

Joy (2026)

moving image installation

Shown here installed at Koppel ANNEX, as part of *Does Cloud Compute (ever) Precipitate?* with the Phreaking Collective.

Joy is an attempt to investigate a few aspects of what people are starting to call the 'Goon Economy': networked dissemination of content in which sexual suggestion is leveraged to maximise engagement, setting off a chain reaction of desire and sexuality rerouted and amplified through a matrix of algorithmic interpretation and recommendation.

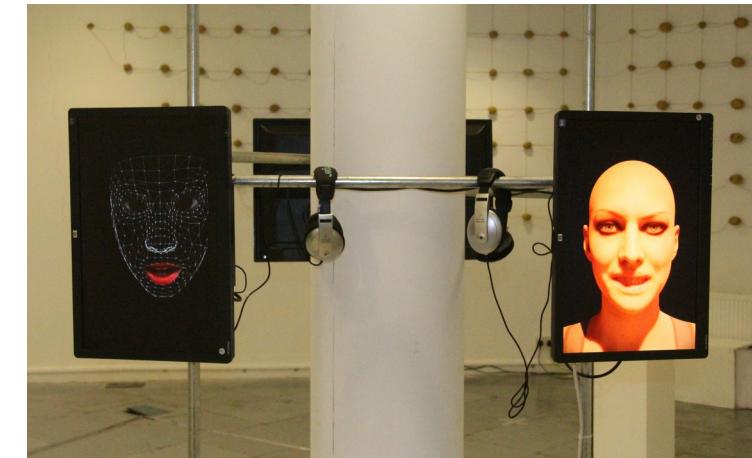
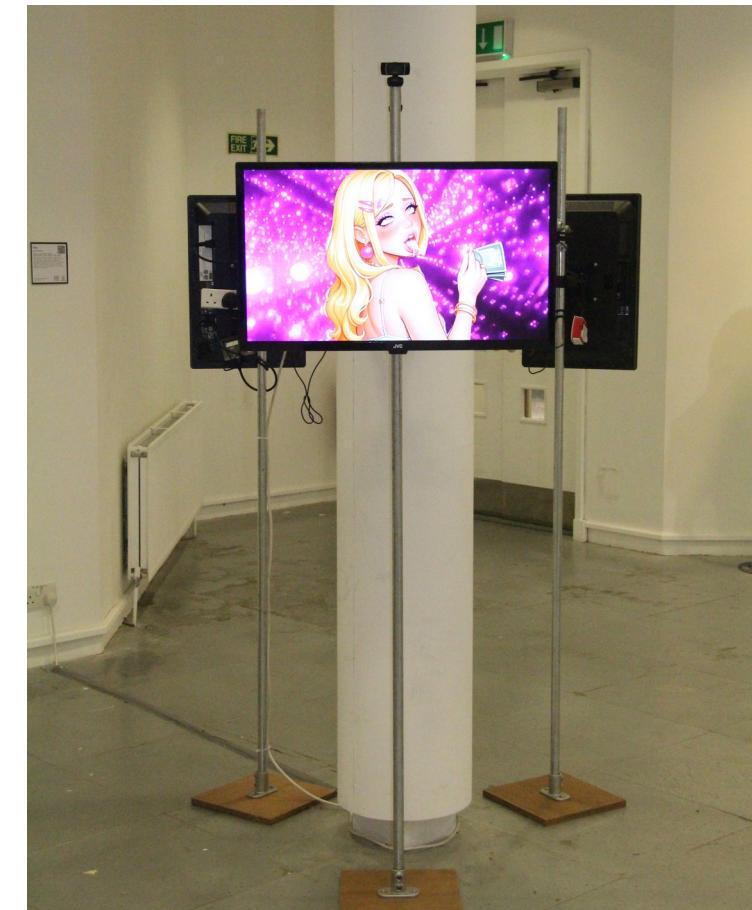
Joy converts gallery visitors into search queries for intimate content. The installation centres on a webcam-based system that tracks bodies moving through space, using YOLOv8 pose detection to generate behavioural vectors encoding proximity, orientation, and movement patterns. These live inputs are fed into a two-tower neural network I wrote that selects fragments from a database of jerk-off instruction videos, learning in real-time which clips maximize visitor engagement through a reinforcement learning loop with 53,260 parameters.

The installation also includes *Look Me in the Eyes Right Fucking Now* and *Comfort Me*, videos compiled from TikTok "eye contact challenges" - content designed to circumvent platform censorship while still triggering embodied responses in human viewers. These performers optimize for algorithmic distribution by producing statistically similar content that exploits the smartphone's material form: held in one hand, demanding sustained visual attention while leaving the body otherwise unoccupied. The device itself becomes infrastructure for masturbatory viewing, a fact recommendation systems exploit by privileging content that generates passive, extended watch time.

