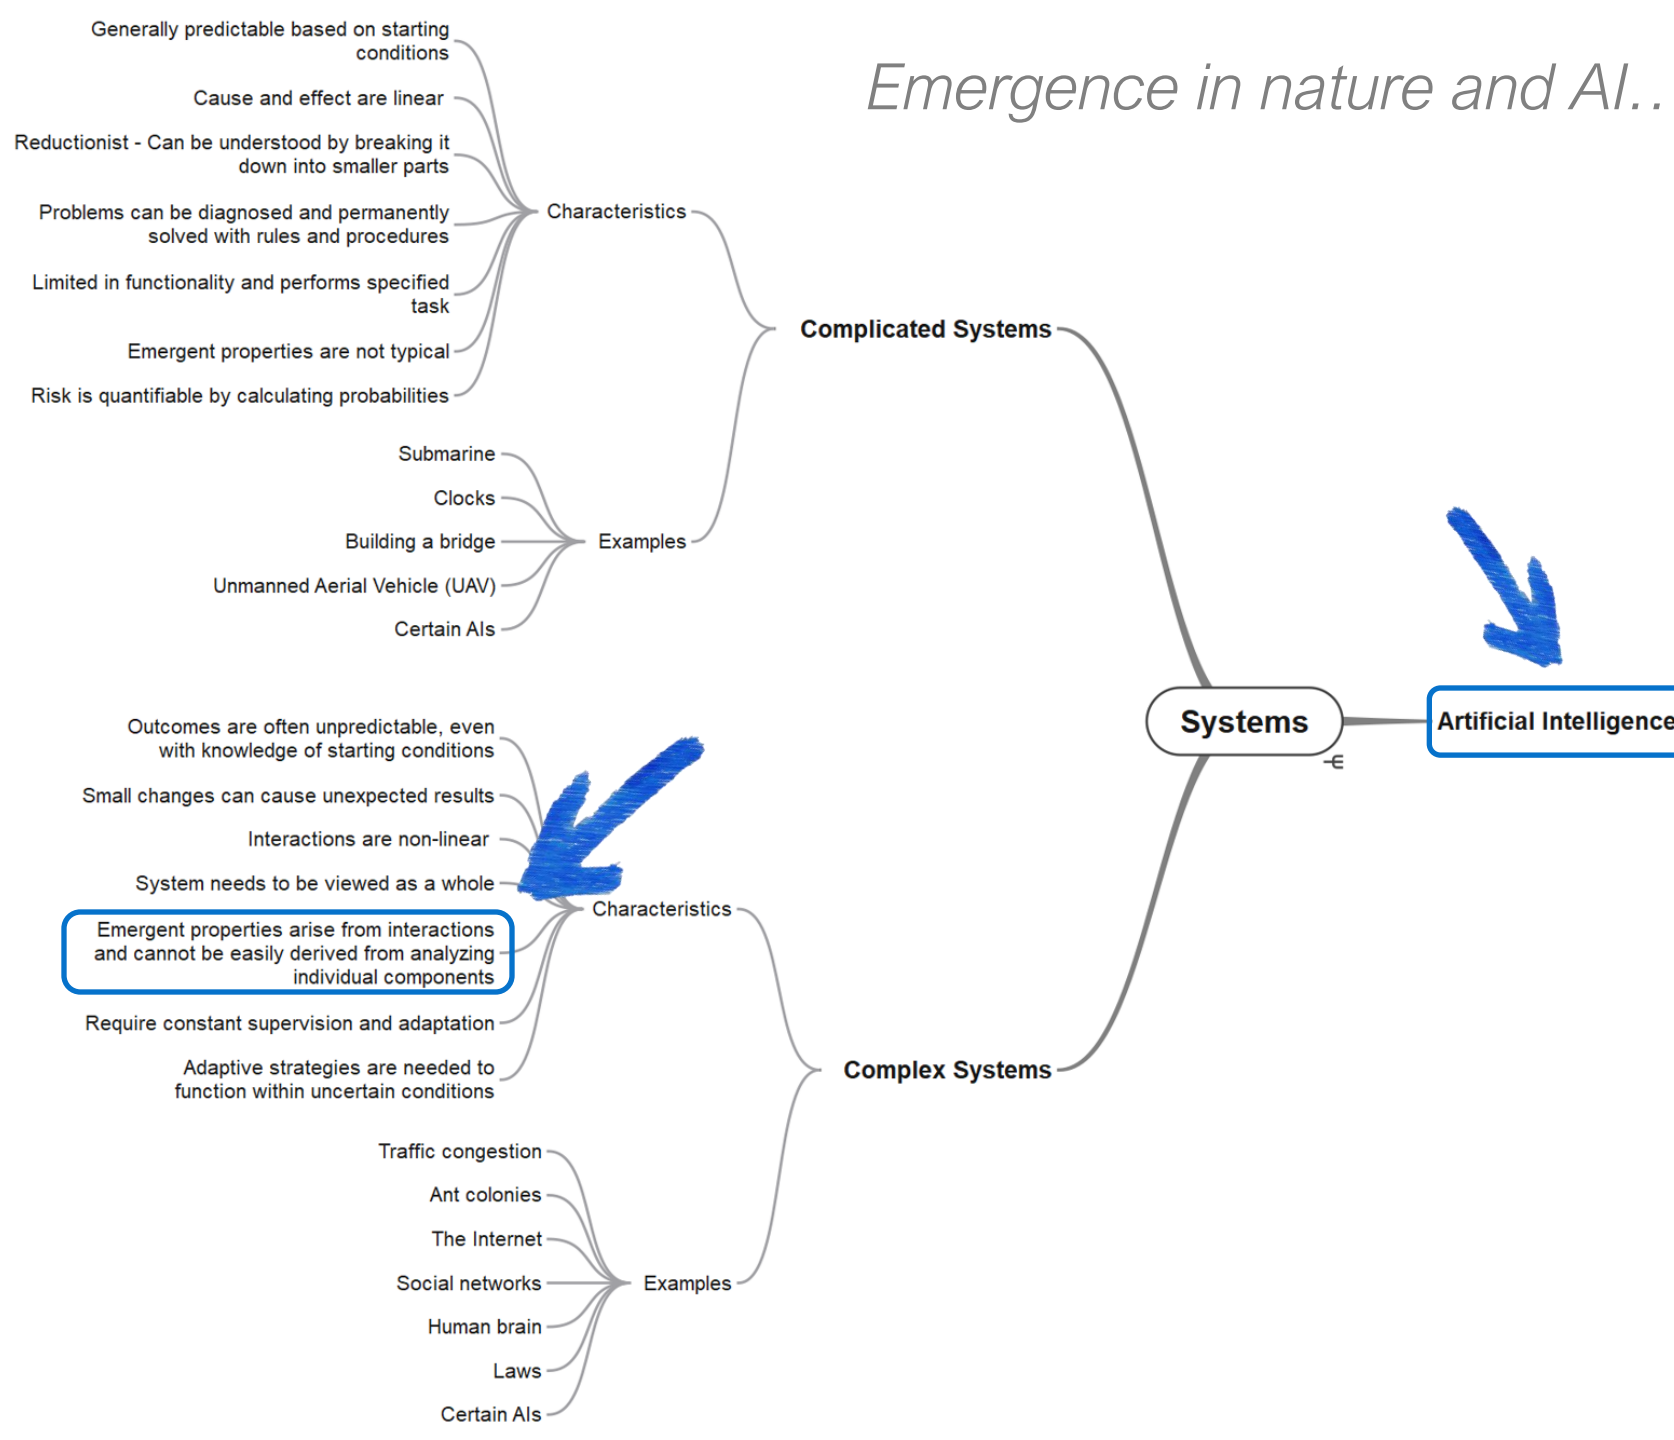


Previously

