Computer Vision for Conservation

Sara Beery

EE/CNS/CS 148 - May 26, 2020

Big goal: monitoring biodiversity, globally and in real time.

Big goal: monitoring biodiversity, globally and in real time.

How can we contribute?




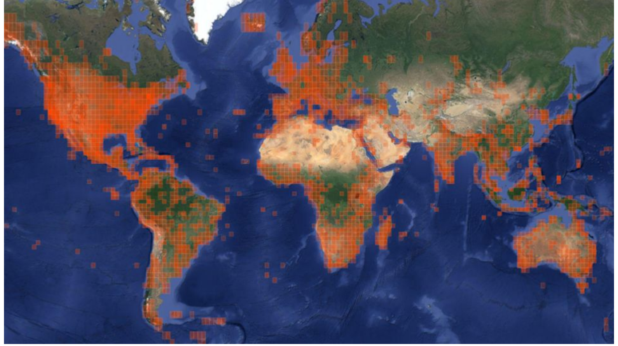

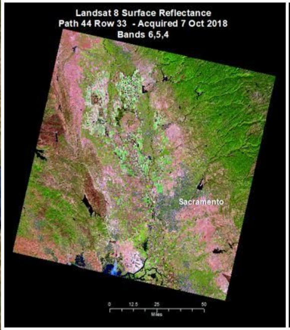
Landsat 8 Surface Reflectance
Path 44 Row 33 - Acquired 7 Oct 2018
Bands 6,5,4
Sacramento
Miles
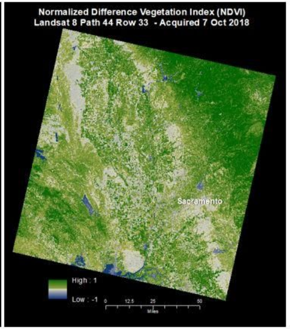
Normalized Difference Vegetation Index (NDVI)
Landsat 8 Path 44 Row 33 - Acquired 7 Oct 2018
Sacramento
High : 1
Low : -1
Miles
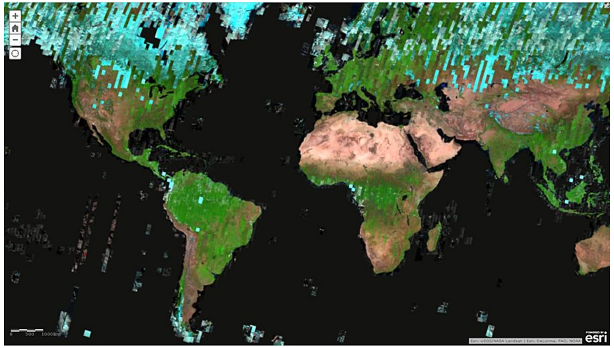
esri


Ltl Acorn 0009
062°F 017°C
09.27.2012

Ltl Acorn 0009
082°F 028°C
10.19.2012 16:57:09
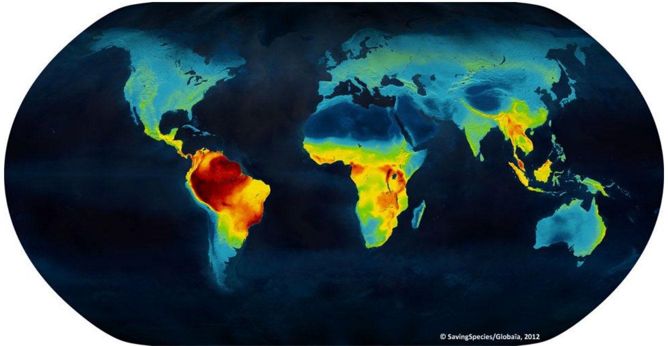
© SavingSpecies/Globaïa, 2012

In-situ Monitoring + Remote Sensing + Citizen Science

In-situ Monitoring

+

Remote Sensing

+

Citizen Science

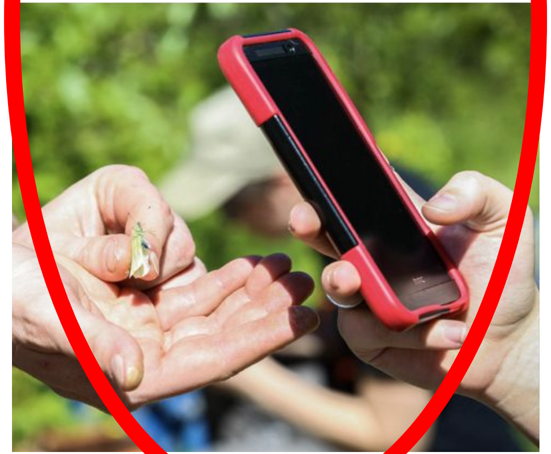

iNaturalist

www.inaturalist.org

CALIFORNIA ACADEMY OF SCIENCES

NATIONAL GEOGRAPHIC

iNaturalist is a joint initiative of the California Academy of Sciences and the National Geographic Society.

How It Works

1 Record your observations

2 Share with fellow naturalists

3 Discuss your findings

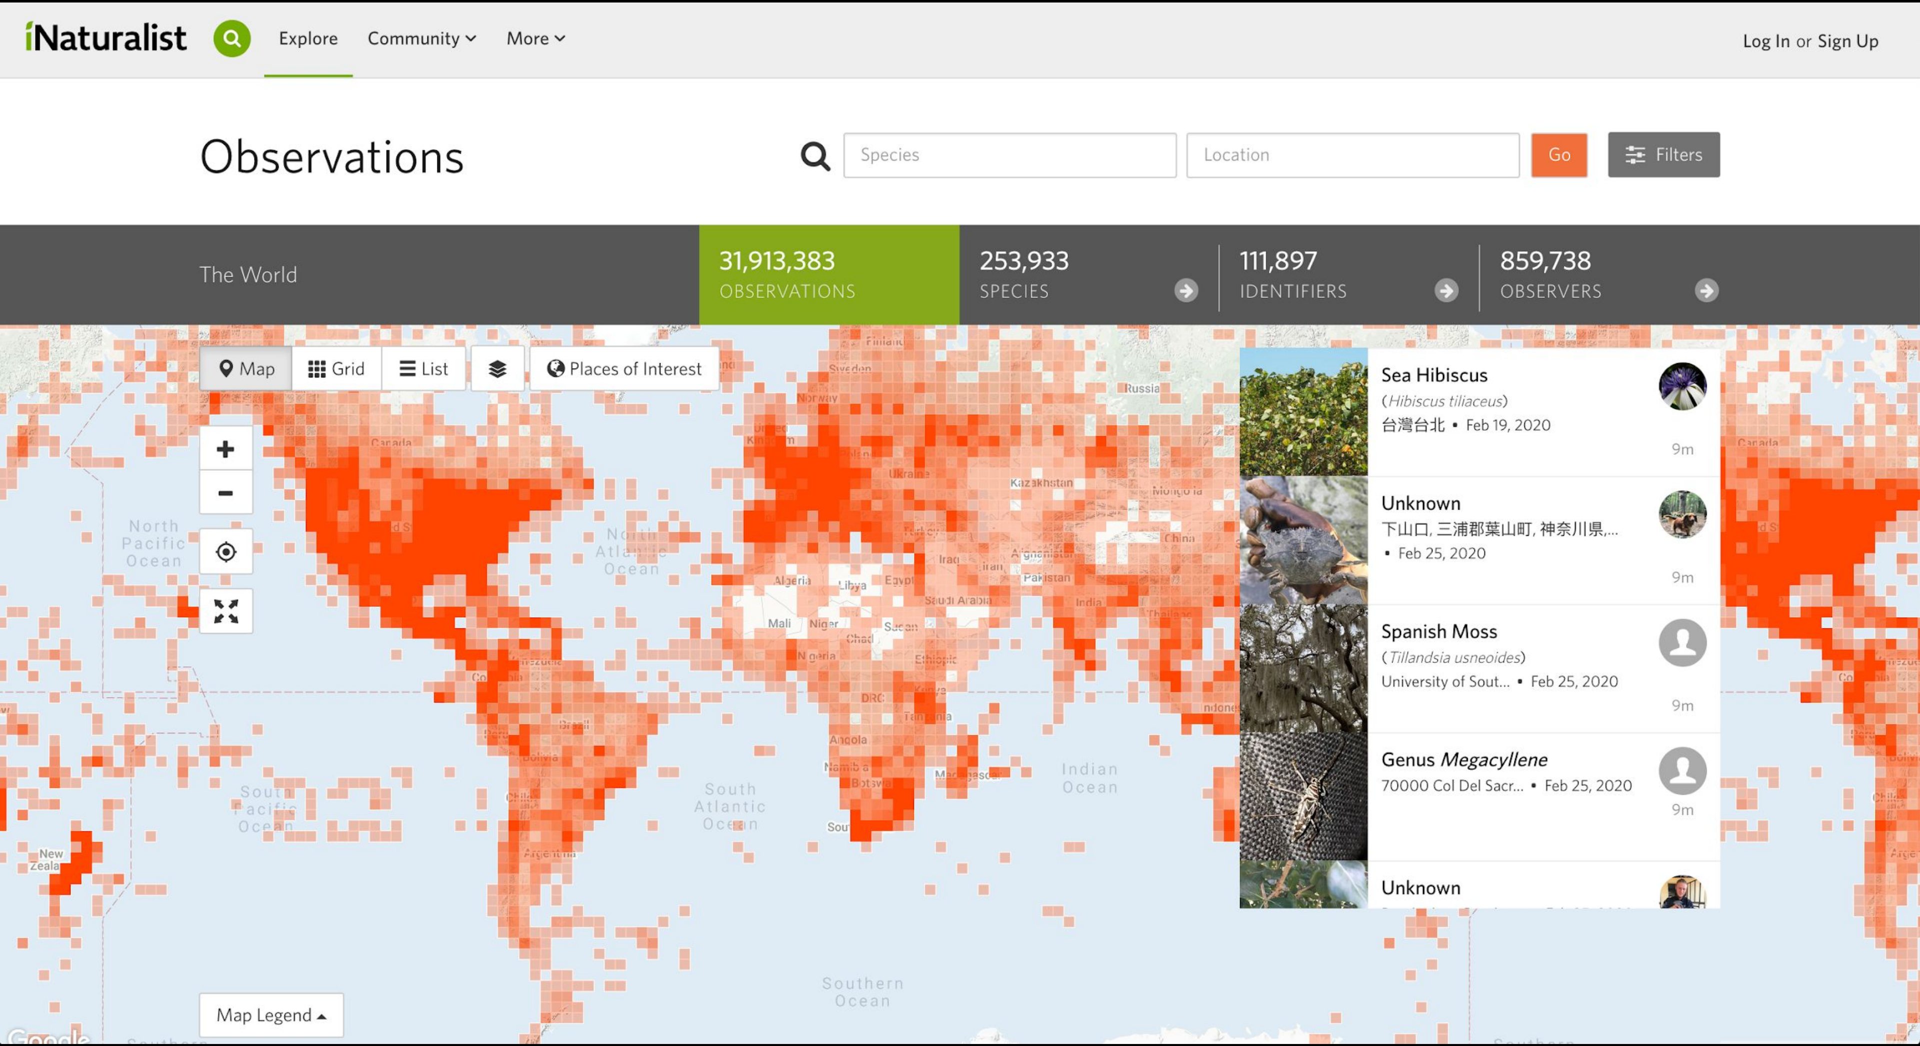
iNaturalist
Explore
Community
More
Log In or Sign Up
Observations
Species
Location
Go
Filters
The World
31,913,383
OBSERVATIONS
253,933
SPECIES
111,897
IDENTIFIERS
859,738
OBSERVERS
Map
Grid
List
Places of Interest
Sea Hibiscus
(Hibiscus tiliaceus)
台灣台北 • Feb 19, 2020
9m
Unknown
下山口, 三浦郡葉山町, 神奈川県,...
• Feb 25, 2020
9m
Spanish Moss
(Tillandsia usneoides)
University of Sout... • Feb 25, 2020
9m
Genus Megacyllene
70000 Col Del Sacr... • Feb 25, 2020
9m
Unknown
North Pacific Ocean
North Atlantic Ocean
South Pacific Ocean
South Atlantic Ocean
Indian Ocean
Southern Ocean
Map Legend

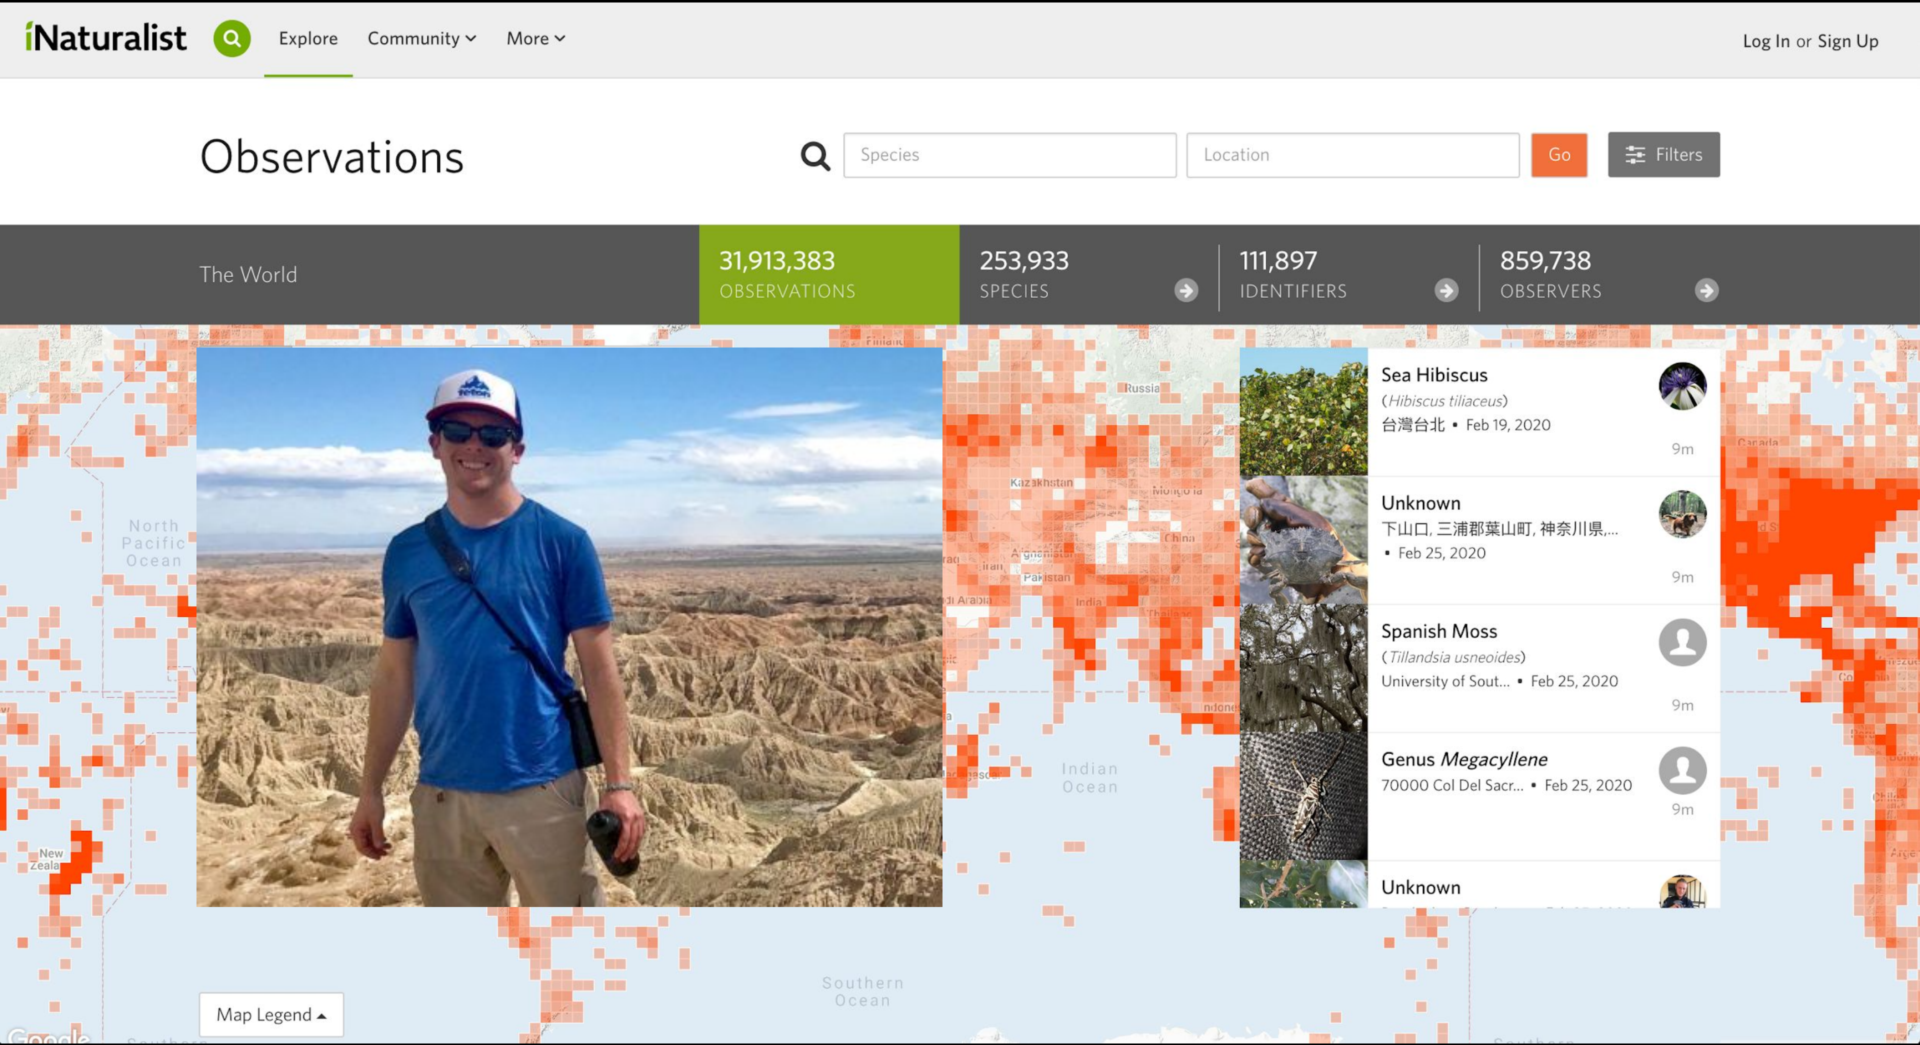

iNaturalist
Explore
Community
More
Log In or Sign Up
Observations
Species
Location
Go
Filters
The World
31,913,383
OBSERVATIONS
253,933
SPECIES
111,897
IDENTIFIERS
859,738
OBSERVERS
Sea Hibiscus
(Hibiscus tiliaceus)
台灣台北 • Feb 19, 2020
9m
Unknown
下山口, 三浦郡葉山町, 神奈川県,...
• Feb 25, 2020
9m
Spanish Moss
(Tillandsia usneoides)
University of Sout... • Feb 25, 2020
9m
Genus Megacyllene
70000 Col Del Sacr... • Feb 25, 2020
9m
Unknown
North Pacific Ocean
Indian Ocean
Southern Ocean
Russia
Kazakhstan
China
India
Pakistan
New Zeala
Map Legend

