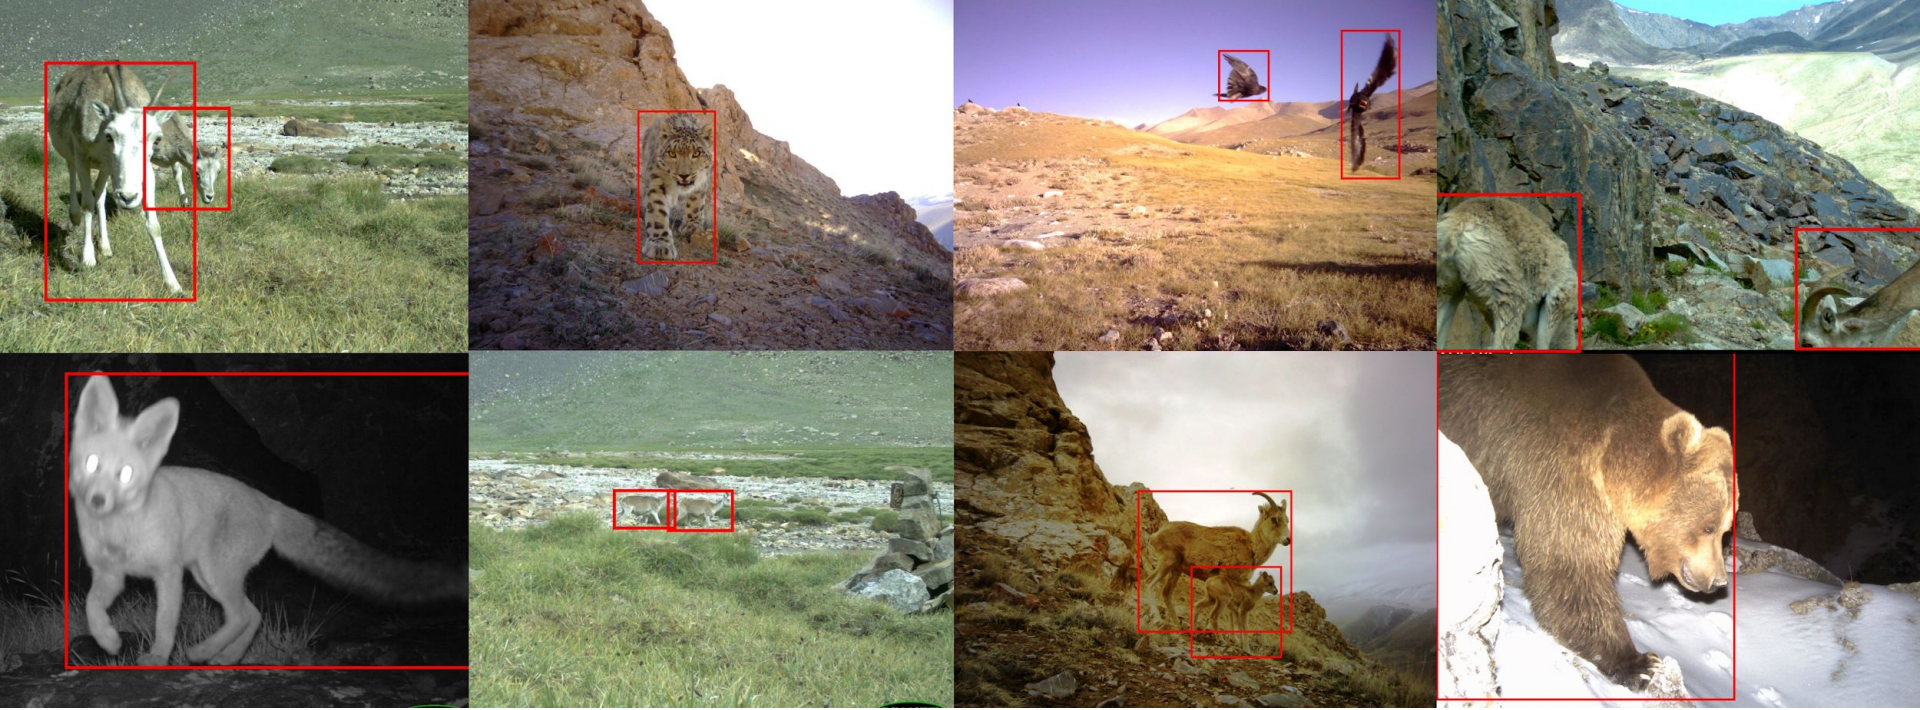

Camera traps: generalization, sample efficiency, best practices, benchmarks, and de-siloing data

