

Echo-dash: Keeping ecologists in the loop with an open source, online ecoacoustic dashboard for interactive exploration of spatiotemporal soundscape data.

Dr Ivor Simpson^{1*}, Kieran Gibb¹, Dr David Kadish², Lukas Gunthermann¹, Dr Chris Andrews³, Dr Jonathan Carruthers-Jones⁴, Dr Jan Dick², Dr Patrice Guyot⁵, Paul Petry⁴, Dr Christopher J Sandom⁶, Dr Nicolas Seymour Smith⁷, Dr James Whitehead⁶, Prof Alice Eldridge⁸

¹ Department of Informatics, University of Sussex, Brighton, UK *I.Simpson@sussex.ac.uk

² School of Arts and Communication, University of Malmo, Sweden

³ UK Centre for Ecology & Hydrology, Edinburgh, Edinburgh, United Kingdom

⁴ Faculty of Agriculture and Forestry, University of Helsinki, Helsinki, Finland.

⁵ EuroMov Digital Health in Motion, University of Montpellier, IMT Mines Ales, Ales, France.

⁶ Department of Evolution and Ecology, University of Sussex, Brighton, UK

⁷ Sussex Digital Humanities Lab, Department of Informatics, University of Sussex, Brighton, UK

⁸ Department of Music, School for Progressive Futures, University of Sussex, Brighton, UK

Abstract

Passive acoustic monitoring (PAM) is being adopted in a range of contexts. Emerging methods facilitate analysis of large-scale data sets, but ecological interpretation of acoustic indices is not straightforward. In addition, the technical and logistical requirements of using emerging AI methods for big data mean that conservation actors increasingly adopt third-party analysis solutions. We argue that these compounding factors undermine robust ecological inference, clouding insight and decision making. To address this, we present *echo-dash*, an accessible, interactive dashboard that facilitates rapid, interactive exploration of spatiotemporal soundscape data by conservation actors. Developed through participatory design, *echo-dash* is built on the simple premise that the potential of PAM can best be realised by keeping human ecological knowledge in the analysis loop. Five key functions to facilitate analysis and interpretation of PAM data were identified and implemented: 1) Calculating soundscape descriptors and probability of species presence; 2) Checking data integrity; 3) Exploring data interactively and in spatiotemporal, environmental contexts; 4) Filtering data by extreme weather, outliers or clusters; 5) Exporting data subsets, plots and code to generate plots. By supporting integration of human, situated ecological knowledge with large scale spatiotemporal data sets, *echo-dash* bolsters PAM's potential to transform ecological monitoring for applied and fundamental ecoacoustics.

1 Introduction

Passive acoustic monitoring (PAM) is being adopted in a range of contexts. Emerging methods facilitate analysis of large-scale data sets, but ecological interpretation of results is not straightforward. In addition, the technical and logistical requirements of using AI methods for big data mean that conservation actors increasingly adopt third-party analysis solutions. We argue that these compounding factors undermine robust ecological inference, clouding insight and decision making, with the potential to reduce confidence in PAM and hamper scientific advances. To address this, we present *echo-dash*, an accessible, interactive dashboard that facilitates rapid, interactive exploration of spatiotemporal soundscape data by conservation actors. Developed through participatory design, *echo-dash* is built on the simple premise that the potential of PAM can best be realised by keeping human ecological knowledge in the analysis loop.

Issues around the interpretability of PAM analysis tools are well recognised. Soundscape indices are notoriously ambiguous: their relationship with classical biodiversity metrics is questioned (Alcocer et al, 2022) and careful interpretation is required (Bradfer-Lawrence, 2023). Emerging AI models promise automated species identification (Kahl et al 2021) and generalised soundscape descriptors (Sethi et al, 2020) but these are not trustworthy due to data biases, potential for

