SAGE
sagecontinuum.org

NSF

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

Sage Digital Continuum

Sage Edge Computing Platform and Self-Learning AI at the Edge

Presenters:
Yongho Kim: Developer of Sage, Assistant computer scientist, Argonne National Laboratory
Dario Dematties: AI researcher in Sage, Postdoc, Northwestern University

Team: Pete Beckman (PI) and Ilkay Altintas, Charlie Catlett, Nicola Ferrier, Scott Collis, Rajesh Sankaran, Eugene Kelly, Jim Olds, Mike Papka, Dan Reed, 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 of THINGS
ARM
neon
National Ecological Observatory Network
HPWREN
SDSC
SAN DIEGO SUPERCOMPUTER CENTER
LINCOLN PARK ZOO
FOR WILDLIFE. FOR ALL.
OHAZ

TAPIA Workshop
Sep 20th, 2024

Historical scientific study and analysis …

Instrument **Data**

Katherine Johnson (née Coleman)

Analysis

Katherine Johnson (née Coleman), one of the first African-American women to work as a NASA scientist - played a key role in the mathematical calculations for John Glenn's orbital mission and made sure that the equations controlling Glenn's capsule were programmed accurately, ensuring a safe lift off and splashdown.

The Digital Continuum

Instrument → HPC/Cloud

IoT Facilities Analysis

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

Why Live on the Edge?

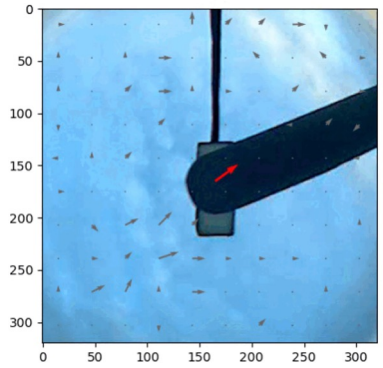

Analyse full resolution data, find highest value data for the science

- 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

Sage Goals

- Build a novel 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
    - Software-defined sensors
  - 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

Put AI@Edge

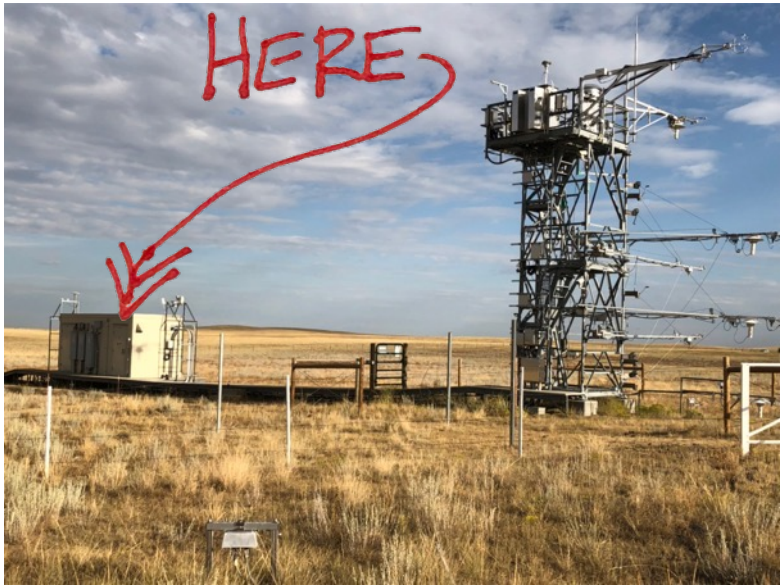

(Sensors sample at 40hz, aggregate to 30min)

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

SAGE
sagecontinuum.org

Building on
NSF Array of Things
(2016-2018)

AI-Based Measurement & Anomaly Detection, & Control

Sage Digital Continuum

What is a "Software Defined Sensor"?

Your software container running here

Analysis produces live results

Plant Species

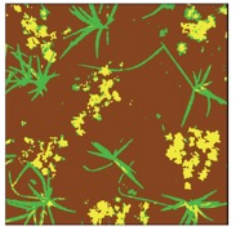

Helmet usage

Bird, Plane or Drone?

Traffic

Social Distancing

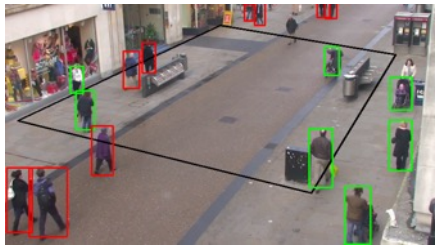

Wildfires: detecting smoke

Flooding / surface water

Many measurements cannot be "sensed" directly but can be computed from image, microphone or other devices or datastreams

SAGE
sagecontinuum.org

Delivering AI@Edge Platforms: Two+ Forms

Wild Sage Node

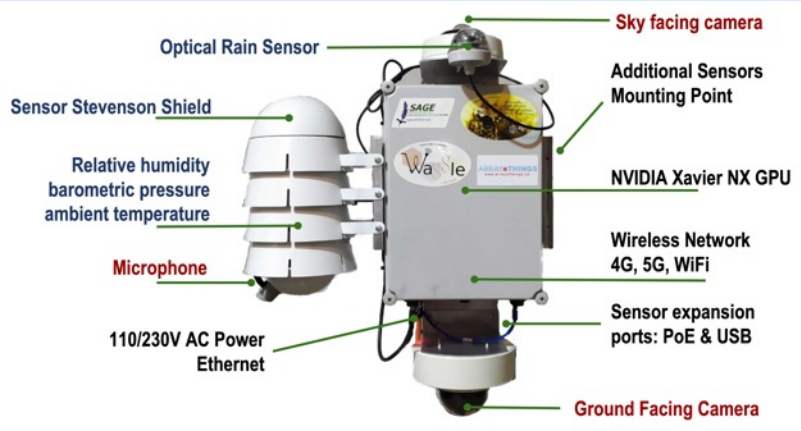

Ready for mounting ***outside***, any PoE sensor can be easily added

Sage Blade

(Sage software stack + pure commodity server)

