SAGE
sagecontinuum.org

NSF

A Software-Defined Sensor Network
Cyberinfrastructure for AI@Edge Computing

Sage Digital Continuum

Scientific Edge Computing: Linking Instruments to Supercomputers

Pete Beckman: Co-Director Northwestern Univ / Argonne Inst. for Science and Engineering
Collaborators: Ilkay Altintas, Charlie Catlett, Scott Collis, Nicola Ferrier, Kate Keahey, Euguene Kelly, Jim Olds, Mike Papka, Dan Reed, Raj Sankaran, Sean Shahkarami, Joe Swantek, Valerie Taylor, Doug Toomey, Frank Vernon, Rommel Zulueta, and many more….

Northwestern University
THE UNIVERSITY OF CHICAGO
Northern Illinois University
Argonne NATIONAL LABORATORY
Colorado State University
UC San Diego
GEORGE MASON UNIVERSITY
THE UNIVERSITY OF UTAH
SAGE sagecontinuum.org
ARRAY THINGS
ARM
neon National Ecological Observatory Network
HPWREN SDSC SAN DIEGO SUPERCOMPUTER CENTER
LINCOLN PARK ZOO. FOR WILDLIFE. FOR ALL.
OHAZ

Instrument

Data

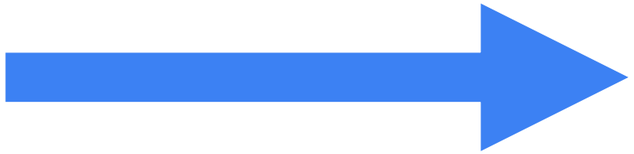

Analysis

The Digital Continuum

Instrument → **HPC/Cloud**

IoT **Facilities** **Analysis**

Analyse full resolution data, find highest valve data for the science

AI@Edge: Digital Continuum

Sensors

LIDAR

Software Defined Radios

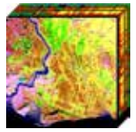

Hyperspectral Imaging

Facilities

Actuators

Servos

Dynamic adaptation

Edge Computing

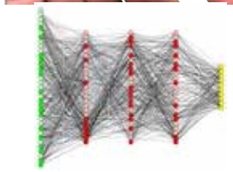

Scientific Data Analysis & Control

Artificial Intelligence
Deep Learning Inference
Lightweight Training
Autonomous Action

Advanced Networking

New inference (model)
Adaptive controls / steering

Computation

Predictive Sim
Digital Twins
Data Analysis
Machine Learning

SAGE
sagecontinuum.org

Why Live on the Edge?

- More data than bandwidth
  - Imaging, LIDAR, SW defined radios, radar, audio, hyperspectral, large facilities, …
- Latency is important
  - Quick local decision, experimental control & actuation; adaptive sensing
- 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-making
- Quiet observation and energy efficiency
  - Vigilant low-power sensors, transmit only essential observations

SAGE

Cyberinfrastructure for
AI at the Edge
sagecontinuum.org

AI@Edge Summer 2022
(Student Outing: June 2022)

Leadership Team


Pete Beckman
(NU: Director)


Nicola Ferrier
(UC: Deputy Dir.)


Scott Collis
(NU: Instruments, Atmos)


Valerie Taylor
UC: Edu, Broader Impacts)


Eugene Kelly
(CSU; Ecosys, NEON)


Mike Papka
NIU: Edu, Broader Impacts


Raj Sankaran
NU: Node Arch


Ilkay Altintas
(SDSC: Data)


Charlie Catlett
(UIllinois: Urban)


Jim Olds
(GMU; Life Sci, Risk)


Dan Reed
(Utah; Architecture)


Kathy Bailey
Proj Mgmt


Helen Taaffe
Proj Mgmt


Joe Swantek
NU: Software


Irene Qualters
(LANL; Advisory Committee Chair)

MSRI-1: 1935984
Start: October 1, 2019

Exciting, Hard, Challenging, CS Problems:

From Instrument to the HPC/Cloud

- Reusable cyberinfrastructure for AI@Edge is new territory
- Digital Continuum programming models largely unexplored
- How can we build triggered simulations and autonomous reactions?
- Edge computing needs multi-tenancy for computation and actuation
- Remote, distributed instruments have unique cyberinfrastructure needs
- New AI methods and algorithms for the Edge
- New resource management for Science Goals
- Experimental cyberinfrastructure must first “do no harm” to operational facilities

Building on concepts from NSF Array of Things

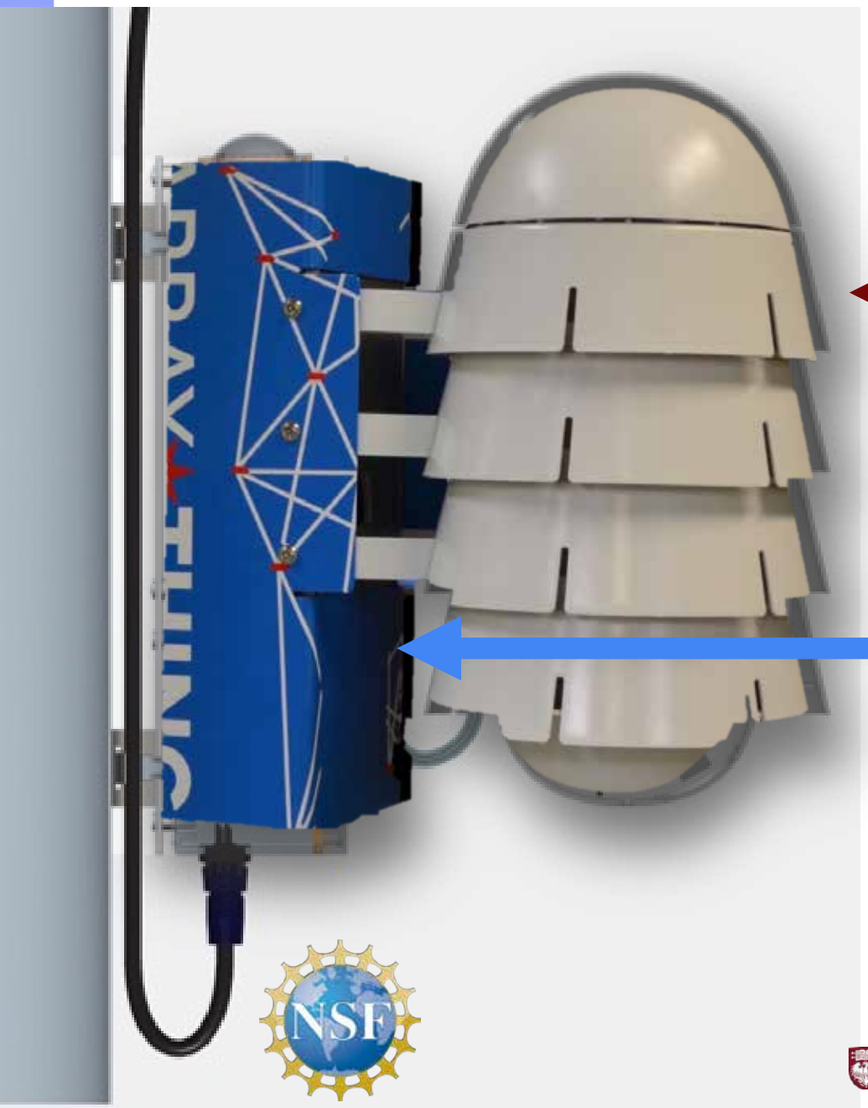

Environmental and Air Quality Sensors

Edge Computing

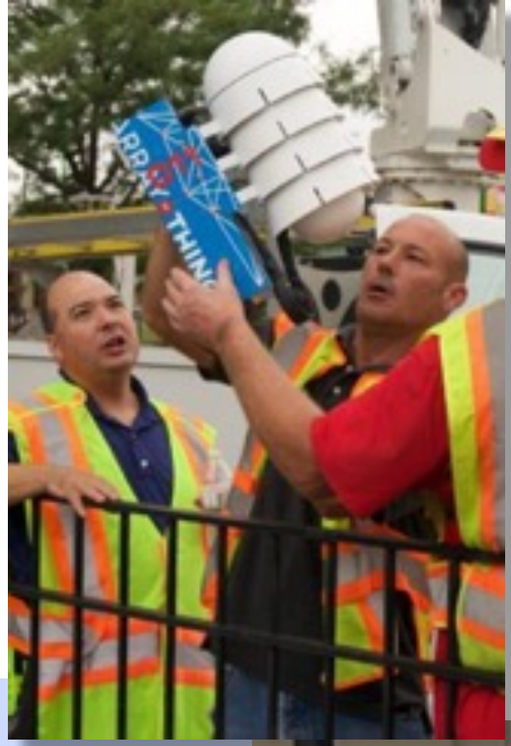

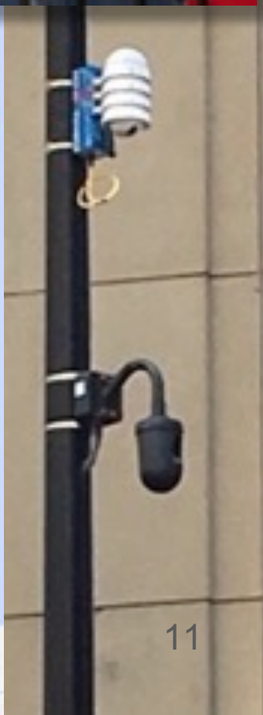

PI: Charlie Catlett, Uchicago
~2016-2018

Sage Digital Continuum

What is a “Software Defined Sensor”?

Your software container running here

Analysis produces live results

AI-Based Measurement & Anomaly Detection

Plant Species

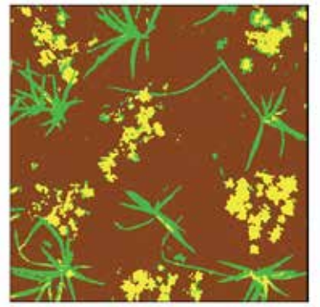

Pedestrian Flow

Birdsong

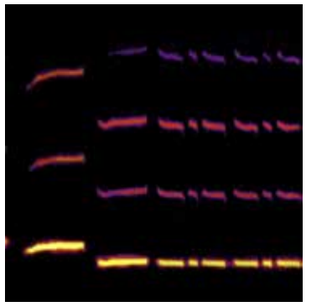

Traffic Flow

Wildlife

Cloud Motion Vectors

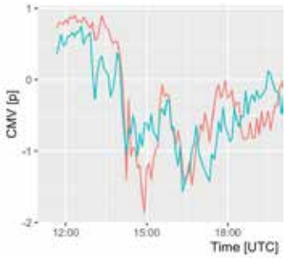

Wildfires: detecting smoke

Flooding / surface water

SAGE
sagecontinuum.org

Sage Cyberinfrastructure: Exciting Goals!

