SAGE

A Software-Defined Sensor Network

Cyberinfrastructure for **AI at the Edge**

sagecontinuum.org

Sage: Cyberinfrastructure for AI@Edge: A Wireless Network for Software Defined Sensors

Pete Beckman, Rajesh Sankaran, Nicola Ferrier, Scott Collis, Charlie Catlett, Eugene Kelly, Valerie Taylor, Mike Papka, Ilkay Altintas, Jim Olds, Kate Keahey, Frank Vernon, Dan Reed, and many more….

Co-Director Northwestern University / Argonne Institute for Science and Engineering (NAISE)

Argonne National Laboratory, Northwestern University, University of Chicago

Northwestern University

THE UNIVERSITY OF CHICAGO

Argonne NATIONAL LABORATORY

Colorado State University

UC San Diego

GEORGE MASON UNIVERSITY

NIU Northern Illinois University

ARRAY of THINGS

ARM

neon National Ecological Observatory Network

HPWREN

SDSC SAN DIEGO SUPERCOMPUTER CENTER

LINCOLN PARK ZOO. FOR WILDLIFE. FOR ALL.

NORTHWESTERN UNIVERSITY QUAECUMQUE SUNT VERA 1851

Argonne NATIONAL LABORATORY

"A supercomputer is a device for turning compute-bound problems into I/O-bound problems."
(Ken Batcher)

Larry Smarr at NCSA, 1986

"An edge computer is a device for turning I/O-bound problems onto compute-bound problems.

AI@Edge
Parallel Computing

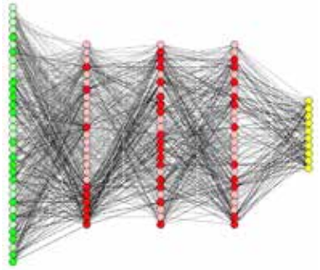

- Artificial Intelligence
- Deep Learning Inference
- Lightweight Edge Learning

AI @ Edge

Sensors

LIDAR

Software Defined Radios

Hyperspectral Imaging

Facilities

Actuators

Servos

Dynamic adaptation

Powerful

Parallel Edge Computing

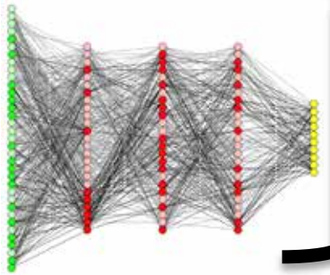

Artificial Intelligence
Deep Learning Inference
Lightweight Training

Edge computing and deep learning with feedback for continuous improvement

Reduced, Compressed data

New inference (model)
Adaptive steering

HPC

Deep Learning Training
Simulation / Forecast

Why Live on the Edge?

- **More data than bandwidth**
  - Spallation neutron source, light source, SW defined radios, HD Cameras, LIDAR, radar, hyperspectral imaging, grid micro-synchrophasors, etc.
- **Latency is important**
  - Quick local decision & actuation; adaptive sensing & control systems
- **Privacy/Security requires short-lived data: process and discard**
  - Compromised devices have no sensitive data to be revealed
- **Resilience requires distributed processing, analysis, and control**
  - Predictable service degradation, autonomy requires local (resilient) decision
- **Quiet observation and energy efficiency**
  - Vigilant sensors, transmit only essential observations, not big data streams