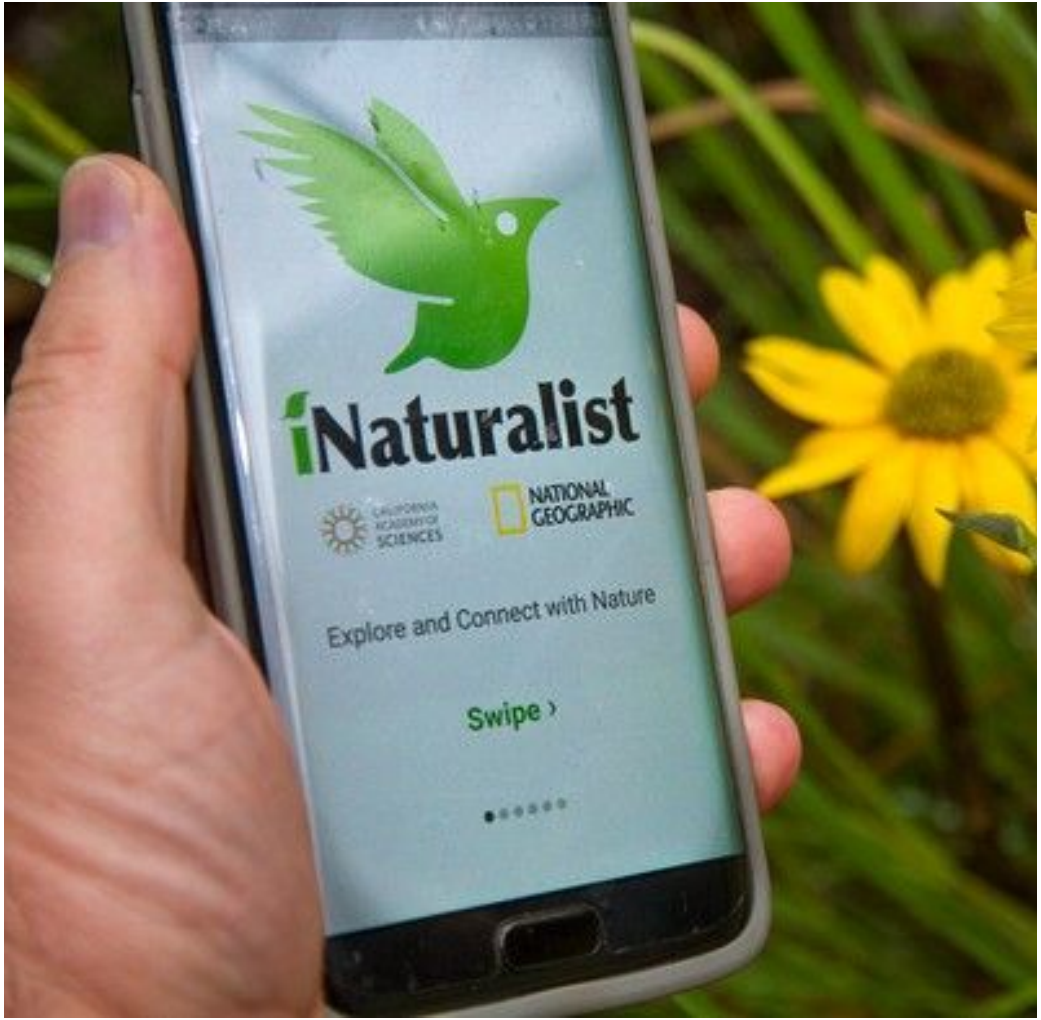
iNaturalist
NATIONAL GEOGRAPHIC
Explore and Connect with Nature
Swipe ›

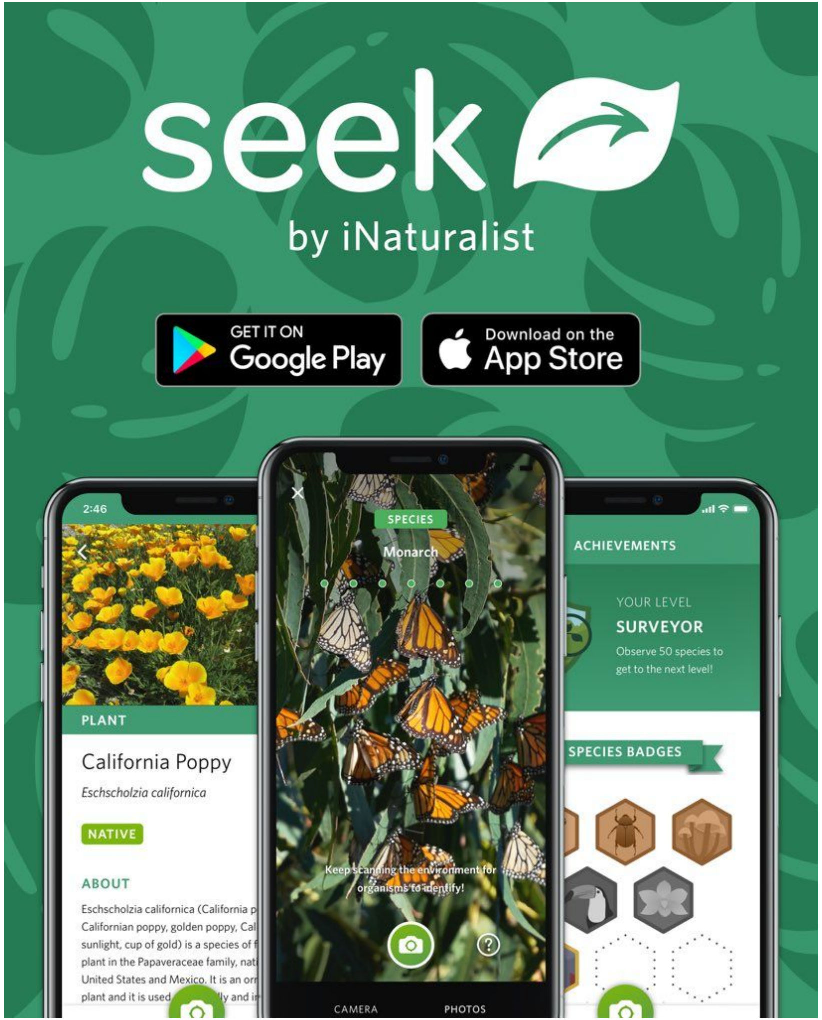
seek
by iNaturalist
GET IT ON Google Play
Download on the App Store
2:46
PLANT
California Poppy
Eschscholzia californica
NATIVE
ABOUT
SPECIES
Monarch
Keep scanning the environment for organisms to identify!
CAMERA
PHOTOS
ACHIEVEMENTS
YOUR LEVEL
SURVEYOR
Observe 50 species to get to the next level!
SPECIES BADGES

iNaturalist 2017

5,089 classes
Classification

iNaturalist 2018

8,142 classes
Taxonomy

iNaturalist 2019

1,100 classes
Similar Species

The iNaturalist Species Classification and Detection Dataset
CVPR 2018
Van Horn, Mac Aodha, Song, Cui, Sun, Shepard, Adam, Perona, Belongie

iNaturalist 2018 Challenge Winner

Classification accuracy across 8K species

Needs Some Improvement

Looking Pretty Good

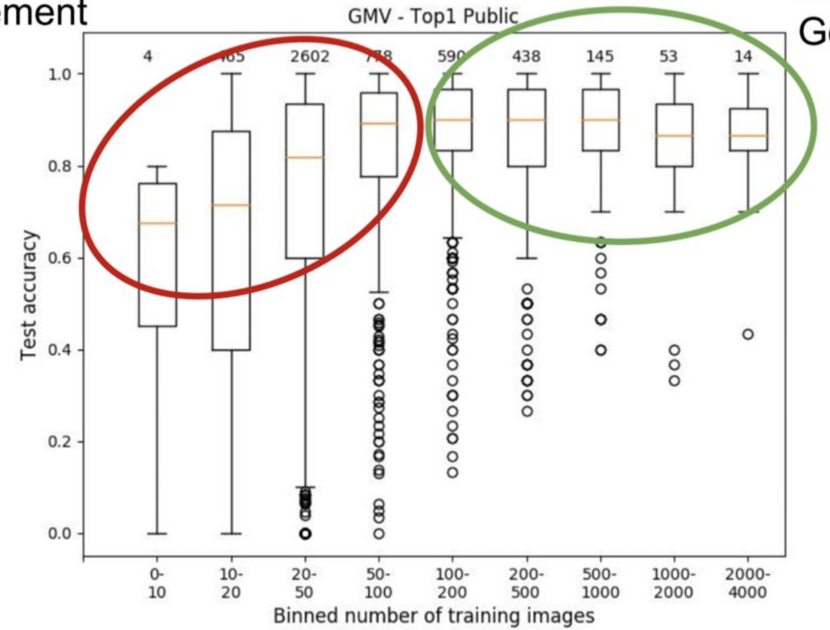

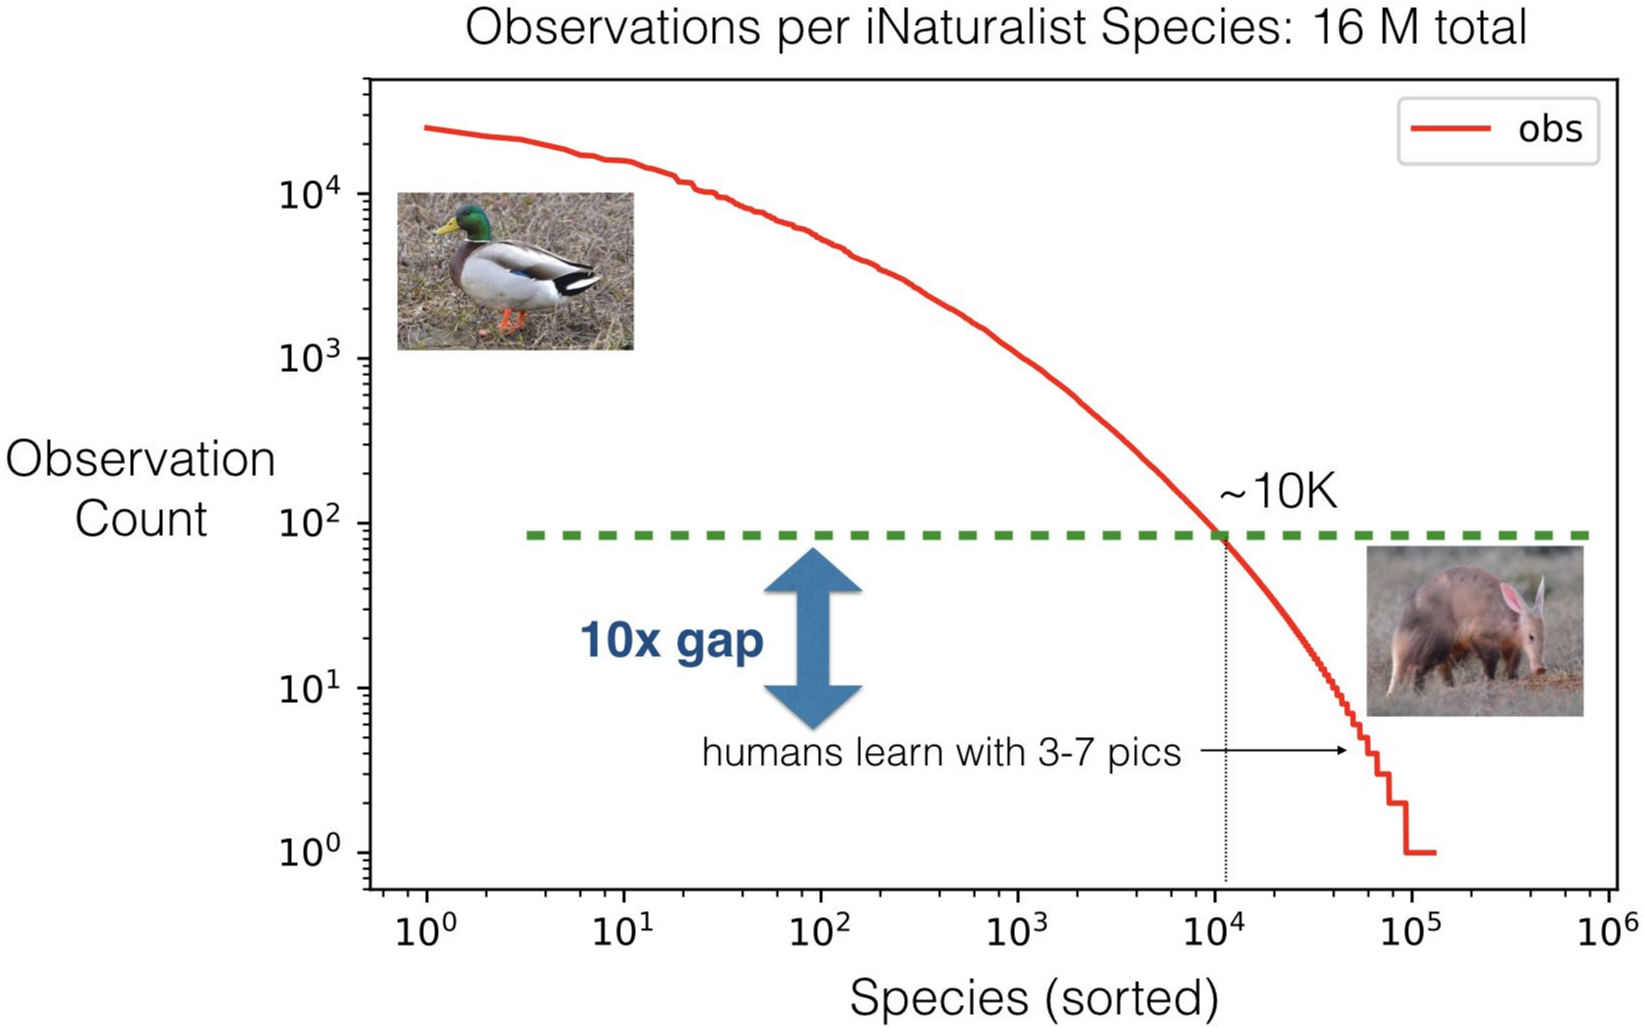

Observations per iNaturalist Species: 16 M total
obs
Observation Count
10^0
10^1
10^2
10^3
10^4
~10K
10x gap
humans learn with 3-7 pics
10^0 10^1 10^2 10^3 10^4 10^5 10^6
Species (sorted)

Can we use information such as where, when, and who captured an image to help determine its class?

