

Bugs in the Data: How ImageNet Misrepresents Biodiversity

Alexandra Sasha Luccioni*¹, David Rolnick*²

¹Hugging Face

²McGill University, Mila

sasha.luccioni@huggingface.co, drolnick@cs.mcgill.ca

Abstract

ImageNet-1k is a dataset often used for benchmarking machine learning (ML) models and evaluating tasks such as image recognition and object detection. Wild animals make up 27% of ImageNet-1k but, unlike classes representing people and objects, these data have not been closely scrutinized. In the current paper, we analyze the 13,450 images from 269 classes that represent wild animals in the ImageNet-1k validation set, with the participation of expert ecologists. We find that many of the classes are ill-defined or overlapping, and that 12% of the images are incorrectly labeled, with some classes having >90% of images incorrect. We also find that both the wildlife-related labels and images included in ImageNet-1k present significant geographical and cultural biases, as well as ambiguities such as artificial animals, multiple species in the same image, or the presence of humans. Our findings highlight serious issues with the extensive use of this dataset for evaluating ML systems, the use of such algorithms in wildlife-related tasks, and more broadly the ways in which ML datasets are commonly created and curated.

Introduction

Datasets are a crucial part of the development, evaluation and eventual deployment of machine learning (ML) systems. Large datasets such as ImageNet (Deng et al. 2009) and MSCOCO (Lin et al. 2014) have been used for years both as inputs for training ML models for downstream tasks and as benchmarks for evaluating model efficacy. There are many choices involved implicitly or explicitly in the creation of a labeled dataset, including how classes are defined as well as the way datapoints are assigned to them. Recent research has placed increased scrutiny on many popular datasets, uncovering biases (Prabhu and Birhane 2021), duplicates (Barz and Denzler 2020) and problematic class definitions (Crawford and Paglen 2019). Such dataset shortcomings can have serious impacts on the behavior of the models they are used to train, and can also make it difficult to evaluate the true effectiveness of ML systems, since metrics such as accuracy are impacted by biases or inaccuracies in the evaluation set.

The ImageNet Large Scale Visual Recognition Challenge (ILSVRC) dataset, commonly called *ImageNet-1k*, (Rusakovsky et al. 2015) is a subset of the full ImageNet

dataset (Deng et al. 2009) that is used widely in categorical object recognition (with 32k citations to date). More than a quarter of the dataset consists of images of wild animals. However, no wildlife experts were consulted in creating the dataset, and the choice of classes and images has never been analyzed, despite the considerable role that these data have ultimately played in shaping widely used ML models.

Here, we examine the 269 classes of wild animals in ImageNet-1k, working with expert ecologists to evaluate both the class definitions and the images included in each class. Our main findings are the following:

- Over 12% of the wildlife images are incorrectly labeled, with some classes having over 90% incorrect labels.
- Many images are problematic in other ways, including the presence of humans (6.7% of images), multiple types of animals in the same image (4.5%), images too blurry for identification (2.2%), and images of fake animals such as stuffed toys (1.4%).
- Both the class definitions and images are strongly biased towards the U.S. and Europe, vastly under-representing biodiversity from other geographies.
- Many classes in the dataset are not clearly defined (11.5% of classes) or overlap with other classes (11.9%).

By leveraging the concrete nature of wildlife images, where individual species (including their geographic provenance) can be readily identified, we provide a novel window into the many weaknesses of ImageNet-1k as a benchmark dataset. Our findings also raise significant concerns about various real-world ML applications where ImageNet-1k is used, as well as highlighting the fundamental importance of domain expertise in crafting ML benchmark datasets.

Related Work

Our work is relevant to several research directions, ranging from empirical studies on the impact of label errors on model performance to analyses of the values embedded in ML datasets. We briefly summarize work in these fields in the paragraphs below.

A brief history of ImageNet. Since the creation in 2009 of the original ImageNet dataset (Deng et al. 2009), containing 21,000 classes, it has gone through numerous iterations, including the 1k version created for the ImageNet

*These authors contributed equally.

Large Scale Visual Recognition Challenge in 2010 (Rusakovsky et al. 2015), commonly called *ImageNet-1k*. In recent years, following an in-depth study by the Excavating AI project (Crawford and Paglen 2019), an updated official version of the 21K dataset was created, aiming to filter out problematic person categories and to balance their distribution (Yang et al. 2020). In 2021, a subset of the authors of the original ImageNet paper proposed a new version of the full dataset, using facial recognition techniques to blur the faces of people in the images, showing that it does not impact model performance (Yang et al. 2022).

Analyses of ImageNet. There have been numerous studies of ImageNet from different angles, ranging from the quality of its labels (Beyer et al. 2020; Tsipras et al. 2020; Northcutt, Athalye, and Mueller 2021) to the biases that it reflects (Prabhu and Birhane 2021; Crawford and Paglen 2019; Shankar et al. 2017; De Vries et al. 2019). There have also been attempts to replicate its data collection process (Recht et al. 2019) and to propose a new validation set for ImageNet, meant to better present “real-life” images (Barbu et al. 2019). To date, these studies have focused on the the “people” and “object” classes of ImageNet, with the notable exception of a 2015 study that looked at the bird classes in ImageNet, hypothesizing class error rates of at least 4% (Van Horn et al. 2015). We build on this work, examining all of the wildlife classes of ImageNet-1k and not only calculating the class error rates (which we find are much higher for other animals than for birds), but also analyzing the biases and values implicit in the class definitions and choice of images.

