

People of Data

Meet the authors: Georgios Rizos, Jenna L. Lawson, and Björn W. Schuller

Georgios Rizos,^{1,2,3,*} Jenna L. Lawson,^{4,5} and Björn W. Schuller^{1,6,7,8}

¹GLAM – Group on Language, Audio, & Music, Department of Computing, Imperial College London, SW7 2RH London, UK

²University of Cambridge, 15 J J Thomson Avenue, CB3 0FD Cambridge, UK

³Imperial College London, 180 Queen's Gate, SW7 2AZ London, UK

⁴Department of Life Sciences, Imperial College London, SL5 7PY Ascot, UK

⁵UK Centre for Ecology and Hydrology, Benson Lane, OX10 8BB Wallingford, UK

⁶EIHW – Chair of Embedded Intelligence for Health Care and Wellbeing, University of Augsburg, 86159 Bavaria, Germany

⁷Technical University of Munich, 21 Arcisstrasse, 80333 Munich, Germany

⁸audEERING GmbH, 1 Friedrichshafener Strasse, 82205 Gilching, Germany

*Correspondence: georgios.rizos12@imperial.ac.uk

<https://doi.org/10.1016/j.patter.2024.100952>

In their recent publication in *Patterns*, the authors proposed a methodology based on sample-free Bayesian neural networks and label smoothing to improve both predictive and calibration performance on animal call detection. Such approaches have the potential to foster trust in algorithmic decision making and enhance policy making in applications about conservation using recordings made by on-site passive acoustic monitoring equipment.

This interview is a companion to these authors' recent paper, "Propagating Variational Model Uncertainty for Bioacoustic Call Label Smoothing".

What motivated you to become a researcher? Is there anyone or anything that helped guide you on your path?

Jenna Lawson: We are facing the dual challenge of mitigating biodiversity loss and reducing the impact of climate change. I am motivated by my passion to help overcome these challenges, which I believe are some of the most pressing challenges facing humanity. Passion for wildlife and our planet starts during childhood, my mother instilled this passion in me and I hope I can pass this on to future generations.

Georgios Rizos: I think I made the switch from a vaguely software engineering-oriented mindset to data science and machine learning during my masters' studies in biomedical engineering. I remember being taken with the range of impactful applications on which such methods can be applied, and really wanted to work hard to become an expert in the methodology. I am still as excitable about new cool applications, though as a postdoc now I tend to temper this excitement with the need to build an expertise. As for guidance in conducting research, academic writing, and charting career course, certainly all of my supervisors so far: Drs. Symeon Papadopoulos

and Yiannis Kompatsiaris from my first research post, my PhD supervisor and co-author of our *Patterns* article Dr. Björn Schuller, and my postdoc advisor Dr. Cecilia Mascolo.

What is the definition of data science in your opinion? What is a data scientist? Do you self-identify as one?

G.R.: *Data science* is a wide umbrella term covering a lot of concepts, I believe. For me, it covers anything pertaining to exploring, cleaning, pre-processing, and learning from data. When I use the term, I tend to mean statistics and machine learning. I am aware that to some people, especially in the industry, it can have connotations related to database systems or data warehousing. As such, the matter of what a data scientist is in each context might have to do with the eye of the beholder. I do self-identify as one, and sometimes introduce myself as one, although I mostly use machine learning researcher.

What attributes, in your opinion, make a data scientist successful?

G.R.: For the most part, whatever makes one successful in most other disciplines. I am very mindful to not encourage any

magical thinking about the requirements of data scientists (or any other role), so I would highlight the basics: knowing how to learn (i.e., slowly), perseverance, application. Effective communication helps, nurturing relationships with excellent colleagues. These are all important, although they probably do not make for an interesting answer. At the risk of possibly glorifying undesirable cognitive patterns, I would say that orientation to minute detail can be quite important, for example, by means of nitpicking, and design of most strict comparisons and ablations. I often ask: "how sure are you that you are not giving unfair advantages to your proposed method"? This is why I often enjoy benchmark, replication, and negative results papers more than a new technical innovation. I suspect my definition of success (and enjoyment) may be idiosyncratic though.

As a data scientist, which of the current trends in data science seem most interesting to you? In your opinion, what are the most pressing questions for the data science community?