Presence-Only Geographical Priors for Fine-Grained Image Classification
ICCV 2019
Mac Aodha, Cole, Perona



Presence-only data:

Northern Mockingbird

- Who
- What
- When
- Where

iNaturalist

Which class $\boldsymbol{y}$ is in image $\boldsymbol{I}$?

$$P(y|I)$$

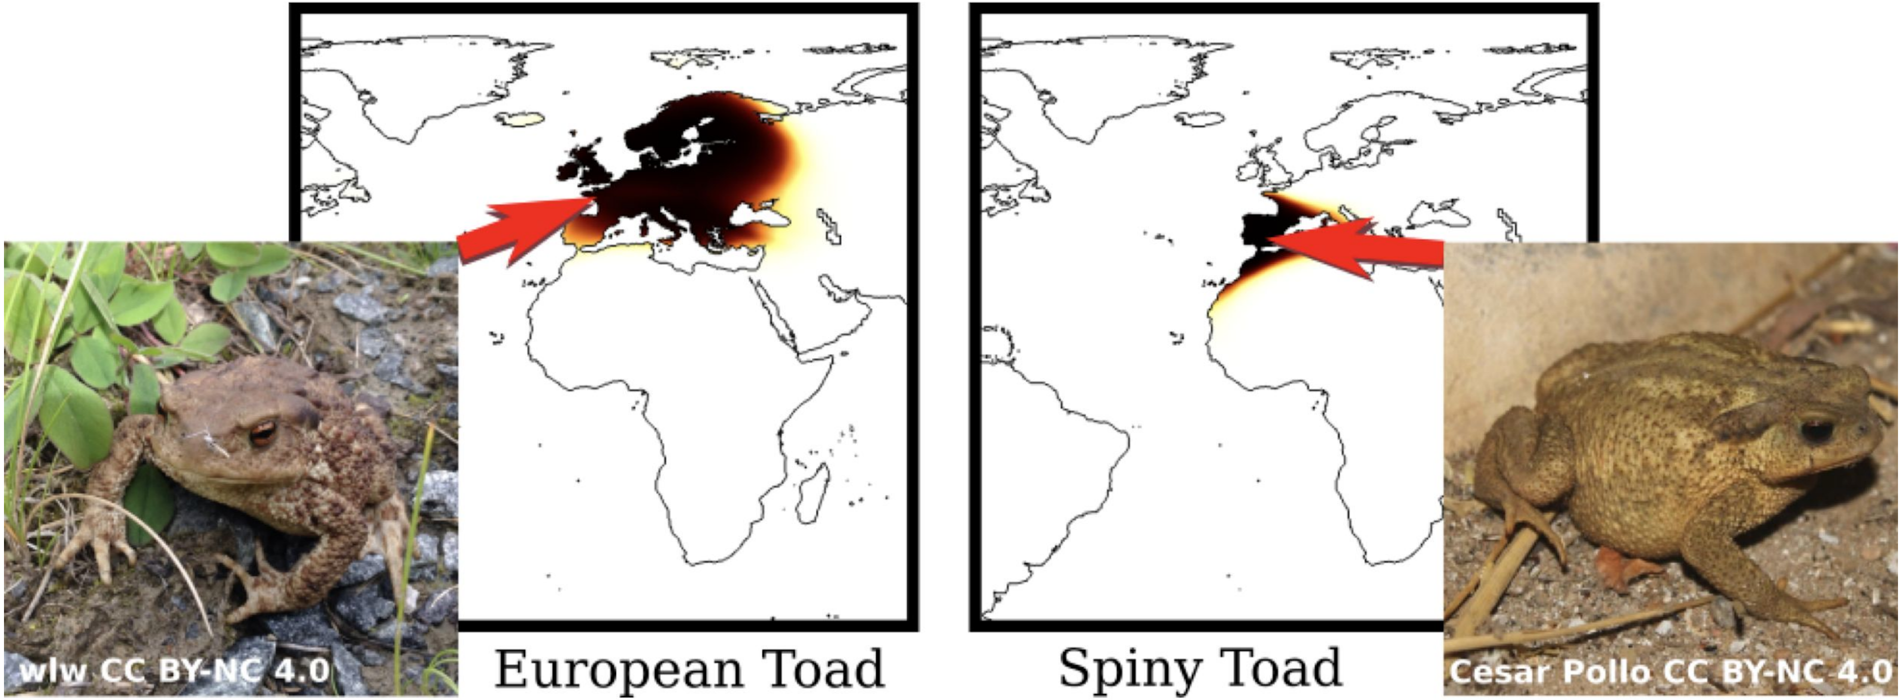
wlw CC BY-NC 4.0
European Toad
Spiny Toad
Cesar Pollo CC BY-NC 4.0

Which class $\boldsymbol{y}$ is in image $\boldsymbol{I}$ at location $\mathbf{x}$?

$$P(y|I,\mathbf{x}) \propto P(y|I)P(y|\mathbf{x})$$

Which class $\boldsymbol{y}$ is in image $\boldsymbol{I}$ at location $\mathbf{x}$?

Image
Classifier

$$P(y|I,\mathbf{x}) \propto \boxed{P(y|I)}P(y|\mathbf{x})$$

Which class y is in image I at location $\mathbf{x}$?

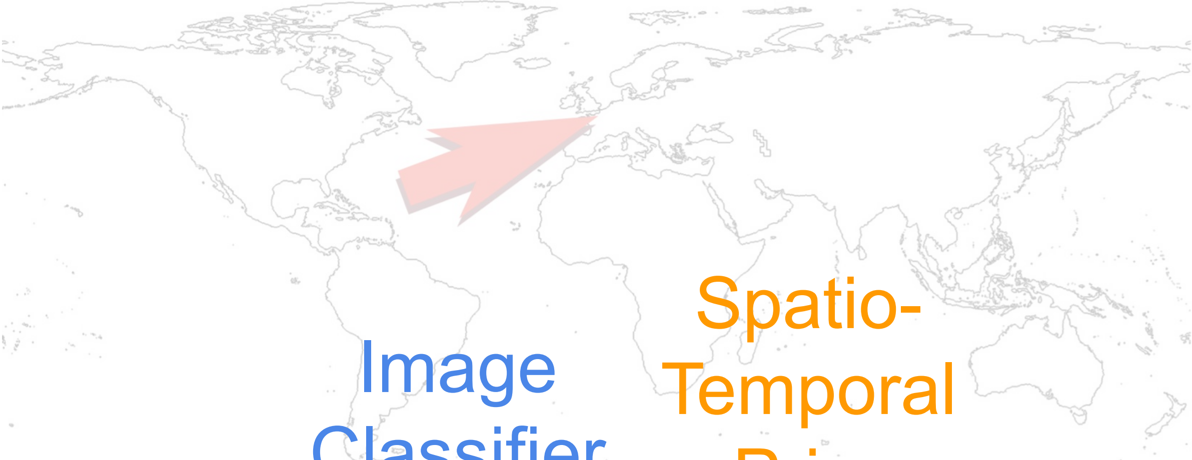

Image Classifier

Spatio-Temporal Prior

$$P(y|I,\mathbf{x}) \propto \boxed{P(y|I)}\boxed{P(y|\mathbf{x})}$$

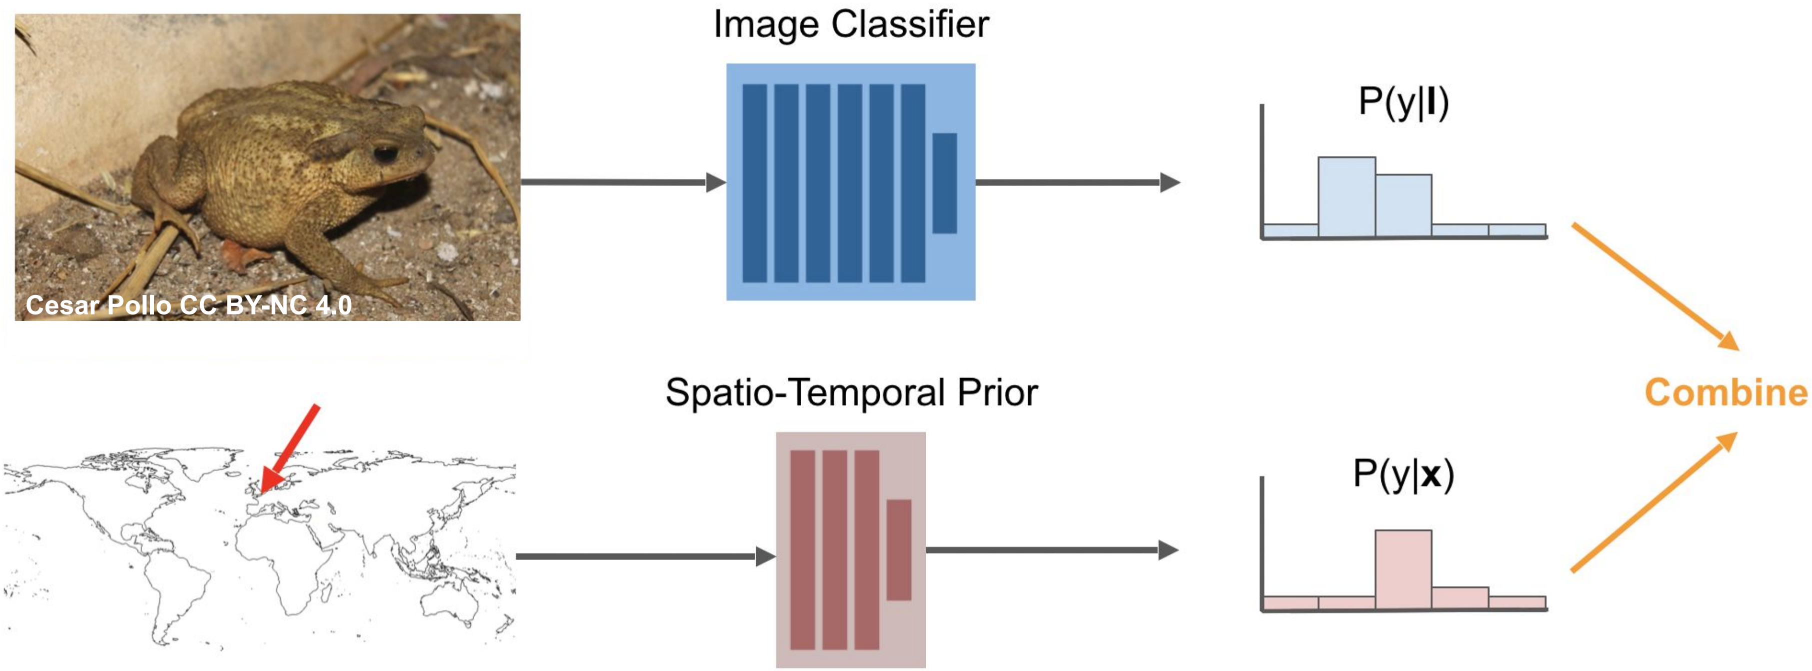

$\mathbf{x}$ = (longitude, latitude, day)

Modular and efficient

Top-1 Classification Results

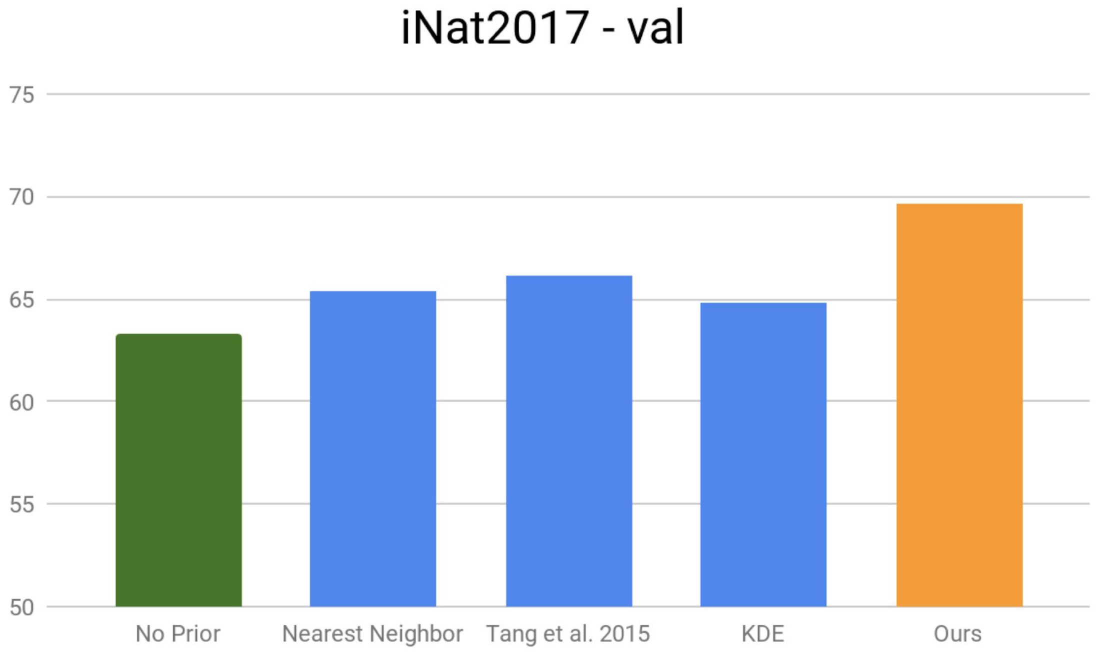

$$P(y|I) \qquad\qquad P(y|I, \mathbf{x})$$

Paper also has results on iNat 2018, NABirds, BirdSnap, YFCC http://www.vision.caltech.edu/~macaodha/projects/geopriors/index.html

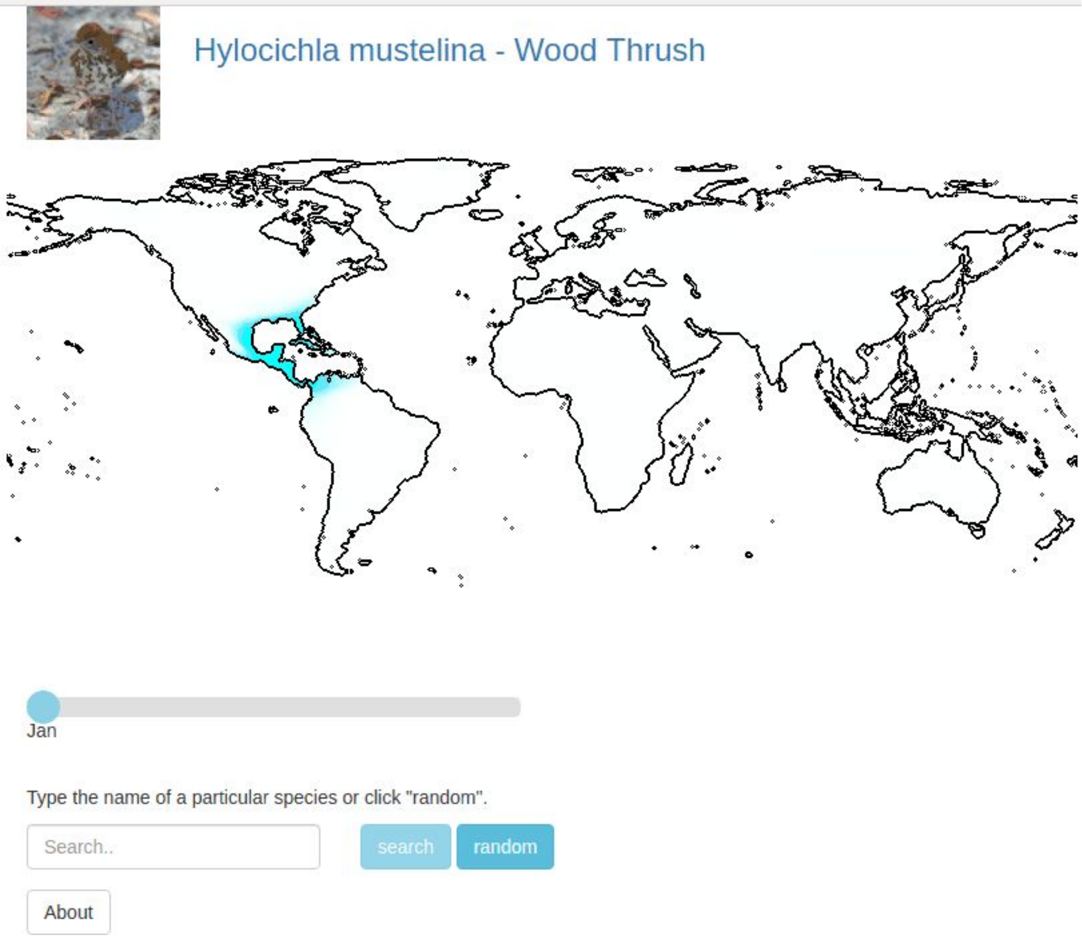

- Trained Models
- Demo
- Code

In-situ Monitoring

Remote Sensing

Citizen Science

Camera traps

- 1,000s of organizations
- 10,000s of projects
- 1,000,000s of camera traps
- 100,000,000s of images

*estimates by Eric Fegraus, Conservation International

Camera traps

- 1,000s of organizations
- 10,000s of projects
- 1,000,000s of camera traps
- 100,000,000s of images

For example: Idaho Department of Fish and Game alone has 5 years of unprocessed, unlabeled data, around 5 million images

