Sage Digital Continuum

From the Edge to the Supercomputer: Analysis along the way

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

Instrument

Data

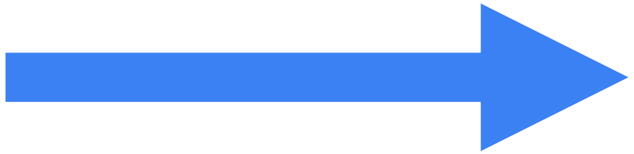

Analysis

The Digital Continuum

Instrument → **HPC/Cloud**

IoT **Facilities** **Analysis**

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

SAGE
sagecontinuum.org

Why Live on the Edge?

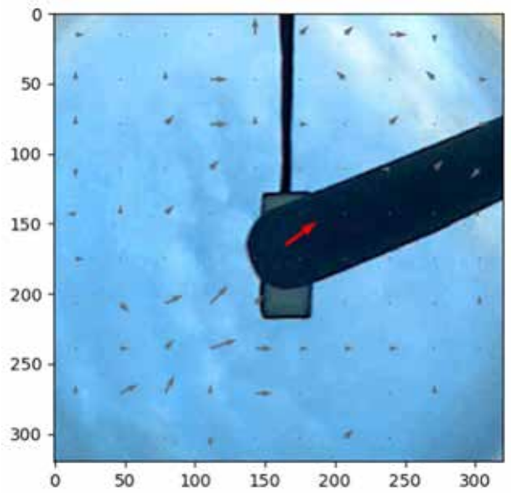

Analyse full resolution data, find highest value data for the science

- 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

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

AI@Edge: Digital Continuum

Sensors

LIDAR

Software Defined Radios

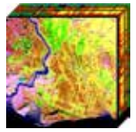

Hyperspectral Imaging

Facilities

Actuators

Servos

Dynamic adaptation

Edge Computing

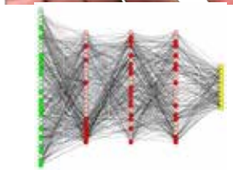

Scientific Data Analysis & Control

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

Sage Digital Continuum

What is a "Software Defined Sensor"?

<> "embedded computing"

Your software container running here

Analysis produces live results

AI-Based Measurement & Anomaly Detection, & Control

Plant Species

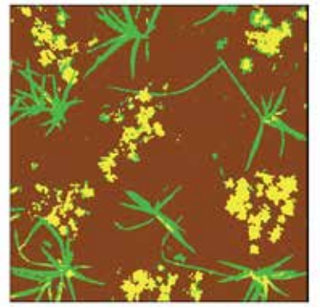

Pedestrian Flow

Birdsong

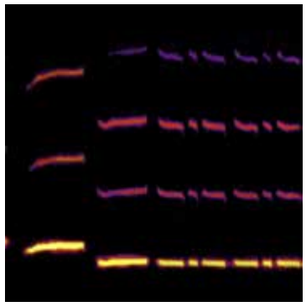

Traffic Flow

Wildlife

Cloud Motion Vectors

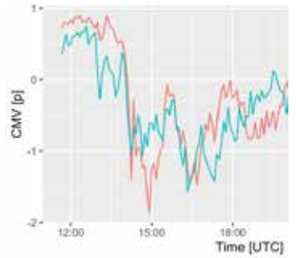

Wildfires: detecting smoke

Flooding / surface water

SAGE
sagecontinuum.org

Delivering AI@Edge: Two Forms

Wild Sage Node

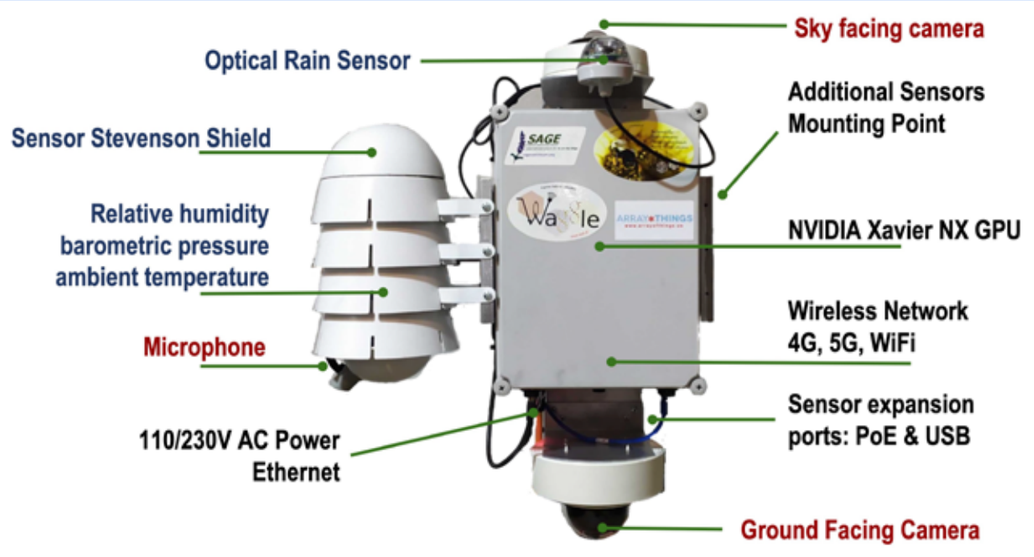

Ready for mounting ***outside***, any PoE sensor can be easily added

Sage Blade

(Sage software stack + pure commodity server)

Rugged server for instrument huts, new sensors easily added

"Waggle" is the core node platform that can be adapted for other uses.

SAGE

Cyberinfrastructure for
AI at the Edge
sagecontinuum.org

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

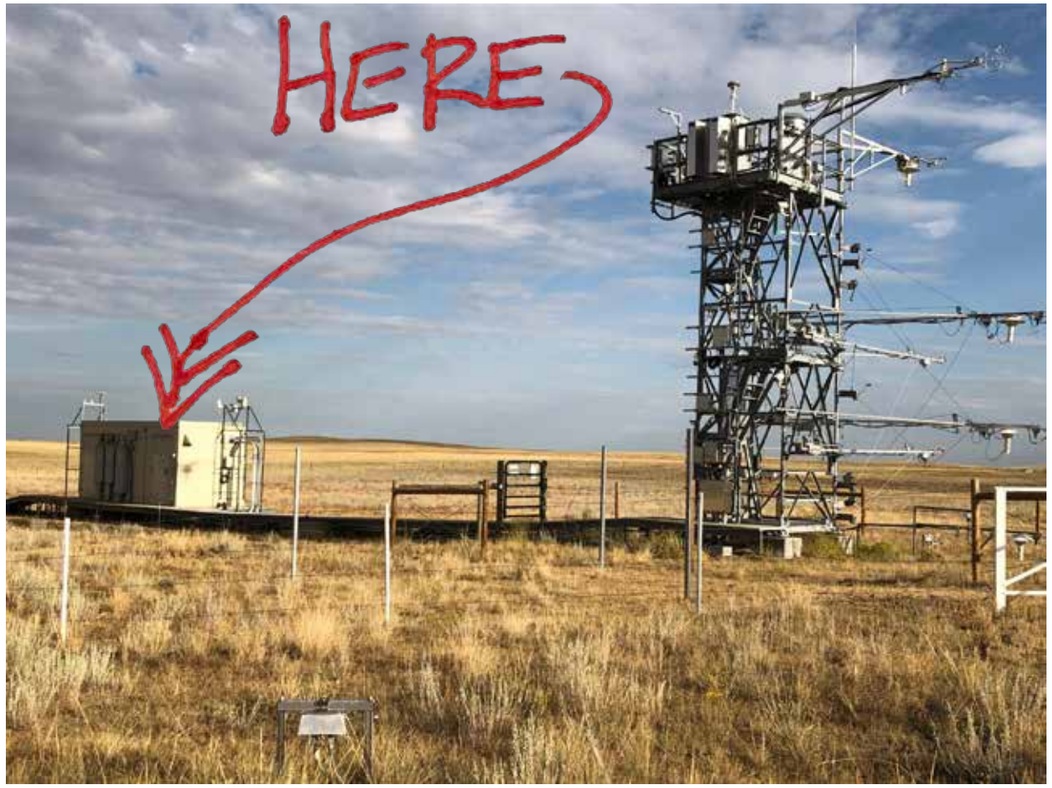

Analyse full resolution data, find highest value data for the science

Sage Software Architecture

Sage Nodes: “Cloud Native”

- Built upon standard AI Stack
- **Containers** on **Kubernetes**
- **Multi-tenancy**
- **“Goal-based” Scheduler**
- Local control for actuation
- Extreme cybersecurity
- Publish data to Beehive

Cloud Infrastructure

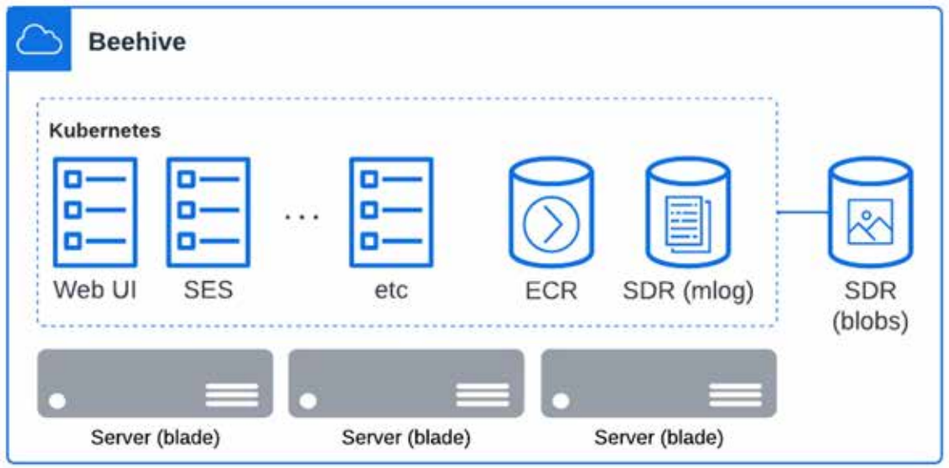

AI@Edge ”Plugin” from Edge Code Repository (ECR) (the “App Store”)