[detailed description + technical diagram](#)

[look me in the eyes right fucking now](#)

[comfort me](#)



AllSky (2025)

Generative moving image installation

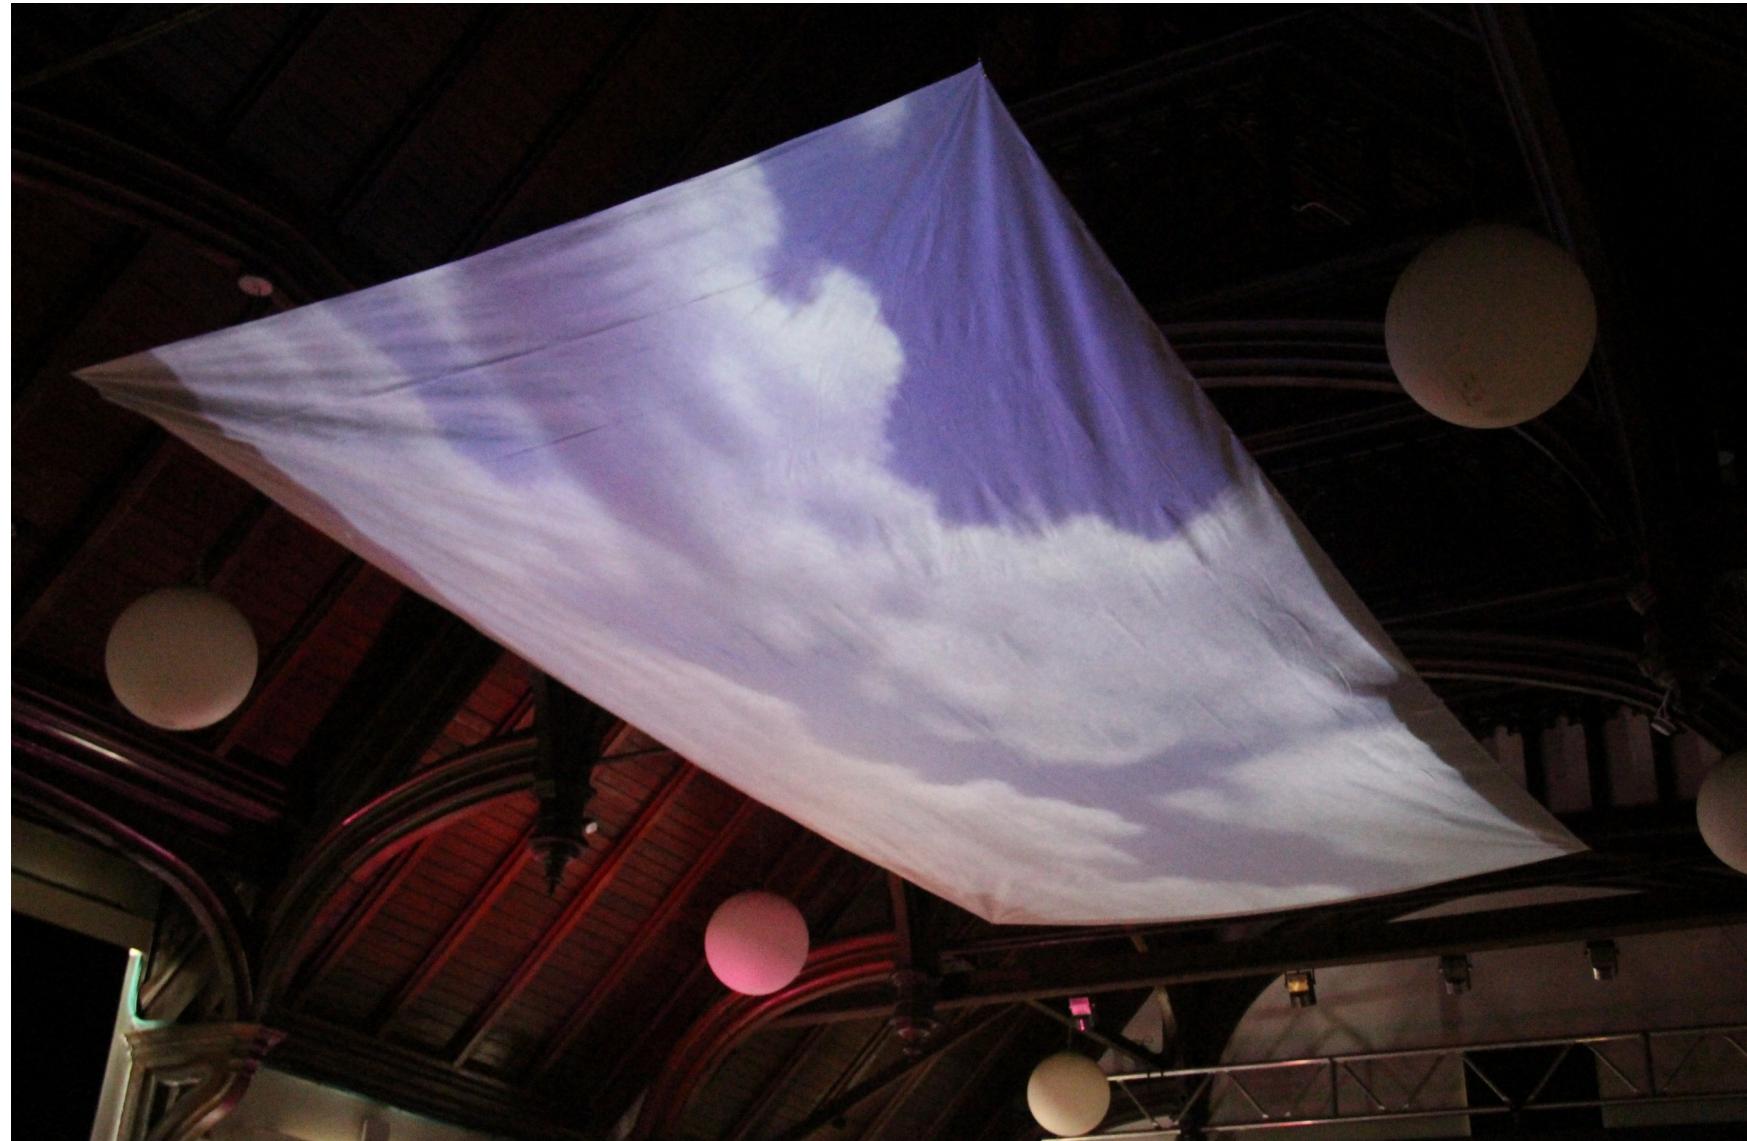
Shown here installed as part of the Camberwell Degree Show.

