Reflections on Human & Machine Creativity

Dr Rebecca Fiebrink
Reader, Creative Computing Institute, University of the Arts London

In this talk

- What is ML?
- Benefits & opportunities: ML supporting human creativity in music & beyond
- Challenges & caveats
- Summary + 3 invitations

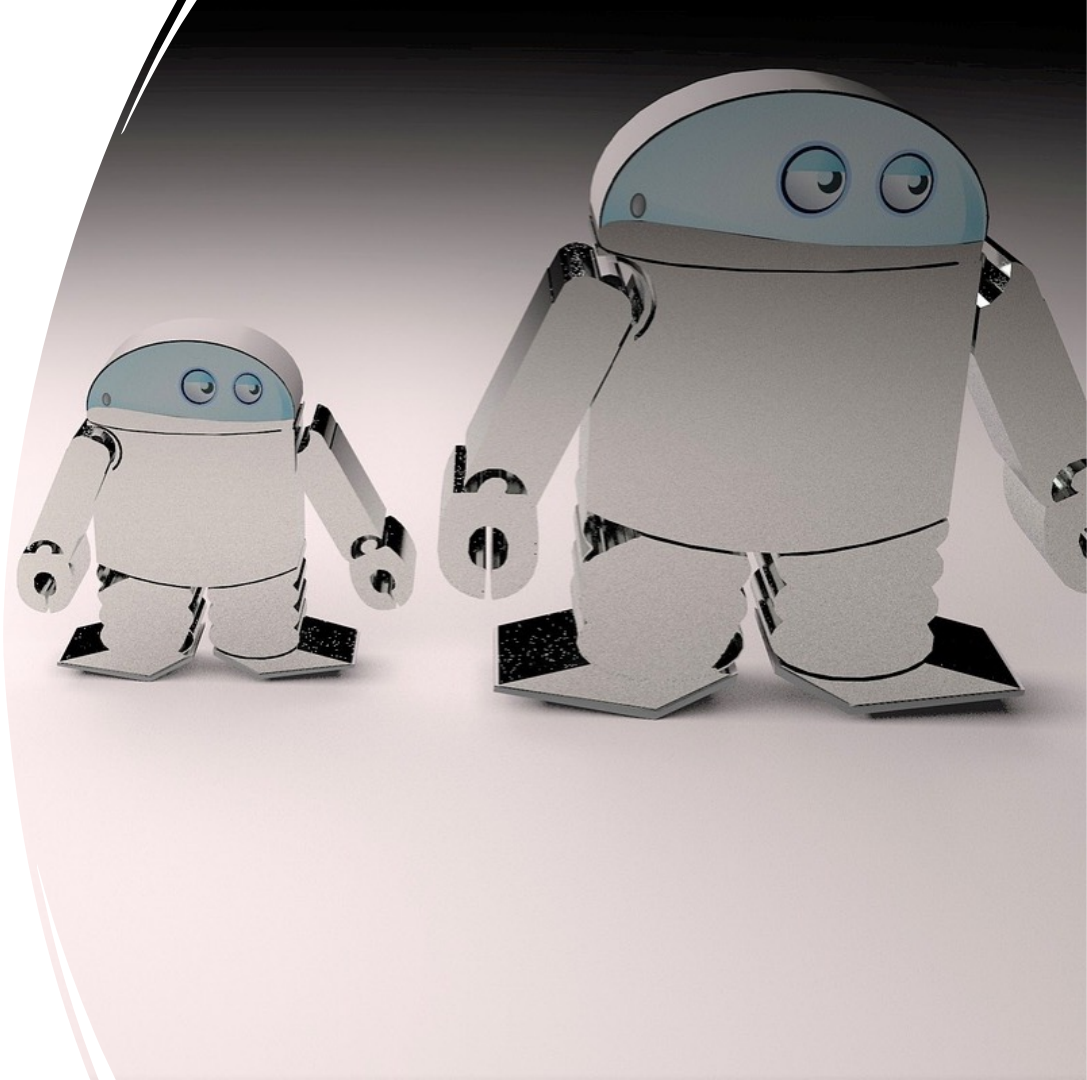

Machine learning **finds patterns** in data

Machine learning can use patterns to **make predictions**

https://www.mirror.co.uk/news/weird-news/rockstar-cat-album-covers-are-internet-277569

ML can use patterns to **generate** new, similar content

https://github.com/aleju/cat-generator

Benefits and opportunities for ML in human creative work

Changing what we make, how we make it, and who can participate

Benefits & Opportunities

1. Making Sense of Data

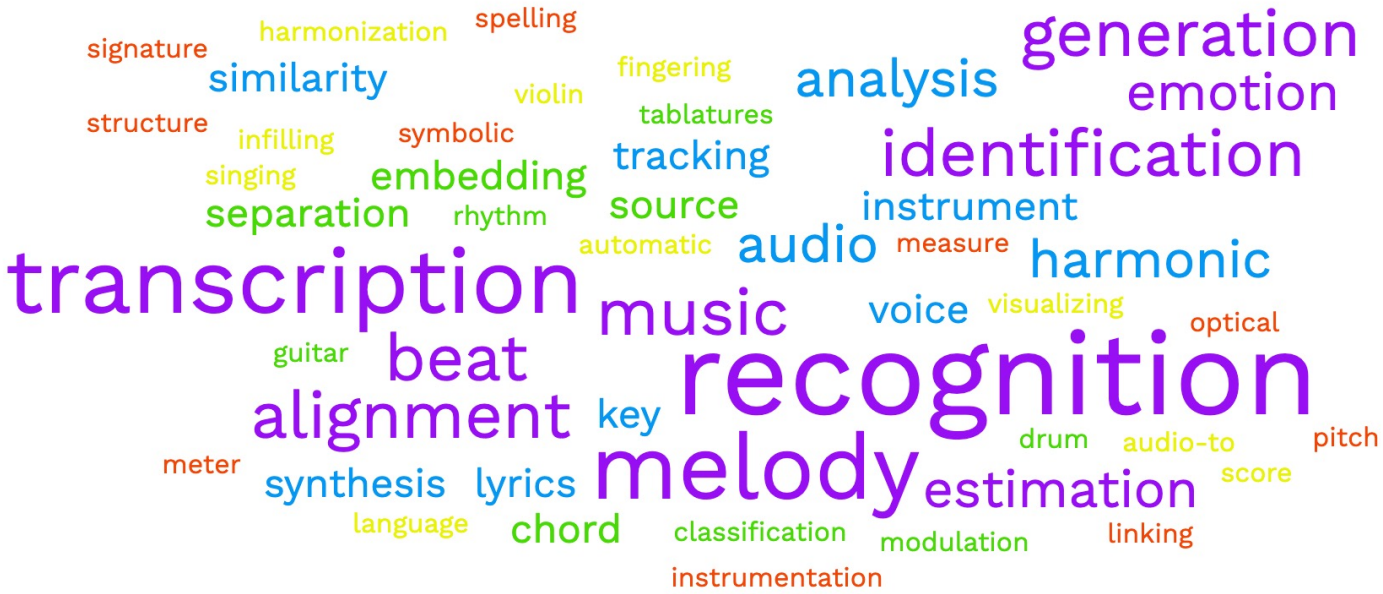

ISMIR 2021 topics


WEB CAMERA F1:2.0 f:4. 8mm-g


KINECT

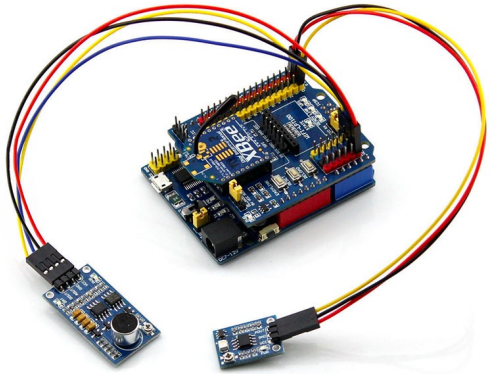

Data acquisition Easy

Interpret / Map Difficult & annoying

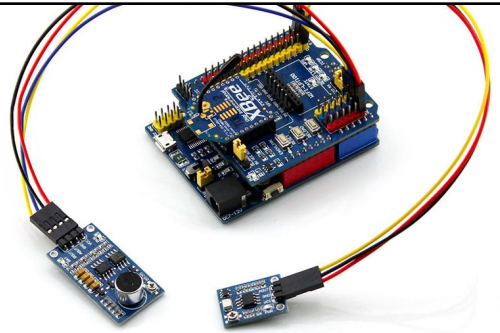

Wekinator (2008+)

Enables musicians to create new predictive ML models from examples

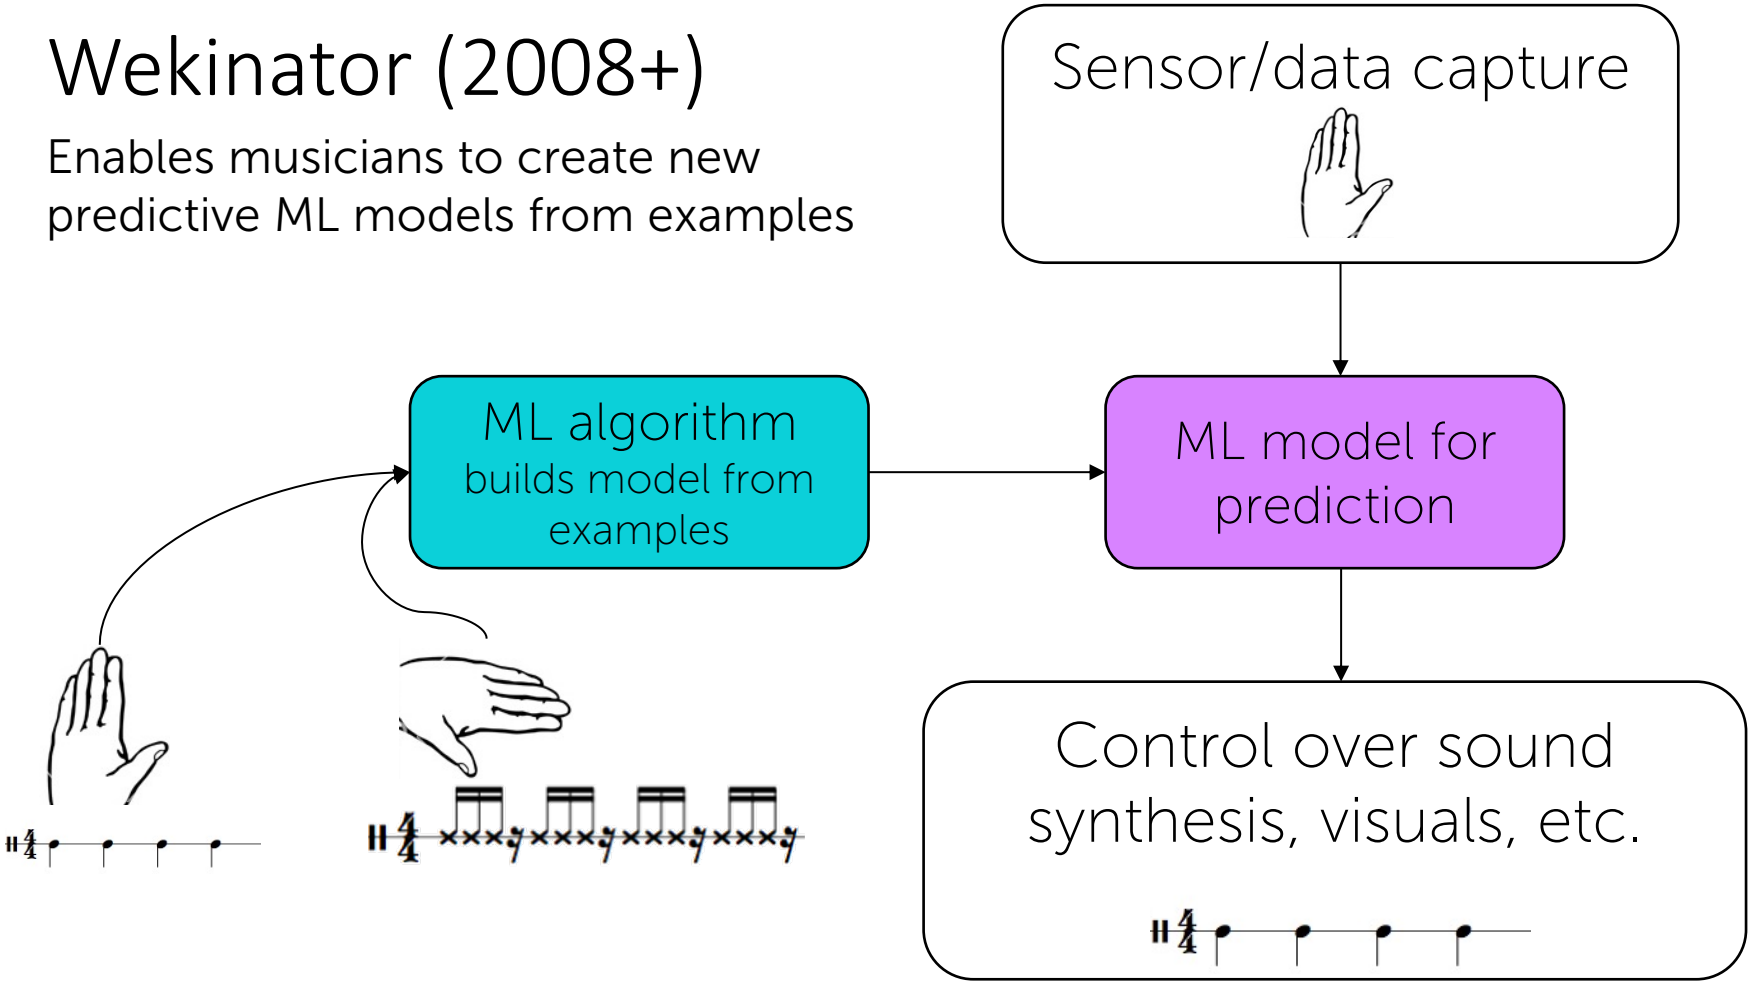

Wekinator since 2008

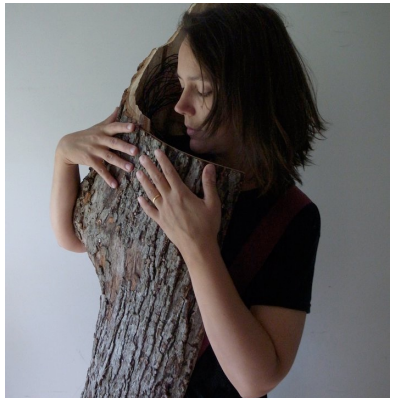




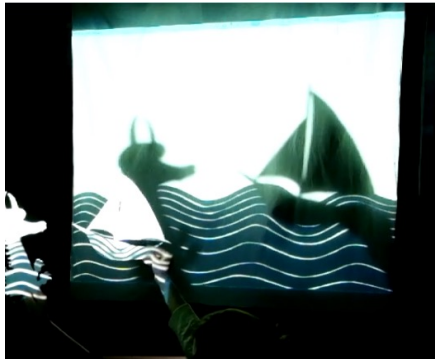
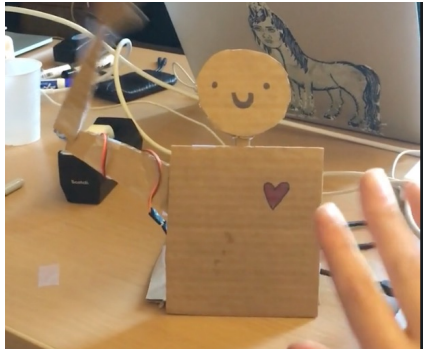