- Build new reusable cyberinfrastructure
  - High-quality, resilient, well-documented software
  - Leverage best Open Source frameworks
    - PyTorch, OpenCV, TensorFlow, Kubernetes, Docker, etc.
- Build community of AI@Edge scientists
  - New AI-based measurements
  - New AI algorithms for edge
- Deploy experimental testbed into production facilities
- Provide new capabilities for live data and triggered responses
- Teach and train students, explore new ideas

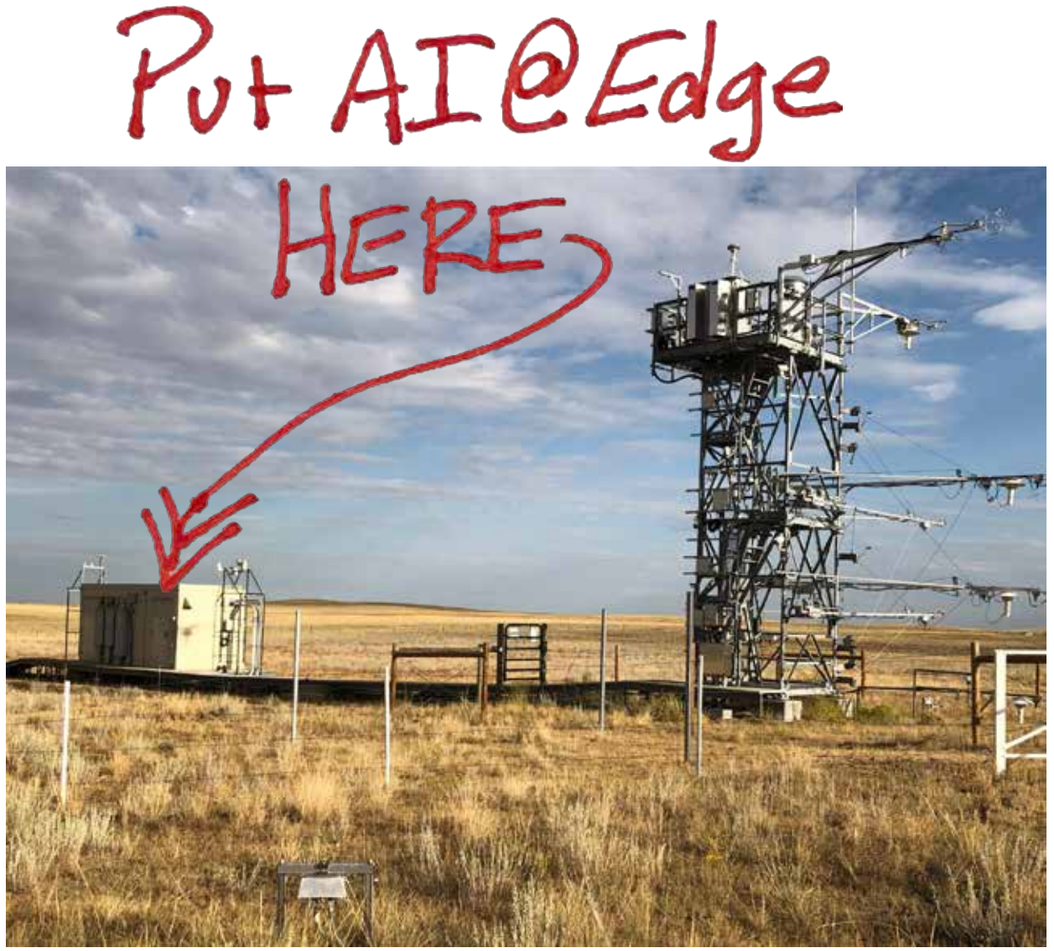

(Sensors sample at 40hz, aggregate to 30min)

Analyse full resolution data, find highest value data for the science

Key Sage Partners & Collaborators

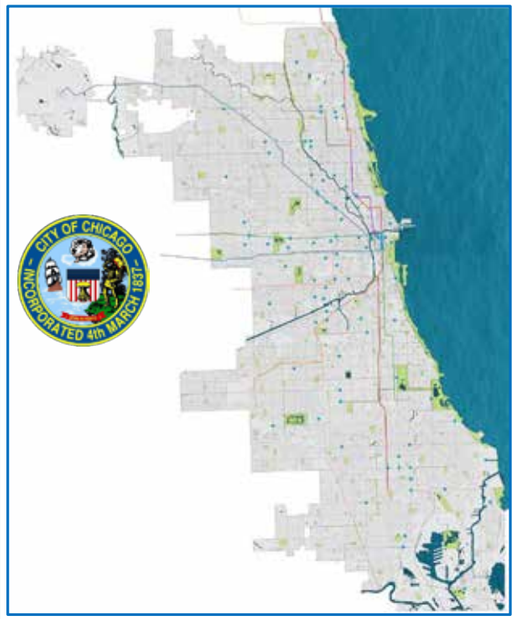

ARRAY✶THINGS

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

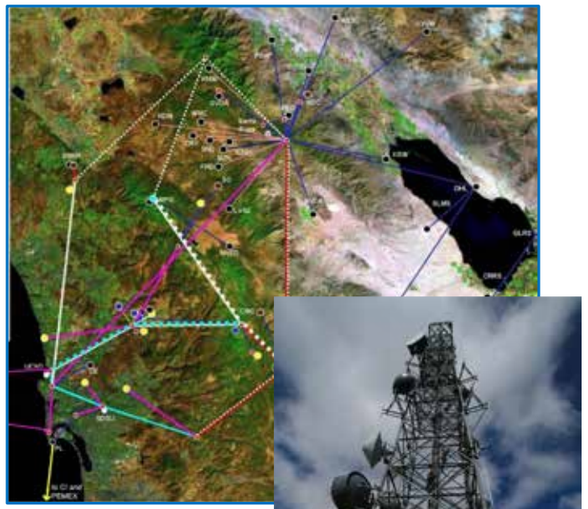

HPWREN/WIFIRE: ***Regional*** Environmental Conditions and Events.

HPWREN

SDSC SAN DIEGO SUPERCOMPUTER CENTER

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

neon

National Ecological Observatory Network

Oregon Hazards Lab: AlertWildfire, ShakeAlert, and flooding hazards

OHAZ

UNIVERSITY OF OREGON

DPI DISCOVERY PARTNERS INSTITUTE

Delivering AI@Edge Platforms: Two Forms

Wild Sage Node

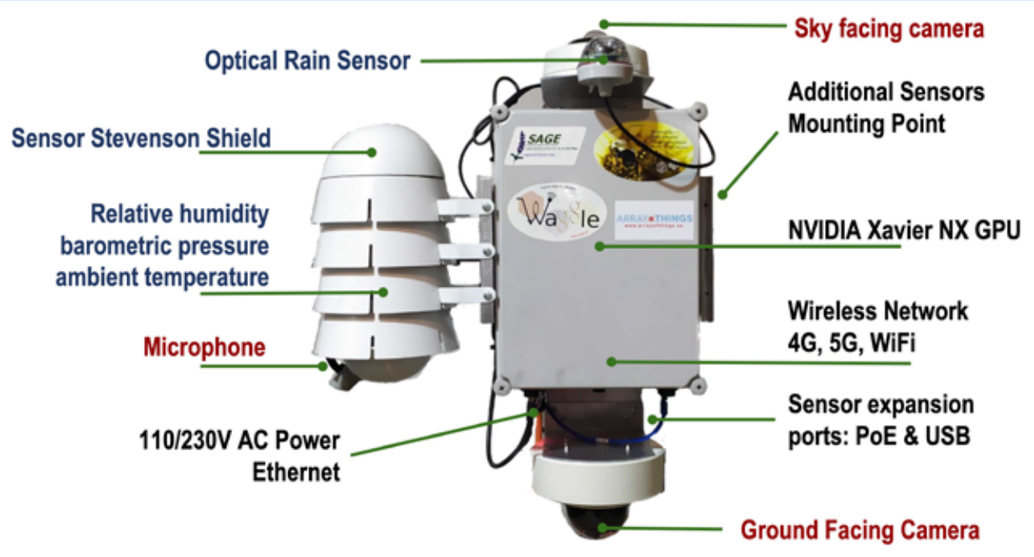

Ready for mounting ***outside***, any PoE sensor can be easily added

Sage Blade

Rugged server for instrument huts, new sensors easily added

Research Credit: Edge Architecture led by Rajesh Sankaran, Northwestern University

Leverages Open Source, Open Hardware, extensible platform called *Waggle* developed at Argonne National Laboratory

SAGE sagecontinuum.org

Sage Digital Continuum

Sage Software Architecture

Sage Nodes

OpenCV docker PYTORCH TensorFlowLite K3S ROS

Cloud Infrastructure

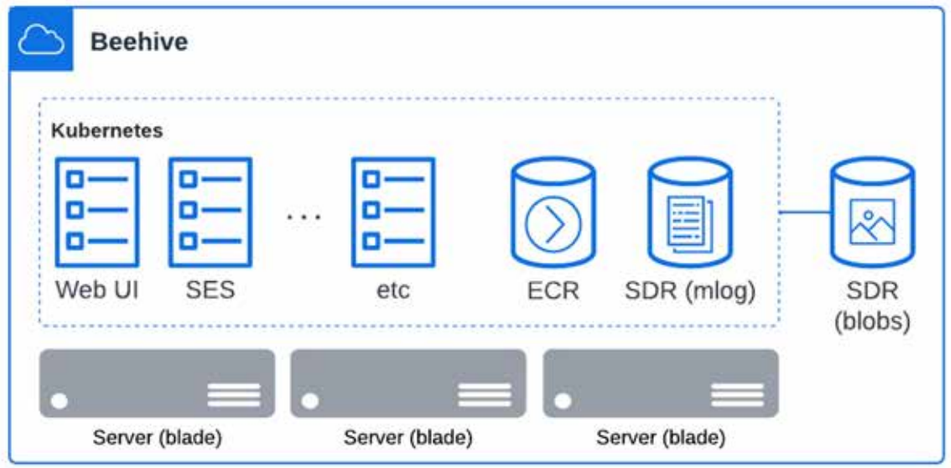

AI@Edge ”Plugin” from Edge Code Repository (ECR) (the “App Store”)

User “Plugins” run in “Waggle Edge Stack” (WES)

- Built on best Open Source AI packages
- Access to sensor and camera streams
- Libraries for efficient GPU usage
- Extreme cybersecurity
- Publish data to Beehive

Beehive manages

- Sage Edge Scheduler (SES)
- Sage Data Repository (log entries)
- Sage Data Repository (binary files)
- User Interface components

SAGE sagecontinuum.org

Research Credit: Software Team led by Joe Swantek, Northwestern University

Sage (beta) | Status | Tests | Factory

89 Nodes | 11:36:45 AM

64 reporting | 2 not reporting | 2 offline

Sanity
55 recent issues

Temps
ok

Search | Status | Project | Focus | Location | Clear | project is SAGE

Health	Type	ID	VSN	Project	Focus	Data	Temp	Last Updated	Uptime	Mem	FS Utilization
pass pass	WSN	000048B02D05A1C2	W02C	SAGE	Urban		33.9°C	nx: 1.9 min ago rpi: 1.8 min ago	162d 19h 3m 6s 162d 19h 1m 23s	6.48 GiB / 7.59 GiB 3.57 GiB / 3.70 GiB	plugins rw ro 10.00% 9.15% 75.29% 55.84% 63.70%
pass	WSN	000048B02D05A1C6	W027	SAGE	Urban		-	nx: 1.1 days ago	74d 15h 14m 3s		plugins rw ro
pass pass	WSN	000048B02D05A1C7	W026	SAGE	Urban		34.6°C	nx: 10.2 secs ago rpi: 1.2 min ago	106d 4h 2m 9s 162d 19h 1m 58s	6.93 GiB / 7.59 GiB 3.51 GiB / 3.70 GiB	plugins rw ro 10.20% 9.28% 75.45% 50.22% 63.46%
	WSN	000048B02D0766D2	W014	SAGE	BadRiver		-	-	-	-	-
	WSN	000048B02D15BC42	W01D	SAGE	NEON-MDP		-	-	-	-	-
	WSN	000048B02D15BC64	W01F	SAGE	NEON-MDP		-	-	-	-	-
pass pass	WSN	000048B02D15BC73	W024	SAGE	Urban		37.8°C	nx: 23.1 secs ago rpi: 43.7 secs ago	33d 12h 4m 45s 162d 18h 53m 4s	6.78 GiB / 7.59 GiB 3.53 GiB / 3.70 GiB	plugins rw ro 10.21% 9.28% 77.71% 33.42% 63.60%

