Job Description

Senior Computer Vision Research Engineer

Portland, Oregon

Job Summary: Bring state-of-the-art machine learning and computer vision algorithms to Wild Me and work with software engineers to advance Wildbook (https://www.wildbook.org) as a world-class platform for wildlife conservation with machine learning.

General Qualifications

- Knowledge and experience in machine learning and computer vision, especially techniques and approaches in individual animal identification. *A minimum qualification of a Master's Degree in a related field is required.*
- Solid mathematical understanding of computer vision techniques, machine learning, mathematical optimization, and hardware acceleration.
- 3-5+ years of experience and/or education in the field, ideally with experience related to wildlife conservation through the application of computer vision or the equivalent.
- Experience in developing clean and maintainable Python. Javascript, Java, and/or C++ code skills a bonus.
- Knowledgeable about web paradigms (including an understanding of REST API design)
- Good communication skills and ability to work within a team and with clients, including proposal and paper writing
- Comfort and experience in Linux environments

We are seeking a team member with:

- Self-direction and clear communication and coordination of active work toward a team goal
- Experience building machine learning models
- Experience with and/or a willingness to learn deep learning frameworks (e.g., using PyTorch / Tensorflow / Caffe / Keras etc). Prefered experience with PyTorch.

Optional bonus skills:

- Experience solving a real-world animal detection or ID problem
- Experience with cloud services (e.g., Azure)
- Experience with crowdsourced data collection and citizen science activities
- Experience building systems at scale
- Experience building web services (JavaScript, Java, HTML, etc.)
- Willingness to help teach others, including code school graduates and Wild Me engineering staff
- Active and productive open source track record

Expectations:

- Integrate the latest computer vision algorithms for animal identification (Wildbook Image Analysis, formerly "IBEIS") into the open-source Wildbook platform
- Work in an open, highly collaborative environment.
- Show enthusiasm for research and application in machine learning, helping us bring machine learning at scale to wildlife biologists across the globe.
- Cross-train in and integrate with Wild Me's existing identification and detection pipeline https://cthulhu.dyn.wildme.io/public/posters/parham_wacv_2018.pdf
- Contribute to writing funding proposals and attend meetings with clients that expect a high degree of competency
- Move to Portland, OR to work in-person at Wild Me's office

Example techniques, repos, and publications we would love to see you familiar with:

- https://github.com/wildbookorg (Wildbook's main repo)
- https://arxiv.org/abs/1902.10847
- https://arxiv.org/pdf/1503.03832.pdf
- https://arxiv.org/abs/1901.08616
- https://arxiv.org/abs/1708.07785
- https://www.kaggle.com/c/humpback-whale-identification
- http://cs.rpi.edu/hotspotter/crall-hotspotter-wacv-2013.pdf
- https://pypi.org/project/graphid/

About Wild Me

As the price of photography and video equipment drops while quality and availability improve (think mobile phones and “GoPro”), images and video from tourism are becoming the most abundant and inexpensive sources of wildlife data. If these images could be widely obtained, rapidly analyzed and combined with related data (e.g., location, date, behavior), then scientists and conservation managers could benefit from larger and broader data sets. An increase in well managed data enables advances in analysis and modeling of animal populations, supporting deeper insight and better methods of protection for endangered animals; this, simply put, allows quicker and more efficient action in regard to wildlife protection.

Wild Me (www.wildme.org) — a 501(c)(3) non-profit organization focused on wildlife conservation — is uniquely comprised of IT professionals and computer scientists, advised by preeminent wildlife biologists. Wild Me has developed the Wildbook platform (see wildbook.org) to help scientists organize wildlife research, collect data from the public (e.g., photos and video), and integrate advanced artificial intelligence to speed data curation. Wildbook takes advantage of tourism and the growing “citizen scientist” movement, bringing the concepts of broad sector inclusion and “Big Data” to wildlife conservation. Data from scientists and citizen scientists in Wildbook lets researchers determine population sizes faster and then adjust ecological policies more quickly with greater accuracy and higher precision. A great example is the Great Grevy’s Rally (http://www.greatgrevysrally.com/), which helps ecologists like Professor Dan Rubenstein of Princeton University utilize photos from the public to count Grevy’s zebras in Kenya.