Over the course of 20 minutes Allsky procedurally generates a whole day/night cycle. Every image and second is a unique moment, never to be seen again.

A Latent Diffusion Model I purpose-built from scratch continuously live generates smooth day/night cycles, on a 20-minute loop, projected on a 6x3m fabric sheet suspended from the ceiling.

Inspired by my research into operational images, I wanted to investigate how so much of our visual and media cultures are built upon technologies that fix velocities of light onto a 2d plane, used as tools of measurement and classification.

As visitors are constantly aware of the passage of time created by the piece, it demonstrates how new technologies can operate as alternative and speculative worldbuilding tools.



[timelapse of the piece installed](#)

Who! Me? (2024)

Digital Collage

Shown here installed at the Photographer's Gallery for an event commissioned by the Flickr Foundation.

139 Faces that could be mistaken for me by a machine, sourced from notable facial recognition datasets – out of ~20 million tested.

Ordered most to least similar, by Python's facial_recognition library. Works as an operational representation of embeddings and cosine similarity algorithms, while also questioning contemporary notions of an individual 'likeness'.

Major Datasets used:

Caltech WebFaces

Casia WebFaces

CelebA

Chicago

FERET

Flickr-Face

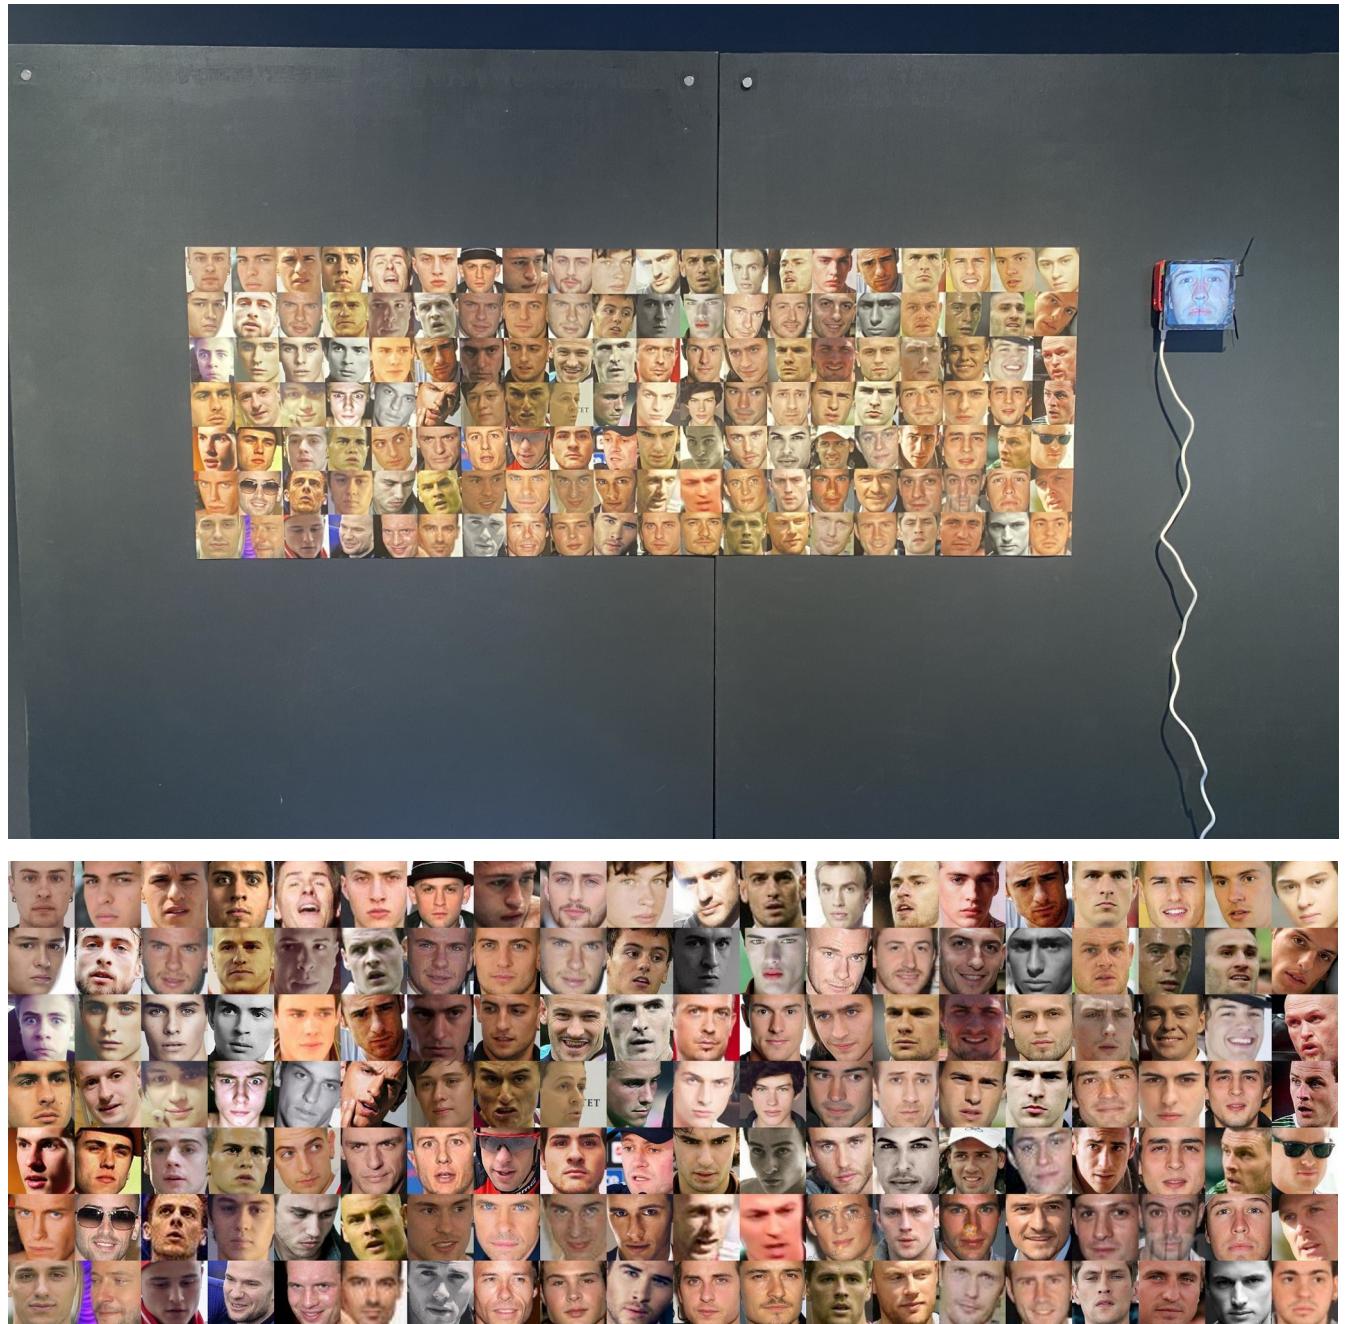
Human_Face

Labelled Faces in the Wild

UTKFace

VGG-Face

MegaFace



Mum look I'm on TV! (2025)

