Understanding Insect Range Shifts with Out-of-Distribution Detection

Yuyan Chen
McGill University
Mila – Quebec AI Institute
Montreal, QC, Canada
`yuyan.chen2@mail.mcgill.ca`

David Rolnick
McGill University
Mila – Quebec AI Institute
Montreal, QC, Canada
`drolnick@cs.mcgill.ca`

Abstract

Climate change is inducing significant range shifts in insects and other organisms. Large-scale temporal data on populations and distributions are essential for quantifying the effects of climate change on biodiversity and ecosystem services, providing valuable insights for both conservation and pest management. With images from camera traps, we aim to use Mahalanobis distance-based confidence scores to automatically detect new moth species in a region. We intend to make out-of-distribution detection interpretable by identifying morphological characteristics of different species using Grad-CAM. We hope this algorithm will be a useful tool for entomologists to study range shifts and inform climate change adaptation.

1 Introduction

Climate change results in increased temperature, more extreme weather events, and changes in precipitation, affecting insect biodiversity [1] and causing insect range shifts in many regions [2]. Understanding how insect distribution and biodiversity change is crucial for habitat conservation and pest management in a changing climate. With more than 160,000 species worldwide [3], moths play critical roles in different ecosystems as pollinators, prey, agricultural pests, and indicator species of habitat disturbance. Due to climate change, some regions have observed a decline in moth population that may threaten the local ecosystem [4], while others have observed distribution shifts or increased populations of moths that lead to more challenges of pest management [5].

Machine learning-based methods including object detection and image classification have been applied for automated camera-assisted insect monitoring, providing valuable large-scale data for quantifying moth abundance and biodiversity [6]. To leverage the camera trap images, we propose to use out-of-distribution detection to automatically detect new species [1] to study moth range shifts. We propose to use Mahalanobis distance-based confidence scores [7] to differentiate in-distribution (ID) and out-of-distribution (OOD) species and interpret the result with Grad-CAM [8], which can provide concept-specific visualization for fine-grained classification.

2 Related work

2.1 Insect monitoring

Automated insect monitoring with machine learning has emerged as a promising area. In agriculture, computer vision has been used to automatically identify insect species with trap images to

---

[1] We use "new species" to refer to any species that has not been documented in a specific region. These can be species that are new to science or well-known species that have only recently colonized the region.

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2023.

facilitate pest management [9, 10]. Meanwhile, an increasing number of studies have also started to focus on monitoring insect biodiversity with machine learning [6, 11, 12]. Bjerge et al. suggest a non-lethal method to monitor moth biodiversity with camera trap images [13]. Jain et al. propose a complete machine-learning pipeline from objection detection to fine-grained species-level classification for automated moth monitoring [14]. Although benchmark datasets including iNat2017 [15] and BIOSCAN-1M insect dataset [16] have spurred interest in both biodiversity and fine-grained classification research, long-term data that cover a wide range of regions and species are still rare [2]. As a result, current studies on insect range shifts [2, 17, 18] still mainly rely on species distribution models instead of directly monitoring their distribution with machine learning and computer vision.

2.2 OOD detection

Most prior works [19–22] on OOD detection focus on common benchmark datasets like CIFAR-10[23], SVHN[24], Texture[25], and iSUN[26], where images in different datasets are from unrelated classes. A few works [27, 28] have considered OOD detection for fine-grained classification datasets including CUB-200-2011[29] and Aircraft 300 [30]. A simple logit-based OOD detection has been applied for coarse-grained insect recognition (a dataset of 9 species for training and 11 species for testing) to detect samples that are "unsure" to the classifier [6]. However, to our knowledge, there is no study on OOD detection that focuses on fine-grained classification of wildlife.

3 Proposed methodology

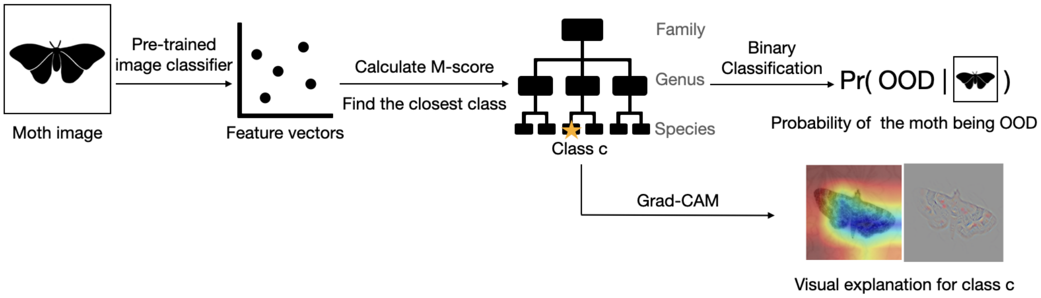

Figure 1: Our proposed workflow (see §3.1 and 3.2).

