Acoustical Society of America, May 26, 2022

ml-audio-dev-tools

Development tools for deep learning models of acoustical signal processing

Scott H. Hawley, Ph.D.

@drscotthawley

- Chem. & Phys. Dept, Belmont U.
- Belmont Data Collaborative
- Harmonai

Brought to you by ASA Technical Committee on Signal Processing

Thanks: Andrew Parker, Brian McFee, Fabian-Robert Stöter, Jesse Engel, Gene Kogan, Dadabots, David Braun, Zach Evans, scart97, Eric Hallahan

ml audio dev tools

…are *quite* few in number

...so we should all build more.

What's too difficult?

Jeremy Howard of fast.ai uses this XKCD cartoon at the start of Lesson 1 in the fast.ai course "Practical Deep Learning for Coders"

...to illustrate the **huge change in the capabilities of Computer Vision systems since the advent of Deep Learning**.

This problem of bird detection is *trivial* nowadays. It is not even worth a homework problem.

What's too difficult?

In 2020, Bruce Denby (Inst. Lagevin, Sorbonne) & I wrote to CA TC with a proposal for a 2021 ASA Student Challenge in Speech-To-Text using video imagery ("Silent Speech"):

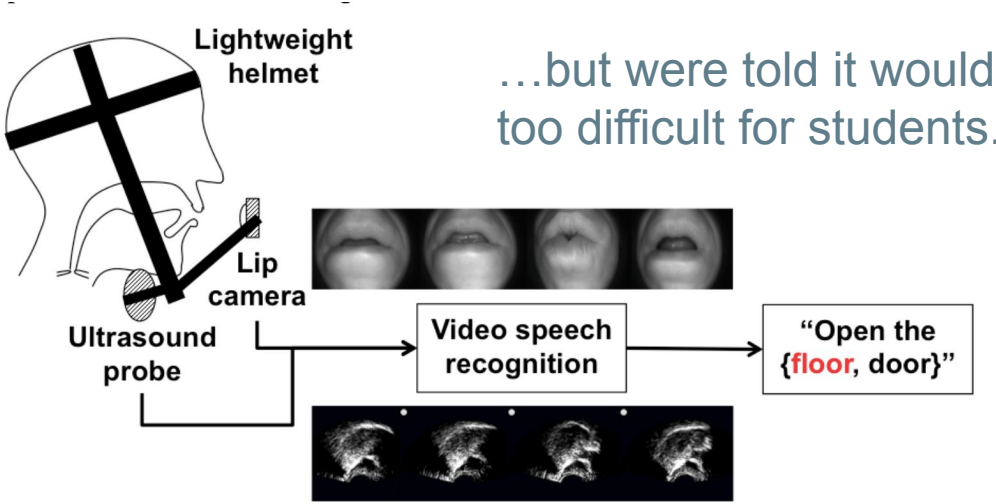

…but were told it would be too difficult for students.

We were *confused*.
With modern DL-CV tools,
This is a *homework problem,*
or at best a minor Kaggle competition.

Some things that were hard become easy with DL

Many human-designed **feature detection** systems can be obviated with Deep Learning systems, which build their own feature detectors as part of the training process.

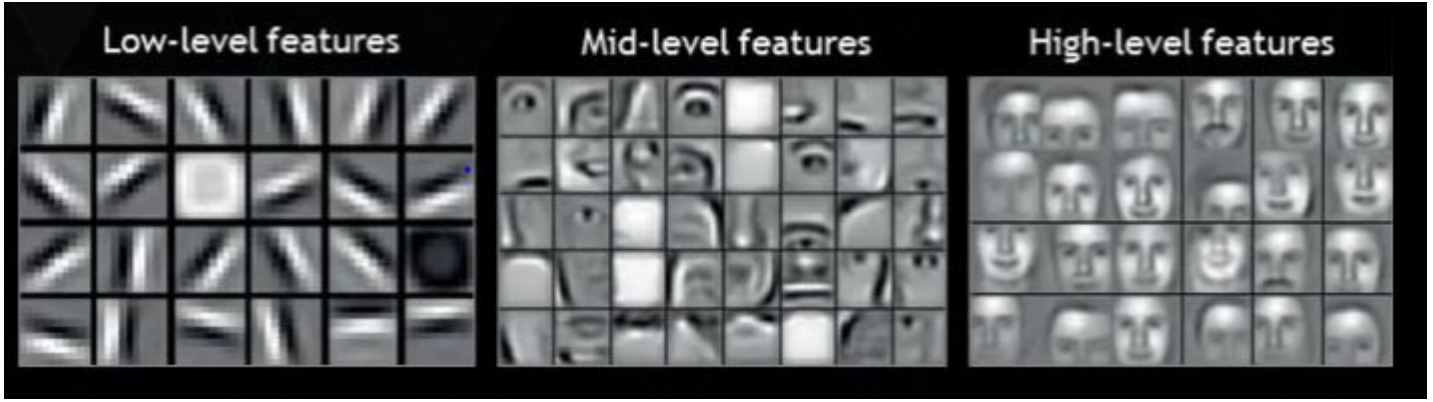

An example hierarchy of feature detectors for a facial recognition system. (Source: Lee et al, 2009.)

Aside: At SP TC meeting Tuesday night, there was talk about a future session on "Feature Extraction and Dimensionality Reduction" – both of these are built in to DL systems.