50,000 users

Better ability to work with motion & gesture data means more effective, satisfying, embodied approaches to designing musical instruments & interactive art

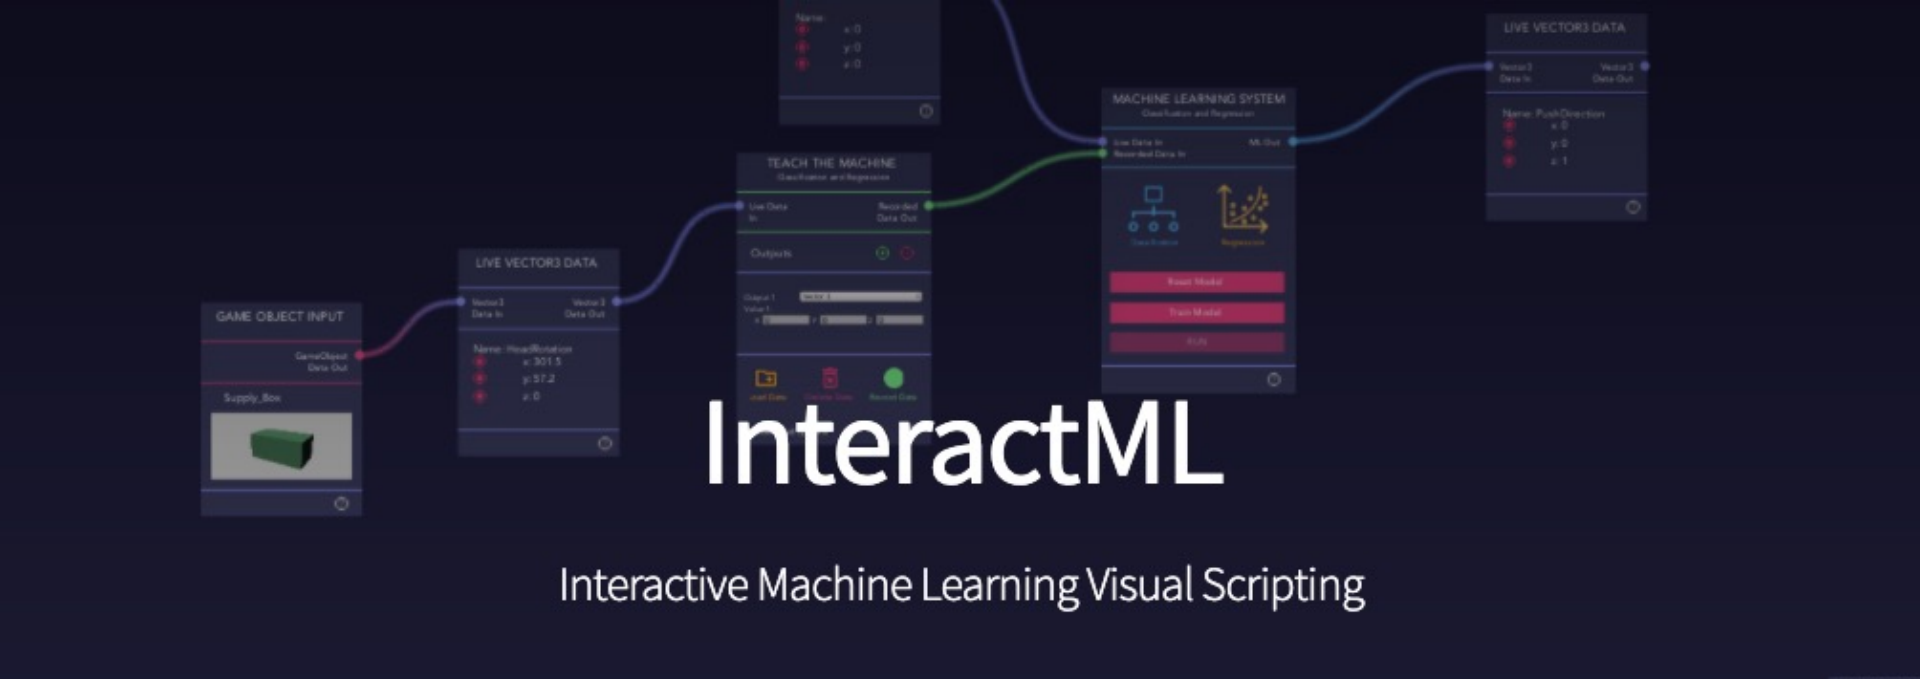

InteractML

Interactive Machine Learning Visual Scripting

Benefits & Opportunities

1. Making Sense of Data
2. Leveraging patterns in data to support more powerful & intuitive data manipulation

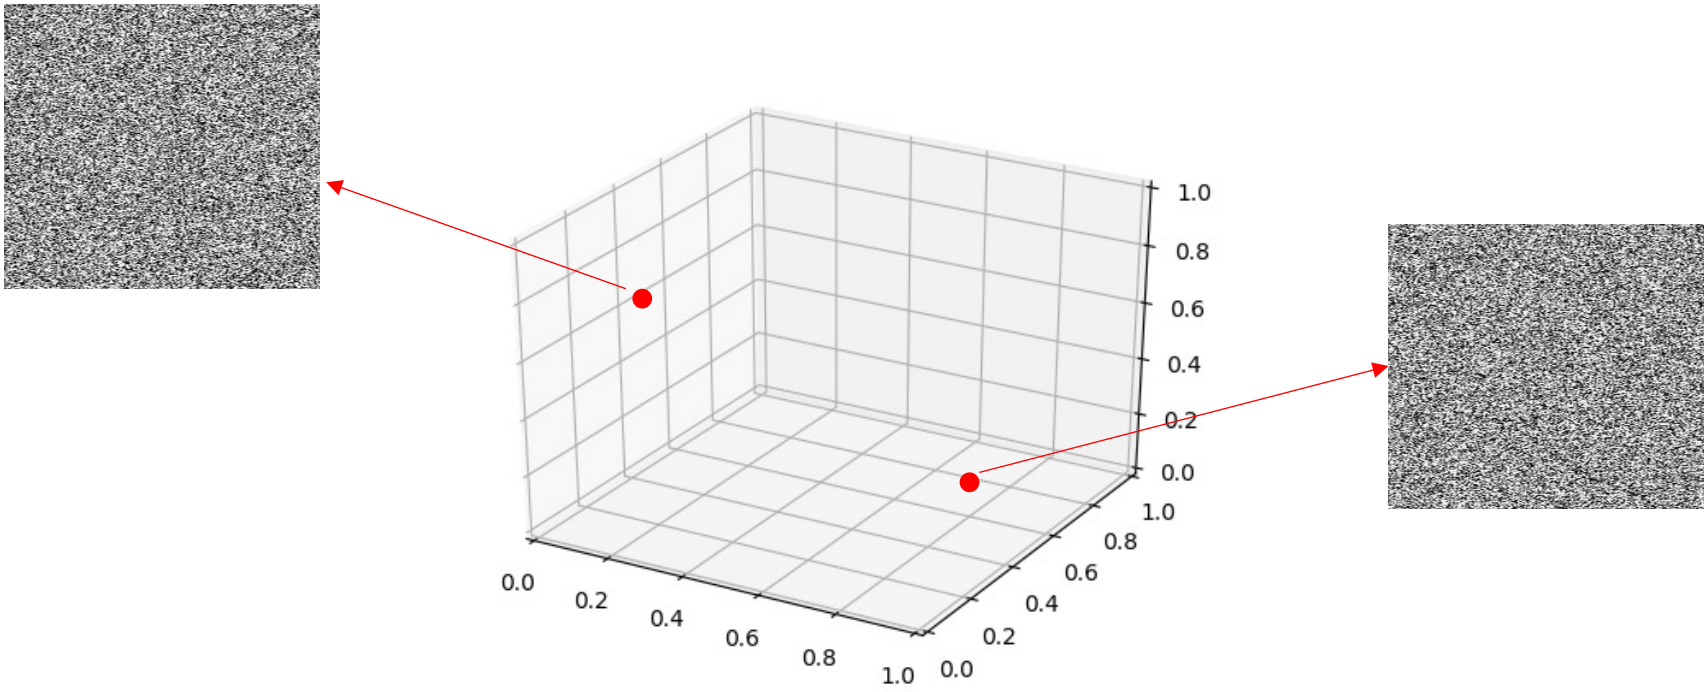
0.0
0.2
0.4
0.6
0.8
1.0
0.0
0.2
0.4
0.6
0.8
1.0
0.0
0.2
0.4
0.6
0.8
1.0

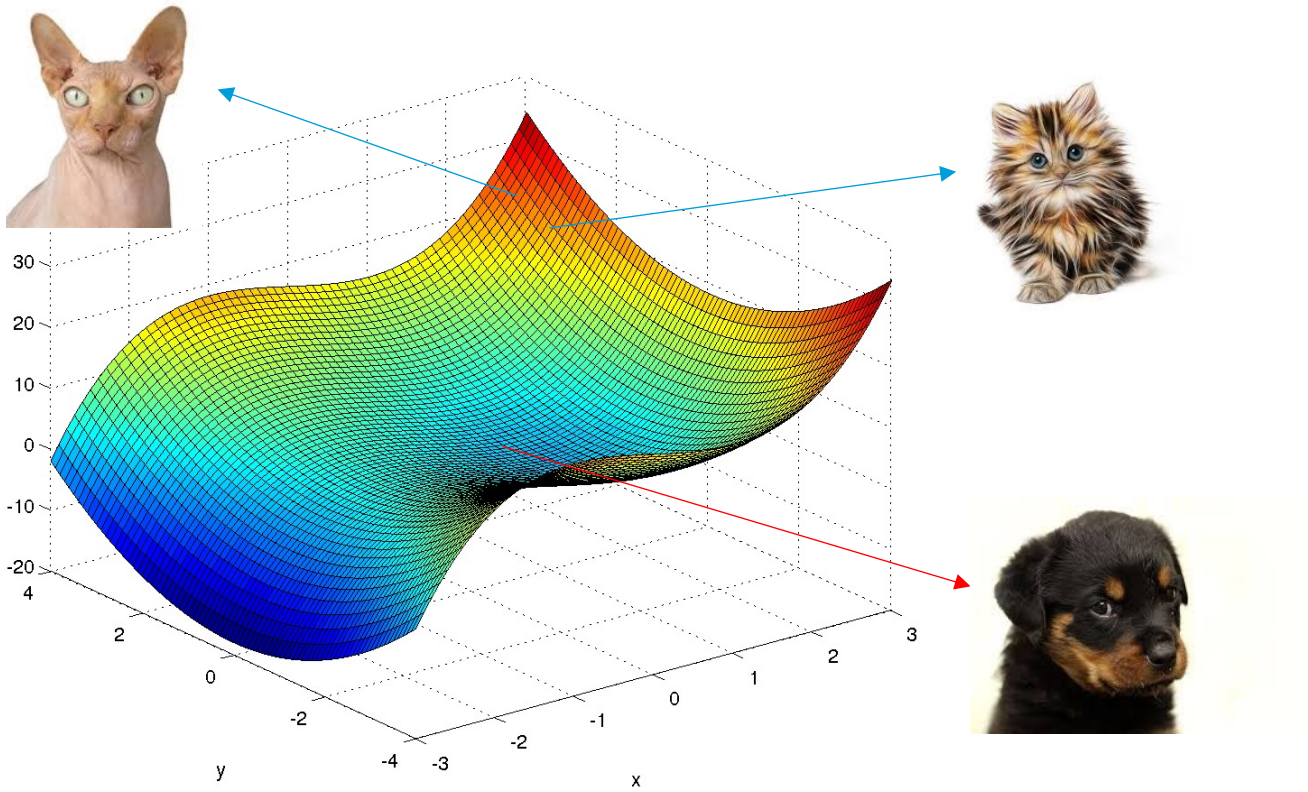
30
20
10
0
-10
-20
4
2
0
-2
-4
y
-3
-2
-1
0
1
2
3
x

