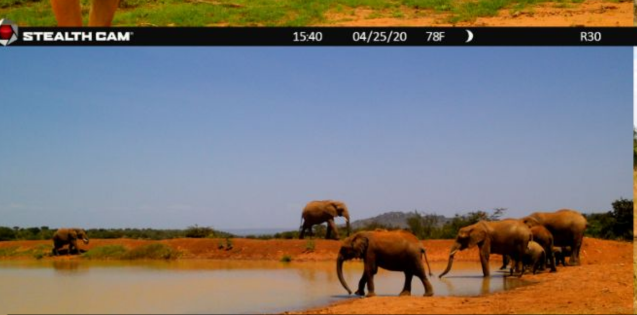

Open Challenges in Generalizable Computer Vision for Ecology



Sara Beery | CamTrap Ecology Meets AI | 9-14-22



Biodiversity is in decline globally


BBC | Sign in | Home | News | Sport | Reel | Worklife | Travel

NEWS

Home | US Election | Coronavirus | Video | World | US & Canada | UK | Business | Tech | Science | Stories

Science

Wildlife in 'catastrophic decline' due to human destruction, scientists warn


WWF

OUR WORK PEOPLE PLACES WILDLIFE About How to help DONATE

PRESS RELEASES

68% Average Decline in Species Population Sizes Since 1970, Says New WWF Report

Declines in monitored populations of mammals, fish, birds, reptiles, and amphibians present a dire warning for the health of people and the planet

Biodiversity data is diverse

Seeing biodiversity: perspectives in machine learning for wildlife conservation, Tuia*, Kellenberger*, Beery*, Costelloe*, et al., Nature Comms (to appear)

Manual data processing doesn’t scale

Camera Traps

One project can collect >10M images/season

Community Scientists

>64M Species observations in iNaturalist

Aerial Surveys

One survey can generate >200TB of video

Scaling Biodiversity Monitoring for the Data Age - ACM XRDS 2021

Manual data processing doesn't scale

Camera Traps

Community Scientists

Aerial Surveys

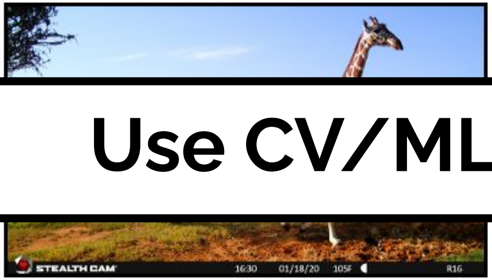

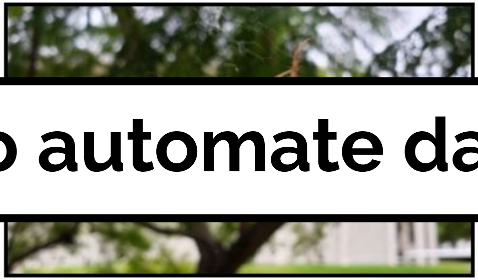

Use CV/ML to automate data processing

One project can collect >10M images/season

>64M Species observations in iNaturalist

One survey can generate >200TB of video

Scaling Biodiversity Monitoring for the Data Age - ACM XRDS 2021

Biodiversity data is noisy

Objects of interest partially observed.

Poor data quality.

Empty data.

Biodiversity data has a long tail

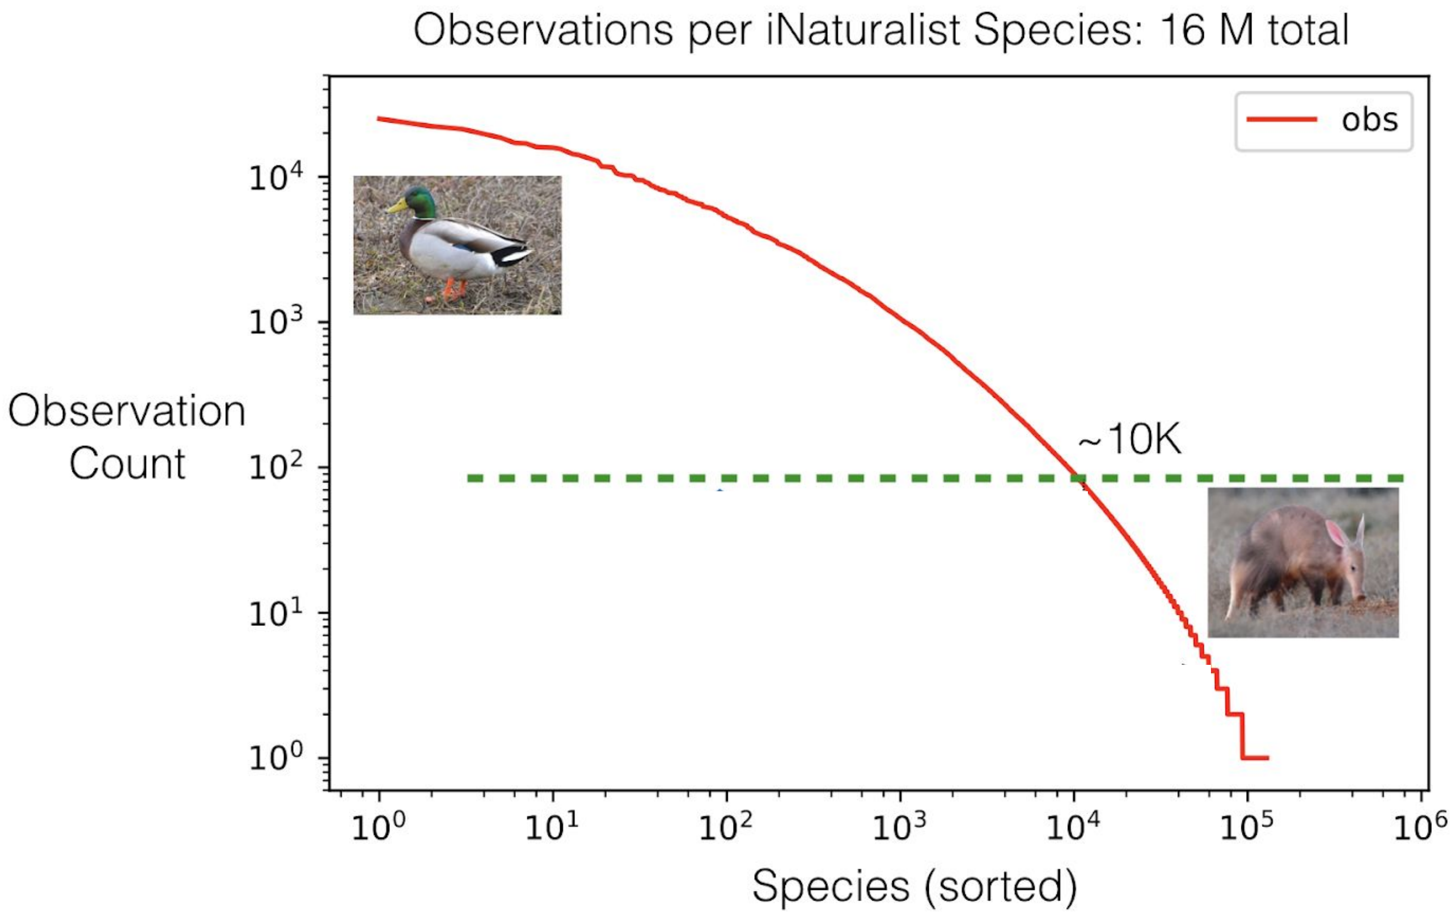

Biodiversity data is not IID

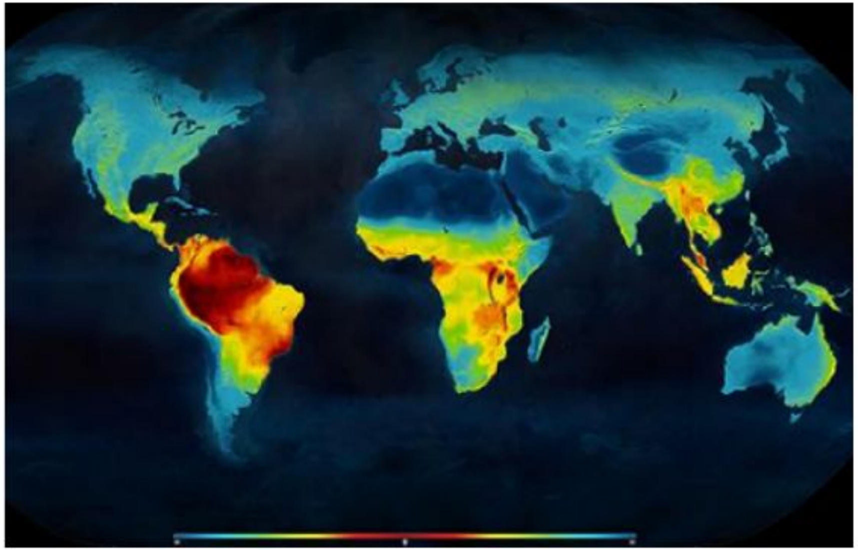

Map of global biodiversity

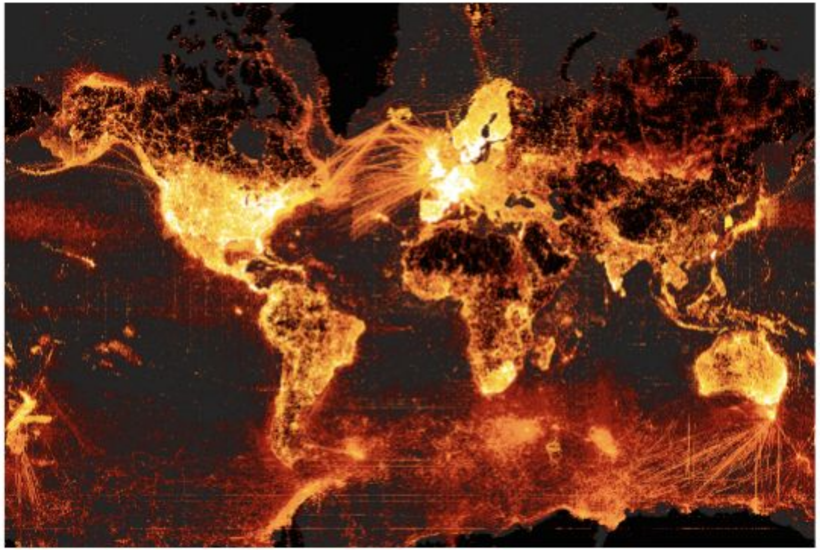

Species occurrence data in GBIF

Scaling Biodiversity Monitoring for the Data Age - ACM XRDS 2021

Distribution shifts are ubiquitous in real-world scenarios: generalization is a key bottleneck to useful, usable CV for ecology

WILDS

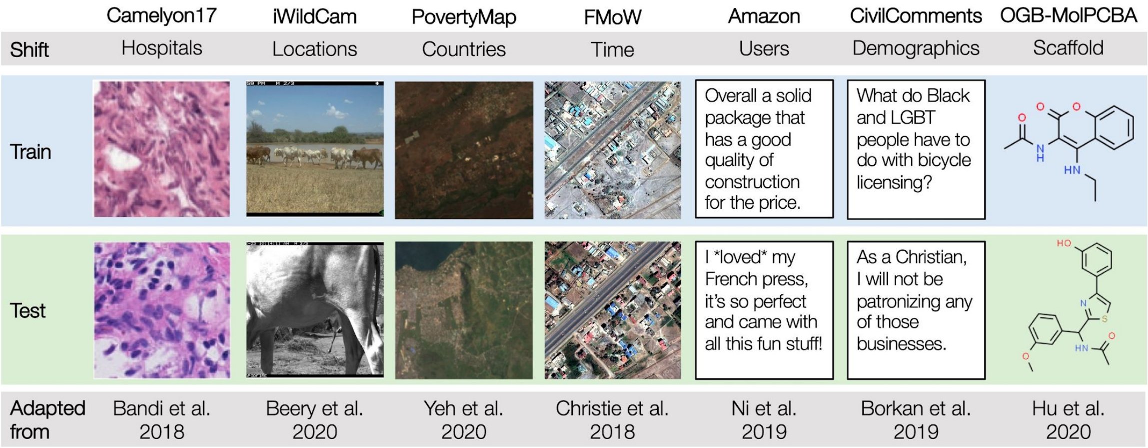

We recently released the first real-world, large-scale, cross-domain benchmarks for domain generalization

WILDS: A Benchmark of in-the-Wild Distribution Shifts, Koh, …, Beery, et al., ICLR 2021
Extending the WILDS Benchmark for Unsupervised Adaptation, Koh, …, Beery, et al., *In Submission*

https://wilds.stanford.edu/

First generalization case study:
Species detection and ID in camera traps

Wildlife Insights

Explore Data Settings

Showing **22,163,565 camera trap records** taken in the whole world between 1990-01-02 and 2021-11-16.