Some things that were hard become easy with DL

Various examples, but here's a couple:

Neural networks are great at *pattern matching* and *denoising*:

I wrote my first denoising autoencoder before I knew what a Wiener filter was!

I wrote an object detector for ellipses (Hawley & Morrison JASA 2021 & JASA Express Letters 2022) without bothering to try Hough Transforms (b/c it was complex)

Which means...

- there's a new generation of coder-scientists (🙋) who ***hope*** that DL will make up for their *lack of signal processing domain knowledge*. **Not without reason:** DL has shown to *beat* former baselines in many fields for many problems.
- there are veterans in the SP field who are perhaps not up to speed on the rapid pace of advancements in DL for "audio", and may be curious about incorporating DL into their work. ...Good news: This is/(can be) "easy"*!

*compared to hard-core SP/math, if you're already good at coding

More good news: "Everything old is new again". Much of classic SP still finds its way into DL systems & helps drive innovation. (e.g., Vector Quantization!)

Who this talk is for:

- Students and ❤️*Esteemed SP Experts*🤦: "Onboarding"
- Experienced DL-Audio researchers: Sharing tips!

Confession/bias of mine:

- When I say "audio" I usually mean "musical" audio:
  - multi-channel
  - high sample rates (44.1kHz+)
  - ...and *not just classification problems*
  - I tend to not even think about *"Speech". No offense*
- For me, DL $\subset$ ML

ml audio dev tools...

...are few in number

...**Compared** to Tools for DL Models for **Images**, **Text**, ...and even ***Speech***

In these other domains, DL is very mature.*

MANY, many, many tools & demos for viz., analysis, saliency,...

*because money

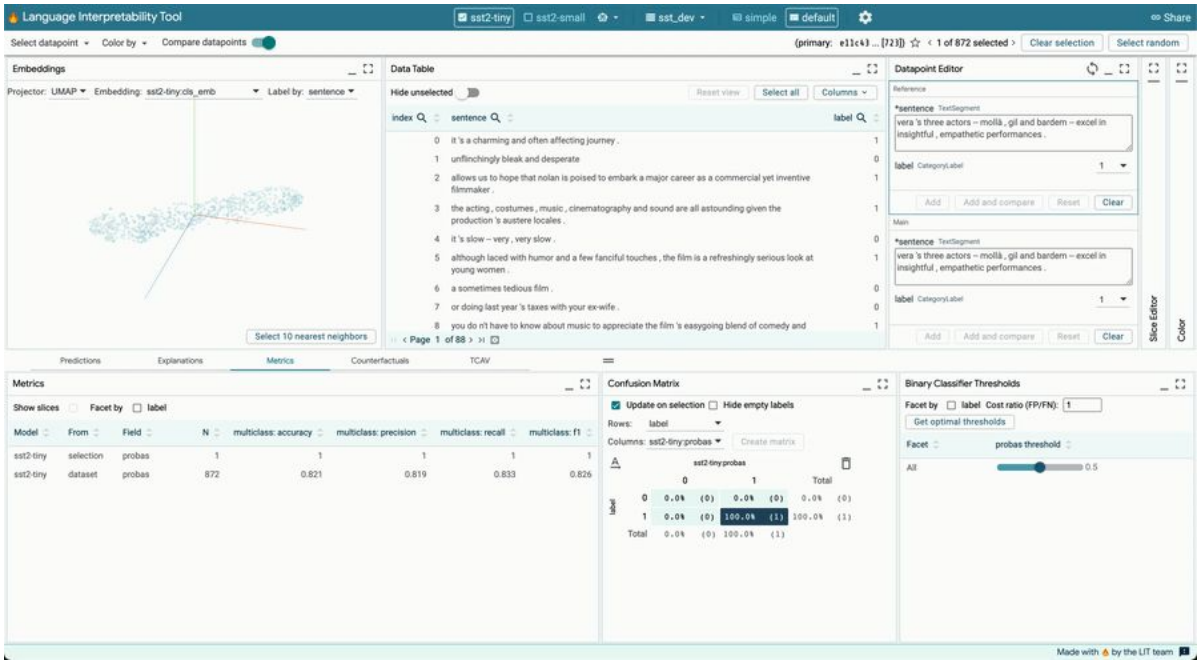

There's no audio tool as cool as Google PAIR's Language Interpretability Tool

Responses I got...

i.e., tools suggested by DL-audio practitioners...

Thanks: Brian McFee, Fabian-Robert Stöter, Jesse Engel, Gene Kogan, Dadabots, Andrew Parker, David Braun, Zach Evans, scart97, Eric Hallahan, & Christian Steinmetz

...were mostly:

I. Basic ML workflow / 101
II. Normal audio workflow
III. A few extra cool tools

I. Basic ML - GPU Computing

Graphics Processing Units (GPUs) are key: 100x+ faster than CPU

You don't need to buy a GPU. Various cloud-based systems let you do GPU computing for (near) free.

Google Colab, Kaggle, Paperspace Gradient, Amazon Sagemaker,...

