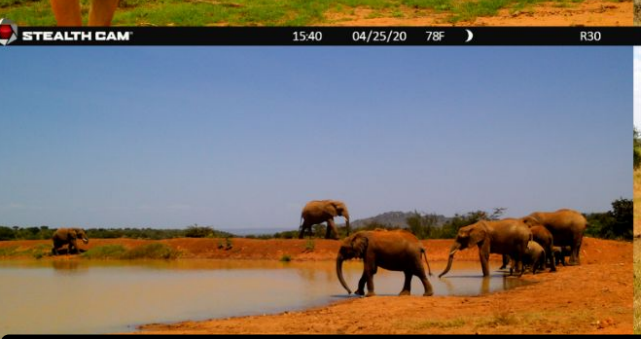

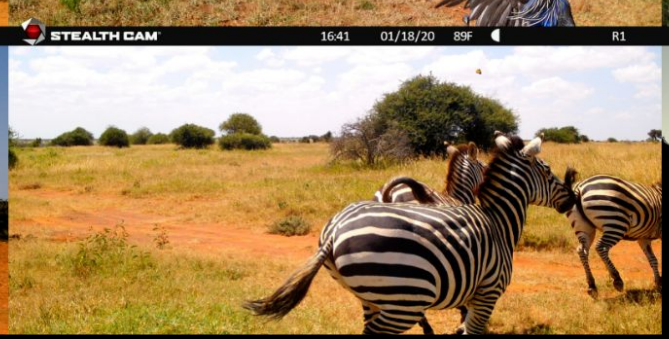

6.S894 - Computer Vision and Planetary Health


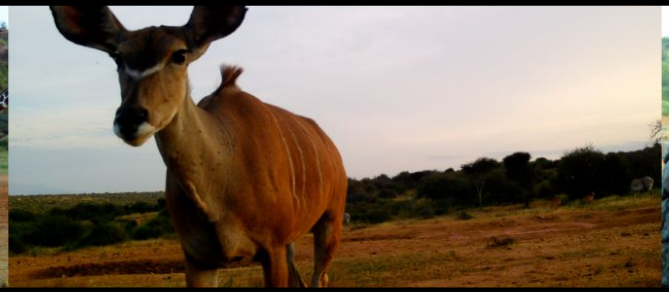


Sara Beery | 2/4/25



Biodiversity is in decline globally

LIVING PLANET REPORT 2020

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

16:3

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 is in decline globally

Biodiversity is intrinsically tied to:

- Climate Change
- Public Health
- Food Security
- Ecosystem Services

Biodiversity data collection is increasing in quantity and diversity

Perspectives in machine learning for wildlife conservation, Tuia*, Kellenberger*, Beery*, Costelloe*, et al., Nature Comms 2022

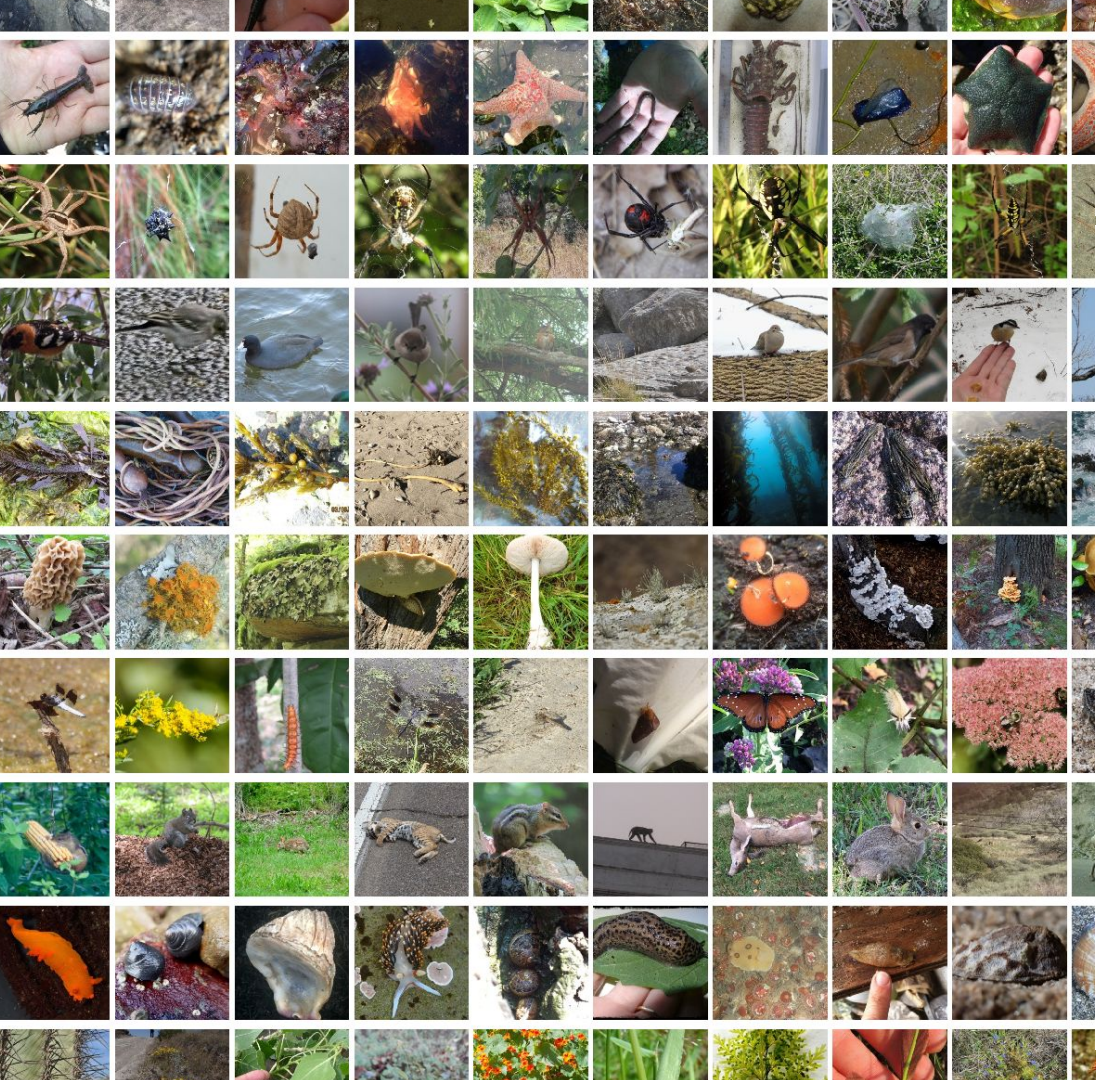

iNaturalist

226,146,504
Observations

506,051
Species

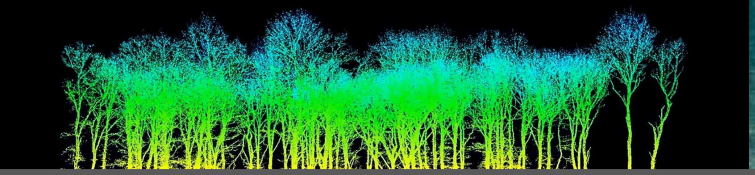

No direct sensor for species across taxa & scale


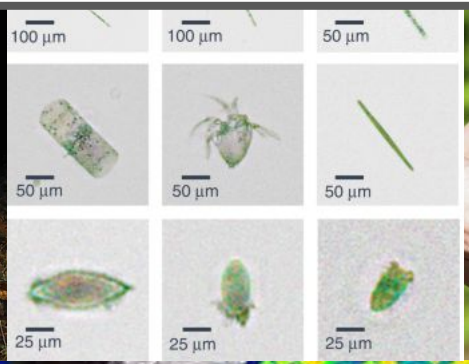


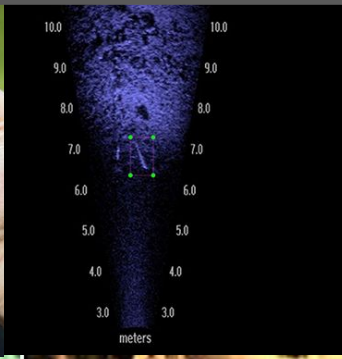

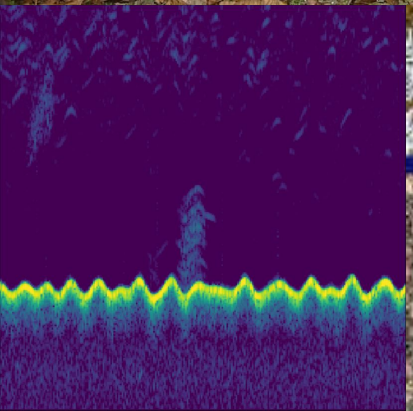
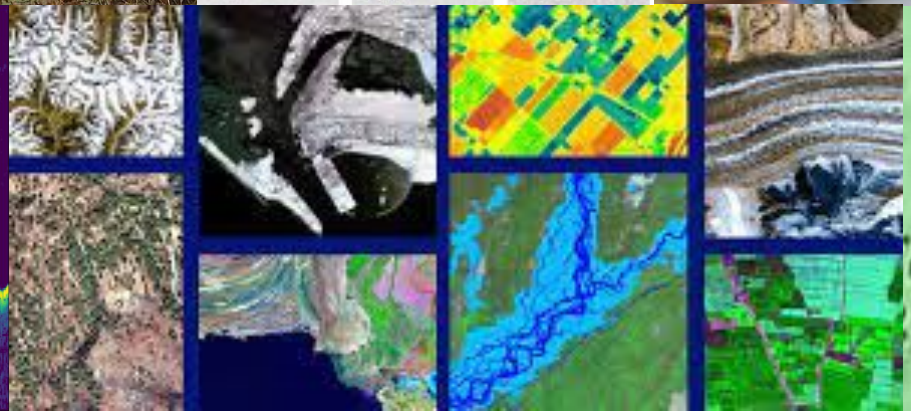
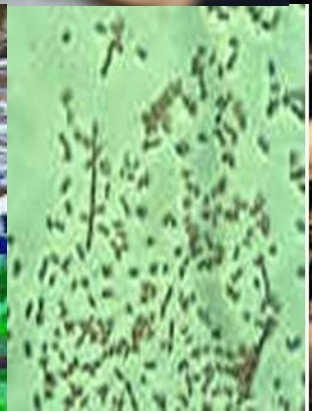


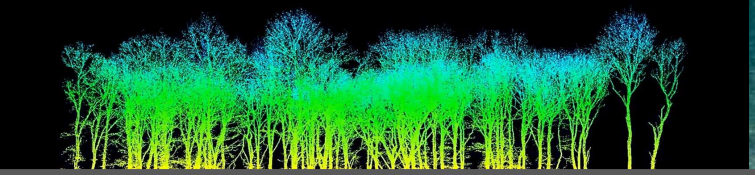

No direct sensor for species across taxa & scale


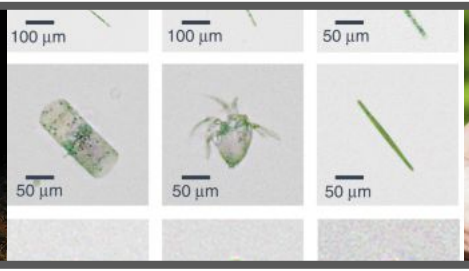


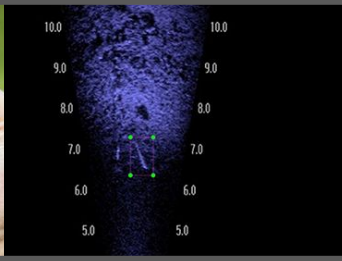

AI tools are necessary to help translate raw data to actionable scientific observations

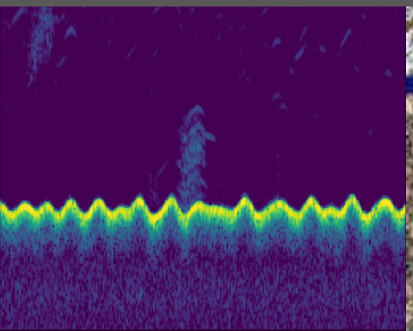
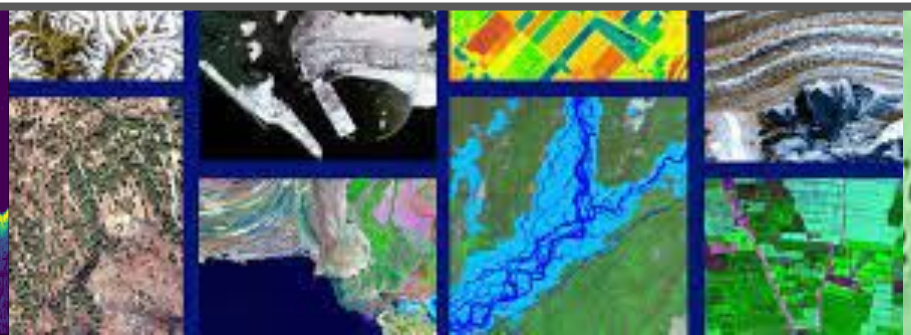
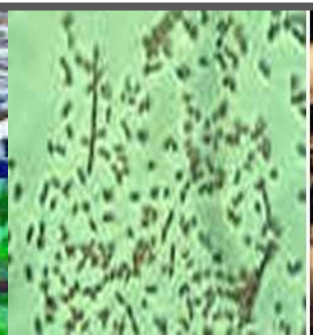


AI can:

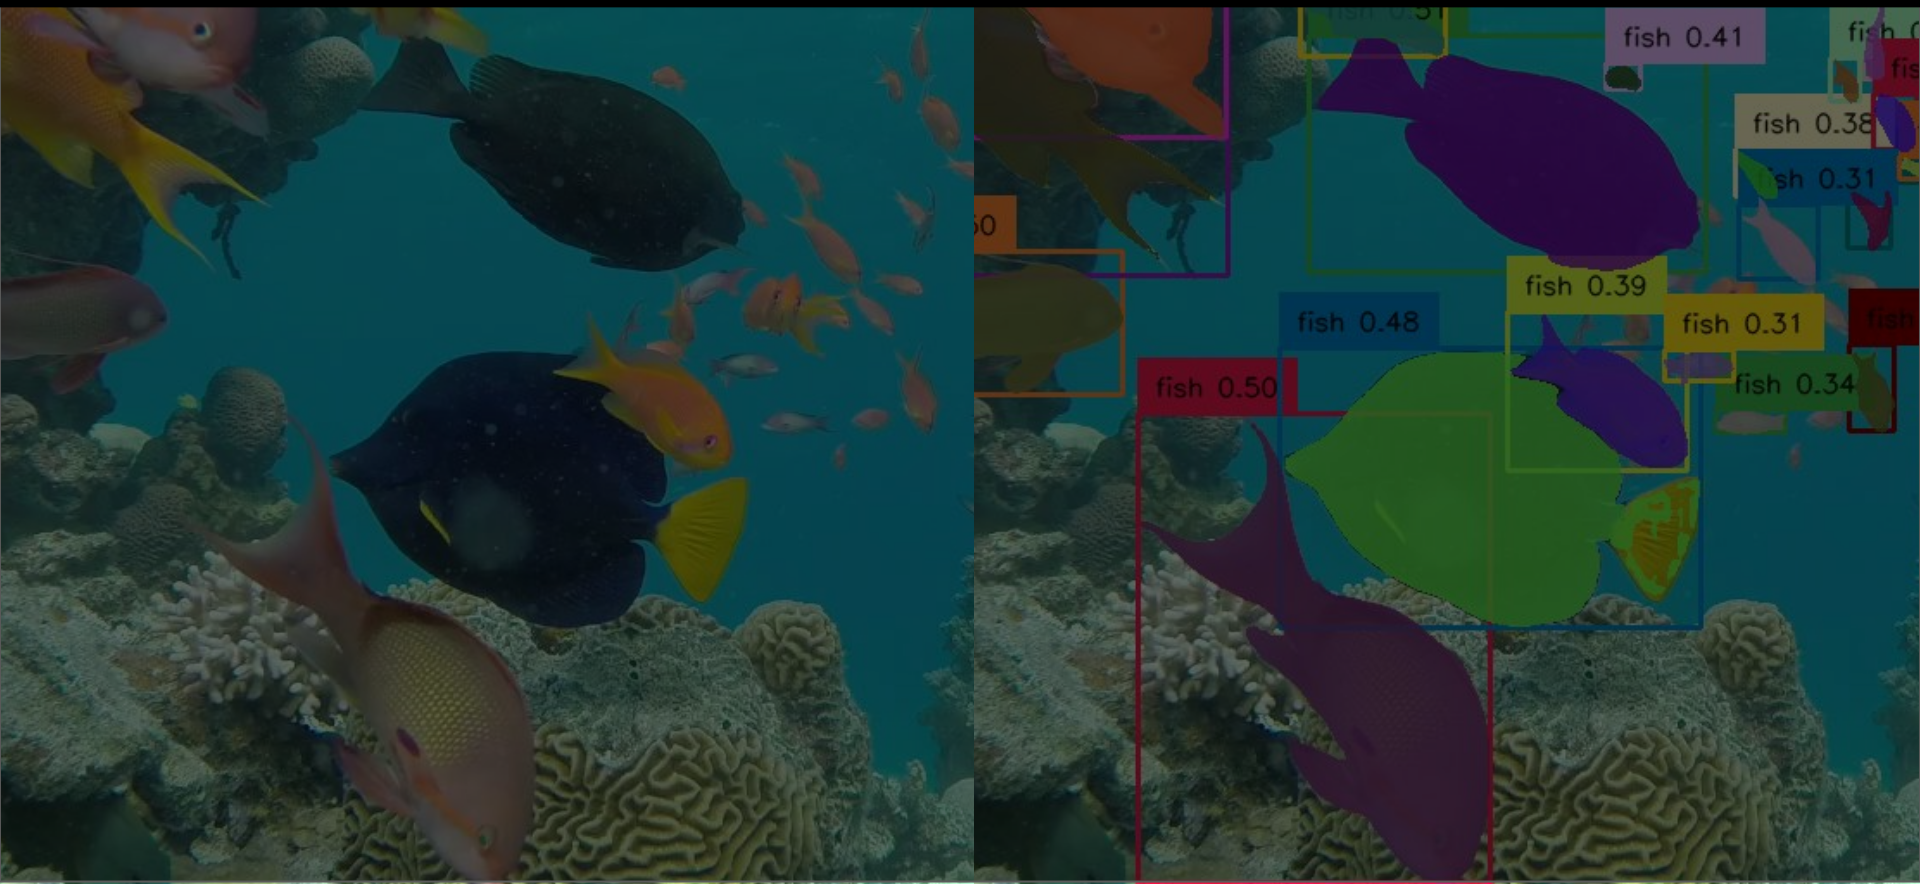

AI can:

- Recognize species

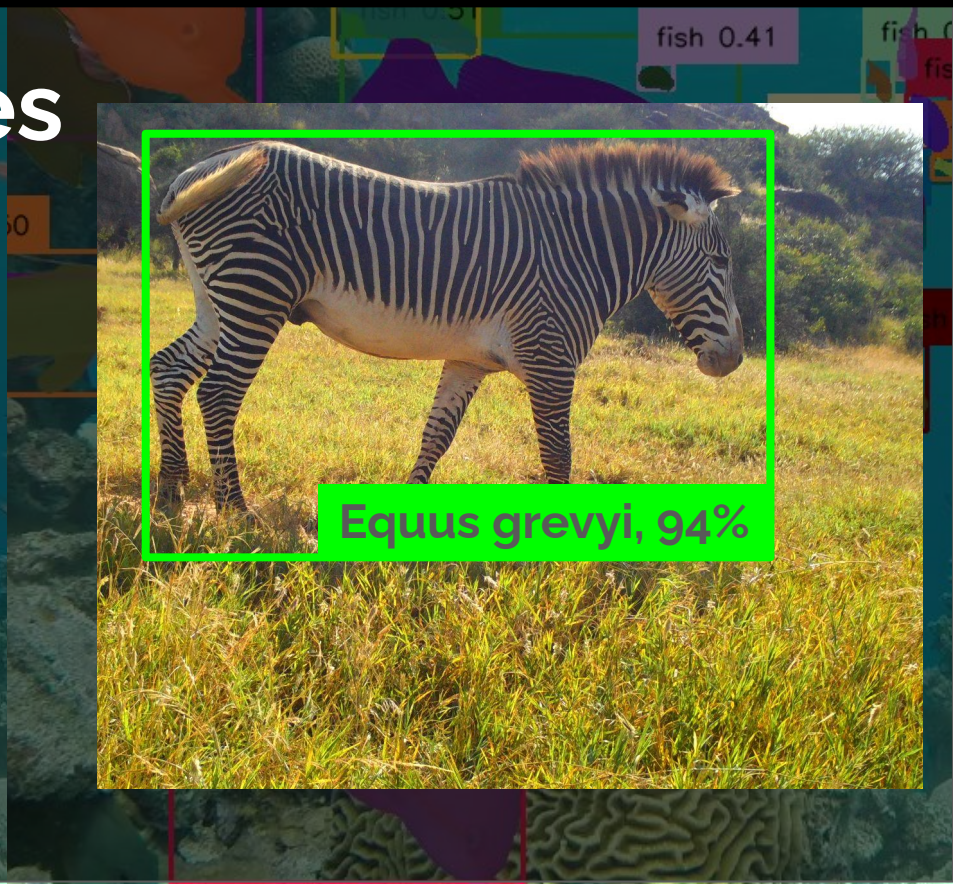

AI can:

- Recognize species
- Identify individuals

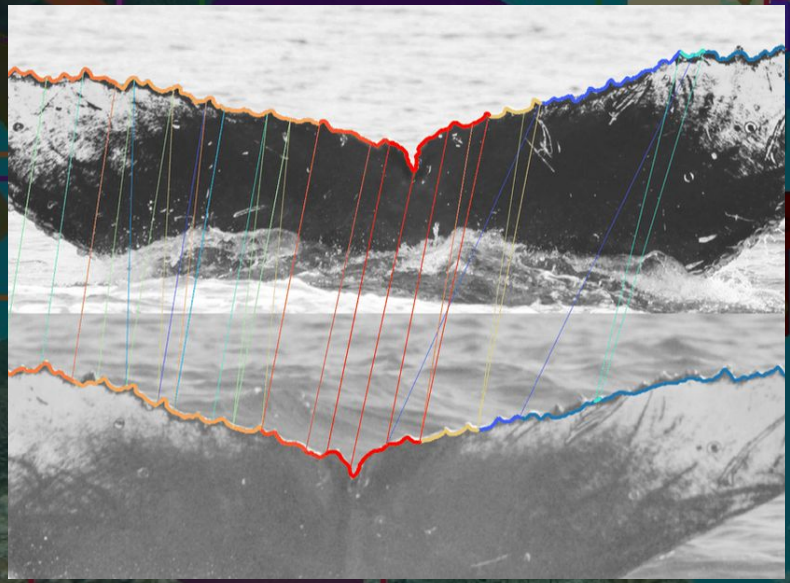

AI can:

- Recognize species
- Identify individuals
- Count large groups

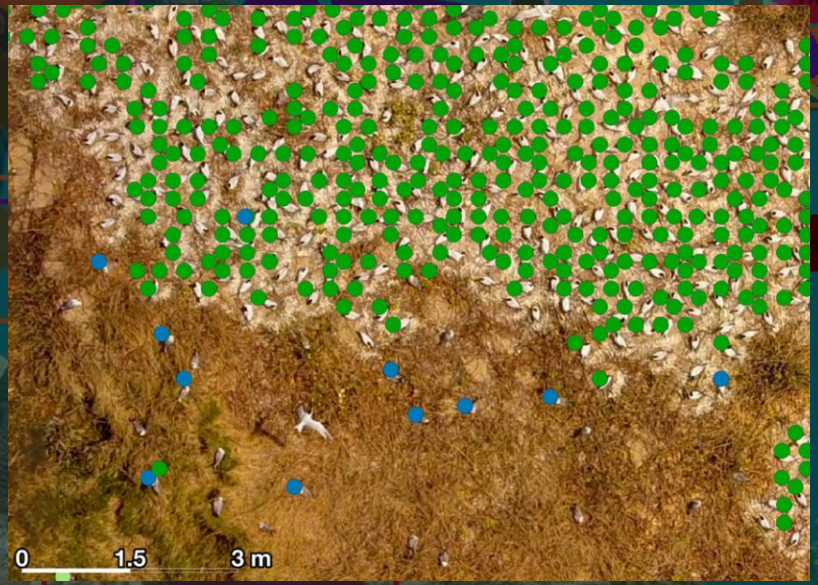

Kellenberger et al., 2021

AI can:

- Recognize species
- Identify individuals
- Count large groups
- Analyze behavior

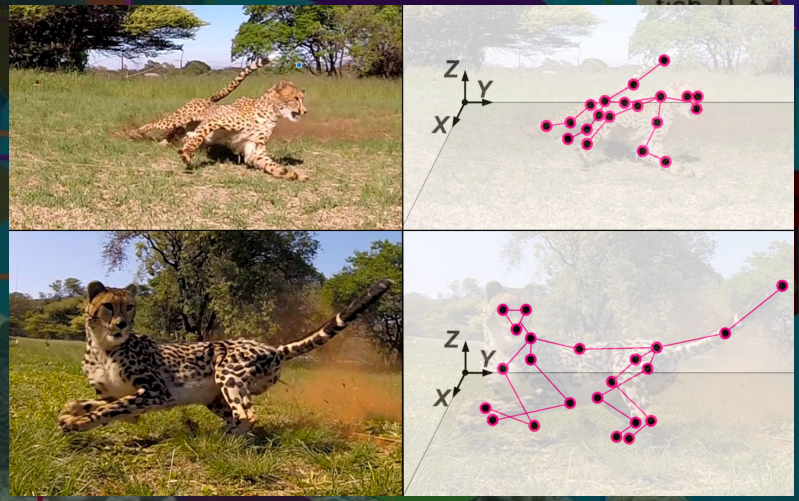

Joska et al., 2021

AI can:

- Recognize species
- Identify individuals
- Count large groups
- Analyze behavior
- Monitor environment

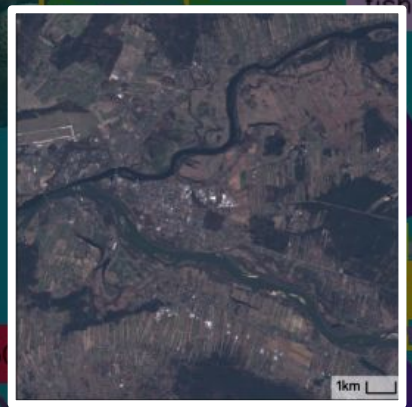

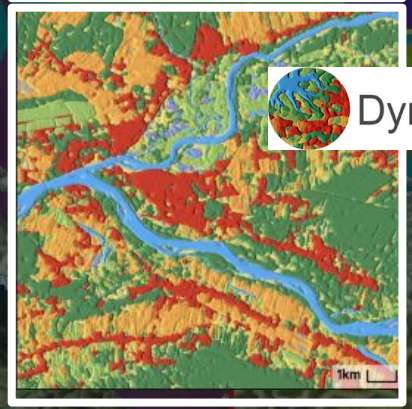

AI can:

- Recognize species
- Identify individuals
- Count large groups
- Analyze behavior
- Monitor environment
- Measure traits

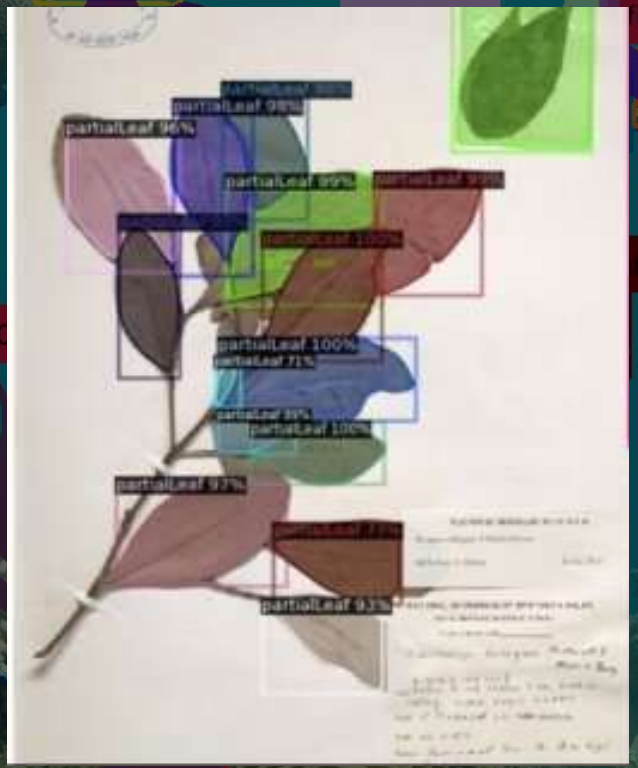

Wilde et al., 2023

AI is used to process data for thousands of conservation organizations globally

The MegaDetector

animal | empty | human | vehicle

Idaho Dept. of Fish and Game

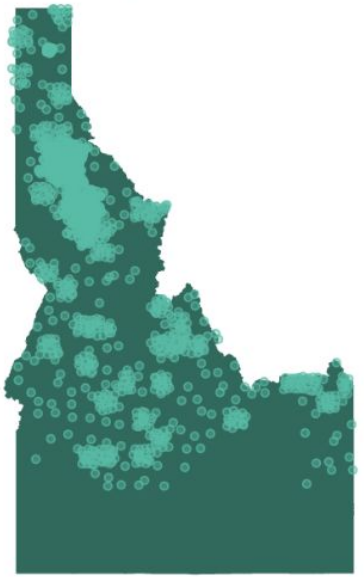

WOLF
pop. mgmt

2,000
cameras

11M
images

Less than 15% of images require human review

Biodiversity data is noisy

Objects of interest partially observed.