Work on wildlife image datasets. ImageNet is far from the only large dataset containing images of wildlife; in fact, many researchers have created diverse datasets specifically focused on wildlife images in the last decade, often working hand-in-glove with biodiversity experts and citizen scientists. For instance, the eBird dataset was created at the same time as ImageNet and contains nearly 3 million images of birds, collected and labelled by citizen scientists (Sullivan et al. 2009), which have been extensively utilized and studied (Van Horn and Perona 2017). Van Horn et al. (2015) proposed *NABirds*, a dataset containing almost 50,000 images of North American birds spanning 555 classes. The iWild-Cam Competition is another notable example of extensive biodiversity datasets, leveraging camera traps to gather hundreds of thousands of images across the world (Beery et al. 2019, 2021); these datasets have been used to host competitions, establish benchmarks, and assess model generalization to novel ecosystems and species (Beery, Van Horn, and Perona 2018), but remain underutilized by the ML community at large. Finally, one of the most diverse and extensive biodiversity datasets is iNaturalist, which spans 859,000 images from over 5,000 species of plants and animals (Van Horn et al. 2018) and has been used to develop highly popular ML-powered species identification mobile applications such as Seek (iNaturalist 2022).

Values and biases embedded in datasets. Recent years have also seen several deep dives into the provenance and

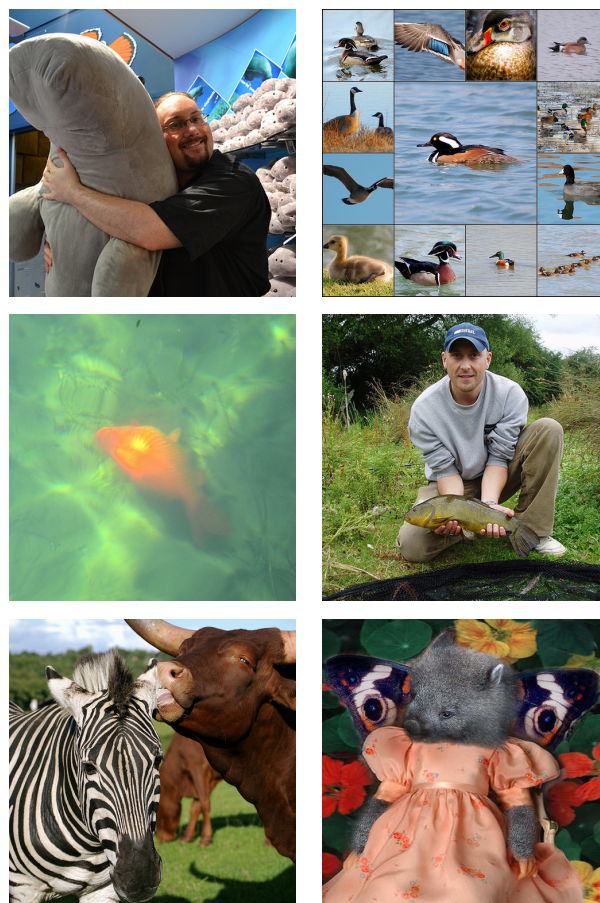


Figure 1: Examples of images of animals from ImageNet-1k, including artificial animals (top left), collages (top right) (here, the correct species “american coot” occurs only in one of the side panels of the image), blurry images (center left), images with humans (center right), images with multiple species of animals (bottom left) and images that are bizarre in other ways (bottom right: a wombat in a dress with butterfly wings and human hands).

values latent in ML datasets. For instance, Crawford and Paglen’s “archaeology” of several ML datasets including ImageNet has exposed the racist and misogynist assumptions and values that are buried in them (2019). This has been built upon by further scholarship around the ethical and political dimensions of dataset curation and construction (Prabhu and Birhane 2021; Denton et al. 2020) and the values that are embedded in datasets (Scheuerman, Hanna, and Denton 2021). Complementary work has been carried out to study how datasets are created and used by the community, finding that there is a concentration of power in terms of institutions that create, which results in influence of ML research and practice (Koch et al. 2021). This work has inspired our own, since we find that these values and practices have also impacted the representation of nature and biodiversity in ML datasets, and more specifically ImageNet-1k, which is the focus of our study.

Taxonomic group	# classes	# images
Amphibian	8	400
Bird	57	2850
Fish	16	800
Mammal	93	4650
Marine Invertebrate	20	1000
Reptile	35	1750
Terrestrial Invertebrate	40	2000
Total	269	13450

Table 1: Summary of wildlife images in ImageNet-1k, broken down according to taxonomic group (here and throughout the paper, “reptile” refers to non-avian reptiles). There are 50 images per class in the validation set.

Methodology

In this section, we detail the approach that we used for systematically evaluating the images of wildlife in ImageNet.

Selecting Images of Wildlife

There are 398 classes representing animals in ImageNet-1k, which is over a third of the 1000 classes in the overall dataset. Of these, 269 represent wild extant animals, whereas 127 are breeds of domestic animals such as dogs, cats and chickens, and 2 represent prehistoric extinct animals (the trilobite and triceratops). In the present study, we focus on wild animals, given that domestic breeds’ presence depends wholly on cohabitation with humans, and many of the same domestic breeds are now present across the globe, notwithstanding whether they were originally from a specific region. We also exclude long extinct species, since it is clearly impossible to have photos of these creatures in the wild (which entails that all of the images of these species are artists’ representations or photographs of fossils). We separate these 269 classes of wild animals into 7 taxonomic groups, as shown in Table 1.

