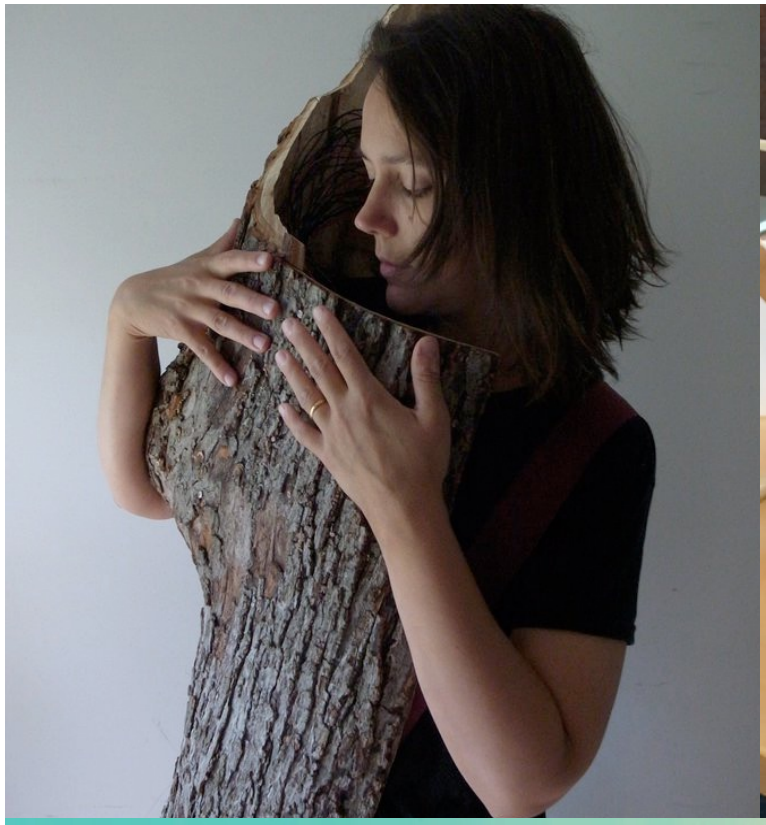
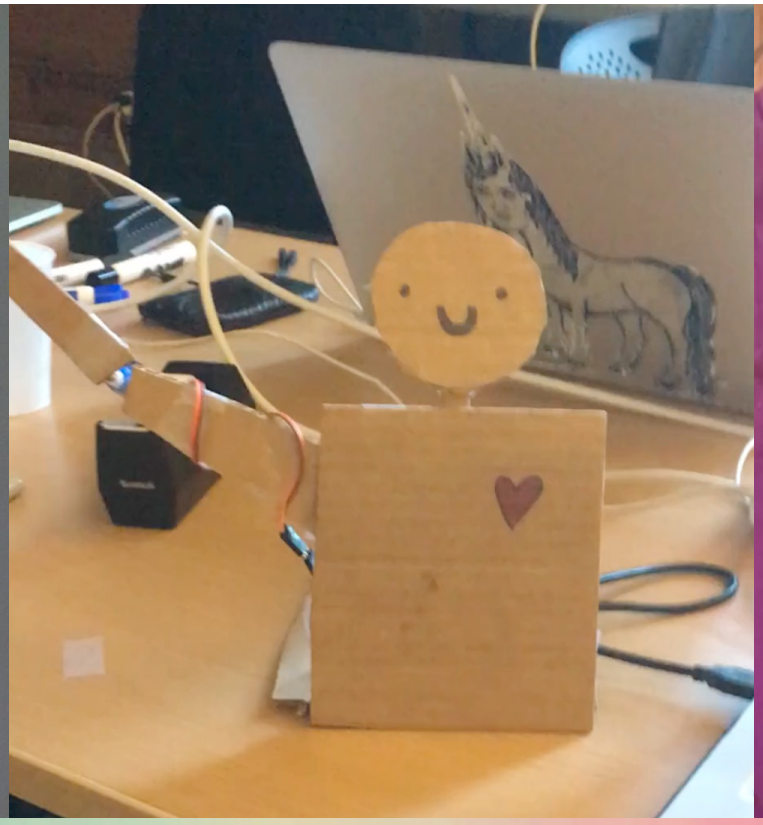
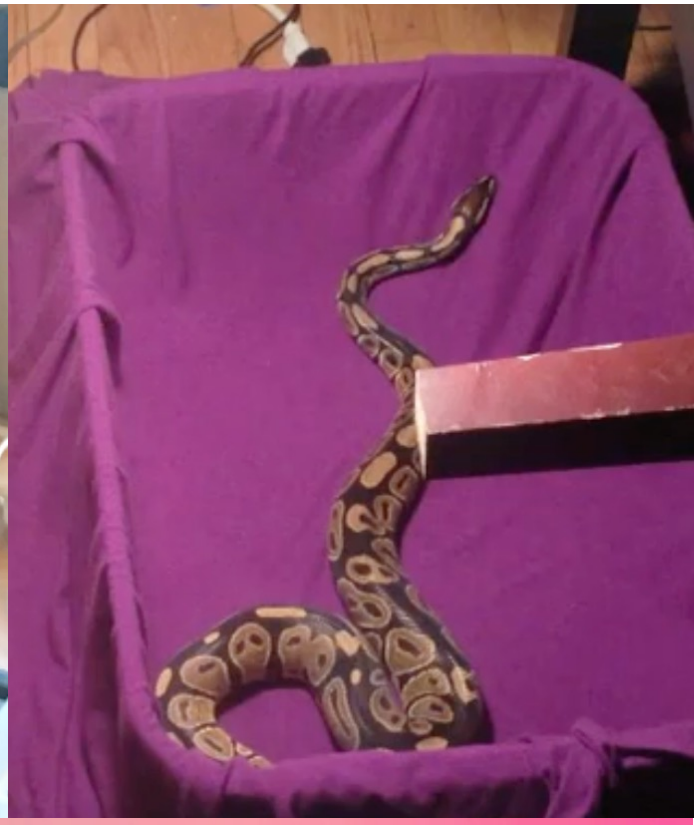

Teaching AI to Creators

Professor Rebecca Fiebrink
Creative Computing Institute, University of the Arts London

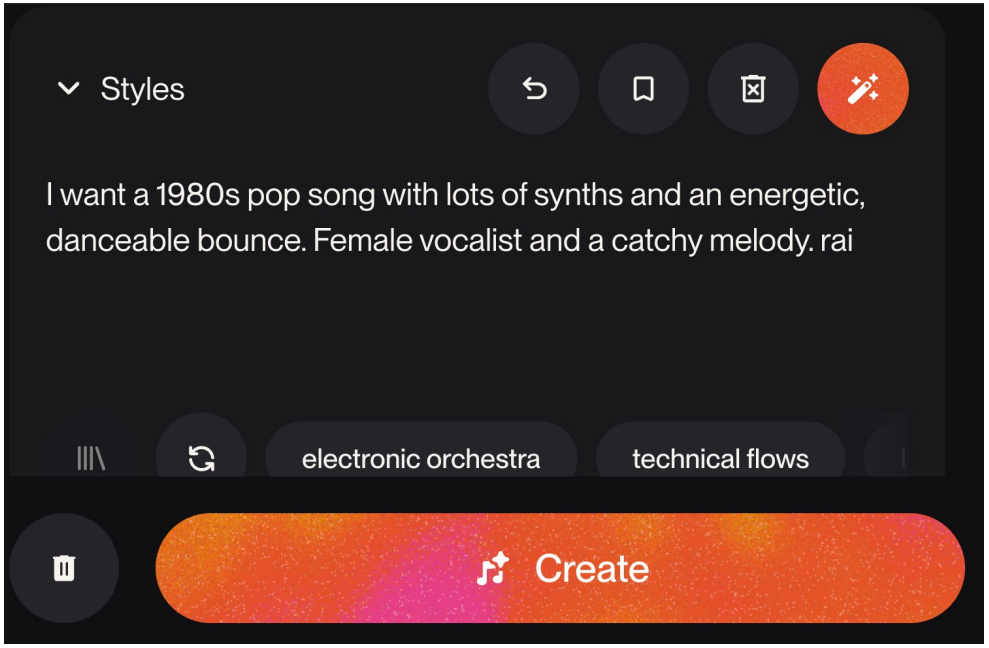

AI in creative work can be—and is—so much more than this!

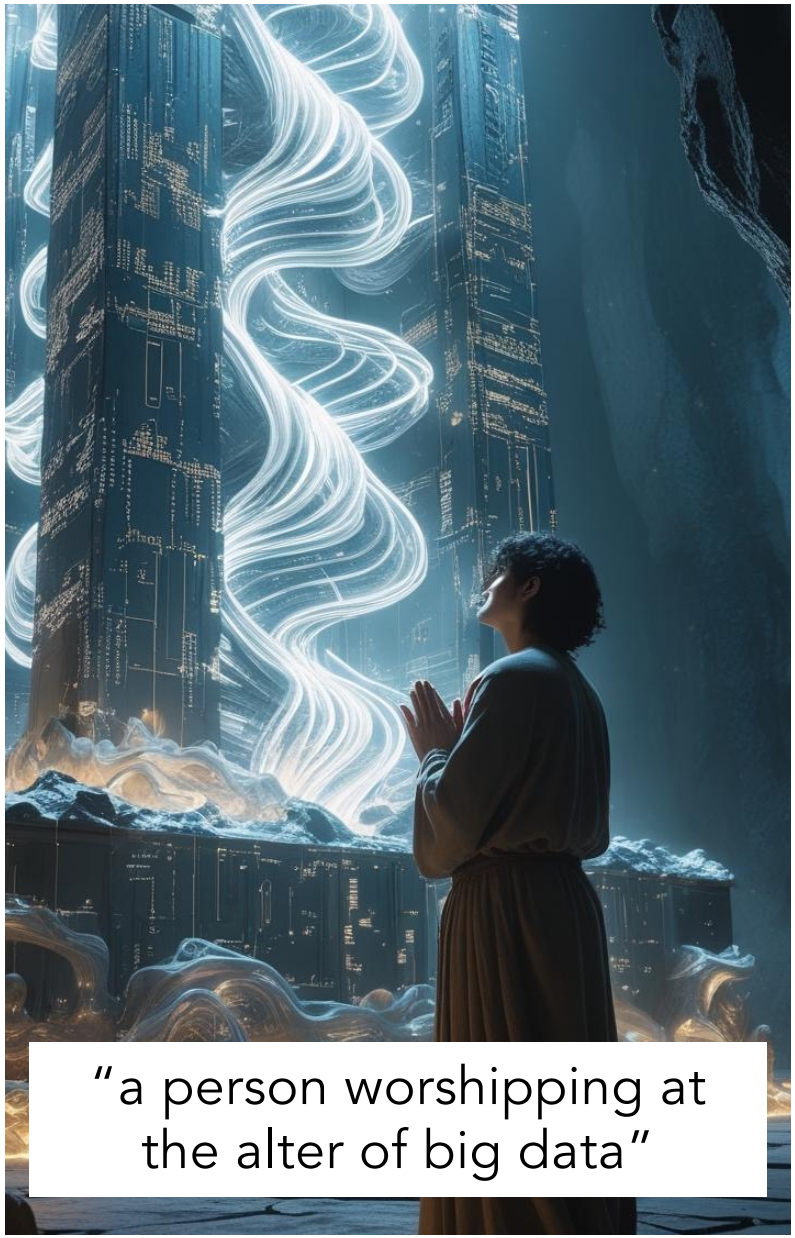

Machine learning is a set of computational tools for

finding patterns in data

and using these patterns to

make **predictions** or **generate** new data

Key questions in my research since ~2008

How can ML—finding and using patterns in data—be of value to musicians, artists, and other creators?

What should ML tools for creators look like?

What do creators need to know about ML to use it effectively, and how can we teach them?

How can ML be of value to musicians and artists?

What might ML tools for creators look like?