See filters and statistics

Each static camera has a distinctive background and class distribution

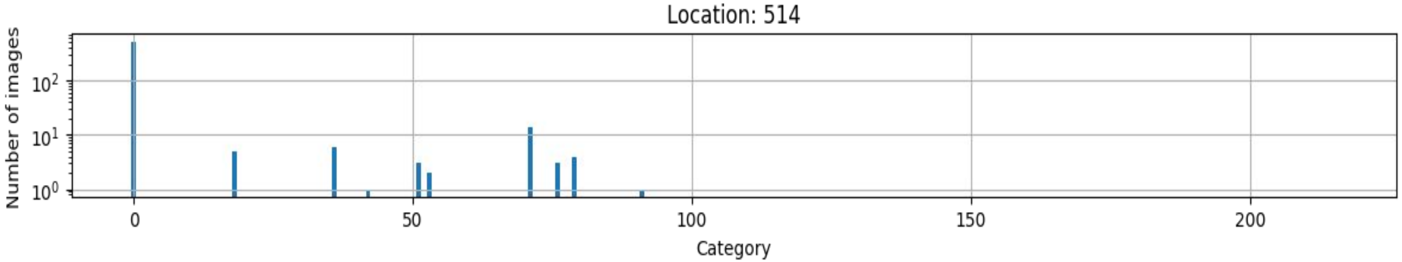

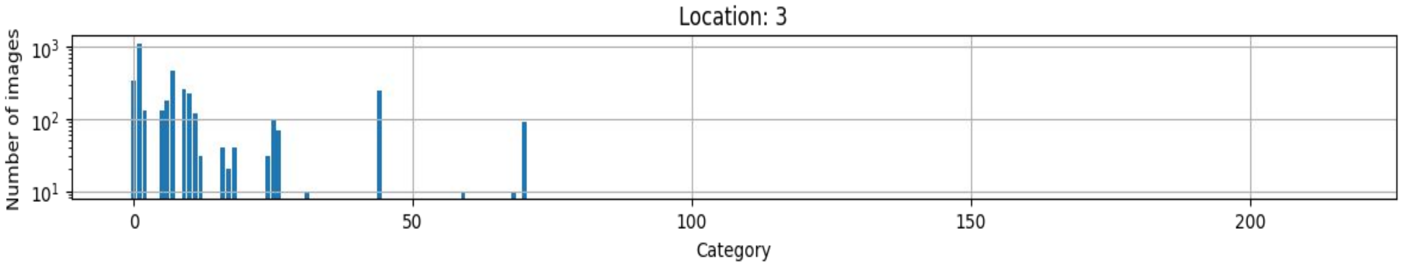

Models don't generalize

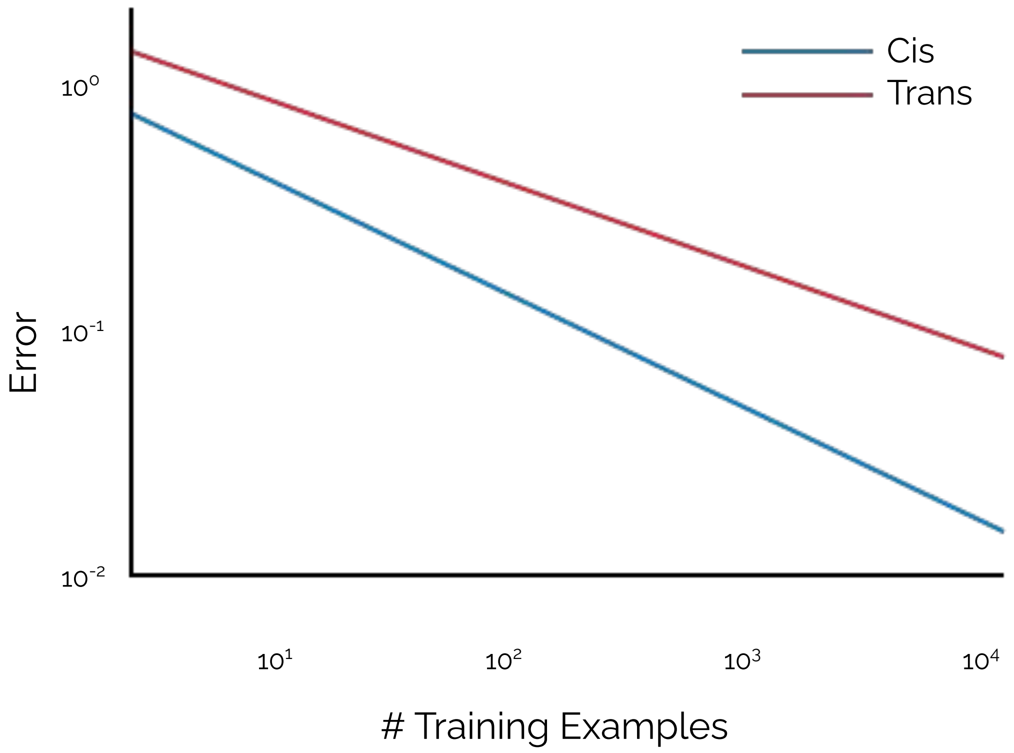

Recognition in Terra Incognita, Beery et al., ECCV 2018

Models don’t generalize

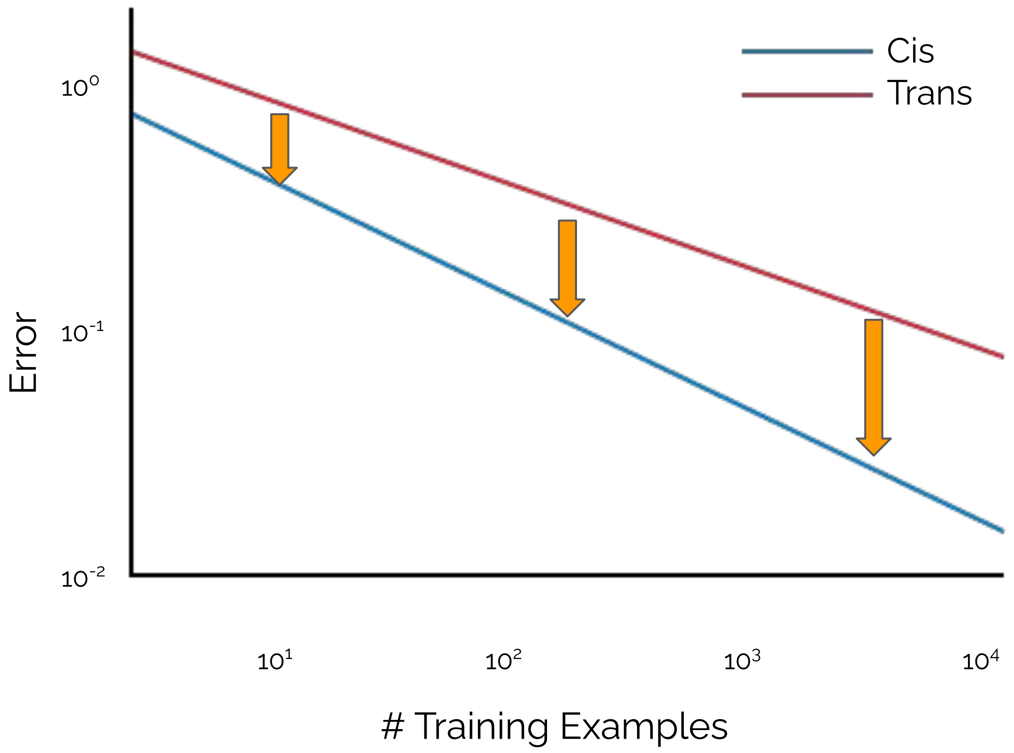

Recognition in Terra Incognita, Beery et al., ECCV 2018

How do we close the gap?

Class-agnostic localization reduces the impact of background, distribution shift, and the long tail

MegaDetector

Microsoft AI for Earth

Efficient Pipeline for Camera Trap Image Review, Beery, et al., DMAIC @ KDD 2019

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

MegaDetector generalizes well to new species, new habitat types, and new parts of the world

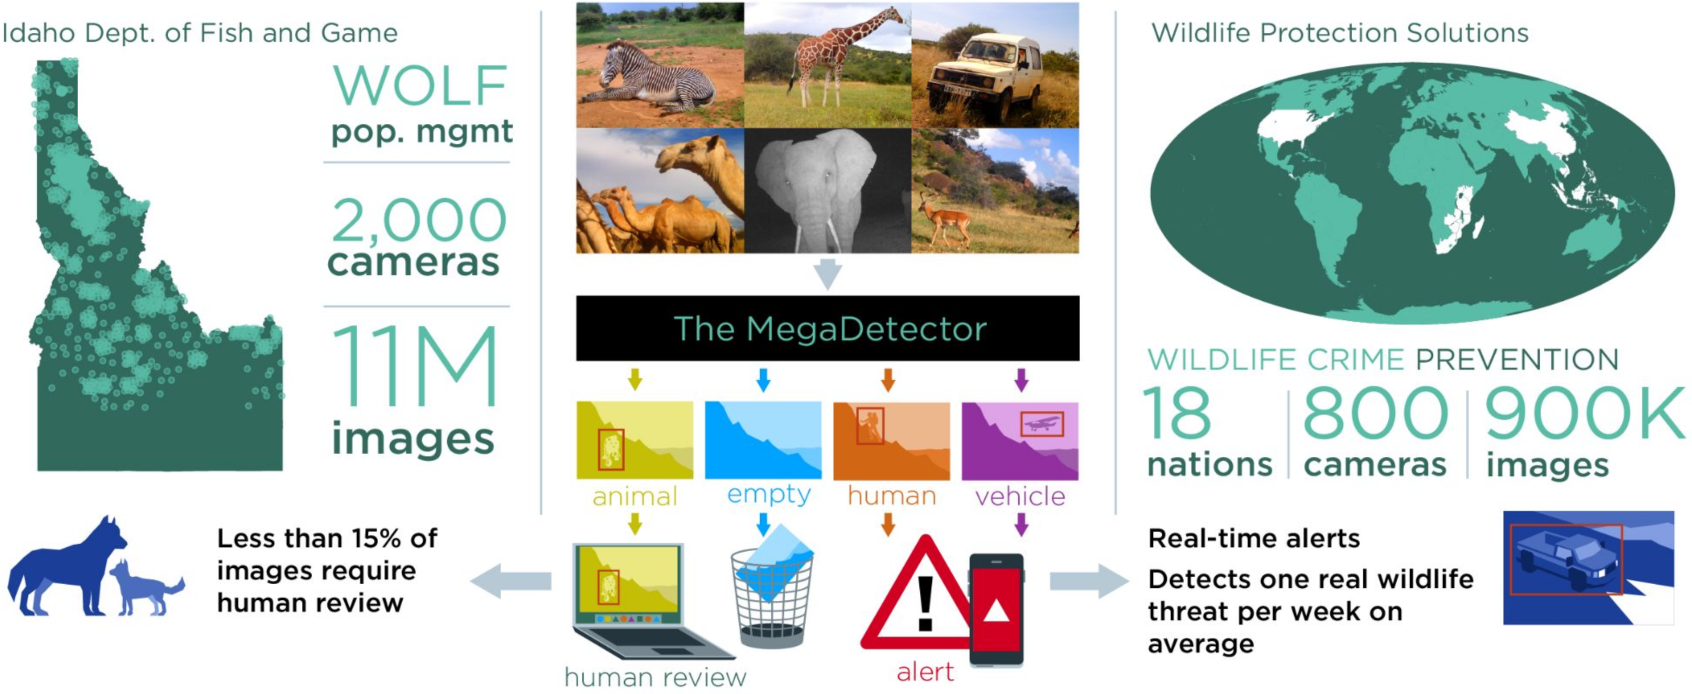

Used to process data for over 40 NGOs and conservation organizations globally, over 100M images last year

Seeing biodiversity: perspectives in machine learning for wildlife conservation, Tuia*, Kellenberger*, Beery*, Costelloe*, et al., Nature Communications (to appear)

Sarah Bassing @S_Bassing · May 19

Thank goodness for the **#MegaDetector** helping me find the ONE animal image mixed in with 170,787 pictures of blowing grass and clouds from this #CameraTrap! Image recognition software is a game changer. #painless #tech4wildlife #WAPredatorPreyProject

Deep active learning to adapt species ID to new projects

- Uses the MegaDetector to crop
- Cluster animals based on visual similarity in new cameras
- Humans ID examples from each cluster (active learning criteria)
- Gets same accuracy with **99.5% fewer labels**

A deep active learning system for species identification and counting in camera trap images, Norouzzadeh, Morris, Beery, Joshi, Jojic, Clune, Methods in Ecology & Evolution, 2021

Learn a spatiotemporal prior to provide context

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

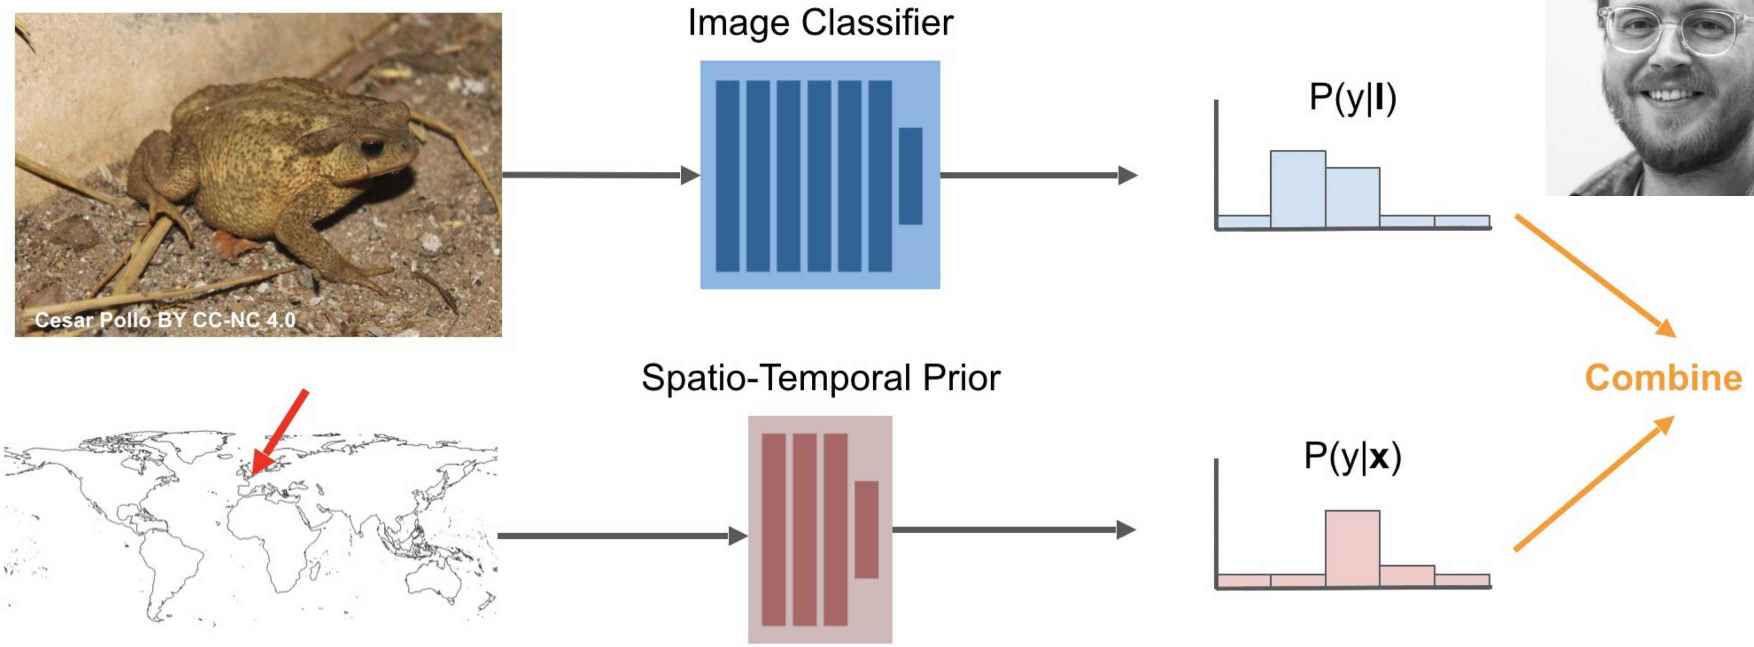

$\mathbf{x}$ = (longitude, latitude, day)