*estimates by Eric Fegraus, Conservation International

Wildlife Insights

CONSERVATION INTERNATIONAL

NORTH CAROLINA Museum of Natural Sciences

Smithsonian Institution

ZSL LET'S WORK FOR WILDLIFE

WCS

WWF

MOL MAP OF LIFE

Google

Camera trap data is challenging

(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

SOA models don't generalize

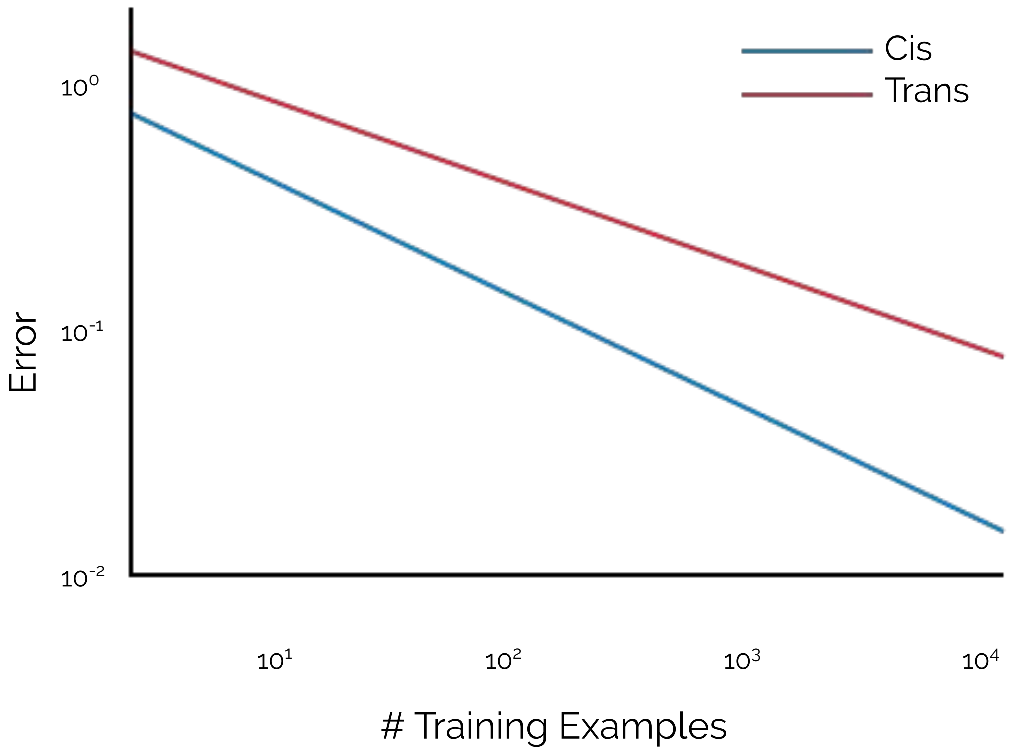

Recognition in Terra Incognita, Beery et al., ECCV 2018

Big increase in error when testing at unseen camera locations

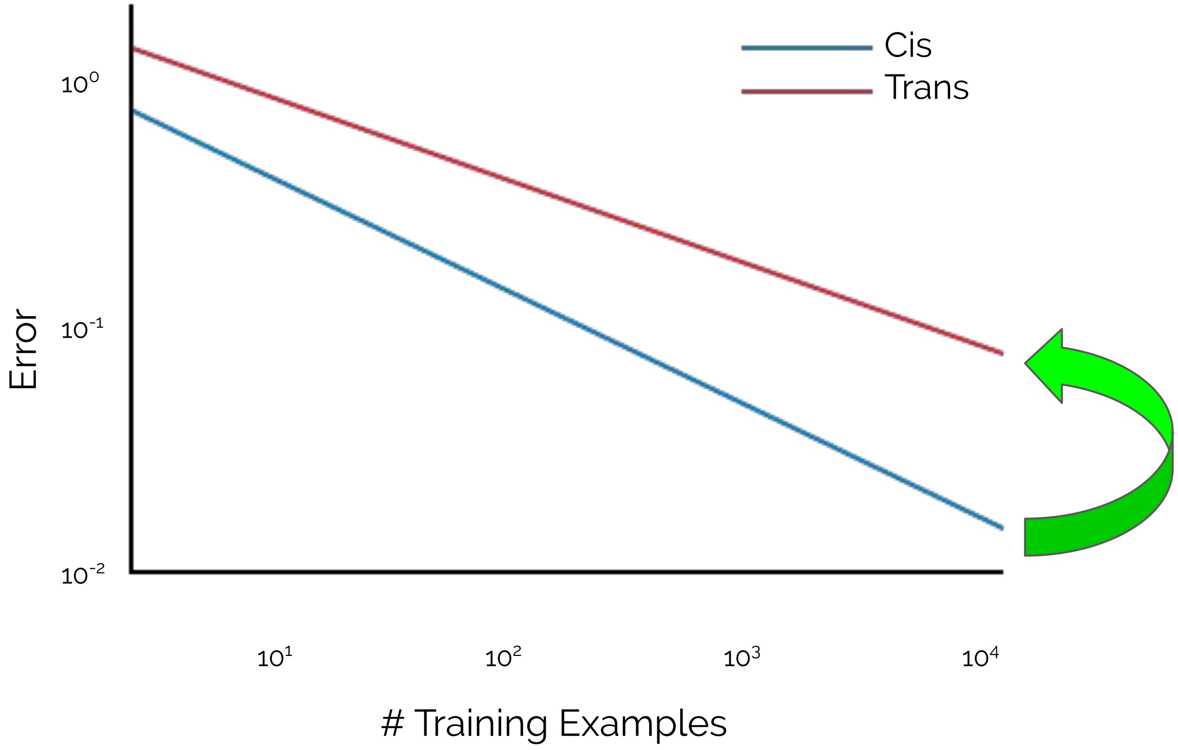

Recognition in Terra Incognita, Beery et al., ECCV 2018

Rare classes are still hard

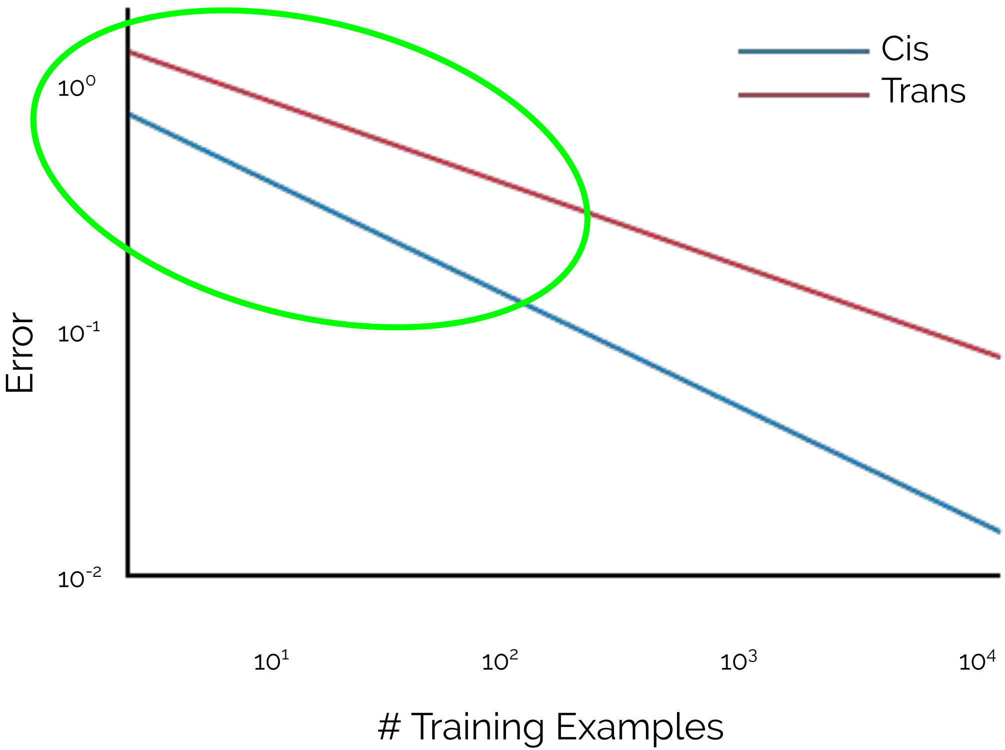

Recognition in Terra Incognita, Beery et al., ECCV 2018

Class-agnostic detectors generalize best

MegaDetector

Microsoft AI for Earth

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

Sorted 4.8 million images in ~2.75 days

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

How do experts label images like this?

Let’s focus on one potential object.

From this image alone, it's impossible to tell if this is foreground or background, let alone what class it is.

Humans look for context in other images from the same camera location.

They often look at many images, spread across a large time horizon.

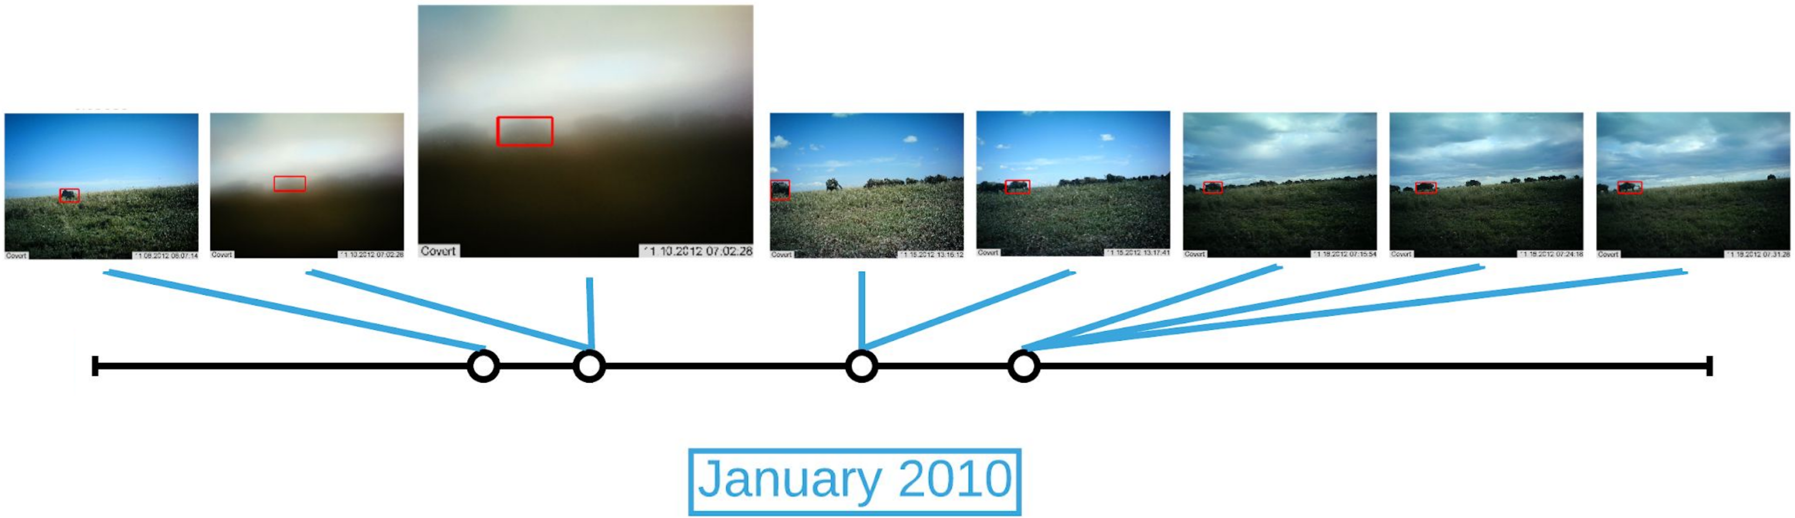

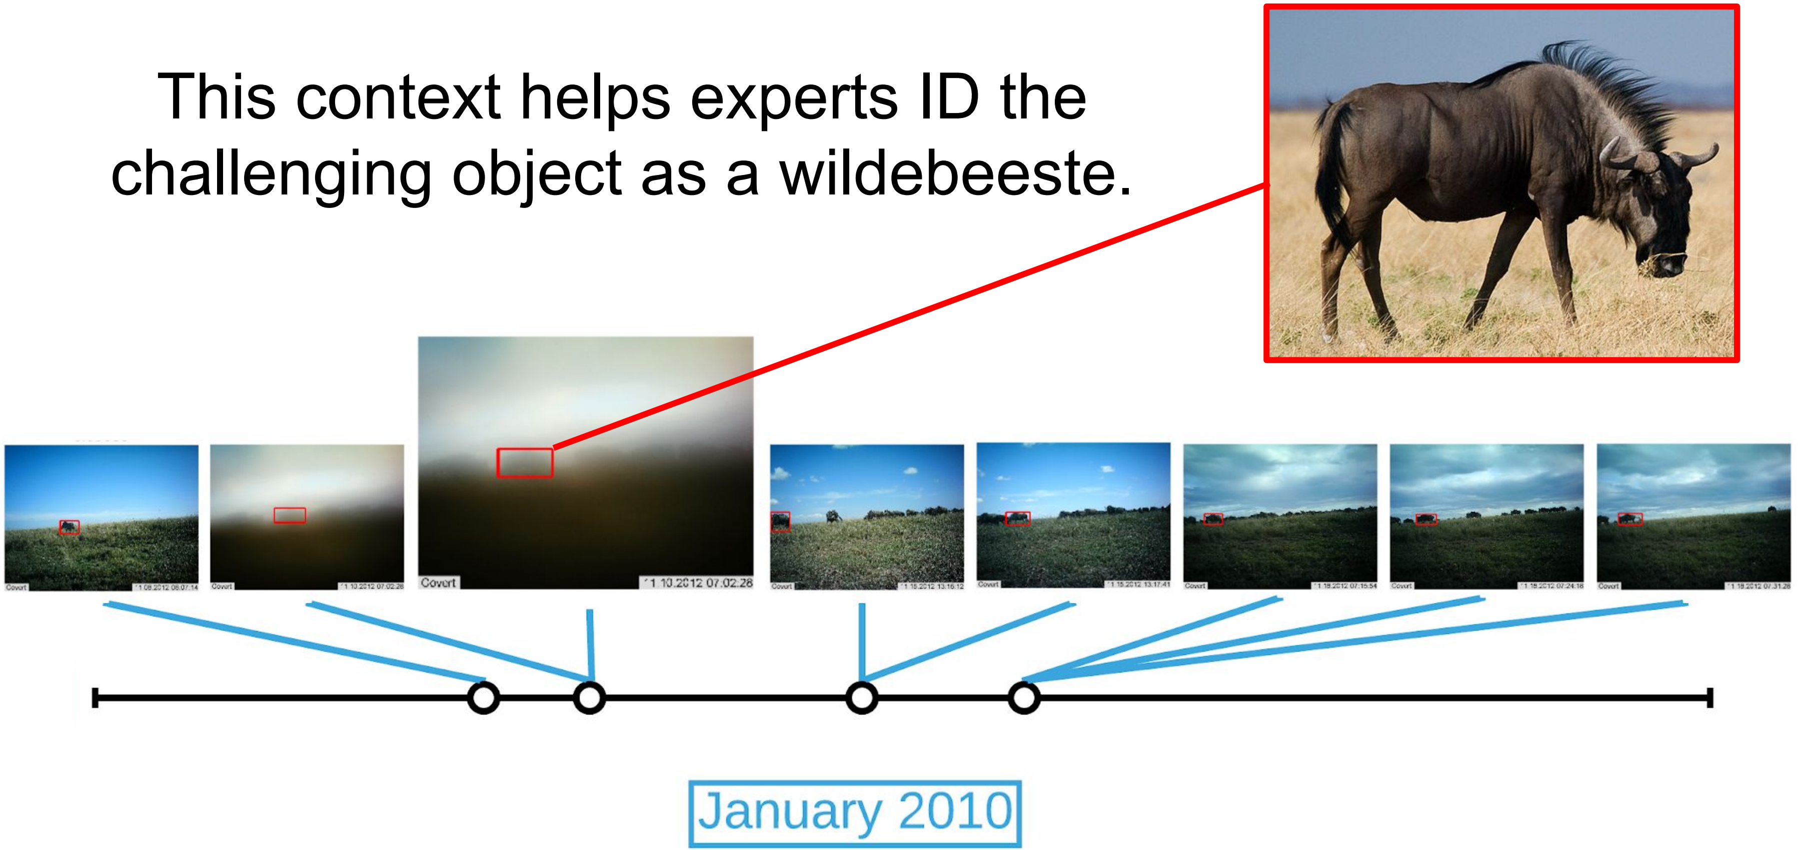
This context helps experts ID the
challenging object as a wildebeeste.
Covert
January 2010

Can we use temporal context over long time horizons, to improve detection and categorization for static cameras?

