CONSERVATION X LABS

Faces, Flukes, Fins, and Flanks: How Multispecies Re-ID Models are Transforming Our Approach to AI for Wildlife

Decoding Communication in Non-Human Species III
6/29/2024

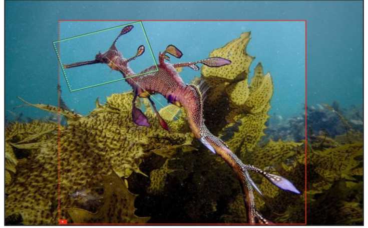

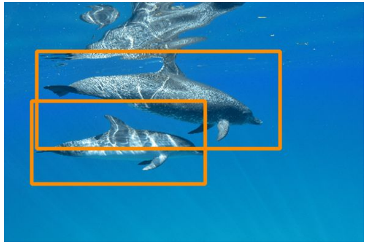

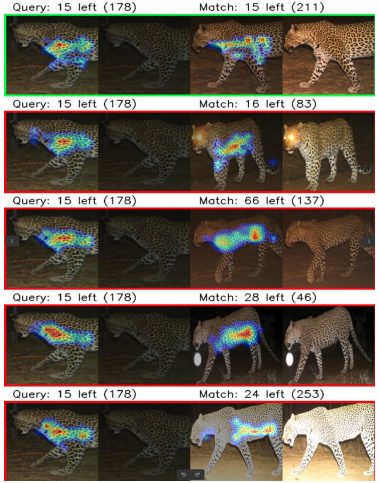

We can use machine learning to scale global conservation

1,900+
Wildlife researchers supported

221k+
Individual animals tracked

1M+
Encounters managed

10M+
Photographs collected

The Challenge

Extinction: Structural Problems to Solve

- ➔ Long gaps between population assessments
- ➔ Slow iteration of protection strategies
- ➔ Gap between promising technology and successful, scalable field application
- ➔ Funding
- ➔ Skills
- ➔ Experiment vs. engineering experience

Our solutions

Develop continuous monitoring and estimation of wildlife populations

Scale, modernize, link, and support front-line conservation efforts

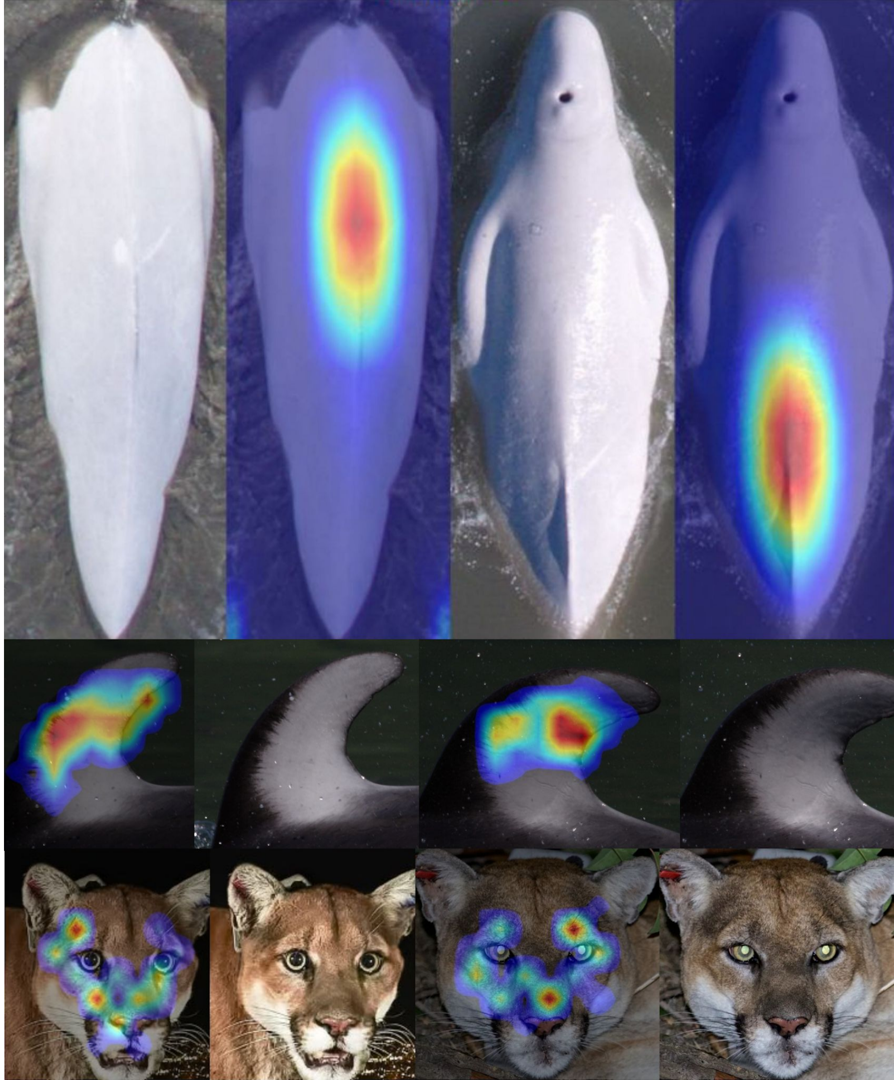

Wild Me Initiatives

Wild Me supports animal population assessment and research with open source platforms that provide long-term data curation, high-speed data processing with AI, and collaboration across borders and regions of study. The initiatives to the right are some exemplars of these efforts.

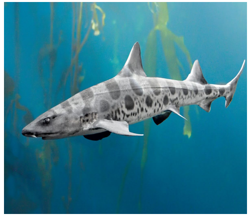

African Carnivore Wildbook

Partner: Technology for Conservation (T4C)

Tracking individual lions, leopards, cheetah, wild dogs, and more across home ranges, borders, and organizations in Africa.

Flukebook

Partner: National Oceanic and Atmospheric Administration (NOAA)

Tracking 50+ species of whales and dolphins across all ocean basins. Flukes, fins, flanks, and more are individually matchable within a single AI model.

Sharkbook.ai

Tracking individual whale sharks, white sharks, sand tiger sharks, leopard sharks, and more.

Amphibian Reptile Wildbook

Tracking multiple life stages of fire salamanders and yellow-bellied toads. Extensible to many more species of reptiles and amphibians globally.

Empowering a global network for our planet.

