tinyML® Talks

Enabling Ultra-low Power Machine Learning at the Edge

"MicroCam: A Low-Power and Privacy Preserving Multi-modal Sensor Platform for Occupancy Detection"

Dr. Senem Velipasalar – Professor, Syracuse University

April 30, 2024

TINY ML

www.tinyML.org

Thank you, **tinyML Strategic Partners**,
for committing to take tinyML to the next Level, together

Executive Strategic Partners

Qualcomm
AI research

Advancing AI research to make efficient AI ubiquitous

Power efficiency

Model design, compression, quantization, algorithms, efficient hardware, software tool

Personalization

Continuous learning, contextual, always-on, privacy-preserved, distributed learning

Efficient learning

Robust learning through minimal data, unsupervised learning, on-device learning

A platform to scale AI across the industry

Perception

Object detection, speech recognition, contextual fusion

Reasoning

Scene understanding, language understanding, behavior prediction

Action

Reinforcement learning for decision making

Edge cloud

Cloud

IoT/IIoT

Automotive

Mobile

Accelerate Your Edge Compute

SYNTIANT

Making Edge AI A Reality

www.syntiant.com

Platinum Strategic Partners

TINY ML TALKS webcast

embedUR

AITRIOS

DEPLOY VISION AI AT THE EDGE AT SCALE

SONY

Gold Strategic Partners

Build the Future of tinyML

on arm

TINY ML TALKS webcast

EDGE IMPULSE

The Leading Development Platform for Edge ML

edgeimpulse.com

Driving decarbonization and digitalization. Together.

Infineon serving all target markets as
Leader in Power Systems and IoT

www.infineon.com

innatera

NEUROMORPHIC INTELLIGENCE FOR THE SENSOR-EDGE

www.innatera.com

Renesas is enabling the next generation of AI-powered solutions that will revolutionize every industry sector.

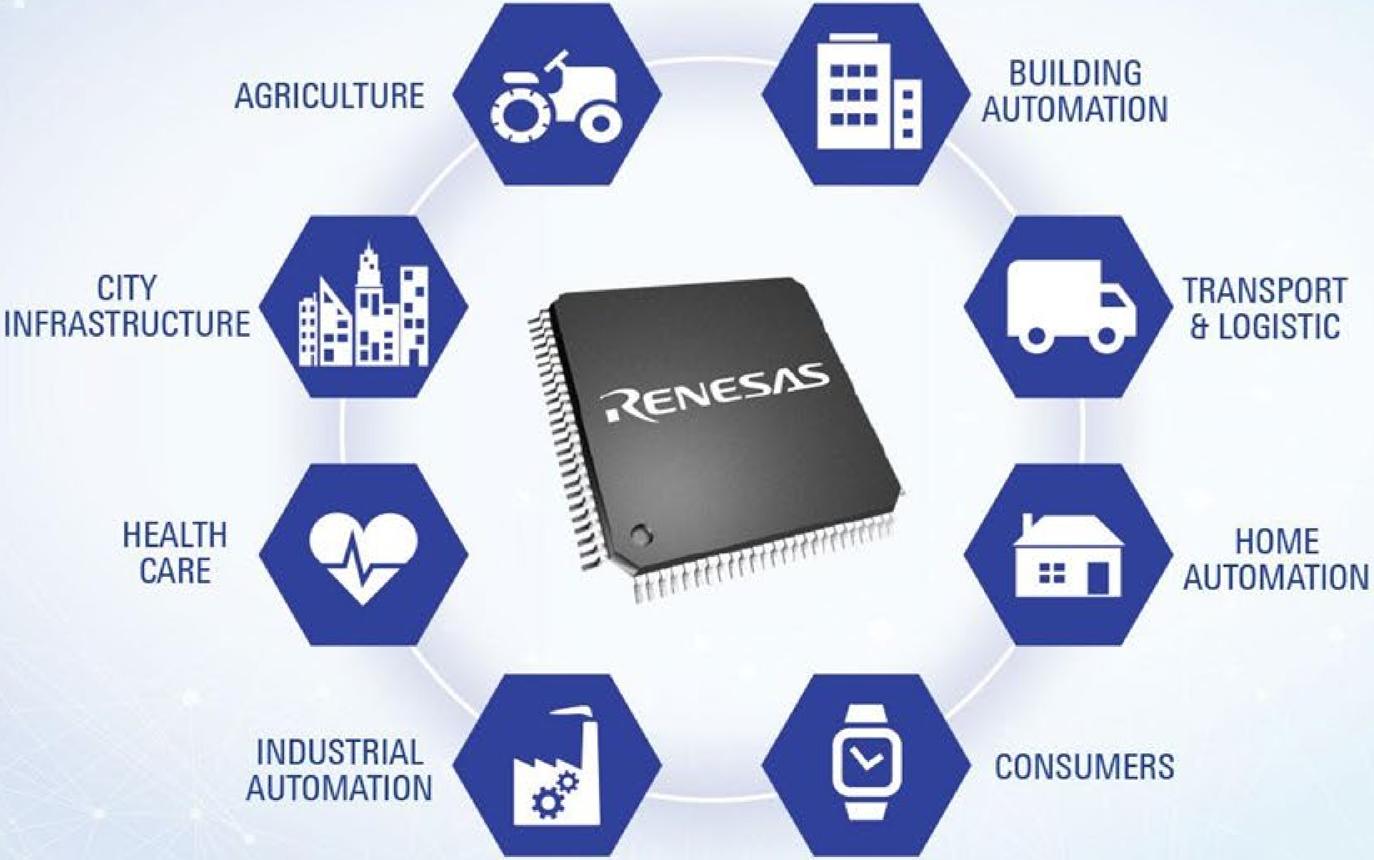

renesas.com

BIG IDEAS FOR EVERY SPACE

STMicroelectronics provides extensive solutions to make tiny Machine Learning easy

life.augmented

www.st.com/ai

Silver Strategic Partners

Join Growing tinyML Communities:

20k members in
50 Groups in 42 Countries

tinyML - Enabling ultra-low Power ML at the Edge

https://www.meetup.com/tinyML-Enabling-ultra-low-Power-ML-at-the-Edge/

4k members
&
15k followers

The tinyML Community

https://www.linkedin.com/groups/13694488/

TINY ML TALKS webcast

YouTube

12.4k subscribers, 671 videos with 453k views

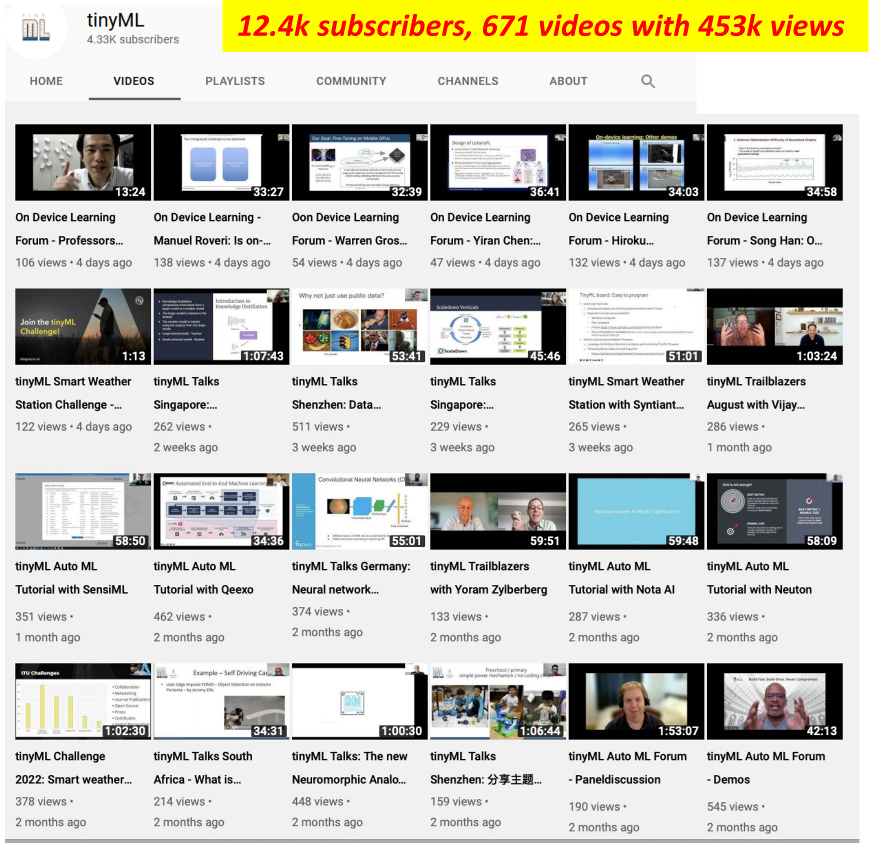

Subscribe to
tinyML YouTube Channel
for updates and notifications
(including this video)

www.youtube.com/tinyML

Reminders

Slides & Videos will be posted tomorrow

tinyml.org/forums

youtube.com/tinym
l

Please use the Q&A window for your questions

Senem Velipasalar Gursoy

Dr. Senem Velipasalar is a Professor in the Department of Electrical Engineering and Computer Science at Syracuse University. She received the Ph.D. and M.A degrees in electrical engineering from Princeton University, and the M.S. degree in electrical sciences and computer engineering from Brown University. The focus of her research has been on machine learning, computer vision, mobile camera applications, wireless embedded smart cameras, multi-camera tracking and surveillance systems. She received a Faculty Early Career Development Award (CAREER) from the National Science Foundation (NSF) in 2011, IEEE Region 1 Technological Innovation (Academic) Award in 2021 and Excellence in Graduate Education Faculty Recognition Award in 2014.

MicroCam: A Low-Power and Privacy Preserving Multi-modal Sensor Platform for Occupancy Detection

Senem Velipasalar
Principal Investigator, Professor
Department of EECS,
Syracuse University
svelipas@syr.edu

Ed Bogucz
T2M Director
Syracuse University
bogucz@syr.edu

Mike Piacentino
Co-Principal Investigator
SRI International
michael.piacentino@sri.com

Outline

- Motivation
- Limitations of Existing Occupancy Sensors
- MicroCam Features and Benefits
- Real-world Challenges
- Evaluation of MicroCam
- Lessons Learned and Conclusion

Motivation

- Heating, ventilation, and air conditioning (HVAC) consumes a significant portion of the energy used in buildings.

- Much of this is wasted energy, used when buildings are either not occupied at all, or occupied well under their maximum design conditions.

- Can we perform occupancy detection reliably and efficiently to autonomously control HVAC systems and save energy?

Existing occupancy sensing solutions

Limitations of existing solutions include one or more of the following:

- They employ sensors or algorithms that are not able to detect stationary occupants;
- They cannot classify the source of the motion (such as a pet);
- Depending on camera resolution and algorithms, they do not allow for embedded or onboard computation, and require external processing;
- Many algorithms developed for camera-based systems are sensitive to lighting changes, and thus prone to missed detections or false alarms;
- Most existing systems depend on adjustment of settings for different scenarios, complicating self-commissioning;
- They cannot provide high enough accuracy;
- They are costly;
- They are not battery-powered, thus limiting ease of use and installation.

MicroCam is a stand-alone, low-cost, multi-modal sensor solution developed to address all these challenges.

Challenges and Goals

- Achieving very high accuracy at a sufficient SWaPC (size, weight, power and cost)
- Not relying on any external computation

- Requirements / Goals:
    - Accuracy > 0.99, number of missed detections ≤ 2/year
    - 30% energy savings
    - Battery life > 3 years on 3 AA batteries
    - Cost: \$0.06/sf
    - Ease of commission, peel and stick solution

MicroCam Overview

- MicroCam consists of low-power, AI-based, IoT platforms
- Each device has multi-modal sensors and can
  - Process motion, audio and video data
  - Send binary occupancy result to the “lead platform”
- Stand-alone solution
  - Battery-powered for easy self-commissioning
  - All processing is performed locally on each platform. No use of external or cloud computing
- Preserves privacy
  - Only the binary occupancy state is shared with the lead platform
  - Preliminary results show promise on images that are modified so that occupants are not discernable