1. Machine learning can help creators work more effectively with data they care about

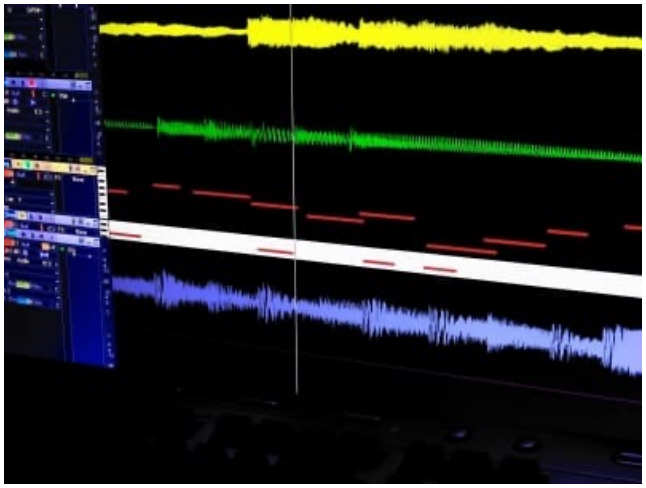


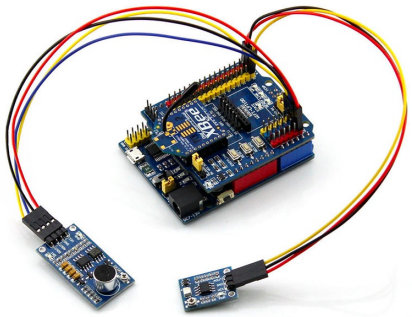


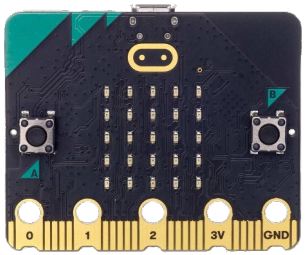


Potential sources of interesting, useful data are are everywhere

Getting this data is often easy; interpreting & using it is hard!

(even for expert programmers)

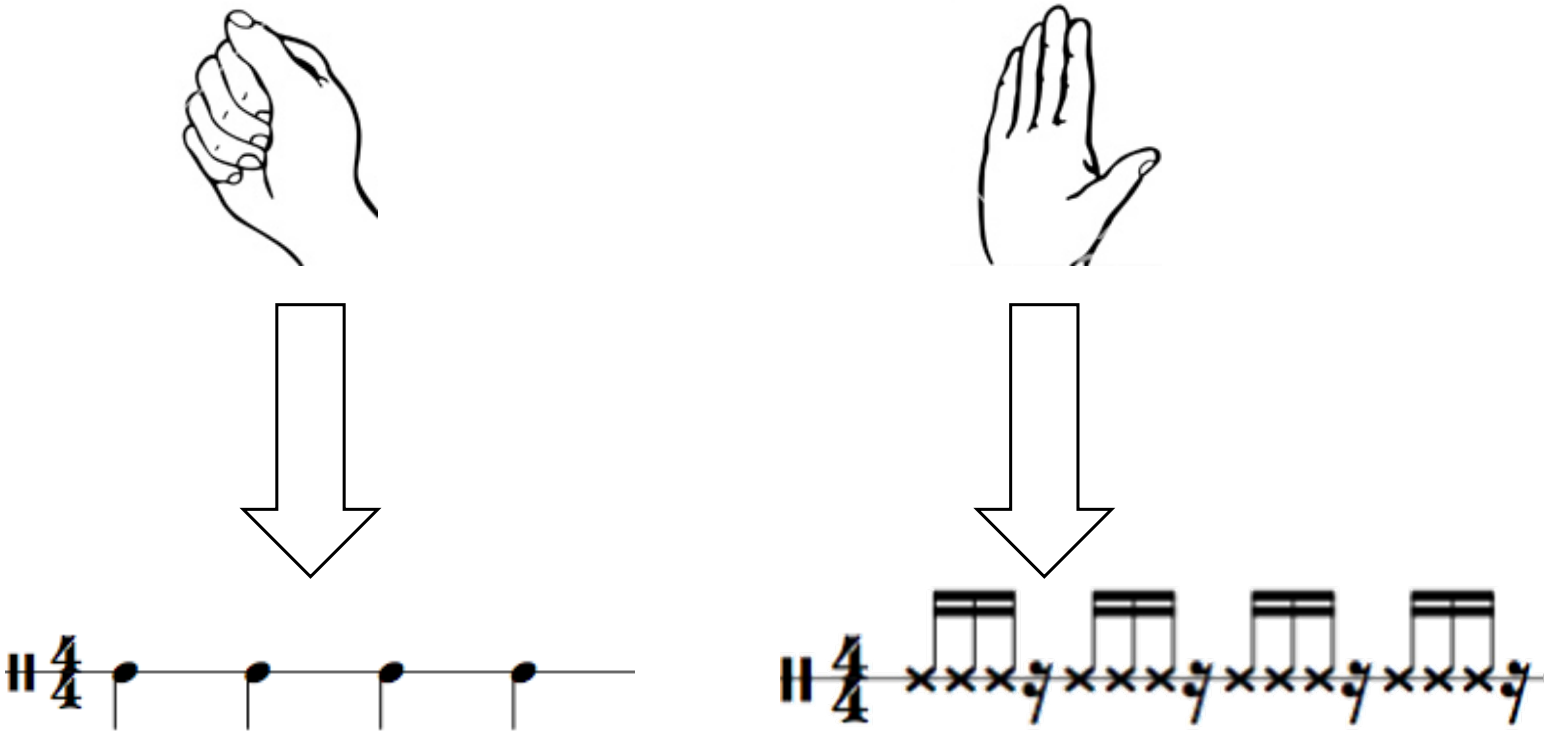

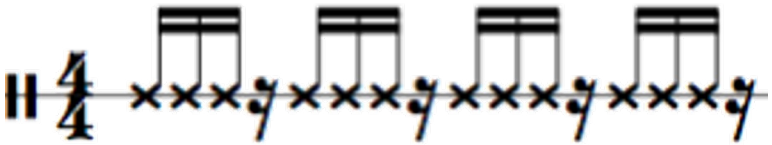

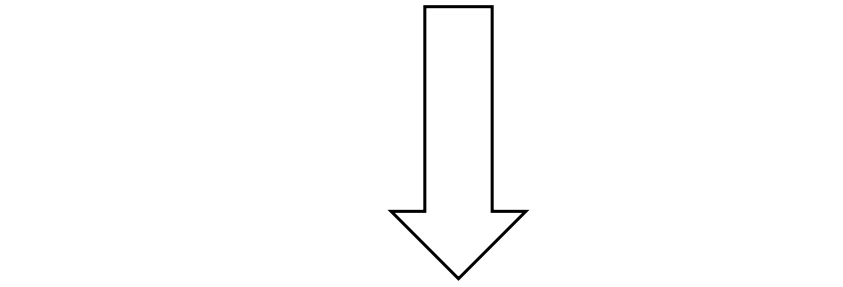

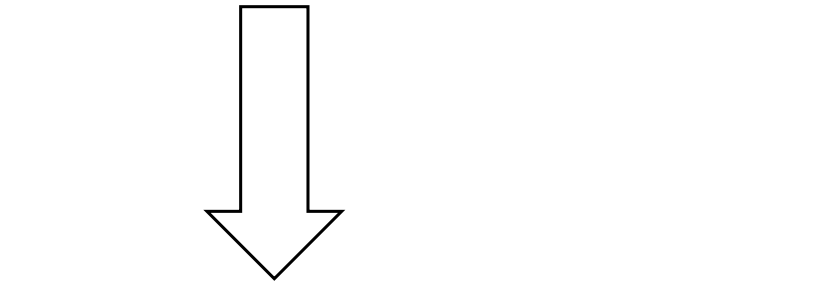

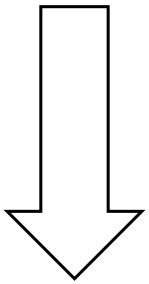

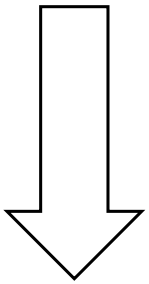

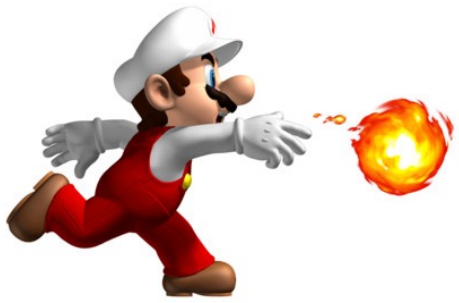

We can use **supervised learning** to infer relationships between human actions and sound / visuals / gameplay / other computer actions

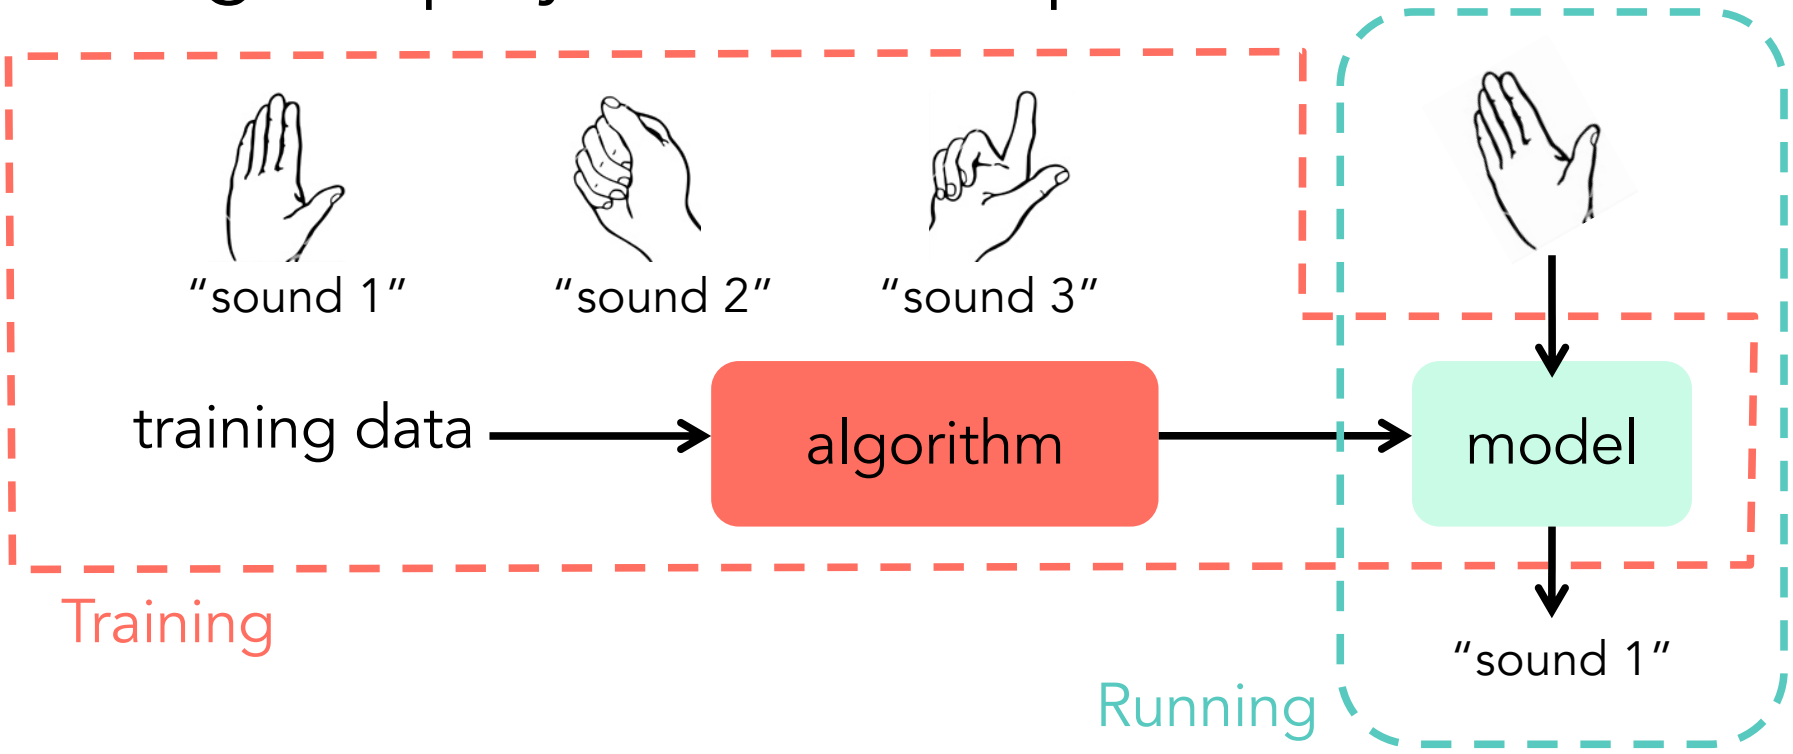

**Making this accessible to creators requires more than a friendly GUI!*

The Wekinator (2008+):

"Interactive machine learning"

