

# From Images to Knowledge with ImageJ & Friends

virtual conference

**Nov 30 - Dec 2, 2020**

Stephan Preibisch, Stephan Saalfeld, Anna Kreshuk,  
Pavel Tomancak and Virginie Uhlmann

hhmi | janelia  
Conferences



## **SongExplorer: A deep learning workflow for classification and unsupervised discovery of animal acoustic communication signals**

**Tutors:** Ben Arthur ([arthurb@janelia.hhmi.org](mailto:arthurb@janelia.hhmi.org))

**Session 1:** 2020-11-30 13:00 UTC – 2020-11-30 17:00 UTC

**Session 2:** 2020-12-02 17:00 UTC – 2020-12-02 21:00 UTC

## **abstract**

SongExplorer: A deep learning workflow for classification and unsupervised discovery of animal acoustic communication signals

Benjamin J. Arthur, Yun Ding, Medhini Sosale, Faduma Khalif, Srini Turaga, David L. Stern

Many animals produce sounds or substrate-borne vibrations that are critical elements of animal communication but that have proved challenging for segmentation by automated algorithms. We have developed SongExplorer, a web-browser based platform that supports an intuitive, interactive workflow for efficient discovery of animal sounds, manual annotation, training of a deep convolutional network, and automated segmentation of recordings. Data can be explored by simultaneously examining song traces, watching synced video, and listening to songs. With sufficient training data, the machine learning algorithm we trained in the current instantiation of SongExplorer displays higher accuracy than existing heuristic algorithms and similar accuracy as two expert human annotators. SongExplorer is a generalizable interface that supports any segmentation algorithm, allowing integration of new methods as they are developed.

## **outline**

We'll start from scratch and train a model to classify audio recordings. Sounds will initially be detected using simple thresholds in the time and frequency domains. We'll cluster these and manually annotate a few representative examples. A model will be trained on this corpus of ground truth and used to make predictions on new recordings. We'll correct a few mistakes and then train a more refined model. Time permitting I'll demonstrate how classes of sounds can be discovered using an unsupervised approach.

Representative video tutorials are hosted on youtube:

<https://www.youtube.com/playlist?list=PLYXyXDkMwZip8x78RAyN6ee9NK42WBbKb>

To follow along users will need a Linux or Apple Macintosh computer. Microsoft Windows nominally works, but the software is not as thoroughly tested there. I can provide sample recordings of *Drosophila melanogaster* (fruit fly) song, but participants are encouraged to bring their own data.

## **bio**

Ben is a software engineer in the Scientific Computing group at the Howard Hughes Medical Institute. In this capacity he develops data acquisition and analysis programs for scientists at the Janelia Research Campus. Since 2012, he has worked on GENIE, WaveSurfer, MouseLight, JAABA, and SongExplorer among other projects.

Previously Ben was a post-doc and research associate at Cornell University in Ron Hoy's lab, where he studied acoustic communication in insects, including mosquitoes. In graduate school he studied sound localization in barn owls with Mark Konishi, and designed analog VLSI models of audition with Carver Mead. Ben's interest in hearing started in college when he was struck by the disparity between artificial speech recognition algorithms and the mammalian auditory pathway. Prior to that he was a computer geek who learned to code at age thirteen.

Ben has a BS in Computer Science from Trinity University (in Texas) and a PhD in Computation and Neural Systems from Caltech.