shortcut learning and uninterpretability (Gibb et al, 2024). As illustrated in Figure 1, issues of uninterpretable AI are compounded because inferences about the *ecological community* are made from analyses of a whole soundscape recording which invariably contains signals other than those emanating from the *acoustic community*. All analysis methods are sensitive to confounds from geophony, technophony as well as device malfunction. When analysis and interrogation of soundscape descriptors or automated species ID are carried out by someone familiar with the ecosystem under study, inference can be supported by ecological knowledge of the site (e.g. habitat, landscape or weather variation during survey, or ethological knowledge of acoustic community). However, the sheer volume of data collected under long term PAM regimes and increasing prevalence of machine-learning based indices create logistical and technical barriers that mean analyses are increasingly being carried out remotely. For example in large programmes such as eLTER¹, the proposed future protocol involves central processing of data with summary results returned to study sites; similarly, conservation technology start-ups (e.g. birdweather², soil acoustics³ etc.) and online services (arbimon⁴) are emerging to meet the needs of private landowners and community projects who lack capacity to store and analyse data. Whilst these solutions are pragmatic, remote analysis of remotely-sensed data without trustworthy analysis algorithms may lead to incorrect ecological inferences and risks undermining reliability of and confidence in PAM for both applied and fundamental ecoacoustics.

2 Design and Implementation

Echo-dash has been co-created through participatory design with a representative range of conservation and research actors to illustrate and provide an alternative. Participatory design (Robertson et al, 2012) is a co-design process that centres users, with the aim of ensuring their needs are met, leading to investment, longevity and productivity. *Echo-dash* was developed through four case studies, each addressing distinct ecological questions working with ecologists with differing degrees of technical expertise: 1) The UK Environmental Change Network site in the **Cairngorms National Park** has used Ecoacoustics to study birds and bats since 2010. More recently acoustic monitoring has been adopted to track changes in the soundscapes of regenerating woodland relative to a moorland area; 2) **Kilpisjärvi Biological Station** supports socio-ecological research in Northern Finland. Acoustic monitoring is adopted to explore interactions between anthropogenic sounds (snow mobiles, tourism) and biological sounds (bird life and other fauna); 3) **Nature Sense** is a nature-monitoring service in south-east England exploring acoustic monitoring alongside other surveys to support evidence-based decision making for nature recovery and food production; 4) The **Sounding Out** project surveyed gradients of human degradation in UK and Ecuador to explore how predictive acoustic indices were of species richness and habitat status..

Through a series of participatory workshops five key functions were identified to facilitate analysis and interpretation of PAM data. Illustrated in Figure 2 these are: 1) **Calculate** soundscape descriptors (acoustic indices and probability of species occurrence); 2) **Check** data integrity; 3) **Explore** data interactively and in spatiotemporal, environmental contexts; 4) **Filter** data by extreme weather, outliers or clusters; 5) **Export** data subsets, plots and code to generate plots. These features differ from the proliferation of services and software that offer species identification, index calculation or clustering by explicitly providing environmental context and supporting active inference by data owners enabling integration of local ecological knowledge. Results can be output directly to reports or statistically analysed locally; data ownership is retained by originators.

¹ <https://elter-ri.eu/> ²<https://www.birdweather.com/>, ³<https://soilacoustics.com/>, ⁴<https://arbimon.org/>

As illustrated in Figure 2 (top), echo-dash² comprises an HPC-ready processing pipeline and web-based front-end. Audio data is processed through a custom library soundADE³, which supports parallel computation of acoustic indices, (pre)processing of audio data according to a user-customisable schema and BirdNET (Kahl et al, 2021) analysis with an interactively-adjustable confidence threshold. An optimised UMAP (McInnes et al, 2018) is precomputed. Further contextual information is provided by weather data (Zippenfenig, 2023) linked to recording times/location, and maps via MapLibre. The front-end (Figure 2 bottom) is built using Dash and Plotly. It is easily configured to run on a server and facilitates playback of user-selected audio.

3 Insights, Discussion and Conclusion

We observed that echo-dash facilitates users to intuitively interrogate their data in depth. The interactive interface, with flexible views of spatial, temporal and weather data, provides key environmental context, filtering and faceting. Together these afford deep insight into the source of variation in soundscape descriptors or patterns of species occurrence, enhancing understanding and supporting ecological inference for users of all technical aptitudes.

