

Enabling Ecology Research Using Hybrid Cloud Technology


Chandra Krintz & Rich Wolski

Dept. of Computer Science,

UC Santa Barbara

ckrintz@cs.ucsb.edu

Aug. 2017



Wildlife Monitoring

- Evaluating diversity, species, and habitat health
- Extracting patterns in activity and behavior of animals
- Monitoring change in land use
- Avoiding dangerous human/animal encounters & overlap
- Educational experiences
- Citizen science



Digital Photography

- Alternative to labor-intensive observation & tracking
- Cost effective and scalable
- Safe and non-invasive
- Increasingly autonomous
 - Motion triggered
 - IoT devices




Image Classification

- Enormous number of images (size and count)
- Classification and identification once done by humans
 - Now automated by data analysis tools
 - Tensorflow: Open source machine intelligence library
 - Other similar packages:
 - Caffe, Torch, Theano, CNTK...



Image Classification: Training

- Process: train a model then use it to classify images
- Training set (large, *manually* labeled, ground truth)
 - Requires lots of compute power and GPUs (thanks *public cloud*!)
 - Can take days to weeks; but can be incremental



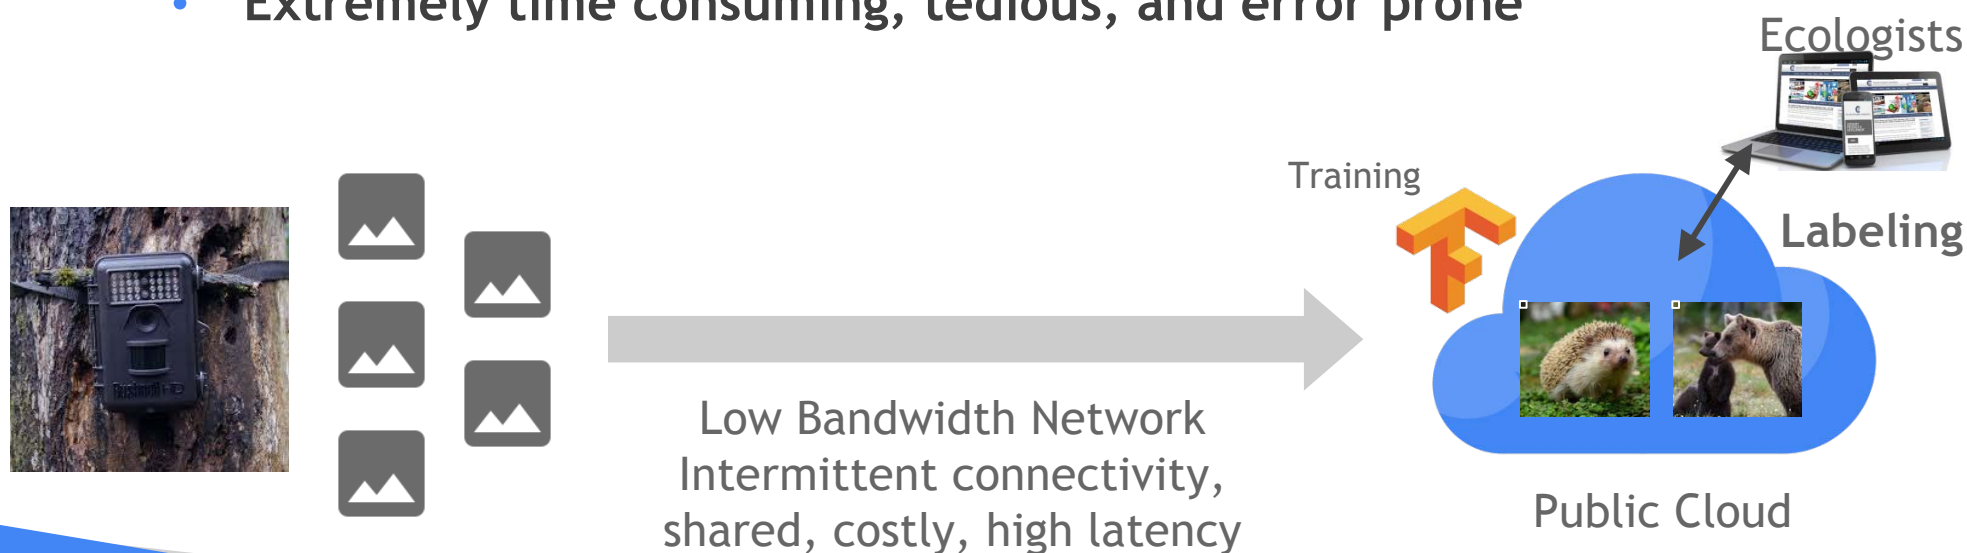
Automatic Image Classification

- Process: train a model then use it to classify images
- Training set (large, *manually* labeled, ground truth)
 - Requires lots of compute power and GPUs (thanks *public cloud!*)
 - Can take days to weeks; but can be incremental
- Classification is **fast and uses many fewer resources**




Challenges With IoT+Cloud for Wildlife Monitoring

- Enormous numbers of images (size and count)
 - Cameras limited storage & processing power
- Automatic classification requires labeling by humans
 - Images must be moved to where they are processed
 - Extremely time consuming, tedious, and error prone