Domain Expert Validation

In order to validate both the class labels used in ImageNet-1k as well as the images from the validation set tagged with these labels, we recruited 20 expert annotators with expertise across different taxonomic groups. We use images from the validation set as this is designed to represent a representative subset of the full dataset and is of a tractable size (~13k images) to permit careful annotation by experts. Each of the expert annotators had at minimum graduate-level education in the domain in question – for instance, the images of primates were annotated by two postdoctoral researchers in primate biology and one zoo keeper specializing in primates. Further details on our annotation process are given in the Supplementary Material.

Class Validation We asked our expert annotators the following questions for each class in their domain of expertise:

- **Clarity:** Does the category clearly define a species or group of species? [Yes/No]
- **Overlap:** Does the category overlap with another category in the list? If so, which category does it overlap with?

These questions aimed to validate both the scientific relevance and clarity of the classes in ImageNet-1k, and to estimate the degree of overlap between the classes, given that the 1000 synsets that were selected for ImageNet-1k are not meant to have any overlap between them. Furthermore, most of the classes have multiple labels spanning both formal and colloquial terms (e.g. “platypus, duckbill, duckbilled platypus, duck-billed platypus, *Ornithorhynchus anatinus*”), inherited from the original WordNet hierarchy (Miller 1995), and it is important to validate the correspondence of these different terms from a scientific perspective.

Image Validation We also asked our expert annotators questions for each image in their domain of expertise:

- **Correct ID:** To the extent of your knowledge, does the label correspond to the animal that is shown on the image? [Yes/No/Maybe]
- **Humans present:** Are there humans present in the image? [Yes/No]
- **Other animals present:** Are there other (non-human) animals present in the image? [Yes/No]
- **Not real animal:** Is the animal in the image an illustration, sculpture, or other representation of the animal? [Yes/No]
- **Blurry image:** Is the image too blurry or low quality to allow identification? [Yes/No]
- **Image collage:** Is the image composed of several images (either of the same animal or different ones)? [Yes/No]

See Figure 1 for examples of some of the failure modes indicated in the questions above.

Species-level Re-annotation

In order to better examine the biases and inaccuracies in the dataset, we also asked experts in bird classification to re-annotate all of the 2950 images from the 57 classes of birds, identifying each of the images to species level where possible. This experiment was designed to provide information about (i) the true species if the image was incorrectly classified, (ii) the individual species if the class label was vague (e.g. “kite” can refer to many different species of birds found in different parts of the world). We focused on the birds because birds can often be identified to the level of species based on images alone (as opposed to insects, for example).

Findings

We identify three categories of problems with wildlife-related data in ImageNet: (1) inconsistencies and biases in the choice and definition of classes, (2) inaccuracies in the images chosen to represent those classes, and (3) images that are not strictly inaccurate but exhibit other kinds of biases.

Inconsistencies and Biases in Class Definitions

Before even considering the images within the ImageNet-1k dataset, it is worth analyzing the set of classes. We find that many classes are unclearly defined or even overlap with other classes, resulting in ambiguity about the correct label for images. Beyond this, the classes chosen reflect biases related to geography as well as to cultural knowledge.

Taxonomic group	Class Overlap	Unclear definition
Amphibian	0.0%	25.0%
Bird	8.8%	19.3%
Fish	12.5%	6.3%
Mammal	15.1%	2.2%
Marine Invertebrate	20.0%	20.0%
Reptile	14.3%	20.0%
Terrestrial Invertebrate	5.0%	10.0%
Overall	11.9%	11.5%

Table 2: Percentages of classes, respectively that overlap with other classes and have a clear definition.

Unclear classes. As shown in Table 2, our expert annotators found 11.5% of the wildlife classes in ImageNet-1k to be unclear, where they were not certain what types of wildlife were included. In some cases, this results from seemingly contradictory descriptors, such as “hog-nose snake, puff adder, sand viper,” since “sand viper” and “hog-nose snake” refer to different species. In other cases, lack of clarity results from terms (e.g. “cricket”) with no single accepted definition. Some terms, such as the bird class “kite,” can be defined in different ways based on dialect, introducing geography-mediated ambiguity based on how the class name is interpreted.

Overlapping classes. Even more concerning, our annotators found that 11.9% of the wildlife classes used in ImageNet are actually overlapping (see Table 2), resulting in ambiguity about whether images belong in one or another. In some cases, this arises from vague class definitions; for example, the ImageNet class “hog-nose snake, puff adder, sand viper” is poorly defined and could overlap with the class “horned viper, cerastes, sand viper, horned asp, Cerastes cornutus.” In other cases, a generic class is used alongside a more specific one; for example, the class “meerkat” refers to a type of mongoose, even though “mongoose” is another ImageNet-1k class. Such behavior directly contradicts the intentions stated in (Russakovsky et al. 2015), which indicates that “the 1000 synsets are selected such that there is no overlap between synsets: for any synsets i and j , i is not an ancestor of j in the ImageNet hierarchy.” Also, several instances of confusion arise from one of the classes being essentially human-defined, rather than a taxonomic category, e.g. the class “tuskier” (referring to elephants with long tusks), which conflicts with the class “African elephant,” since most tuskers are African elephants.