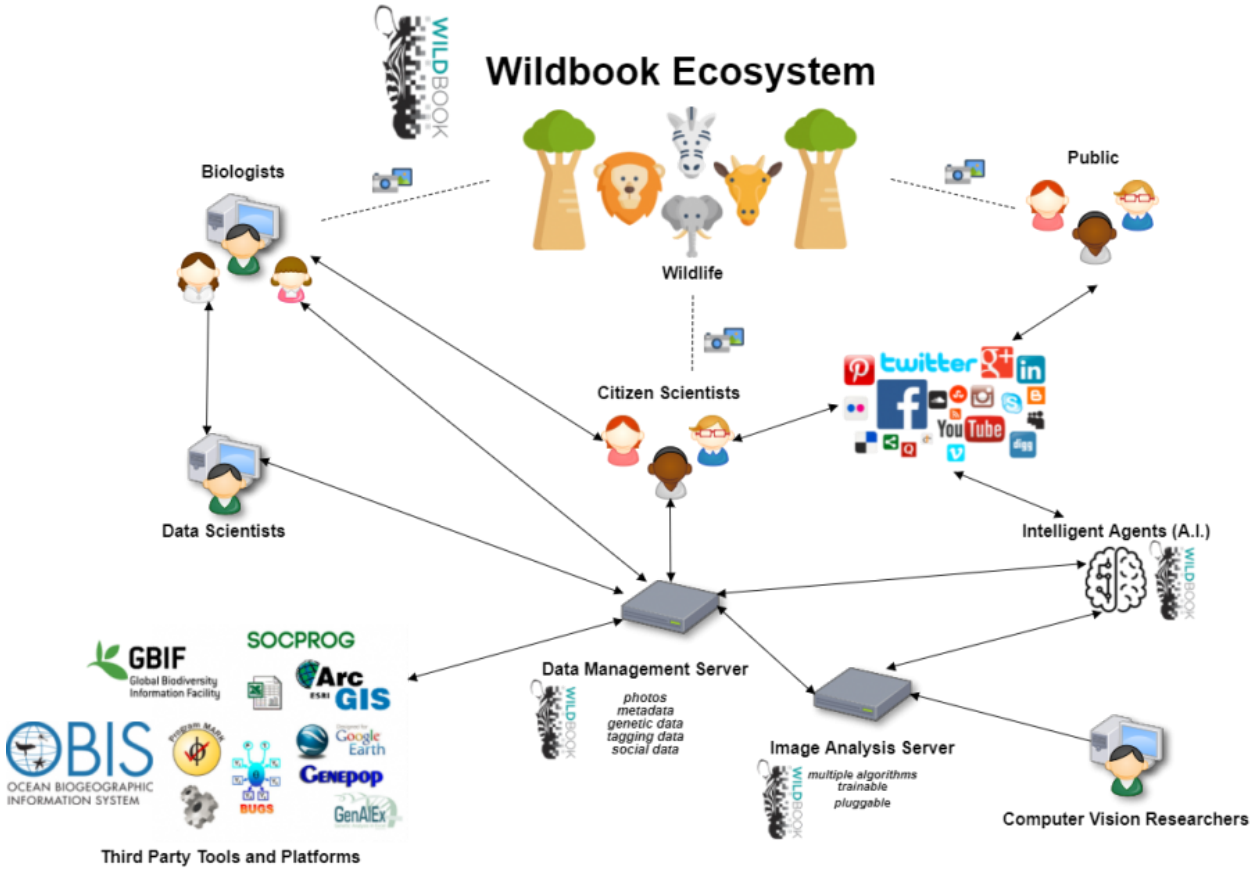

Figure 1a. The Wildbook Ecosystem of Scientists, Citizen Scientists, and A.I.

Integrating Artificial Intelligence with Data Management

Wildbook includes a two-part, multi-species computer vision pipeline to find and identify individual animals in photos collected under real-world conditions, especially with citizen science contribution.

Detection

Our detection pipeline is a cascade of deep convolutional neural networks (DCNNs) that applies a fully-connected classifier on extracted features. Three separate networks produce: (1) whole-scene classifications looking for specific species of animals in the photograph, (2) object annotation bounding box localizations, and (3) the viewpoint, quality, and final species classifications for the candidate bounding boxes.

Figure 1b. A sperm whale (Physeter macrocephalus) fluke is detected automatically with machine learning in Flukebook.

Identification

The second major computer vision step is identification, which assigns a name label to each annotation from detection. To do this, multiple algorithms and machine learning models may be employed. Scores from the query that match the same individual are accumulated to produce a single potential score for each animal. The animals in the database are then ranked by their accumulated scores. A post-processing step spatially verifies the descriptor matches and then re-scores and re-ranks the database individuals.

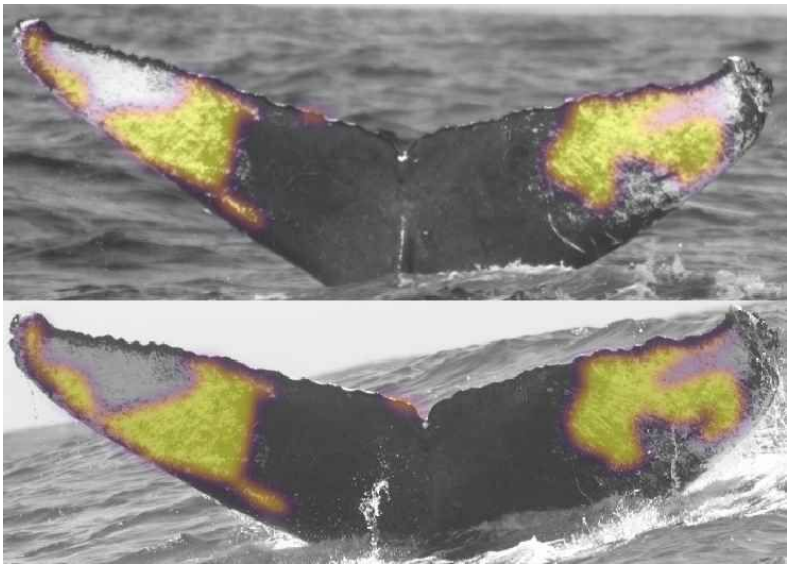

Figure 1c. A humpback whale (Megaptera novaeangliae) fluke is matched using a comparison of white and black contrast of its fluke in Flukebook. Wildbook provides multiple techniques for matching several species using an automated computer vision pipeline.