Rugged server for instrument huts, new sensors easily added

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

Wild Sage Node: Manufacturing

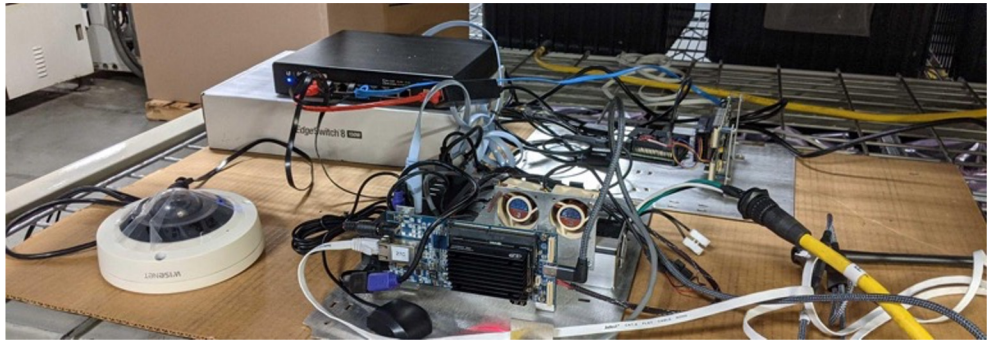

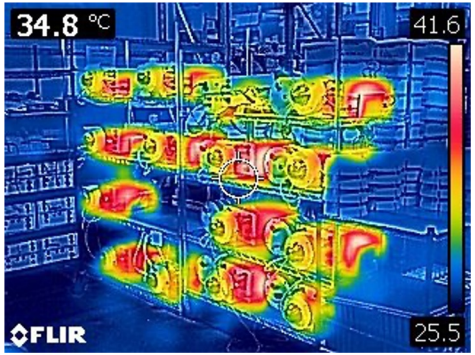

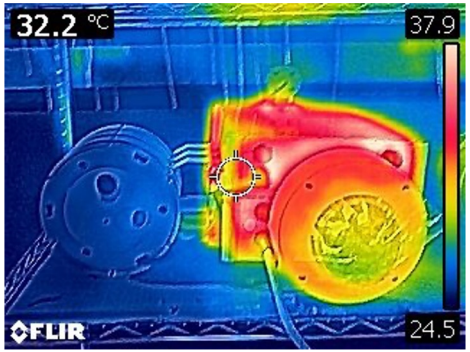

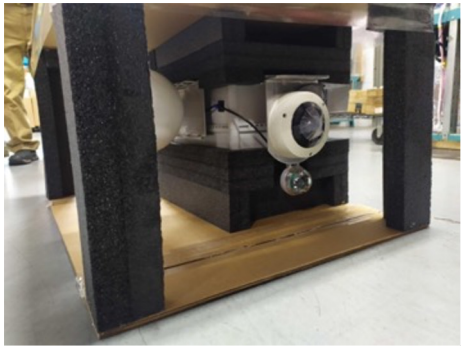

Sage Digital Continuum

SAGE
sagecontinuum.org

Wild Sage Node: Design Qualification

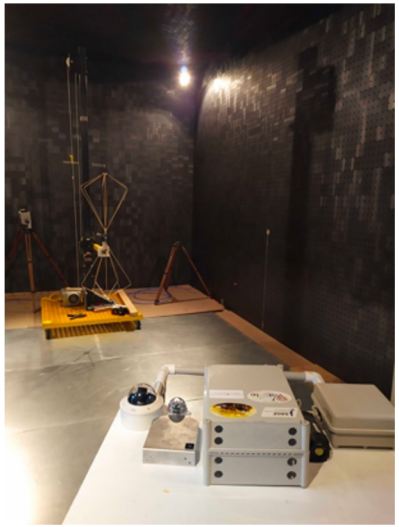

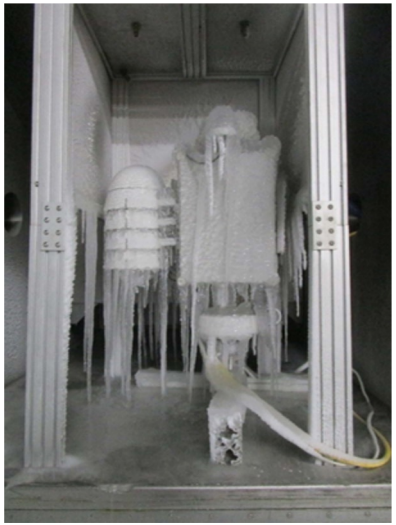

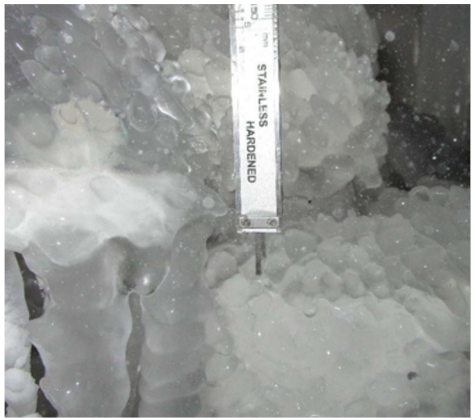

Electrical, EMI, and physical environmental tests to qualify the design.

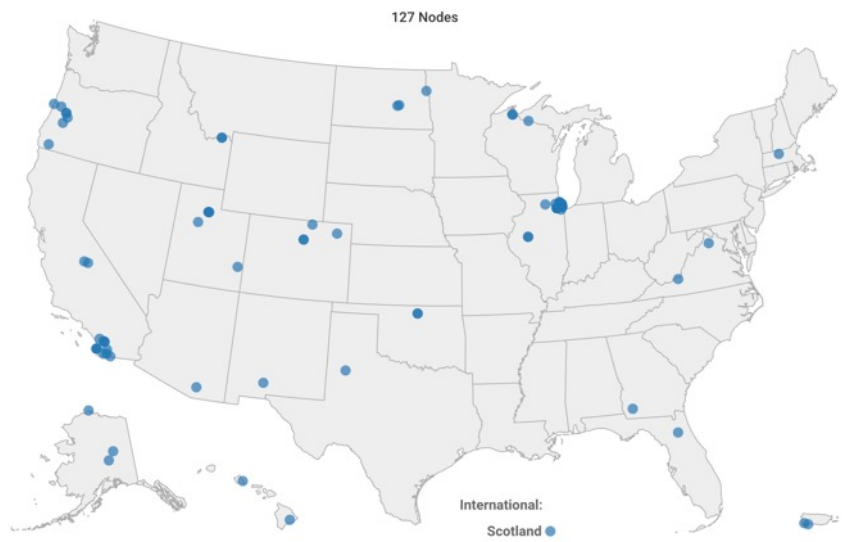

Nodes
127

Active Jobs
93

Recent Apps
36

App Data
2,147,782
Records in the last 24 hours*

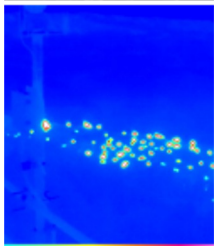

Wild Sage Node Deployment: University of Utah's Taft-Nicholson Center in Montana

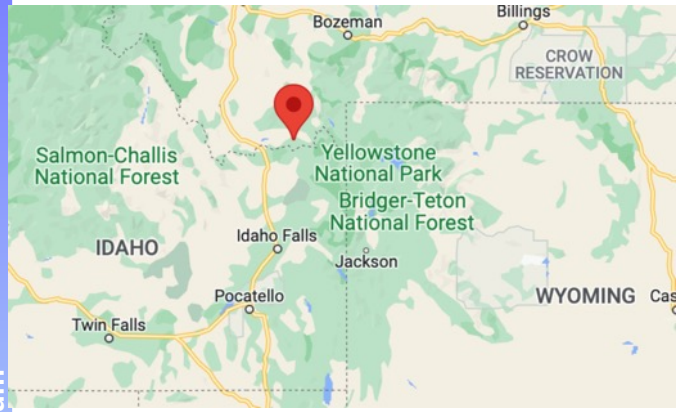

Motivated by the success of a deployment with NEON at a controlled burn on the Konza praire, we are planning deployment of ~5 mobile Sage towers. The first deployment was at a remote site in Montana last fall.