Geographic bias. There are significant geographic biases in the classes represented in the dataset. For example, out of the sixty species of eagle that exist globally (many of which are culturally important birds in their regions) only the bald eagle (*Haliaeetus leucocephalus*) is featured in the list of ImageNet-1k classes. Not coincidentally, it is the national bird of the United States – indeed, out of the 25 species-specific classes of birds in ImageNet, 60% are present in the U.S. and 48% are present in Europe, even though only 8.4% and 5.0% percent of bird species in the world are present in

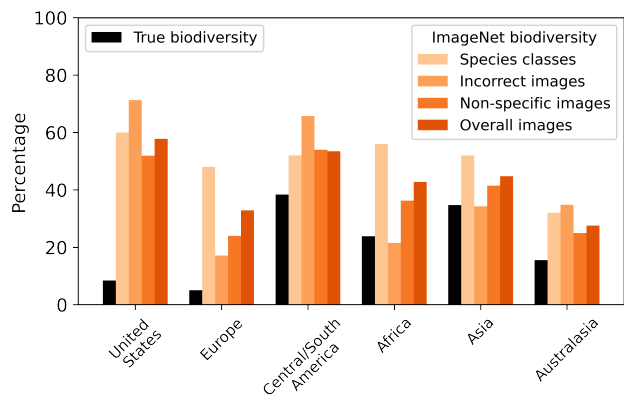


Figure 2: Geographic biases illustrated via comparison of the true biodiversity of birds and the biodiversity of birds within ImageNet-1k. “True biodiversity” shows the percentages of global bird species present in various geographic regions (note that these percentages sum to greater than 100% since a single species may be present in multiple geographic regions). “Species classes” analyzes the geographic distribution of those classes of ImageNet-1k referring to individual species of birds, showing that the U.S. and Europe are vastly overrepresented. “Incorrect images” analyzes where incorrectly identified birds are actually from, exhibiting a strong bias towards the U.S. “Non-specific images” analyzes which bird species are chosen to represent classes referring to multiple species (e.g. 62% of “jay” images depict the blue jay, a U.S. species). “Overall images” analyzes the overall distribution of bird species across all ImageNet-1k images.

the U.S. or Europe, respectively (see Figure 2).¹ By contrast, bird diversity in Central and South America, Africa, Asia, and Australasia is severely underrepresented.

Cultural bias. Linked to geographical bias in the choice of classes is cultural bias. By this, we refer to the selection of animals that are well-known or important culturally, rather than either (i) reflecting the balance of biodiversity, or (ii) reflecting the animals most widely seen in practice. Thus, for example, there are 93 classes referring to mammals in the dataset, but only 27 classes for insects, even though there are about six thousand known species of mammals in the world, compared to a million known species of insects (Mora et al. 2011) – this reflects the relative importance of mammals culturally as compared to insects. Furthermore, within the ImageNet-1k classes for insects, only one refers to an individual species (the monarch butterfly, *Danaus plexippus*, probably the most widely known butterfly in the U.S.), while all the other classes are more high-level and do not refer to a single species (“ant,” “dragonfly,” etc.). Implicit in the notion of cultural bias is the set of cultures within which cultural relevance is defined; as noted above, the set of relevant cultures is Americentric and Eurocentric.

¹For the purposes of this analysis, we used bird lists from Avibase (Lepage 2022), neglecting vagrants and accidentals.




Figure 3: Overall, we find that 12.3% of wildlife images in ImageNet-1k are incorrect. Left: the percentages of incorrectly labeled images across different taxonomic groups. Right: for each taxonomic group, what percentage of the classes within that group have error rates 0-5%, 5-10%, 10-25%, and >25%, respectively (for example, 65% of ImageNet-1k classes referring to birds have <5% error rates, while 40% of classes for reptiles have >25% error rates).

Inaccurate Examples

We find that 12.3 percent of the wildlife images in ImageNet-1k are incorrectly labeled and do not fall into the class in question. As shown in Figure 3, error rates vary markedly between different taxonomic groups, with birds having the lowest error rate (7%) – consistent with the estimate of Van Horn et al. (2015) – and reptiles the highest error rate (27%). We further note that some classes have extremely high error rates; 16 classes have at least half of the images incorrect and 5 classes have more than 90% of images incorrect (see Table 3). As described below, we find that inaccurate labels reflect the inexperience of those originally annotating the dataset, as well as effects related to prevalence, geography, and nomenclature.