Context R-CNN:
Long Term Temporal Context for Per-Camera Object Detection
CVPR 2020
Beery, Wu, Rathod, Votel, Huang

Contextual memory strategy

- Extract features offline
- Reduce feature size
- Curate features
- Maintain spatiotemporal information

Use attention to incorporate context

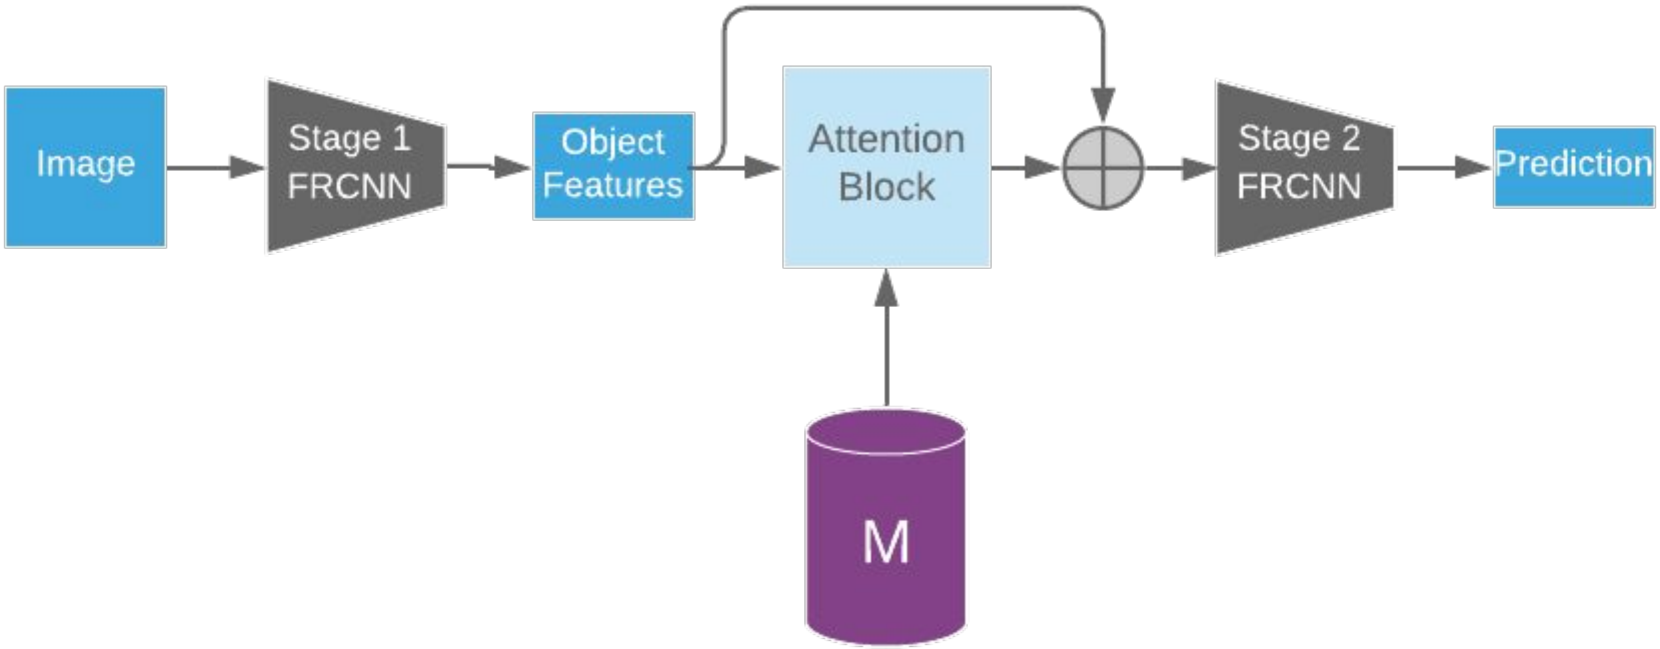

Context is incorporated based on relevance

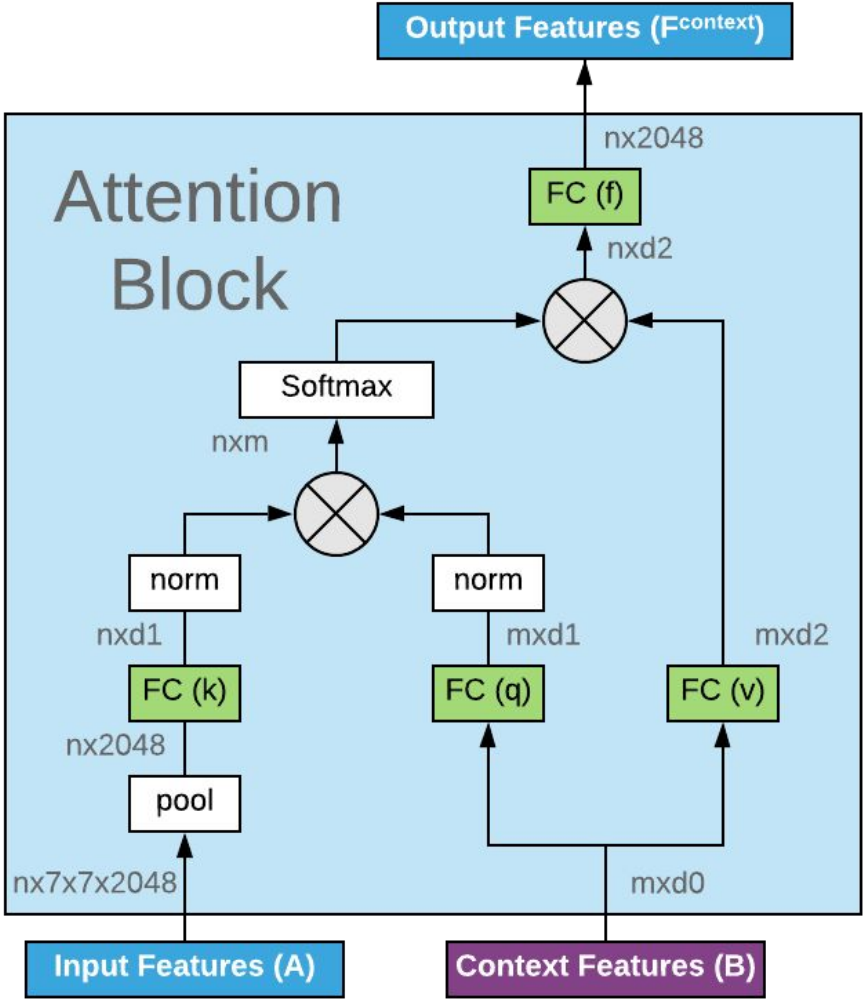

Datasets

- **Snapshot Serengeti (SS):** 225 cameras, 3.4M images, 48 classes, Eastern African game preserve
- **Caltech Camera Traps (CCT):** 140 cameras, 243K images, 18 classes, American Southwestern urban wildlife
- **CityCam (CC):** 17 cameras, 60K images, 10 vehicle classes, traffic cameras from NYC

Results

SS: Snapshot Serengeti
CCT: Caltech Camera Traps
CC: CityCam

	SS		CCT		CC	
Model	mAP	AR	mAP	AR	mAP	AR
Single Frame	37.9	46.5	56.8	53.8	38.1	28.2
Context R-CNN	**55.9**	**58.3**	**76.3**	**62.3**	**42.6**	**30.2**

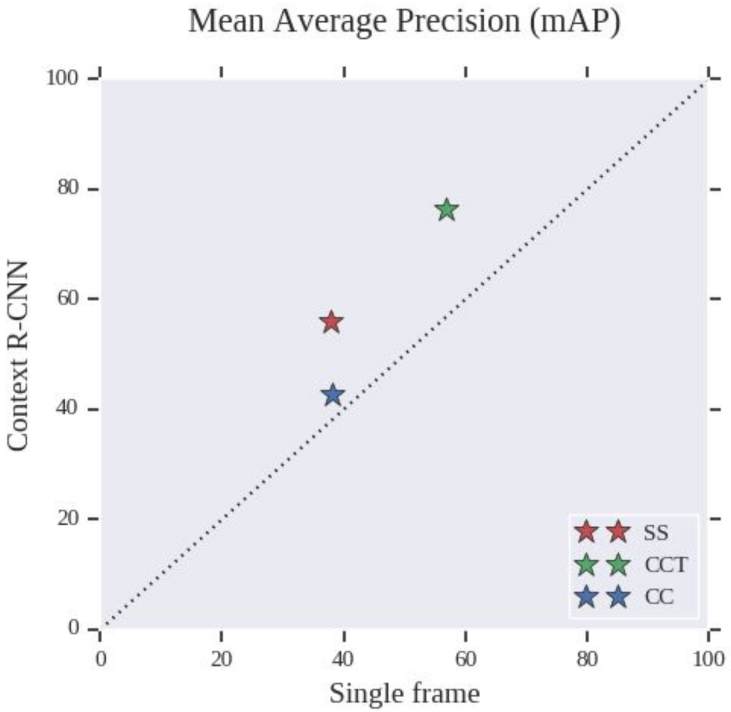

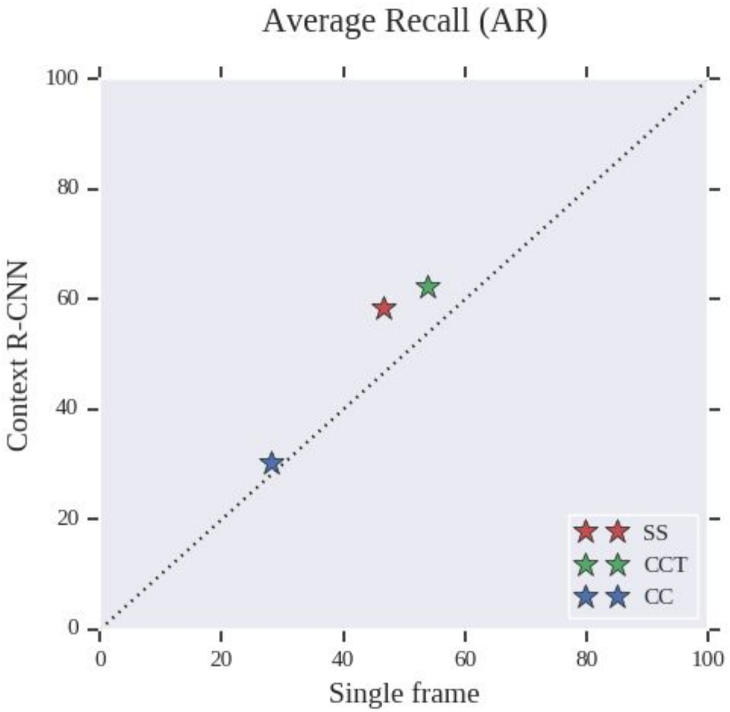

mAP improves for all classes (shown on SS*)

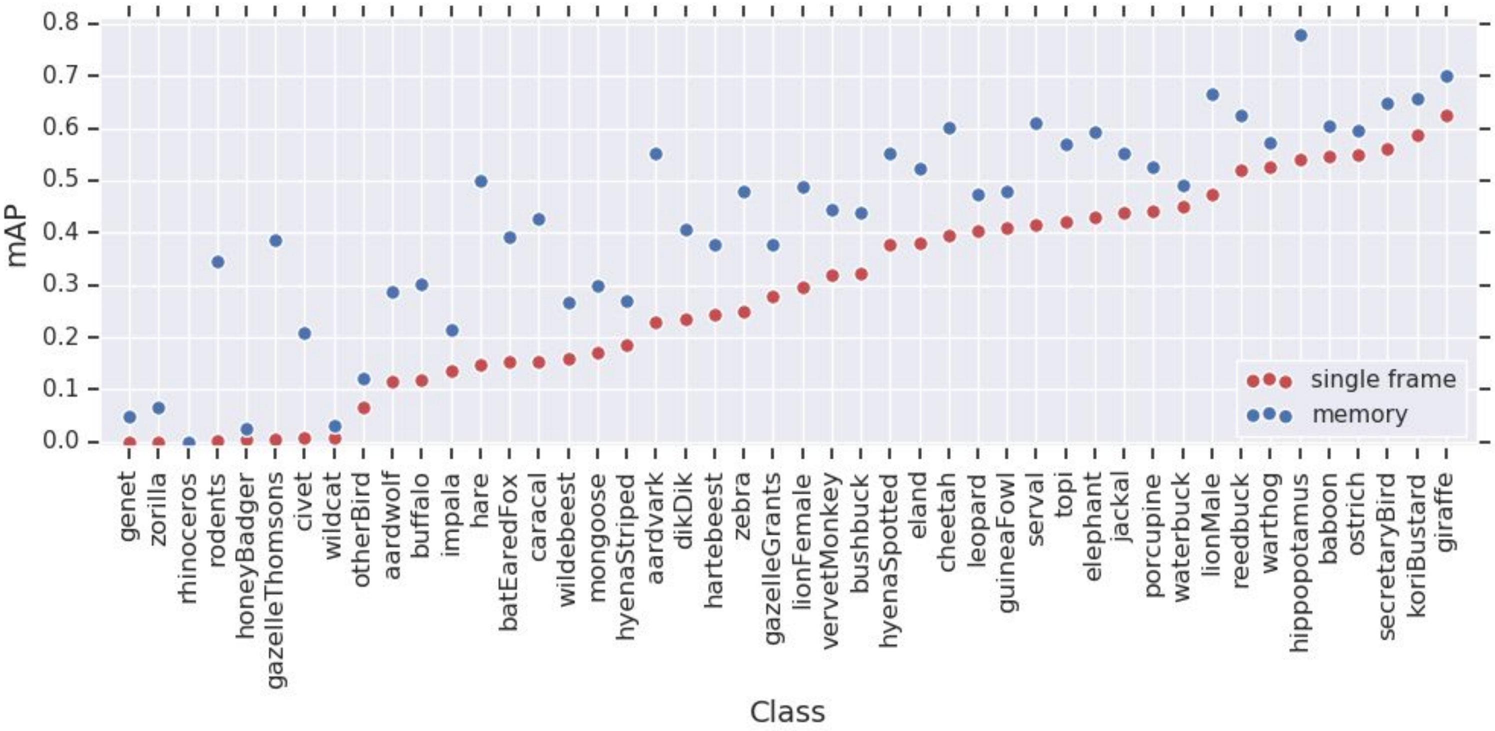

*See Supplementary Material for similar results on other datasets

Improves predominantly on challenging cases

(a) Object moving out of frame.

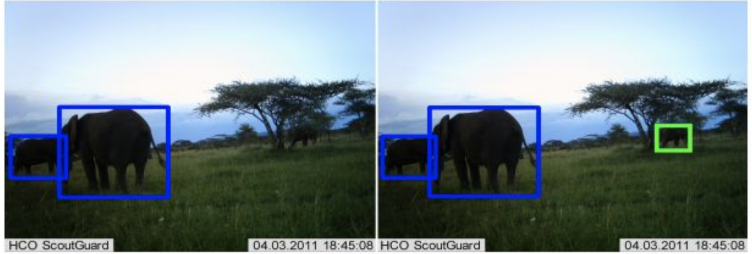

(b) Object highly occluded.

(c) Object far from camera.

(d) Objects poorly lit.

(e) Background distractor.

Correctly labels objects in challenging images

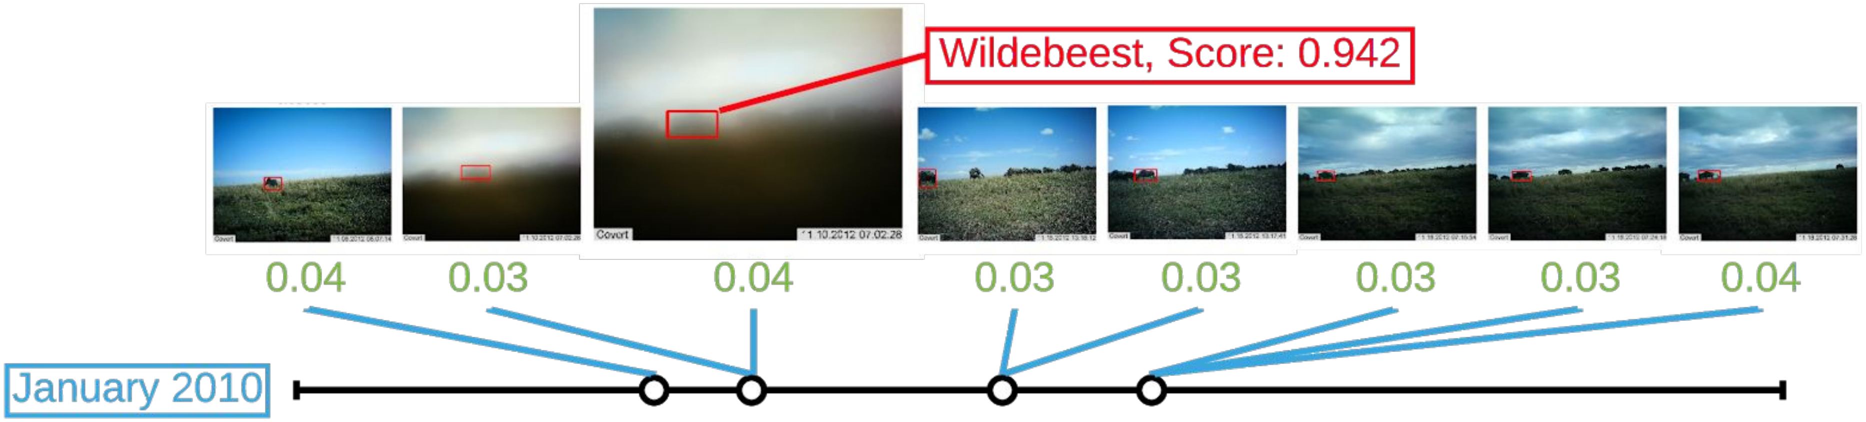

Able to categorize wildebeest through severe fog. The green scores are the corresponding contextual attention weights for each boxed feature.

