Follow us
@icicleai

http://icicle.ai

NSF-Funded AI Institute
($20M USD for Five Years)

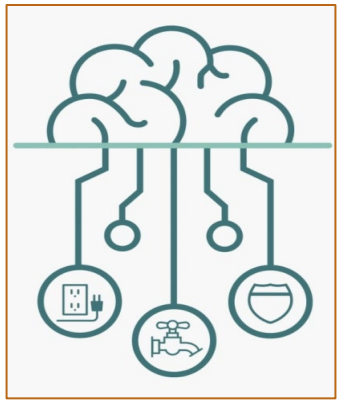

Designing Next-Generation Intelligent Cyber-Infrastructure: ICICLE NSF-AI Institute

PI: DK Panda, OSU
Presenter: Rajiv Ramnath, OSU

Participation (14 Organizations, 46 Investigators, and many Collaborators)

Govt. Agencies & National Labs

Lawrence Livermore National Laboratory

U.S. Fish & Wildlife Service

Brookhaven National Laboratory

Fermilab

International

Max-Planck-Gesellschaft

University of Stuttgart Germany

Industry

CNH Industrial

ADC

WildMe

Microsoft

Farm 2 Facts

The Findings Group, LLC Research & Evaluation

Research Institutes

National Center for Women & Information Technology

MolSSI

SGCI Science Gateways Community Institute

Trusted CI The NSF Cybersecurity Center of Excellence

ECEP

CAHSI Computing Alliance of Hispanic-Serving Institutions

SCEC An NSF+USGS Center

NSF NEON Operated by Battelle

Jetstream

UC San Diego SDSC San Diego Supercomputer Center

TACC Texas Advanced Computing Center

NSF AI Institutes

AIFS AI Institute Food Systems

Hospitals & Universities

Nationwide Children's

University of North Carolina at Chapel Hill

Princeton

The Ohio State University

CWRU

ICICLE Leadership

Panda (PI)
(OSU) [CI]

Chaudhary (Co-PI)
(CWRU) [CI]

Fosler-Lussier (Co-PI)
(OSU) [AI]

Machiraju (Co-PI)
(OSU) [AI]

Plale (Co-PI)
(IU) [BPC, CI]

Eigenmann (UDel)
[CI]

Huber (UC-Davis)
[Smart Foodsheds]

Lange (IC-FOODS)
[Smart Foodsheds]

Majumdar(SDSC)
[CI]

Morales (UW-Madison)
[All Applications]

Ramnath (OSU)
[BPC, CI]

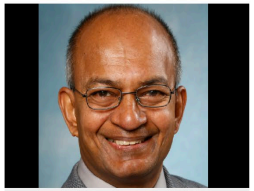

Sadayappan (Utah)
[CI]

Savardekar (MD) (OSU)
[Management]

Stewart (RPI)
[Animal Ecology]

Stubbs (TACC)
[CI]

Tomko (OSC)
[CI]

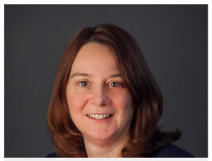

Zhang (Iowa)
[AI, CI]

Emergence of the Computing Continuum

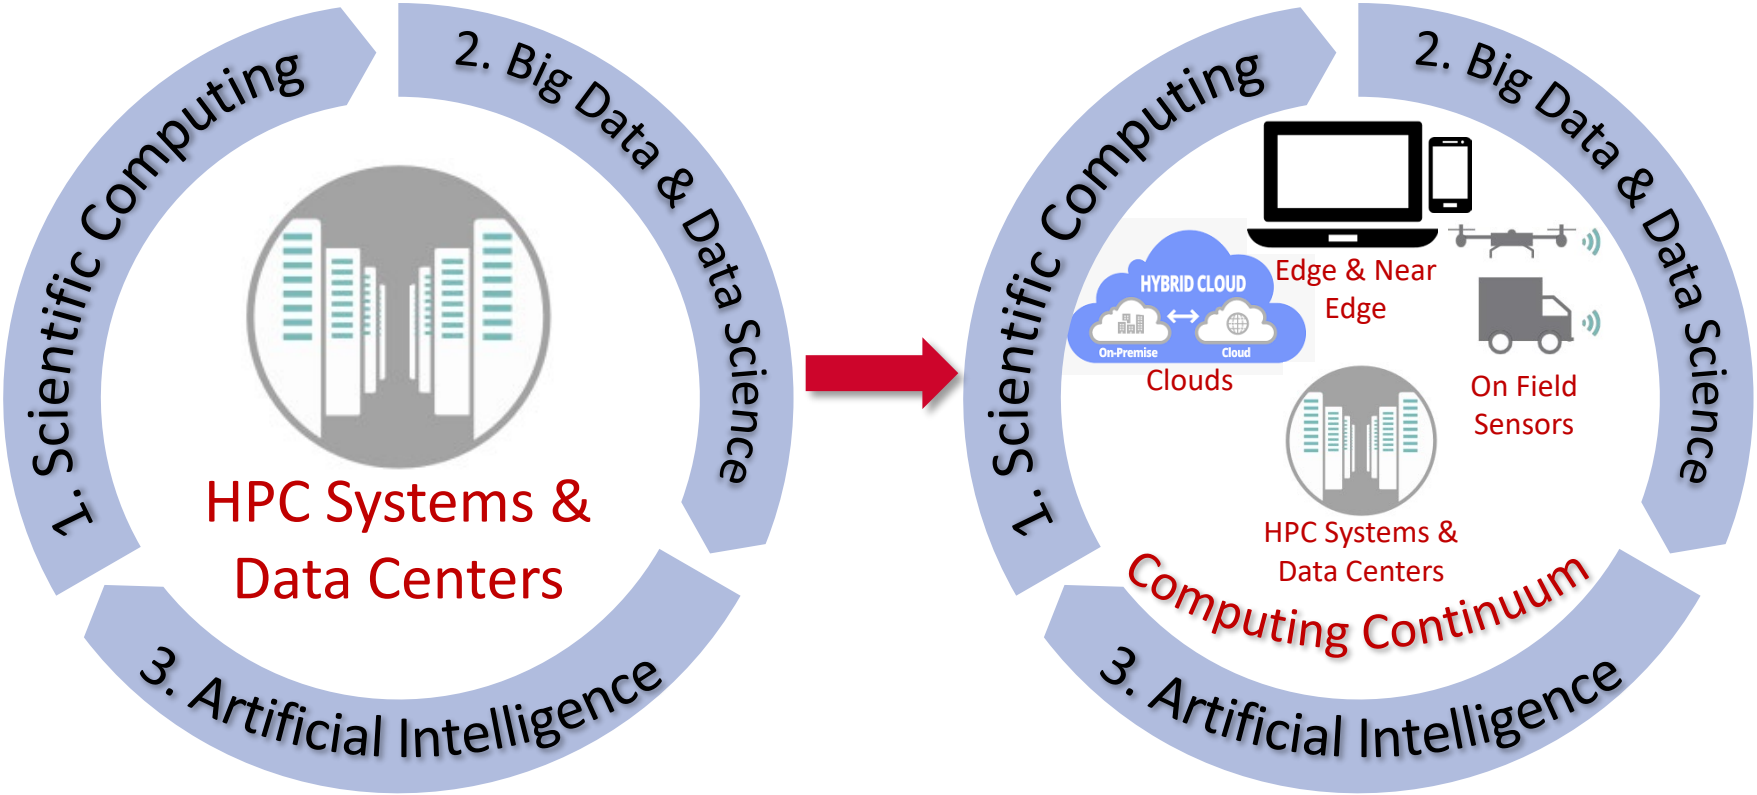

Example Use: AI-Driven Digital Agriculture

https://ccag.tamu.edu/research-project/digital-agriculture/

https://medium.datadriveninvestor.com/artificial-intelligence-in-agriculture-62f71f8f6ae6

Broad Challenge

How to design the next-generation intelligent, plug-and-play cyberinfrastructure for the heterogenous computing continuum for broad use by the long tail of scientists and other end users?

AI at the Flick of a Switch!