Beehive manages

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

SAGE
sagecontinuum.org

Building Community for AI@Edge Applications

The Edge Code Repository

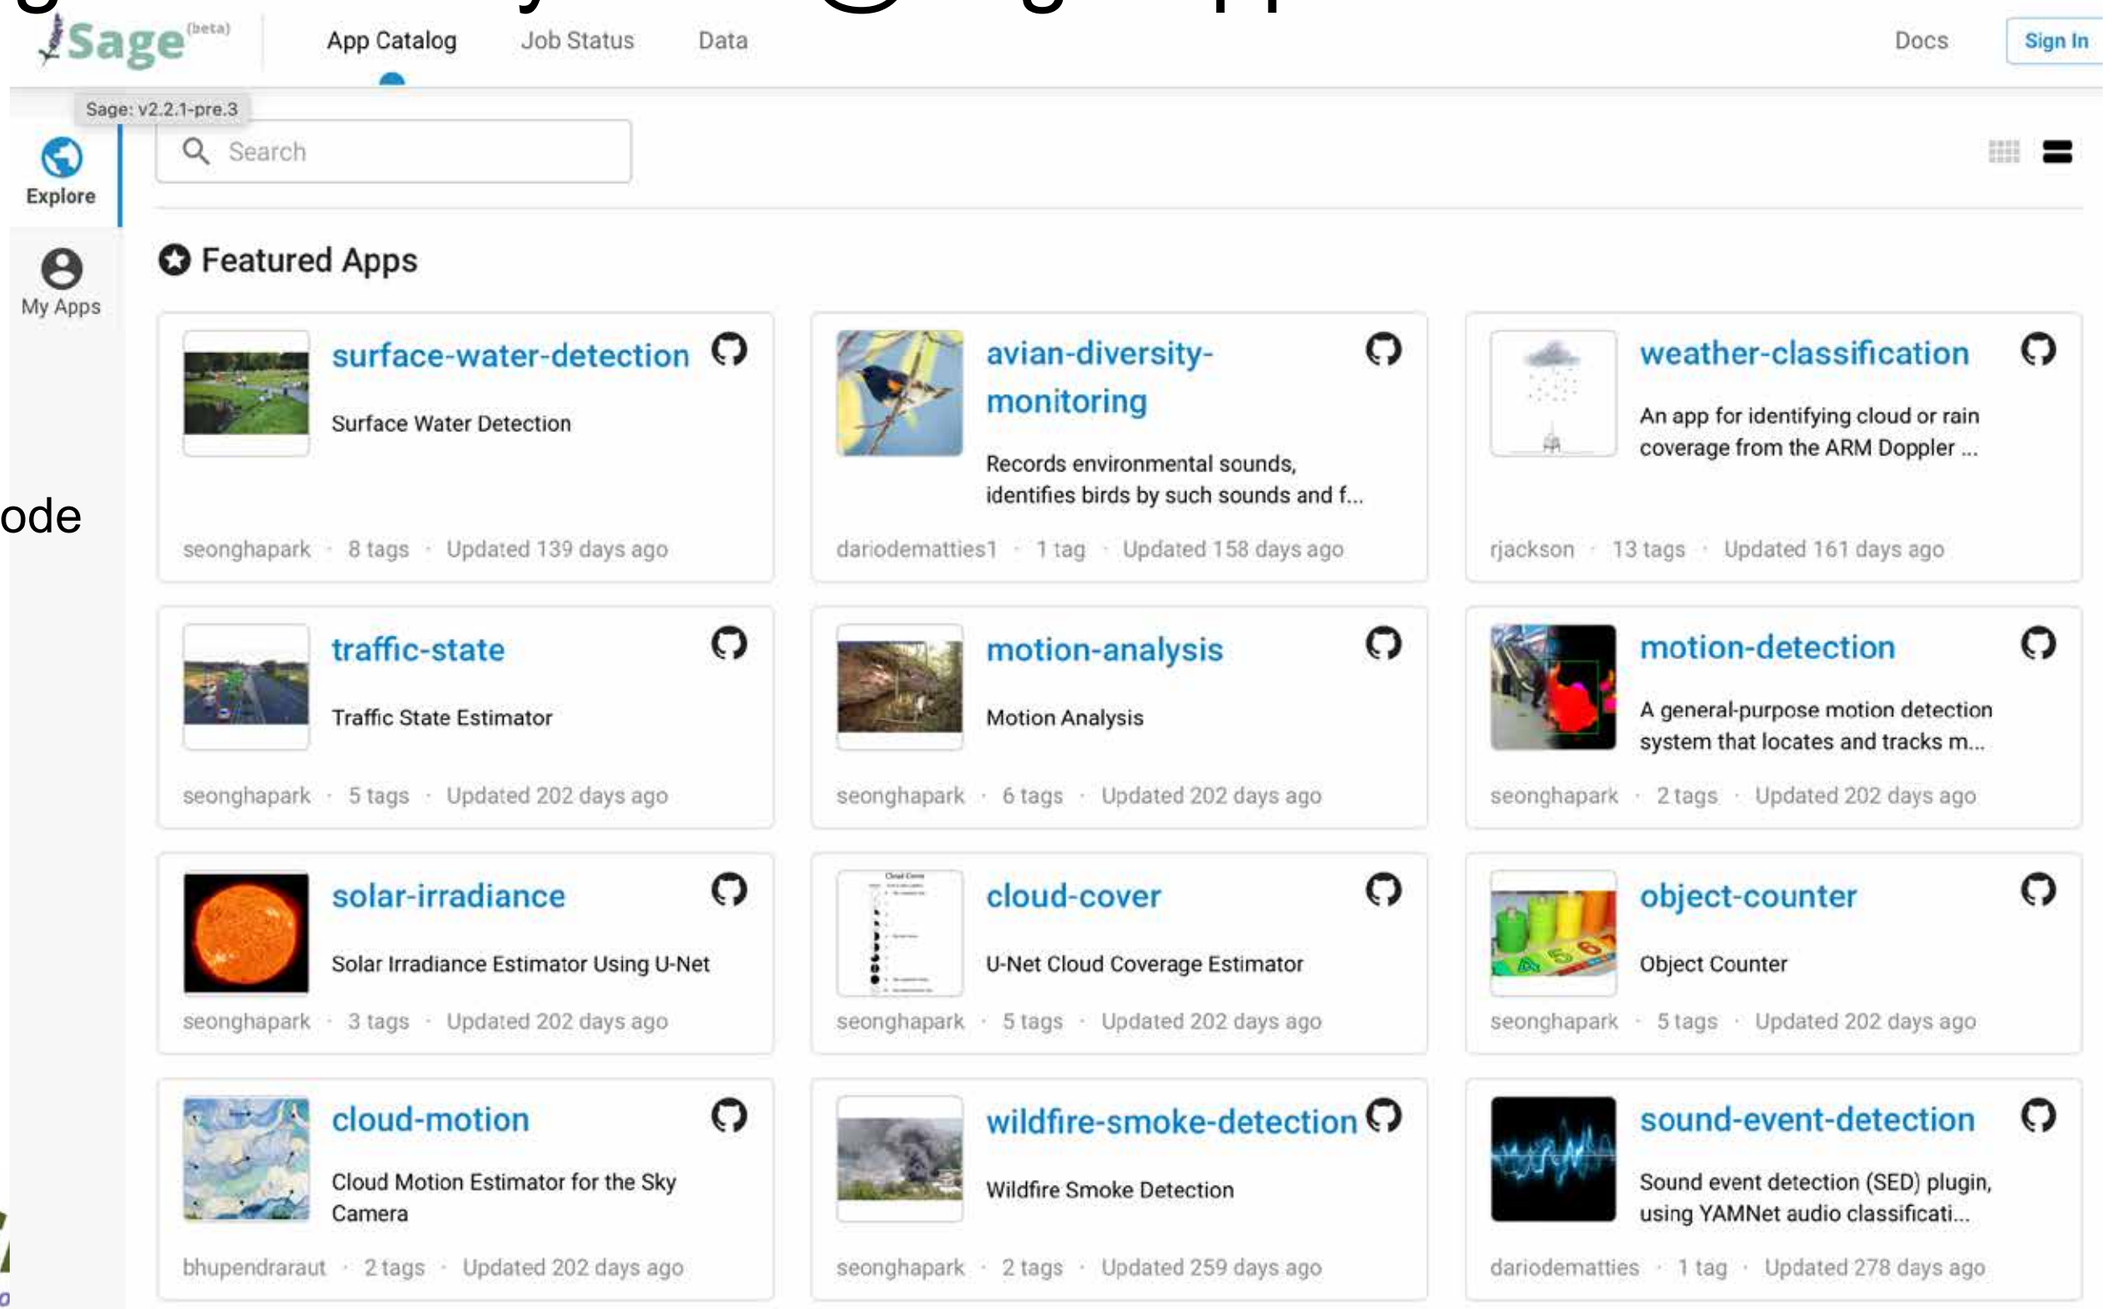

Sage Digital Continuum

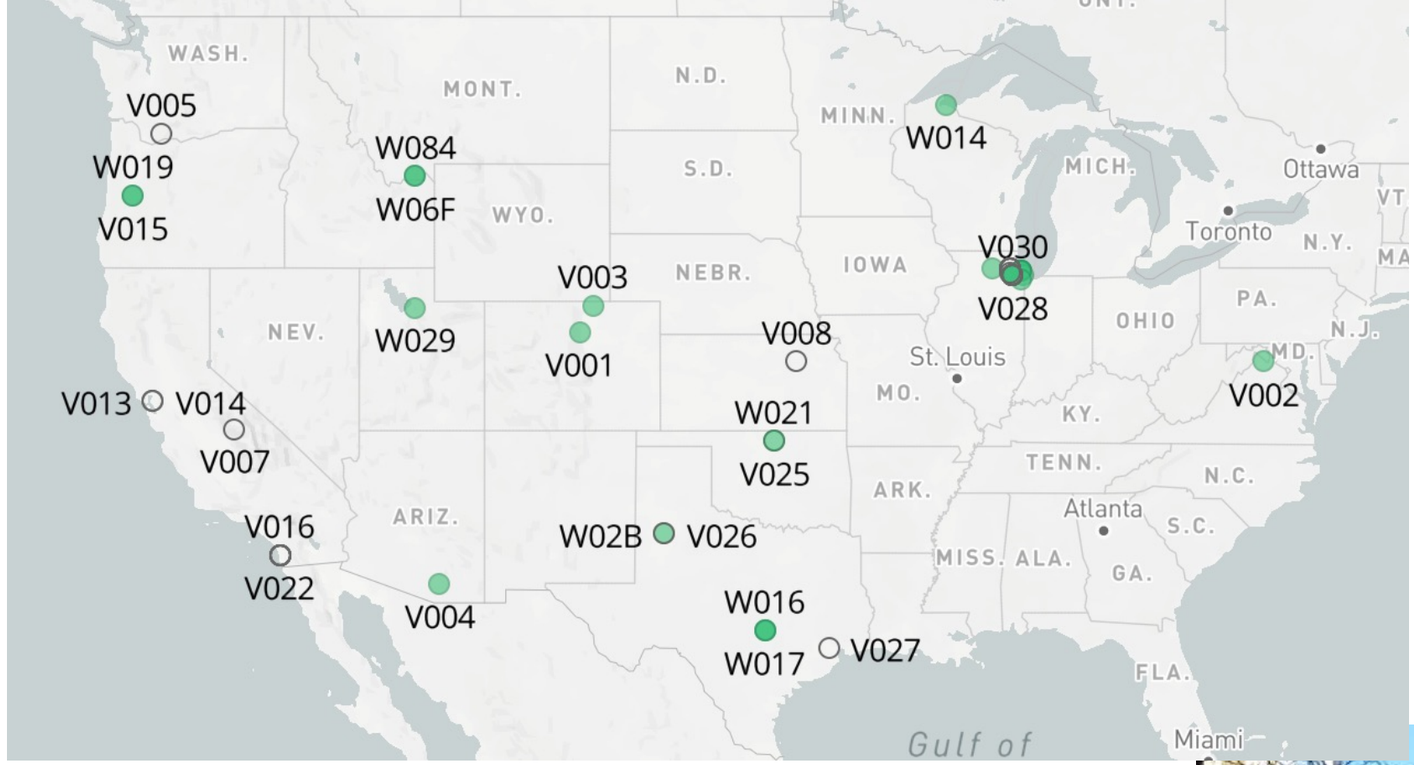

W015

E. 87th & S. Cottage Grove, Chicago, IL

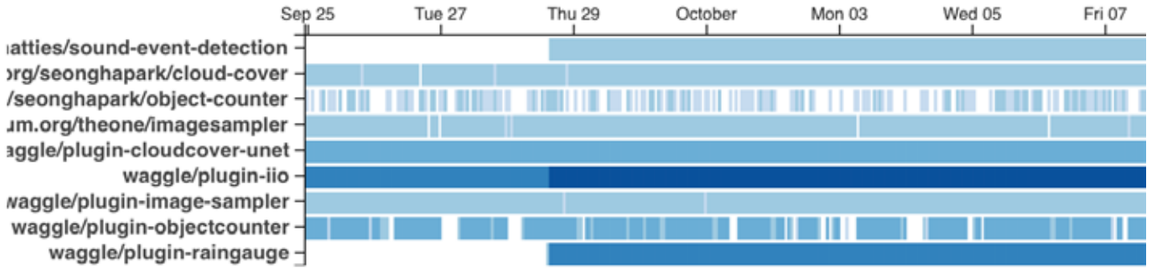

SAGE
sagecontinuum.org

Avian diversity monitoring

- 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

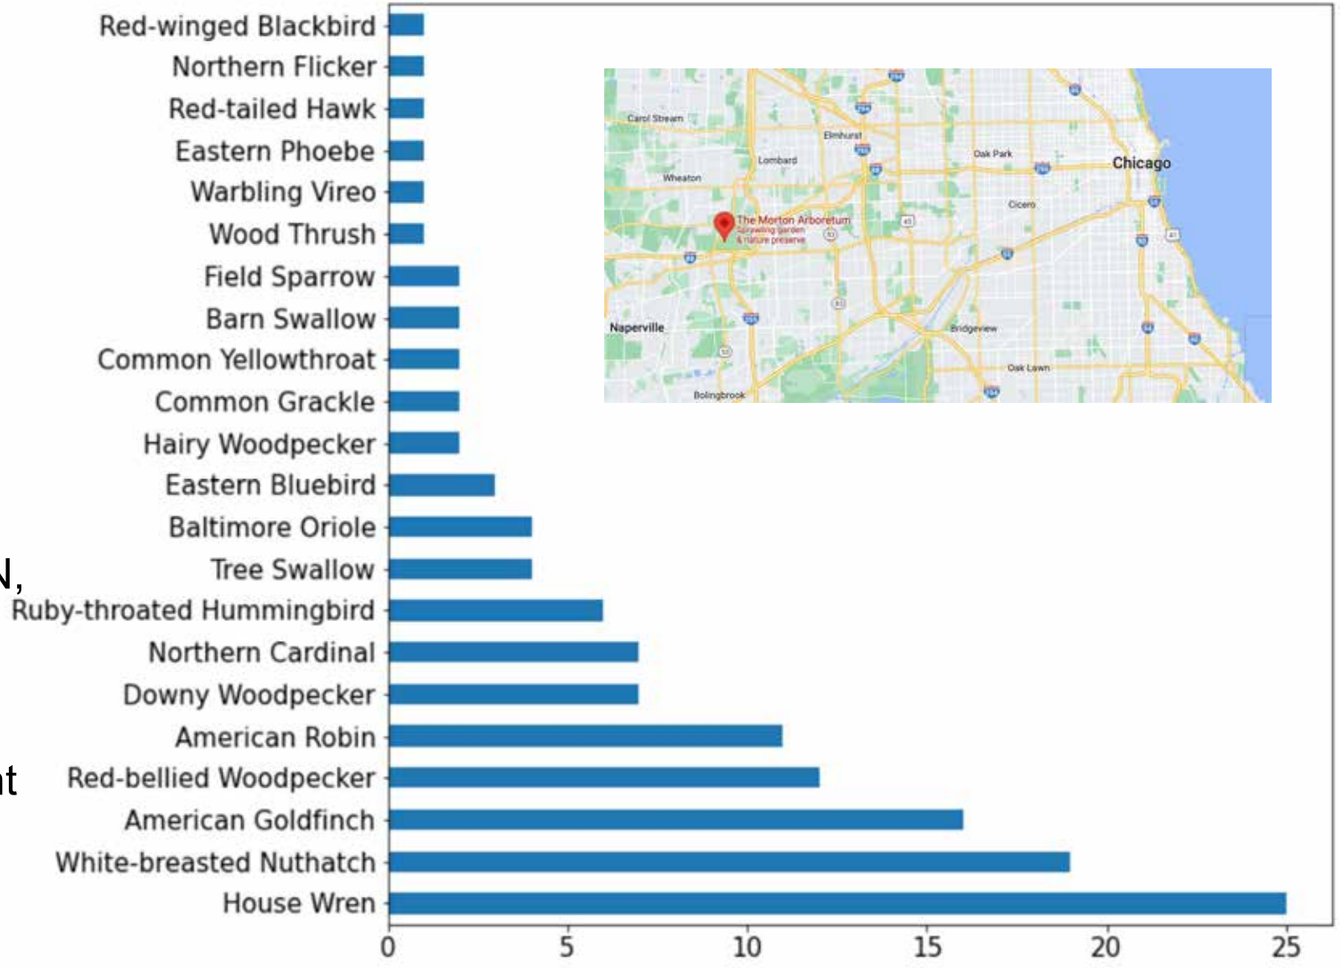

