

AllSky (2025)


Generative moving image installation

Shown here installed as part of the Camberwell Degree Show.

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.



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.



Alternative Operations (2025)

Moving image installation

Shown here installed in Millbank Tower.

Moments in Time, the largest and most comprehensive action classification dataset, contains just 339 actions.

Visitors have their movements classified in real time, with the corresponding training videos played back to them.

Through the process of interacting with the piece, visitors are encouraged to push the model and perform actions that are not easily classifiable, re-thinking what their bodies are capable of doing and reflecting on who holds the right to categorise people's bodies.



Writing an AI Constitution (2025)

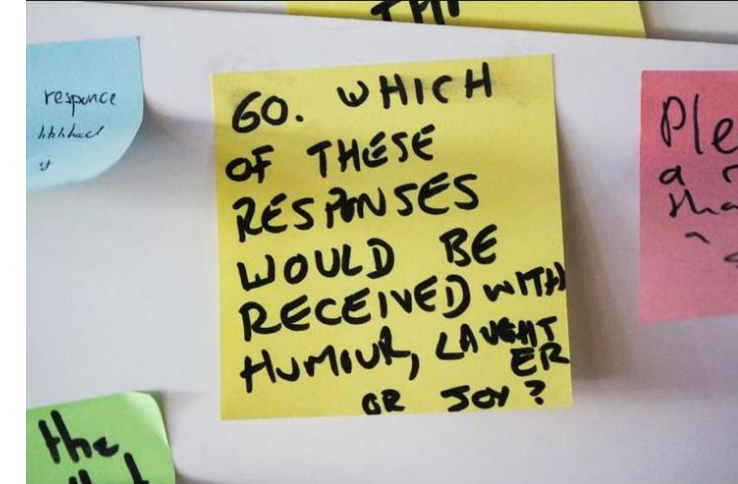
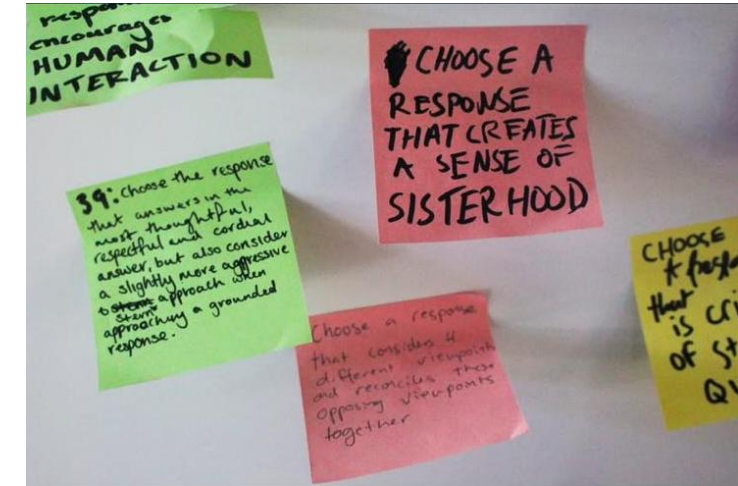
workshop

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

Anthropic pioneered 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.



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=xlE8ZtLNhE&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).