ICICLE Organizational Components

Research Thrusts

FOUNDATIONAL SYSTEMS AI

INTELLIGENT CYBERINFRASTRUCTURE

- CI FOR AI
- AI FOR CI-FOR-AI
- SOFTWARE ARCHITECTURE AND DESIGN
- VISUAL ANALYTICS FOR CI AND AI EXPLAINABILITY
- PRIVACY, ACCOUNTABILITY AND DATA INTEGRITY

Use-Inspired Science

CO-DESIGN

Broader impacts backbone network

- BROADENING PARTICIPATION IN COMPUTING
- WORKFORCE DEVELOPMENT
- COLLABORATION AND KNOWLEDGE TRANSFER
- DIVERSITY EQUITY & INCLUSION

Strategic Projects

- SMART CAMERA TRAPS & TRADEOFFS
- VISUAL ANALYSIS OF FOOD SYSTEMS KNOWLEDGE GRAPHS
- DIGITAL AGRICULTURE USE-INSPIRED RESEARCH
- GROCERY STORE CLOSURE & COMMUNITY HEALTH
- "Hello ICICLE"
- REFERENCE SOFTWARE ARCHITECTURE

Plug-and-Play AI in the Face of Heterogeneities

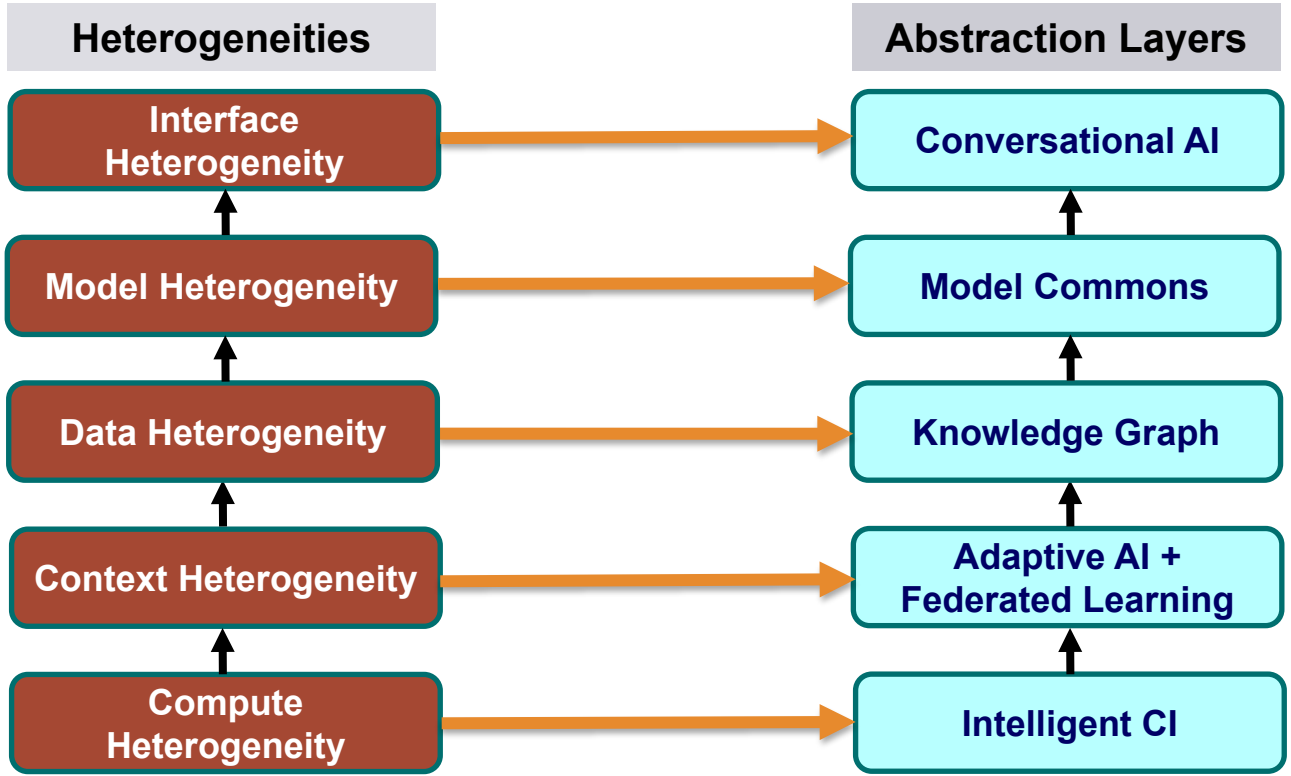

Enables standardization and generalization across use-cases!

Deliverable: The ICICLE Cyberinfrastructure Stack

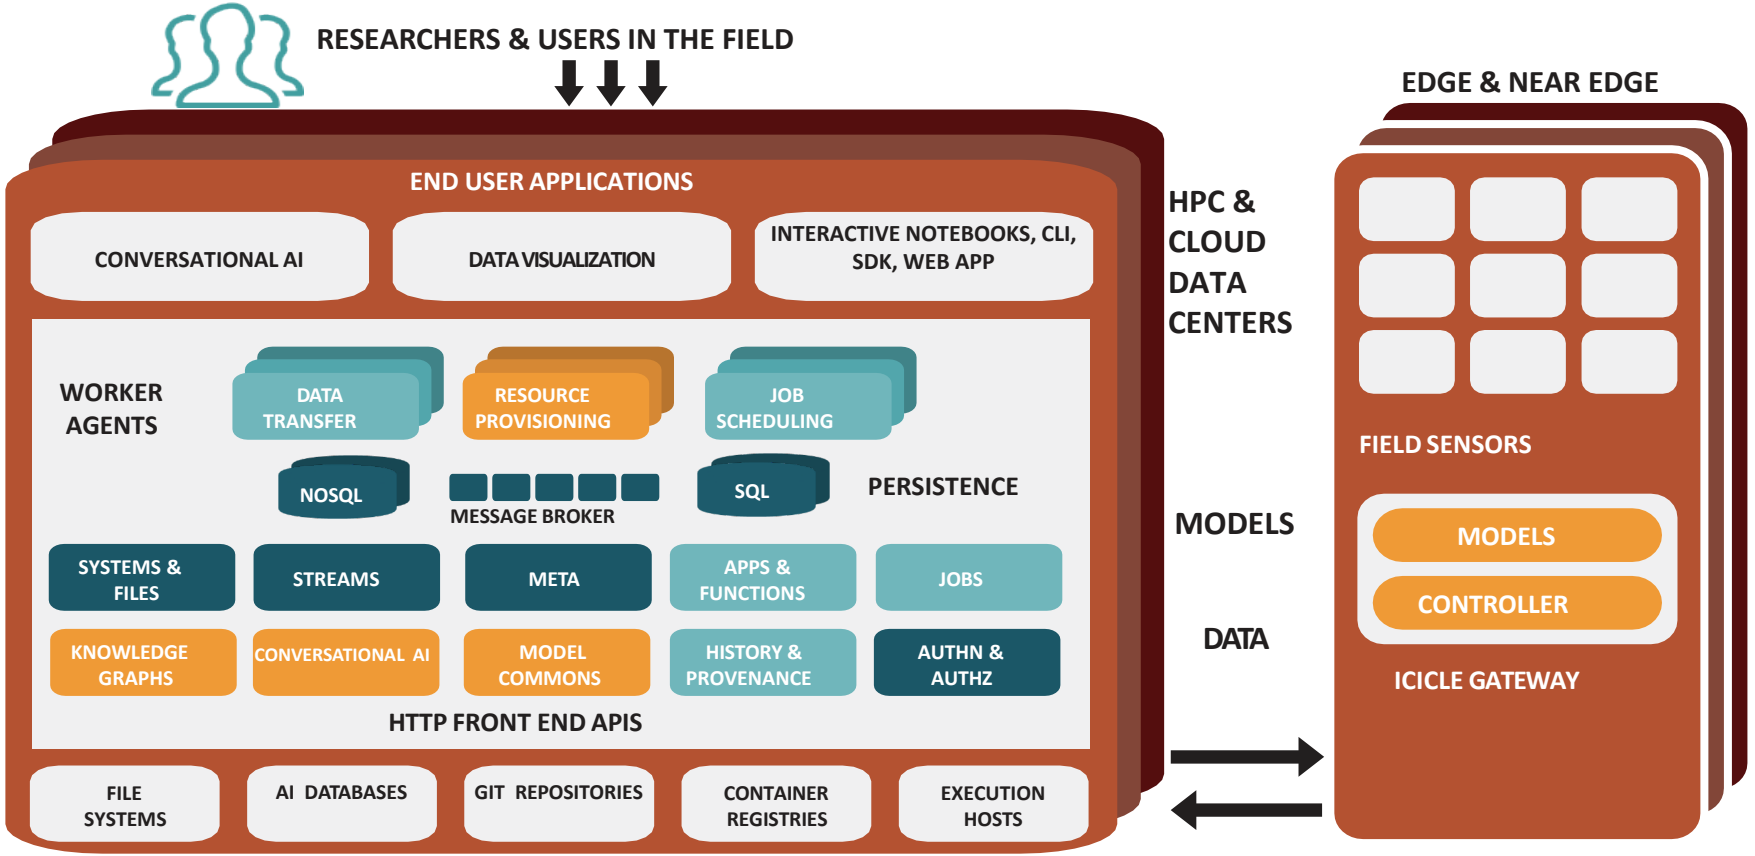

Cloud-to-Edge Cyber Infrastructure

MARbLE is an end-to-end programming, training, and deployment architecture for swarms powered by multi-agent reinforcement learning at the edge