Morton Arboretum Avian Detection, June 28, 2021 (24 hour)

[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.

Research Credit: Dario Dematties, Bhupendra Raut, Nicola Ferrier

Self-supervised Avian Diversity Monitoring Joint-Embedding Architecture

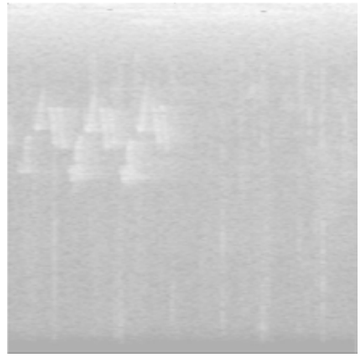

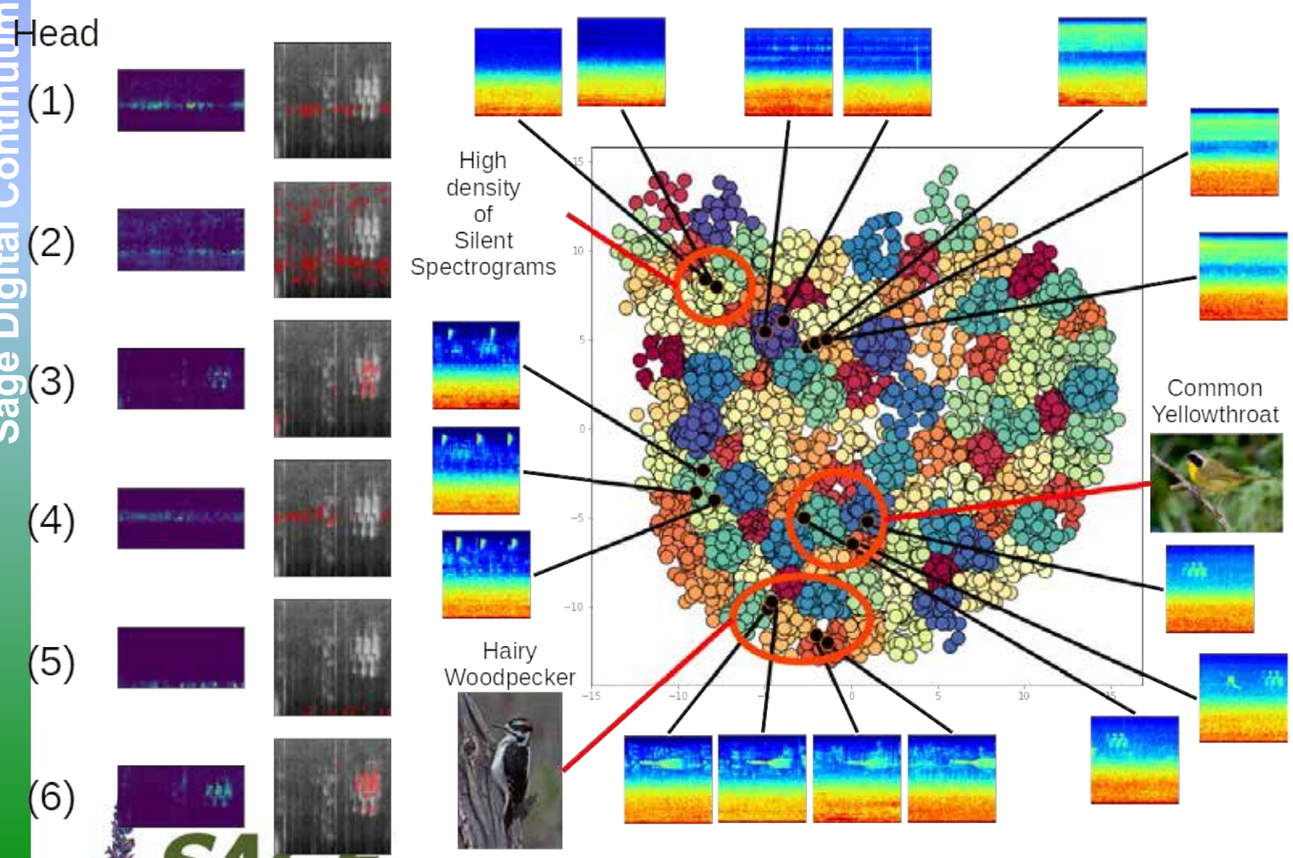

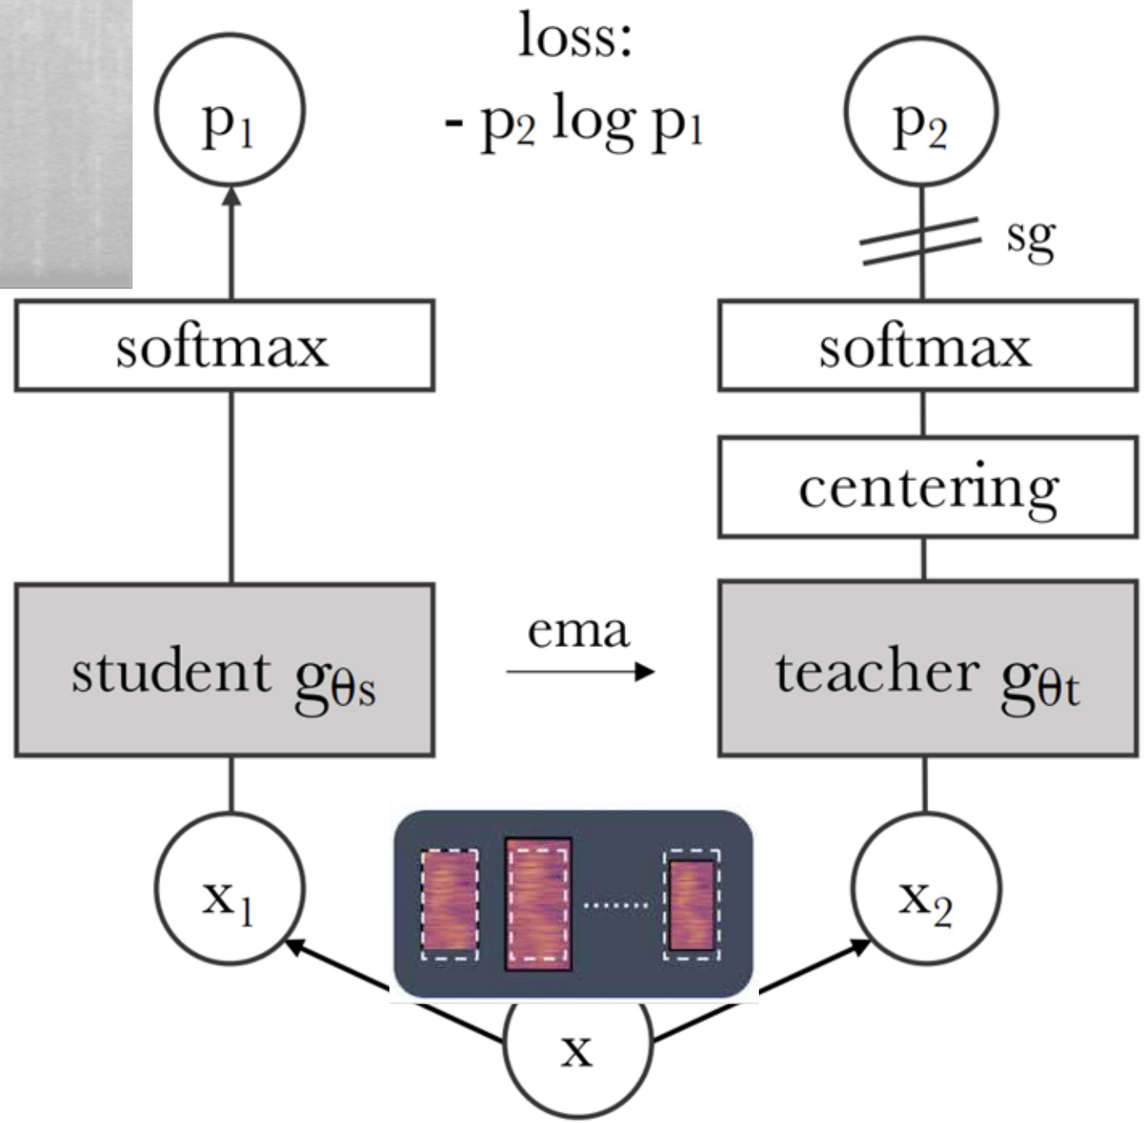

Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P. and Joulin, A., 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 9650-9660).

