

A Framework for Autonomic Computing for In Situ Imageomics

Jenna Kline, Christopher Stewart, Tanya Berger-Wolf, Michelle Ramirez, Samuel Stevens, Reshma Ramesh Babu, Namrata Banerji, Alec Sheets, Sowbaranika Balasubramaniam, Elizabeth Campolongo, and Matthew Thompson

The Ohio State University

Charles V. Stewart and Maksim Kholiavchenko
Rensselaer Polytechnic Institute

Daniel I. Rubenstein
Princeton University

Nina Van Tiel
EPFL

Jackson Miliko
Mpala Research Centre

Abstract—*In situ imageomics* is a new approach to study ecological, biological and evolutionary systems wherein large image and video data sets are captured in the wild and machine learning methods are used to infer biological traits of individual organisms, animal social groups, species, and even whole ecosystems. Monitoring biological traits over large spaces and long periods of time could enable new, data-driven approaches to wildlife conservation, biodiversity, and sustainable ecosystem management. However, to accurately infer biological traits, machine learning methods for images require voluminous and high quality data. Adaptive, data-driven approaches are hamstrung by the speed at which data can be captured and processed. Camera traps and unmanned aerial vehicles (UAVs) produce voluminous data, but lose track of individuals over large areas, fail to capture social dynamics, and waste time and storage on images with poor lighting and view angles. In this vision paper, we make the case for a research agenda for *in situ imageomics* that depends on significant advances in autonomic and self-aware computing. Precisely, we seek autonomous data collection that manages camera angles, aircraft positioning, conflicting actions for multiple traits of interest, energy availability, and cost factors. Given the tools to detect object and identify individuals, we propose a research challenge: *Which optimization model should the data collection system employ to accurately identify, characterize, and draw inferences from biological traits while respecting a budget?* Using zebra and giraffe behavioral data collected over three weeks at the Mpala Research Centre in Laikipia County, Kenya, we quantify the volume and quality of data collected using existing approaches. Our proposed autonomic navigation policy for *in situ imageomics* collection has an F1 score of 82% compared to an expert pilot, and provides greater safety and consistency, suggesting great potential for state-of-the-art autonomic approaches if they can be scaled up to fully address the problem.

Index Terms—autonomous flight, UAVs, ecology, machine learning, computer vision, imageomics

I. INTRODUCTION

Many disciplines within the field of biology— colloquially, those ending with ‘omics’— seek to characterize biological structure, function, and dynamics via observation [1]. *Imageomics* is a new scientific field that seeks to extract biological traits of individual organisms, species, populations, and whole ecosystems from images using computational approaches [2].

Large sub-fields of biology, particularly ecology and conservation science, focus on studying organisms in the context of their environment, *in situ*. Recent sensing and imaging technology developments have accelerated the ability to gather biological data in the field at scale and machine learning techniques can extract insights from this data, also at scale. However, existing systems are not yet matching the needs of *in situ* scientific inference at large ecosystem scales and high resolutions down to the individual organism simultaneously. For example, understanding the collective dynamics of behavior of individual animals and groups and their response to the habitat and environmental conditions requires fine grain individual to population scale observations over seconds to hours of time, moving over potentially large distances and changing habitats. Here, we propose a framework for behavioral biological trait observations and inference (called behavioral *in situ imageomics*) that matches such biological need in the field.

The remainder of the paper is organized as follows. We provide the necessary biological background and relevant related work in Section II. Section III provides the motivation for *in situ imageomics* and the questions and challenges it raises. Section IV details *in situ imageomics* computational workloads and the technical challenges that our framework proposes. Section V presents an *in situ behavioral imageomics* case study of Kenyan animal behavior data collection with UAVs. Finally, Section VI contains a road map to pursue the research agenda laid out in this vision paper which we hope the community will be interested in addressing.