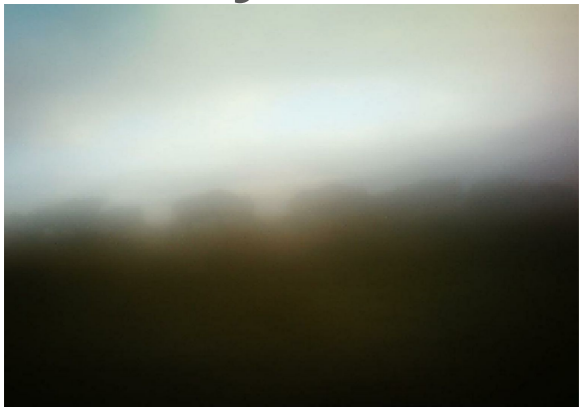

Poor data quality.

Empty data.

Biodiversity data has a long tail

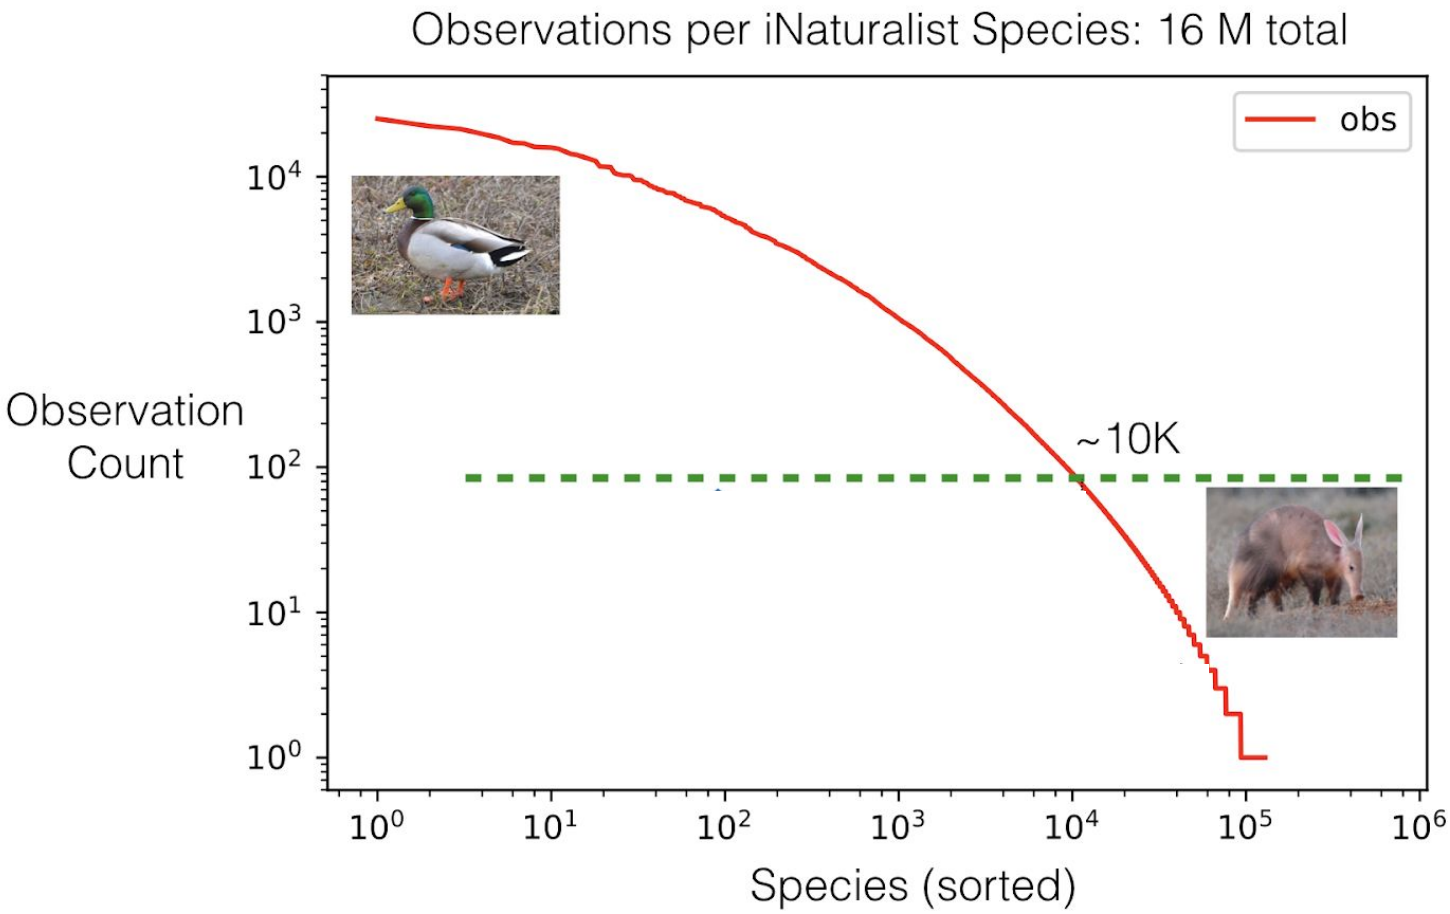

Biodiversity data is not IID

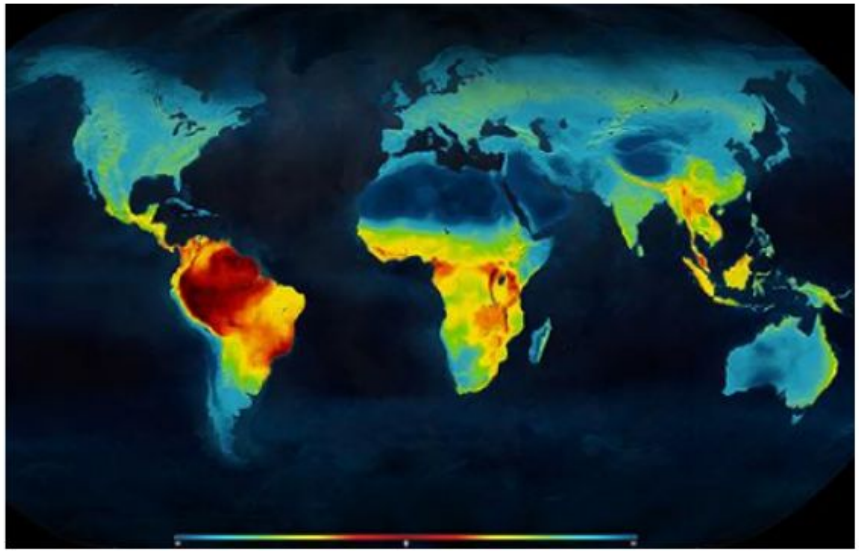

Map of global biodiversity

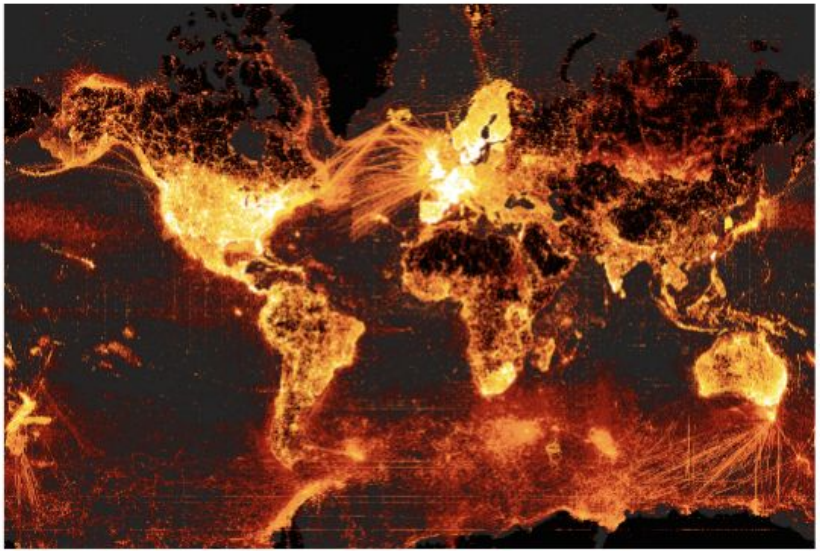

Species occurrence data in GBIF

Scaling Biodiversity Monitoring for the Data Age - ACM XRDS 2021


Each modality captures complementary but biased aspects of the taxonomic tree

e.g. camera trap PIR detection rates vary per-species based on size and temperature