Two phase deployment:
Initial Deployment on campus with line power and Starlink Internet (university network as back up).

Final deployment on a hilltop powered by solar and wind, and Starlink.

Sensors: *Sage node with cameras, microphone, TPH, precipitation, dust and thermal camera.*

44.59784787095959, -111.8116293810192

Architecture

Architecture and Software Stack for AI@Edge Computing & Sensing

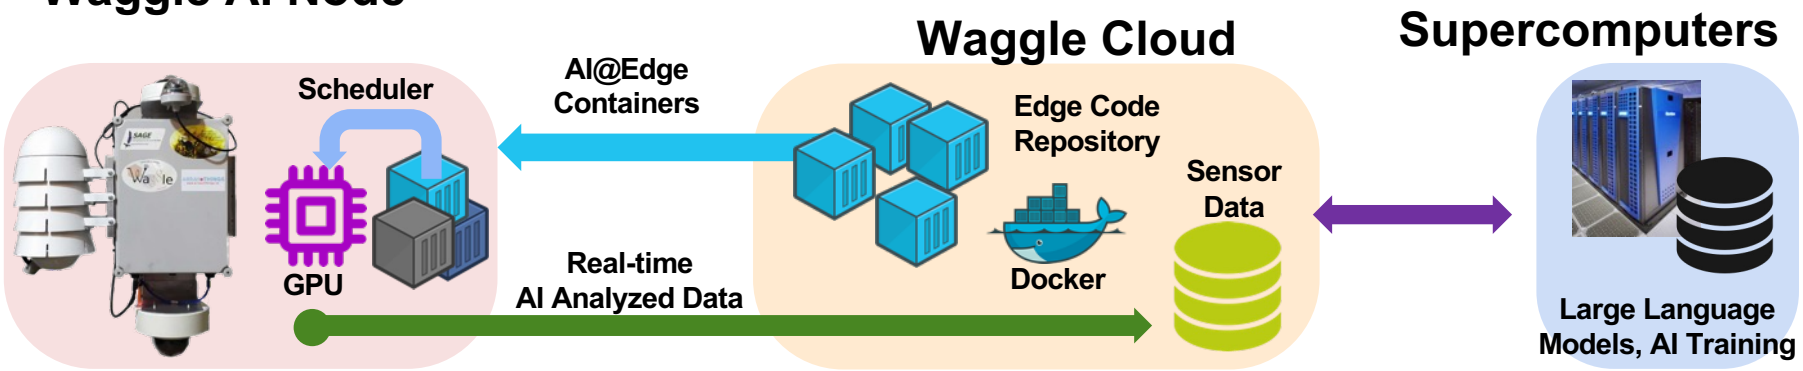

AI toolchain for secure, real-time, distributed AI

Built on industry components...

Waggle AI@Edge Software Stack

- AI toolchain for the edge
- Goal-based scheduler & resource manager
- Fully containerized AI@Edge applications
- Support for multi-tenancy
- Extreme cybersecurity
- Resilient data movement
- Cloud-based data store and management

docker SAGE sagecontinuum.org TensorFlow TensorRT NVIDIA CUDA K3S PYTORCH TensorFlowLite OpenCV ROS

Software Team led by Sean Shahkarami

A National AI@Edge Resource for the Community

The Edge Code Repository

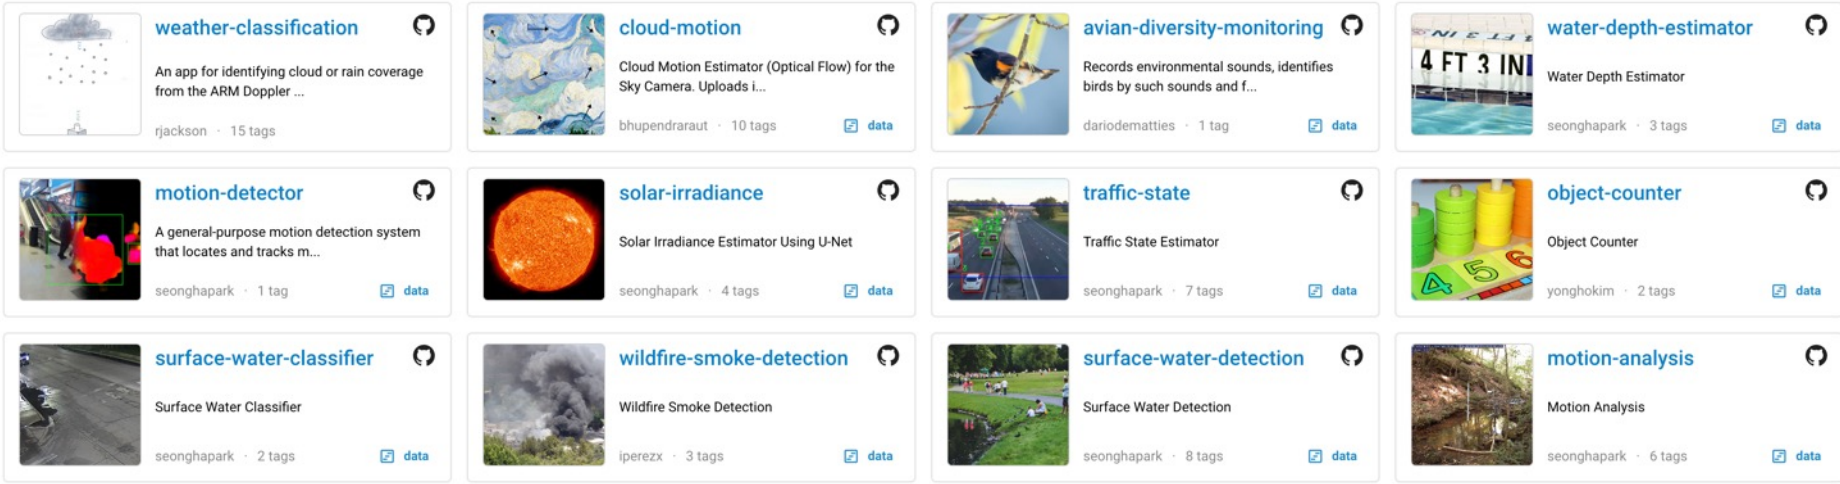

https://portal.sagecontinuum.org/apps/explore

Community contributions will increase software defined sensing capabilities of the nodes – additions to Edge Code Repository benefit SAGE users.

www.sagecontinuum.org

SAGE
sagecontinuum.org

Sage Digital Continuum

Viewing data from Cloud (Waggle Beehive)

