

WILDPROP: Visual Estimation of Wildlife Body Proportions at Scale

Mustafa Chasmai, Aaron Sun, Subhransu Maji

{mchasmai, aaronsun, smaji}@umass.edu

Manning College of Information and Computer Sciences
University of Massachusetts, Amherst



Fig. 1: Estimating wildlife body proportions. Given a single annotated reference image and a set of specified body parts, WILDPROP estimates the population-level distribution of their length ratios from a large, uncuration image collection. We show exemplar species and body parts used in our experiments. All images are sourced from iNaturalist [27]; observer IDs are listed below each image.

Abstract. Population-level morphometric measurements are central to ecological and evolutionary studies, but traditionally require controlled imaging or physical specimen handling, limiting their scalability. We present WILDPROP, a framework for estimating wildlife body-proportion distributions directly from large-scale, unconstrained image repositories. We cast morphometric estimation as a retrieval-driven correspondence problem: given a single user-annotated canonical image, WILDPROP performs pose-aware retrieval using foundation model features, transfers part endpoints via dense patch-level matching, filters predictions using geometric consistency, and aggregates measurements across retrieved images to estimate length-ratio distributions. Unlike supervised keypoint pipelines, our approach adapts to arbitrary species and user-defined body parts without per-species training. Evaluations on three large morphometric datasets spanning birds and amphibians show median relative errors of 10–20%. We further demonstrate broad applicability through case studies measuring proportions across diverse taxa, including birds, frogs, insects, and flowers. Ablations show that pose-aware retrieval is critical for stable estimation, while robust aggregation mitigates keypoint and pose noise. Our results suggest that carefully filtered 2D correspondences over web-scale imagery can provide scalable morphometric proxies for comparative analyses across taxa, geography, and seasonality.

Keywords: Vision at scale · Ecology · Morphometrics · Pose estimation

1 Introduction

Quantitative morphometric measurements—such as ratios of limb, wing, or head dimensions—are central to ecology and evolutionary biology (see Fig. 1). Large-scale efforts such as AVONET [65] have demonstrated how curated trait datasets can inform macroevolution [14], community assembly [37], ecosystem modeling [21], and population health assessment [24, 32, 44]. Morphological variation within species across geography, and convergence across species within shared environments, are often diagnosed through such proportional measurements. These patterns have long served as the empirical basis for fundamental biological frameworks such as the theory of evolution [9].

Collecting morphometric measurements at scale, however, remains difficult. Traditional pipelines rely on physical specimen handling or controlled photography with standardized poses and rulers (see Fig. 2). These approaches can be precise, but are labor-intensive, disruptive to wildlife and ecosystems, and limited in coverage; many datasets contain measurements for only a small number of individuals per species. In contrast, public repositories such as iNaturalist [27] contain millions of wildlife images spanning taxa, geography, and time. These collections offer an appealing opportunity for scalable morphometric analysis, but geometric measurement from arbitrary photographs is ill-posed due to viewpoint variation, articulation, occlusion, scale ambiguity, and perspective distortion.

Our key observation is that although most images in such repositories are unsuitable for direct measurement, web-scale collections often contain a substantial subset of images captured in near-canonical poses where relative geometric relationships are approximately preserved. This suggests a different formulation: rather than reconstructing geometry from arbitrary views, we can selectively retrieve measurement-friendly instances and aggregate their statistics.

We introduce WILDPROP, a scalable, training-free framework for estimating population-level body-proportion distributions from large, unconstrained image collections (see Fig. 3). Given a single reference image in a canonical pose with user-defined part endpoints, WILDPROP: (1) retrieves pose-aligned images from the desired population using foundation-model representations; (2) transfers part endpoints via dense patch-level matching with geometric consistency filtering; and (3) aggregates measurements through iterative refinement to estimate stable ratio distributions. This formulation is user-friendly and flexible: new species and traits require only a single annotated query image, without per-species training or task-specific keypoint labels. Recent advances in self-supervised visual representations, such as DINO [62], and foundation segmentation models, such as SAM [55], make such training-free correspondence robust across diverse taxa.

We evaluate WILDPROP on three large morphometric datasets spanning birds and amphibians, mimicking the original measurement protocols and comparing estimated population ratios to physical ground truth. Given only a single annotated image, WILDPROP achieves 10–20% median relative error on various part-length ratio estimation tasks using large-scale image collections. Ablations demonstrate that viewpoint-aware retrieval, geometric verification, and large retrieval sets are critical for estimation accuracy. We further present case studies




Fig. 2: Morphometry paradigms¹. The choice between the measurement method is driven by a tradeoff between measurement costs and estimation errors.

illustrating broad applicability across taxa and subgroup analyses across geography and seasonality. We conclude with a discussion of limitations and directions for future work. Code available at [github:cvl-umass/wildprop](https://github.com/cvl-umass/wildprop).

2 Related Work

Morphometrics. Measuring morphological traits has a long history in ecology and evolutionary biology, where the shape, size, and structure of body parts are used to study species and their interactions within ecosystems [2, 43, 60]. A classic example comes from Darwin’s theory of evolution [9]: variation in finch beak morphology in the Galápagos Islands provided key evidence for natural selection. Morphometric traits can reveal evolutionary adaptations [6], help identify gender in some species [7, 49], and distinguish visually similar species [38]. Large-scale studies have measured morphological traits across diverse taxa, including plants [15, 30], mammals [25, 28, 29, 45, 72], amphibians [36], insects [16], and birds [56, 59, 65]. Global coverage of such measurements could enable the development of general ecosystem models [21], but the scale and diversity of species make this task extremely challenging. The lack of standardized measurement frameworks across taxonomic groups further exacerbates the problem. Beyond broad coverage, stratification is another key aspect of morphometric studies. Repeated measurements over time can reveal ecological change [17, 18, 51, 77], while analyses across demographic or geographic subgroups can provide insights into species variation [58] and population health [24, 32, 44]. However, such studies are typically limited to specific niches and are difficult to scale globally. Our approach, WILDPROP, scales to global taxonomic diversity, readily adapts to user-specified body proportions, and enables flexible stratification by appropriately restricting the retrieval set.