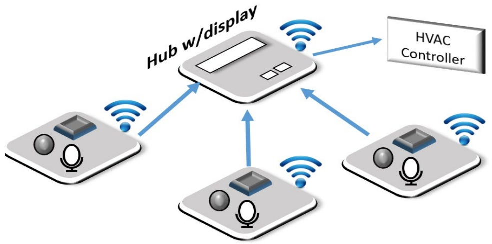

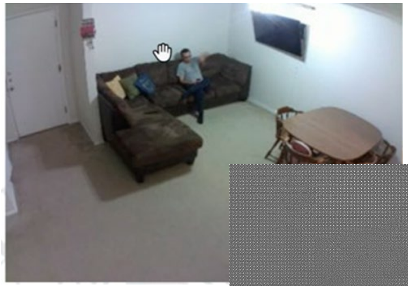

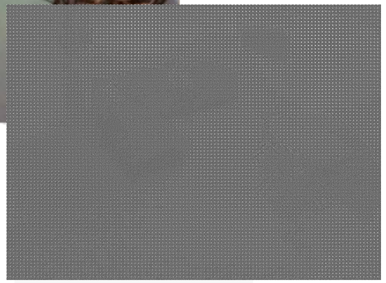

MicroCam Benefits and Performance

- Can operate under daylight, low-light and no-light conditions

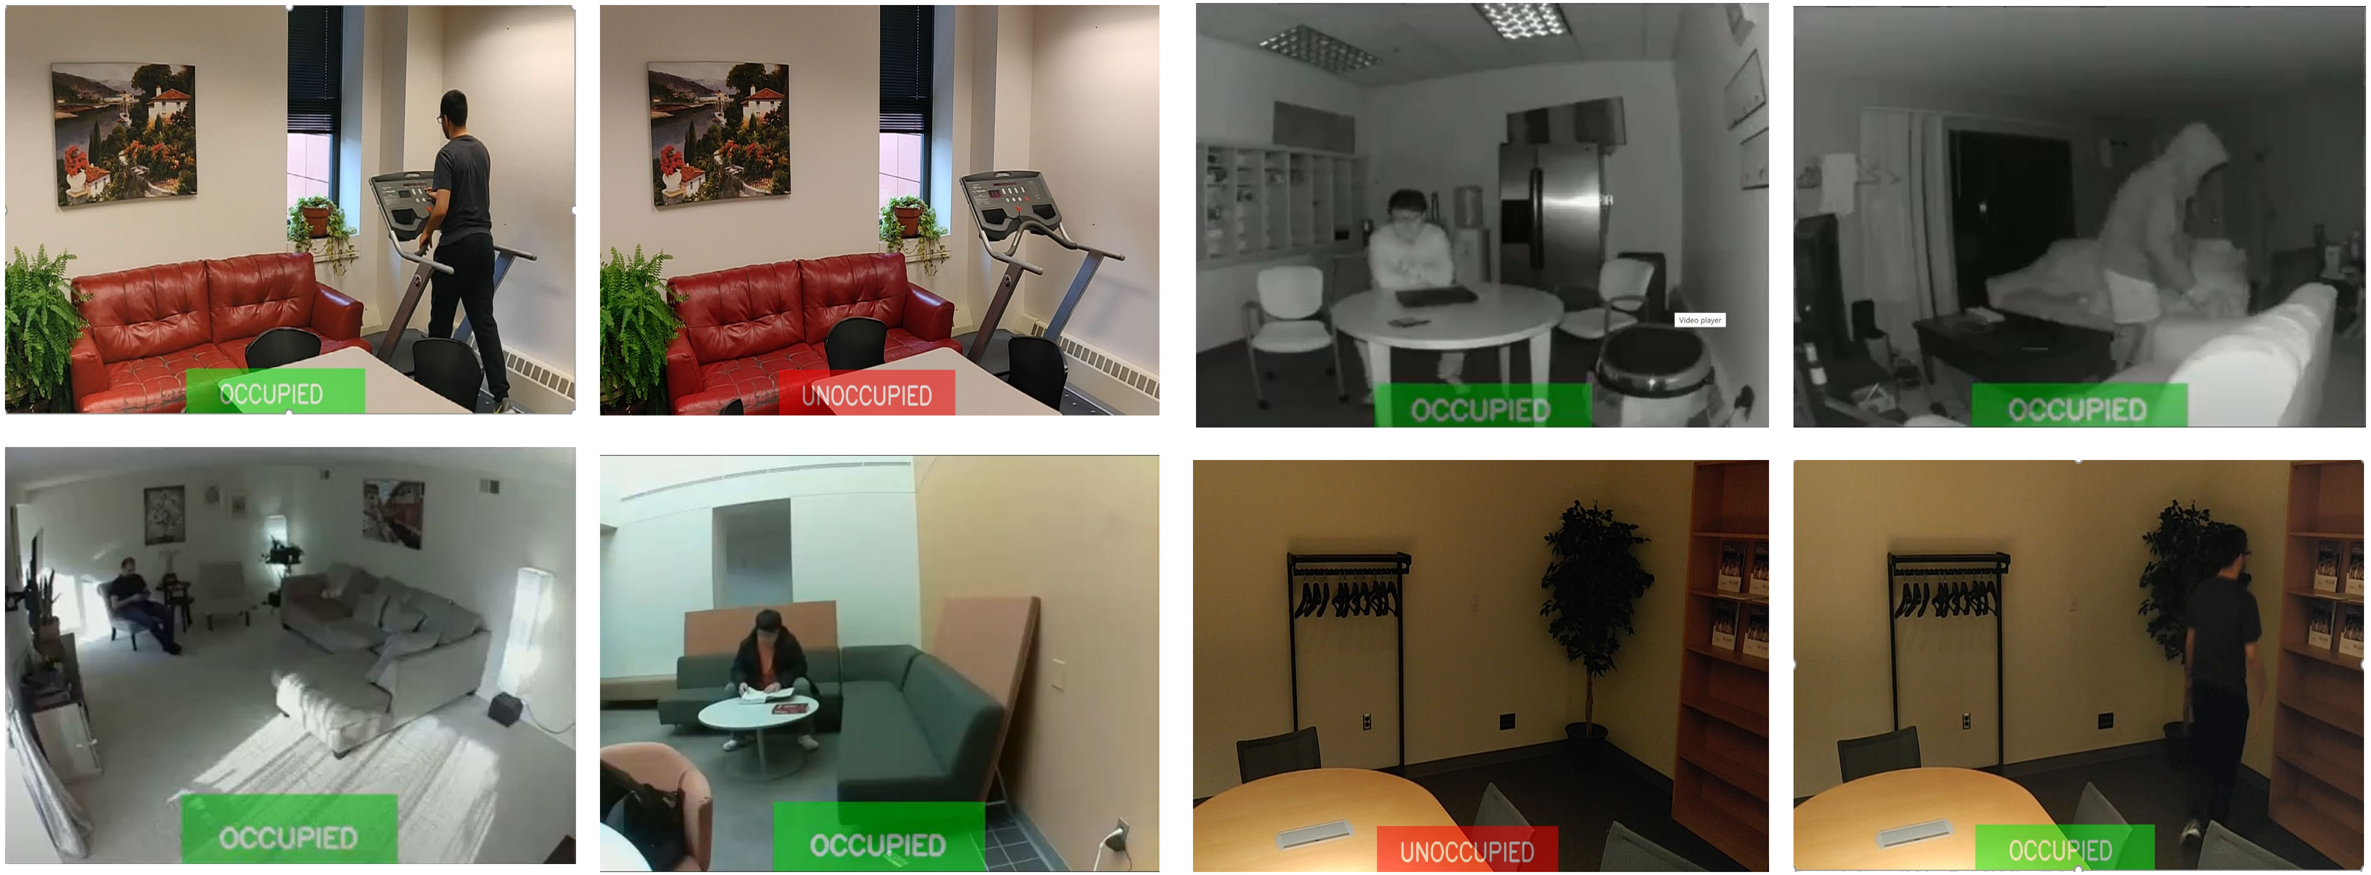

MicroCam Benefits and Performance

- Can detect stationary occupants

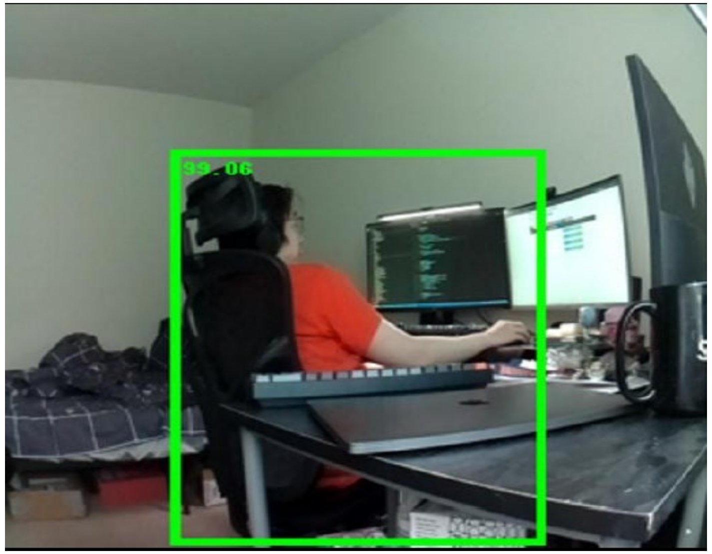

SYRACUSE UNIVERSITY ENGINEERING & COMPUTER SCIENCE
SVLS LABORATORY
SRI International
NEW YORK STATE OF OPPORTUNITY. NYSERDA
arpa·e CHANGING WHAT'S POSSIBLE

MicroCam Benefits and Performance

- Can differentiate people from pets and other sources of motion (e.g., robot vacuums, fans, blowing curtains, etc.)