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

Second generalization case study:
Monitoring salmon escapement in static sonar

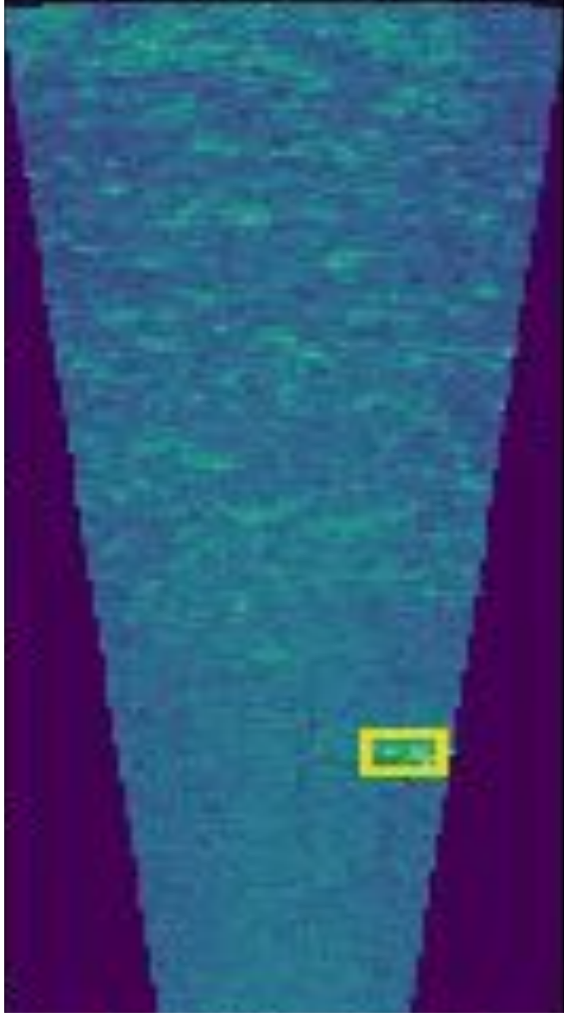

@ECCV 22 with Justin Kay, Peter Kulits, Suzanne Stathatos, Erik Young, Siqi Deng, Grant van Horn, Pietro Perona

Sonar deployment to monitor salmon returns

Photos courtesy of ADFG

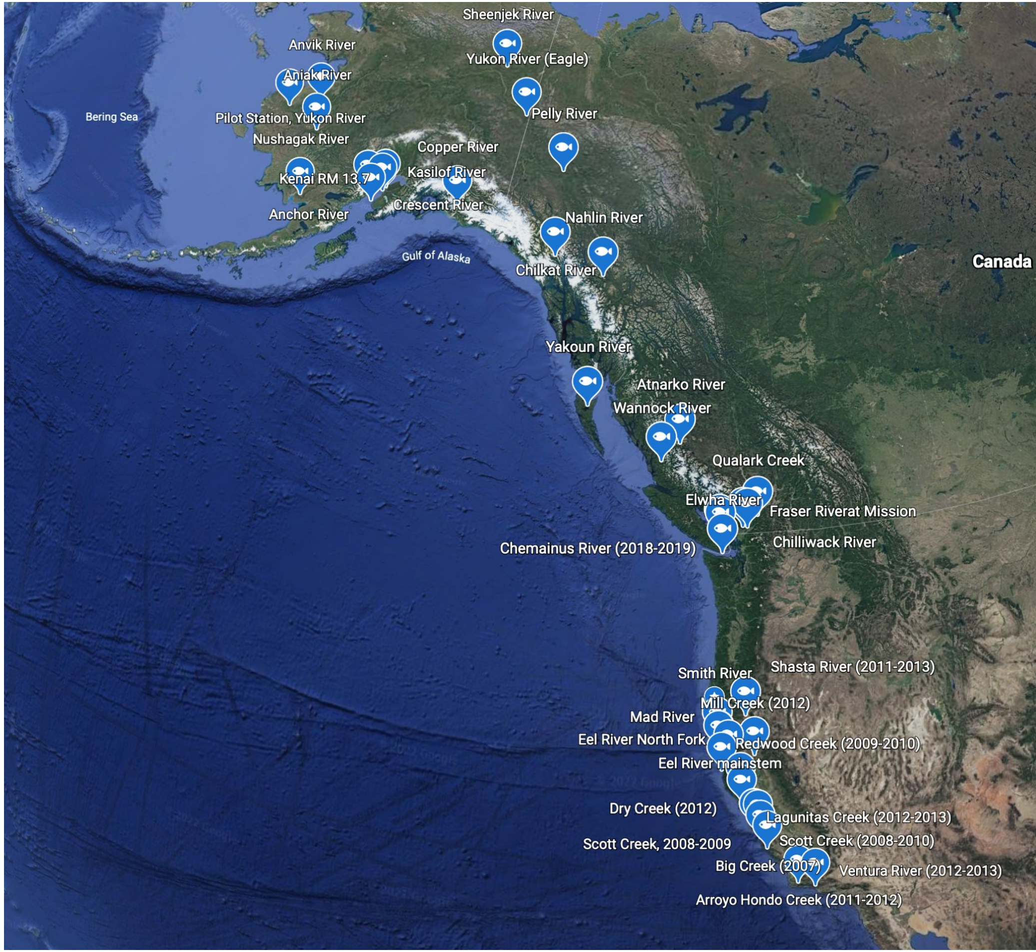

Sonar deployment to monitor salmon returns

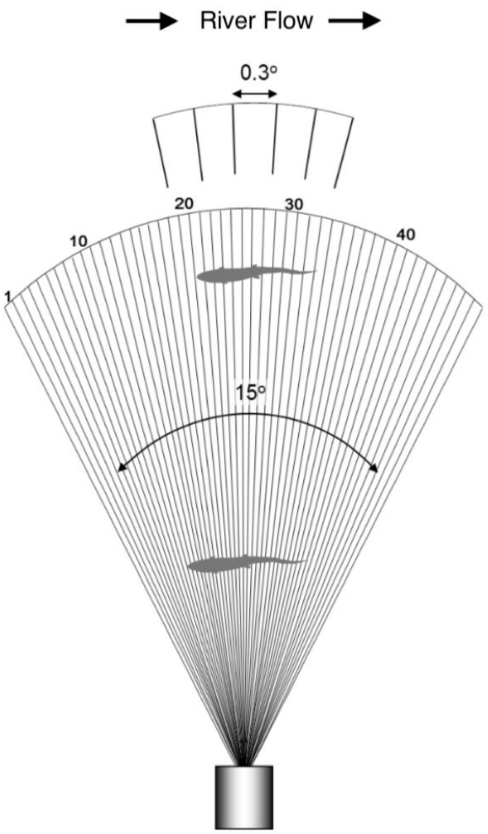

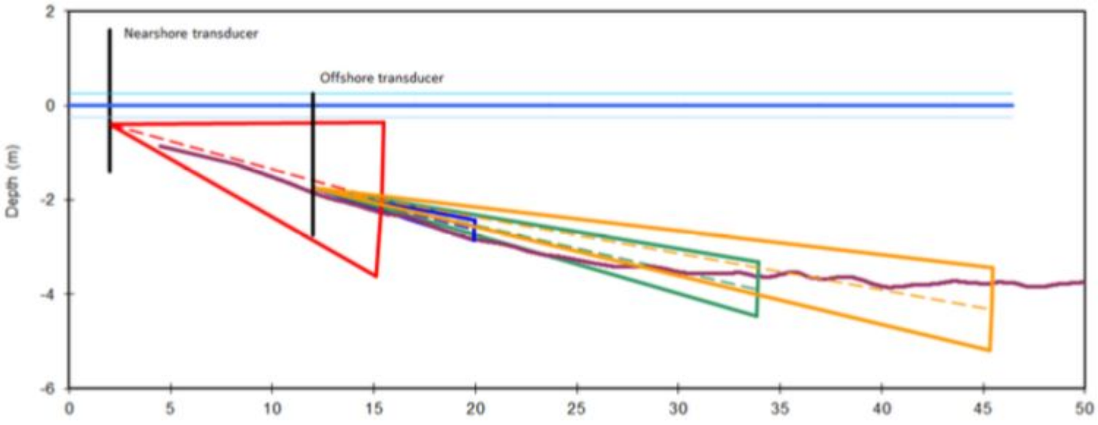

Source: ADFG

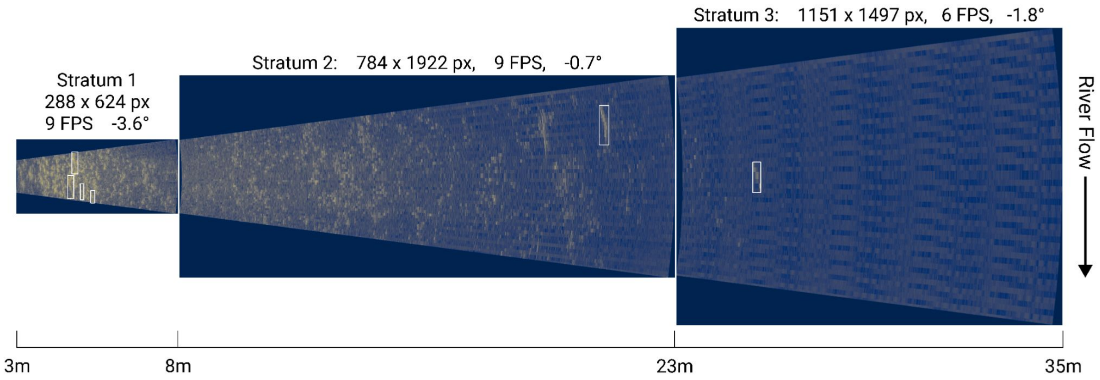

Source: ADFG

Source: Caltech

Manual Processing

Source: ADFG

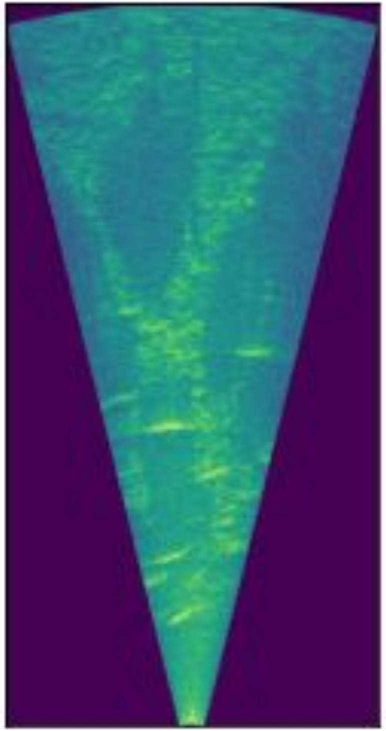

```
Total Fish      =      67
Upstream        =      2
Downstream      =      50

Total Frames     = 7193
Expected Frames  = -1
Total Time       = 0:29:58
Expected Time    = 0:00:00

Upstream Motion  = Right To Left

Count  File Name:   N/A
Editor ID          = N/A
Intensity          = 0.0 dB
Threshold          = 0.0 dB
Window Start       = 1.00
Window End         = 40.00
Water Temperature  = 18 degC


*** Manual Marking (Manual Sizing:  Q = Quality,

File  Total  Frame#  Dir   R (m)    Theta   L(cm)
Motion           Q    N  Comment
------------------------------------------------
------------------------------------------------
  1      1      44  Down   17.19    -0.2   105.2
Running <-->    -1  -1
  1      2     191  Down    7.30     0.5    89.5
Running <-->    -1  -1
  1      3     277  Down    8.41     0.2   120.0
Running <-->    -1  -1
  1      4     548  Down   19.63    -0.2    96.3
Running <-->    -1  -1
  1      5     752  Down   27.13     0.2    98.9
Running <-->    -1  -1
  1      6     826  Down   12.86     0.2    94.8
Running <-->    -1  -1
  1      7    1071  Down   10.77     0.2    93.2
Running <-->    -1  -1
  1      8    1238  Down   13.62    -0.2    86.7
Running <-->    -1  -1
  1      9    1353   N/A   22.04     5.2   105.9
Running <-->    -1  -1
  1     10    1471   N/A   25.45     6.4    61.2
Running <-->    -1  -1
  1     11    1521  Down   34.80    -0.2   123.7
Running <-->    -1  -1
```