ECN site managers from the **Cairngorms** reported that echo-dash provided transformative accessibility to their data: *“this is an incredibly powerful tool for long-term monitoring site managers. It will save us hours of work and allow us to integrate the data quickly whilst keeping ownership of the data”*. For Nature Sense, echo-dash enabled ecologists to visualise trends within county-wide datasets. This expedited deeper understanding of the ecological patterns at study sites than would otherwise be possible. For the **Kilpisjarvi** dataset rapid spatiotemporal visualisation was valuable for checking sensor failure and data coverage; weather-based filtering proved particularly valuable in this remote region: *“Combined, this allows us to clearly identify and explore the more advanced visualisations for filtered time windows where we have full spatial and temporal data coverage with characteristic weather patterns”*. Even when exploring the previously-analysed **Sounding Out** dataset, new insights were reached by aggregating data in differing spatial and temporal views (see Figure 3, top).

By providing an intuitive, interactive interface, echo-dash enables users to rapidly explore long term spatiotemporal PAM data. Insights from case studies suggest that this supports a form of rapid hypothesis generation and testing which bolsters ecological inference. Enabling visualisation of spatiotemporal variation in soundscape data helps interpret results and drive understanding. Filtering by extreme weather events builds confidence in results.

4 Conclusion and Future Work.

There is growing interest in PAM for long term environmental monitoring, but analysis and interpretation are not straightforward. We believe that we have an ethical responsibility to ensure that evidence-based decision-making is robust. Echo-dash democratises analysis, enabling data originators to easily stay in the analysis loop, such that they can more directly examine the evidence basis for inference, ask questions and highlight potential issues, which is less viable when analysis is outsourced. This connection fosters a greater understanding of the benefits and limits of PAM in those that will use the results. Echo-dash is designed in a modular fashion; future work will support next generation soundscape descriptors, including self-supervised soundscape

² <https://github.com/ecolistening/echo-dash>

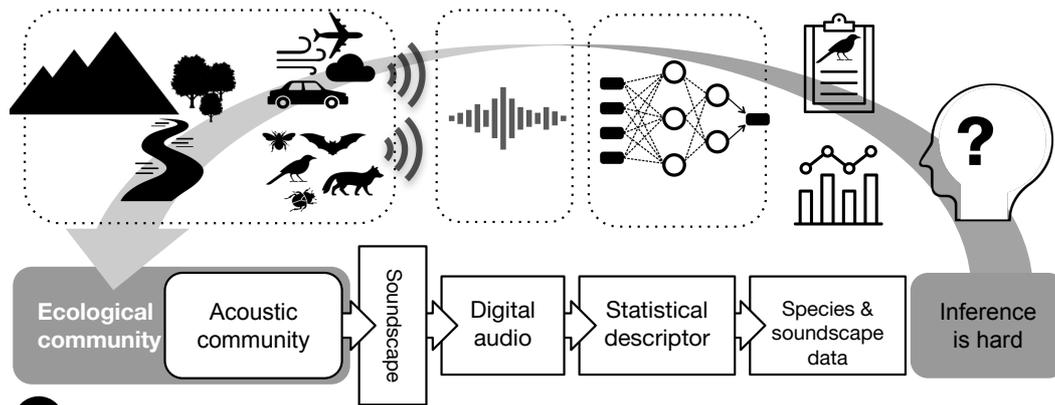
³ <https://github.com/ecolistening/SoundADE>

models (Gibb et al 2024), wider species data and multimodal data (e.g. camera trap and eDNA) as well as providing more expressive spatiotemporal visualisations and analysis tools.

