Peer Community In Ecology

Democratising AI for Ecology: MEWC Delivers Custom Wildlife Classification to All

Cédric Sueur based on peer reviews by *Timm Haucke* and 1 anonymous reviewer

Barry W. Brook, Jessie C. Buettel, Peter van Lunteren, Prakash P. Rajmohan, R. Zach Aandahl (2025) MEWC: A user-friendly AI workflow for customised wildlife-image classification. EcoEvoRxiv, ver. 3, peer-reviewed and recommended by Peer Community in Ecology. https://doi.org/10.32942/X2ZW3D

Submitted: 24 February 2025, Recommended: 16 May 2025

Cite this recommendation as:

Sueur, C. (2025) Democratising Al for Ecology: MEWC Delivers Custom Wildlife Classification to All. *Peer Community in Ecology*, 100787. https://doi.org/10.24072/pci.ecology.100787

Published: 16 May 2025

Copyright: This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit https://creativecommons.org/licenses/by/4.0/

Although artificial intelligence (AI) offers a powerful solution to the bottleneck in processing camera-trap imagery, ecological researchers have long struggled with the practical application of such tools. The complexity of development, training, and deploying deep learning models remains a significant barrier for many conservationists, ecologists, and citizen scientists who lack formal training in computer science. While platforms like Wildlife Insights (Ahumada et al. 2020), MegaDetector (Beery et al. 2019), and tools such as ClassifyMe (Falzon et al. 2020) have laid critical groundwork in AI-assisted wildlife monitoring, these solutions either remain opaque, lack customisability, or are often locked behind commercial or infrastructural limitations. Others, like Sherlock (Penn et al., 2024), while powerful, are not always deployable without significant local expertise. A notable example of a more open and collaborative approach is the DeepFaune initiative, which provides a free, high-performance tool for the automatic classification of European wildlife in camera-trap images, highlighting the growing importance of locally relevant, user-friendly AI solutions developed through broad partnerships (Rigoudy et al. 2022).

It is in this context that Brook et al. (2025) makes a compelling and timely contribution. The authors present an elegant, open-source AI pipeline—MEWC (Mega-Efficient Wildlife Classifier)—that bridges the gap between high-performance image classification and usability by non-specialists. Combining deep learning advances with user-friendly software engineering, MEWC enables users to detect, classify, and manage wildlife imagery without advanced coding skills or reliance on high-cost third-party infrastructures.

What makes MEWC particularly impactful is its modular and accessible architecture. Built on Docker containers, the system can run seamlessly across different operating systems, cloud services, and local machines. From

image detection using MegaDetector to classifier training via EfficientNet or Vision Transformers, the pipeline maintains a careful balance between technical flexibility and operational simplicity. This design empowers ecologists to train their own species-specific classifiers and maintain full control over their data—an essential feature given the increasing scrutiny around data sovereignty and privacy.

The practical implications are impressive. The case study provided—focused on Tasmanian wildlife—demonstrates not only high accuracy (up to 99.6%) but also remarkable scalability, with models trainable even on mid-range desktops. Integration with community tools like Camelot (Hendry and Mann 2018) and AddaxAI (Lunteren 2023) further enhances its utility, allowing rapid expert validation and facilitating downstream analyses.

Yet the article does not shy away from discussing the limitations. As with any supervised system, MEWC's performance is only as good as the training data provided. Class imbalances, rare species, or subtle morphological traits can challenge even the best classifiers. Moreover, the authors caution that pre-trained models may not generalise well across regions with different fauna, requiring careful curation and expert tagging for local deployments.

One particularly exciting future direction briefly mentioned—and worth highlighting—is MEWC's potential application to behavioural and cognitive ecology (Sueur et al. 2013; Battesti et al. 2015; Grampp et al. 2019). Studies in these domains underscore the need for scalable tools to quantify social dynamics in real time. By assisting with individual identification and the detection of postures or spatial configurations, MEWC could significantly enhance the throughput, reproducibility, and objectivity of such research.