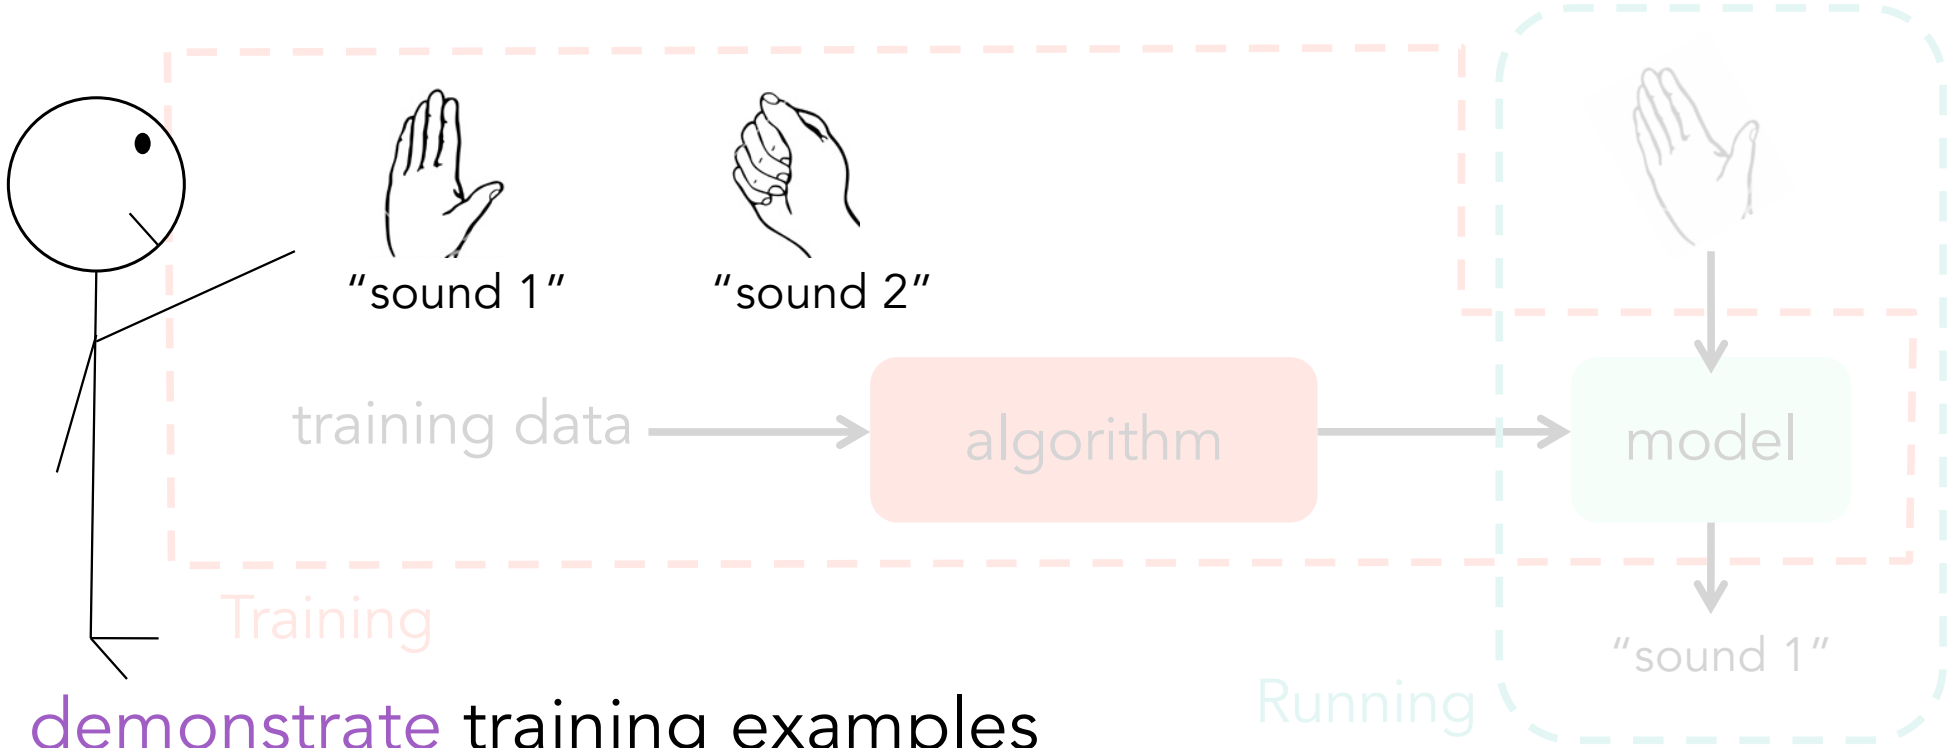

demonstrate training examples

The Wekinator (2008+):

"Interactive machine learning"

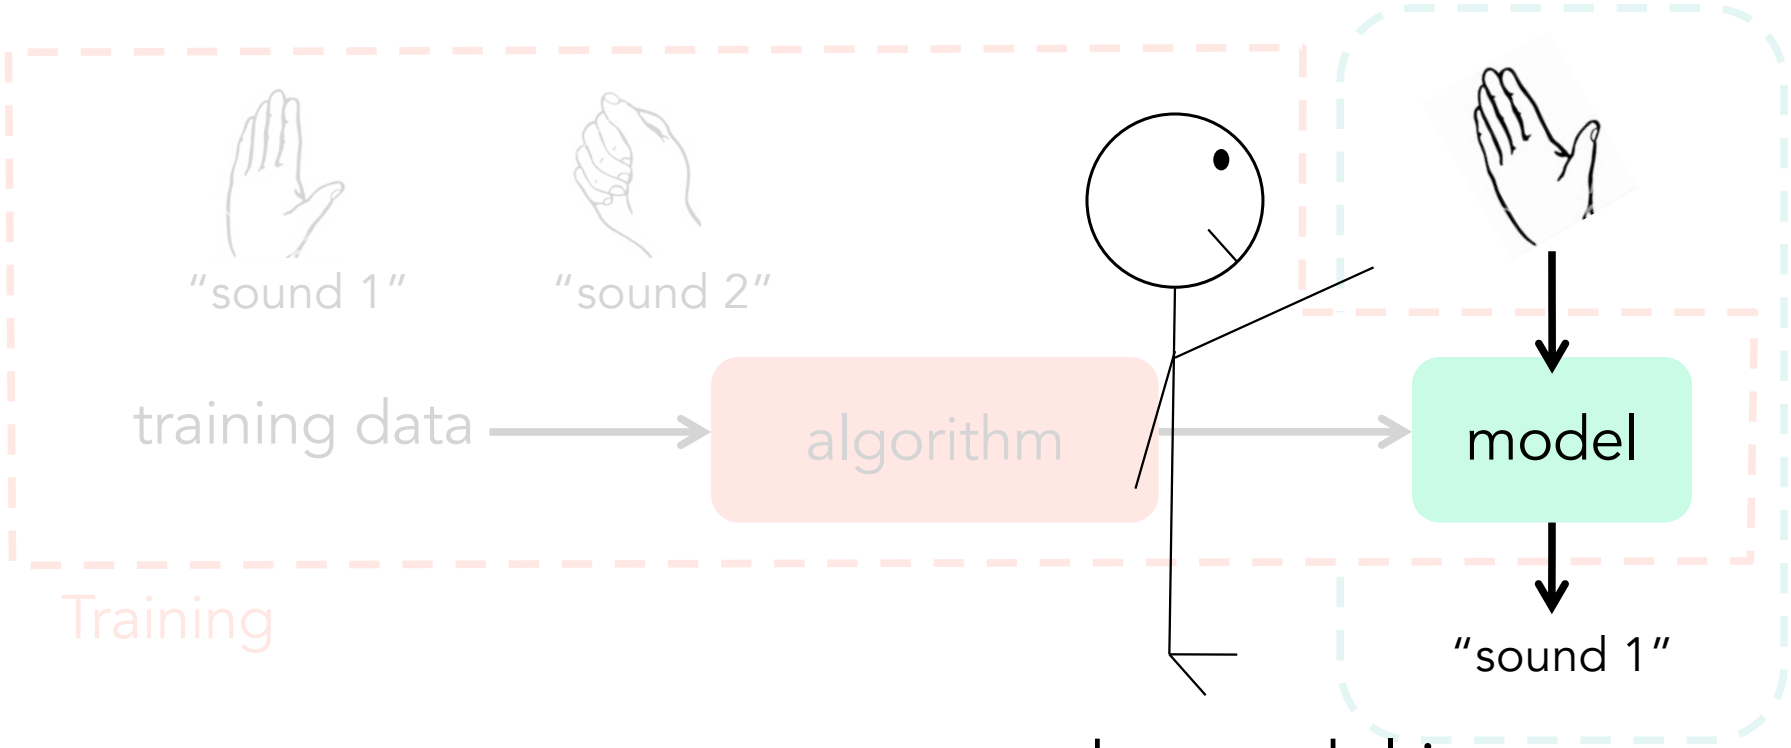

test out the model in real-time

The Wekinator (2008+):

"Interactive machine learning"

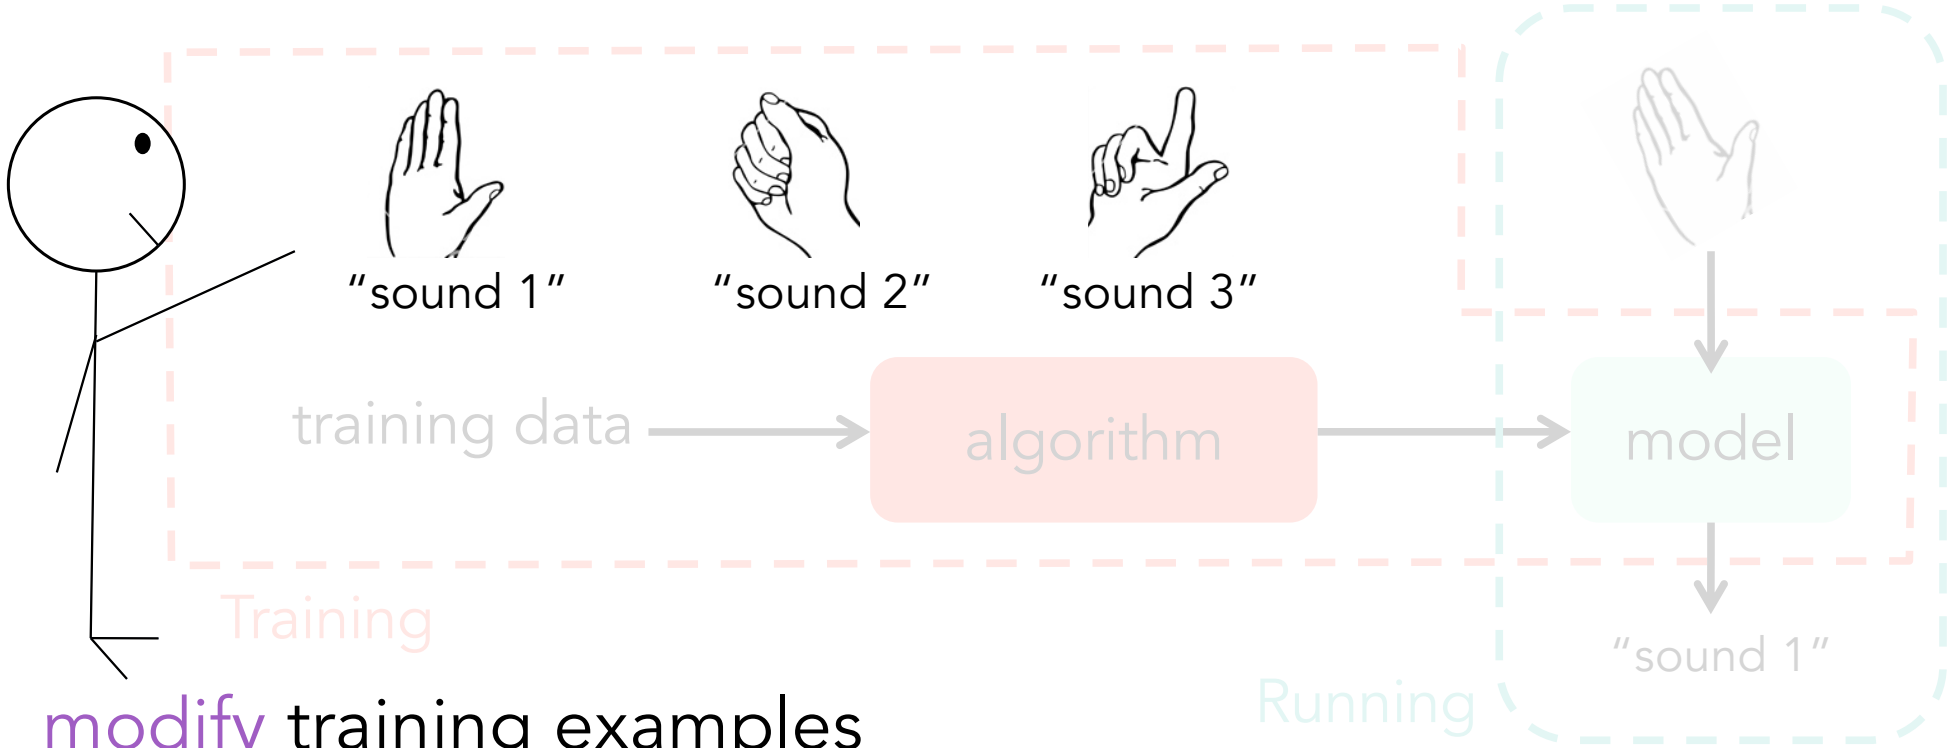

modify training examples
(and repeat)