MicroCam Benefits and Performance

➢ Can preserve privacy

MicroCam Features

- Up to 8 platforms can operate concurrently, one being the lead
- Member platforms are designed for 3-year (target) no-maintenance operation with 3AA batteries. Current battery life > 1 year
- Only the lead platform needs to be wired
- Platforms communicate via Bluetooth
- No access to cloud is needed
- Platforms have around 20-22 feet and 110° x 70° detection range
- The height of each platform should be between 7-8 feet from floor

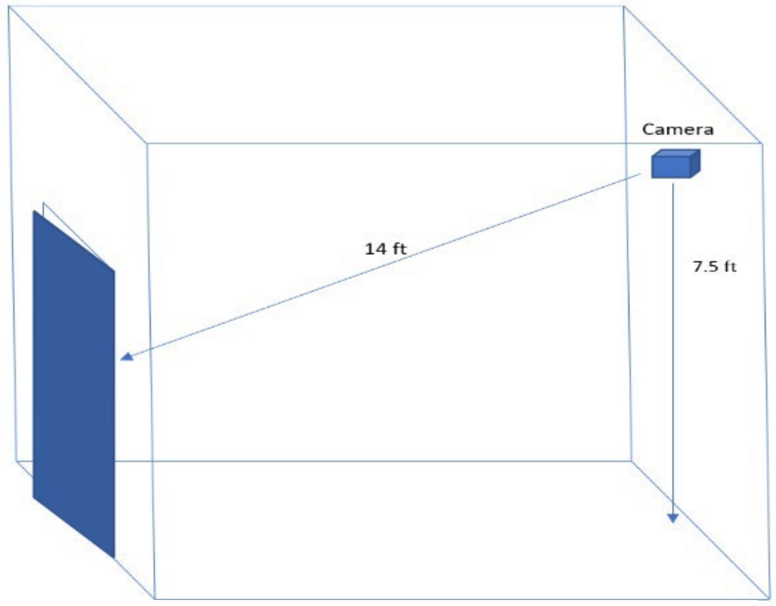

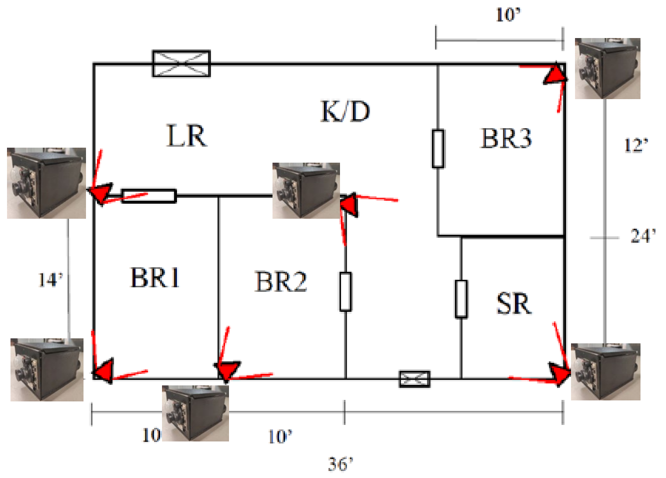

Test layout courtesy of Dr. Kristen Cetin, Yiyi Chu and Debrudra Mitra from Michigan State University

MicroCam Features

- Implemented an Application Programming Interface (API) and a web-based graphical user interface (GUI)
- The API lets users and other devices interact with the MicroCam system. It is the infrastructure that notifies the HVAC controller system about human presence
- It is a highly scalable RESTFul API, which is invoked with simple HTTP GET calls
- The MicroCam system can be configured through API calls when needed

MicroCam Features

- The GUI provides information about the state of the whole system as well as individual platforms. It runs together with the API on the lead platform.

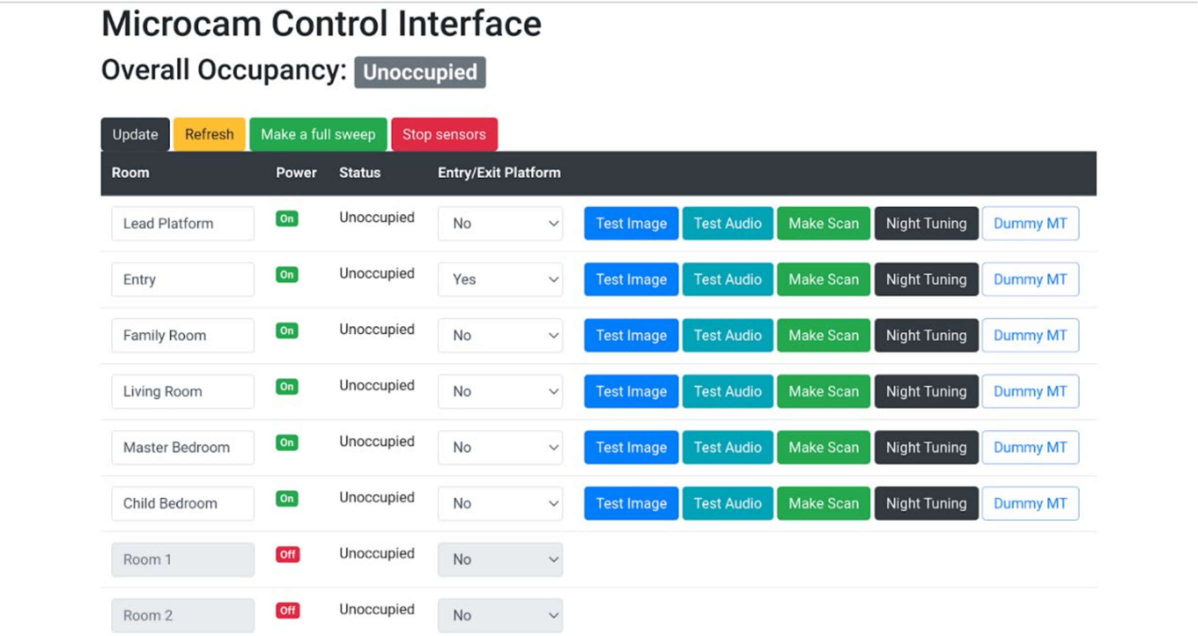

- Platforms start to appear one by one as they are connected to the Lead Platform via Bluetooth.