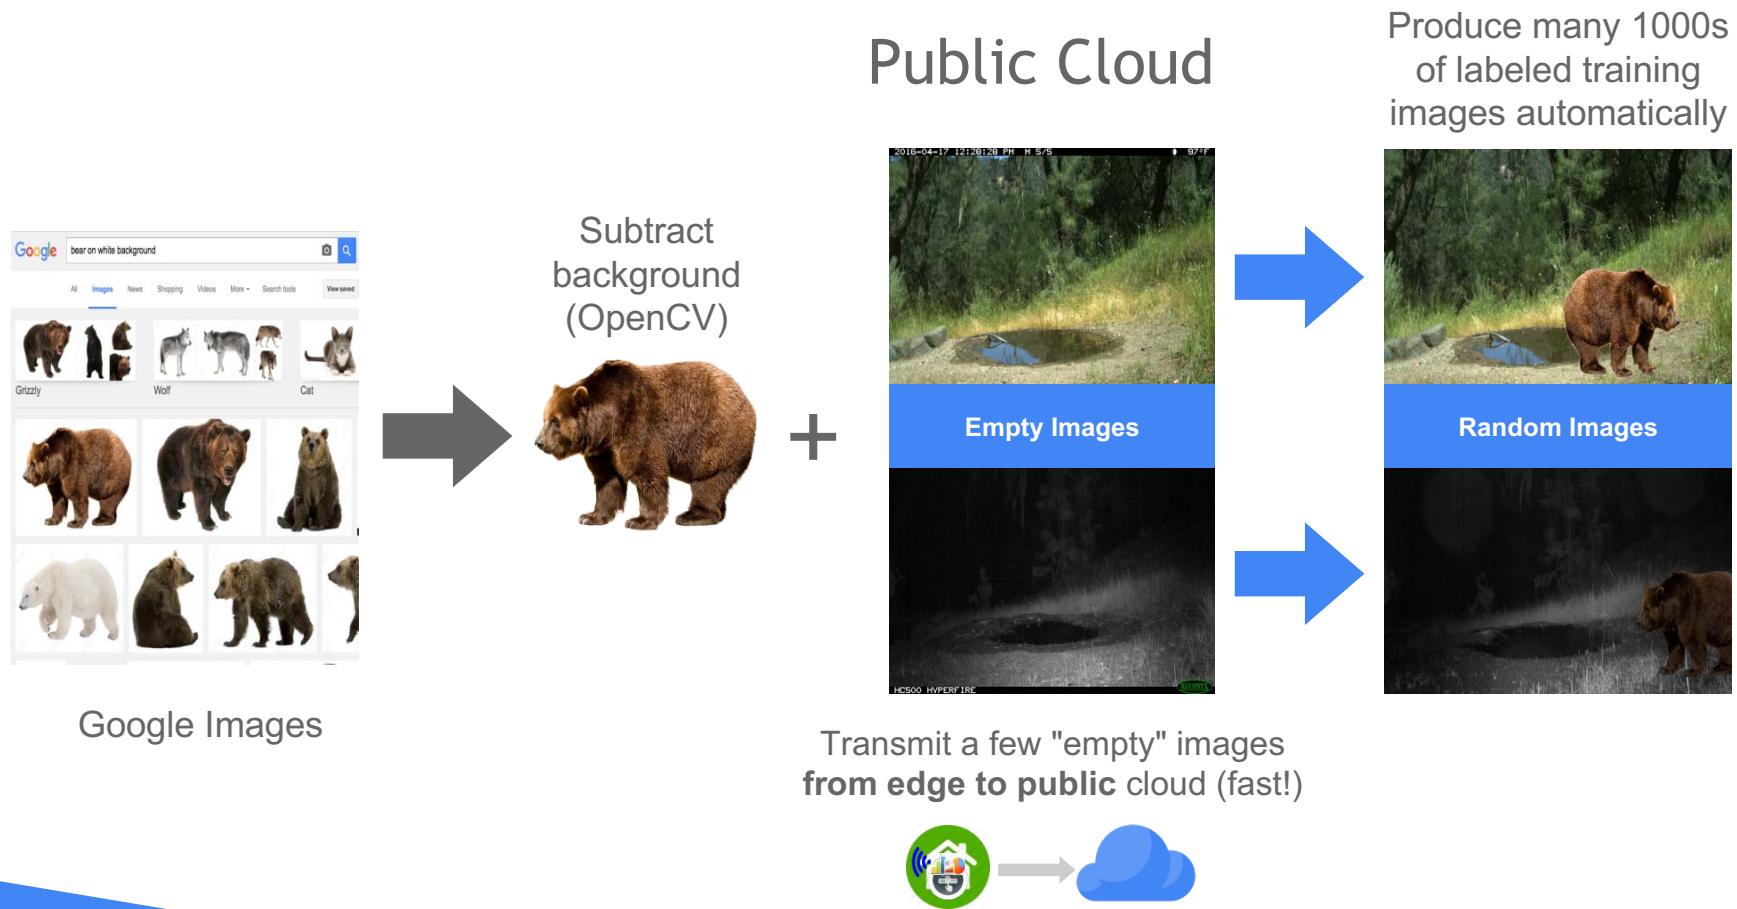


Where's The Bear (WTB)

- Multi-tier IoT system
 - Move the *code to the data (images)* not vice versa
 - Via **Edge Clouds**: robust, self-managing appliances, on-site
 - Low latency, high-bandwidth direct connectivity to cameras
 - Local image classification
 - The need for training (large number of images needed) defeats the purpose/benefit!