Edge computing unleashed on understanding climate

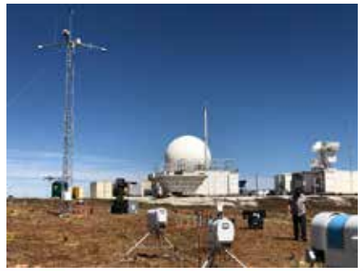

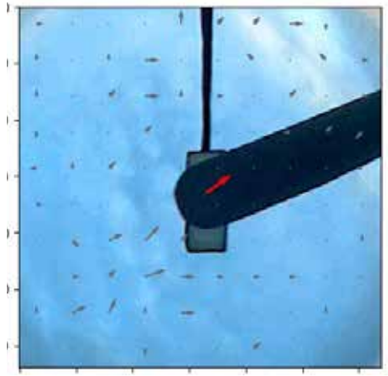

Analyzing cloud patterns

Self-learning Cloud Patterns

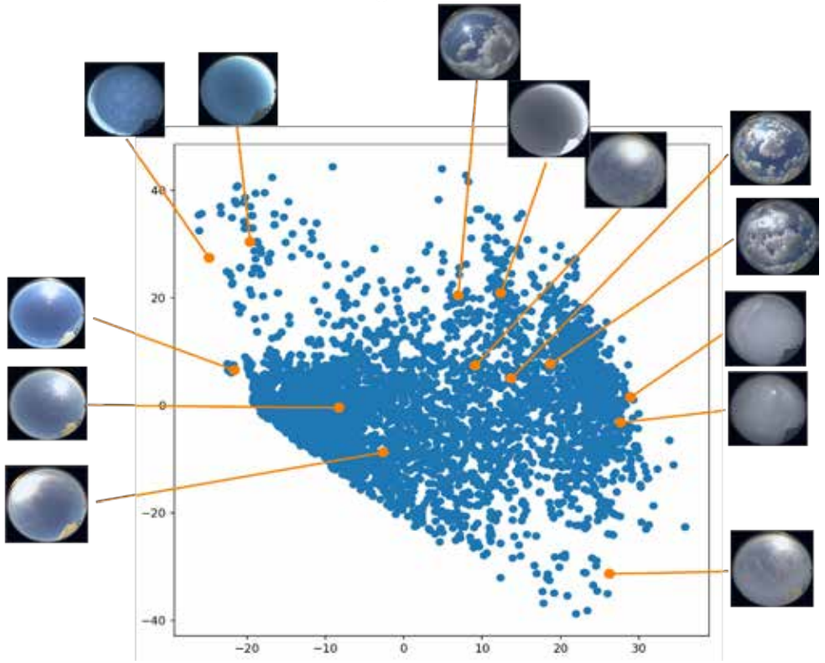

Cloud Motion Vector

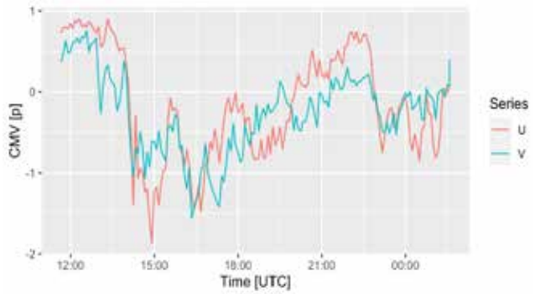

Cloud Cover Estimation

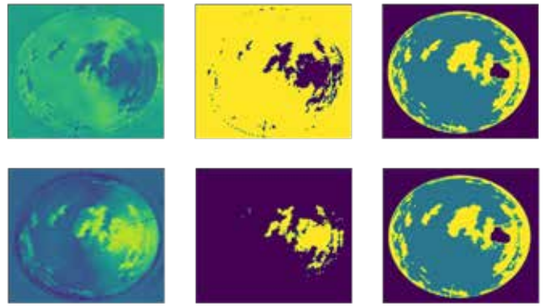

Automated LIDAR particle identification

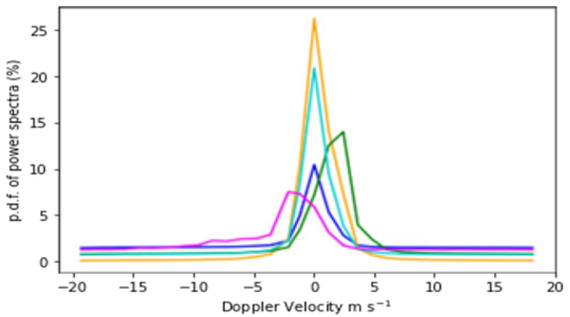

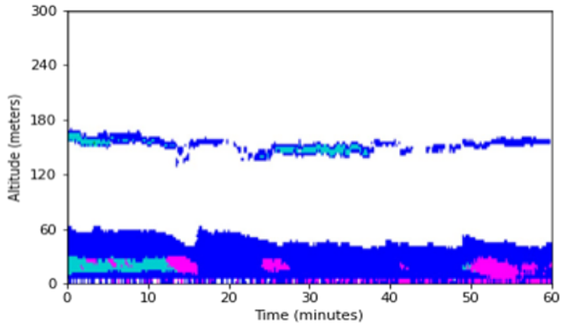

Solar energy estimation

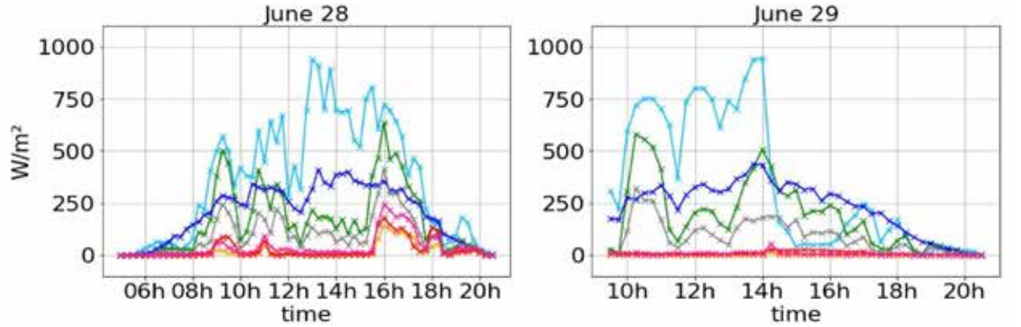

Sage Digital Continuum

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

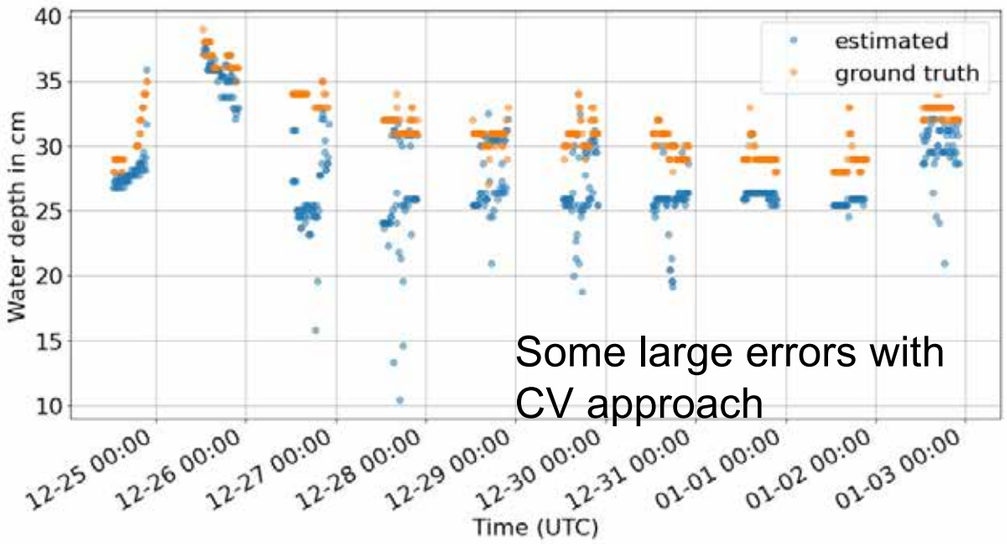

Human annotation using Labelbox

U-Net (segmentation)

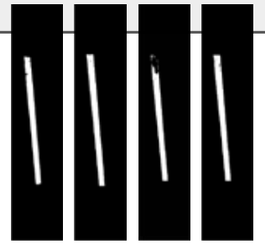

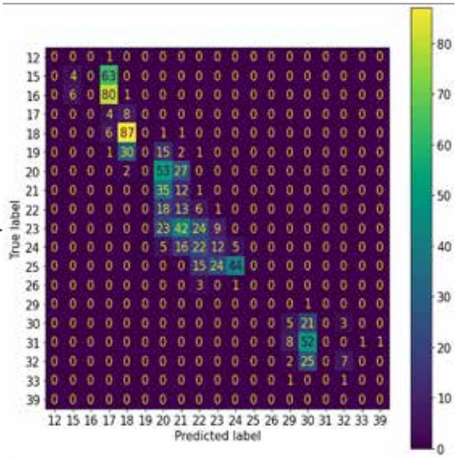

ResNet32 based Regression

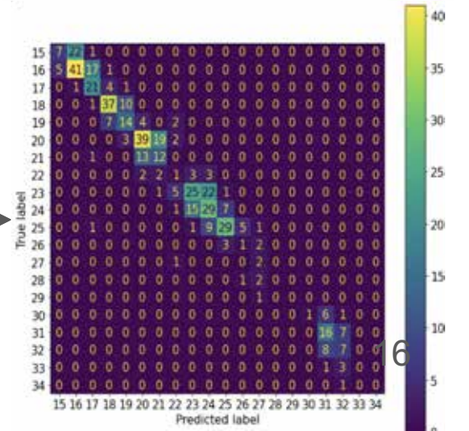

Measuring Water and Snow Depth


RGB image

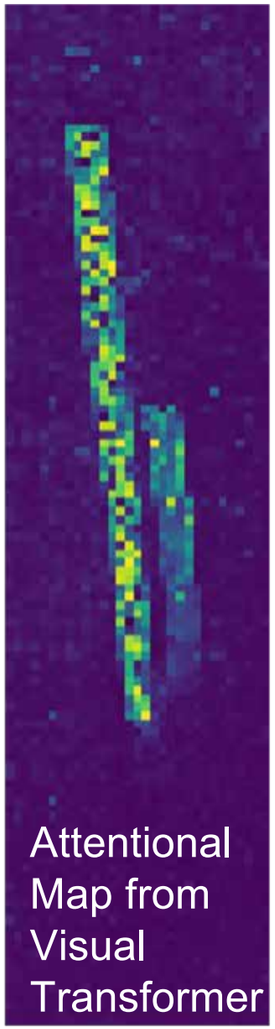
Attentional Map from Visual Transformer


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.

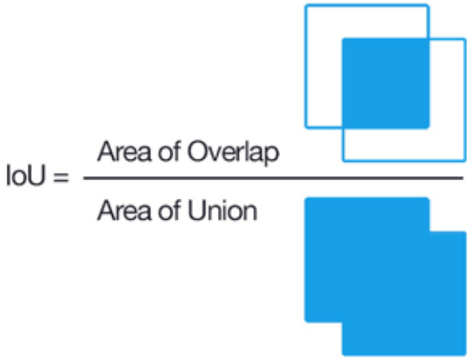

Intersection/Union (IoU) = 0.729


Inters.