Vision for morphometrics. Prior efforts to measure body size and proportions tend to rely on physical, in-situ studies or controlled photography (Fig. 2), followed by digital measurement in an effort to reduce the required domain expertise and maintain persistent records for cross-verification. Compared to traditional physical measurement protocols, these approaches report significantly

¹ The first 3 images are borrowed from AVONET [65], [71] and [10], respectively.

reduced specimen handling time [71], which can enable faster data collection and improved scalability. This is a practical and scalable setup for livestock [41, 50] and museum specimens [22, 23, 70], but is quite cost-intensive for wildlife in natural ecosystems [7, 10, 34, 48, 71]. Efforts towards non-invasive alternatives, such as camera traps [12] and aerial imagery [1], have generally been limited to large mammals. In contrast, our approach is highly scalable, completely non-invasive and broadly applicable to diverse taxonomic groups.

Animal pose estimation. Animal pose estimation has received significant attention in recent years, building on advances in computer vision [3, 19, 40, 54, 75]. However, these methods typically assume fixed keypoint definitions and require extensive annotations. A notable example is DeepLabCut [42], which has been widely used for detection and tracking of animals but still relies on manual annotation and training on the order of 100–200 images per category. We note that even perfect 2D localization does not resolve measurement ambiguity due to scale and foreshortening effects arising from projection. While monocular depth estimation has improved substantially [31, 52, 76], existing models are optimized for large rigid objects and background geometry, and are unreliable for fine-grained articulated structures typical in wildlife imagery. General image-to-3D reconstruction has also advanced rapidly in recent years, including foundation models (e.g., SAM3D [4], Trellis [73, 74]) and recent extensions to animal 3D reconstruction [26]; however, these methods have not yet been widely studied for precise morphometric measurement. Instead of explicitly solving the 3D measurement problem, we mitigate perspective effects by conditioning on pose through retrieval, effectively restricting viewpoint variability before estimation. This also reduces the burden on the keypoint matching step, allowing us to generalize from a single example using features from modern foundation models [47, 62].

CV for ecology. Computer vision has been widely applied to ecological problems [66], driven in part by large community-driven platforms such as iNaturalist [5, 27, 67], eBird, and others. Recent work has also explored the use of large vision–language models [68] to enable ecologists to query arbitrary attributes beyond species identity. Automated vision approaches have already seen some success in large scale investigations of qualitative traits, such as species coloration and plumage [8, 20]. Our work on quantitative morphological measurements is complementary to these efforts. For example, attributes such as age, sex, color, or geographic location could be used to further stratify population-level estimates beyond species. We investigate several such examples in our case studies.

3 Methodology

The overall approach is illustrated in Fig. 3. The user provides a single annotated image (query) with multiple specified body parts (at least two). We assume that these parts are approximately fronto-parallel in the image, so that foreshortening effects are negligible. The goal is to estimate the ratio of these body parts from a large collection of relevant images (e.g., images of the same species or further




Fig. 3: Method overview. The proposed approach consists of three key components: (1) pose-aware retrieval to identify images from a large collection (*e.g.*, iNaturalist) where the desired measurements can be reliably estimated; (2) keypoint matching and geometric alignment to transfer part annotations; and (3) iterative refinement of keypoint appearance and correspondences to estimate population-level body proportions. In this example, the user provides a canonical image with annotated keypoints defining the ratio of wing width to length (left). The top three pose-aligned retrieved images and their estimated correspondences are shown in the middle panels, from which the distribution of wing width-to-length ratios is computed. Also, note that in this case, due to symmetry, each instance provides two measurements of width and length.

stratified subsets). Our approach consists of three steps: (1) retrieving images from the collection with object poses similar to the query image; (2) performing keypoint matching and alignment; and (3) iteratively refining appearance and alignment to estimate population-level distributions of part-length ratios.

3.1 User Query

The user provides a query image I_q , a set of keypoints $K = \{(x_1, y_1), \dots, (x_k, y_k)\}$, and a set of parts $P = \{(i_1, j_1), \dots, (i_n, j_n)\}$, where $1 \leq i_\ell, j_\ell \leq k$ and $n \geq 2$. Each part (i_ℓ, j_ℓ) specifies the two keypoints defining a body segment, and (x_i, y_i) denotes the image location of the i -th keypoint. The goal is to estimate ratios between the lengths of the specified parts over a target image collection.

Discussion. WILDPROP assumes that the annotated parts are approximately fronto-parallel in the chosen view. This assumption may not hold when the parts needed for a ratio are best observed from different viewpoints (*e.g.* bill length from a side view and bill width from a frontal view). In such cases, the ratio can still be estimated by using a common anchor part visible across views (*e.g.* wing or body length) and applying WILDPROP with multiple query images. For example, one can estimate each part relative to the shared anchor in its appropriate view, and then combine these estimates to recover the desired cross-view ratio.

3.2 Pose-aware Retrieval

We assume access to a large image collection representing the target population. In most experiments, this collection consists of high-resolution iNaturalist

images [27] of the same species, though we also perform analyses stratified by season and geographic region. The images are captured in the wild and therefore vary substantially in pose, viewpoint, background, scale, and image quality. We do not assume access to camera calibration information.

We first retrieve images whose poses are similar to that of the query image. Patch-level DINOv3 embeddings are effective for pose-aware retrieval, but can be sensitive to background features. To reduce background influence, we segment the target object using SAM2 [55] together with Grounding-DINO [39], using simple text prompts such as “bird” or “frog,” and zero-out the background. Each segmented image is resized to 518×518 , and DINOv3 patch embeddings are extracted and concatenated into a single image descriptor. We then use cosine similarity between descriptors as the retrieval criterion. We explore alternative retrieval variants in Tab. 4b. In our experiments, we retain the top 100 matches for each query. Although this is only a small subset of the full collection, it is often substantially larger than the number of individuals in traditional morphometric datasets. For example, AVONET [65] reports fewer than nine individuals per species on average.

Motivation. The goal of this stage is to efficiently retrieve a small set of candidate images from a large collection, typically containing 10^4 – 10^5 images. We therefore use global image-level descriptors for fast retrieval, and defer fine-grained keypoint alignment and geometric verification to the next stage. This coarse-to-fine design parallels classical SIFT-based instance retrieval pipelines in computer vision [46, 63], where efficient global retrieval is followed by local feature matching and geometric verification.