II. BACKGROUND AND RELATED WORK

Biologists are interested in collecting a variety of data about the phenotype and functioning of organisms, particularly trait data, ranging from morphology (physical characteristics) to behavior, in the environmental context, *in situ*. The massive amount of biological data captured by experts and citizen scientists (e.g., iNaturalist [3]), can provide new insight into the continuing evolution of species and habitats as well as the actions of individuals over time [4]. Modern machine learning

and computer vision techniques are required to process the vast amount of biological data available and these methods demand massive volumes and quality of data beyond what is typically curated for biological analysis. Adaptive, data-driven approaches to collect such datasets are hamstrung by the speed at which data can be captured and processed, especially behavioral data. Existing field data collection techniques are time-consuming and expensive, or infeasible for a human observer. Even images and videos collected in the field require labor-intensive manual labeling and pre-processing before they can be analyzed using machine learning and computer vision tools [5]–[7]. Emerging technology, including camera traps (motion or heat activated stationary cameras) and UAVs can produce voluminous data of animals in their natural habitat. Camera traps mitigate observer bias and disturbance from human presence, but are constrained to data from a fixed location. Small unmanned aerial systems (sUAS) consisting of one or more UAVs and their control systems, such as those available commercially from manufacturers such as DJI and Parrot, can track animals dynamically and can traverse remote terrain more quickly and with less disturbance to wildlife than heavy-duty sport utility vehicles (SUVs) typically used in field work that requires off-roading in remote terrain. Such sUAS missions can capture fine-grained details, such as animal behaviors within their social and environmental context, previously unavailable to researchers. These missions require trained pilots who can manually conduct missions tailored to the geographic region and species. Replacing manually controlled sUAS with an autonomously controlled sUAS would significantly reduce the barriers for ecological research. Yet, current sUAS technology may lose track of individuals over large areas, fail to capture social dynamics, and waste time and storage on images with poor lighting and viewing angles. Noisy datasets require more memory to store, are more cumbersome transport, and require more pre-processing and cleaning before they can be analyzed. Edge computing techniques that allow for more precise data collection and pre-processing at the data source can overcome these challenges.

Identifying common and disparate behaviors between individuals, demographic classes, and species within their common social and environmental contexts is a challenging biological problem because it requires repeated, fine-grained observations of behavior across multiple settings. Collecting behavior data *in situ* allows animal ecologists to capture the behavior and social interactions of all members of a group of individuals using sUAS, thus providing the data for each individual in order to compare similarities and differences among group members. Recording videos or images using UAVs provides the fine-grained detail of focal sampling with the breath of scan sampling, giving researchers the advantages of both techniques [5], [8]. As commercial UAVs become cheaper and more readily available, efforts have been made to use them to automatically capture video and photo data for animal ecology, conservation, and agriculture applications. Computer vision techniques can automatically track the location and postures of wildlife from the nadir-angle (i.e. bird's eye view

and even classify individuals by species and age-sex class from this data [9]. This approach can successfully reconstruct detailed models of the landscape, which can aid biologists in understanding how animals' movement and decision-making is shaped by their environment and social context. UAVs have been used to fight poaching of endangered species, especially if UAVs are equipped with long wave thermal infrared cameras to identify poachers and animals at night [7]. In [10], UAVs were used to count and individually identify animals based on their unique morphology. The animals of interest in this study could be individually identified using photos collected from the nadir-angle. This approach is not suitable for all species, such as zebras and giraffes, but demonstrates the effectiveness of pairing UAVs with computer vision techniques to detect and identify individual animals. Several computer vision algorithms have been proposed to automatically detect and track objects with sUAS. However, industry and academia have primarily focused on automatically tracking individual humans or vehicles. In [11], a convolutional neural network is used to detect objects in real-time, and adjust the UAV's movements and position to keep the focal object in view. A method utilizing inertial measurement unit data, GPS data, and the moving object detector to calculate the distance between the UAV and the object of interest to adjust the position of the UAV is proposed in [12]. A multi-agent approach for cooperative object localization and tracking is described in [13]. Finally, [14] proposes a design for UAVs to simultaneously complete real-time target tracking and path planning.