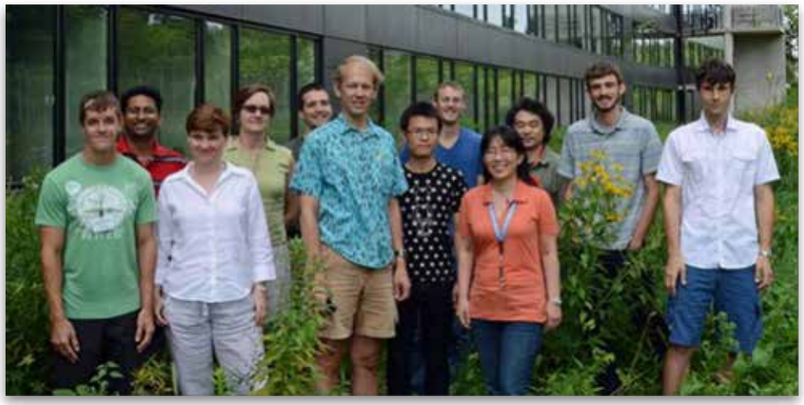

2013

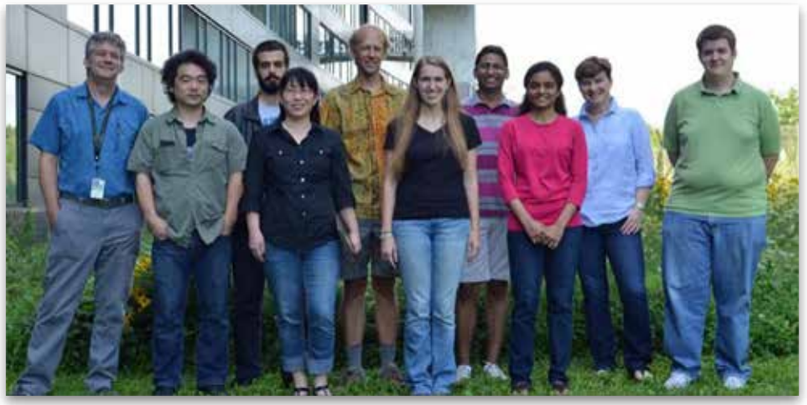

2014

2015

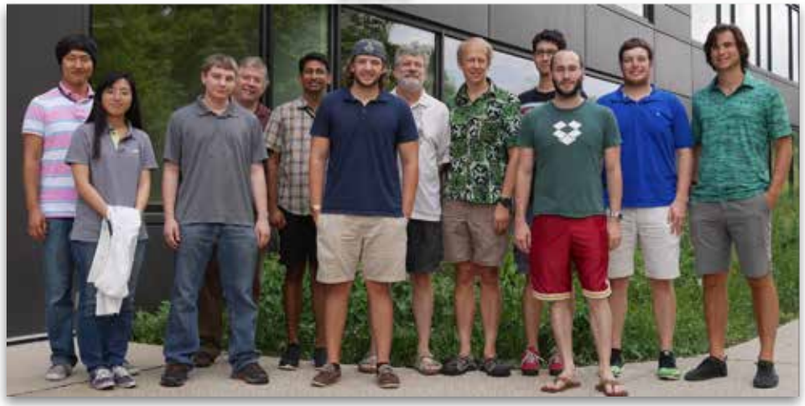

2016

2017

2018

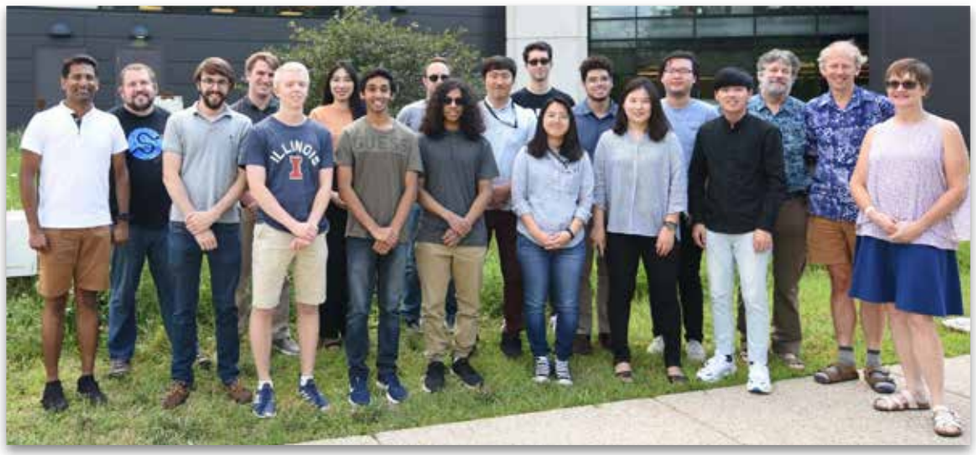

2019

2020

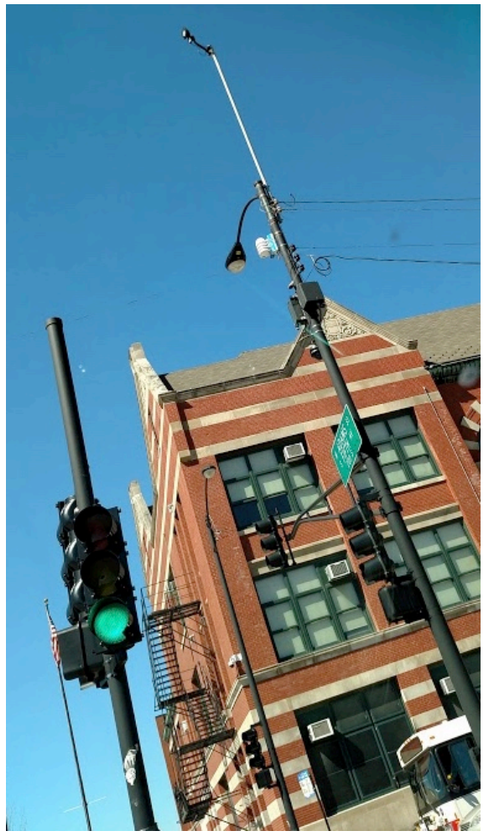

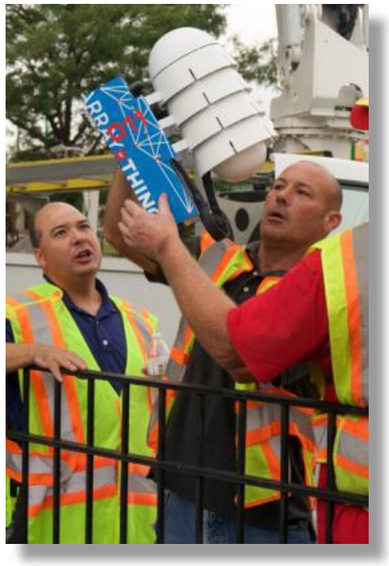

Array of Things:
The first *Edge Computing* deployment

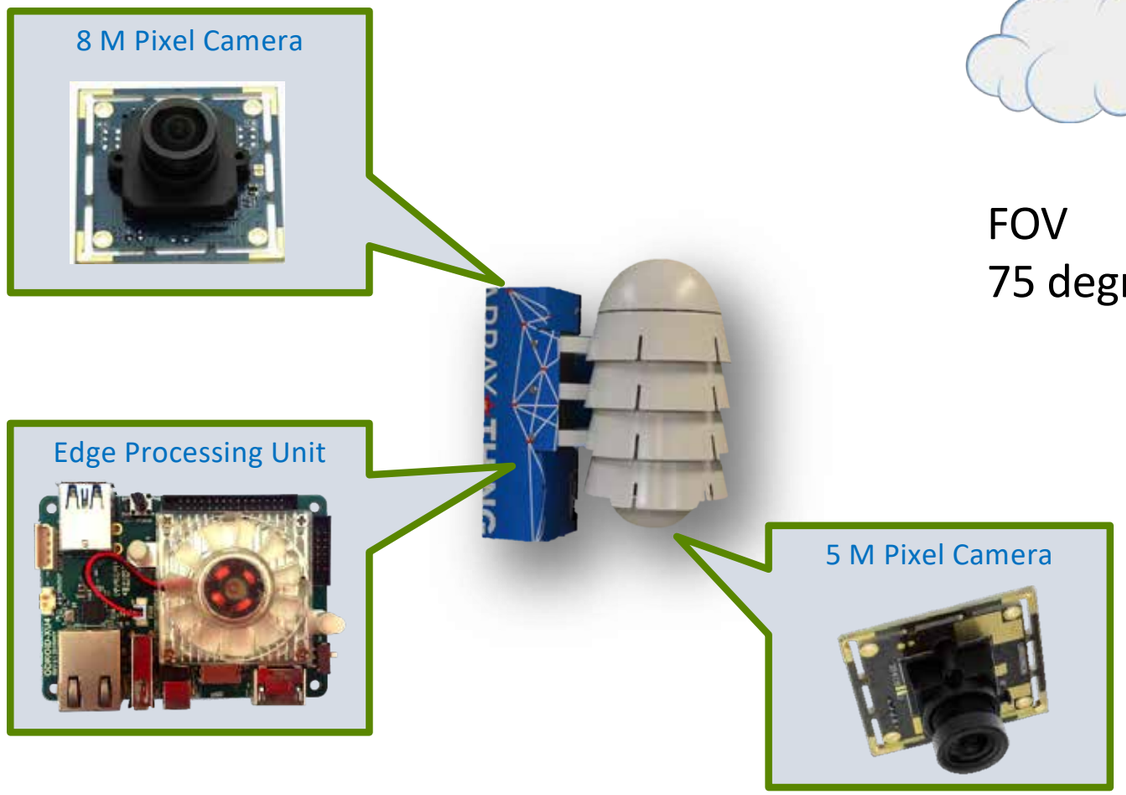

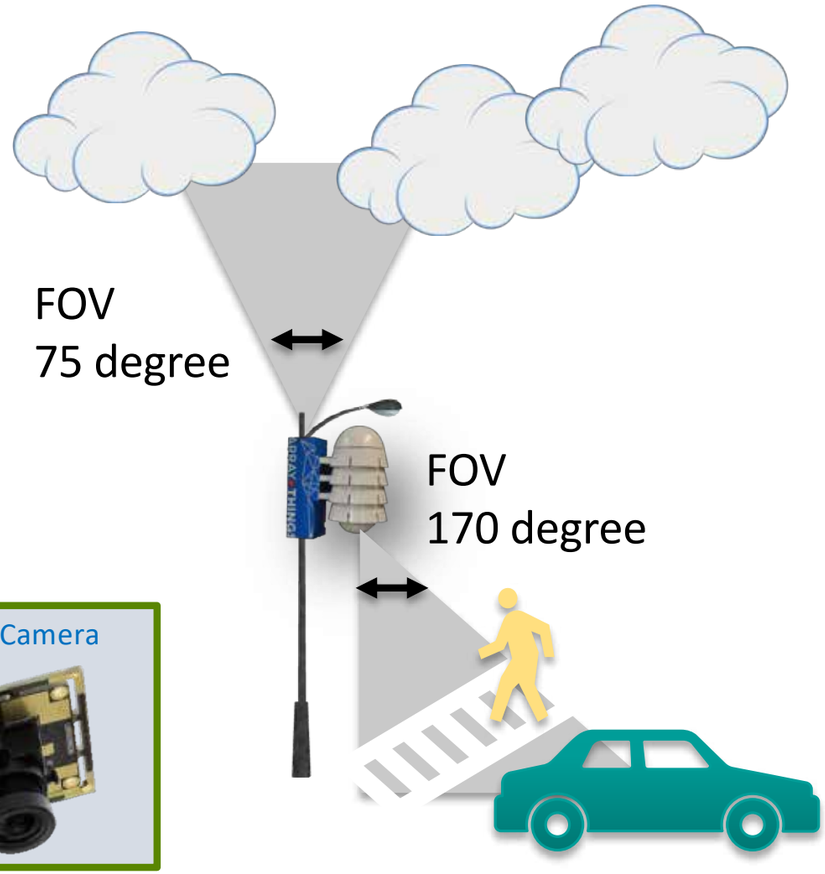

NSF: 2015

What does AI@Edge Enable?

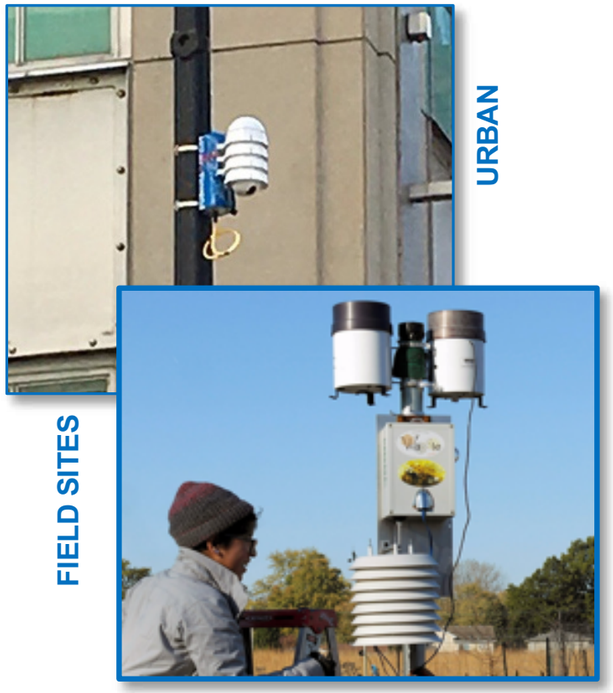

Software-Defined Sensors

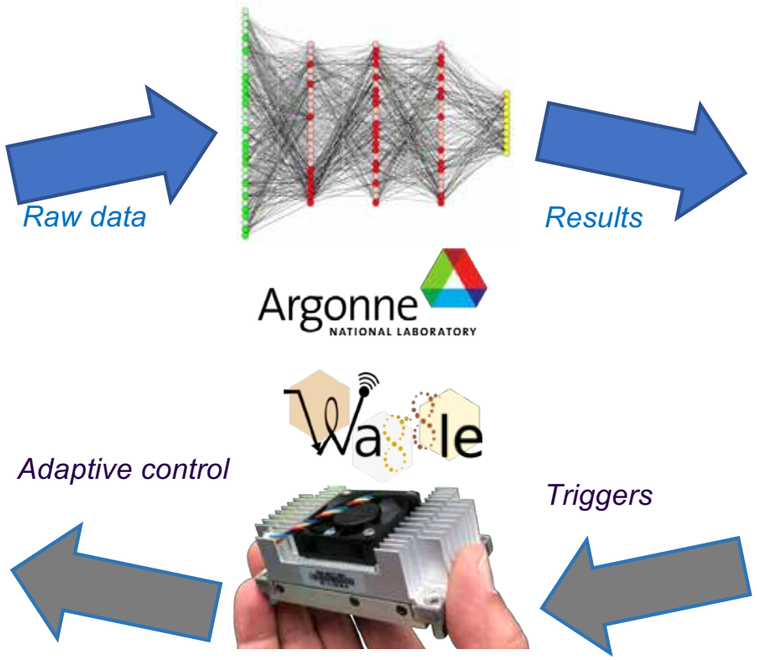

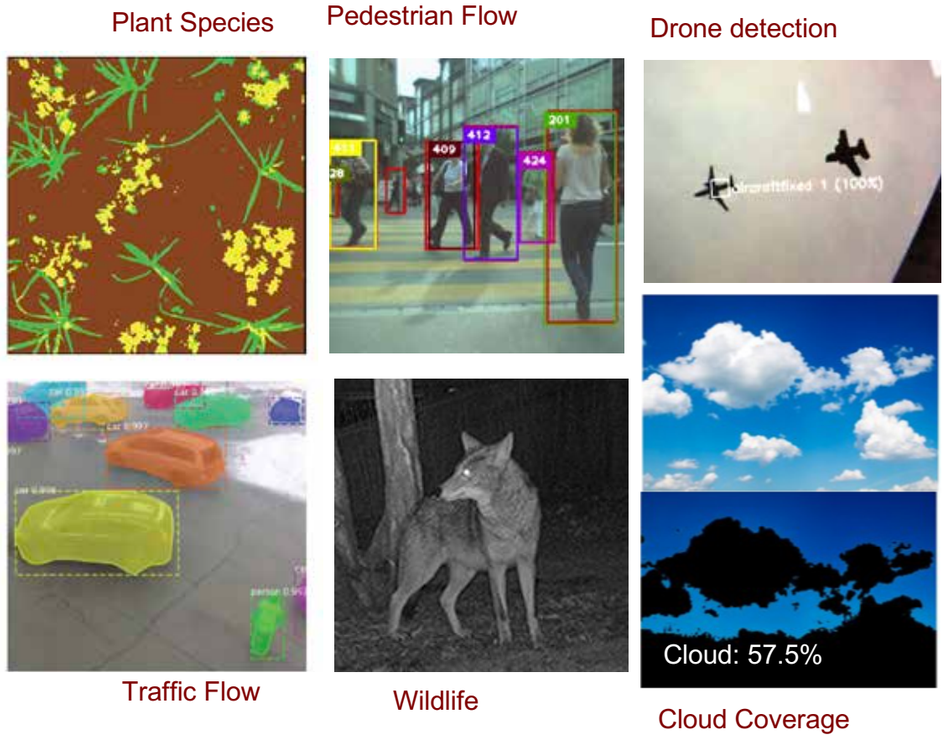

Autonomous Adaptive Sensing

SAGE

Cyberinfrastructure for
AI at the Edge
sagecontinuum.org

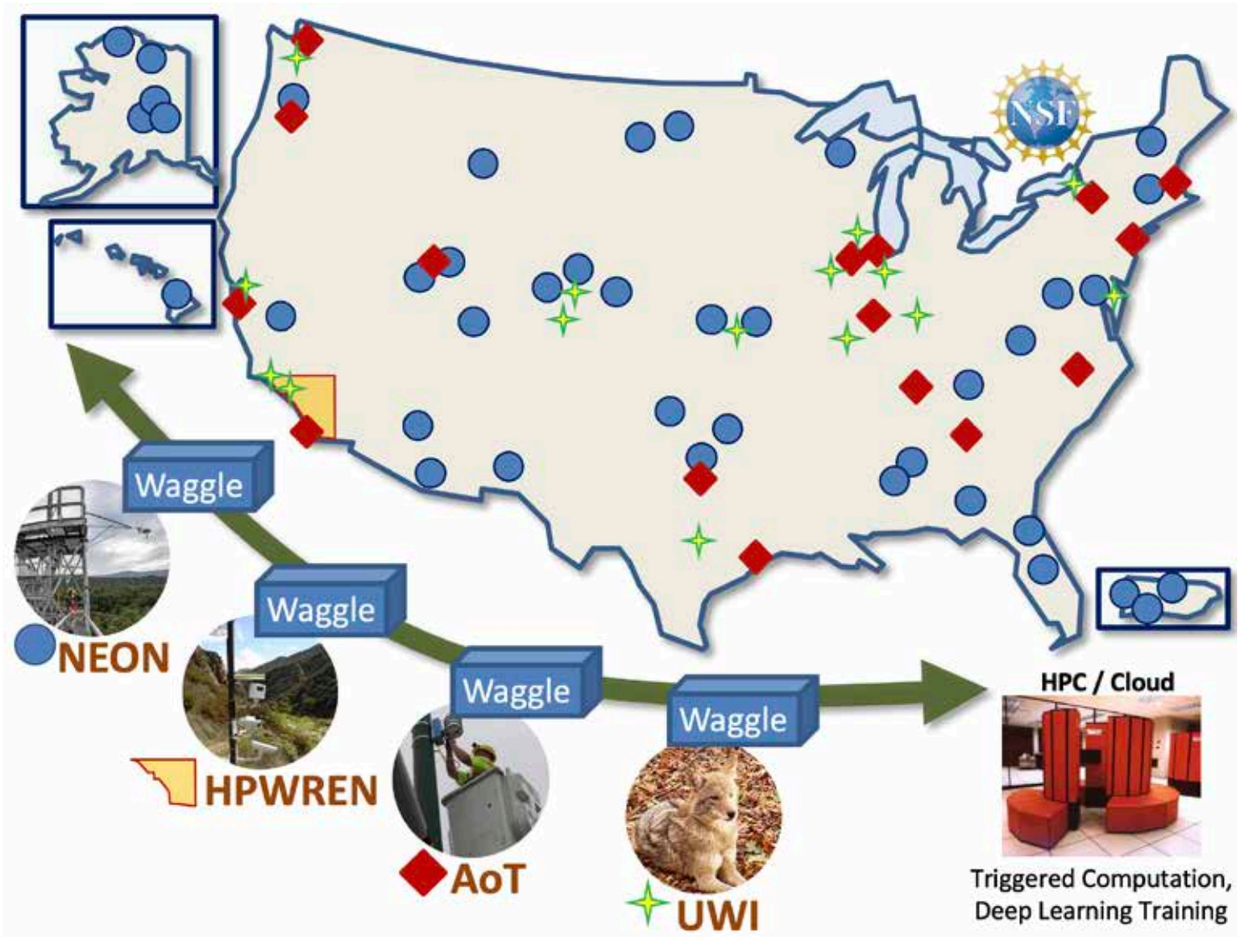

Education & Training

Leadership Team

Pete Beckman
(NU; Director)

Nicola Ferrier
(NU; Deputy Dir.