3.3 Keypoint Matching and Alignment

For each retrieved image, we transfer the part annotations from the query image through patch-level keypoint matching. We construct background-suppressed images by halving the intensity of pixels identified as background by SAM2 [55] and extract their DINOv3 [62] patch embeddings. For each query keypoint (x_i, y_i) , we extract the 768-dimensional embedding of the nearest 16×16 patch in the query image I_q . We then find the patch in the retrieved image with the highest cosine similarity and use its center as the predicted corresponding location (x'_i, y'_i) . In our ablation studies, we find that higher-resolution inputs improve matching accuracy, but the gains tend to saturate beyond 518×518 .

Individual matches can be error-prone, especially for symmetric or visually similar parts such as wing tips, legs, or repeated body structures. We therefore apply a geometric consistency check. Given the predicted correspondences, we use RANSAC [13] to fit a rigid transformation between the query keypoints and their predicted locations in the retrieved image. Images with more than a specified number of outlier correspondences are excluded from estimation. We use a relatively permissive outlier threshold to account for natural variation in shape and pose across individuals.

3.4 Iterative Refinement and Ratio Estimation

Keypoint appearance features extracted from a single query image may not capture the visual diversity of the target population, leading to alignment errors. To address this, we introduce an iterative refinement step that adapts the keypoint descriptors to the retrieved image set. Given the correspondences estimated in the previous step, we update each keypoint descriptor by averaging the corresponding patch features from the top K retrieved images together with the original query feature. The parameter K controls the degree of adaptation to the target distribution. This refinement is particularly useful when transferring annotations across species or across visually diverse subpopulations.

The updated keypoint descriptors are then used to re-estimate correspondences in the retrieved images, and the process is repeated for a fixed number of iterations. After refinement, we compute the length of each specified part in every retained image and form the corresponding part-length ratios. These per-image ratios are then aggregated to estimate population-level statistics such as the mean, median, and variance.

3.5 Implementation Details

We use the ViT-B/16 DINOv3 backbone [62] for both retrieval and keypoint matching, and the ViT-H variant of SAM2 [55] for masking. For iterative refinement, we perform three iterations. At each iteration, we update keypoint descriptors using the top 20 retrieved images together with the original query descriptors, aggregated by a weighted average with weights proportional to feature similarity. These hyperparameters were selected based on relative error on the AVONET California Gull validation setting [65] and then applied unchanged to all other species and datasets.

For RANSAC, we use a fitting threshold of 5% of the image diagonal and discard a retrieved image if more than 20% of the predicted keypoints are classified as outliers. Additional hyperparameter details are provided in Appendix A1. For the CLIP variants explored in Sec. 4.3, we use the ViT-B/16 backbone. For Stable Diffusion [11], we use the v3.5-medium model. All experiments were conducted on a single NVIDIA A100 GPU.

4 Experiments and Results

4.1 Datasets and Evaluation Metrics

We validate this approach on three large morphometric datasets spanning birds and amphibians, as well as several case studies involving other taxa, and show that large image collections on platforms such as iNaturalist, combined with representations from existing foundation models, enable scalable, population-level estimation of body proportions.

For each experiment, we perform retrieval on full-resolution images from iNaturalist [27] that (1) belong to the search corpus, (2) were captured on or

before 31 December 2025, and (3) are licensed for research use. We also manually search for suitable query images on iNaturalist and label keypoints for the various body parts using Labelbox [35]. Example species and body parts used in our experiments are shown in Fig. 1.

For our quantitative experiments, part specifications vary across datasets, and we annotate keypoints according to the prescribed measurement protocols for each dataset. Consider AVONET [65], for example (see Tab. 1, right): we measure bill length by placing one keypoint at the tip of the bill and another at the anterior edge of the nostrils (nares). Some parts are not feasible to measure visually—for example, AVONET includes another bill measurement from the tip to the base of the skull, which is typically covered by feathers—and we therefore exclude such measurements from our experiments. Each dataset reports linear measurements for individual specimens across various species, and we use proportions derived from these measurements as references for our estimates.

We generally compute all part proportions relative to the longest part (*e.g.* wing length for birds) for numerical stability. Factors such as age, gender, and health introduce variance in measurements across individuals of a species, and we aim to capture the overall distribution of each body proportion. In the absence of individual correspondence, we compare summary statistics. Concretely, we compute the median proportion estimated by WILDPROP ($\hat{\mu}$) and the median proportion measured in the dataset (μ), and report the relative estimation error as $E = |\hat{\mu} - \mu| / \mu$.

For our qualitative case studies, we explore diverse taxonomic groups and select species with distinct morphological patterns that can be qualitatively verified. For example, we analyze butterfly species with substantially different wing dimensions and evaluate whether the estimated proportions separate them. We use kernel density estimation [61] to visualize the estimated distributions.

4.2 Quantitative Results

Tabs. 1 to 3 show quantitative comparisons with existing physical morphometric studies on birds [65], shorebirds [59], and frogs [36], respectively. Query images used for the different canonical poses are shown next to each table, with all measured parts visualized as green lines. We also visualize selected model predictions from these datasets in Fig. 6 and the Appendix A5.

AVONET. Tab. 1 shows results on AVONET across five species spanning several bird families. Following the measurement protocol outlined by the authors, we estimate four parts across two poses: bill length, tarsus (leg length), wing width, and wing length. On average, our measurements have about 11% relative error compared to the ground truth median estimates. Fig. 4 shows the median and the 25–75 quantiles estimated using our approach and the ground truth across species. While the medians follow similar trends, our estimates have larger interquartile ranges. However, given the small sample sizes ($n \leq 11$) in these studies, it is unclear whether the physical measurements capture the full variation in the population.

Table 1: AVONET [65]. Relative error (%) in median proportions. Lower is better. To the right are example query images and keypoints: side view for Cape May Warbler (bill, tarsus, and wing); and bottom view for California Gull (wing length and width).