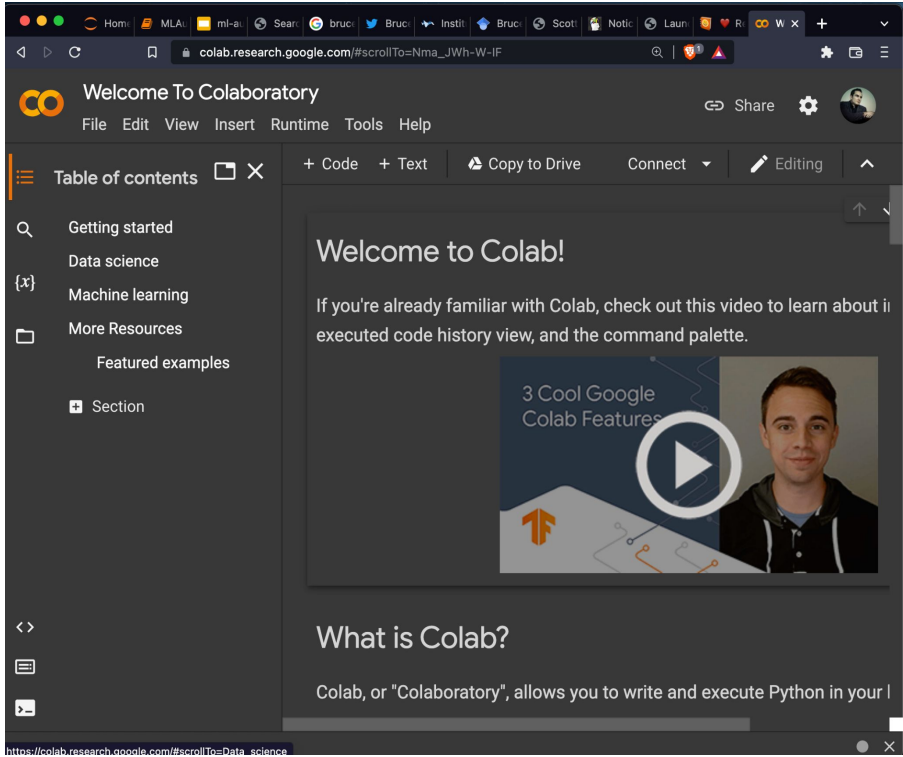

These typically make use of Jupyter notebooks.

(For awesome lib-dev via Jupyter notebooks, check out nbdev.)

I. Basic ML - Python Programming

Python is overwhelmingly the most popular & well-supported programming language for ML. (Much more so than MATLAB, C++, Julia, R, JS,...)

There's usually a library/package that does what you want.

Which DL Library? ➡ **PyTorch** ⬅, Tensorflow, JAX,..

Lightning

I hate `conda` and only use venv+ `pip`

I. Python Packages for Audio and/or DL

- librosa: "Swiss army knife": (Also by Brian McFee: mir_eval)
- torchaudio: GPU processing
- auraloss: Loss functions (Christian Steinmetz)
- Pedalboard: Data augmentation (Peter Sobot/Spotify)
- DawDreamer: Python DAW (David Braun)
- ONNX: Export to JUCE / C++ / JS

II. Normal Audio Workflow

- Audacity ← We'll come back to this one!
- Sonic Visualizer:
- Reaper
- Logic

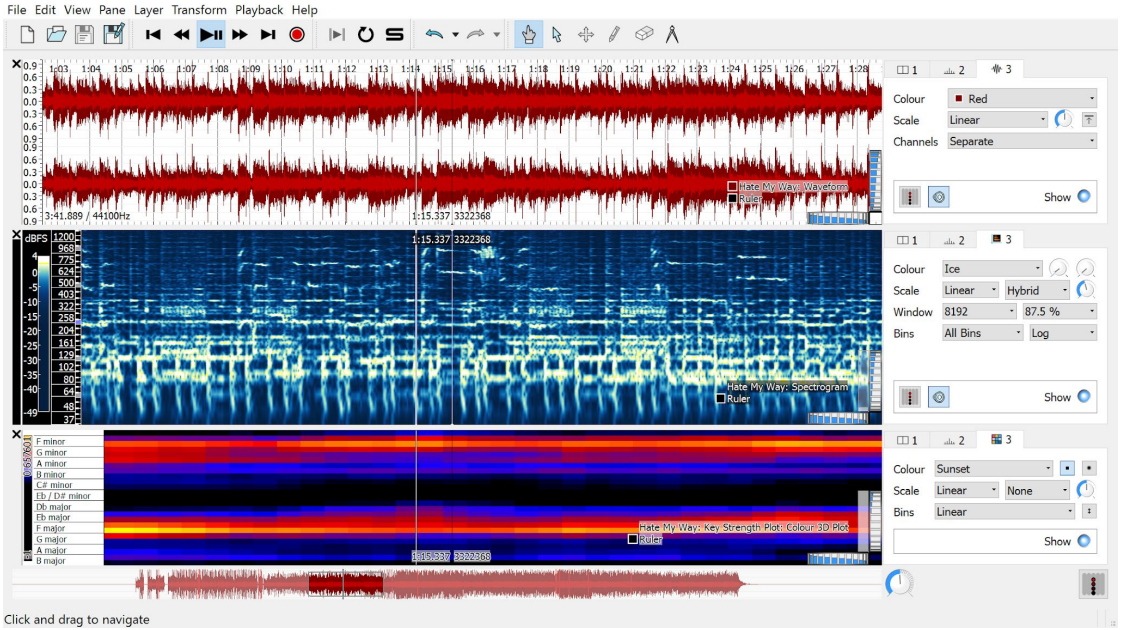

Suggested by Jesse Engel

II. Audio workflow - VSCode audio-preview

(Suggested by
Fabian-Robert Stöter)