SAGE
sagecontinuum.org

admin.sagecontinuum.org

Building Community for AI@Edge Sage Applications
The Edge Code Repository

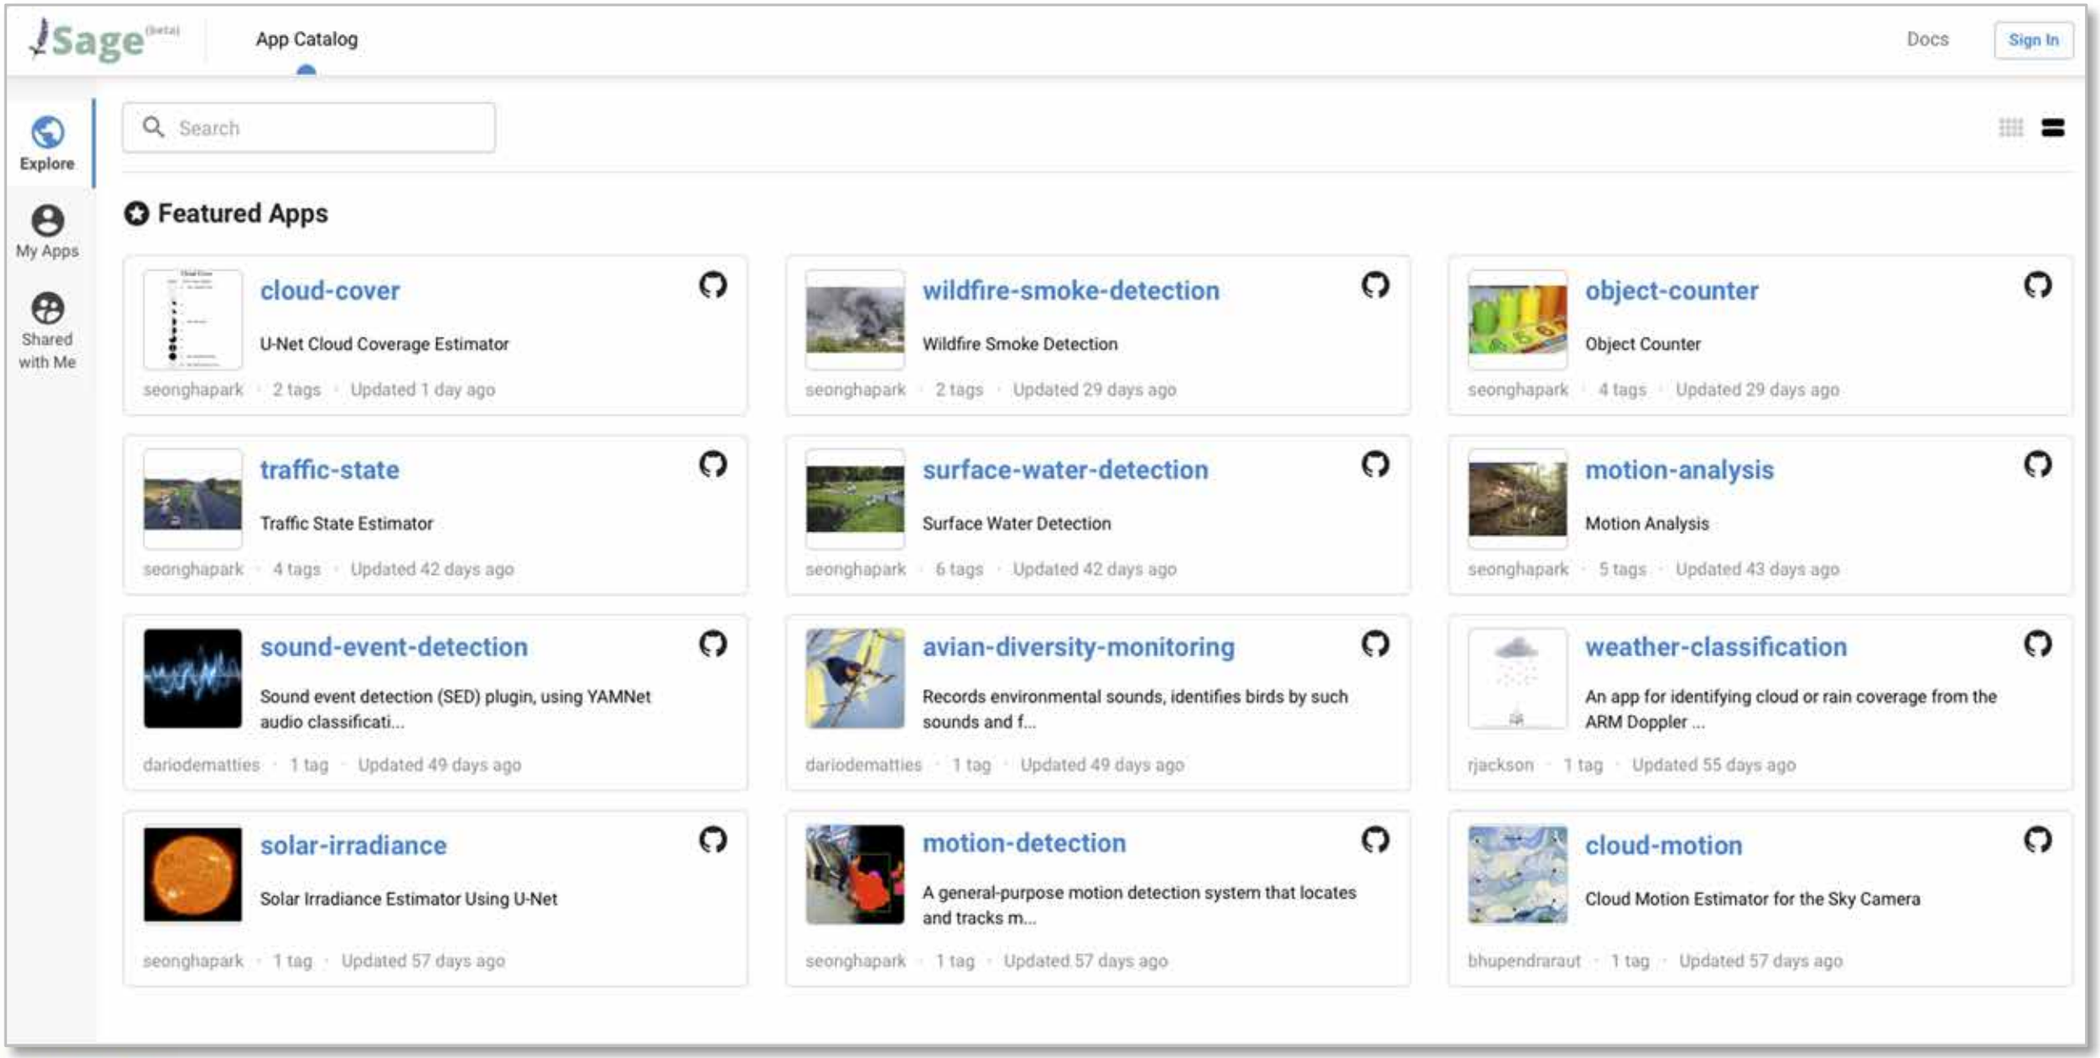

Sage Digital Continuum

SAGE
sagecontinuum.org

portal.sagecontinuum.org

AI Portfolio led by Sage Deputy Director: Nicola Ferrier, UChicago

Deploying Wild Sage Nodes in Chicago

Sean Shahkarami, Uchicago, overseeing installation
(in brown jacket, by corner of building)

Undergraduate Research: Pedestrian Detection and Paths

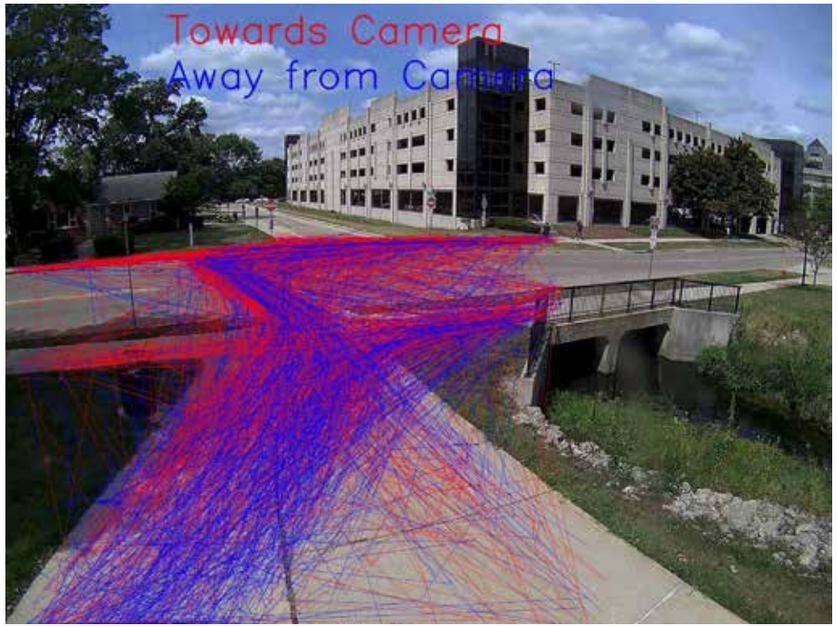

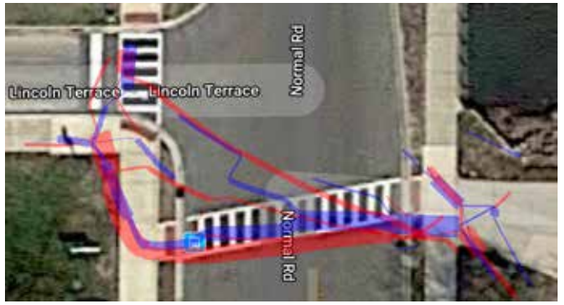

NIU experimental node with wired network connection

- Experiment with sampling rate and resolution
- Work is now being ported to Sage node

YOLO based model for identifying people and to check for use of crosswalk

Pedestrian data processed to understand patterns and transformed for top-down view then bundled to highlight patterns

Justin Derus, **Wesley Kwiecinski**, Pratool Bharti, Michael Papka: Northern Illinois University

Avian diversity monitoring

Image Creator: Becky Matsubara

Waggle Microphone

- Bird diversity changes as a metric to track the current environmental conditions
- We automate Avian Diversity Monitoring by using a DNN, called BirdNET [1], capable of identifying 984 North American and European bird species by sound. Weekly cumulative detections of non-migratory species occurrence was highly correlated with human point count observations
- It will be possible to get exposure to many organisms occupying diverse areas without needing to detect them during demanding and expensive human fieldwork

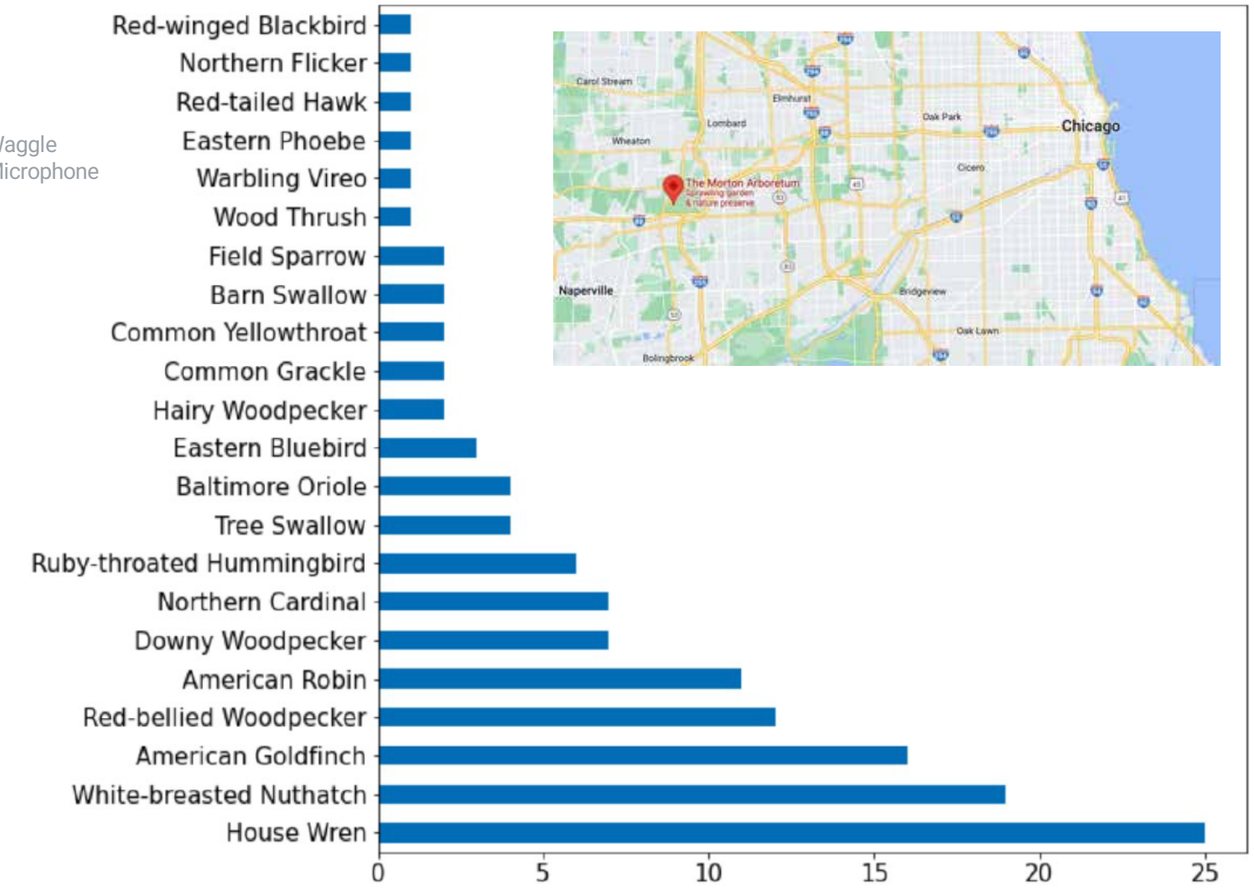

Morton Arboretum Avian Detection, June 28, 2021 (24 hour)

[1] Stefan Kahl, Connor M. Wood, Maximilian Eibl and Holger Klinck. BirdNET: A deep learning solution for avian diversity monitoring. Ecological Informatics Volume 61, March 2021.