Wekinator Demo

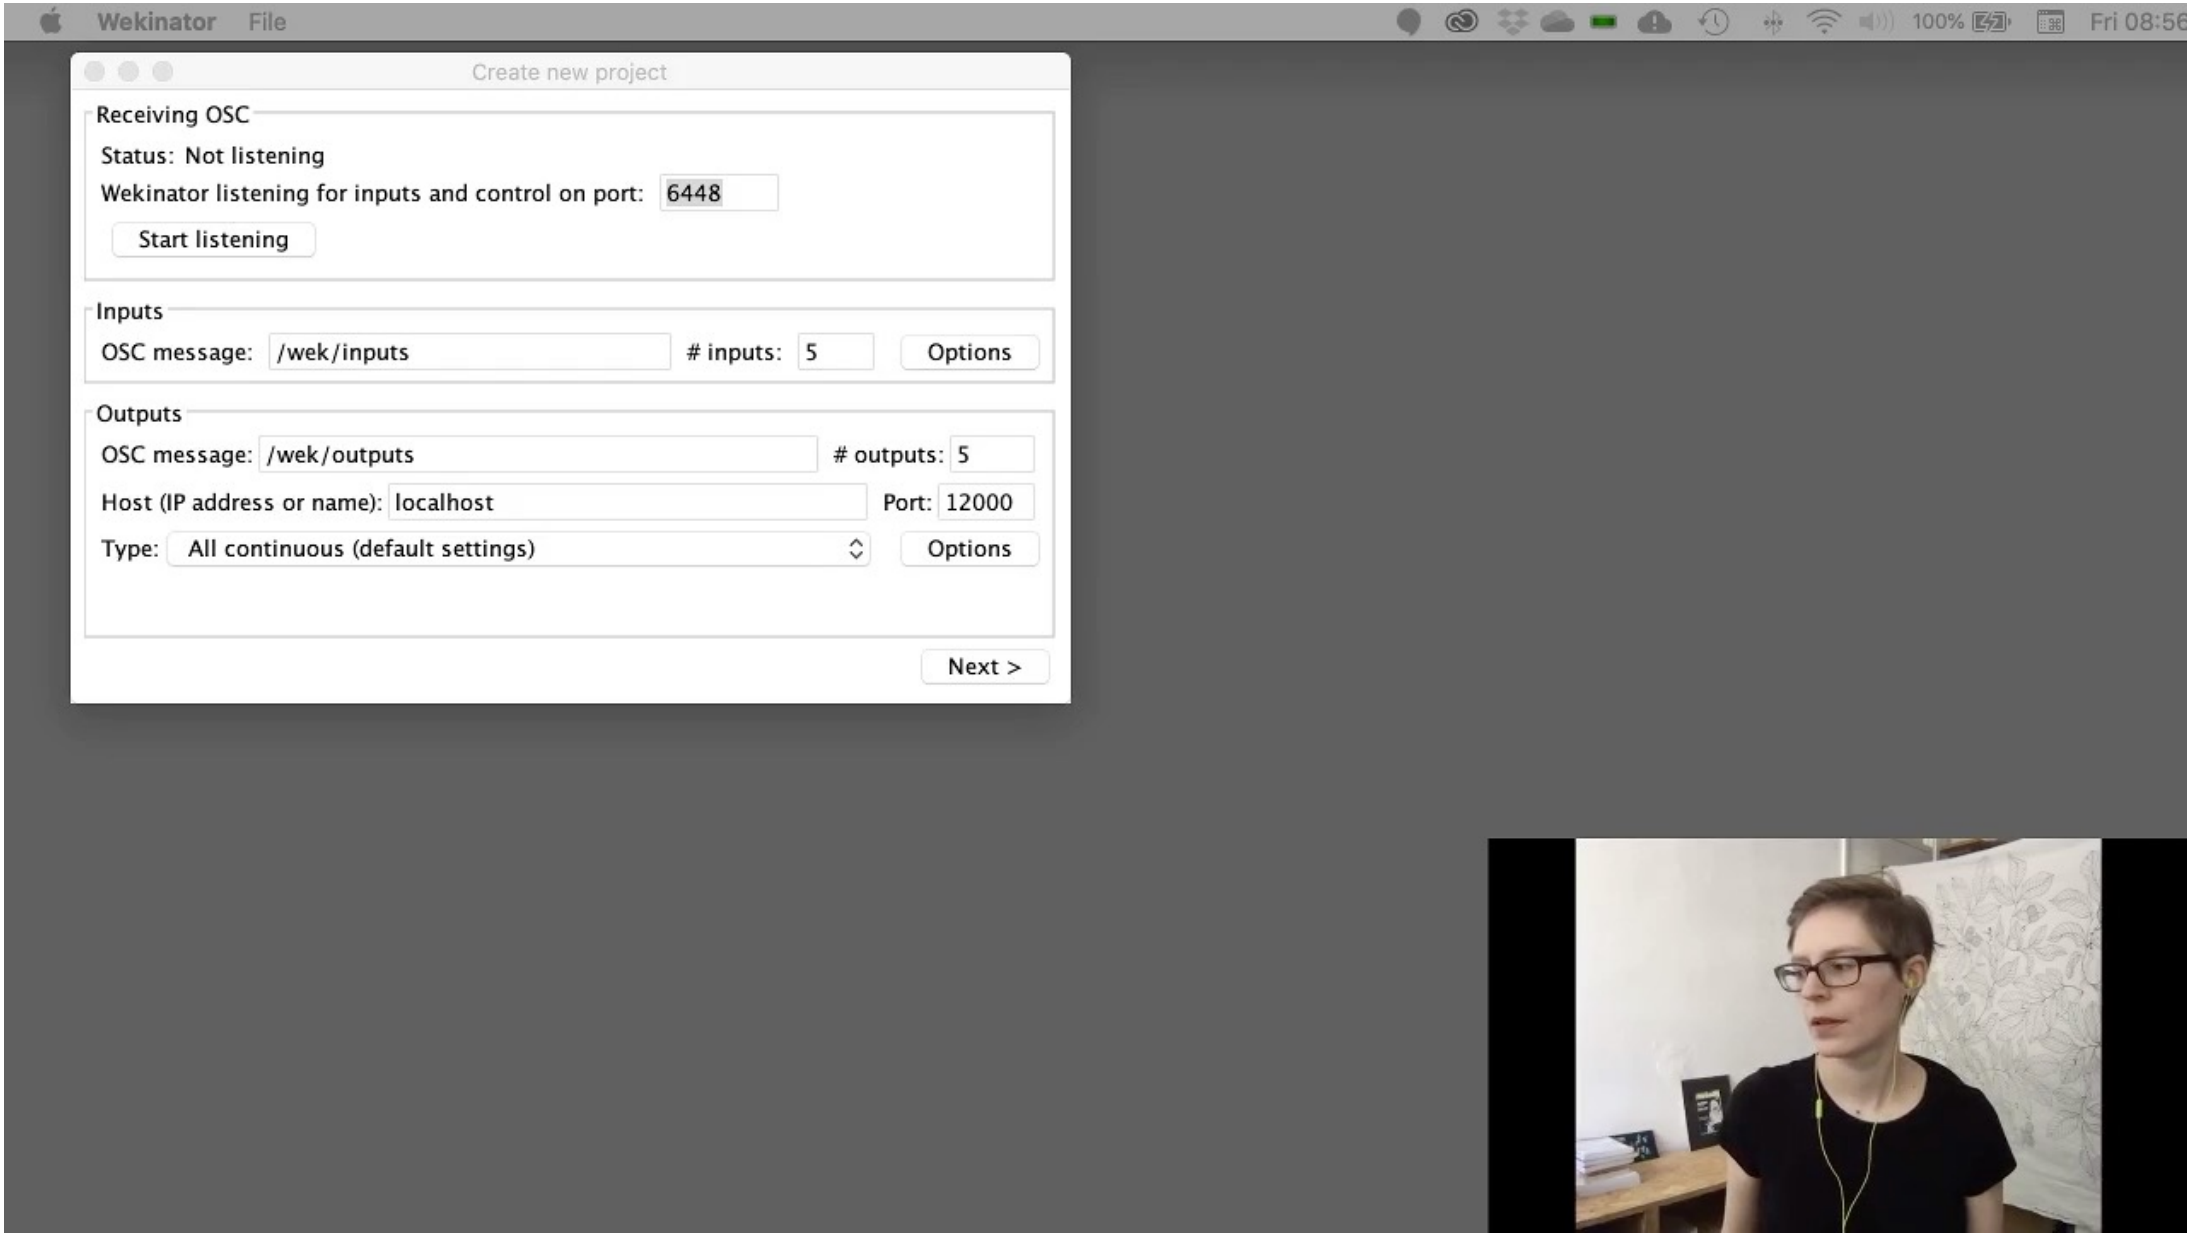
Wekinator File
100%
Fri 08:56
Create new project
Receiving OSC
Status: Not listening
Wekinator listening for inputs and control on port: 6448
Start listening
Inputs
OSC message: /wek/inputs
inputs: 5
Options
Outputs
OSC message: /wek/outputs
outputs: 5
Host (IP address or name): localhost
Port: 12000
Type: All continuous (default settings)
Options
Next >

Wekinator since 2008

- 50,000+ users
- Used in world's first MOOC on ML for musicians and artists (on Kadenze), and in artist workshops worldwide (e.g., at Barbican, BBC Radio 1 Academy)
- Used in teaching at UAL, Stanford, NYU, CalArts, Columbia, Creative Coding Amsterdam, many other universities and institutions

Anne Hege, From the Waters


Michelle Nagai, MARtLET

MACHINES

Wanna Stop Staring at Screens? Try on These Glasses

CAMILO SALAS
Jun 1 2016, 4:40pm

The glasses become opaque, effectively blinding you until you look away.

They look like something out of *Black Mirror*, only there's no mechanism or sound or movement behind the scenes. But try on these glasses and every time you look at a screen your vision gets blocked.

In other words, they're glasses that screen screens, according to Chino Kim, the 28-year-old behind Screeners, a project that began as a personal quest to rid his life of screens.

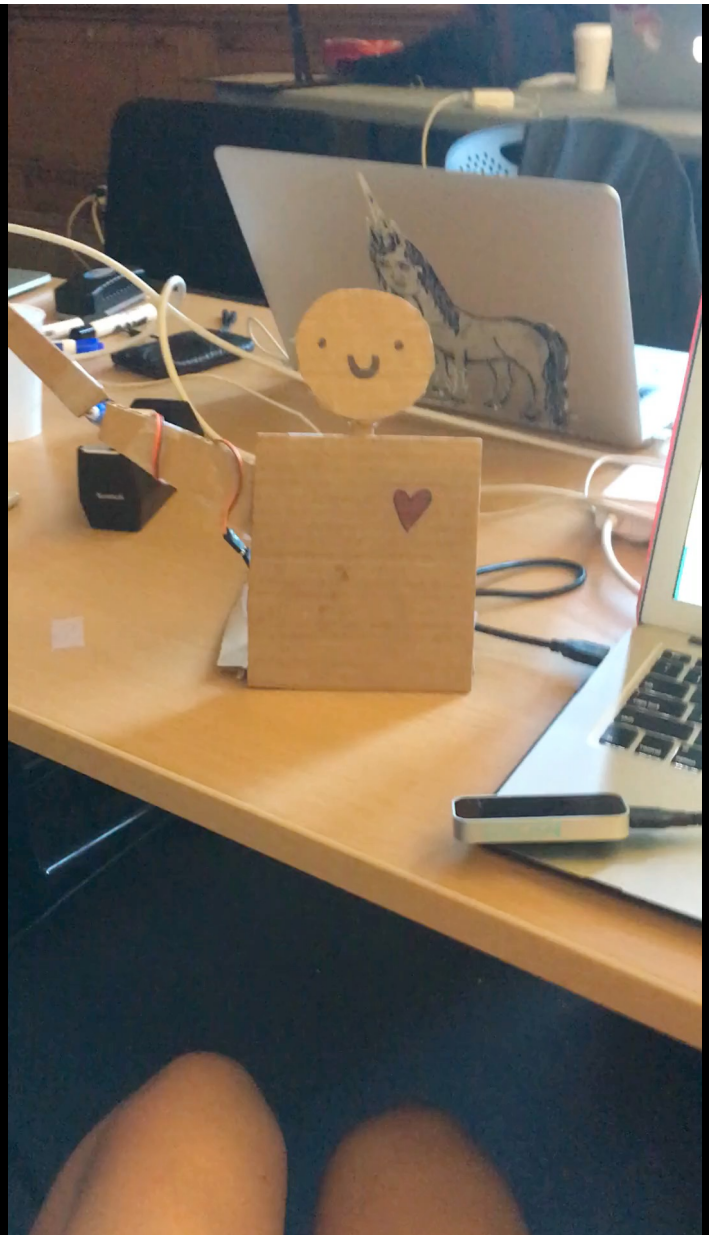

Other Tools for Wekinator-style Interactive ML

By my team:

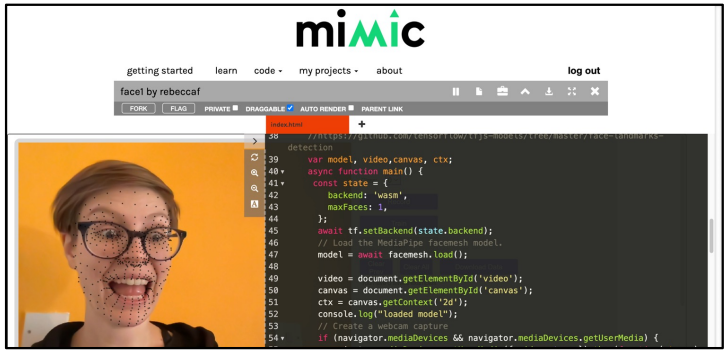

For Web + creative coding

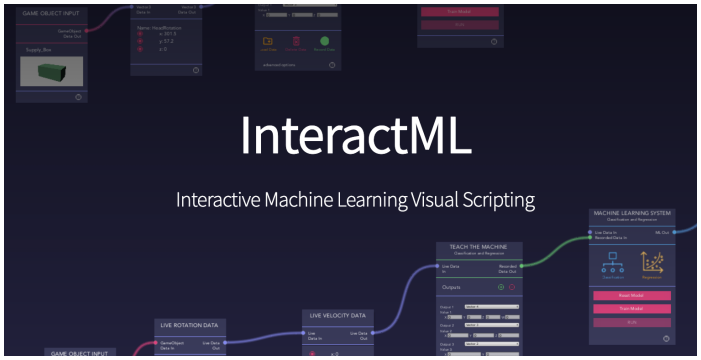

For VR + Games

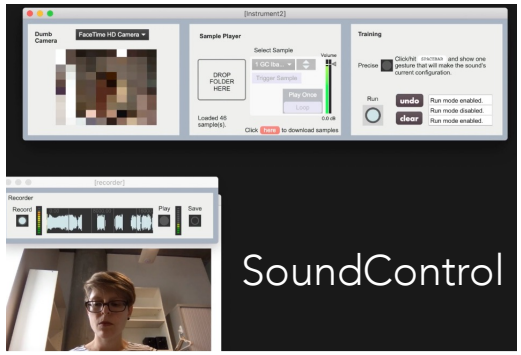

For kids, music teachers & therapists

By others:

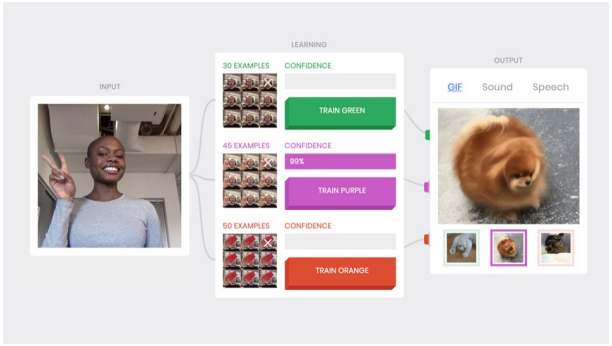

Teachable Machine

Micro:bit CreateAI

MiMu gloves

2. Machine learning allows people to **communicate** ideas and intentions to computers through **data/examples**


Michelle Nagai, MARtLET

Memo Akten: Learning to See

Akten, Fiebrink, and Grierson, SIGGRAPH Art Papers 2019
https://www.memo.tv/works/learning-to-see/

R-VAE

Gabriel Vigliensoni

- Construct a generative model from a very small dataset (e.g. dozens of examples) in a small set of genres/styles
  - Models are quite specific!
- Provide a performance interface for navigating a 2D representation of the generative model's "latent space" in realtime

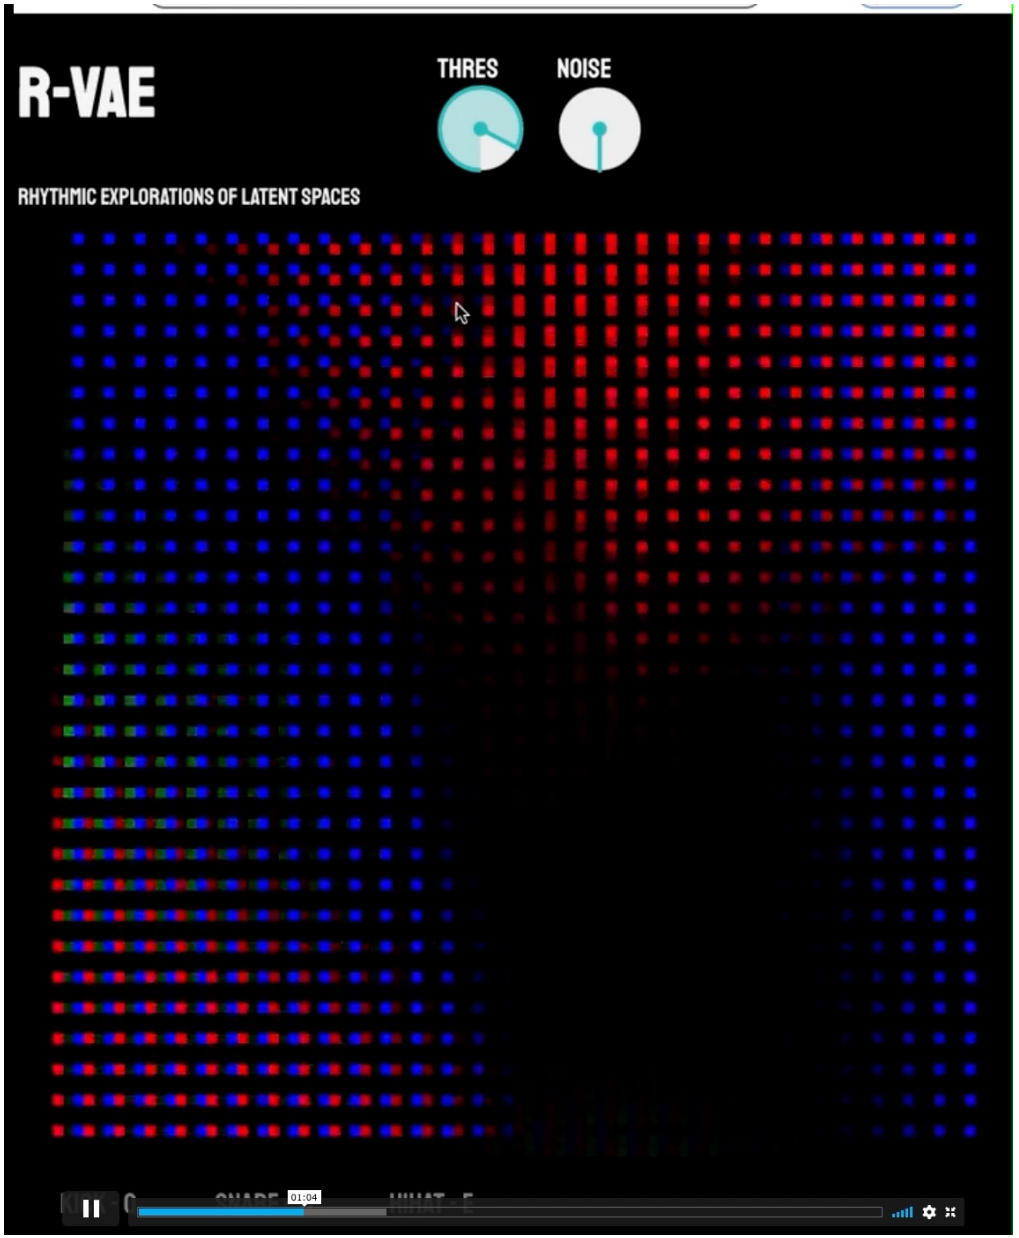

https://github.com/vigliensoni/R-VAE

Vigliensoni et al. 2022. Journal of Creative Music Systems

vs

```
int note_package[NUM_VOICES][2];
int cmd_package[NUM_COMMANDS][2];
int drum_package[NUM_DRUMS][2];
// Whether or not to process packages in the given timeslot
int should_process_packages;

// Prepare note / command / drum packages for
// this timeslot across all 8 voices
0 => should_process_packages;
for (0 => int i; i < master_loop_num_tracks(); i++) {

    master_loop[i][master_loop_index] @=> int noteslot[];
    master_loop[i][master_loop_index][0] => int pitch_copy;
    master_loop[i][master_loop_index][1] => int duration_copy;

    if (!is_empty(noteslot)) {
        if (noteslot[0] > 0) {  // melodic note
            noteslot[0] => note_package[i][0];
            noteslot[1] => note_package[i][1];
            1 => should_process_packages;

            // Consume the note.
            0 => master_loop[i][master_loop_index][0];
            0 => master_loop[i][master_loop_index][1];
            // Do note tracking logistics.
            // When we consume a note, we should
            // increment the notes_added starting
            // index so we don't waste time trying
            // to zero that note out on a "stop"
```

Data/examples can be better than math, code, or language for communicating tacit knowledge and embodied practices to a computer.

3. Designing with ML can enable more people to become creators

Large, one-size-fits-all models can make it easier for novices to conform to a standard

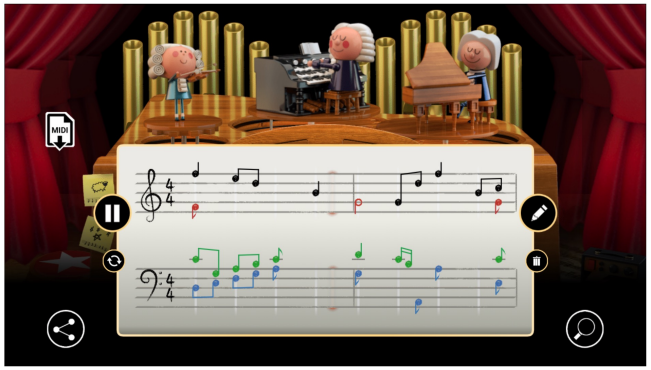

Write music like Bach!

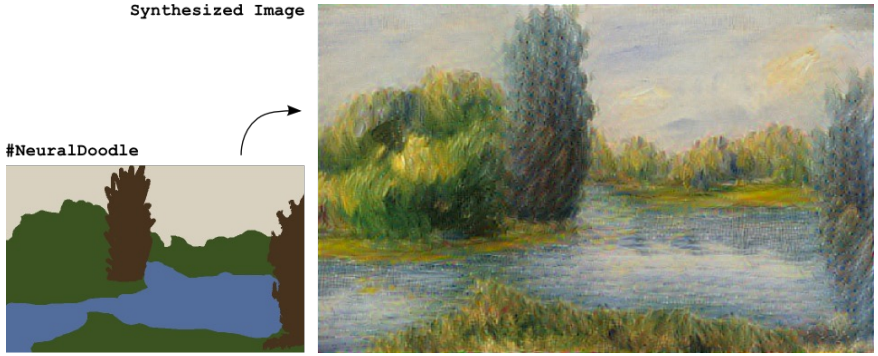

Make an image like Monet!

Make a video like Wes Anderson!