III. MOTIVATION FOR IN SITU IMAGEOMICS

Recent research has explored *in situ* imageomics, i.e., gathering images in the field and linking them to knowledge bases to characterize wildlife populations. For example, Wildbook harvests photos taken by citizen scientists to identify individual animals [15] to accurately estimate species' population levels. Obtaining accurate census information is crucial to secure legal protections and funding to preserve at-risk species [16], such as the highly endangered Grevy's zebra. Biologists believe that *in situ* imageomics could provide individual identification for a wider range of animal species, provide insight into individual and social behavior, intra- and inter-species dynamics, and deeper understanding of ecological processes. *In situ* data can provide environmental and social context, and capture traits that morph over time and biogeographic contexts, providing an opportunity for new, data-driven approaches to scientific insight, conservation, and sustainable biodiversity management [4], [17], [18]. We quantify the volume and quality of *in situ* imageomics computational workloads, including the data and autonomous navigation models, by describing a recent *in situ* imageomics data collection project at the Mpala Research Centre in Laikipia County, Kenya in January 2023. We collected behavioral data under a variety of social and environmental contexts of three species: Plains zebra (*Equus quagga*), Grevy's zebra (*Equus grevyi*), and reticulated giraffes (*Giraffa reticulata*) [5]. We conducted manual UAV missions to capture video footage of these species, along with

scan and focal sampling of their behavior in the field [8], employing current state-of-the-art data collection techniques. Using this dataset, we present initial steps in deploying autonomic computing system for in situ imageomics: (1) We propose a novel navigation methodology to collect animal behavior data autonomously, and demonstrate how this approach can be used to collect additional in situ imageomics datasets. (2) We convert real-world UAV video datasets to flexible data structures that can be used to train and test new autonomous navigation models.

IV. IN SITU IMAGEOMICS COMPUTATIONAL WORKLOADS: DATA & AUTONOMOUS NAVIGATION MODELS

A. *In situ* imageomics sUAS data collection missions

The most useful animal ecology data from sUAS is obtained from off-nadir views, where the UAV is high enough to avoid surface-level obstructions, such as vegetation shown in Fig. 1, but still able to capture fine details, such as behavior and unique morphological markings, as shown in Fig. 3. Table I summarizes the results produced by the different data collection angles shown in Figs 1, 2, 3. Several techniques are available to researchers to collect animal behavior data from a non-nadir angle. The first technique is to simply manually pilot the UAV to collect photos and videos. A manual approach does not require any computationally complex navigation models and allows the pilot to tailor their mission to collect specific animal behavior or individually identifiable markings. A manual approach produces a high percentage of flight time capturing usable data. For animal ecology applications, usable data is defined as frames containing images of animals with sufficient resolution to detect behavior, observe markings for individual identification or both. State of the art computer vision techniques used to process video and images typically require at least 30 pixels for object detection, and 700 pixels for individual detection [19]. However, manual approaches require an expert pilot which are expensive to hire and may not be available in all areas. It also introduces human error and makes consistent data collection difficult. Automating UAV missions resolves several problems associated with a manual approach. Previous work has found that automatic missions are 3x less expensive, are safer, and produce better spatial accuracy than manually piloted flights [20]. Many commercially available sUAS offer built-in waypoint mission planning, sometimes referred to as 'lawn-mower missions' in which the UAV automatically flies a grid-based pattern. These automatic flights visit waypoints to collect videos or photos from a nadir angle, which can capture data faster and with greater accuracy than manually piloted missions [21]. Such missions have been successfully implemented in agriculture applications [20], [22], forestry [23], [24], and animal ecology [9], [10]. This approach may not always be suitable for animal ecology applications since as noted above, it is often difficult to detect behavior from a nadir-angle. For some species, such as zebras and giraffes, the individually identifiable morphological features are not visible from above. Object detection algorithms are typically trained on images