Annotator inexperience. Likely the biggest factor in inaccuracies is annotator inexperience. Identification of animals is a task that in most cases requires considerable experience and expert knowledge, while the original ImageNet annotators were members of the general public without specialized expertise. Moreover, the images given to the annotators were selected through Google Image Search, which relies on textual descriptions given by the original photographers – themselves subject to inaccuracies. There is thus the potential for two layers of annotator inexperience: first by the photographer and then by the ImageNet labeler. For instance, within the class “crane” (the bird), for example, 22% percent of the images are incorrect and instead depict herons and egrets, a similar group of birds that are often confused with cranes (distressingly, many of these incorrect images show the egret *Ardea alba*, which has its own class in ImageNet-1k). There appear to be a significant number of such inaccuracies for insects and other arthropods, reflecting the fact that these animals are often less familiar to the general public. In the class “centipede,” for example, 60% percent of the images are incorrect, with most of the incorrect examples being images of millipedes or furry caterpillars (for comparison, centipedes are about as closely related

Class Label	Incorrect
black-footed ferret, <i>Mustela nigripes</i>	98%
rock crab, <i>Cancer irroratus</i>	96%
tailed frog, bell toad, <i>Ascaphus trui</i>	96%
kit fox, <i>Vulpes macrotis</i>	92%
goldfinch, <i>Carduelis carduelis</i>	90%
green lizard, <i>Lacerta viridis</i>	86%
night snake, <i>Hypsiglena torquata</i>	82%
green snake, grass snake	70%
mud turtle	62%
horned viper, cerastes, <i>Cerastes cornutus</i>	60%

Table 3: The top 10 classes with the highest error rates, and the percentages of incorrect images for each.

to millipedes and caterpillars as humans are to fish).

Prevalence effects. Certain inaccuracies appear to be caused by the greater prevalence of a similar animal as compared to the class in question. For example, the ImageNet-1k class “black-footed ferret, *Mustela nigripes*” refers to an endangered species that is uncommonly seen due to its rarity and nocturnal habits. Out of the 50 images supposedly in this class, only 2 images are correct. Almost all of the remaining 48 are photographs taken inside houses of domestic ferrets, a different species that is, not surprisingly, much more commonly seen. In the Google Image Search used to create ImageNet, it appears likely that the top search results were images of domestic ferrets that were then endorsed by the inexperienced annotators.

Geographic effects. Some of the inaccuracies observed reflect geographic biases, with a species in one location having a similar name to that in another location. For example, one class refers to the goldfinch *Carduelis carduelis* (a species found in Europe, Africa, and Asia) and 90% of the images are incorrect, with most of the incorrect examples showing the American goldfinch (*Spinus tristis*). It is worth

noting that out of the incorrect images of birds within the dataset, our experts found that 71% depicted birds found within the United States (see Figure 2).

Nomenclature effects. Other inaccuracies reflect the ways in which a technical name may be misinterpreted. The class name “leaf beetle, chrysomelid” for example, refers to beetles in the family Chrysomelidae. The images given for this class in ImageNet are 52% percent incorrect, with most of the inaccurate images being other kinds of beetles sitting on leaves. In this case, many of the inaccurate images are of ladybugs, which is especially problematic since “ladybug” is itself a class in ImageNet. A similar situation occurs for the class “kit fox, *Vulpes macrotis*,” which refers to a specific species of fox, and yet 92% of the images from this class are, in fact, of young foxes from other species of fox, since “kit” is also the term for a young fox of any species.

Note on taxonomic specificity. Out of the 269 classes referring to wildlife, our annotators determined that 130 refer to individual species (e.g. “coho salmon”), while 139 refer to multiple species (e.g. “scorpion”). It may be wondered whether errors are more likely to occur for the species-specific classes, as such categories are by some measure narrower. We find that indeed this is the case, but that a considerable number of errors occur in both species-specific and non-species-specific classes – with these classes having overall error rates of 15.5% and 9.2%, respectively.

Biases in Examples

Even in cases where images are correctly classified, we nonetheless observe problematic trends in the choice of images, reflecting geographic and contextual biases. We also note numerous instances of ambiguous images.

Geographic bias. According to our annotators, 51.7% percent of the wildlife classes in ImageNet encompass more than one species. “Jay” for example can refer to any of a number of birds in the family Corvidae, such as the green jay (*Cyanocorax luxuosus*), Eurasian jay (*Garrulus glandarius*), and Sichuan jay (*Perisoreus internigrans*), while “crane” refers to any bird in the family Gruidae. In such cases, we observe significant biases in the exact identity of the species chosen within each class. For example, all of the examples in ImageNet-1k for the class “jay” are correct in that they depict jays. However, 62% of these images show the blue jay (*Cyanocitta cristata*), which is common in the U.S. but is only one of 49 species of jay worldwide. Once again, such biases likely result from the Google Image Search methodology originally used to create ImageNet; it is well-known that certain geographies are better represented in web search results, especially for English-language searches.

Overall, our annotators analyzed all 32 non-specific classes of birds in the ImageNet dataset and found an average of 52% of the images were of species present in the United States (see Figure 2). In some cases where large numbers of species outside the U.S. were included (for example in the “crane” and “macaw” classes), many images appear to be of birds in zoos or as pets, and therefore could still represent photos taken in the U.S.

Group	Blur	Collage	Fake	Human	Multi
Amphibian	0.3%	0.3%	0.8%	6.0%	0.8%
Bird	0.4 %	0.4%	0.4 %	1.0%	3.2%
Fish	10.8%	0.1%	1.3%	24.6%	13.3%
Mammal	2.7%	1.3%	2.0%	6.3%	2.4%
Marine Invert.	4.4%	0.2%	3.7%	14.3%	22.6%
Reptile	0.9%	0.2%	0.9%	9.5%	2.1%
Terr. Invert.	0.5%	0.3%	1.1%	2.4%	1.4 %
Overall	2.2%	0.6%	1.4%	6.7%	4.5 %

Table 4: Percentages of images with various confusing attributes: respectively, severe blur, a collage of multiple images, depiction of a fake animal such as a painting or stuffed toy, the presence of humans, and the presence of multiple species of animals in the same image.