- Users can check for proper installation and define room aliases and designate them as Entry/Exit platforms

MicroCam Features

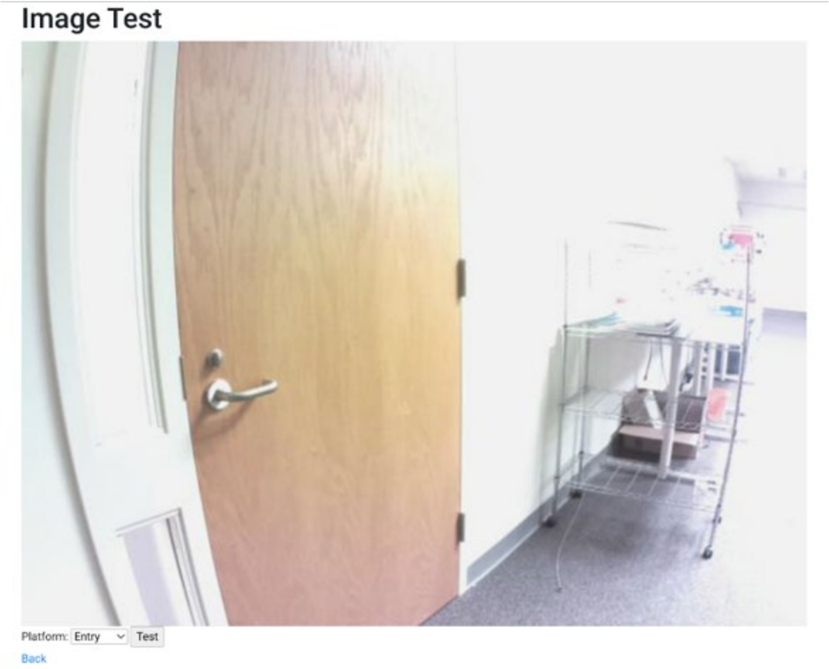

Figure: GUI Camera Testing

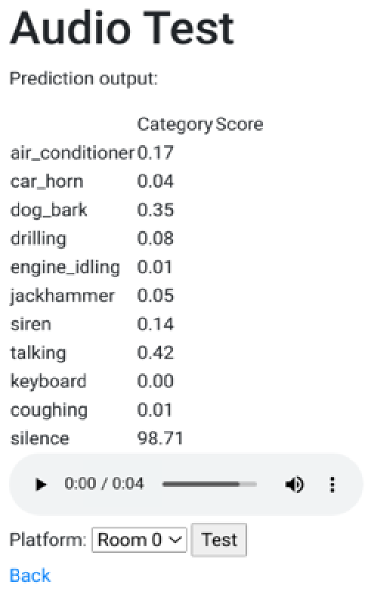

Figure: GUI Microphone Testing

- Users can test the camera sensor on individual platforms
- It is especially useful during commissioning for better placement
- The GUI also provides audio testing for sensor validation

MicroCam Features

- Each MicroCam platform consists of a Raspberry Pi processor, NoIR (Visible/NIR) camera, microphone, motion sensor and light source
- Processor: ARMv8 64-bit SoC @ 1.4GHz
- Memory: 1GB LPDDR2 SDRAM
- Wireless Connectivity: Bluetooth 4.2/BLE
- A Near-IR Light Emitting Diode (LED), operating at 850nm, is used to illuminate the room without disturbing occupants when the lights are turned off

Autonomous Occupancy Detection

- We have modified a lightweight neural network, which is run on the MicroCam platforms without relying on any external processors.
- Parameter size is 22MB
- Around 1.3 GFLOPs
- The accuracy of the original model on the first and second test sets we collected was only 87.56% and 77.7%, respectively.

Values	Original / pre-trained model First Test Set	Original / pre-trained model Second Test Set
TP	1018	1898
TN	1109	1109
FP	19	19
FN	283	844
Precision	98.16%	99.00%
Recall	78.24%	69.21%
Accuracy	87.56%	77.70%

First Test Set
1301 captured images + 1128 Negatives (2429 images)

Second Test Set
2742 captured images + 1128 Negatives (3870 images)

Example Images from the First Test Set

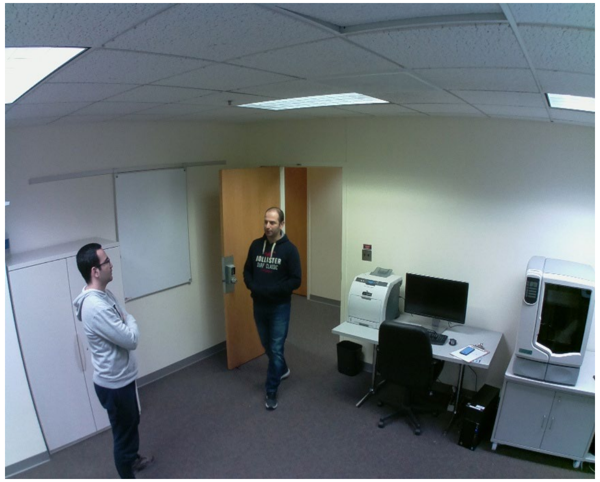

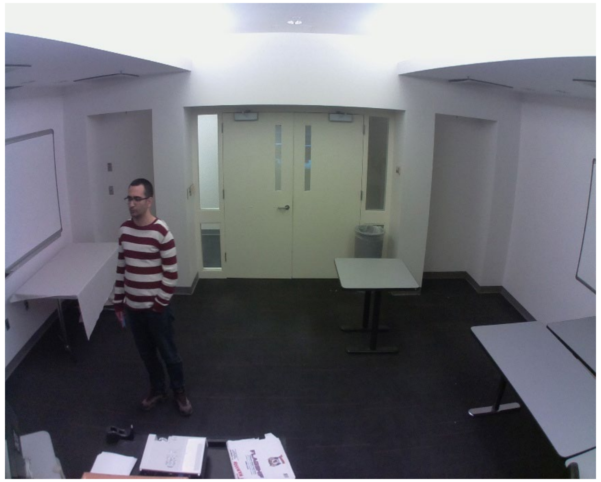

Example Images from the Second Test Set

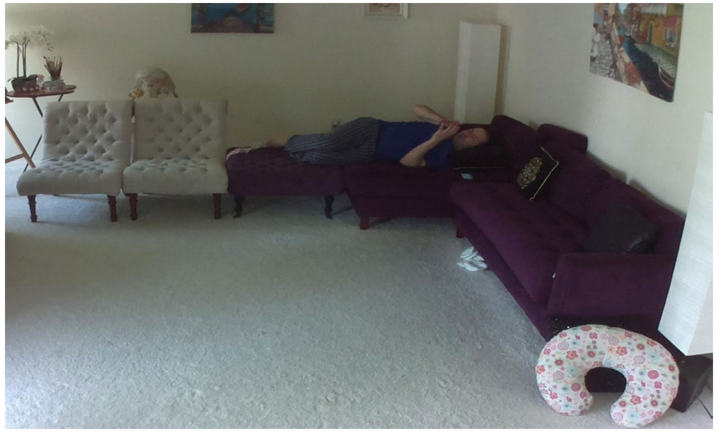

Autonomous Occupancy Detection

- In addition to a neural network model for visible images, we also trained a neural network for low-light and no-light conditions for data from the NoIR sensor
- We fine-tuned the models to increase accuracy
- 24,940 and 23,336 images were used to train on NoIR images and RGB images, respectively
- In addition, we performed data augmentation to cover size, rotation and distance variances