TABLE I
COMPARISON OF PHOTO AND VIDEO COLLECTION ANGLES

Evaluation Criteria	Data Collection Angle		
	<i>Ground</i>	<i>Nadir</i>	<i>Non-nadir</i>
Occlusion by vegetation?	Yes	No	No
Biometric markings visible?	At times	No	Yes
Behavior visible?	At times	At times	Yes



Fig. 1. Ground-level view, animals' distinguishing markings and behaviors obscured by foliage.

of animals taken from the ground-level, like iNaturalist [3], and require extensive retraining to detect and classify animals from above [9], [10]. One solution is to automate non-nadir flights by instructing the UAV to fly to a waypoint, and rotate about the vertical axis 360 degrees while gathering photos or videos. However, it is difficult to pre-plan waypoint missions to capture sufficient usable data since animals' movements are unpredictable. Pausing the UAV at each waypoint to capture data reduces the area that can be covered in one battery charge thus reducing the mission throughput. Rotational maneuvers generate extra noise which may cause the animals to move out of the pre-planned mission coverage.

Commercially available UAVs equipped with 'follow-me' technology, in which the UAV 'locks-in' on an object of interest and adjusts its actions to keep the focal object in view, can be modified to track individuals animals and, by doing so, track that individual's interactions in a herd from a non-nadir angle. Problems arise when an individual separates itself from the herd. For example, if the UAV is 'locked in' on that individual, it will lose sight the main group and not be able to collect the social interactions and behaviors of the other individuals. As another problem, an object trackers can accidentally switch focal animals, especially in tight groups, leading to inconsistent data collection. These challenges necessitate a novel approach to object tracking with sUAS, in which the group as a whole is the priority for tracking. The group 'follow-me' navigation model optimizes for the best view of all individuals in the group, not one individual. Prioritizing the center of mass for the group allows for individuals to deviate their behavior without negatively impacting the navigation policy for the group. Tracking groups from the non-nadir angle produces more usable data than tracking and collecting photos and videos from the nadir-angle, but requires addition computation to translate images taken from this angle to navigation. The computation required to detect each individual in the group increases as the number of animals in the group increases, which increases task execution time. Increased execution time will degrade throughput, and in the worst case, may cause the UAV to lose sight of the group. This challenge is an opportunity for autonomic and adaptive computing to allocate resources intelligently to scale compute



Fig. 2. Nadir-view, animals' distinguishing markings on hip & shoulder not visible and behavior difficult to identify



Fig. 3. Off-Nadir View, not obscured by foliage; behavior and distinguishing markings easily seen.

to meet the demands of this application.

B. System for autonomous *in situ* imageomics data collection

We propose a navigation algorithm to allow sUAS to autonomously track groups of individuals, and a data-driven approach to test the performance of the navigation algorithm.

1) *Autonomous navigation policy*: The group 'follow-me' functionality uses an object detector to identify an object of interest—a herd of zebras or a tower of giraffes—and automatically adjusts the UAV's movements to keep that object within its field of view. This navigation policy, shown in Fig. 4, is optimized to identify and track zebras for behavioral data collection and individual identification, but this approach has broad applicability across different animal ecology applications, stealth tracking, and search and rescue UAV missions. Similar to existing 'follow-me' policies, this navigation policy assumes the human pilot is able to navigate the UAV towards the object of interest, in this case, a herd of zebras, until the UAV is sufficiently close to detect the animals of interest, at which point of the autonomous navigation policy takes control with human supervision. This approach uses YOLO pre-trained on the COCO dataset for object detection and classification [25]. Fig. 5 shows the calculations performed by the autonomous navigation policy in steps 3 and 4 of Fig. 4. If the centroid of the herd is sufficiently far from the centroid of the camera, or passes into the outer region of the frame, the UAV will adjust its position left or right, and forward or back, to return the herd centroid to the center region of the camera. The proposed navigation policy in Fig. 4 is the first step in the development of a sequence of maneuvers for sUAS to gather *in situ* imageomics data autonomously. Algorithm IV-B1 shows how a generic *in situ* imageomics data collection mission will be executed end-to-end, with Fig. 4 giving a flow chart of steps 3-6 of the algorithm. With Algorithm IV-B1 as a starting point, the autonomous navigation policy can be augmented with additional maneuvers tailored to each biological question.