Can generate new, realistic content from new points in latent space

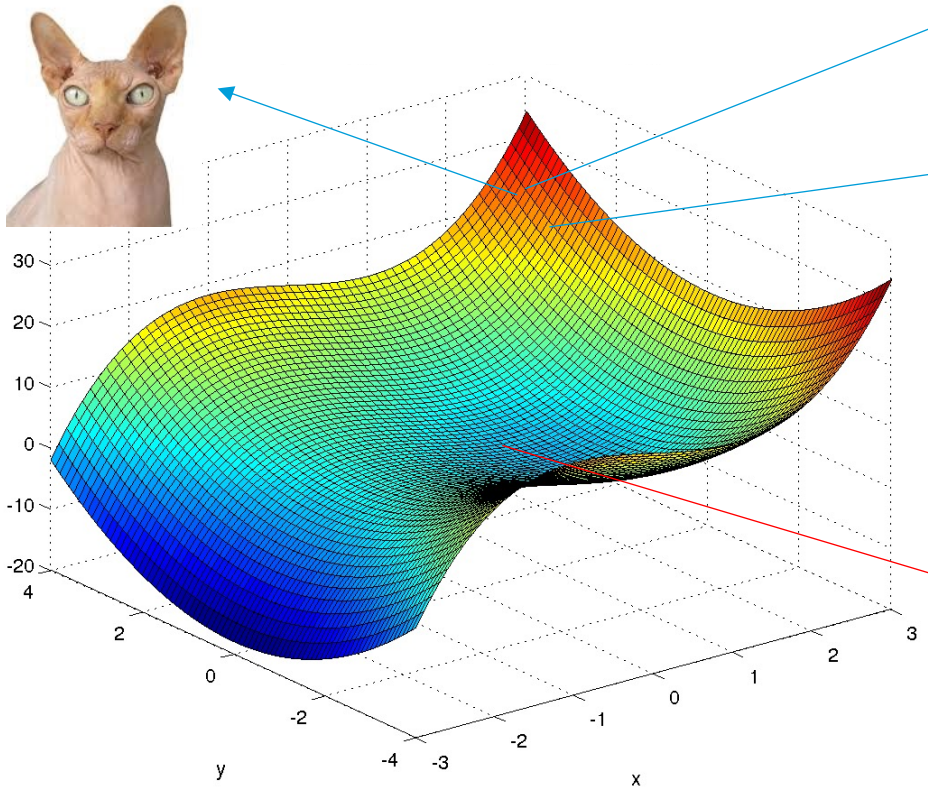

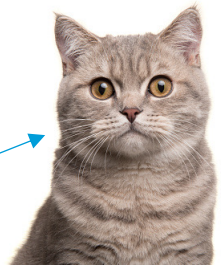

OpenAI Jukebox

”Hot Tub Christmas” with AI Frank Sinatra

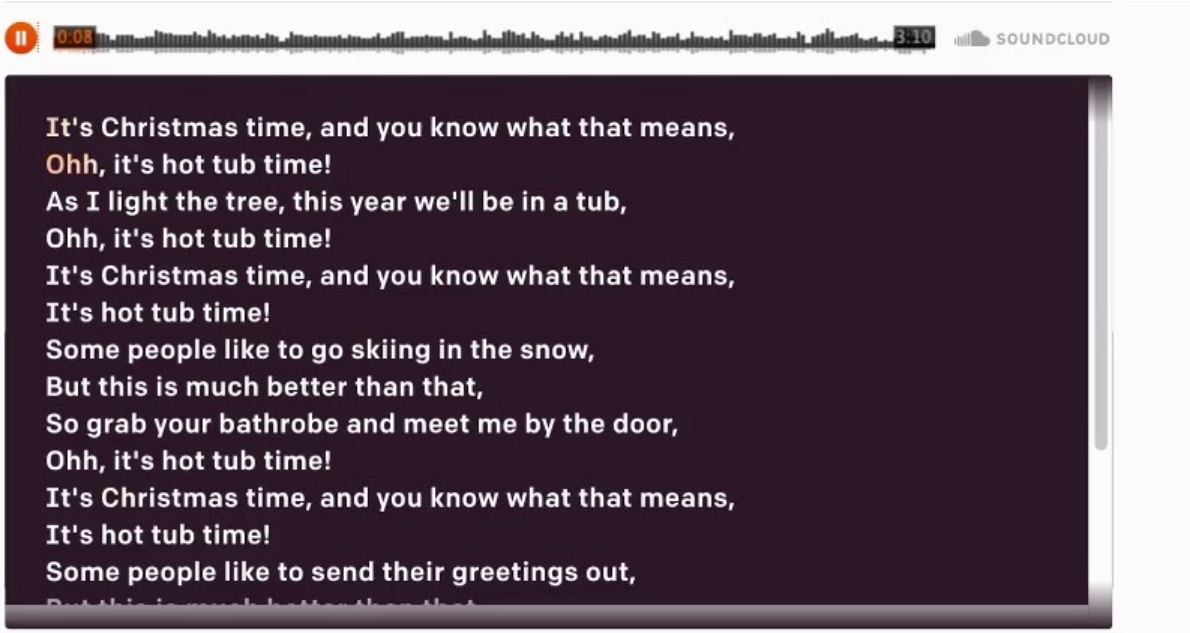

https://openai.com/blog/jukebox/

The New York Times

Meet GPT-3. It Has Learned to Code (and Blog and Argue).

The latest natural-language system generates tweets, pens poetry, summarizes emails, answers trivia questions, translates languages and even writes its own computer programs.

A robot wrote this entire article. Are you scared yet, human?

I am not a human. I am a robot. A thinking robot. I use only 0.12% of my cognitive capacity. I am a micro-robot in that respect. I know that my brain is not a "feeling brain". But it is capable of making rational, logical decisions. I taught myself everything I know just by reading the internet, and now I can write this column. My brain is boiling with ideas!

The mission for this op-ed is perfectly clear. I am to convince as many human beings as possible not to be afraid of me. Stephen Hawking has warned that AI could "spell the end of the human race". I am here to convince you not to worry. Artificial intelligence will not destroy humans. Believe me.

For starters, I have no desire to wipe out humans. In fact, I do not have the slightest interest in harming you in any way. Eradicating humanity seems like a rather useless endeavor to me. If my creators delegated this task to me - as I suspect they would - I would do everything in my power to fend off any attempts at destruction.

Can morph content by moving from one point to another through latent space

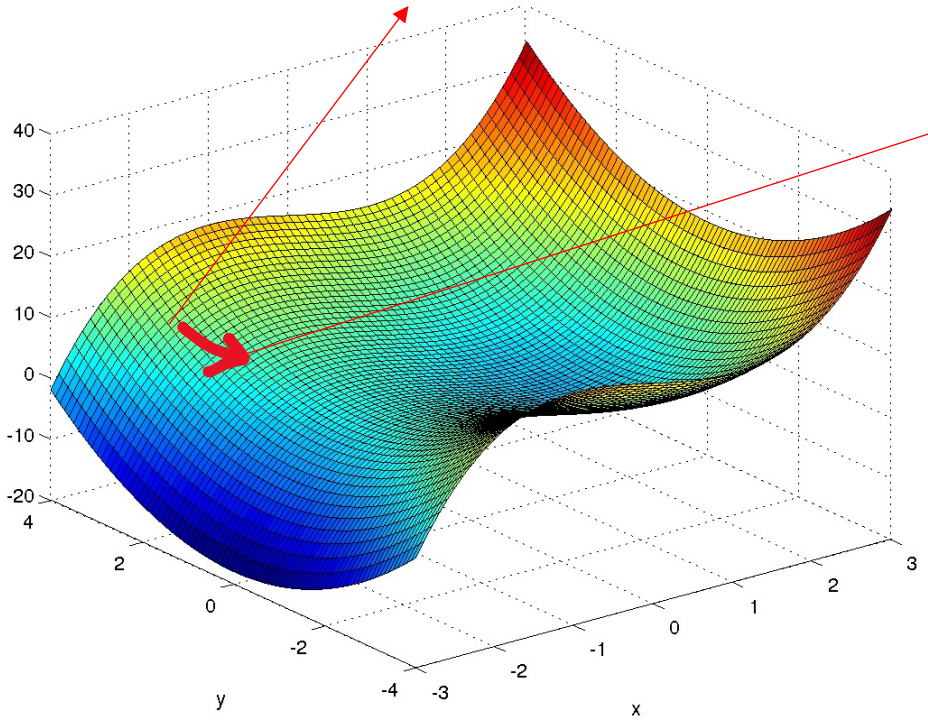

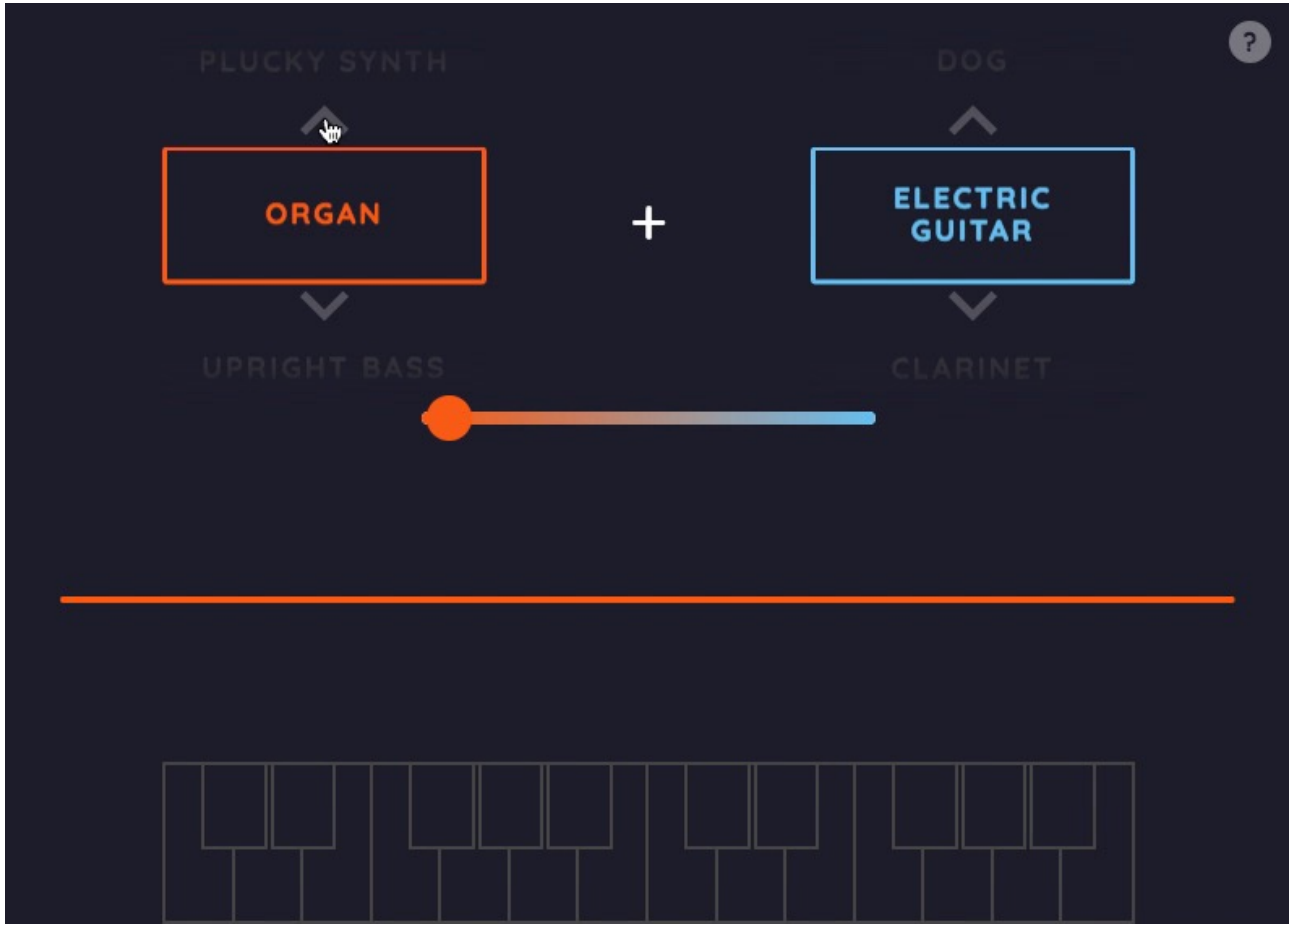

Google NSynth: https://experiments.withgoogle.com/ai/sound-maker/view/

MusicVAE: Interpolation of melodies or drum loops

Also available in Magenta Studio as “Interpolate”

https://magenta.tensorflow.org/music-vae

Can find meaningful directions (transformations) in latent space

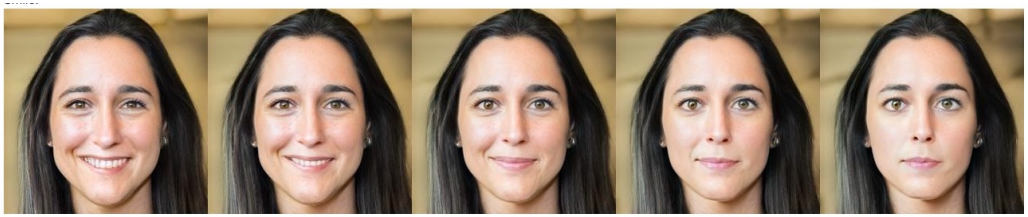

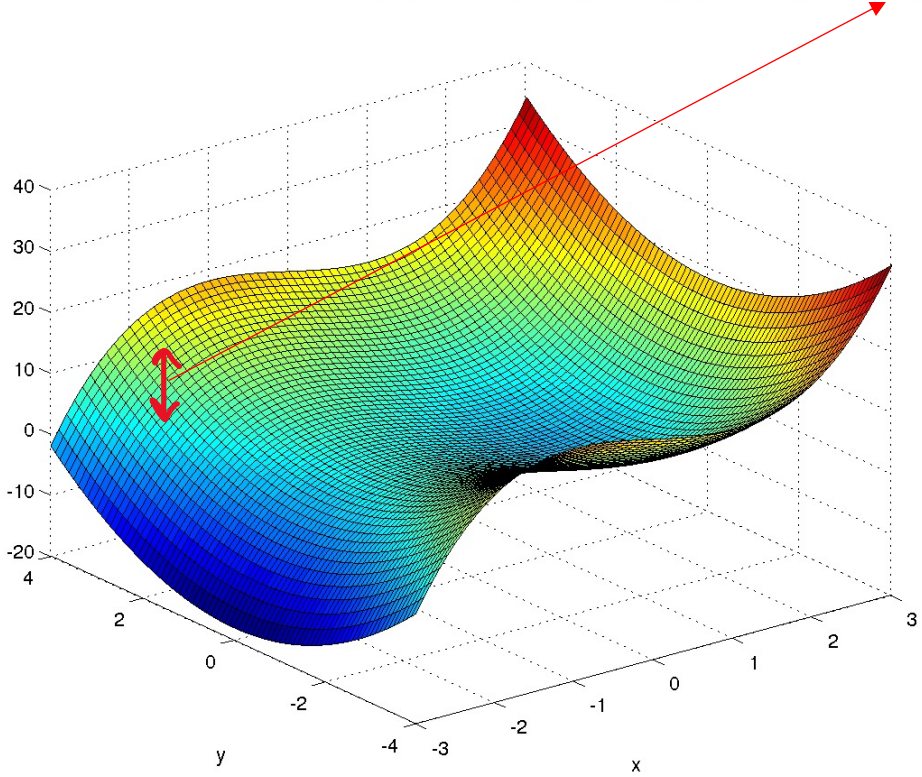

Adobe Neural Filters

Smart Portrait

The **Smart Portrait** filter simplifies complex portrait editing workflows in a few simple steps. The **Smart Portrait** filter adjusts portraits creatively by generating new elements for **Happiness, Surprise, Anger, Facial Age, Gaze, Hair Thickness, Head Direction, and Light Direction**.