Training new models
with new data
allows easy
personalization and
flexibility

Katan, Grierson & Fiebrink, CHI 2015

Empowering creators to build for themselves enables more useful outcomes, unanticipated by “experts”

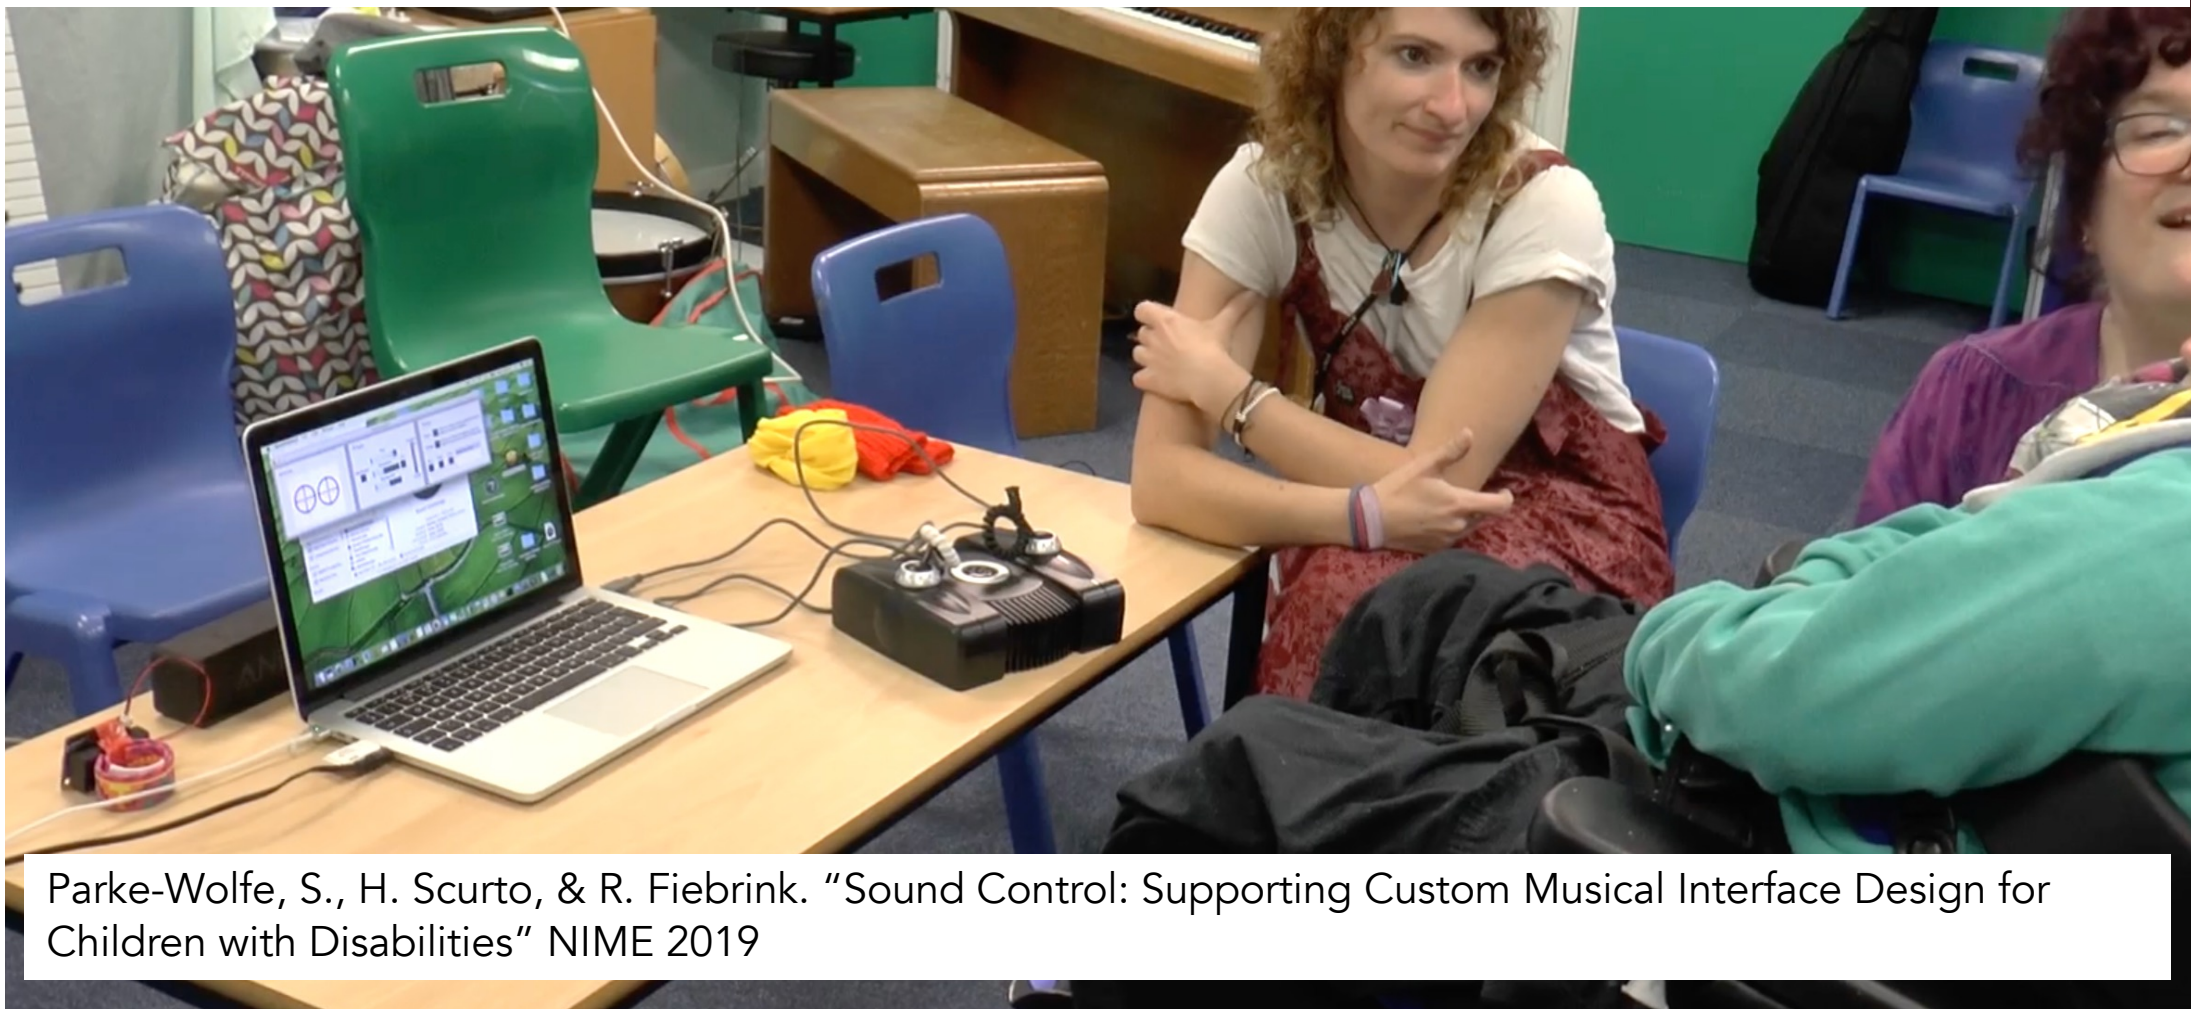

Parke-Wolfe, S., H. Scurto, & R. Fiebrink. “Sound Control: Supporting Custom Musical Interface Design for Children with Disabilities” NIME 2019

Sometimes, helping novices conform to a standard can facilitate other creative tasks

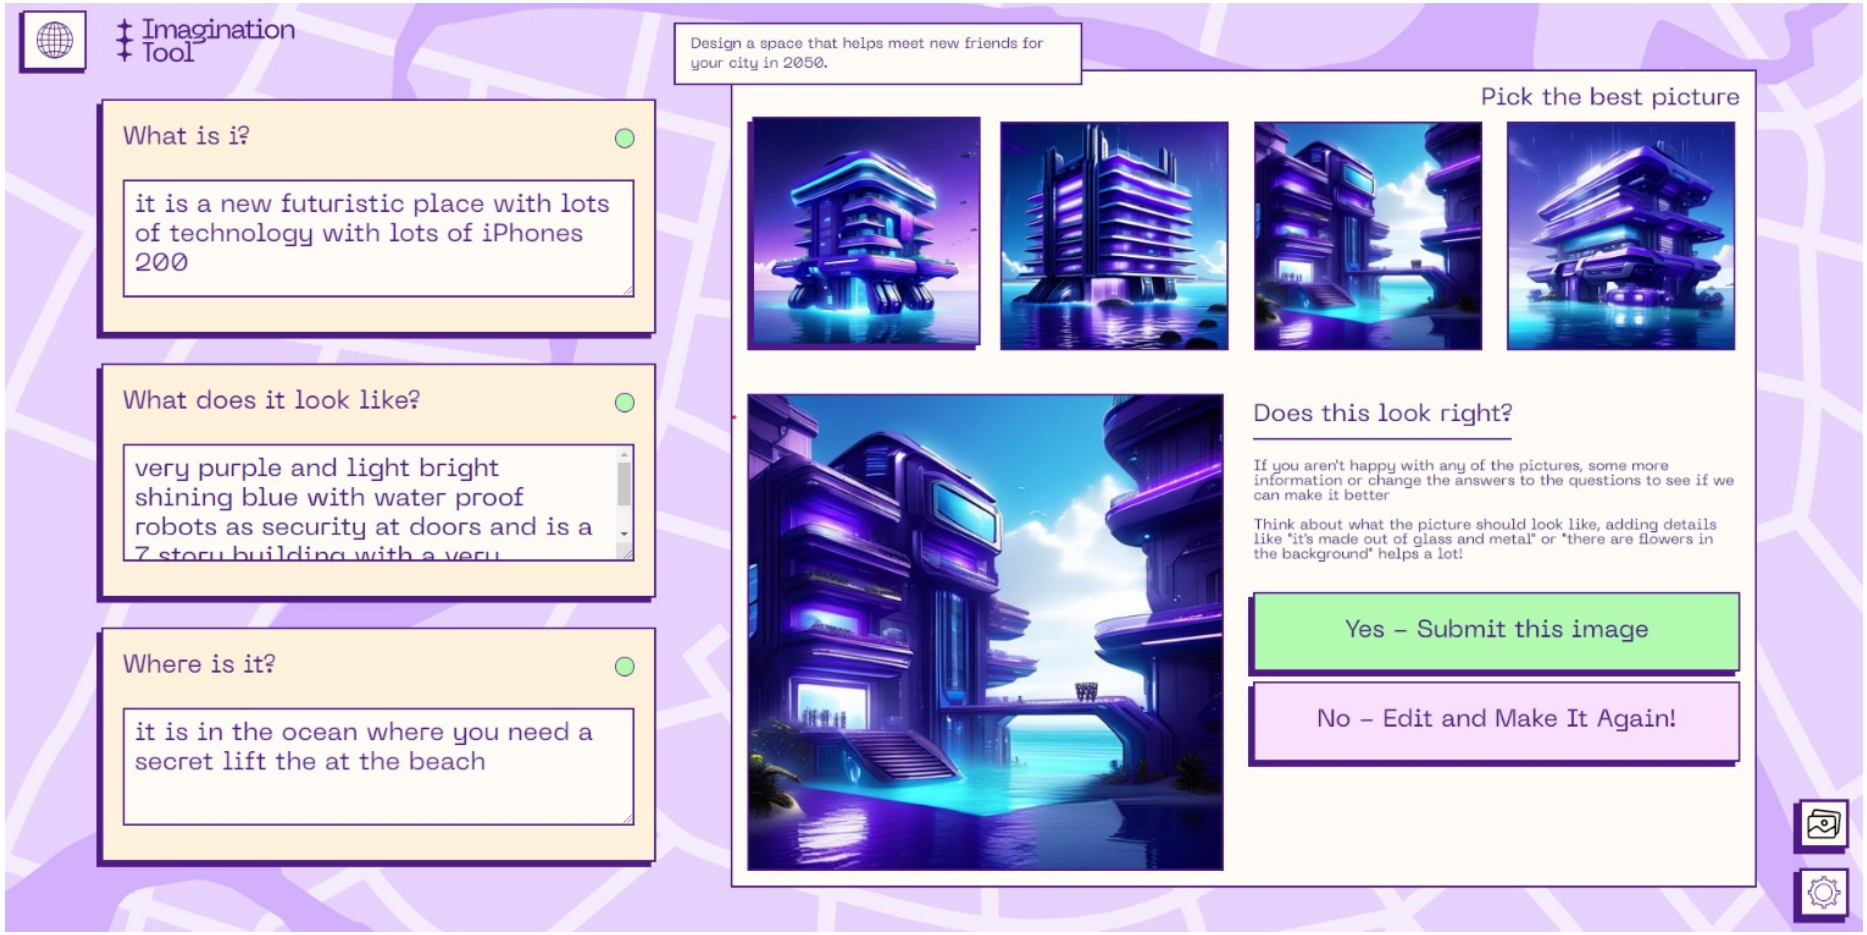

Meriç, White, Suárez Zapico, Yanovich, Koby-Hirschmann & Fiebrink. ICCC 2024.