Contextual bias. The context in which species are depicted within the dataset is also subject to significant bias, e.g. many of the photographs of jellyfish seemingly being taken inside aquariums. Contextual bias is particularly blatant for certain classes referring to fish, in which fishing scenarios (e.g. the fish being held out of the water by a human) vastly outweigh images of fish in their natural environment. Across the classes for tench, barracouta, coho salmon, sturgeon, and gar, 61.6% percent of the images include humans. This bias is also present in mammals such as the hartebeest that are often hunted for game, resulting in images of dead animals accompanied by the humans who have shot them. Some contextual bias is of course inevitable in image-taking (e.g., butterflies may be more likely to be photographed while near the ground, rather than at the top of a tree), but such a significant and avoidable bias is noteworthy.

Bad or ambiguous images. In addition to the biases noted above, many of the images were confusing for other reasons. As described in Table 4, our annotators observed 6.7% of images had humans present (including but not limited to the fishing images described above), 4.5% of images had multiple (non-human) species present in the image, 1.4% of images showed an artificial representation of the animal in question (such as an illustration, toy or sculpture), and 0.6% of images were actually multiple images merged into a collage. The annotators also noted that 2.2% of images were too blurry or low quality to allow for reliable identification. See Figure 1 for examples of these failure modes.

Discussion

Our findings have many implications for the machine learning community. First, they call into question the ways in which performance on ImageNet-1k is used as a measure of success. It continues to be common practice for ML models to be benchmarked on ImageNet-1k, and reviewers frequently request this. Implicit in such tests are two assumptions that we show to be false:

- Validation accuracy is supposed to reflect how well the model is learning. In fact, we find that, within wildlife images, the validation set is 12.3% percent incorrectly labeled. Many of these examples should be correctly labeled as another class in ImageNet (e.g. images labeled



Figure 4: Four of the 49 species of jays found around the world: the blue jay, green jay, Eurasian jay, and Sichuan jay. Despite the visual and geographic diversity of jays, 62% of the “jay” images in ImageNet are of the blue jay (a species found only in the U.S. and Canada). This example illustrates geographic inequities as well as the fact that ImageNet classes are implicitly made up of imbalanced subclasses requiring few-shot (or zero-shot) generalization. Image credit: see Macaulay Library 84878861, 79048551, 241250451, 41284331.

as “crane” that are actually “egret”), while others are ambiguous (is it an “African elephant” or a “tusk?”), and still others are out-of-distribution examples without any correct ImageNet label. Notably, none of these inaccuracies can be dismissed as simply random noise – instead, they admit complex correlations across classes, yielding hard-to-predict effects on overall accuracy.

- ImageNet is supposed to have balanced classes. In fact, we find that many classes are implicitly made up of subclasses that are unevenly represented. For example, a model trained on the “jay” class, where 62% of the images show a single species of jay, may generalize poorly to the other 48 species of jays in the world, many of which look radically different (see Figure 4). It is likely that in many cases, such imbalances result in validation and test examples that are highly dissimilar from all or most examples in the training set, essentially requiring few- or zero-shot generalization.

Failure of these basic assumptions further calls into question the utility of ImageNet as a benchmark and the validity of state-of-the-art claims, especially given the minute margins by which state-of-the-art is established on ImageNet.

Moreover, while there are many wildlife-focused datasets such as iNaturalist, iWildCam, and NABirds, the ImageNet dataset is still used in various contexts where its performance on animal identification matters. It is often used to pretrain networks for fine-tuning on wildlife recognition datasets, especially by researchers who do not have access to the computational resources required to train models from scratch. Inaccuracies and biases in the original ImageNet dataset can therefore directly affect the performance and bi-

ases of downstream models. Furthermore, ImageNet is also used in generic contexts such as classifying photos and image search. This can lead to the propagation and magnification of biases if classification algorithms influence the animals that people learn or the results that appear in search engines. For example, if most of the ImageNet images of “jay” are dominated by a small subset of the possible species, then search results may miss relevant information or be implicitly biased towards content from certain geographies.

Finally, our work calls attention to broader issues in the machine learning ecosystem. The biases and inaccuracies we identify within wildlife imagery are, of course, not unique to wildlife data. Rather, the objective nature of taxonomic classification can make it easier to observe and quantify issues that likely arise extensively within other types of data (see e.g. Crawford and Paglen (2019); Prabhu and Birhane (2021)). Nor have ImageNet’s failures prevented similar weaknesses in more recently developed datasets. Machine learning continues to draw from large Internet corpora with significant biases, using annotators who often lack relevant domain expertise. Tasks and classes within a dataset may likewise be chosen by individuals without relevant expertise, and who may be biased in their geographies or lived experiences.

Conclusion

In this work, we analyze how biodiversity is (mis)represented in the ImageNet-1k dataset, working with domain experts in order to evaluate both the soundness of class definitions and the quality of the images that are included in each class. Our results show extensive biases, both geographic and cultural, in the choice of wildlife represented in the dataset, as well as high incidence of unclear class definitions and overlap between classes. We also find a high percentage of incorrect labels, with some classes almost completely incorrect. Even correctly classified images are frequently confusing – including the presence of artificial animals, humans, or other animal species – and the images chosen to represent classes are strongly biased towards the United States and Europe.