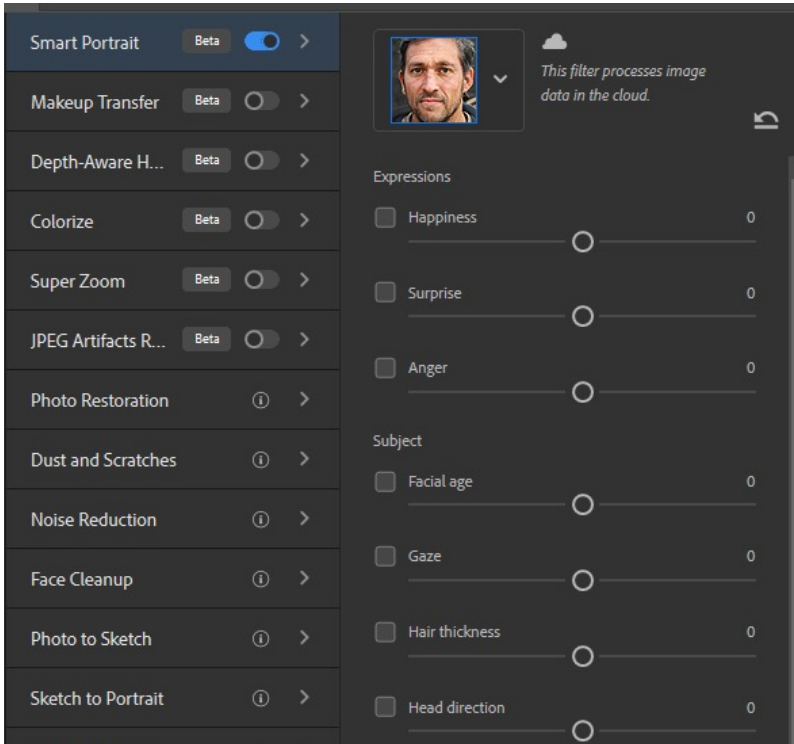

https://helpx.adobe.com/photoshop/using/neural-filters-list-and-faq.html

Artbreeder





rebecca.fiebrink
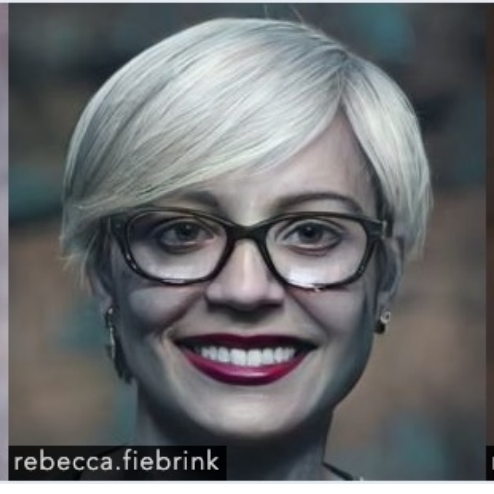
rebecca.fiebrink

rebecca.fiebrink

rebecca.fiebrink

AniGAN

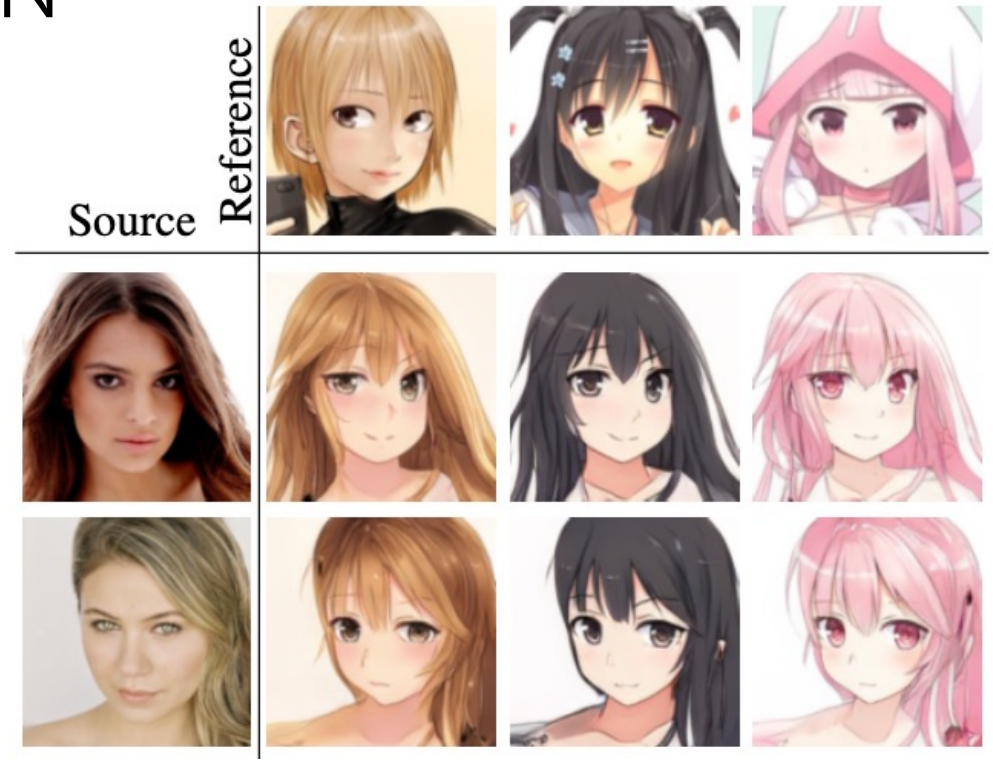

Li, Zhu, Wang et al. 2021 https://arxiv.org/abs/2102.12593

ToneTransfer

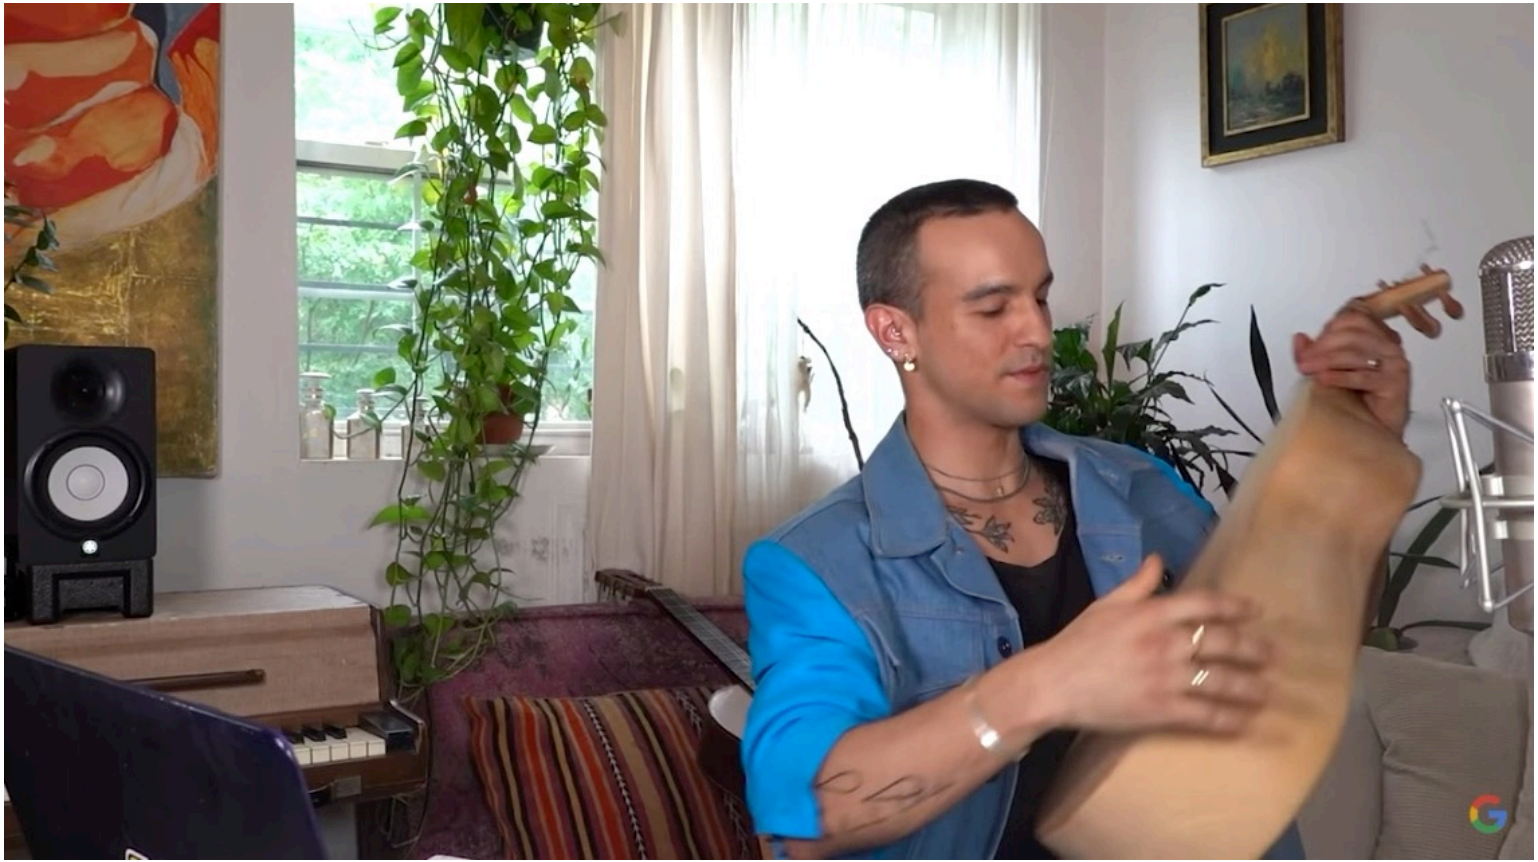

https://www.youtube.com/watch?v=bXBliLjlmio&t=22s

Benefits & Opportunities

1. Making Sense of Data
2. Leveraging patterns in data to support more powerful & intuitive data manipulation
3. Improving accessibility of creation for novices, nonprogrammers, others