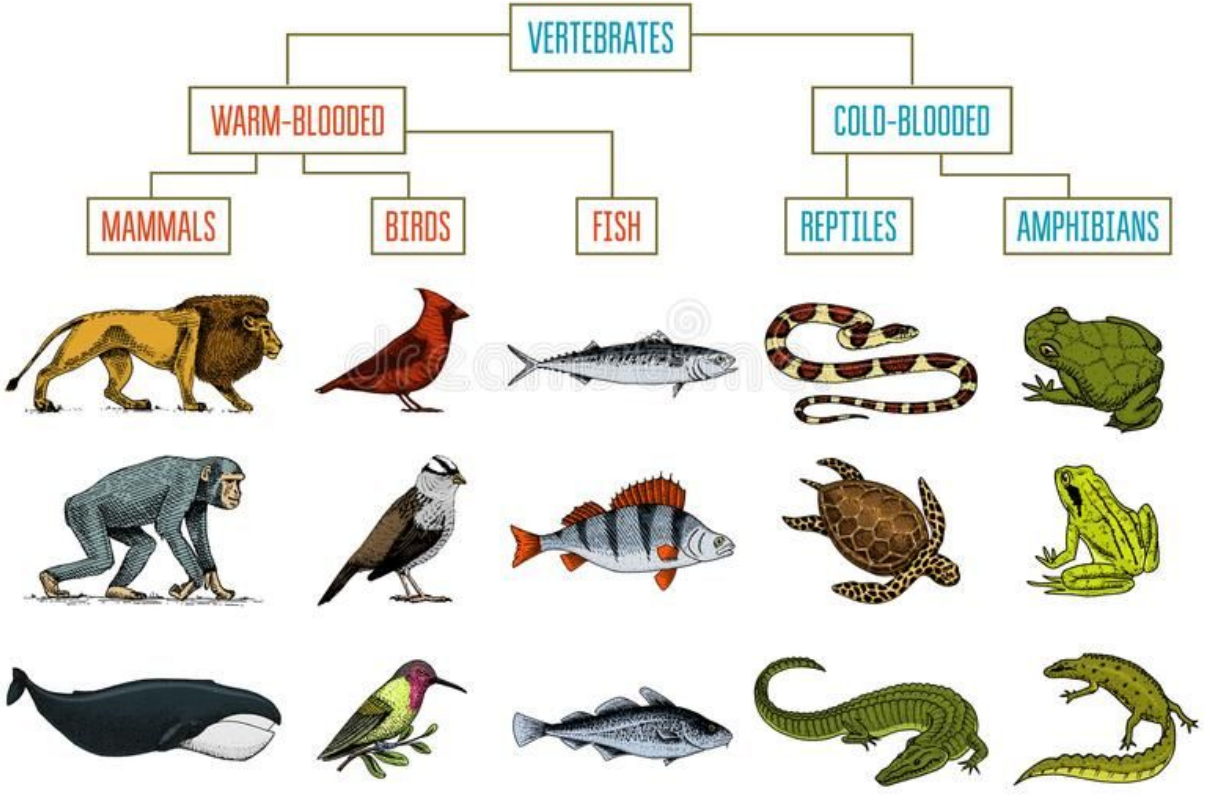

Detecting individual tree crowns

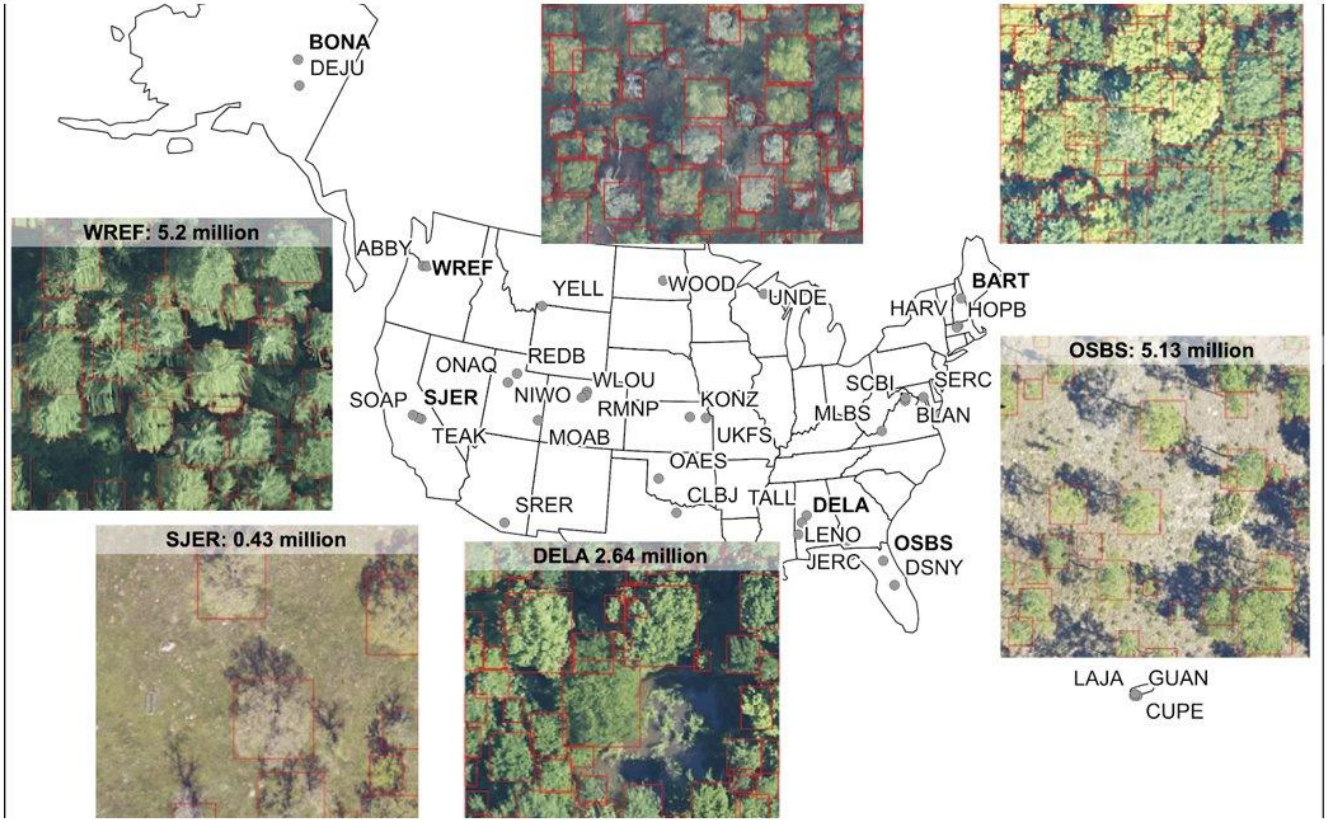

NEONCROWNS Dataset

104,675,304 trees

http://visualize.idtrees.org/

Weinstein et al., 2020

Detecting and counting salmon in static sonar

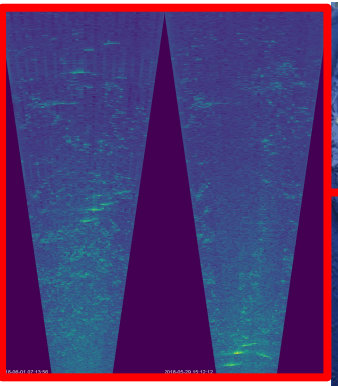

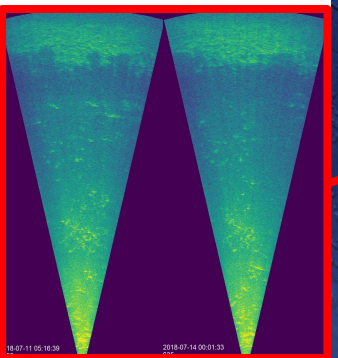

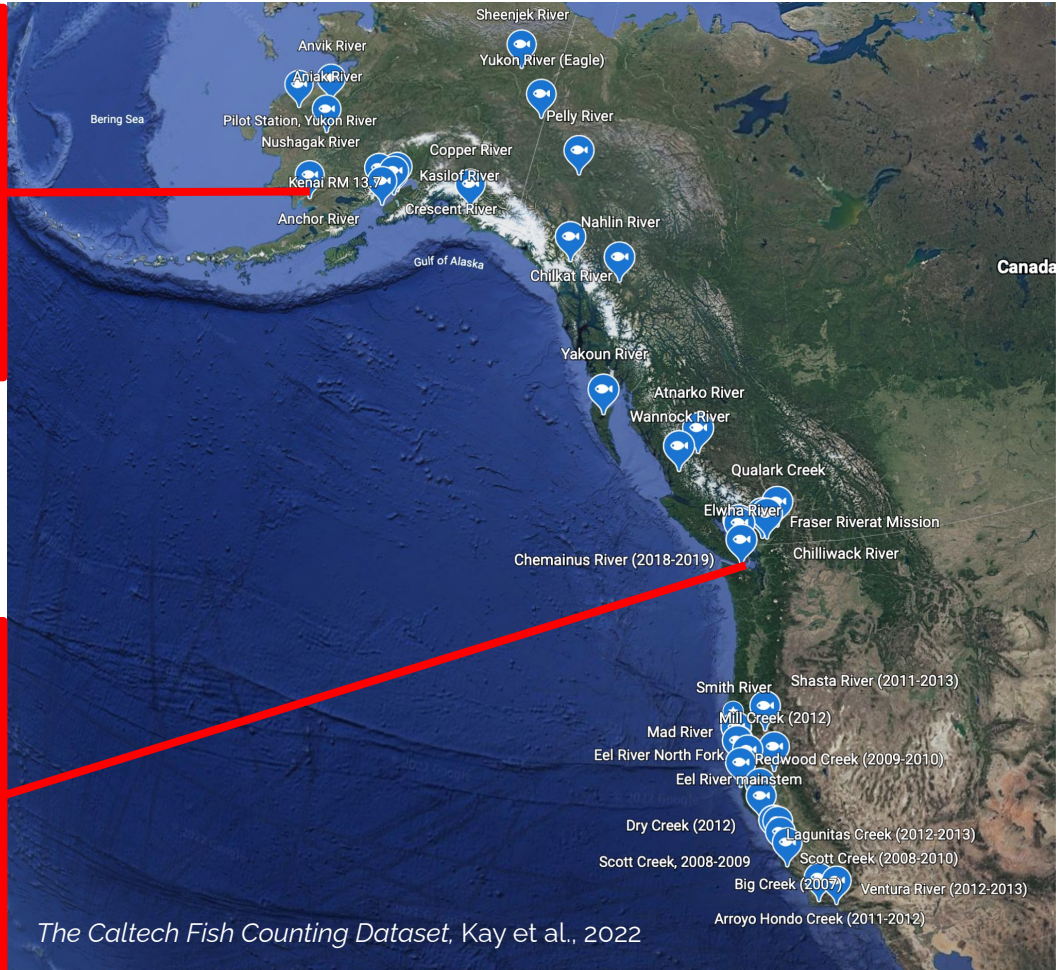

The Caltech Fish Counting Dataset, Kay et al., 2022

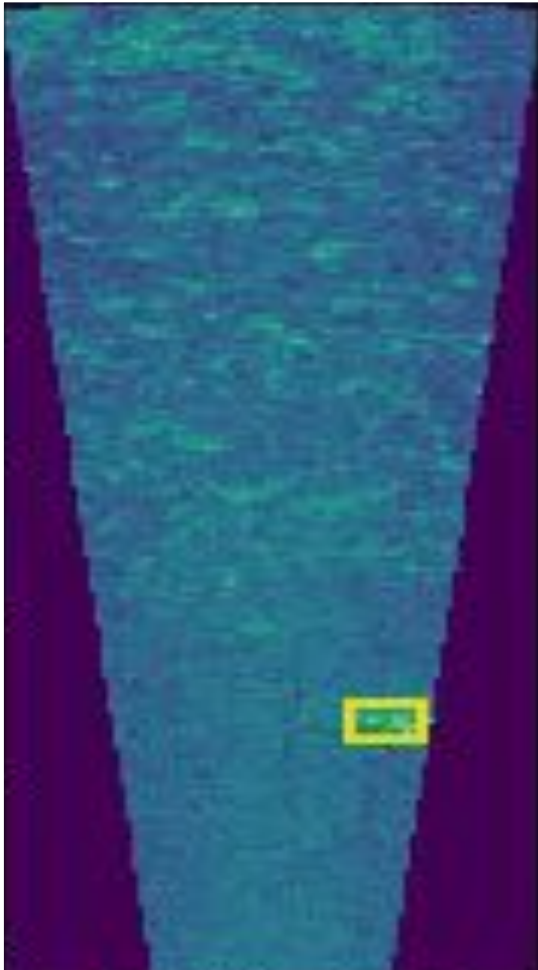

Static Cameras

Each static camera has a distinctive background and class distribution

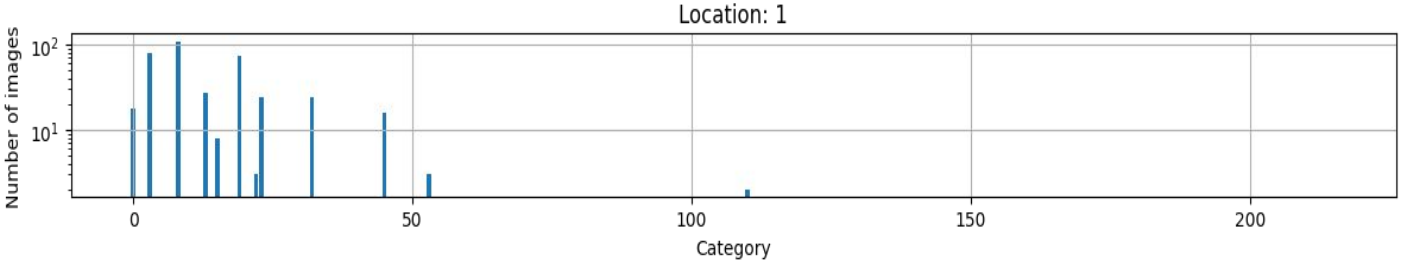

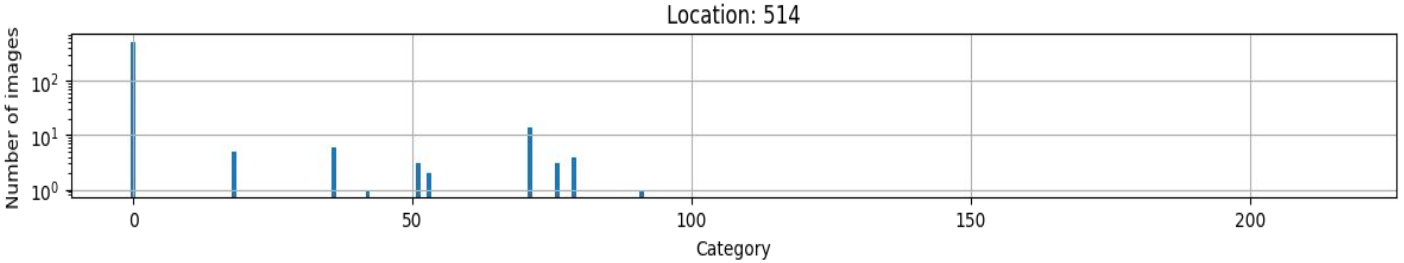

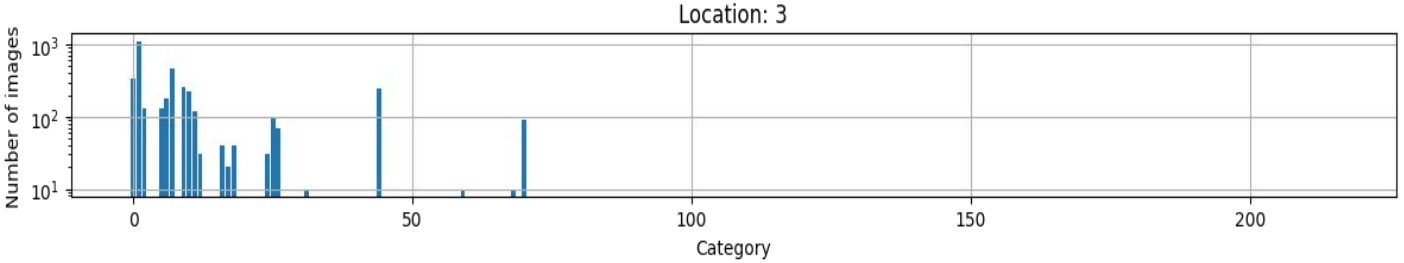

Models don't generalize

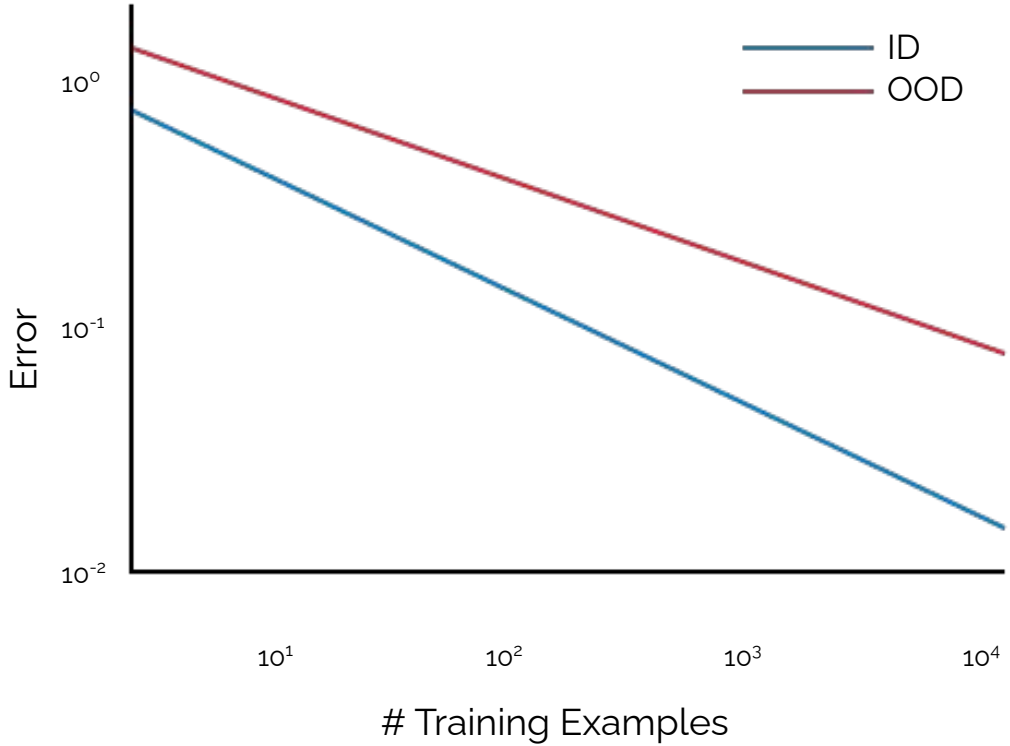

Recognition in Terra Incognita, Beery et al., ECCV 2018

Distribution shifts are ubiquitous in real-world

WILDS

https://wilds.stanford.edu/

Pang Wei Koh*, Shiori Sagawa*, Henrik Marklund, Sang Michael Xie, Marvin Zhang,
Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Sara Beery,
Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang

	Camelyon17	iWildCam	PovertyMap	FMoW	Amazon	CivilComments	OGB-MolPCBA
Shift	Hospitals	Locations	Countries	Time	Users	Demographics	Scaffold
Train	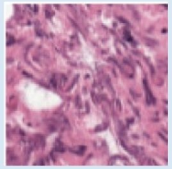				Overall a solid package that has a good quality of construction for the price.	What do Black and LGBT people have to do with bicycle licensing?	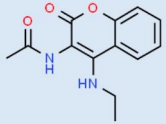
Test	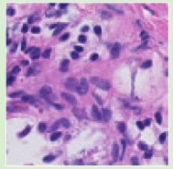				I *loved* my French press, it's so perfect and came with all this fun stuff!	As a Christian, I will not be patronizing any of those businesses.	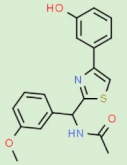
Adapted from	Bandi et al. 2018	Beery et al. 2020	Yeh et al. 2020	Christie et al. 2018	Ni et al. 2019	Borkan et al. 2019	Hu et al. 2020