SAGE sagecontinuum.org

Research Credit: Dario Dematties, Bhupendra Raut, Nicola Ferrier

Paintbrush Prairie Bird Detection

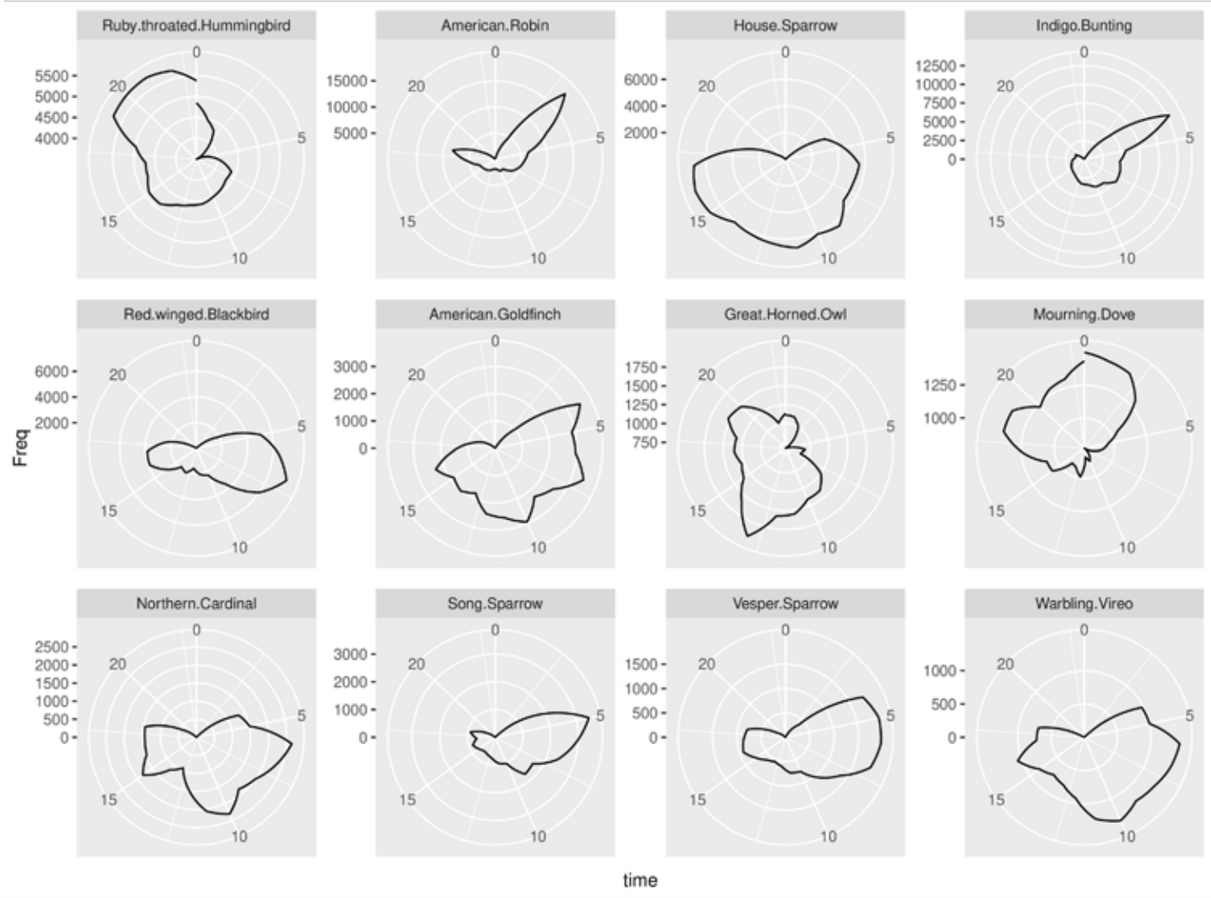

Total calls, for top 12 species, as a function of the hour of the day (UTC-06).

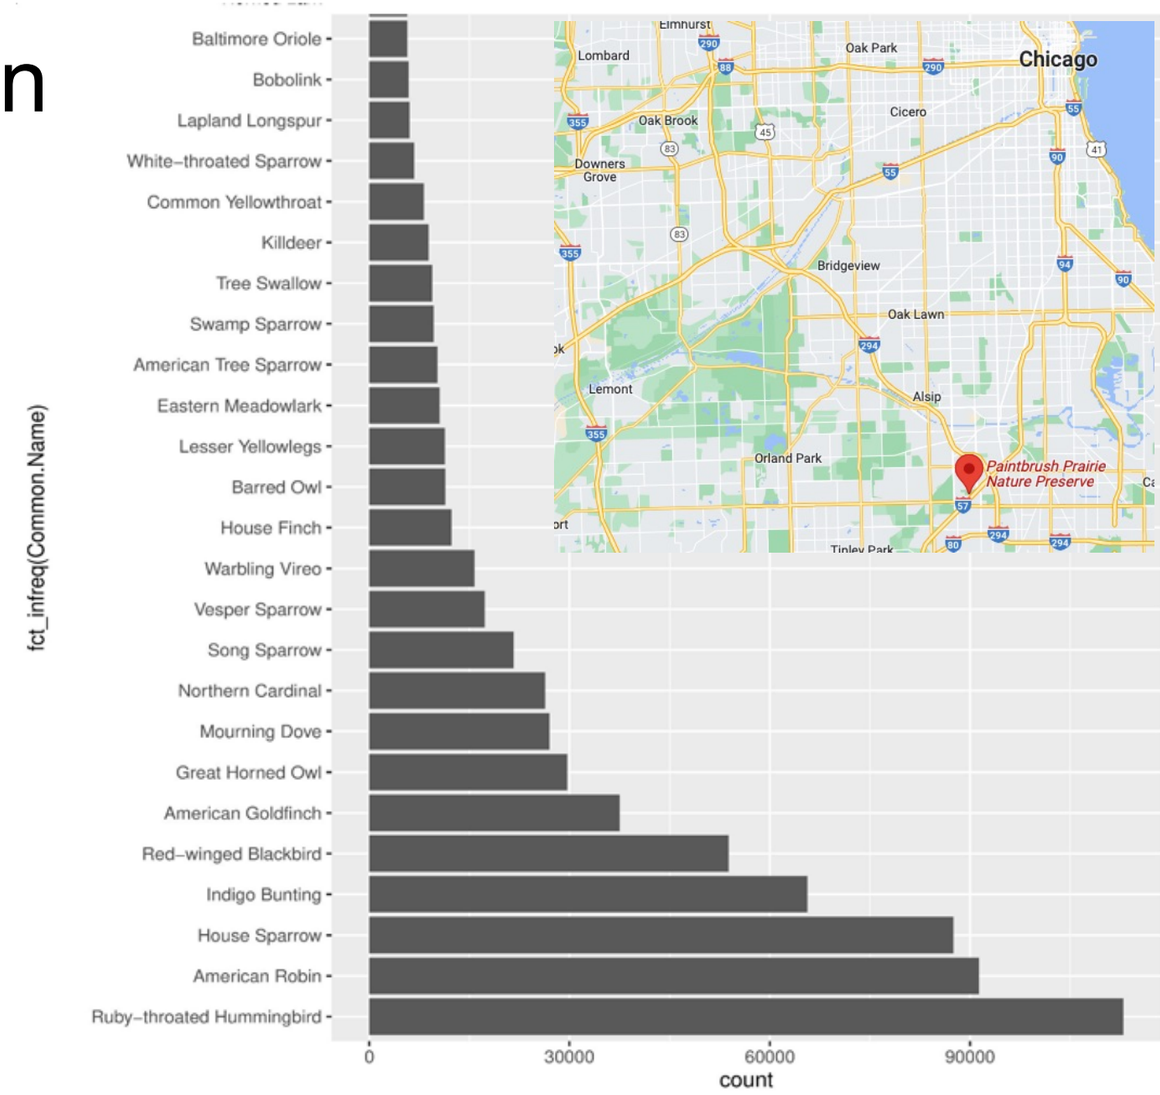

Top 25 bird calls recorded at the Paintbrush Prairie Natural Preserve (Nature Conservancy Site) from Sep 2020 to Dec 2021

Research Credit: Dario Dematties, Bhupendra Raut, Nicola Ferrier

Cloud Motion Vectors with Hemispheric Sky Camera

- Real-time cloud motion vectors from sky images have applications in meteorological analysis, nowcasting, and short-term prediction of solar irradiance.

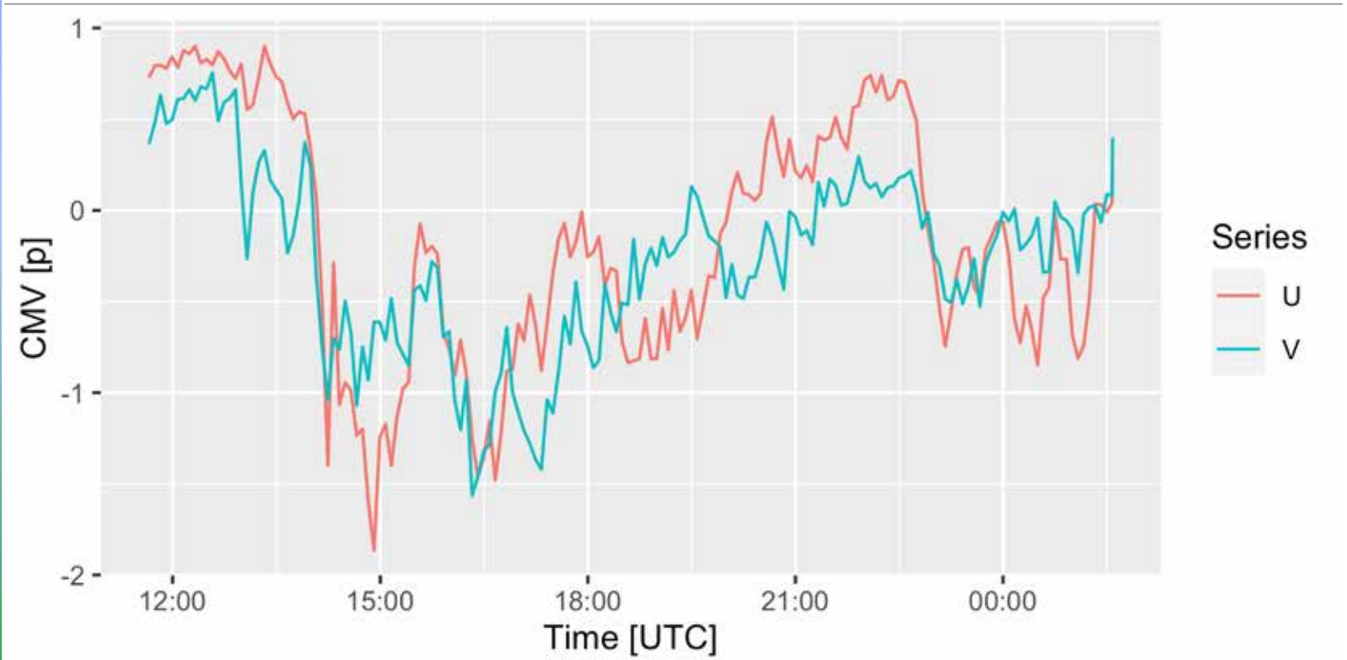

Future: Instrument steering and Nowcasting

- Camera contamination by rain and snow is identified by the ML algorithm and reported.
- The use of AI/ML with the other ssnsors produce valuable products (e.g. Solar irradiance, nowcasting, locale weather reporting).

Scott Collis, SAGE Co-PI, Northwestern University

Bhupendra Raut, Northwestern University

Cloud M… Surprise Training Data …ky Camera

- Real-tim…
  images …
  meteoro…
  short-ter…

…ination by rain and snow is identified …ithm and reported.

…L with the other ssnsors produce …cts (e.g. Solar irradiance, nowcasting, …reporting).

