Interactive design of interpretable features for marine soundscape data annotation

Thiago S. GOUVÊA, Ilira TROSHANI, Marc HERRLICH, Daniel SONNTAG
Interactive Machine Learning, DFKI, Germany

Abstract. Machine learning (ML) is increasingly used in different application domains. However, to reach its full potential it is important that experts without extensive ML training be able to create and effectively apply models in their domain. This requires forms of co-learning that need to be facilitated by effective interfaces and interaction paradigms. Inspired by the problem of detecting and classifying sound events in marine soundscapes, we are developing Seadash. Through a rapid, iterative data exploration workflow, the user designs and curates features that capture meaningful structure in the data, and uses these to efficiently annotate the dataset. While the tool is still in early stages, we present the concept and discuss future directions.

Keywords. interactive machine learning, marine research, passive acoustic monitoring, sound classification, data programming

1. Introduction

AI systems are deeply integrated with our daily lives, and that includes our work lives. Professionals of all domains of expertise expect to take part in the benefits of the ongoing ML revolution—AI systems are expected to act as work assistants, helping humans to better acquire new insights and learn new knowledge more effectively [1,2]. Great hope is placed on such systems to help tackling climate change [3] and promoting sustainable development [4]. In practice, realisation of these expectations is often slowed down by lack of familiarity of domain experts with ML theory and models, as well as low availability of domain-specific annotated datasets [5]. We propose that these problems can be mitigated by interactive machine learning (IML) systems for designing interpretable models through interacting with the data graphically. Inspired by one particular domain, namely marine acoustic monitoring, we are developing Seadash.

Human activity is changing the face of our planet. Ecosystems everywhere are heavily impacted, and marine ecosystems are no exception. The oceans are home to a large diversity of life forms, and they are of importance also as a source of wealth and resources for the economy [6]. The sea soundscape has changed dramatically as a consequence of economic activity, posing dangers for marine species [7]. Monitoring of ocean soundscapes is important both for long term ecological studies (e.g. population monitoring) as well as for mitigating impact of specific anthropogenic sound sources (e.g. construction of off-shore wind farms; oil surveys). Given the magnitude of the oceans, monitoring poses coverage challenges, and underwater autonomous vehicles (UAVs) such as buoyancy gliders are instrumental [8]. These are small ($< 2m$), acoustically quiet and low-power UAVs housing hydrophones and multiple other oceanographic sensors that are capable of long-term deployments during which they log large amounts of data with minimal disruption of marine habitats. While several of the sound-events of interest are well characterised (section **??**), novel sound-events are likely to arise with each new deployment, and future-proofing the technology requires a method for efficient manual annotation and subsequent automatic detection of custom sound-events.

We are developing Seadash as an interactive tool for designing interpretable models for detection and classification of sound events in datasets such as those generated from hydrophone recordings carried out by UAVs. Seadash development is still in an early stage. We will present the current state along with design ideas and future directions.

2. Concept and Implementation

Seadash is designed as a ML-based data analysis tool that privileges the knowledge and intuition of domain experts who are not ML experts. While ML algorithms in general are powerful tools for pattern recognition (or data hungry "curve fitting", as described by some notable critics [9]), human intelligence seeks to explain the world through models that mimic its structure [10]. The latter approach, when applied to ML models, leads to improved data efficiency [11,12] and interpretability [13]. Seadash fosters human-AI collaboration precisely at this point. Its core concept is to empower users with graphical data exploration tools to iteratively design features that capture the underlying factors of variation in the data. In the present context, these factors are the entities that compose the marine soundscape (e.g., marine mammals, ships, UAVs) and the sound events produced