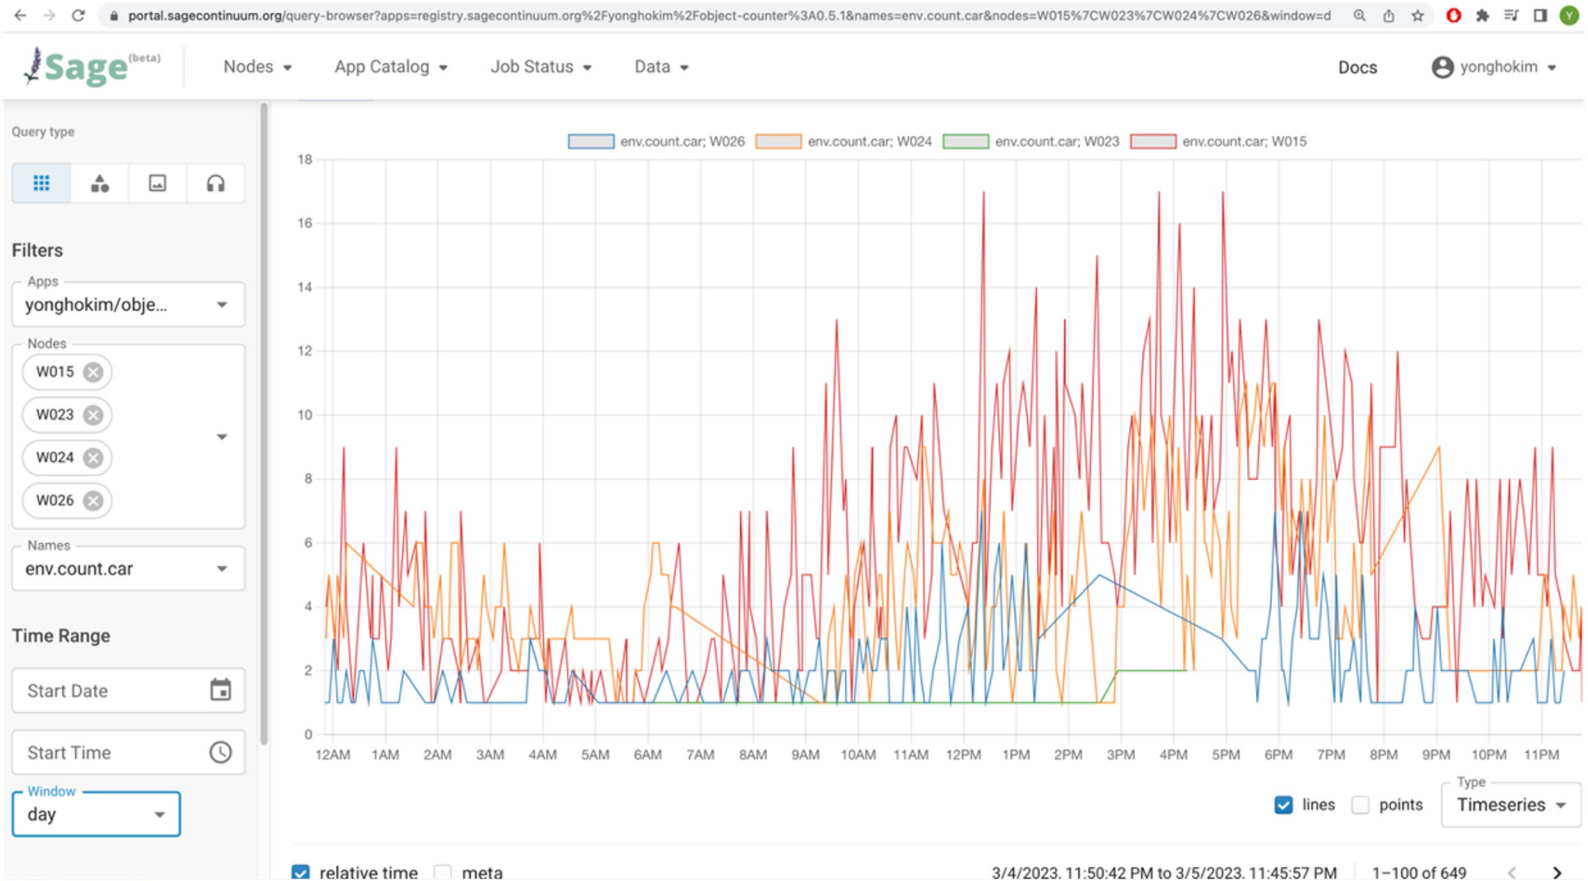

Credit for the Web UI design and implementation: Neal Conrad, Argonne

Demo: Analyzing the local temperatures from nodes

- https://github.com/sagecontinuum/sage-data-client
- https://github.com/sagecontinuum/sage-data-client/blob/main/examples/contrib/geospatial_mapping_example_v2.ipynb

Multi-tenancy with Sage Edge Scheduler

- Jobs include a “Science Goal”
- Science rules are evaluated to schedule applications for different scientific studies