Ilkay Altintas
(SDSC; Data)

Charlie Catlett
(UIllinois; AoT)

Scott Collis
(NU; ARM)

Valerie Taylor
UChicago; Broader Impacts)

Jim Olds
(GMU; Life Sci, Risk)

Dan Reed
(Utah; Architecture)

Eugene Kelly
(CSU; NEON)

Irene Qualters
(LANL; Advisory Committee Chair)

VERSITY OF CAGO

Argonne NATIONAL LABORATORY

SAGE Partner Instruments

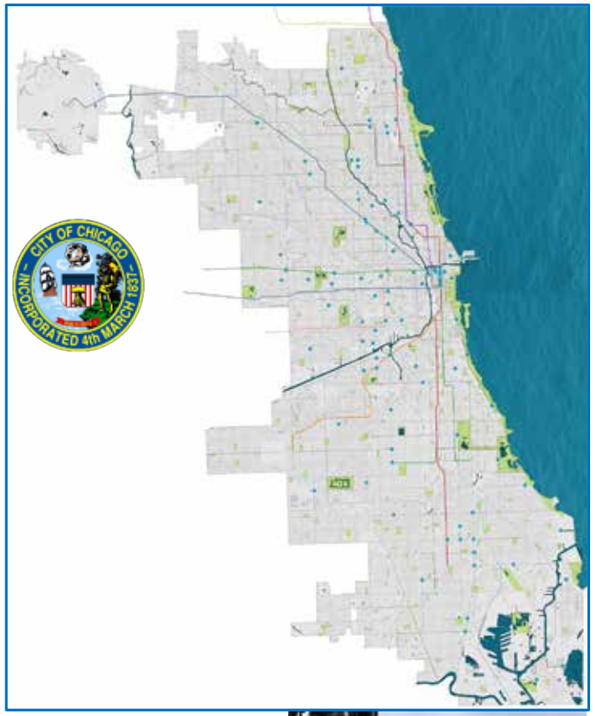

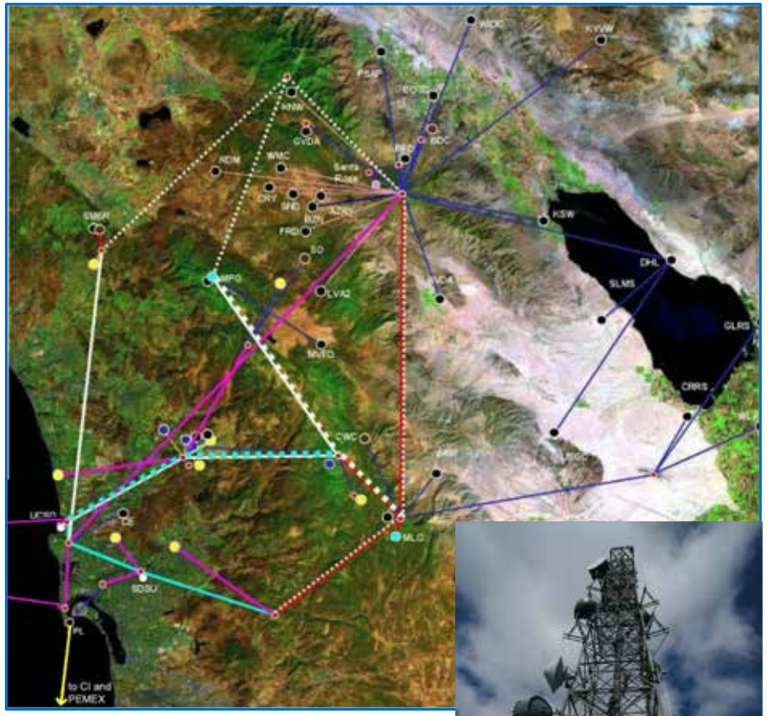

AoT: ***Neighborhood*** scale urban environment and activity.

HPWREN/WIFIRE: ***Regional*** Environmental Conditions and Events.

NEON: ***Continental*** scale ecology and environment.

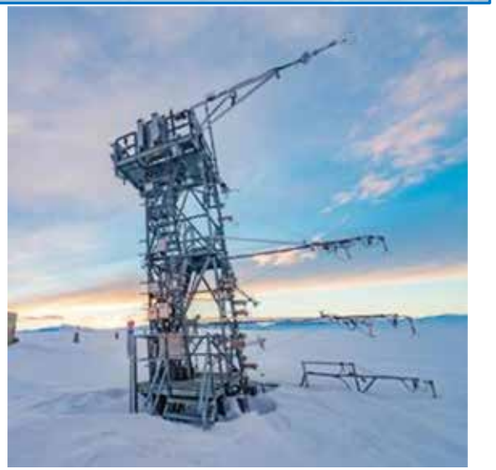

How Will Scientists Use Sage?
A User-Driven, Science Architecture:

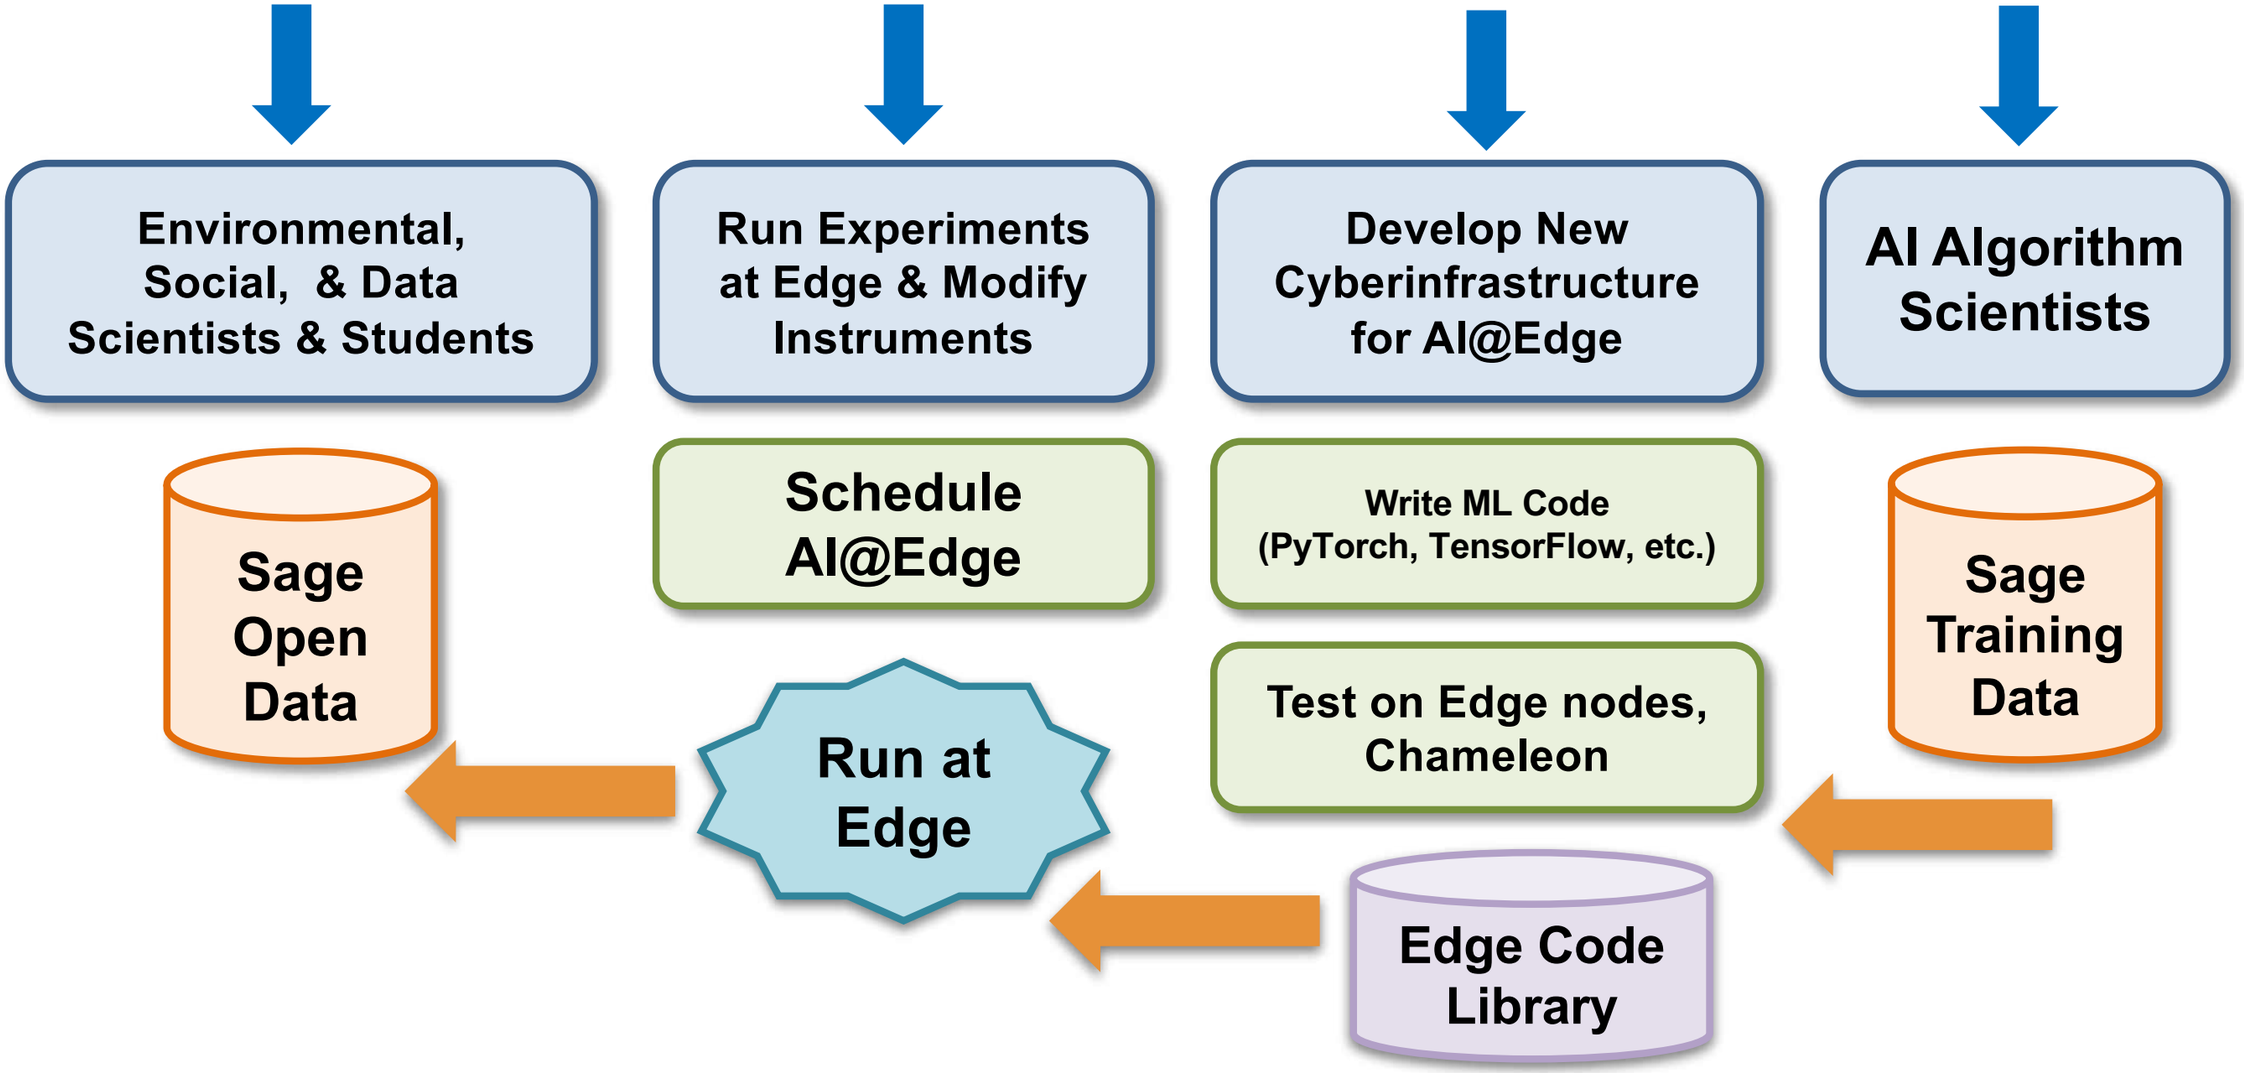

SAGE Technical Architecture

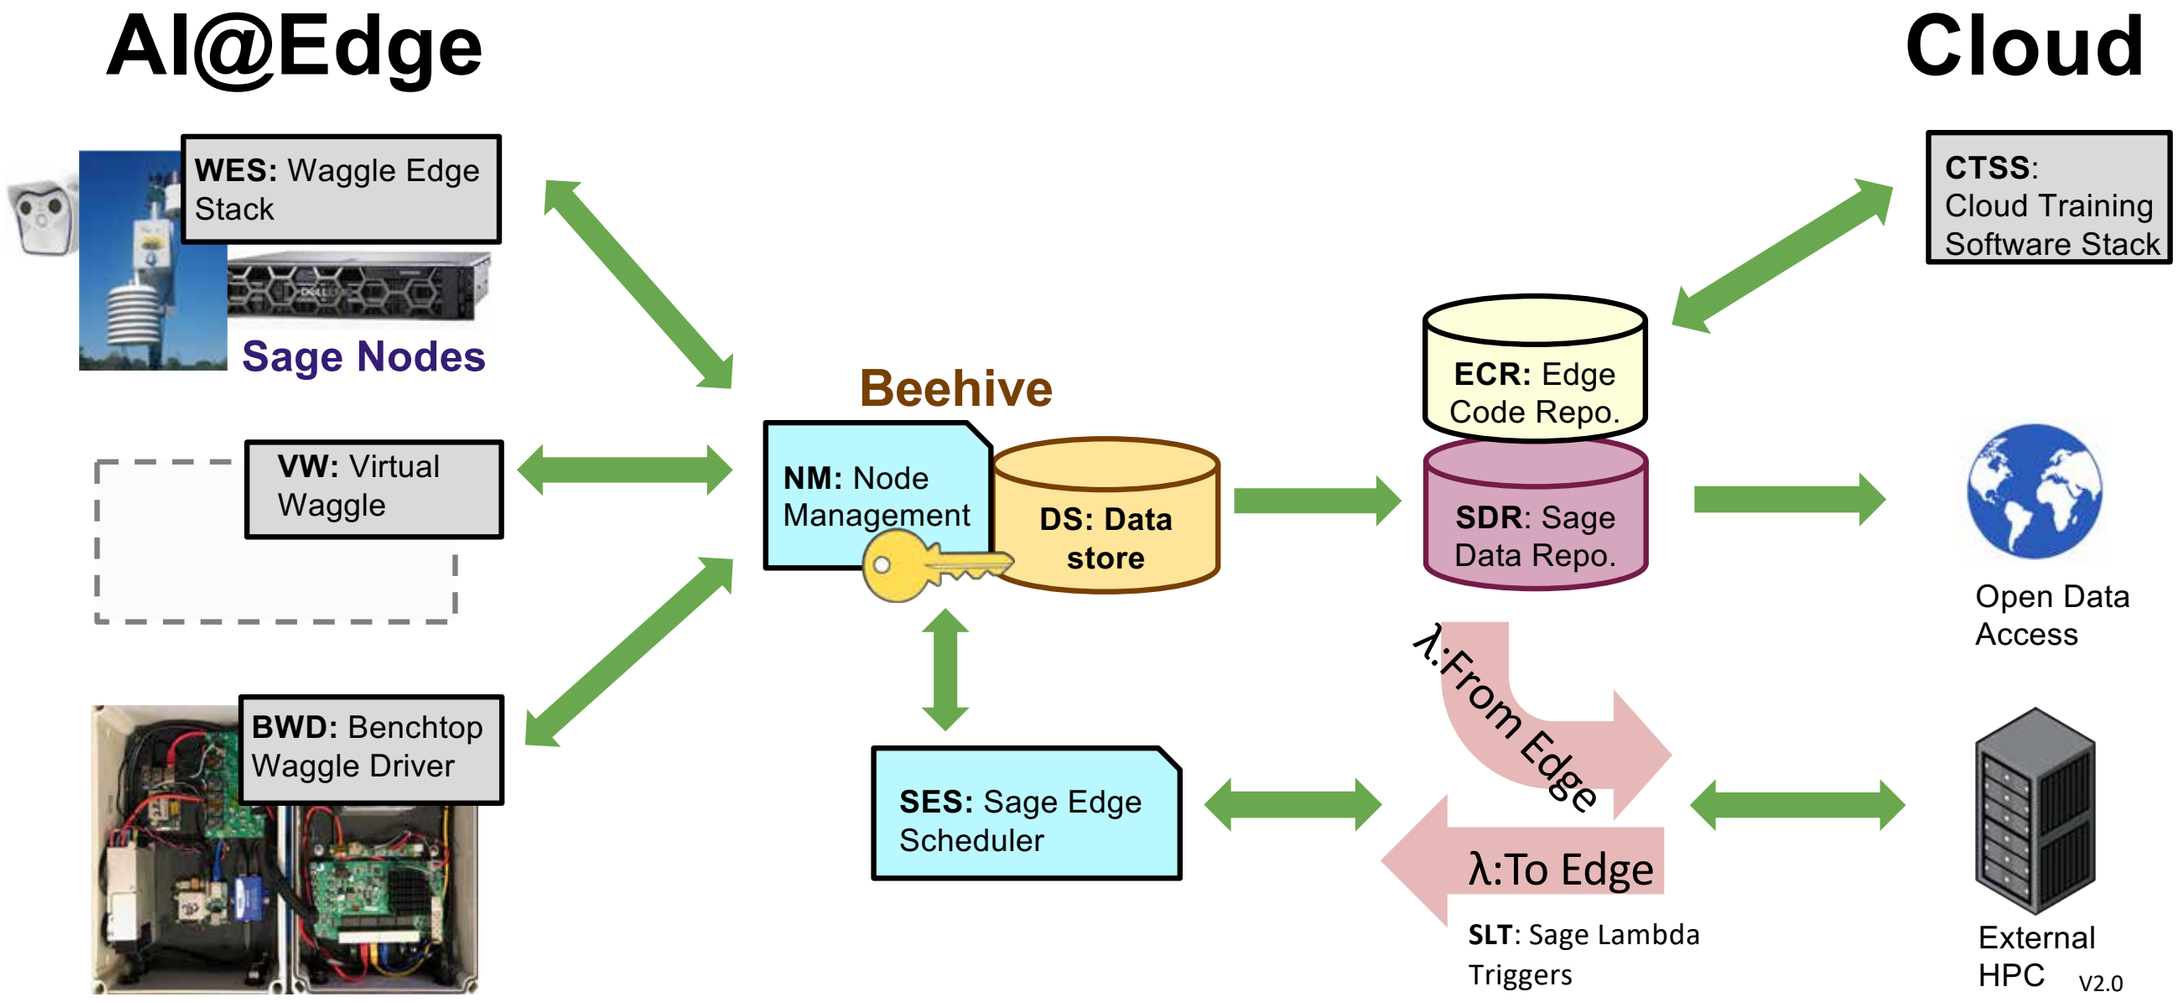

Ecology: NSF NEON & Sage

AI@HPC + AI@Edge = Intelligent Forecast & Sensing

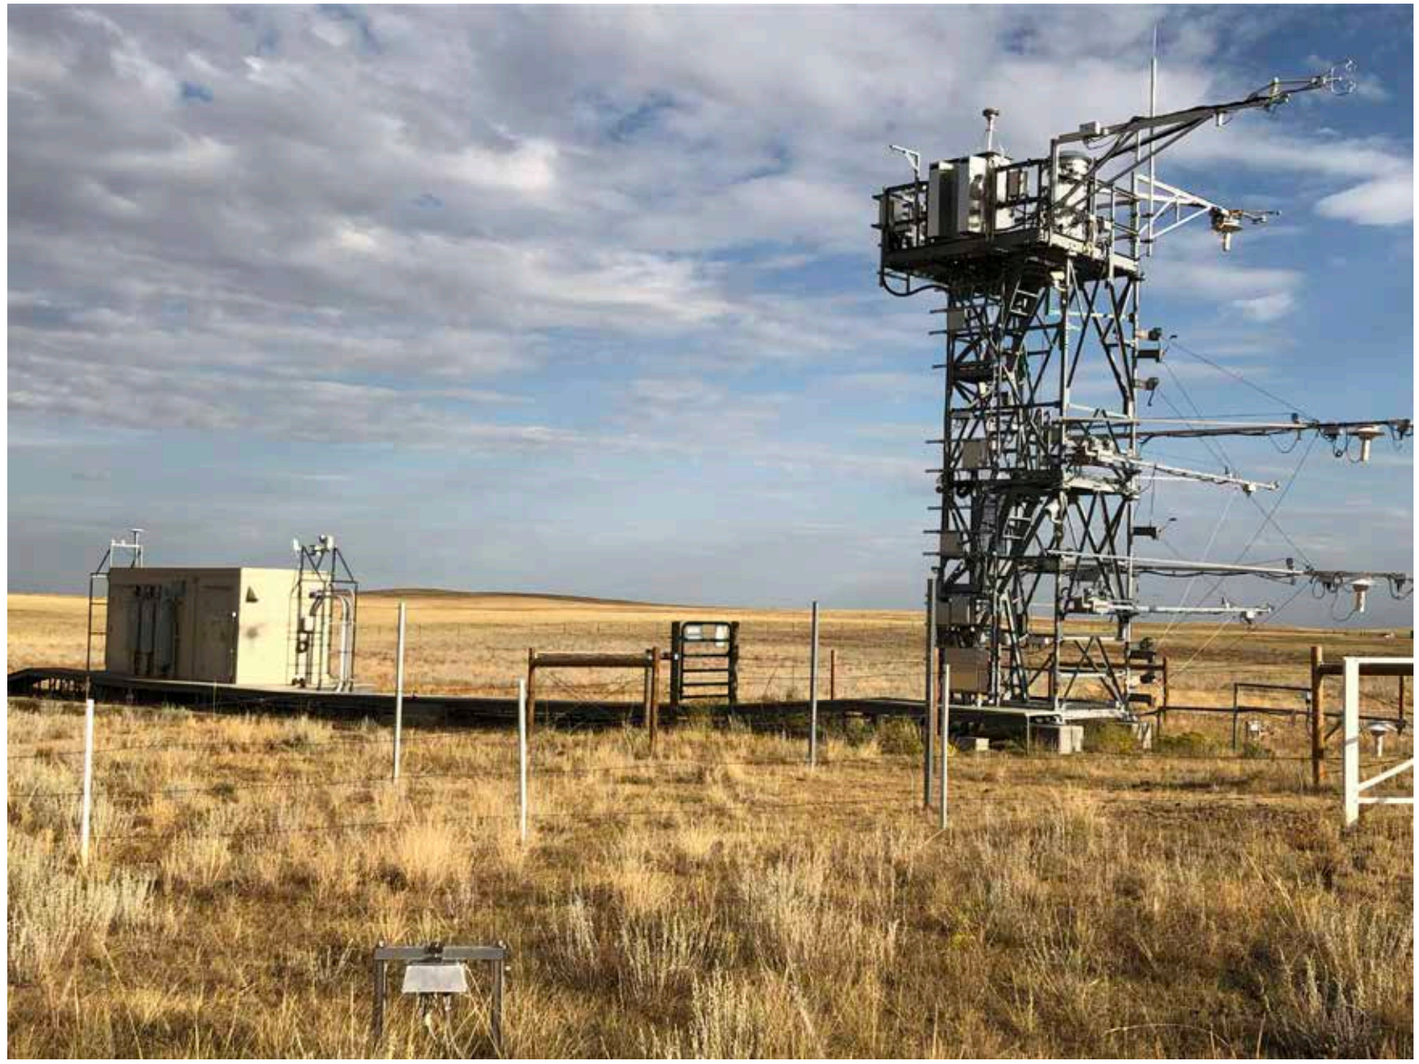

NEON: National Ecological Observatory Network. Multi-decade project to understand changing ecosystem

81 field sites, 100K data samples each year.

Sage will deploy AI@Edge to link with AI@HPC and detect interesting phenomenon and notify scientists in real time

From bats to migrating animals to clouds…

Water water everywhere

- Using AoT camera images of street flooding for confirmation and verification of flood presence and height
- Using Flood Models and Machine Learning to predict flood height based on rain amount

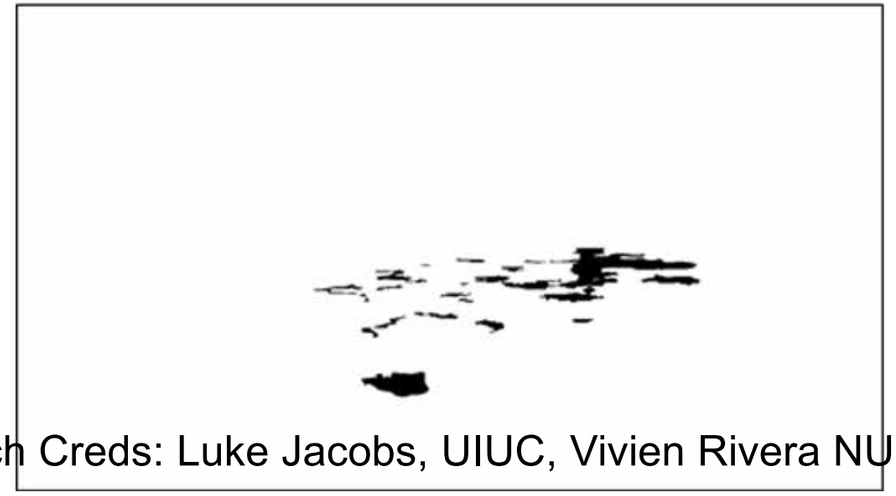

Research Creds: Luke Jacobs, UIUC, Vivien Rivera NU


NEON.D16.MART.DP1.20002 - NetCam SC IR - Thu Nov 21 2019 17:00:06 UTC
Camera Temperature: 28.5
Exposure: 1995

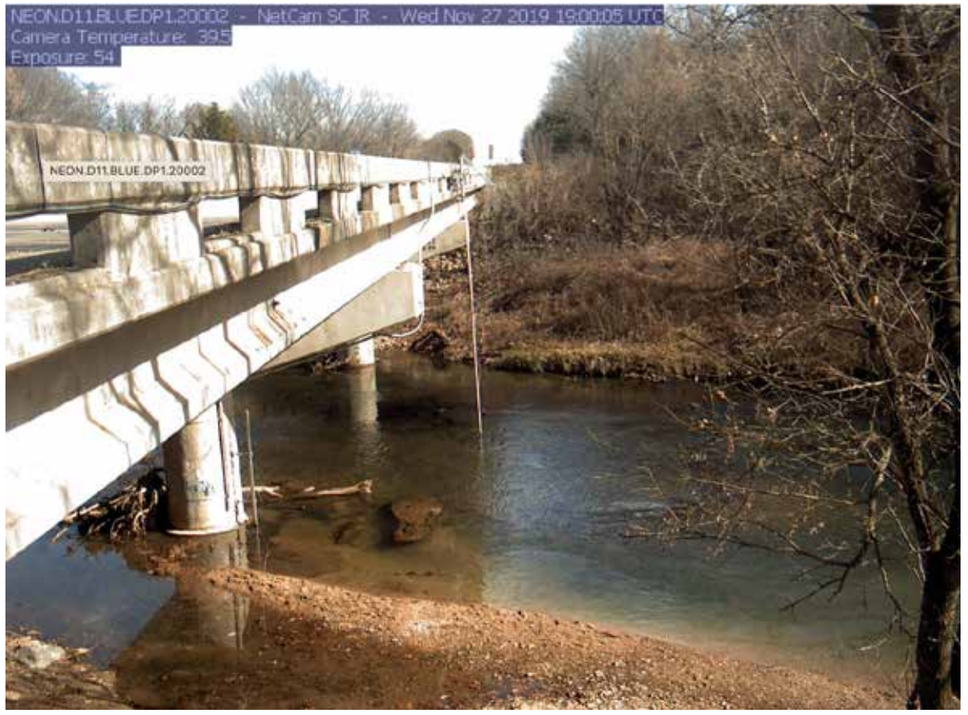
NEON.D11.BLUE.DP1.20002 - NetCam SC IR - Wed Nov 27 2019 19:00:05 UTC
Camera Temperature: 39.5
Exposure: 54
NEON.D11.BLUE.DP1.20002