Counting Baseline

1. Detect

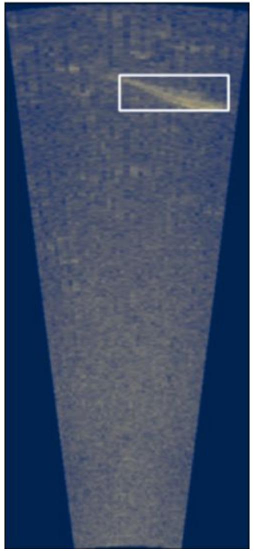

Counting Baseline

1. Detect

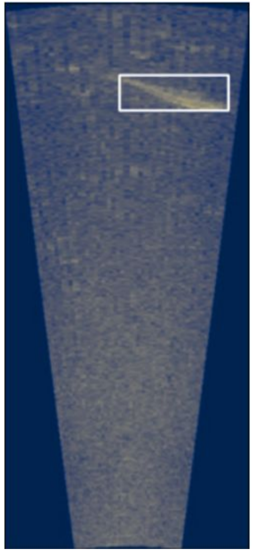

2. Track

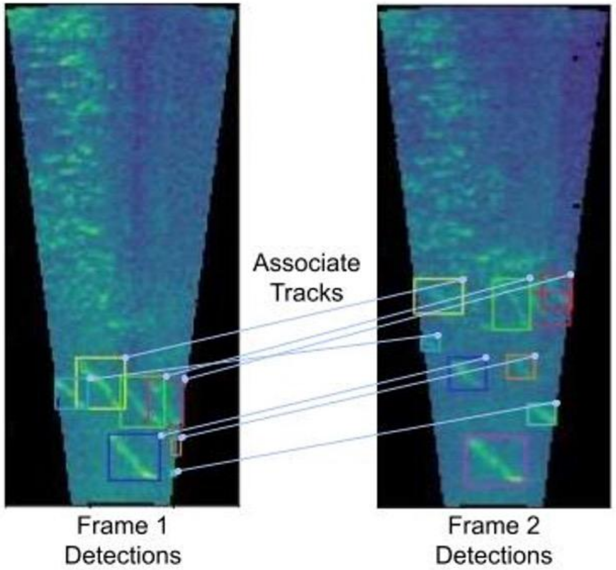

Counting Baseline

1. Detect

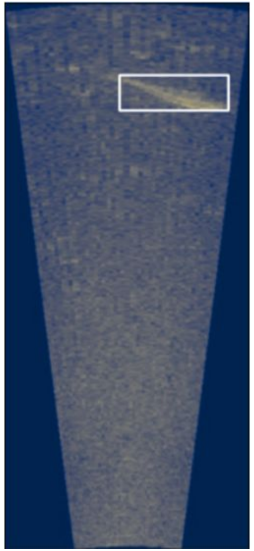

2. Track

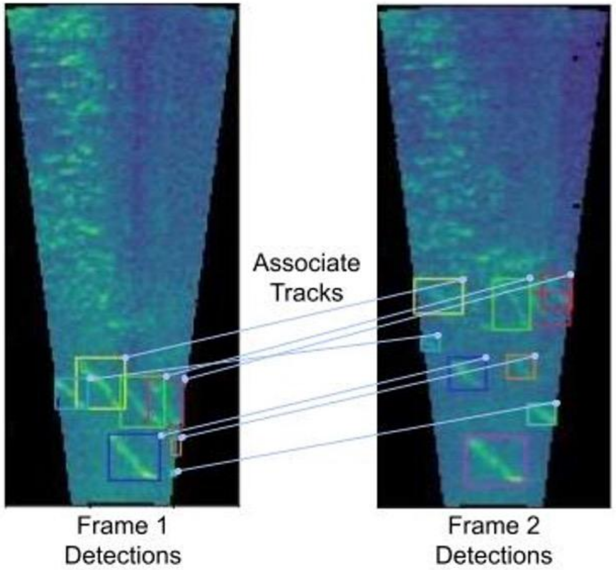

3. Count

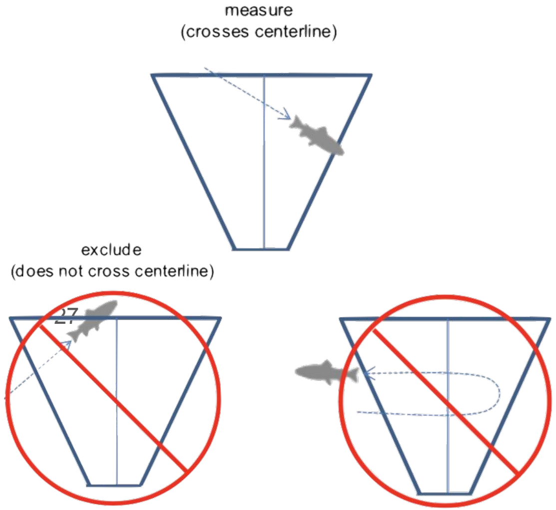

Source: Key et al. Operational Plan: Kenai River Chinook Salmon Sonar Assessment at River Mile 13.7, 2020–2022

Challenges

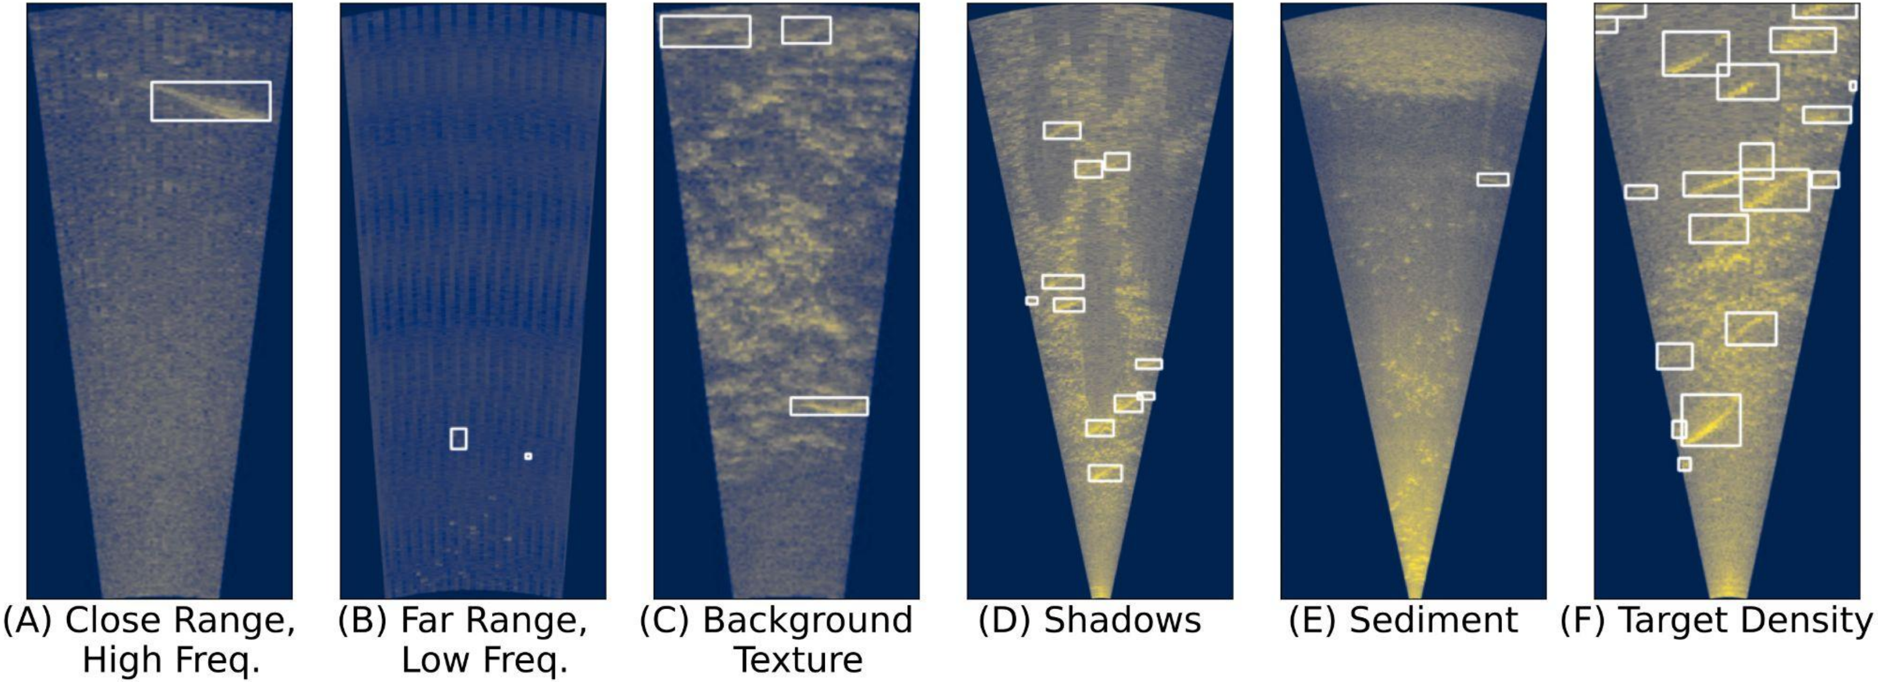
(A) Close Range,
High Freq.
(B) Far Range,
Low Freq.
(C) Background
Texture
(D) Shadows
(E) Sediment
(F) Target Density