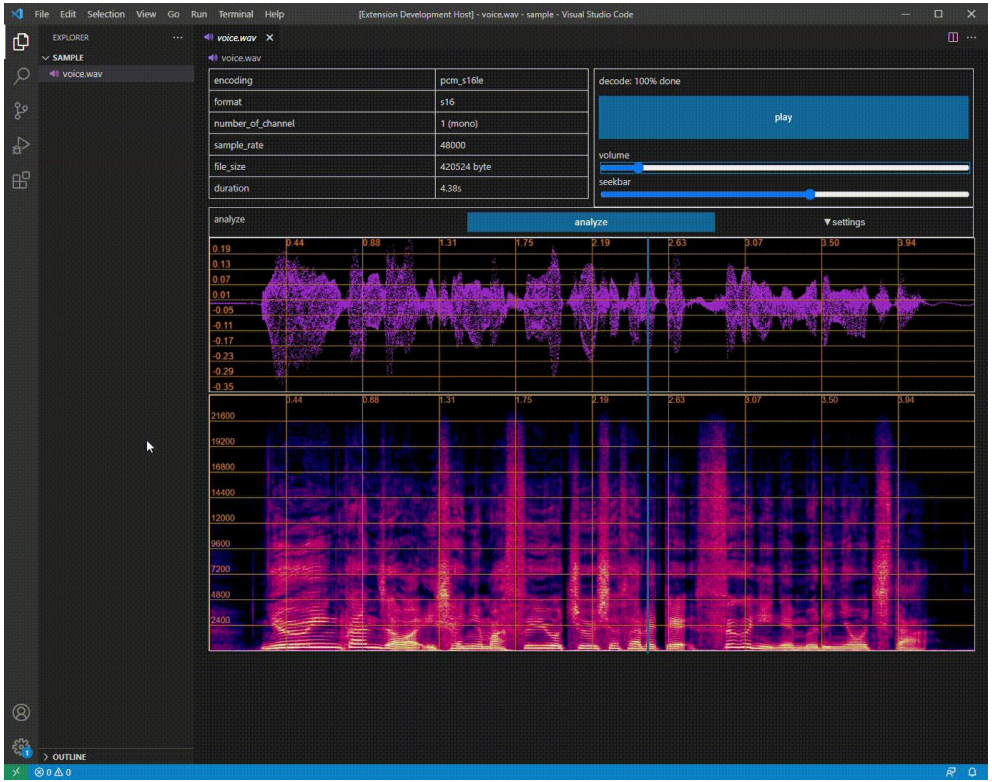

Works in Remote mode!
i.e., play on your laptop
the files on your server.

II.5 Trick: (Mel-)Spectrograms + Image-based DL

- Many posts on "The remarkable effectiveness of Convolutional Neural Networks on (Mel-)Spectrograms"
- Translation equivariance of CNNs fits well with phase (+ pitch) translation invariance of human auditory system.
- Upshot: Just using images of spectrograms with an image-based code can work surprisingly well.
- Makes a great baseline before going for full end-to-end audio DL

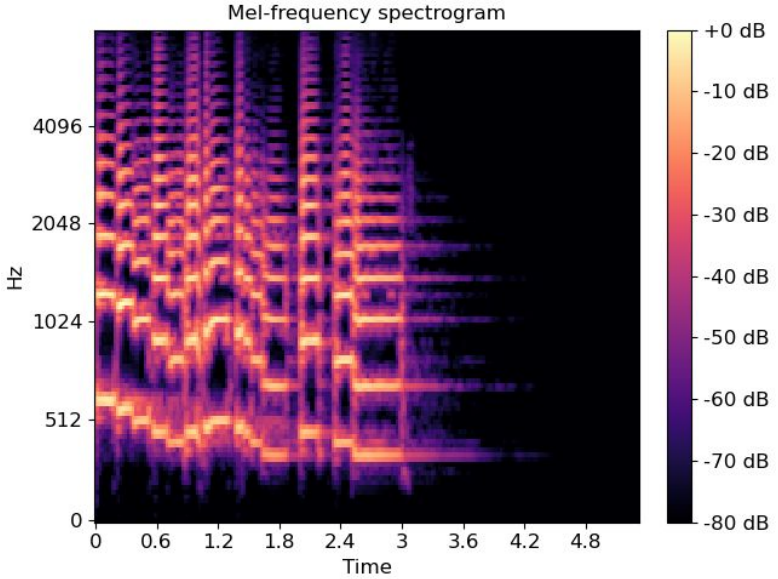

Source: LibRosa (McFee et al)

If you're just starting out, probably start with this method

III. Extra Cool - WandB (Audio) Callbacks!

- Weights & Biases ("wandb") is a cloud-based data-logging service you can use for free.

  (Hawley & Morrison JASA-EL 2022 found it "essential" for keeping track of many, many runs.)