G.R.: I am invested in any work that advances trustworthiness in artificial intelligence (AI). This can mean uncertainty-awareness and model calibration, as





Figure: From left to right Georgios Rizos, Jenna L. Lawson, and Björn W. Schuller.

explored in our *Patterns* article,¹ as well as concepts like explainability and interpretability. Of particular interest to me is fair AI, which aims to reduce disparity of model performance across demographic groups. That being said, there are multiple definitions of what constitutes fairness that are arguably conflicting. Also, the fit of each definition is possibly dependent on the nature of the task and the label distribution. The discipline of fair AI leads us to understand implicit biases in commonly adopted deep model architectures, training, and evaluation processes, as well as entire experimental designs. Such work can turn quite philosophical, and philosophy is one thing, I believe, that is needed more in the field to craft better definitions and targets. You can't solve AI inequity with more compute.

What barriers have you faced in pursuing data science as a career?

J.L.: A career in conservation is a hard one to pursue. There are few jobs and positions are often poorly paid. A higher degree is often needed, which is not often accessible to everyone, which is why it's important to support those with a passion for conservation but without the opportunities for study. It is also important to recognize those that are trained in informal settings, as their knowledge often exceeds a university level degree. As ecologists we often lack the data science skills needed in favor of developing the theory and fieldwork skills for our roles, however I see this shifting now with many data scientists working in ecology and less and less people coming through with field-based skills.

What is the role of data science in your field? What advancements do you expect in data science in this field over the next 2–3 years?

J.L.: I am most interested in how we can use technology to increase the spatial and temporal scale of monitoring, reduce the invasiveness of data collection, and automate methods to allow us to do more with the available time and funds. I think one of the most pressing questions for data science is how we can create species classifiers for the sheer number of species known to science and, indeed, those not known to science. For example, a classifier might be able to pick out a species it has never seen before, which might be a newly encountered species. A major challenge is also ensuring that classifiers function across different areas. For example, a classifier trained on data from Brazil might not function for the same species in Costa Rica. Is it a matter of collecting training data for every region or can data science solve this challenge? Data scientists and ecologists need to work together to develop accurate classifiers for those species that can tell us most about an ecosystem, by identifying umbrella and keystone species, pollinators, predators. We need to work together from the project design all the way through to the outputs to maximize the benefits of our specialist skills. In the next 2–3 years I expect there to be classifiers available for many more species and for progress in statistical models to follow new automated methods of monitoring.

Let's talk about the work you published at *Patterns*. How did this project you wrote about come to be?

G.R.: My involvement with animal call detection came after my PhD supervisor, Björn Schuller, introduced me via email to co-author Jenna Lawson, who was recording spider monkey calls using passive acoustic monitoring as part of her own PhD. This led to the discovery—on my part—of a fascinating application, meeting several biology/conservation experts and academics, and a very fruitful interdisciplinary collaboration. The importance of the problem, and my need to explore uncertainty-aware deep learning methods on exciting new data for my PhD motivated me to get involved. A couple of publications later, our *Patterns* article¹ was conceptualized, in part, by our need to evaluate the generalization of our methods to more datasets with different acoustic properties, locations, and species, as well as by the desire to explore models that exhibit high calibration performance on acoustic animal detection.

Did you encounter any particular difficulties, or were there any specific challenges about data, data management, or FAIR data sharing that you dealt with? How did you overcome them? Can others use the solutions you used to overcome these challenges?

G.R.: This should not be surprising to anyone, but data organization, cleaning, and pre-processing is complicated—beyond, even, the great effort that had

already gone into fieldwork, recordings, and manual annotation before even my involvement. Even then, there was a need to further transform the annotations into a format more amenable to processing by a range of baselines, while at the same time making sure that the model to be learned will be of use to the conservation experts. In terms of partitioning to training, validation, and test sets, we needed to avoid doing ourselves any favors by ensuring that the sets were independent with respect to recording location or, at the very least, significant time distance. This required cross-referencing with certain reports on the fieldwork. Since at each time point in a recording, multiple species could be heard, there was detailed pre-processing required in audio clip segmentation in order to generate a good quality multi-label classification dataset. The help of MSc students in this and previous projects, like co-author Xin Wen, was invaluable. The data are now publicly available, as well as the scripts to perform such pre-processing and replication of experiments.

What's next for the project? What's next for you?

G.R.: On the project, there are a lot of ideas circulating, based on methods that have been shown to perform well on similar problems and data. Since passive acoustic monitoring makes available far larger quantities of data than can realistically be manually annotated, self-supervised learning techniques to learn robust acoustic representations that can help generalization to new locations, species, or animal dialects seems promising. Such an approach entails challenges as well; there is a need to account for drastic domain shifts pertaining to acoustic environment and background noise distribution. This can be a step toward the development of easy-to-use tools for non-machine learning experts and along with human-in-the-loop annotation solutions can lead to automation of the data science pipeline. As for me personally, I am continuing research on the methods of our *Patterns* article¹ but also in the context of automated diagnostics on resource-restricted mobile devices as part of my postdoc position.