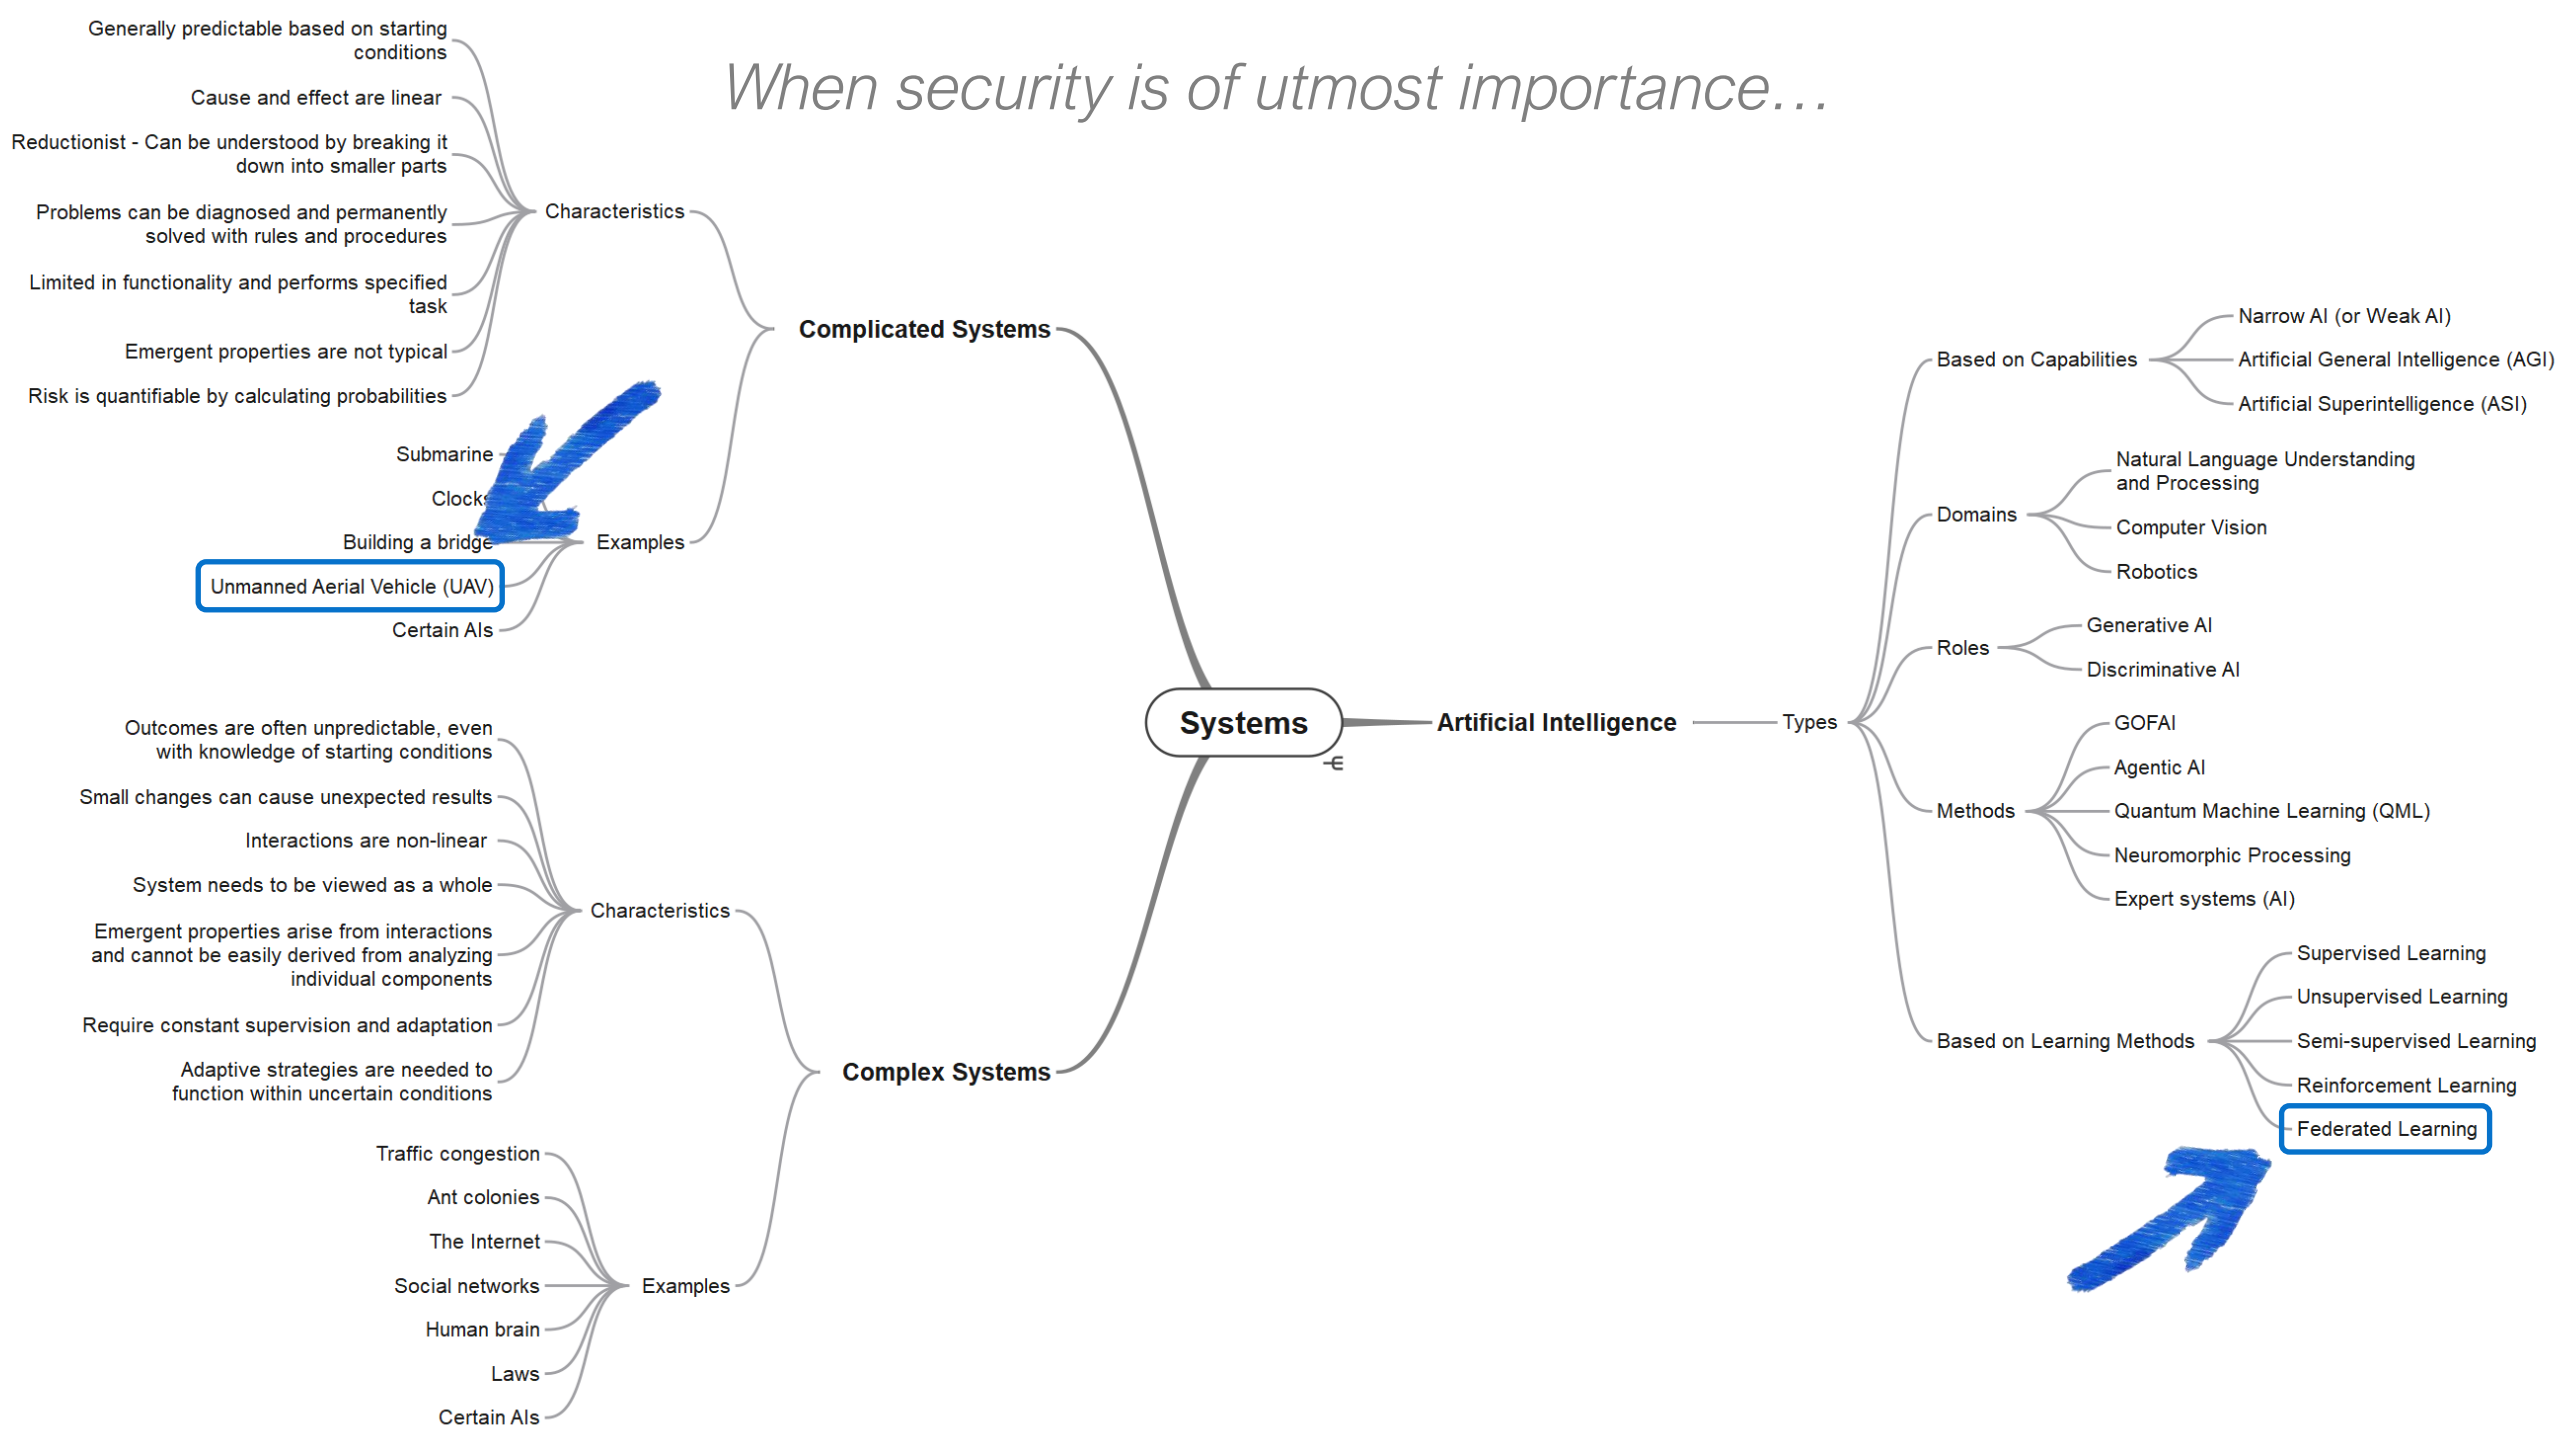
When ChatGPT 3.5 became a household name...