Union

Dario Dematties, Northwestern University

PANDA (LBNL) and DAWN (ANL)
(testing network of AI@edge – collaborative decisions)

Wildebeest 5G project exploring integration

Science begins at the instrument

- **"Spectral Triage"** separates radon progeny spectra from static background
- NMF background model learned ab-initio
- Model re-trained upon request

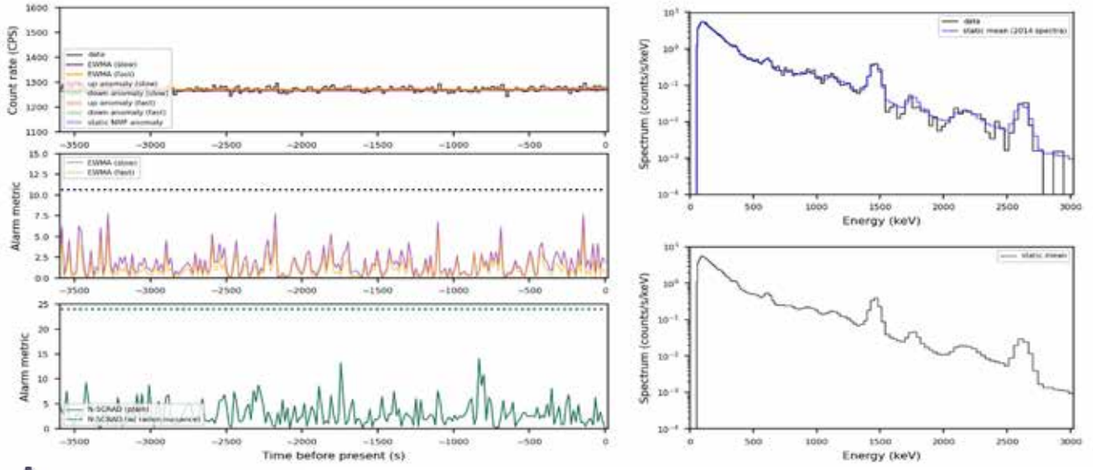

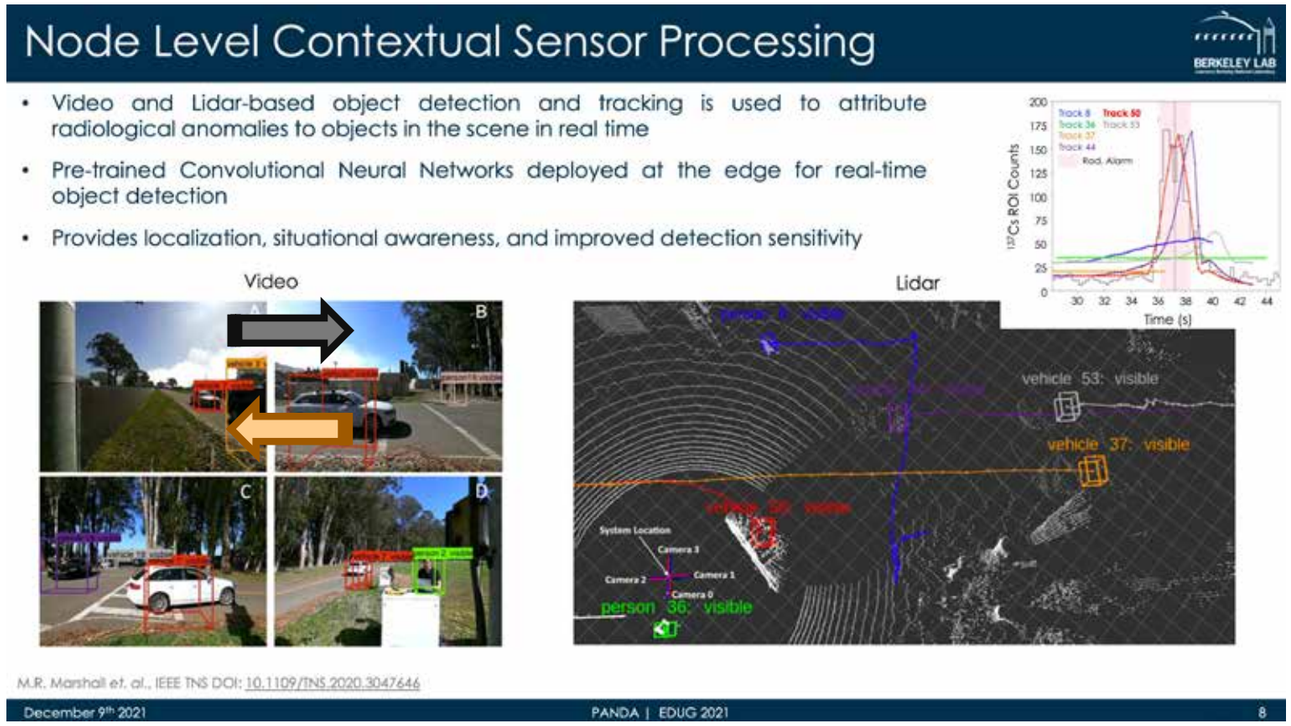

SAGE
sagecontinuum.org

Ren Cooper <RJCooper@lbl.gov>

Sage Digital Continuum

NNSA: 2 DAWN Nodes @ RUSH Hospital (briefings with stakeholders every 2 weeks)