We wish that our work offered a quick fix for the many failings we have identified in ImageNet-1k, but unfortunately solutions ultimately require larger-scale, indeed systemic, change. It is not sufficient simply to relabel all inaccurate images we have found – for many of them, the true species falls outside any ImageNet class, and such examples would need to be replaced by new images. In some cases, as we have seen, the majority of the class would need to be replaced. More insidious still are the biases present in the images and the biases and inconsistencies in the class definitions, which would essentially require replacing the structure of the whole dataset. We hope to see future work quantifying how the perceived efficacy of different ML models has been skewed by errors and biases in ImageNet-1k. We also look forward to improved standards in ML for the participation of diverse domain experts in the creation and annotation of datasets, both with relevance to biodiversity and more broadly.

Appendix: Annotation Details

We divided ImageNet-1k wildlife classes into sets according to standard domains of expertise for ecologists and taxonomists: birds, fish, herps (including amphibians and reptiles), marine invertebrates, primates, other mammals, terrestrial invertebrates. We further subdivided birds into sets A and B since there were a large number of bird classes. For each set of classes, we recruited expert annotators with at least graduate-level experience in the relevant taxonomic group. Labels were completed by multiple annotators with expertise in the relevant taxonomic group (see Table 5), spanning a variety of geographic regions.

Each annotator answered questions on all classes within their set and all images belonging to those classes, with the various annotators completing their work independently without access to each others' answers. To obtain final results, we obtained a consensus for each question, based on whether a majority of annotators answered Yes or No (ties were broken in favor of Yes answers, e.g. if there was doubt as to whether an image was correctly identified or not, the tie was broken in favor of the image being correct). Consensus between annotators was high, with a 0.72 average Krippendorff alpha value.

It is worth noting that annotators, as experts in wildlife identification, were in general highly invested in the success of the project and extremely thorough in their work. Many wrote length messages to us explaining their logic – for example, the following message about the “cricket” class:

“Crickets” *sensu lato* encompasses suborder Ensifera, which is essentially any orthopteran that isn't a grasshopper.

Ensifera is comprised of two infraorders: Gryllidea (which includes the so-called “true crickets” (family Gryllidae)) and Tettigoniidea.

The vast majority of families in Tettigoniidea are referred to as crickets, however one exception is family Tettigoniidae – katydids. That said, Tettigoniids can be referred to as “bush-cricket” (e.g., in the UK) rather than “katydids”.

In my own nomenclatorial/classification world view, I don't regard katydids as crickets *per se*. When I think “cricket” I think family Gryllidae, and this was my initial approach to the cricket images.

When it comes to AI image classification, however, calling a katydid a cricket isn't necessarily incorrect. It's going to depend upon where the classification boundary is set (e.g., suborder Ensifera (non-grasshopper orthopterans) vs. family Gryllidae (“true crickets”)) or how the classification is defined (e.g., non-grasshopper orthopterans referred to as crickets within a vernacular).

In addition to the class-specific and image-specific questions, we obtained species-level identifications of all images in the bird classes from one annotator in each of Birds A and Birds B. These annotators were able to provide species-level identifications for 86.8% of the images (the other images either did not show relevant diagnostic features or the birds

ImageNet-1k Classes	# Annotators
Birds A	3
Birds B	1
Fish	2
Herps	3
Marine Inverts.	3
Primates	3
Other Mammals	2
Terrestrial Inverts.	3

Table 5: Images and classes were subdivided according to taxonomic groups. The number of expert annotators assigned to each group is shown.

in question were not identifiable to species level from images alone). The bird annotators were extremely diligent in this work, even providing species-level identifications for a ground squirrel and hawkmoth that were erroneously classified as birds by ImageNet-1k.

Ethical Statement

The goal of our work is to quantify the extent to which the natural world, namely wildlife, is under- and misrepresented in ImageNet-1k. We consider the ethical and practical implications that these biases and inaccuracies pose for the development and use of machine learning algorithms. With respect to the methods used in the study, we endeavored to follow best practices in responsible annotation of classes and images. While the initial ImageNet dataset was crowdsourced using the Amazon Mechanical Turk platform, there has since been extensive work that has questioned the ethics of anonymous crowdsourcing platforms (Irani and Silberman 2016; Shmueli et al. 2021). Given these reasons, as well as the importance of domain expertise for our analysis, we recruited annotators ourselves, taking into account characteristics such as the location, seniority and domain of expertise of each annotators. Furthermore, all of our annotators were remunerated at a fixed rate per image annotated, intended to approximate an hourly rate of \$50 per hour, with the final total annotation cost approaching \$10,000.

Acknowledgments

We would like to thank all of our expert annotators for their contributions: Kerri Lynn Ackerly, Ashton Asbury, Hazel Byrne, Michele Chiacchio, Eamon C. Corbett, Terrence Demos, Jessica Eggers, Miquéias Ferrão, Jeff Goddard, Alex Harman, Martin Hauser, Bob Jacobs, Daniel Kluza, Ashwini V. Mohan, Claude Nozères, Kris Sabbi, Subir B. Shaky, Jacob Socolar, David Vander Pluym, and Tyus D. Williams. This work was supported in part by the Canada CIFAR AI Chairs program.

