MVAPICH
MPI, PGAS and Hybrid MPI+PGAS Library

HiDL
High-Performance Deep Learning

HiBD
High-Performance Big Data

Designing Next-Generation Intelligent CyberInfrastructure: An Overview of the NSF-AI ICICLE Institute

NIRTD MAGIC Seminar Series (June '22)

by

Dhabaleswar K. (DK) Panda

The Ohio State University

E-mail: panda@cse.ohio-state.edu

http://www.cse.ohio-state.edu/~panda

Follow us on

https://twitter.com/mvapich

High-End Computing (HEC) has been evolving over the last three decades with multiple **stages**

Stage 1 (1975 -): Scientific Computing with Supercomputing/High-Performance Computing (HPC)

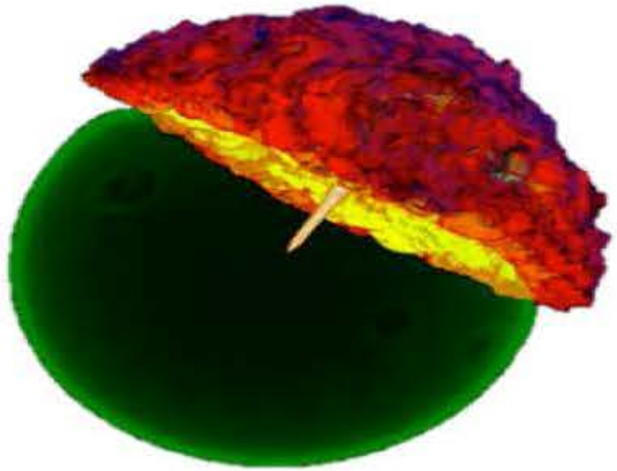

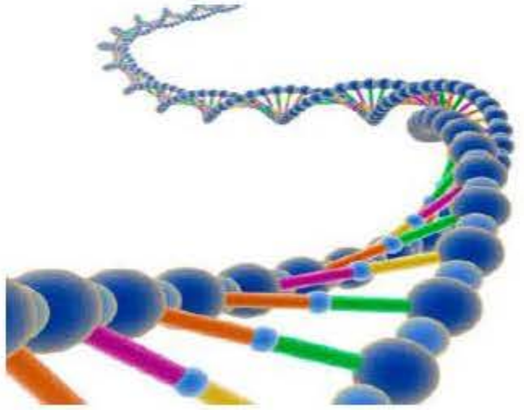

Stage 2 (2000 -): HPC + Big Data Analytics

- Big Data changed the way people understand and harness the power of data, both in the business and research domains
- Big Data and High-Performance Computing (HPC) started converging to meet large scale data processing challenges
- Running High Performance Data Analysis (HPDA) workloads in the cloud has been gaining popularity
  - According to the latest OpenStack survey, 27% of cloud deployments are running HPDA workloads
- Has evolved into Data Science

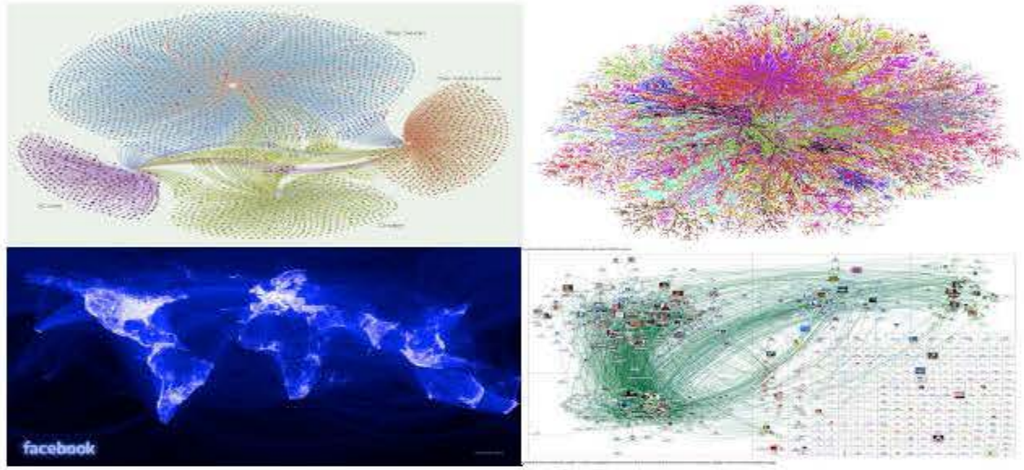

http://www.coolinfographics.com/blog/tag/data?currentPage=3

http://www.climatecentral.org/news/white-house-brings-together-big-data-and-climate-change-17194

Stage 3 (2010 -): HPC + AI (Machine Learning/Deep Learning)

- Machine Learning (ML)
  - "the study of computer algorithms to improve automatically through experience and use of data"
- Deep Learning (DL) – a subset of ML
  - Uses Deep Neural Networks (DNNs)
  - Perhaps, the most revolutionary subset!
- Based on learning data representation
- DNN Examples: Convolutional Neural Networks, Recurrent Neural Networks, Hybrid Networks
- **AI-Enabled Science, Arts, Health, and Business**

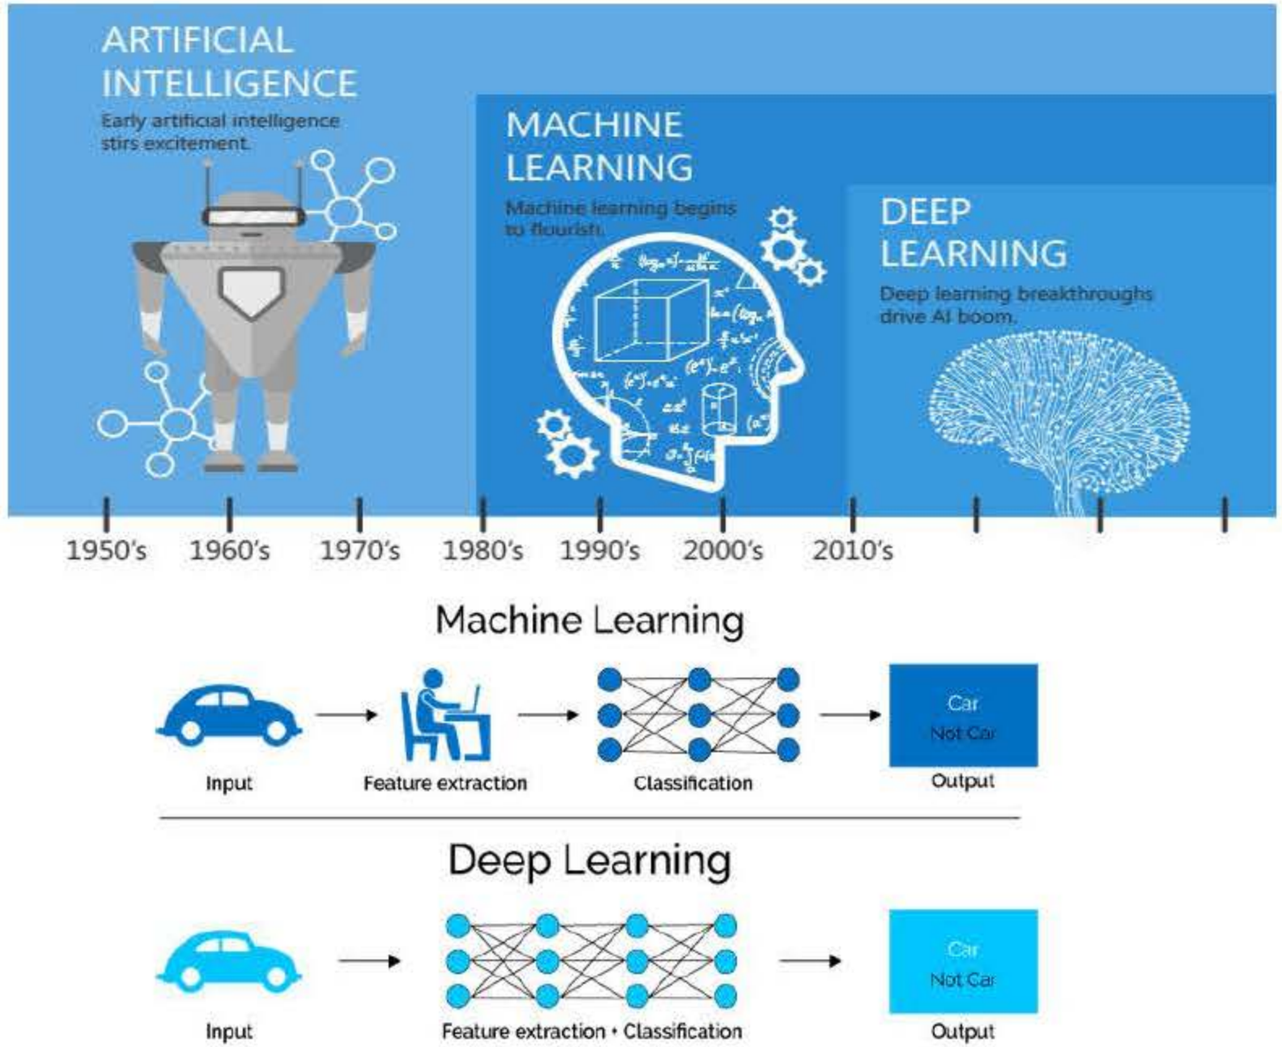

Courtesy: https://hackernoon.com/difference-between-artificial-intelligence-machine-learning-and-deep-learning-1pcv3zeg, https://blog.dataiku.com/ai-vs.-machine-learning-vs.-deep-learning, https://en.wikipedia.org/wiki/Machine_learning