3.1 Dataset and pre-trained classifier

Since we do not have information on new species that might occur in a region, it is impractical to include OOD samples during training. Hence, we will use post-hoc OOD detection methods which use features extracted by pre-trained classifiers to differentiate ID and OOD samples. Pre-trained classifiers were trained with GBIF, an open-source biodiversity dataset [31]. We use trained classifiers for Quebec-Vermont, the UK-Denmark, and Panama (following methods detailed in [14]). Classes included in the training set were determined by regional species checklists provided by entomologists. For OOD detection training and evaluation, we leverage images from both GBIF and camera traps, using known moth species that are not in the target region to test our OOD detection algorithms.

3.2 OOD detection and visual explanation

We use Mahalanobis distance-based confidence score (M-score) [7] for OOD detection since we can obtain an in-distribution class c that is closest to the out-of-distribution image $\mathbf{x}$. This allows us to use class activation maps to highlight the regions corresponding to the selected class to make OOD detection more interpretable and facilitate the validation of model prediction by entomologists.

Mahalanobis distance measures how many standard deviations a point P is away from the center of a distribution D in high dimension [32]. We first compute the layer-wise empirical class means $\hat{\mu}_{c,\ell}$ and covariance $\hat{\mathbf{\Sigma}}_\ell$ for the training set $\{(\mathbf{x}_1, y_1), \cdots, (\mathbf{x}_n, y_n)\}$:

$$\hat{\mu}_{c,\ell} = \frac{1}{n_c} \sum_{i:y_i=c} f_\ell(\mathbf{x}_i), \; \hat{\mathbf{\Sigma}}_\ell = \frac{1}{n} \sum_c \sum_{i:y_i=c} \left(f_\ell(\mathbf{x}_i) - \hat{\mu}_{c,\ell}\right) \left(f_\ell(\mathbf{x}_i) - \hat{\mu}_{c,\ell}\right)^T , \tag{1}$$

where $f_\ell(\mathbf{x})$ denotes the feature extraced from the ℓ-th layer of the pre-trained classifier f for input $\mathbf{x}$ and n_c denotes the number of samples in class c. We then compute the M-score from the ℓ-th layer of the pre-trained classifier as

$$M_\ell(\mathbf{x}) = \max_c \left[-\left(f_\ell(\mathbf{x}) - \hat{\mu}_{c,\ell}\right)^T \hat{\mathbf{\Sigma}}^{-1} \left(f_\ell(\mathbf{x}) - \hat{\mu}_{c,\ell}\right) \right] , \tag{2}$$

where $\mathbf{x}$ is the input vector. The advantage of using the M-score over other OOD detection methods is that it computes the closest class to a given sample, which allows us to interpret the result later. To utilize both low-level and high-level features, we combine confidence scores from different layers of the network and train a binary classification model using these scores as features.

After obtaining the closest class c, we use Grad-CAM [8] to visualize the result. Grad-CAM can compute a class-discriminative localization map for any class c, making it suitable to be combined with M-score to make OOD detection more interpretable. Grad-CAM can be further improved by multiplying the localization map and guided backpropagation for class c to output concept-specific visualization for fine-grained classification.

4 Preliminary results

The proposed OOD detection method was tested with a ResNet50 model trained with 2,530 species in the UK and Denmark. All training and testing images are from GBIF [31]. ID samples are from the validation set of the pre-trained classifier and consist of 125,230 images.

4.1 OOD detection

For a preliminary test, we constructed the OOD samples by randomly choosing six moth families and filtering out all species included in the training set of the pre-trained classifier, which resulted in a total of 444 remaining species and 142,613 images. We calculated five M-scores for each image based on features extracted from five layers of ResNet50 and trained a Random Forest classifier. We trained the classifier with 80% of the ID samples and OOD samples from five families. Then, we evaluated the model using the previously excluded family and the rest 20% of the ID samples. We evaluated the model's performance with the Area Under the Receiver Operating Characteristic (AUROC), the Area Under the Precision-Recall Curve (AUPR), and the False Positive Rate when the true positive rate is at 95% (FPR95), as they are commonly used in OOD detection literature.