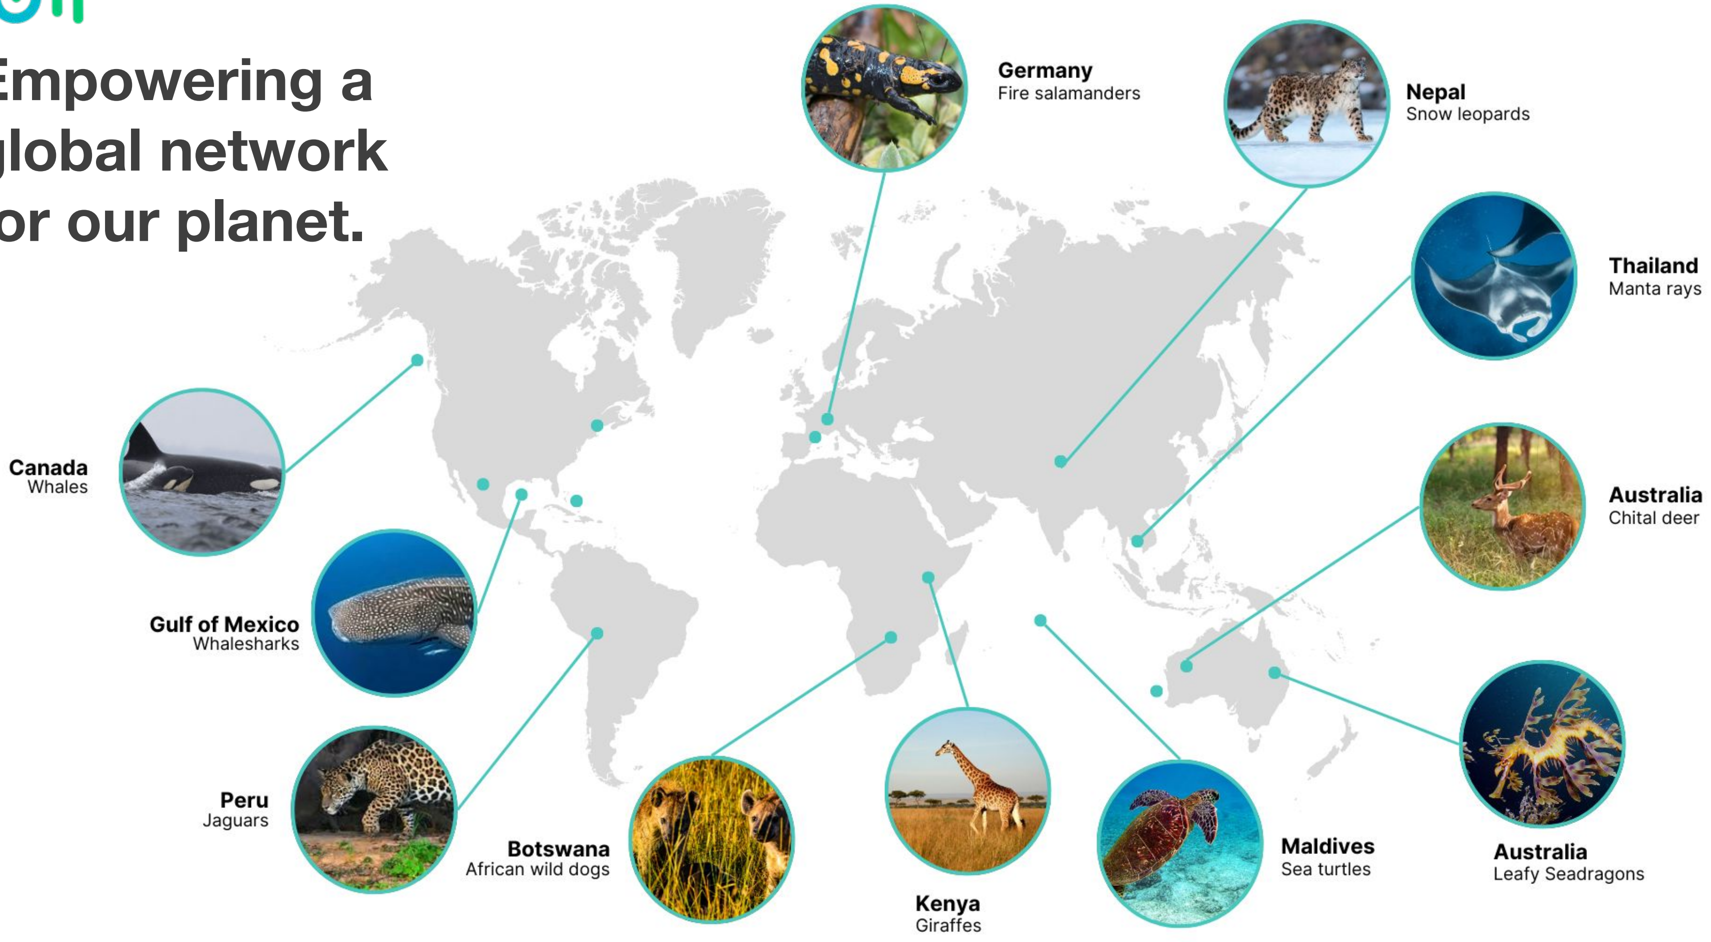

Four Basic Vision Tasks for Wildlife

1. Ignore empty imagery (whole-image classification)
    - Aerial surveys
    - Camera traps
    - Biggest immediate value to users
2. Localize (bounding box) animals
    - Pop. estimation (counts)
3. Label bounding boxes
    - Viewpoint
    - Species
    - Health
    - Behavior
4. Re-ID of individuals (real world; open world)
    - Pop. estimation
    - Social grouping and dynamics
    - Movement