Stage 4 (2015 -): Emergence of the Computing Continuum

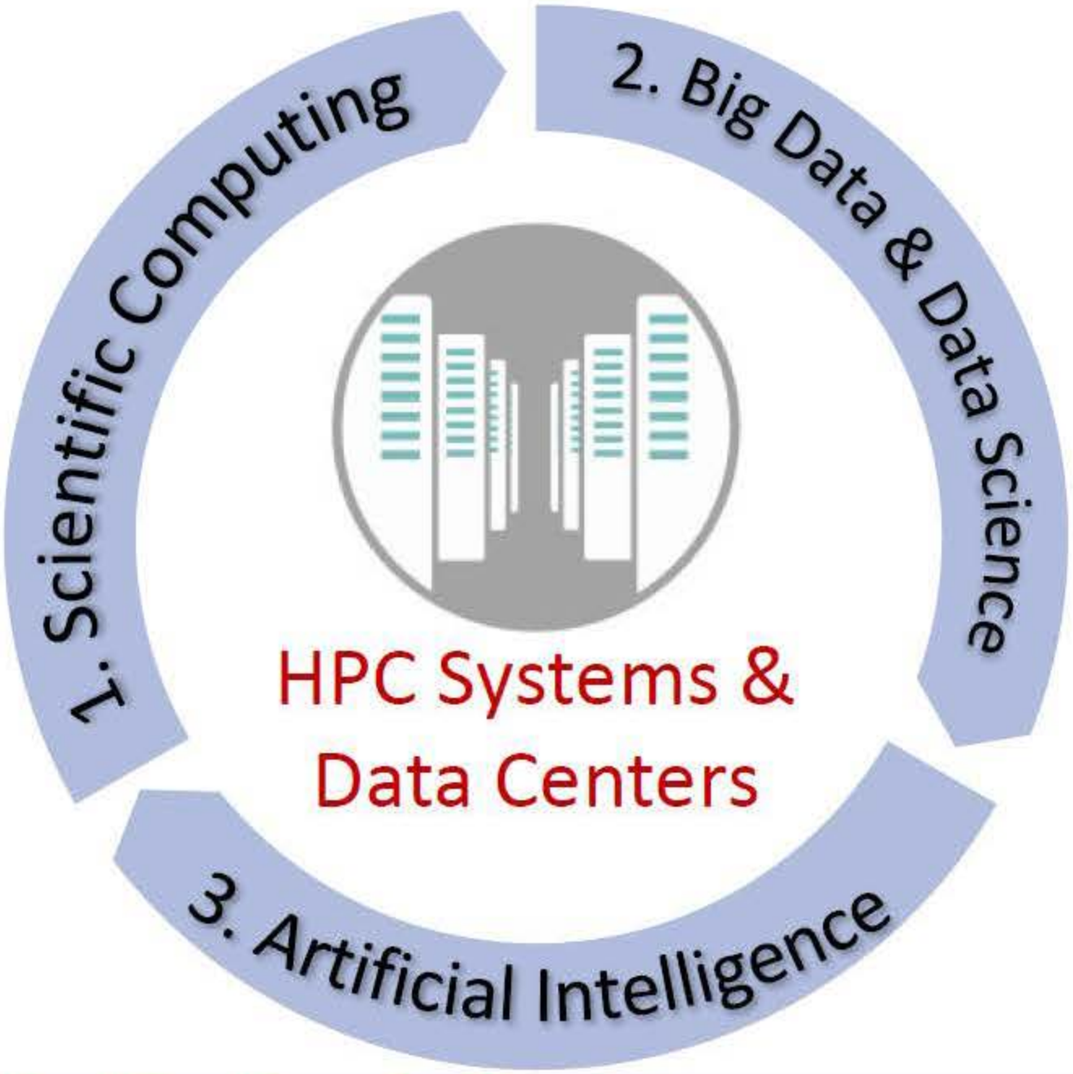

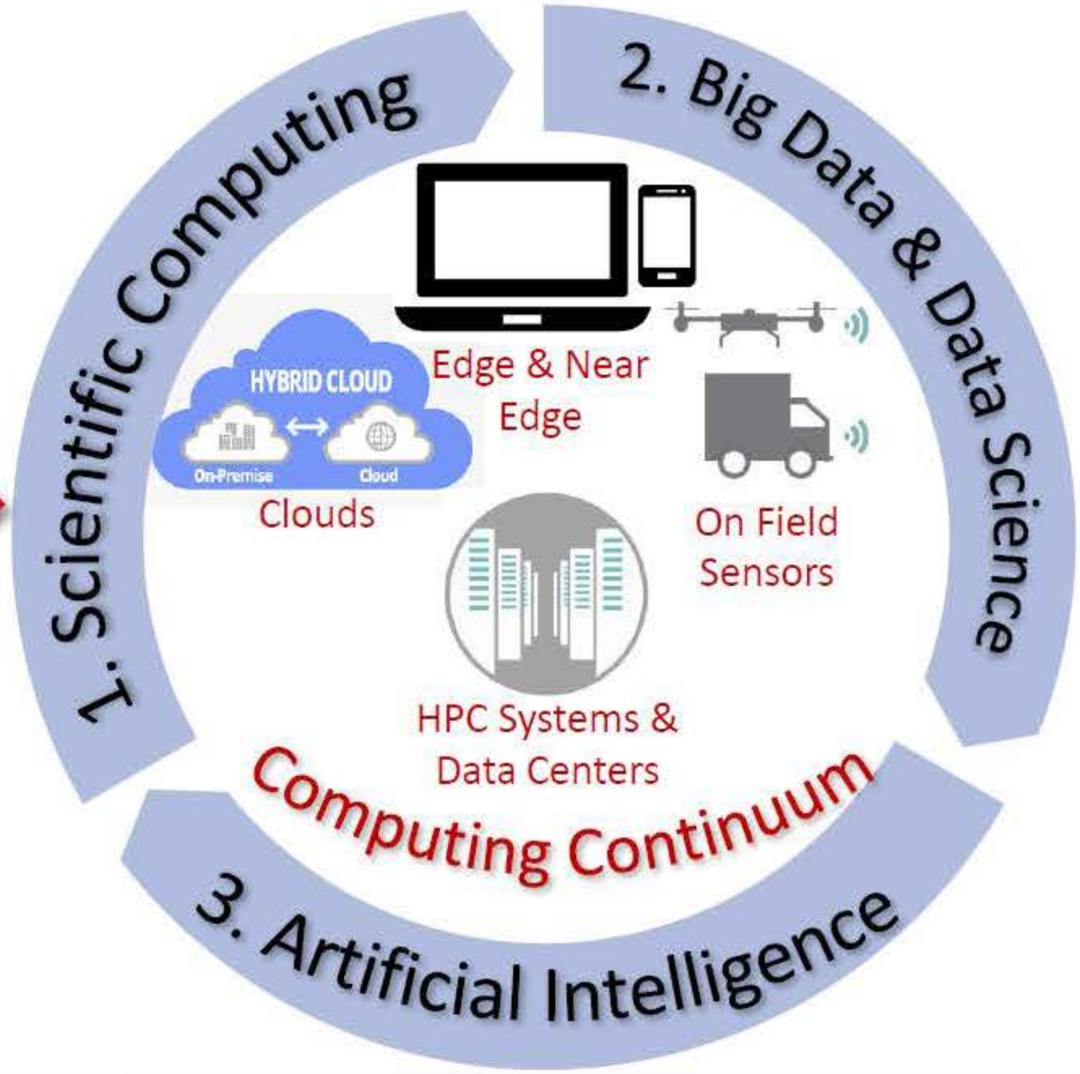

AI-Driven Precision Agriculture

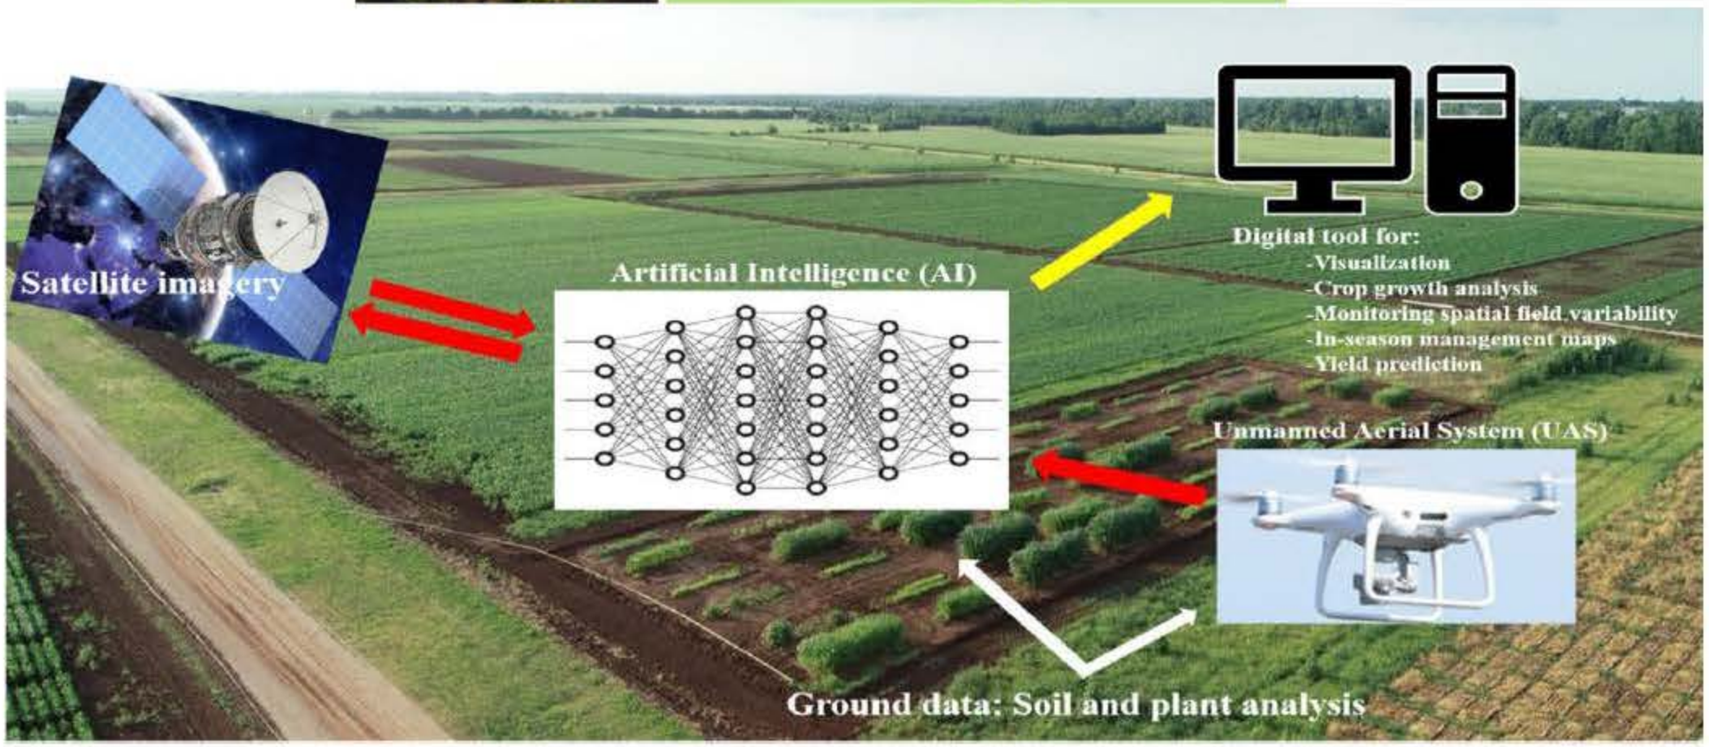

https://ccag.tamu.edu/research-project/digital-agriculture/ https://medium.datadriveninvestor.com/artificial-intelligence-in-agriculture-62f71f8f6ae6

Increasing Usage of HPC, AI, and Data Science in multiple Disciplines with Distributed Data and Heterogeneous Computing