J.L.: We will continue to work together to develop models for some of the key species used in this analysis, such as the spider monkey, to ensure these

models function across regions and countries where the species are present and even across different sub-species. For me, I am now working on developing new automated monitoring methods for moths, using a camera system, and acoustics for a wide variety of species, as well as developing partnerships with local organizations across tropical ecosystems.

Professor Schuller, what drew you to this area of research? How has the research focus of your team evolved over the years?

Björn Schuller: I love audio and listen to a lot of music. Saving our planet and habitat seems a key priority and animal monitoring seems to be one of the opportunities for artificial intelligence to contribute. Being an audiophile, this seemed the obvious choice.

Where is the team currently based, and how long have you been there? What kind of atmosphere do you look to foster in your team?

B.S.: The teams are based in London and Munich since more than a decade. The split works surprisingly well. I look for high diversity in all aspects. We have expertise in computer science, languages, signal processing and engineering, and even sound arts to psychology and physics. The team comes from all over the world and highly benefits from the different viewpoints and expertise.

Which achievement or discovery in your career are you most proud of? Looking back, what advice would you have given yourself at the start of the career?

B.S.: I really like having organized research challenges for highly reproducible and comparable research already in a time long before Kaggle and the likes. We were faced with almost incomparable studies and moved to being able both to compare in terms of working on the same data, but also standardized audio features by our openSMILE toolkit and feature specifications. I would have advised myself to start publishing in journals earlier on and start right away to challenge myself more by always going one notch up immediately.

A lot of data scientists continue their career outside of academia; what is your view on that? Do you encourage your students and postdocs to continue their careers in academia and establish their own teams? Are you supportive of careers outside of academia?

B.S.: In today's computer science landscape, it seems normal to do both—sequentially or even in parallel. Many colleagues have double appointments keeping a position in academia while joining industry. Many also started up—including myself. I think this is a good solution, as you can gain experience from the “real world” and lead your research questions by that. Hence, I think, ultimately it is in one's heart where to go—industry, academia, or even both. For me, ultimately, academia is a dream place I can highly recommend working hard for.

DECLARATION OF INTERESTS

The authors declare no competing interests. Georgios Rizos is also affiliated with the University of Cambridge. This work was performed during his PhD candidacy at Imperial College London. Jenna Lawson is also affiliated with the UK Centre for Ecology and Hydrology. Pranay Shah is now affiliated with Advai Ltd. Pranay Shah and Xin Wen worked on this study as MSc students at Imperial College London. Bjoern W. Schuller is also affiliated with the Technical University of Munich and audEERING GmbH.

REFERENCE

1. Rizos, G., Lawson, J.L., Mitchell, S., Shah, P., Wen, X., Banks-Leite, C., Ewers, R., and Schuller, B.W. (2024). Propagating variational model uncertainty for bioacoustic call label smoothing. *Patterns*, 5, 100932.

About the authors

Dr. Georgios Rizos is currently a research associate at the Department of Computer Science and Technology, University of Cambridge, UK. He did his PhD at the Department of Computing, Imperial College London, UK. Before that, he was a research assistant at the Centre for Research and Technology Hellas, Thessaloniki, Greece. His research interests include uncertainty quantification, Bayesian neural networks, and self-supervised learning applied to bioacoustic and healthcare-related audio data, as well as fair machine learning.

Dr. Jenna L. Lawson is a biodiversity scientist specializing in the use of acoustics, cameras, and robotics for monitoring and conservation. Jenna completed her PhD at Imperial College London in 2021, where she used bioacoustic machine learning to monitor Geoffroy's spider monkey and investigated the effects of palm and teak plantations on biodiversity. She then spent 2 years

working for the Robotics department at EPFL in Switzerland and Imperial College London. Currently, at UK CEH, Jenna manages an international network of monitoring equipment to understand the effects of anthropogenic change within tropical ecosystems and the impacts of sustainable farming on biodiversity in the UK.

Dr. Björn W. Schuller received his diploma, doctoral degree, and teaching qualification from the Technical University of Munich where he is now a professor of health informatics. He is also a professor of artificial intelligence at Imperial College London, and co-founder, CEO, and current CSO of audEERING, among other professional

affiliations. Previous academic appointments include a professorship at the Universities of Augsburg and Passau, in Germany, key researcher at Joanneum Research in Graz, Austria, and the CNRS-LIMSI in Orsay, France. He is a fellow of the ACM, IEEE, BCS, ELLIS, ISCA, and AAAC. He has (co-)authored 1,400+ publications.