Emergence in nature and AI...



When security is of utmost importance...





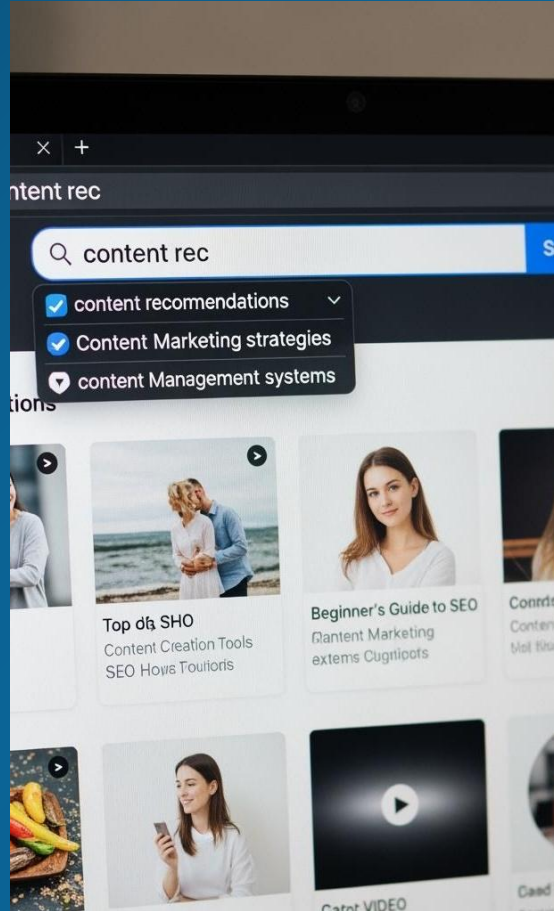
Intro and Motivation

What do these scenarios have in common?