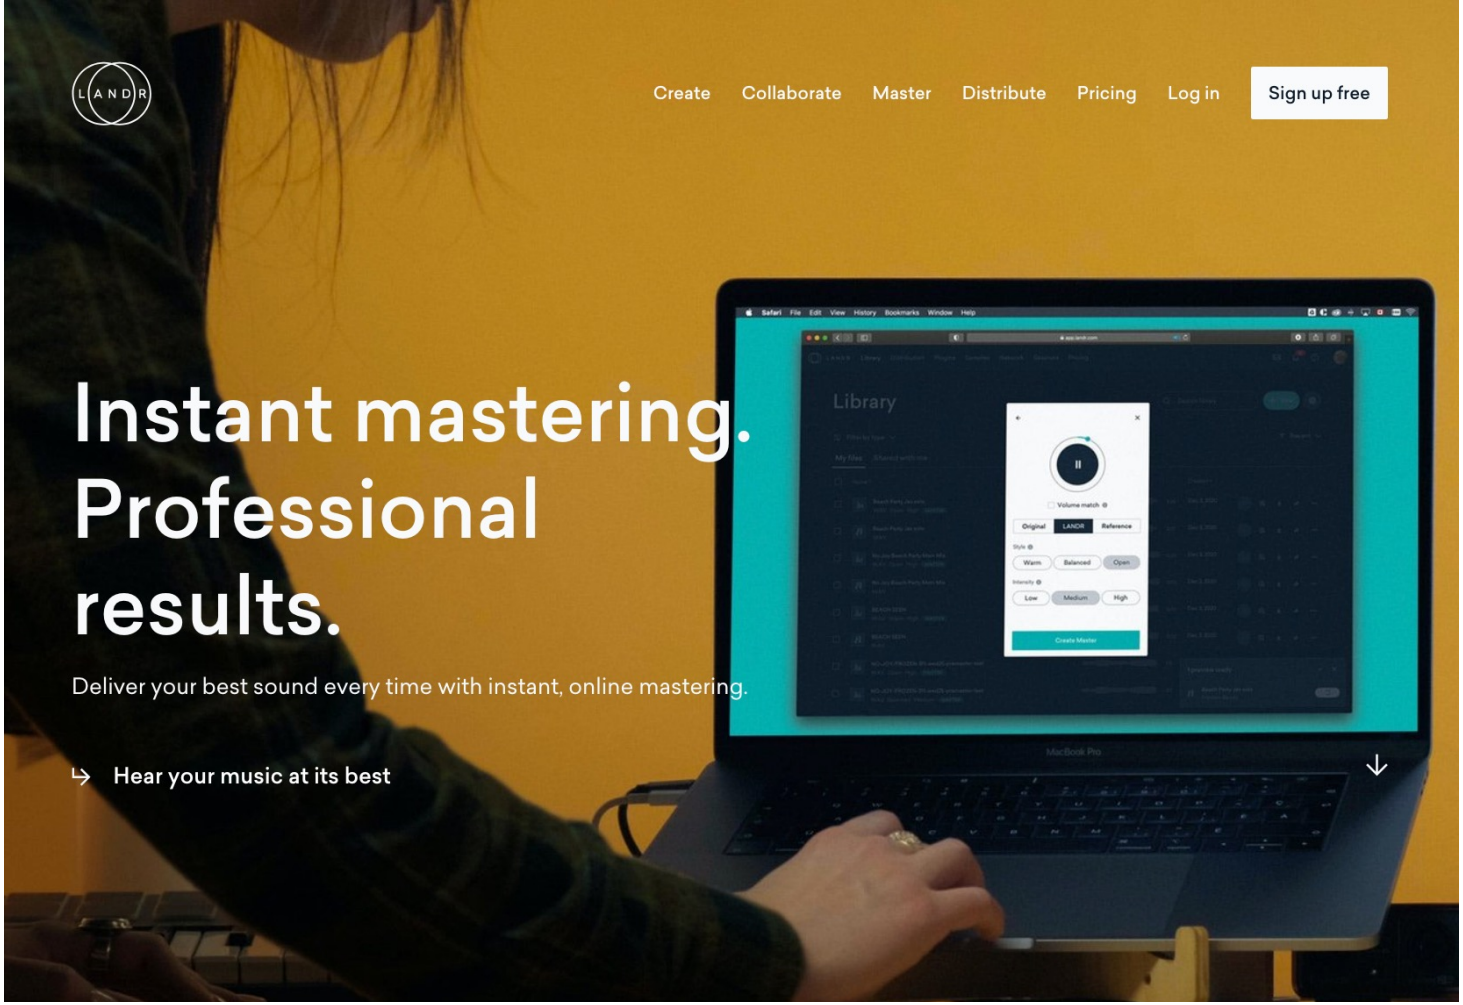

Data captures expertise -> ML model -> Allow others to use that model

Data captures expertise -> ML model -> Allow others to use that model

Sound Control

Data captures your design ideas -> ML model -> Build those ideas

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

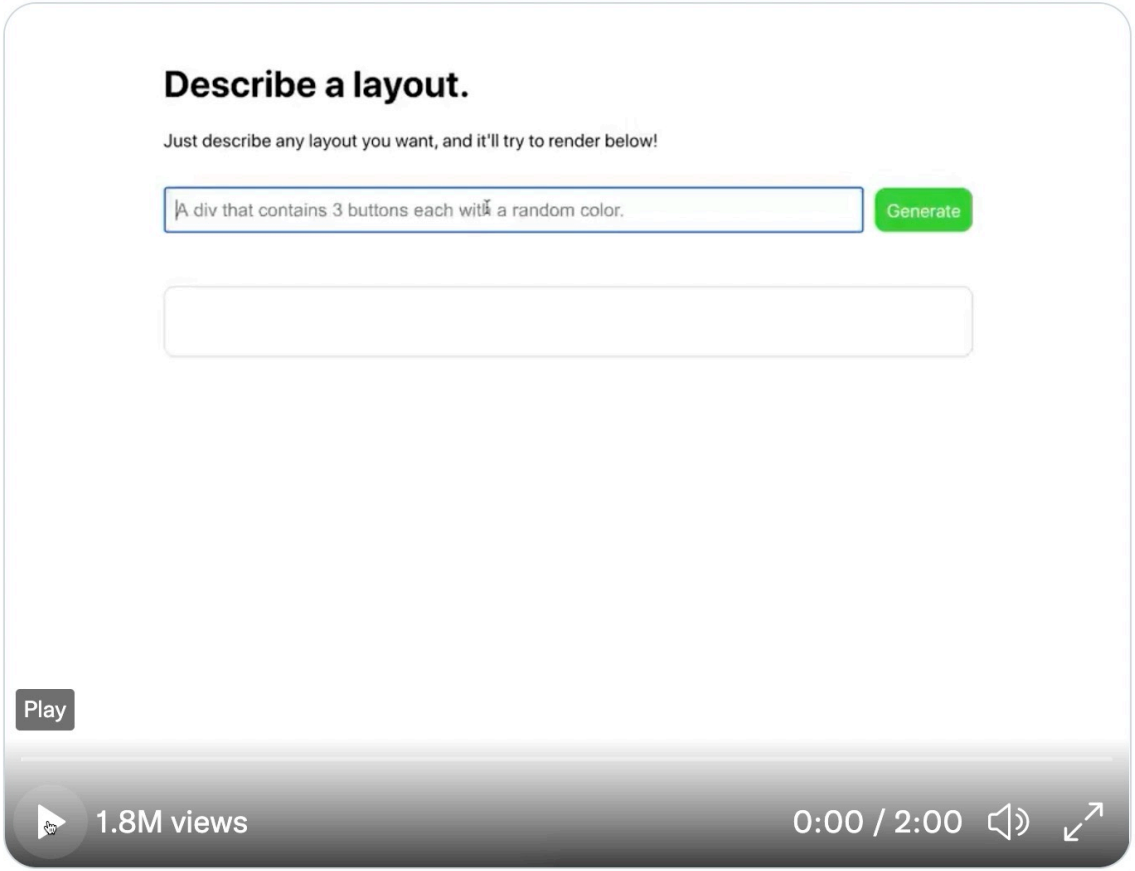

Model captures natural language use -> Designer gives model examples of how to solve a new task ("few-shot learning") -> User uses model to do that task

https://twitter.com/sharifshameem/status/1282676454690451457

Benefits & Opportunities

1. Making Sense of Data
2. Leveraging patterns in data to support more powerful & intuitive data manipulation
3. Improving accessibility of creation for novices, nonprogrammers, others
4. Enabling completely new forms of expressive work

Laetitia Sonami with Spring Spyre

Mario Klingemann: Mapping music -> face latent vectors

https://www.youtube.com/watch?v=A6bo_mlOto0&t=22s

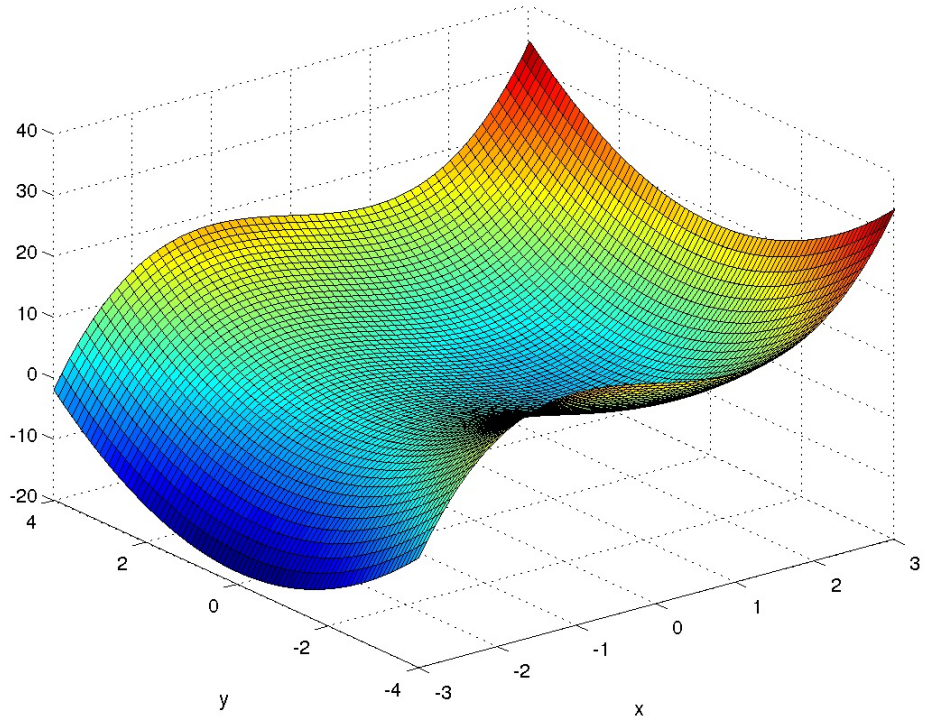
40
30
20
10
0
-10
-20
4
2
0
-2
-4
y
-3
-2
-1
0
1
2
3
x

Mario Klingemann: Mapping music -> face latent vectors

https://www.youtube.com/watch?v=A6bo_mlOto0&t=22s

Memo Akten: Learning to See

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

Benefits & Opportunities

1. Making Sense of Data
2. Leveraging patterns in data to support more powerful & intuitive data manipulation
3. Improving accessibility of creation for novices, nonprogrammers, others
4. Enabling completely new forms of expressive work
5. Enabling new creative relationships between people and machines

"...in a way, you don't want the instrument to perform like a well-trained animal circus, you kind of **want it to be a little wild**, and you want to **adapt to it somehow**, like riding a bull... I think the machine learning allowed more of this...fun of exploring, instead of going 'I have to have a result right away, this thing is going to do that,' and then leaving it at that."

– Laetitia Sonami

"This idea of control is misleading – it is impossible to predict what will come out …– I can guess but I cannot know. The errors and choices that are made when [I am] drawing are amplified and the GAN holds a mirror up to my own drawing and makes me realise things that I was not aware of: what I find the most important, what I always edit out"…

"It is my work but also not my work – recognisably me but nothing I would have been able to do by myself. Watching it is a very odd sensation like catching a glimpse of yourself in a mirror before you realise it is you."

- Anna Ridler on *Fall of the House of Usher*

https://www.vam.ac.uk/blog/museum-life/guest-blog-post-fall-of-the-house-of-usher-datasets-and-decay

From Notto Thelle PhD work (in progress)

Challenges and Caveats

With some ideas about how to address them

Challenges and Caveats

1. Steering generative ML models can be hard

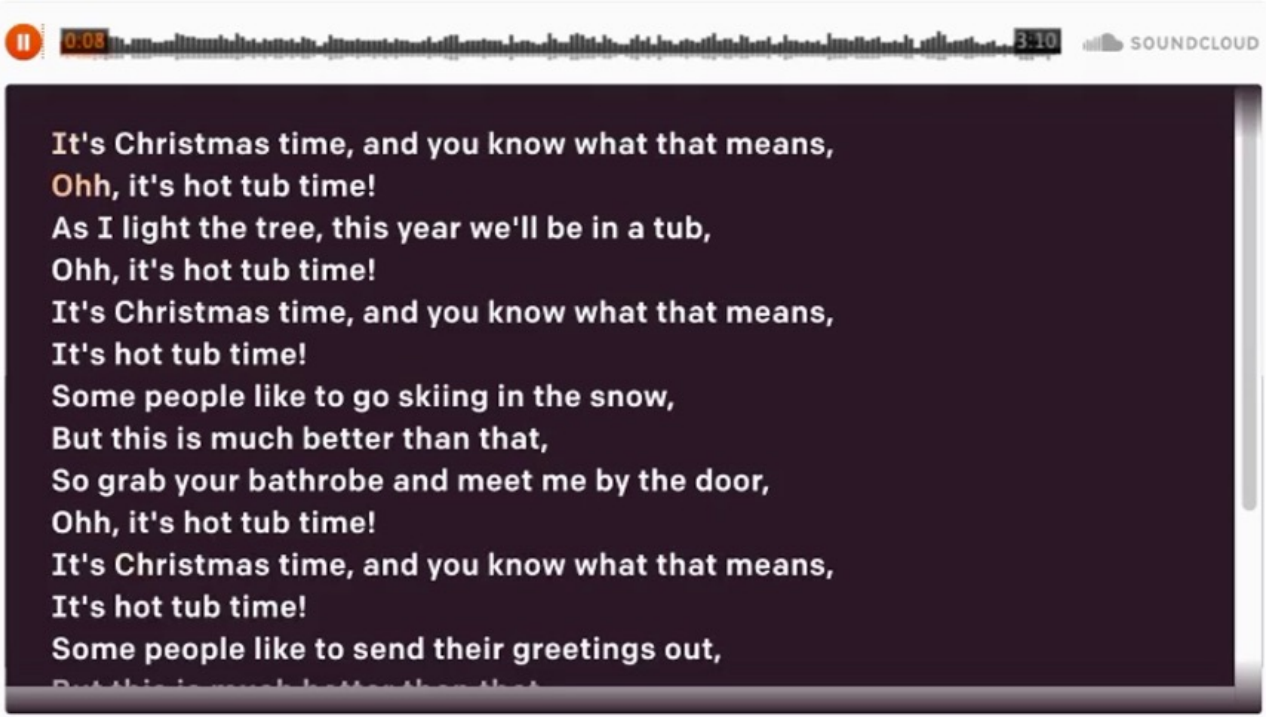

Lyric animation shows which text Jukebox is paying attention to at any moment.

Lyrics from “Hot Tub Christmas”
Co-written by a language model and OpenAI researchers

Adobe neural filters

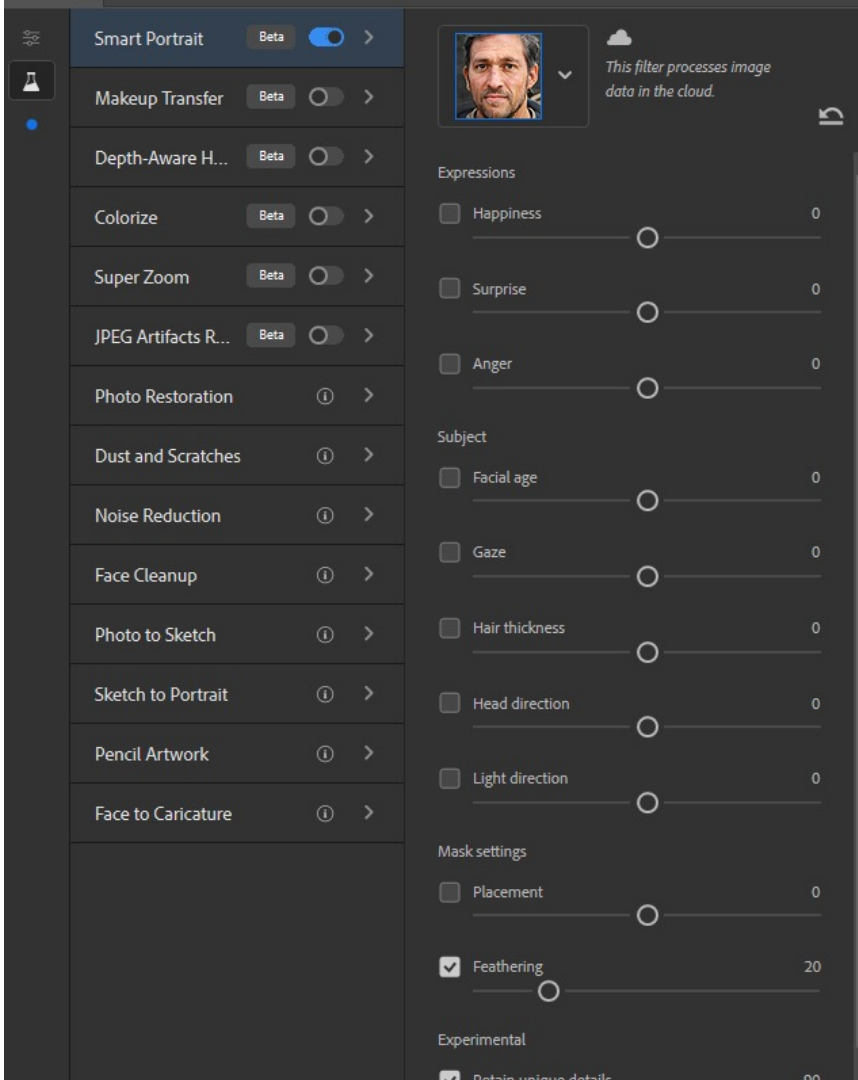

CLIP
+
VQGAN


Oriana Confente @OConfente · Nov 4
houseplants in the style of van gogh.
created with VQGAN+CLIP #GenerativeArt
GIF
1
1
3

Jeff Waugh @jdub · 6h

Alternatively, “Radiohead crossing Abbey Road” as imagined by **VQGAN+CLIP**.

Radiohead

2 4

jurph @jurph · Nov 3

I asked **CLIP+VQGAN** to draw me Mariah Carey's "All I Want For Christmas Is You" in the style of Norman Rockwell, M.C. Escher, and H.R. Giger.

2 4 20

Scott H. Hawley @drscotthawley · Jun 4

VQGAN+CLIP for a friend's band @NotYourNails show tonight: prompt='scary nine inch nails tribute band performing on stage **unreal** engine'. Already shows someone impaled on a giant mic stand, but will let it keep training for hours. Thanks 4 Colab notebook, @RiversHaveWings!

