


















Insect Identification in the Wild: The AMI Dataset

Aditya Jain^{1*}, Fagner Cunha^{2*}, Michael James Bunsen^{1*}, Juan Sebastián Cañas¹, Léonard Pasi¹, Nathan Pinoy³, Flemming Helsing³, JoAnne Russo⁴, Marc Botham⁵, Michael Sabourin⁶, Jonathan Fréchette¹², Alexandre Ancil⁷, Yacksecari Lopez⁸, Eduardo Navarro⁸, Filonila Perez Pimentel⁸, Ana Cecilia Zamora⁸, José Alejandro Ramirez Silva⁸, Jonathan Gagnon¹², Tom August⁵, Kim Bjerger³, Alba Gomez Segura⁵, Marc Bélisle⁹, Yves Basset⁸, Kent P. McFarland⁴, David Roy⁵, Toke Thomas Høye³, Maxim Larrivé¹⁰, and David Rolnick^{1,11}

¹ Mila - Quebec AI Institute

² Federal University of Amazonas

³ Aarhus University

⁴ Vermont Center for Ecostudies

⁵ UK Centre for Ecology & Hydrology

⁶ McGuire Center

⁷ Centre de données sur le patrimoine naturel du Québec

⁸ Smithsonian Tropical Research Institute

⁹ Université de Sherbrooke

¹⁰ Montreal Insectarium

¹¹ McGill University

¹² Independent Researcher

Abstract. Insects represent half of all global biodiversity, yet many of the world’s insects are disappearing, with severe implications for ecosystems and agriculture. Despite this crisis, data on insect diversity and abundance remain woefully inadequate, due to the scarcity of human experts and the lack of scalable tools for monitoring. Ecologists have started to adopt camera traps to record and study insects, and have proposed computer vision algorithms as an answer for scalable data processing. However, insect monitoring in the wild poses unique challenges that have not yet been addressed within computer vision, including the combination of long-tailed data, extremely similar classes, and significant distribution shifts. We provide the first large-scale machine learning benchmarks for fine-grained insect recognition, designed to match real-world tasks faced by ecologists. Our contributions include a curated dataset of images from citizen science platforms and museums, and an expert-annotated dataset drawn from automated camera traps across multiple continents, designed to test out-of-distribution generalization under field conditions. We train and evaluate a variety of baseline algorithms and introduce a combination of data augmentation techniques that enhance generalization across geographies and hardware setups. The dataset is made publicly available¹.

* Denotes equal contribution. Correspondence to: moth-ai@mila.quebec

¹ <https://github.com/RolnickLab/ami-dataset>

1 Introduction

Computer vision algorithms for species identification have become important tools in studying and protecting the Earth’s wildlife. Biodiversity applications have also been beneficial to computer vision, serving as a fruitful source of new benchmarks and algorithms [11, 23, 68, 75]. However, research in computer vision for biodiversity has often focused on large animals, not reflecting the full set of challenges faced in ecology and conservation.

There are over a million known species of insects on Earth, representing nearly one half of all known organisms (by comparison, there are only 10,000 species of birds, and even fewer mammals) [51, 62]. Insects play essential roles in nearly all terrestrial ecosystems, as well as providing ecosystem services such as pollination that are critical for humanity [60], yet a looming “insect apocalypse” threatens both their abundance and diversity [72]. Despite the crisis, data on insect populations is shockingly sparse [27], as a result of less resources, a plethora of species, and the high level of expertise typically required to study insects. Amid this data gap, computer vision tools for automated monitoring have the potential to radically reshape the study and conservation of insects. In particular, recent breakthroughs in camera traps for insects have made it possible to automatically photograph large numbers of individuals in the wild [38, 41, 58], opening up the the potential for image-processing algorithms to parse this raw data.

Species recognition from camera traps is a well-studied problem in the computer vision community [55], with common challenges including poor lighting, occlusion of individuals, camouflage, and blur [14]. However, working with insects presents distinct challenges from those faced in traditional camera traps, which are often designed for large animals. For example, while traditional camera trap images may only rarely contain a target species, almost every image from an insect camera trap includes insects – often hundreds of them. Where traditional camera traps may have a short list of species of interest, insect monitoring requires considering thousands of species at any given location, with many of them being extremely similar, and most with very little labeled data.

While recent authors have proposed the use of computer vision tools for detecting and identifying insects in camera trap data [8, 17, 20, 32, 43, 63], such studies have involved small datasets with only a few species represented. There has to date been no work assessing the performance of algorithms in identifying large number of insect species, or addressing the challenges of generalization across in-the-wild deployments. Our work here offers the following contributions:

- We introduce the **AMI** (Automated Monitoring of Insects) dataset, consisting of two parts: 1) **AMI-GBIF**, a dataset of ~ 2.5 M human-captured images curated from citizen science platforms and museum collections, 2) **AMI-Traps**, an expert-annotated dataset of 2,893 insect camera trap images (representing 52,948 labeled insects) collected from a global network of automated camera traps, designed to test in-the-wild performance (Fig. 1).

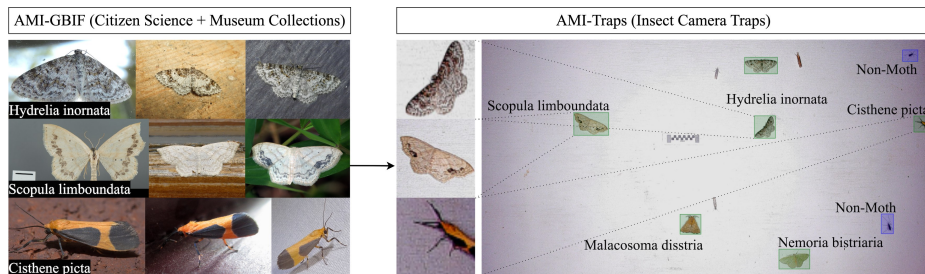


Fig. 1: Illustration of the difference between insect species identification from human-captured images and in the wild. On left are shown images of three moth species from our AMI-GBIF dataset, curated from citizen science data and museum collections. On right are shown images of the same species (along with other insects) within a photograph taken by an automated camera trap in the wild from our AMI-Traps dataset. One of the challenges of the AMI dataset is generalizing from AMI-GBIF to AMI-Traps without additional labeled data.

- We frame the problem of practical insect monitoring in the field as a set of computer vision tasks, entailing very fine-grained classification from long-tailed data, with the additional challenge of generalization across multiple kinds of imagery.
- We test a variety of strong baseline models for the proposed tasks, including a combination of data augmentation approaches that significantly improve performance under field conditions.