Attention is temporally adaptive to relevance

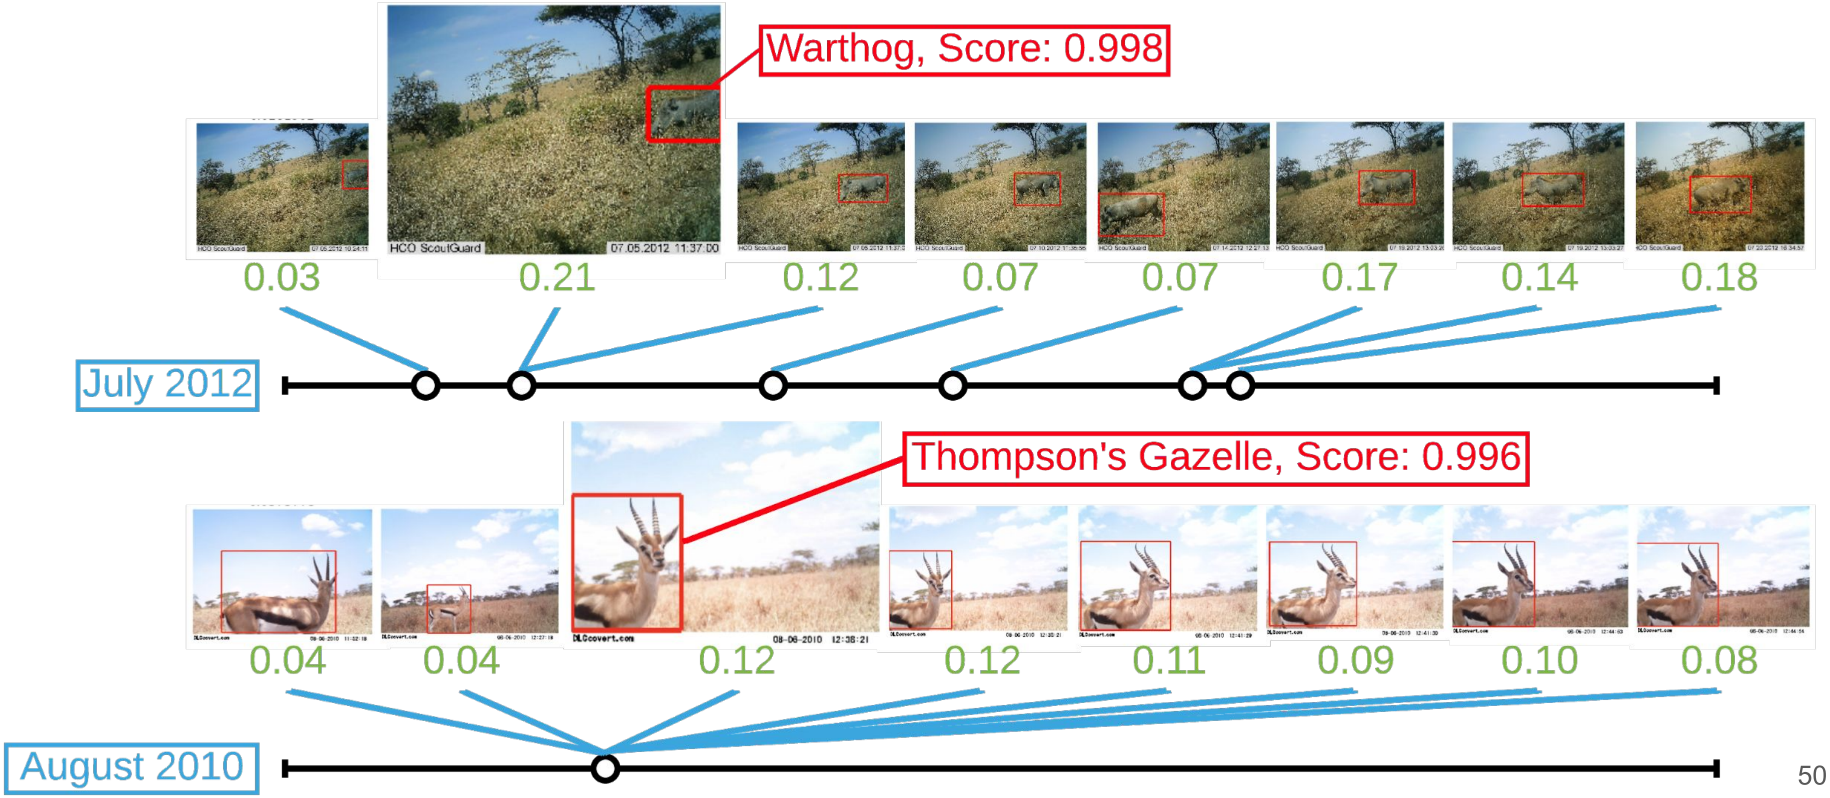

Bigger (memory) is better

SS	mAP	AR
One minute	50.3	51.4
One hour	52.1	52.5
One day	52.5	52.9
One week	54.1	53.2
One month	**55.6**	**57.5**

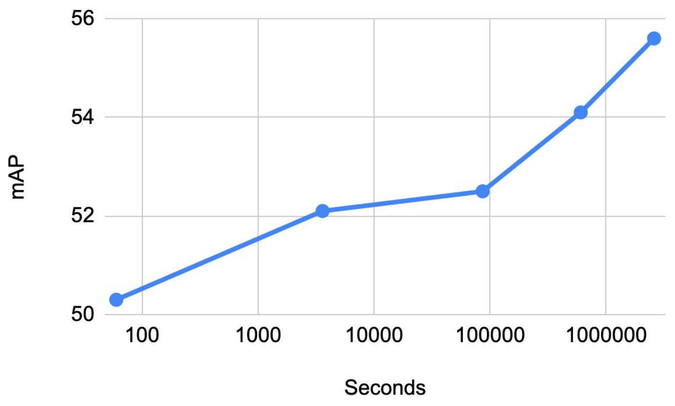

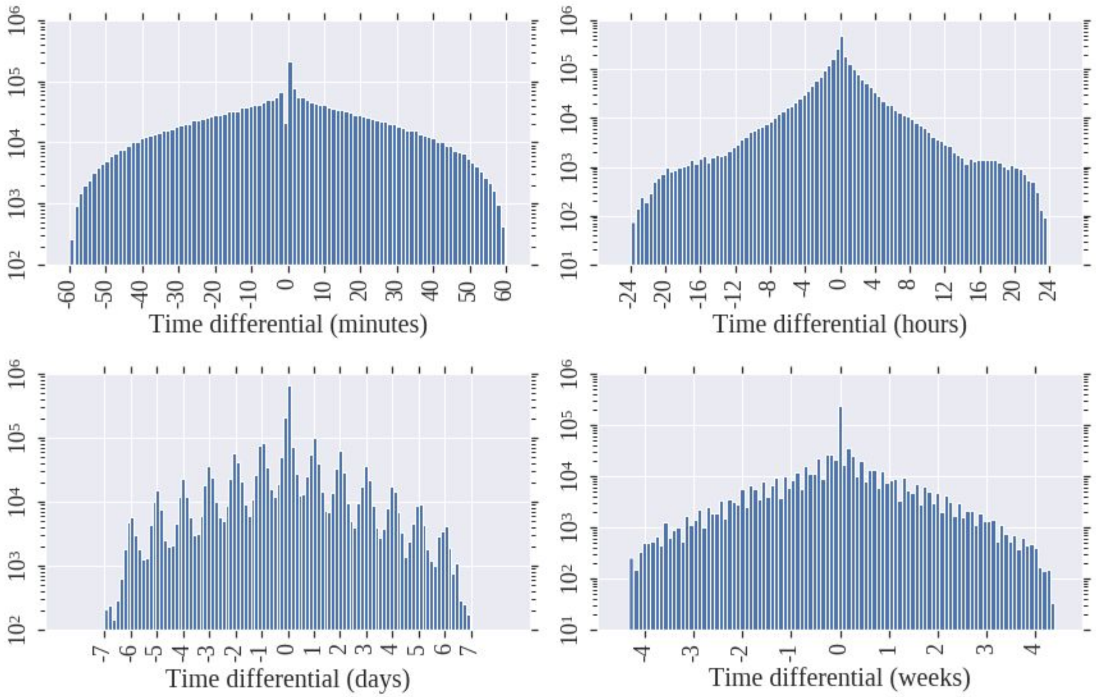

Histogram of time differentials from the highest-scoring object in the keyframe to the attended frames for varied time horizons.

Bigger (memory) is better

SS	mAP	AR
One minute	50.3	51.4
One hour	52.1	52.5
One day	52.5	52.9
One week	54.1	53.2
One month	**55.6**	**57.5**

day/night periodicity!

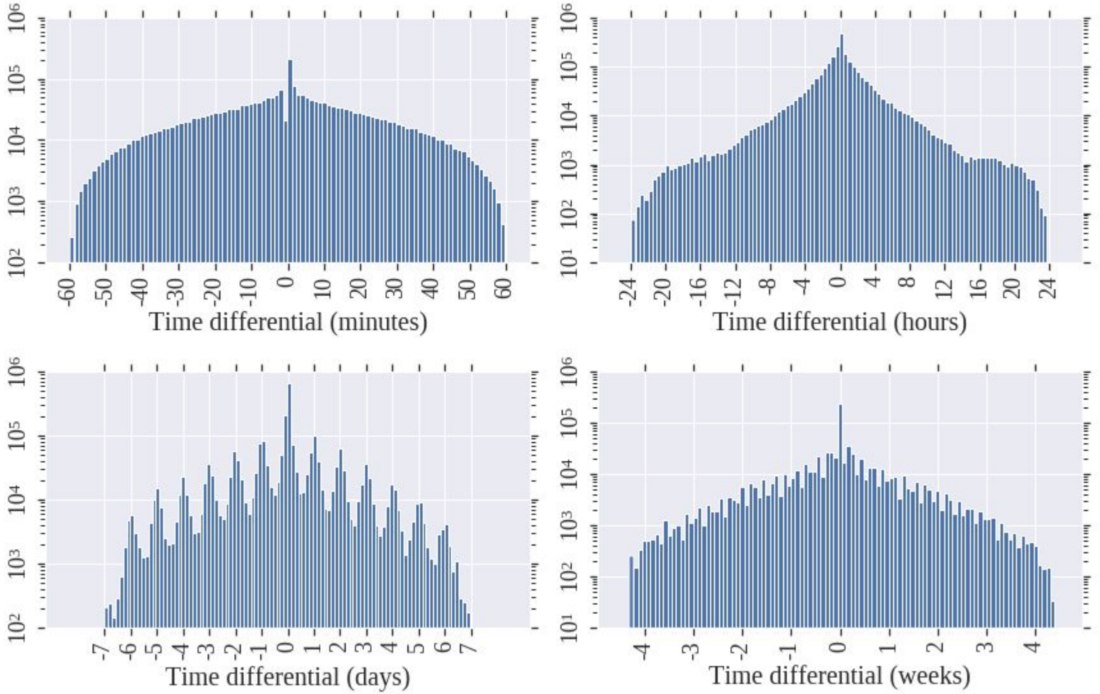

Histogram of time differentials from the highest-scoring object in the keyframe to the attended frames for varied time horizons.

Background classes are learned without supervision

Adding features from empty images reduces false positives

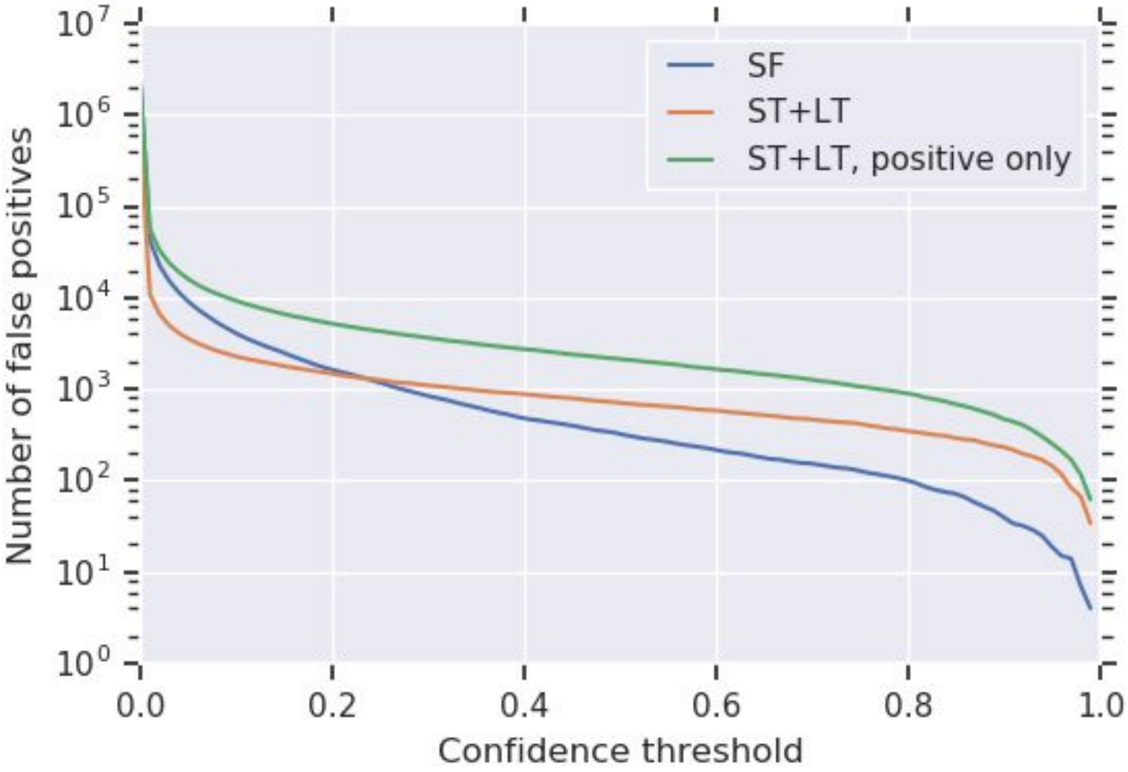

Of the 100 most confident “false positives” returned by our ST+LT model, 97/100 were in fact mis-annotated.

Can we leverage camera trap data to monitor populations via re-ID?

The Great Grevy's Rally: an animal re-ID success story

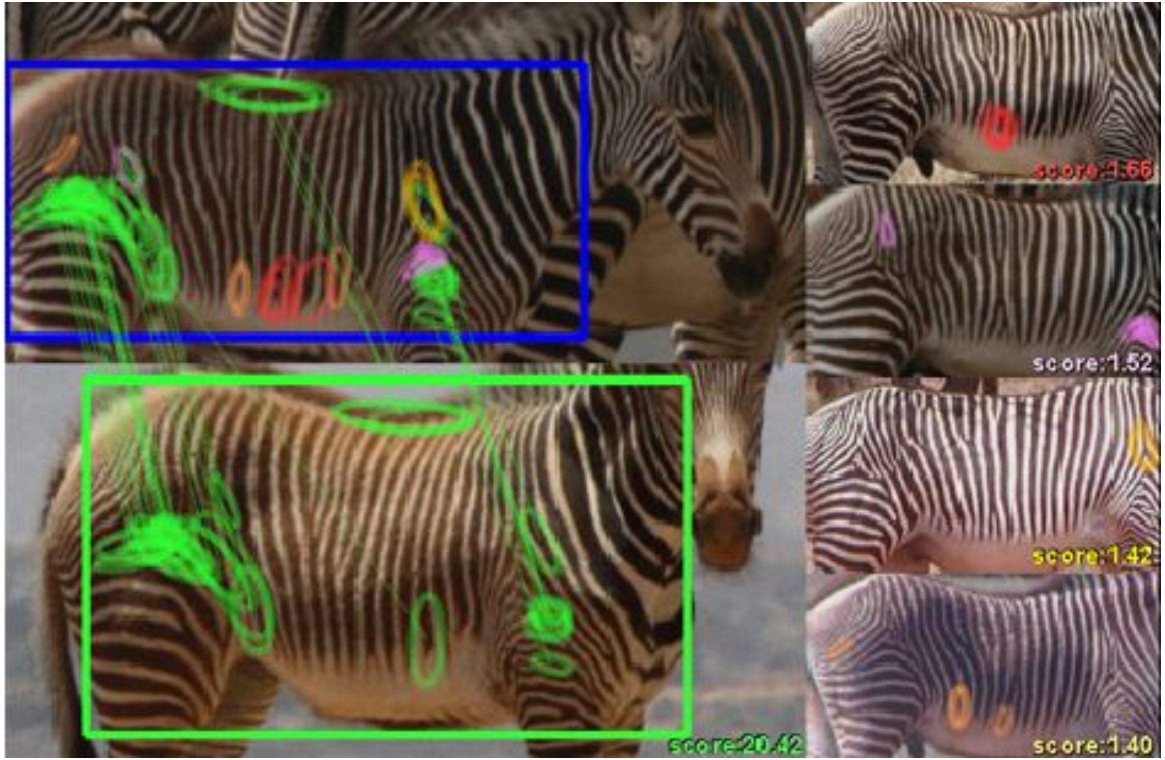

WILD ME

THE GREAT GREVY'S RALLY

WILDBOOK

Camera Trap Data Collection at GGR 2020

- Mpala Research Center in Laikipia
- 100 camera traps with 3 spatial sampling strategies
- We want to compare capture-recapture using the camera trap data and using the citizen science data

Placement Strategies

12 "magnet" cameras

47 roadway cameras

10 random roadway cameras

21 random grid cameras

5 paired timelapse/video cameras at magnet sites

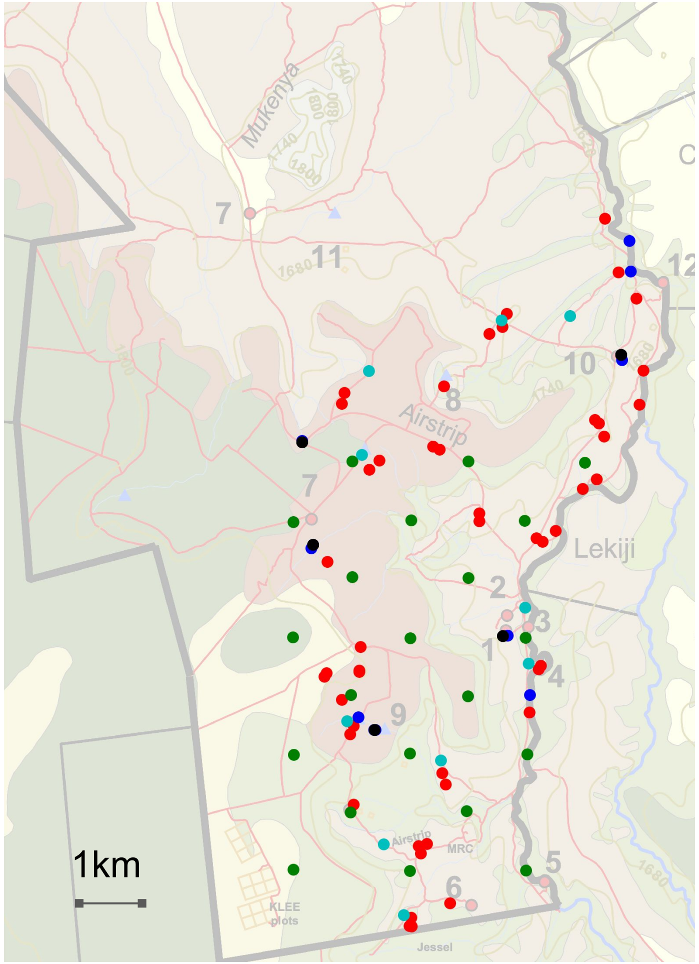

Good news! Some images get matched right away.