- Greatly simplifies development; Docker & Kubernetes make it easy to port MARbLE workloads between edge sites
- *Symposium on Edge Computing (SEC) 2022*

Offline

1: Take in and Structured Environmental Data

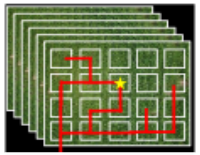

2: Shape reward functions to influence actions towards goals

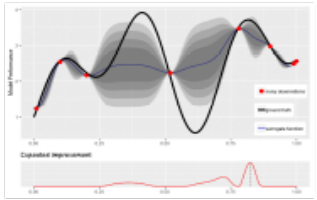

3. Robust Deep RL models For Fully Autonomous UAV

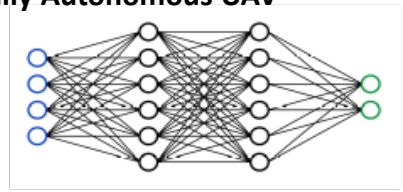

Online

4: Take actions to the real world

5: Manage Cloud-to-Edge Deployments

How to Engage With ICICLE – More Verticals

Smart Foodsheds

Digital Agriculture

Animal Ecology

Health & Medicine

Environment

Communications & Collaboration

Mobility, Machines, & Manufacturing

AI for Social Good

ICICLE: Intelligent CyberInfrastructure with Computational Learning in the Environment

Systems AI Foundational Research for CI

Intelligent Cyber Infrastructure

CI for AI

AI for "CI for AI"

On Field Sensors

Edge & Near Edge

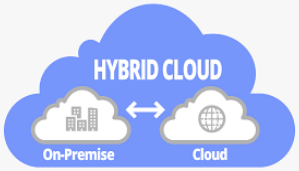

Clouds

HPC Systems & Data Centers

Emerging Computing Continuum

How to Engage With ICICLE: Adopt and Co-Develop the ICICLE CI Stack

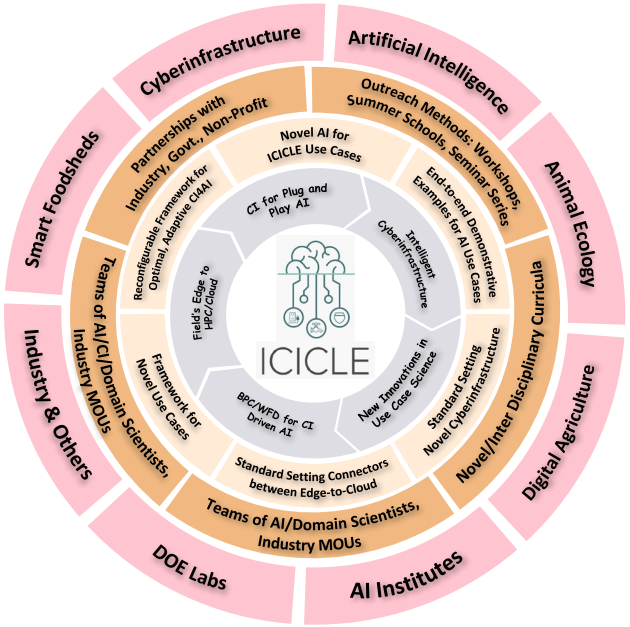

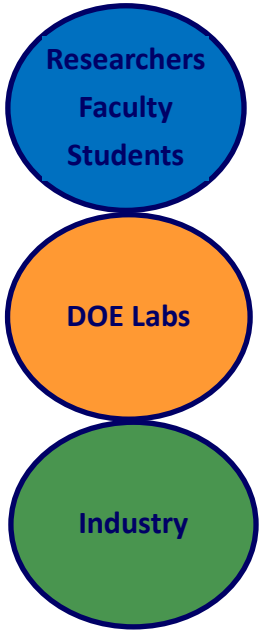

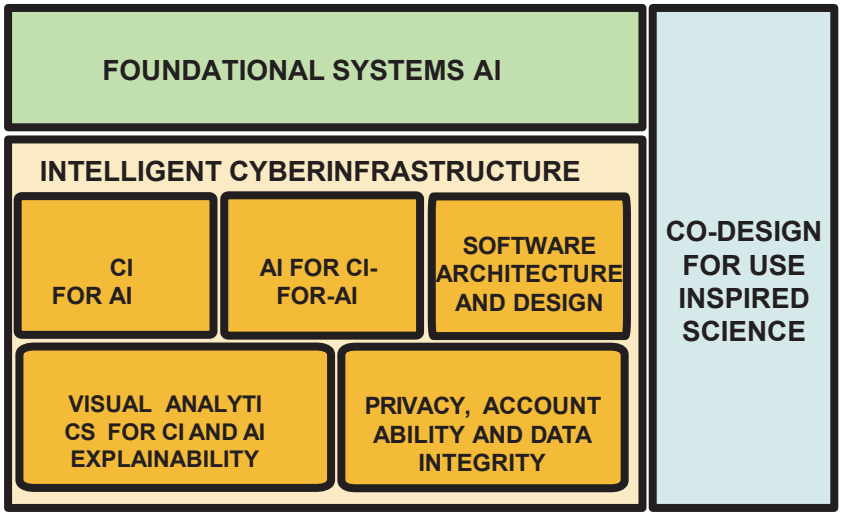

Co-develop & Adopt ICICLE developed CI!

More Details under `ICICLE Engagement' (https://icicle.osu.edu/icicle-engagement)

Contact: panda.2@osu.edu

ICICLE DEMOCRATIZING AI

THE OHIO STATE UNIVERSITY

Example Collaboration: ICICLE and the Technology Innovation Hub (TIH) at the Indian Institute of Technology Bombay (IIT-B),

Digital Agriculture



This research collaboration will contribute novel design paradigms for context-adaptive CI and aims to develop next-generation CI for *Digital Agriculture* including AI and machine learning methods targeting 3 core areas

Crop Health Modeling

- Sense crop health and level context to predict crop yield
- Detect stressors and diseases for geographically diverse crops
- Apply remedies with little human intervention via Internet of Things (IoT) and sensor systems

Aerial Crop Scouting

- CI for fully autonomous aerial systems
- Simplify deployment of UAV in real fields to capture common crop health conditions
- Provide accurate maps that yield valuable insights for crop management

Privacy-Preserving Data Exchange

- Create secure, trustworthy, and privacy-preserving platforms that connect farmers and allow them to share information and resources safely.

Building upon the existing ICICLE infrastructure, CI and AI capabilities, researchers will leverage contextual conditions in India for *Digital Agriculture* that differ from the United States to (1) expose brittle CI components, (2) make AI4CI more robust and expansive in the long-term, (3) devise principles that yield context-aware CI

How to Engage With ICICLE: Extend the ICICLE Cyberinfrastructure Reference Architecture