Algorithm 1 In situ imageomics data collection mission

```

1: while sUAS has sufficient return-to-home (RTH) battery do
2:   Execute scanning pattern to locate small local herds
3:   if Herd is identified then
4:     Approach herd and reduce altitude
5:     Detect individual animals
6:     for All detected instances do
7:       Acquire individual ID photo
8:     end for
9:   end if
10:  Return to wider view of herd to record behavior
11: end while
12: Return to home to recharge

```

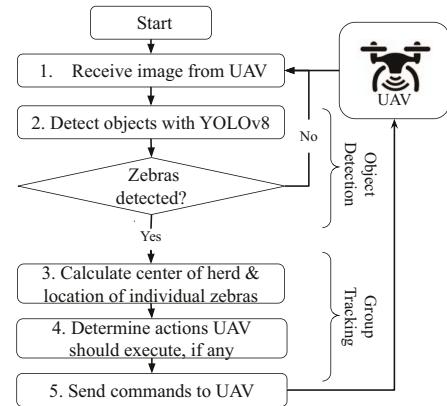


Fig. 4. Autonomous navigation policy for individual ID

2) *In situ imageomics autonomy cubes*: To test the navigation policy proposed here, as well as future autonomous maneuvers, we propose a novel methodology to simulate off-nadir view missions using real-world UAV footage. Autonomy cubes are a specialized data structure containing all UAV-sensable points in a discretized environment as an n -dimensional hypercube, introduced in [26]. The path of an autonomous sUAS cannot be predicted prior to the mission execution in a real environment which makes such systems difficult to build. To test these systems, simulations or traces are commonly used. However, simulations are synthetic and traces can only provide a single path of execution. Autonomy cubes allow the autonomous sUAS to query for real sensed data that would result from an action taken by the sUAS in a real environment. As shown in Table II, *in situ* imageomics autonomy cubes include a temporal element, since the state of the environment changes quickly over time and location as the animals move throughout the landscape. These autonomy cubes also include a heading descriptor—the direction the UAV is facing, since the sensed data will be different depending on the heading of the UAV at a non-nadir angle. The images collected during the real-world mission are manipulated using cropping and zoom to simulate sensed data from a 12-meter radius around the original flight path. For example, if the model recommended a deviation from the original flight path that was 3 meters forward and 3 feet to the left, the model would receive a zoomed in version of the original image with the right third cropped out. The wide-angled UAV camera

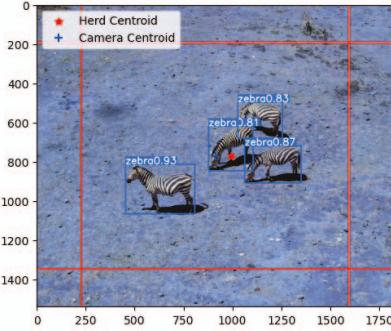


Fig. 5. Centroids and right, left, top, and bottom ranges of frame

TABLE II
IN SITU IMAGEOMICS AUTONOMY CUBE

Element	Description
Lat	GPS location, latitude
Long	GPS location, longitude
Altitude	Height above launch point in meters
Heading	Degrees, generalized to 8 cardinal directions
Time	Time elapsed since mission start in seconds
Photo	Video frame at location, altitude, heading, & time

lens generated a large image allowing for this manipulation while maintaining data integrity. If the model recommended an action that placed the UAV outside the 12-meter radius of sensed data, or facing the away from the animals, the model would receive an image of an empty savanna landscape.