4. Machine learning can enable exciting new types of creative work

New instruments

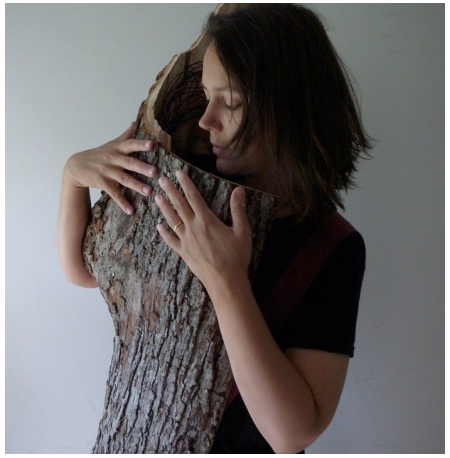

Michelle Nagai

Gabriel Vigliensoni

Marije Baalman

Laetitia Sonami

Chicks on Speed

New performance techniques

New interactions, including with audiences

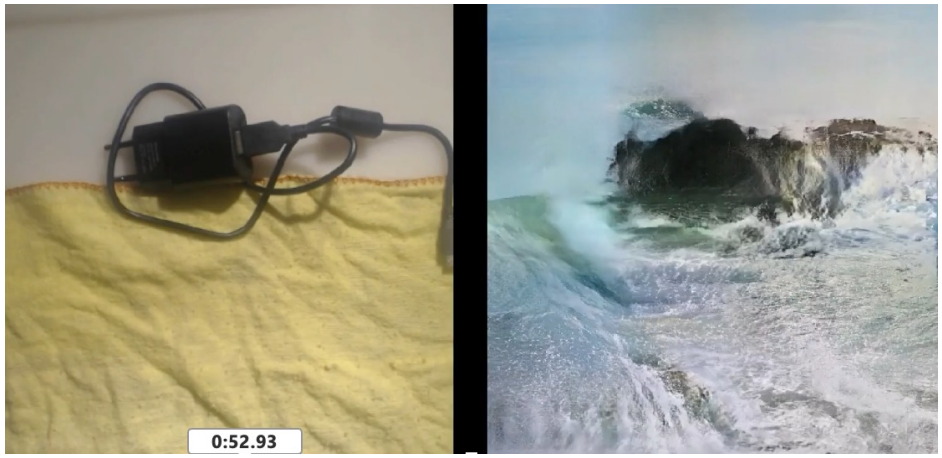

Memo Akten

Snakinator

Control Google maps with food

Stephanie Dinkins

Holly Herndon

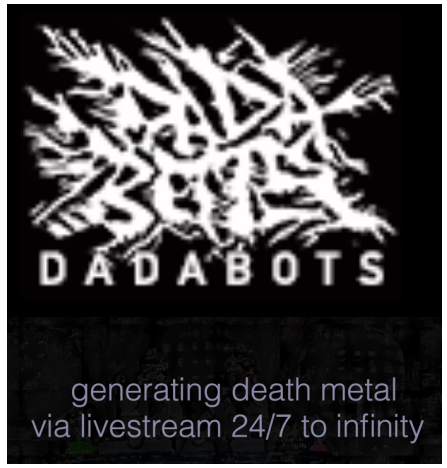

Dadabots

Teaching ML to creators: Principles and challenges

Principle 1:

Teach ML at a high level with minimal math, but not as a black box

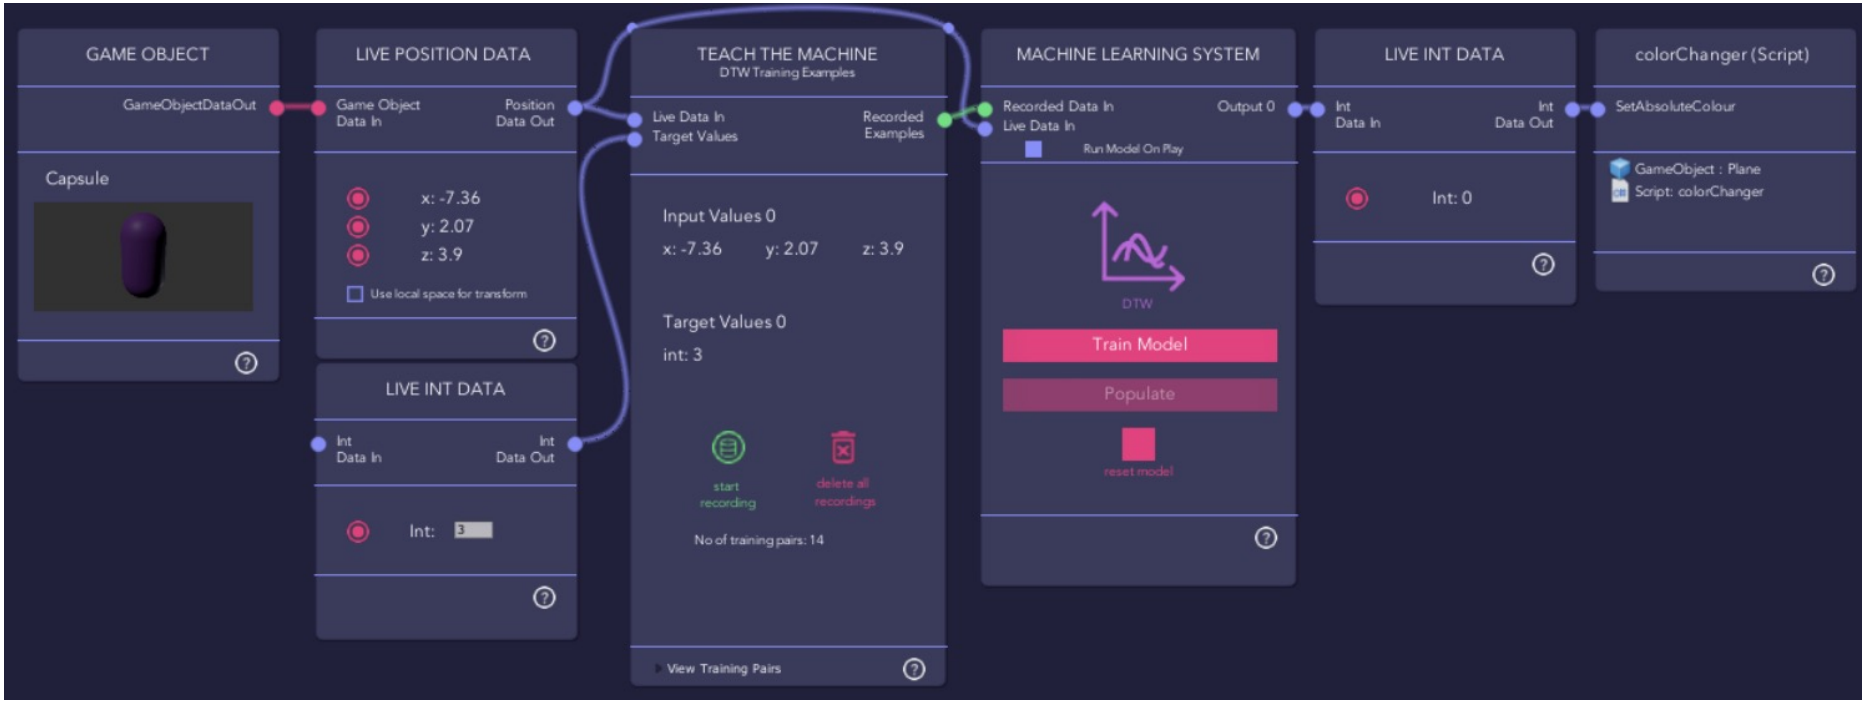

Interactml.com: Interactive ML for VR in Unity & Unreal

Principle 2:
Small data and interactive ML can be very effective

- Quick experimentation
- Creatively useful
- Allows building intuition about big-data systems too
  - E.g., outliers, regularisation, features; bias in data and data as a *choice*

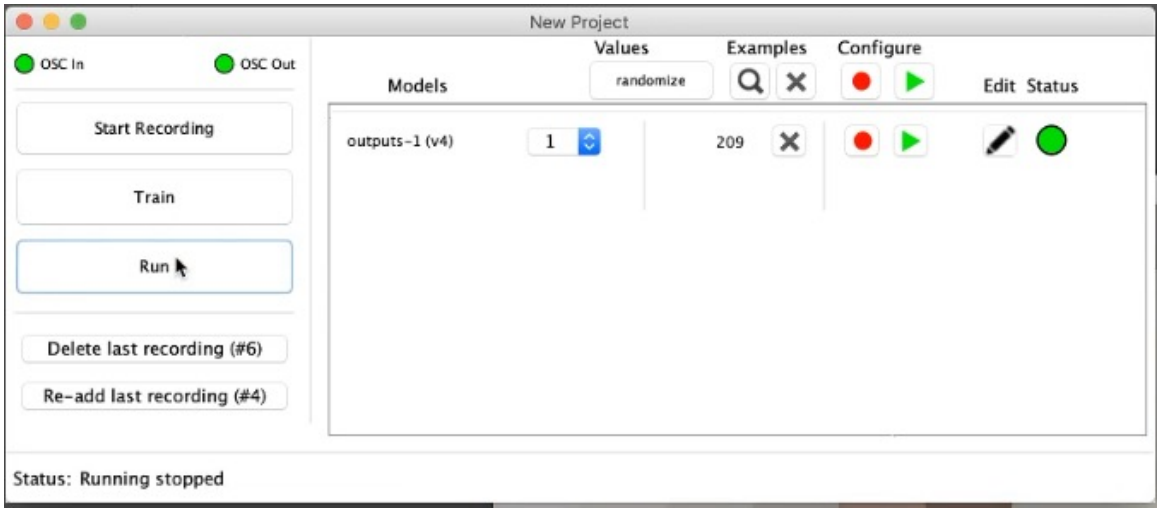

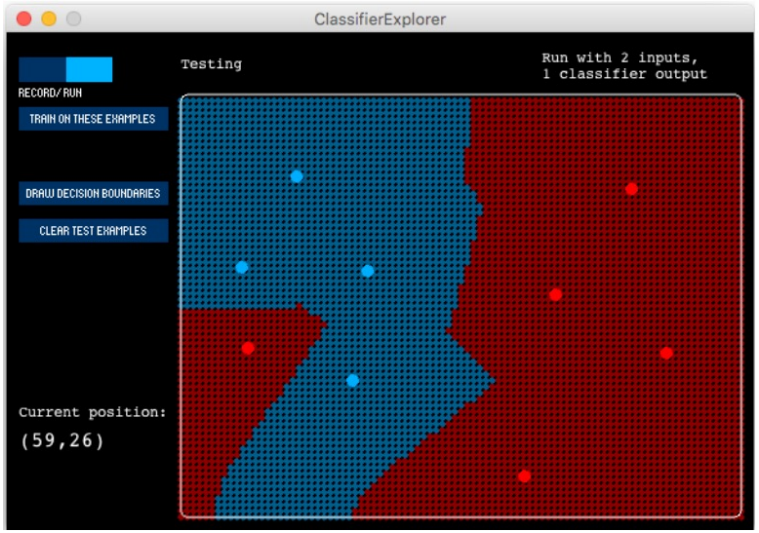

Fiebrink. 2019. "Machine learning education for artists, musicians, and other creative practitioners." ACM Trans. Computing Education.