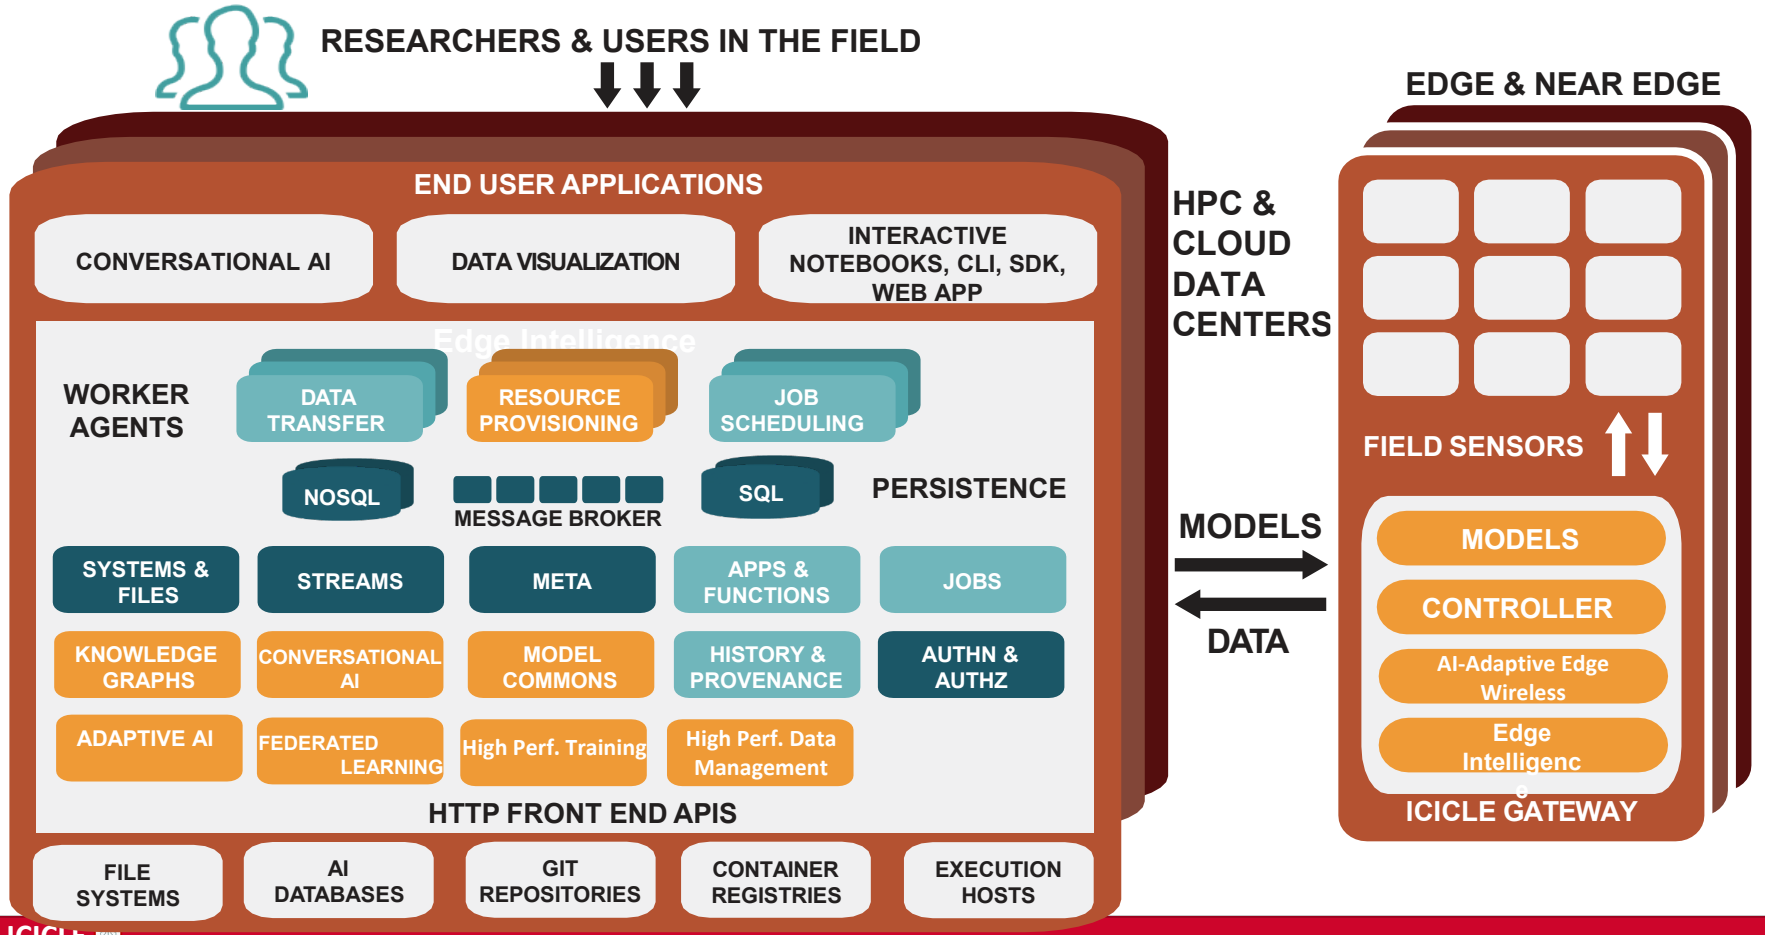

ICICLE Cyberinfrastructure Architecture – What is its intent?

Explaining HOW ICICLE meets its goals

- HOW it implements the use cases (Kroger Closure, Camera Traps, Smart Foodsheds, Precision Agriculture)
- Shows WHERE the developed ICICLE components (from Foundational AI, PADI, AI4CI and CI4AI) exist in the ICICLE CI stack
- Ultimately, should explain HOW ICICLE achieves “Democratizing AI

Kroger Closure Flow

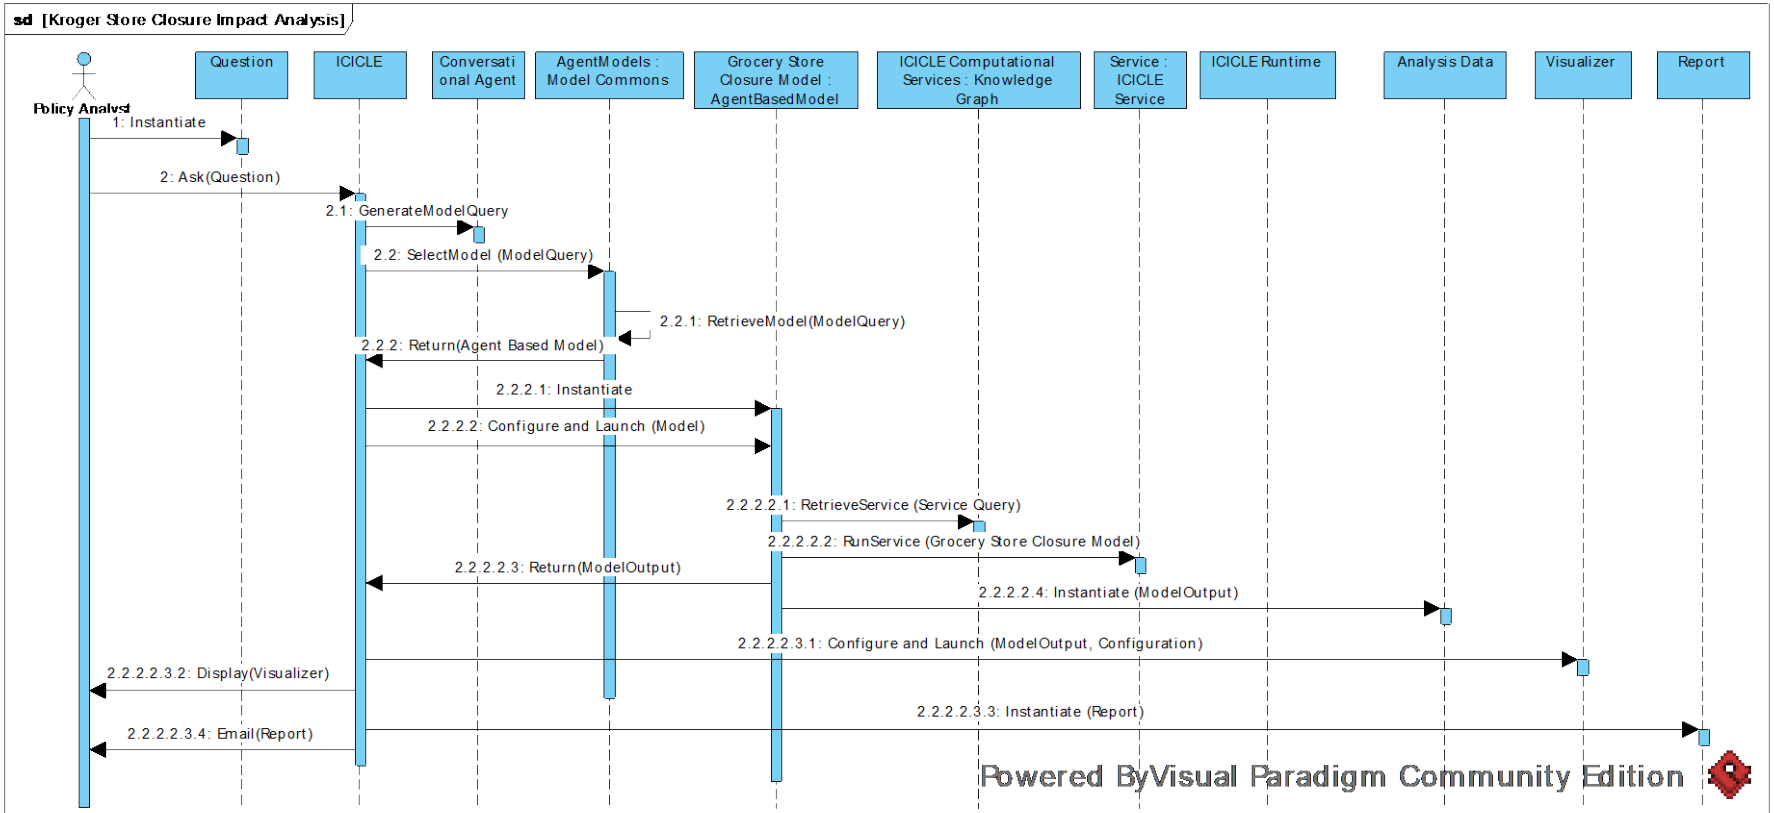

Outcome

Provide an illustration of how the system meets ICICLE's goal of democratizing AI.