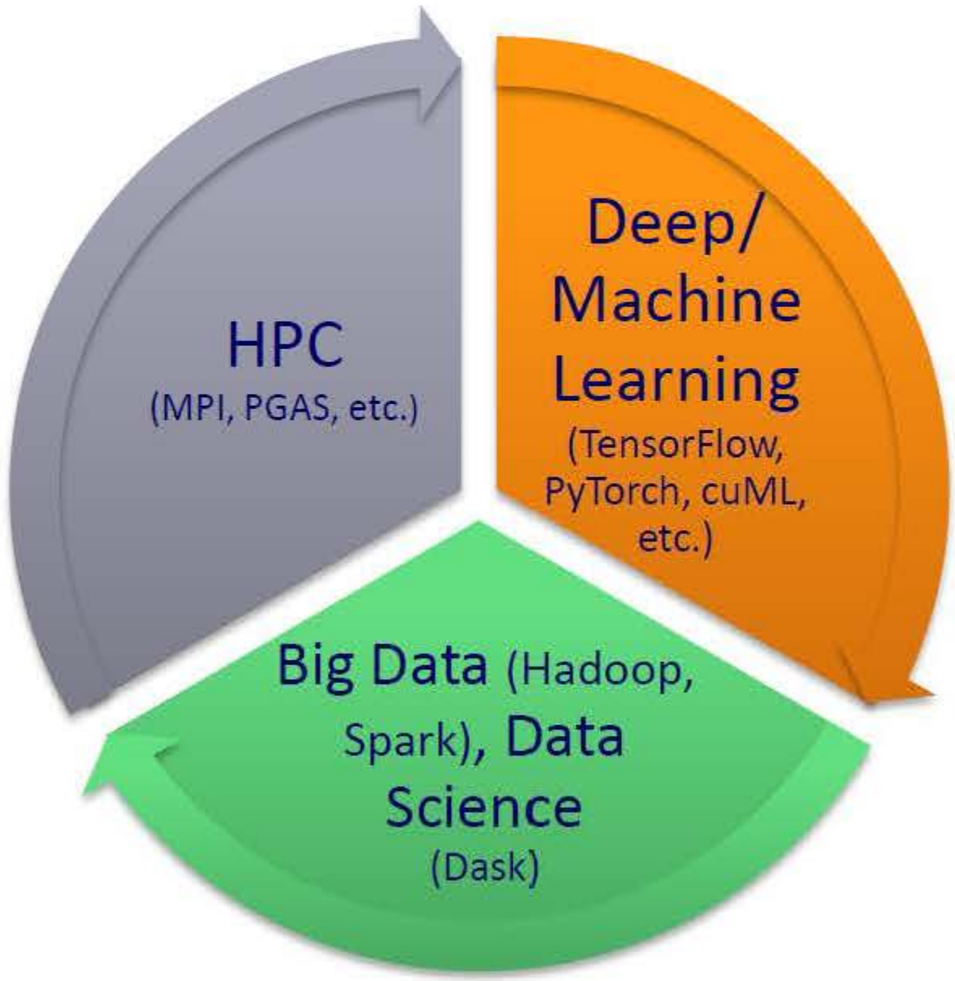

Convergence of HPC, Deep/Machine Learning, and Data Science!

Increasing Need to Run these applications on the Cloud!!

Broad Challenge:

How to design next-generation intelligent cyberinfrastructure with plug-and-play capabilities to handle societal problems while taking advantage of heterogeneous (CPU + GPU) High Performance Computing (HPC) and Cloud resources?

Follow us
@icicleai

http://icicle.ai

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

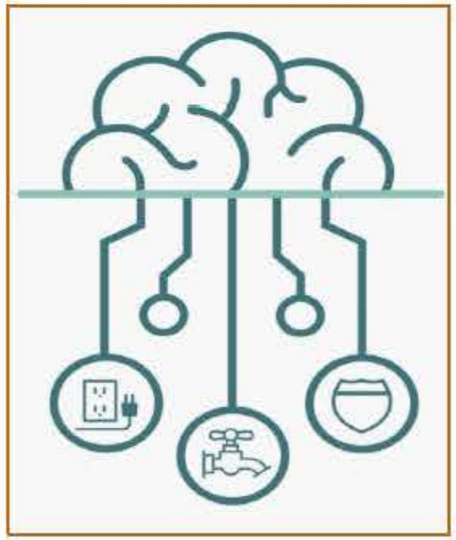

Intelligent CyberInfrastructure With Computational Learning in the Environment (ICICLE)

The Vision

A **national infrastructure** that enables AI at the flick of a switch, ICICLE will:

- **Democratize AI** through **integrated plug-and-play AI**.
- Catalyze **foundational AI/CI** and transform application domains.
- **Transparent and trustworthy** infrastructure for AI-enabled future,
- Address **societal problems** (conservation, food insecurity) and **national priorities**
- Grow **new generations of workforce** and Incubate **sustainable and inclusive** communities

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

Govt. Agencies & National Labs

Lawrence Livermore National Laboratory

BROOKHAVEN NATIONAL LABORATORY

Fermilab

International

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

SCiEC AN NSF+USGS CENTER

NSF | neon Operated by Battelle

Jetstream

NSF AI Institutes

AIFS AI INSTITUTE FOOD SYSTEMS

Hospitals & Universities

NATIONWIDE CHILDREN'S

ICICLE Leadership Team

OVERALL LEADS


Fosler-Lussier (Co-PI) (OSU)


Machiraju (Co-PI) (OSU)


Panda (PI) (OSU)


Chaudhary (Co-PI) (CWRU)


Plale (co-PI) (IU)


Savardekar (MD) (OSU)

EDUCATION & WORKFORCE DEVELOPMENT


Lange (IC-FOODS)

BROADENING PARTICIPATIO N IN [illegible]PUT


Biggers (IU)


Thomas (SDSC)

FOUNDATIONAL SYSTEMS AI


Fosler-Lussier (OSU)


Su (OSU)

INTELLIGENT CYBERINFRASTRUCTURE

CI FOR AI


Blanas (OSU)


Zhang (TACC)

AI FOR CI-FOR-AI


Chaudhary (CWRU)


Eigenmann (UDel)

SOFTWARE ARCHITECTURE AND DESIGN


Stubbs (TACC)


Zhang (TACC)

VISUAL ANALYTICS FOR CI AND AI EXPLAINABILITY


Machiraju (OSU)


Shen (OSU)

PRIVACY, ACCOUNTABILITY AND DATA INTEGRITY


Ayday (CWRU)


Plale (IU)

CO-DESIGN FOR USE INSPIRED SCIENCE


Berger-Wolf (OSU)


Hoy (OSU)


Shearer (OSU)

COLLABORATION & KNOWLEDGE TRANSFER


Lange (IC-FOODS)


Ramnath (OSU)

CI DEVELOPMENT FOR INTEGRATION WITH SHARED RESEARCH FACILITIES


Majumdar (SDSC)


Tomko (OSU)

Objectives: Intelligent CyberInfrastructure for Computing Continuum

Use Inspired Science Domains

Smart Foodsheds

Animal Ecology

Digital Agriculture

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"**

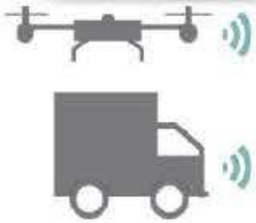

On Field Sensors

Edge & Near Edge

Clouds

HPC Systems & Data Centers

Emerging Computing Continuum

Integrating a broad range of

- Scientists-in-the-field
- Engineers
- Educators
- Collaborative partners
- Institutions

under one roof enables **democratized, adaptable, plug-and-play AI** and **long-tail science.**

Research Plan: Overall Vision

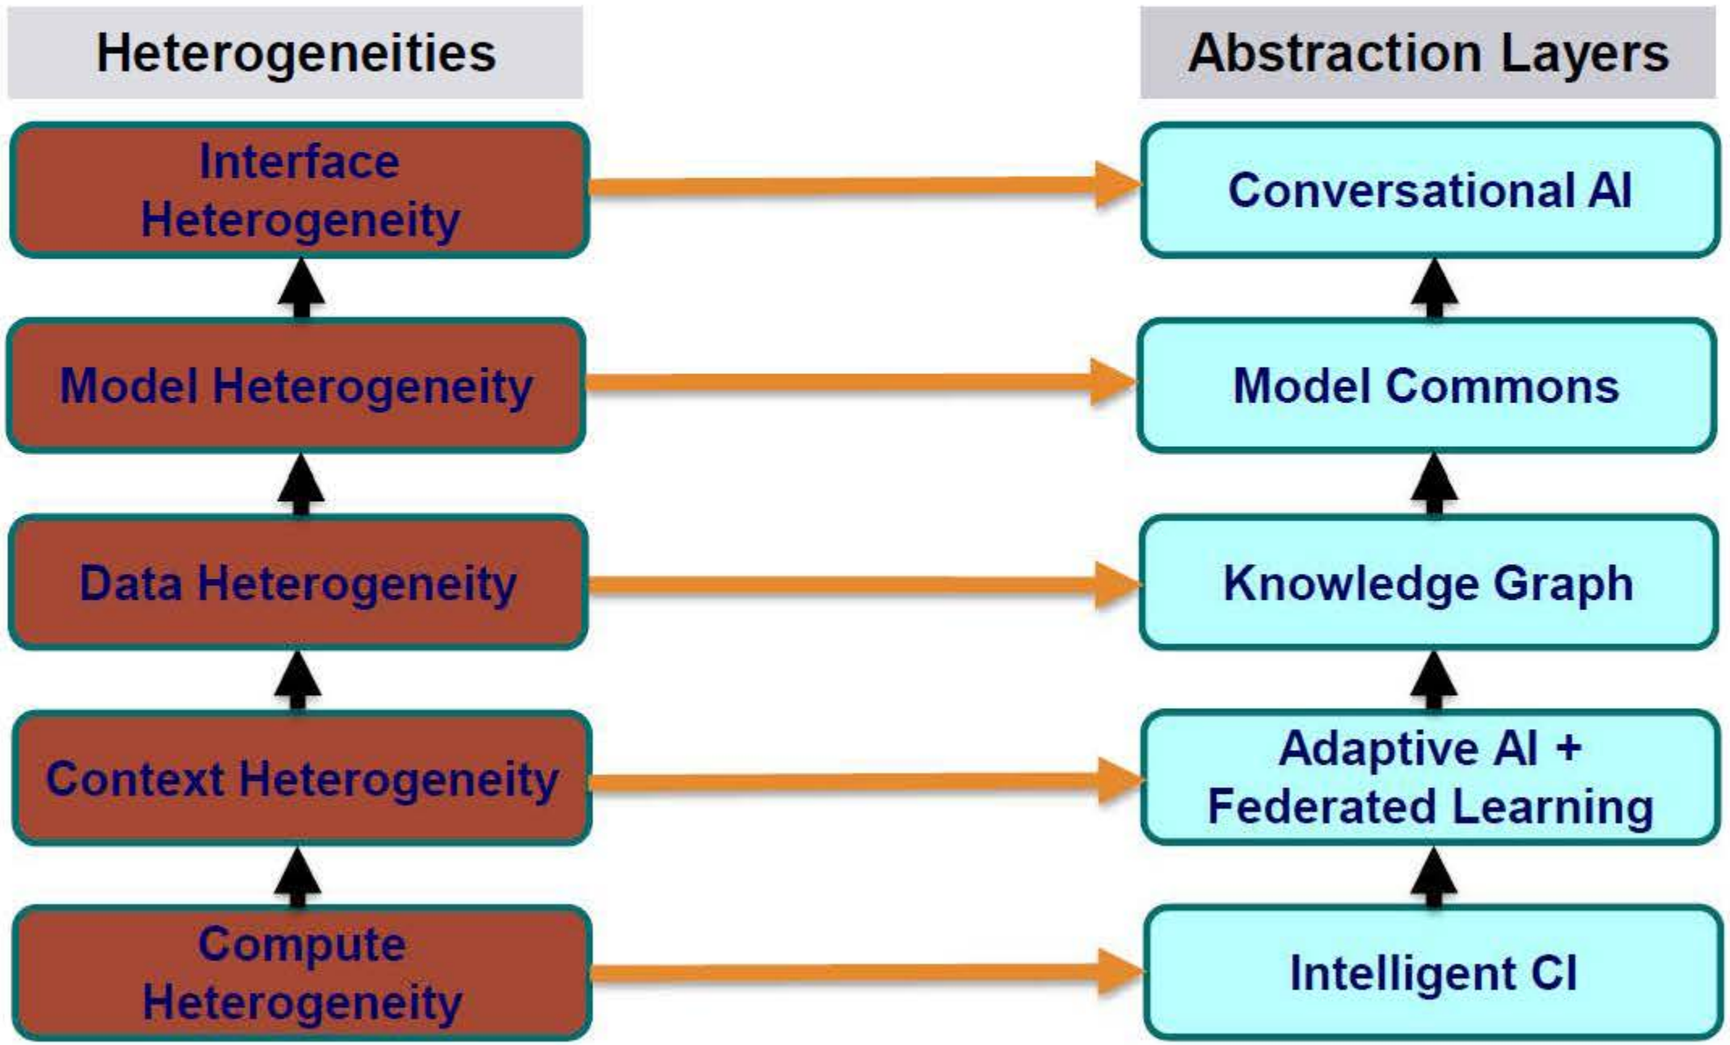

The DNA: Foundational Systems AI

Knowledge Graphs

- Multimodal
- Spatio-temporal
- Auto construction
- Knowledge-based reasoning and pre-training