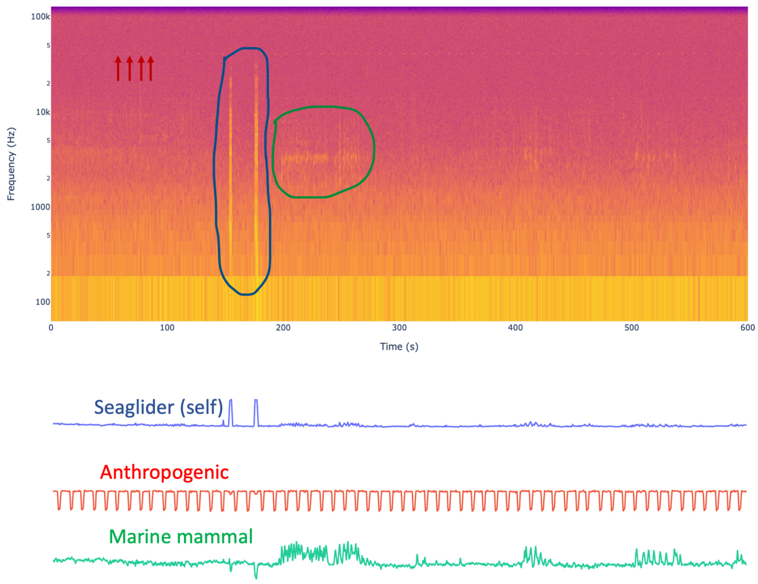

Figure 1. Features built with Seadash disentangle marine soundscape into its composing entities. (**top**) Spectrogram of sounds recorded by an UAV. Includes sounds of the robot itself, unknown human sources, and marine mammals. (**bottom**) Features built with graphical data exploration tools in Seadash are interpretable.

by them (figure 1). These features will then serve as the basis for downstream tasks such as data annotation or automated detection and classification.

The current Seadash implementation is based on Dash [14], a Python framework for creating web-based applications; signal processing and ML elements are implemented with SciPy and scikit-learn [15]. The main elements of the UI are a file input element, a list of entities, and the data visualisation panels (figure 2). The data visualisation panel is vertically split in two halves, each with its own auxiliary control panel. On the top half, an always visible spectrogram constitutes a least-processed visualisation of the selected file, and serves the purpose to ground the users' intuition during data exploration. The bottom half is the feature design panel. The control panel associated with the spectrogram allows zooming on a region of interest (ROI) and adjusting resolution, as well as playing the sound file cropped and filtered to match the time interval and frequency band delimiting the ROI. At this stage, the user should be able to identify sound events of interest and to attribute them to an entity. Entities are thought of as the actors in the data generative process responsible for generating each of the event types, and will serve as classification labels. After adding or selecting an entry on the entities list, the user

Figure 2. A screenshot of the Seadash application. Its main user interface (UI) elements are a file input element (top), a list of entities (left), and the data visualisation panels (centre-right, taking the largest part of the UI area). These are split between a spectrogram (top) and designed features (bottom). Each feature is a sub-tab of an entity.

proceeds to the feature design step. Features are designed by applying transformations to the spectrogram, and are displayed on the bottom half of the data visualisation panel, always time-aligned with the spectrogram above it.

Seadash offers interactive graphical tools for signal processing, unsupervised ML, and dimensionality reduction that allow rapid, iterative experimentation with the aim of producing features whose activation can be interpreted as the presence of the marked entity. One of the simplest tools builds features by computing the average power withing a user selected frequency band; that would suffice to capture the broadband noise produced by the UAV itself in figure 1. The features depicted in figure 1 were produced with a four-click procedure implementing a processing pipeline consisting of frequency band selection, Gaussian convolution along the temporal dimension, and principal components analysis (PCA). The user is encouraged to design multiple, redundant features using distinct methods. Features will serve as a representational basis for downstream tasks (e.g. data annotation by setting thresholds on the features). At the time of writing, we are developing new feature design methods based on embeddings in deep neural networks, as well as closing the loop between data annotation and feature design by introducing modules implementing supervised ML methods to the toolbox (figure 3).

3. Conclusion and Future Work

In this paper, we present Seadash: an interactive data annotation and model building tool inspired by ocean research. Building on the concept of data programming and extending it to the graphical domain, Seadash already allows domain experts to identify and create interpretable features that can be used for downstream ML tasks. As a next step, we will systematically study domain experts' experiences and usage. We are working on alternative embeddings and selection methods (e.g. feature pairs could be plotted against each other to form a bi-dimensional state space, and a lasso selection tool could be used to demarcate regions to be labelled.) In the process of aiding the user in annotating the dataset, Seadash will incorporate methods of topological data analysis (TDA) [16] and learn a set of features that constitute a disentangled representation [11,12] of the data.