Species	Bill L. / Wing L.	Tarsus / Wing L.	Wing W. / Wing L.	Mean
Great Egret	16.7	8.6	9.0	11.4
Blue Jay	4.0	12.9	8.8	8.6
White-tailed Kite	15.9	15.7	7.7	13.1
Cape May Warbler	14.1	8.1	15.2	12.5
California Gull	0.4	14.9	18.0	11.1
Mean	10.2	12.0	11.7	11.3



Shorebirds. Tab. 2 shows results on five species from the Alaska Science Center adult shorebird dataset [59]. Using a single reference image per species, we measure four parts: exposed culmen (an alternative measure of bill length), tarsus, head length, and wing length. On this dataset, our measurements have about 11% relative error compared to the ground truth. We observe more accurate results for tarsus and culmen ratios than for head length. We speculate that keypoint matching is less accurate for head length because the back-of-head keypoint is less distinctive than the keypoints used for other parts, such as the bill tip or wing tip.




Fig. 4: Tarsus:Wing-length ratio medians and quartiles across AVONET species. Sample sizes for both WILDPROP estimations and physical measurements are indicated at the base of each bar.

Frogs. Tab. 3 shows results on five frog species covered in the University of Michigan Deep Blue Data [36]. Using two reference images per species, we measure four distances: eye-eye, eye-nostril, head width, and nostril-nostril. Similar to the head length for shorebirds, we attribute the poor head width estimate to indistinct keypoints for this measure compared to other parts. Despite the large number of available images, the canonical poses required for our analysis were relatively rare for the Pacific Chorus Frog, which likely explains its worse performance across parts.


Discussion of quantitative results. We do not observe a consistent pattern across all species or parts, as estimation error is influenced by the abundance of target poses, the distinctiveness of the keypoints, and the size of the part relative to the image. Barring a few exceptions, the estimated proportions are within 10–20% of the actual physical measurements, and inter-species trends align with expectations (*e.g.* longer tarsus for Egret than Warbler in Fig. 4).

Table 2: ShoreBirds [59]. Relative error (%) in median proportions. To the right is an example query: Ruddy turnstone with culmen, head, tarsus, and wing keypoints.


Species	Tarsus / Wing L.	Culmen / Wing L.	Total Head / Wing L.	Mean
Ruddy Turnstone	5.2	7.3	13.9	8.8
Dunlin	12.8	7.5	9.0	9.8
Western Sandpiper	15.4	15.3	23.1	17.9
Bar-tailed Godwit	11.0	5.9	17.6	11.5
Lesser Yellowlegs	5.4	6.8	9.0	7.1
Mean	10.0	8.6	14.5	11.0

**Table 3: Frogs [36].** Relative error (%) in median proportions. Lower is better. To the right are example query images and keypoints: front view for Emerald Glass Frog (eyes, nostrils, and head width); and top view for Red-eyed Tree Frog (eyes and nostrils).

Species	Eye-Nostril / Eye-Eye	Head Width / Eye-Eye	Nostril-Nostril / Eye-Eye	Mean
Red eyed Tree Frog	18.9	12.6	12.6	14.7
Woodhouse’s Toad	8.3	30.0	18.5	18.9
Emerald Glass Frog	12.6	21.2	11.4	15.1
Cuban Tree Frog	7.5	22.8	17.1	15.8
Pacific Chorus Frog	34.4	22.3	25.3	27.3
Mean	16.3	21.8	17.0	18.4

**Pose availability & model performance.**

The availability of images in suitable poses is an important factor influencing WILDPROP performance and likely explains several trends observed in our experiments (*e.g.*, Pacific Chorus Frog in the front-view setting, and Warbler and Blue Jay in the bottom-view setting). To investigate this effect, we quantify pose availability using the mean similarity of the top 100 retrieved images and compare it with estimation error in Fig. 5. The observed negative correlation suggests that retrieval similarity can serve as a useful predictor of performance: when more pose-aligned images are available, WILDPROP tends to produce more accurate estimates.

**Fig. 5: Pose availability.** Relative error vs mean retrieval similarity for each species and pose.**4.3 Ablations**

We ablate the key components of WILDPROP in Tab. 4. First, we explore the contribution of each module in the overall pipeline (Tab. 4a). Using estimates




Fig. 6: Prediction Visualizations. Sample queries and model predictions. We demonstrate the pose-similarity of retrieved images, and matched keypoints on them.


from the query image alone results in a 12.2% AVONET relative error. However, WILDPROP also provides an estimate of the full distribution, capturing intra-species variation, and supporting cross-species queries (*e.g.* birds of prey in Fig. 8). Matching keypoints on randomly sampled images yields an exorbitant 216% error, highlighting the difficulty of measuring parts from arbitrary poses. Restricting our analysis to pose-retrieved images reduces the errors to 40.0%. Using background-suppressed images based on SAM [55] segmentation reduces the error to 16.1%, while RANSAC-based filtering further reduces it to 12.2%. Finally, iterative refinement leads to a modest 1% drop in error.

Next, we compare different model variants for retrieval (Tab. 4b) as well as for keypoint matching and iterative refinement (Tab. 4c). CLIP-based models used for retrieval instead of DINO lead to worse performance, underscoring the importance of features that encode geometry and pose. Among the DINO variants, using the CLS token yields roughly the same performance as using a concatenation of all patch tokens, while switching to DINOv2 [47] improves AVONET performance to 10.1% relative error but worsens validation performance (hence we use DINOv3 instead). For keypoint matching, using a simple unweighted average leads to a modest increase of 0.7%, and removing the restriction to top- k retrieved images during aggregation does not affect the relative error on AVONET. Among the model variants, using DINOv2 features for keypoint matching yields a 2.2% higher error. Interestingly, while Stable Diffusion [11] features worsen performance to 25.2% relative error, concatenating DINOv3 and Stable Diffusion features performs best, with an 11.2% error.

Next, we investigate the influence of retrieval sample size (Fig. 7a) on estimation errors. As the number of images available for a species increases, so does the likelihood of finding images with poses suitable for measurement, and consequently, measurement accuracy. We simulate different sample sizes by randomly sampling fractions of images and observe the resulting trend for the Great Egret, which has approximately 150K images. For smaller sample sizes, we observe retrieved images where parts of the legs are submerged in water, which likely explains the sharp increase in errors for the tarsus proportion. In contrast,

Table 4: Ablations. For each experiment, we report relative errors in medians for AVONET and highlight the row indicating the configurations used in WILDPROP.