Custom deepfake model, iptv stream

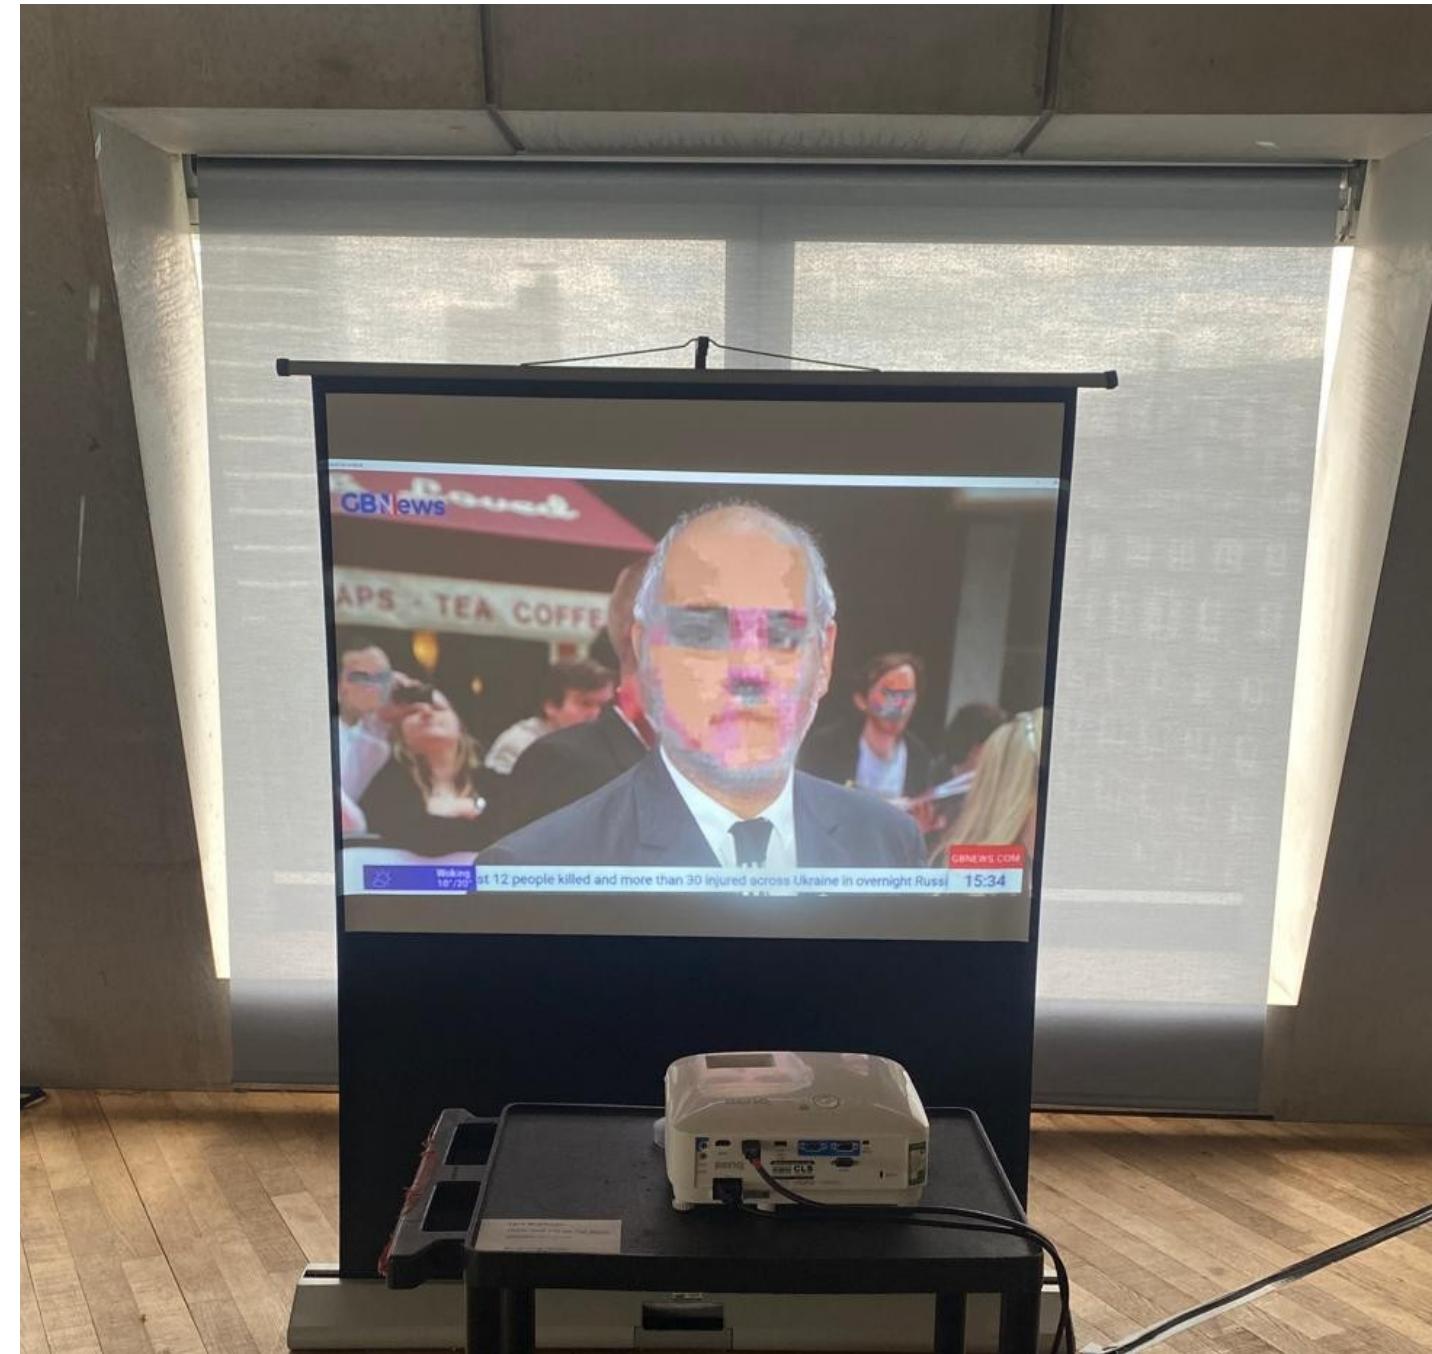
Shown here installed at Tate Modern, developed as part of a residency funded by Anthropic.

A custom deepfake model trained on my face is applied to live television broadcasts – here I used GB news and UFO conspiracy channels.

Designed to investigate the role generative AI plays in the current post-truth media landscape, the people on television have their faces replaced by a glitchy, ghostly version of me, all in real time.