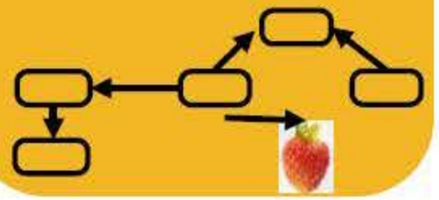

Model Commons

- KG-supported
- Precise profiling
- Flex Composition
- Versioning and provenance

Crop Yield Model
Midwest
Corn, Strawberry…

Adaptive AI

- Context-aware
- Interactive
- Continual learning
- Distillation-based compression

Federated Learning

- Heterogeneity
- Applicability to a variety of models
- Context-aware
- Privacy-preserving and robustness

Conversational AI

- KG- and model-commons-aware
- Bootstrapping and adaptivity
- Multimodal contextual response

https://pixabay.com/photos/strawberry-field-fruit-plant-3630814/
https://commons.wikimedia.org/wiki/File:Strawberry_field_at_Bedugul.jpg

The Enabler@Edge: CI4AI

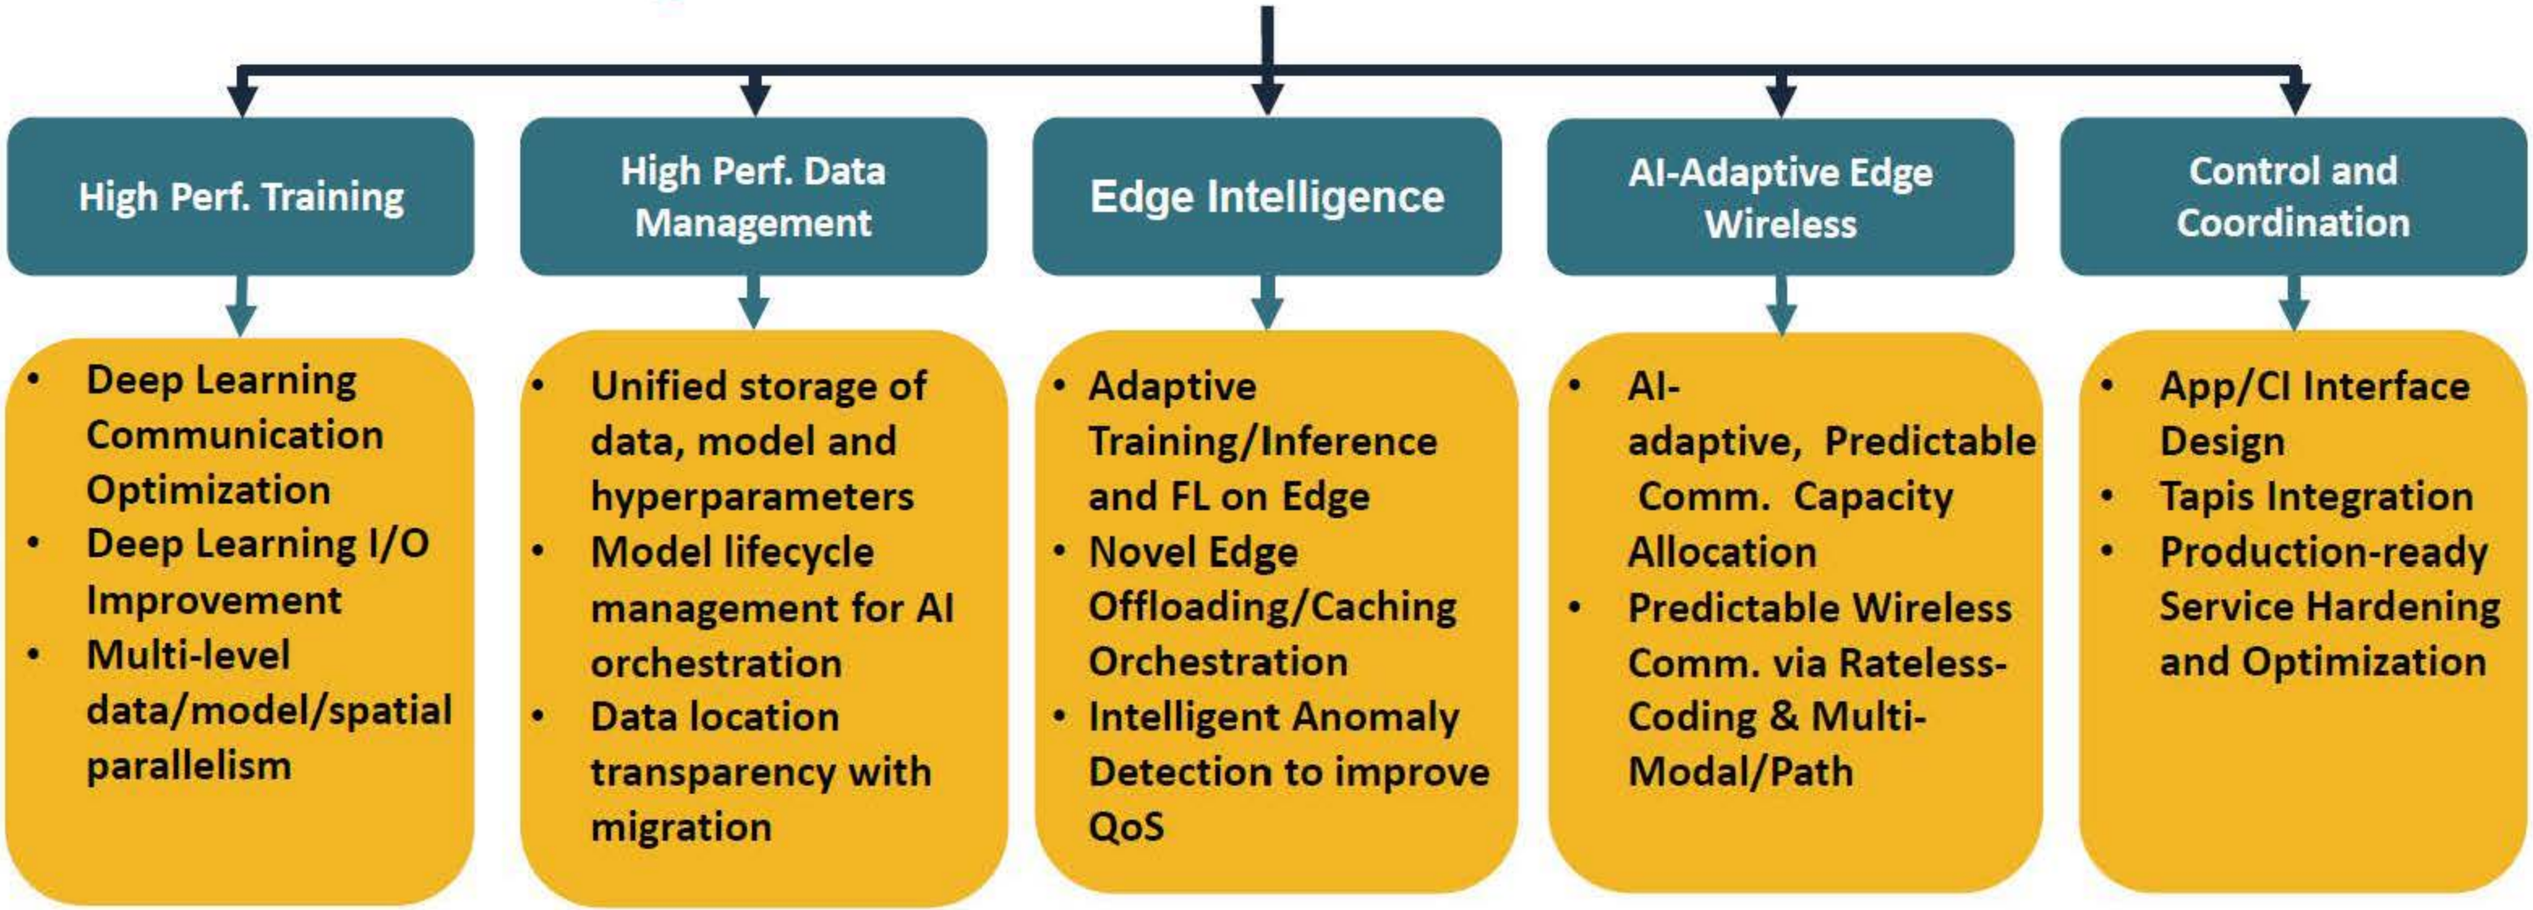

The Enabler@Scale-and-@Edge: AI4CI

[redacted]

Intelligent Primitives (sparse/dense)
Product: High Performance Library
(Powered by AI to maximize the utilization of CI)

Innovation: **Learn** from hardware (e.g., SIMD width), network (e.g., bandwidth/latency), and data sparsity to **extract best attainable performance**

Utility: (1) **Portable** high performance on diverse HPC systems (2) **Usable** at the backend of any other AI system

Heterogeneous and rapidly evolving platform (CPUs, GPUs, Supercomputers, edge devices)