Good performance on realistic benchmarks ≠ impact

Model

Good performance on realistic benchmarks ≠ impact

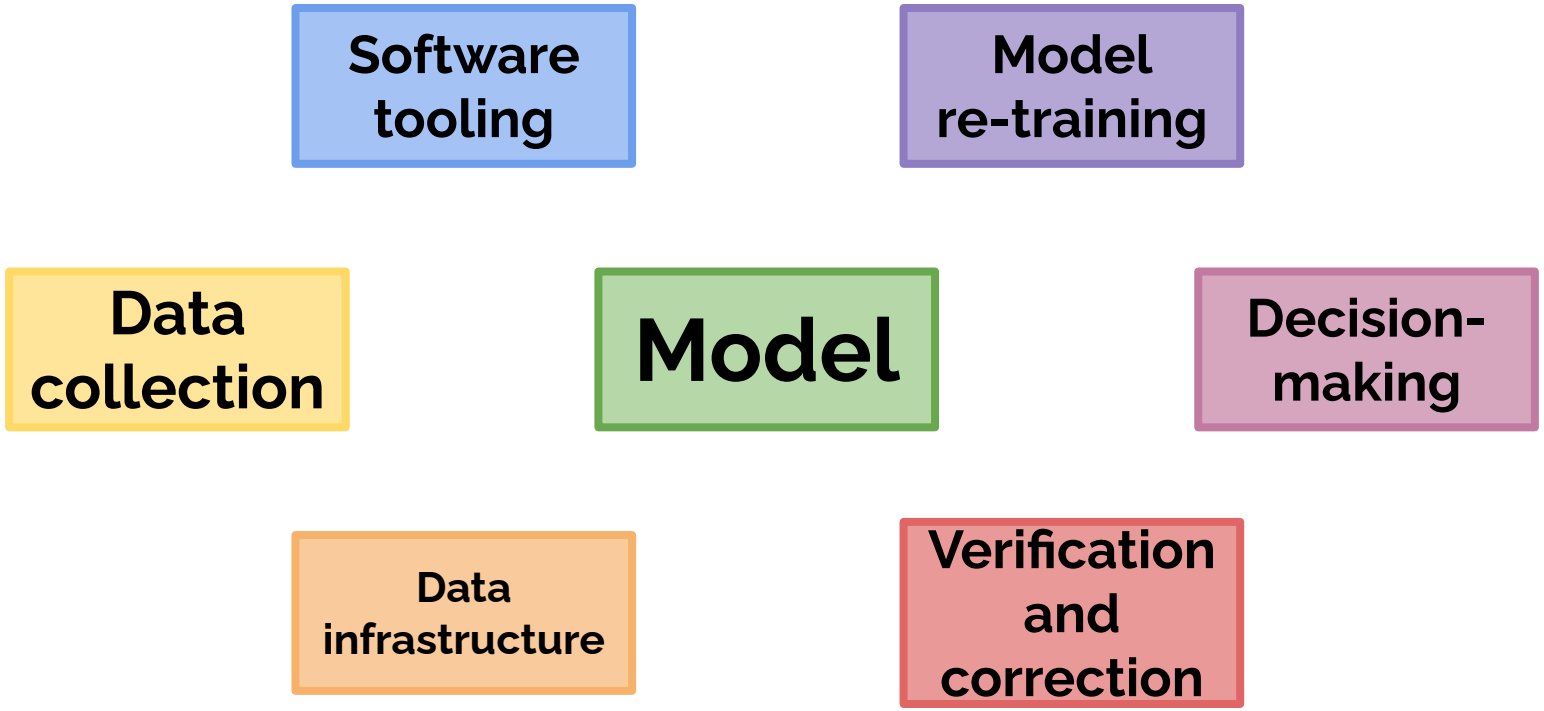

Good performance on realistic benchmarks ≠ impact

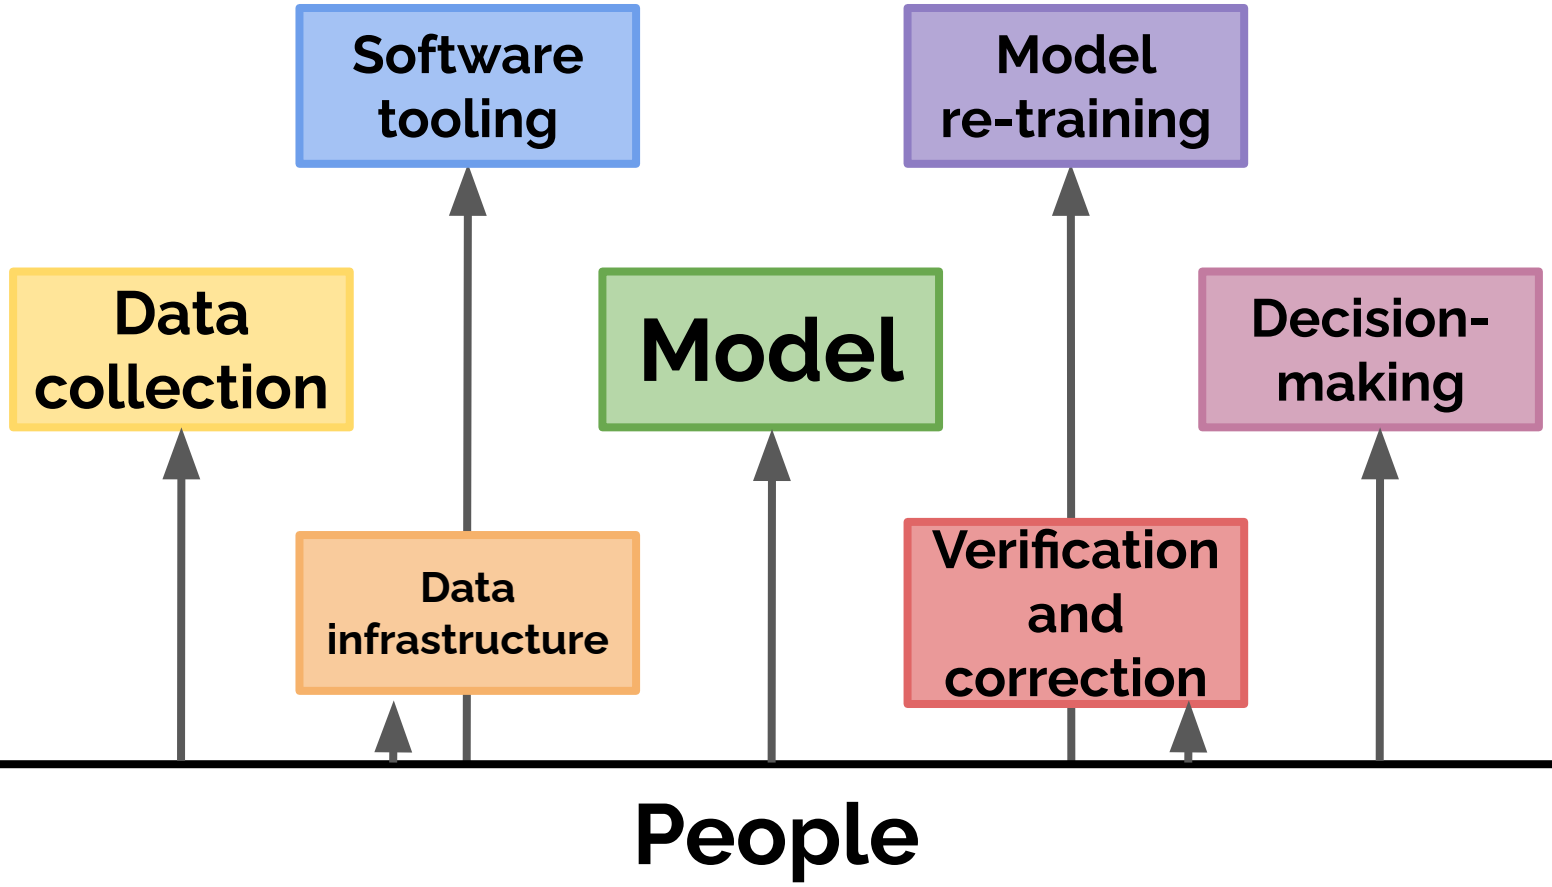

Impactful AI systems are:

- Useful (not perfect!)
- Accessible
- Collaborative
- Well-communicated

Let's look at this with the MegaDetector:

https://github.com/agentmorris/MegaDetector

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

Thanks Aaron Greenville for the sweet gif!

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

Useful: used to process data for hundreds of NGOs, agencies, and conservation organizations globally

Idaho Dept. of Fish and Game

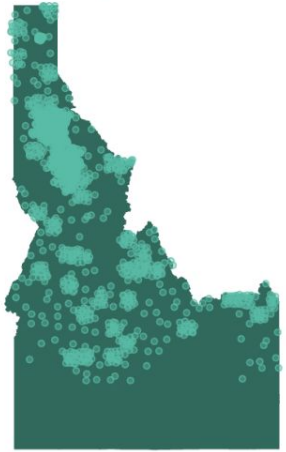

WOLF
pop. mgmt

2,000
cameras

11M
images

The MegaDetector

animal | empty | human | vehicle

Less than 15% of images require human review

human review

alert

Wildlife Protection Solutions

WILDLIFE CRIME PREVENTION

18 nations | 800 cameras | 900K images

Real-time alerts

Detects one real wildlife threat per week on average

Use of object detection in camera trap image identification: assessing a method to rapidly and accurately classify human and animal detections for research and application in recreation ecology

Mitchell Fennell, Christopher Beirne, A. Cole Burton

“In our application, MegaDetector detected human and animal images with 99% and 82% precision, and 95% and 92% recall respectively, at a confidence threshold of 90%. **The overall time required to process the dataset was reduced by over 500%, and the manual processing component was reduced by 840%.** The index of human detection events from MegaDetector matched the output from manual classification, with a mean 0.45% difference in estimated human detections across site-weeks.”

Accessible: Hosted via an open-source API, and integrated into existing tools

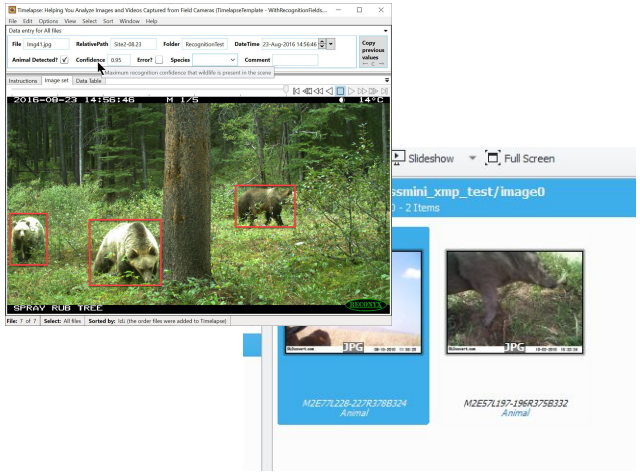

Desktop
(Timelapse, digiKam)

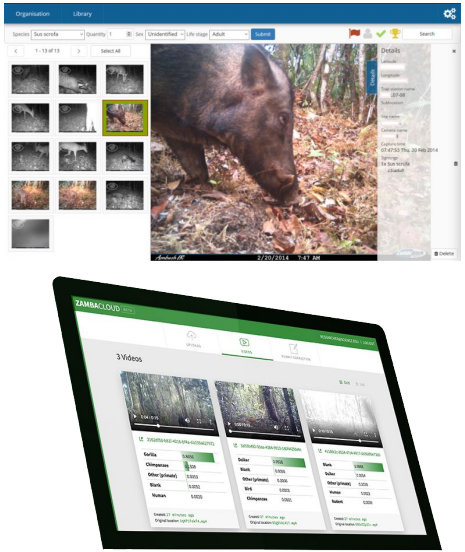

Cloud
(Camelot, Zamba Cloud)

It's complicated
(Zooniverse, eMammal)

Collaborative: Clear path for feedback, iterative improvements

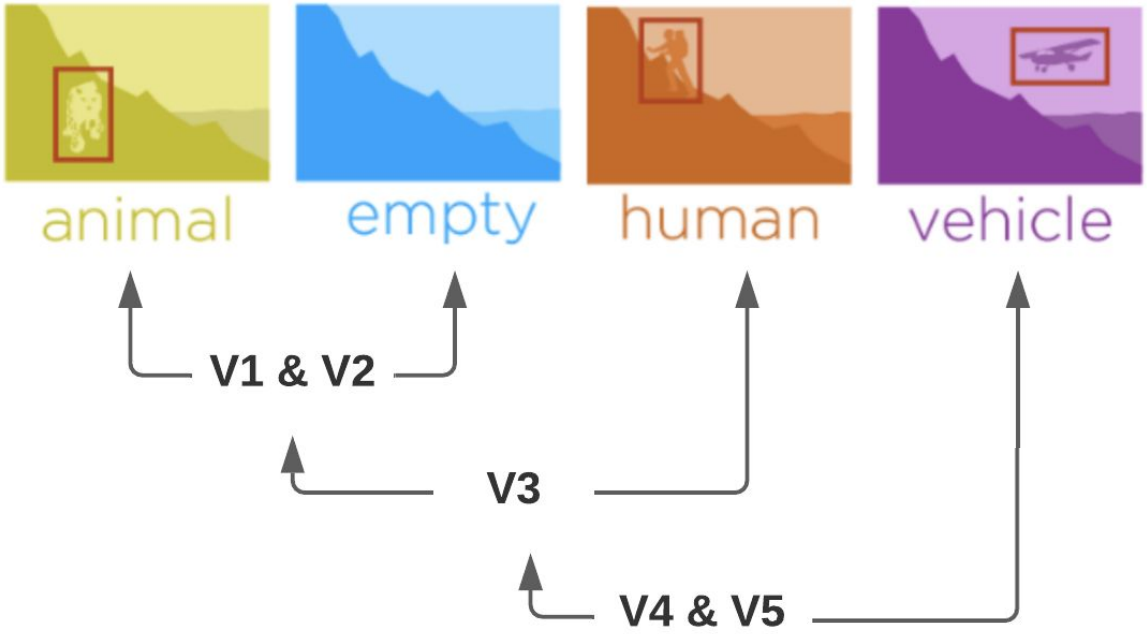

Well-communicated: define risks, known failure modes, and best practices for validation and use

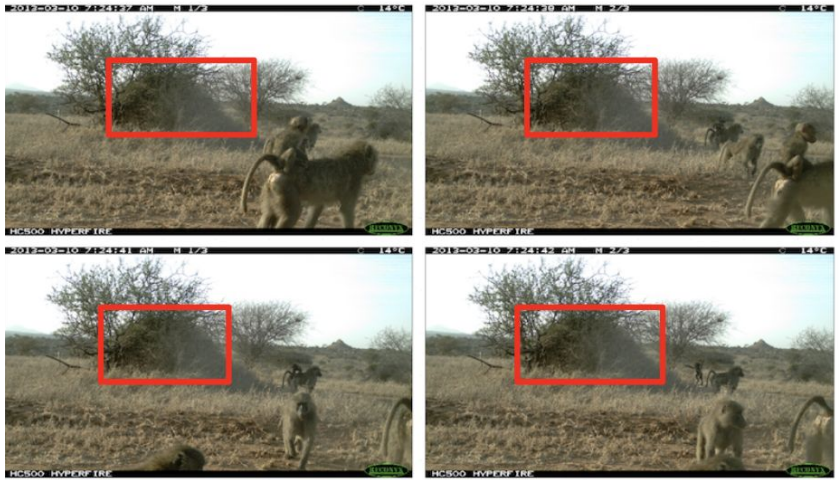

Remove salient & static false positives

Sorting == accelerating, i.e. annotating this:

...is faster than annotating this:

Use humans to catch errors efficiently

Well-communicated: define risks, known failure modes, and best practices for validation and use

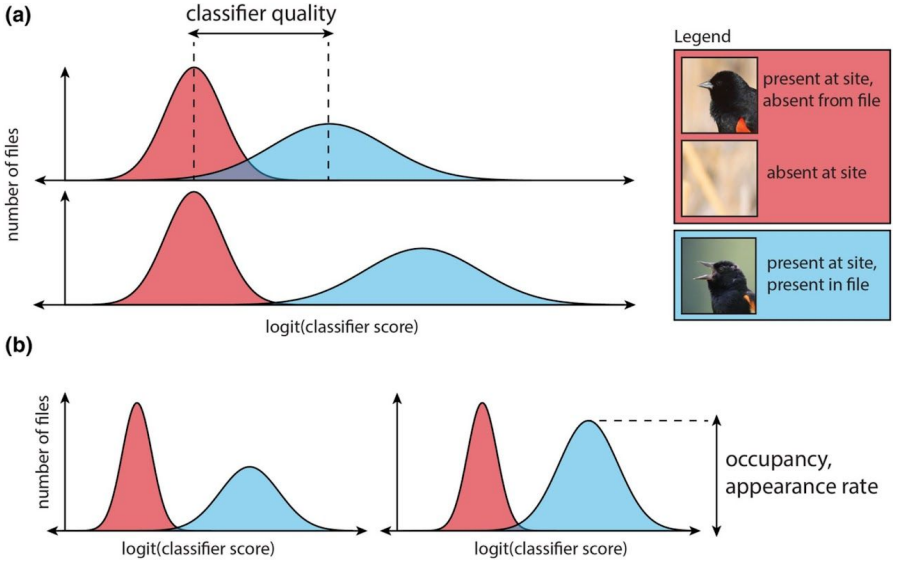

Select a detection threshold based on risk and resources

Account for false positives and negatives in downstream modeling

Rhinehart et al.

In this course we will systematically explore these challenges and methods to address them

Course overview

Meets T/Th, 9:30-11am in 32-124

Short introductory lectures each week, main content of course is focused around reading, presenting, and discussing foundational and recent papers at the intersection of CV and planetary health

Course staff

Instructor Sara Beery
beery at mit dot edu

TF Justin Kay
kayj at mit dot edu

Office hours TBD,
We will each hold one per week

Important links

Webpage: https://cv4planet.github.io/
Syllabus, slides, reading assignments

Piazza: https://piazza.com/class/m6ic2i3cah83l7
Presentation sign-up link

Canvas: https://canvas.mit.edu/courses/31089/
Lecture recordings

Course Content

- One 30-min overview lecture per week introducing the topic
- There will be **2-5 required readings each week**
- Student presentations in class
  - Each class will consist of 2-3 paper presentations by student groups. Each group will consist of three students, each with a different role:
    - **Summarization:** Summarize the paper (5 minutes)
    - **Critique:** Discuss limitations of the paper (5 minutes)
    - **Extension:** Discuss 2-3 possible extensions of the paper (5 minutes)
- 10 minutes of additional class discussion per paper
- Group research proposal (details to be announced in March)

We reserve the right to update this as needed!

Grading Policy

- **60% class presentations**
  - 10% per role/paper (each student will present six times, twice in each role)
- **20% class participation**
- **20% final project**

Week 1

Course overview, introduction to the history and current use of CV in ecology and the environment

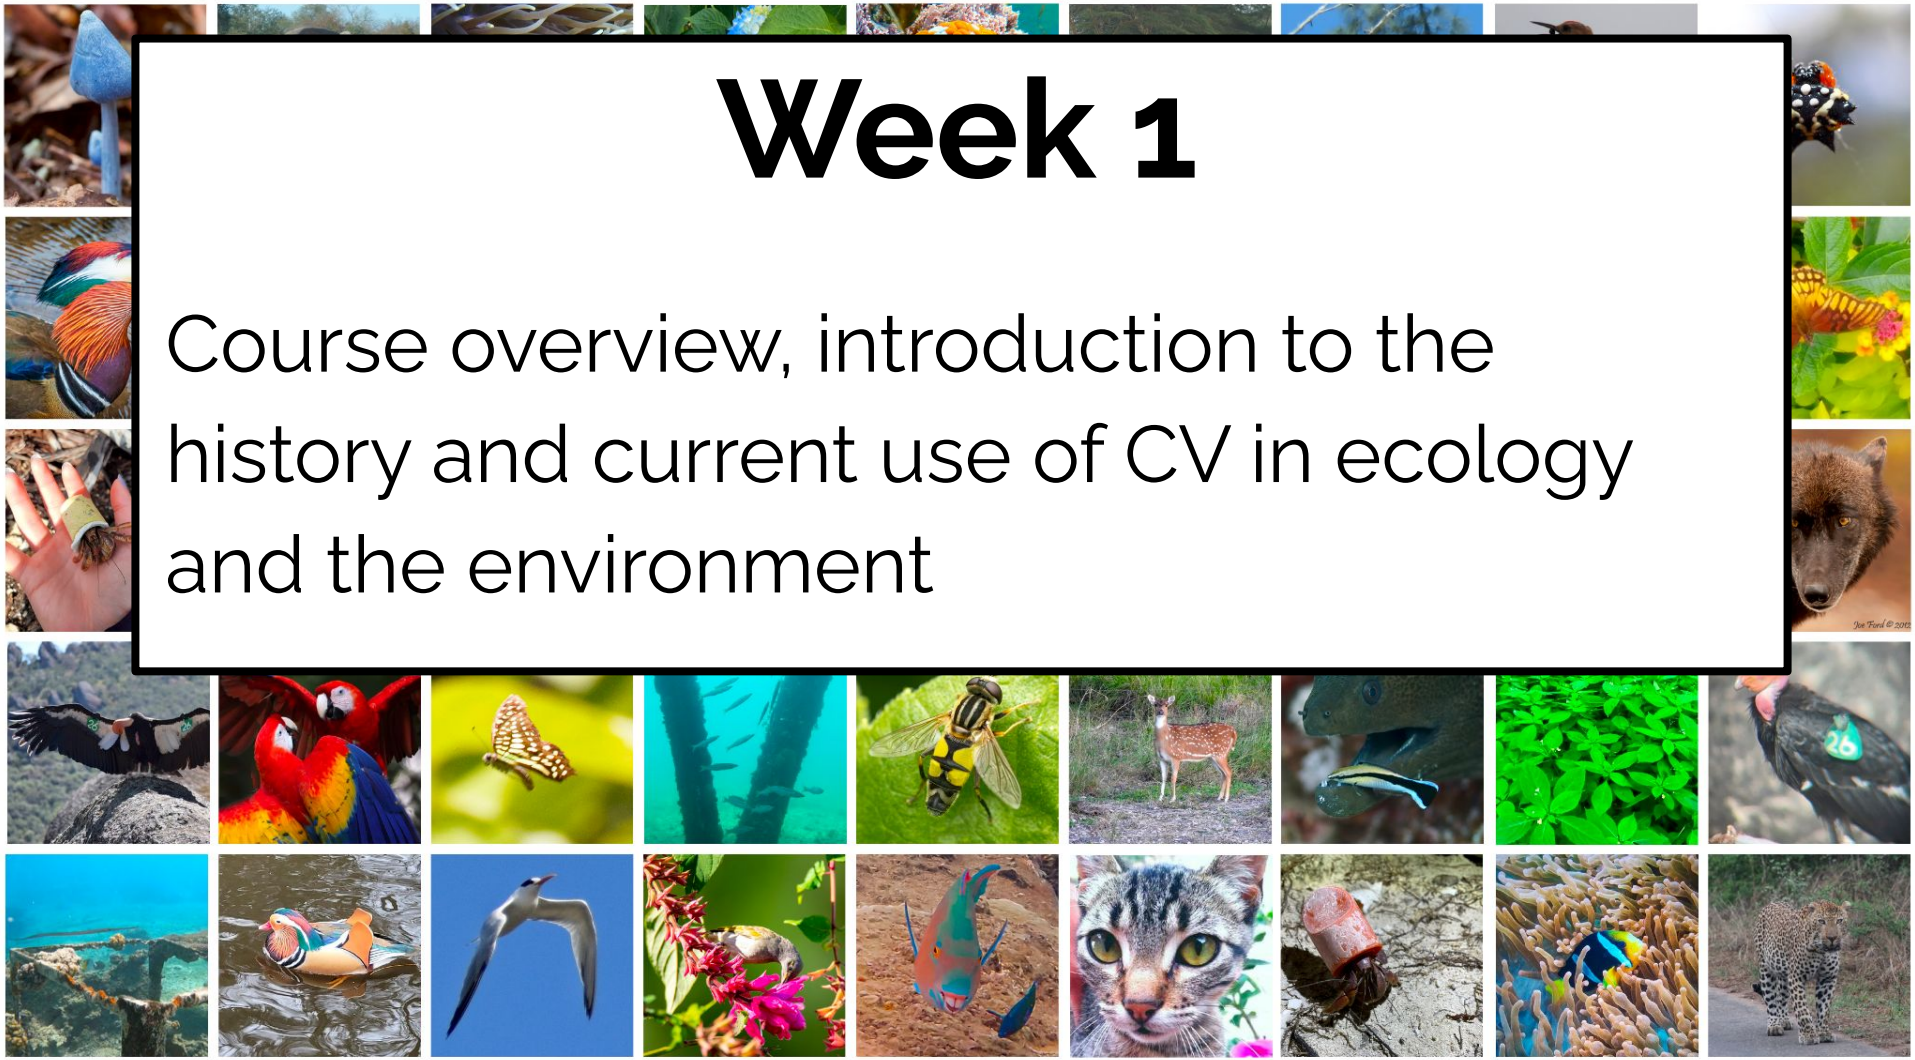

Week 2

Overview of planetary crises including climate change, biodiversity loss, and human-wildlife conflict, current global policies and goals including EBVs, 30 by 30, what is needed to make progress, how we measure