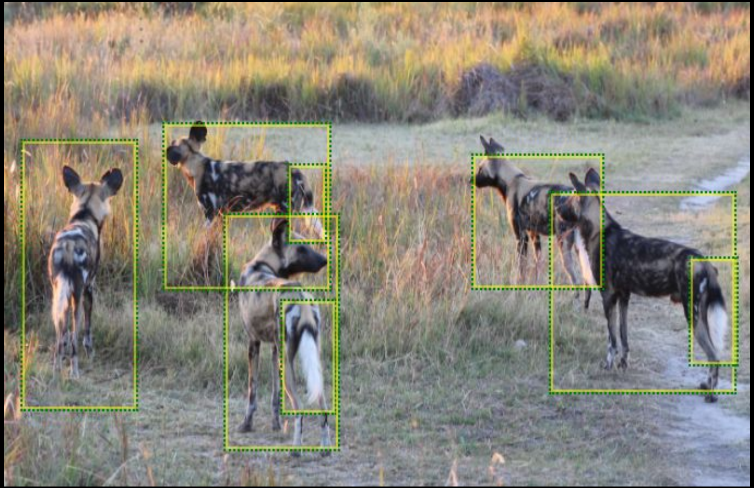

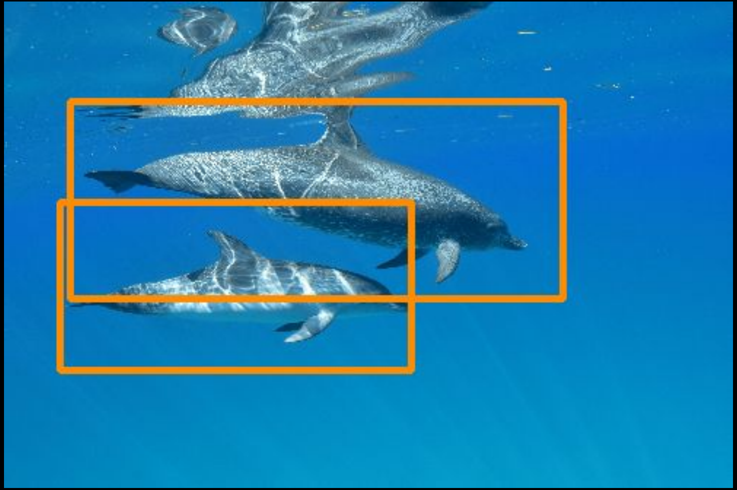


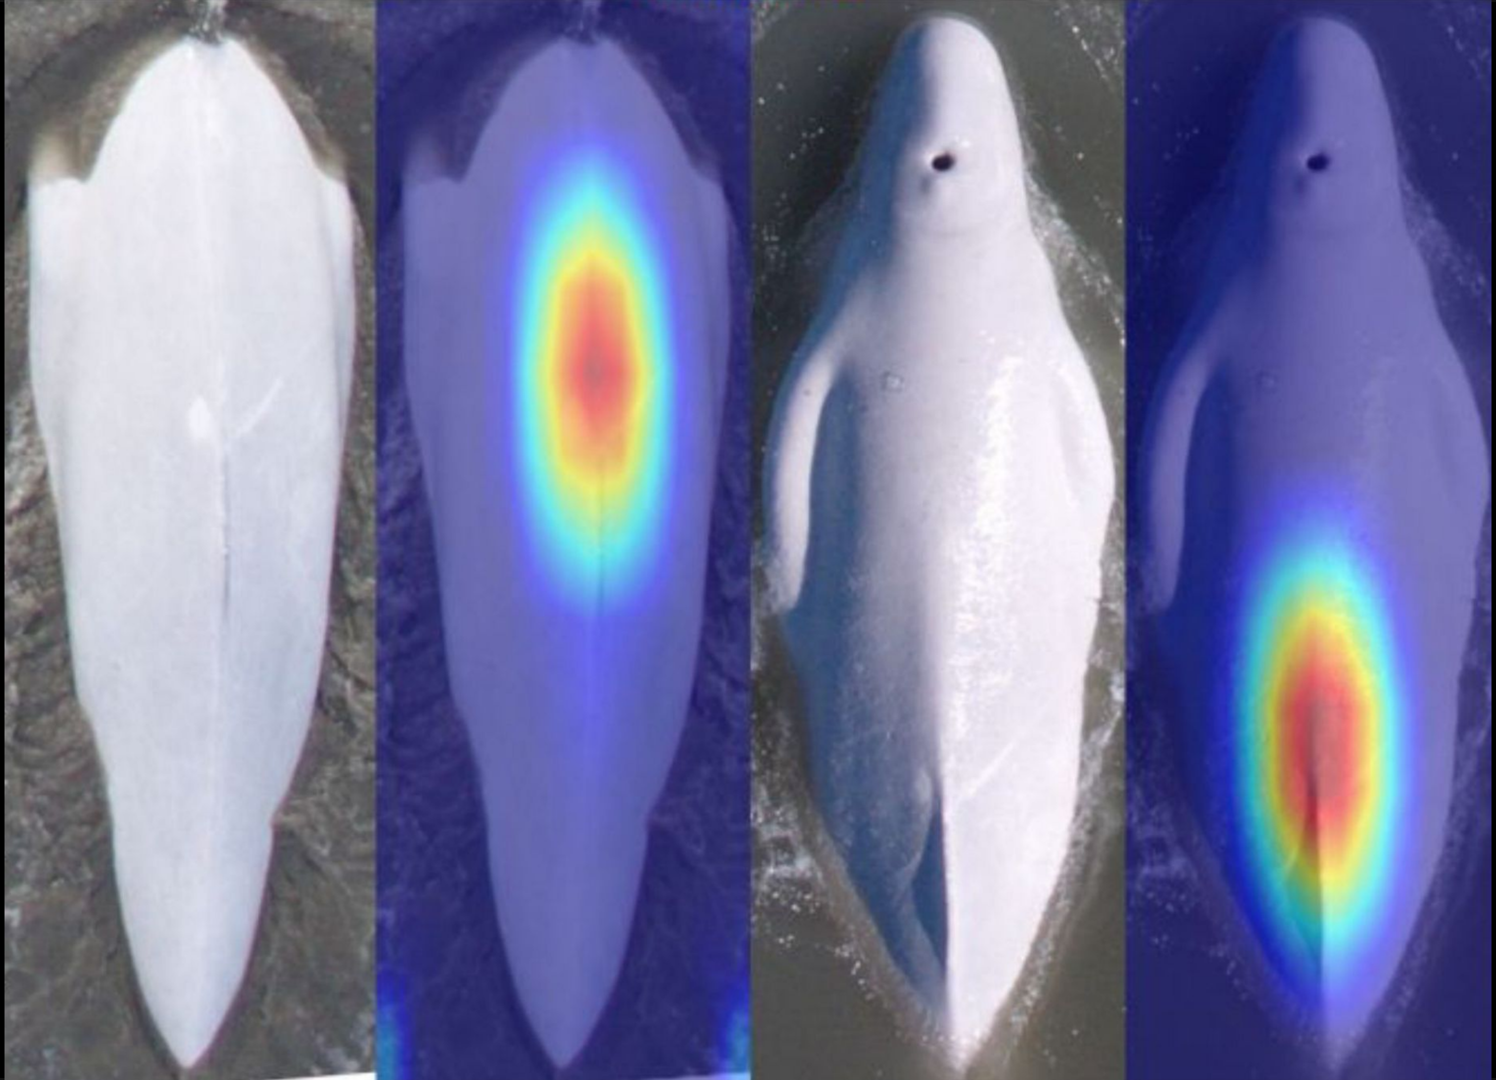

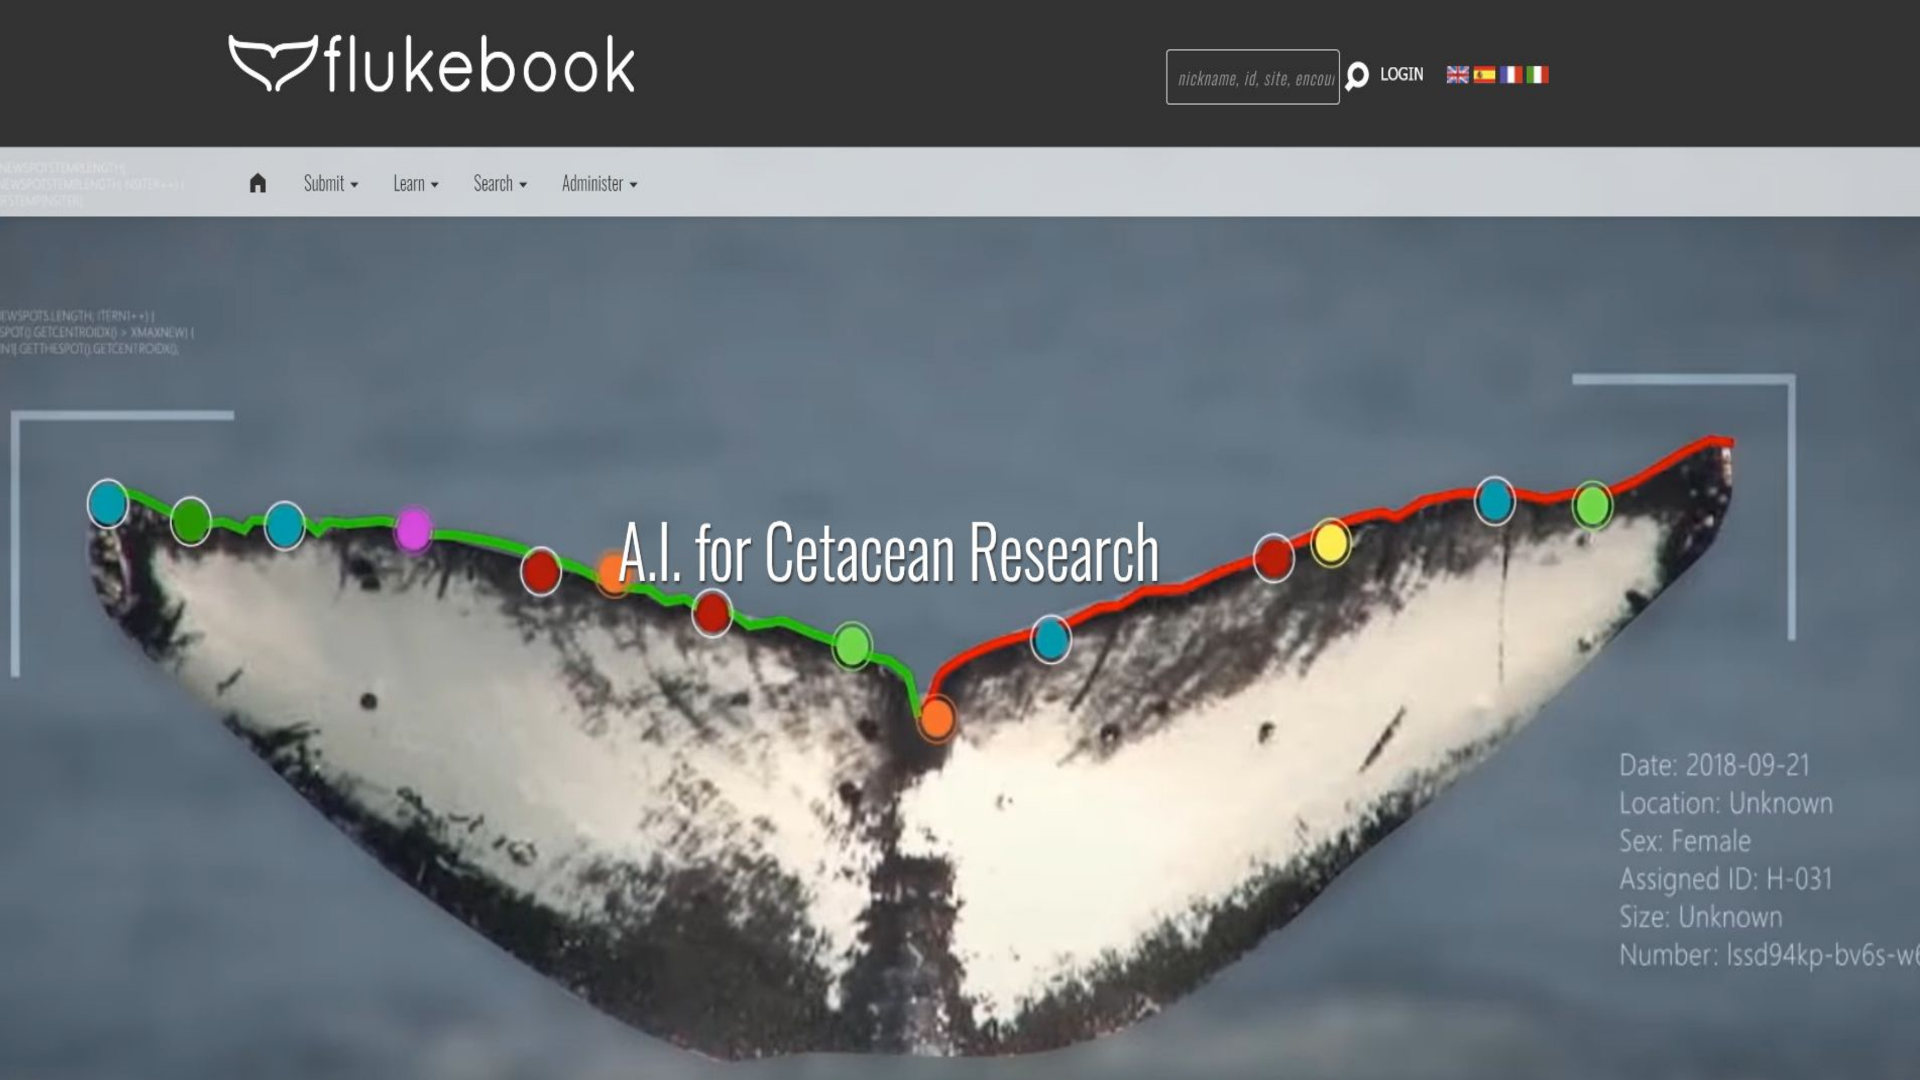
flukebook
nickname, id, site, encou
LOGIN
Submit
Learn
Search
Administer
A.I. for Cetacean Research
Date: 2018-09-21
Location: Unknown
Sex: Female
Assigned ID: H-031
Size: Unknown
Number: lssd94kp-bv6s-w6

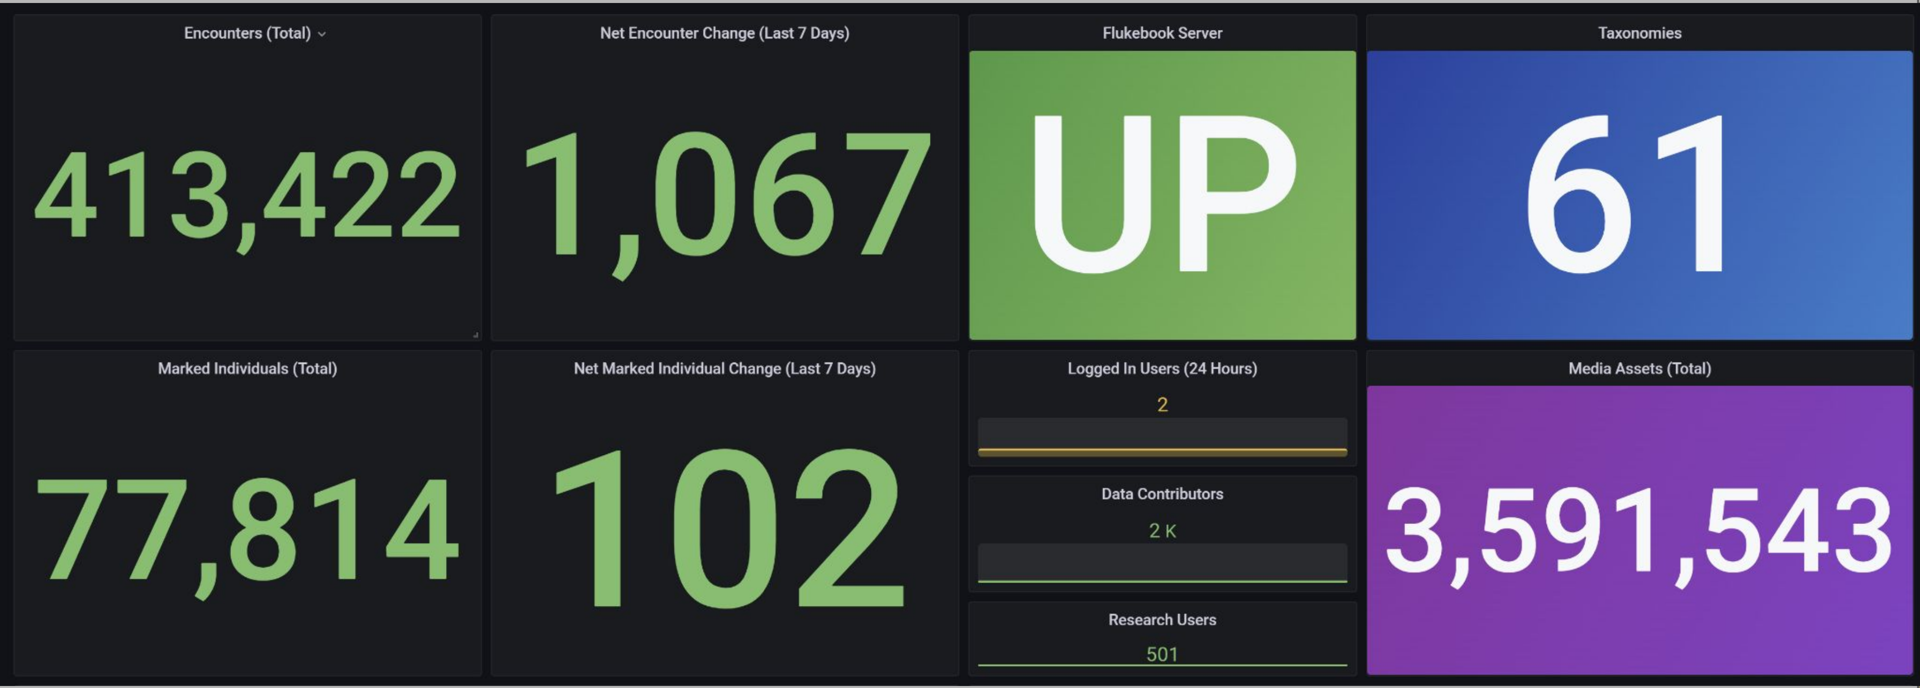
Encounters (Total)
413,422
Net Encounter Change (Last 7 Days)
1,067
Flukebook Server
UP
Taxonomies
61
Marked Individuals (Total)
77,814
Net Marked Individual Change (Last 7 Days)
102
Logged In Users (24 Hours)
2
Data Contributors
2 K
Research Users
501
Media Assets (Total)
3,591,543