This opens the door to even richer applications. Behavioural ecologists might use MEWC for fine-grained detection tasks such as individual grooming interactions, kin proximity analysis, or identification of tool-use sequences in wild primates. Similarly, for within-species classification (e.g. sex, reproductive state, or disease symptoms), MEWC's modular backbone and compatibility with transfer learning architectures like EfficientNet or ViT make it a suitable candidate for expansion (Ferreira et al. 2020; Clapham et al. 2022).

In conclusion (Brook et al. 2025) have delivered more than a tool—they've designed an ecosystem. MEWC lowers the technical barrier to AI in ecology, promotes open science, and enables tailored workflows for a wide variety of conservation, research, and educational contexts. For anyone interested in democratising ecological AI and reclaiming control over wildlife-monitoring data, this article and its associated software are essential resources.

References:

Ahumada JA, Fegraus E, Birch T, et al (2020) Wildlife Insights: A Platform to Maximize the Potential of Camera Trap and Other Passive Sensor Wildlife Data for the Planet. Environmental Conservation 47:1–6. https://doi.org/10.1017/S0376892919000298

Battesti M, Pasquaretta C, Moreno C, et al (2015) Ecology of information: social transmission dynamics within groups of non-social insects. Proceedings of the Royal Society of London B: Biological Sciences 282:20142480. https://doi.org/10.1098/rspb.2014.2480

Beery S, Morris D, Yang S, et al (2019) Efficient pipeline for automating species ID in new camera trap projects. Biodiversity Information Science and Standards 3:e37222. https://doi.org/10.3897/biss.3.37222

Barry W. Brook, Jessie C. Buettel, Peter van Lunteren, Prakash P. Rajmohan, R. Zach Aandahl (2025) MEWC: A user-friendly AI workflow for customised wildlife-image classification. EcoEvoRxiv, ver.3 peer-reviewed and recommended by PCI Ecology https://doi.org/10.32942/X2ZW3D

Clapham M, Miller E, Nguyen M, Van Horn RC (2022) Multispecies facial detection for individual identification of wildlife: a case study across ursids. Mamm Biol 102:921–933. https://doi.org/10.1007/s42991-021-00168-5 Falzon G, Lawson C, Cheung K-W, et al (2020) ClassifyMe: A Field-Scouting Software for the Identification of Wildlife in Camera Trap Images. Animals 10:58. https://doi.org/10.3390/ani10010058

Ferreira AC, Silva LR, Renna F, et al (2020) Deep learning-based methods for individual recognition in small birds. Methods in Ecology and Evolution 11:1072–1085. https://doi.org/10.1111/2041-210X.13436

Grampp M, Sueur C, van de Waal E, Botting J (2019) Social attention biases in juvenile wild vervet monkeys: implications for socialisation and social learning processes. Primates 60:261–275 https://doi.org/10.1007/s10329-019-00721-4

Hendry H, Mann C (2018) Camelot—intuitive software for camera-trap data management. Oryx 52:15–15. https://doi.org/10.1017/S0030605317001818

Lunteren P van (2023) AddaxAI: A no-code platform to train and deploy custom YOLOv5 object detection models. Journal of Open Source Software 8:5581. https://doi.org/10.21105/joss.05581

Rigoudy N, Dussert G, Benyoub A, et al (2022) The DeepFaune initiative: a collaborative effort towards the automatic identification of the French fauna in camera-trap images. bioRxiv 2022–03

Sueur C, MacIntosh AJJ, Jacobs AT, et al (2013) Predicting leadership using nutrient requirements and dominance rank of group members. Behav Ecol Sociobiol 67:457–470. https://doi.org/10.1007/s00265-012-1466-5

Reviews

Evaluation round #1

DOI or URL of the preprint: https://doi.org/10.32942/X2ZW3D Version of the preprint: 2

Authors' reply, 16 May 2025

Download author's reply Download tracked changes file

Decision by Cédric Sueur ^(D), posted 03 April 2025, validated 03 April 2025

The manuscript presents a highly valuable and well-executed contribution to the conservation technology field through the MEWC system. Both reviewers express strong support for the paper, noting its clarity, relevance, and practical implementation. The workflow is well-justified and the case study is convincing.