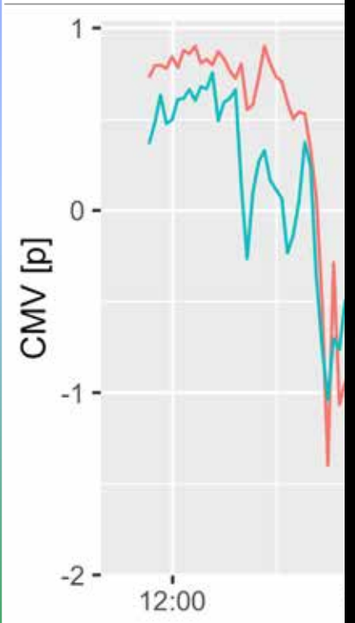

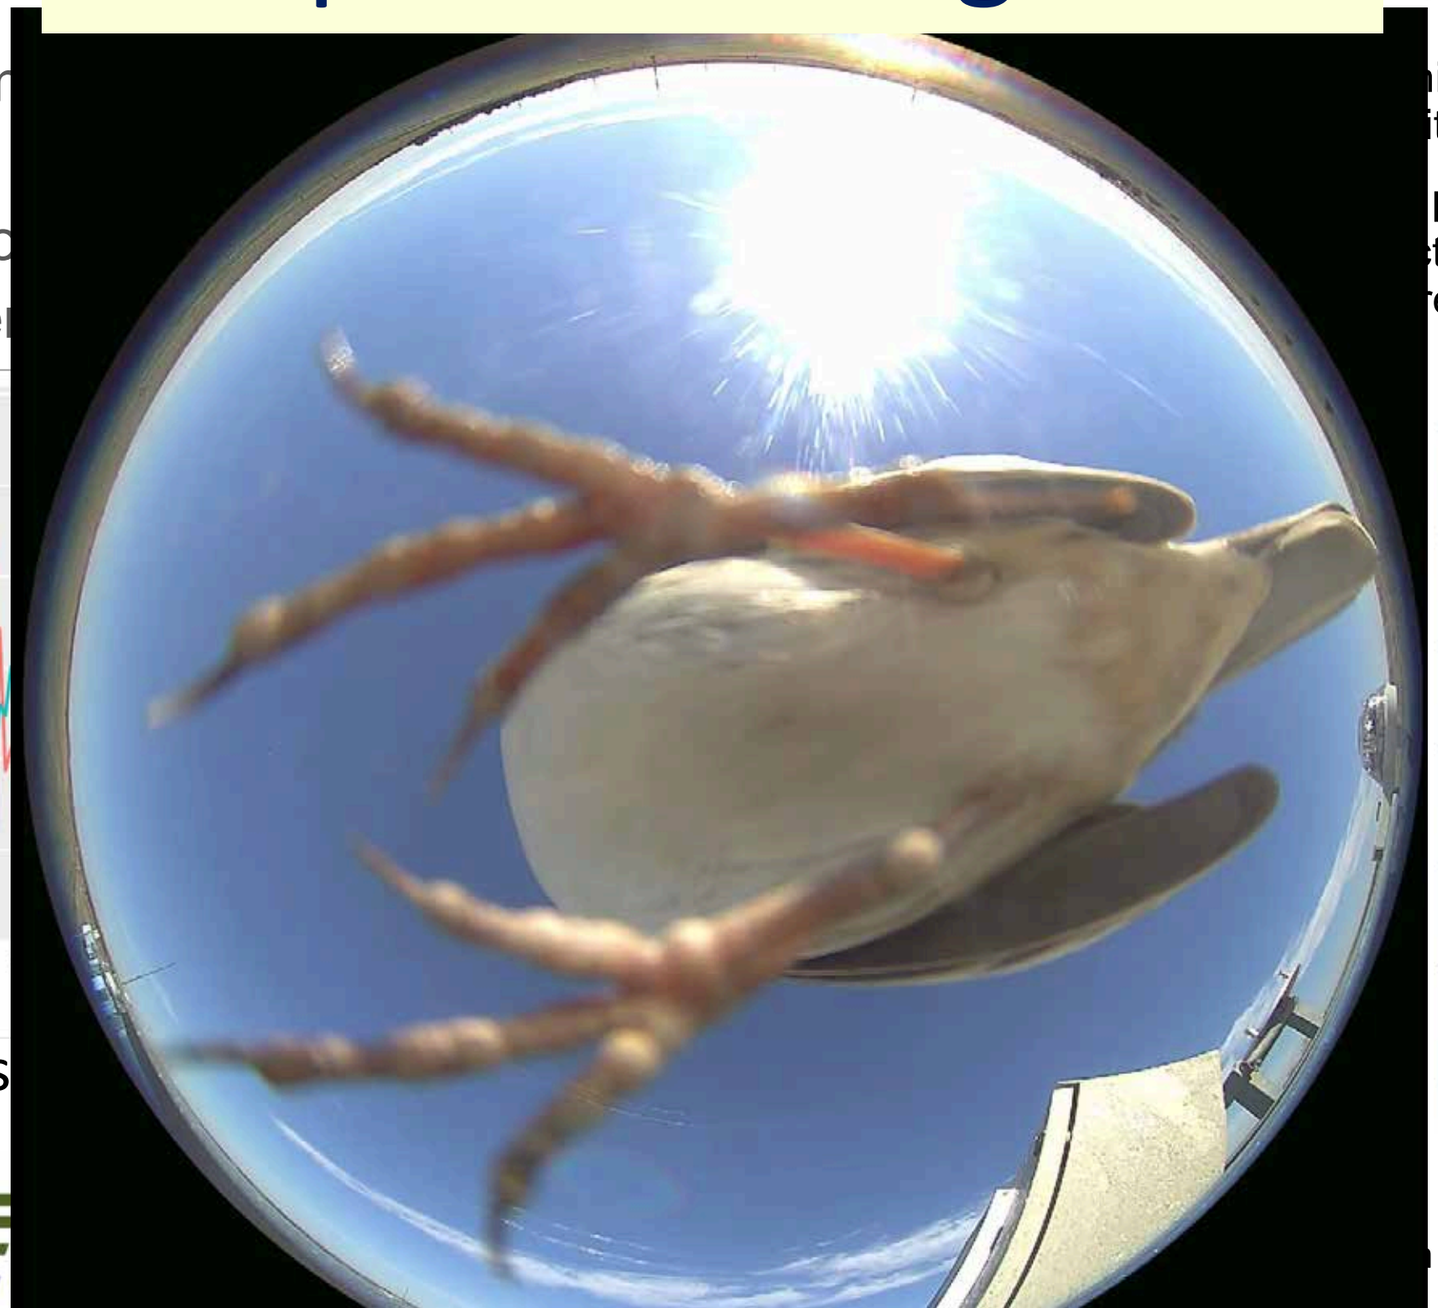

Future: Ins…

SAGE
sagecontinuum.org

…n Raut, Northwestern University

Measuring Water and Snow Depth

We are evaluating multiple approaches to estimate the water (or snow) level from images of rulers (in of a stream at a NEON site)

- Computer vision (CV) based
- Machine Learning algorithms
  - U-Net, ResNet
  - Self-supervised Learning

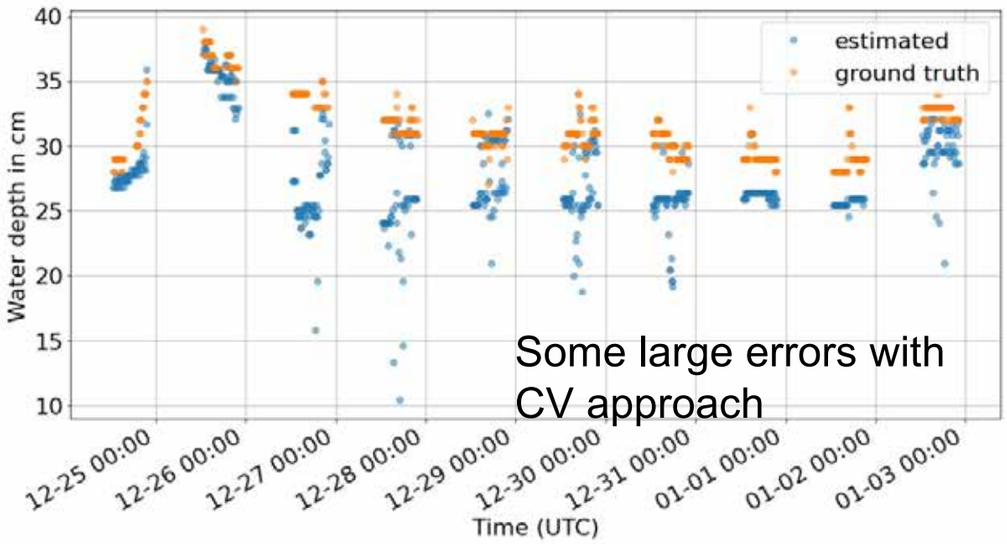

Human annotation using Labelbox

U-Net (segmentation)

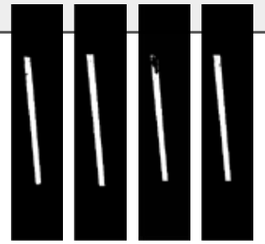

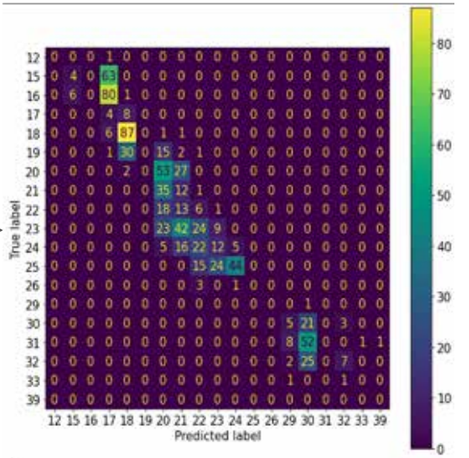

ResNet32 based Regression

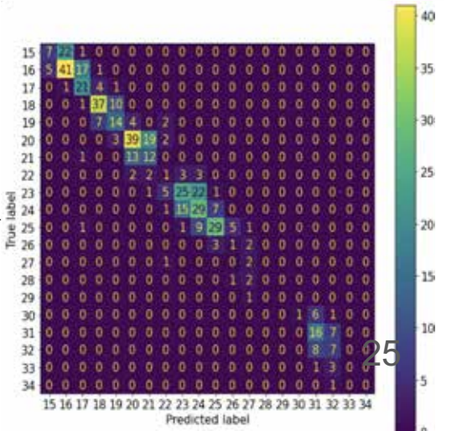

Measuring Water and Snow Depth


RGB image

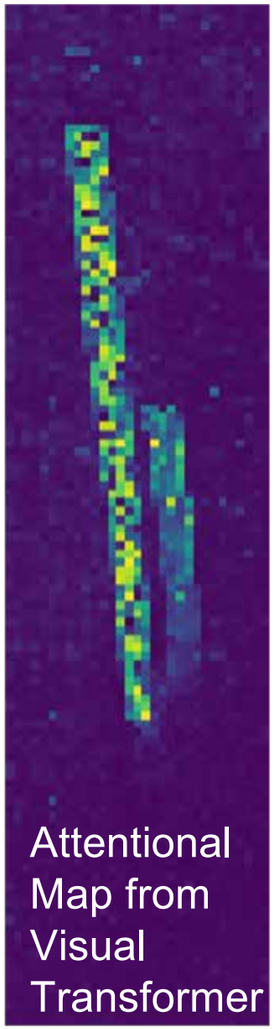
Attentional Map from Visual Transformer


Mask generated from the attentional map


Labels produced by humans

Self-supervised Segmentation

- Exploring visual transformer ML
- ML model was trained using only images from IMAGENET (no labels and no NEON data)
- An Intersection over Union score > 0.5 is normally considered a “good” prediction.

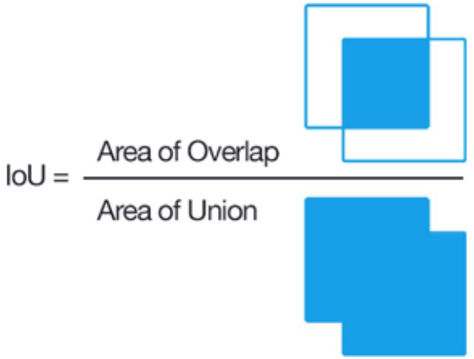

Intersection/Union (IoU) = 0.729


Inters.