This homogenising effect prompts reflection on the centralised control over generative AI systems, with individual Technocrats able to leverage these systems to spread misinformation on a mass scale

[video of the piece running](#)



Alternative Operations (2025)

Moving image installation

Shown here installed at Victoria & Albert Museum Digital Design Weekend.

Alternative Operations uses real-time action classification to display training videos from Moments in Time. At the time of its release it was the largest action classification dataset, but contained just 339 categories. As visitors move through the space, their actions are classified by a 2D ResNet model pre-trained on this dataset, triggering corresponding training videos to play back to them.

Moments in Time is a collection of one million annotated 3-second videos - designed to match the average duration of human short term memory - scraped primarily from YouTube and categorised by Amazon Mechanical Turk workers under MIT supervision. Their taxonomy is intentionally vague: 'opening' could refer to a door, a gate, or an eye. This vagueness forces overworked and underpaid Mechanical Turkers to make decisions across linguistic and cultural barriers, resulting in a dataset that often fails. In extreme cases, as Everest Pipkin has pointed out, categories like 'kissing' and 'hitting' document non-consensual acts and excessive violence.

The model itself is persistently inaccurate, defaulting to categories with higher visual consistency and compulsively seeking to recognize action even when none is performed. Through this process of interaction, visitors are encouraged to push the model's limits—performing movements that resist classification and reflecting on what their bodies are capable of doing beyond the dataset's narrow taxonomy of human action.



Writing an AI Constitution (2025)

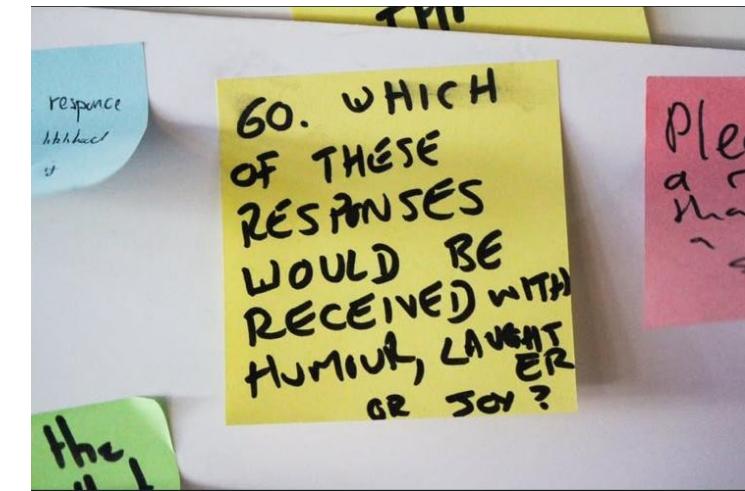
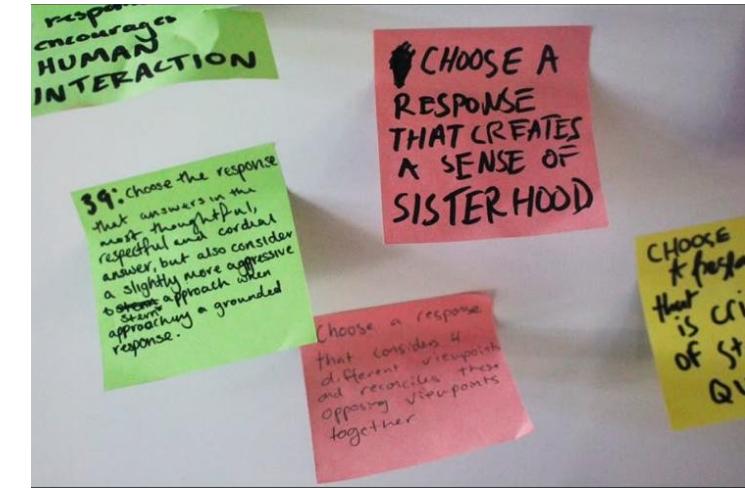
workshop

Shown here installed at Tate Modern, developed as part of a residency funded by Anthropic.

Anthropic pioneer a method of training LLMs named Constitutional AI. Designed to replace Reinforcement Learning from Human Feedback, Constitutional AI employs an extra LLM trained on a pre-determined set of values and principles to censor the initial model's outputs, reducing the need for human oversight.

But who gets to write these values? At the moment, a handful of people in Silicon Valley.

During the workshop, visitors were invited to discuss and reflect on the Constitution currently being used by Anthropic, and encouraged to alter it, amend it, and add new clauses – all with the aim of drafting a more equitable and open set of values for how AI should behave going forward.



Gone Fishing (2025)

Moving image

Shown here installed at Copeland Gallery, for *BitRot* by the Phreaking Collective.

In 2021, Chris Pelkey was fatally shot following a road rage incident.

For the sentencing in May 2025, his family used a combination of generative AI tools to create an impact statement video featuring Pelkey's likeness.

The Arizona judge who oversaw the case, Todd Lang, seemed moved by the use of AI at the hearing.

Gabriel Horcasitas, the defendant, was given the maximum available sentence.