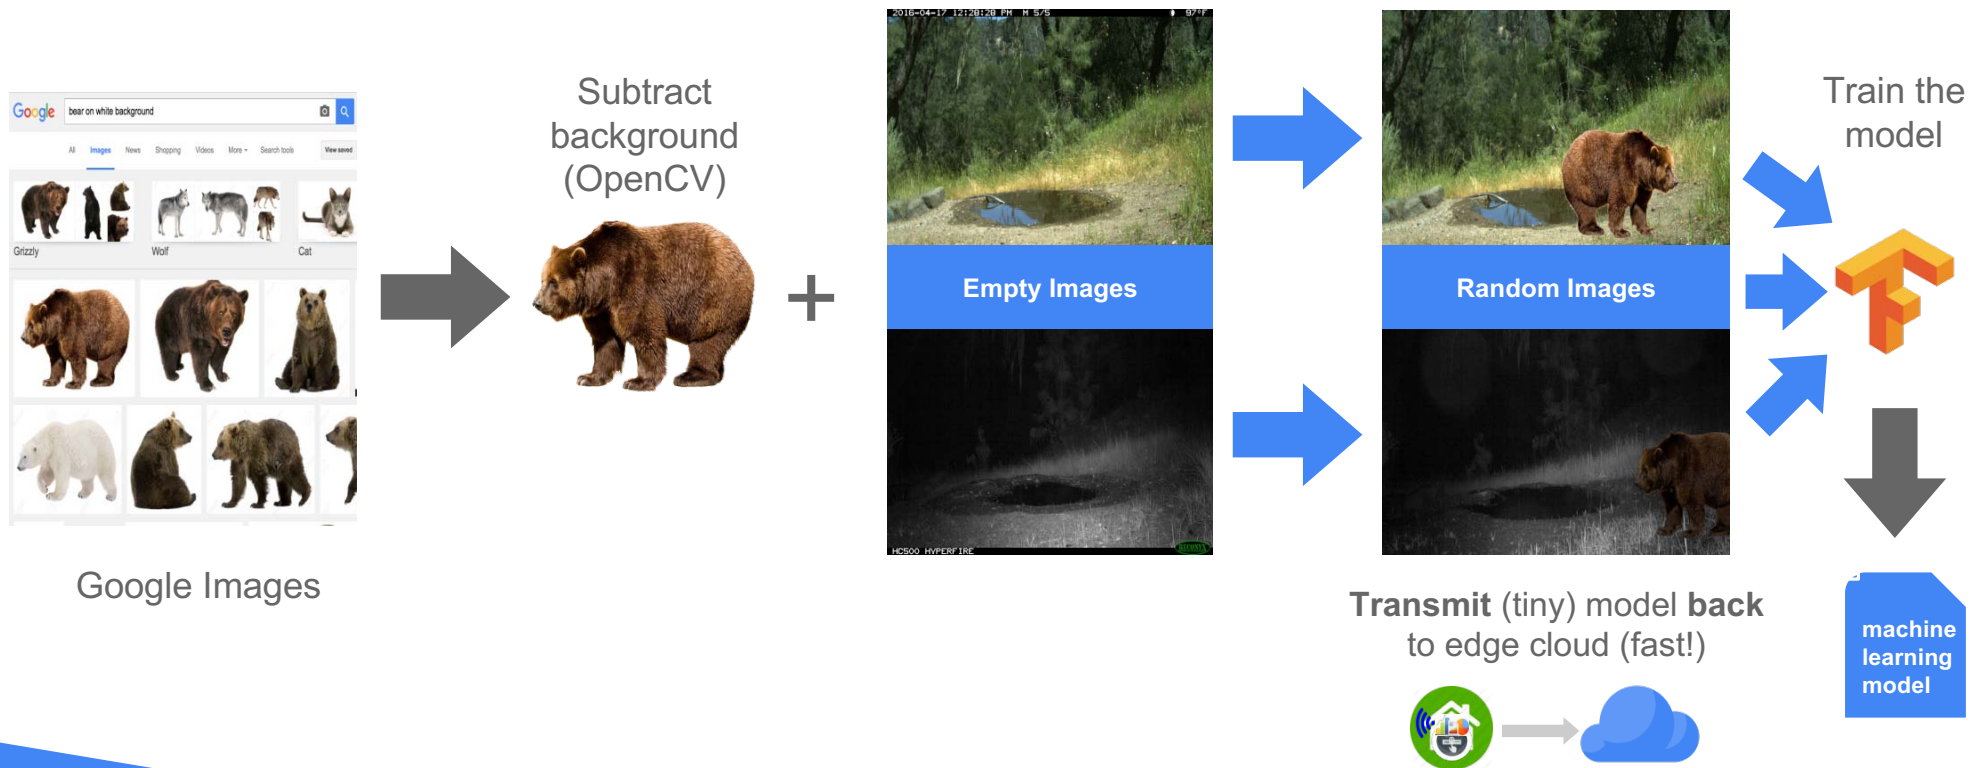


WTB: Train Model In Public Cloud Using "Fake" Images