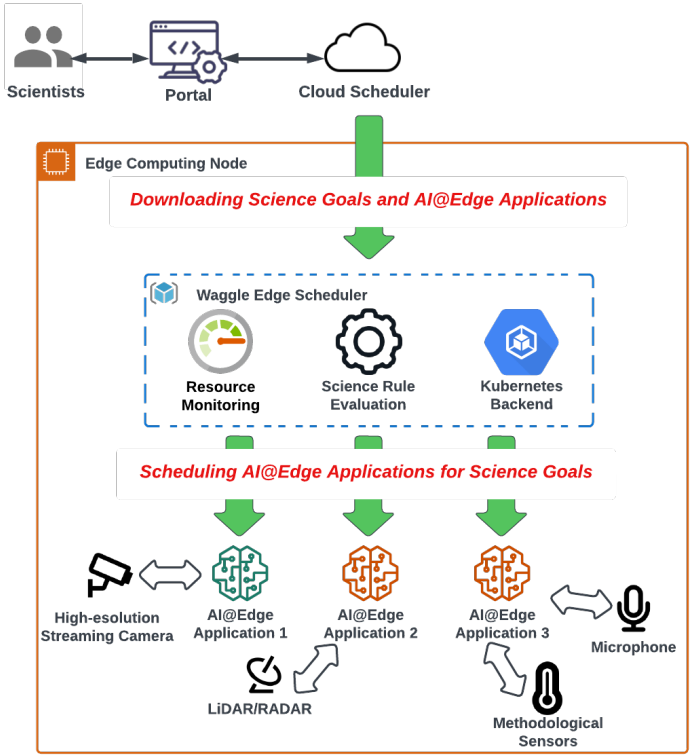

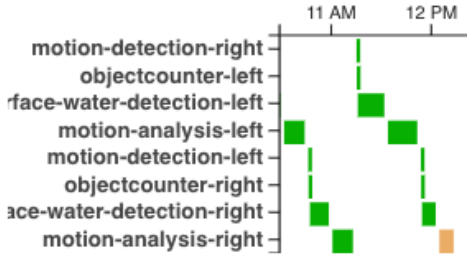

Round-robin scheduling policy

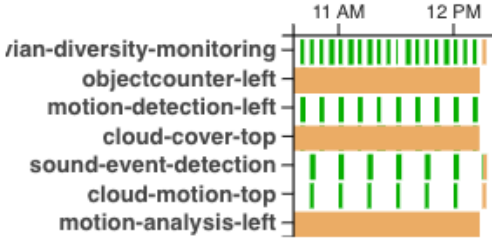

GPU-aware scheduling policy

CS and Systems: Enhancing Resource Monitoring

Edge Monitoring Framework:
System resource monitoring

Application Sidecar:
Application resource
and performance monitoring

We are currently testing the framework and moving it into production

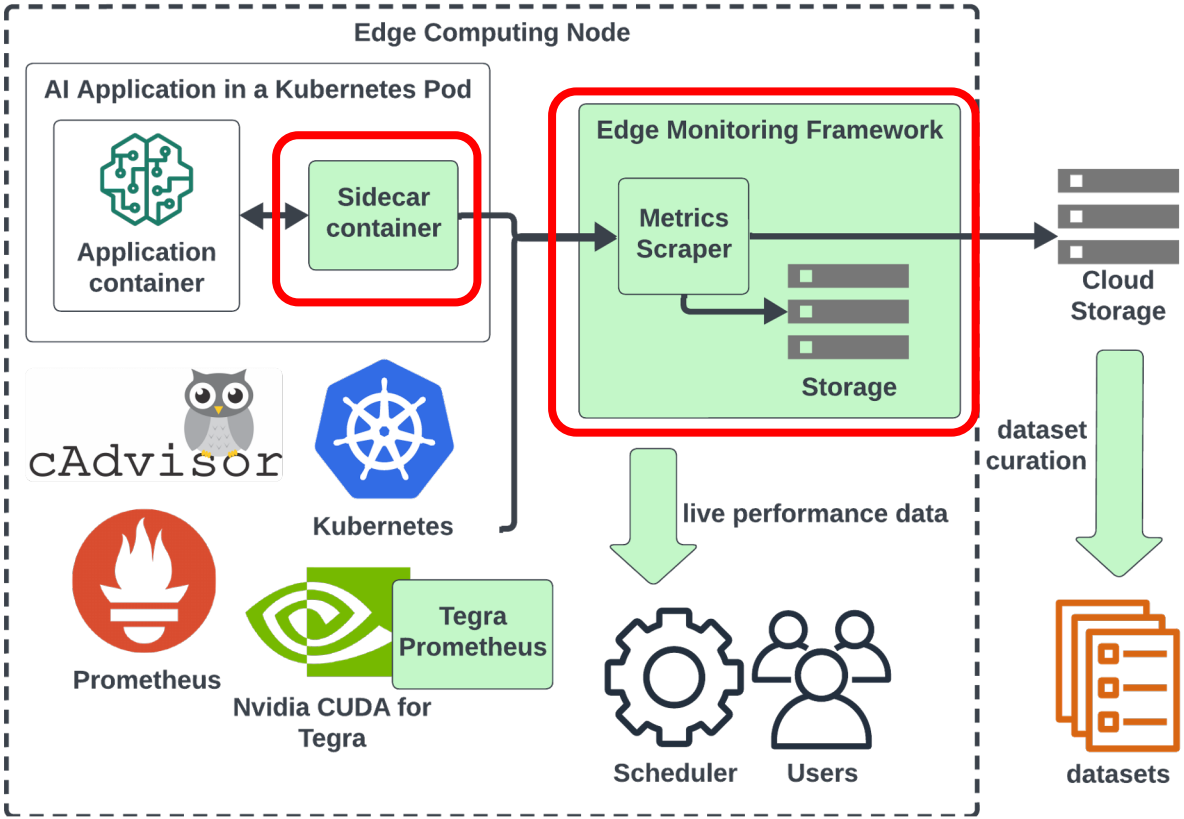

Demo2: Pulling node performance data

- https://github.com/waggle-sensor/edge-scheduler
- https://github.com/waggle-sensor/edge-scheduler/blob/main/scripts/analysis/analyze_node_performance.ipynb