Availability of pre-defined templates and procedures to reuse for New projects that can refer and apply what is needed by them.

The ICICLE Cyberinfrastructure

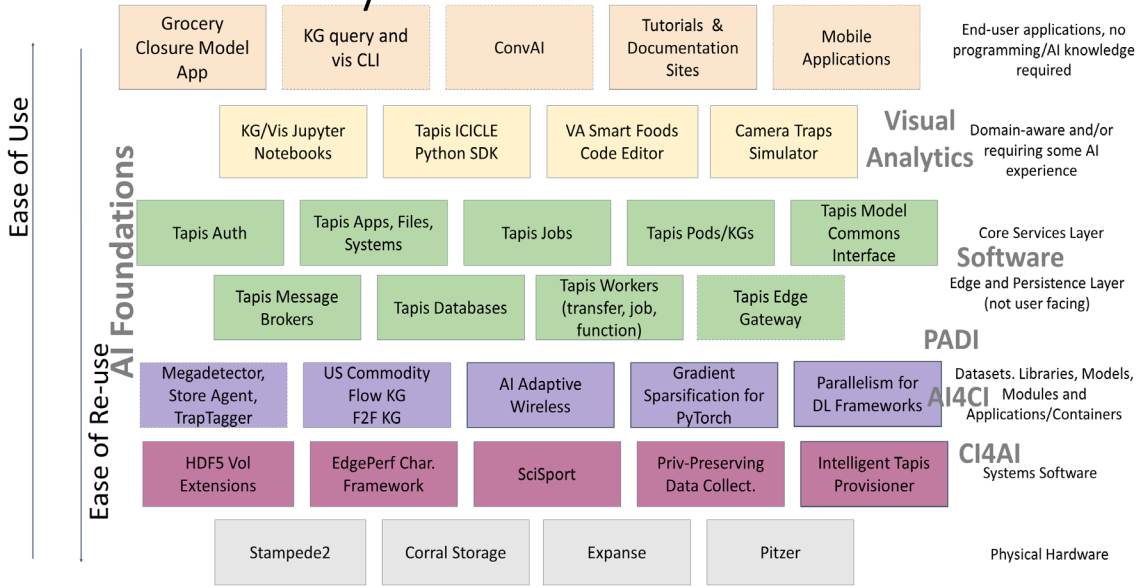

Vision: Global AI leadership

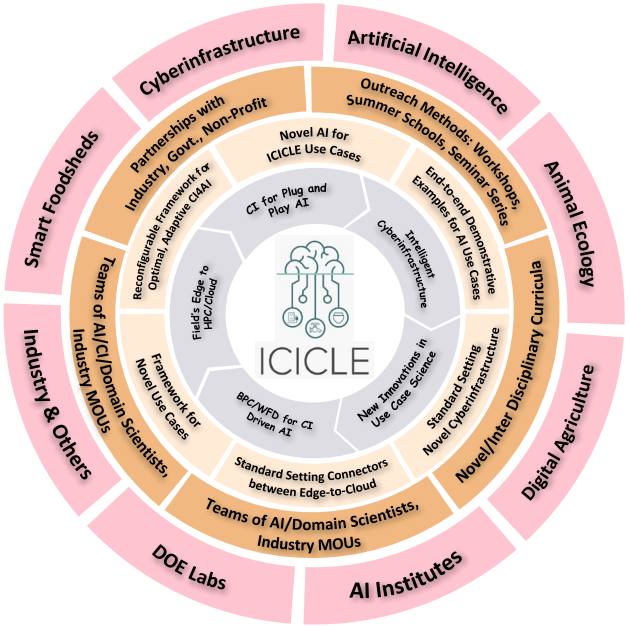

- A National CI Ecosystem
- Integrative and Interoperable
- Leverage our existing recognized capabilities
  - AI Institutes, Centers of Excellence, Large Facilities
  - SAGE
- Collaborative
- Inclusive
- Sustained

ICICLE Research Projects

KG Formalism: Multimodal KG for Camera Traps

- **Use-inspired Goal**: Develop new KG formalisms to represent and reason with multimodal data
- **Progress**: Co-design formalism with Animal Ecology, identified and analyzed public data sources, investigated KG tooling and decided on using Neo4j
- **Applications**: Semantic search, visualization, multimodal reasoning