2 Related Work

Biodiversity Image Datasets. Advancements in computer vision have been closely intertwined with the creation of curated datasets such as ImageNet [26, 59], Pascal VOC [30], Open Images [44], and MS-COCO [46]. In recent years, the computer vision and machine learning communities have increasingly directed their focus towards addressing pressing human threats, including biodiversity loss [68] and climate change [57]. These domains not only hold significant potential for societal impact but also offer fertile ground for exploring intriguing machine learning questions such as domain generalization [77], long-tailed distributions [71], and fine-grained classification [74]. Such challenges have found especially fruitful testbeds in the classification of species, where datasets such as iNaturalist [70] have proven useful for both the computer vision and ecology communities. While some challenges have considered plant identification [15, 45, 53, 56], the majority have focused on animals, including datasets focused on birds [16, 69, 73], mammals [9, 14, 64], and fish [22, 25, 39, 40]. Such datasets have often leveraged concurrent advancements in sensor technology, including in satellite imagery and drones. In particular, camera traps have been widely adopted by biologists as an in situ and cost-effective means of wildlife

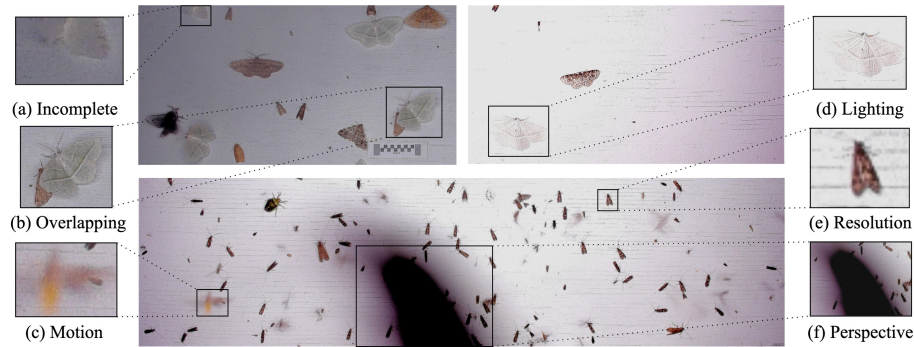


Fig. 2: Examples of challenges in the AMI-Traps dataset. (a) **Incomplete:** In some occasions, the body of an insect is incomplete due to its position in the automatically taken image. (b) **Overlapping:** Sometimes insects occlude each other. (c) **Motion:** Moving insects can be blurred and hard to classify. (d) **Lighting:** The lighting and camera exposure of the images is variable and can lead to poor contrast that masks some or all detail. (e) **Resolution:** Some insects are very small, leading to low-resolution images. (f) **Perspective:** While insects are generally positioned at a fixed distance, sometimes insects appear in the air or perched on the camera itself. Additionally, unexpected objects such as spider webs, dirt, and vegetation can be captured by the cameras and may be confused with objects of interest.

monitoring [55], and have in turn presented fascinating challenges for the computer vision community [14], including in classification [12, 29], object detection [13, 31], out-of-distribution generalization [34, 42, 50], and active learning [54]. Such algorithms have to date, however, been designed mostly for large animals such as mammals.

Insect Image Datasets. As algorithms and datasets for species recognition have proliferated, insects have traditionally been underrepresented. Within ImageNet1k [59], only 27 classes depicted insects, as compared to 57 for birds and 93 classes for wild mammals [49]. Some popular fine-grained classification datasets such as NABirds and CUB have focused exclusively on birds [69, 73], while camera trap datasets such as iWildCam have been focused on mammals [12]. The iNat2017 dataset [70] included 1,021 classes of insects out of 5,089 total classes in the dataset. By largely focusing on other taxonomic groups, prior works have neglected the challenges posed by insect data, such as extremely similar classes (reflecting far greater diversity than any other group of organisms), small size, hyperabundance, etc. In recent years, however, some works have introduced insect-focused benchmark datasets. The newly released BIOSCAN-1M dataset includes images and DNA sequences, with a focus on fly identification [33]. Despite the very large size of the dataset (1.1 million datapoints), the number of object classes involved is relatively small, with the focus being to distinguish between 40 taxonomic families of flies, reflecting the challenging nature of insect identification. Other recent datasets for computer vision include one focused on 8 classes of pollinators across Europe [61], and the Indian Butterflies dataset,

encompassing 30k images of butterflies from India across 315 species [67]. A variety of publications have used computer vision algorithms for detection of pest insects in specific agricultural applications [10, 52, 66].

Insect Camera Traps. The use of camera traps for studying insect biodiversity is relatively recent. Geissmann et al. [32] introduced the Sticky Pi device, which uses sticky cards to capture (and kill) insects, which are then automatically photographed. Bjerge et al. [21] introduced an automated camera trap using UV light to attract and photograph moths and other insects (without killing them), which was built on by subsequent works [18, 20]. A similar approach focused on diurnal pollinators via a camera trap installed above flowers [17, 19]. These works also propose deep learning methods for processing the resulting images, considering both object detection and classification of a small number of insect categories. The lack of fine-grained classification is largely due to the time required to annotate images from the traps, making it difficult to obtain a large enough labeled dataset from trap data alone.

3 Benchmark Design

In this paper, we propose a benchmark dataset for insect species identification that addresses the challenges faced in real-world use cases. These includes both (1) challenges that distinguish insect identification from other kinds of species identification, notably the massive diversity of often-similar species, and (2) the particular challenges associated with deploying these algorithms “in the wild”.

To address the first family of challenges, we introduce a dataset for fine-grained insect classification that (to our knowledge) is by far the largest yet released, with ~ 2.5 M images across over 5,000 classes. We concentrate on moths, which are readily drawn to UV camera traps and can often be visually identified to species (unlike some other insects) [21].

The second family of challenges is more complex – involving the gap between field conditions and the high-quality images typically used for training species identification algorithms. Figure 2 illustrates some of the challenges that computer vision algorithms face on data from automated insect camera traps, in contrast to images taken by humans. Insects may occlude each other or be partially “out of frame”. Motion blur and lighting issues may occur, along with objects in unexpected positions. The resolution of images may be very low, owing to the small size of insects at a stereotyped distance from the camera. Also, insect camera traps typically use webcams and PCB camera modules due to their cost and software integration capabilities, and the resulting images are often both structurally different and overall lower quality than those taken by human photographers.

As described in the preceding section, previous authors working with insect camera traps have addressed the problem of challenging camera trap data by training algorithms on annotated data from the particular camera trap in question. However, manual annotation is extremely time-intensive, since it must be performed by expert entomologists, of whom there are very few in any given

region. Insect species vary greatly across regions, meaning that fine-grained annotations in one location are generally not useful in another. Furthermore, the hardware used in insect camera traps also varies greatly (with considerable differences in lighting and background, as well as image sharpness, contrast, noise reduction settings, and compression artifacts, which are exacerbated on small scale crops), meaning that even two datasets from the same region may not be comparable.

In this paper, we remove the problem of annotation by solving a more challenging computer vision problem. Namely, we propose to train algorithms on abundant labeled images of insects taken by human photographers, then use these algorithms on out-of-distribution camera trap images. Our benchmark dataset is designed to test this functionality, by including both a large training set of curated photographs taken by humans and an out-of-distribution test set drawn from a geographically diverse set of insect camera traps.