- Among the things you can log & playback are audio examples:

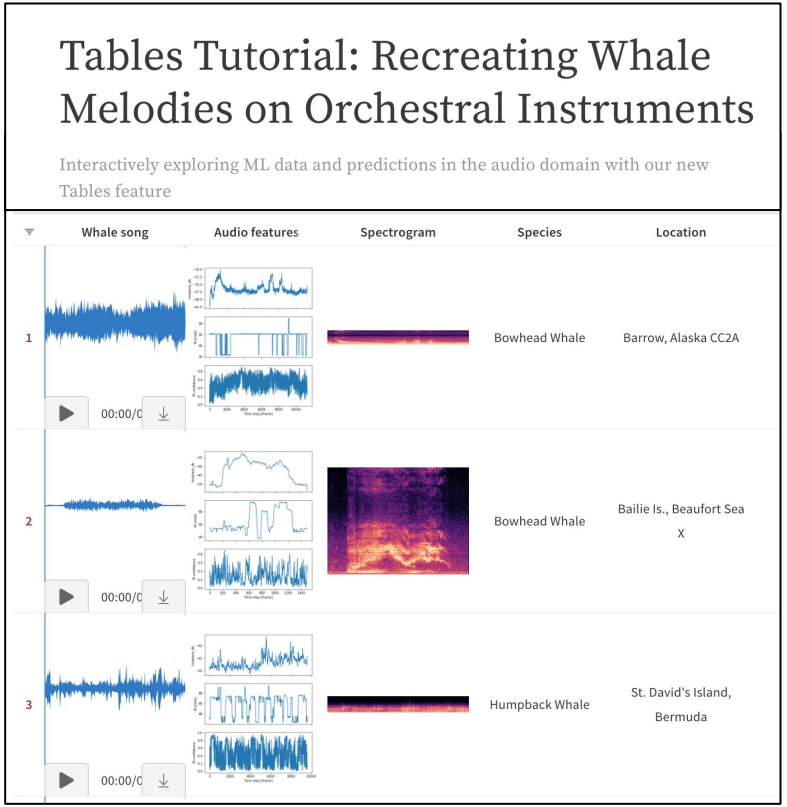

III. Communication: Demo Hosting via Gradio.app

Demo for [Christian Steinmetz & Josh Reiss, NeurIPS 2021]