Union

Dario Dematties, Northwestern University

Surface Water Detection

Linked with HPC, can be used to build hydrology models and predictive capabilities

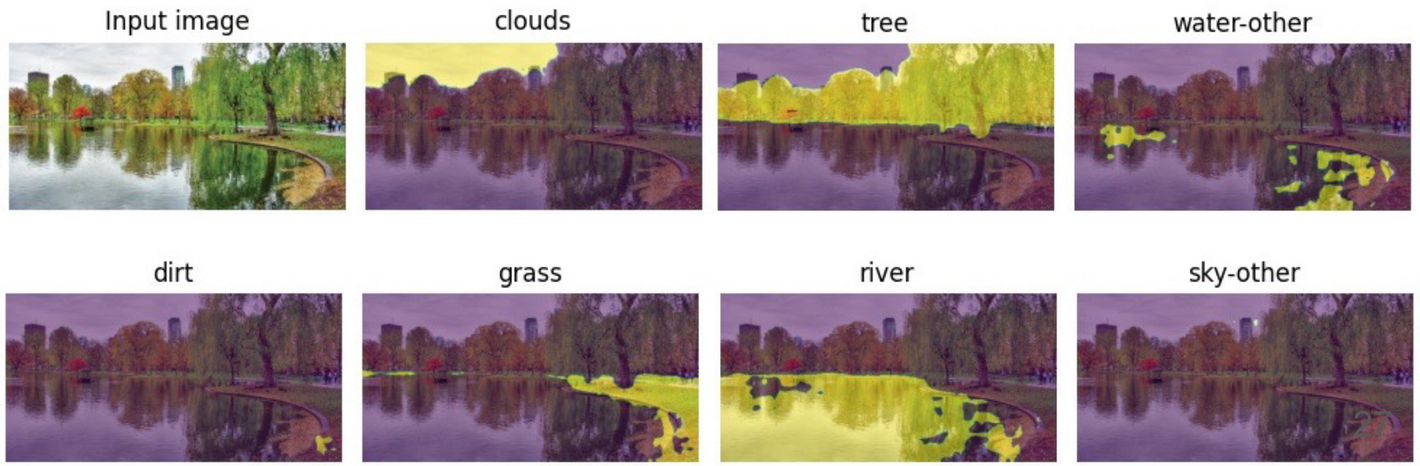

SAGE
sagecontinuum.org

Nicola Ferrier, Uchicago

Sage Digital Continuum

Sage and POWDER: Next Generation Wireless

ES-642 Dust Sensor

Mobotix M16 on PT Unit

USRP B210 SDR

POWDER

Powder (the **P**latform for **O**pen **W**ireless **D**ata-driven **E**xperimental **R**esearch) is flexible infrastructure enabling a wide range of software-defined experiments on the future of wireless networks.

Powder supports software-programmable experimentation on 5G and beyond, massive MIMO, ORAN, spectrum sharing and CBRS, RF monitoring, and anything else that can be supported on software-defined radios.

Utah Field Museum

SAGE
sagecontinuum.org

Sage Digital Continuum

Wildfire Detection and Prediction

Ilkay Altintas, UCSD, Co-PI for SAGE

WIFIRE

Exploring wildfire detection at the edge linked to HPC simulations

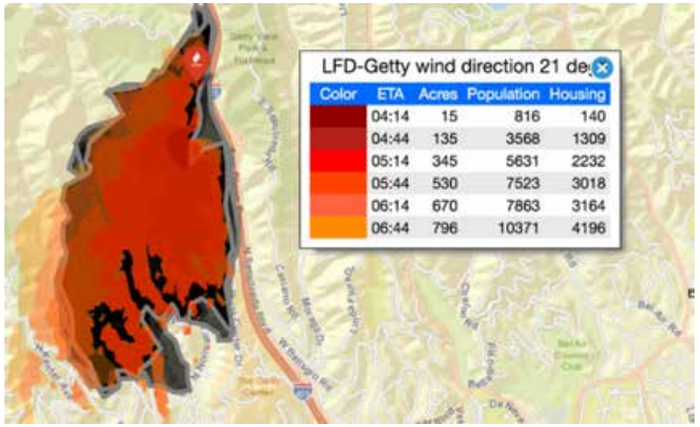

ALERTWildfire: A unique wildfire detection and monitoring system

Collaboration: Doug Toomey, UOregon

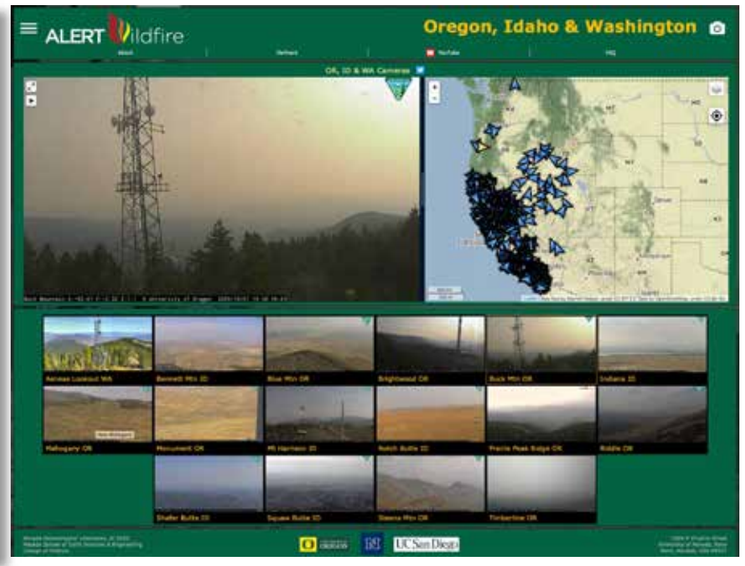

UNIVERSITY OF OREGON

OHAZ

Frank Vernon, UCSD, HPWREN
Mt Wilson Fire

Exploring New Methods for Wildfire Smoke Detection

Two approaches to improve the predictions are

1. Use of thermal IR camera, and
2. Incorporating motion of the smoke in the DL models.

- Prescribed burns and real wildfire data is needed to train the AI models.
- Cloud temperatures can be used to estimate cloud-base heights and cloudiness.

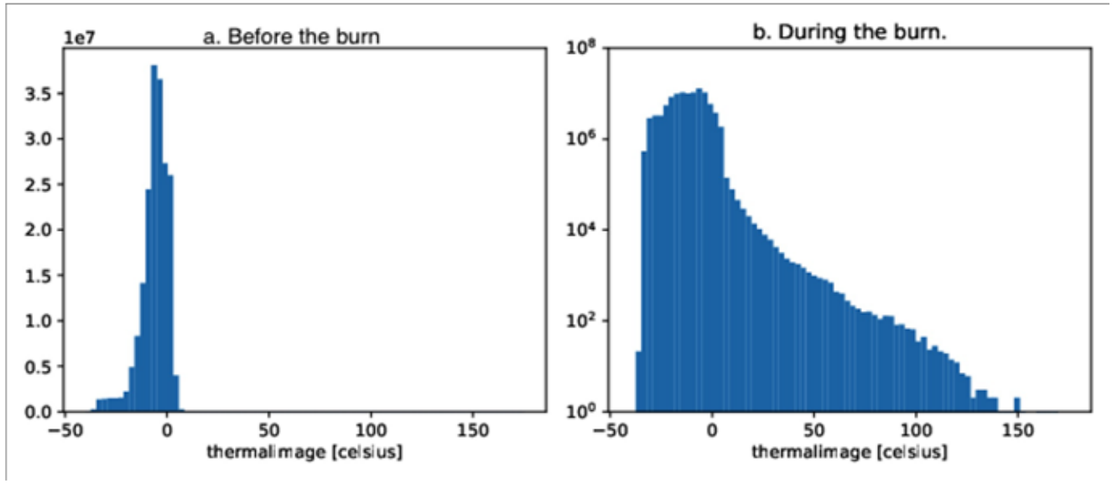

SAGE
sagecontinuum.org

Bhupendra Raut, Northwestern University

Sage Co-PI: Eugene Kelly, Colorado State
eugene.kelly@colostate.edu

NEON Mobile Deployment Platform (MPD) with Sage

Konza Prairie for controlled burn: April 2022.

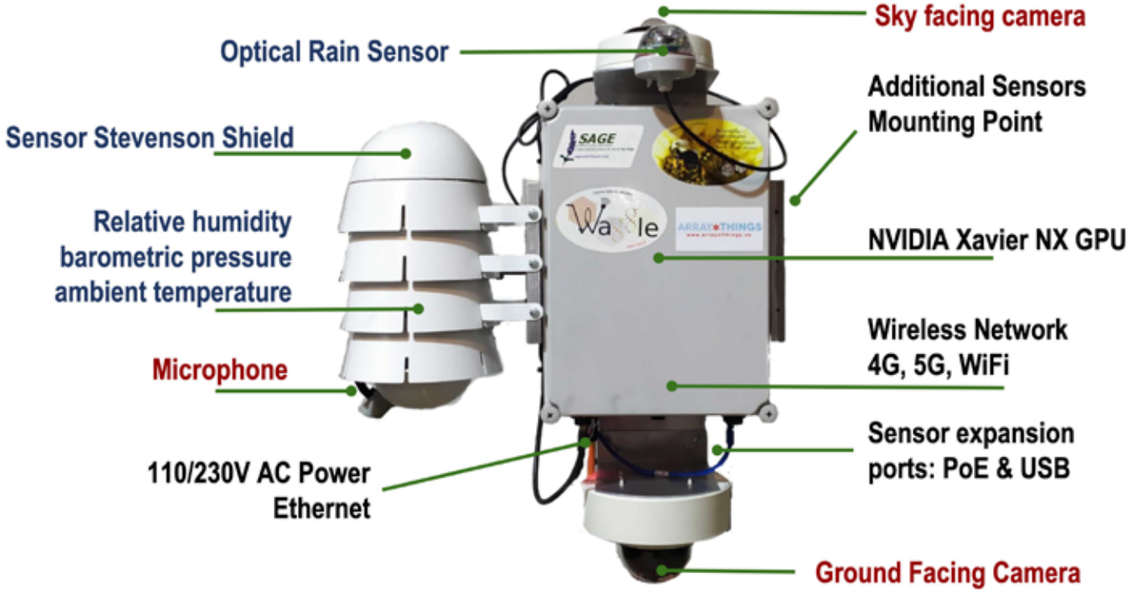

Special Thanks: NEON Team!
Rommel Zulueta @ Battelle

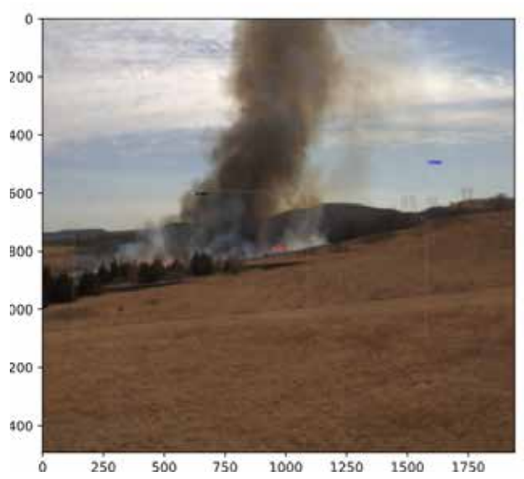

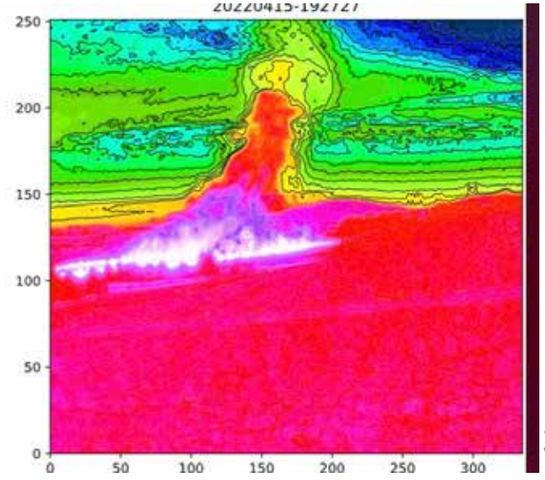

SAGE
sagecontinuum.org

Sage NEON NSF Controlled Burn
Konza Prairie, Kansas

Data from the experiment already available to the community!