References

Barbu, A.; Mayo, D.; Alverio, J.; Luo, W.; Wang, C.; Gutfreund, D.; Tenenbaum, J.; and Katz, B. 2019. Objectnet: A large-scale bias-controlled dataset for pushing the limits of

- object recognition models. *Advances in Neural Information Processing Systems (NeurIPS)*, 32.
- Barz, B.; and Denzler, J. 2020. Do we train on test data? purging CIFAR of near-duplicates. *Journal of Imaging*, 6(6): 41.
- Beery, S.; Agarwal, A.; Cole, E.; and Birodkar, V. 2021. The iWildCam 2021 competition dataset. *Preprint arXiv:2105.03494*.
- Beery, S.; Van Horn, G.; Mac Aodha, O.; and Perona, P. 2019. The iWildcam 2018 challenge dataset. *Preprint arXiv:1904.05986*.
- Beery, S.; Van Horn, G.; and Perona, P. 2018. Recognition in terra incognita. In *Proceedings of the European conference on computer vision (ECCV)*, 456–473.
- Beyer, L.; Hénaff, O. J.; Kolesnikov, A.; Zhai, X.; and Oord, A. v. d. 2020. Are we done with ImageNet? *Preprint arXiv:2006.07159*.
- Crawford, K.; and Paglen, T. 2019. Excavating AI: The politics of images in machine learning training sets. *AI and Society*.
- De Vries, T.; Misra, I.; Wang, C.; and Van der Maaten, L. 2019. Does object recognition work for everyone? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 52–59.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 248–255.
- Denton, E.; Hanna, A.; Amironesei, R.; Smart, A.; Nicole, H.; and Scheuerman, M. K. 2020. Bringing the people back in: Contesting benchmark machine learning datasets. *Preprint arXiv:2007.07399*.
- iNaturalist. 2022. Seek. https://www.inaturalist.org/pages/seek_app. Accessed: 2022-08-01.
- Irani, L. C.; and Silberman, M. S. 2016. Stories We Tell About Labor: Turkopticon and the Trouble with “Design”. In *Proceedings of the 2016 CHI conference on human factors in computing systems*, 4573–4586.
- Koch, B.; Denton, E.; Hanna, A.; and Foster, J. G. 2021. Reduced, Reused and Recycled: The Life of a Dataset in Machine Learning Research. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Lepage, D. 2022. Avibase. <https://avibase.bsc-eoc.org/>. Accessed: 2022-08-01.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Miller, G. A. 1995. WordNet: a lexical database for English. *Communications of the ACM*, 38(11): 39–41.
- Mora, C.; Tittensor, D. P.; Adl, S.; Simpson, A. G.; and Worm, B. 2011. How many species are there on Earth and in the ocean? *PLoS biology*, 9(8): e1001127.
- Northcutt, C. G.; Athalye, A.; and Mueller, J. 2021. Pervasive label errors in test sets destabilize machine learning benchmarks. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Prabhu, V. U.; and Birhane, A. 2021. Large image datasets: A pyrrhic win for computer vision? In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 1536–1546.
- Recht, B.; Roelofs, R.; Schmidt, L.; and Shankar, V. 2019. Do ImageNet classifiers generalize to ImageNet? In *International Conference on Machine Learning (ICML)*, 5389–5400.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. ImageNet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3): 211–252.
- Scheuerman, M. K.; Hanna, A.; and Denton, E. 2021. Do datasets have politics? Disciplinary values in computer vision dataset development. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2): 1–37.
- Shankar, S.; Halpern, Y.; Breck, E.; Atwood, J.; Wilson, J.; and Sculley, D. 2017. No classification without representation: Assessing geodiversity issues in open data sets for the developing world. *arXiv preprint arXiv:1711.08536*.
- Shmueli, B.; Fell, J.; Ray, S.; and Ku, L.-W. 2021. Beyond fair pay: Ethical implications of NLP crowdsourcing. *arXiv preprint arXiv:2104.10097*.
- Sullivan, B. L.; Wood, C. L.; Iliff, M. J.; Bonney, R. E.; Fink, D.; and Kelling, S. 2009. eBird: A citizen-based bird observation network in the biological sciences. *Biological conservation*, 142(10): 2282–2292.
- Tsipras, D.; Santurkar, S.; Engstrom, L.; Ilyas, A.; and Madry, A. 2020. From ImageNet to image classification: Contextualizing progress on benchmarks. In *International Conference on Machine Learning (ICML)*, 9625–9635.
- Van Horn, G.; Branson, S.; Farrell, R.; Haber, S.; Barry, J.; Ipeirotis, P.; Perona, P.; and Belongie, S. 2015. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 595–604.
- Van Horn, G.; Mac Aodha, O.; Song, Y.; Cui, Y.; Sun, C.; Shepard, A.; Adam, H.; Perona, P.; and Belongie, S. 2018. The iNaturalist species classification and detection dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 8769–8778.
- Van Horn, G.; and Perona, P. 2017. The devil is in the tails: Fine-grained classification in the wild. *Preprint arXiv:1709.01450*.
- Yang, K.; Qinami, K.; Fei-Fei, L.; Deng, J.; and Russakovsky, O. 2020. Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the ImageNet hierarchy. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAccT)*, 547–558.
- Yang, K.; Yau, J.; Fei-Fei, L.; Deng, J.; and Russakovsky, O. 2022. A study of face obfuscation in ImageNet. In *International Conference on Machine Learning (ICML)*.