Variant	AVONET	Variant	AVONET	Variant	AVONET
Query Only	12.2	CLIP [53]	17.1	DINOv3	11.3
Random Samp.	216.1	SigLIP [78]	11.8	w/o weigh. mean	12.0
+ Retrieval	40.0	BioCLIP [64]	14.8	w/o top-k kp filter	11.3
+ SAM	16.1	DINOv3-cls	11.3	DINOv2	13.5
+ RANSAC	12.2	DINOv2-patch	10.1	Stable Diff. [11]	25.2
+ Iterative	11.3	DINOv3-patch	11.3	DINOv3+SD [11]	11.2
(a) WILDPROP components		(b) Retrieval variants		(c) Kp. Matching & Refinement	

**Fig. 7: Ablations.** We present errors for the Great Egret. (a-b): 10 runs mean $\pm$ std.

the bill-to-wing-length proportion is more stable—as images with these parts clearly visible are abundant—and only begins improving near 100K images.

Next, we investigate the influence of query keypoint noise (Fig. 7b). Small errors in the query keypoints are natural for user-provided annotations and may arise from ambiguity in part endpoints (*e.g.* the wing base covered by feathers), resolution constraints, or human error. We simulate such noise by adding Gaussian perturbations of varying relative magnitudes to the pixel coordinates of the query keypoints. For noise levels below 1% of the image diagonal, the estimates remain stable for all three proportions. For context, in the query images, 1% of the image diagonal corresponds to 14.9%, 7.6%, and 7.7% of bill, tarsus, and wing-width, respectively. We observe sharp increases in measurement errors around 2–3% noise, most notably for the tarsus-to-wing estimate, which increases from about 10% to more than 40% relative error.

Lastly, we investigate the effect of image resolution (Fig. 7c). We exclude RANSAC filtering for a clearer comparison, and vary image resolution only in the keypoint matching stage. At very low resolutions, parts may be compressed into a single patch, leading to a sharp increase in error. This is particularly evident for small parts such as the bill. Higher resolutions tend to yield better results, with diminishing returns beyond 512.

4.4 Qualitative Case Studies

Fig. 8 explores the capabilities and limitations of WILDPROP through several qualitative case studies motivated by our interactions with ecologists and na-




Fig. 8: Case studies. WILDPROP exhibits broad applicability across taxa and parts. Each panel here presents the estimated distribution of particular body proportions across a few species. Exemplar images and parts are shown for each case study.

ture enthusiasts. Each case study focuses on a small set of species within the same taxonomic group and estimates body proportions that distinguish them. Additional case studies are presented in Appendix A6.

WILDPROP estimations distinguish species across diverse taxa. WILDPROP is broadly applicable to vastly different species and parts. We attempt measurement of upper and lower wing proportions for butterflies, base and petal proportions for flowers, and antler proportions for deer. In each example, we find our measurements separate three distinct species within the taxonomic group and align with our expectations.

WILDPROP requires only a single query image for similar species. We analyze the ratio of wingspan to body proportion of various birds of prey in Fig. 8. In this case study, we use only a single query image for all species and find that pose and keypoint matching is fairly accurate across species. Unlike most other families, birds of prey are frequently seen soaring overhead in harsh backlit conditions, requiring identification using only silhouette shape and flight style [57]. While identification can sometimes be possible using these proportions alone, we observe significant overlap in some species, particularly Peregrine Falcon and Bald Eagle. Other attributes which WILDPROP can not assess, such as absolute size and flight patterns, are likely required in these situations.




Fig. 9: Geographic and Temporal Stratification. WILDPROP allows measurement of specific sub-populations based on criteria such as geographic location or time-of-year.

WILDPROP can identify geographical and temporal variation. The search corpus in WILDPROP can be modified to obtain various stratified distributions, and Fig. 9 presents two such case studies. First, we find WILDPROP is able to identify geographic variations in Dunlin. Wintering Dunlin populations in western Europe, and western and eastern North America represent three distinct subspecies populations: *schinzii*, *pacifica*, and *hudsonia*, respectively. By stratifying our search corpora by location, we find the European population to have a median lower bill : head ratio than the North American populations, which are quite similar. This matches reported Dunlin subspecies studies, which find *pacifica* and *hudsonia* are the largest and longest-billed [69]. Although these studies indicate *hudsonia* is slightly shorter-billed on average than *pacifica* [69], we observe that WILDPROP is unable to distinguish these very similar subspecies.

Next, we examine seasonal change in antler length for Western Roe Deer. Rather than at the species level, we filter images based on month and estimate measurements for each month separately. We observe the expected trend in antler heights when stags grow their antlers in early spring and shed them in October–December [33]. Amusingly, WILDPROP selects ear tips as the closest match to the query antler tips when antlers are completely absent. Filtering out images that do not have antlers allows us to overcome this limitation and makes the trend even more apparent.

4.5 Discussions and Limitations

Rather than replacing physical measurement, we envision WILDPROP as a tool for rapid hypothesis generation and exploratory analysis before committing to more extensive measurement studies. While broad applicability across taxonomic groups and anatomical parts is a key strength of WILDPROP, it also comes with important limitations. First, estimates from WILDPROP are generally noisier than physical measurements, limiting its use in high-precision applications. Second, the image-based measurement paradigm has inherent constraints: scale ambiguity limits estimation to proportions rather than absolute metric lengths, and non-visual traits such as body mass or skeletal dimensions are not directly measurable from images alone. Third, because WILDPROP samples from citizen-

science image collections, its estimates may reflect biases toward human population centers, observer preferences, visually distinctive species, and popular taxa.

Evaluation also remains challenging. Existing physical morphometric datasets are often small, unevenly sampled, or collected under protocols that differ from the available image data, making direct comparisons imperfect. In addition, certain viewpoints and poses can be difficult to photograph for some species (*e.g.* bottom views or flight poses for small birds such as warblers and hummingbirds), limiting performance for proportions visible only under such conditions. Finally, beyond basic geometric heuristics, WILDPROP does not currently have a principled abstention mechanism. As a result, it may produce unrealistic keypoints or proportions when retrieval fails or when the specified body parts are absent.