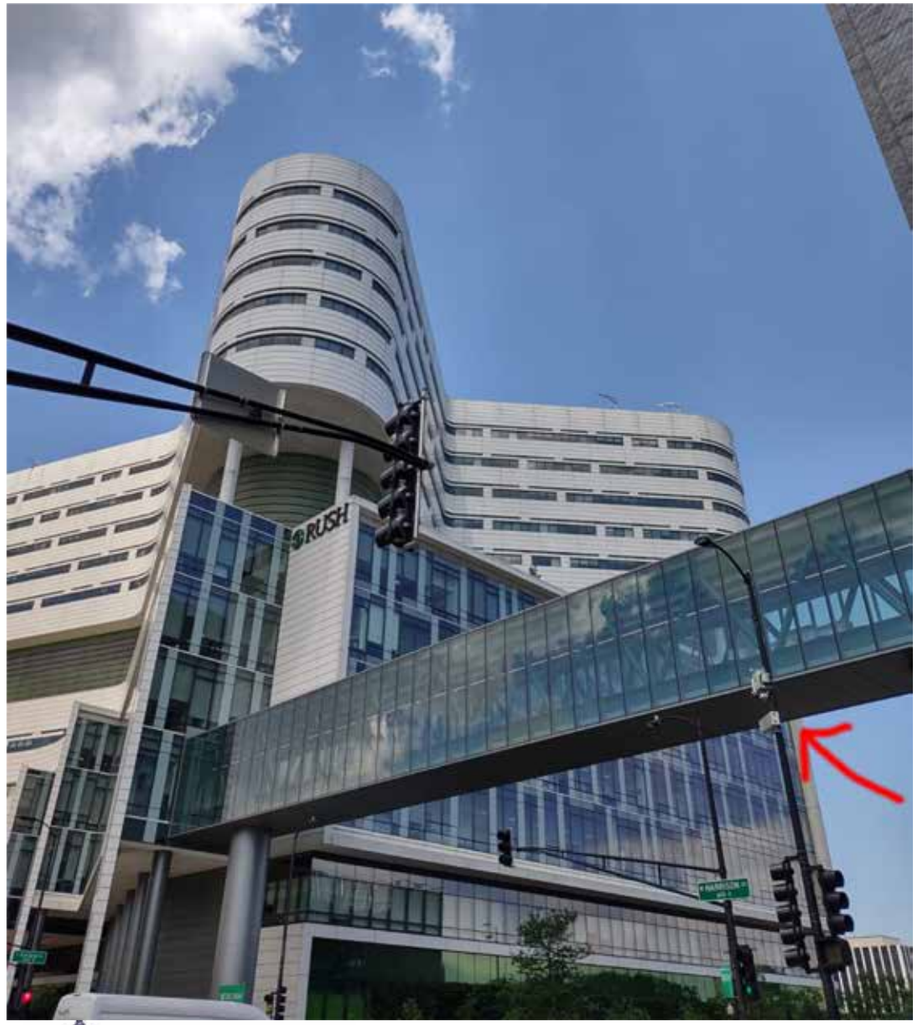

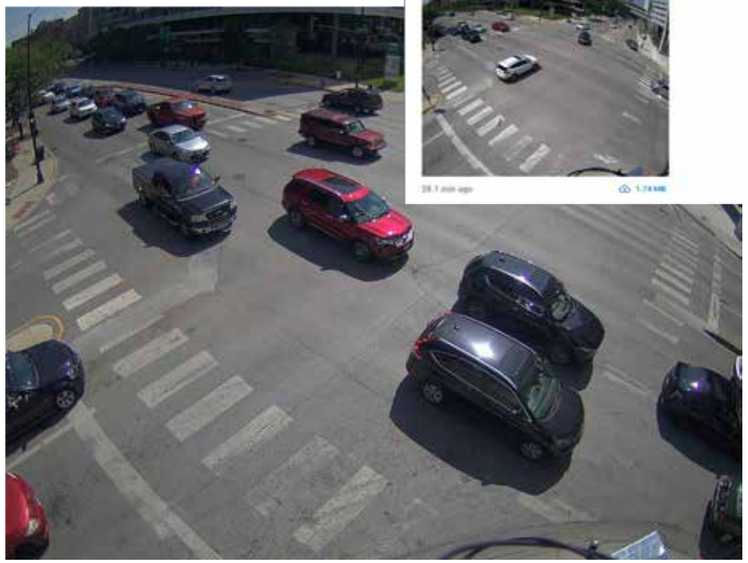

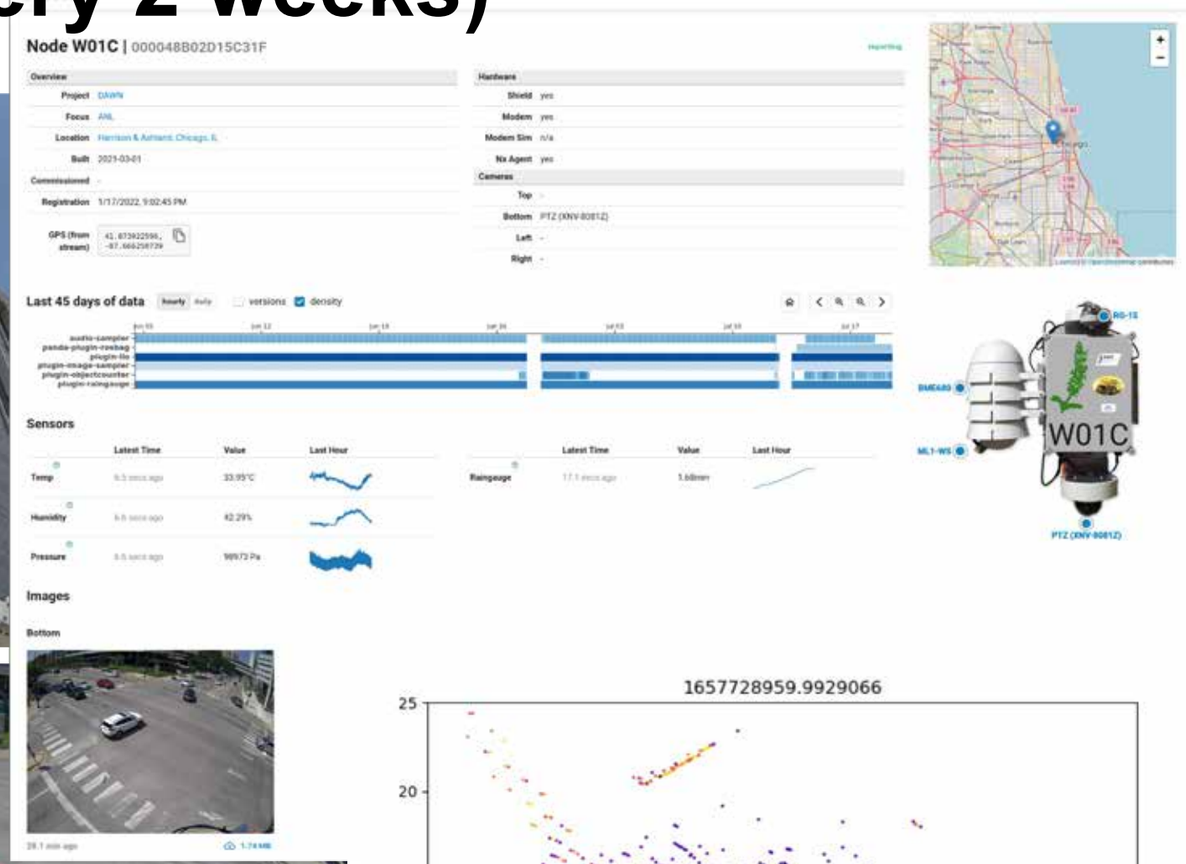

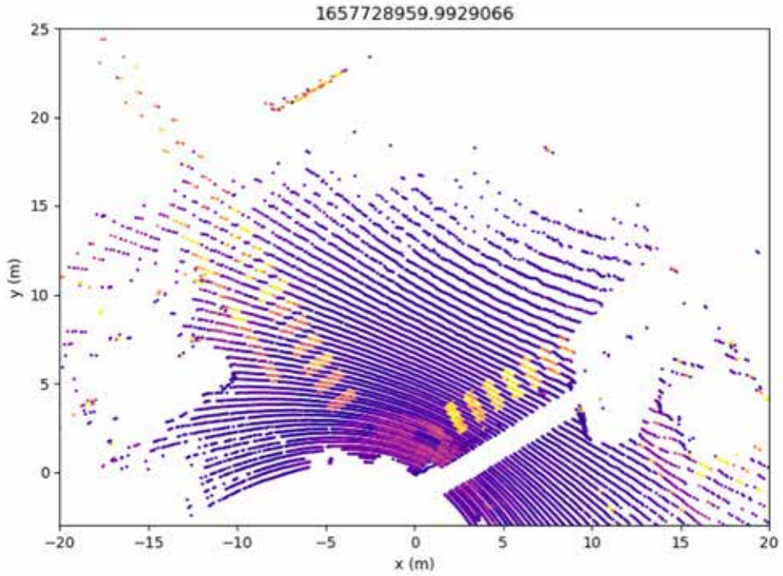

SAGE
sagecontinuum.org

Sage Digital Continuum

Surface Water Detection

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

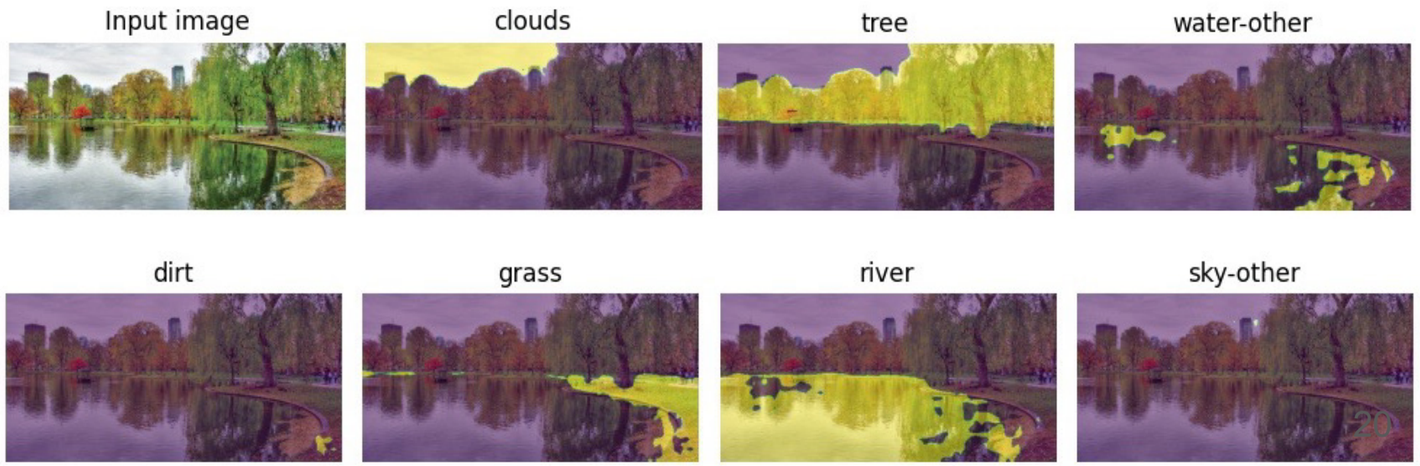

SAGE
sagecontinuum.org

Nicola Ferrier, Uchicago

Sage Digital Continuum

Ilkay Altintas, UCSD, Co-PI for SAGE

Wildfire Detection and Prediction

Exploring wildfire detection at the edge linked to HPC simulations

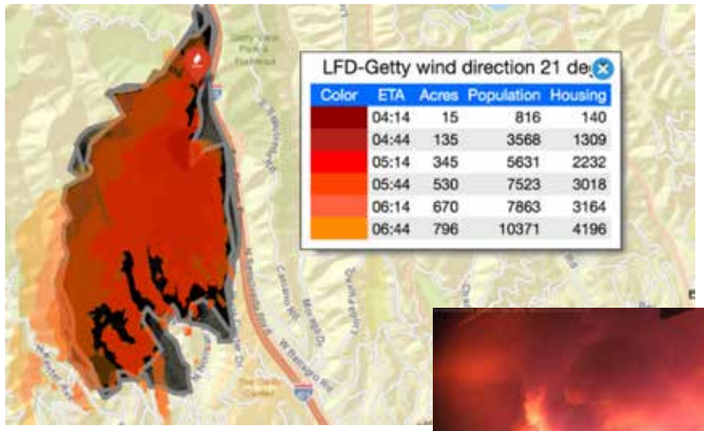

ALERTWildfire: A unique wildfire detection and monitoring system

UNIVERSITY OF OREGON

Collaboration: Doug Toomey, UOregon

OHAZ

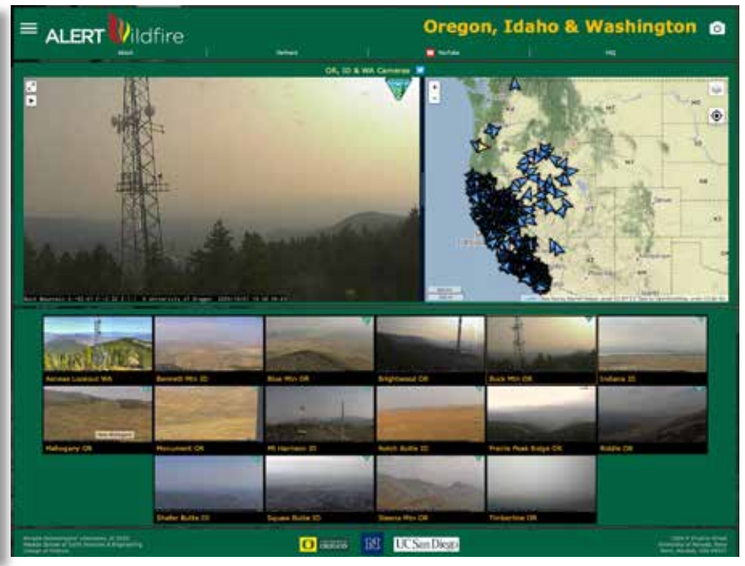

Sage Co-PI: Eugene Kelly, Colorado State
eugene.kelly@colostate.edu

NEON Mobile Deployment Platform (MPD) with Sage Konza Prairie for controlled burn: April 2022.

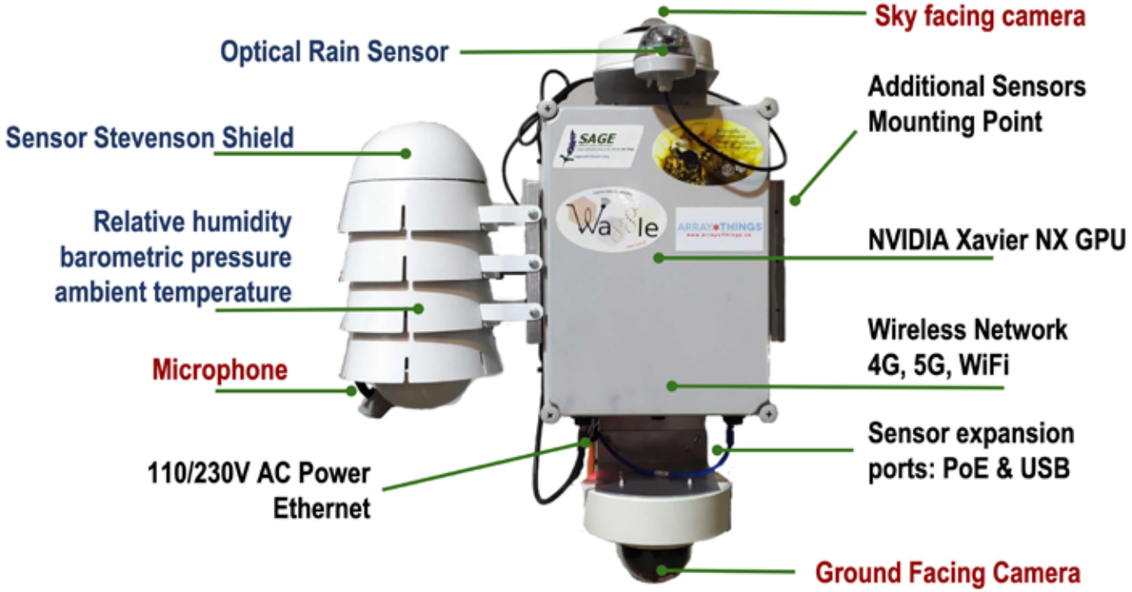

Special Thanks: NEON Team!
Rommel Zulueta @ Battelle

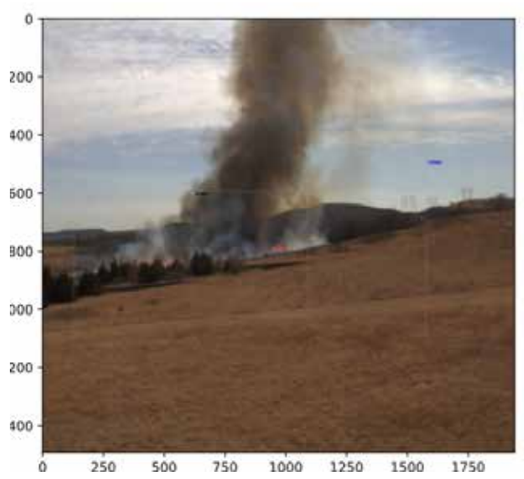

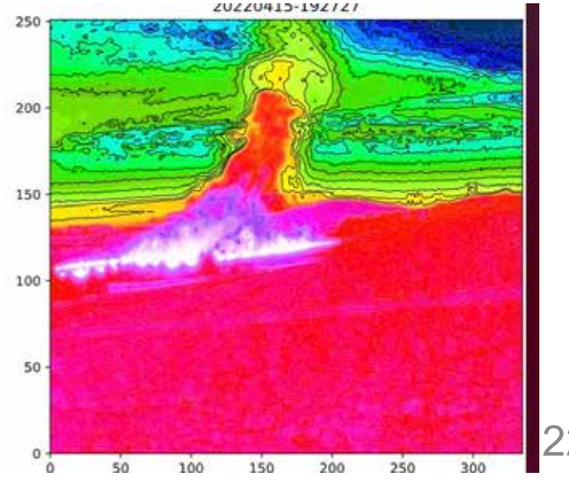

SAGE
sagecontinuum.org

Sage NEON NSF Controlled Burn

Konza Prairie, Kansas


HONDA EU7000is

NSF Coastlines and People

Strengthening Resilience of Manoomin, the Sentinel Species of the Great Lakes, with Data-Science Supported Seventh Generation Stewardship
PI: Josiah Hester