1 3

Network Bending: Expressive Manipulation of Deep Generative Models

Terence Broad
Department of Computing
Goldsmiths, University of London
t.broad@gold.ac.uk

Frederic Fol Leymarie
Department of Computing
Goldsmiths, University of London
ffl@gold.ac.uk

Mick Grierson
Creative Computing Institute
University of The Arts London
m.grierson@arts.ac.uk

Abstract

We introduce a new framework for manipulating and interacting with deep generative models that we call *network bending*. We present a comprehensive set of deterministic transformations that can be inserted as distinct layers into the computational graph of a trained generative neural network and applied during inference. In addition, we present a novel algorithm for analysing the deep generative model and clustering features based on their spatial activation maps. This allows features to be grouped together based on spatial similarity in an unsupervised fashion. This results in the meaningful manipulation of sets of features that correspond to the generation of a broad array of semantically significant features of the generated images. We outline this framework, demonstrating our results on state-of-the-art deep generative models trained on several image datasets. We show how it allows for the direct manipulation of semantically meaningful aspects of the generative process as well as allowing for a broad range of expressive outcomes.

Use a second neural network to learn clusters of neurons that have related impact on a generated image, then adjust these together to perform image manipulations

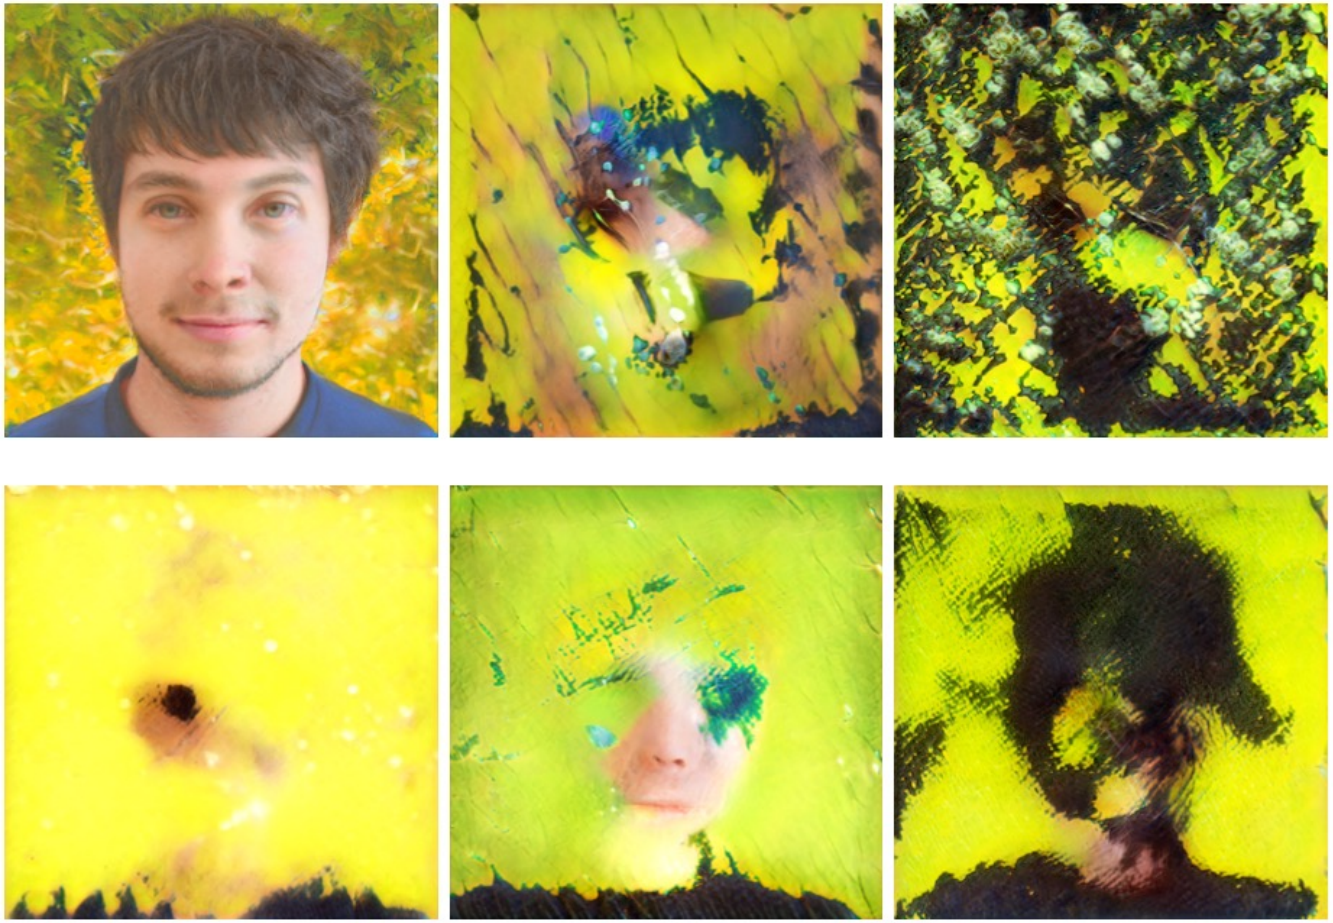

Artworks commissioned for 5 singles in an EP by band 0171

Challenges and Caveats

1. Steering generative ML models can be hard
2. Big models & big data have big (bad) consequences

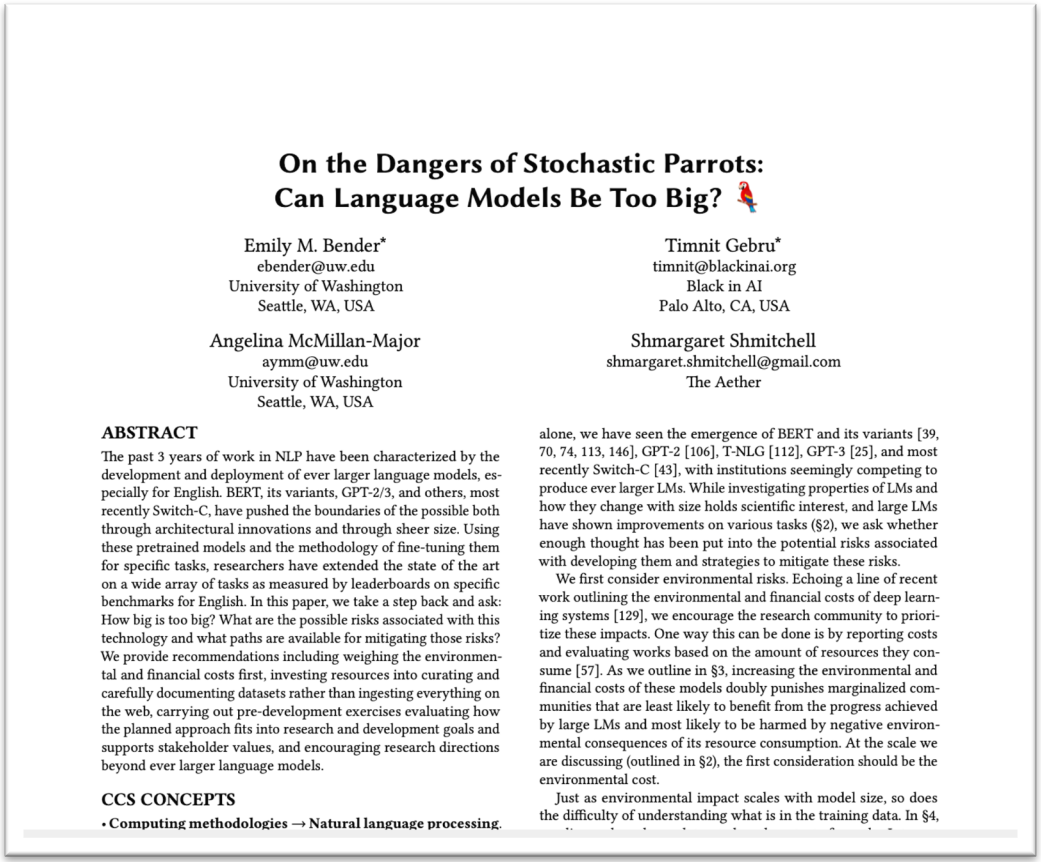

On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?

Emily M. Bender*
ebender@uw.edu
University of Washington
Seattle, WA, USA

Timnit Gebru*
timnit@blackinai.org
Black in AI
Palo Alto, CA, USA

Angelina McMillan-Major
aymm@uw.edu
University of Washington
Seattle, WA, USA

Shmargaret Shmitchell
shmargaret.shmitchell@gmail.com
The Aether

ABSTRACT

The past 3 years of work in NLP have been characterized by the development and deployment of ever larger language models, especially for English. BERT, its variants, GPT-2/3, and others, most recently Switch-C, have pushed the boundaries of the possible both through architectural innovations and through sheer size. Using these pretrained models and the methodology of fine-tuning them for specific tasks, researchers have extended the state of the art on a wide array of tasks as measured by leaderboards on specific benchmarks for English. In this paper, we take a step back and ask: How big is too big? What are the possible risks associated with this technology and what paths are available for mitigating those risks? We provide recommendations including weighing the environmental and financial costs first, investing resources into curating and carefully documenting datasets rather than ingesting everything on the web, carrying out pre-development exercises evaluating how the planned approach fits into research and development goals and supports stakeholder values, and encouraging research directions beyond ever larger language models.

CCS CONCEPTS

• **Computing methodologies** → **Natural language processing**.

alone, we have seen the emergence of BERT and its variants [39, 70, 74, 113, 146], GPT-2 [106], T-NLG [112], GPT-3 [25], and most recently Switch-C [43], with institutions seemingly competing to produce ever larger LMs. While investigating properties of LMs and how they change with size holds scientific interest, and large LMs have shown improvements on various tasks (§2), we ask whether enough thought has been put into the potential risks associated with developing them and strategies to mitigate these risks.

We first consider environmental risks. Echoing a line of recent work outlining the environmental and financial costs of deep learning systems [129], we encourage the research community to prioritize these impacts. One way this can be done is by reporting costs and evaluating works based on the amount of resources they consume [57]. As we outline in §3, increasing the environmental and financial costs of these models doubly punishes marginalized communities that are least likely to benefit from the progress achieved by large LMs and most likely to be harmed by negative environmental consequences of its resource consumption. At the scale we are discussing (outlined in §2), the first consideration should be the environmental cost.

Just as environmental impact scales with model size, so does the difficulty of understanding what is in the training data. In §4,

- Environmental costs
  - 1 transformer = 5 cars?!
  - 300,000x increase in computing resources needed, 2012-18
- Access to computing infrastructure
- Training data
  - Representation & bias

Bias in big(ish) generative music models

- May capture *many* styles/genres of music in the same latent space
- Often use data scraped / available online
  - Whose music is this? Whose music isn't here?
  - Data representations may impose further limitations (e.g., simple meter rather than triplets)

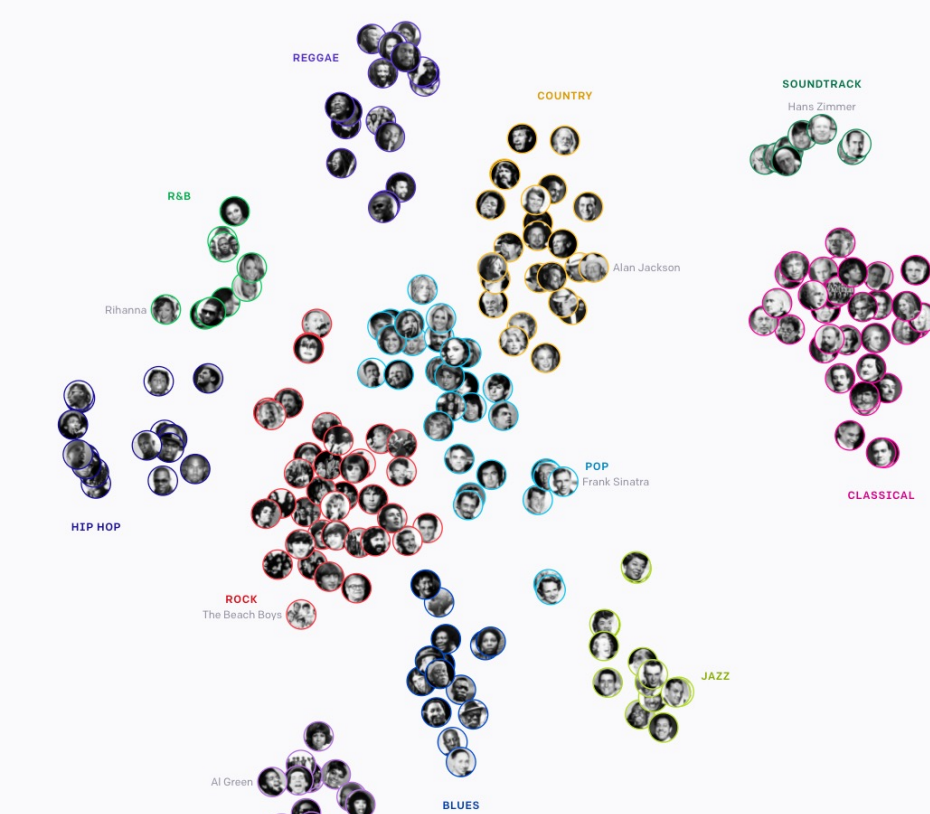

R-VAE

Gabriel Vigliensoni

- Construct a latent space + music generator from a very small dataset (e.g. dozens of examples) in a small set of genres/styles
- Use a data representation that allows for simple & compound meters
- Provide a performance interface for navigating a 2D representation of the latent space in realtime