The AMI dataset is therefore designed to test fine-grained classification of moth species both in in-distribution and out-of-distribution contexts. It also includes the task of binary classification of insects as moth or non-moth, which is a prerequisite for camera traps since moths are not the only insects to appear. Note that we do not consider the task of object detection, which in applications is a precursor to insect identification. This is because object detection against a relatively plain background is a simpler task, where good results have already been obtained in the insect camera trap literature [18, 20, 21, 32].

4 The AMI Dataset

In this section, we describe the construction of the datasets that we release in this work: **AMI-Traps** and **AMI-GBIF**. For AMI-Traps, we collected insect camera trap data by field deployments in three regions: NE-America (Northeast of North America), which includes data from Quebec in Canada and Vermont in the USA; W-Europe (Western Europe), which includes data from the UK and Denmark; and C-America (Central America), comprising data from Panama. For AMI-GBIF, we compiled public data for the same three regions using the Global Biodiversity Information Facility (GBIF) [7]. The GBIF platform provides access to biodiversity data aggregated from multiple sources, such as museums, research institutions, and citizen science platforms such as iNaturalist [4], Observation.org [5], and BugGuide [1]. For brevity, we sometimes refer to AMI-Traps and AMI-GBIF as AMI-T and AMI-G, respectively. Figure 3 shows examples of images from the two datasets and Tab. 1 compares the number of images and coverage of classes at different taxonomic levels.

4.1 AMI-Traps Dataset

Data Collection. The source images in the AMI-Traps dataset were collected from a variety of different insect camera traps deployed across 22 monitoring stations in the three regions previously described. The cameras were configured

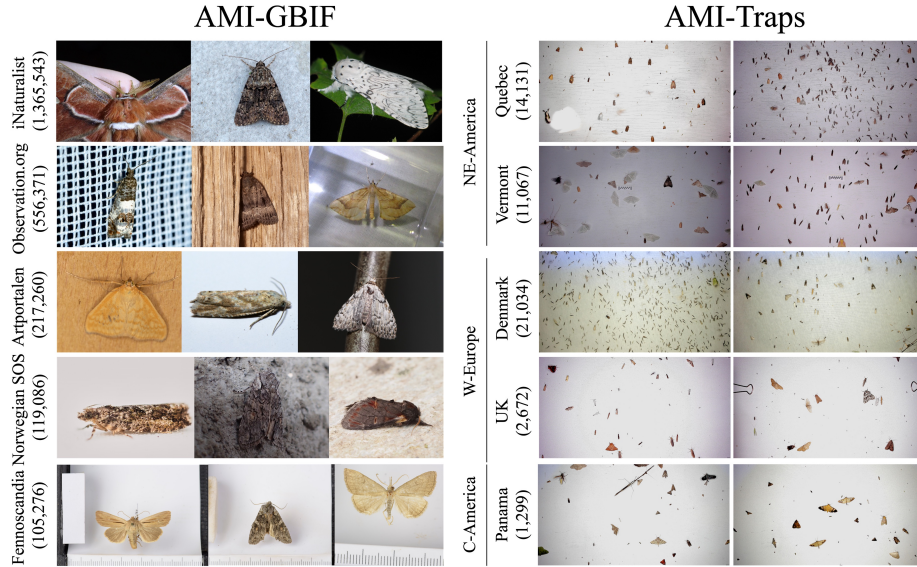


Fig. 3: Sample images from the AMI Dataset. The AMI Dataset is composed of (1) AMI-GBIF, curated from a number of sources with imagery from citizen science and museum collections, and (2) AMI-Traps, drawn from automated camera traps for insects across five countries in three regions. The number of individual insects annotated per sub-dataset is denoted in parentheses.

to be motion-sensitive, as well as taking photos on a fixed interval of 10 minutes; we selected a subset of images for annotation so as to avoid highly similar images. In total, 2,893 trap images were sampled from all cameras and made available on the Label Studio annotation platform (see Appendix B for more details).

Annotation Process. Images were annotated by entomologists with expertise in the species of the particular region. Labeling was conducted according to a three-step process. 1) We designed a custom object detection model to suggest boxes around objects in the image (refer to Appendix B for details on this model), to save annotators time. Annotators were asked to correct errors in the proposed boxes as well as drawing any missing boxes. Insects appearing smaller than 1 centimeter were ignored. 2) Annotators were asked to label image crops according to one of three coarse categories: moth, non-moth, or unidentifiable. Some insects could not be labeled due to motion blur, occlusion, or other factors. 3) For insects labeled as moths, annotators were asked to select the narrowest taxonomic group that could be determined from the image - identifying moths to species where possible, if not species then genus, and if not genus then family.

Data Statistics. 2,893 source images from the traps were annotated (see Tab. B.2 in the supplementary for further details) which includes Vermont (900), Denmark (892), Quebec (533), UK (446), and Panama (122). A total of 52,948 insects were labeled within these images, including 37,105 (70.08%) non-moths,

14,105 (26.64%) moths, and 1,738 (3.28%) unidentifiable. Out of the 14,105 moth crops, 10,854 (76.95%), 7,065 (50.09%) and 5,374 (38.10%) crops were labeled at the family, genus, and species level respectively.

4.2 AMI-GBIF Dataset

AMI-GBIF has two sub-datasets: binary and fine-grained classification. The fine-grained classification dataset contains images for moth species in each of the regions covered by AMI-Traps. The binary classification dataset contains images of moths aggregated across all regions and non-moths from anywhere on earth.

Dataset Construction. First, in collaboration with regional moth specialists, we created “checklists” of species that could occur in each of the three regions of interest. It is worth noting that some species are included in more than one checklist, though most are unique to a single region. Next, we compared the checklists against the GBIF taxonomic backbone [2], merged synonym names (since some species have multiple accepted names), and manually investigated doubtful or unmatched names. Following that, we used the GBIF occurrence search tool to download metadata for all observations from the order Lepidoptera with images, which returned 15.1 million observations [3]. We then fetched images from observations in which the species is included in our regional checklists. The number of training images per species were limited to 1,000, though for many species it was considerably lower due to the long-tailed nature of the data. All images were resized to have the smallest dimension of 450px and packaged using the WebDataset format [6].

Quality Control. After downloading the images, we applied a set of cleaning steps to ensure the quality of the dataset. First, we removed duplicate URLs, i.e., images linked to more than one observation. This may occur when an image contains multiple individual insects, or when a placeholder is used for an observation (for example, Fig. A.1b in Appendix A shows a picture used as a placeholder for more than 100,000 observations). Next, we excluded images from specific datasets (refer to Tab. A.1 in Appendix A) that we found to have a high percentage of images unsuitable for training our models, for instance, images of body parts or text only. Another issue we identified is that some observations contained thumbnail images that are too small to provide sufficient information to train our models; we removed all images with a width or height less than 64 pixels. Finally, as our dataset targets only adult insects, we removed images containing non-adult individuals such as eggs and caterpillars. For observations where the life stage was reported by the observer, approximately 60% of the cases, we removed all occurrences tagged as non-adult. For the remaining images with a missing life stage tag, we designed a classification model to predict the life stage (refer to Appendix A for details about this model) and removed all images classified as non-adult.