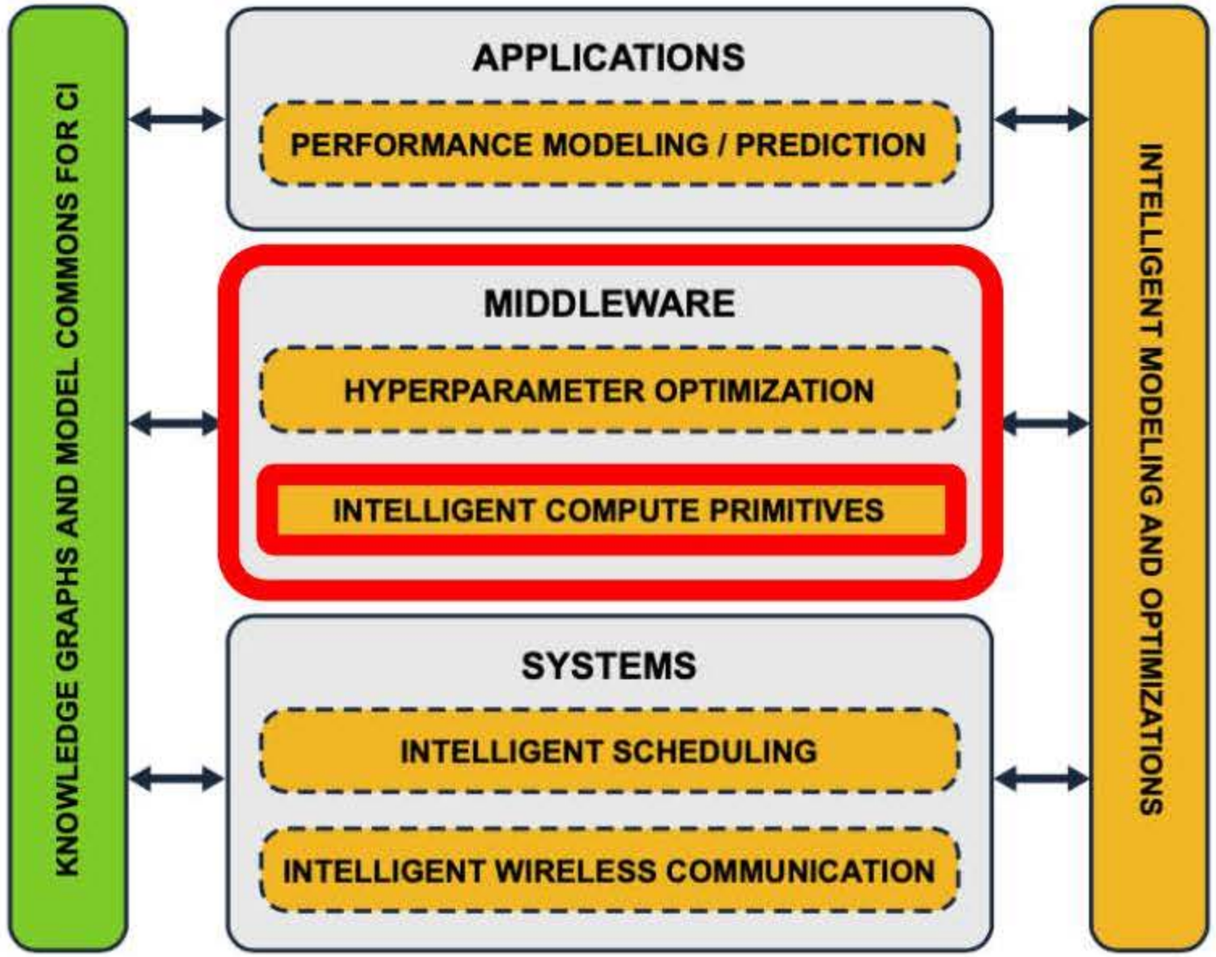

Research Plan: Use-Inspired Science

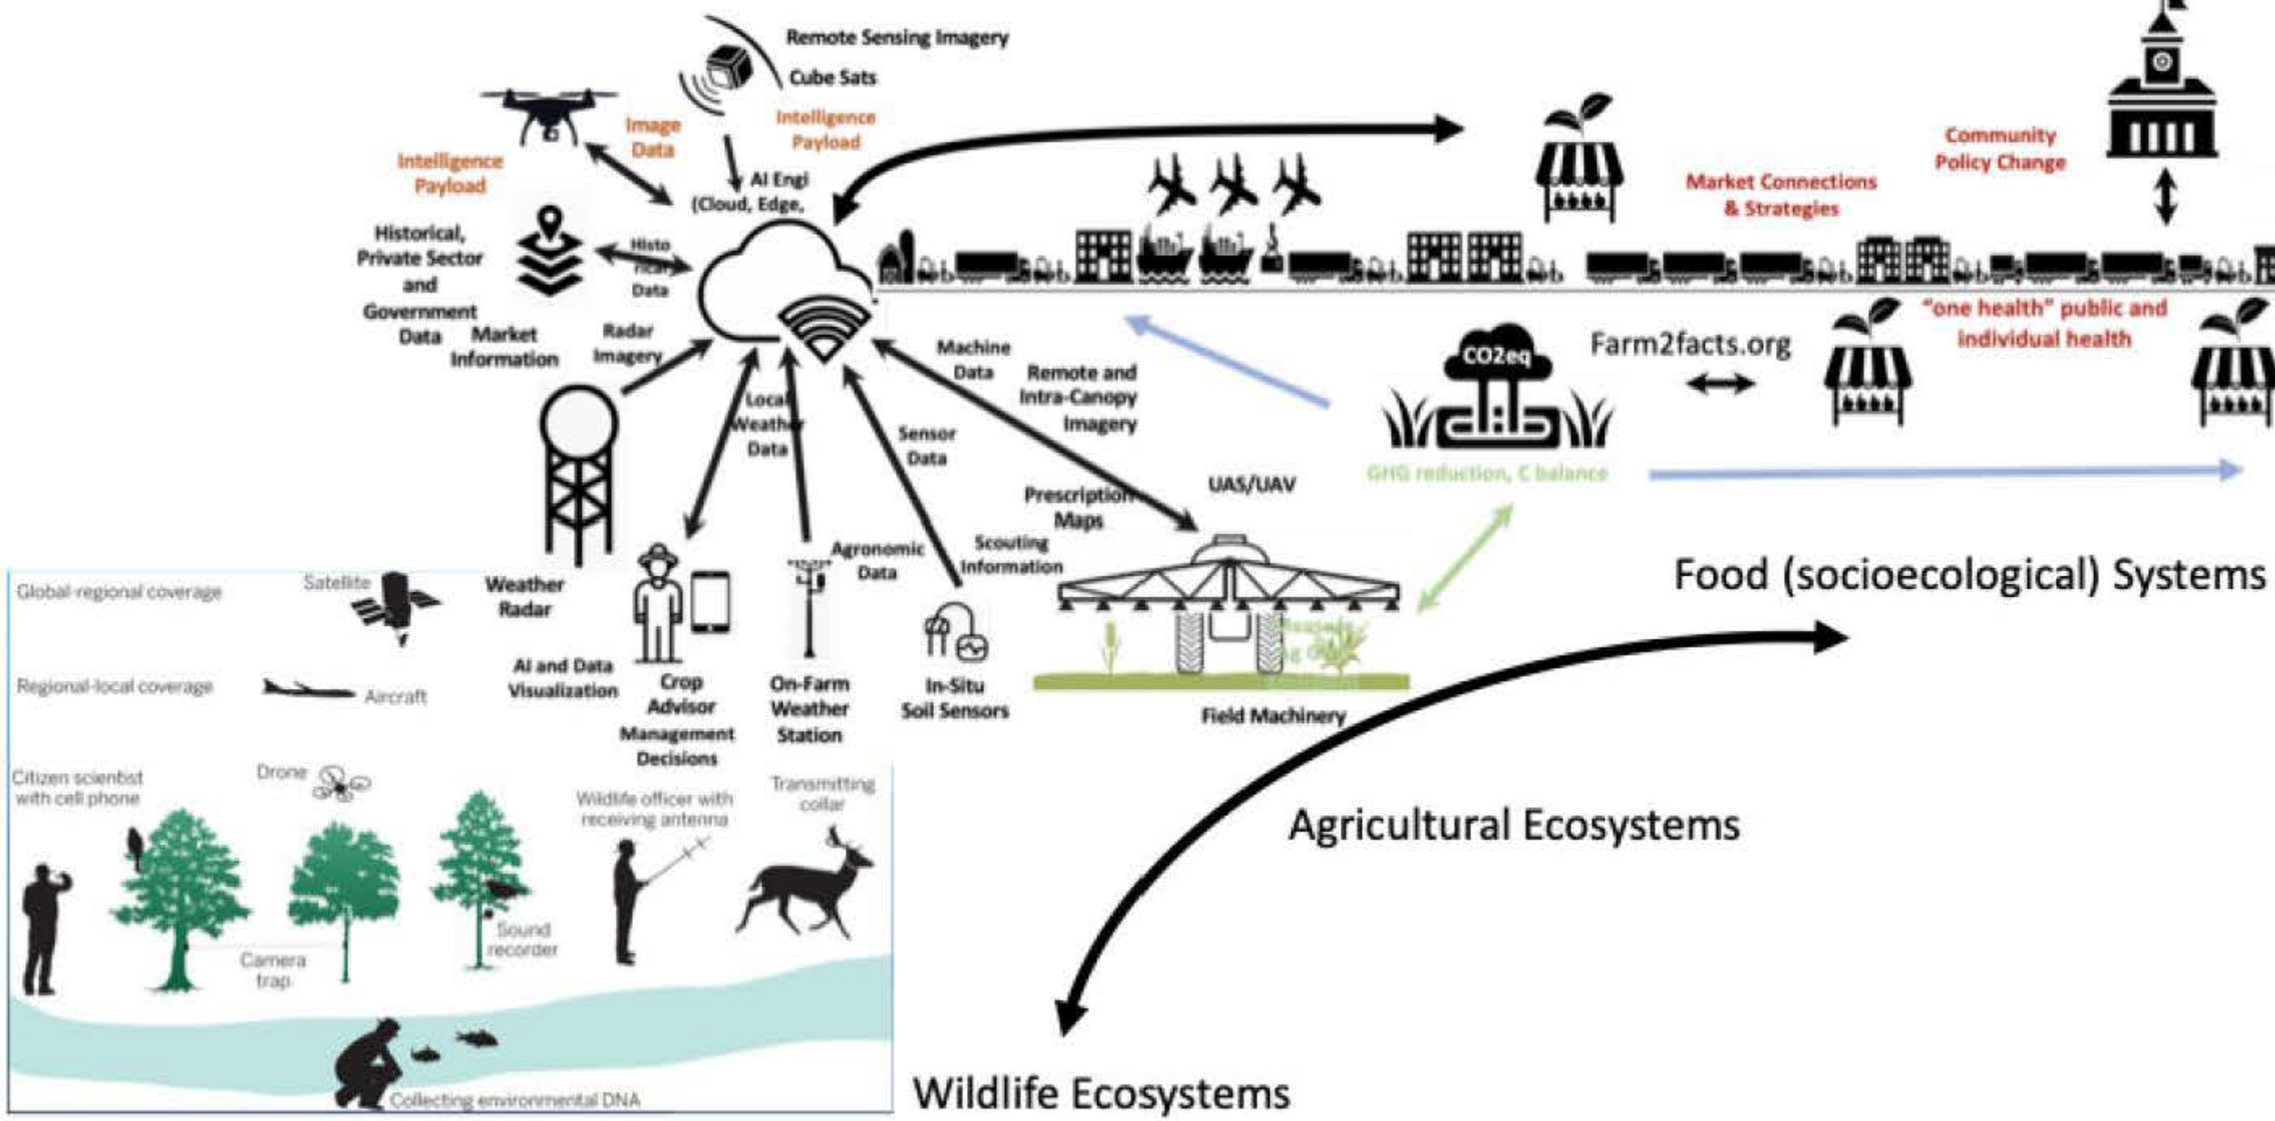

Crop Care: Nutrient and Pest Management

Demonstrate swarms of small unmanned aerial systems to study crop stressors, such as insect infestations, nutrient stress, disease, etc. and produce crop and soil maps that improve agricultural productivity

- **Edge computing setup to speed up AI model training (NSF Funded Tapis@Texas Adv. Comp. Center)**
- **SoftwarePilot, open-source software developed for ICICLE at OSU, will be used to pilot sUAS**
- **Produce crop and soil maps that improve agricultural productive**
- **Demonstrate use-cases on multiple crops (e.g., corn, soybean, wheat) via novel neural network models**
- **Feedback control for automation of agricultural field machinery (e.g., targeted pesticide application etc.).**
- **Extraction of actionable information from machine and agronomic data to support cropping decisions.**

Crop Care: Water Management and Quality

Capture nutrient stress and moisture deficit in corn and soybean field and provide feedback control for automation of agricultural field machinery, e.g., controlled drainage, robotic irrigation, etc.