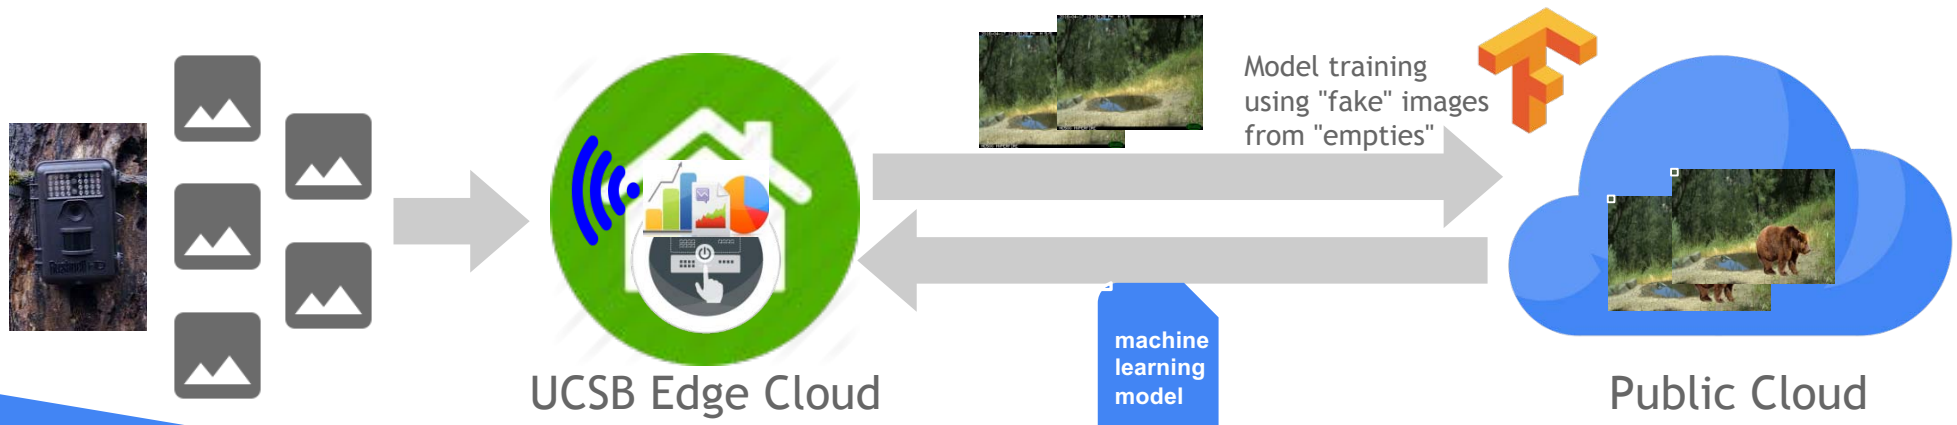
WTB: Train Model In Public Cloud Using "Fake" Images