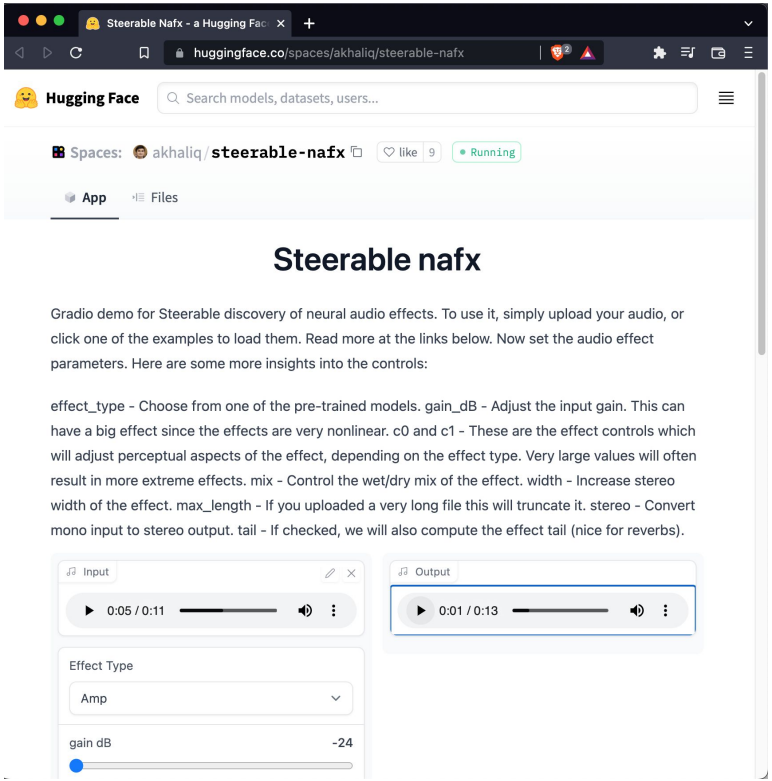

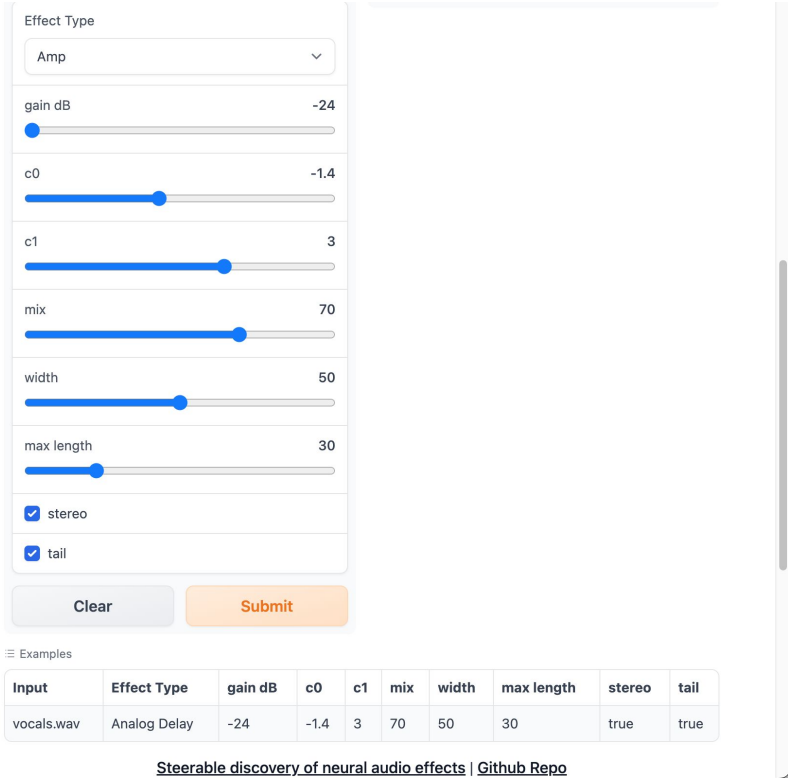

(Suggested by Fabian-Robert Stöter & scart97)

III. Label Studio

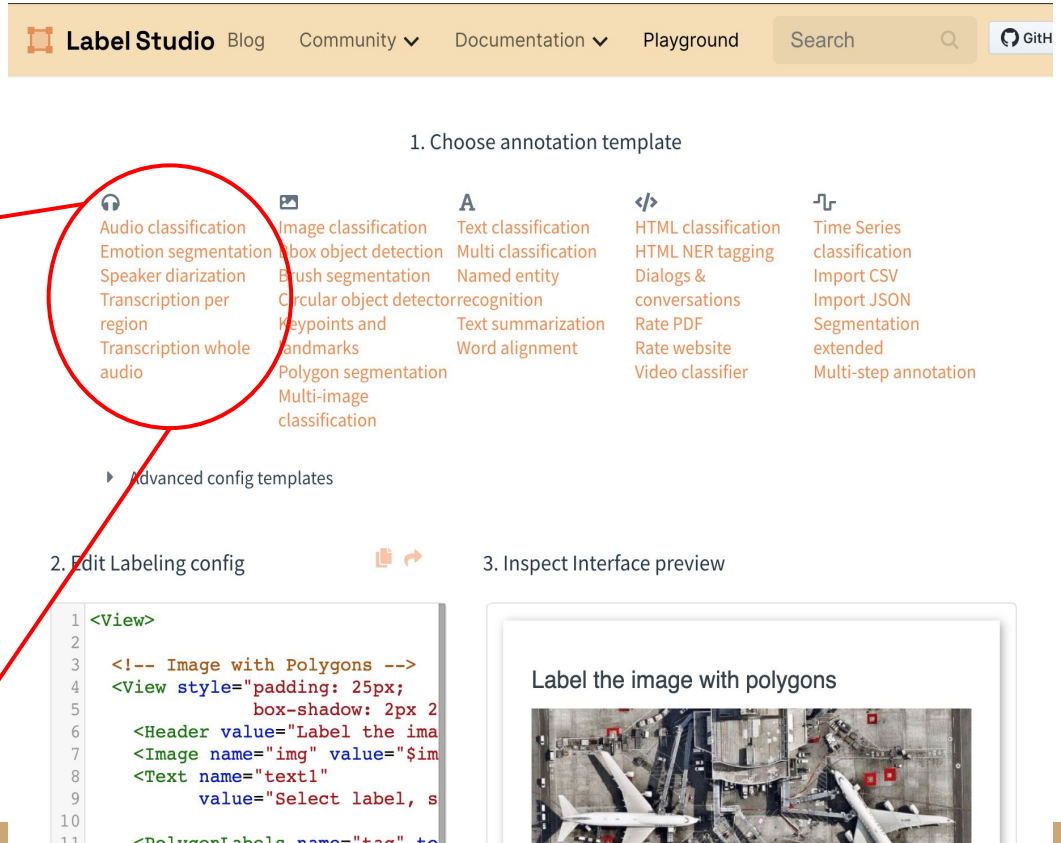

Audio classification
Emotion segmentation
Speaker diarization
Transcription per region
Transcription whole audio

III. Extra Cool - Audacity DL Models!

- Download models off 🤗

  (& upload your own first!)

- Run model as audio effect

- So far only mono, no knobs

Note: This is a custom Audacity build, download from here

Help them add sliders!