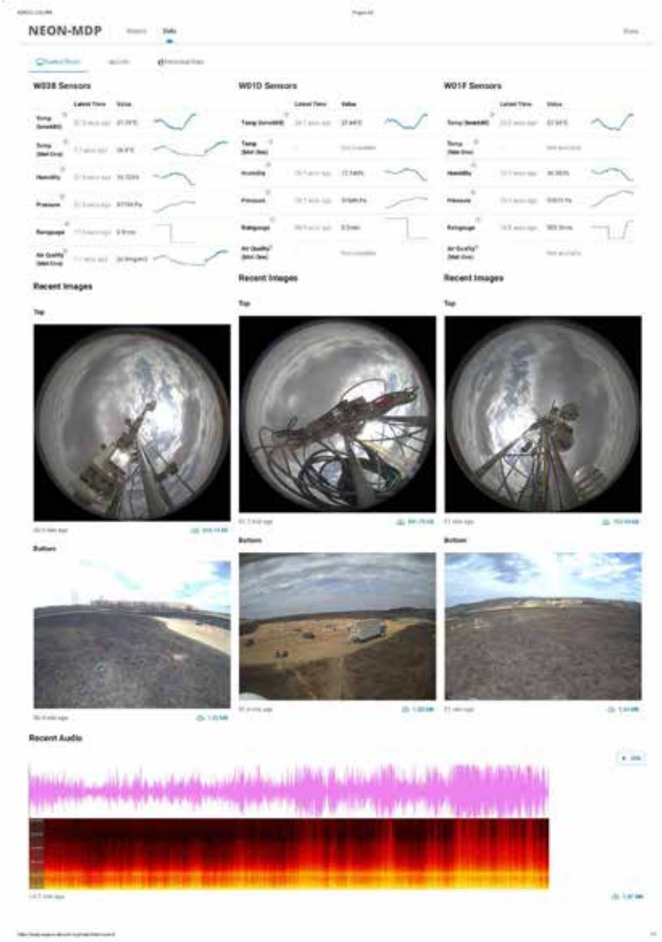

WIFIRE Commons

DATASETS ORGANIZATIONS ABOUT Search

Organizations / SAGE - NEON / NEON MDP / Sage / WIFIRE ...

NEON MDP / Sage / WIFIRE BP3d: Konza Prairie Burn Experiment

Dataset extent

Map data © OpenStreetMap contributors
Tiles by Stamen Design (CC BY 3.0)

Organization

SAGE
A Software-Defined Sensor Network
Cyberinfrastructure for AI at the Edge
NSF
neon
Operated by Battelle

SAGE - NEON

The Sage project is designing and building a new kind of national-scale reusable cyberinfrastructure to enable AI at the edge. https://sagecontinuum.org/ The National Science... read more

License

Creative Commons Attribution 4.0

OPEN DATA

NEON MDP / Sage / WIFIRE BP3d: Konza Prairie Burn Experiment

The Konza Prairie Biological Station, located in the Flint Hills of northeastern Kansas, is one of the last native tallgrass prairies. Working with the Konza Prairie Station, NEON and the Sage Project have collaborated to deploy a NEON mobile deployment platform (MDP) augmented with Sage artificial intelligence (AI) deployed to the edge. The "Wild Sage Nodes" and "Sage Blades" provide advanced computation and instrumentation to help study a controlled burn of the prairie. Sage AI@Edge algorithms have provided breakthrough analysis of instruments, from LIDAR and thermographic cameras to air quality and scintillation detectors. Some of the AI algorithms already developed for Sage are available in the Edge Code Repository (https://portal.sagecontinuum.org/apps/explore) – from analysis of bird species and flooding to wildfire detection and measuring cloud dynamics.

Data collected on April 15, 2022 include images from a thermographic camera, RGB cameras, particle sensors, and more. AI algorithms analyzed some of the data streams in real time, while other data streams logged the events and will be used later with advanced self-supervised AI algorithms to improve algorithms, build training data sets, and help scientists better understand the earth's atmospheric and environmental processes.

See the following jupyter notebook as a reference for accessing the data:
https://github.com/iperezx/sage-smoke-detection/blob/master/post-processing/sage-data-client.ipynb

Data and Resources

reading.sensor.csat3.pkl	3D wind speed, direction and sonic temperature	Explore
reading.sensor.g2131i_raw.pkl	Atmospheric CO2 isotopes	Explore
reading.sensor.hfp01sc.pkl	Soil heat flux plate	Explore
reading.sensor.hmp155.pkl	Relative humidity	Explore
reading.sensor.l2130i_raw.pkl	Atmospheric H2O isotopes	Explore
reading.sensor.li191r.pkl	Photosynthetically active radiation (quantum line)	Explore
reading.sensor.li7200_raw.pkl	CO2 and H2O concentrations turbulent	Explore

Ismael Perez and Ilkay Altintas, UCSD

Collaboration with CSIRO in Australia

City Environment Sensor Network

Gold Coast has rapid population growth and population densification, an existing city owned network of IoT sensors (LoRaWAN) and is an existing member of Open and Agile Smart Cities group (OASC).

Key Science objectives

- Model-data fusion to increase spatial and temporal resolution of modeled local weather (and climate)
  - integration with physically based models like CCAM and Spark
  - for applications with machine learning techniques
- Calibration and validation to ensure accuracy of data
- Application of sensor data in disaster management and for city resilience
- Provide data to inform policy and strategy decisions on environmental and urban growth

Collaboration with Mahesh Prakash and Nikhil Garg, CSIRO

10 Waggle nodes will be deployed here

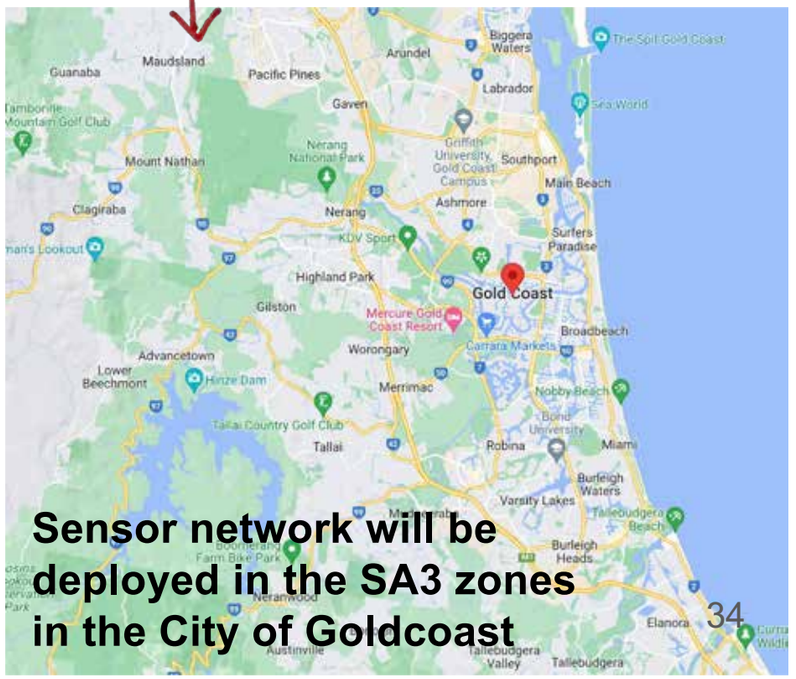

Sensor network will be deployed in the SA3 zones in the City of Goldcoast

Partnership with Ojibwe Nation
Study Climate Change Impact on Manoomin (Wild Rice)

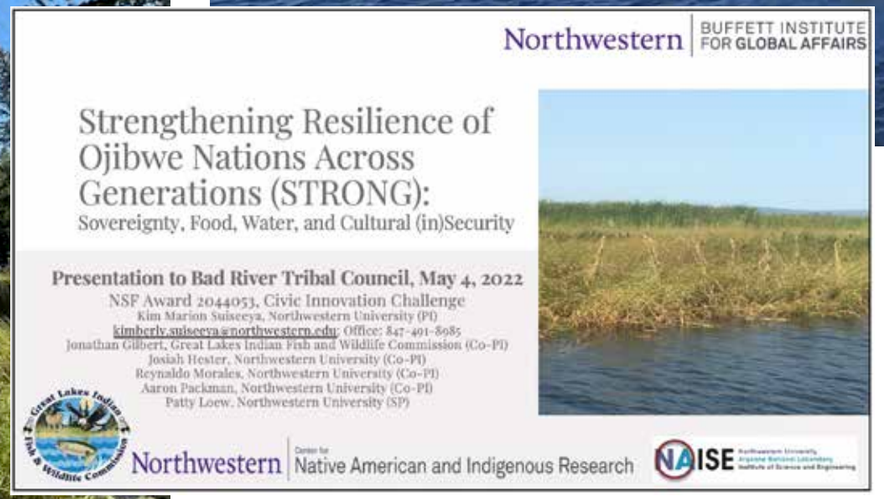

- **Sage node at Bad River Fish Hatchery**
- **Deployment approved by Tribal Council**