Scaling Problems with Real World Re-ID

- Per-species projects are a lot of work
  - Limited data limits accuracy, generalization
  - Different visual features require different approaches
  - Custom user engagements and data schemas
  - Retraining often required for expansion and cross-application
  - *Don't scale to solving the extinction crisis*

- Multiple computer vision algorithms and architectures
  - $ to support
  - Many things to debug/optimize/synchronize
  - Fight for computer resources

From Single to Multispecies Re-ID: Intellectual Heritage

- 2014-2020: HotSpotter, finFindR, CurvRank, DTW
- 2021: Pose Invariant Embeddings
- 2022: HappyWhale Whale and Dolphin Kaggle Competition
- 2023: DrivenData Where’s Whale-do Beluga Competition
- 2023: MiewID (single species; open world)
- 2023-2024: MiewID (multi-species; open world)
  - Whale/dolphins (23 species)
  - Terrestrial carnivores (9 species)
  - Face ID (13 species)
  - Sea turtles (3 species)
- 2024: MiewID (multi-species; marine/terrestrial; open world)
  - v0: 54 Species/61 classes - *Current*
  - v1: 66+ species/70+ classes - *Underway*

Project Time

Years

Days

What are deep learning models / embeddings?

- Old-generation models rely on us pre-configuring features
  - Sensitive to preprocessing
  - Not robust to perspective shifts and occlusions
  - Possibly leaving out some of the useful signal

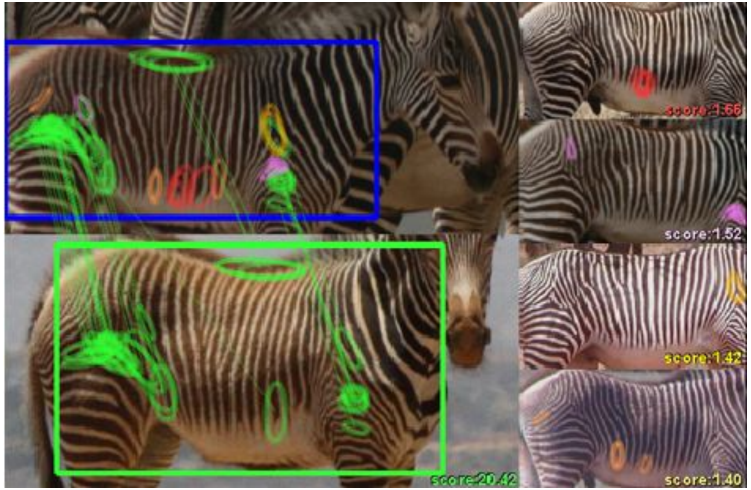

- Deep learning / multispecies
  - The model 'figures it out' from all included pixels/features
  - Shared experience for generalization

Fundamental problem: data-hungry!