However, several clarifications and improvements are requested: 1. Discussion of MEWC's reliance on MegaDetector and its limitations. 2. Clarification on confidence thresholds, blank images, and classifier behavior. 3. Clarification of the "Train Base" vs. "Train Final" distinction. 4. Minor corrections to terminology (e.g., "segmentation" vs. "object detection") and copy edits.

The reviews do not raise concerns that question the core validity or significance of the work, but addressing the detailed points raised will strengthen clarity and usability for a broader audience. Thus, minor revisions are recommended.

Reviewed by anonymous reviewer 1, 09 March 2025

This submission describes the MEWC system for training and running custom species classifiers for camera trap images. The paper is an excellent overview of the authors' motivation for developing MEWC, and an excellent high-level overview of the architecture.

Because this submission is about a system, not about an experiment or specific innovation, my review of the submission is as much about the system as it is about the paper itself. The MEWC repos are among the most mature and well-documented OSS repos in the conservation technology space; I was able to access and run the relevant tools. Exactly how easy they are to use is outside the scope of this review, but the repos are complete and purposeful, and provide access to everything described in this paper. The authors appear to have made wise decisions about integration with other tools (Camelot, AddaxAI), allowing this repo to be part of a larger OSS ecosystem, rather than just an island.

So my review of the submission - encompassing the paper and the repos - is quite positive. I'm not 100% sure who the audience is for the paper itself: much of the technical detail will be uninformative for ecologists, e.g. it seems neither here nor there for *users* that the authors used Terraform and Ansible, and for potential *contributors*, this would be immediately clear (and more up to date) from the repo itself. But that's not really a criticism of the submission, just a note that the repo is a more significant contribution than the paper.

If the authors were going to make any significant edits to the paper, my top request would be clarifying the difference between the "Train Base" and "Train Final" steps; I'm not clear on the difference. If I train just once (which seems like a common scenario), it would seem that my base model *is* my final model. If I train, then fine-tune, then fine-tune again, was the second training base "base" or "final"? This seems like a confusing distinction, unless there is something fundamentally different between these stages that I'm not seeing. Is there is, consider clarifying. If there's not, consider dropping this distinction from the paper, replacing this with a single "train" stage, and mentioning in the text that the (one and only) "train" stage can be repeated as many times as the user wants to repeat it.

The remainder of this review will be minor copy edits that did not impact my overall review.

_

Inconveniently for the authors, EcoAssist was renamed to "AddaxAI", likely just *after* this draft was submitted. I think it's worth making this change throughout the paper; it will be confusing to readers if they search for "EcoAssist" in the future.

Boxes are hard to see in Figure 3; consider re-rendering.

-

"...which leverages deep learning for wildlife-image classification"

I would recommend against hyphenating "wildlife image" here.

-

"It can be run using simple command-line prompts or via a user-friendly Graphical User interface"

I would recommend against capitalizing "Graphical User", but if you do, I recommend capitalizing "Interface" as well.

-

"An increasingly used tool for this purpose is the 'camera trap'."

If you put "camera trap" in quotes, I would immediately define it, e.g.:

"An increasingly used tool for this purpose is the 'camera trap', a camera designed to be triggered based on motion or time that can be deployed for long periods in harsh environments."

If you want to assume that the reader knows what a camera trap is, IMO that's also fine, but in that case I wouldn't put it in quotes.

"In this scenario data management, rather than data collection becomes the limiting factor in the completion of research projects"

Missing comma after "collection".

"...is the commercial website 'Wildlife Insights', sponsored by Google"

Wildlife Insights is not a commercial product; it's run by a group of non-profits, and it's free for the vast majority of users. The Google sponsorship feels like an odd/informal thing to mention here, given that sponsorship was not mentioned for any other projects.

"Leverage cutting-edge developments in computer vision, but leave the details behind-the-scenes..." "behind the scenes" should not be hyphenated here.

"Make the workflow easy for non-specialists to use, completely reproducible, and yet ensure that it is powerfully flexible for expert-level fine-tuning or expansion."