Artificial Intelligence (AI) and Sage: Examples

- https://sagecontinuum.org/science/category/recent-projects

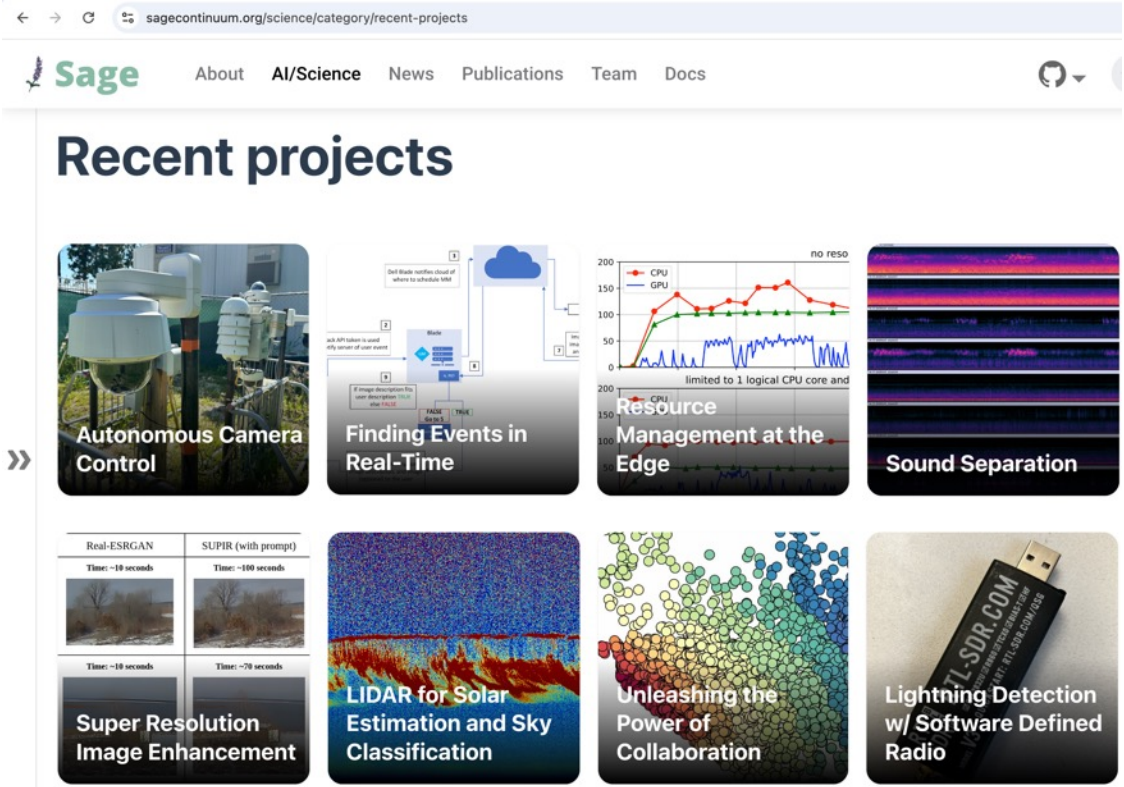

Paintbrush Prairie Bird Detection

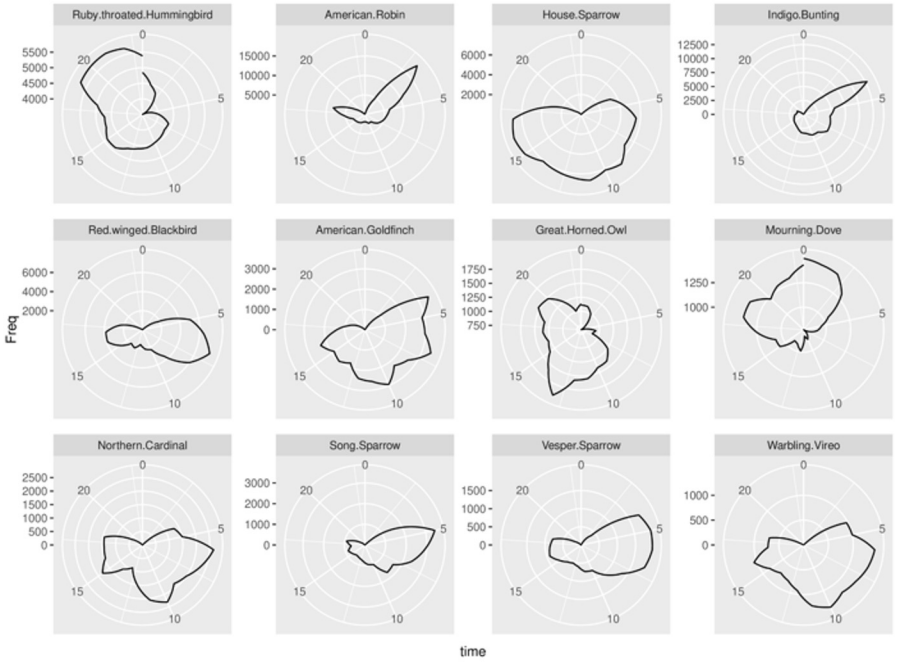

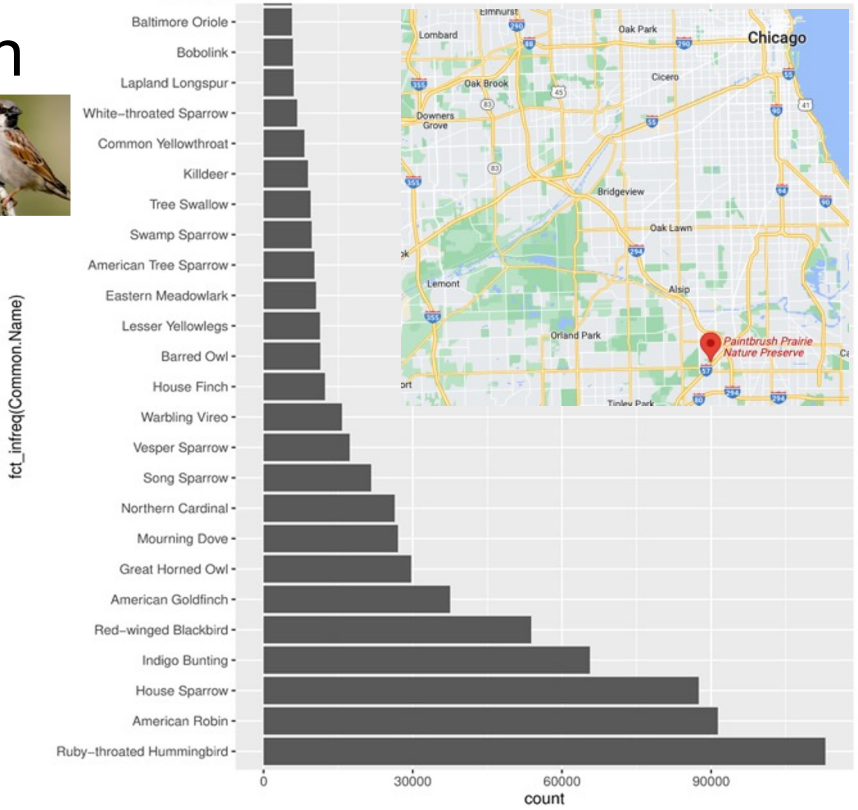

Total calls, for top 12 species, as a function of the hour of the day (UTC-06).

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

BirdNet from Cornell University

Research Credit: Dario Dematties, Bhupendra Raut, Nicola Ferrier

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

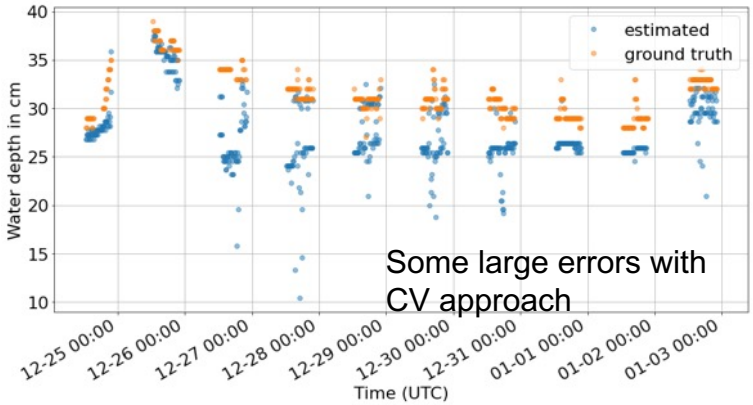

Human annotation using Labelbox

U-Net (segmentation)

ResNet32 based Regression

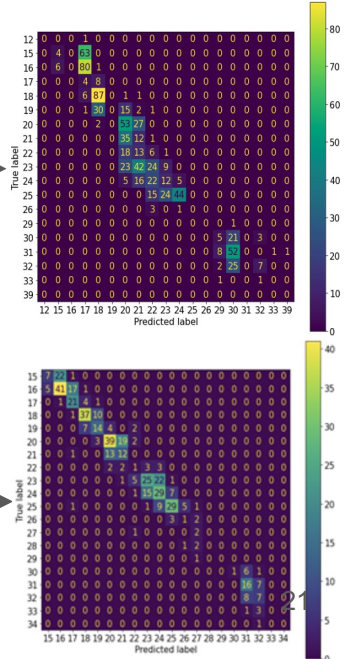

Wildfire Detection and Prediction

Ilkay Altintas, UCSD, Co-PI for SAGE

AI@Edge for wildfire detection (data used in HPC simulations)

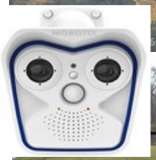

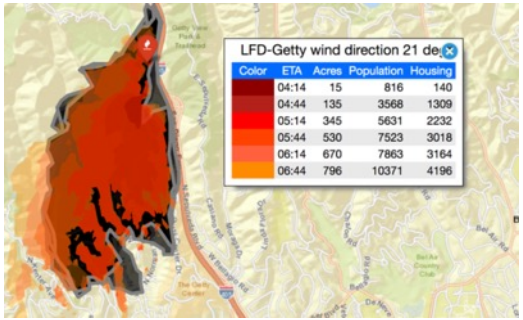

Sage project will move Pan-Tilt-Zoom cameras toward suspected outbreaks, and use infrared cameras to build self-supervised AI training