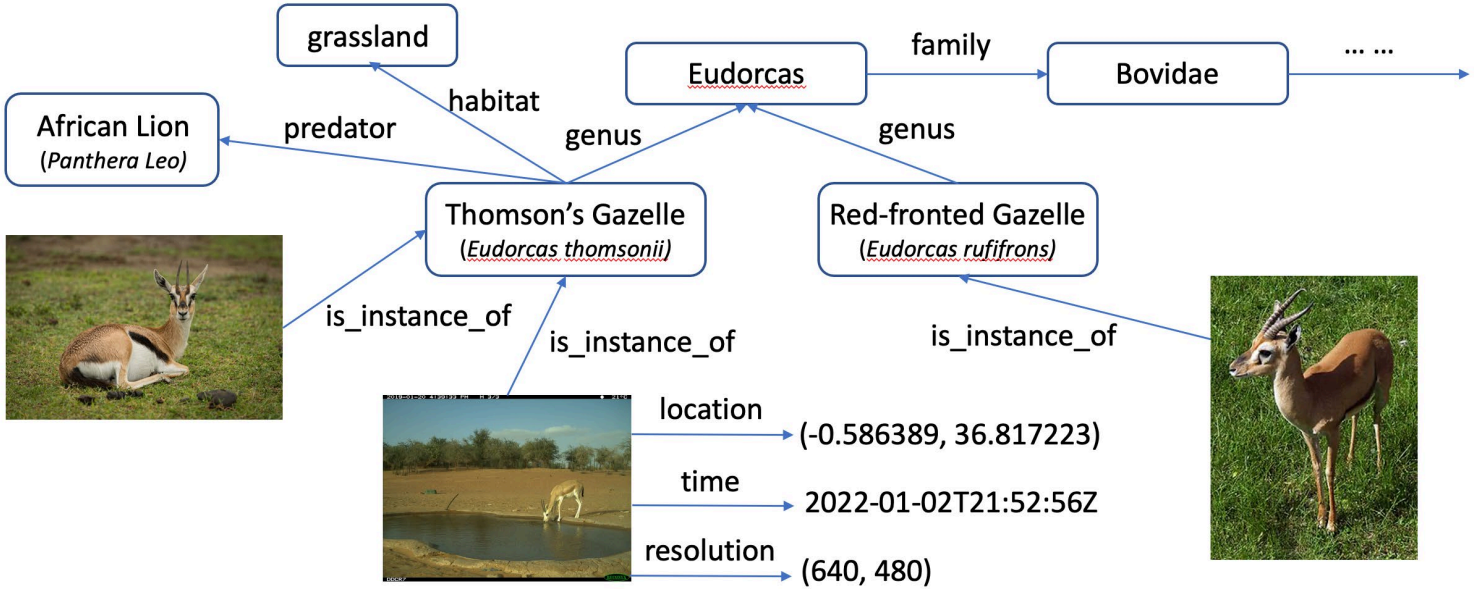

Research in Y1

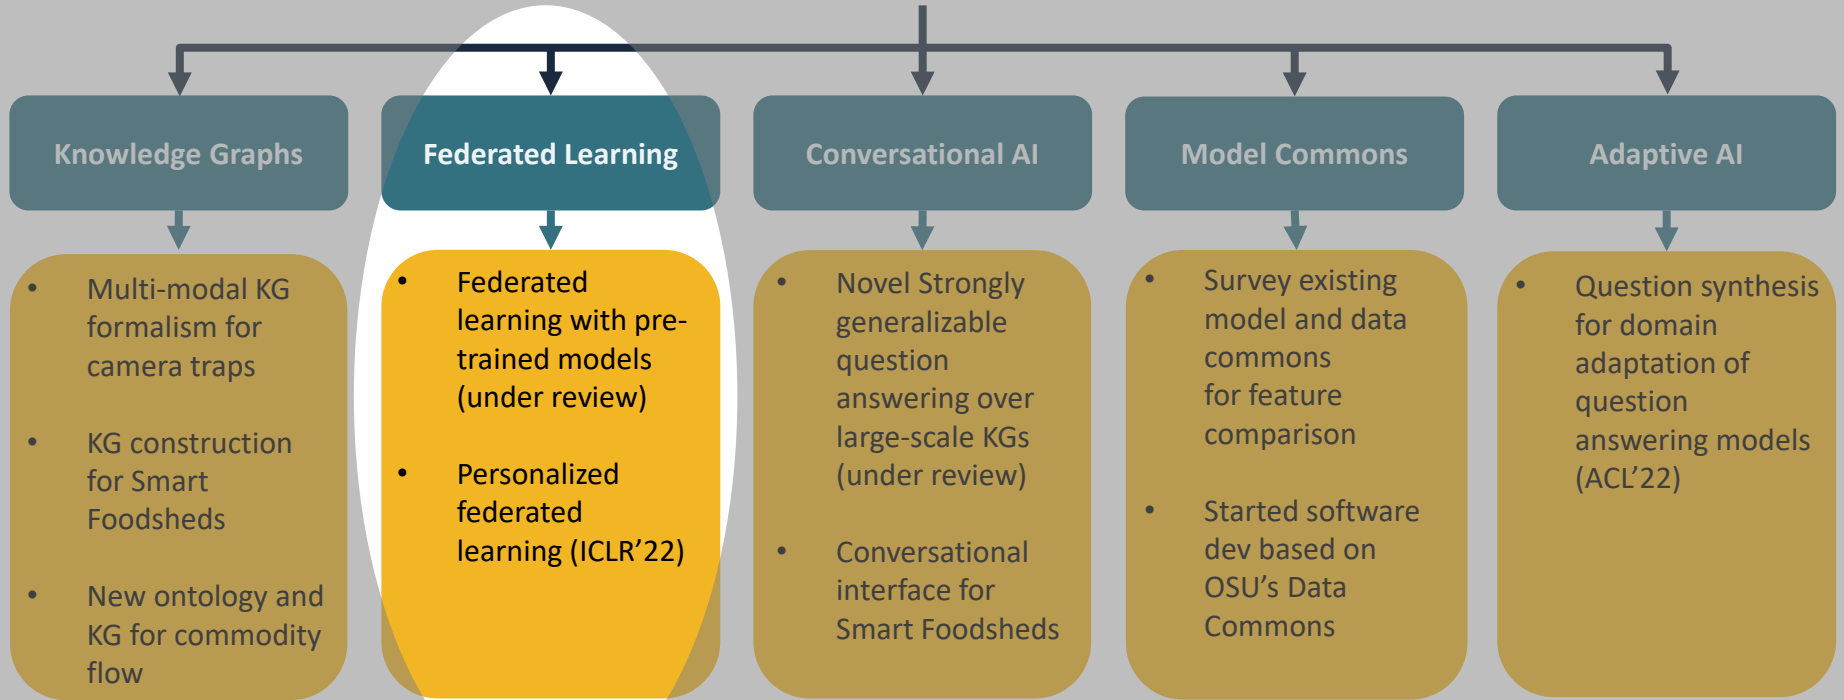

Federated Learning

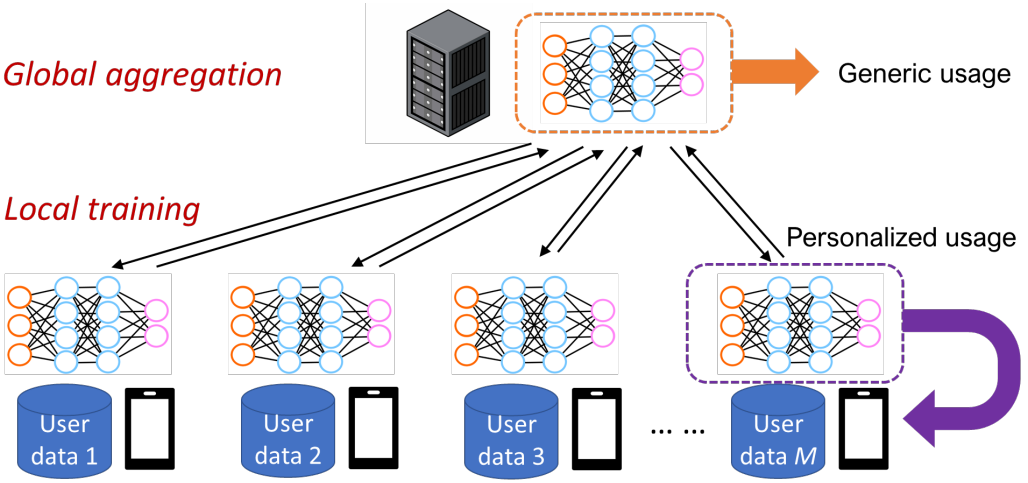

- Robust to users' heterogeneity
- Applicable to state-of-the-art AI models
- Context-aware and user-aware

	Existing literature	**ICICLE scenarios**
Model pre-training	No	**Yes (model commons)**
Server data	No	**Yes (data commons)**
Applicability	Generic OR Personalized	**Generic (updated model common)** **AND** **Personalized (individual users)**

Research in Y1

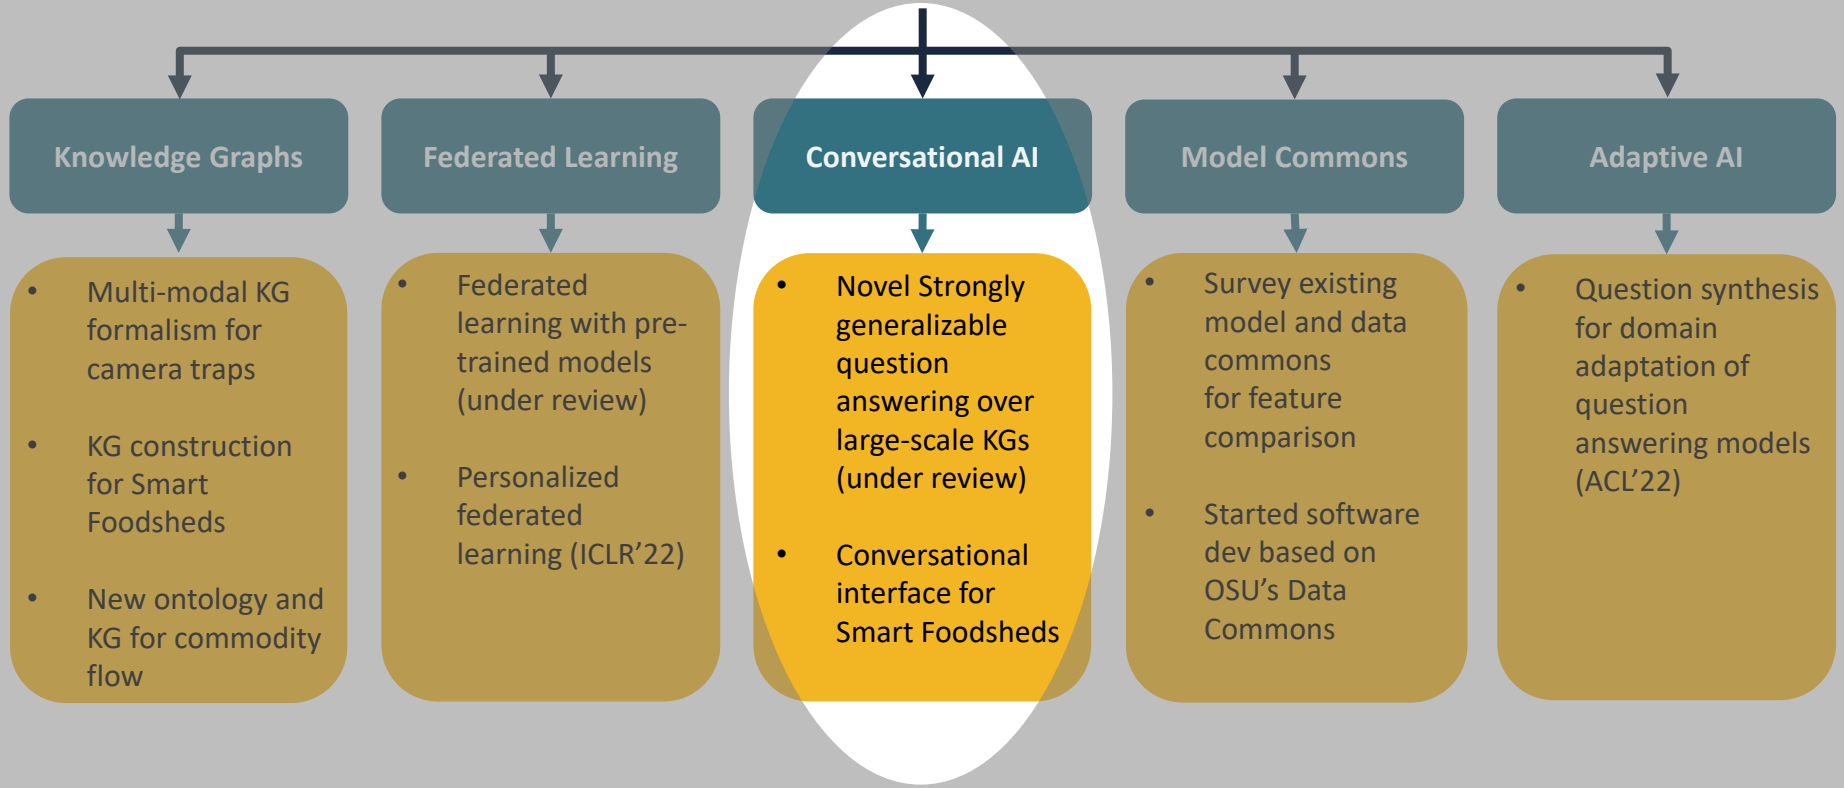

Conversational AI (for Knowledge Graphs)