Generalizable detection is the bottleneck

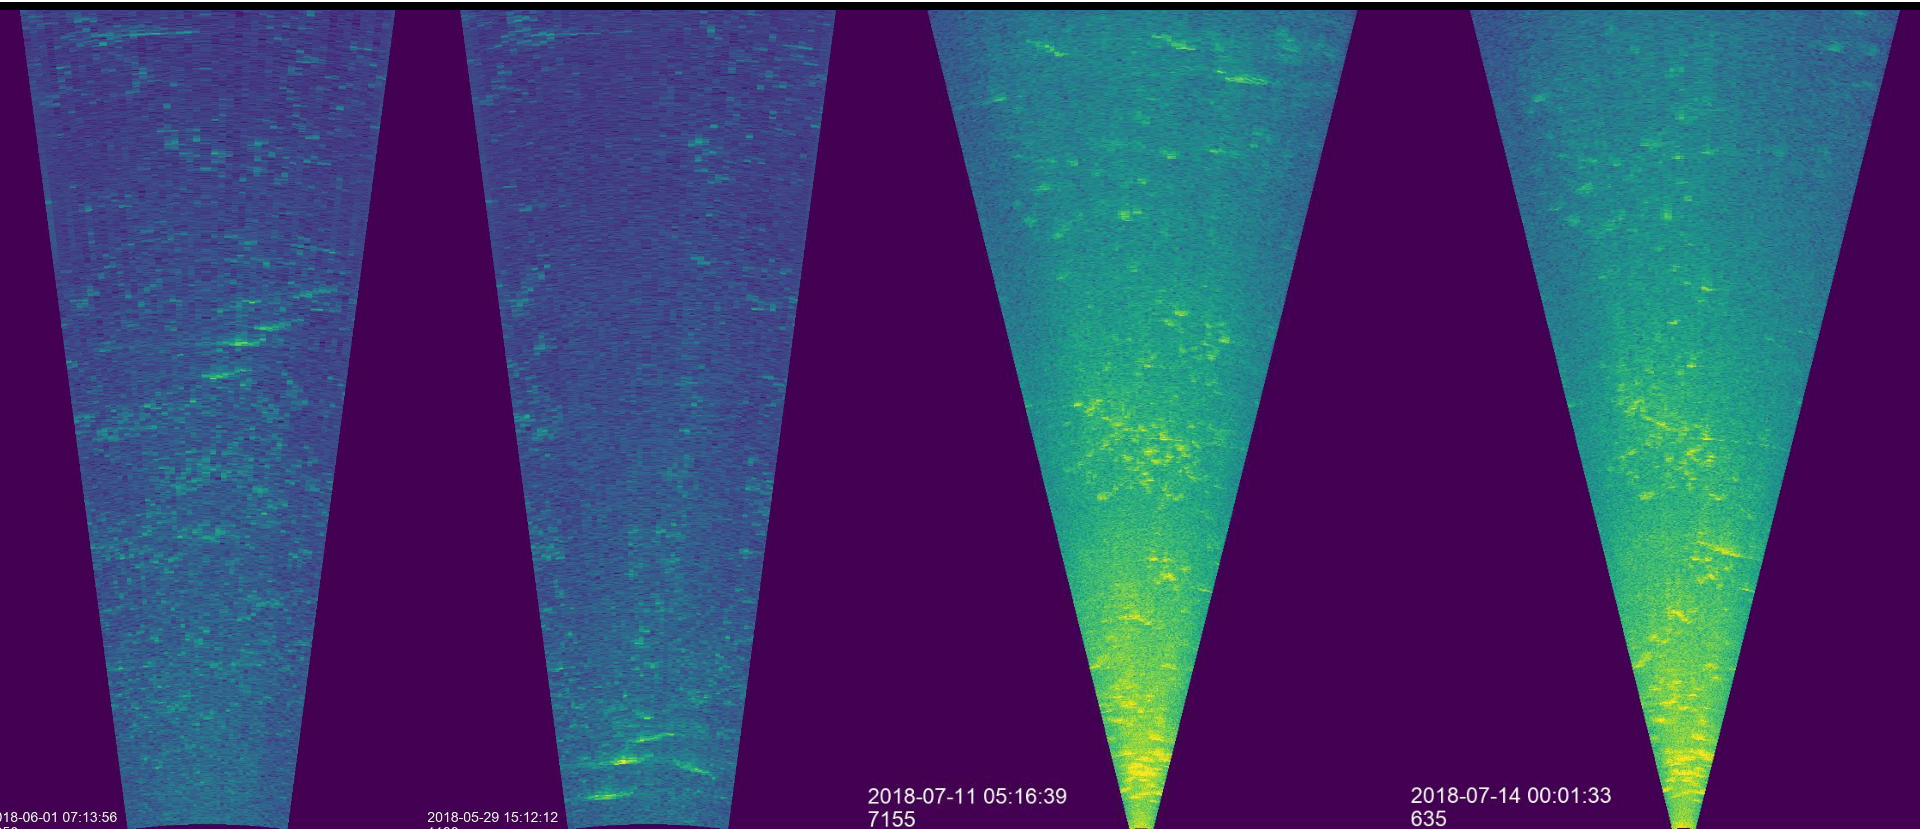

Third generalization case study:
Multiview Urban Forest Monitoring

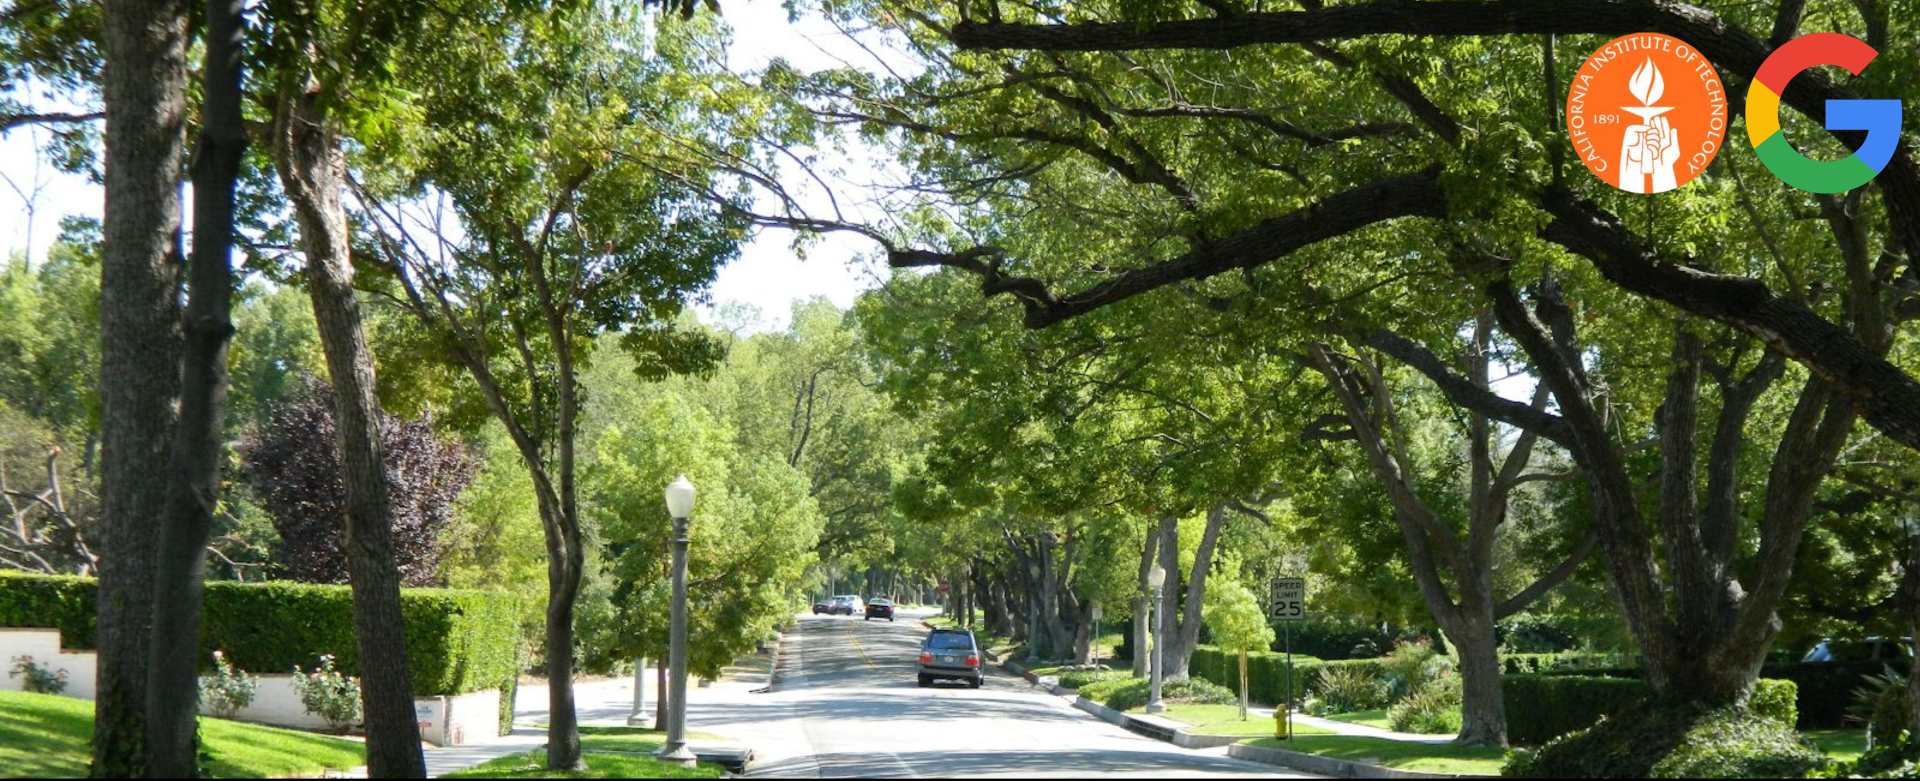

Auto Arborist

@CVPR22 with Guanhang Wu, Trevor Edwards, Filip Pavetic, Bo Majewski, Shreyasee Mukherjee, Stanley Chan, John Morgan, Vivek Rathod, Jonathan Huang

Benefits of the Urban Forest

Biodiversity

Cities support regional biodiversity

Large trees and a diverse, connected urban forest supporting a rich array of wildlife, particularly birds

Reduces Air Pollution

Removes some 784k tons of air pollution annually

Implied global value: $15-20B/yr
Potential impact: $1.5B-$5B/yr

Carbon Sequestration

Total opportunity for additional carbon sequestration ranges from 1GT to 2.4GT

At $50/ton, that's a value of $50B-$120B, cumulatively (i.e., not annually)

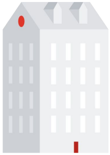

Reduced Energy Use

Trees reduce building energy use and avoided pollutant emissions ($8B+ value in U.S. alone)

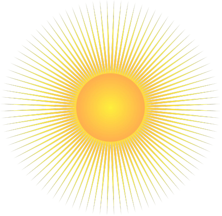

Extreme Heat Islands

Lowers surface and air temperatures by providing shade and through evapotranspiration

Physical + Mental Health

Trees in a community correlate with lower asthma rates, reduced hospital visits during heat waves and improved mental health

Slide thanks to Google AI for Nature and Society

These benefits are not accessible to all

Pasadena

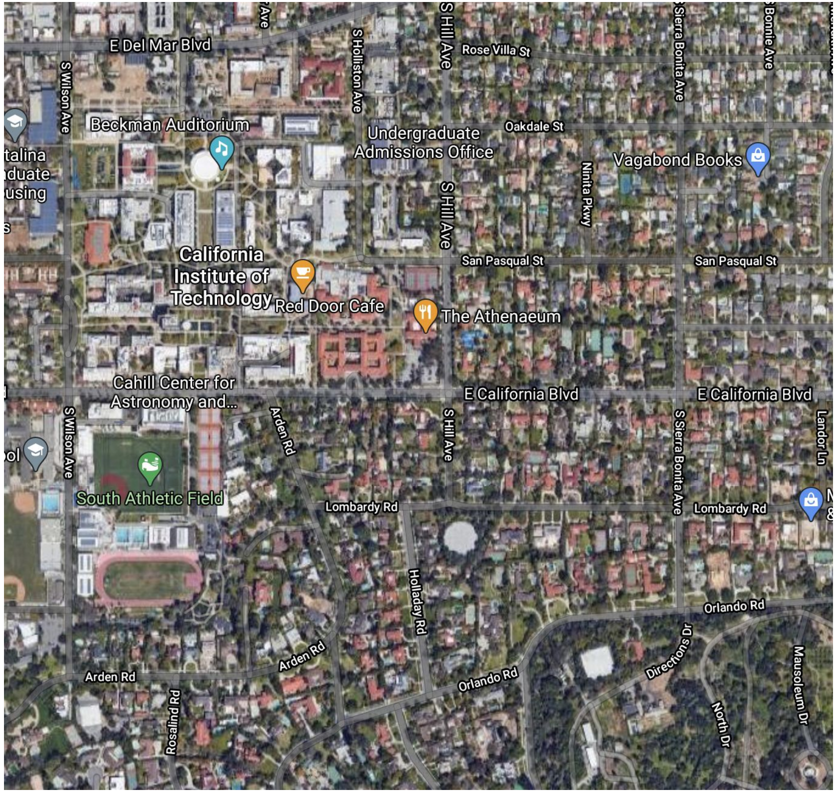

Carson

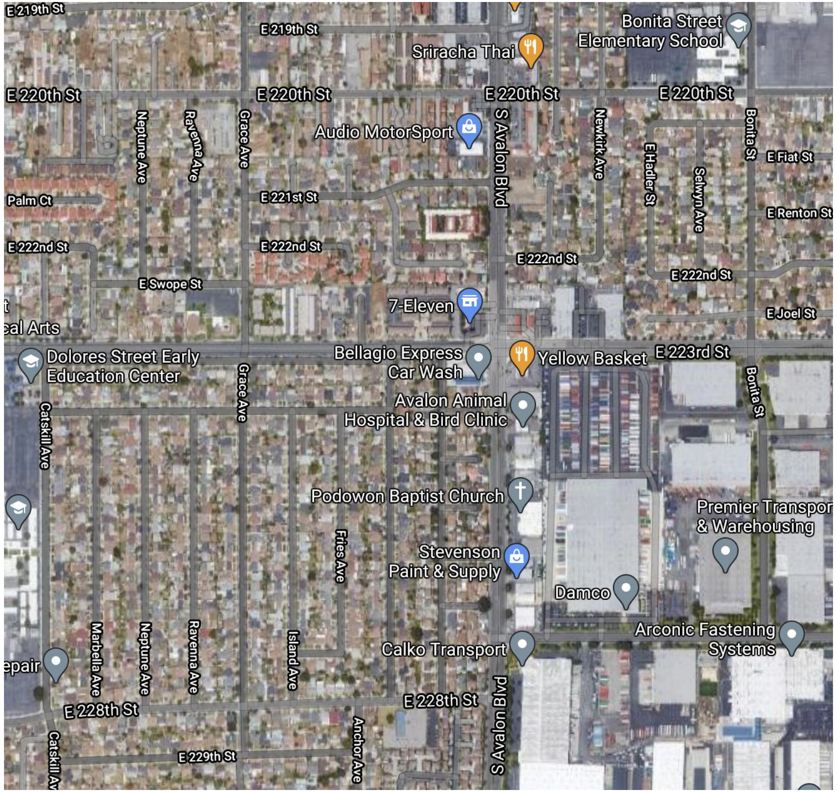

Tree inventories are $$$

A single traditional census costs ~$10M

- **Inequitable**
- **Out of date**
- **Limited scope**

Tree canopy prediction in LA via Urban Ecology Team

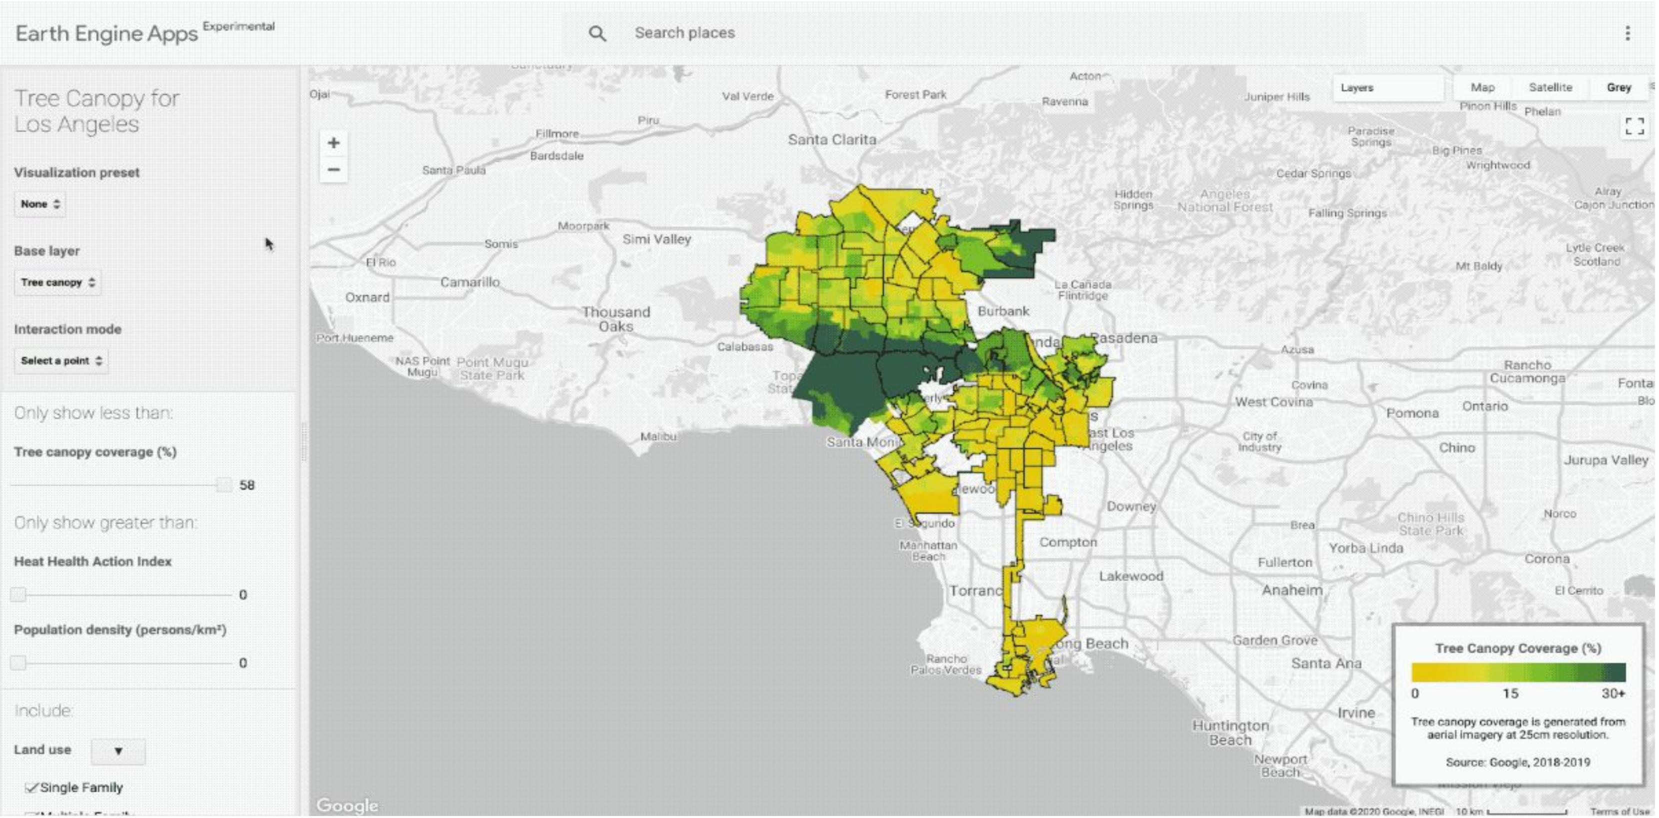

https://insights.sustainability.google/labs/treecanopy

Tree canopy prediction is not enough

Instance locations and species identification is needed to:

- Estimate water retention
- Estimate carbon sequestration
- Estimate potential heat reduction
- Monitor species' reaction and resilience to our changing climate at scale
- Strategically plan planting to maximize biodiversity

The Auto Arborist Dataset:

23 cities, 344 genera, 2.6M tree records, >1M trees w/ imagery

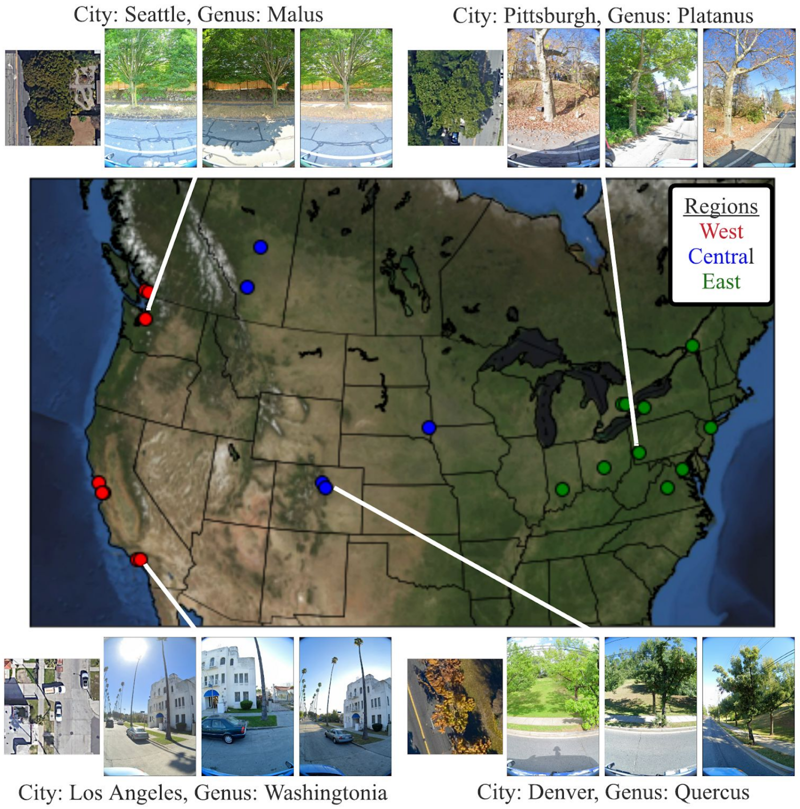

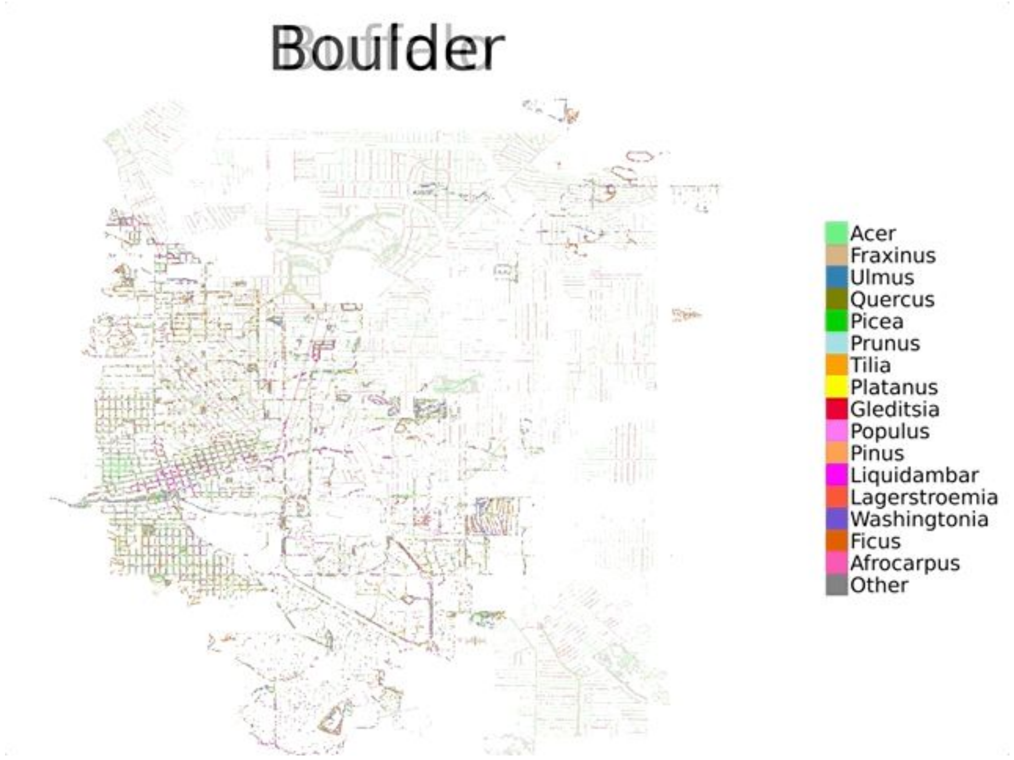

Long tailed and fine-grained, with real-world spatiotemporal and taxonomic structure capturing natural domain shifts across cities

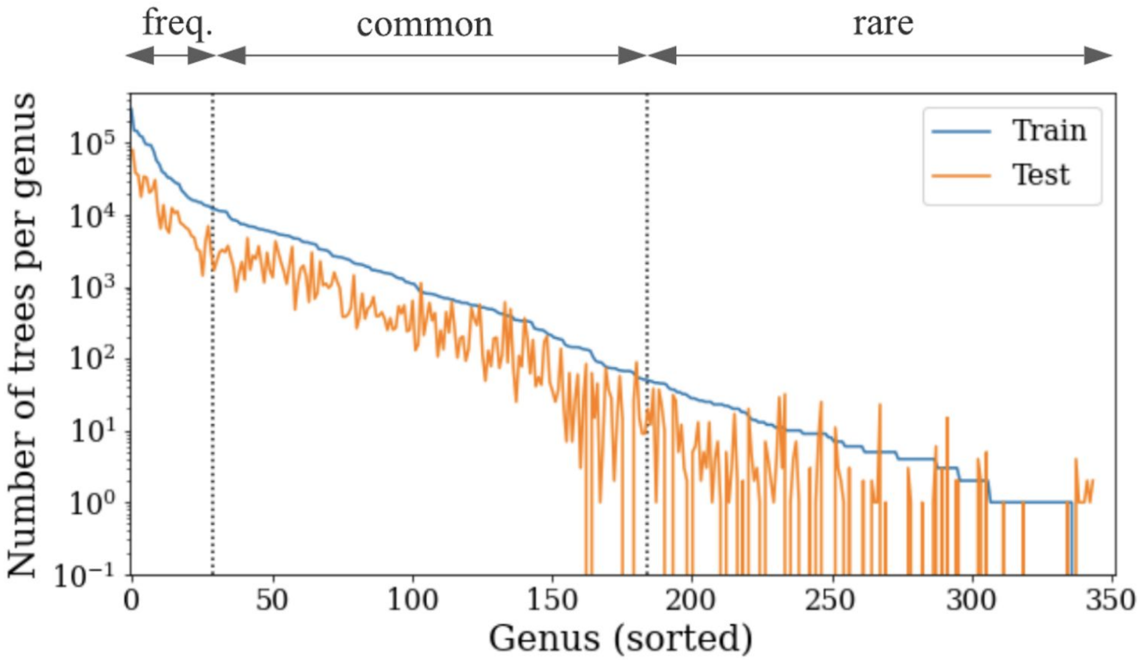

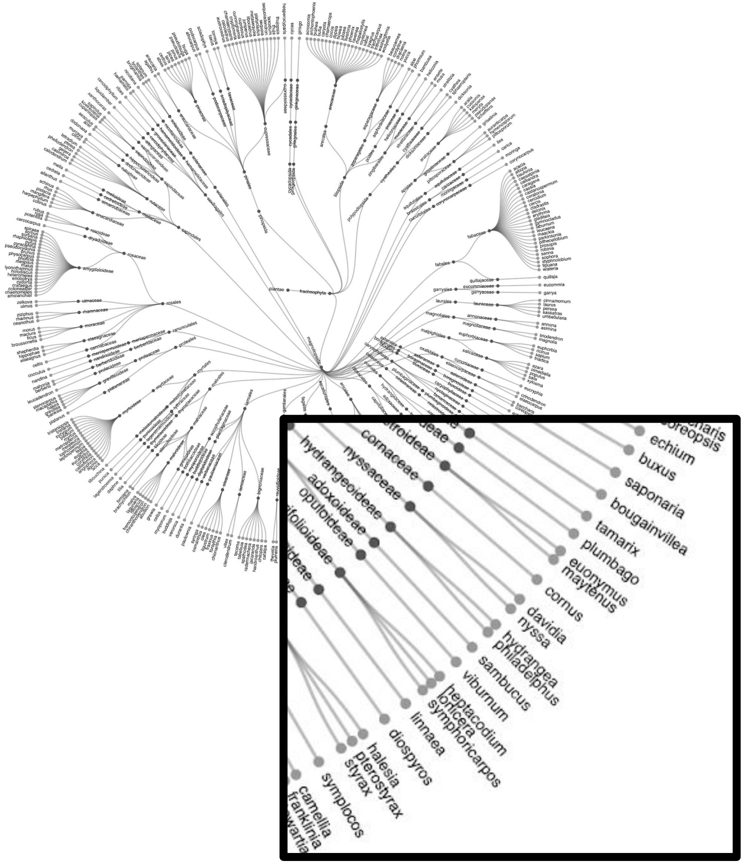

Multiview aerial and street level imagery for the same tree instance

Sioux City, Fraxinus (Ash)

Sioux City, Tilia

Pittsburgh, Taxodium

Buffalo, Cercis (Redbud)

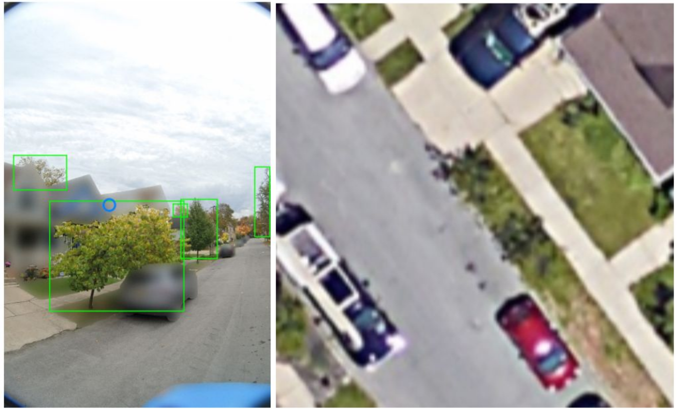

Combining information across views achieves best results

Train Set	Aerial	1 SL	3 SL	A+SL
Region W	20.63	41.53	45.12	**46.07**
Region C	18.8	44.77	46.91	**47.12**
Region E	17.54	43.25	45.13	**46.21**
Full	18.7	46.13	49.0	**49.23**
Full w/ Regional MoE				**49.96**

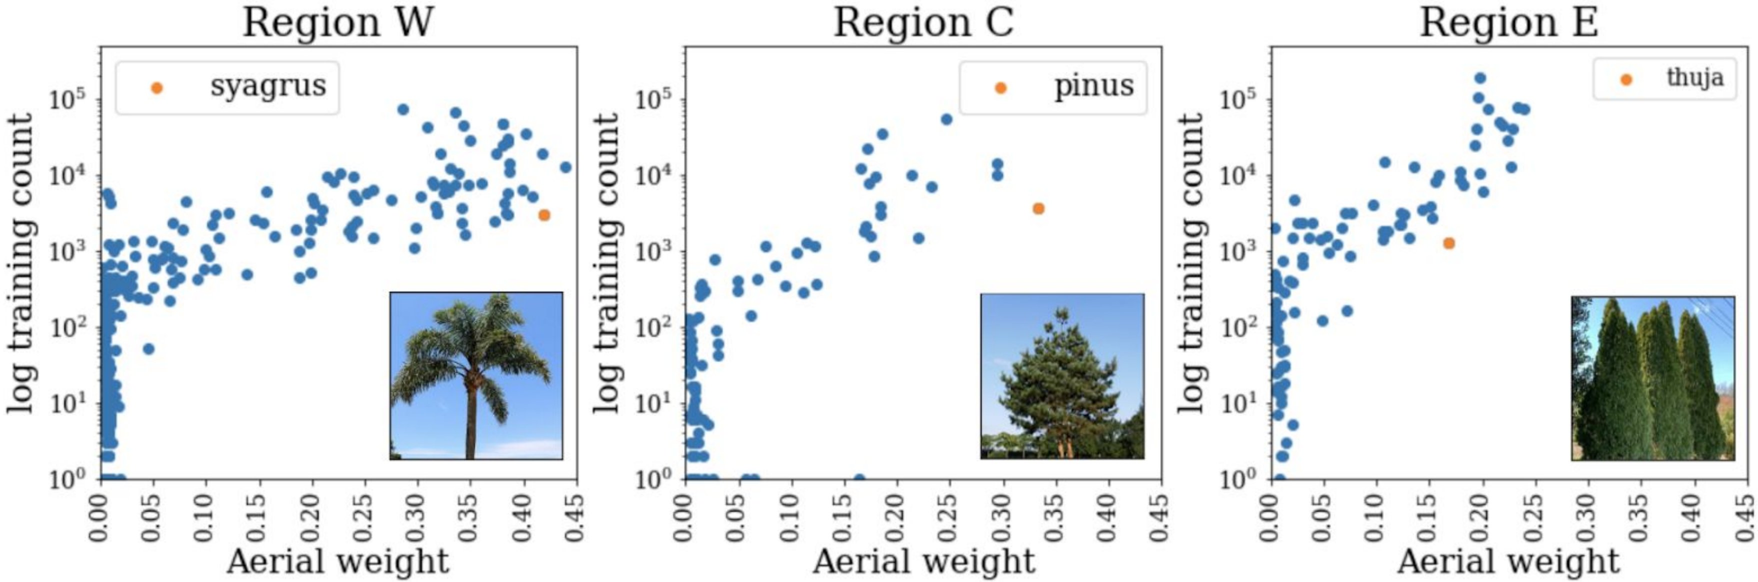

Models trained on the full dataset outperform city-specific or region-specific models