- **Extraction of actionable information from machine and agronomic data to support water management**
- **Adaptive AI deployed on edge, IoT and in-field devices will boost efficacy**
- **Target irrigation scheduling and water table level**
- **Discharge water quality via controlled drainage**

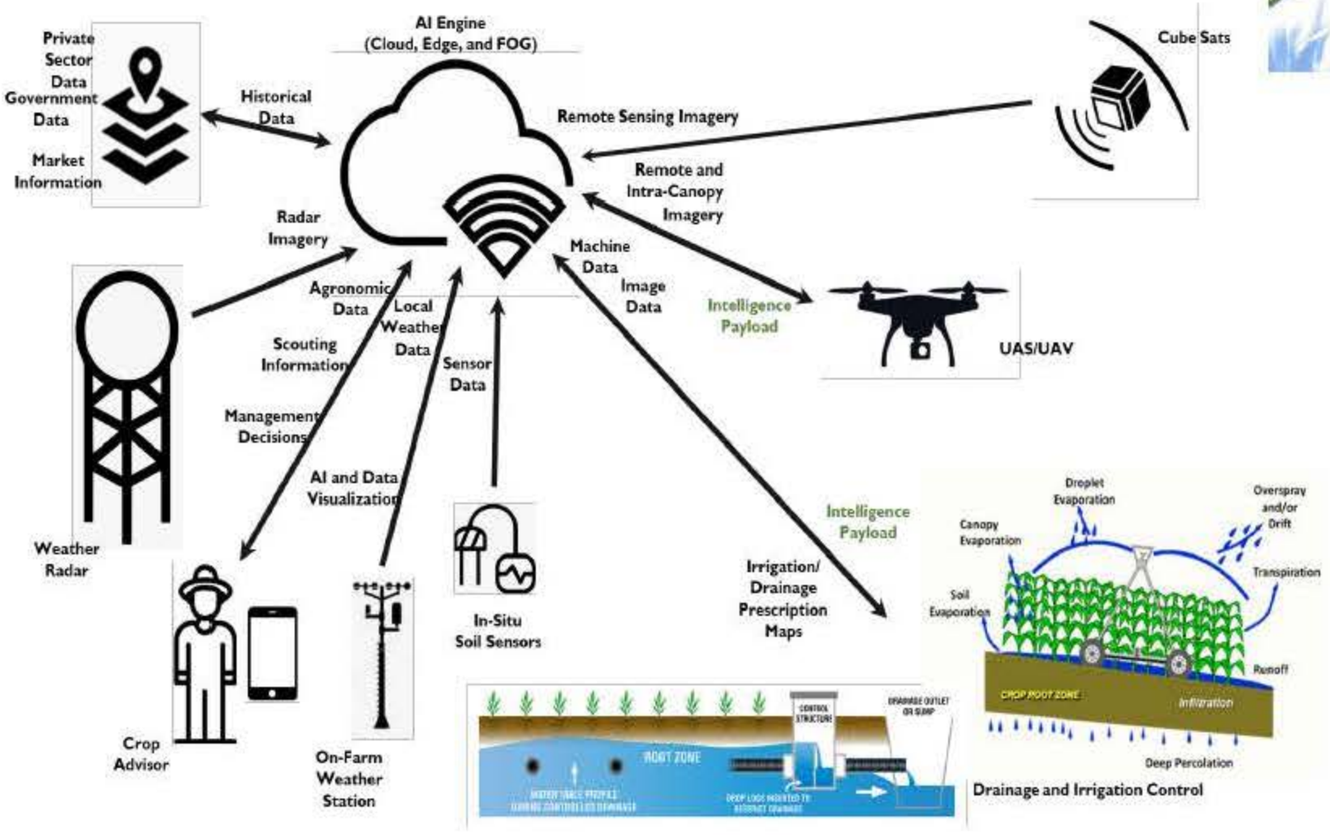

A Deliverable: The ICICLE Software Architecture

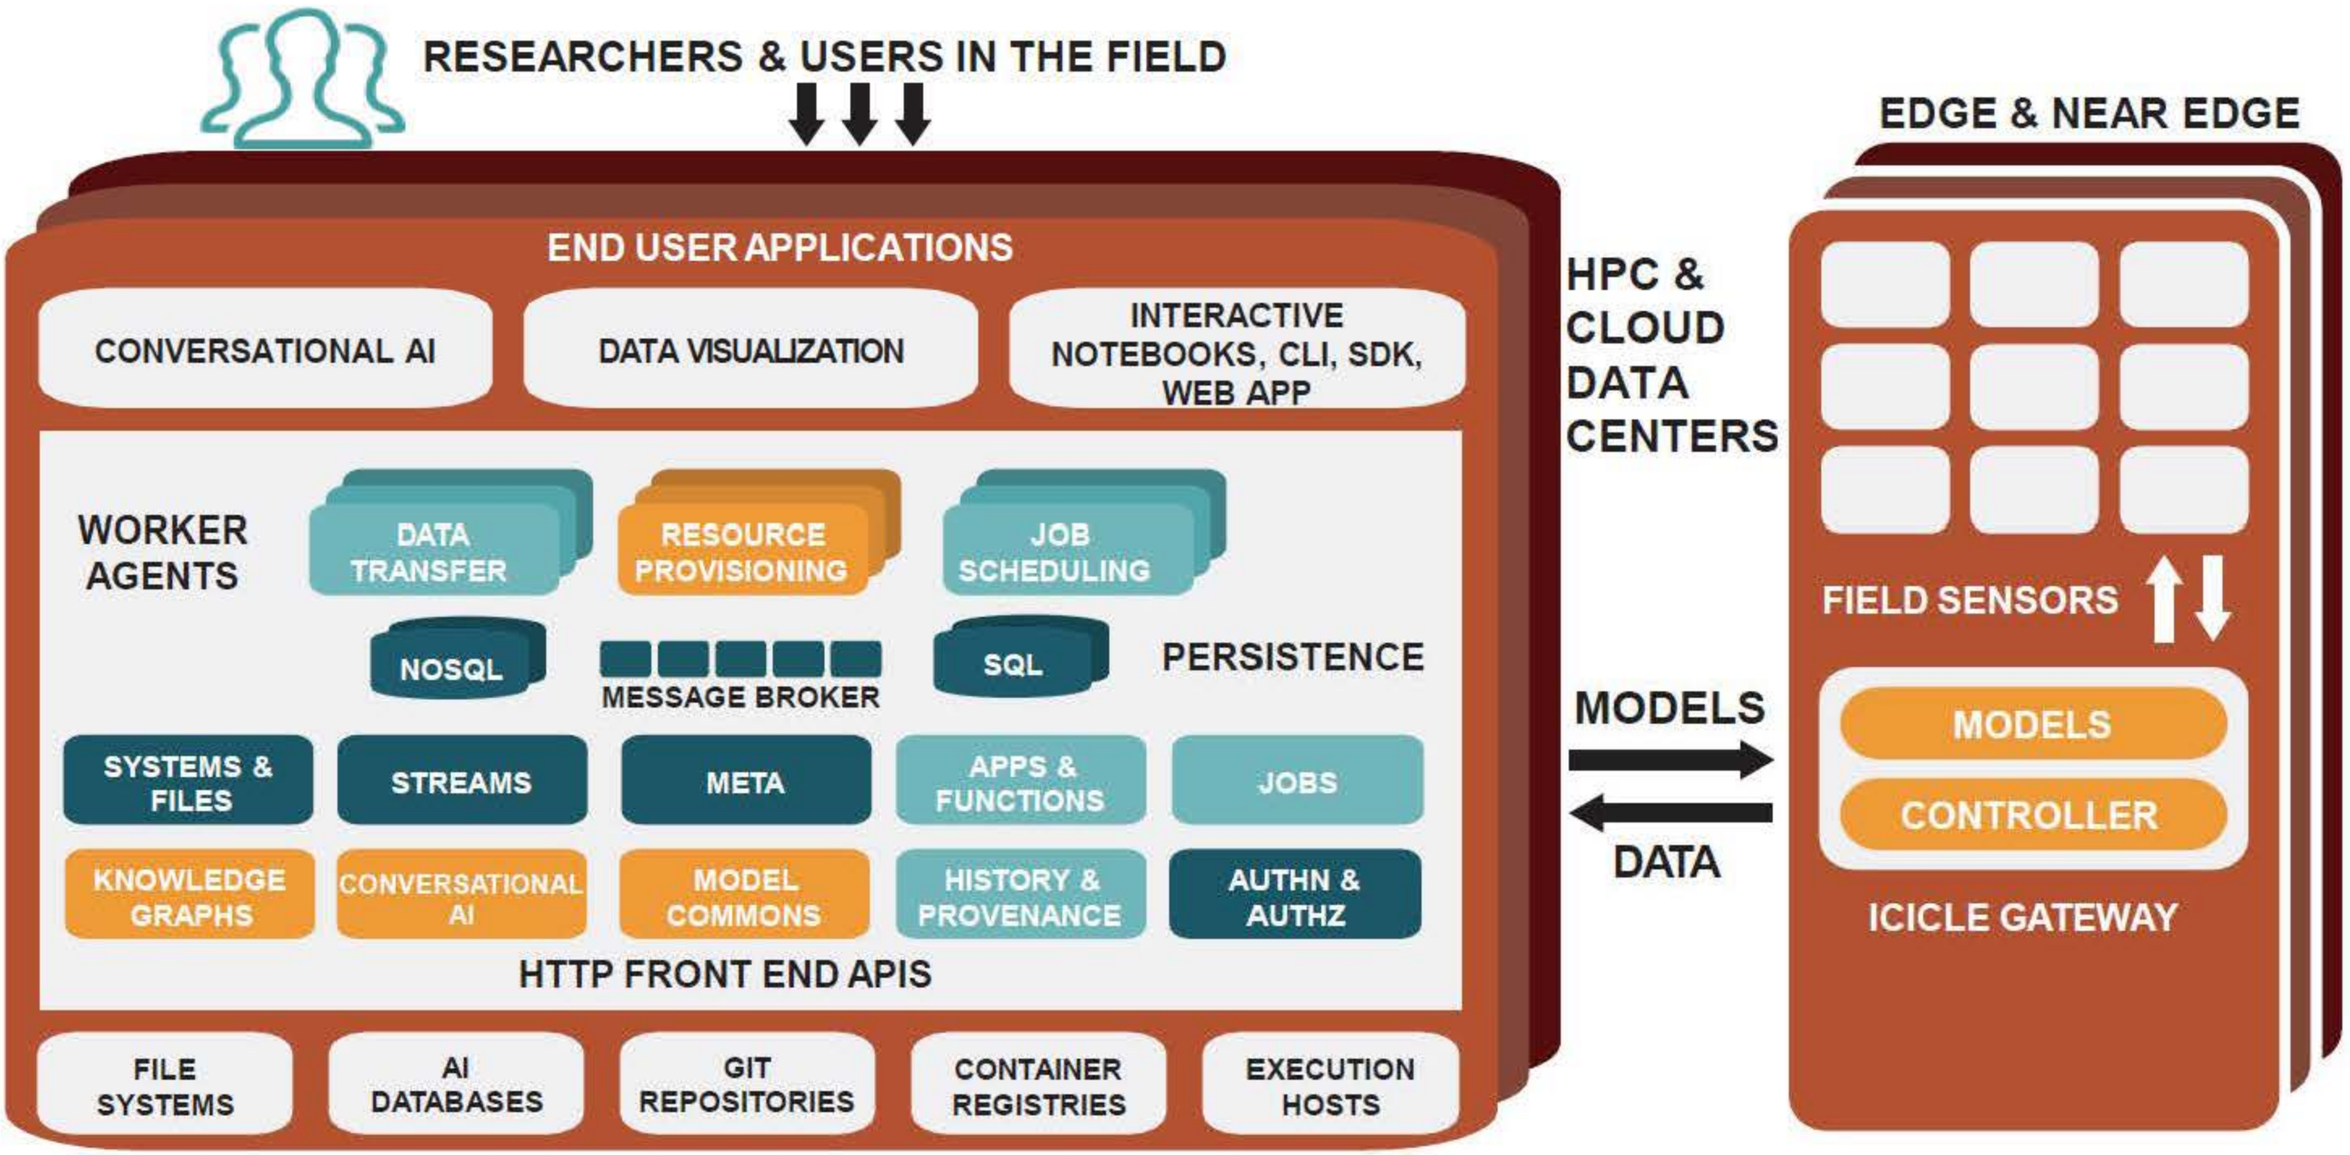

ICICLE As A Whole

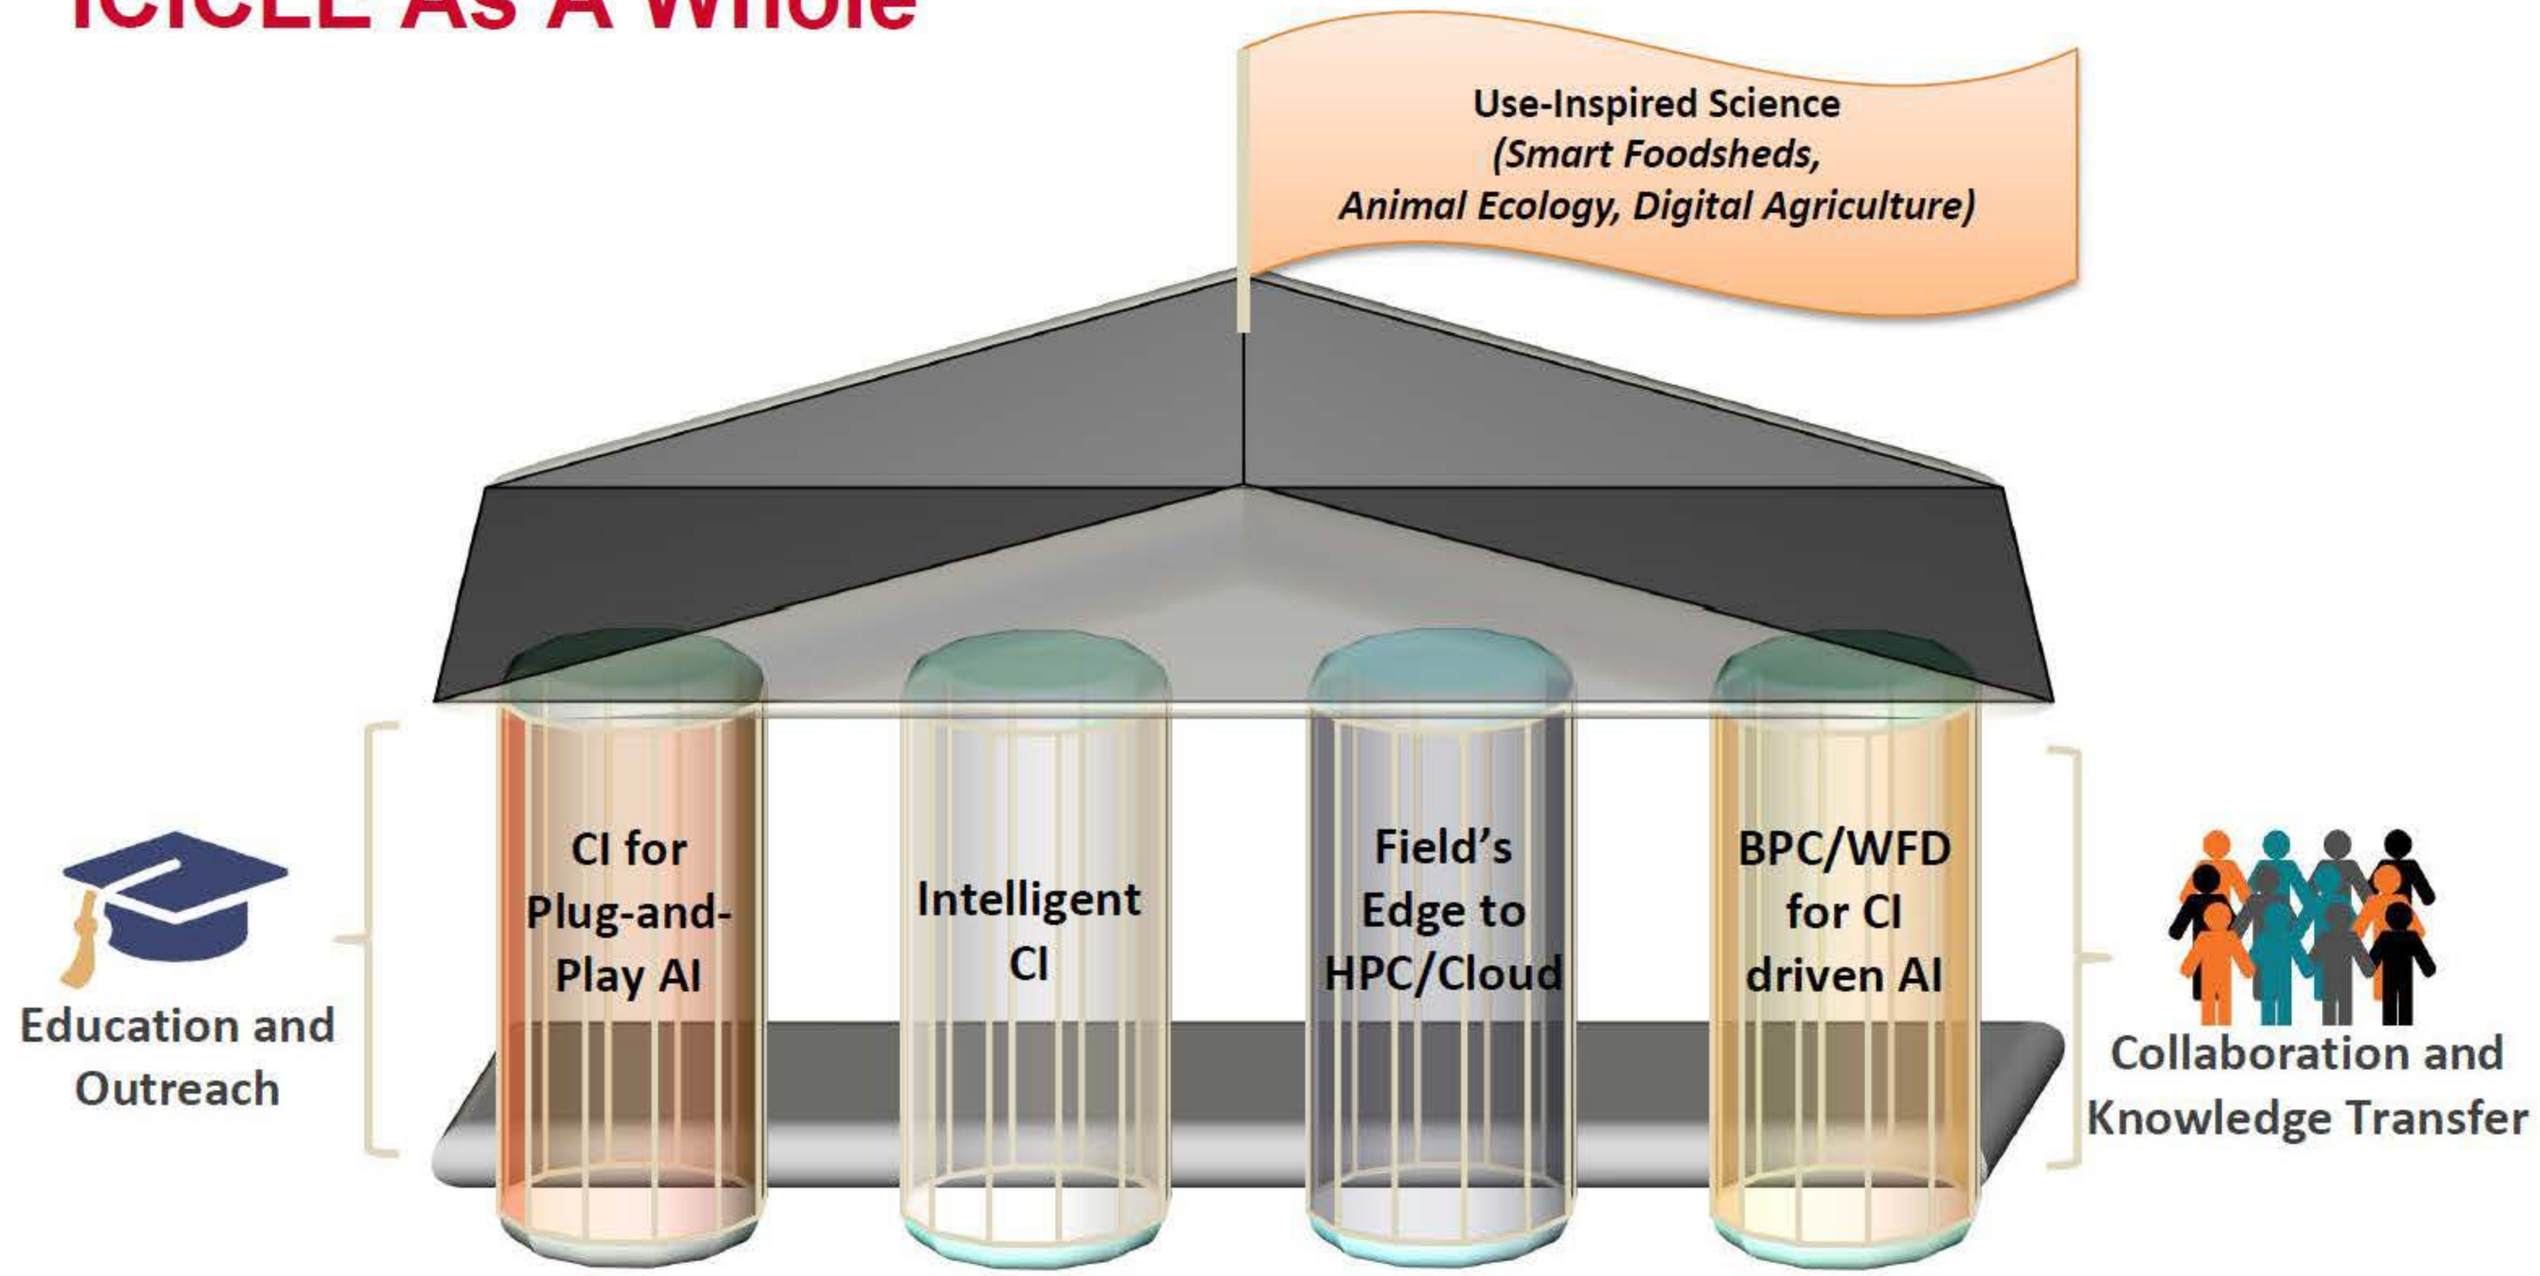

ICICLE Enables Global AI leadership

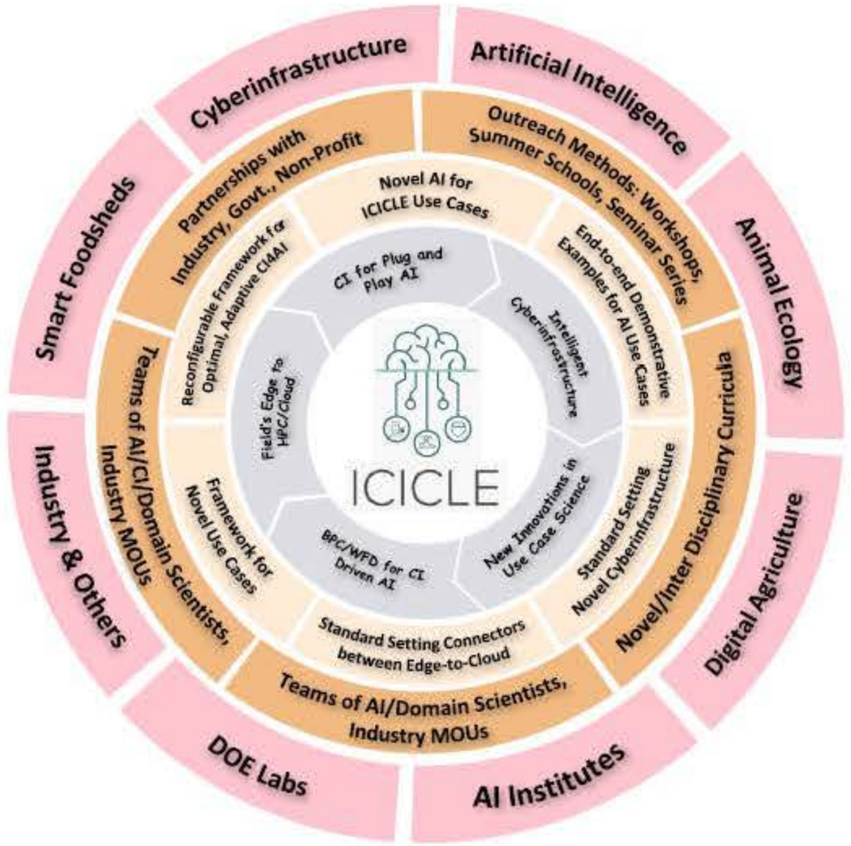

- Integrate into the National CI Ecosystem
- Integrative and Interoperable
- Leverages existing recognized capabilities
  - Centers of Excellence, AI Institutes, Large Facilities