Multispecies re-ID: Initial Experiments

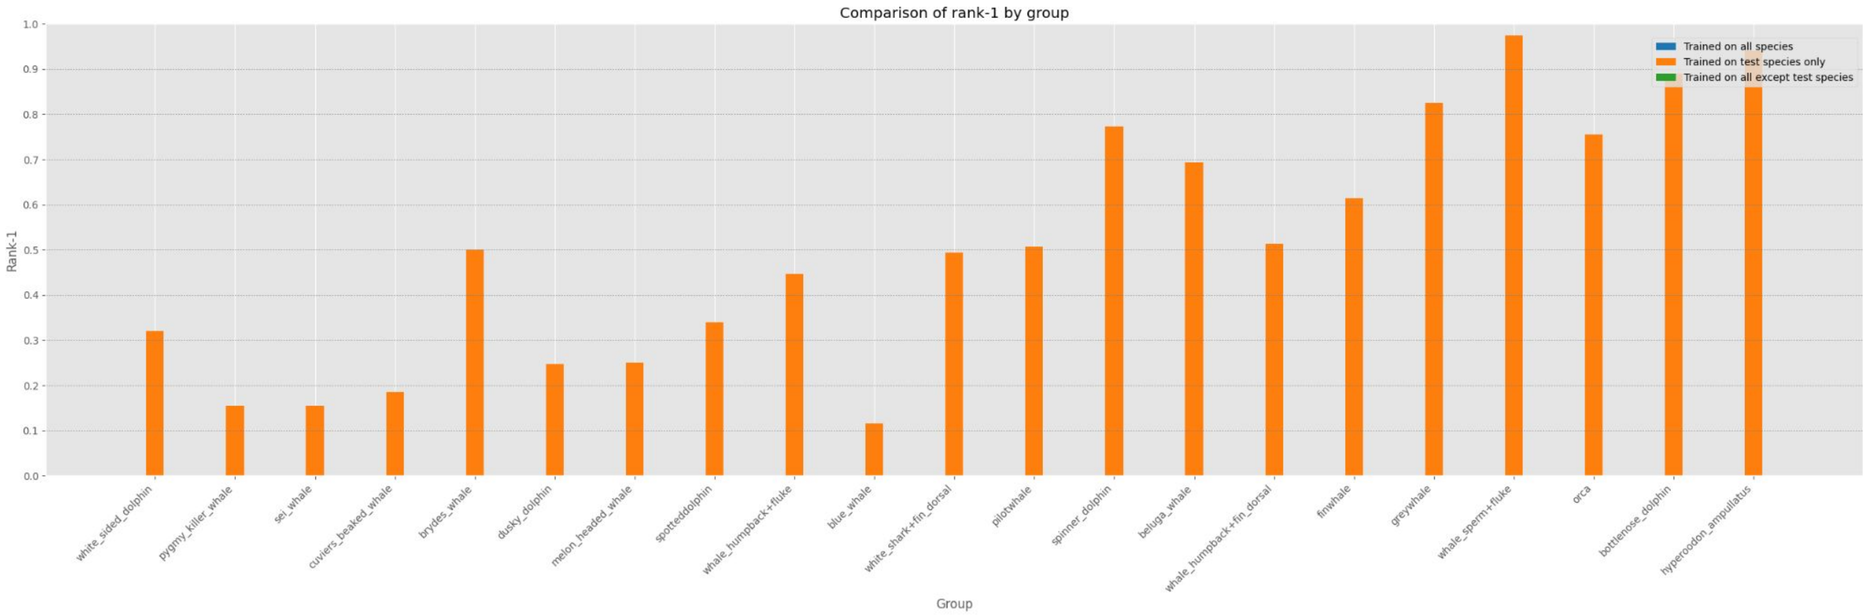

More data

Multispecies re-ID: Initial Experiments

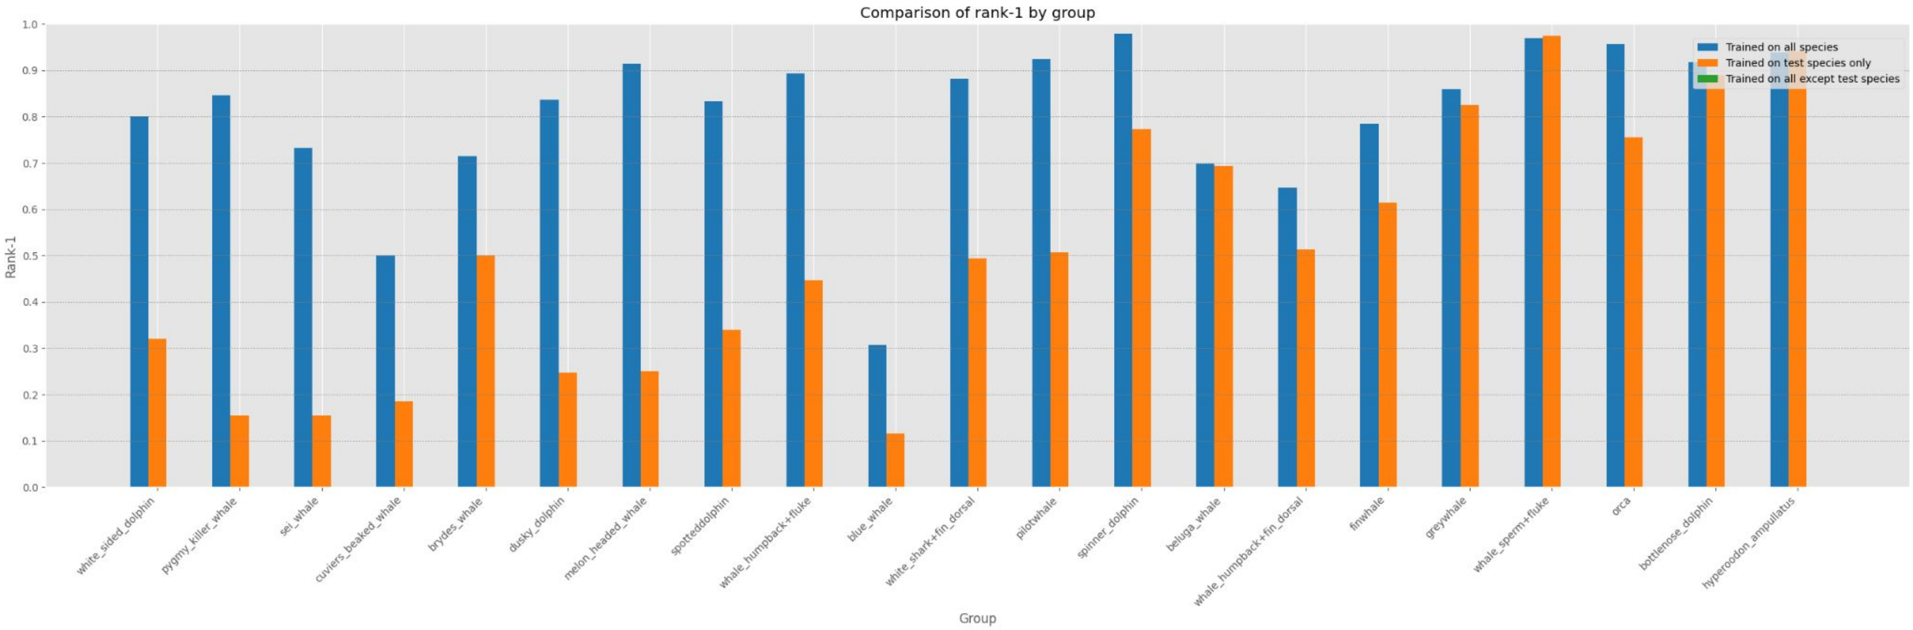

More data

Multispecies re-ID: Initial Experiments

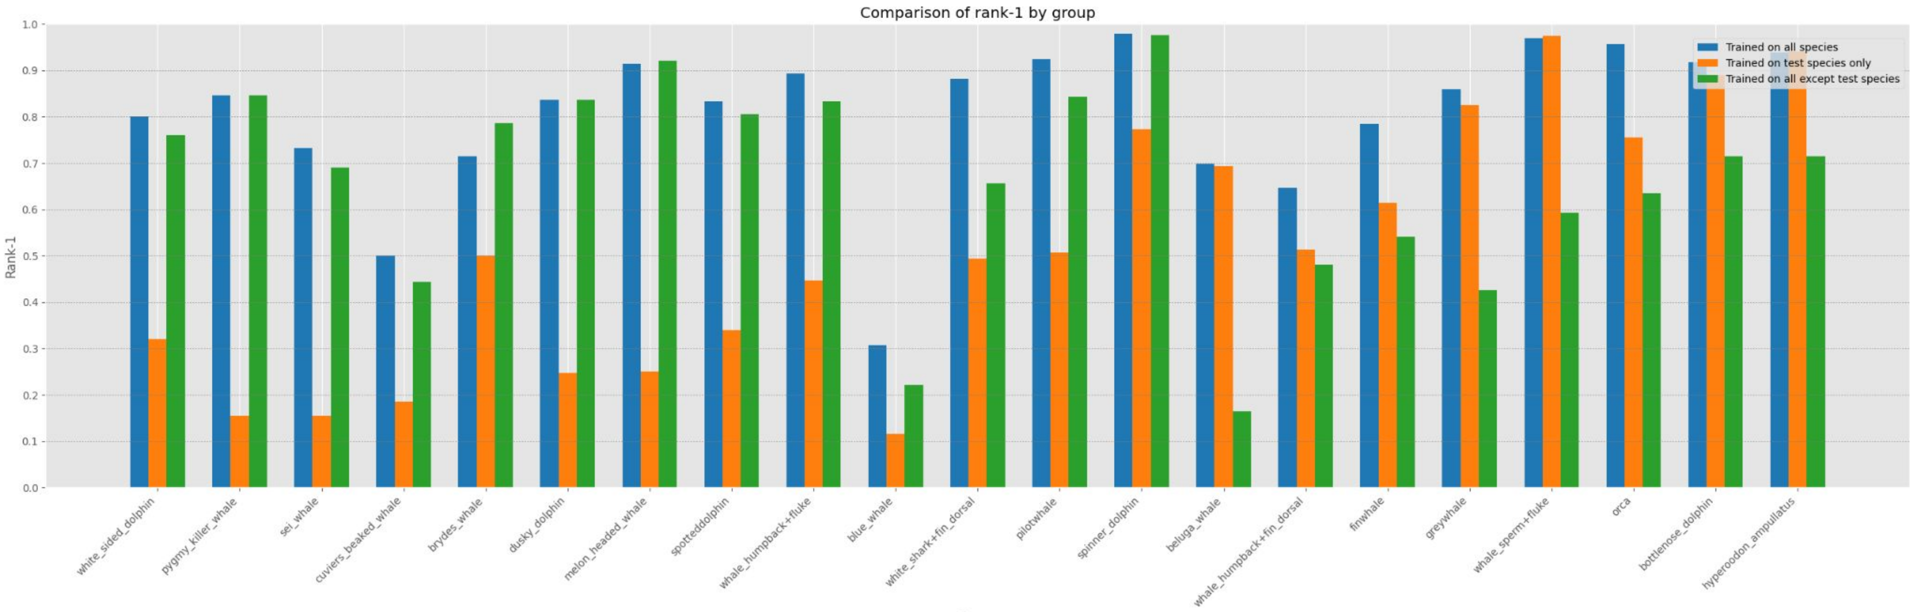

More data

Gain

- Accuracy
- Speed 30x+
- Scope
- Scale
- Collaboration
- Cross-application to new species
- ‘Acquired robustness’