Dataset Split. Each GBIF observation may contain more than one image, but each image is considered an independent instance of the species. However, as images belonging to the same observation may be very similar, we keep images from the same observation in the same data split. To do so, we split the dataset

using the observation ID as reference, with the proportion of 70%, 10%, and 20% for the train, validation and test sets respectively. We apply a stratified split, i.e., we guarantee that this proportion is also maintained at the species level. We exclude species with fewer than 5 images, so the minimum number of images for a species in the training set is 4, and in the test set is one. Due to the data split based on observation ID, the validation set might lack images for some species. As some species may belong to multiple checklists, we keep the train/val/test splits consistent across the regions to prevent data leakage when training a unified model for all regions.

Binary Classification. Since moths are not the only insects that appear on the camera trap screen, for our proposed benchmark task of binary classification, we curate a dataset comprising 350,000 images for each of the moth and non-moth categories. The moth class consists of adult moth images randomly sampled from the AMI-GBIF dataset. For the non-moth category, we leveraged expert knowledge of the taxonomic groups likely to appear on the screen, fetching images of adult individuals from GBIF of the following groups: Diptera (85,474), Hemiptera (75,053), Odonata (68,770), Coleoptera (45,954), Araneae (21,321), Orthoptera (21,261), Formicidae (12,206), Ichneumonidae (6,581), Trichoptera (5,316), Neuroptera (4,880), Opiliones (1,880), Ephemeroptera (1,304) — this includes mostly insects, with some Arachnids. The dataset is split into train, validation, and test sets with proportions of 70%, 10%, and 20%, respectively. Specifically for the moth category, the splits are kept consistent with the splits for the fine-grained species classification task.

Statistics. The final dataset contains ~ 2.5 M images from 5,364 species, 1,734 genera and 77 families of moths, along with 350,000 additional images of non-moth groups. The binary dataset contains 700,000 images, composed of 350,000 moth images and the 350,000 non-moth images. The fine-grained dataset contains 2,564,392 moth images, and Tab. 1 shows the number of images and species for each region. Though species are selected based on their presence in regions of interest, we include all GBIF observations of these species from anywhere on Earth (see Fig. A.2 in supplementary for the global distribution of images). As in other biodiversity datasets, the number of images per species in AMI-GBIF follows a long-tailed distribution (see Fig. C.3 in supplementary). Specifically, only 3,099 species have more than 100 images and are considered under the “many-shot” bucket, 1,358 species contains between 20 and 100 images and are considered “medium-shot”, and there are 907 species with less than 20 training images, considered as “few-shot” categories.

5 Methods

In this section, we describe the methods used to assess the difficulty of the proposed benchmark tasks. We evaluate a variety of high-performing image classification architectures and use extensive data augmentation to address the challenges of domain shift between the two datasets. Motivated by the particular structure of the AMI-Traps dataset, we add a mixed-resolution (MixRes) data

Table 1: Comparison between the AMI-GBIF and AMI-Traps dataset for the fine-grained classification task. Images for the traps dataset denote the individual insect crops and not the raw trap images. The last three rows show the unique classes for each of the three taxonomic levels for the two datasets for different regions. It is also important to note that regions have species in common.

	All Regions		NE-America		W-Europe		C-America	
	AMI-G	AMI-T	AMI-G	AMI-T	AMI-G	AMI-T	AMI-G	AMI-T
Images	2,564,392	14,105	1,179,943	9,066	1,579,333	4,066	99,405	973
Families	77	43	68	34	65	26	27	11
Genera	1,734	344	887	204	957	192	385	12
Species	5,364	516	2,497	277	2,603	244	636	6

augmentation technique. As we describe in the following section, our approaches perform quite well, with MixRes leading to an especially significant performance improvement.

Models. We work with five computer vision models widely used in identification of wildlife: ResNet50 [35], MobileNetV3-Large [37], ConvNeXt-T [47], ConvNeXt-B [47], and ViT-B/16 [28]. We select ResNet50 because it is the most popular backbone used as a baseline for almost any computer vision task. Considering the possibility that the trained models may need to run on hardware with limited resources, such as edge devices and personal computers, for in-field predictions, we choose MobileNetV3Large as a lightweight model alternative. Furthermore, considering that vision transformers have surpassed convolutional neural networks and become the state-of-the-art models for image classification, we select the ViT-B/16 model to provide a baseline from this family of architectures. Finally, for representing modern convolutional neural networks, we choose the recently introduced ConvNeXt architecture, which has proven to be competitive with transformers using only standard ConvNet modules. From this family of models, we select two versions: the base version (ConvNeXt-B) and, considering lightweight usage scenarios, we also include the tiny version (ConvNeXt-T). Unless otherwise stated, we use a default input resolution of 128x128 for all models.

Data Augmentation Approaches. One characteristic of the in-the-wild images is the large variety of crops due to the diversity of insect sizes. A considerable number of them have small resolutions, below 64×64 . At this scale, some important visual features that the model may have learned from high-resolution GBIF images might not be identifiable, leading to inaccurate predictions. This issue can be further amplified when these low-resolution crops are scaled up to the input model size. To compel the model to learn features for low-resolution inputs and avoid bias towards visual features only distinguishable in high-resolution images from other domains, we use a data augmentation technique we refer to as Mixed Resolution (MixRes) that simulates this low-resolution input during training (see Fig. D.4 in supplementary). This approach consists of randomly resizing the training input image to 25% or 50% of the model input size with a

probability ρ for each case. In our experiments ρ is set to 0.25, maintaining the original input size with probability 0.5. We employ MixRes in all experiments unless otherwise stated.

In addition to MixRes, we also use a set of standard data augmentation operations: random resize crop with scale of 0.3 to 1.0 and an aspect ratio between 3/4 and 4/3, random horizontal flip with 50% probability, and RandAugment [24] with $N = 2$ operations and magnitude $M = 9$. Finally, the image is normalized using the mean and standard deviation of ImageNet and scaled to the input network size.

Implementation Details. We train all models with the AMI-GBIF training sets using the ImageNet-1K pretrained weights from the timm library [76] (for ConvNext and ViT we used ImageNet-22K weights). Each model is trained with a batch size of 128 and the AdamW optimizer [48] for 30 epochs. The initial learning rate is set to 0.001, with a weight decay of $1e-5$. The learning rate is linearly warmed up from 0 to the initial value for 2 epochs and then decayed using cosine scheduling [36]. We also apply label smoothing [65] with a value of 0.1.