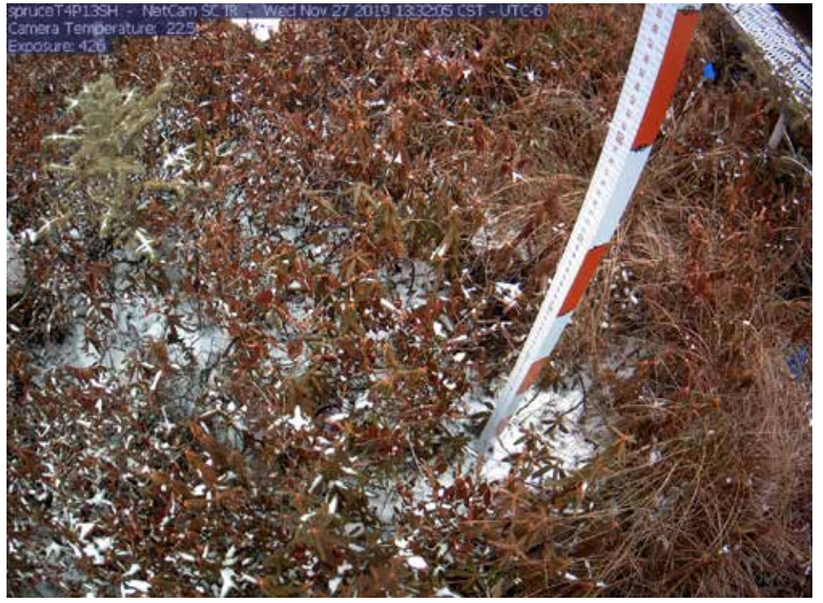
spruceT4P13SH - NetCam SC IR - Wed Nov 27 2019 13:32:05 CST - UTC-6
Camera Temperature: 22.5
Exposure: 426


NEON.D16.MART.DP1.20002 - NetCam SC IR - Tue Nov 19 2019 16:00:06 UTC
Camera Temperature: 36.0
Exposure: 1037


NEON.D13.NIWO.DP1.00042 - NetCam SC IR - Wed Nov 27 2019 19:00:06 UTC
Camera Temperature: 22.0
Exposure: 2
NEON.D13.NIWO.DP1.00042


NEON.D18.TOOL.DP1.00042 - NetCam SC IR - Fri Nov 15 2019 01:00:05 UTC
Camera Temperature: -1.5
Exposure: 1532

Social Distancing Detector

Ori Zur, Northwestern University ’22
Nicola Ferrier, Scott Collis, Argonne National Laboratory

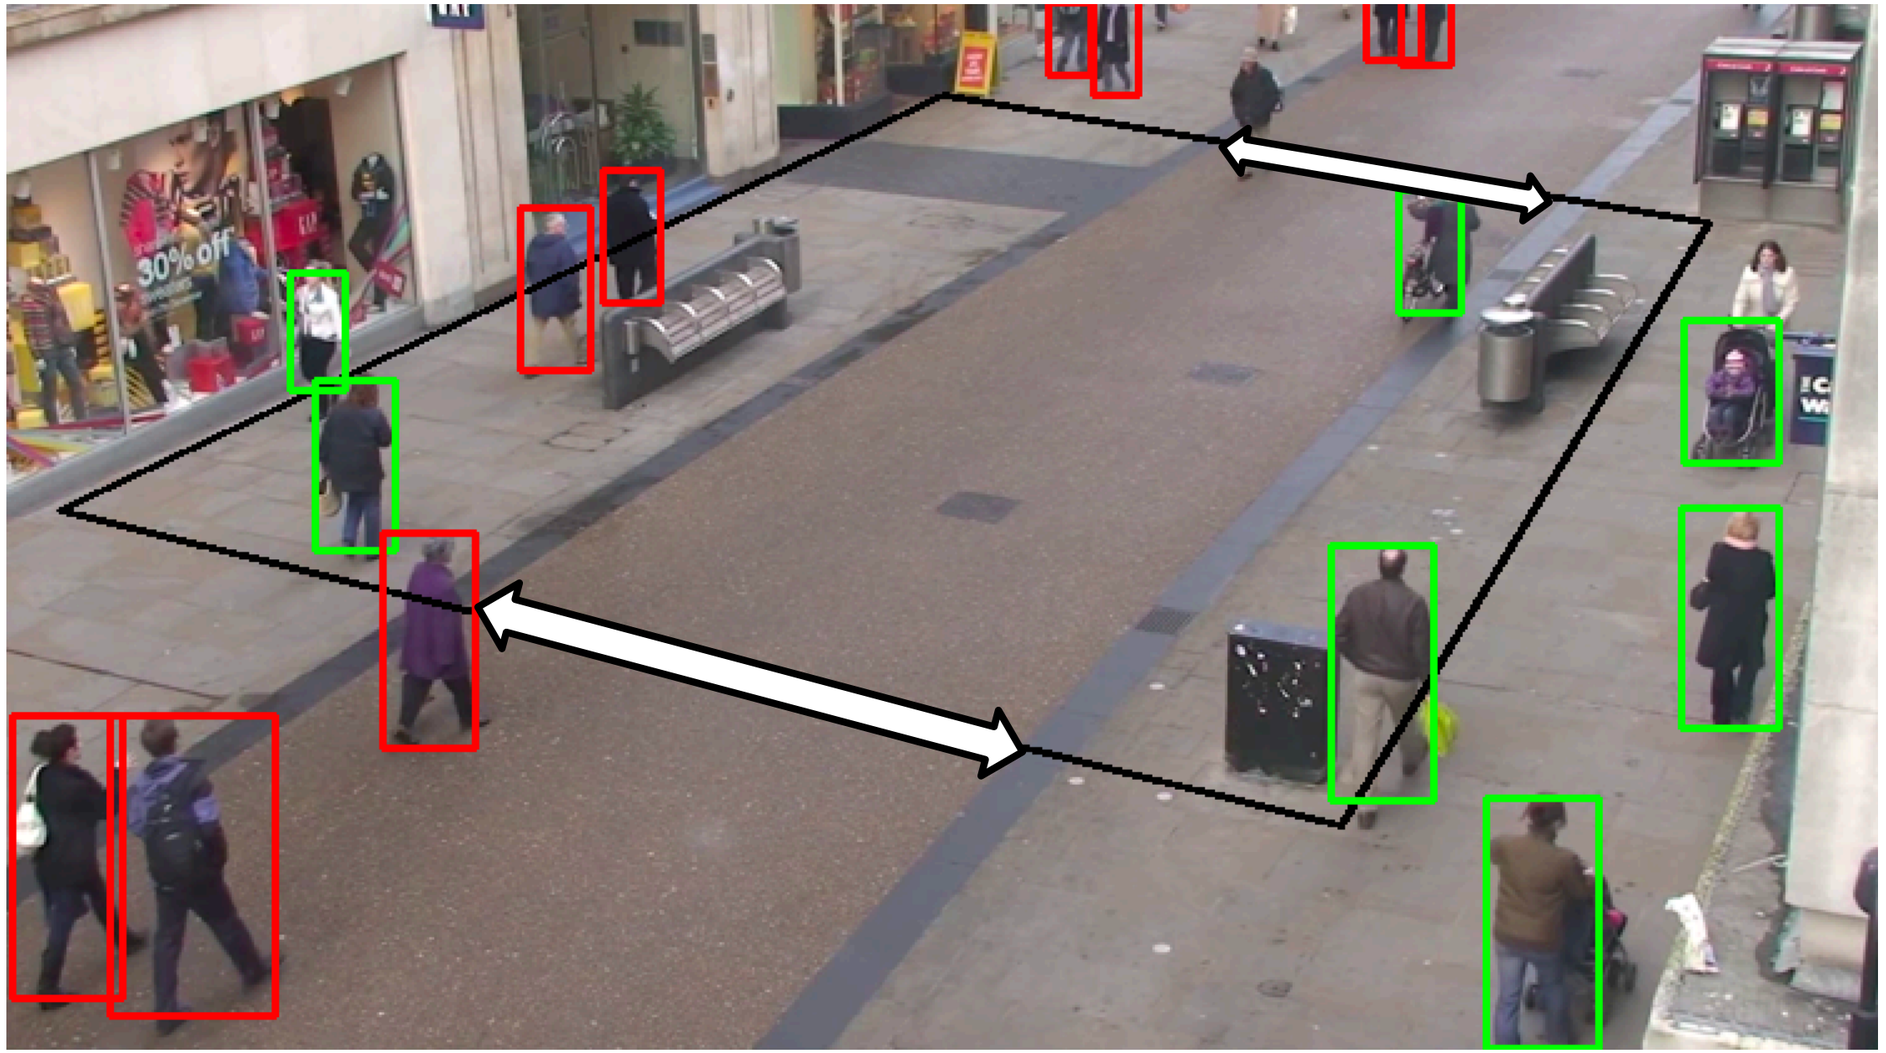

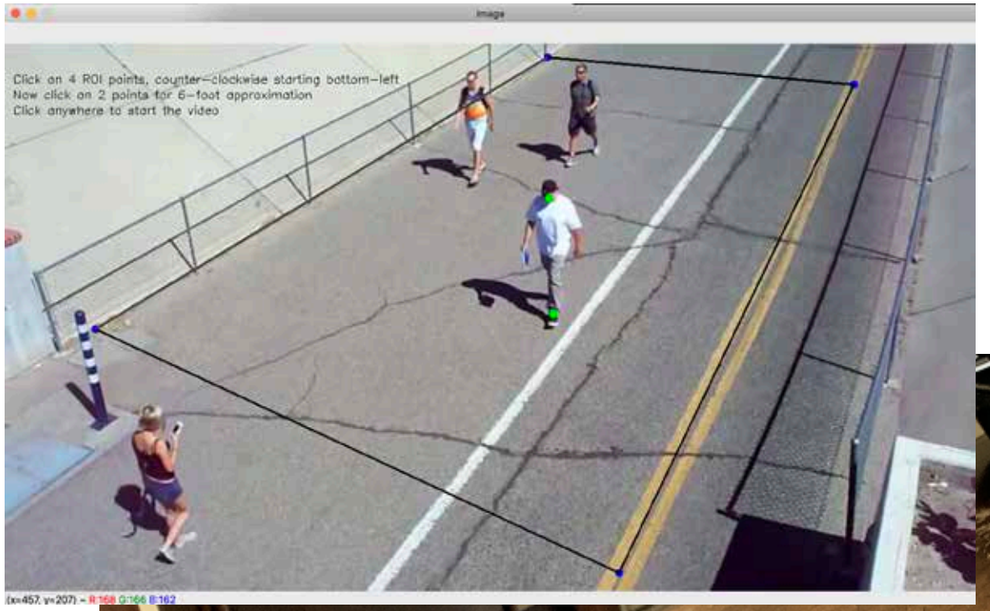
Image
Click on 4 ROI points, counter-clockwise starting bottom-left
Now click on 2 points for 6-foot approximation
Click anywhere to start the video


1 square = 1/4" x 1/4"

Operation– Person Detection (YOLOv3)