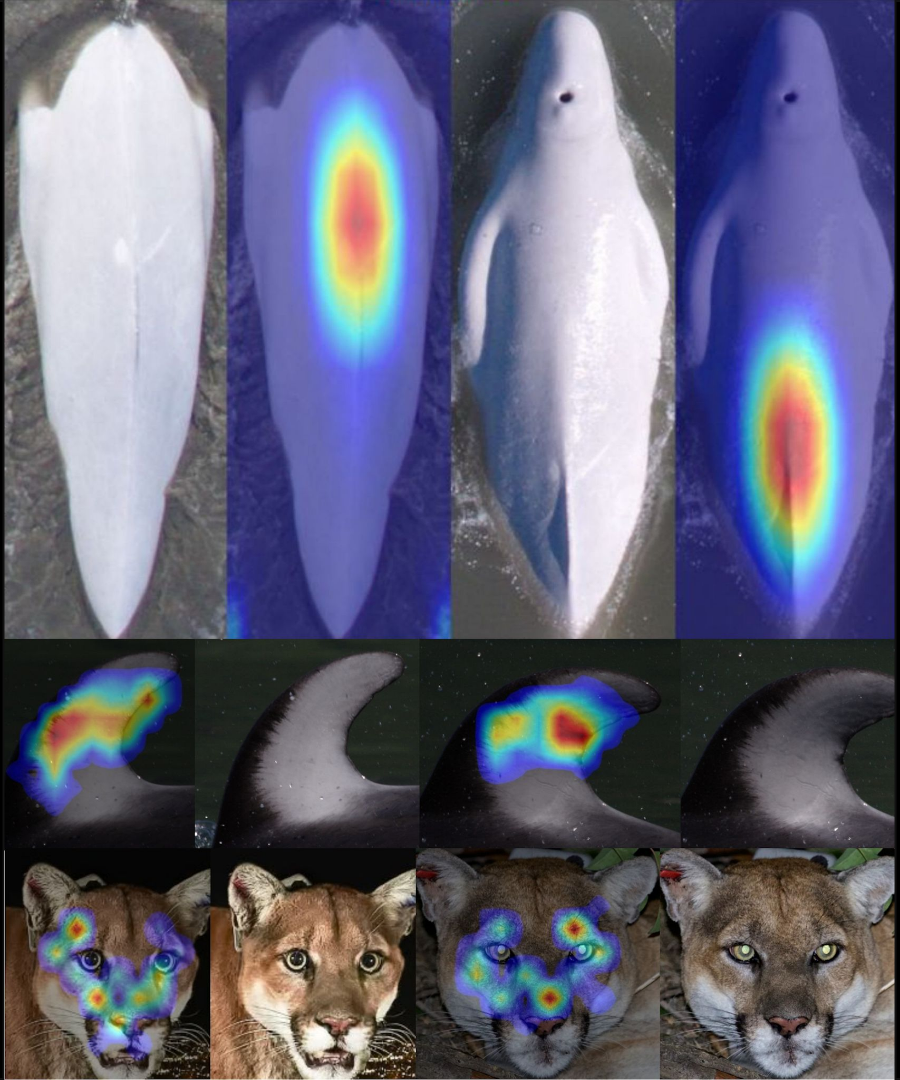

Loss

- Explainability

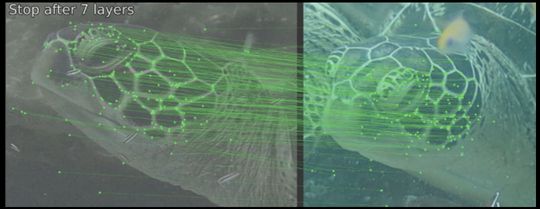

Open set performance

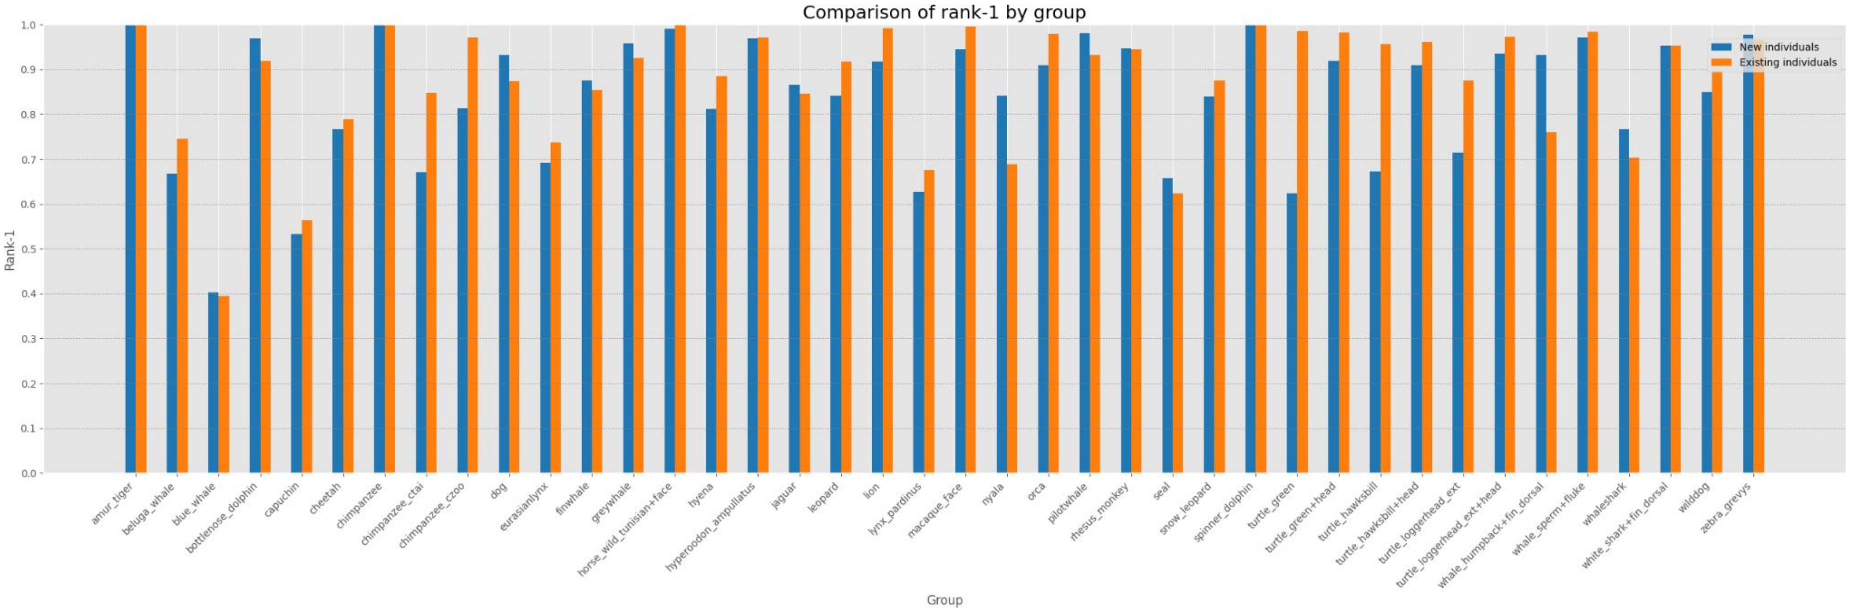

Cross-viewpoint matching - examples


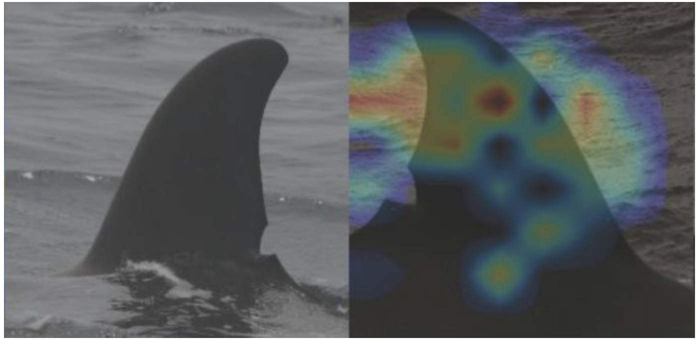
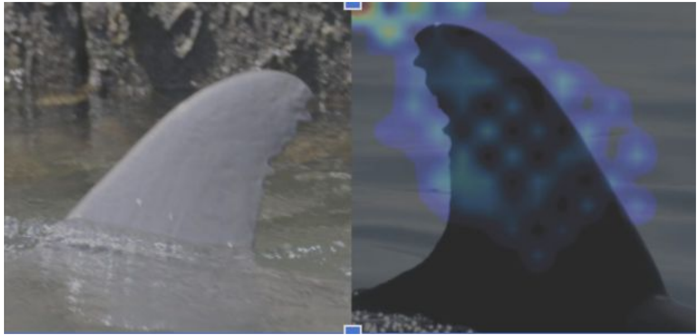