Collaboration with Kim Marion Suiseeya, Northwestern University

Partnership with Native Hawai'ian Community Study Climate Change Impact

Collaboration with Josiah Hester, Northwestern University

Sage Cyberinfrastructure: Key Point

Your science, your sensors. Sage is the cyberinfrastructure

Sage can interface with any instrument or sensor*

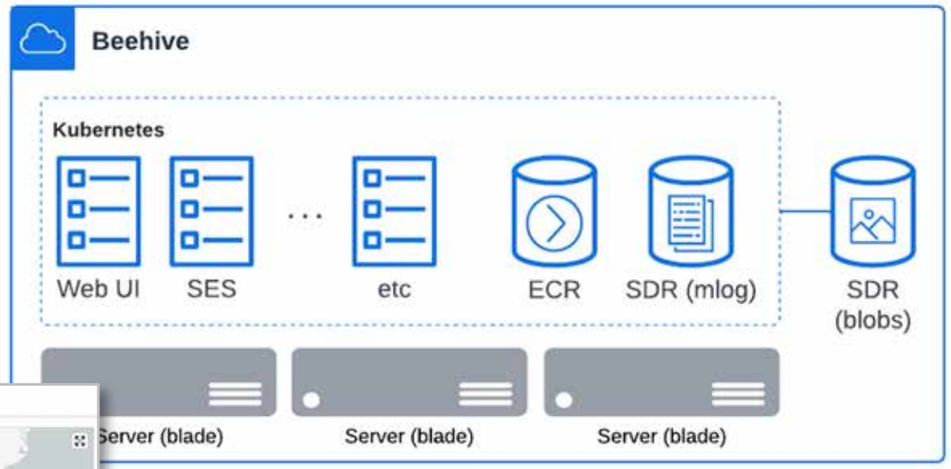

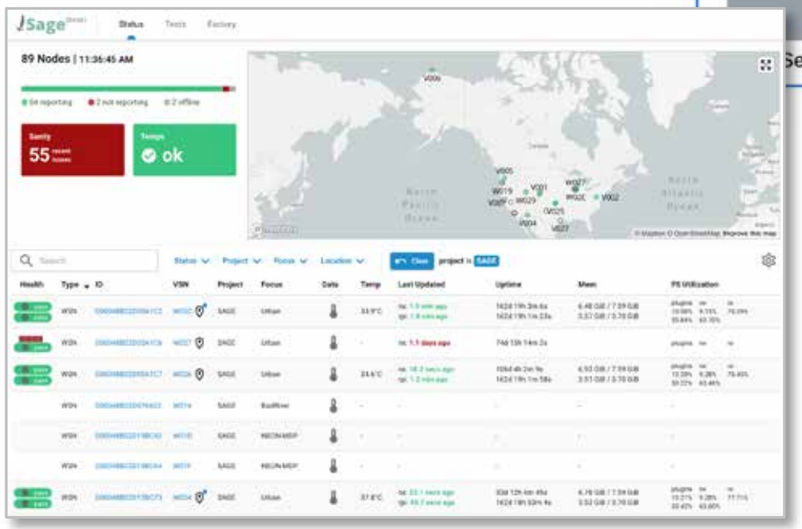

OpenCV docker PYTORCH TensorFlowLite K3S ROS

* We don't do Windows

SAGE sagecontinuum.org

Sage Portal Views

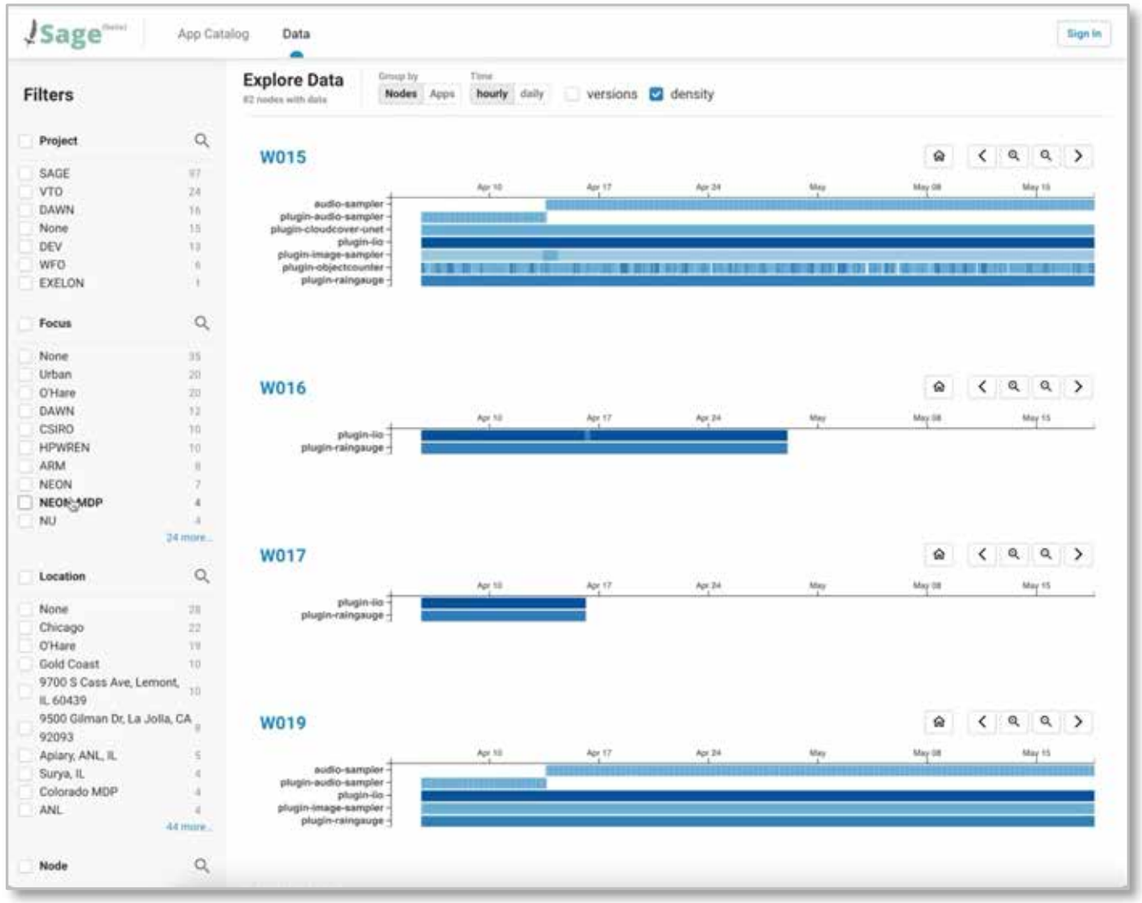

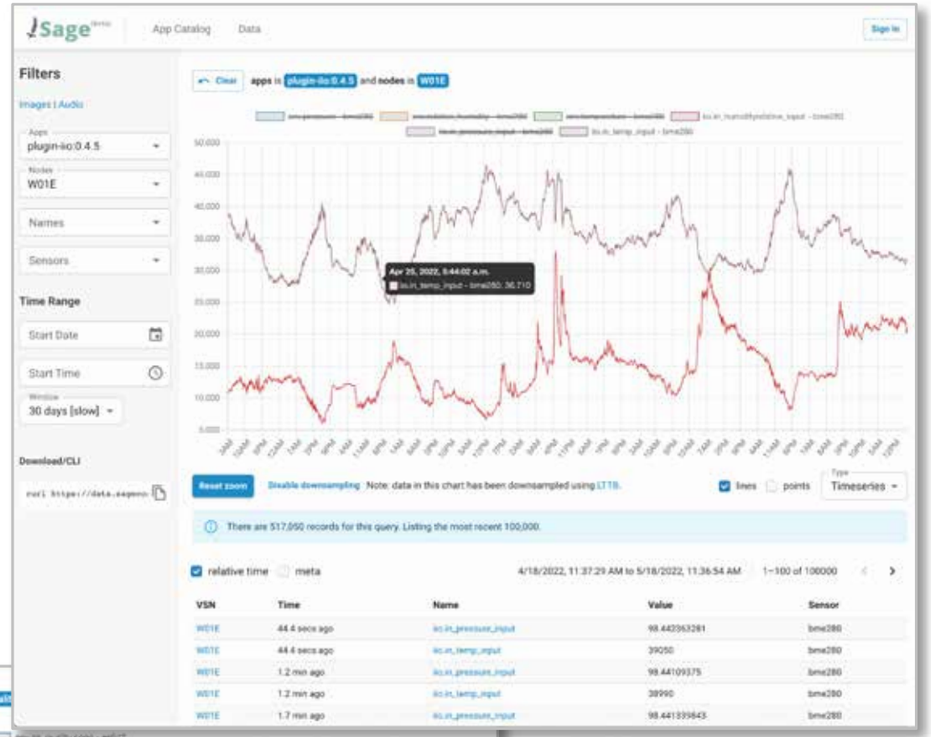

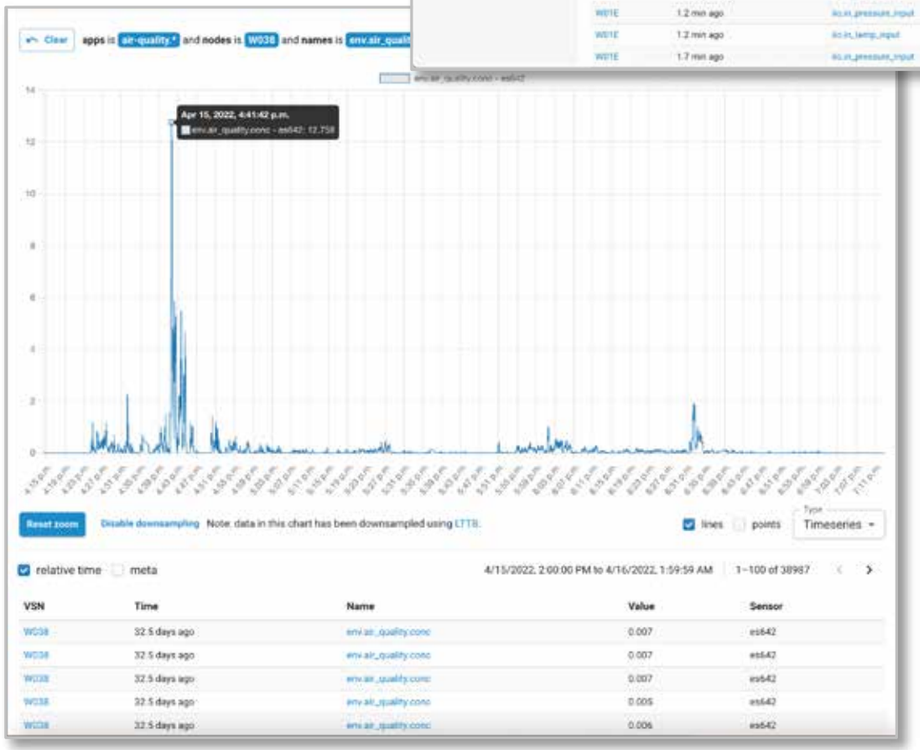

Data can be downloaded live via API interface, downloaded as CSV TGZ Bundles, or browsed

Credit: Neal Conrad, UChicago

Sage Resources

Getting started with Sage! - https://docs.sagecontinuum.org

Sage AI@Edge Apps - https://portal.sagecontinuum.org/apps/explore

Sage Data - https://portal.sagecontinuum.org/data

Sage Konza MDP Campaign - https://mdp.sagecontinuum.org

Overall Sage system status - https://admin.sagecontinuum.org/status

Portal showing the current set of science jobs executing on the various nodes will be available to the public soon…

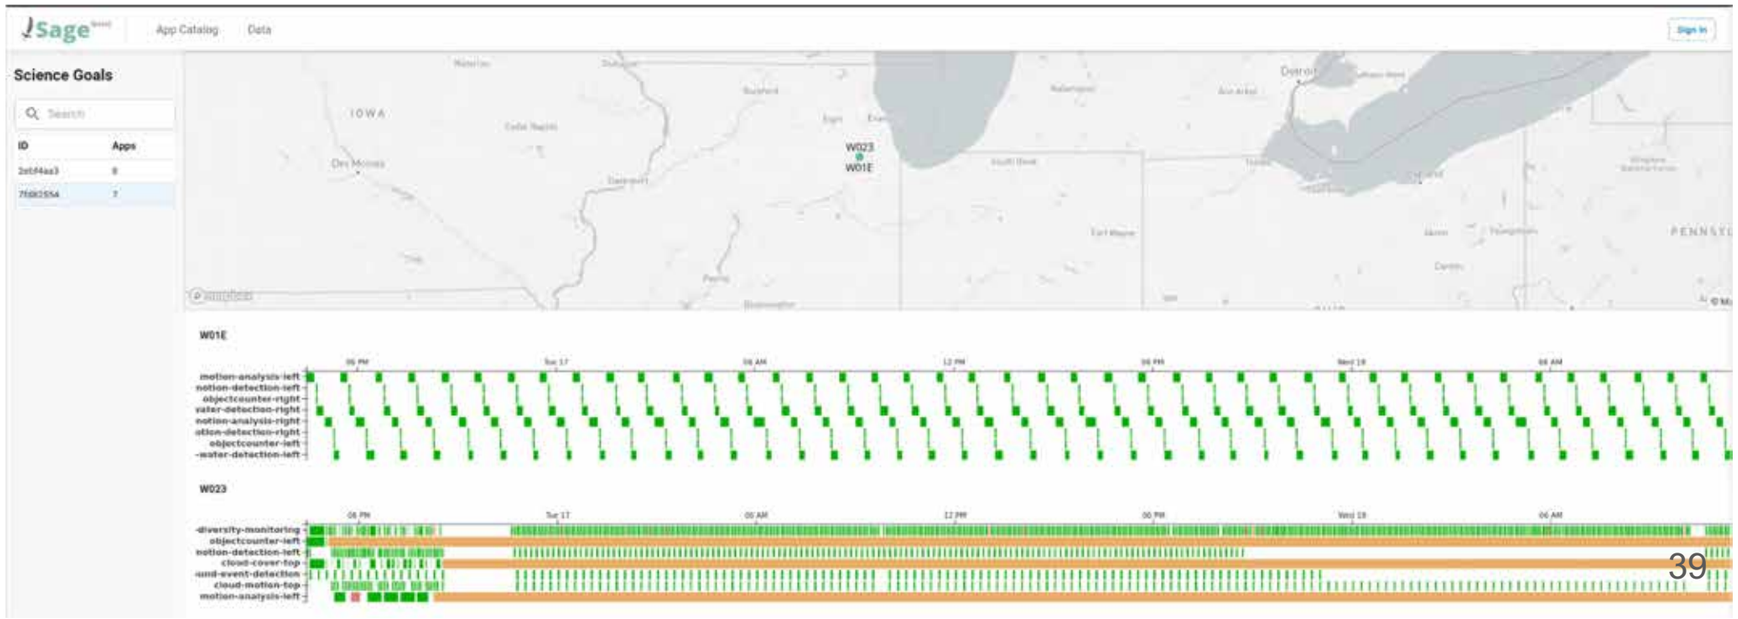

Questions?

Leadership Team


Pete Beckman
(NU: Director)


Nicola Ferrier
(UC: Deputy Dir.)


Scott Collis
(NU: Instruments, Atmos)


Valerie Taylor
UC: Edu, Broader Impacts)


Eugene Kelly
(CSU; Ecosys, NEON)


Mike Papka
NIU: Edu, Broader Impacts


Raj Sankaran
NU: Node Arch


Ilkay Altintas
(SDSC: Data)


Charlie Catlett
(UIllinois: Urban)


Jim Olds
(GMU; Life Sci, Risk)


Dan Reed
(Utah; Architecture)


Kathy Bailey
Proj Mgmt


Helen Taaffe
Proj Mgmt


Joe Swantek
NU: Software


Irene Qualters
(LANL; Advisory Committee Chair)

Professors Aaron Packman and William Miller, Northwestern University
Gensburg-Markham Prairie, The Nature Conservancy
Photo Credits: Liliana Hernandez-Gonzalez, Northwestern University
Dec 2015

SAGE
sagecontinuum.org

Sage Digital Continuum

Special Thanks

Students!

Ilkay Altintas
Kathy Bailey
Daniel Balouek-Thomert
Pete Beckman
John Blair
Eric Bruning
Adam Brust
Charlie Catlett
Scott Collis
Neal Conrad
Geoff Davis
Dario Dematties
Nicola Ferrier
Jannick Fischer
Larry Hartman
Robert Jackson
Euguene Kelly
Yongho Kim
Nick Maggio
Seth Magle
Bill Miller
Patrick O'Neal
Jim Olds
Aaron Packman
Mike Papka
Seongha Park
Ismael Perez
Bhupendra Raut
Dan Reed
Mike SanClements
Raj Sankaran
Sean Shahkarami
Sergey Shemyakin
Joe Swantek
Helen Taaffe
Valerie Taylor
Doug Toomey
Frank Vernon
Rommel Zulueta