Sara Beery
Caltech/Microsoft AI for Earth/Google Perception & Geo-Conservation
CVWC @ ICCV 2019




iNaturalist


Susan Elliott

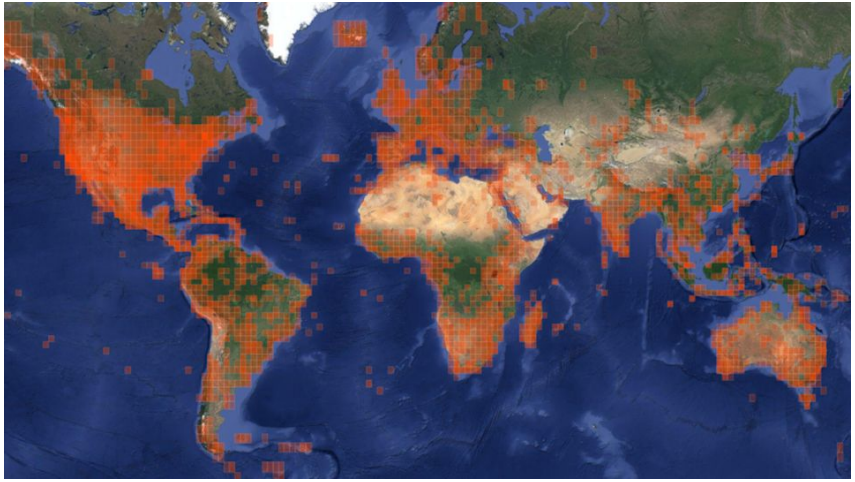

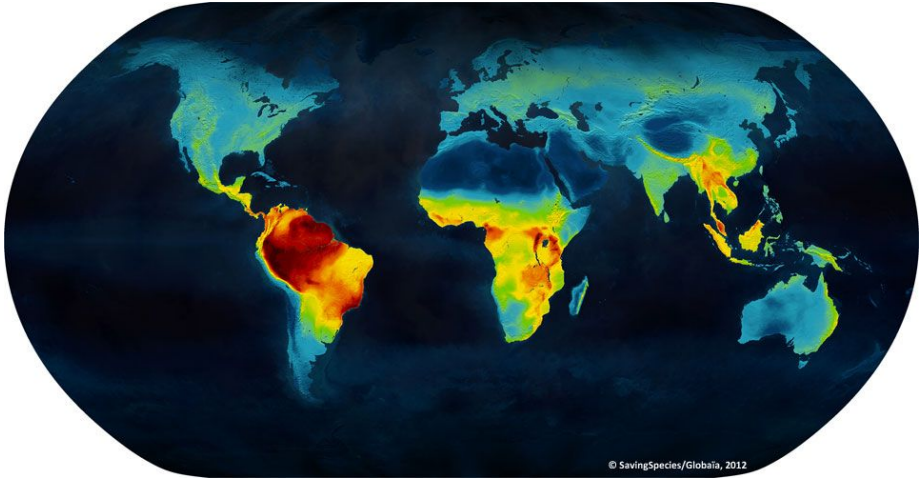
© SavingSpecies/Globaia, 2012


Susan Elliott


Ltl Acorn 0009
062°F 017°C


09.27.2012 Ltl Acorn 0009
082°F 028°C


10.19.2012 16:57:09
B 1:07:42 AM M 3/3
54°F
ERT PRO
RECONYX

Data Challenges

(1) Illumination