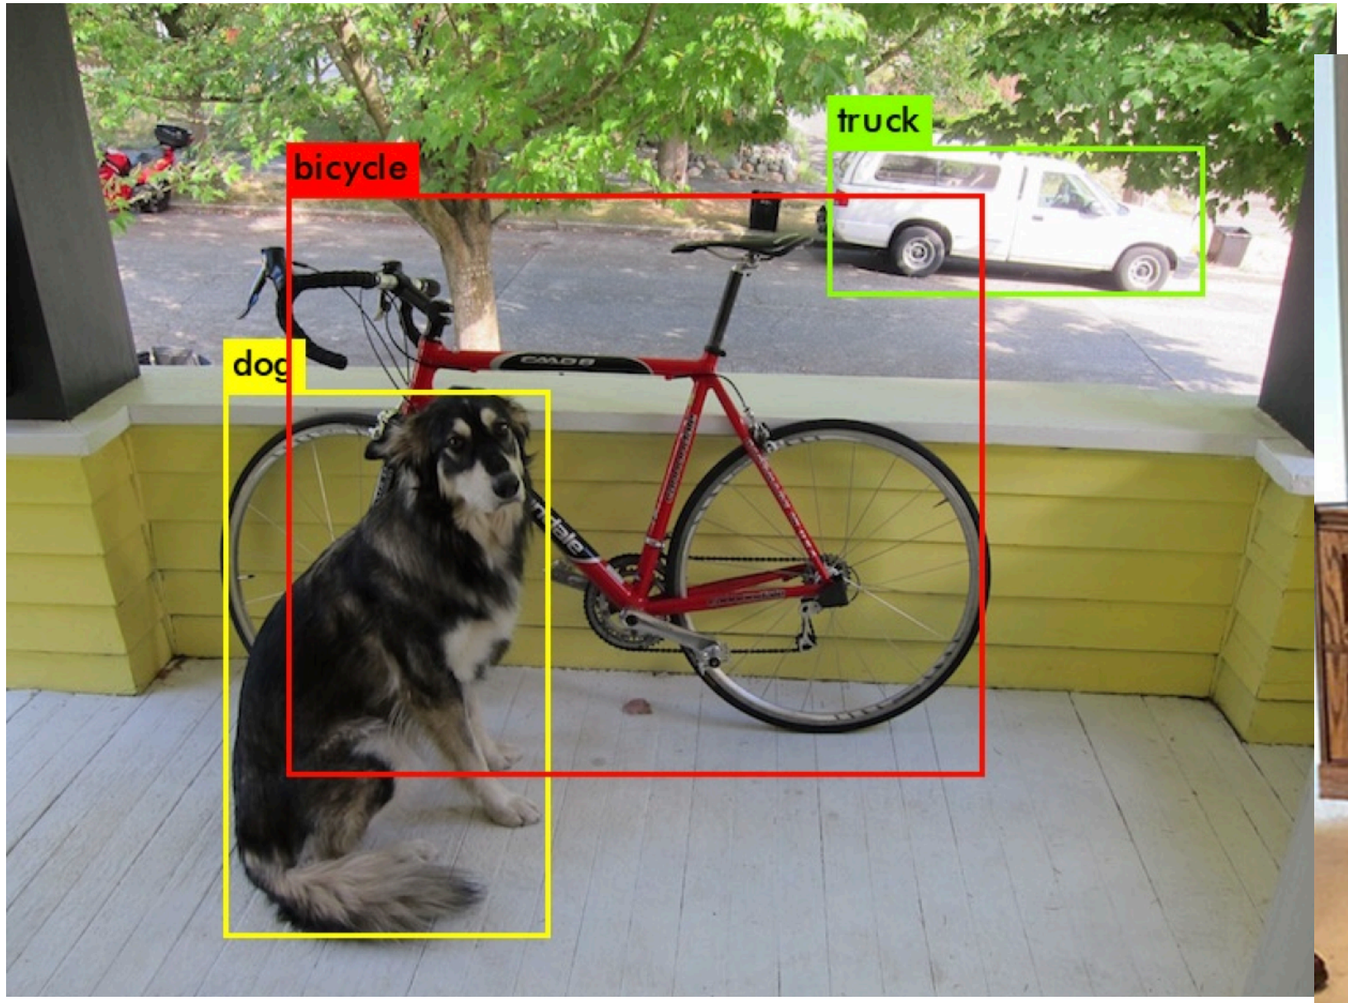

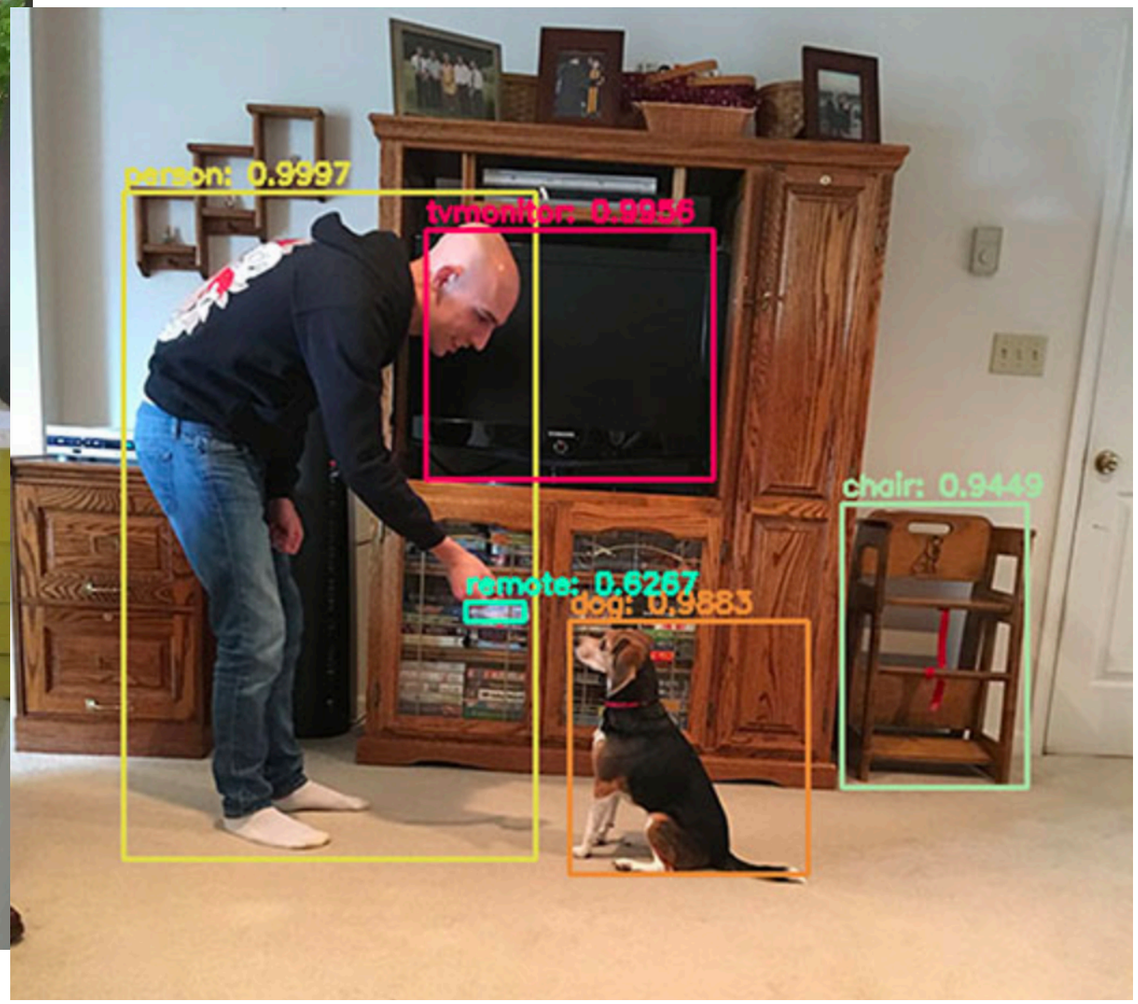

Detection of Wildfire

Wildfire: AI@HPC + AI@Edge = Intelligent Forecast & Sensing

Image from HPWREN tower, where Sage will deploy AI@Edge for real-time fire detection

SDSC WIFIRE project runs HPC simulation to predict wildfire

Sage project will move Pan-Tilt-Zoom cameras toward suspected outbreaks and run AI@Edge to search for fire


20200906.060052 Mt. Wilson Observatory East, http://hpwren.ucsd.edu a4
HPWREN

Early detection of smoke from wildfires is critical for societal and environmental well-being

Aristana Scourtas, Northwestern University, Masters AI
Nicola Ferrier, Argonne National Laboratory
Ilkay Altintas, UCSD

(Alkhatib, 2014) *A Review on Forest Fire Detection Techniques*

Our data

HPWREN	number of datasets	percent of datasets	number of images	percent of images
total	232	100	18215	100
before fire	NA	NA	9031	49.58001647
after fire	NA	NA	9184	50.41998353

How do humans process smoke?

Look for movement!

Combined CNN-LSTM

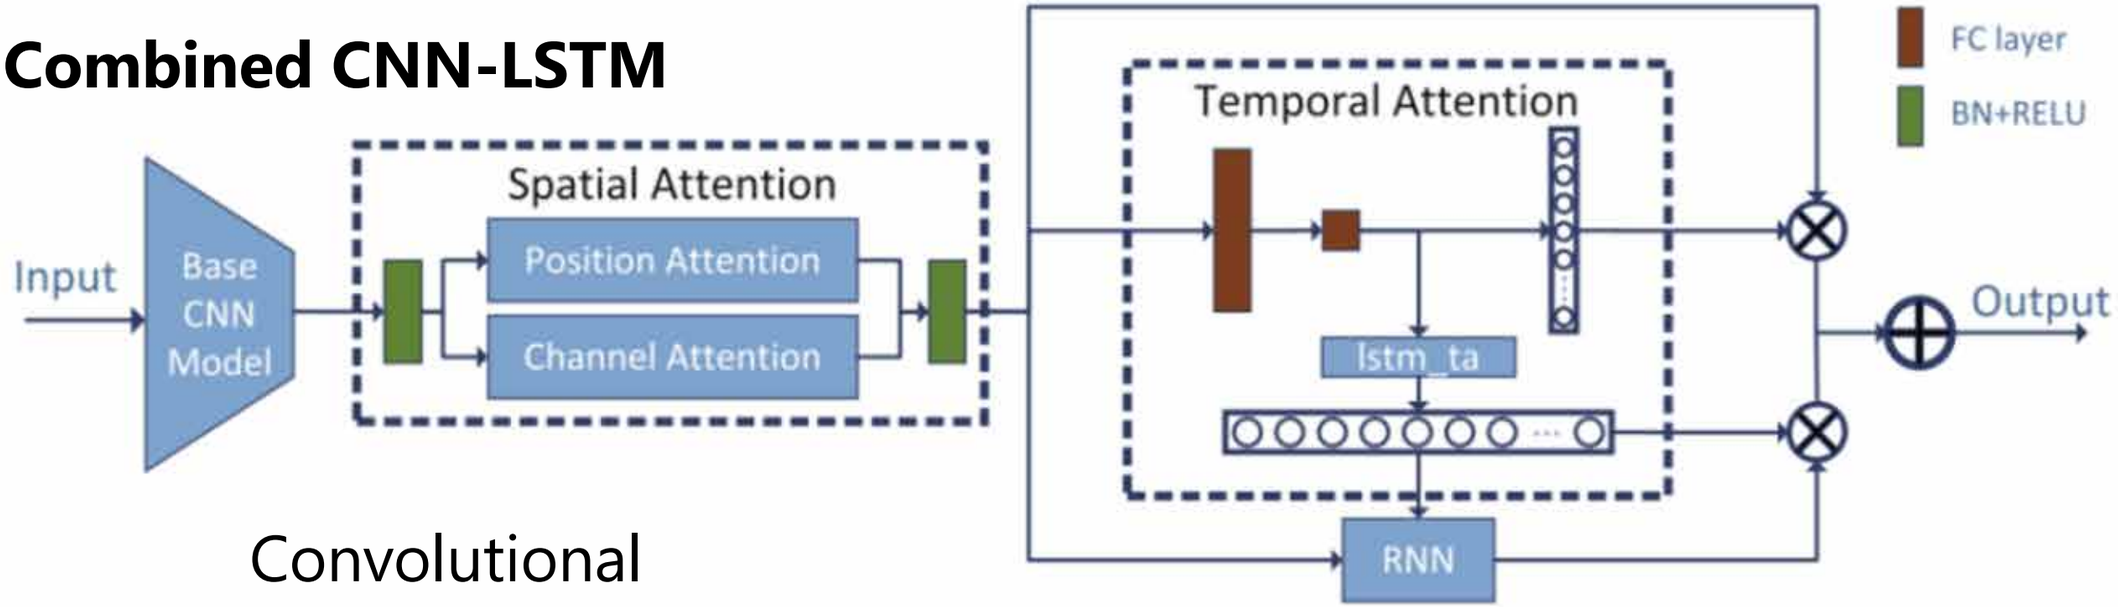

Convolutional Neural Network

Long Short-Term Memory network

Fig. 2. The illustration of our proposed spatial-temporal fusion network based on residual learning and spatial-temporal attention. The spatial-temporal fusion network consists of four main parts: Base CNN Model, spatial attention mechanism, temporal attention mechanism and residual learning of spatial-temporal features.