Public Cloud



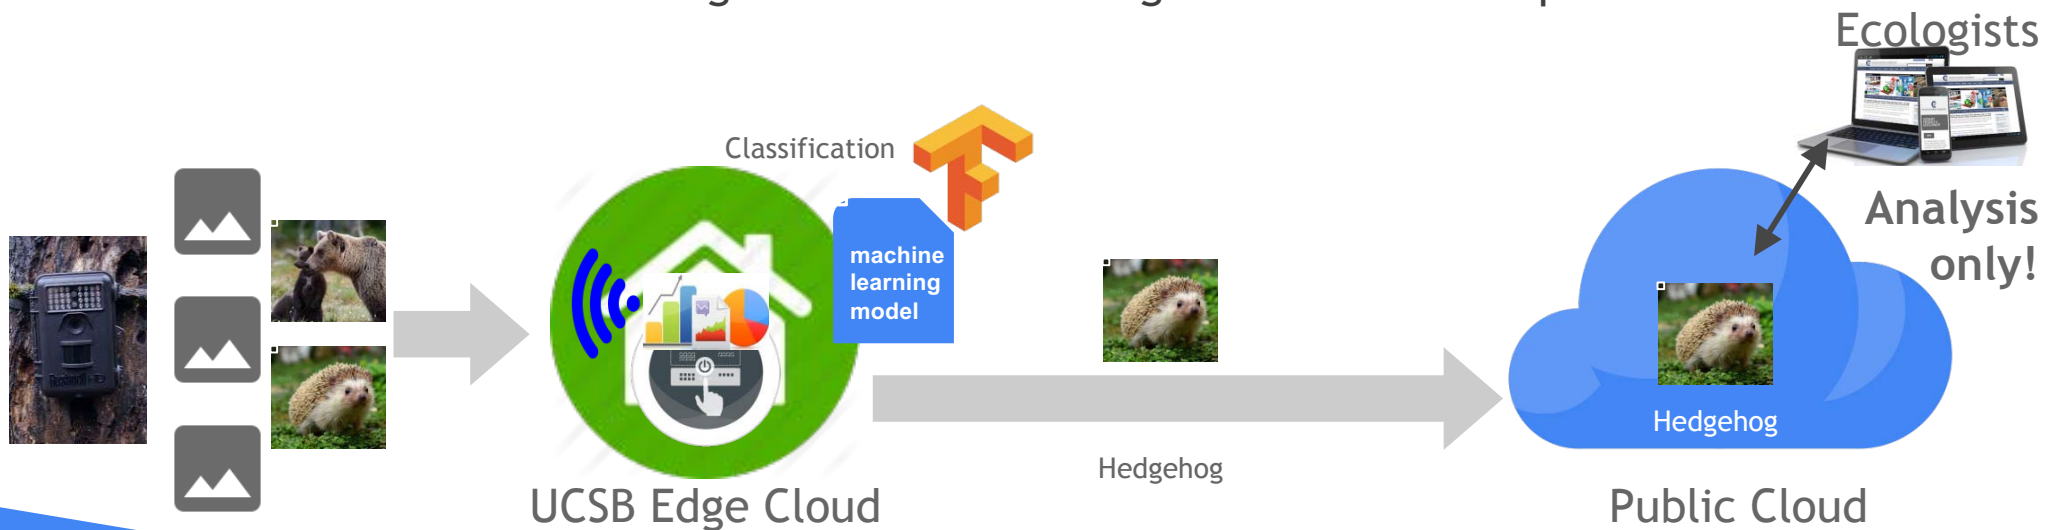
Where's The Bear (WTB)

- Multi-tier IoT system
 - Move the *code to the data (images)* not vice versa
 - Via **Edge Clouds**: robust, self-managing appliances, on-site