Autonomous Occupancy Detection

Values	Original / pre-trained model First Test Set	Original / pre-trained model Second Test Set	After fine-tuning First Test Test	After fine-tuning Second Test Set
TP	1018	1898	1255	2718
TN	1109	1109	1086	1085
FP	19	19	42	43
FN	283	844	46	24
Precision	98.16%	99.00%	96.76%	98.44%
Recall	78.24%	69.21%	96.46%	99.12%
Accuracy	87.56%	77.70%	**96.37%**	**98.26%**

Syracuse University Engineering & Computer Science · SVLS Laboratory · SRI International · New York State of Opportunity NYSERDA · arpa-e Changing What's Possible

New Data Collection for Edge Cases

- We collected more images for edge cases, such as people lying on sofa, sleeping in bed, wearing hats and sunglasses indoors, while considering different camera perspectives, distances and so on.
- First, we collected 420 images from the Internet and evaluated performance
- Person was missed in 13 out of 420 images. **Observations:**

  a) When face is visible,
    - if the model sees any additional body part (e.g. arm, leg, chest), it detects the person in general
    - if only face is visible, it is usually not enough for detecting a person.

  b) When face is not visible,
    - if the body is in a torso-up position and at least %30 of the body (other than head) is visible, then person is detected in most cases
    - If person lies down and most of the body is not visible than person cannot be detected.

New Data Collection for Edge Cases

- We collected additional 215 images (from internet and captured by ourselves) for extremely hard and relatively rare cases, such as people almost completely covered by blanket and extreme occlusion.

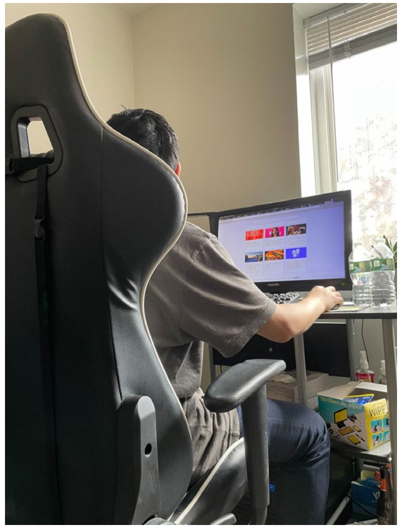

Issues During Integrated Testing

- Motion sensor being unexpectedly triggered due to the WiFi signal
- Adjusting the exposure and ISO settings for the low light and no-light modes, pitch-black images when captured in no-light despite the Near IR LED

Before

After

Issues During Integrated Testing

- LED is a high-power element, requires an external power
- A lab equipment was being used to manually power-up the LED
  - Bulky and not stable solution
  - Prone to human error

Solution: We have designed and built a power module ourselves for all platforms that provides power from the same power line as Pi automatically on-demand

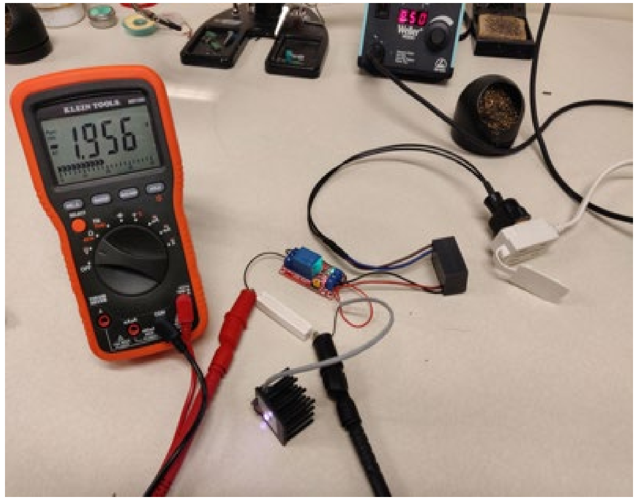

Issues During Integrated Testing

- Flickering Light Sensor - at certain illumination levels, the light sensor can behave unreliably, causing the mechanical NIR filter of the camera and the NIR LED to turn on and off too frequently
  - May damage the filter motor
  - May damage the NIR power module

Solution:

Use Pi to filter the output of the light sensor
Pi will drive mechanical filter of the camera

MicroCam Energy Calculations

System Configuration	Processor (idle power)	Processor (processing added)	Camera (idle)	Camera (processing)	Microphone (base)	Microphone (record)	IR Light (off at idle)	Op Days
1 (Pi 3+, NOIR, Mic, motion)	1300mW	1200mW	582mW	257mW	91mW	85mW	2300mW	367
2 (Pi Zero, PIM*, Mic, motion)	100mW	10mW	582mW	257mW	91mW	85mW	1300mW***	851
3 (Pi Zero, PIM, Mic, motion) Devices off in idle**	100mW	10mW	0	829mW	0	176mW	1300mW	1956

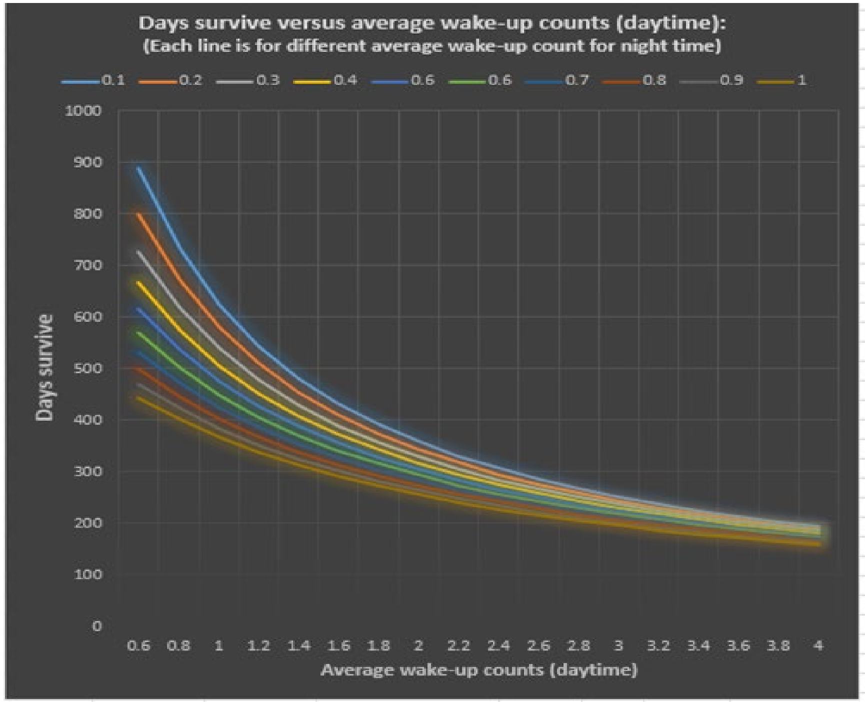

The bulk of the consumption is related to Raspberry Pi

Carefully designed scheduling mechanisms, and sensing strategies to minimize the energy consumption of individual platforms