Table 1: OOD detection performance comparison using test sets with different OOD families

Family	Adelidae	Apatelodidae	Crambidae	Erebidae	Oecophoridae	Sphingidae
AUROC	92.94	94.09	92.60	92.43	92.60	91.34
AUPR	18.40	53.07	95.44	96.98	58.72	88.30
FPR95	25.92	23.47	32.16	31.99	32.68	34.15

4.2 Visualization

We show the result of the visualization method with an OOD species *Nemophora bellella* in Fig. 2. We use c_i to denote the closest class to the input image based on features extracted from layer i of the ResNet50 model. Based on the features extracted from layer 2 and layer 3 of our pre-trained classifier, *Paraswammerdamia albicapitella* and *Stigmella centifoliella* are considered the closest class c_2 and c_3 to this species by our OOD detection algorithm. For layer 2, Grad-CAM and Guided Grad-CAM highlighted the white stripe, and *P. albicapitella* also has a white part, although it is not on the wing. For layer 3, Grad-CAM and Guided Grad-CAM highlighted more high-level features including its head, thorax, and the white stripe, and *S. centifoliella* also has a yellow head and a white stripe.

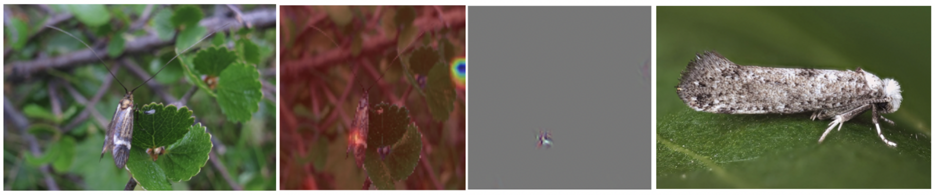

(a) Input image (b) Grad-CAM Layer 2 (c) Guided Grad-CAM Layer 2 (d) The closest species c_2 in the training set

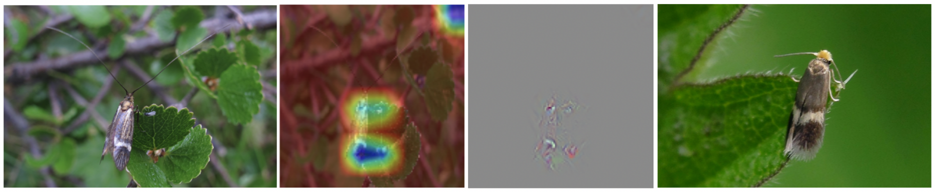

(e) Input image (f) Grad-CAM Layer 3 (g) Guided Grad-CAM Layer 3 (h) The closest species c_3 in the training set

Figure 2: (a,e) The input image that contains an OOD species *N. bellella*. (b,c) Grad-CAM and Guided Grad-CAM visualization highlighting support for the closest class c_2 based on features from layer 2 of ResNet50. (d) An image of *P. albicapitella*. (f,g) Grad-CAM and Guided Grad-CAM visualization highlighting support for the closest class c_3 based on features from layer 3 of ResNet50. (h) An image of *S. centifoliella*. In (b, f), regions with colder colors (i.e. blue) correspond to a higher score for the given class.

5 Future work and pathway to impact

We will further test our method on the camera trap images to see if the OOD detection method can generalize well to images from different sources. We will also investigate the features extracted by the pre-trained classifier to understand why the performance of the method is different when tested on different families. Continued collaboration with entomologists is key for this research to be impactful. Entomologists can decide the ideal location for camera traps for us to obtain high-quality training and validation images. We also need to rely on their domain expertise to validate the results when the method is deployed in real life for new species detection – the goal is not to replace experts but to help them filter data, and the visual explanation generated by Grad-CAM and Guided Grad-CAM can facilitate the validation process. We expect our research to provide long-term large-scale data for entomologists to better study insect range shifts and accordingly inform climate adaptation measures.

Acknowledgement

This research was enabled in part by computing resources and technical help provided by Mila - Quebec AI Institute.

References

[1] David L Wagner. Insect declines in the anthropocene. *Annual review of entomology*, 65:457–480, 2020.

[2] Laura H Antão, Benjamin Weigel, Giovanni Strona, Maria Hällfors, Elina Kaarlejärvi, Tad Dallas, Øystein H Opedal, Janne Heliölä, Heikki Henttonen, Otso Huitu, et al. Climate change reshuffles northern species within their niches. *Nature Climate Change*, 12(6):587–592, 2022.

[3] David Carter. *Butterflies and moths*. Penguin, 2023.

[4] Richard Fox. The decline of moths in Great Britain: a review of possible causes. *Insect conservation and diversity*, 6(1):5–19, 2013.

[5] Sandra Skendžić, Monika Zovko, Ivana Pajač Živković, Vinko Lešić, and Darija Lemić. The impact of climate change on agricultural insect pests. *Insects*, 12(5):440, 2021.

[6] Kim Bjerge, Quentin Geissmann, Jamie Alison, Hjalte MR Mann, Toke T Høye, Mads Dyrmann, and Henrik Karstoft. Hierarchical classification of insects with multitask learning and anomaly detection. *Ecological Informatics*, page 102278, 2023.