Inference Settings. At test time, the insect crops from AMI-Traps are padded with black pixels to make them square, preventing the distortion of the shape of the insect before prediction. Next, the images are normalized with the mean and standard deviation of ImageNet and, finally, resized to the model input size. No padding is done for the test images from AMI-GBIF.

Higher Taxonomic Rank Prediction. We consider the accuracy of our baseline models both with respect to the exact insect species predicted and with respect to higher level taxonomic groups such as genus and family. Indeed some images in our dataset cannot be labeled at the species level at all due to the quality or perspective of the image, and our annotators instead provided labels at the genus or family level. In our experiments, to make predictions at a higher taxonomic level, we simply sum together species-level predictions to get genus results, and similarly, the genus predictions are grouped together to get family predictions.

Evaluation Metrics. During the binary classification step, it is important to achieve high precision by filtering out as many non-moths as possible to prevent the propagation of errors in the pipeline. At the same time, it is also important to assess recall because if too many moths are discarded at this step, it may impact the utility of the data for ecological studies. Therefore, for the binary classification task, along with accuracy, we also consider precision, recall, and F1-score as evaluation metrics. For the fine-grained classification task, we evaluate the model in terms of micro- and macro-averaged accuracy at the species, genus and family levels. Micro-averaging aggregates all class samples to calculate the accuracy while the macro-average is the average of per-class accuracy. Unless otherwise stated, the reported accuracy is micro-averaged. We conduct each experiment for three independent runs and report the mean and standard deviation.

6 Experiments and Results

Binary Classification Results. This task involves distinguishing moths from non-moths. Here we train a ResNet50 model on the AMI-GBIF dataset and evaluate performance both on the AMI-GBIF held-out test set and on the entirety of the AMI-Traps dataset. On AMI-GBIF (in-distribution), the model achieves a test accuracy of $97.77\% \pm 0.02\%$. Table 2 shows the results for the AMI-Traps dataset in more detail. We observe that the model has a high recall but a lower precision, which means it is over-predicting the moth class. Also, the bigger the image crop, the better the accuracy of the insect classification; many non-moths where the model errs are extremely small and are therefore highly blurred in the trap images. Table E.3 in the supplementary illustrates the improvement in model performance when crops are square padded (our default at inference time) versus an ablation where padding is not performed.

Table 2: Binary classification model evaluation results on the AMI-Traps dataset after training on AMI-GBIF, showing that recall is high throughout, while precision drops on very small insects, which are predominantly non-moths. Crops $>100\text{px}$ or $>150\text{px}$ represent that atleast one side of the box is greater than the threshold.

	Moths	Non-moths	Total	Accuracy	Precision	Recall	F1-score
All crops	14,105	37,105	51,120	86.48 ± 1.00	68.31 ± 1.70	95.03 ± 0.39	79.48 ± 1.27
Crops $> 100\text{px}$	11,630	16,534	28,164	88.39 ± 0.97	79.46 ± 1.43	96.99 ± 0.20	87.35 ± 0.94
Crops $> 150\text{px}$	6,446	5,961	12,407	92.95 ± 0.22	89.83 ± 0.43	97.48 ± 0.12	93.70 ± 0.30

Fine-grained Classification Results. The fine-grained classification task aims to identify moths to the species level (and secondarily to genus or family level). In our experiments, we train models on AMI-GBIF data for each of the three regions of the AMI dataset, evaluating both on AMI-GBIF (in-distribution) and AMI-Traps (out-of-distribution) data (see Tab. 3). Not all moth species present in a region appear regularly at camera traps; hence, the list of species in AMI-Traps is a subset of species in AMI-GBIF for a given area (see Tab. 1). Therefore, we also evaluate the models on an in-distribution subset (Test \ddagger) from AMI-GBIF containing only the species present in AMI-Traps. Meanwhile, for AMI-Traps, there are 16 species not present in AMI-GBIF that were removed during our evaluation procedure (these represent unexpected species not included in the regional checklists used to define AMI-GBIF). Overall, the best model at the species level is ConvNeXt-B for both AMI-GBIF and AMI-traps. For C-America, annotations are mostly provided at the genus or family level, due to the challenging nature of visual identification of moths in this biodiversity-rich region. Furthermore, while the genus and family accuracies are similar between N-America and W-Europe, they are much lower for C-America. This can be attributed to the fact that there are 10 times fewer images in the training set for this region. Note that for genus and family evaluation, there are some cases where the accuracy is lower than the species level. This occurs because those

Table 3: Accuracy results for the fine-grained classification task. Overall, ConvNeXt-B achieves the best results for both AMI-GBIF and AMI-traps, with ViT-B/16 achieving slightly better results for some cases at the genus and family levels. Test \ddagger represents a subset of the AMI-GBIF test set, containing only the species present in AMI-Traps. Species-level results on AMI-Traps for the C-America region are not available since most annotations were provided only at higher taxonomic levels.

Region	Model	AMI-GBIF				AMI-Traps					
		Species		Species		Genus			Family		
		Micro Test	Micro Test \ddagger	Micro Top1	Macro Top1	Micro Top1	Macro Top1	Macro Top1	Micro Top1	Macro Top1	Macro Top1
NE-America	ResNet50	87.56±0.12	89.39±0.14	71.86±1.02	64.96±2.00	80.07±0.49	74.66±0.35	90.44±0.41	69.30±4.15		
	MBNetV3L	83.73±0.17	85.80±0.29	67.41±0.72	61.21±1.20	74.83±1.30	71.20±2.47	89.31±0.62	64.48±1.87		
	ConvNeXt-T	89.81±0.05	91.21±0.06	74.89±0.05	68.00±0.76	83.26±1.22	77.26±0.85	90.61±0.53	74.60±0.64		
	ConvNeXt-B	90.45±0.03	91.66±0.03	76.75±1.52	68.56±0.62	85.46±0.83	78.28±0.14	91.07±0.29	76.87±1.10		
	ViT-B/16	86.39±0.05	87.94±0.10	77.43±0.79	68.02±0.34	85.57±0.98	78.97±0.72	91.45±0.34	73.66±1.49		
W-Europe	ResNet50	86.45±0.08	90.25±0.03	79.39±0.39	72.84±0.21	78.95±0.26	73.22±0.60	89.13±0.42	75.24±1.85		
	MBNetV3L	82.20±0.05	86.47±0.13	76.30±1.29	70.35±0.93	76.35±1.19	72.88±1.28	88.33±0.47	72.21±0.38		
	ConvNeXt-T	88.82±0.03	92.05±0.05	82.44±0.89	74.59±0.62	81.86±0.79	74.15±0.59	90.08±0.59	74.58±2.93		
	ConvNeXt-B	89.35±0.01	92.48±0.05	82.87±0.39	76.41±0.50	82.45±0.44	75.98±1.05	91.27±0.60	79.39±2.70		
	ViT-B/16	85.45±0.03	89.22±0.11	82.44±0.36	75.25±0.98	83.01±0.26	75.96±0.35	91.68±0.17	75.81±0.79		
C-America	ResNet50	87.46±0.23	95.92±0.68	-	-	47.44±1.11	66.59±0.60	69.18±0.24	42.40±0.31		
	MBNetV3L	83.68±0.24	92.89±0.27	-	-	41.03±2.22	59.32±1.18	68.34±1.69	41.42±2.98		
	ConvNeXt-T	87.65±0.11	95.46±0.46	-	-	55.77±3.33	70.68±1.58	69.46±0.42	45.35±0.67		
	ConvNeXt-B	88.62±0.09	96.22±0.26	-	-	51.92±3.85	68.63±1.94	72.24±1.47	44.31±1.32		
	ViT-B/16	79.80±0.26	88.36±1.02	-	-	47.44±4.44	64.70±4.07	69.87±1.51	42.00±3.60		