There is a grammar issue here, perhaps replace with:

"Make the workflow easy for non-specialists to use, completely reproducible, and flexible for expert-level fine-tuning or expansion."

-

"This overturns the perception by many ecologists as it being an 'arcane art' needing specialist data-science and programming skills to implement"

There is a grammar issue here, perhaps replace with:

"This overturns the perception by many ecologists that it is an 'arcane art' needing specialist data-science and programming skills to implement"

Reviewed by Timm Haucke, 02 April 2025

The paper proposes a user-friendly workflow (MEWC) for training custom camera trap image classifiers and running inference on those models. It relies on a general animal detection model (MegaDetector) to detect and localize animals in camera trap imagery. This localization is then used to crop images to the area of interest, which are subsequently classified by fine-tuned classification models.

Overall, I really like this paper. I think that MEWC fills an important gap in the camera trapping ecosystem, especially since it provised a mechanism for users to fine-tune their own, custom models instead of just running inference on a provided, fixed model. The case study is convincing and the drawn conclusions are overall adequately supported by the results.

However, I think there are some points that could be clarified:

- If I understand the paper correctly, the detection model (MegaDetector) is not fine-tuned. This means that MEWC is only feasible for scenarios that MegaDetector works well in (so, for example, it might not work well with camera traps that are setup to capture images from a top-down view). I think this is an important limitation that should be discussed more thoroughly, even if MegaDetector generalized well in many cases.
- 2. Related to point 1, how is the MegaDetector confidence threshold chosen? Are users able to customize it?
- 3. Related to point 2, how are blanks that slip through MegaDetector handled? Is this an additional class in the classifier? This is related to points 1 & 2 in that, depending on how well MegaDetector performs in a given scenario, you might need to reduce the confidence threshold so much that a lot of blank images land above that threshold and are processed by the classifier.

- 4. This paper [1] might provide some additional motivation for cropping before classifying.
- 5. The paper refers to MegaDetector as a "segmentation" model. In a computer vision context, this could be confused with actual pixel-wise segmentation models. Instead, the right term is "object detection" model.
- 6. Machine learning architectures are constantly evolving. How easy is it to add new architectures to MEWC? Does one only have to specify the corresponding kimm implementation?
- 7. In Table 2, I find the test accuracy to be hard to interpret, especially since per-class accuracy might vary strongly between different classes (even if the dataset is class-balanced). Instead, it might be worth including micro- and macro-averaged precision and recall.
- 8. On page 19 the paper says: "typical camera study amassing 100-500 thousand (K) images monthly". I am sure there are studies amassing this amount of images, but I am not sure those are "typical" studies.
- 9. Is Docker an inherent part of MEWC or is MEWC compatible with other Open Container Initiative (OCI) runtimes? There has recently been a shift from Docker to other runtimes. Furthermore some supercomputing environments only support specific runtimes like Singularity.

Review checklist:

- Title and abstract
 - Does the title clearly reflect the content of the article? [x] Yes, [] No (please explain), [] I don't know
 - Does the abstract present the main findings of the study? [x] Yes, [] No (please explain), [] I don't know
- Introduction
 - Are the research questions/hypotheses/predictions clearly presented? [x] Yes, [] No (please explain),
 [] I don't know
 - Does the introduction build on relevant research in the field? [x] Yes, [] No (please explain), [] I don't know
- Materials and methods
 - Are the methods and analyses sufficiently detailed to allow replication by other researchers? [] Yes,
 [x] No (please explain), [] I don't know
 - Are the methods and statistical analyses appropriate and well described? [x] Yes, [] No (please explain), [] I don't know
- Results
 - In the case of negative results, is there a statistical power analysis (or an adequate Bayesian analysis or equivalence testing)? [] Yes, [] No (please explain), [] I don't know
 - Are the results described and interpreted correctly? [x] Yes, [] No (please explain), [] I don't know
- Discussion
 - Have the authors appropriately emphasized the strengths and limitations of their study/theory/methods/argument? [x] Yes, [] No (please explain), [] I don't know
 - Are the conclusions adequately supported by the results (without overstating the implications of the findings)? [x] Yes, [] No (please explain), [] I don't know