We have matches to nighttime data!

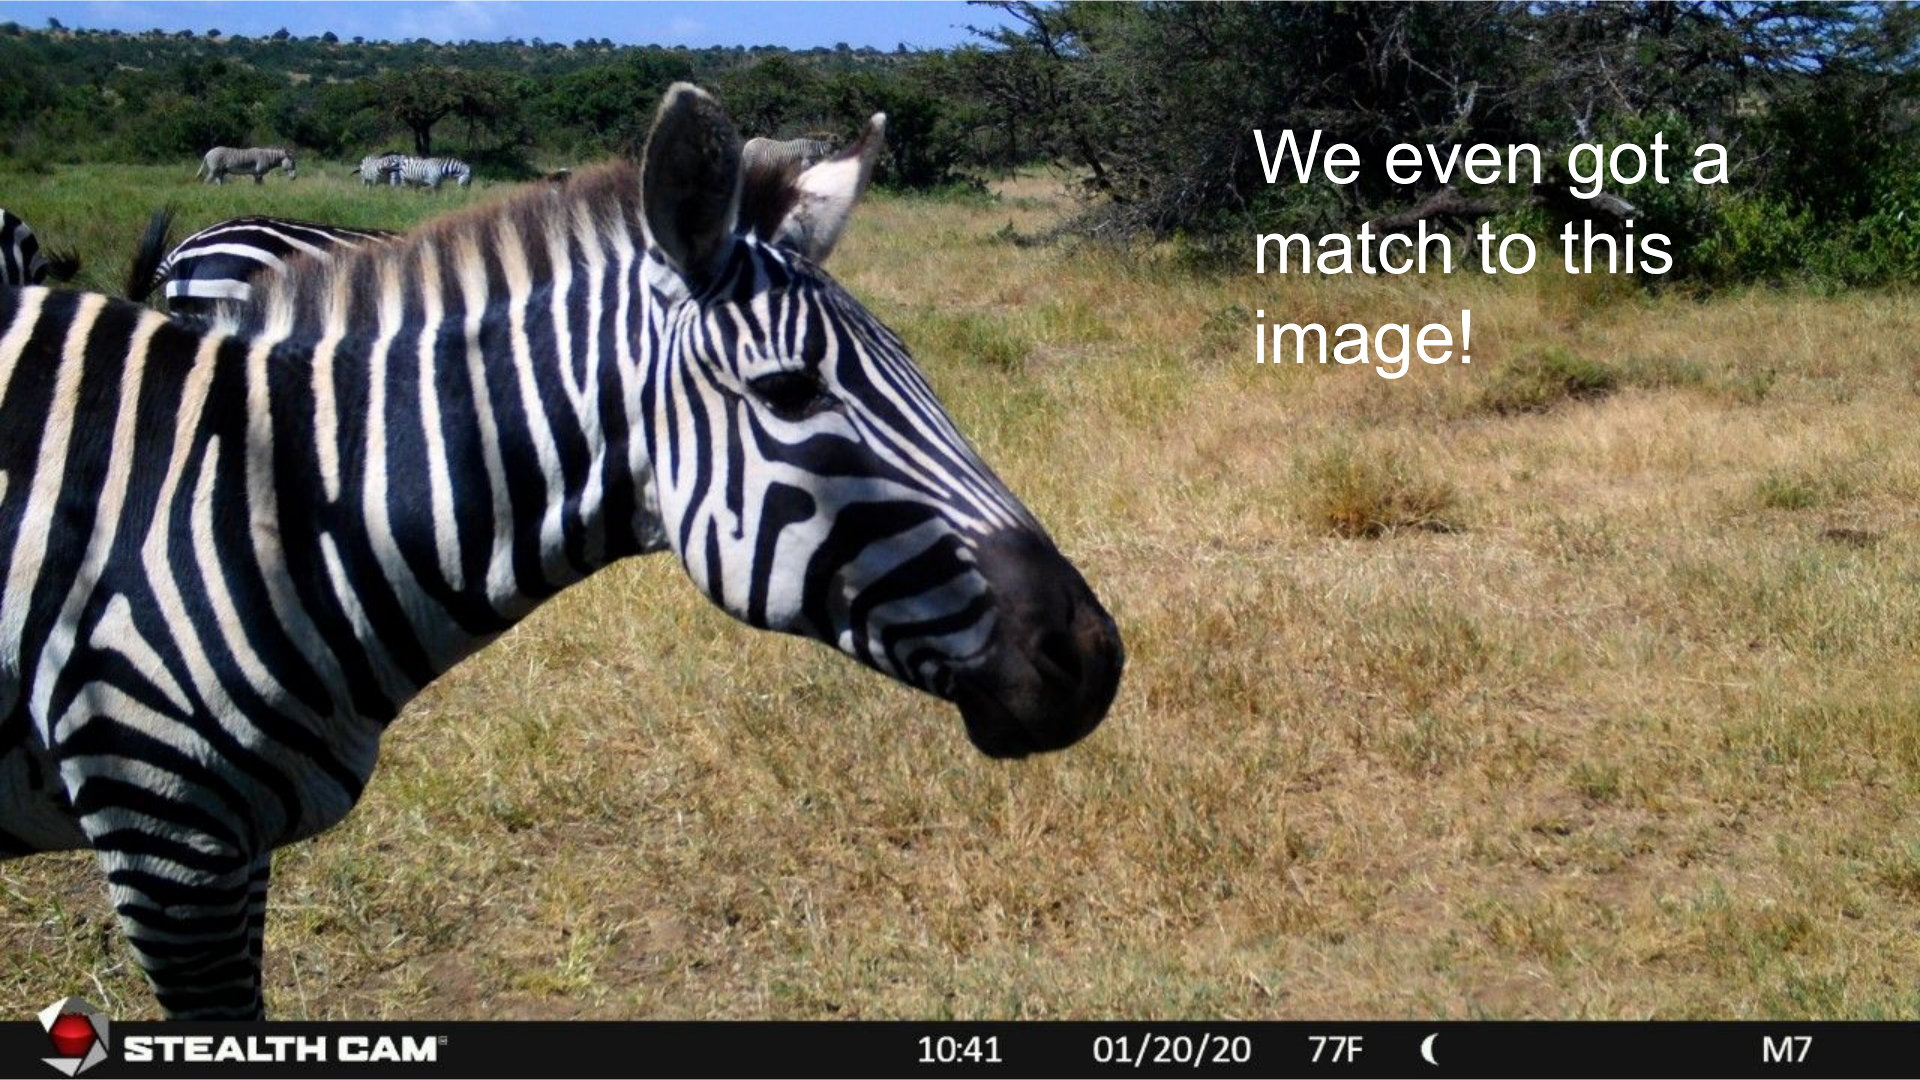
We even got a match to this image!
STEALTH CAM
10:41
01/20/20
77F
M7

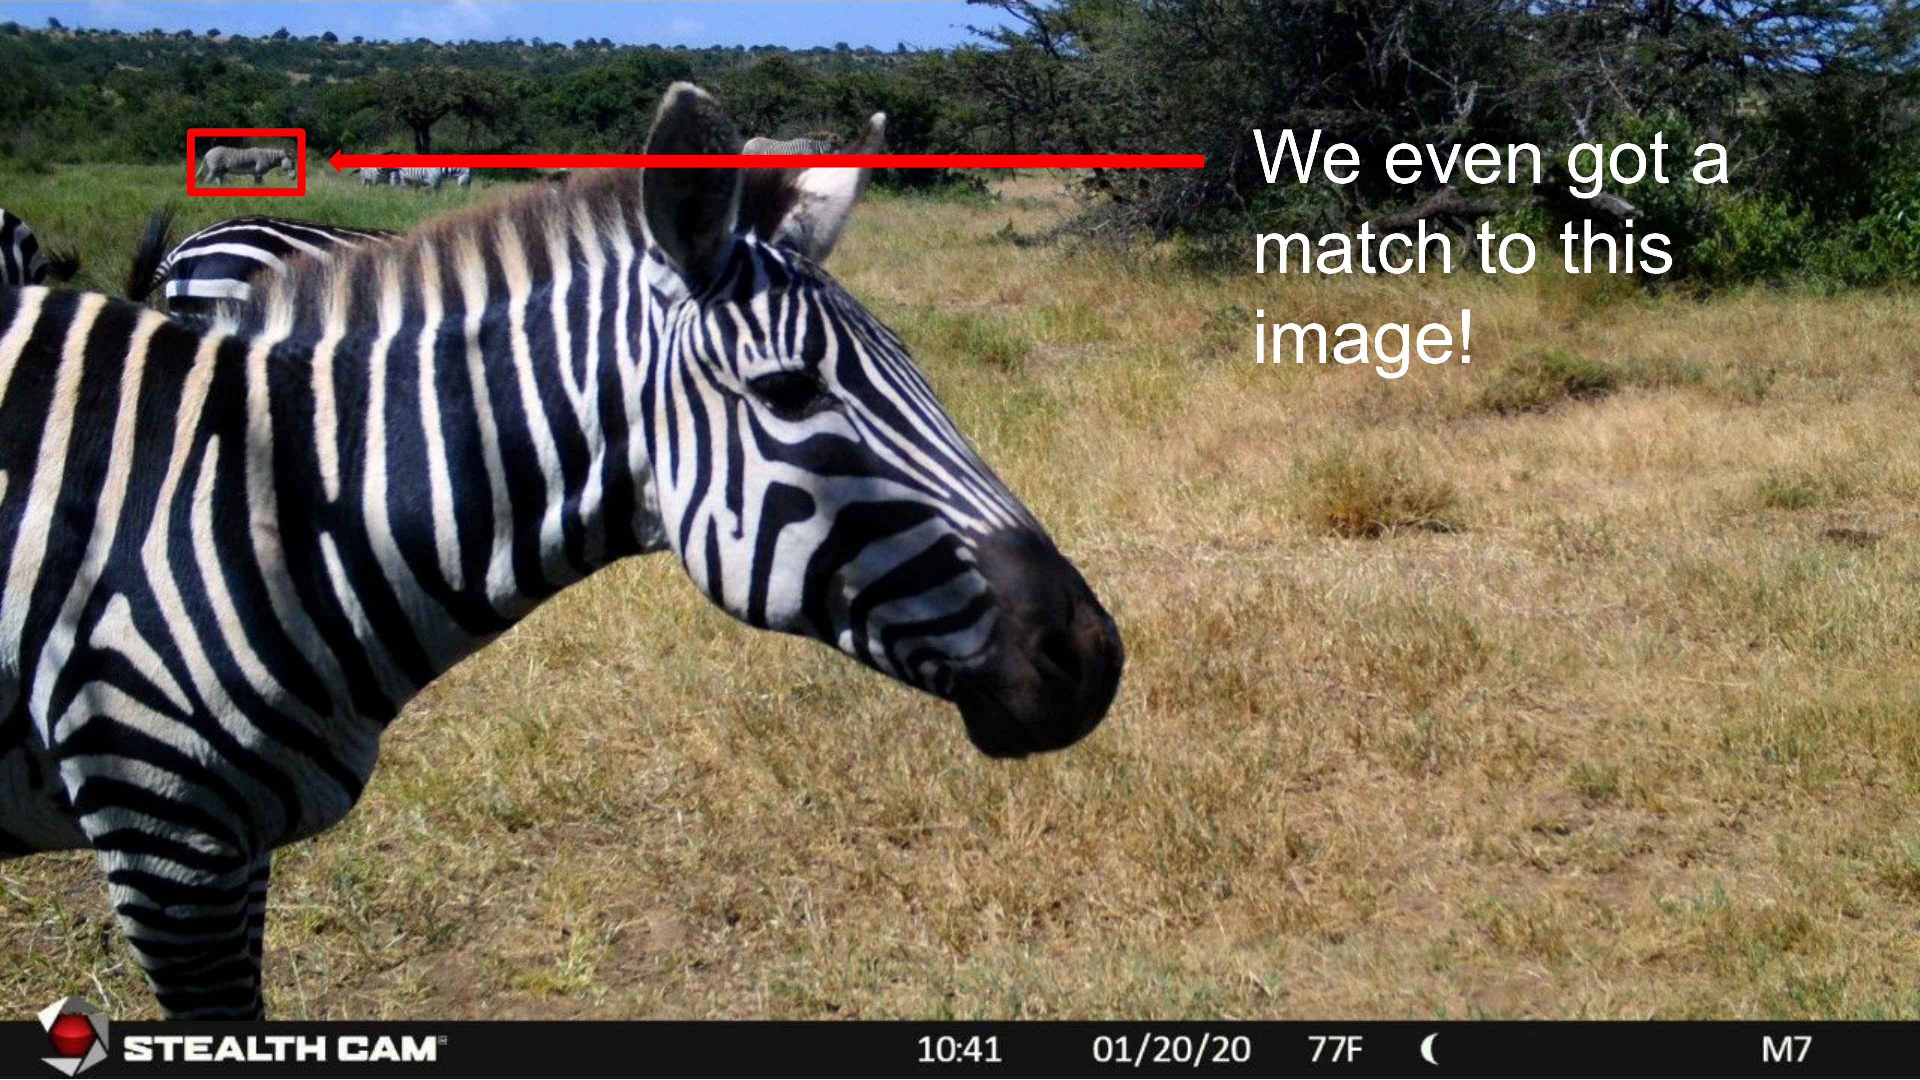
We even got a match to this image!
STEALTH CAM
10:41
01/20/20
77F
M7

Pretty cool

But we've seen a lot of zebras that would currently be unidentifiable….