(2) Blur

(3) ROI Size

(4) Occlusion

(5) Camouflage

(6) Perspective

All these images have an animal in them

(1) Illumination

(2) Blur

(3) ROI Size

(4) Occlusion

(5) Camouflage

(6) Perspective

On average 70% of the images from each camera are empty

Classification Accuracy

Best-fit line through per-species error

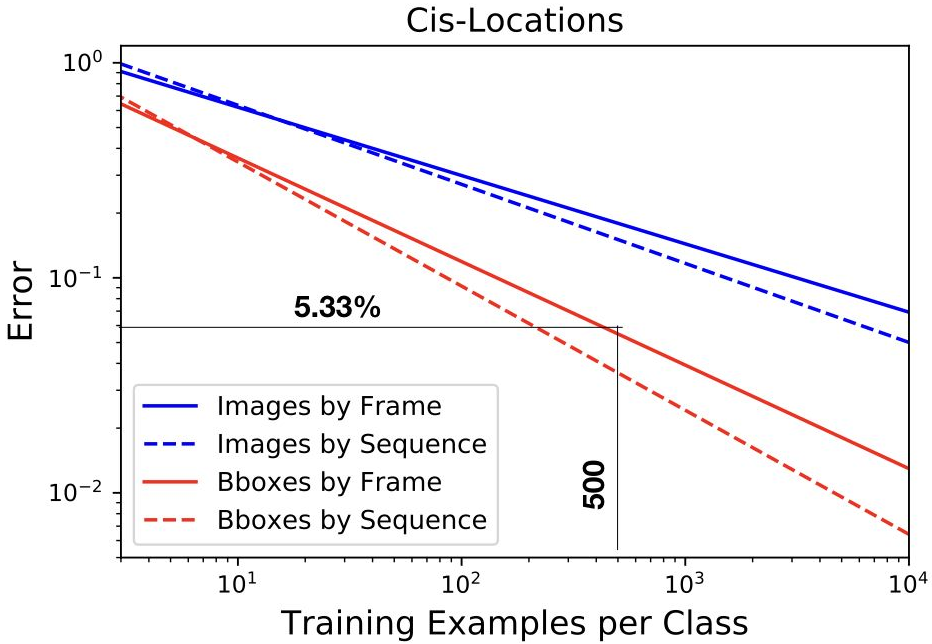

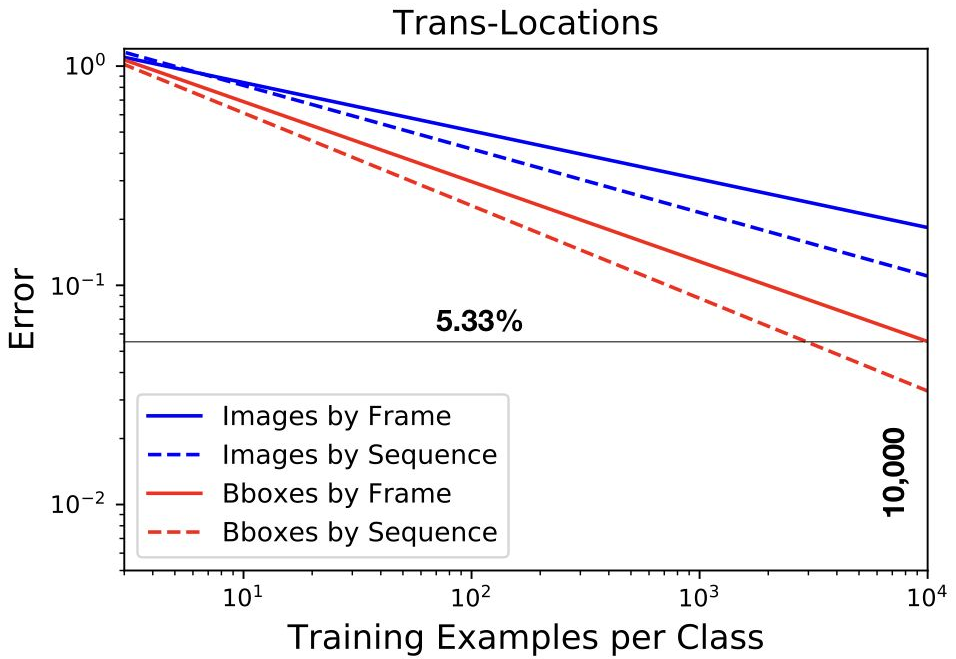