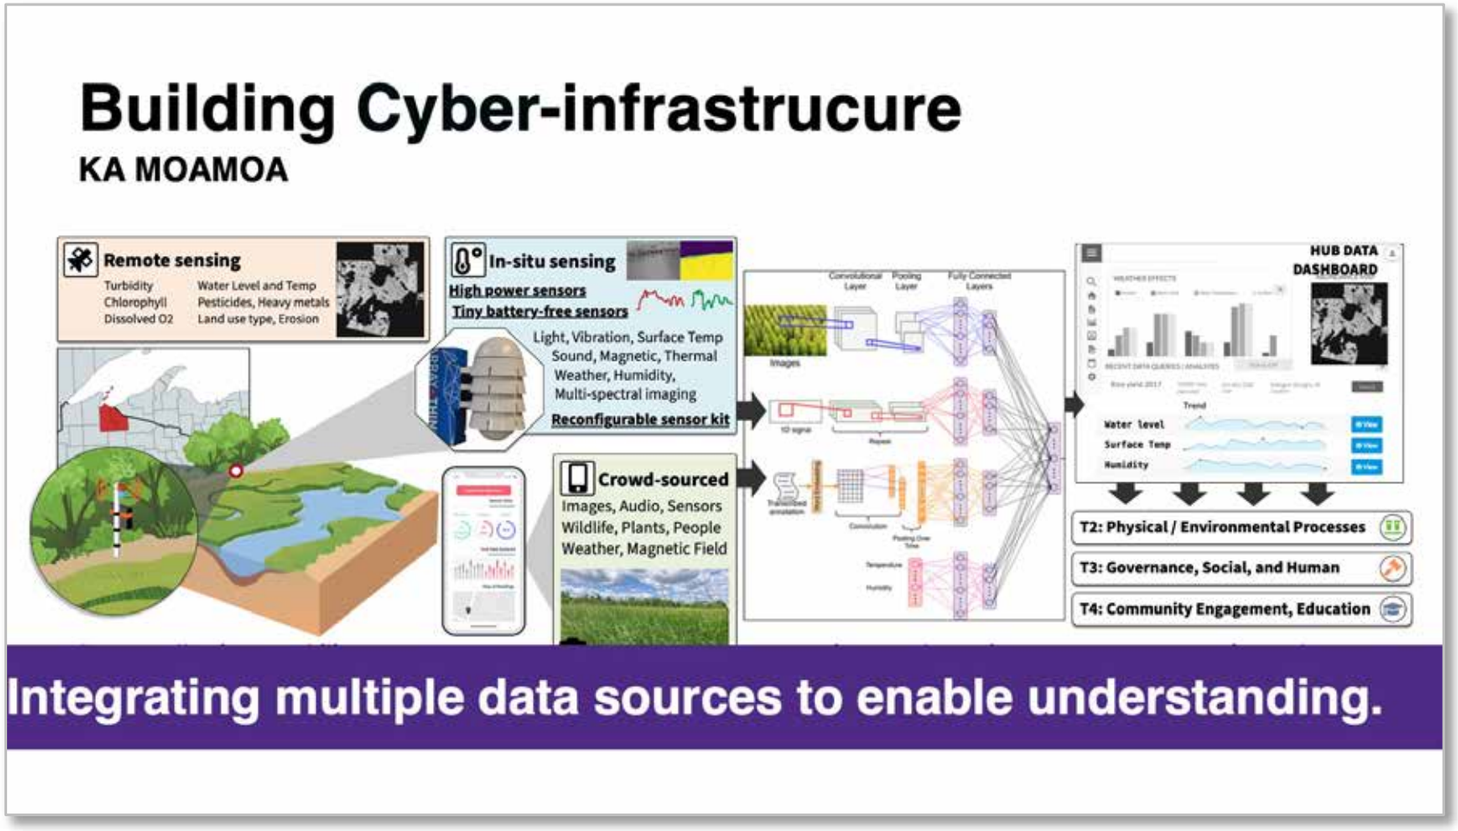

Jonathan Gilbert, Biological Services Director, Great Lakes Indian Fish & Wildlife Commission (GLIFWC)

Sensor/Instrument — HPC/Cloud

Can we leverage the tools and architectures from decades of experience building supercomputers for the continuum?

- AI@Edge open source software
- Edge resource allocation
  - Node sharing & ”Will it fit?”
- Goal-based edge scheduler
- Operating System management
- Resilient data movers

- HPC open source software stack
- HPC allocation
  - CPU hours, storage, etc.
- Job scheduler for shared access
- Operating system management
- Storage services & data movers

SAGE
sagecontinuum.org

A Challenge:
We Need a Programming Model…. (think shmem :-)

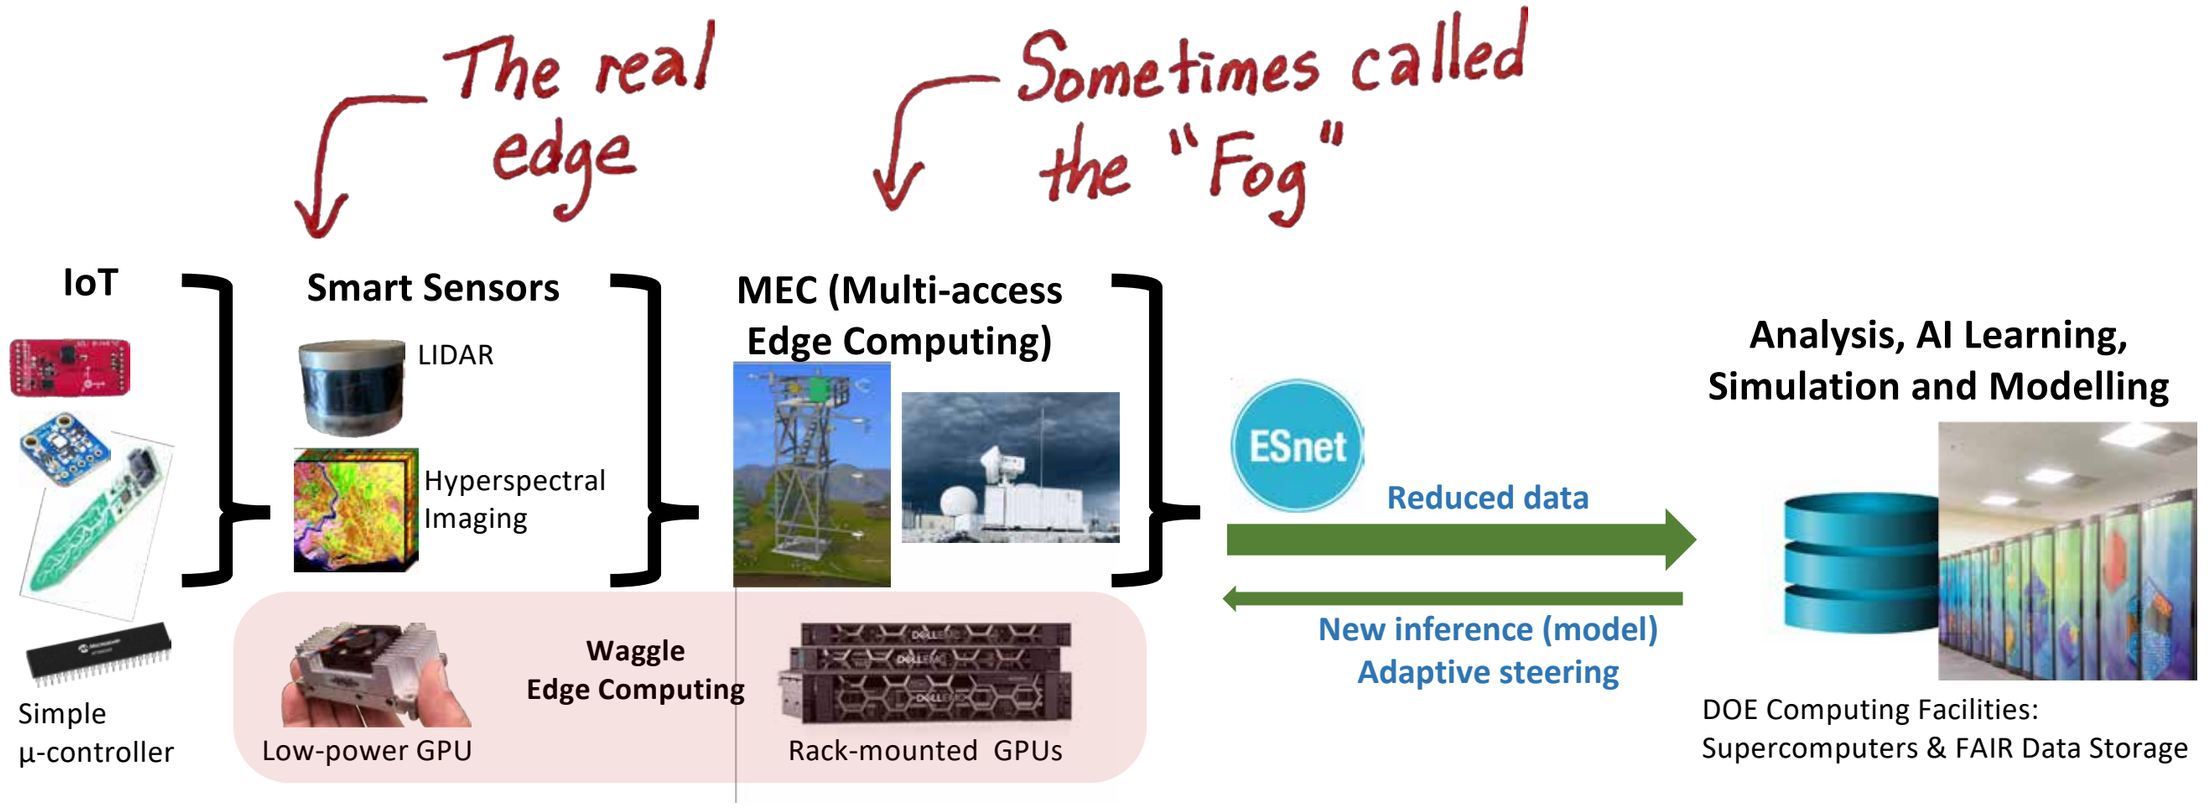

SAGE
sagecontinuum.org

Multi-tenancy with Sage Edge Scheduler

- Jobs include a “Science Goal”
- Examples of science rules for different science studies
  - sampling: *cronjob(“15m”), cronjob(“14:00”)*
  - trigger-based: *value(“temperature”) > 32*, *value(“raingauge.uint”) > 3*
  - event-based: *event(“big-storm”), event(“fire-season”)*