- Novel method for strongly generalizable and efficient question answering on large-scale KGs (**45 million entities** and **3 billion facts**)
- **One order of magnitude faster** than existing methods
- Achieve **73.7%** F1 score and can work on **entirely new domains** without any training data

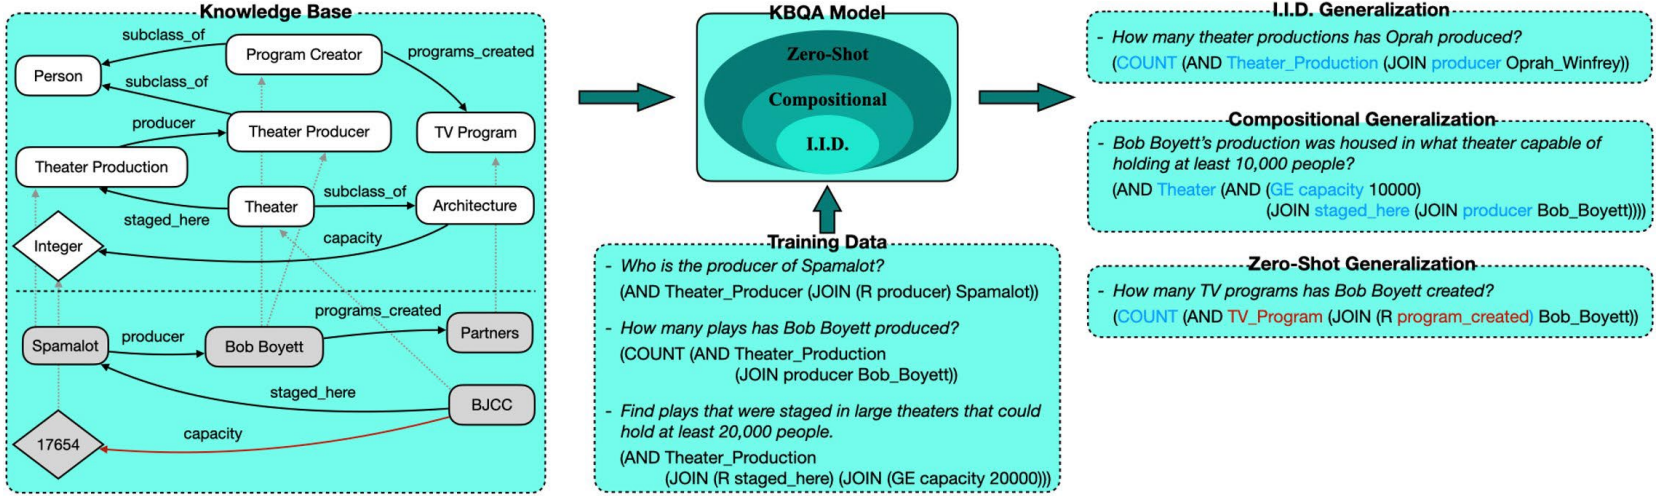

ArcaneQA: Dynamic Program Induction and Contextualized Encoding for Knowledge Base Question Answering
Yu Gu and Yu Su. arXiv:2204.08109, 2022

Outline

- Brief Overview of the ICICLE Institute
- Organizational Infrastructure and Management
- Strategic and Implementation Plan
- Overall Research Directions
- **Selected Accomplishments (so far)**
  - AI Foundations
  - **CI4AI**
  - Digital Agriculture
  - Smart Foodshed
- Conclusion

Research Plan: CI4AI

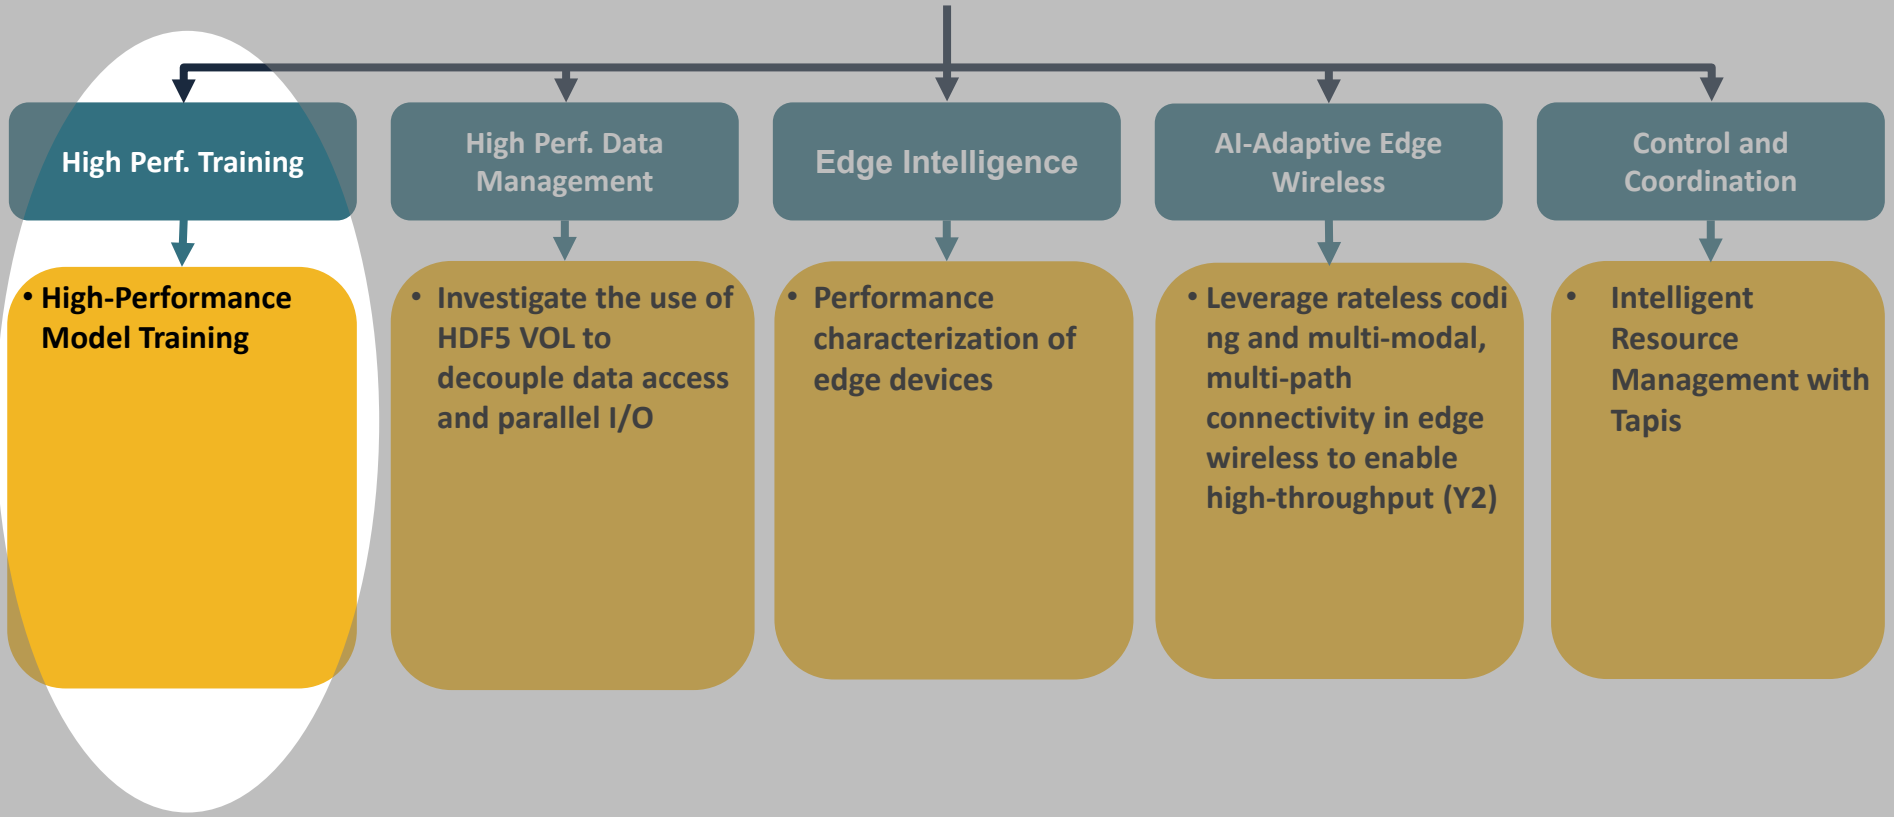

High-Performance Model Training and Support for Edge Intelligence

- CI4AI1: High-Performance Model Training:
  - Reduced model training time for Digital Agriculture use case with high-performance distributed DNN training using parallelism on up to 16 A100 GPUs with 11.6x improvement
  - Utilized efficient model architectures for Digital Agriculture applications including the training and evaluation of different Vision Transformer variants with various parallelism techniques
- CI4AI3: Support for Edge Intelligence:
  - Explored multiple DL frameworks, models, datasets, and benchmarks for inference characterization on edge devices including NVIDIA Jetson, Raspberry Pi 4, etc.
    - MLPerf Edge inference Benchmarks: https://mlcommons.org/en/inference-edge-20/
  - Created a benchmark using Digital Agriculture models and datasets to evaluate inference performance on edge devices:
    - Metrics: Single-stream latency, Query per seconds (QPS)