Wildfire Management



Personalized Recommendations

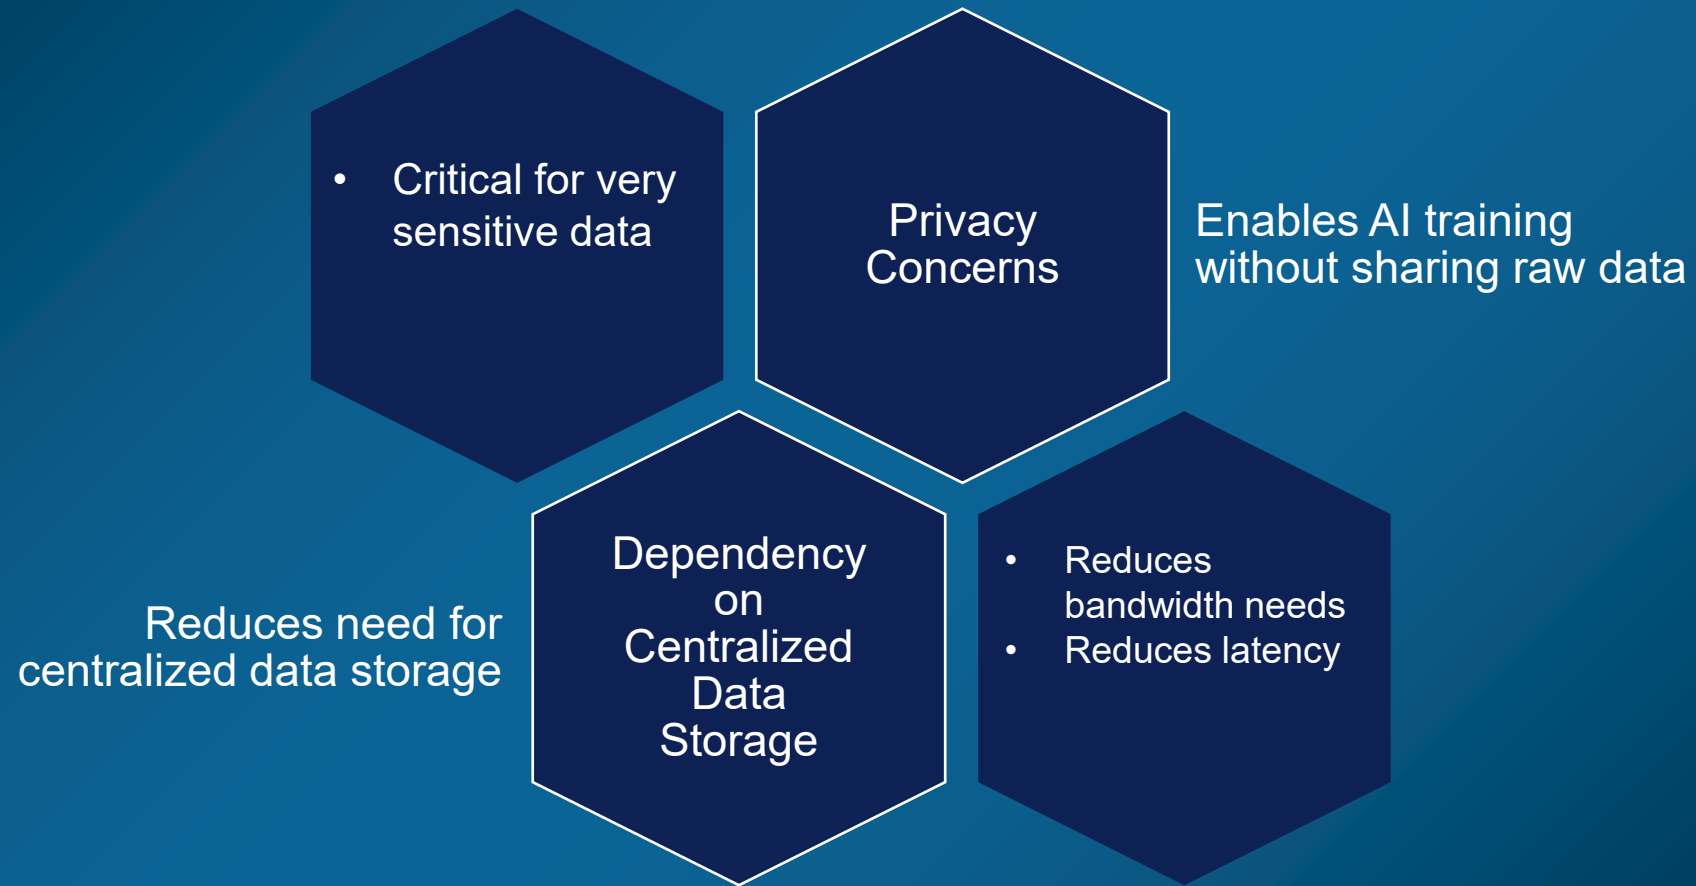


Disease Spread Mitigation



When data sharing is less than ideal...

Why Federated Learning?



Why Now?



- Rise of edge computing

- phones, UAVs, sensors
- data is coming from everywhere

- Data Breaches:

- Data at rest: 60–80% of breaches
- Data in motion: 10–25% of breaches
- Data in use: 15–30% of breaches

- Privacy and regulatory drivers

- e.g., GDPR, HIPAA

- Growing demand for on-device intelligence

Data Glut

Sensitive Data

Limited Connectivity

What is Federated Learning?

Federated learning is a machine learning paradigm across decentralized devices while keeping data local.






Types of Machine Learning

Discriminative vs Generative AI



Supervised ML


Predict outcomes from labeled data



- Spam Detection
- House Price Prediction
- Medical Diagnosis

Unsupervised ML


Discover hidden patterns in unlabeled data



- Customer Segmentation
- Anomaly Detection
- Topic Modeling

Semi-supervised ML

Use both labeled and unlabeled data to improve learning



- Email classification with few labels
- Image recognition with limited labels
- Speech-to-text with small transcriptions

Self-supervised ML

Generate labels from data itself to learn representations



- Next-word prediction
- Filling missing image parts
- Solving image rotations or puzzles

Reinforcement Learning

Train agents to make sequential decisions



- Game AI
- Autonomous Driving
- Robotics

Transfer Learning

Leverage pretrained models for new tasks

Pretraining

Use a large dataset to train a base model



Feature Extraction

Remove the final layers to extract useful features



Fine-tuning

Train the model on the new domain/task with a smaller dataset



Validation

Evaluate the fine-tuned model's performance



Deployment

Deploy the fine-tuned model for the specific application




- Image Classification
- Natural Language Processing
- Speech Recognition



Federated Learning


Federated Learning

Train models collaboratively without sharing raw data



- Predictive Text
- Healthcare Diagnostics
- Smart Home Devices

Federated Learning



1. Setup – We have a server and several edge devices
2. Server sends the model to the edge devices
3. Edge devices collect data from their environment
4. Edge devices generate gradients (deltas) comparing model prediction with reality
5. Edge devices send gradients to the server
6. Server collects all the gradients and updates the model
7. Repeat #2




$$w_{global} = \frac{\sum_{k=1}^K n_k \cdot w_k}{\sum_{k=1}^K n_k}$$

Federated Averaging
McMahan et al. (2017)