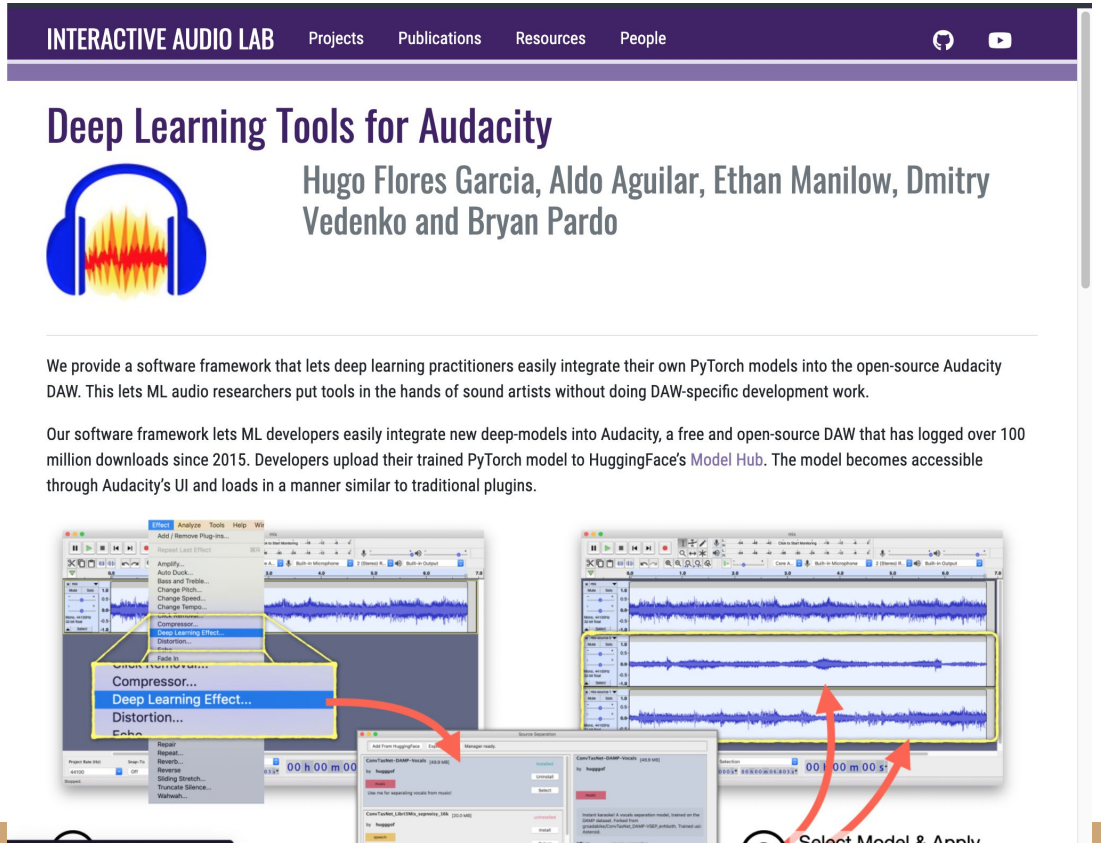

III. MPF tools like Essentia(.js)

Similarity

Analyze audio and compute features to find similar sounds or music tracks.

Classification

Classify sounds or music based on computed audio features.

Deep learning inference

Use data-driven TensorFlow models for a wide range applications from music annotation to synthesis.

Mood detection

Find if a song is happy, sad, aggressive or relaxed.

Key detection

Find a key of a music piece.

Onset detection

Detect onsets (and transients) in an audio signal.

Segmentation

Split audio into homogeneous segments that sound alike.

Beat tracking

Estimate beat positions and tempo (BPM) of a song.

Melody extraction

Estimate pitch in monophonic and polyphonic audio.

Audio fingerprinting

Extract fingerprints from any audio source using the Chromaprint algorithm.

Cover song detection

Identify covers and different versions of the same music piece.

Spectral analysis

Analyze spectral shape of an audio signal.

Loudness metering

Use various loudness meters

Audio problems detection

Identify possible audio quality

Voice analysis

Voice activity detection and

Synthesis

Analyze, transform and

III. A Plug for nbdev

Python library development via Jupyter notebooks

By Jeremy Howard & fast.ai crowd, but doesn't require fastai

is "literate programming" instantiated: code, docs & tests *are one*

Built-in CI via GitHub Actions

Hawley & Morrison (JASA-EL 2022) found it ***essential* for** working efficiently (**staying sane**)