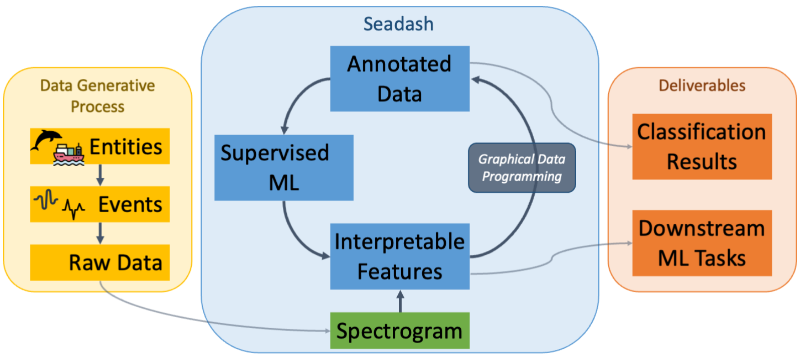

Figure 3. In Seadash, users build interpretable features from which sound events produced by identifiable entities can be recovered. The data thus embedded can be used in downstream tasks such as annotation through graphical data programming.

Acknowledgement

TSG acknowledges Gustavo B. M. Mello (OsloMet University), Luca Tassara (Akvaplan-niva), and Ehsan Abdi (Cyprus Subsea) for inspiring ideas and use cases.

References

[1] Dengel A, Devillers L, Schaal LM. Augmented Human and Human-Machine Co-evolution: Efficiency and Ethics. In: Braunschweig B, Ghallab M, editors. Reflections on Artificial Intelligence for Humanity. Cham: Springer International Publishing; 2021. p. 203-27. Available from: https://doi.org/10.1007/978-3-030-69128-8_13.

[2] Sejnowski T. Artificial intelligence will make you smarter; 2018. Available from: http://theconversation.com/artificial-intelligence-will-make-you-smarter-101296.

[3] Rolnick D, Donti PL, Kaack LH, Kochanski K, Lacoste A, Sankaran K, et al. Tackling Climate Change with Machine Learning. ACM Computing Surveys. 2022 Feb;55(2):42:1-42:96. Available from: https://doi.org/10.1145/3485128.

[4] Gomes C, Dietterich T, Barrett C, Conrad J, Dilkina B, Ermon S, et al. Computational sustainability: computing for a better world and a sustainable future. Communications of the ACM. 2019 Aug;62(9):56-65. Available from: https://doi.org/10.1145/3339399.

[5] Simard PY, Amershi S, Chickering DM, Pelton AE, Ghorashi S, Meek C, et al. Machine Teaching: A New Paradigm for Building Machine Learning Systems. arXiv. 2017 Aug. ArXiv: 1707.06742. Available from: http://arxiv.org/abs/1707.06742.

[6] OECD, editor. The ocean economy in 2030. Paris: OECD Publishing; 2016.

[7] Duarte CM, Chapuis L, Collin SP, Costa DP, Devassy RP, Eguiluz VM, et al. The soundscape of the Anthropocene ocean. Science. 2021;371(6529):eaba4658.

[8] Verfuss UK, Aniceto AS, Harris DV, Gillespie D, Fielding S, Jiménez G, et al. A review of unmanned vehicles for the detection and monitoring of marine fauna. Marine Pollution Bulletin. 2019 Mar;140:17-29. Available from: https://www.sciencedirect.com/science/article/pii/S0025326X19300098.

[9] Pearl J. To Build Truly Intelligent Machines, Teach Them Cause and Effect; 2018. Available from: https://www.quantamagazine.org/to-build-truly-intelligent-machines-teach-them-cause-and-effect-20180515/.

[10] Lake BM, Ullman TD, Tenenbaum JB, Gershman SJ. Building machines that learn and think like people. Behavioral and Brain Sciences. 2017;40. Publisher: Cambridge University Press. Available from: https://www.cambridge.org/core/journals/behavioral-and-brain-sciences/article/building-machines-that-learn-and-think-like-people/A9535B1D745A0377E16C590E14B94993.

[11] Higgins I, Amos D, Pfau D, Racaniere S, Matthey L, Rezende D, et al. Towards a definition of disentangled representations. arXiv preprint arXiv:181202230. 2018.

[12] Schölkopf B, Locatello F, Bauer S, Ke NR, Kalchbrenner N, Goyal A, et al. Toward causal representation learning. Proceedings of the IEEE. 2021;109(5):612-34.

[13] Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence. 2019 May;1(5):206-15. Number: 5 Publisher: Nature Publishing Group. Available from: https://www.nature.com/articles/s42256-019-0048-x.

[14] Dash. Plotly; 2022. Available from: https://dash.plotly.com/.

[15] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine Learning in Python. MACHINE LEARNING IN PYTHON:6.

[16] Wasserman L. Topological Data Analysis. Annual Review of Statistics and Its Application. 2018;5(1):501-32. Available from: https://doi.org/10.1146/annurev-statistics-031017-100045.