Deep learning training pipeline

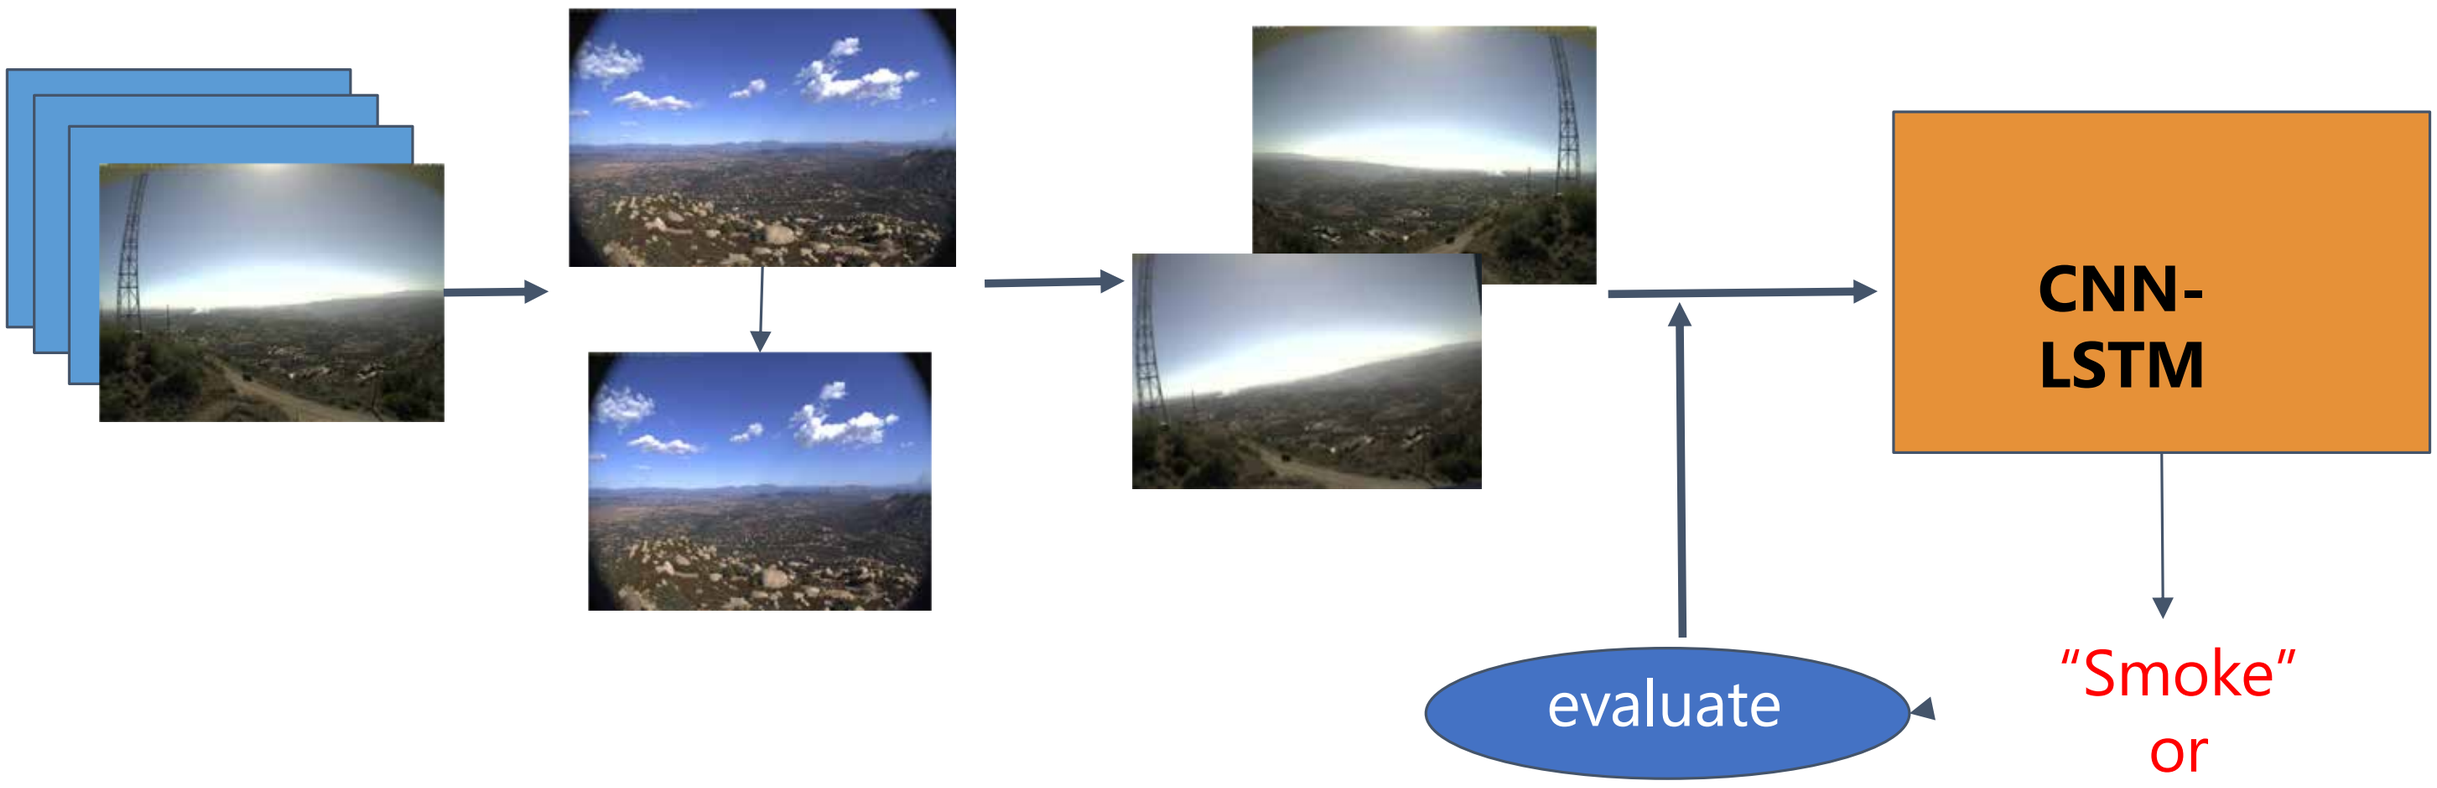

Snowflake Classification

Neelanshi Varia, Northwestern, AI Masters
Aaron Kennedy (UND), Scott Collis, Argonne National Laboratory

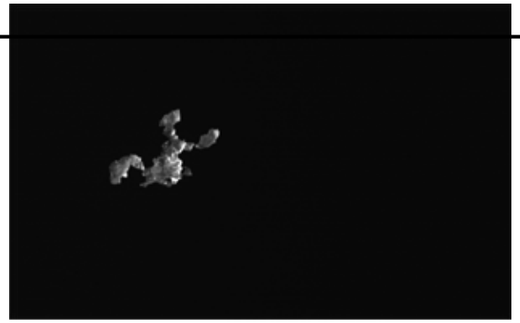

Type: Aggregate
images: 67

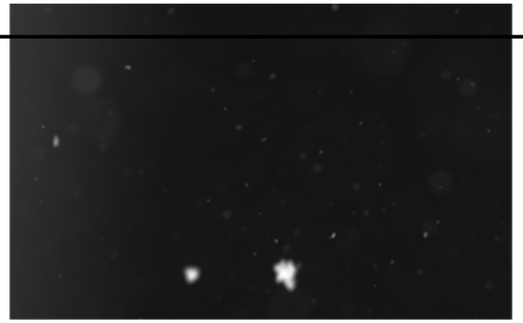

Type: Blowing snow chaos
images: 10

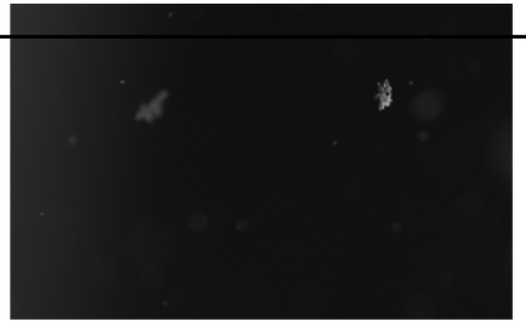

Type: Dendrite stellar
images: 9

Type: Frozen drop weird melty
images: 7

Type: Heavy rime flakes
images: 9

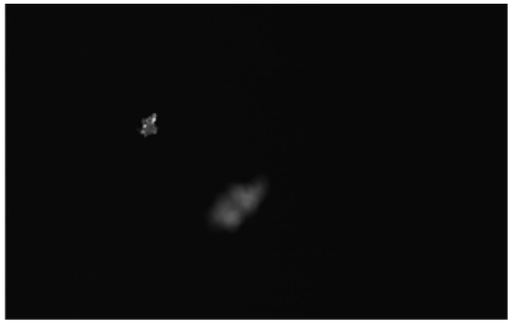

Type: Melted flake
images: 28

Model selection and overview

- Tried on current data:
  - SVM
  - Neural network with 3 hidden layers
  - CNN (inspired by A. Hicks paper) => Best accuracy so far

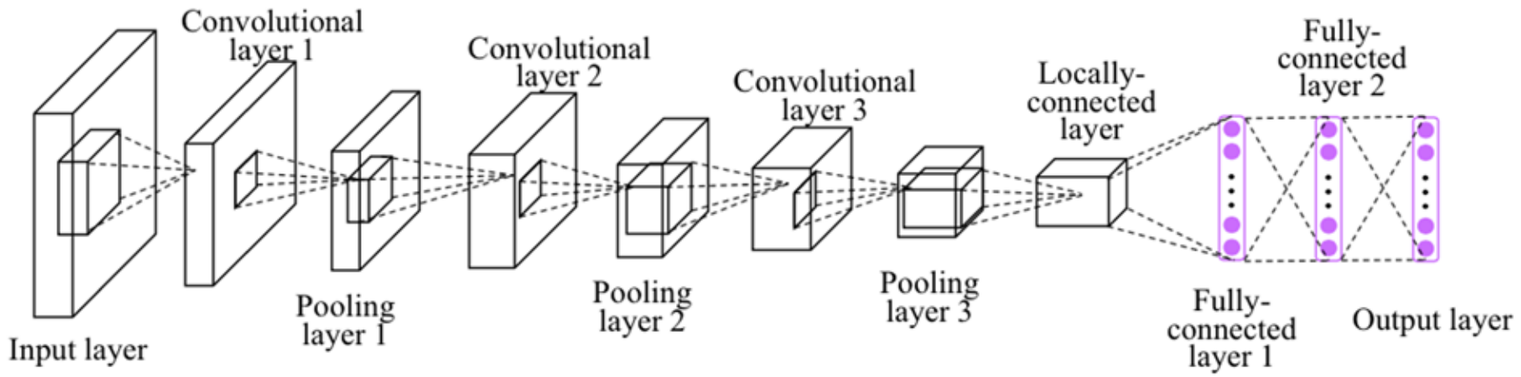

- To try on new data (for multi-class classification):
  - Deep-CNN
  - FCN

Science Use Case: ARM's Doppler LIDARs

- The DOE ARM Program uses scanning Doppler LIDARs which measure backscatter off aerosols, clouds (where it extinguishes fast) and rain drops.
- This is intended to document the "clear air" dynamics.
- At the raw level the data collected is "in phase and quadrature components" or IQ data. This is used to derive a spectrum of the doppler motions.
- The on-board processing picks the highest peak to return moments. This makes a lot of often invalid assumptions about phenomena.