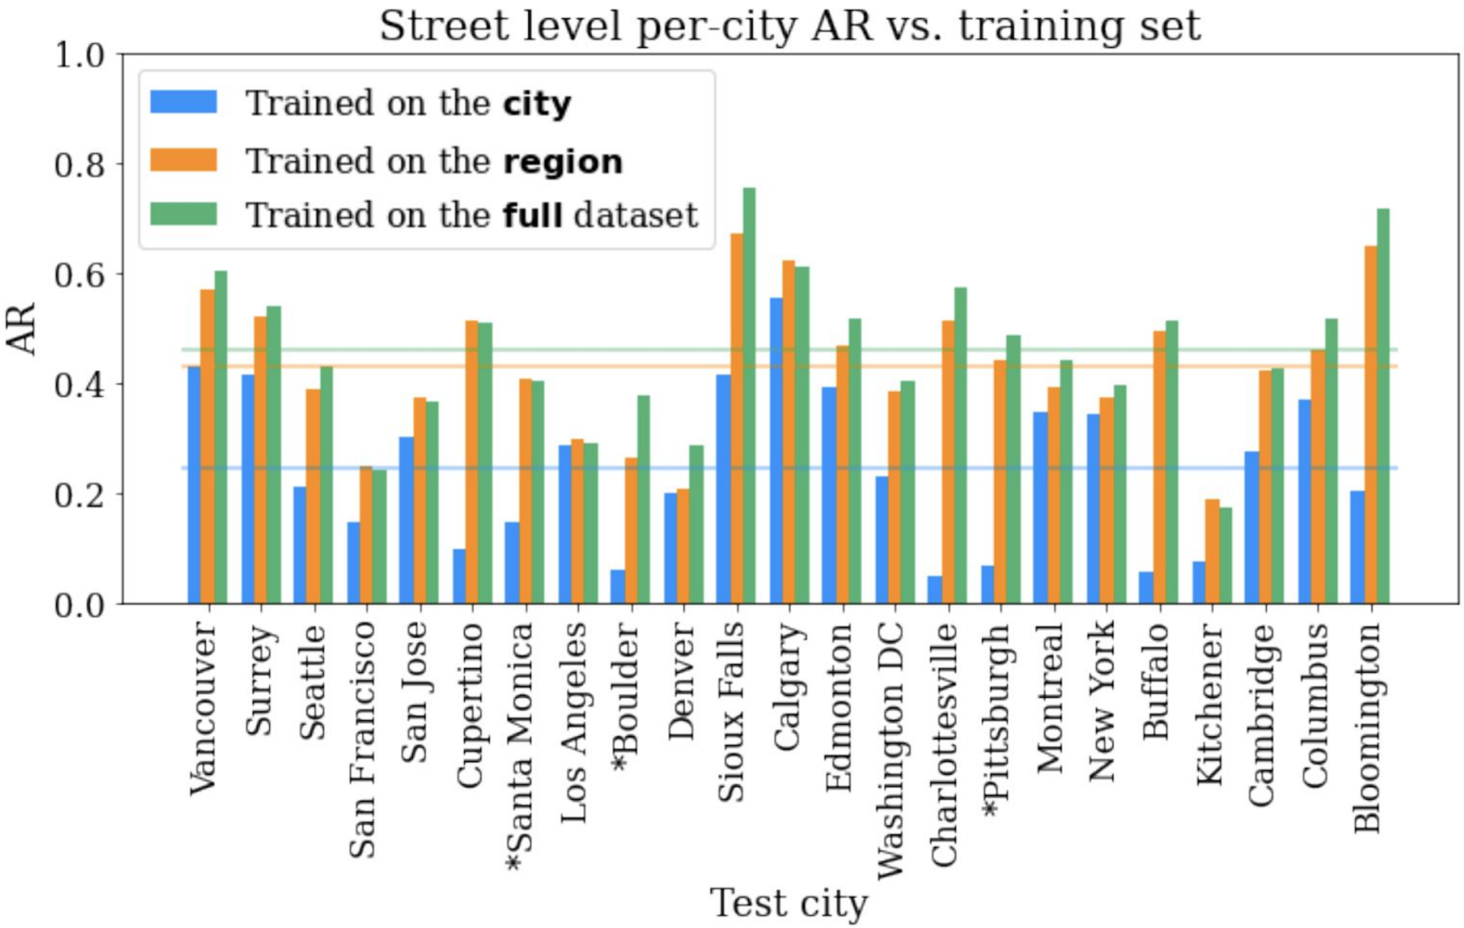

Diverse, large-scale data curation is valuable: some cities generalize better than others, but the full dataset generalizes best to all cities

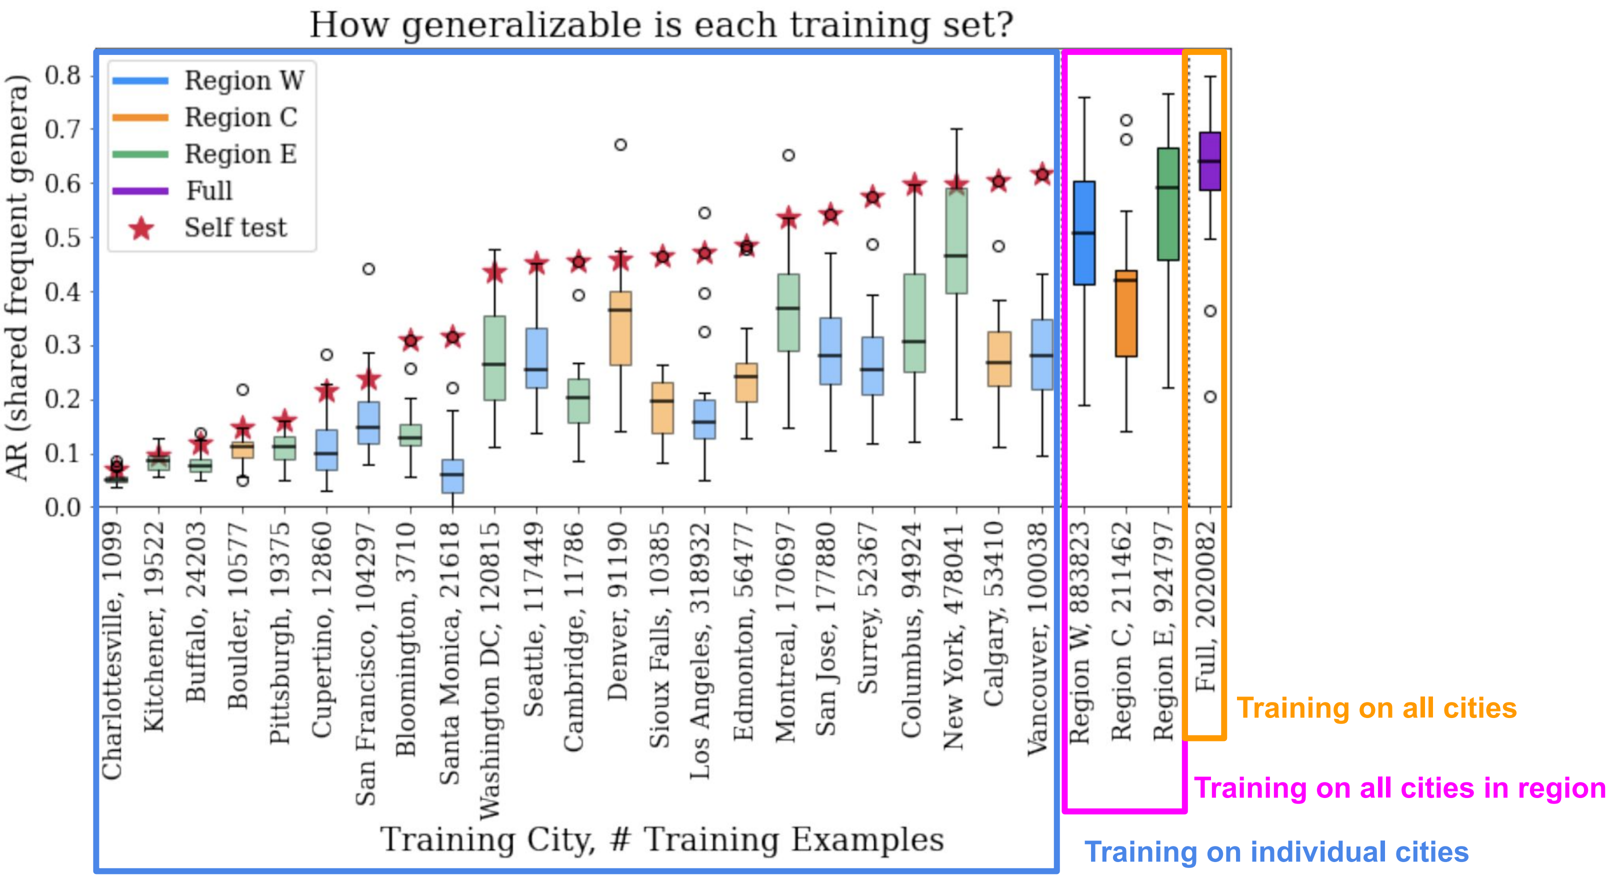

Distribution Shifts Across Cities

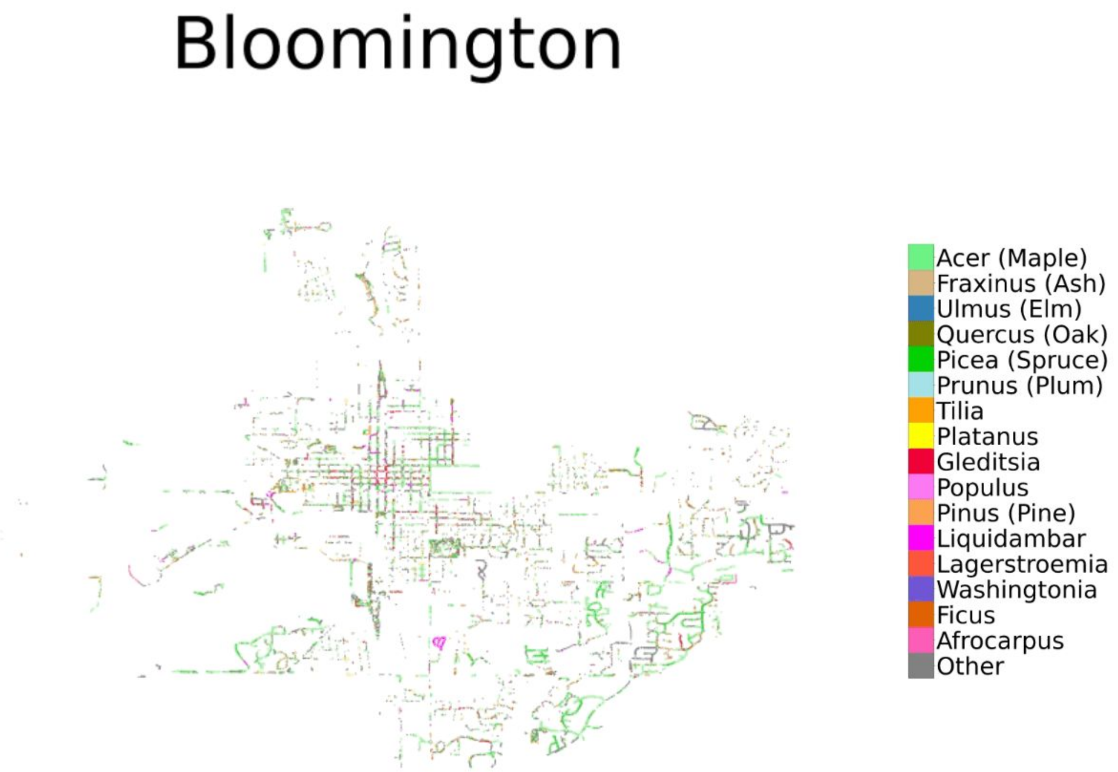

Distribution Shifts Across Cities

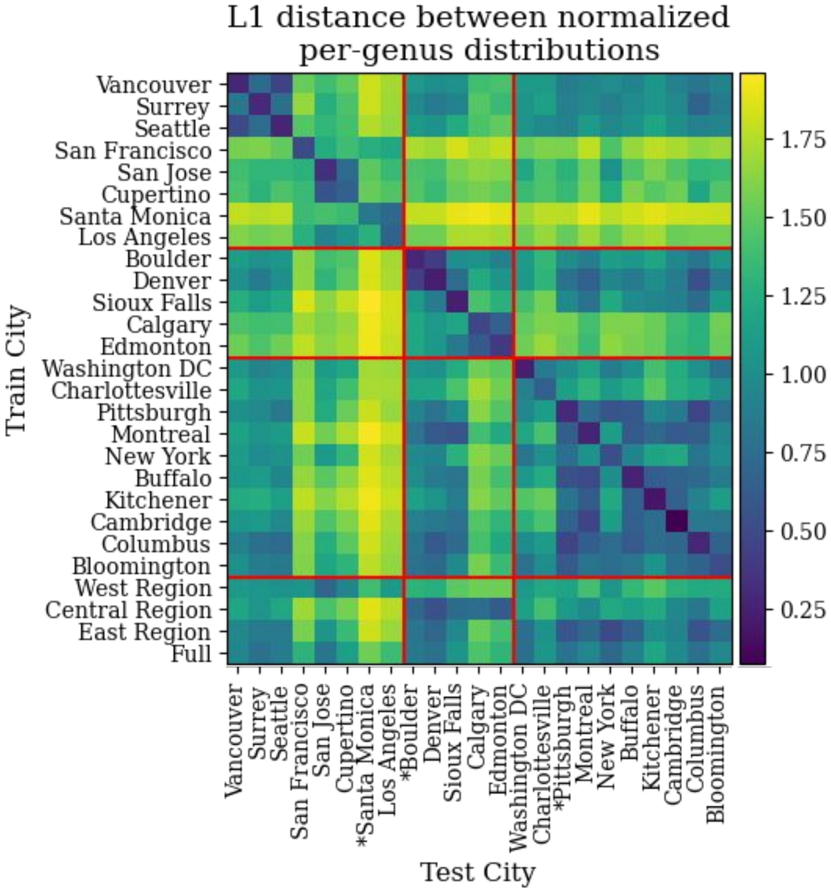

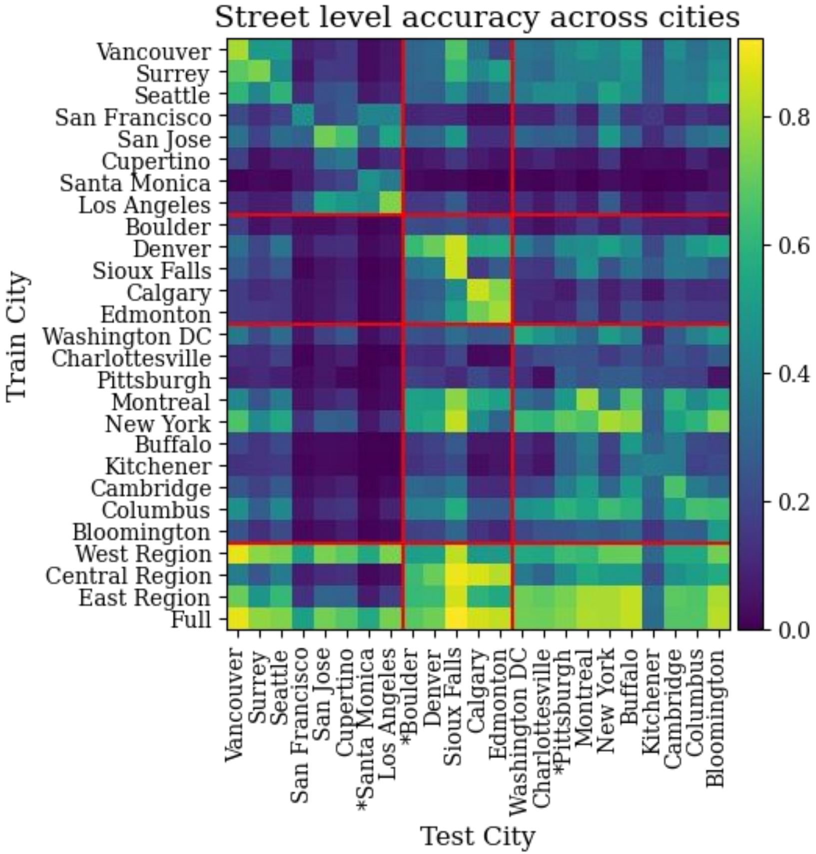

Dataset Release

https://google.github.io/auto-arborist/

Open challenges in CV4Ecology

- Global and Local Domain shift
- Long-tailed distributions
- Sparse, low-quality, multimodal data
- Interactive ecologist-AI systems
- Equitable access to technology
- Limited Interdisciplinary capacity