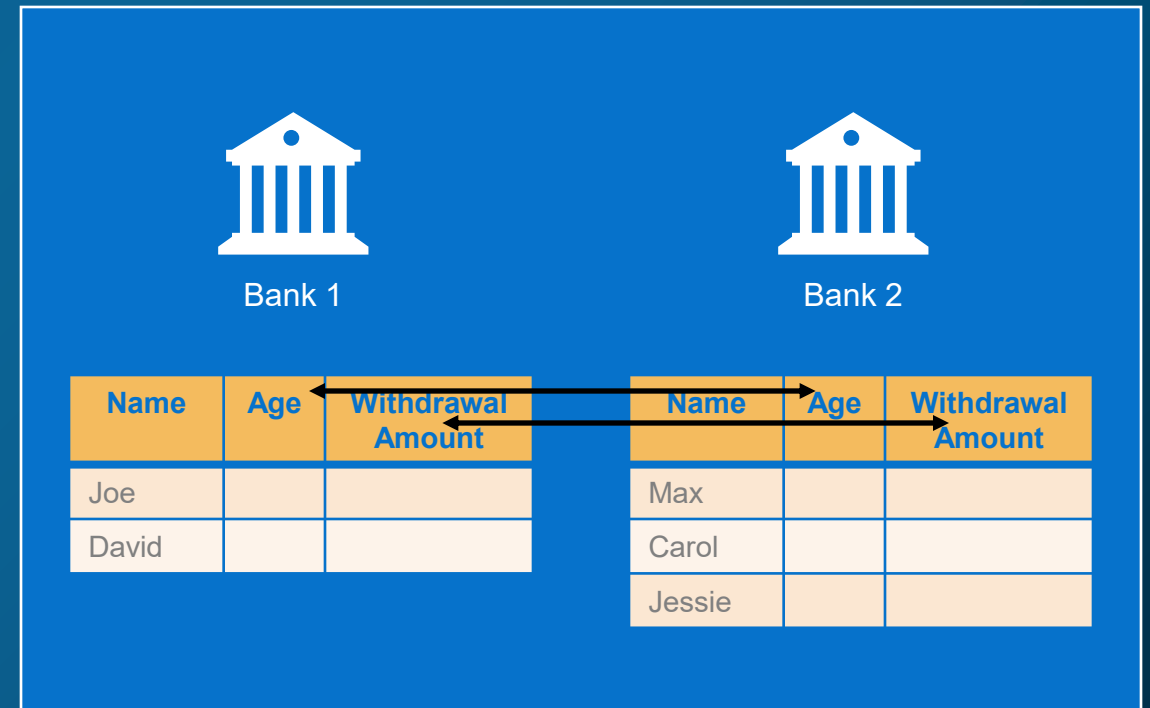
Types of Federated Learning

Horizontal Federated Learning (HFL)




- Devices or organizations have similar feature spaces (e.g., same types of data) but different samples
 - Example: Banks in different regions collaborate on a fraud detection model using transaction data with the same features

Modality	Example	Key Feature
Text	Predictive text on smartphones	Same features, different samples
Audio	Speech recognition across smart speakers	Common audio features, unique users
Images	Facial recognition in security systems	Same features, different individuals
Sensor Data	Smart meters predicting energy usage	Similar features, household-specific data

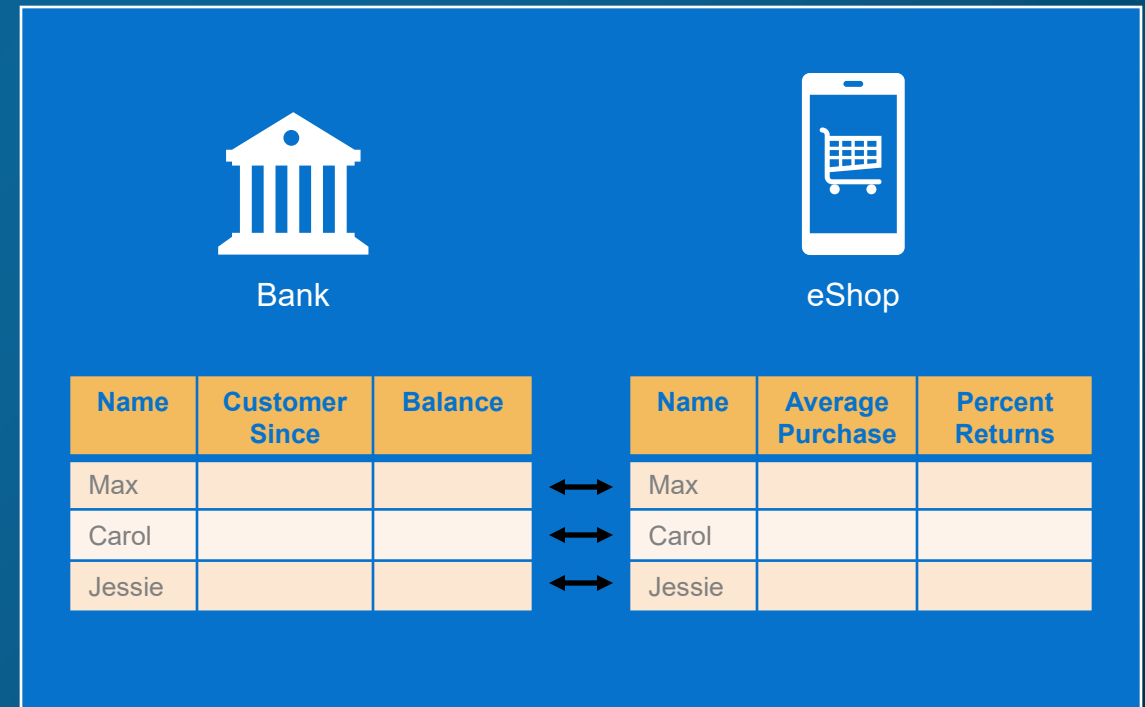


Vertical Federated Learning (VFL)




- Devices or organizations have different feature spaces but share overlapping samples (e.g., same users)
 - Example: A bank and an e-commerce platform collaborate to predict user creditworthiness using complementary data

Modality	Example	Key Feature
Text	Bank + e-commerce for loan prediction	Different features, same users
Audio	Mental health analysis using stress-level recordings and transcripts	Different features, overlapping samples
Images	Hospital + pharmacy collaboration (medical images + prescriptions)	Different data types, shared patients
Sensor Data	Traffic cameras + roadside sensors for congestion prediction	Different sensors, same vehicles