Cross-viewpoint matching

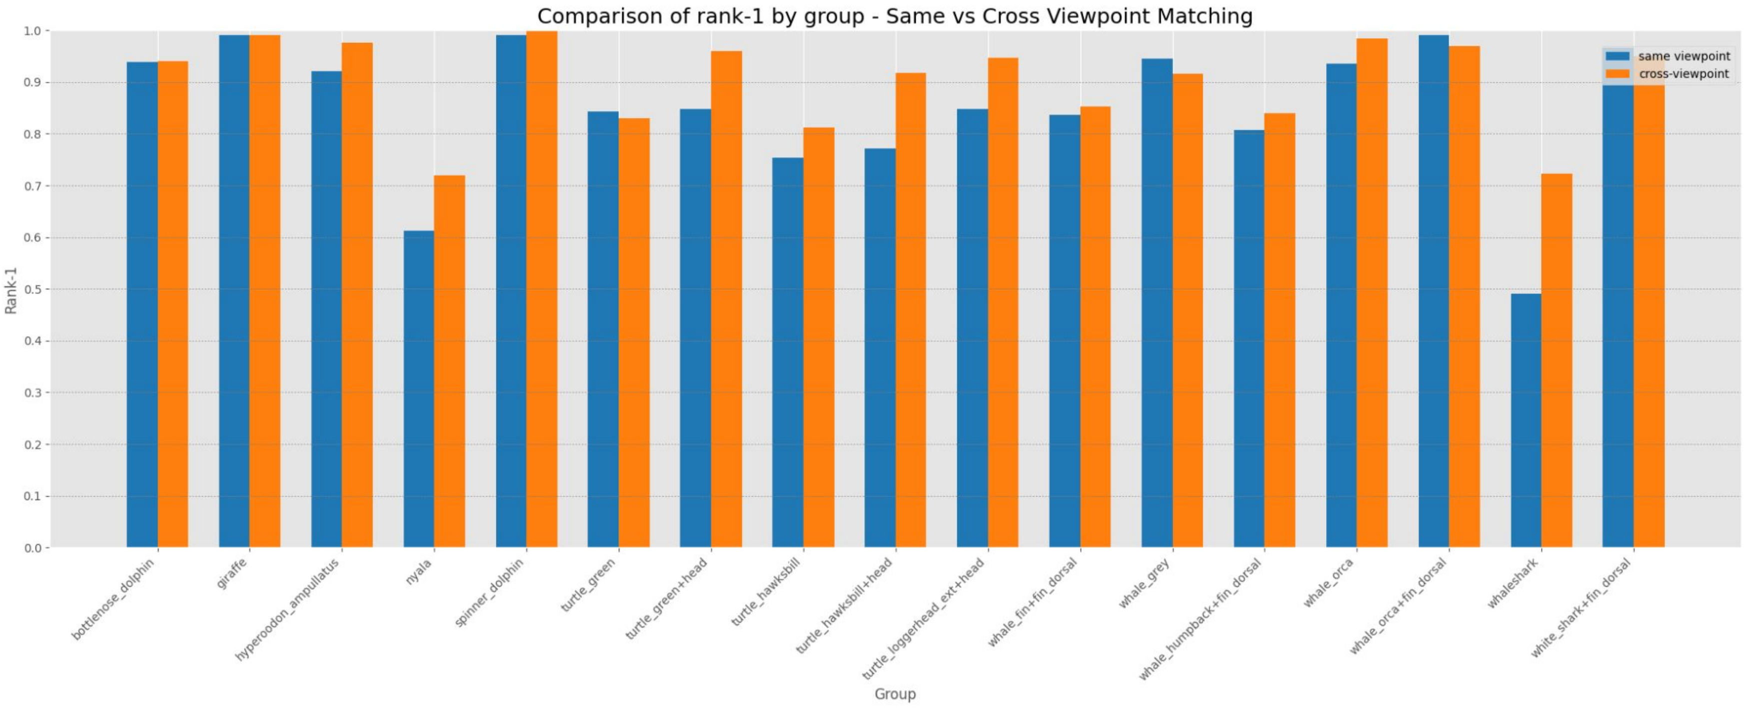

‘Hard’ matches

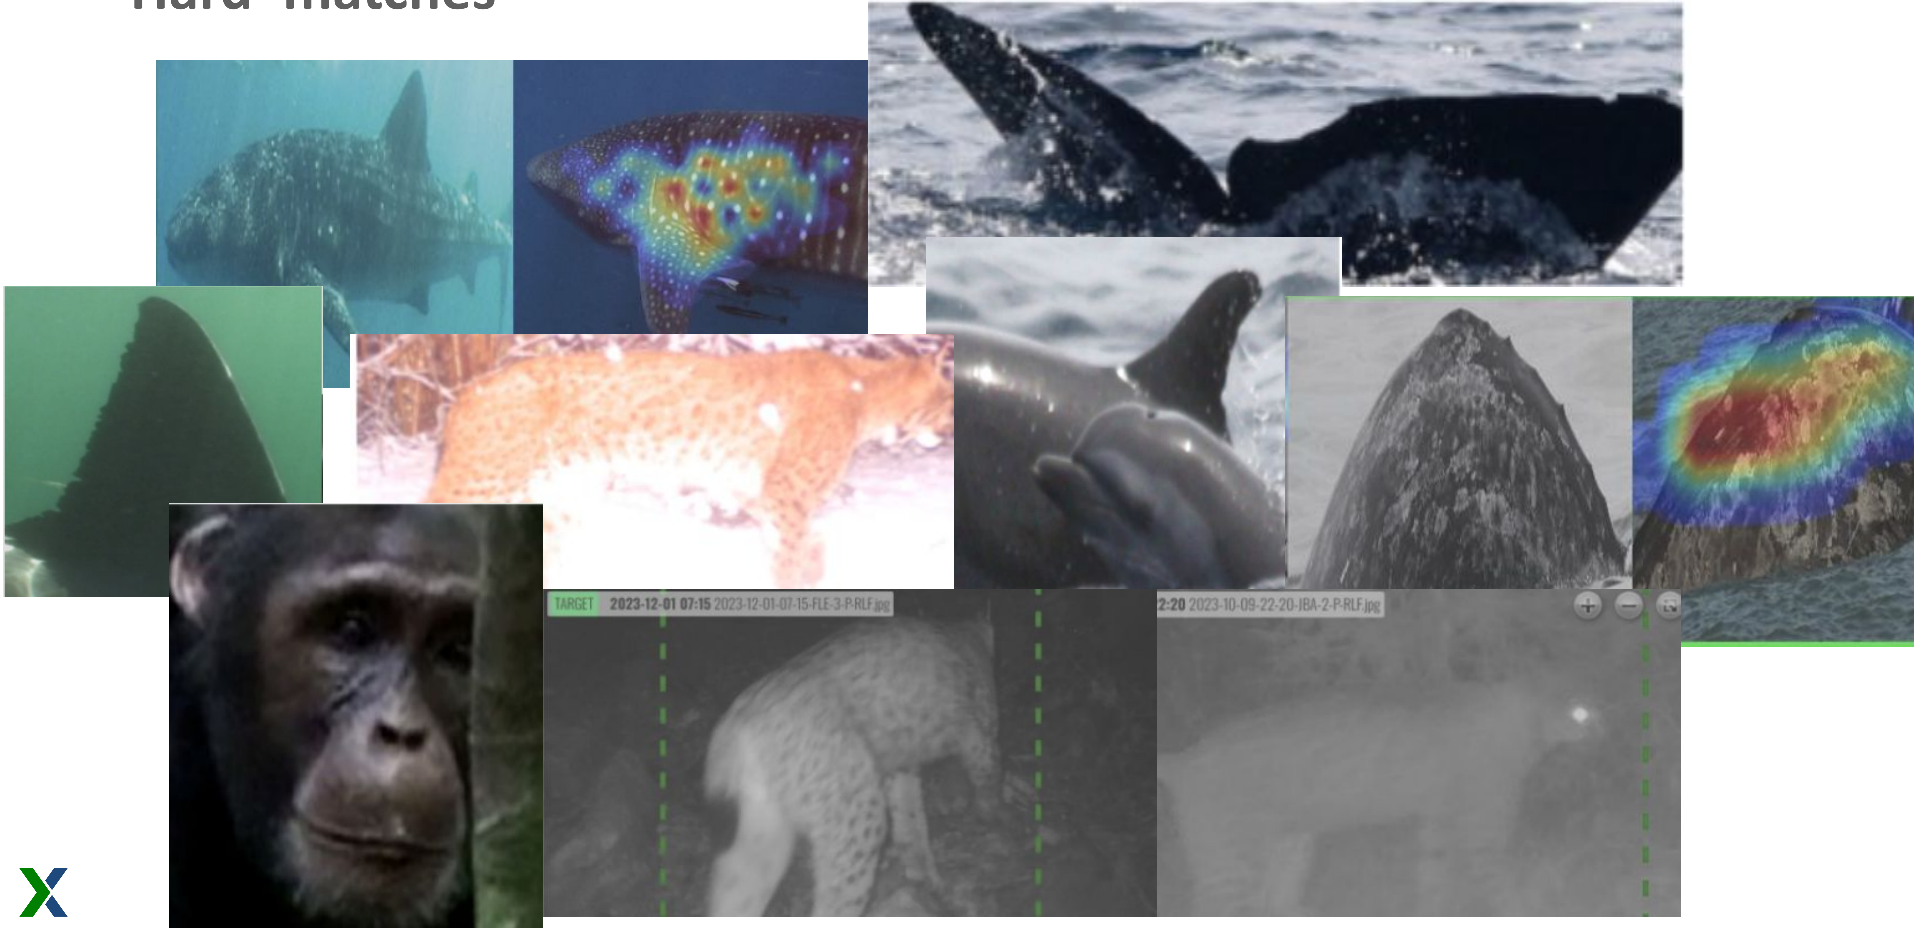

How do we push this further?

Bootstrapping - solving the cold-start problem

Cluster Visualization

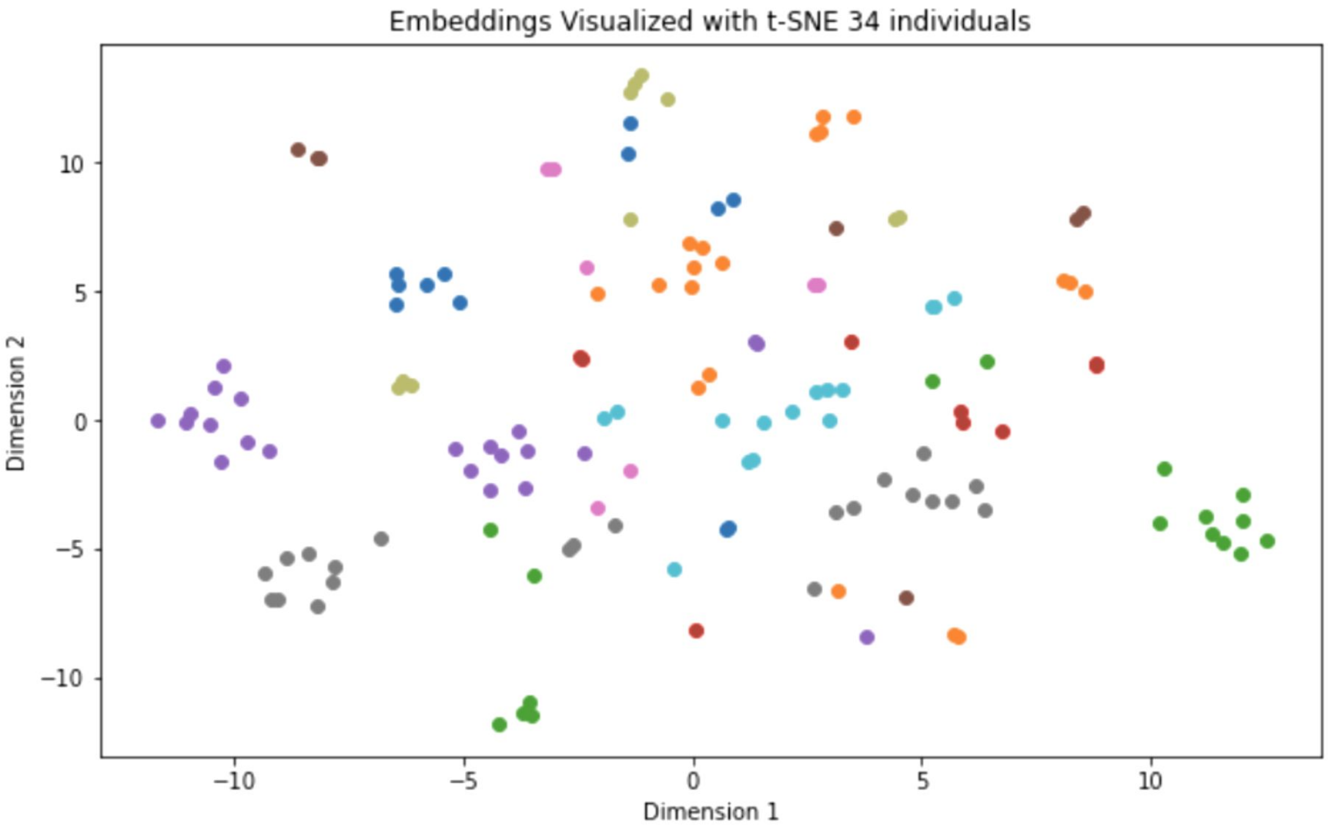

- Data:

  Sperm whales

  34 individuals
  174 samples
- Species:
  In-domain

  Location:
  Out-of-domain

Unsupervised Clustering

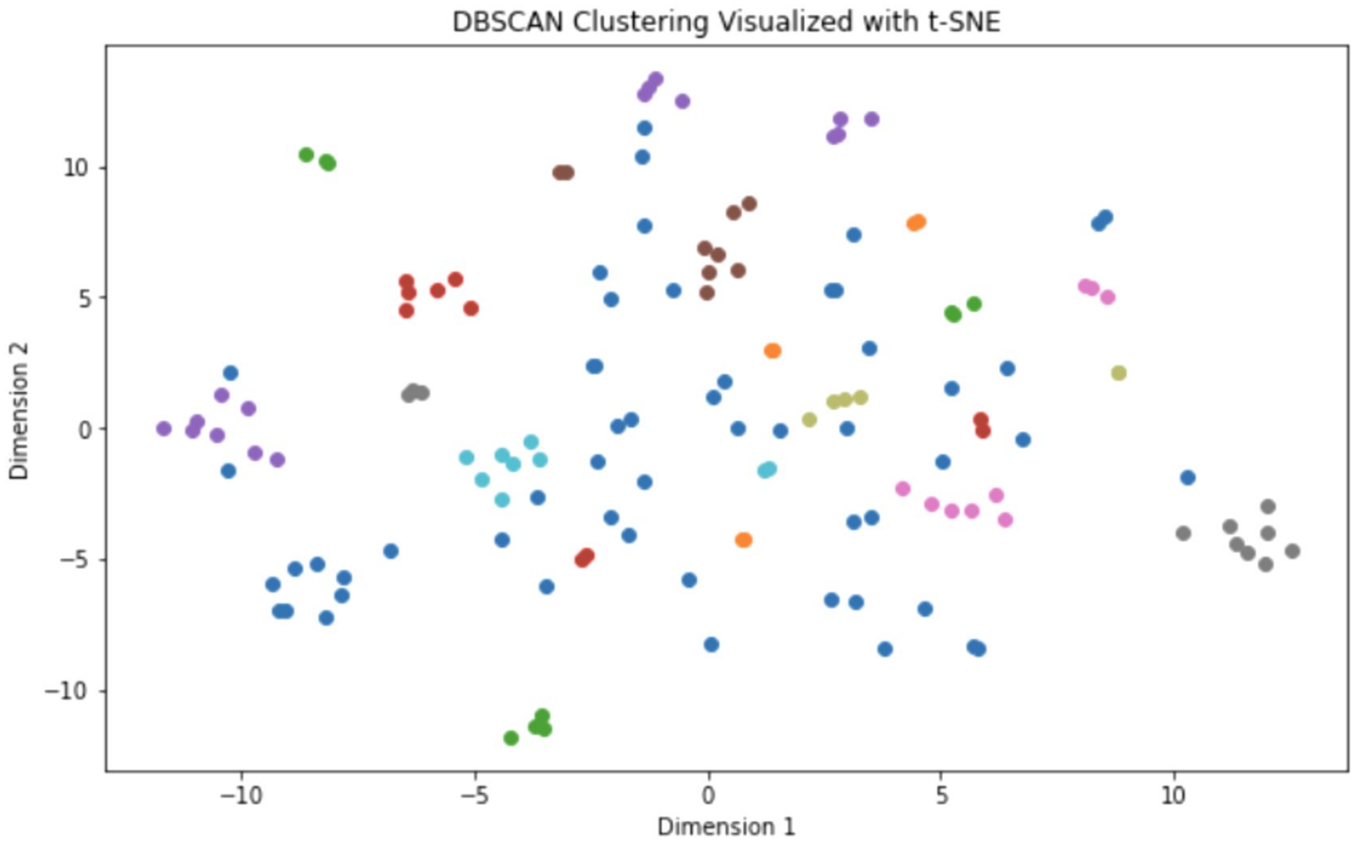

- 70% of samples assigned @>95% precision

Clustering - Quality Check

- Weak labels confidently predicted for 379 / 557 annotations - 61 / 86 individuals

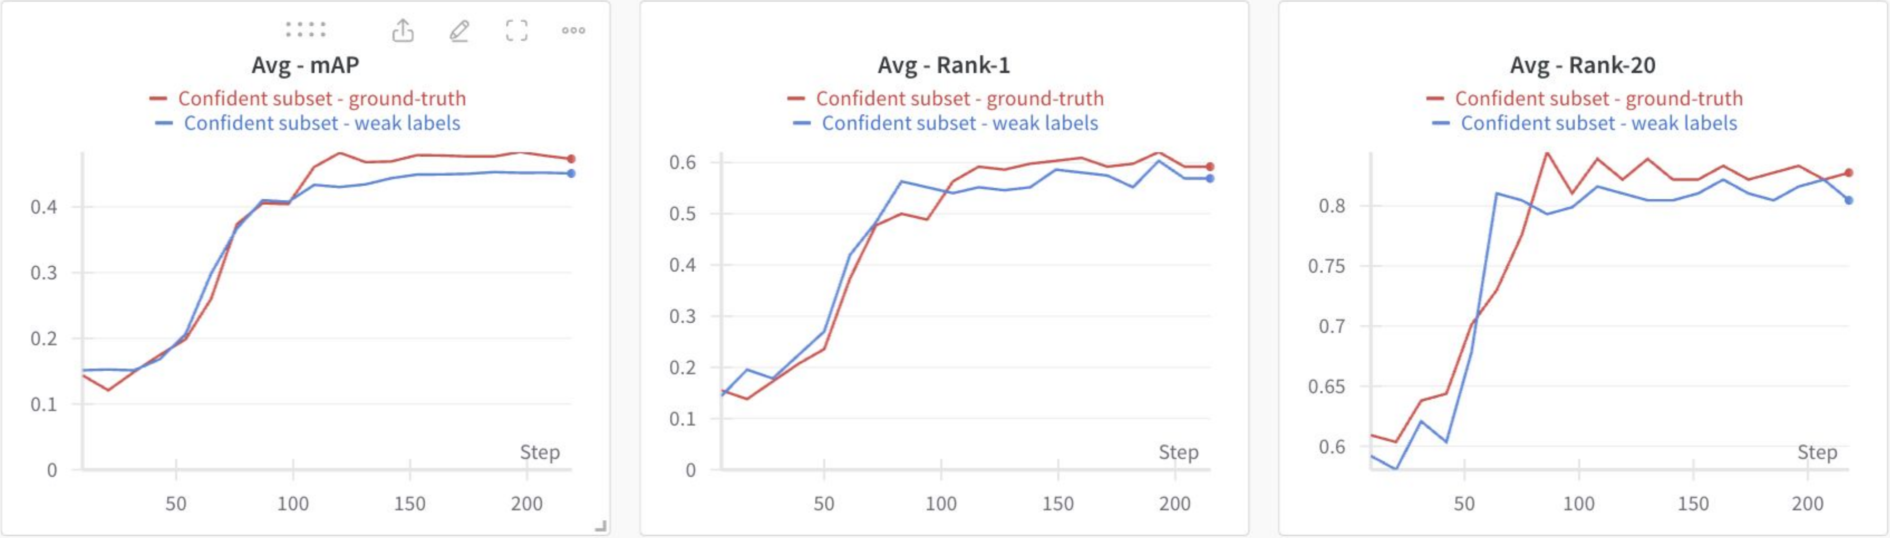

Faster Dataset Bootstrapping with Multispecies

Easy Case: Sperm Whales

Use-case:

1. Bring in a catalog of new species.
2. Cross-apply the multi-species model to help with the cold start.
3. Once a subset of data is curated and reviewed. Retrain the model with this small dataset.
4. Repeat until complete.

Faster Dataset Bootstrapping with Multispecies

Harder Case: Grey Whales

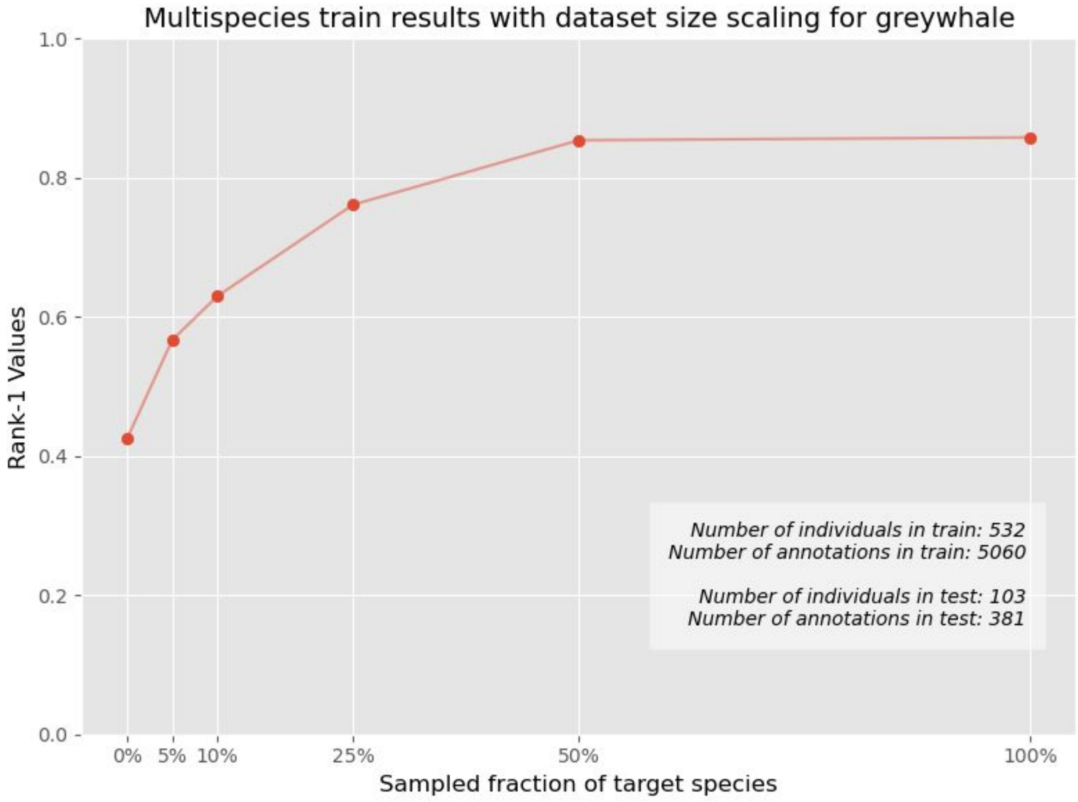

Samples:

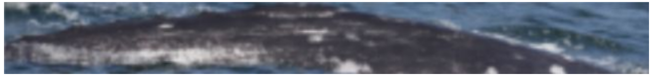

What’s next?

Move from a mono-functional model to multi-functional agent

- Foundation: How do we prepare and review large datasets of images of mixed quality?
- Community: How do we let many humans inject new knowledge into the model as their data curation progresses?
- xAI for Human Acceptance
- Support embeddings-based (not ID) population estimation
- Embeddings-based re-ID on device
- LLM/VLM Agents: assist in the full process of curation
    - Inspect, interact, and support better human ID decisions
    - Auto-curate new data batches
    - Retrieve and curate new data from social media
    - Summarize trends and speed up analysis

Thank You.

Jason Holmberg
Chief Data Officer
jason@conservationxlabs.com

Lasha Otarashvili
Sr. Computer Vision Research Engineer
lasha@conservationxlabs.org

Wild Me Lab, Conservation X Labs
wildme.org
Conservationxlabs.org