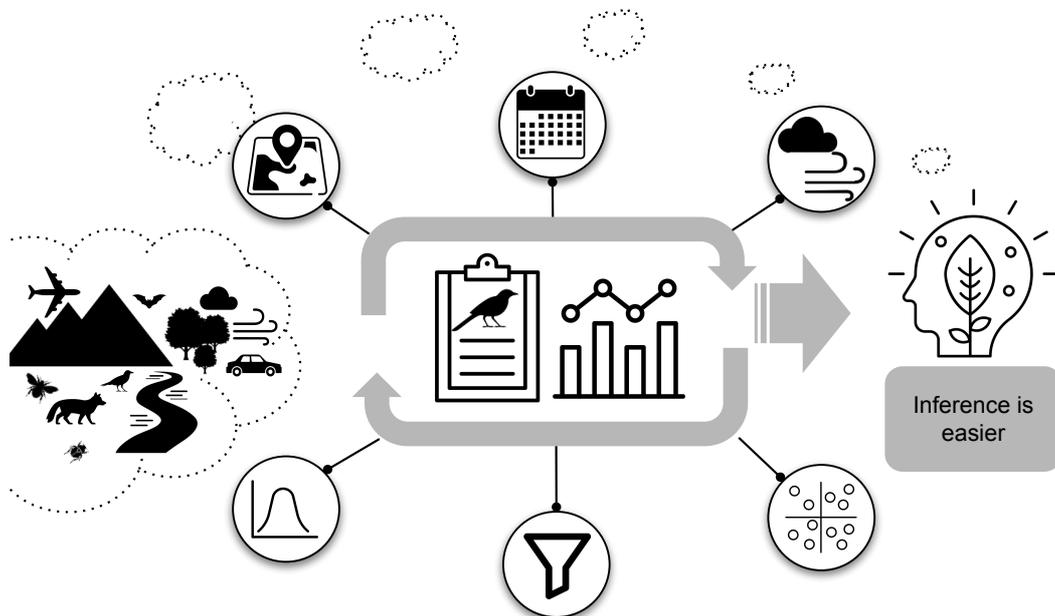
Acknowledgements: The development of echo-dash has been funded by Higher Education Impact Fund and Mark Leonard Trust. The input of JD and AC was funded by The National Capability for UK Challenges (NC-UK) programme, a UKRI/NERC initiative to integrate monitoring and modelling data for the UK environment.

5 References

- Alcocer, I., Lima, H., Sugai, L. S. M., & Llusia, D. (2022). Acoustic indices as proxies for biodiversity: a meta-analysis. *Biological Reviews*, 97(6), 2209-2236.
- Bradfer-Lawrence, T., Desjonqueres, C., Eldridge, A., Johnston, A., & Metcalf, O. (2023). Using acoustic indices in ecology: Guidance on study design, analyses and interpretation. *Methods in Ecology and Evolution*, 14(9), 2192-2204.
- Gibb, K. A., Eldridge, A., Sandom, C. J., & Simpson, I. J. (2024). Towards interpretable learned representations for ecoacoustics using variational auto-encoding. *Ecological Informatics*, 80, 102449.
- Kahl, S., Wood, C. M., Eibl, M., & Klinck, H. (2021). BirdNET: A deep learning solution for avian diversity monitoring. *Ecological Informatics*, 61, 101236.
- McInnes, L., Healy, J., & Melville, J. (2018). UMAP: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:1802.03426.
- Robertson, T. & Simonsen, J. Participatory Design. In Simonsen, J., & Robertson, T. (Eds.). (2012). *Routledge International Handbook of Participatory Design* (1st ed.). Routledge. <https://doi.org/10.4324/9780203108543>
- Sethi, S. S., Jones, N. S., Fulcher, B. D., Picinali, L., Clink, D. J., Klinck, H., ... & Ewers, R. M. (2020). Characterizing soundscapes across diverse ecosystems using a universal acoustic feature set. *Proceedings of the National Academy of Sciences*, 117(29), 17049-17055.
- Zippenfenig, P. (2024). Open-Meteo.com Weather API (1.4.0). Zenodo. <https://doi.org/10.5281/zenodo.145824>



1 Ecological inference from the analysis of passive acoustic surveys is hard because there are multiple stages of abstraction from ecological community to species and soundscape data, inherent confounds within the soundscape, and ambiguity in most analysis methods.