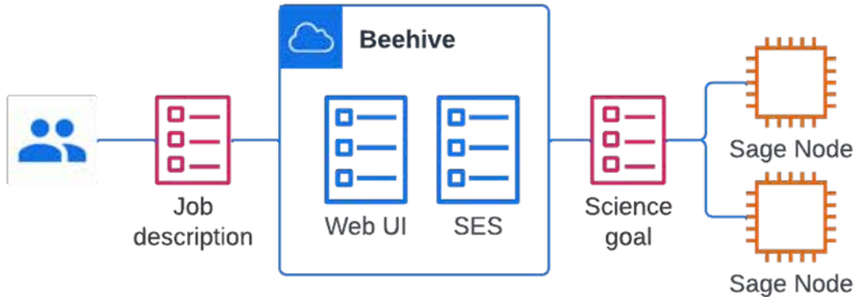

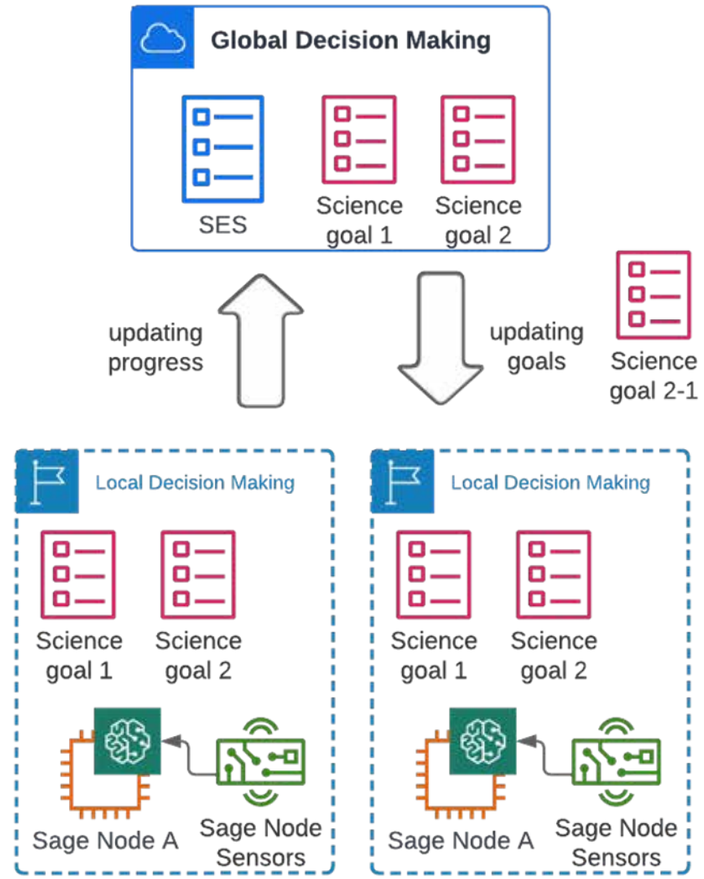

For more detailed design, see paper:
https://www.sciencedirect.com/science/article/pii/S0743731522001009

```
yongho.kim@node1:~$ sesctl stat
JOB_ID  NAME                    USER      STATUS     AGE
=================================================================
9       sage-ttu-job            waggle    Submitted  138h47m28s
10      sage-arm-ok-job         waggle    Submitted  138h47m25s
11      sage_mt_image_sampler   yongho    Submitted  114h58m26s
2       vto-traffic-state       yongho    Submitted  239h49m40s
6       sage-uoregon-job        waggle    Submitted  138h47m39s
1       vto-video-sampler       yongho    Submitted  305h59m36s
5       vto-job                 waggle    Submitted  138h47m46s
7       sage-urban-job          waggle    Submitted  138h47m35s
8       sage-utah-job           waggle    Submitted  138h47m31s
```

Research Credits: Yongho Kim, Northwestern University

Sharing the Node: Multi-tenancy

“Will it Fit” and automatic performance tuning
Collaborative with with UOregon TAU team

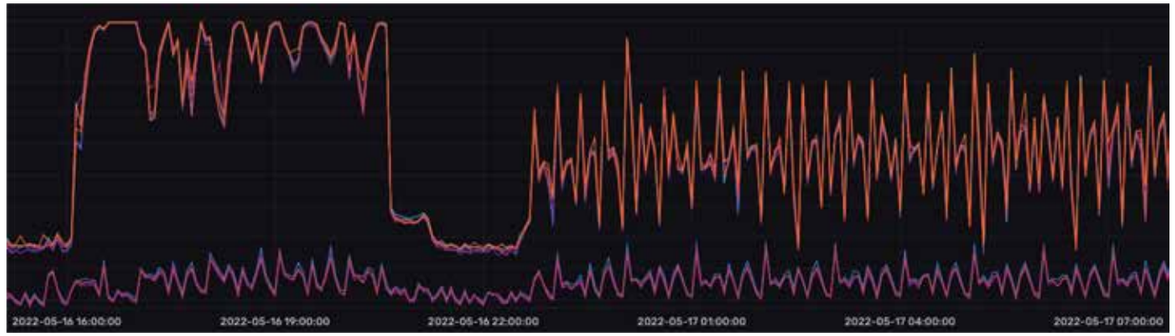

W019

1264 Franklin Blvd, Eugene, OR

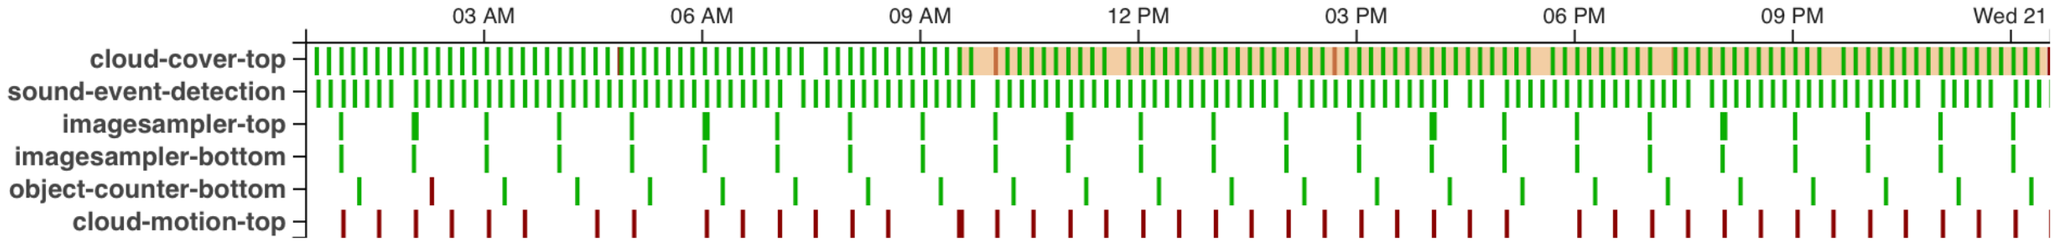

W02B

3003 15th St, Lubbock, TX

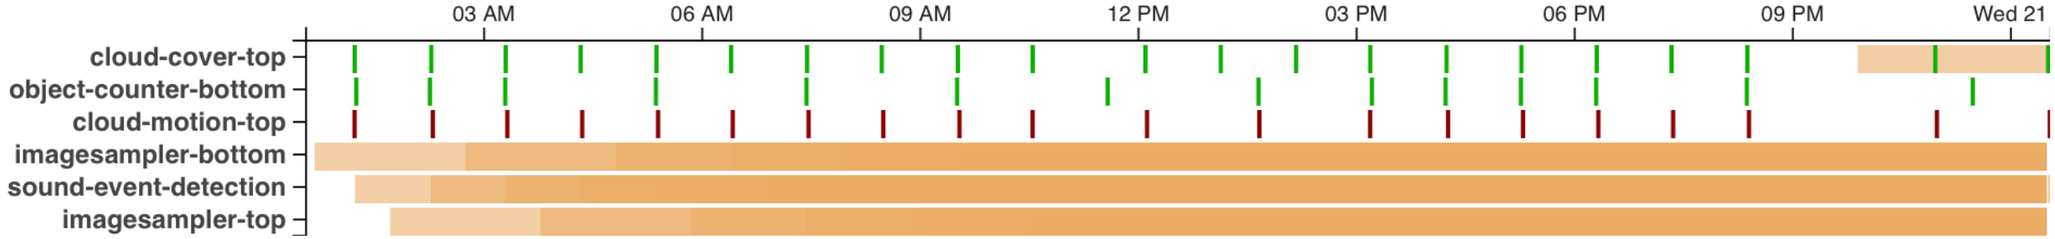

Tinkering with CLIP?

There's been a lot of excitement in the community about image *generation* using DALL-E 2 and Stable Diffusion…



CLIP is a building block used internally by some generative models to drive training by scoring how well an image matches a text description!

It's trained on the *many* text-image pairs available to provide meaning to text.

How can we best use this? “Query@Edge”?

CLIP is useful for quickly experimenting with new kinds of detection without having to collect data and train a model - think minutes instead of hours or days!

```
python3 main.py \
    "a person drinking coffee" \
    "a person making a call" \
    "a person jogging" \
    "a construction crew fixing the road" \
    "a red sports car" \
    "a busy intersection"
```

```
12.499 0.000   a red sports car
16.413 0.001   a busy intersection
17.943 0.006   a construction crew fixing the road
20.251 0.065   a person jogging
21.546 0.237   a person making a call
22.612 0.690   a person drinking coffee
```

AI@Edge science problems for students.

- Measuring river depth against graduated marker
- Auto-steering of PTZ cameras based on local AI
- Measuring snow depth against graduated marker
- Measuring vegetative states, growth rates
- Self-supervised learning: IR, LiDAR, audio, and RGB
- Vehicle types and flow speeds
- Quantify flower blooming (color, count)
- Outlying conditions from previous sensor data
- Calculating biodiversity based on audio
- Measuring surface water coverage
- Measuring lightning via RF (software defined radios)
- Measuring visibility across a field
- Measuring rime ice thickness
- Measuring ice coverage on a large body of water
- Measuring water flow speed
- Classifying wildlife behaviors
- Improved wildfire detection algorithms
- Wildlife tracking in open fields (speed, direction, count)
- Ultrasonic bat detection
- Measuring pedestrian movement dynamics
- Measuring land changes (riverbeds, plant coverage)
- Measuring water turbidity, debris movement, floating waste
- Measuring vehicle dynamics: identification of sliding, crashes, mishaps
- Measuring bike usage, bike lane dynamics
- Identifying urban "near misses"
- Measuring bird flocks and dynamics

Exciting, Hard, Challenging, CS Problems:

From Instrument to the HPC/Cloud

- **Programming model for the Digital Continuum**
- Lightweight AI training, federated learning at edge
- Self-supervised learning with multiple instruments
- Microelectronics for low-power AI@edge
- Container technology for Cloud/HPC and the edge
- Cooperative sharing: multi-tenancy
- Control loops for actuation, steering
- Movement (drones, robots)
- Digital twin / MODEX for setting local edge goals
- Advanced networking: wireless and satellite

Special Thanks

Students!

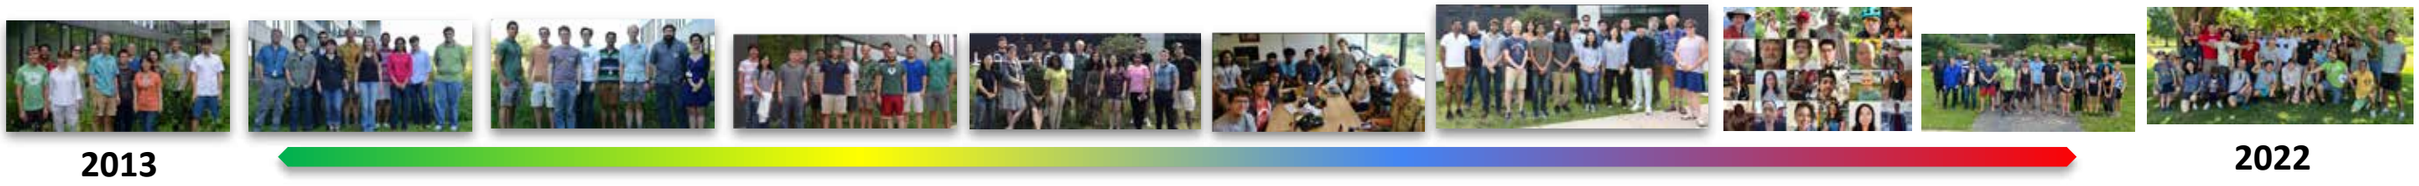

Ilkay Altintas
Kathy Bailey
Daniel Balouek-Thomert
Pete Beckman
John Blair
Eric Bruning
Adam Brust
Charlie Catlett
Scott Collis
Neal Conrad
Geoff Davis
Dario Dematties
Nicola Ferrier
Jannick Fischer
Larry Hartman
Robert Jackson
Eugene Kelly
Yongho Kim
Nick Maggio
Seth Magle
Bill Miller
Patrick O'Neal
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

Questions

Sage Digital Continuum

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
Waggle Github - https://github.com/waggle-sensor
Sage Continuum Github - https://github.com/waggle-sensor

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