Beery *et al.*, *Recognition in Terra Incognita,* ECCV18

iNaturalist 2018 Challenge Winner

Classification accuracy across 8K species

Needs Some Improvement

Looking Pretty Good

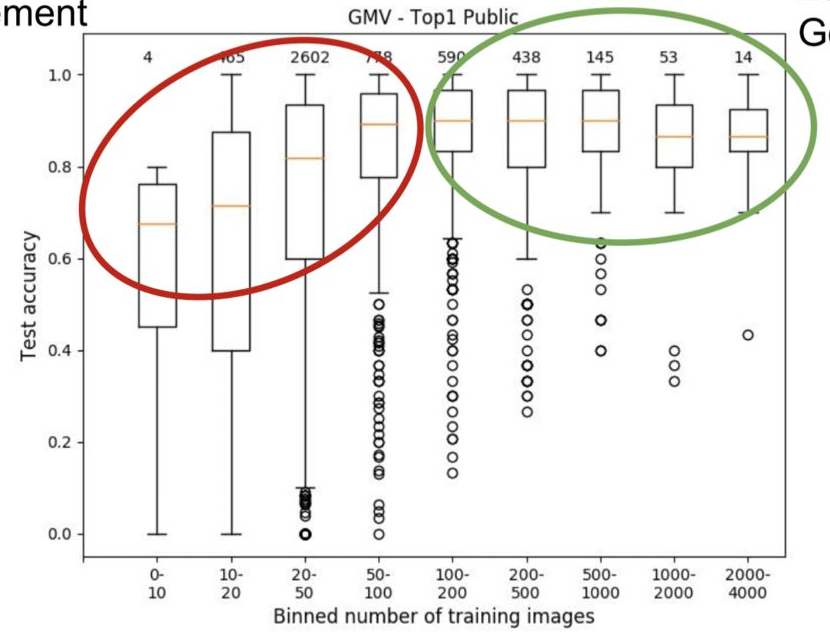

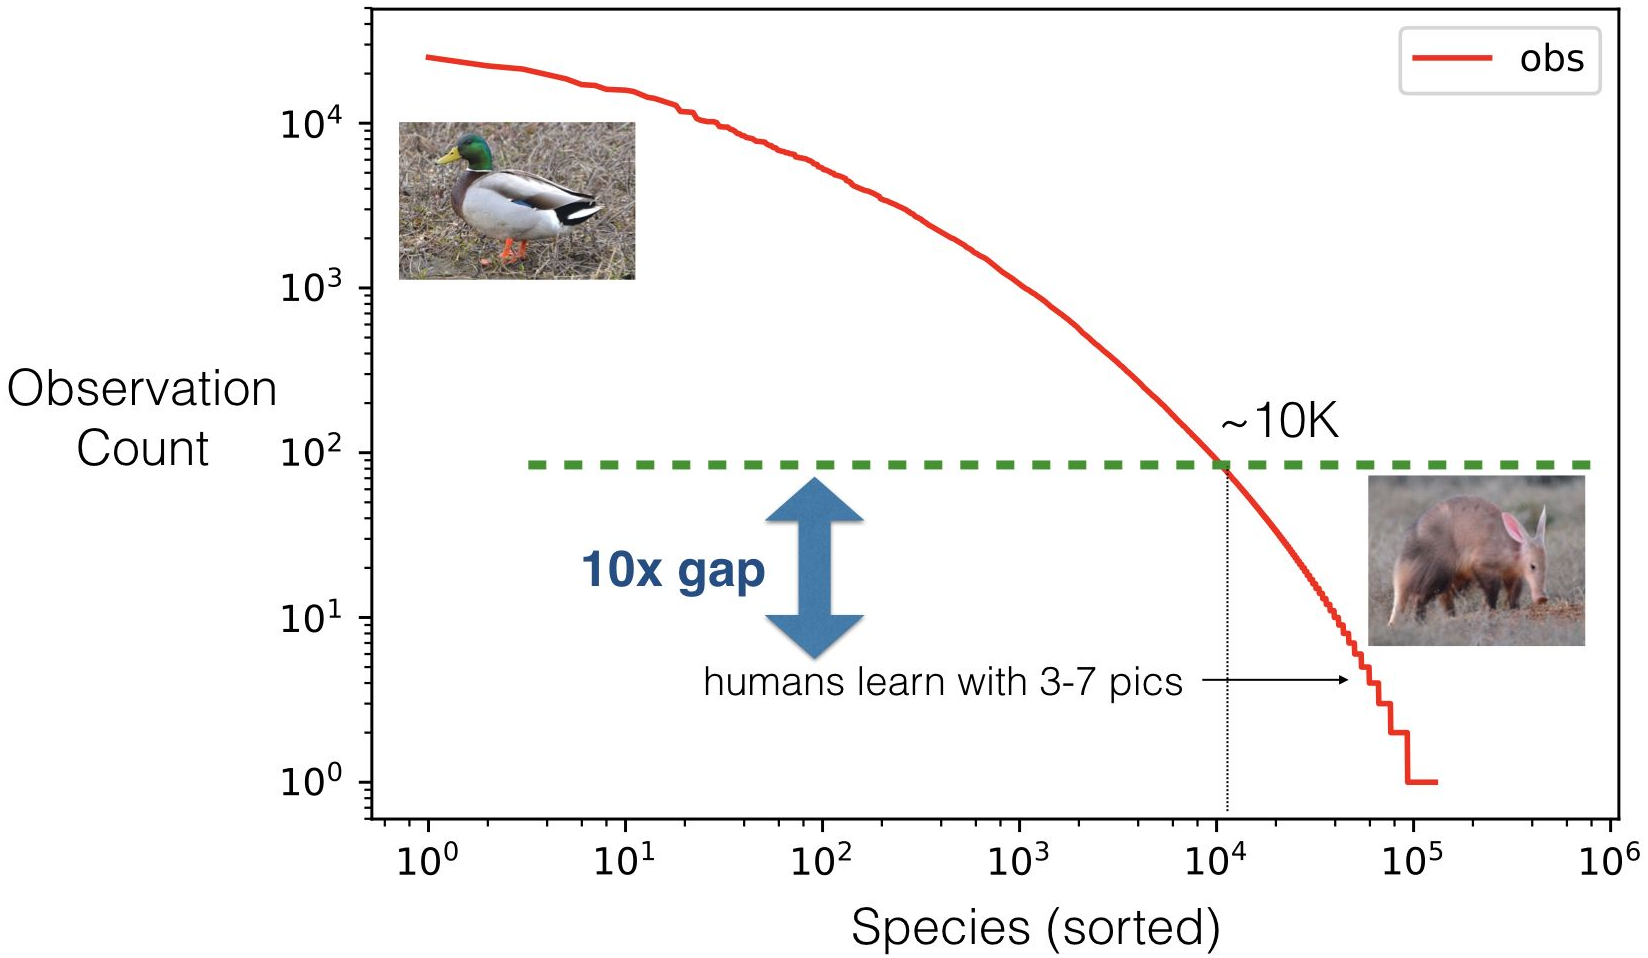
Observations per iNaturalist Species: 16 M total
obs
Observation Count
10^4
10^3
10^2
10^1
10^0
~10K
10x gap
humans learn with 3-7 pics
10^0
10^1
10^2
10^3
10^4
10^5
10^6
Species (sorted)