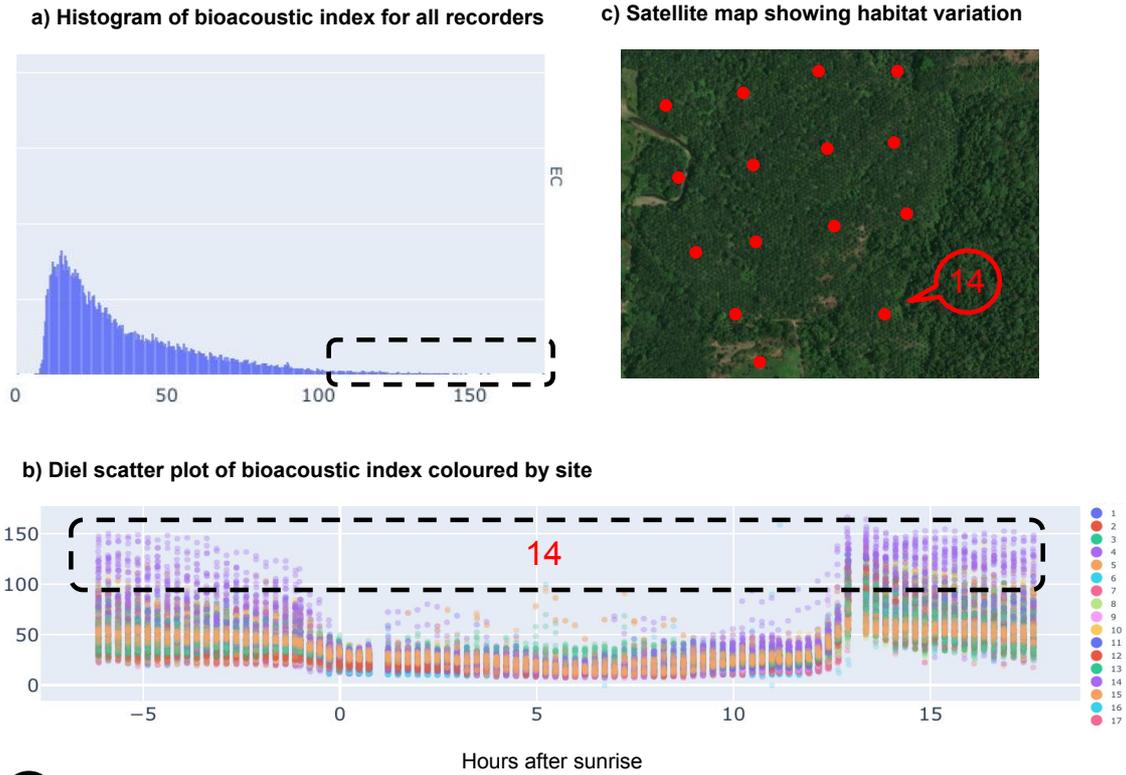


2 Rapid, interactive, iterative exploration of analysis results in flexible views and environmental contexts can facilitate interpretation by providing insight into sources of variation, leading to robust ecological inference.