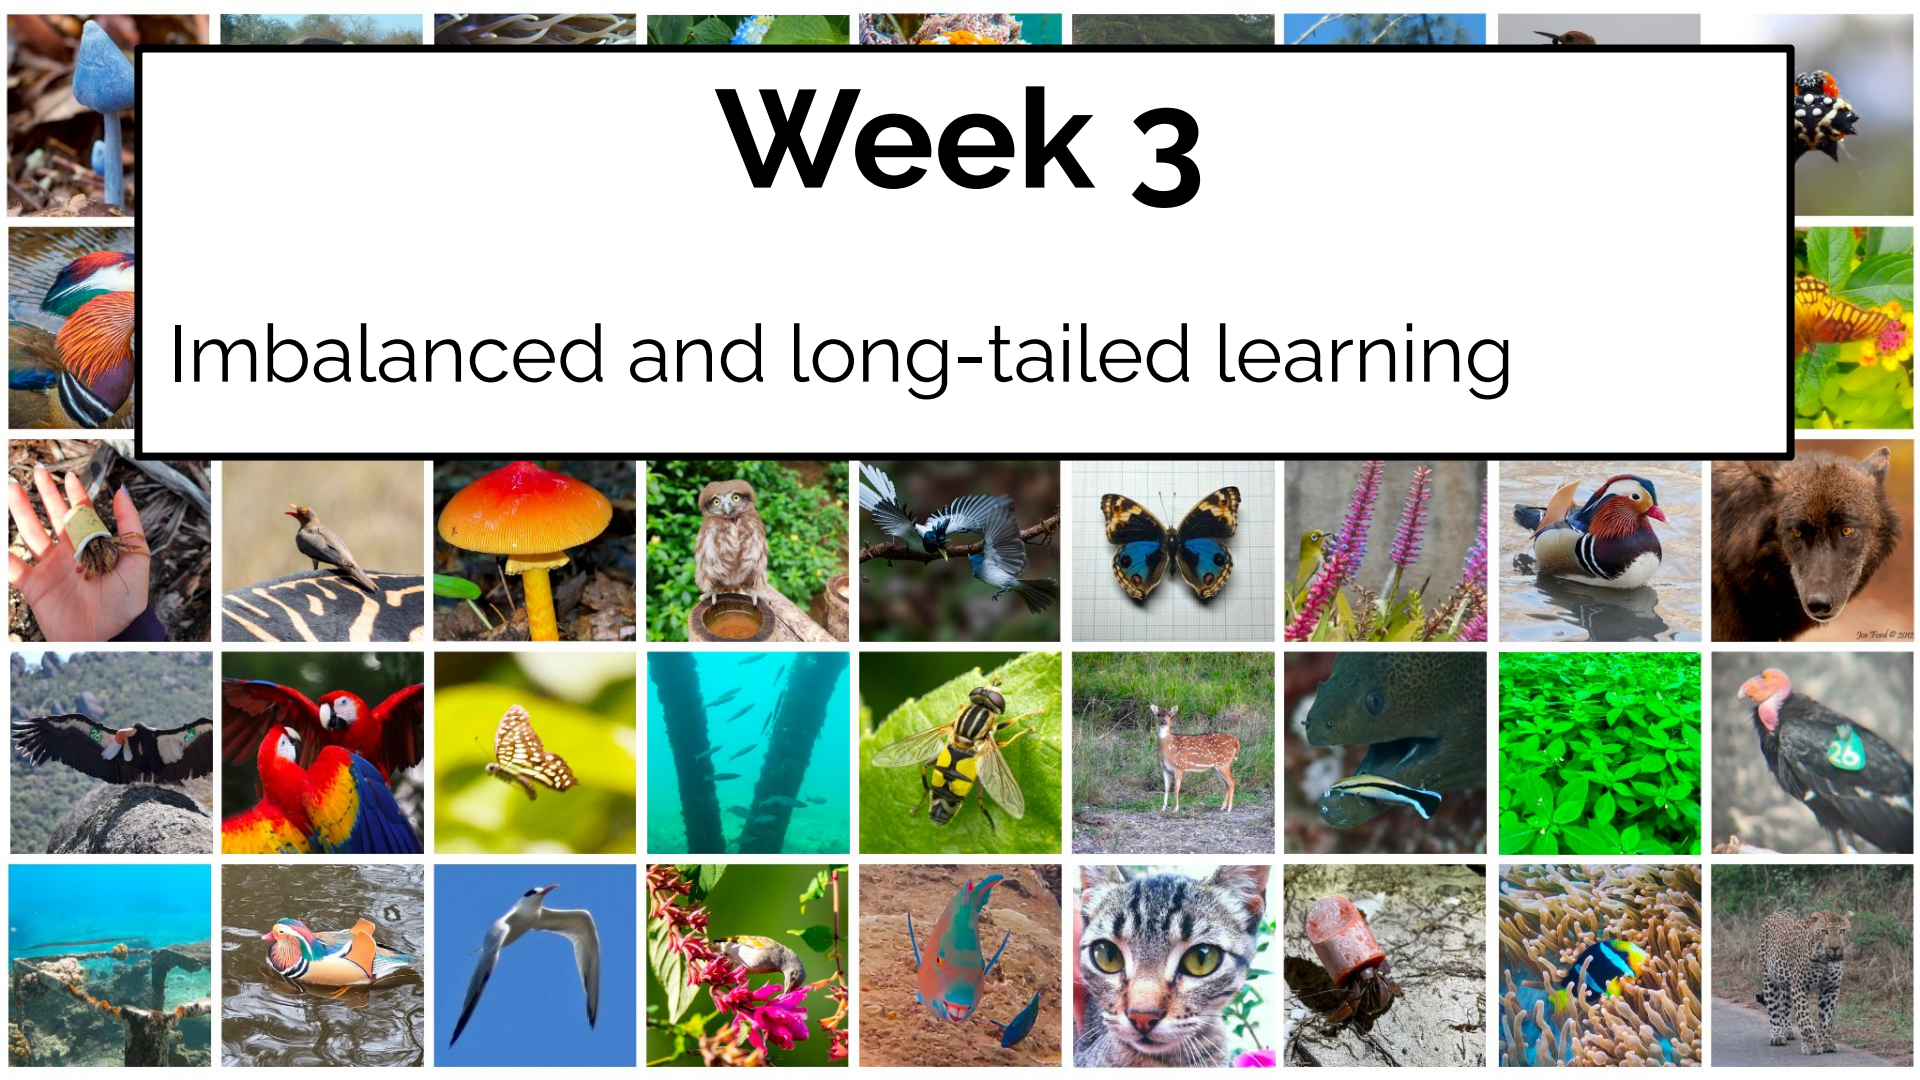

Week 3

Imbalanced and long-tailed learning

Week 4

Fine-grained recognition

Week 5

Open-set/open-world learning

Week 6

Distribution shift and distributional robustness

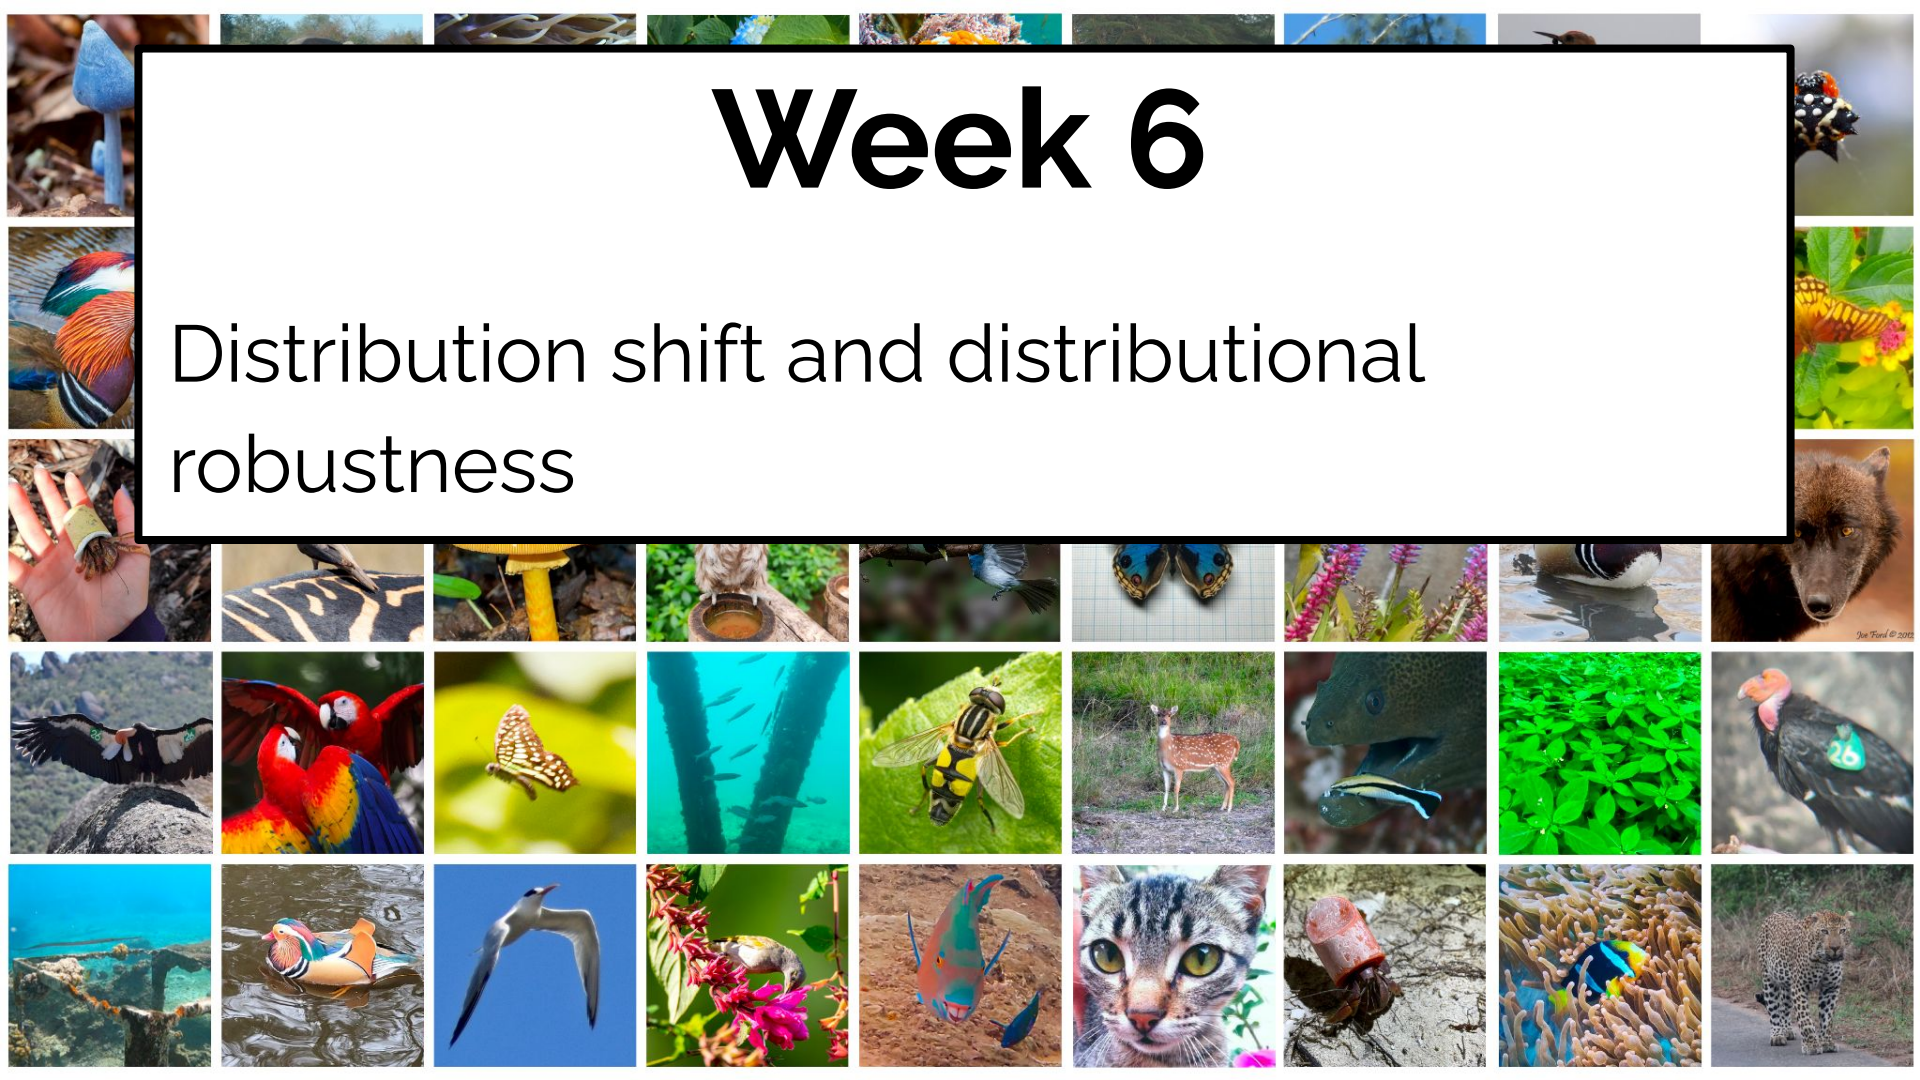

Week 7

Domain adaptation and specialization

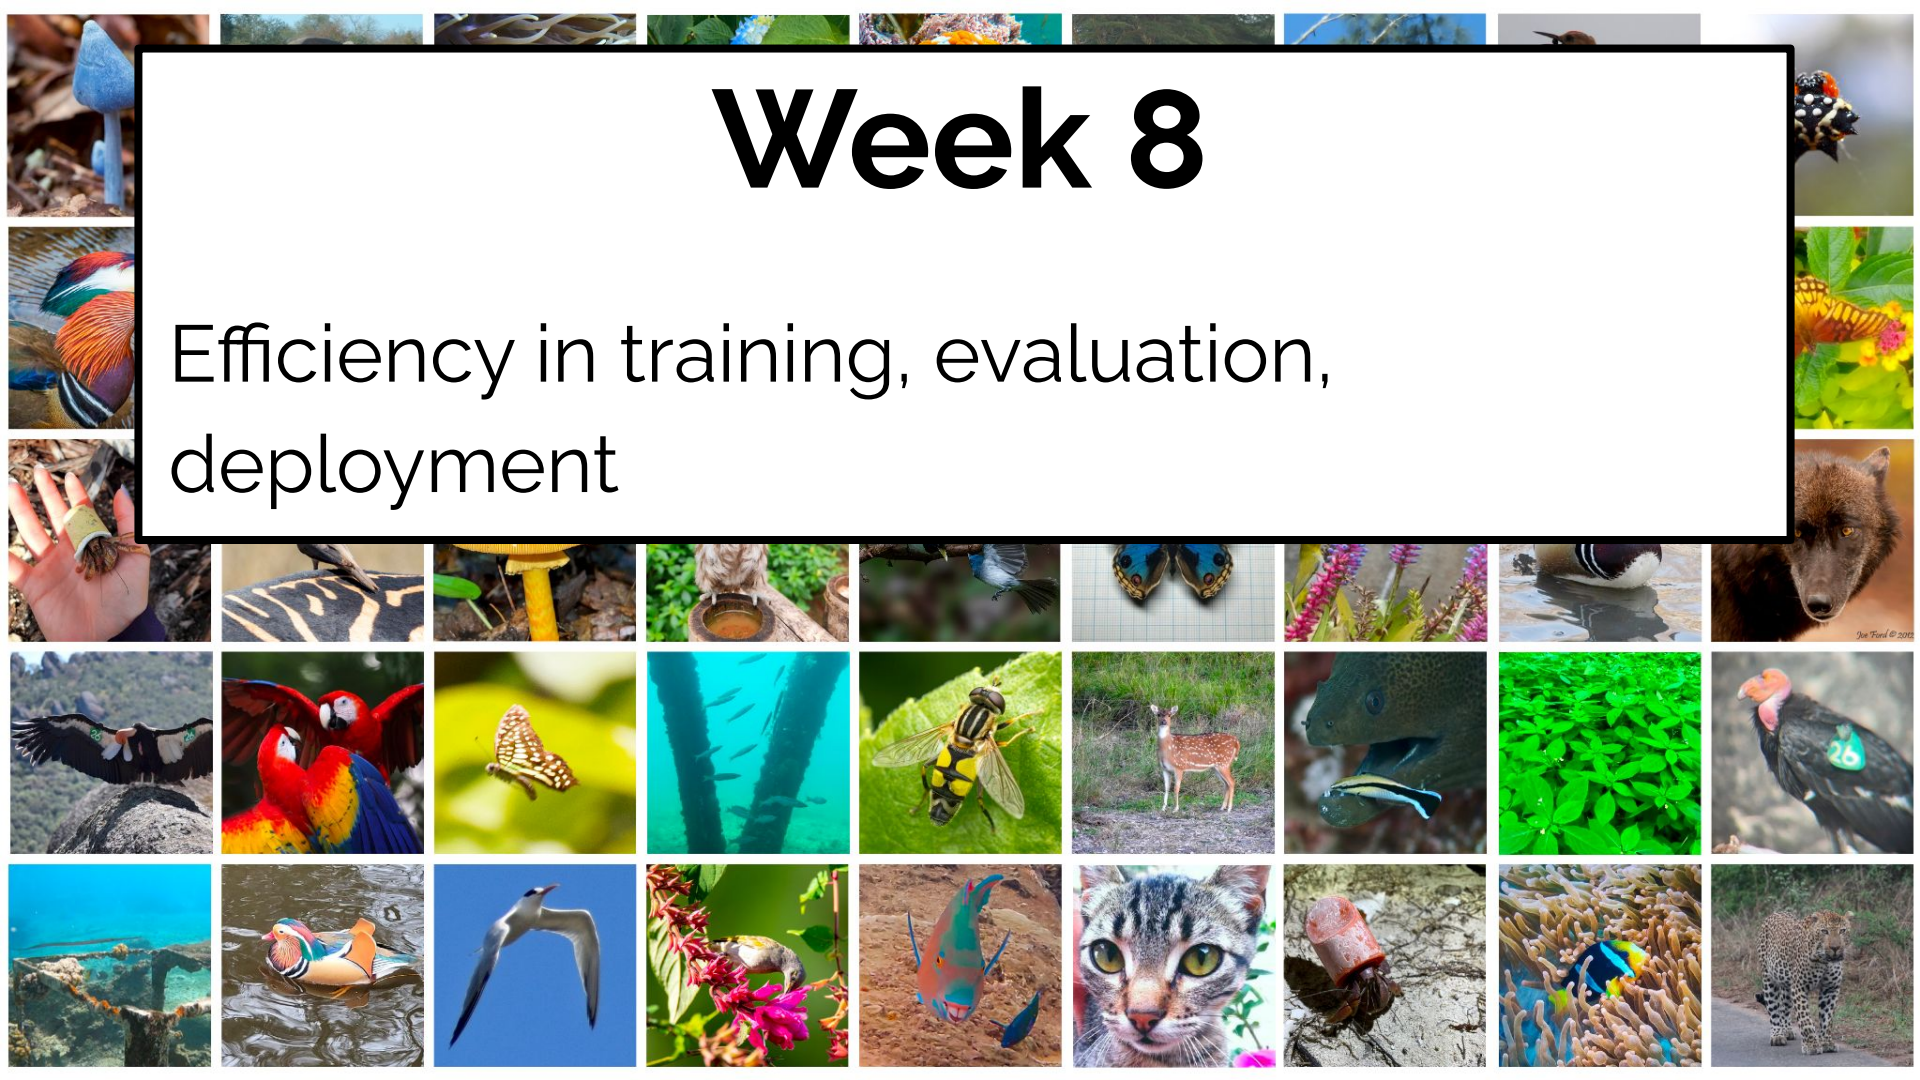

Week 8

Efficiency in training, evaluation, deployment

Week 9

Human-AI systems: Active learning

Week 10

Human-AI systems: Selective prediction

Week 11

Human-AI systems: Active inference and decision support

Week 12

Multimodality: Vision and language

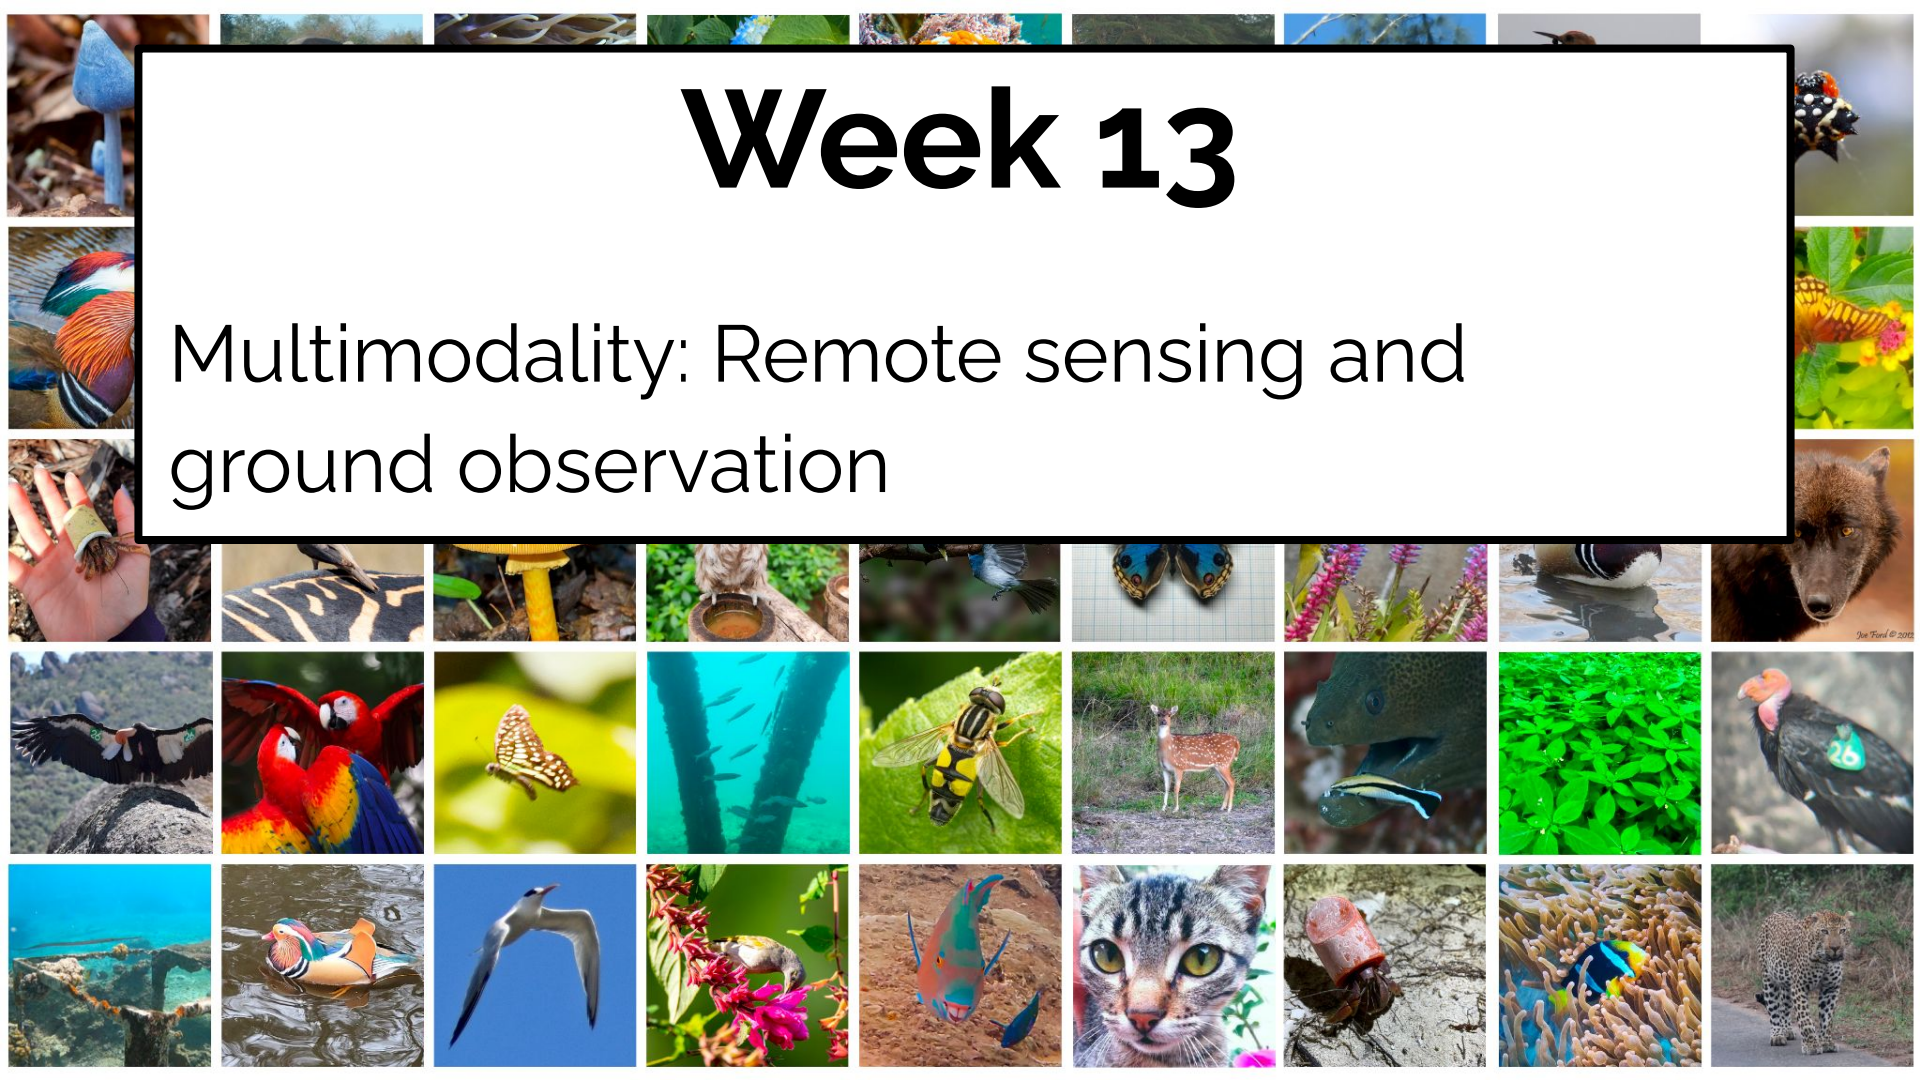

Week 13

Multimodality: Remote sensing and ground observation