[7] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31, 2018.

[8] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.

[9] Weiguang Ding and Graham Taylor. Automatic moth detection from trap images for pest management. *Computers and Electronics in Agriculture*, 123:17–28, 2016.

[10] Querriel Arvy Mendoza, Lester Pordesimo, Mitchell Neilsen, Paul Armstrong, James Campbell, and Princess Tiffany Mendoza. Application of machine learning for insect monitoring in grain facilities. *AI*, 4(1):348–360, 2023.

[11] Brian J Spiesman, Claudio Gratton, Richard G Hatfield, William H Hsu, Sarina Jepsen, Brian McCornack, Krushi Patel, and Guanghui Wang. Assessing the potential for deep learning and computer vision to identify bumble bee species from images. *Scientific reports*, 11(1):7580, 2021.

[12] Dimitri Korsch, Paul Bodesheim, and Joachim Denzler. Deep learning pipeline for automated visual moth monitoring: insect localization and species classification. *arXiv preprint arXiv:2307.15427*, 2023.

[13] Kim Bjerge, Jakob Bonde Nielsen, Martin Videbæk Sepstrup, Flemming Helsing-Nielsen, and Toke Thomas Høye. An automated light trap to monitor moths (lepidoptera) using computer vision-based tracking and deep learning. *Sensors*, 21(2):343, 2021.

[14] Aditya Jain, Fagner Cunha, Michael Bunsen, Léonard Pasi, Anna Viklund, Maxim Larrivée, and David Rolnick. A machine learning pipeline for automated insect monitoring. In *NeurIPS 2023 Workshop on Tackling Climate Change with Machine Learning*, 2023.

[15] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The iNaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8769–8778, 2018.

[16] Z. Gharaee, Z. Gong, N. Pellegrino, I. Zarubiieva, J. B. Haurum, S. C. Lowe, J. T. A. McKeown, C. Y. Ho, J. McLeod, Y. C. Wei, J. Agda, S. Ratnasingham, D. Steinke, A. X. Chang, G. W. Taylor, and P. Fieguth. A step towards worldwide biodiversity assessment: The BIOSCAN-1M insect dataset. In *Advances in Neural Information Processing Systems (NeurIPS) Datasets & Benchmarks Track*, 2023.

[17] Henri Vanhanen, Timo O Veteli, Sonia Paivinen, Seppo Kellomaki, and Pekka Niemela. Climate change and range shifts in two insect defoliators: gypsy moth and nun moth-a model study. *Silva Fennica*, 41(4):621, 2007.

[18] Juha Pöyry, Miska Luoto, Risto K Heikkinen, Mikko Kuussaari, and Kimmo Saarinen. Species traits explain recent range shifts of finnish butterflies. *Global Change Biology*, 15(3):732–743, 2009.

[19] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.

[20] Yen-Chang Hsu, Yilin Shen, Hongxia Jin, and Zsolt Kira. Generalized odin: Detecting out-of-distribution image without learning from out-of-distribution data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10951–10960, 2020.

[21] Weitang Liu, Xiaoyun Wang, John Owens, and Yixuan Li. Energy-based out-of-distribution detection. *Advances in neural information processing systems*, 33:21464–21475, 2020.

[22] Haoqi Wang, Zhizhong Li, Litong Feng, and Wayne Zhang. Vim: Out-of-distribution with virtual-logit matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4921–4930, 2022.

[23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

[24] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.

[25] Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014.

[26] Pingmei Xu, Krista A Ehinger, Yinda Zhang, Adam Finkelstein, Sanjeev R Kulkarni, and Jianxiong Xiao. Turkergaze: Crowdsourcing saliency with webcam based eye tracking. *arXiv preprint arXiv:1504.06755*, 2015.

[27] Guangyao Chen, Peixi Peng, Xiangqian Wang, and Yonghong Tian. Adversarial reciprocal points learning for open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):8065–8081, 2021.

[28] Sagar Vaze, Kai Han, Andrea Vedaldi, and Andrew Zisserman. Open-set recognition: A good closed-set classifier is all you need? *arXiv preprint arXiv:2110.06207*, 2021.

[29] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.

[30] Guangyao Chen, Limeng Qiao, Yemin Shi, Peixi Peng, Jia Li, Tiejun Huang, Shiliang Pu, and Yonghong Tian. Learning open set network with discriminative reciprocal points. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 507–522. Springer, 2020.

[31] GBIF.org. (1 February 2023) GBIF Occurence Download `http://doi.org/10.15468/dl.7zxw9y`.

[32] Prasanta Chandra Mahalanobis. On the generalized distance in statistics. *Sankhyā: The Indian Journal of Statistics, Series A (2008-)*, 80:S1–S7, 2018.