 - Low latency, high-bandwidth direct connectivity to cameras
 - Local image classification using model trained in public cloud



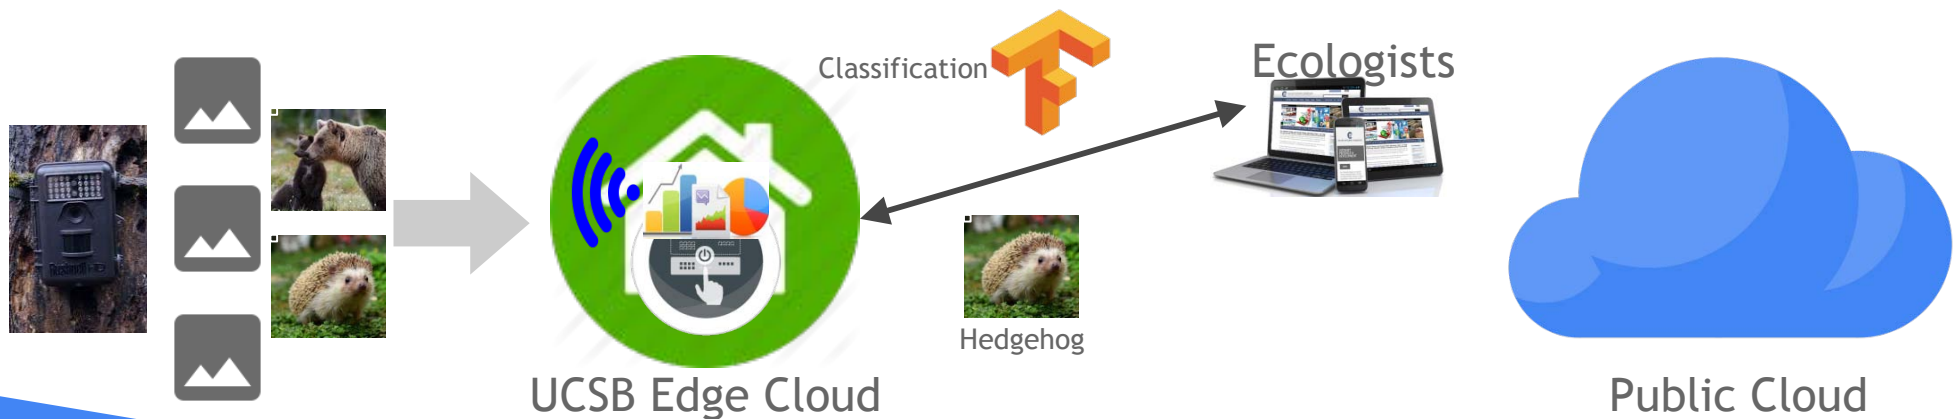
Where's The Bear (WTB)