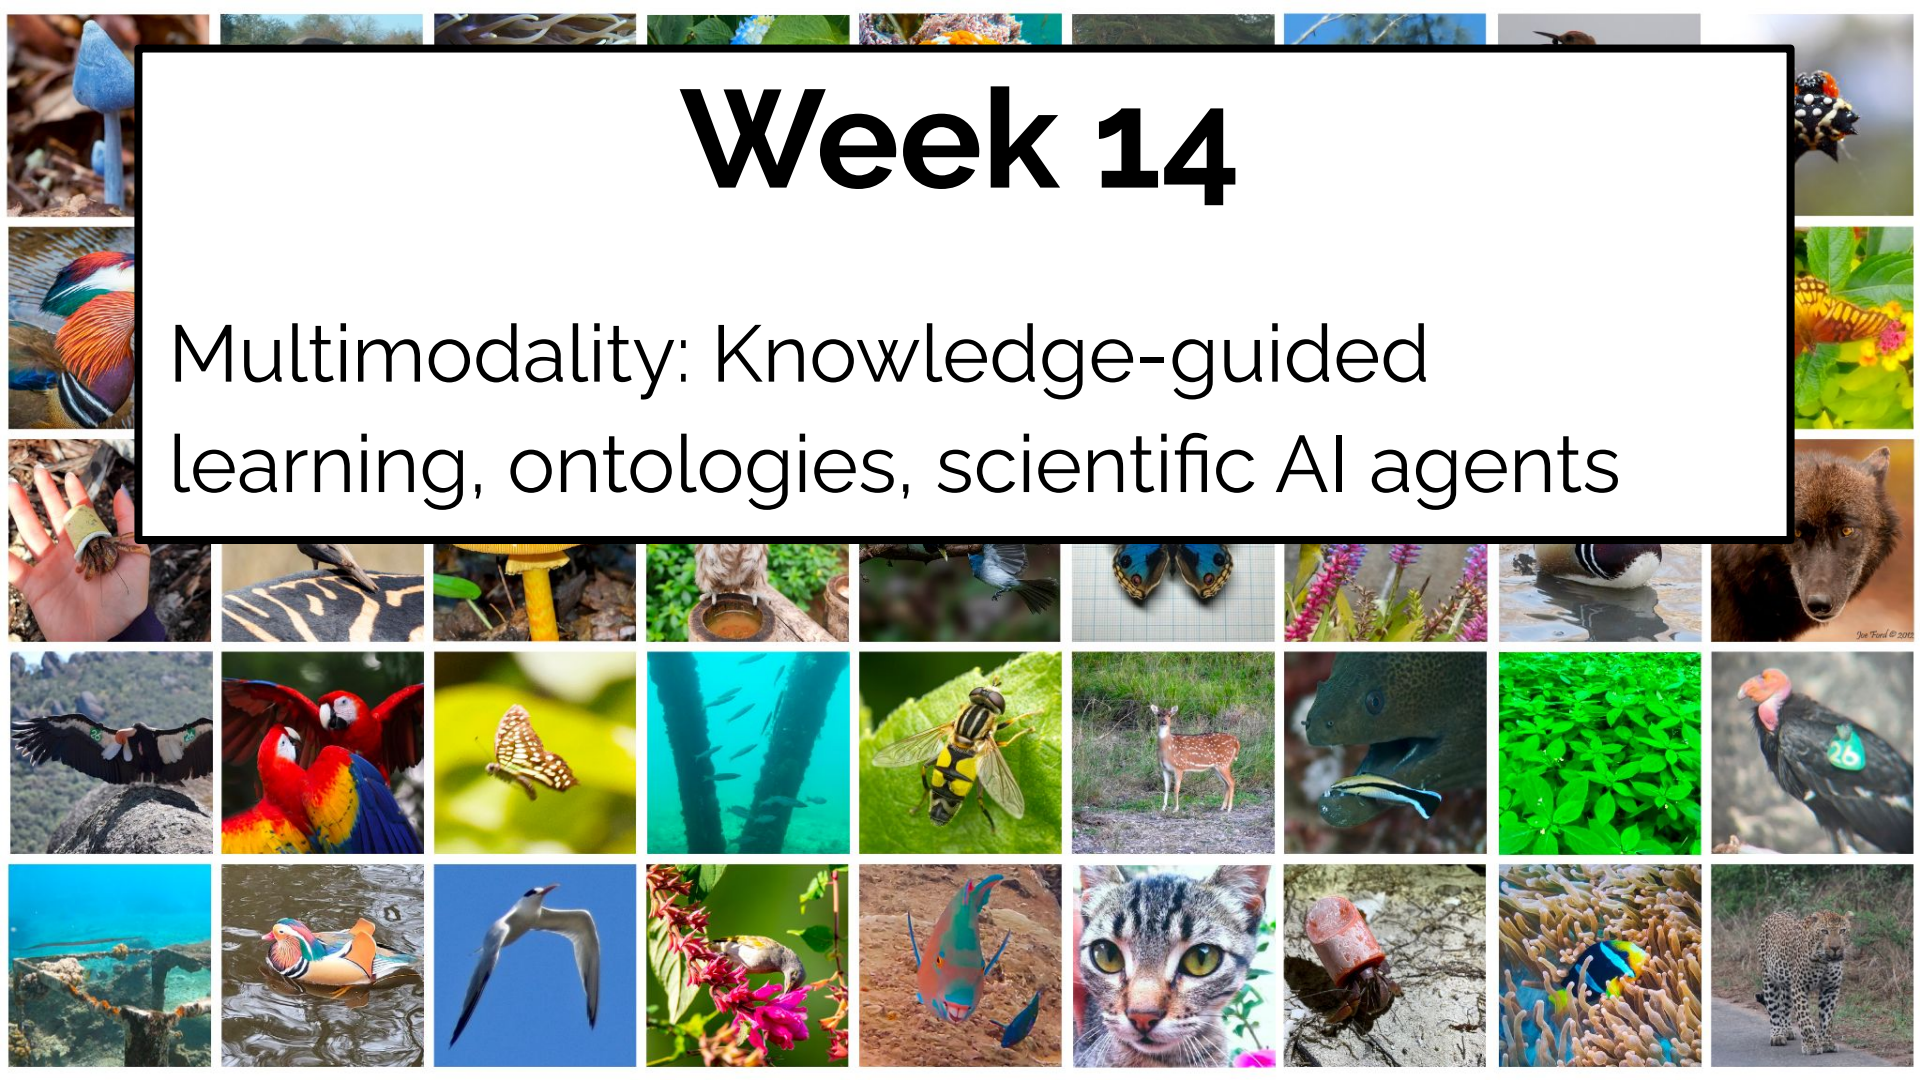

Week 14

Multimodality: Knowledge-guided learning, ontologies, scientific AI agents

Week 15

Final project presentations

In-class presentations

Each student will do each role >=twice

- Summarization
  - Summarize the main contributions of the paper and contextualize them
- Critique
  - Point out limitations or flaws of the paper and suggest how they might be remedied
- Extension
  - Discuss and map out 2-3 possible extensions of the paper

Final project

- Groups of 1-3
- Propose a research project
- Submit a <=4 page proposal and present the project to the class
  - Additional details to come in March when the project is assigned

Questions?

Info at
cv4planet.github.io

Questions on Piazza