higher taxonomic levels include crops that have labels only at higher levels. In the supplementary material, we include additional results for AMI-GBIF long-tail accuracies (Tab. E.4), a pre-training ablation study (Tab. E.5), and a study analyzing classification accuracy as a function of prediction confidence (Fig. E.5).

Data augmentation approaches. To assess the effectiveness of the mixed resolution data augmentation technique, we conduct an ablation study using ResNet50 for the NE-America region. As the effects of MixRes could already be covered by other data augmentation operations used in our preprocessing, such as random resize crop or one of the operations of RandAugment, we analyze multiple combinations of these operations. Additionally, we evaluate a higher input resolution, using the original 224 input size of ResNet50, to determine whether this leads to better predictions. Table 4 shows the results. While applying MixRes results in a slight decrease in accuracy within the in-distribution AMI-GBIF test set, for the in-the-wild AMI-Traps test set there is a substantial improvement, ranging between 4% and 7% when compared to models not trained with MixRes. Another important observation is that, although RandAugment and random resize together contribute to an overall improvement of around 16%, MixRes remains additionally beneficial. Finally, despite an increase of more than 4% in accuracy within the original domain when using a higher resolution, there is a marginal decrease in the accuracy of the in-the-wild results.

7 Conclusion

Many ecological applications require species-level image classification to meet their goals. Despite advances in computer vision, there remain challenges in

Table 4: Ablation study on data augmentation techniques using a ResNet50 for the NE-America region evaluated on micro-averaged top-1 species accuracy. “RandAug” denotes the combination of the RandAugment and random resize crop data augmentation techniques. AMI-GBIF[‡] test set contains only the species present in AMI-Traps.

Resolution	RandAug	MixRes	AMI-GBIF [‡]	AMI-Traps
128	✗	✗	89.18±0.08	51.46±2.73
128	✓	✗	89.71±0.16	67.88±1.27
128	✗	✓	87.41±0.06	58.41±0.30
128	✓	✓	89.39±0.14	71.86±1.02
224	✓	✓	92.97±0.17	68.41±1.06

implementing fine-grained classification that is effective in real-world scenarios. Here we outline the challenges of insect identification in the wild and propose a benchmark dataset to tackle them. We present AMI-GBIF, a dataset with 2.5 million labeled photographs of 5,000 moth species. We also introduce AMI-Traps, a dataset of images captured by automated insect monitoring systems across five countries, labeled by regional experts. We evaluate multiple classification models both on the AMI-GBIF dataset and in out-of-distribution generalization to AMI-Traps, with ConvNeXt-B generally performing the best. We also evaluate a range of data augmentation techniques, finding that a combination of simple approaches, notably mixed-resolution augmentation, considerably improves generalization under field conditions.

Several limitations of this work should be noted. First, while annotating AMI-Traps already occupied a considerable time for expert entomologists, further expanding it to a greater number of species and geographic regions would be helpful in further stress-testing field deployments of fine-grained classification algorithms. Second, while we focus on camera traps for nocturnal insects, it would be useful also to consider diurnal insect camera traps, for example those aimed at pollinators, which pose some additional challenges, such as multiple perspectives and a diversity of backgrounds. We hope furthermore that future works will consider more sophisticated domain adaptation approaches building on the baselines we present here.

This work represents a step towards computer vision algorithms for species recognition that are effective in the ultra-challenging regime of insect identification and work effectively in field conditions. The approach we propose should make it possible in particular to generalize algorithms for insect camera traps rapidly and cost-effectively across different hardware setups and regions of interest. We hope that the datasets, challenges and strategies presented here will propel further research in computer vision for insects and help accelerate initiatives to understand and protect our planet’s biodiversity.

Acknowledgements

We thank Annie-Shan Morin, Anna Viklund, and Jamie Alison for their helpful input and support to the project.

This research was generously funded by Mila Quebec AI Institute and Ministère de l'Économie, de l'Innovation et de l'Énergie. The work was additionally supported in part by the Canada CIFAR AI Chairs program (NSERC 585136), the Global Climate Center on AI and Biodiversity Change (ABC) (NSF OISE-2330423), the Montreal Insectarium, the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES-PROEX) - Finance Code 001, the Amazonas State Research Support Foundation - FAPEAM - through the POS-GRAD project 2024/2025, the AMBER project funded by the abrdn Charitable Foundation, the Aage V. Jensen Naturfond grant no. N2024-0013, the European Union's Horizon Europe Research and Innovation programme under grant agreement no. 101060639 (MAMBO), and a Discovery Grant by NSERC. The computing resources were provided by Mila Quebec AI Institute, the Digital Research Alliance of Canada, and NVIDIA Corporation.

A.J., F.C., and M.J.B. led dataset creation and experiments. J.S.C. and L.P. contributed to machine learning pipeline development. N.P., F.H., J.R., M.B., M.S., J.F., A.A., Y.L., E.N., F.P.P., A.C.Z., J.A.R.S., and J.G. performed insect identification for the AMI-Traps dataset. T.A., K.B., A.G.S., M.B., Y.B., K.P.M., D. Roy, T.T.H., M.L., and D. Rolnick guided and advised the project.