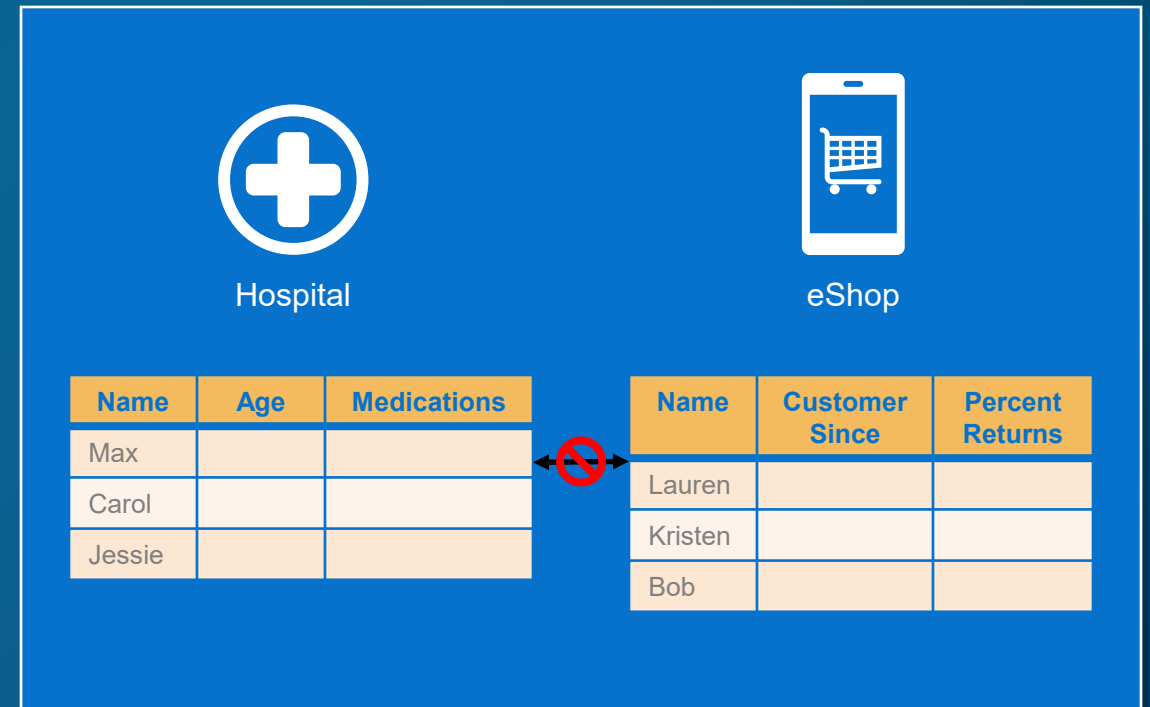


Federated Transfer Learning (FTL)



- Devices or organizations have different feature spaces and few (or no) overlapping samples, leveraging transfer learning to share knowledge
 - Example: A hospital and a retail store collaborate to improve prediction models by transferring learned patterns across domains

Modality	Example	Key Feature
Text	News agency + retail platform for cross-domain sentiment analysis	Knowledge transfer, no overlapping users
Audio	Cross-language speech recognition	Transfer learning for less-resourced data
Images	Medical X-rays + wildlife photos for injury detection	Transferable patterns, different domains
Sensor Data	Weather + soil moisture sensors for crop yield prediction	Different sensors, no direct overlap

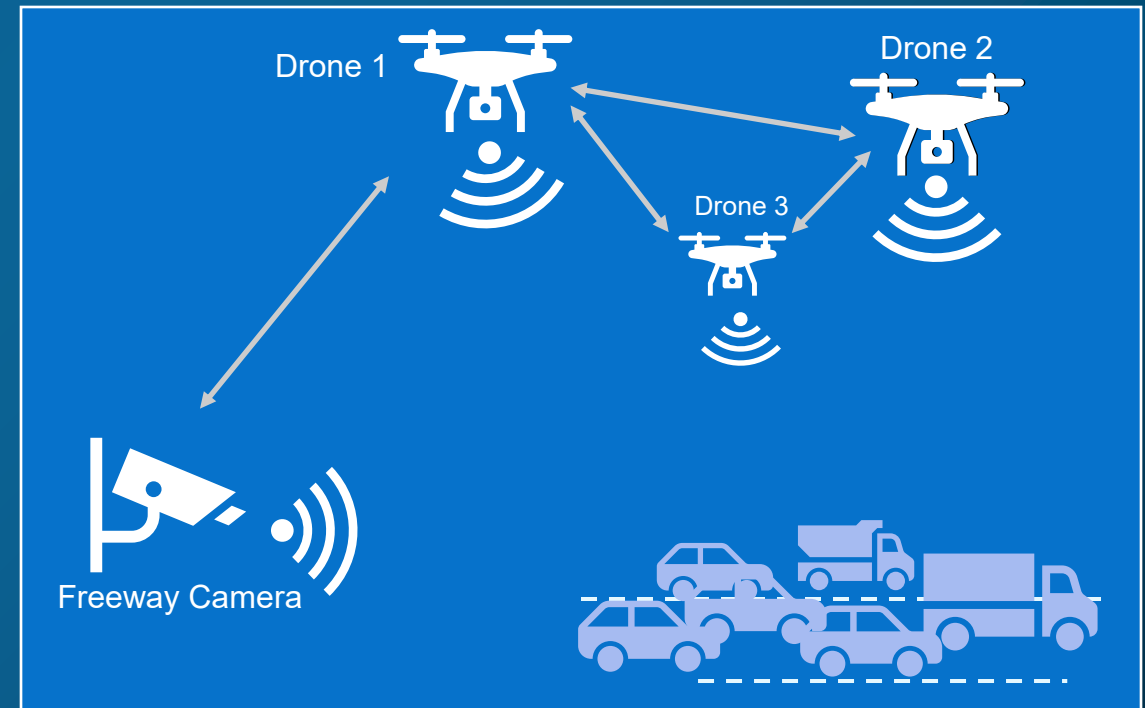


Peer-to-peer Federated Learning (P2P-FL)