MicroCam Energy Calculations

System Configuration	Processor (idle power)	Processor (processing added)	Camera (idle)	Camera (processing)	Microphone (base)	Microphone (record)	IR Light (off at idle)	Op Days
1 (Pi 3+, NOIR, Mic, motion)	1300mW	1200mW	582mW	257mW	91mW	85mW	2300mW	367
2 (Pi Zero, PIM*, Mic, motion)	100mW	10mW	582mW	257mW	91mW	85mW	1300mW***	851
3 (Pi Zero, PIM, Mic, motion) Devices off in idle**	100mW	10mW	0	829mW	0	176mW	1300mW	1956

Average Pi wake times	2	3	4	5	6	7	8	9	10
Battery life (days) with #3	1956	1562	1361	1205	1043	949	846	783	711
Power Consumption over 3 years (Wh)	5.88	7.36	8.45	9.54	11.02	12.12	13.59	14.68	16.17

10 wake-ups result in approximately 16 Wh over 3 years

SYRACUSE UNIVERSITY ENGINEERING & COMPUTER SCIENCE

Real-world Testing of MicroCam

Real-world testing was performed covering various difficult scenarios (aka "edge cases"), including people lying down, people in chairs, scenes with cats, very low-light and no-light conditions.

Integrated long-term laboratory tests

- Daytime tests lasted over 111 hours
- The overall accuracy of daytime tests is 100% with no false positives
- No-light tests lasted over 10 hours
- The overall accuracy of the no-light tests is >99%

Tests in different apartments

- Additional real-life tests in three different NY state apartments for a total duration of around 412 hours
- Average accuracy is 99.37%

3rd Party Testing

- Performed under the leadership of Dr. Kristen Cetin at Michigan State University

The accuracy is defined as the ratio of number of correct decisions to the number of all samples