References

1. BugGuide.Net: Identification, Images, & Information For Insects, Spiders & Their Kin For the United States & Canada, <https://bugguide.net/> accessed Feb 12 2024
2. GBIF Backbone Taxonomy, <https://doi.org/10.15468/39omei> accessed Nov 13 2023
3. GBIF Occurrence Download, <https://doi.org/10.15468/d1.4ua6wg> accessed Oct 19 2023
4. iNaturalist, www.inaturalist.org, accessed Feb 9 2024
5. Observation.org, <https://observation.org>, accessed Feb 12 2024
6. WebDataset, <https://github.com/webdataset/webdataset>, accessed on April 2 2024
7. What is GBIF?, <https://www.gbif.org/what-is-gbif> accessed Sep 28 2023
8. Alison, J., Alexander, J.M., Diaz Zeugin, N., Dupont, Y.L., Iseli, E., Mann, H.M., Høye, T.T.: Moths complement bumblebee pollination of red clover: a case for day-and-night insect surveillance. *Biology Letters* **18**(7), 20220187 (2022)
9. Anton, V., Hartley, S., Geldenhuis, A., Wittmer, H.U.: Monitoring the mammalian fauna of urban areas using remote cameras and citizen science. *Journal of Urban Ecology* **4**(1), juy002 (2018)
10. Badgujar, C.M., Armstrong, P.R., Gerken, A.R., Pordesimo, L.O., Campbell, J.F.: Real-time stored product insect detection and identification using deep learning: System integration and extensibility to mobile platforms. *Journal of Stored Products Research* **104**, 102196 (2023)

11. Beery, S.: Scaling biodiversity monitoring for the data age. *XRDS: Crossroads, The ACM Magazine for Students* **27**(4), 14–18 (2021)
12. Beery, S., Agarwal, A., Cole, E., Birodkar, V.: The iwildcam 2021 competition dataset. arXiv preprint arXiv:2105.03494 (2021)
13. Beery, S., Morris, D., Yang, S.: Efficient pipeline for camera trap image review. arXiv preprint arXiv:1907.06772 (2019)
14. Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: Proceedings of the European conference on computer vision (ECCV). pp. 456–473 (2018)
15. Beery, S., Wu, G., Edwards, T., Pavetic, F., Majewski, B., Mukherjee, S., Chan, S., Morgan, J., Rathod, V., Huang, J.: The auto arborist dataset: A large-scale benchmark for multiview urban forest monitoring under domain shift. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21294–21307 (2022)
16. Berg, T., Liu, J., Woo Lee, S., Alexander, M.L., Jacobs, D.W., Belhumeur, P.N.: Birdsnap: Large-scale fine-grained visual categorization of birds. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2011–2018 (2014)
17. Bjerge, K., Alison, J., Dyrmann, M., Frigaard, C.E., Mann, H.M., Høye, T.T.: Accurate detection and identification of insects from camera trap images with deep learning. *PLOS Sustainability and Transformation* **2**(3), e0000051 (2023)
18. Bjerge, K., Frigaard, C.E., Karstoft, H.: Object detection of small insects in time-lapse camera recordings. *Sensors* **23**(16), 7242 (2023)
19. Bjerge, K., Geissmann, Q., Alison, J., Mann, H.M., Høye, T.T., Dyrmann, M., Karstoft, H.: Hierarchical classification of insects with multitask learning and anomaly detection. *Ecological Informatics* **77**, 102278 (2023)
20. Bjerge, K., Mann, H.M., Høye, T.T.: Real-time insect tracking and monitoring with computer vision and deep learning. *Remote Sensing in Ecology and Conservation* **8**(3), 315–327 (2022)
21. Bjerge, K., Nielsen, J.B., Sepstrup, M.V., Helsing-Nielsen, F., Høye, T.T.: An automated light trap to monitor moths (lepidoptera) using computer vision-based tracking and deep learning. *Sensors* **21**(2), 343 (2021)
22. Boom, B.J., He, J., Palazzo, S., Huang, P.X., Beyan, C., Chou, H.M., Lin, F.P., Spampinato, C., Fisher, R.B.: A research tool for long-term and continuous analysis of fish assemblage in coral-reefs using underwater camera footage. *Ecological Informatics* **23**, 83–97 (2014)
23. Borowiec, M.L., Dikow, R.B., Frandsen, P.B., McKeeken, A., Valentini, G., White, A.E.: Deep learning as a tool for ecology and evolution. *Methods in Ecology and Evolution* **13**(8), 1640–1660 (2022)
24. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 702–703 (2020)
25. Cutter, G., Stierhoff, K., Zeng, J.: Automated detection of rockfish in unconstrained underwater videos using haar cascades and a new image dataset: Labeled fishes in the wild. In: 2015 IEEE Winter Applications and Computer Vision Workshops. pp. 57–62. IEEE (2015)
26. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 248–255. IEEE (2009)

27. Didham, R.K., Basset, Y., Collins, C.M., Leather, S.R., Littlewood, N.A., Menz, M.H., Müller, J., Packer, L., Saunders, M.E., Schönrogge, K., et al.: Interpreting insect declines: seven challenges and a way forward. *Insect Conservation and Diversity* **13**(2), 103–114 (2020)
28. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net (2021), <https://openreview.net/forum?id=YicbFdNTTy>
29. Dubey, A., Ramanathan, V., Pentland, A., Mahajan, D.: Adaptive methods for real-world domain generalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14340–14349 (2021)
30. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. *International Journal of Computer Vision* **88**, 303–338 (2010)
31. Fennell, M., Beirne, C., Burton, A.C.: Use of object detection in camera trap image identification: Assessing a method to rapidly and accurately classify human and animal detections for research and application in recreation ecology. *Global Ecology and Conservation* **35**, e02104 (2022)
32. Geissmann, Q., Abram, P.K., Wu, D., Haney, C.H., Carrillo, J.: Sticky pi is a high-frequency smart trap that enables the study of insect circadian activity under natural conditions. *PLoS Biology* **20**(7), e3001689 (2022)
33. Gharaee, Z., Gong, Z., Pellegrino, N., Zarubiieva, I., Haurum, J.B., Lowe, S.C., McKeown, J.T., Ho, C.C., McLeod, J., Wei, Y.Y.C., et al.: A step towards worldwide biodiversity assessment: The bioscan-1m insect dataset. arXiv preprint arXiv:2307.10455 (2023)
34. Gulrajani, I., Lopez-Paz, D.: In search of lost domain generalization. arXiv preprint arXiv:2007.01434 (2020)
35. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
36. He, T., Zhang, Z., Zhang, H., Zhang, Z., Xie, J., Li, M.: Bag of tricks for image classification with convolutional neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 558–567 (2019)
37. Howard, A., Sandler, M., Chu, G., Chen, L.C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al.: Searching for mobilenetv3. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1314–1324 (2019)
38. Høye, T.T., Årje, J., Bjerger, K., Hansen, O.L., Iosifidis, A., Leese, F., Mann, H.M., Meissner, K., Melvad, C., Raitoharju, J.: Deep learning and computer vision will transform entomology. *Proceedings of the National Academy of Sciences* **118**(2), e2002545117 (2021)
39. Katija, K., Orenstein, E., Schlining, B., Lundsten, L., Barnard, K., Sainz, G., Boulais, O., Cromwell, M., Butler, E., Woodward, B., et al.: Fathomnet: A global image database for enabling artificial intelligence in the ocean. *Scientific reports* **12**(1), 15914 (2022)
40. Kay, J., Kulits, P., Stathatos, S., Deng, S., Young, E., Beery, S., Van Horn, G., Perona, P.: The caltech fish counting dataset: A benchmark for multiple-object tracking and counting. In: European Conference on Computer Vision. pp. 290–311. Springer (2022)