C. Edge network requirements for *in situ* imageomics

Autonomous navigation and data management for *in situ* imageomics applications require autonomic computing solutions to manage the resources across the edge network. For sUAS, the edge network includes UAVs as sensors, as well as mobile devices, laptops, and remote controllers. *In situ* imageomics applications are characterized by resource-scarce, remote environments. Under such resource constraints, the edge network should make efficient use of edge and sensor power and compute resources while minimizing latency. Our proposed system makes use of the limited compute resources available on UAVs by running a light-weight object detector model locally. In addition to efficient use of edge resources, the system must conduct near real-time analysis of image data. This is required for the computer vision-driven autonomous navigation, and to make decisions about what data to send to the cloud for further analysis. Our proposed Algorithm IV-B1 allows for more precise data collection by positioning the UAV optimally depending on the mission phase: individual ID or behavior data collection. This approach produces a cleaner raw dataset compared to existing methods, which speeds up data transportation and analysis.

V. IN SITU IMAGEOMICS CASE STUDY: KENYAN ANIMAL BEHAVIOR DATA COLLECTION WITH SUAS

A. *In situ* imageomics autonomy cube

Using a manually piloted DJI Air 2S UAV, we collected behavioral video data of three species: Plains zebra (*Equus*

quagga), Grevy's zebra (*Equus grevyi*), and reticulated giraffes (*Giraffa reticulata*) over a period of three weeks at the Mpala Research Center in Laikipia, Kenya. We collected over 17 hours of video footage in 4K or 5.4K resolution which resulted in a raw dataset of 1.1 terabytes. Using this raw footage, we curated the raw data to create a dataset that serves as an initial benchmark for autonomic computing workloads for *in situ* imageomics, i.e. the computational and performance of self-adaptive systems for this application. We selected a subset of the raw video footage to create *in situ* imageomics autonomy cubes that can be utilized to design, train, and test new navigation policies. We ran experiments to determine the computational overhead required to create autonomy cubes using a single CPU on the Ohio Supercomputer. The CPU is an Intel Xeon E5-2680 V4 (Broadwell), with 28 cores per nodes, 128 GB memory, and 1.5 TB local disk space. For a three minute manually piloted UAV mission, it takes 109 seconds to create a corresponding autonomy cube. Each battery charge for a DJI Air 2S has a flight time of approximately 15-20 minutes for each *in situ* imageomics mission, so each will require approximately 9-12 minutes to create the cube.

B. Computer vision autonomous navigation for data collection

The navigation model described in Fig. 4 was tested on several *in situ* imageomics autonomy cubes described above. We tested the performance of the navigation policy using three pre-trained YOLO models: version 5 (YOLOv5su), and version 8 nano (YOLOv8n) and medium (YOLOv8m) [25]. The image data was pre-processed by cropping the images to eliminate the extra background generated by the wide-angle view. As shown in Table III, all three models performed well in identifying frames that contained zebras. YOLOv5su and YOLOv8m were able to detect approximately 80% of the herd on average, while YOLOv8n detected 66% of the herd on average. Next, we conducted an experiment to examine the impact of the YOLO object detector's output on the decisions made by the navigation model on where to fly next. All three models performed the same actions as the original mission more than two-thirds of the time as shown in Table IV. Despite having the highest levels of accuracy in detecting zebras, YOLO5su performed the worst in determining next steps to take. All models performed similarly in regards to precision and accuracy with an F1-score of approximately 80%. Interestingly, despite detecting 11% fewer animals compared to YOLOv8m, YOLOv8n had the highest level of accuracy in determining the next action to take. This suggests that our proposed navigation model is effective in tracking the herd as a whole, even when the object detector may miss some individuals in the group. These results suggest that smaller, lightweight object detection and classification models are well suited for conditions with limited resources.