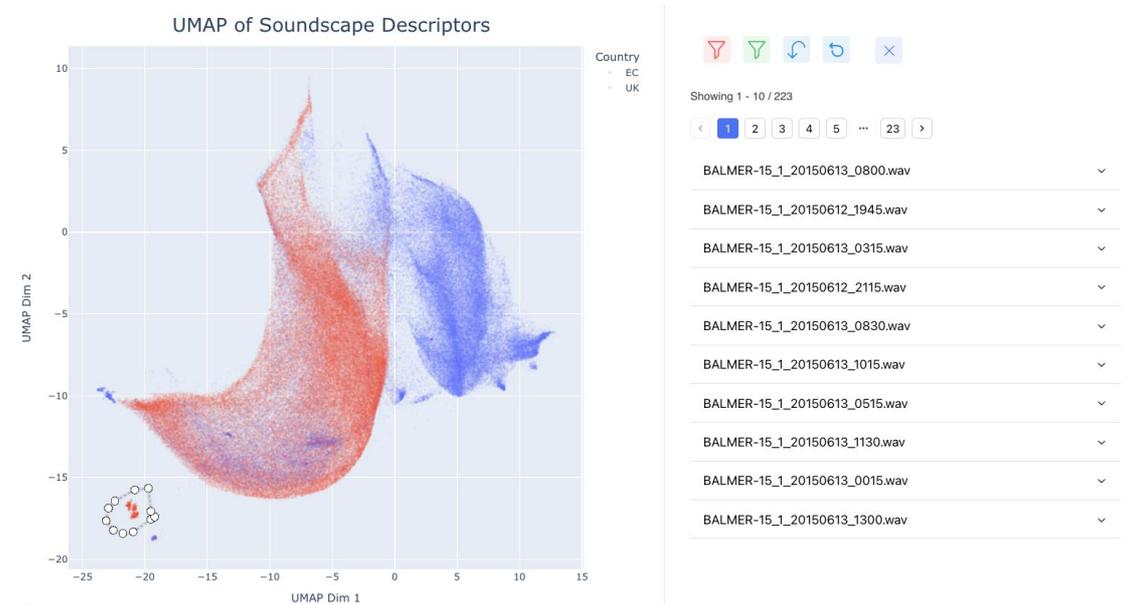
Figure 1. Schematic illustration of the challenge of robust ecological inference in passive acoustic monitoring (1) and the interactive approach which motivates echo-dash (2). When making inferences (1) regarding the ecological community, we consider the biophony of the acoustic community, the signals of which are entangled with sounds of geophony and technophony; the resulting soundscape is digitally recorded and statistically described. When these descriptors are ambiguous (as in soundscape indices) or inherently biased, uninterpretable and prone to shortcut learning (as in BirdNET or machine learning embeddings) this adds additional layers of complexity. As such, interpreting meaning from acoustic data without wider ecological context is hard. Echo-dash (2) provides flexible, accessible, interactive data-views, that enable users to explore automated species data and soundscape descriptors in spatial, temporal and environmental context. Data owners and users familiar with the survey sites are able to draw on their tacit and explicit ecological knowledge of the local environment, helping to identify confounds. Customisable filtering and grouping facilitate rapid formation and testing of hypotheses about the ecological meaning of descriptors and model outputs, supporting more robust ecological inference. Results can be output directly for reporting or further analysis; with data retained by the originator.



Figure 2. Echo-dash workflow features and functions. Echo-dash comprises a parallelised data-processing pipeline feeding into a dashboard (1). The dashboard affords interactive perceptualisation, filtering and faceting of data and views and exports data subsets, plots and code to generate plots. Panel 2 illustrates the multiple, flexible views that support users in: checking meta-data for integrity (spatial locations and temporal recording schedules, flagging file corruption); exploring data interactively, through plots and selected audio playback, and in wider environmental contexts (descriptor distributions, UMAP, diel plots, species lists and species richness, weather data). All views can be filtered to explore observed patterns and clusters (by time, date, extreme values, weather events or selecting clusters in the UMAP). All plots can be interactively faceted and flexibly aggregated. Species data views include functional and habitat groupings and species richness polar plots. Known species lists can be uploaded for comparison. BirdNET detection thresholds can be interactively updated. Results and curated data sets, plots and code to generate them can be exported in a range of formats.



1 Interactive exploration of different data aggregations bolsters insight into source of variation that is obscured under standard approaches



2 Interactive filtering of UMAP reveals source of variation previously unapparent

Figure 3. Examples of insight driven by interactive exploration. Under standard spatiotemporal aggregations, the variance observed across recorders at a particular site are hard to interpret (1a). By visualising the same data as a scatter plot (1b), coloured by location and cross-checking with satellite map (1c), it is observed that bioacoustic index values above 100 come from recorder 14 and that this is located on the edge of the palm oil plantation, adjacent to an area of secondary forest. The hypothesis that this might lead to differences in acoustic community can be tested by examining species probability results (if avian) or auditioning the relevant samples (for other species). Interactive graphical filtering (2) of the UMAP enables inspection of outlying clusters and auditioning can be used to check hypothesis of device error.