- Multi-tier IoT system
 - Move the *code to the data (images)* not vice versa
 - Via **Edge Clouds**: robust, self-managing appliances, on-site
 - Low latency, high-bandwidth direct connectivity to cameras
 - Local image classification using model trained in public cloud



Where's The Bear (WTB)

- Multi-tier IoT system
 - Move the ***code to the data (images)*** not vice versa
 - Via ***Edge Clouds***: robust, self-managing appliances, on-site
 - Low latency, high-bandwidth direct connectivity to cameras
 - Local image classification using model trained in public cloud
 - Works even when the Internet isn't



Deployment and Empirical Methodology



SEDGWICK RESERVE
UC SANTA BARBARA NATURAL RESERVE SYSTEM

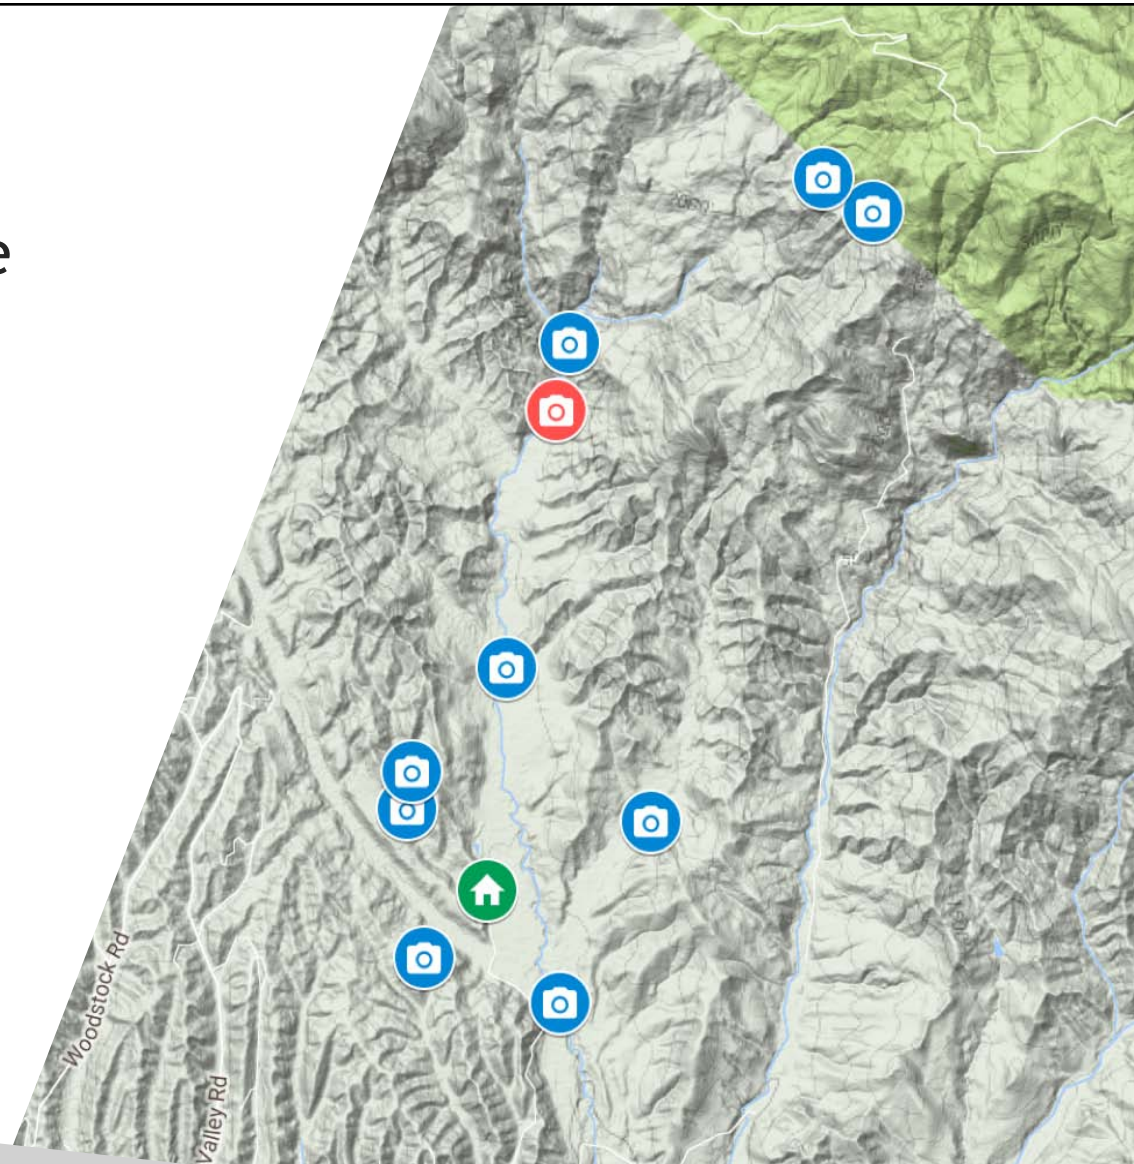


UCSB Sedgwick Reserve

- Motion-triggered Wildlife Camera Traps
- 9 sq. miles

Total: 1.12 Million
Images: 714GB

This study: 238K
Images: 93GB



Where's the Bear?



Results:

4890

randomly selected
images

vs.

1473

images of interest
transmitted

5.3 hrs

vs.

1.6 hrs

Researcher time saved!

< 1% error for bear & coyote; ~10% error for deer

Challenging Images



WTB Findings and Future Work