$$\frac{TP+TN}{TP+TN+FP+FN}$$

Tests in Different Apartments

Apartment I

Camera 1 view

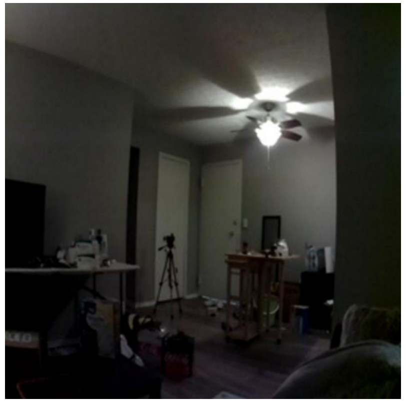

Camera 2 View

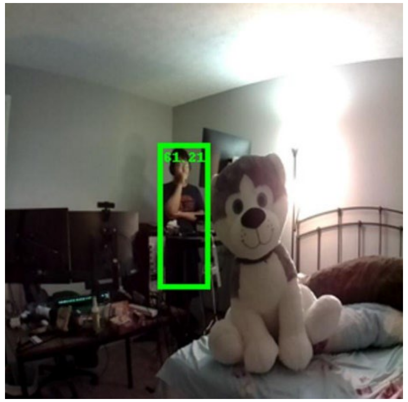

Camera 3 View

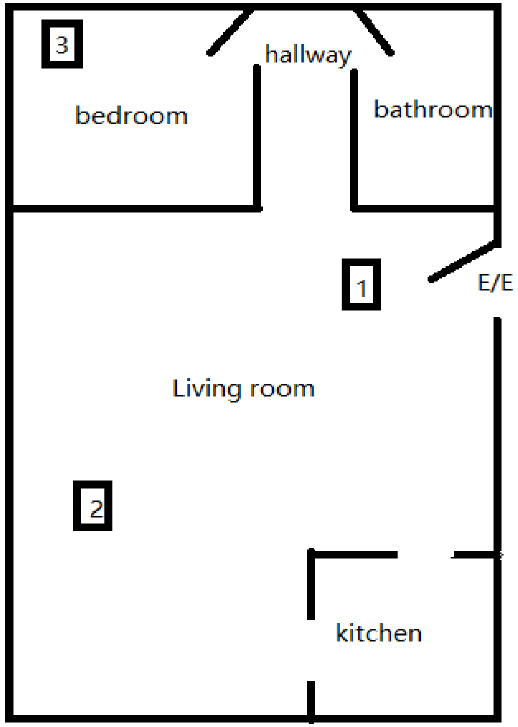

Total time tested: 2810 mins

TN : 50 mins
TP : 2760 mins
FP : 0
FN : 0

Precision = 100%
Recall = 100%
Accuracy = **100%**

Tests in Different Apartments

Apartment 2

View of the E/E platform

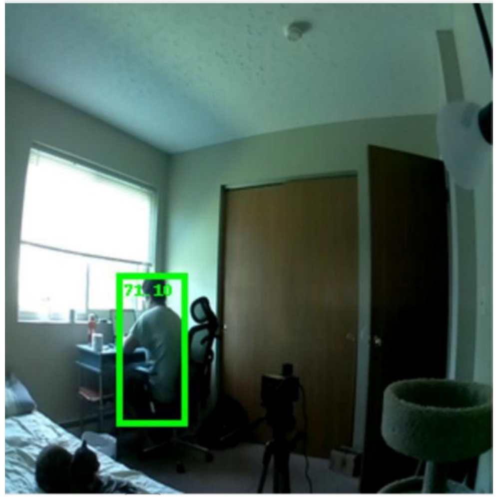

View of Camera 1

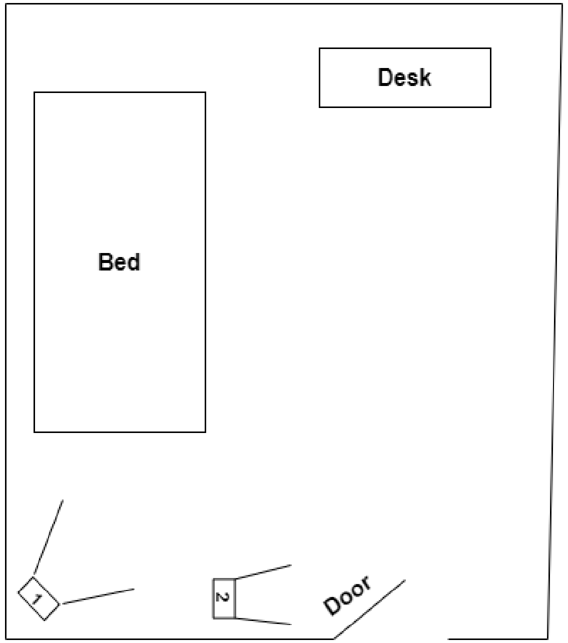

Total time tested: 2850 mins

TN: 148 mins

TP: 2692 mins

FP: 10 mins

FN: 0

Precision = 99.62%

Recall = 100%

Accuracy = **99.64%**

Tests in Different Apartments

Apartment 3

Camera 1 view

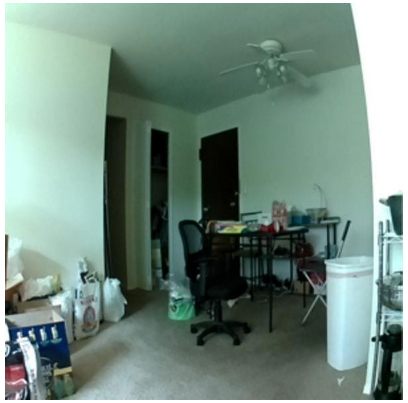

Camera 2 View

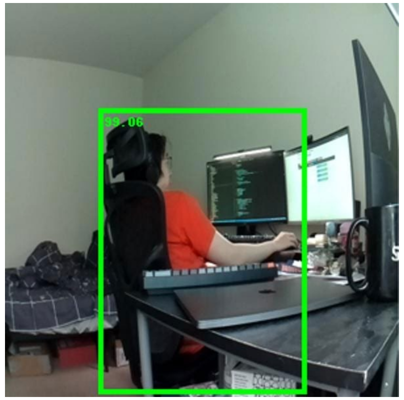

Camera 3 View

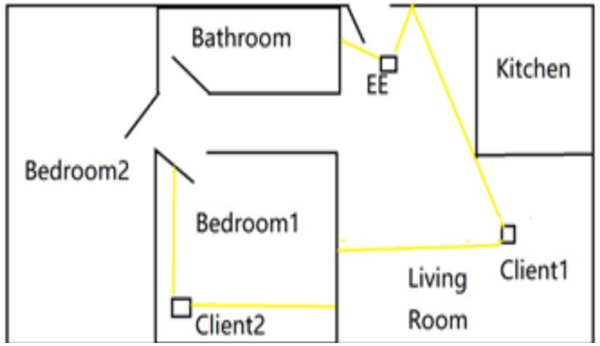

1. E/E platform
2. Client platform in the living room.
3. Client platform in the bedroom.

Total time tested: 88 hours

Precision = 100%

Recall = 98.95%

Accuracy = **99.07%**

A false negative happened when Subject 2 left the apartment while subject 1 was sleeping. The sleeping person was not detected until they changed their sleeping position.

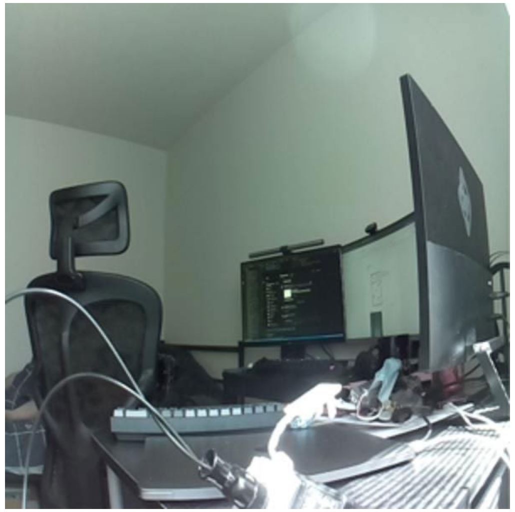

Third Party Testing

- Since a great emphasis is put on the battery life, we developed a duty cycle/decision mechanisms that try to minimize the number of full scans.
- Analysis of the 3rd part test results mainly provides good news, since
  - none of the missed detections was due to camera sensors or detection algorithms, but due to scheduling mechanisms
  - These issues can be solved with more frequent scans or by updating the decision mechanisms

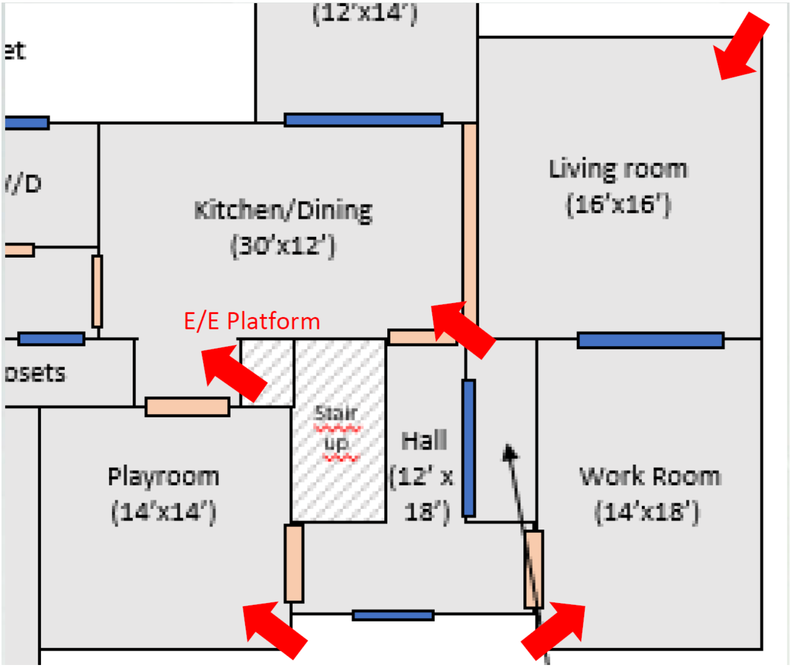

Courtesy of Dr. Kristen Cetin, Yiyi Chu and Debrudra Mitra from Michigan State University

MicroCam Demo

Energy Saving Analysis and Effect of False Positives

- Dr. Tarek Rakha and his team at Georgia Tech performed simulations with preliminary occupancy models, generating as high as 27 false positives/week while saving up to 30% energy weekly.

10 full scans per day	6 full scans per day
Number of Weekly Scans: 70 Expected Duration of Error = 78 minutes Hourly Energy Loss per Error = 0.05% - 0.4%	Number of Weekly Scans: 42 Expected Duration of Error = 100 minutes Hourly Energy Loss per Error = 0.05% - 0.4%
Range of allowable weekly number of false positives for 30% benchmark energy savings = 3 - 27 false positives	Range of allowable weekly number of false positives for 30% benchmark energy savings = 2 - 21 false positives

Product Design

- Worked on designing and 3D-printing an "anticipated product", which has more attractive and futuristic visuals
- Current size: 2.5 x 4 x 3.25 inches
  Target size: 2 x 2 x 2 inches
- Peel-and-stick or magnetic attachment

Special thanks to

Current and former Ph.D. students

- Fatih Altay
- Weiheng Chai
- Jiajing Chen
- Burak Kakillioglu
- Jiyang Wang

SRI and Latent AI Colleagues

- Sek Chai
- Mike Piacentino
- David Zhang

Sponsors

- ARPA-E
- NYSERDA

SU, Georgia Tech and Michigan State Colleagues

- Ed Bogucz
- Chetna Chianese
- Don Carr
- Tarek Rakha
- Kristen Cetin

Copyright Notice

www.tinyml.org

Copyright Notice