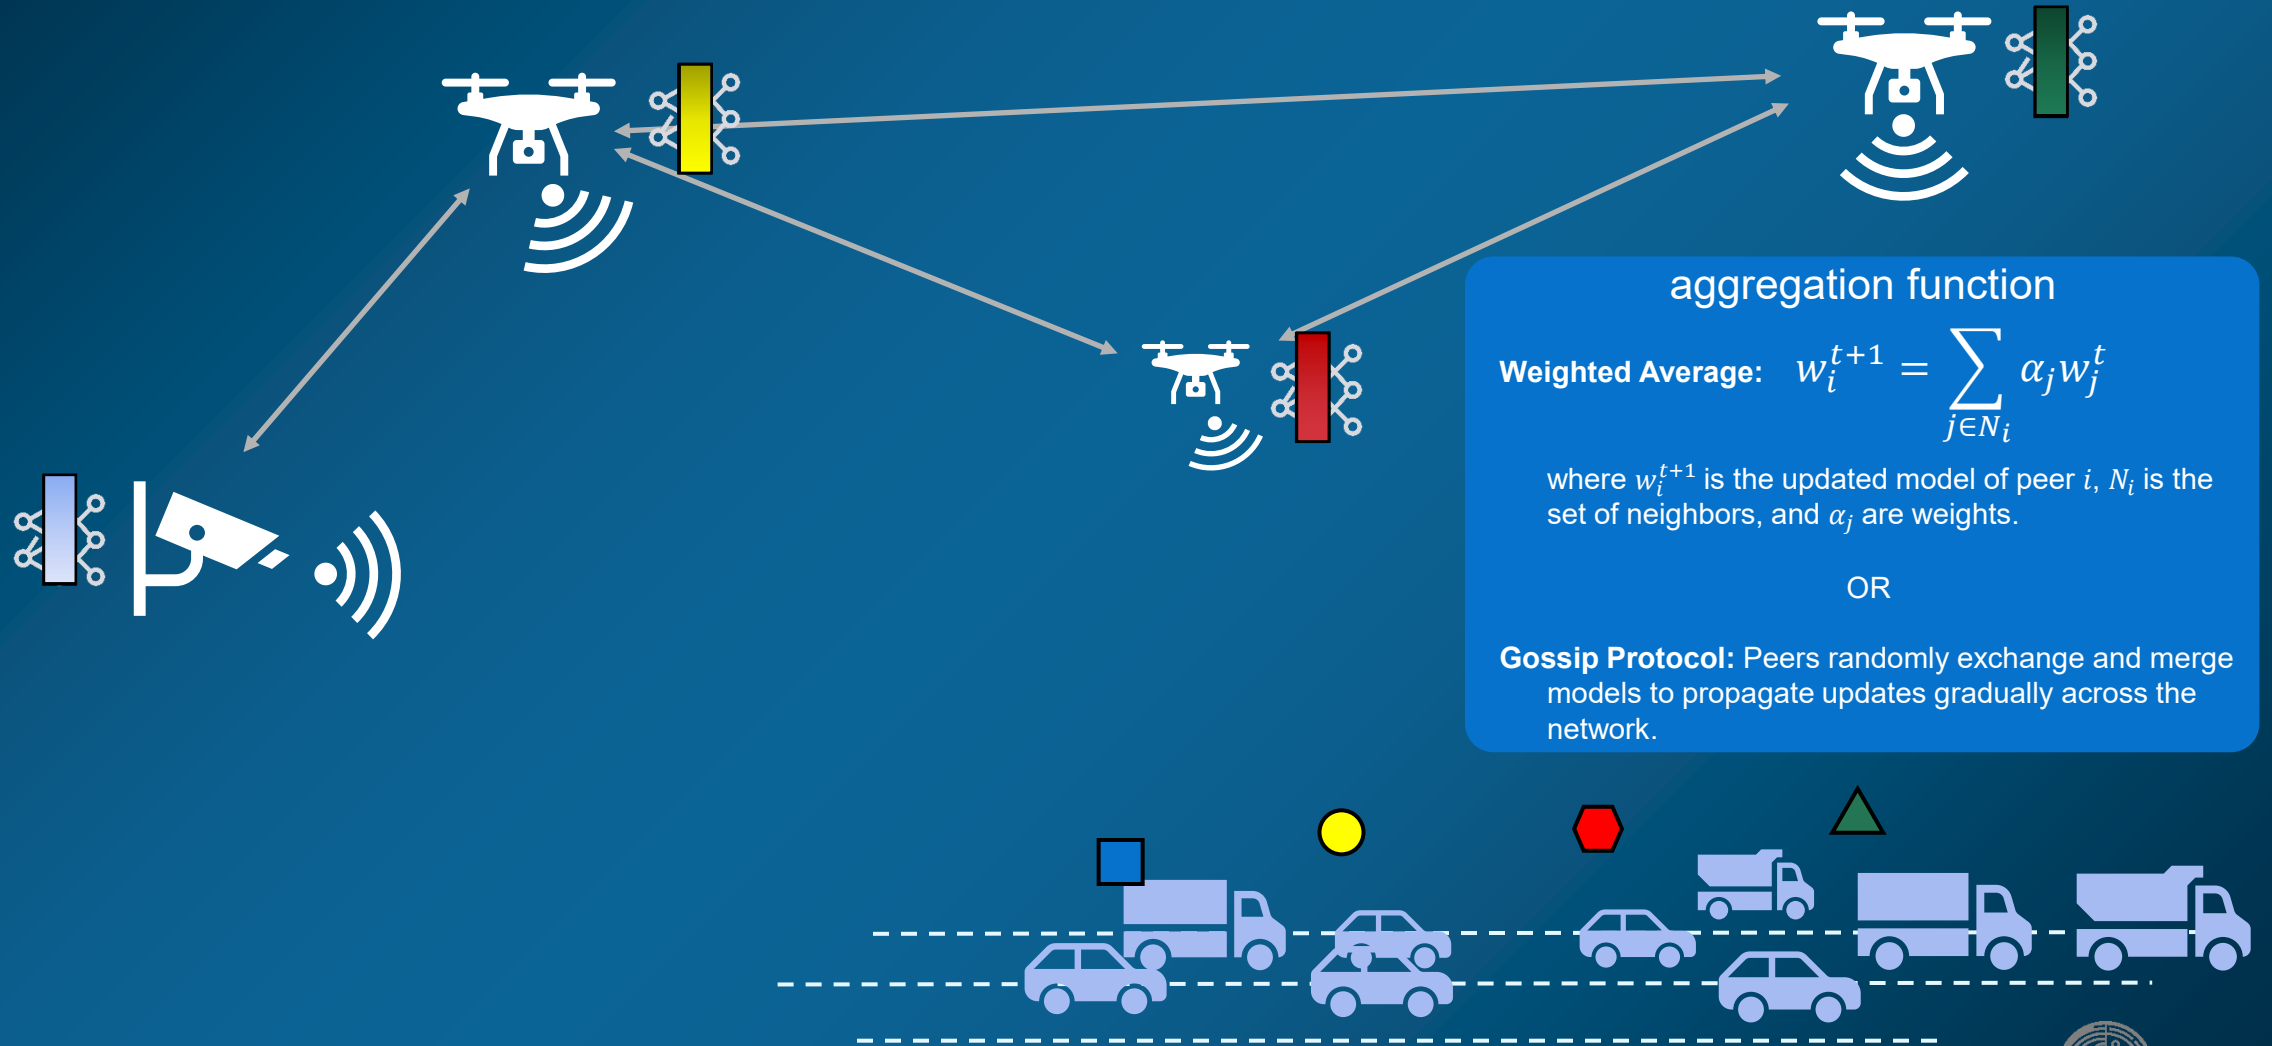


- Autonomous devices (or nodes) communicate directly with each other to collaboratively train a global model without a central server
 - Example: Drones and a freeway camera collaborate without the need for a central server

Modality	Example	Key Feature
Text	Decentralized spam detection using email clients	Peer-to-peer collaborative training
Audio	Wake-word detection training across voice assistants	Decentralized audio recognition
Images	Traffic monitoring with smart cameras	P2P collaboration for image analysis
Sensor Data	Air quality sensors collaboratively predicting pollution levels	Sensor-based peer-to-peer learning



Peer-to-peer Federated Learning






Real-Life Examples of Federated Learning

Application: Virtual Assistants

- Voice Assistants
 - e.g., Siri, Alexa, Google Assistant
- Smart Keyboards
 - AI-enhanced input tools designed to improve text input, offering features like predictive typing, autocorrect, and voice typing
 - e.g., Gboard
- Productivity Tools:
 - e.g., scheduling, note-taking, or quick searches



Application: Wildfire Detection with UAVs



- Drones analyze local fire patterns
- Privacy: No raw images shared
- Collaboration: Improves detection across the fleet

Application: Environmental Monitoring

- Aggregate data from distributed sensors to identify patterns
- Applications include pollution mapping and climate-change monitoring




Application: Personalized Healthcare

- Hospitals collaboratively train models to detect cancer or rare diseases without sharing sensitive patient data



Application: Smart Cities

- Sensors collect local traffic data to optimize traffic signals
- FL enables cities to collaborate across regions for better models
- Privacy: Individual movements aren't exposed



Application: Financial Services



- Banks train models collaboratively to detect fraudulent transactions
- Sensitive financial data stays within each institution



Swarm Intelligence

Swarm Intelligence

- Definition

- A collective behavior of decentralized, self-organized systems that solve problems through simple interactions

- Purpose

- To achieve **complex problem-solving** or **optimization** by mimicking the behavior of natural swarms

Aspect	Federated Learning (FL)	Swarm Intelligence (SI)
Primary Goal	Train machine learning models collaboratively	Solve optimization problems (e.g., routing, resources)
Data Handling	Data remains local; only model updates are shared	Not focused on data privacy; agents share local information
Communication	Can be via server or fully decentralized	Local
Inspiration	Rooted in machine learning and privacy preservation.	Inspired by biological systems (e.g., ants, bees)
Emergent Behavior	Not a focus; behavior is pre-defined by algorithms	Emergent behavior arises from local interactions



Ant bridge from their live bodies
(e)Science News

- Minimizing travel time (path optimization)
- Resource allocation
- Dynamic adaptation
- Energy efficiency
- Emergent behavior

Combining Swarm Intelligence and Federated Learning

- SI optimizes traffic
- FL learns patterns from local data.