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

Summer School on Computer Vision Methods for Ecology

CALTECH RESNICK SUSTAINABILITY INSTITUTE

http://cv4ecology.caltech.edu/

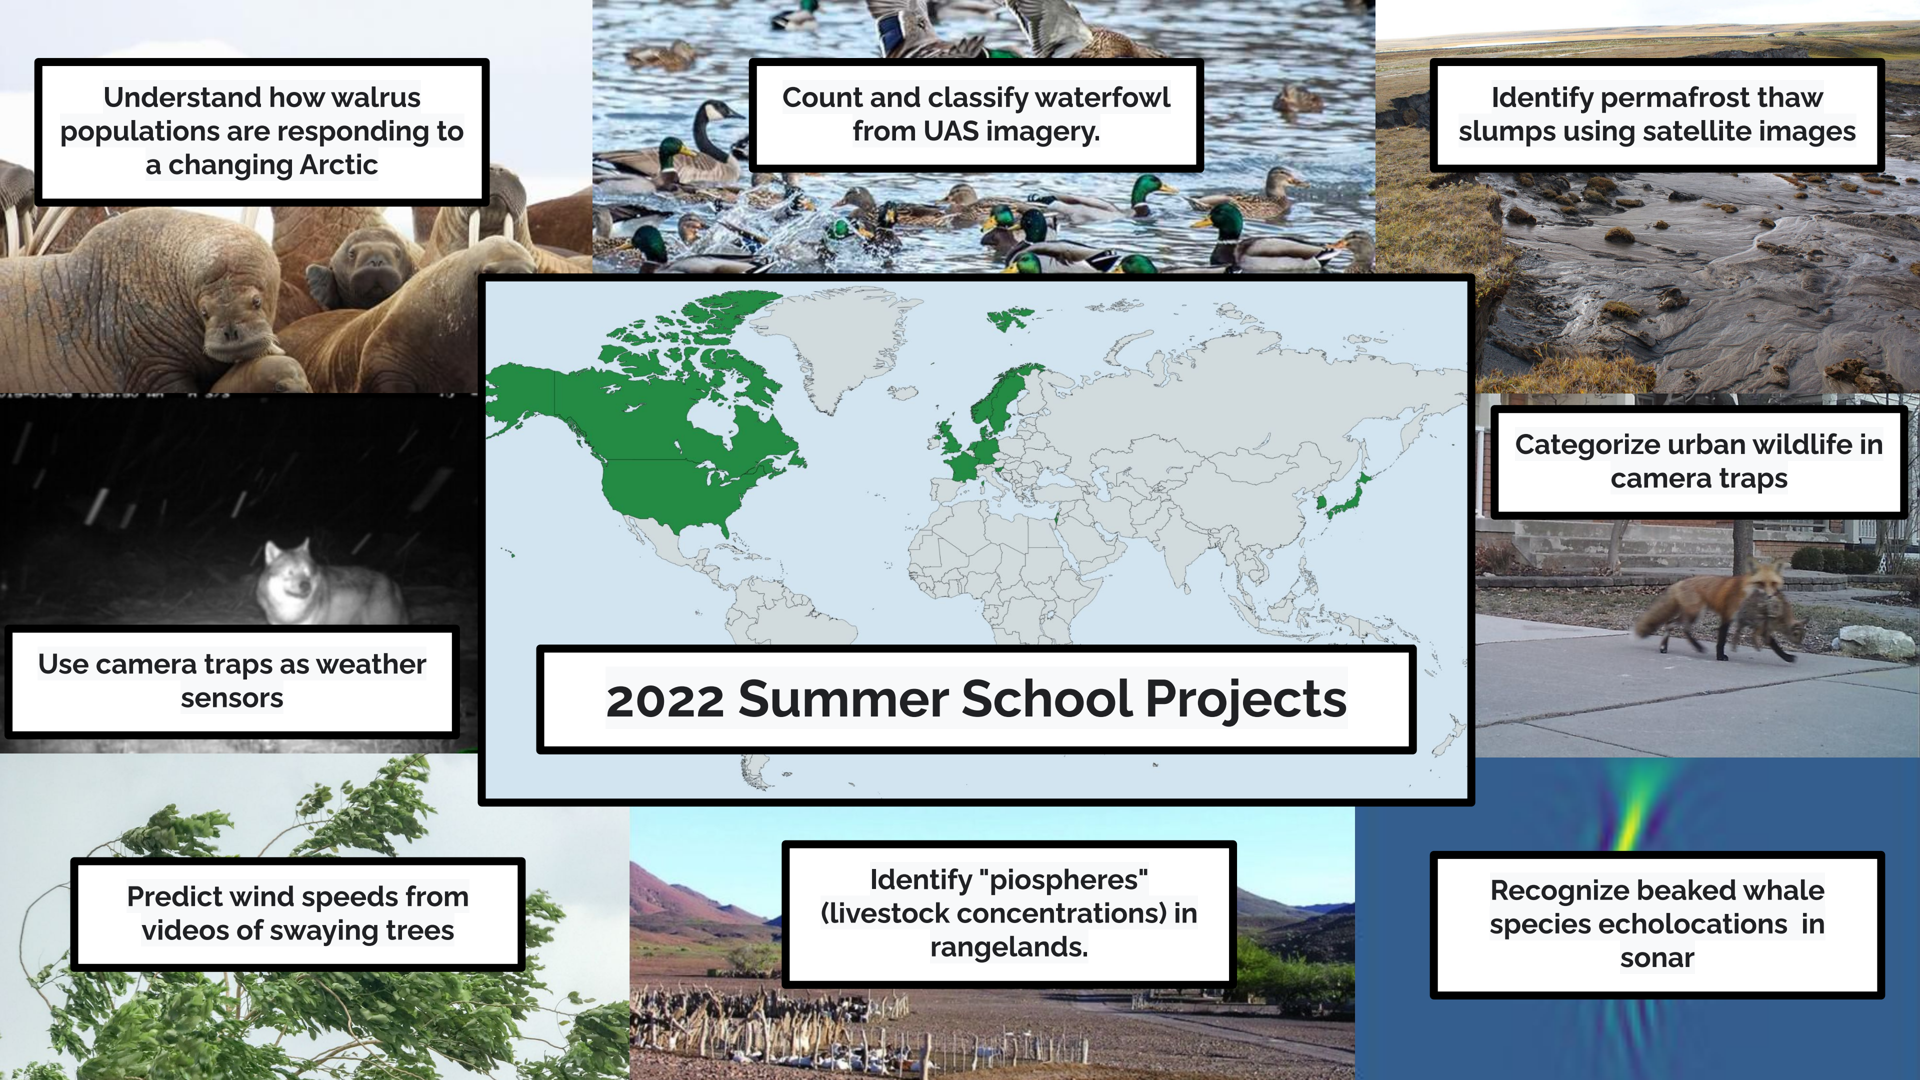
Understand how walrus populations are responding to a changing Arctic
Count and classify waterfowl from UAS imagery.
Identify permafrost thaw slumps using satellite images
Categorize urban wildlife in camera traps
Use camera traps as weather sensors
2022 Summer School Projects
Predict wind speeds from videos of swaying trees
Identify "piospheres" (livestock concentrations) in rangelands.
Recognize beaked whale species echolocations in sonar

Thanks!

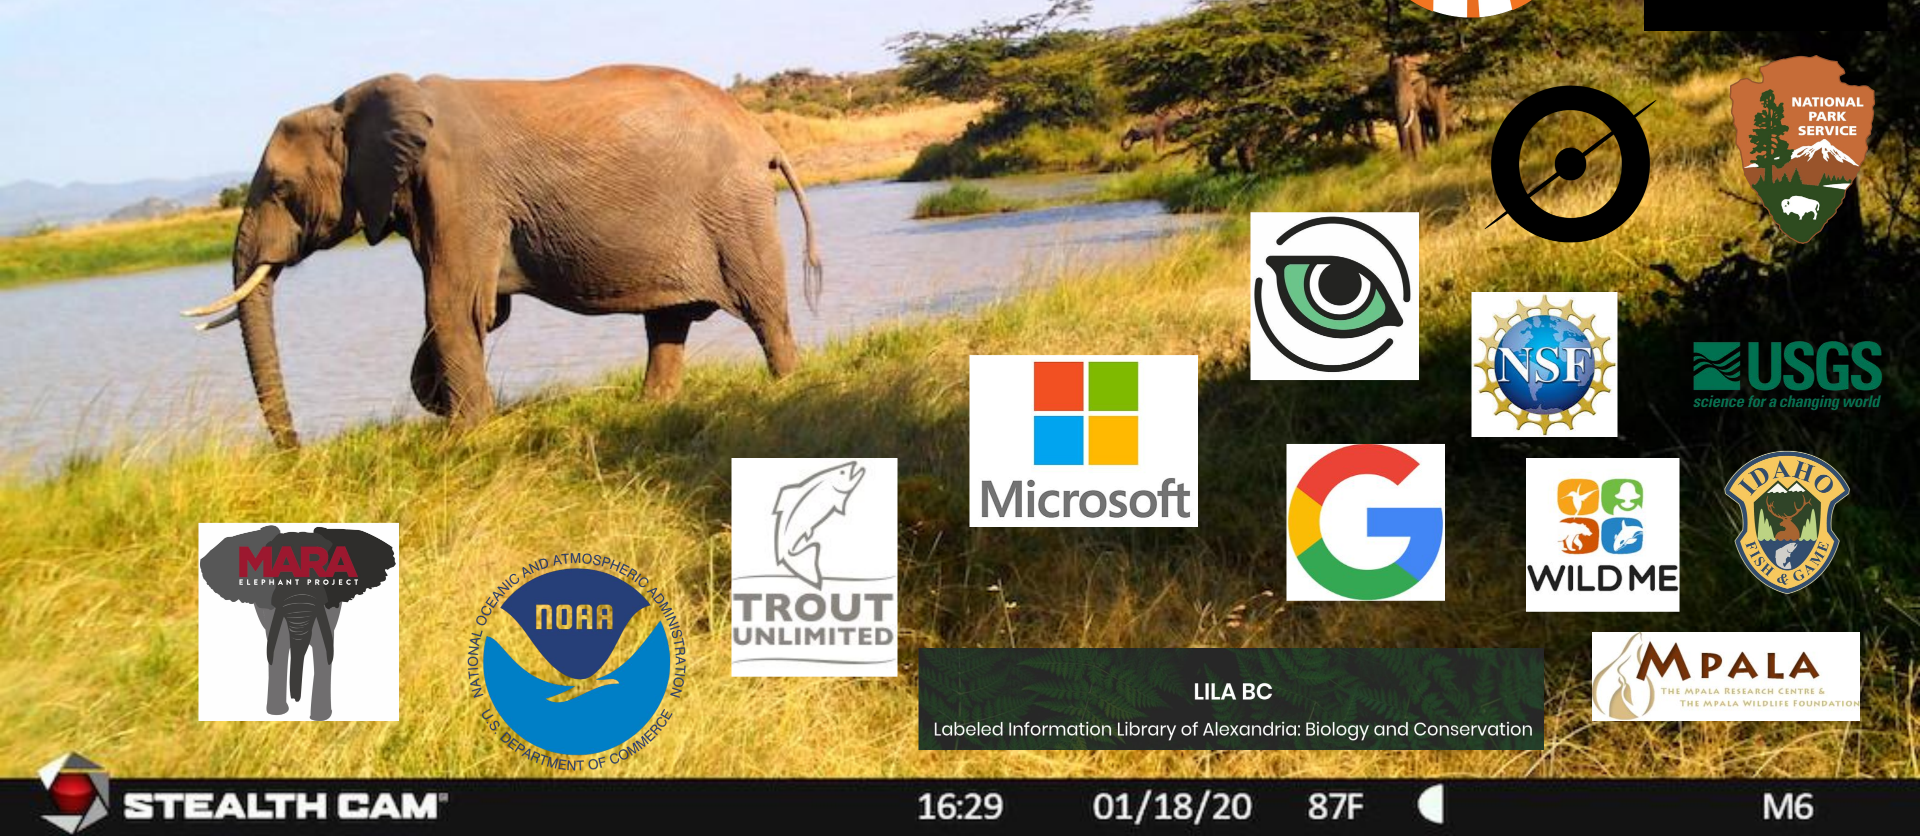
CALIFORNIA INSTITUTE OF TECHNOLOGY
1891
NATIONAL PARK SERVICE
NSF
USGS
science for a changing world
Microsoft
TROUT UNLIMITED
WILD ME
IDAHO FISH & GAME
MARA ELEPHANT PROJECT
NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION
NOAA
U.S. DEPARTMENT OF COMMERCE
LILA BC
Labeled Information Library of Alexandria: Biology and Conservation
MPALA
THE MPALA RESEARCH CENTRE & THE MPALA WILDLIFE FOUNDATION
STEALTH CAM
16:29
01/18/20
87F
M6