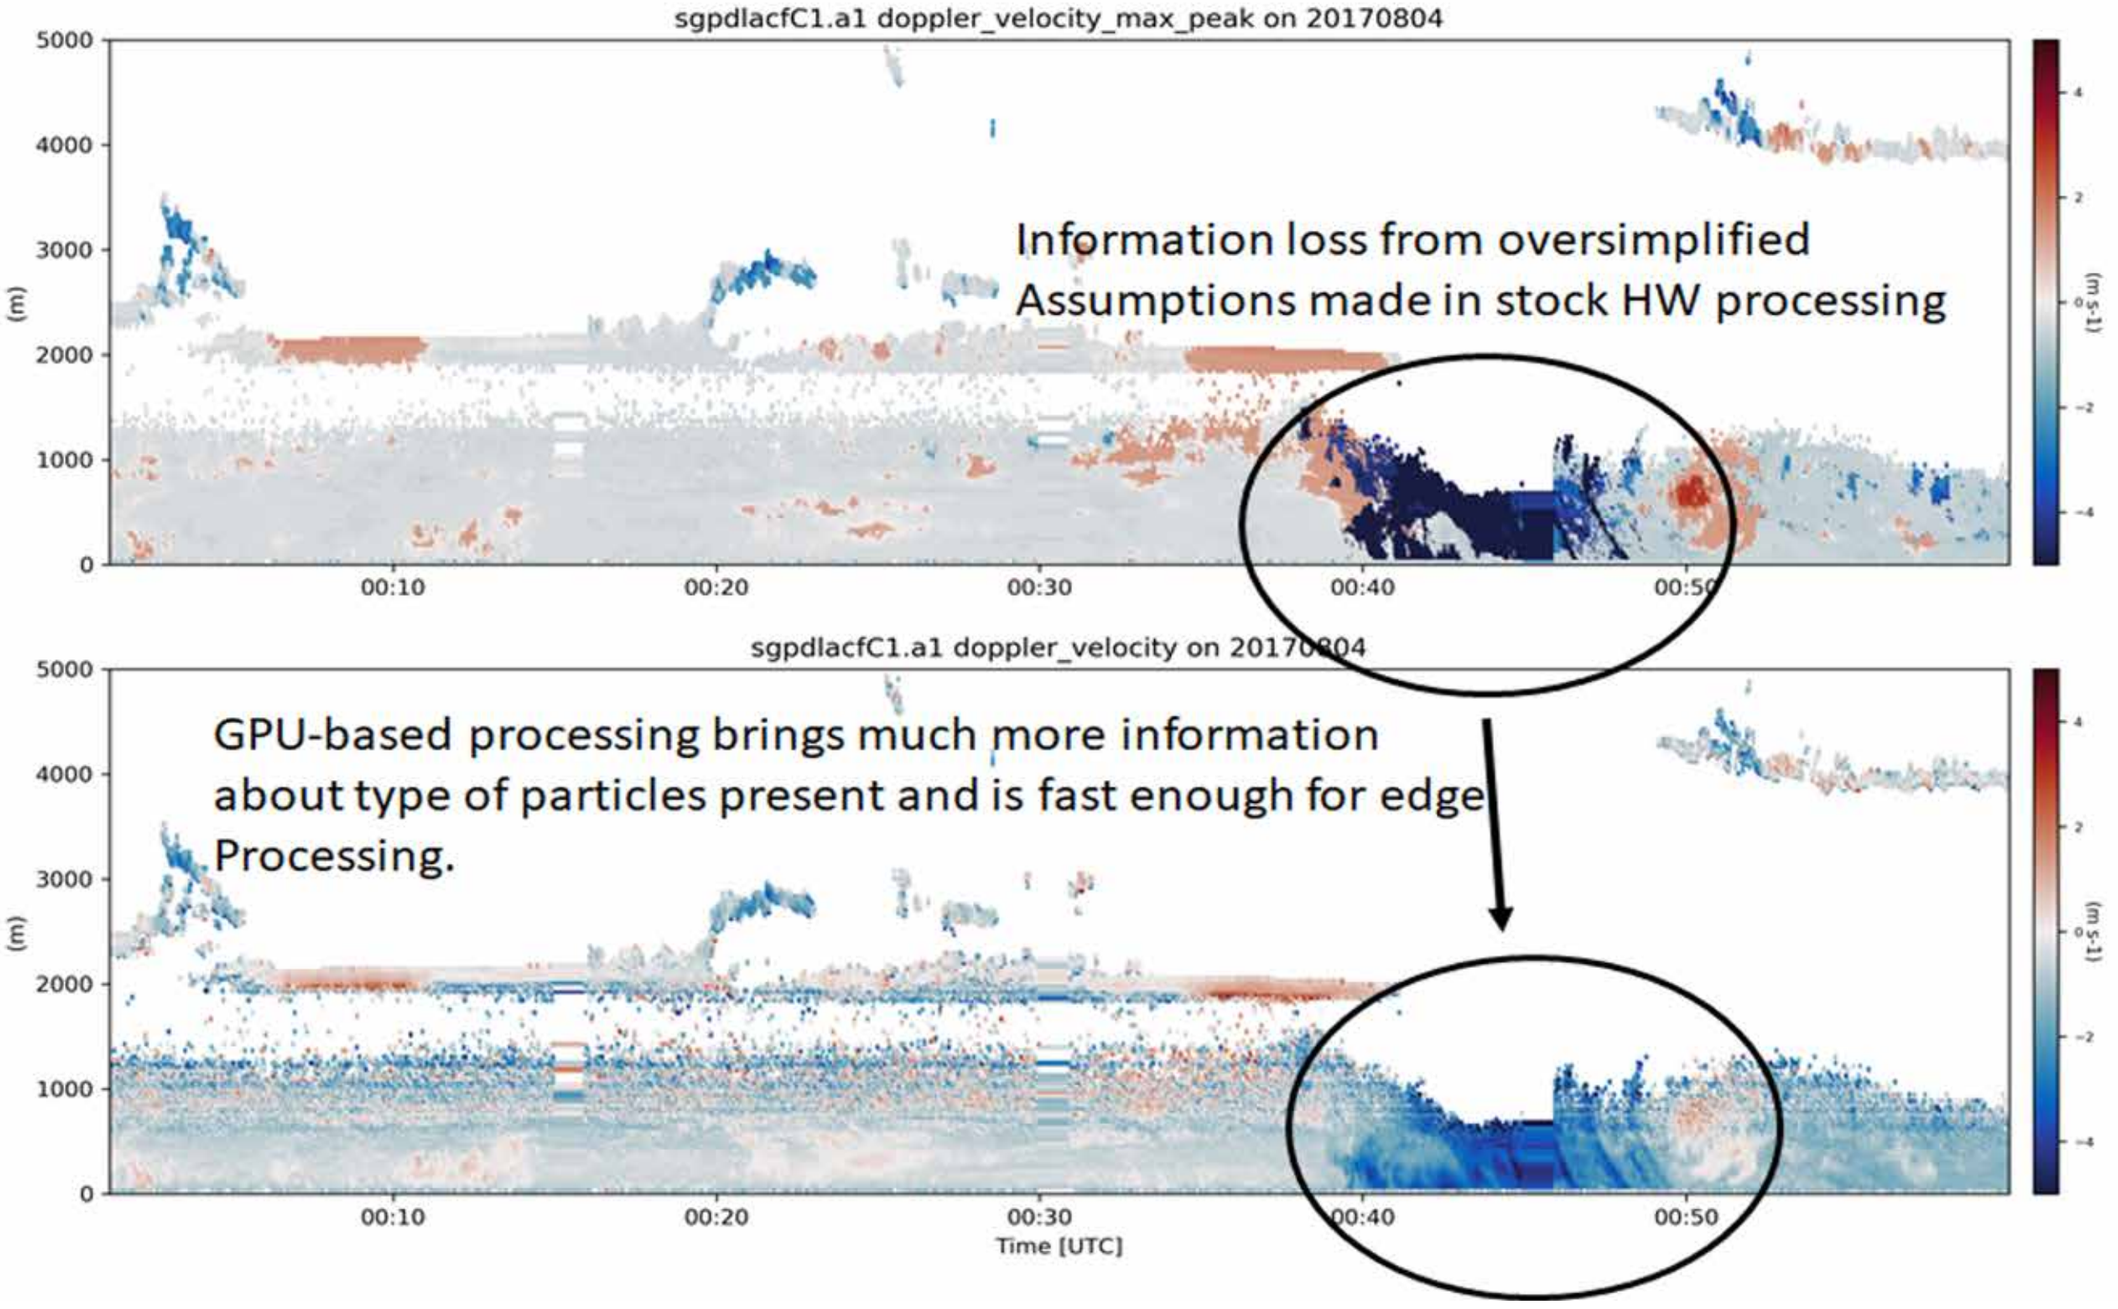
sgpdlacfC1.a1 doppler_velocity_max_peak on 20170804
Information loss from oversimplified
Assumptions made in stock HW processing
sgpdlacfC1.a1 doppler_velocity on 20170804
GPU-based processing brings much more information
about type of particles present and is fast enough for edge
Processing.
(m)
Time [UTC]
(m s-1)

Bonus Example
Pedestrian Flows

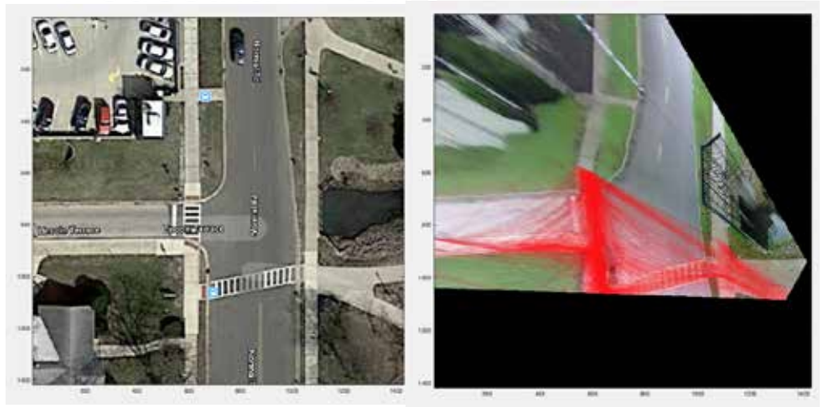

National Chung Hsing University
Taichung City, Taiwan

Source: Mike Papka, Argonne National Laboratory and Northern Illinois University,
Pratool Bharti (Northern Illinois University), and David Koop (Northern Illinois University)

SAGE
A Software-Defined Sensor Network
Cyberinfrastructure for AI at the Edge
sagecontinuum.org

NIU Northern Illinois University
Northwestern University
THE UNIVERSITY OF CHICAGO
Argonne NATIONAL LABORATORY

Many Science Problems….

Wildfires: detecting smoke

Urban flooding

Plant Species

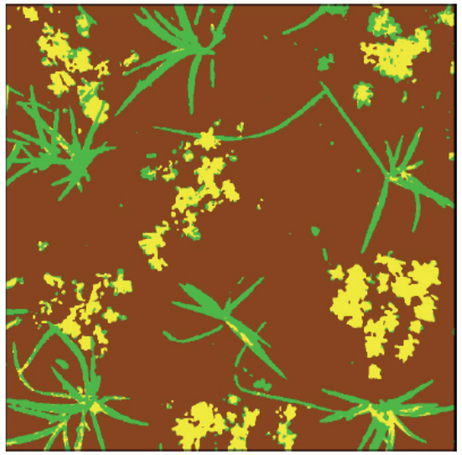

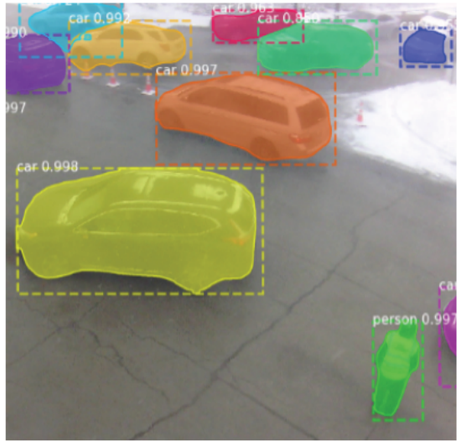

Traffic Flow

Pedestrian Flow

Wildlife

Drone detection

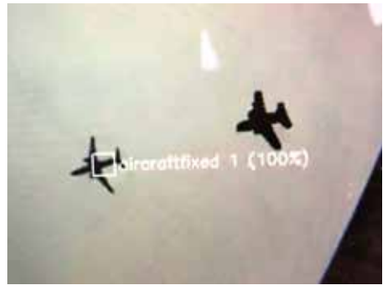

Cloud Coverage

Snow Depth

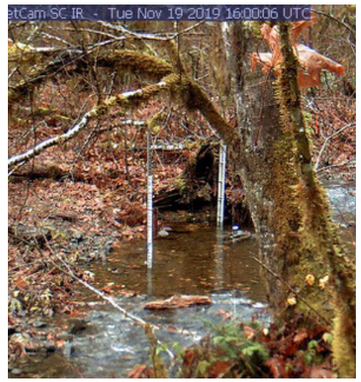

Water Depth

SAGE
A Software-Defined Sensor Network
Cyberinfrastructure for AI at the Edge
www.sagecontinuum.org

The Future: Programming The Computing Continuum

IoT/Edge ⟶ **HPC/Cloud**

Size	Nano	Micro	Milli	Server	Fog	Campus	Facility
Example	Adafruit Trinket	Particle.io Boron	Array of Things	Linux Box	Co-located Blades	1000-node cluster	Datacenter
Memory	0.5K	256K	8GB	32GB	256G	32TB	16PB
Network	BLE	WiFi/LTE	WiFi/LTE	1 GigE	10GigE	40GigE	N*100GigE
Cost	$5	$30	$600	$3K	$50K	$2M	$1000M

Count = 10^9
Size = 10^1




Count = 10^1
Size = 10^9

Continuum Abstract Model & Runtime

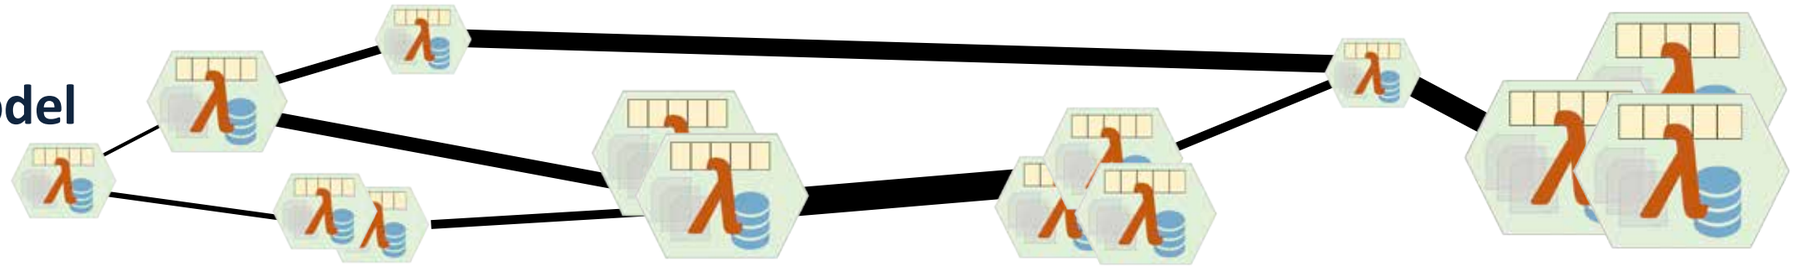

Linking Autonomic and Self Organizing at the Edge to Managed HPC Resources


Many measurements cannot be “sensed” directly but can be computed from image, microphone or other devices
How can we build intelligent systems that pursue goals?

Questions?

SAGE

A Software-Defined Sensor Network

Cyberinfrastructure for **AI at the Edge**

sagecontinuum.org