Outline

- Brief Overview of the ICICLE Institute
- Organizational Infrastructure and Management
- Strategic and Implementation Plan
- Overall Research Directions
- **Selected Accomplishments (so far)**
  - AI Foundations
  - CI4AI
  - **Digital Agriculture**
  - Smart Foodshed
- Conclusion

ICICLE Cyber Infrastructure for Digital Agriculture

What, How, and Why

- Food security is a wicked, salient, and global grand challenge
  - Population growth, environmental sustainability, ecosystems, resource sharing and world peace
- Digital agriculture, the use of computational technologies like sensors, autonomy, robotics, and AI to aid food production, is transforming challenges related to food security
  - Orients food producers toward efficiency and automation

ICICLE research vision: Digital agriculture warrants (inter-) national cyber infrastructure tailored to its unique resource demands

- *What does CI for digital agriculture look like?*
- *How can we scale digital agriculture CI via collaboration?*
- *Why national CI will enable huge leaps forward in digital agriculture?*

Digital Agriculture: Aerial Crop Scouting

- **Crop Scouting:** Develop *heat maps* that describe crop health across a field
  - Inform self-driving tractors and sprayers to reduce field inputs
  - Predict crop yields for harvest and market timing
  - Identify trends across farms, such as introduction of resistant weeds

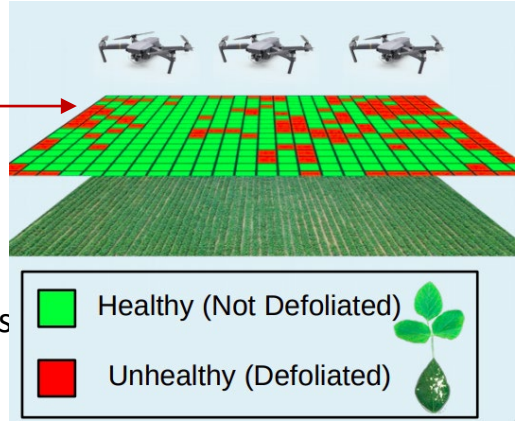

- **Technology:** Unmanned aerial vehicles (UAV) capture high resolution images
  - Flying low (10 ft above ground): 1 pixel -> millimeters

- **Next-gen AI:** Transform high-resolution images to detailed models of crop health
  - AI models mimic the viewpoint of agricultural experts walking through the field

- **Next-gen AI:** Reinforcement learning allows UAV to autonomously decided where to fly and when to land
  - Motivation: Pilots, recharging UAV batteries, transporting equipment == *$$$*
  - Solution: Use a few carefully selected images to extrapolate accurate heat maps

- **Next-gen CI:** Cloud to Edge Middleware
  - High-Performance, Distributed Model Training
  - Edge Resource Management for Swarms of Autonomous UAVs

Digital Agriculture Cyber Infrastructure

Data Scientists in Agriculture

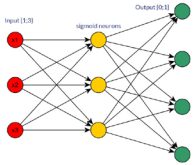

AI & Systems Researchers

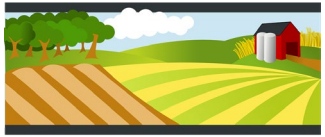

Farmers in the field

ICICLE
DEMOCRATIZING AI

Digital Agriculture Cloud Services

- Digital Ag Front-End
- Data Commons
- Model Commons
- Knowledge Graph

Edge Middleware

ICICLE Enables Novel Crop Health Models

Data Scientists in Agriculture

ICICLE provides CI for storing and labeling large agricultural data sets, making it easy for experts to encode crop-health models for everyone to use

ICICLE also provides Model Commons which stores and profiles AI models (in terms of accuracy, memory, training cost, etc.)

DefoNet (CompAg, 2022) is a custom neural network architecture designed to detect severe leaf defoliation despite overlapping leaves and shadows

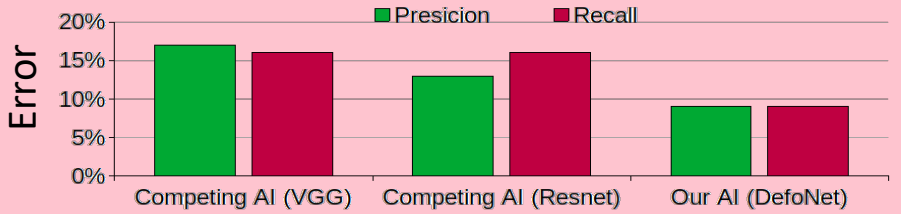

Digital A

Data Commons

Model Commons

Knowledge Graph

High-Performance Model Training in ICICLE

Next-Gen CI

- **High Performance, Distributed Model Training for Ag Data**

Exploit high-performance distributed/parallel training algorithms for ag datasets and models (e.g., DefoNet)

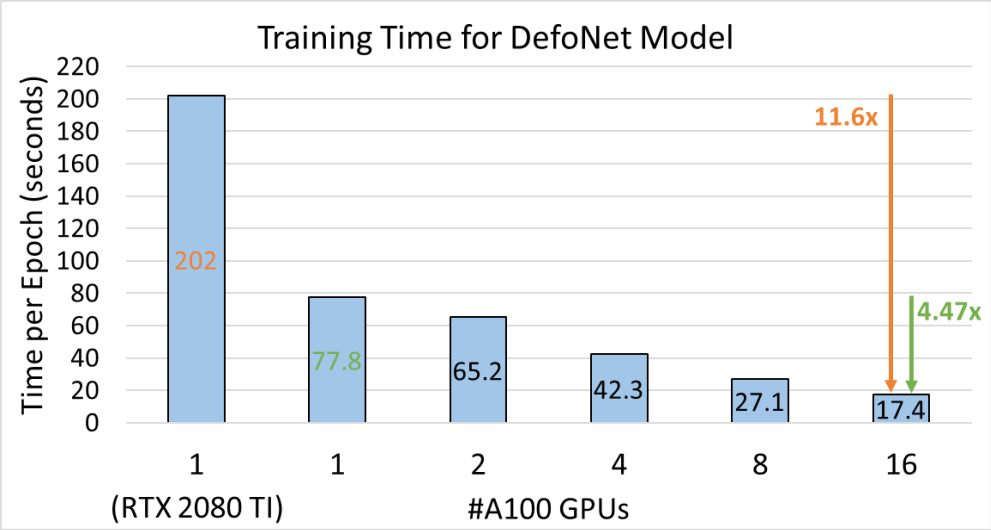

High-Performance Model Training in ICICLE

Data Scientists in Agriculture

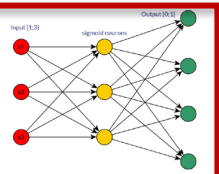

AI & Systems Researchers

Digital agriculture has created a groundswell in data.

Which AI models should be used for any given data set? **Which AI models are cheapest to deploy?**

Are human concerns in agricultural data unlike other data processing? **Is privacy a showstopper?**

How should the CI change to geographically diverse contexts, e.g., corn in US versus onion in India versus rice in Taiwan

ICICLE
DEMOCRATIZING AI

Digital Agriculture Cloud Services

- Digital Ag Front-End
- Data Commons
- Model Commons
- Knowledge Graph

Summary

- ICICLE Goal: Democratizing AI and advanced CI
- ICICLE is complementary to SAGE
- ICICLE is a driver of NSF's vision of an accessible National CI Ecosystem
- Outcomes from this workshop:
  - Access to new capabilities especially on the edge
  - Progress towards "plug and play" across collaborating projects
  - A CI reference architecture
  - Integrated datasets - domain + infrastructure
  - Integrative Projects

Q&A