Principle 3:

Modular ML tools allow customisation and remixing

OpenSoundControl input
from sensors/video/audio/…

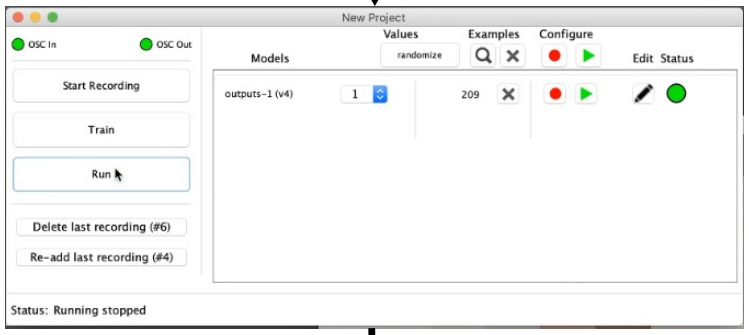

OpenSoundControl output
to sound/animation/game/…

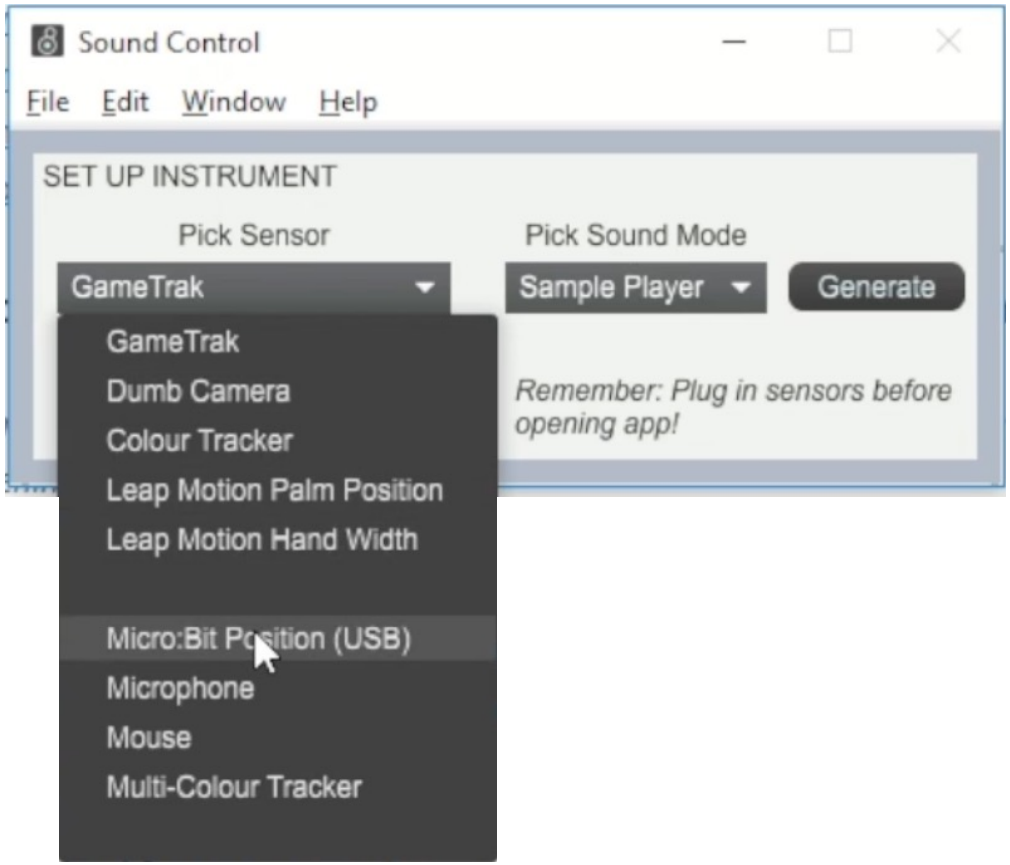

Principle 4:

Many ML learning objectives can be framed in terms of the creative design process

Design operation	Learning objectives
Establishing design requirements, Formulation	Understand the structure of ML problems and algorithms' capabilities Identify feasible uses for ML in new projects, and map them onto ML structures
Synthesis	Reason about properties of learning algorithms, data, and problem domains to make good choices Apply knowledge of ML workflows and skill with tools to build and use ML models Pass data/signals between ML tools and other project components
Analysis, Evaluation	Choose and apply appropriate ML evaluation methods
Reformulation	Reason about and exercise appropriate ways to improve an ML system
Other	Understand ways ML has been used in other creative work, and draw on this to contextualise one's own work

As described in: Fiebrink. 2019. "Machine learning education for artists, musicians, and other creative practitioners." *ACM Trans. Computing Education.*
Draws on design operations from: Howard, Culley, and Dekoninck. 2008. "Describing the creative design process by the integration of engineering design and cognitive psychology literature." *Design studies.*

Challenge 1:
Commercial tools are not always great for scaffolding learning or genuine creative work

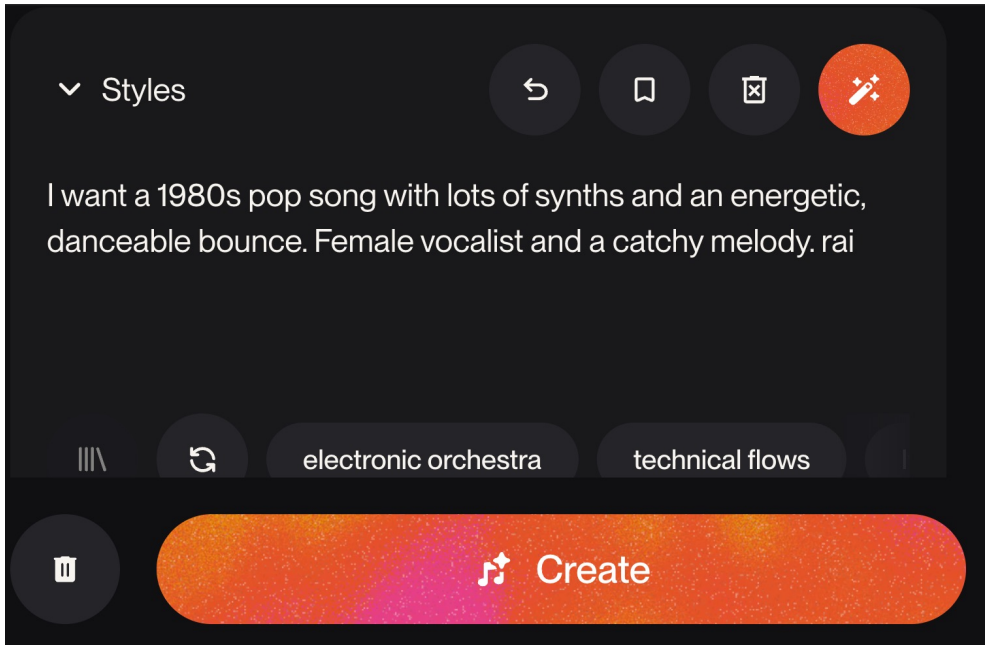


```python
best_loss = 100000

num_epochs = 100
for epoch in range(num_epochs):
    train_loss = 0.0
    test_loss = 0.0

    # Training loop
    for i, data in enumerate(train_loader, 0):
        inputs, labels = data
        inputs = inputs.to(device)
        labels = labels.to(device)
        optimizer.zero_grad()
        outputs = model(inputs)
        loss = criterion(outputs, labels.unsqueeze(1))
        loss.backward()
        optimizer.step()
        train_loss += loss.item() / len(train_loader)

    # After 10 epochs
    if (epoch+1) % 10 == 0:
        # Test loop
        with torch.no_grad():
            for i, data in enumerate(test_loader, 0):
                inputs, labels = data
                inputs = inputs.to(device)
                labels = labels.to(device)
                outputs = model(inputs)
                loss = criterion(outputs, labels)
                test_loss += loss.item() / len(test_loader)

        print(f'Epoch {epoch + 1}, train loss: {train_loss:.3f}, test loss: {test_loss:.3f}')
        if test_loss < best_loss:
            best_loss = test_loss
            torch.save(model.state_dict(), 'best_dog_rating_network.pt')
```

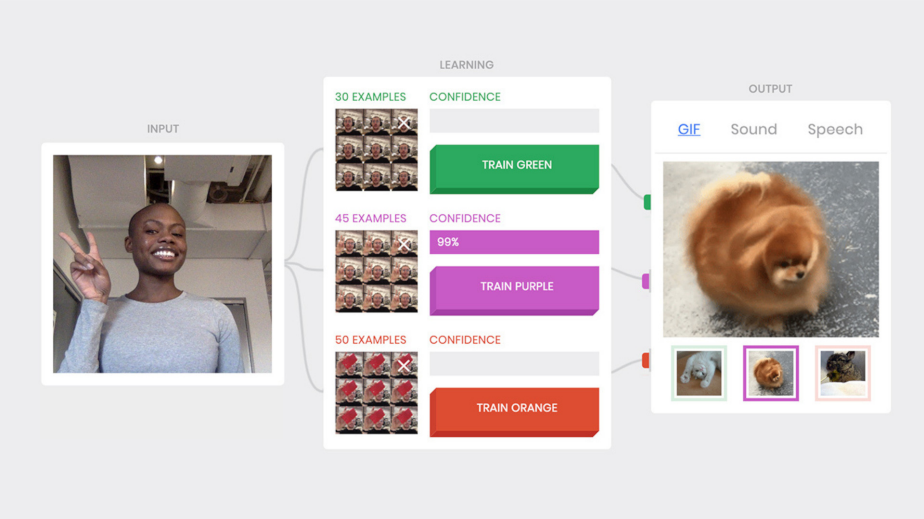

Teachable Machine

Micro:bit CreateAI

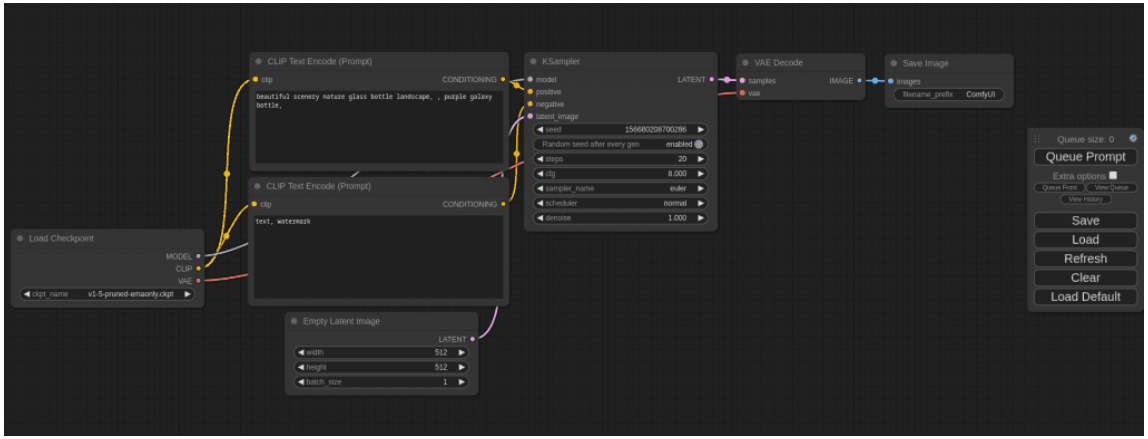

ComfyUI

Challenge 2:
Making generative AI tools safe for kids is incredibly difficult

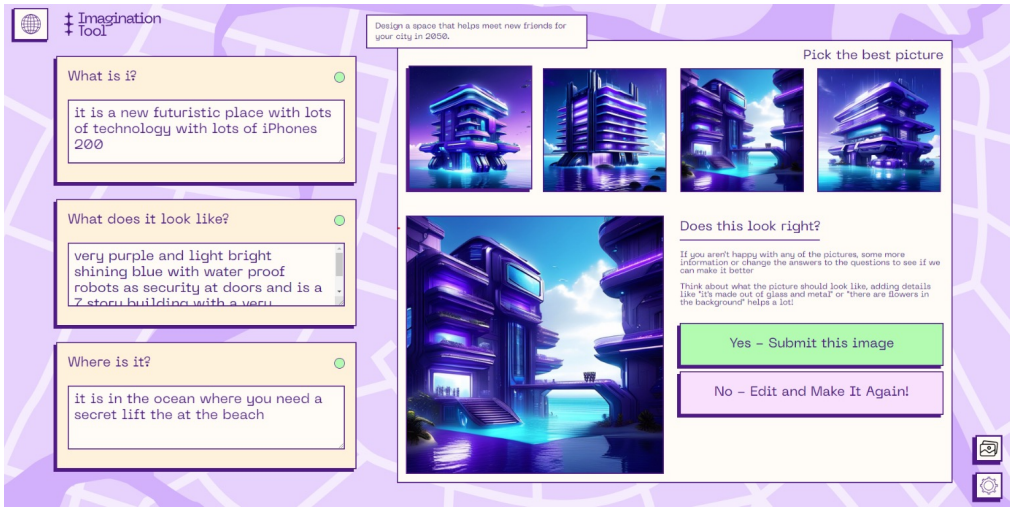

- Chose a model with built-in filter
- Bad/safe words lists
- Toxicity detection
- Human (teacher) moderation
- Prompt engineering
- Extensive testing in advance

Meriç, White, Suárez Zapico, Yanovich, Koby-Hirschmann & Fiebrink. ICCC 2024.

Final Remarks

Learning and teaching creative ML is worth it!

- ML can help creators work more effectively with data they care about
- ML allows people to communicate ideas and intentions to computers through examples, which can be especially helpful for embodied practices
- Designing with data can enable more people to become creators—and not just by helping them to conform to existing standards of media creation
- ML can enable exciting new types of creative work

Learning and teaching creative ML is possible!

- A growing set of tools allow meaningful experimentation and original creative work, without (much/any) math or coding
  - Wekinator, MIMIC (browser/JS), InteractML (Unreal)
  - Teachable Machine, Flucoma (music), Microbit CreateAI, TouchDesigner, ComfyUI
- Creative applications can be a fun jumping-off point for broader ML engagement, and a meaningful destination in themselves

Thanks!

www.wekinator.org

mimicproject.com

interactml.com

r.fiebrink@arts.ac.uk