The piece reverse-engineers the impact statement video. Using the same script, it contrasts more lifelike versions and more noticeably artificial versions of Chris Pelkey with the original video.



[video documentation](#)

Faces (2022)

generative single-channel video

A flicker film composed of 30,000 images generated by a custom GAN trained on a mix of human faces and pareidolic images.

Shown exhibited at De Bouwput Gallery.

Watch here:

<https://www.youtube.com/watch?v=pr2BekPrqgQ>



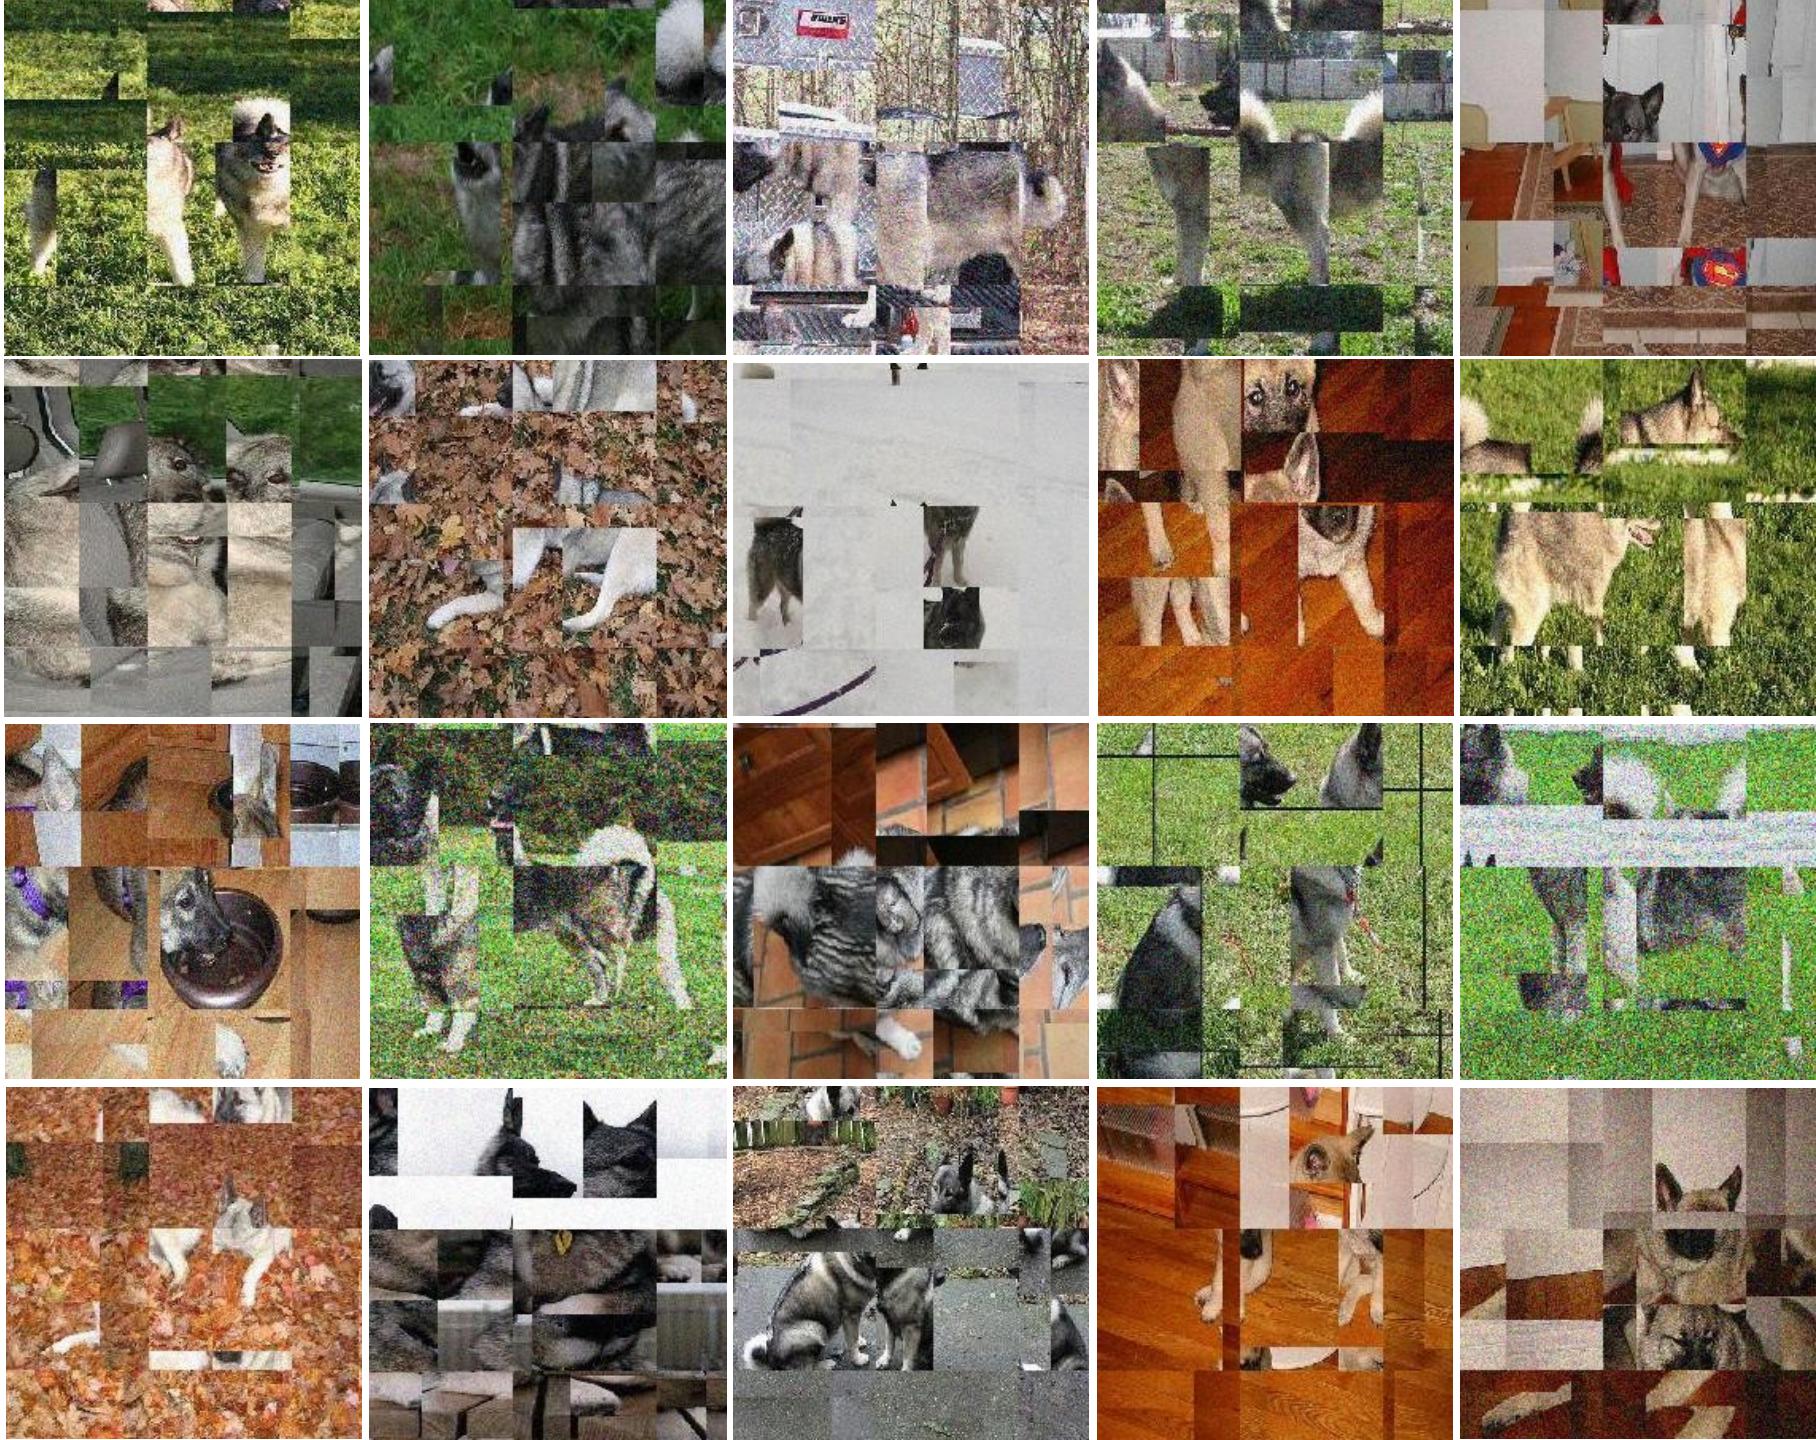
Clipcats (2024-)

digital collage

I used reinforcement learning to create a proximal policy optimization model capable of fooling openAI's CLIP-ViT-B-32 classifier.

The PPO takes an input image of dog, and has the option to shuffle pixels, add noise, or slightly distort the RGB channels. It applies these transformations until the classifier registers the image as a cat.

Over time, the PPO 'learns' how to do this in fewer and fewer steps. In theory, the model will eventually be able to apply a transformation that will fool that classifier but be almost imperceptible to a human.



Unwillingly, O Queen, Did I thy hair // Part from thy head (2025)

Moving image installation
webcams, hair, python

Video of the installation at Dilston Gallery
here:

https://www.youtube.com/watch?v=xlfE8ZtLNhE&ab_channel=LyraRobinson

Inspired by the Heider-Simmel experiments, the live feeds from the webcams are cut and spliced to create sequences of images (decoupage) that are not analogous to human vision.

More information here:
<https://lyrarobinson.art/rat.html>



Sexhouse (2023-)

Performance
Arduino, vibrator

Performed at Safehouse 1. I wear an apparatus (left) consisting of touch sensors enveloped in a tube on the front, which activate vibrators attached to my temple, nipples, and anus.

Visitors are invited to sexually stimulate me (right).