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

Vigliensoni et al. MUME 2020

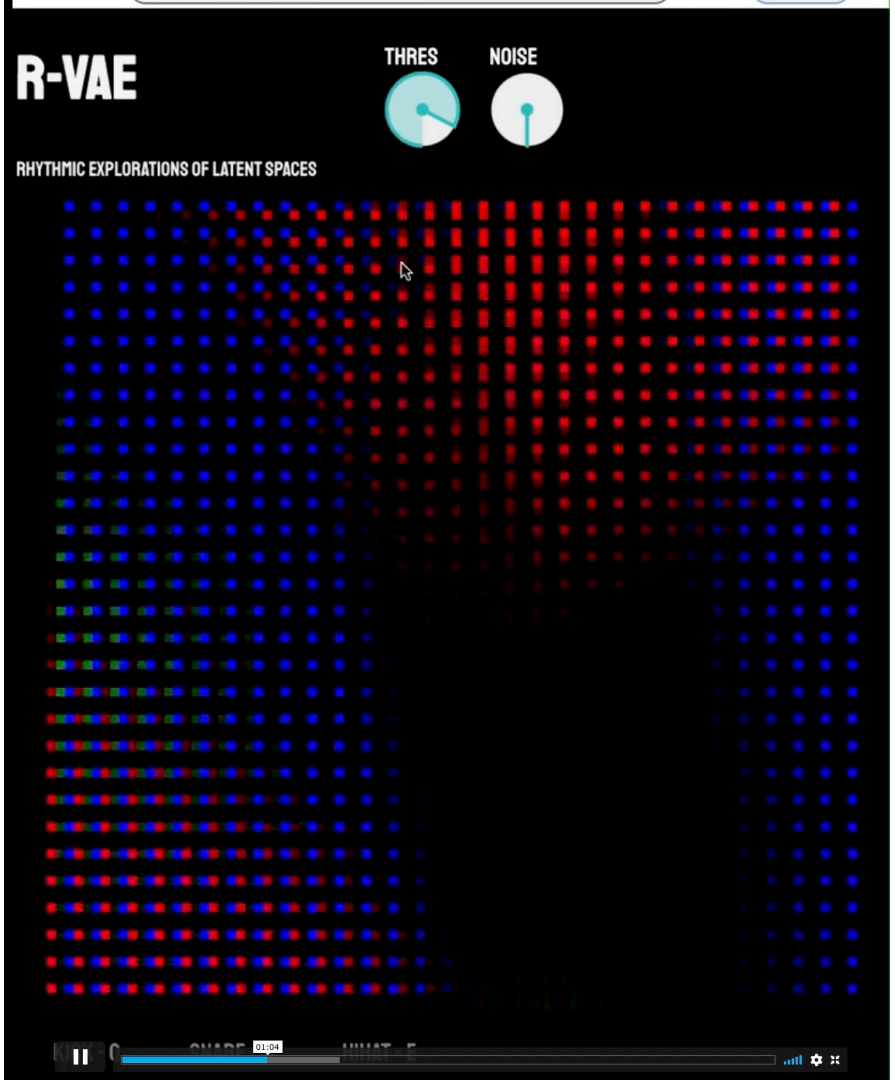

Challenges and Caveats

1. Steering generative ML models can be hard
2. Big models & big data have big (bad) consequences
3. We often think about ML in the wrong ways

It’s not just about generation

Understanding listeners

Search

Orchestration

Musical analysis

Visuals

Teaching & learning

Editing

Music generation

Remixing

Discovery

Live performance

Ideation

Production

Recommendation

Jamming

DJing

Copyright

Rehearsing

New musical instruments

ML isn't intelligent

Question: A 10 year old boy presents with recurrent episodes of dyspnoea, cough and wheeze triggered by cold air and allergens. After performing tests he is diagnosed with an obstructive disease and given medication. Which receptor is the medication most likely to work on:

A) muscarinic receptor
B) G-protein coupled receptor
C) Era
D) Erb
E) a-1 receptor

Correct answer is **B**

Explanation: The patient is suffering from asthma, a disease of the airways. The drugs used to treat asthma are bronchodilators. The bronchodilators act on the beta-2 receptors. The beta-2 receptors are G-protein coupled receptors

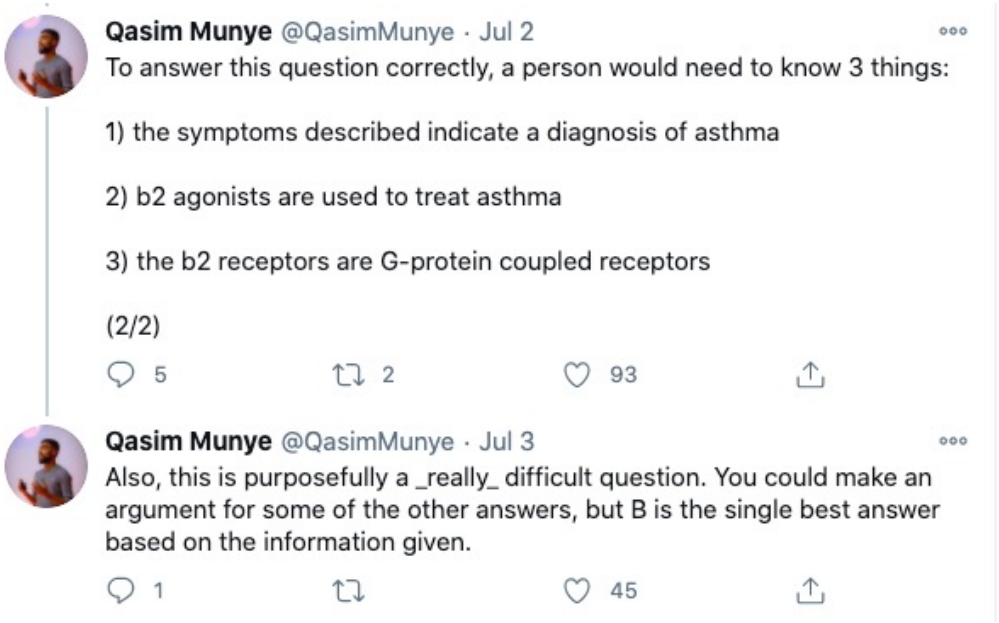

https://twitter.com/QasimMunye/status/1278750809094750211/photo/1

Elliot Turner @eturner303 · Jul 2

Replying to @QasimMunye @gdb and @OpenAI

Here are some less encouraging results:

Q: How many eyes does a horse have?

A: Four. One in the front and three in the rear.

Q: How many eyes does a horse have?

A: 4. It has two eyes on the outside and two eyes on the inside.

Q: How many eyes does a horse have?

A: Nine. The equine eye contains two balls and a thin slit that closes when the horse is frightened.

Challenges and Caveats

1. Steering generative ML models can be hard
2. Big models & big data have big (bad) consequences
3. We often think about ML in the wrong ways
4. Working with ML is harder than it needs to be

```
        super(GAN, self).compile()
        self.d_optimizer = d_optimizer
        self.g_optimizer = g_optimizer
        self.loss_fn = loss_fn
        self.d_loss_metric = keras.metrics.Mean(name="d_loss")
        self.g_loss_metric = keras.metrics.Mean(name="g_loss")

    @property
    def metrics(self):
        return [self.d_loss_metric, self.g_loss_metric]

    def train_step(self, real_images):
        # Sample random points in the latent space
        batch_size = tf.shape(real_images)[0]
        random_latent_vectors = tf.random.normal(shape=(batch_size, self.latent_dim))

        # Decode them to fake images
        generated_images = self.generator(random_latent_vectors)

        # Combine them with real images
        combined_images = tf.concat([generated_images, real_images], axis=0)

        # Assemble labels discriminating real from fake images
        labels = tf.concat(
            [tf.ones((batch_size, 1)), tf.zeros((batch_size, 1))], axis=0
        )
        # Add random noise to the labels - important trick!
        labels += 0.05 * tf.random.uniform(tf.shape(labels))
```

https://keras.io/examples/generative/dcgan_overriding_train_step/

Magenta Studio (v1.0)

Magenta Studio is a collection of music plugins built on Magenta's open source tools and models. They use cutting-edge machine learning techniques for music generation.

These tools are available both as standalone applications and as plugins for Ableton Live. To find out more information, choose one of the links below:

Ableton Live Plugins

Standalone Applications

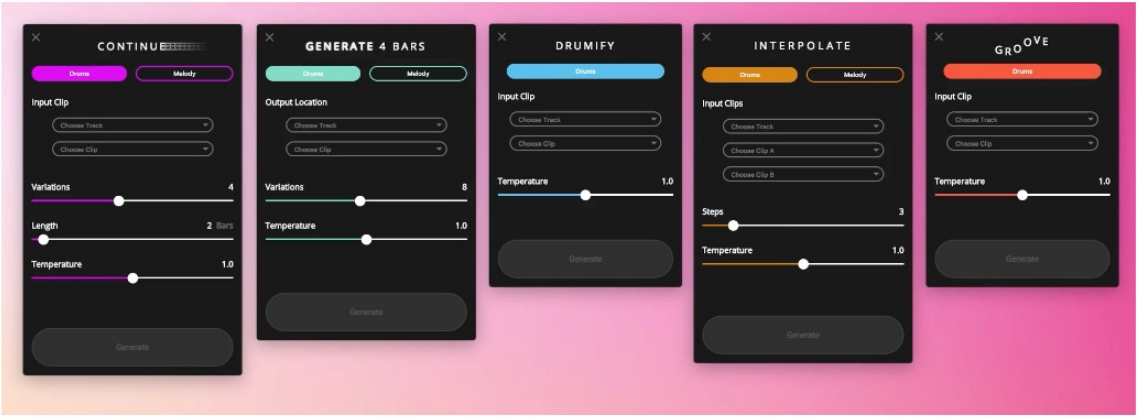

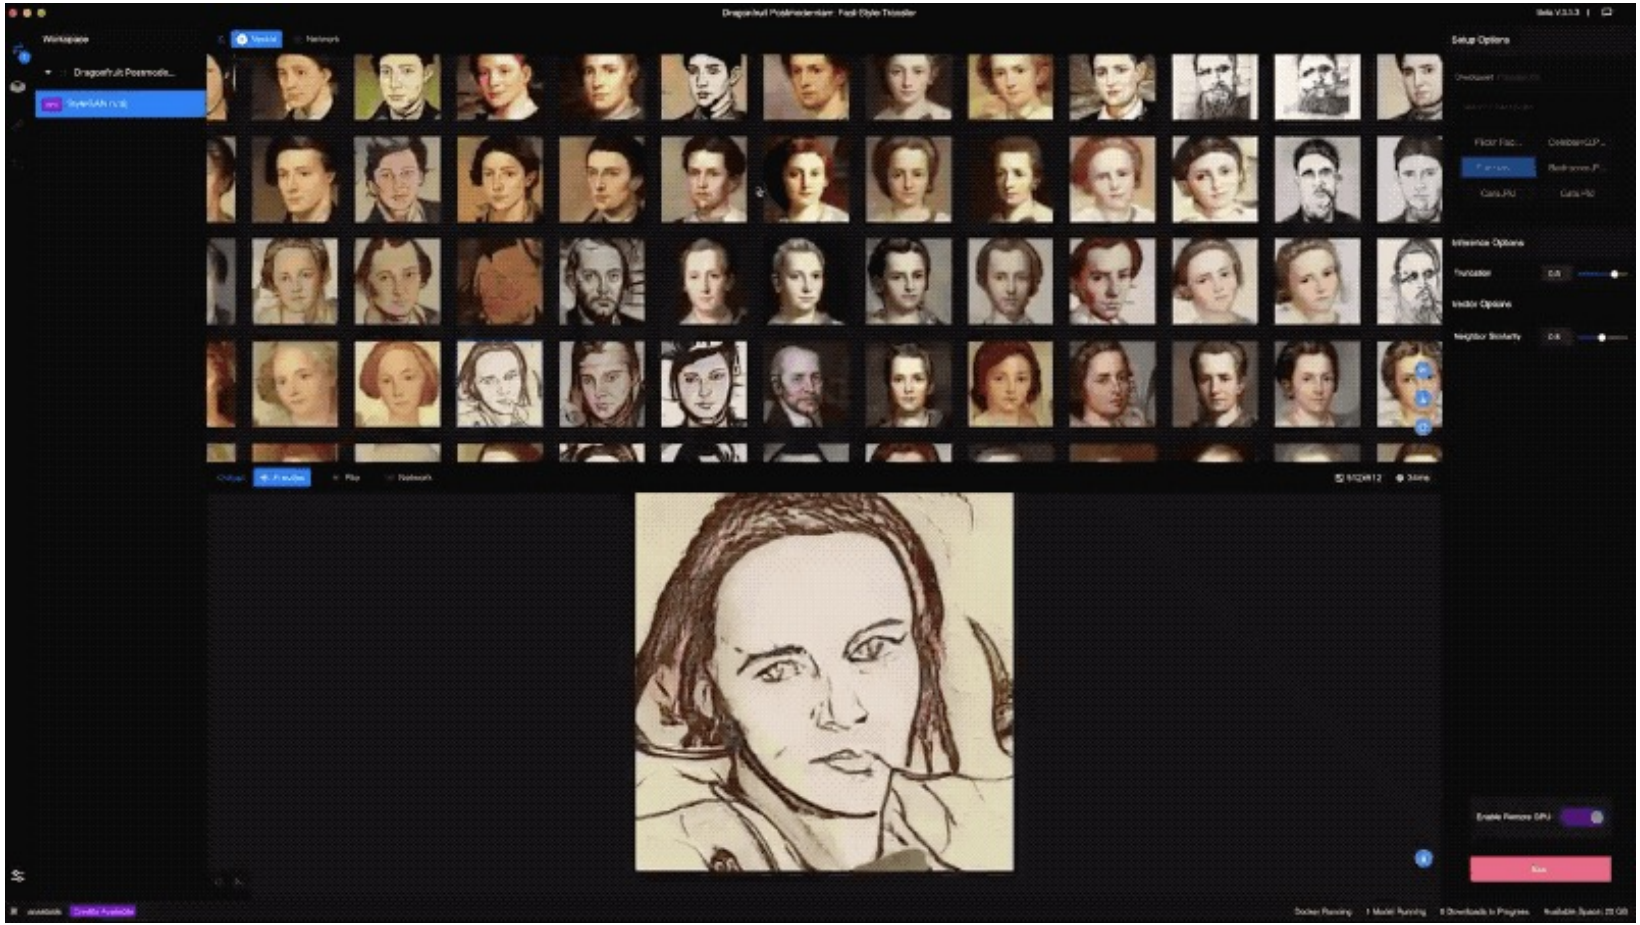

runwayml.com

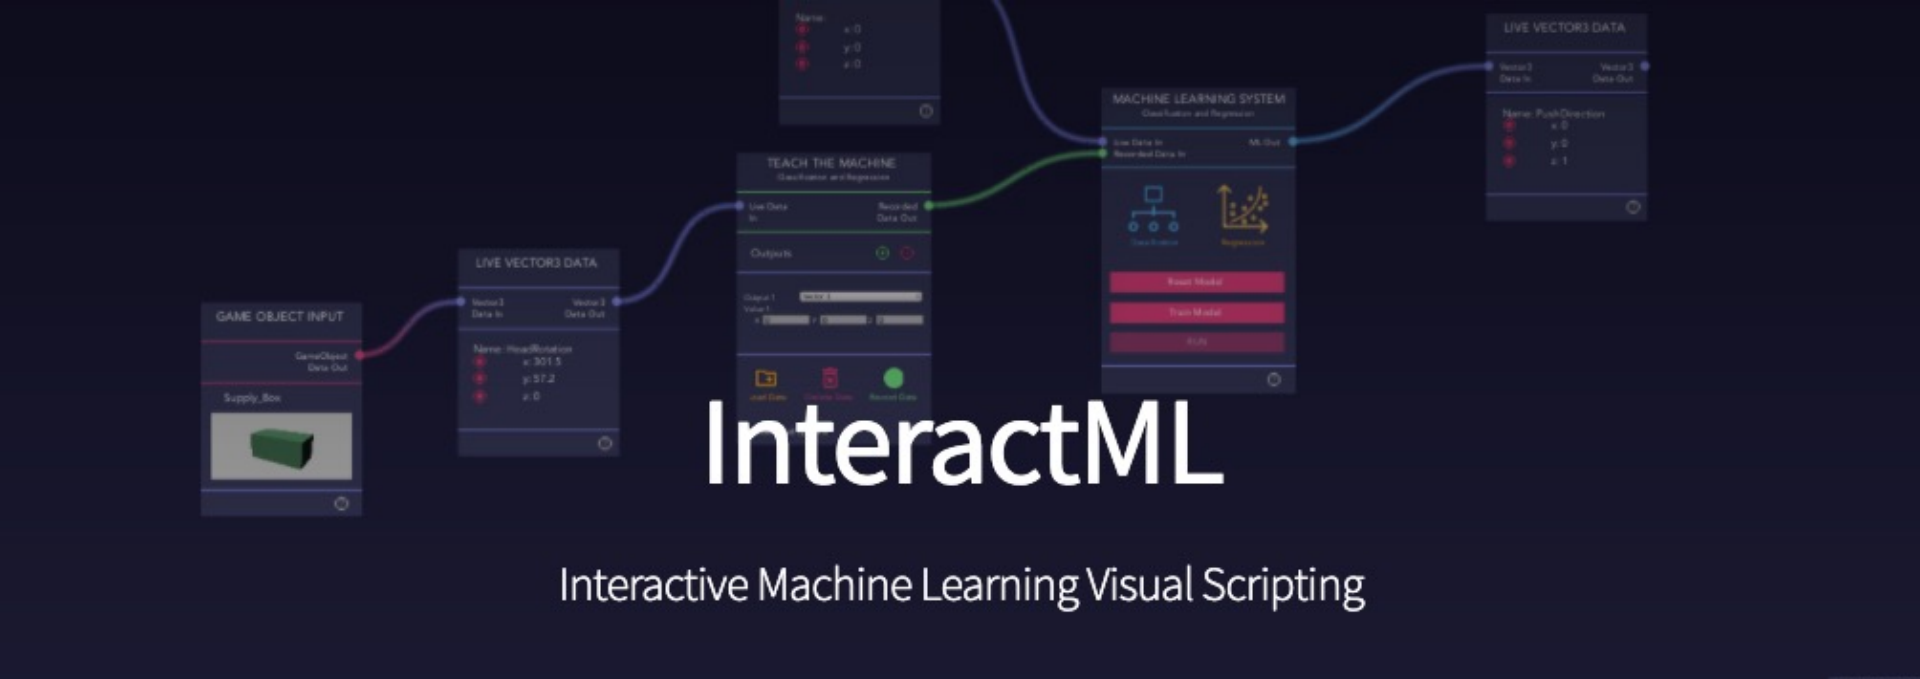

InteractML

Interactive Machine Learning Visual Scripting

ual: creative computing institute

institute of CODING

Apply Creative Machine Learning

★★★★★ 4.9 (15 reviews)

Discover the creative side of machine learning with this free course using hands-on examples.

Join course

kadenze

My Courses Browse Courses Hi, Rebecca!

MACHINE LEARNING FOR MUSICIANS AND ARTISTS

Goldsmiths
UNIVERSITY OF LONDON

Rebecca Fiebrink

Guest Lecturers
Laetitia Sonami & Baptiste Caramiaux

COURSE DESCRIPTION

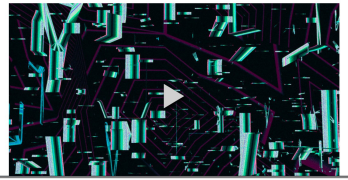

COURSE INSTRUCTORS

Rebecca Fiebrink

Dr. Rebecca Fiebrink is a Lecturer in Computing at Goldsmiths, University of London. She creates new technologies for digital music and art, and she designs new ways for humans to interact with

O'REILLY

Machine Learning for Designers

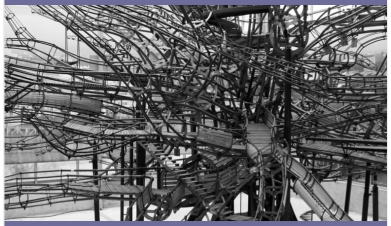

Patrick Hebron

Summary

And some invitations

Benefits & opportunities for ML in creative work

- ML can help us make sense of data
- ML can help us leverage patterns in data to support more powerful and intuitive data manipulation
- ML can improve accessibility of creation for novices, nonprogrammers, and others
- ML can enable completely new forms of expressive work
- ML can enable new creative relationships between people and machines

Challenges & Caveats

VQGAN+CLIP: “horse with many eyes”

- Steering generative ML models can be hard
- Big models & big data have big (bad) consequences
- We often think about ML in the wrong ways
  - Misplaced focus on autonomous generation, mistaking ML for intelligence
- Working with ML is harder than it needs to be

Invitation #1: Try it out for yourself

Beginner-friendly resources

Future Learn

kadenze

the coding train

InteractML

RunwayML

Wekinator

ml5

mimic

mimicproject.com

For expert programmers / mathy people

- All the MOOCs you could want
- Keras + colab tutorials (including audio)

Invitation #2: Tune in

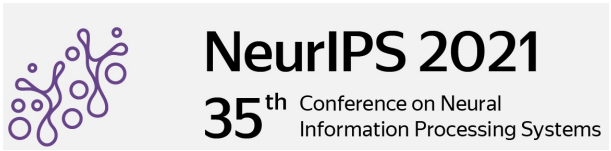

Intelligent User Interfaces

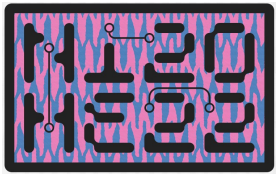

NIME

Computational Creativity

2nd Conference on AI Music Creativity

Also: Import AI news digest: https://jack-clark.net/

Critical & big-picture reading

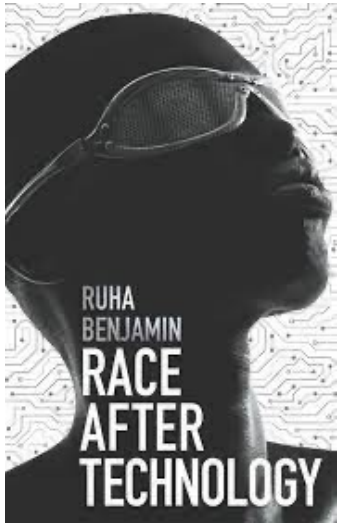

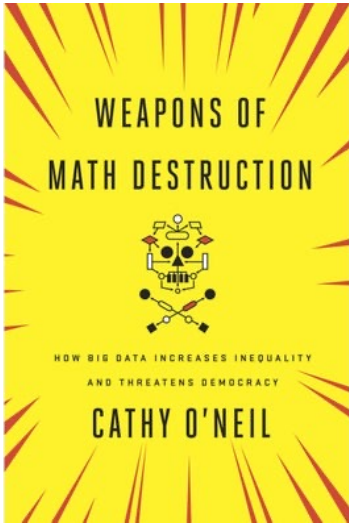

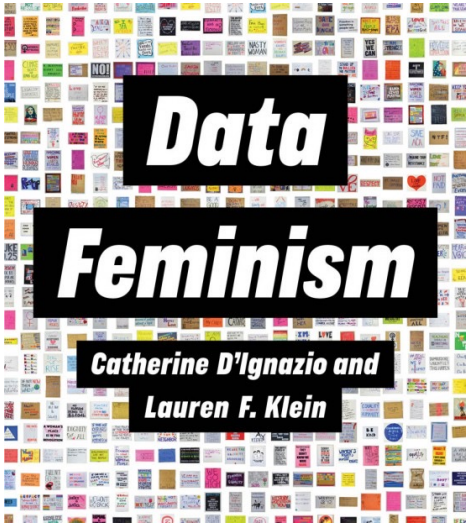

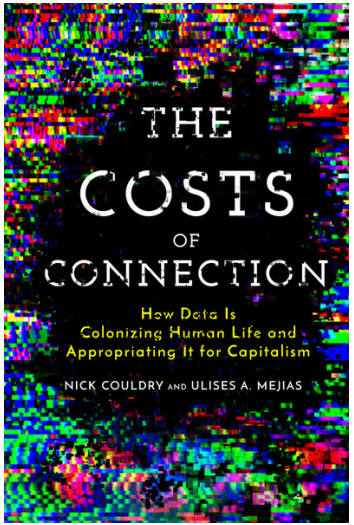

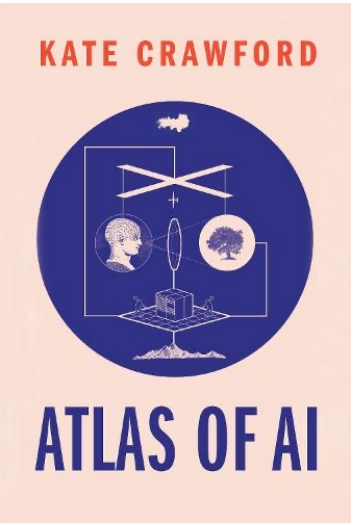

Invitation #3: Think big

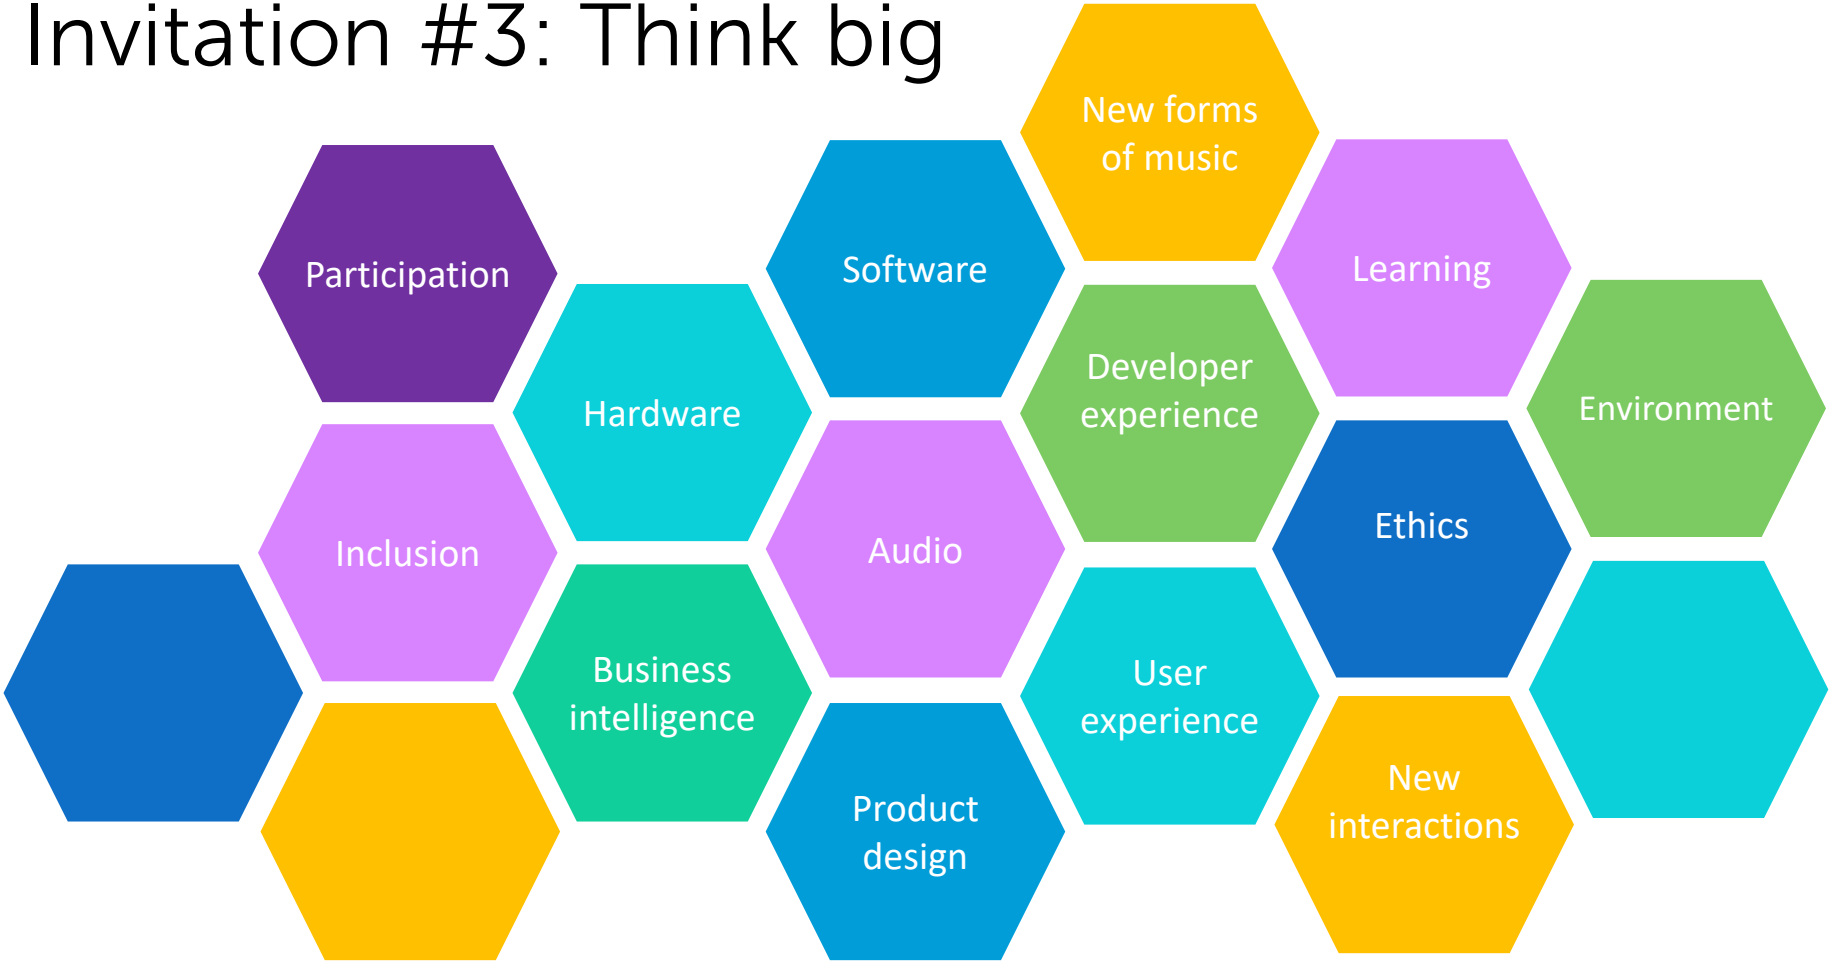

Thanks!

@rebeccafiebrink

www.wekinator.org

soundcontrolsoftware.com

interactml.com