41. van Klink, R., August, T., Bas, Y., Bodesheim, P., Bonn, A., Fossøy, F., Høye, T.T., Jongejans, E., Menz, M.H., Miraldo, A., et al.: Emerging technologies revolutionise insect ecology and monitoring. *Trends in ecology & evolution* (2022)
42. Koh, P.W., Sagawa, S., Marklund, H., Xie, S.M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R.L., Gao, I., Lee, T., David, E., Stavness, I., Guo, W., Earnshaw, B.A., Haque, I.S., Beery, S., Leskovec, J., Kundaje, A., Pierson, E., Levine, S., Finn, C., Liang, P.: WILDS: A benchmark of in-the-wild distribution shifts. In: *International Conference on Machine Learning (ICML)* (2021)
43. Korsch, D., Bodesheim, P., Denzler, J.: Deep learning pipeline for automated visual moth monitoring: insect localization and species classification. *arXiv preprint arXiv:2307.15427* (2023)
44. Krasin, I., Duerig, T., Alldrin, N., Ferrari, V., Abu-El-Haija, S., Kuznetsova, A., Rom, H., Uijlings, J., Popov, S., Veit, A., et al.: Openimages: A public dataset for large-scale multi-label and multi-class image classification. Dataset available from <https://github.com/openimages> **2**(3), 18 (2017)
45. Kumar, N., Belhumeur, P.N., Biswas, A., Jacobs, D.W., Kress, W.J., Lopez, I.C., Soares, J.V.: Leafsnap: A computer vision system for automatic plant species identification. In: *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part II* 12. pp. 502–516. Springer (2012)
46. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V* 13. pp. 740–755. Springer (2014)
47. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11976–11986 (2022)
48. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: *International Conference on Learning Representations* (2019)
49. Luccioni, A.S., Rolnick, D.: Bugs in the data: How imagenet misrepresents biodiversity. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 37, pp. 14382–14390 (2023)
50. Ming, Y., Yin, H., Li, Y.: On the impact of spurious correlation for out-of-distribution detection. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 36, pp. 10051–10059 (2022)
51. Mora, C., Tittensor, D.P., Adl, S., Simpson, A.G., Worm, B.: How many species are there on earth and in the ocean? *PLoS biology* **9**(8), e1001127 (2011)
52. Nawoya, S., Ssemakula, F., Akol, R., Geissmann, Q., Karstoft, H., Bjerge, K., Mwirize, C., Katumba, A., Gebreyesus, G.: Computer vision and deep learning in insects for food and feed production: A review. *Computers and Electronics in Agriculture* **216**, 108503 (2024)
53. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: *2008 Sixth Indian conference on computer vision, graphics & image processing*. pp. 722–729. IEEE (2008)
54. Norouzzadeh, M.S., Morris, D., Beery, S., Joshi, N., Jovic, N., Clune, J.: A deep active learning system for species identification and counting in camera trap images. *Methods in ecology and evolution* **12**(1), 150–161 (2021)
55. Oliver, R.Y., Iannarilli, F., Ahumada, J., Fegraus, E., Flores, N., Kays, R., Birch, T., Ranipeta, A., Rogan, M.S., Sica, Y.V., et al.: Camera trapping expands the view into global biodiversity and its change. *Philosophical Transactions of the Royal Society B* **378**(1881), 20220232 (2023)

56. Ouaknine, A., Kattenborn, T., Laliberté, E., Rolnick, D.: Openforest: A data catalogue for machine learning in forest monitoring. arXiv preprint arXiv:2311.00277 (2023)
57. Rolnick, D., Donti, P.L., Kaack, L.H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A.S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., et al.: Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)* **55**(2), 1–96 (2022)
58. Roy, D., et al.: Towards a standardised framework for ai-assisted, image-based monitoring of nocturnal insects. *Philosophical Transactions of the Royal Society B* (2024). <https://doi.org/10.1098/rstb.2023.0108>, (Forthcoming)
59. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* **115**(3), 211–252 (2015)
60. Seibold, S., Gossner, M.M., Simons, N.K., Blüthgen, N., Müller, J., Ambarlı, D., Ammer, C., Bauhus, J., Fischer, M., Habel, J.C., et al.: Arthropod decline in grasslands and forests is associated with landscape-level drivers. *Nature* **574**(7780), 671–674 (2019)
61. Stark, T., Ştefan, V., Wurm, M., Spanier, R., Taubenböck, H., Knight, T.M.: Yolo object detection models can locate and classify broad groups of flower-visiting arthropods in images. *Scientific Reports* **13**(1), 16364 (2023)
62. Stork, N.E., et al.: How many species of insects and other terrestrial arthropods are there on earth. *Annual review of entomology* **63**(1), 31–45 (2018)
63. Suto, J.: Codling moth monitoring with camera-equipped automated traps: A review. *Agriculture* **12**(10), 1721 (2022)
64. Swanson, A., Kosmala, M., Lintott, C., Simpson, R., Smith, A., Packer, C.: Snapshot serengeti, high-frequency annotated camera trap images of 40 mammalian species in an african savanna. *Scientific data* **2**(1), 1–14 (2015)
65. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2818–2826 (2016)
66. Teixeira, A.C., Ribeiro, J., Morais, R., Sousa, J.J., Cunha, A.: A systematic review on automatic insect detection using deep learning. *Agriculture* **13**(3), 713 (2023)
67. Theivaprakasham, H.: Identification of indian butterflies using deep convolutional neural network. *Journal of Asia-Pacific Entomology* **24**(1), 329–340 (2021)
68. Tuia, D., Kellenberger, B., Beery, S., Costelloe, B.R., Zuffi, S., Risse, B., Mathis, A., Mathis, M.W., van Langevelde, F., Burghardt, T., et al.: Perspectives in machine learning for wildlife conservation. *Nature communications* **13**(1), 792 (2022)
69. Van Horn, G., Branson, S., Farrell, R., Haber, S., Barry, J., Ipeirotis, P., Perona, P., Belongie, S.: 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*. pp. 595–604 (2015)
70. 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)
71. Van Horn, G., Perona, P.: The devil is in the tails: Fine-grained classification in the wild. arXiv preprint arXiv:1709.01450 (2017)
72. Wagner, D.L.: Insect declines in the anthropocene. *Annual review of entomology* **65**, 457–480 (2020)

73. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset. Tech. rep. (2011)
74. Wei, X.S., Song, Y.Z., Mac Aodha, O., Wu, J., Peng, Y., Tang, J., Yang, J., Belongie, S.: Fine-grained image analysis with deep learning: A survey. *IEEE transactions on pattern analysis and machine intelligence* **44**(12), 8927–8948 (2021)
75. Weinstein, B.G.: A computer vision for animal ecology. *Journal of Animal Ecology* **87**(3), 533–545 (2018)
76. Wightman, R.: Pytorch image models. <https://github.com/rwightman/pytorch-image-models> (2019)
77. Zhou, K., Liu, Z., Qiao, Y., Xiang, T., Loy, C.C.: Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022)