Smart Cities



- SI coordinates drones
- FL trains models for victim detection
-

Disaster Response



- SI collects data
- FL improves predictions

Environmental Monitoring



- SI manages resources
- FL trains private predictive models

Healthcare



- SI optimizes tasks
- FL improves decisions from local conditions

Agriculture



- SI optimizes routes
- FL refines forecasts

Logistics





Technical Challenges

Technical Challenges

Challenge	Description	Mitigation
Communication Efficiency	Potentially high communication overhead on server or P2P participants	Optimize communication protocols and reduce the frequency or size updates: <ul style="list-style-type: none">• Compress model updates• Increase local training epochs• Use partial client participation.
Privacy and Security	There is still a risk of leaking sensitive information through transmitted model updates or gradients	<ul style="list-style-type: none">• Encrypt updates using secure aggregation• Detect and filter malicious updates
System and Device Heterogeneity	Variability in capabilities can affect training efficiency and model performance, especially when some devices may drop out or contribute unreliable updates	<ul style="list-style-type: none">• Use lightweight or pruned models• Dynamically adjust workloads by device• Prioritize updates from reliable devices
Model Convergence and Robustness	Asynchronous updates, varying participation rates, and the impact of heterogeneous data may affect reliability in convergence of the global model. Risk of biased models that perform poorly on underrepresented data segments or in dynamic environments	<ul style="list-style-type: none">• Use robust aggregators (e.g., Krum)• Adjust learning rates dynamically• Handle asynchronous updates with care


Risk Summary

PROBABILITY		Communication Efficiency	Statistical Heterogeneity
		System & Device Heterogeneity	Privacy & Security
			Model Convergence and Robustness
	IMPACT		

Statistical Heterogeneity (Non-IID Data)

- Scenario:
 - Devices typically hold non-IID data

What is non-IID



- **Independent:** data samples are not affected by one others
 - Ex: Each temperature sensor in a geographical area is not affected by other temperature sensors
- **Identically Distributed:** data samples are drawn from the same probability distribution
 - Ex: All sensors measure temperature under similar conditions, so their readings follow the same distribution

Statistical Heterogeneity (Non-IID Data)

- Widely recognized as the most critical challenge in federated learning
 - Causes biased updates
 - Slower convergence
 - Poor generalization
 - Makes it difficult to aggregate models effectively
- Currently a focus of research




But, Art, dude, you
said **FEDERATED
TRANSFER
LEARNING** deals
with this stuff?



It can only
help so much.
*Stercus intrat,
stercus exit.*

A cartoon illustration of Homer Simpson, a yellow-skinned man with a large nose and a worried expression, wearing a white polo shirt and blue trousers. He is standing on a grey surface.

What can we
do?

A cartoon illustration of Marge Simpson, a yellow-skinned woman with a large, pointed yellow hairdo, wearing a red dress and a pearl necklace. She is pointing her right index finger towards Homer Simpson.

You can personalize models
for each client, cluster clients
with similar data, and apply
regularization (e.g., FedProx)



Image generated by Google Imagen.
I'm not violating copyright, am I?



Conclusion

Summary

Paradigm	Data Dependency	Learning Objective	Key Applications
Supervised Learning	Requires fully labeled data	Predict outputs from labeled inputs	Image classification, spam detection, regression tasks
Unsupervised Learning	Works with unlabeled data	Discover hidden patterns or structures	Clustering, dimensionality reduction, anomaly detection
Semi-Supervised Learning	Combines labeled and unlabeled data	Improve performance using limited labeled data	Text classification with limited labels, medical image analysis
Self-Supervised Learning	Uses unlabeled data and generates labels from the data itself	Learn representations for downstream tasks	Pretraining models (e.g., GPT, BERT), masked autoencoders
Reinforcement Learning	Interacts with an environment	Optimize actions to maximize long-term rewards	Robotics, game-playing AI, autonomous vehicles
Federated Learning	Decentralized data across devices	Train models collaboratively without sharing raw data	Personalized recommendations, healthcare diagnostics, edge AI

Key Takeaways

- FL enables collaborative AI without data sharing
- From UAVs to hospitals, from phones to traffic sensors, FL is an invaluable systems architecture for AI
- Still many open challenges: privacy, robustness, optimization, and especially, statistical heterogeneity
- FL combined with SI is very powerful




Demonstration


Demonstration: Flower

<https://colab.research.google.com/github/adap/flower/blob/main/examples/flower-in-30-minutes/tutorial.ipynb>


- Uses MNIST data set:
 - 70,000 28x28 handwritten digits
 - 60,000 for training
 - 10,000 for test



Centralized



Distributed
(first 30 of 100 clients shown)





Question-Answer-Discussion

Unlock the Power of AI for Your Organization

- Gain a competitive edge
- Improve efficiency and automate routine tasks
- Enhance decision-making with data-driven recommendations



Expert AI Services

- Strategic AI
- Market research and analysis
- Grant research
- AI-assisted chatbots
- AI-driven proposal writing
- Predictive analytics and forecasting
- AI Education

Unlock the true potential of AI. While 95% of companies struggle to generate returns from their AI investments, our AI consulting services provide a clear path to success.

We specialize in aligning AI strategies with your business goals, implementing tailored solutions, and optimizing AI performance to deliver measurable results.

Art Villanueva, DEng, ESEP
Principal
art@phronos.com



PHRONOS
PRACTICAL AI

Federated Learning:

A Systems Engineering Approach to Smarter AI

In the first presentation of this evolving series, Dr. Villanueva explored the convergence of Systems Engineering and Artificial Intelligence. In the second, he examined the phenomenon of emergence – how complex system-level behavior arises from simple, local interactions. This third installment delves into federated learning (FL), a distributed machine learning paradigm that applies principles of emergence at scale, enabling intelligent global behavior to arise from decentralized, local interactions.

Federated Learning enables edge devices to collaboratively train shared models without transferring raw data. Through decentralized computation and global aggregation, intelligent behavior emerges from the collective experience of autonomous participants – from smartphones and hospitals to fleets of unmanned aerial vehicles (UAVs). This talk will examine the algorithmic foundations of FL, its systemic challenges (non-IID data, intermittent connectivity, and adversarial threats), and its compelling applications in systems where privacy, resiliency, and autonomy converge.

From collaborative healthcare systems to adaptive environmental monitoring to smart infrastructure and UAV networks, federated learning provides a powerful framework for enabling distributed intelligence across diverse domains. Dr. Villanueva will examine how FL represents not just a technical tool, but a systems architecture for building complex, adaptive, and mission-aware AI at the edge, where privacy, resilience, and autonomy are paramount.