E.g. learning pose variability












Camera traps are static, and animals are habitual.

Sample efficiency from within a single camera is very low.

Classification Accuracy

Best-fit line through per-species error

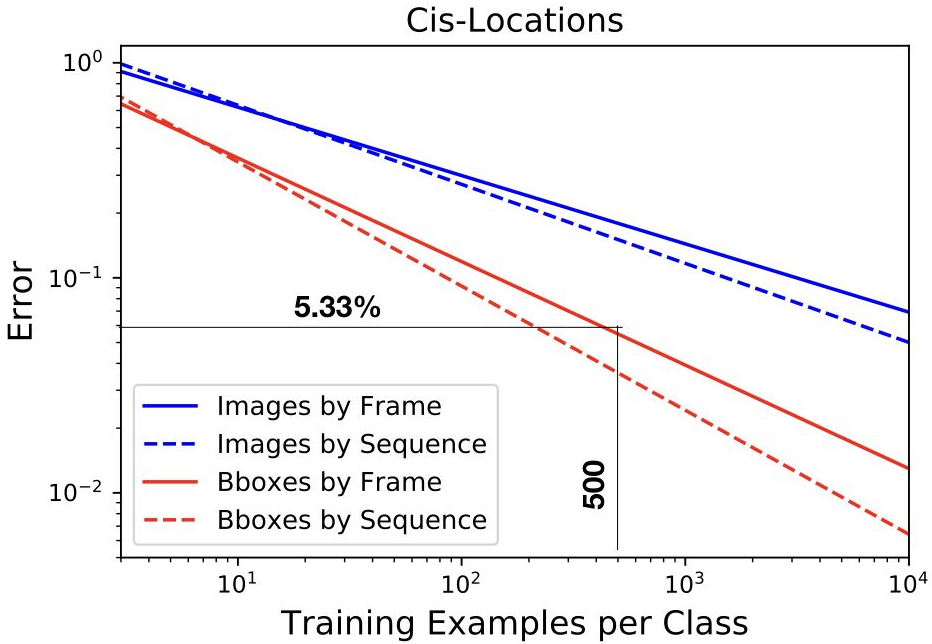

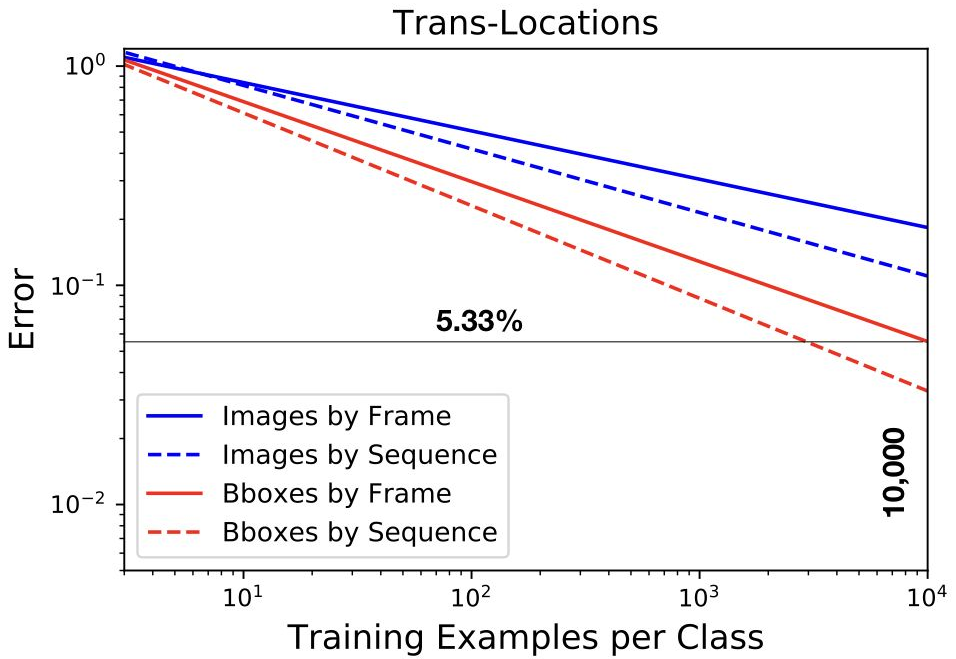

Beery *et al.*, *Recognition in Terra Incognita,* ECCV18

Blank classification via detection

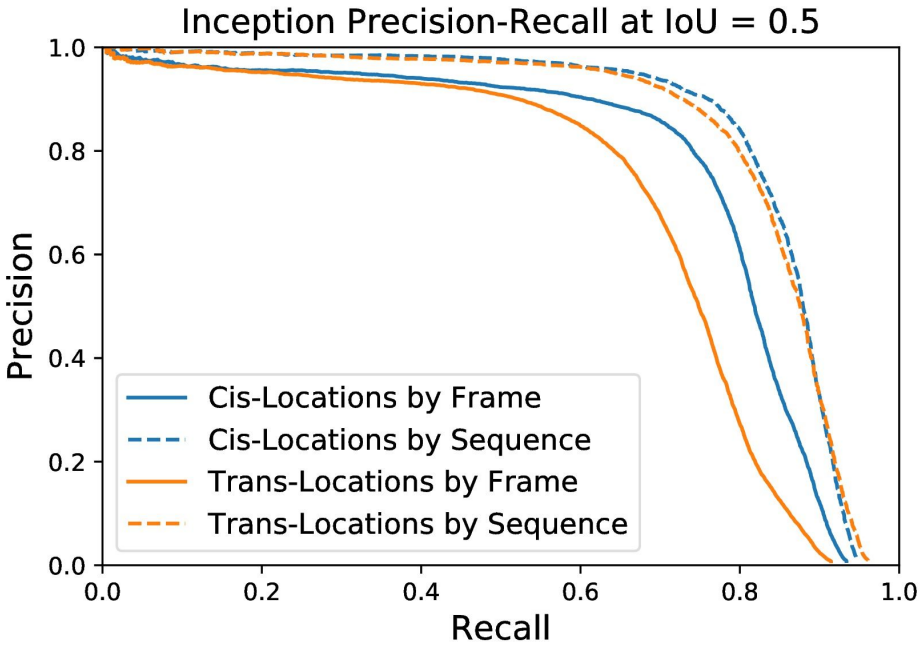

The MEGA DETECTOR

Microsoft AI for Earth

https://github.com/microsoft/CameraTraps/blob/master/megadetector.md

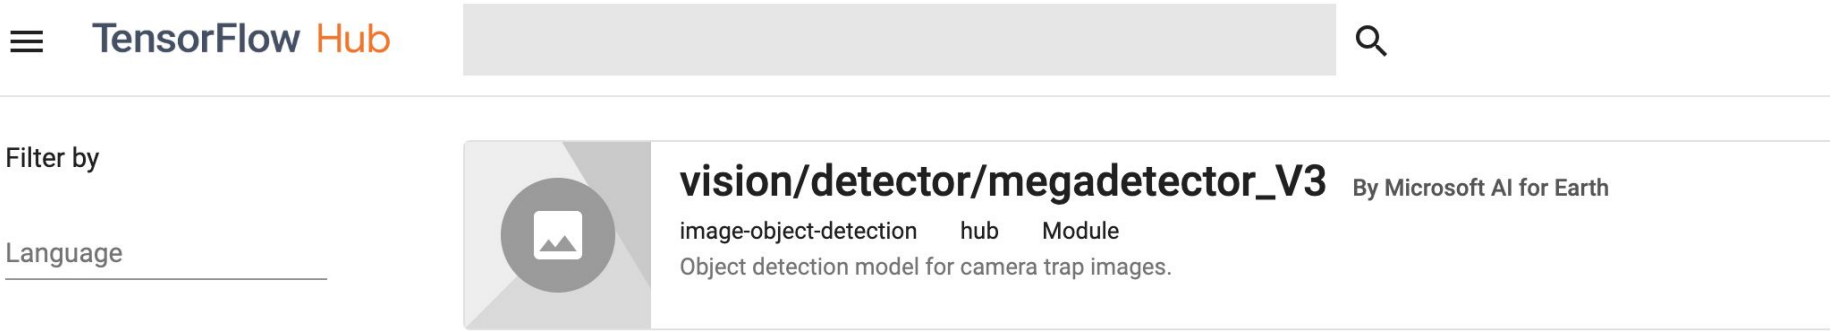

https://overlay.sandbox.google.com/embed?overlay_name=megadetector_v3

IDAHO FISH & GAME

Sorted 4.8 million images in ~2.75 days

This would have taken 10 people working full-time 40 weeks to complete

Survey	Number of labeled images	Average precision	Precision at 90% recall
St. Joe elk	239,006	0.967	0.871
Beaverhead elk	239,910	0.885	0.316
Clear Creek mustelids	199,954	0.988	0.989

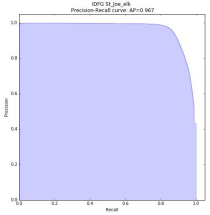

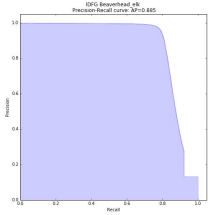

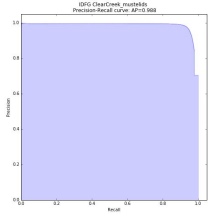

Proposed pipeline

Mega-detector

No boxes? → Empty

Boxes?

Project-Specific Multiclass Detector or Bbox Classifier → Antelope

Combine our class-agnostic detector with project- or region-specific classifiers.

What can we do about the long tail?






(f) Real Camera Traps


(g) TrapCam-Unity


(h) TrapCam-AirSim


(i) Sim on Empty


(j) Real on Empty

Beery, S., Liu, Y., et al., *Synthetic Examples Improve Generalization for Rare Classes*, WACV 2020

How best to leverage temporal signal?