VI. FUTURE RESEARCH AGENDA AND ROAD MAP

In situ imageomics depends on significant advances in autonomic and self-aware computing. We present the following challenges to the community:

TABLE III
YOLO OBJECT DETECTION PERFORMANCE

	yolov5su	yolov8n	yolov8m
% of frames with zebras detected	98.7	95.9	99.0
% of avg total animals detected	78.1	66.4	77.6

TABLE IV
AUTONOMOUS NAVIGATION POLICY RESULTS

	yolov5su	yolov8n	yolov8m
# of actions that differ	43	39	41
% of actions matching original flight	66.4	68.8	68.0
F1 Score	81.0	82.1	82.5

- Can self-aware sUAS avoid disturbing wildlife?** As sUAS approach, herds move away. Self-aware systems can now learn and adapt to internal constraints, e.g., mission goals and power with reinforcement learning, but learning herd-specific sensitivities to noise and movement introduces a challenging stochastic element to the problem.
- Can we develop what-if models for dynamic scenes?** Self-aware sUAS may employ multiple maneuvers to approach wildlife. Choosing the right model may require simulations that infer wildlife behaviors.
- Can autonomic computing systems adapt processing and storage capabilities quickly enough?** In the presence and absence of predators, herds disperse and reassemble rapidly. Herd size corresponds with computational demand for object detection and individual identification. In addition, many social behaviors emerge in large herds, further increasing computational load and memory requirements for complex models of social behaviors. Future research must explore adaptive edge architecture, possibly with cloud offload, for in situ imageomics.

Acknowledgments: This material is based upon work supported by the National Science Foundation under Award No. 2118240 and 2112606. The data used in this work was collected in accordance with Kenya's National Commission for Science, Technology & Innovation research license NA-COSTI/P/22/18214 and Princeton University's Institutional Animal Care and Use Committee (IACUC) 1835F.