- **For IoT, public cloud == many tools & services**
 - Public cloud and edge cloud work together to solve problems
 - Can save significant network use (image transfer avoidance)
 - Can save significant researcher time (automatic classification)
 - While maintaining high accuracy
- **WTB Advance: Able to classify animals that appear rarely**
 - In insufficient numbers to train with
- **Next steps**
 - Small animals vs empty images, improving accuracy
 - Identifying features, counting
 - Making edge cloud robust to outage and faults

Thanks!

Students



William Berman



Kyle Carson



Stratos Dimopoulos



Jonathan Easterman



Angad Gill



Nevena Golubovic



Wei-Tsung Lin



Benji Lampel



Kevin Malta



Andy Rosales Elias



Michael Zhang



Chandra Krintz



Rich Wolski

Collaborators:

UCSB, CalPoly, UCDavis, Fresno State, Powwow Energy,
Sedgwick Reserve,
UC Extension, Private Growers

Support:

Google, Huawei, IBM Research, Microsoft Research,
NSF, NIH, California Energy Commission

ckrintz@cs.ucsb.edu

<http://www.cs.ucsb.edu/~ckrintz/racelab.html>

UCSB RACELab

The Lab for Research on Adaptive Computing Environments
Computer Science Department, Harold Frank Hall (E-5), Santa Barbara, CA