5 Conclusion and Future Work

We present WILDPROP, a scalable approach for estimating wildlife body proportions from large, opportunistic image collections. Compared to physical measurement, WILDPROP provides a noisier but substantially more scalable alternative that can support ecological, biological, and evolutionary analyses when precise individual-level measurements are unavailable or impractical. Across three large-scale morphometric studies, our population-level body-proportion estimates are typically within 10–20% of physical measurements. Our estimates also follow expected trends in several qualitative case studies spanning diverse taxonomic groups and body parts. In addition, WILDPROP enables measurements under different stratifications, supporting analyses such as geographic variation in Dunlin bill proportions and seasonal trends in deer antlers. Together, these results suggest that retrieval-conditioned correspondence offers a practical pathway for estimating geometric population statistics from opportunistic imagery.

Several directions remain for future work. While WILDPROP uses off-the-shelf foundation models for broad applicability, dedicated animal pose or correspondence models could improve robustness, especially for taxa, viewpoints, or body parts with limited pose-aligned imagery. Incorporating calibrated depth estimation or explicit 3D keypoint detection could further extend the approach to a broader range of poses and viewpoints, moving toward more reliable individual-level proportion estimates across large populations. Incorporating a small number of physical measurements, when available, could also help calibrate image-based estimates and partially alleviate scale ambiguity. Future work could explore richer anatomical representations, such as part segmentations, curves, or point clouds, which provide additional geometric context beyond endpoints. Finally, extending the framework to additional morphological traits, including inferred body mass, coloration, and shape descriptors, could broaden the scope of image-based morphometric analysis.

Acknowledgments. The project was funded in part through grants #2329927 and #2504073 from the National Science Foundation, USA. We thank Feipeng Huang and Mark Titus for case study recommendations.

References

1. Bagchi, C., Medina, J., Irschick, D.J., Maji, S., Christiansen, F.: Automated extraction of right whale morphometric data from drone aerial photographs. *Remote Sensing in Ecology and Conservation* **11**(4), 438–453 (2025)
2. Bartomeus, I., Gravel, D., Tylianakis, J.M., Aizen, M.A., Dickie, I.A., Bernard-Verdier, M.: A common framework for identifying linkage rules across different types of interactions. *Functional ecology* **30**(12), 1894–1903 (2016)
3. Cao, J., Tang, H., Fang, H.S., Shen, X., Lu, C., Tai, Y.W.: Cross-domain adaptation for animal pose estimation. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9498–9507 (2019)
4. Chen, X., Chu, F.J., Gleize, P., Liang, K.J., Sax, A., Tang, H., Wang, W., Guo, M., Hardin, T., Li, X., et al.: Sam 3d: 3dfy anything in images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7220–7232 (2026)
5. Cole, E., Wilber, K., Van Horn, G., Yang, X., Fornoni, M., Perona, P., Belongie, S., Howard, A., Aodha, O.M.: On label granularity and object localization. In: *European Conference on Computer Vision*. pp. 604–620. Springer (2022)
6. Conklin, J.R., Verkuil, Y.I., Riegen, A.C., Battley, P.F.: How wry is a wrybill. *Wader Study* **126**(3), 228–235 (2019)
7. Cooke, R., Rendall, A.R., Weston, M.A., Porch, N., Bradsworth, N., White, J.G.: Photography can determine the sex of a predator with limited sexual dimorphism: A case study of the powerful owl. *Global Ecology and Conservation* **22**, e00959 (2020)
8. Cooney, C.R., He, Y., Varley, Z.K., Nouri, L.O., Moody, C.J., Jardine, M.D., Liker, A., Székely, T., Thomas, G.H.: Latitudinal gradients in avian colourfulness. *Nature Ecology & Evolution* **6**(5), 622–629 (2022)
9. Darwin, C.: *On the Origin of Species by Means of Natural Selection*. John Murray, London (1859)
10. Delfino, H.C., Carlos, C.J.: Photogrammetry as a tool for sex identification and sex ratio estimation in chilean flamingos. *Journal of Field Ornithology* **97**(1) (2026)
11. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: *Forty-first international conference on machine learning* (2024)
12. Fergus, P., Chalmers, C., Matthews, N., Nixon, S., Burger, A., Hartley, O., Sutherland, C., Lambin, X., Longmore, S., Wich, S.: Towards context-rich automated biodiversity assessments: deriving ai-powered insights from camera trap data. *Sensors* **24**(24), 8122 (2024)
13. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* **24**(6), 381–395 (1981)
14. FitzJohn, R.G.: Quantitative traits and diversification. *Systematic biology* **59**(6), 619–633 (2010)
15. Funk, J.L., Larson, J.E., Ames, G.M., Butterfield, B.J., Cavender-Bares, J., Firn, J., Laughlin, D.C., Sutton-Grier, A.E., Williams, L., Wright, J.: Revisiting the holy grail: using plant functional traits to understand ecological processes. *Biological Reviews* **92**(2), 1156–1173 (2017)
16. Gao, S., Yu, W., Tian, T., Lu, Z., Zhang, X., Li, Q., Chen, Y.: A morphological traits dataset of heteroptera sampled in biodiversity priority areas of southwest china. *Scientific Data* **11**(1), 694 (2024)