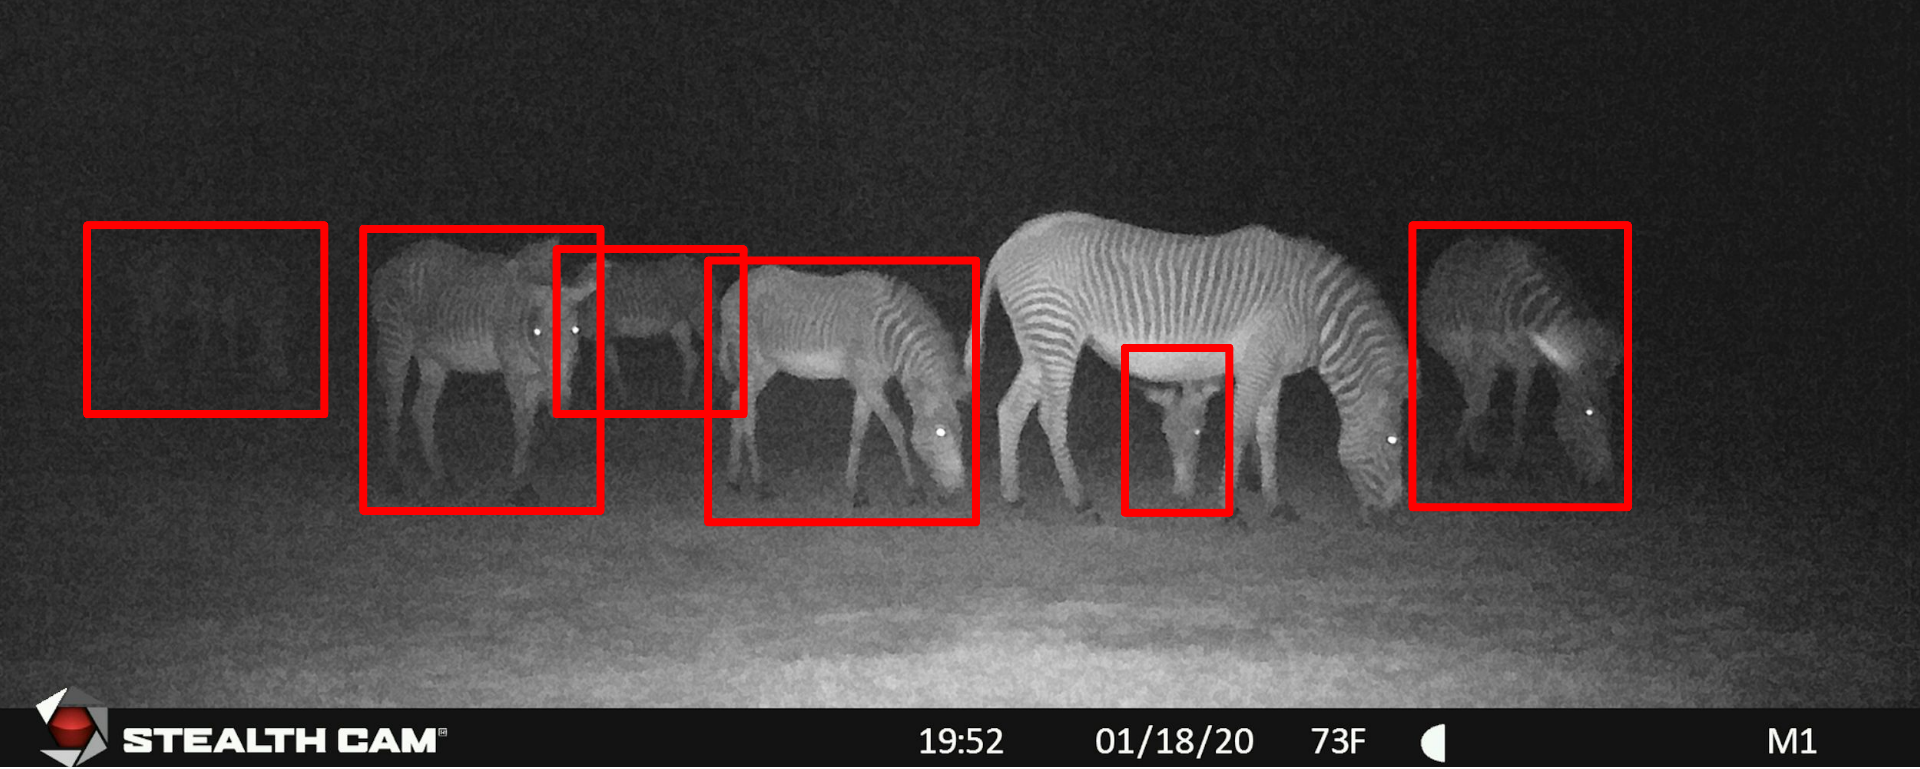

Extend context-based approach to re-ID?

Open questions

- Species modeling
  - Multiple spatial sampling strategies
- Combining data streams
  - iNaturalist/eBird
  - Satellite imagery
  - Aerial drone imagery
  - Social media data

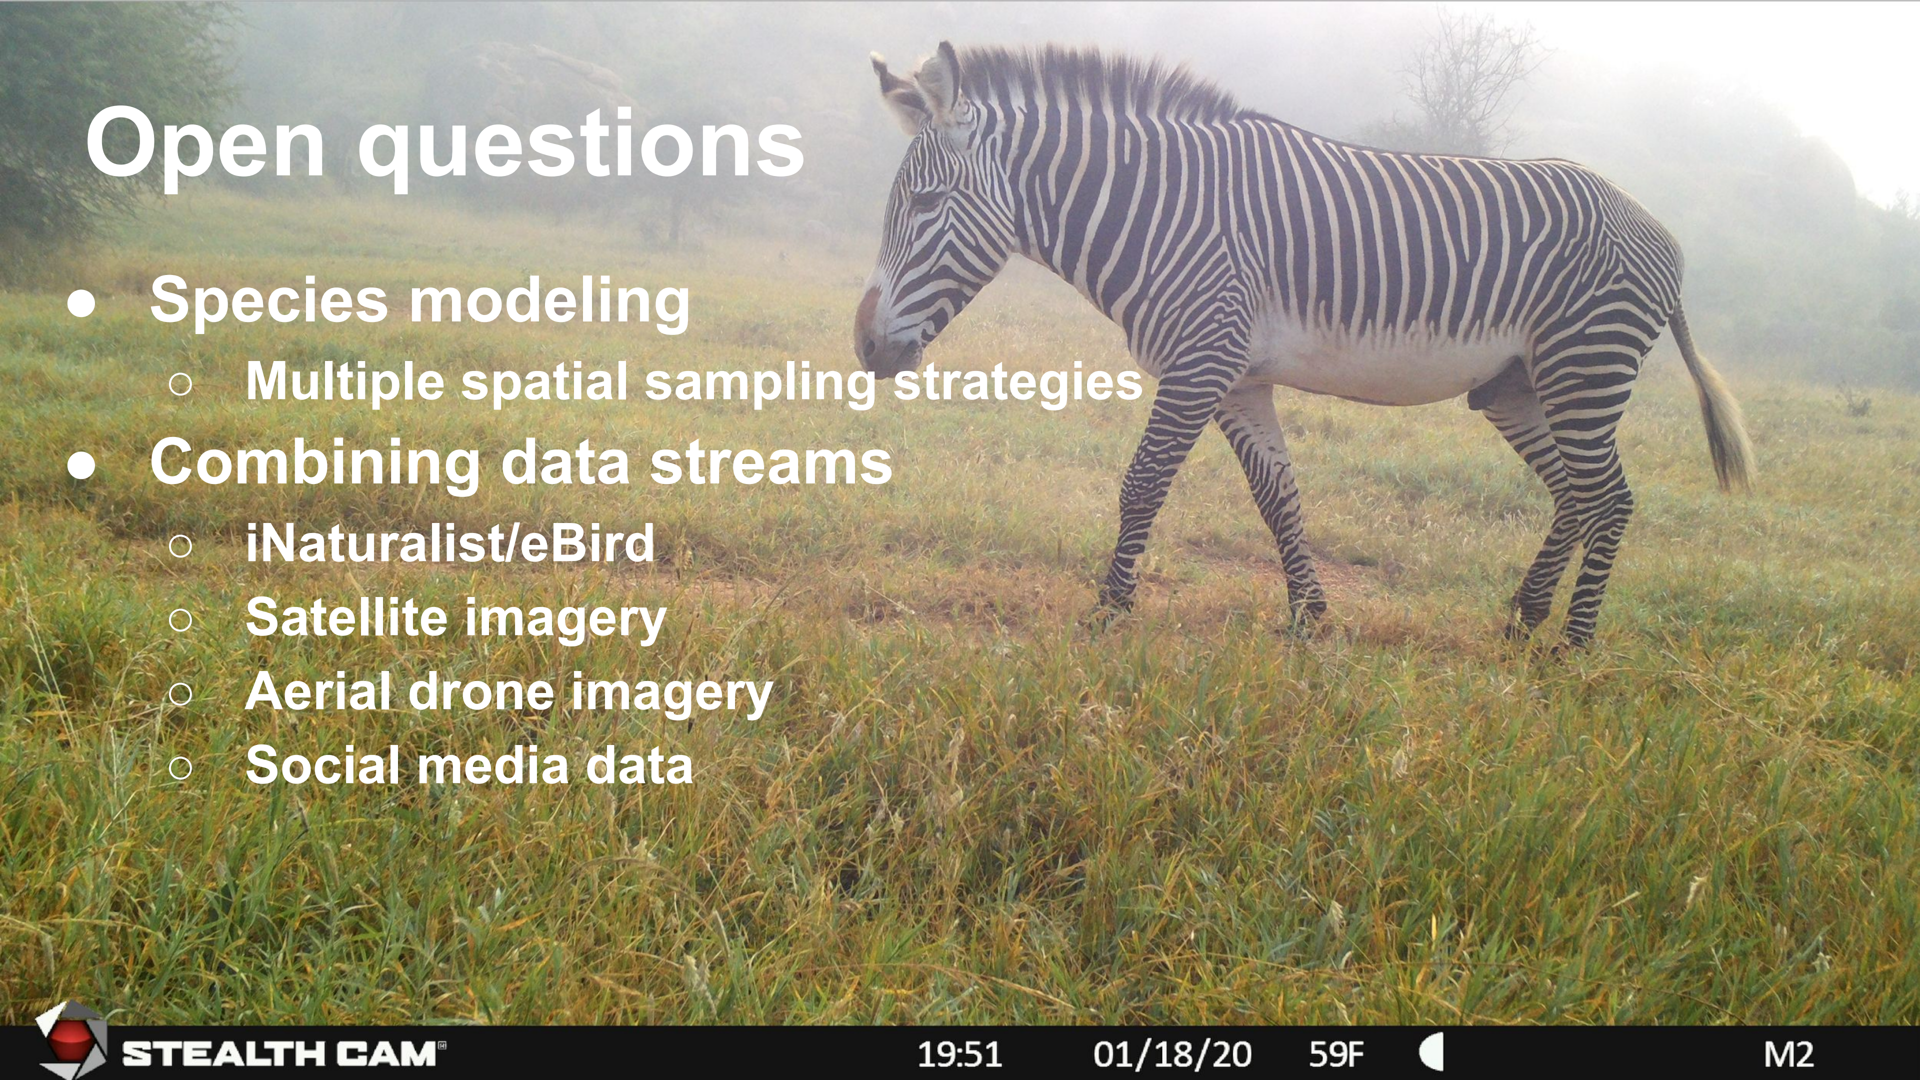

Biodiversity-focused competitions

Global camera traps (WCS) + RS
Data Release: March

https://www.kaggle.com/c/iwildcam-2020-fgvc7

GeoLifeCLEF 2020

Location-Based
Species
Recommendation

2M Species Observations + RS + LC + Covariates
Data Release: March

https://www.imageclef.org/GeoLifeCLEF2020

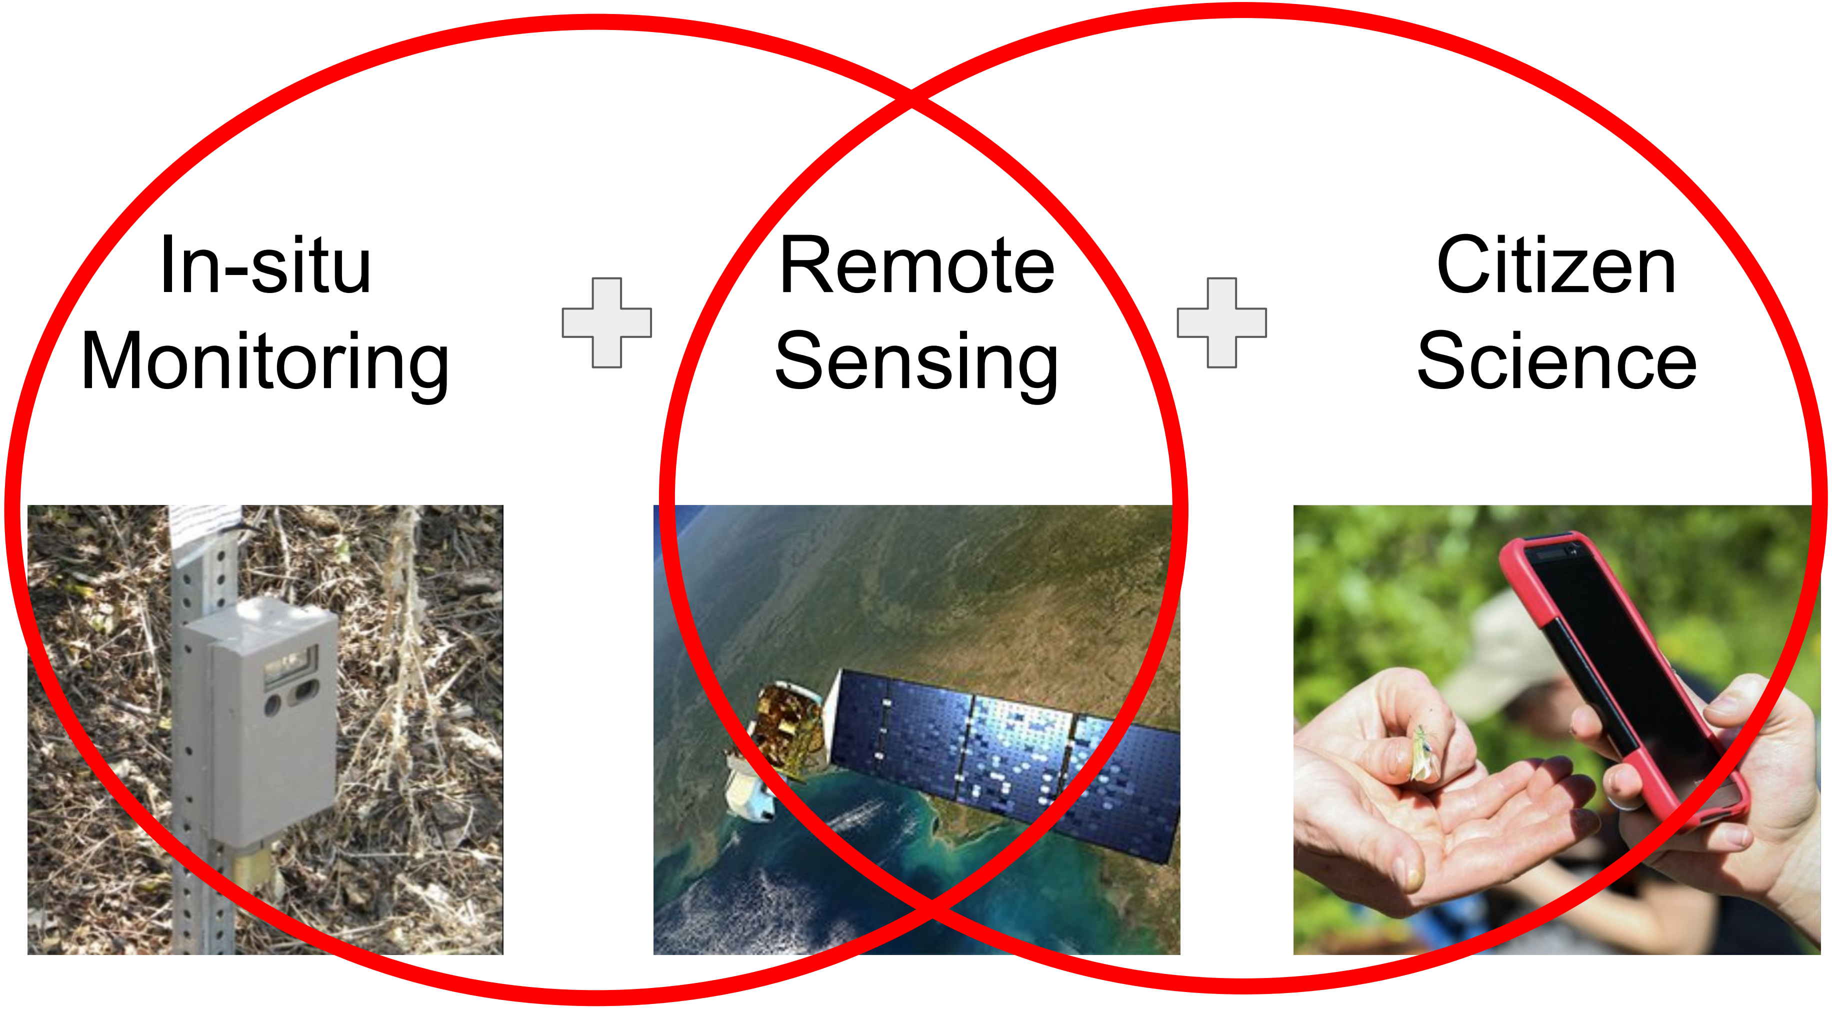
In-situ Monitoring
Remote Sensing
Citizen Science

Big challenges

- Long-tailed distribution
- Sparse, low-quality data
- Global generalization

Interested? Join our slack channel by emailing aiforconservation@gmail.com

Acknowledgements

Microsoft

AI for Earth

Caltech
Vision Lab

WILD ME

CALIFORNIA INSTITUTE OF TECHNOLOGY 1891

NSF

NATIONAL PARK SERVICE

IDAHO FISH & GAME

USGS
science for a changing world

MPALA
THE MPALA RESEARCH CENTRE &
THE MPALA WILDLIFE FOUNDATION

LILA BC

Labeled Information Library of Alexandria: Biology and Conservation