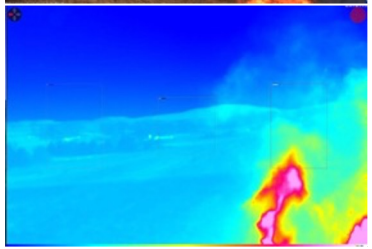

Exploring 2 approaches to improve predictions:

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

SAGE
sagecontinuum.org

Sage Digital Continuum

AI to Detect Wildfires in Real Time

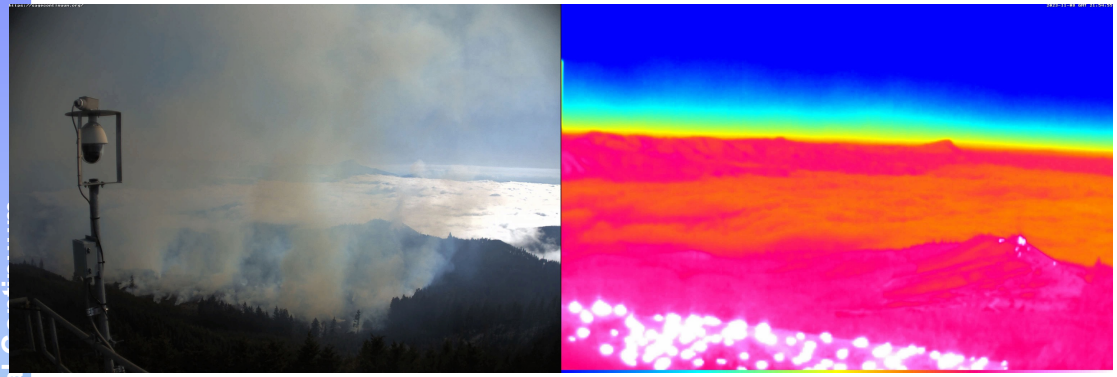

WIFIRE

OHAZ

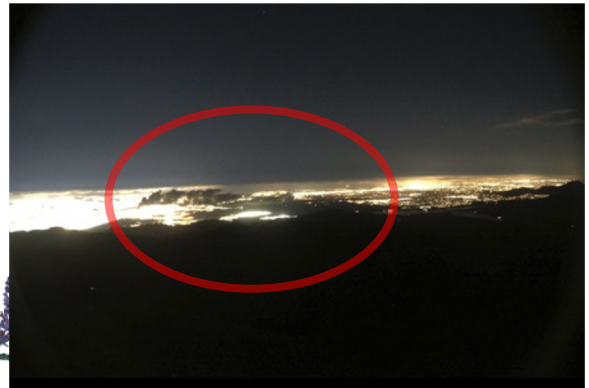

Sage/Waggle Detects Kilauea Volcanic Eruption

- June 3, 2024 at 2am: eruption occurred 5 miles away from Sage/Waggle sensor.
- Sage captures eruption on camera & while testing image against several AI Visual Language Models (Moondream, Yolo, ChatGPT, Claude) for fire detection, our AIs were able to identify the event as a “volcanic eruption”.
- Since then, we have built a workflow to keep the AI running to detect future potential events, and to test AI improvements.
- On Oct 1 we will finally install the sensor in Lahaina/Maui.

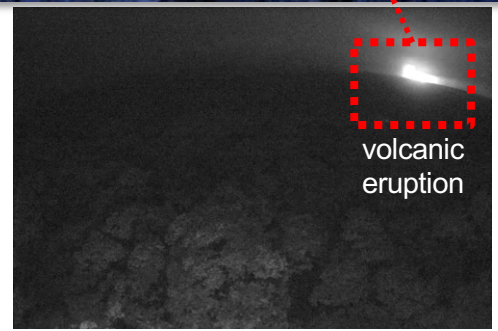

AI fire/lava detection service
AIs being tested:
Moondream
Yolo7
ChatGPT
Claude

SAGE3 Visualization of Sage/Waggle telemetry

Moondream AI is small enough to potentially put inside Waggle node

Research Credit: Dr. Leigh, Hawaii Data Science Institute, RAPID – an Evaluation of an Artificial Intelligence-enhanced Edge Sensor System for Multi-Hazard Monitoring and Detection (ACI 2346568)

SAGE
sagecontinuum.org

Sage Digital Continuum

Undergraduate Research: Pedestrian Detection and Paths

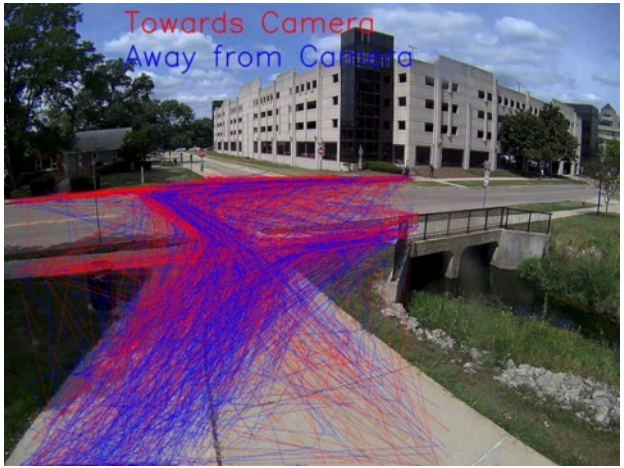

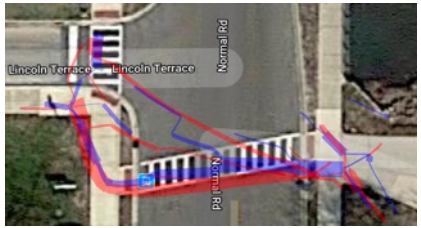

NIU experimental node with wired network connection

- 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

Future direction: Generative AIs + Actuators

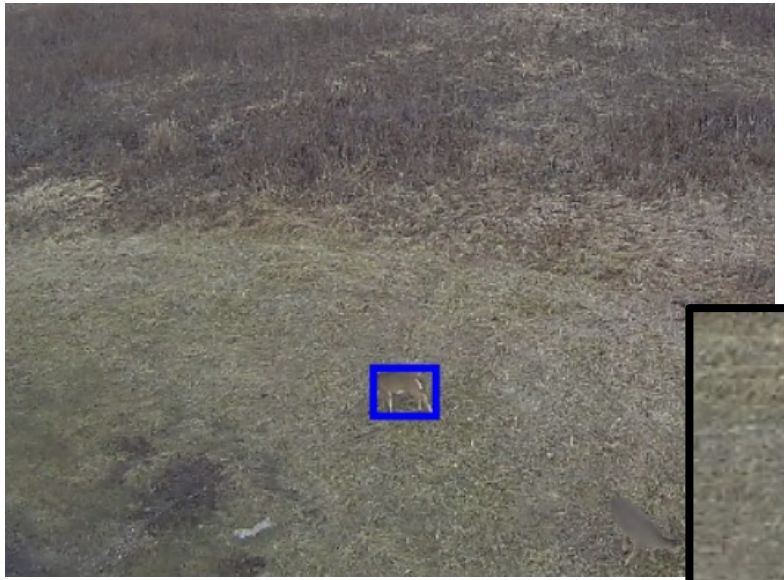

Zooming in and follow!