- Collaborative
- Sustainable
  - Workforce Development, Broadening Participation, Collaboration and Knowledge Transfer
  - Benefits other institutes, large facilities, and all sciences beyond lifetime of award

Engaging With ICICLE

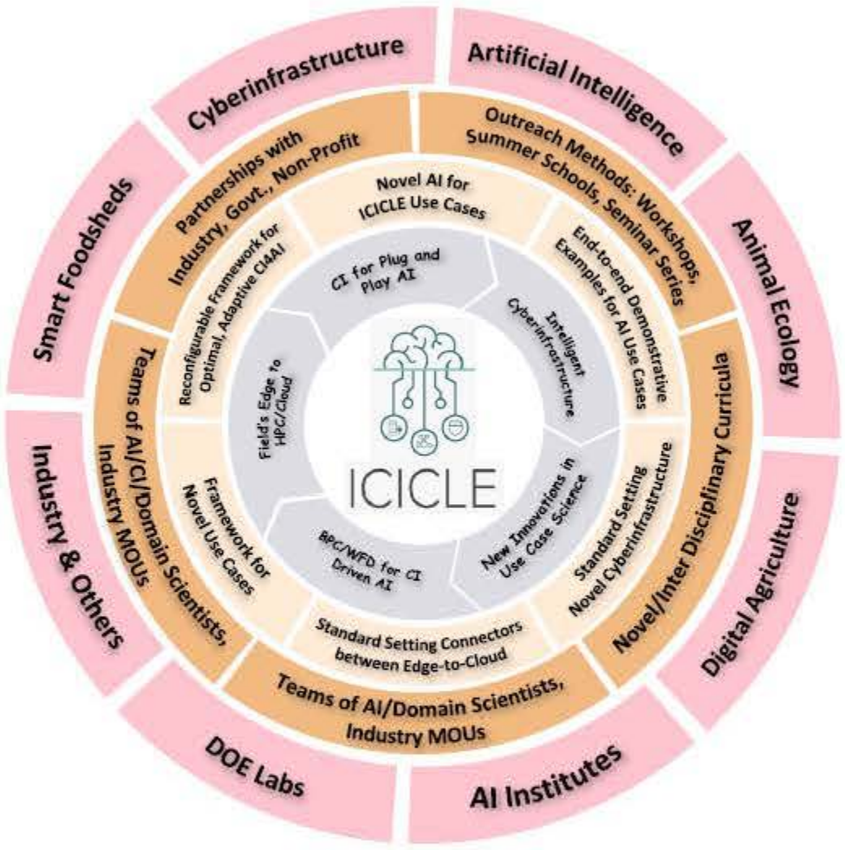

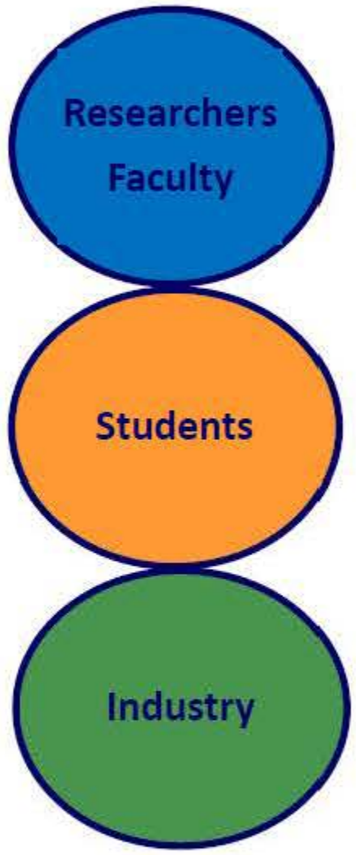

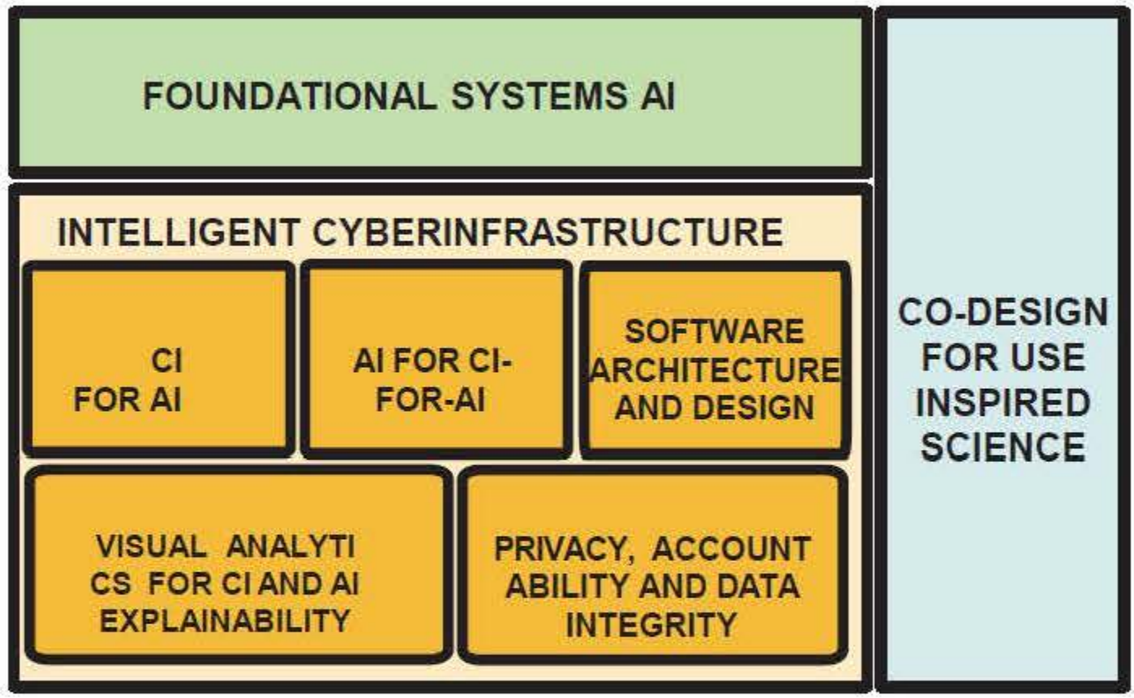

Co-develop & Adopt ICICLE developed CI!

Contact: panda@cse.ohio-state.edu

Conclusions

- AI solutions can help to solve many societal problems
- Increasing use of HPC, AI and Data Science with heterogenous resources
- Need for plug-and-play-based AI solutions which can democratize AI
- The new ICICLE NSF-AI Institute aims to establish next-generation cyberinfrastructure to provide comprehensive AI-driven solutions to many societal problems

Thank You!

panda@cse.ohio-state.edu

Network-Based Computing Laboratory

http://nowlab.cse.ohio-state.edu/

MVAPICH
MPI, PGAS and Hybrid MPI+PGAS Library

The High-Performance MPI/PGAS Project
http://mvapich.cse.ohio-state.edu/

HiBD
High-Performance Big Data

The High-Performance Big Data Project
http://hibd.cse.ohio-state.edu/

HiDL
High-Performance Deep Learning

The High-Performance Deep Learning Project
http://hidl.cse.ohio-state.edu/

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development (NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674, Fax: 202-459-9673, Email: nco@nitrd.gov, Website: https://www.nitrd.gov