17. Geladi, I., De León, L.F., Torchin, M.E., Hendry, A.P., González, R., Sharpe, D.M.: 100-year time series reveal little morphological change following impoundment and predator invasion in two neotropical characids. *Evolutionary Applications* **12**(7), 1385–1401 (2019)
18. Gomez, M.L., Benham, P.M., Bowie, R.C.: Evidence for temporal morphological change in northern but not southern breeding crowned sparrows (*zonotrichia* spp.). *Evolutionary Ecology* **39**(5), 455–473 (2025)
19. Graving, J.M., Chae, D., Naik, H., Li, L., Koger, B., Costelloe, B.R., Couzin, I.D.: Deepposekit, a software toolkit for fast and robust animal pose estimation using deep learning. *elife* **8**, e47994 (2019)
20. Hantak, M.M., Guralnick, R.P., Zare, A., Stucky, B.J.: Computer vision for assessing species color pattern variation from web-based community science images. *Iscience* **25**(8) (2022)
21. Harfoot, M.B., Newbold, T., Tittensor, D.P., Emmott, S., Hutton, J., Lyutsarev, V., Smith, M.J., Scharlemann, J.P., Purves, D.W.: Emergent global patterns of ecosystem structure and function from a mechanistic general ecosystem model. *PLoS biology* **12**(4), e1001841 (2014)
22. He, Y., Cooney, C.R., Maddock, S., Thomas, G.H.: Using pose estimation to identify regions and points on natural history specimens. *PLOS Computational Biology* **19**(2), e1010933 (2023)
23. He, Y., Cooney, C.R., Maddock, S., Thomas, G.H.: Phenolearn: a user-friendly toolkit for image annotation and deep learning-based phenotyping for biological datasets. *Journal of Evolutionary Biology* **38**(8), 1152–1162 (2025)
24. Hodgson, J.C., Holman, D., Terauds, A., Koh, L.P., Goldsworthy, S.D.: Rapid condition monitoring of an endangered marine vertebrate using precise, non-invasive morphometrics. *Biological Conservation* **242**, 108402 (2020)
25. Honsey, A.E., Tingley III, R.W., Anweiler, K.V., Brant, C.O., Farha, S.A., Fedorowicz, P.W., Leonhardt, B.S., Eshenroder, R.L.: Demographic, morphometric, and meristic data describing cisco (*coregonus artedi*) captured in the spanish river, ontario, canada, 15–16 november 2022. U.S. Geological Survey data release (2024). <https://doi.org/10.5066/P1N2WFEJ>, <https://doi.org/10.5066/P1N2WFEJ>
26. Hu, X., Lyu, J., Liu, J., Liu, Y., Zuffi, S., An, L., Goetz, S.: Sam 3d animal: Promptable animal 3d reconstruction from images in the wild. arXiv preprint arXiv:2605.07604 (2026)
27. iNaturalist. <https://www.inaturalist.org> (2026), accessed on Jan 1, 2026
28. Johnson, N.A., Smith, C.H., Pfeiffer, J.P., Randklev, C.R., Williams, J.D., Austin, J.D.: Molecular and morphological data on two species complexes in the freshwater mussel genus *cyclonaias*. U.S. Geological Survey data release (2018). <https://doi.org/10.5066/P9SRSHV2>, <https://doi.org/10.5066/P9SRSHV2>
29. Jones, K.E., Bielby, J., Cardillo, M., Fritz, S.A., O'Dell, J., Orme, C.D.L., Safi, K., Sechrest, W., Boakes, E.H., Carbone, C., et al.: Pantheria: a species-level database of life history, ecology, and geography of extant and recently extinct mammals: Ecological archives e090-184. *Ecology* **90**(9), 2648–2648 (2009)
30. Kattge, J., Bönisch, G., Díaz, S., Lavorel, S., Prentice, I.C., Leadley, P., Tautenhahn, S., Werner, G.D., Aakala, T., Abedi, M., et al.: Try plant trait database—enhanced coverage and open access. *Global change biology* **26**(1), 119–188 (2020)
31. Ke, B., Obukhov, A., Huang, S., Metzger, N., Daudt, R.C., Schindler, K.: Repurposing diffusion-based image generators for monocular depth estimation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9492–9502 (2024)

32. Kershaw, J.L., Sherrill, M., Davison, N.J., Brownlow, A., Hall, A.J.: Evaluating morphometric and metabolic markers of body condition in a small cetacean, the harbor porpoise (*phocoena phocoena*). *Ecology and Evolution* **7**(10), 3494–3506 (2017)
33. Kierdorf, U., Gomez, S., Stock, S.R., Antipova, O., Kierdorf, H.: Bone resorption and formation in the pedicles of european roe deer (*capreolus capreolus*) in relation to the antler cycle—a morphological and microanalytical study. *Journal of Anatomy* **243**(5), 842–859 (2023)
34. Kristiansen, H.H., Metz, M., Silva-Garay, L., Jutfelt, F., Leeuwis, R.H.: Husmorph: a simple machine learning app for automated morphometric landmarking. *Conservation Physiology* **13**(1), coaf073 (2025)
35. Labelbox. Available: <https://labelbox.com> (2025), online
36. Larson, J.G., Weiner, A.: Linear morphological measurements of frogs. University of Michigan - Deep Blue Data (2022). <https://doi.org/10.7302/abtx-c461>, <https://doi.org/10.7302/abtx-c461>
37. Laughlin, D.C.: The intrinsic dimensionality of plant traits and its relevance to community assembly. *Journal of ecology* **102**(1), 186–193 (2014)
38. Lee, C.T., Birch, A.: Identification of willow and alder flycatchers by primary-tip spacing: The p6: 7 ratio. *Western Birds* **56**(1), 21–44 (2025)
39. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Li, C., Yang, J., Su, H., Zhu, J., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499* (2023)
40. Lu, C., Koniusz, P.: Few-shot keypoint detection with uncertainty learning for unseen species. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 19416–19426 (2022)
41. Ma, W., Qi, X., Sun, Y., Gao, R., Ding, L., Wang, R., Peng, C., Zhang, J., Wu, J., Xu, Z., et al.: Computer vision-based measurement techniques for livestock body dimension and weight: A review. *Agriculture* **14**(2), 306 (2024)
42. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., Bethge, M.: Deeplabcut: markerless pose estimation of user-defined body parts with deep learning. *Nature neuroscience* **21**(9), 1281–1289 (2018)
43. Mayfield, M.M., Levine, J.M.: Opposing effects of competitive exclusion on the phylogenetic structure of communities. *Ecology letters* **13**(9), 1085–1093 (2010)
44. Morfeld, K.A., Lehnhardt, J., Alligood, C., Bolling, J., Brown, J.L.: Development of a body condition scoring index for female african elephants validated by ultrasound measurements of subcutaneous fat. *PloS one* **9**(4), e93802 (2014)
45. Myhrvold, N.P., Baldridge, E., Chan, B., Sivam, D., Freeman, D.L., Ernest, S.M.: An amniote life-history database to perform comparative analyses with birds, mammals, and reptiles: Ecological archives e096-269. *Ecology* **96**(11), 3109–3109 (2015)
46. Nister, D., Stewenius, H.: Scalable recognition with a vocabulary tree. In: *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*. vol. 2, pp. 2161–2168. Ieee (2006)
47. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193* (2023)
48. Panda, A.K., Verma, V., Srivastav, A., Badola, R., Hussain, S.A.: Digital image processing: A new tool for morphological measurements of freshwater turtles under rehabilitation. *PloS One* **19**(3), e0300253 (2024)
49. Pazvant, G., İnce, N.G., Özkan, E., Gündemir, O., Avanus, K., Szara, T.: Sex determination based on morphometric measurements in yellow-legged gulls (*larus michahellis*) around istanbul. *BMC zoology* **7**(1), 35 (2022)