REFERENCES

- [1] Omics- wikipedia. <https://en.wikipedia.org/wiki/Omics>, 2023.
- [2] Imageomics Institute. About the imageomics institute (nsf oac-2118240).
- [3] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *CVPR*, 2018.
- [4] Devis Tuia, Benjamin Kellenberger, Sara Beery, Blair R. Costelloe, Silvia Zuffi, Benjamin Risse, Alexander Mathis, Mackenzie W. Mathis, Frank Van Langerveld, Tilo Burghardt, Roland Kays, Holger Klinck, Martin Wikelski, Iain D. Couzin, Grant Van Horn, Margaret C. Crofoot, Charles V. Stewart, and Tanya Berger-Wolf. Perspectives in machine learning for wildlife conservation. *Nature Communications*, 13(1):792, Feb 2022.
- [5] Maksim Kholiavchenko, Jenna Kline, Michelle Ramirez, Samuel Stevens, Alec Sheets, Reshma Ramesh Babu, Nina Van Tiel, Namrata Banerji, Elizabeth Campolongo, Matthew Thompson, and et al. Kabr: High-quality dataset for kenyan animal behavior recognition from drone videos. *In submission*, 2023.
- [6] William Andrew, Colin Greatwood, and Tilo Burghardt. Deep learning for exploration and recovery of uncharted and dynamic targets from uav-like vision. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, page 1124–1131, Oct 2018.
- [7] Elizabeth Bondi, Fei Fang, Mark Hamilton, Debarun Kar, Donnabell Dmello, Jongmoo Choi, Robert Hannaford, Arvind Iyer, Lucas Joppa, Milind Tambe, and Ram Nevatia. Spot poachers in action: Augmenting conservation drones with automatic detection in near real time. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(11), Apr 2018.
- [8] Jeanne Altmann. Observational study of behavior: Sampling methods. *Behaviour*, 49(3–4):227–266, 1974.
- [9] B. Koger, A. Deshpande, J. Kerby, J. Graving, B. Costelloe, and I. Couzin. Quantifying the movement, behaviour and environmental context of group-living animals using drones and computer vision. <https://onlinelibrary.wiley.com/doi/abs/10.1111/1365-2656.13904>, 2023.
- [10] William Andrew, Colin Greatwood, and Tilo Burghardt. Aerial animal biometrics: Individual friesian cattle recovery and visual identification via an autonomous uav with onboard deep inference. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, page 237–243, Nov 2019.
- [11] Ali Rohan, Mohammed Rabah, and Sung-Ho Kim. Convolutional neural network-based real-time object detection and tracking for parrot ar drone 2. *IEEE Access*, 7:69575–69584, 2019.
- [12] Peng Chen, Yuanjie Dang, Ronghua Liang, Wei Zhu, and Xiaofei He. Real-time object tracking on a drone with multi-inertial sensing data. *IEEE Transactions on Intelligent Transportation Systems*, 19(1):131–139, Jan 2018.
- [13] Aamir Ahmad, Guilherme Lawless, and Pedro Lima. An online scalable approach to unified multirobot cooperative localization and object tracking. *IEEE Transactions on Robotics*, 33(5):1184–1199, 2017.
- [14] Yisha Liu, Qunxiang Wang, Huosheng Hu, and Yuqing He. A novel real-time moving target tracking and path planning system for a quadrotor uav in unknown unstructured outdoor scenes. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(11):2362–2372, Nov 2019.
- [15] Tanya Y. Berger-Wolf, Daniel I. Rubenstein, Charles V. Stewart, Jason A. Holmberg, Jason Parham, Sreejith Menon, Jonathan Crall, Jon Van Oast, Emre Kiciman, and Lucas Joppa. Wildbook: Crowdsourcing, computer vision, and data science for conservation, 2017.
- [16] Daniel I. Rubenstein. *Ecology, Social Behavior, and Conservation in Zebras*, volume 42, page 231–258. Elsevier, 2010.
- [17] Ethan R. Deyle, Robert M. May, Stephan B. Munch, and George Sugihara. Tracking and forecasting ecosystem interactions in real time. *Proceedings of the Royal Society B: Biological Sciences* 283(1822):20152258, Jan 2016.
- [18] Marc Besson, Jamie Alison, Kim Bjerge, Thomas E. Gorochowski, Toke T. Høye, Tommaso Jucker, Hjalte M. R. Mann, and Christopher F. Clements. Towards the fully automated monitoring of ecological communities. *Ecology Letters*, 25(12):2753–2775, 2022.
- [19] Jonathan P. Crall, Charles V. Stewart, Tanya Y. Berger-Wolf, Daniel I. Rubenstein, and Siva R. Sundaresan. Hotspotter — patterned species instance recognition. In *2013 IEEE Workshop on Applications of Computer Vision (WACV)*, pages 230–237, 2013.
- [20] Jayson Boubin, John Chumley, Christopher Stewart, and Sami Khanal. Autonomic computing challenges in fully autonomous precision agriculture. In *2019 IEEE International Conference on Autonomic Computing (ICAC)*, page 11–17, Jun 2019.
- [21] DJI. 6 best follow me drones (with video comparisons) 2020, Jun 2018.
- [22] Zichen Zhang, Jayson Boubin, Christopher Stewart, and Sami Khanal. Whole-field reinforcement learning: A fully autonomous aerial scouting method for precision agriculture. *Sensors*, 20(2222):6585, Jan 2020.
- [23] Linyuan Li, Jun Chen, Xihan Mu, Weihua Li, Guangjian Yan, Donghui Xie, and Wuming Zhang. Quantifying understory and overstory vegetation cover using uav-based rgb imagery in forest plantation. *Remote Sensing*, 12(22):298, Jan 2020.
- [24] Sebastian Egli and Martin Höpke. Cnn-based tree species classification using high resolution rgb image data from automated uav observations. *Remote Sensing*, 12(2323):3892, Jan 2020.
- [25] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. Yolo by ultralytics, Jan 2023.
- [26] Jayson G. Boubin, Naveen T. R. Babu, Christopher Stewart, John Chumley, and Shiqi Zhang. Managing edge resources for fully autonomous aerial systems. In *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing, SEC '19*, page 74–87, New York, NY, USA, 2019. Association for Computing Machinery.