Work done at Google Research, submitting to CVPR 2020

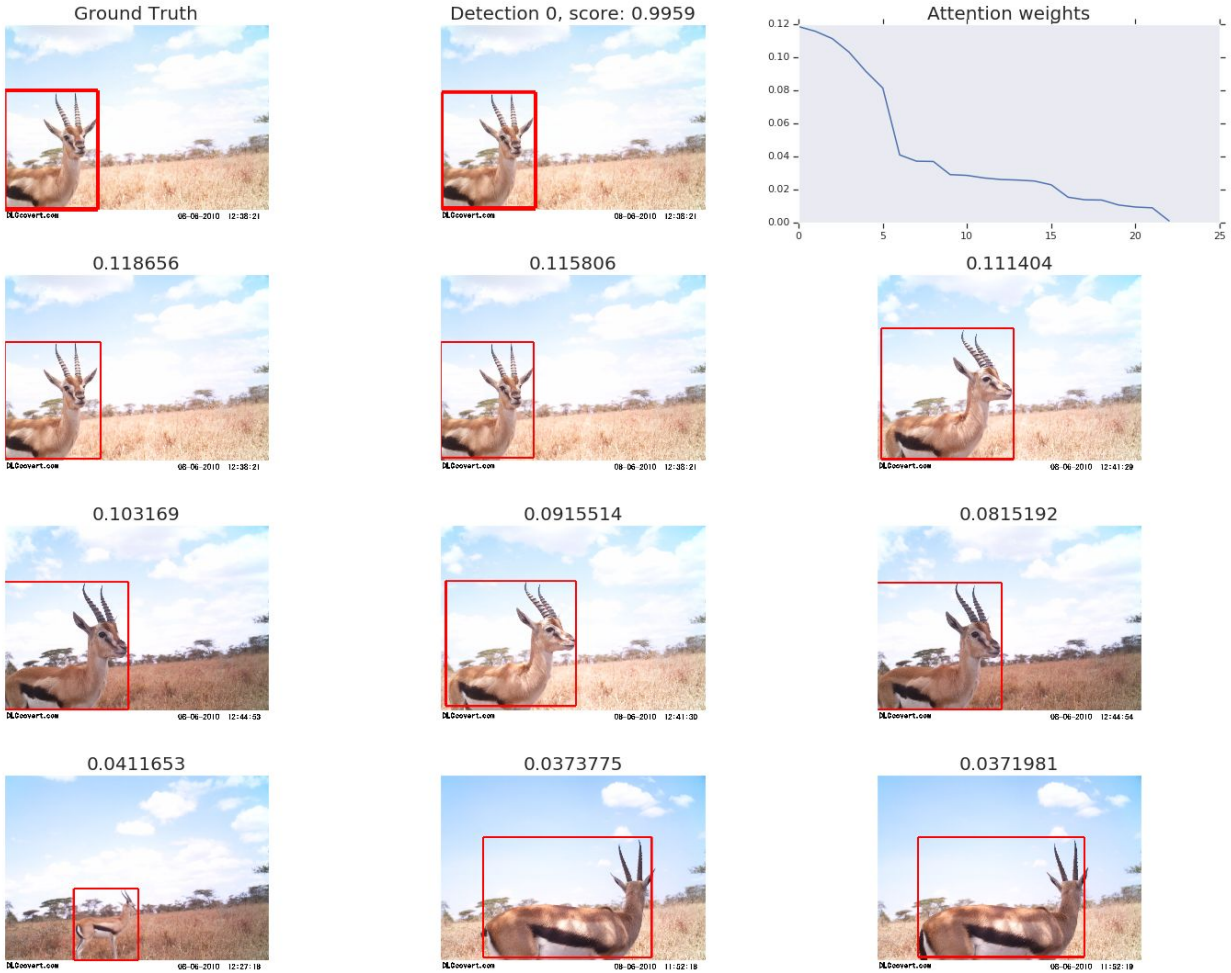

Benchmarks and metrics for the camera trap community

+

Diverse but lightweight, with human-labeled class and bbox data.
The benchmark dataset is available on https://LILA.science

https://www.wildlifeinsights.org/try-ai

Acknowledgements

Microsoft

AI for Earth