50. Peng, C., Cao, S., Li, S., Bai, T., Zhao, Z., Sun, W.: Automated measurement of cattle dimensions using improved keypoint detection combined with unilateral depth imaging. *Animals* **14**(17), 2453 (2024)
51. Phillips, J.G., Hagey, T.J., Hagemann, M., Gering, E.: Analysis of morphological change during a co-invading assemblage of lizards in the hawaiian islands. *Evolutionary Biology* **51**(2), 257–268 (2024)
52. Piccinelli, L., Sakaridis, C., Yang, Y.H., Segu, M., Li, S., Abbeloos, W., Van Gool, L.: Unidepthv2: Universal monocular metric depth estimation made simpler. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025)
53. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PmLR (2021)
54. Rao, J., Zhao, B.N., Wang, Y.: Probabilistic prompt distribution learning for animal pose estimation. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 29438–29447 (2025)
55. Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al.: Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714* (2024)
56. Ricklefs, R.E., Travis, J.: A morphological approach to the study of avian community organization. *The Auk* **97**(2), 321–338 (1980)
57. Roberts, P.M.: Books and other resources on hawk identification, hawkwatching, and migration. *Bird Observer* **29**(4), 311–317 (2001)
58. Russell, G., Christiansen, F., Colefax, A., Sprogis, K.R., Cagnazzi, D.: Comparisons of morphometrics and body condition between two breeding populations of australian humpback whales. *Wildlife Research* **51**(1), WR23026 (2023)
59. Ruthrauff, D.R., Tibbitts, T.L., Gill, Robert E., J., Handel, C.M.: Adult shorebird morphological measurement data (ver. 3.0, november 2024). U.S. Geological Survey data release (2023). <https://doi.org/10.5066/P9KNRWXB>, <https://doi.org/10.5066/P9KNRWXB>
60. Schleuning, M., Fründ, J., García, D.: Predicting ecosystem functions from biodiversity and mutualistic networks: an extension of trait-based concepts to plant–animal interactions. *Ecography* **38**(4), 380–392 (2015)
61. Silverman, B.W.: *Density Estimation for Statistics and Data Analysis*. CRC Press (1986)
62. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khali-dov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. *arXiv preprint arXiv:2508.10104* (2025)
63. Sivic, Zisserman: Video google: A text retrieval approach to object matching in videos. In: *Proceedings ninth IEEE international conference on computer vision*. pp. 1470–1477. IEEE (2003)
64. Stevens, S., Wu, J., Thompson, M.J., Campolongo, E.G., Song, C.H., Carlyn, D.E., Dong, L., Dahdul, W.M., Stewart, C., Berger-Wolf, T., et al.: Bioclip: A vision foundation model for the tree of life. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 19412–19424 (2024)
65. Tobias, J.A., Sheard, C., Pigot, A.L., Devenish, A.J., Yang, J., Sayol, F., Neate-Clegg, M.H., Alioravainen, N., Weeks, T.L., Barber, R.A., et al.: Avonet: morphological, ecological and geographical data for all birds. *Ecology letters* **25**(3), 581–597 (2022)

66. Tuia, D., Beery, S., Costelloe, B.R., Oliver, R.Y., Lecomte, N.: Towards ‘digital ecology’: Advances in integrating artificial intelligence from data generation to ecological insight (2026)
67. Van Horn, G., Mac Aodha, O., Song, Y., Cui, Y., Sun, C., Shepard, A., Adam, H., Perona, P., Belongie, S.: The inaturalist species classification and detection dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8769–8778 (2018)
68. Vendrow, E., Pantazis, O., Shepard, A., Brostow, G., Jones, K.E., Mac Aodha, O., Beery, S., Van Horn, G.: Inquire: A natural world text-to-image retrieval benchmark. *Advances in Neural Information Processing Systems* **37**, 126500–126514 (2024)
69. Warnock, N., Gill, R.E.: Dunlin systematics. *Birds of the World* (2023)
70. Weeks, B.C., Zhou, Z., O’Brien, B.K., Darling, R., Dean, M., Dias, T., Hassena, G., Zhang, M., Fouhey, D.F.: A deep neural network for high-throughput measurement of functional traits on museum skeletal specimens. *Methods in Ecology and Evolution* **14**(2), 347–359 (2023)
71. Williams, H.M., Wilcox, S.B., Patterson, A.J.: Photography as a tool for avian morphometric measurements. *Journal of Ornithology* **161**(1), 333–339 (2020)
72. Wilman, H., Belmaker, J., Simpson, J., De La Rosa, C., Rivadeneira, M.M., Jetz, W.: Eltontraits 1.0: Species-level foraging attributes of the world’s birds and mammals: Ecological archives e095-178. *Ecology* **95**(7), 2027–2027 (2014)
73. Xiang, J., Chen, X., Xu, S., Wang, R., Lv, Z., Deng, Y., Zhu, H., Dong, Y., Zhao, H., Yuan, N.J., et al.: Native and compact structured latents for 3d generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14419–14429 (2026)
74. Xiang, J., Lv, Z., Xu, S., Deng, Y., Wang, R., Zhang, B., Chen, D., Tong, X., Yang, J.: Structured 3d latents for scalable and versatile 3d generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 21469–21480 (2025)
75. Xiong, T., Tan, D., Tian, W.: Diffpose-animal: A language-conditioned diffusion framework for animal pose estimation. *arXiv preprint arXiv:2508.08783* (2025)
76. Yang, L., Kang, B., Huang, Z., Xu, X., Feng, J., Zhao, H.: Depth anything: Unleashing the power of large-scale unlabeled data. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10371–10381 (2024)
77. Youngflesh, C., Saracco, J.F., Siegel, R.B., Tingley, M.W.: Abiotic conditions shape spatial and temporal morphological variation in north american birds. *Nature Ecology & Evolution* **6**(12), 1860–1870 (2022)
78. Zhai, X., Mustafa, B., Kolesnikov, A., Beyer, L.: Sigmoid loss for language image pre-training. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 11975–11986 (2023)