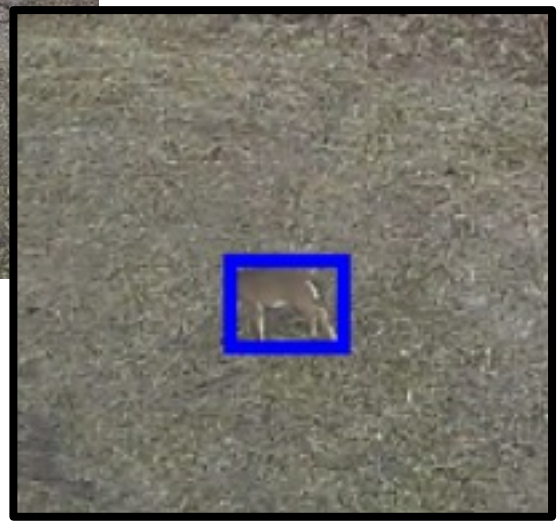

```
  warnings.warn(
{'action': {'noop': [0.0, 0.0, 0.0], 'short': {'left': [-0.1, 0.0, 0.0], 'right': [0.1, 0.0
, 0.0], 'right_down': [0.1, -0.1, 0.0], 'up': [0.0, 0.1, 0.0], 'down': [0.0, -0.1, 0.0], 'zo
, 'right': [1.0, 0.0, 0.0], 'up': [0.0, 1.0, 0.0], 'down': [0.0, -1.0, 0.0], 'zoom_in': [0.0
5.0, 0.0, 0.0], 'up': [0.0, 3.0, 0.0], 'down': [0.0, -3.0, 0.0]}}}
A new version of the following files was downloaded from https://huggingface.co/microsoft/Fl
- configuration_florence2.py
. Make sure to double-check they do not contain any added malicious code. To avoid download
A new version of the following files was downloaded from https://huggingface.co/microsoft/Fl
- modeling_florence2.py
. Make sure to double-check they do not contain any added malicious code. To avoid download
/usr/local/lib/python3.10/dist-packages/torch/cuda/__init__.py:118: UserWarning: CUDA initia
e your GPU driver by downloading and installing a new version from the URL: http://www.nvidi
version that has been compiled with your version of the CUDA driver. (Triggered internally a
  return torch._C._cuda_getDeviceCount() > 0
A new version of the following files was downloaded from https://huggingface.co/microsoft/Fl
- processing_florence2.py
. Make sure to double-check they do not contain any added malicious code. To avoid download
A new version of the following files was downloaded from https://huggingface.co/microsoft/Fl
- configuration_florence2.py
. Make sure to double-check they do not contain any added malicious code. To avoid download
First image in the episode
/imgs/310.0,1.0,3.0_17_2024-09-16_21:38:59.037565.jpg
<<<<<<>>>>>>----------<<<<<<>>>>>>
<<<<<<>>>>>>----------<<<<<<>>>>>>
        following a control panel object
<<<<<<>>>>>>----------<<<<<<>>>>>>
<<<<<<>>>>>>----------<<<<<<>>>>>>
x1: 93.1199951171875
y1: 185.22000122070312
x2: 360.0
y2: 373.1400146484375
image_width: 1920
image_height: 1080
zoom_level: 3.0
current_h_fov:  24.27496484576456
current_v_fov:  13.61218322426082
Move the camera to center the object
Pan: -9.273036601949276
Tilt: -3.2873421486526686
```

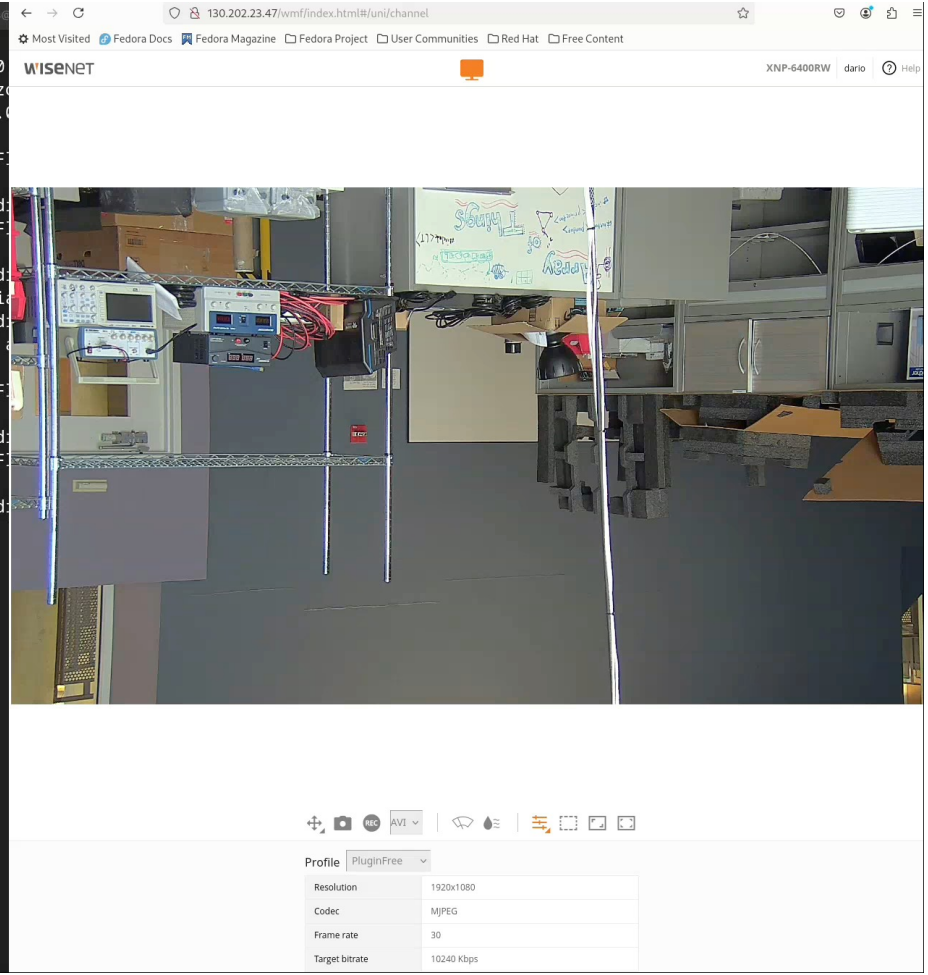

Special Thanks

Students!

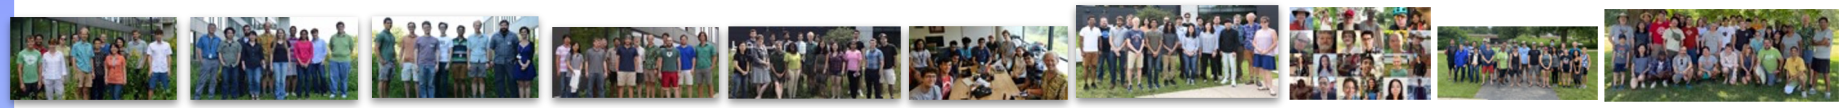

2013 — 2022

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