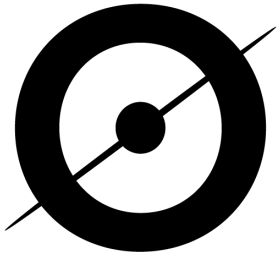

LILA BC

Labeled Information Library of Alexandria: Biology and Conservation

If you would like to stay connected:

aiforconservation@gmail.com

Email here to be invited to the AI for Conservation slack channel

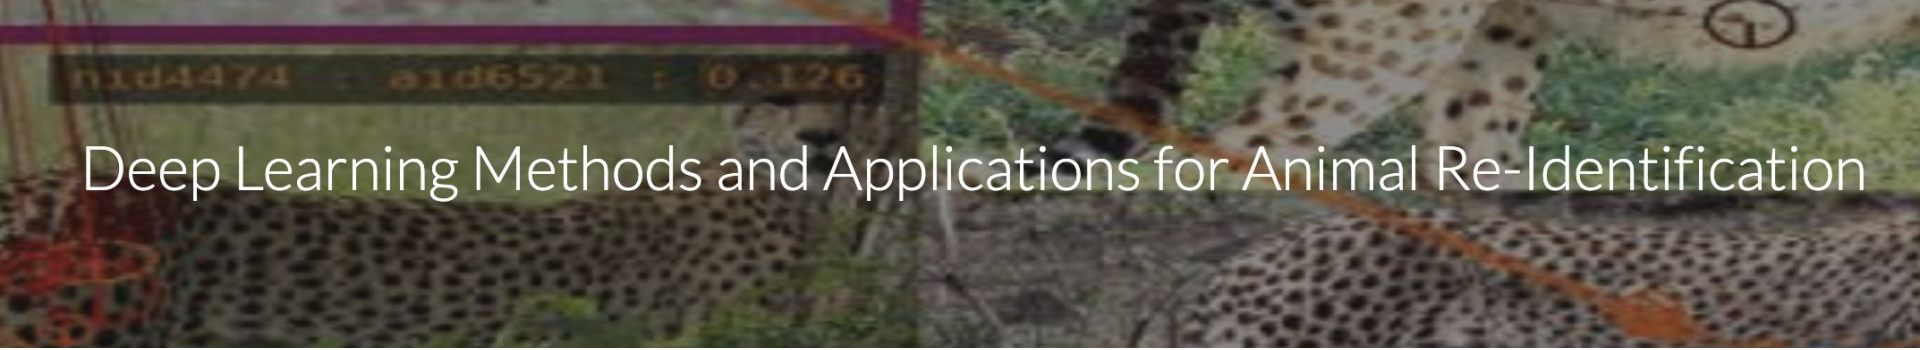

WACV 2020

March 2-5

Aspen, Colorado

https://sites.google.com/view/wacv2020animalreid/