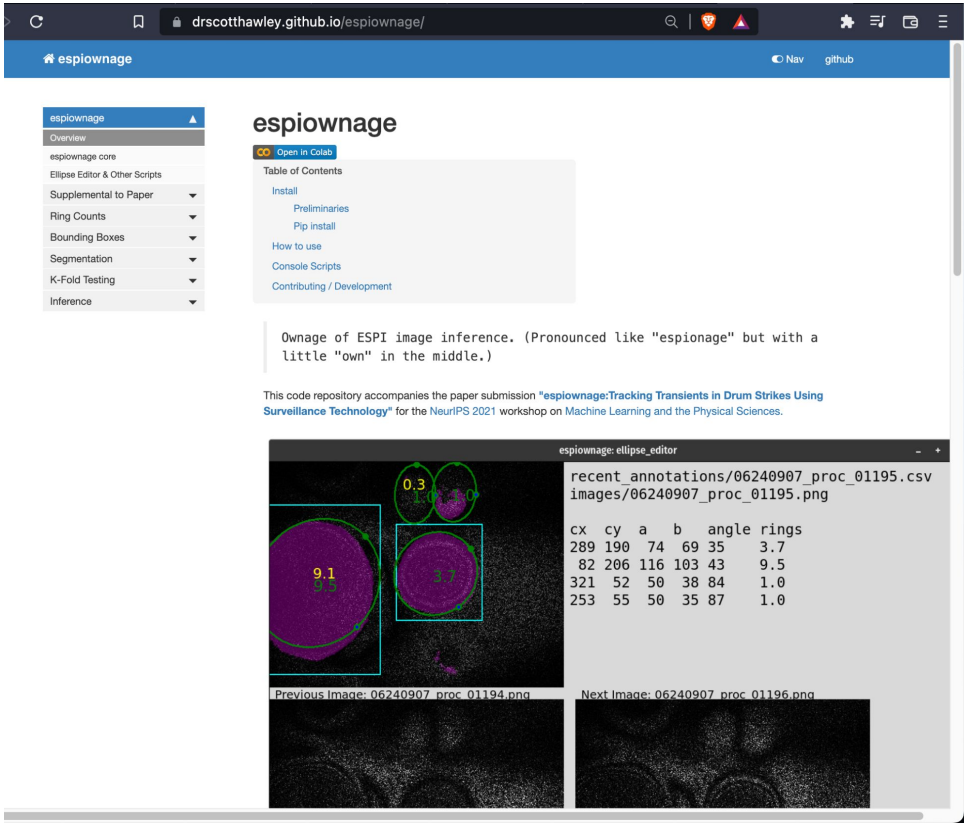

III. WE ALL SHOULD BUILD MORE

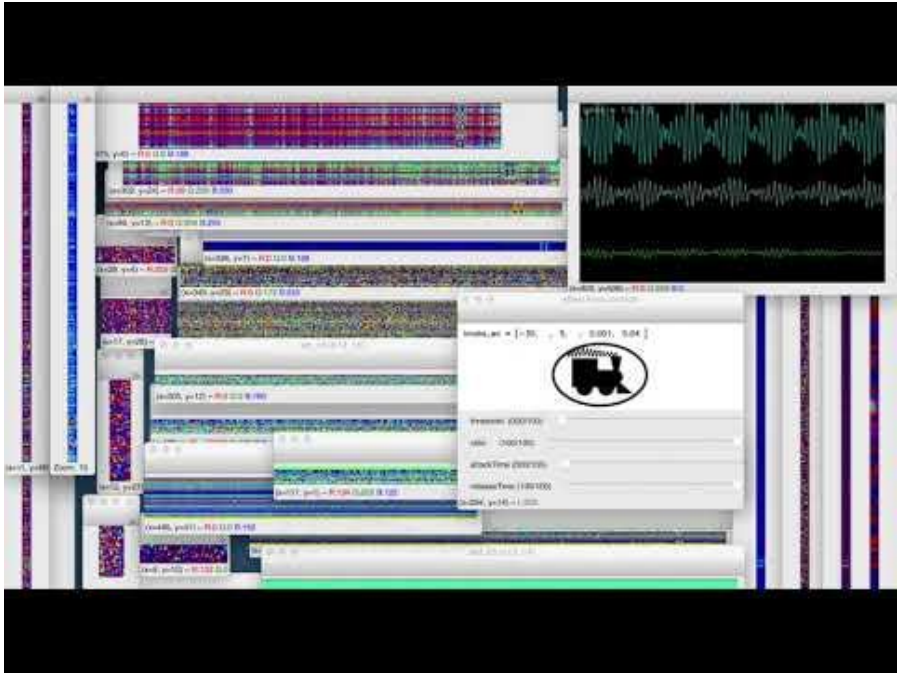

My attempt at interactive viz of PyTorch layer activations: images + oscilloscope.
To some student: *Please* fork this and make it *your own* & make it *good*.

III. Post-Talk: New One! Neutone by Quosmo

"AI audio plugin & community, Bridging the gap between AI research and creativity"

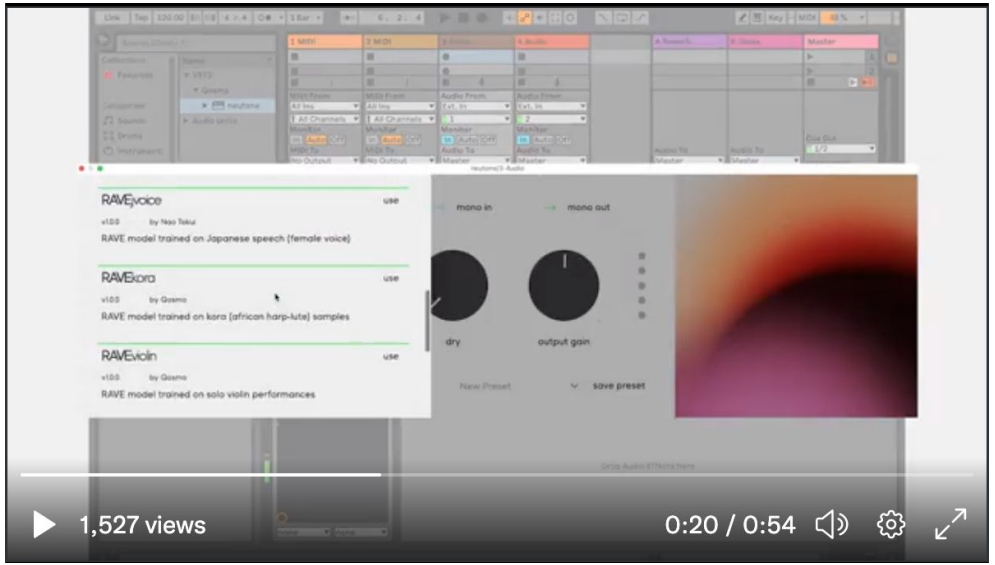

For audio creators

Neutone makes AI technologies accessible for all to experiment with. You'll find transformative AI audio instruments that will spark endless creative possibilities.

For AI researchers

Neutone is a go-to platform for you to share real-time AI audio processing models with potential users in the audio production community.

Download Plugin

Submit Your Model