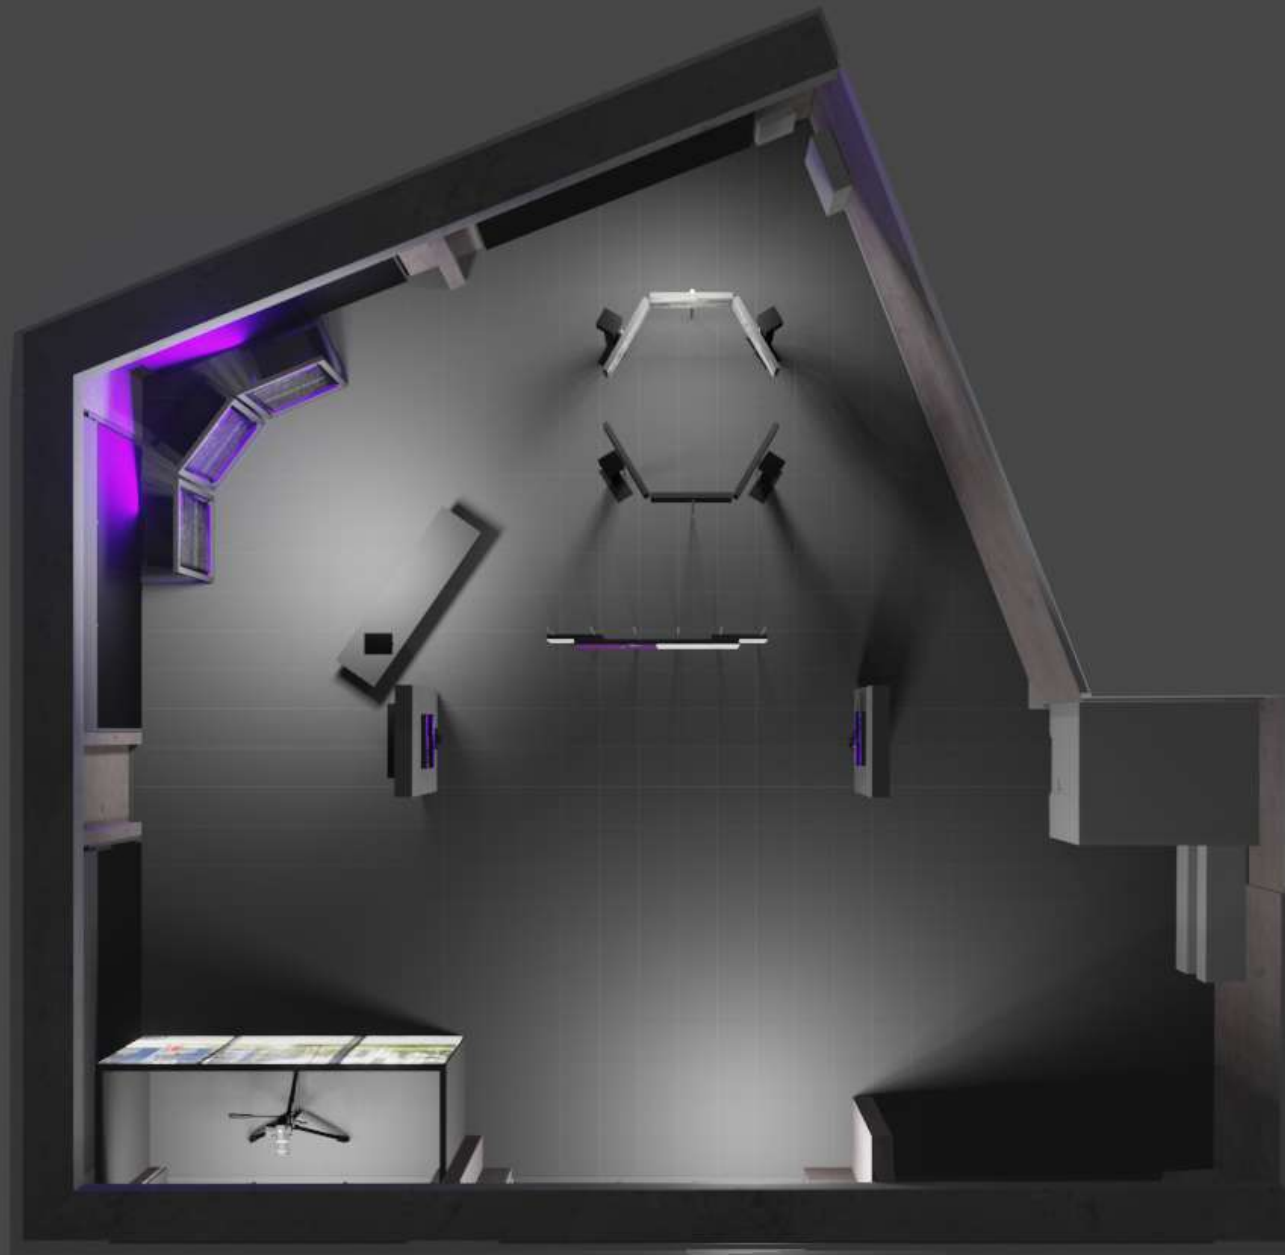
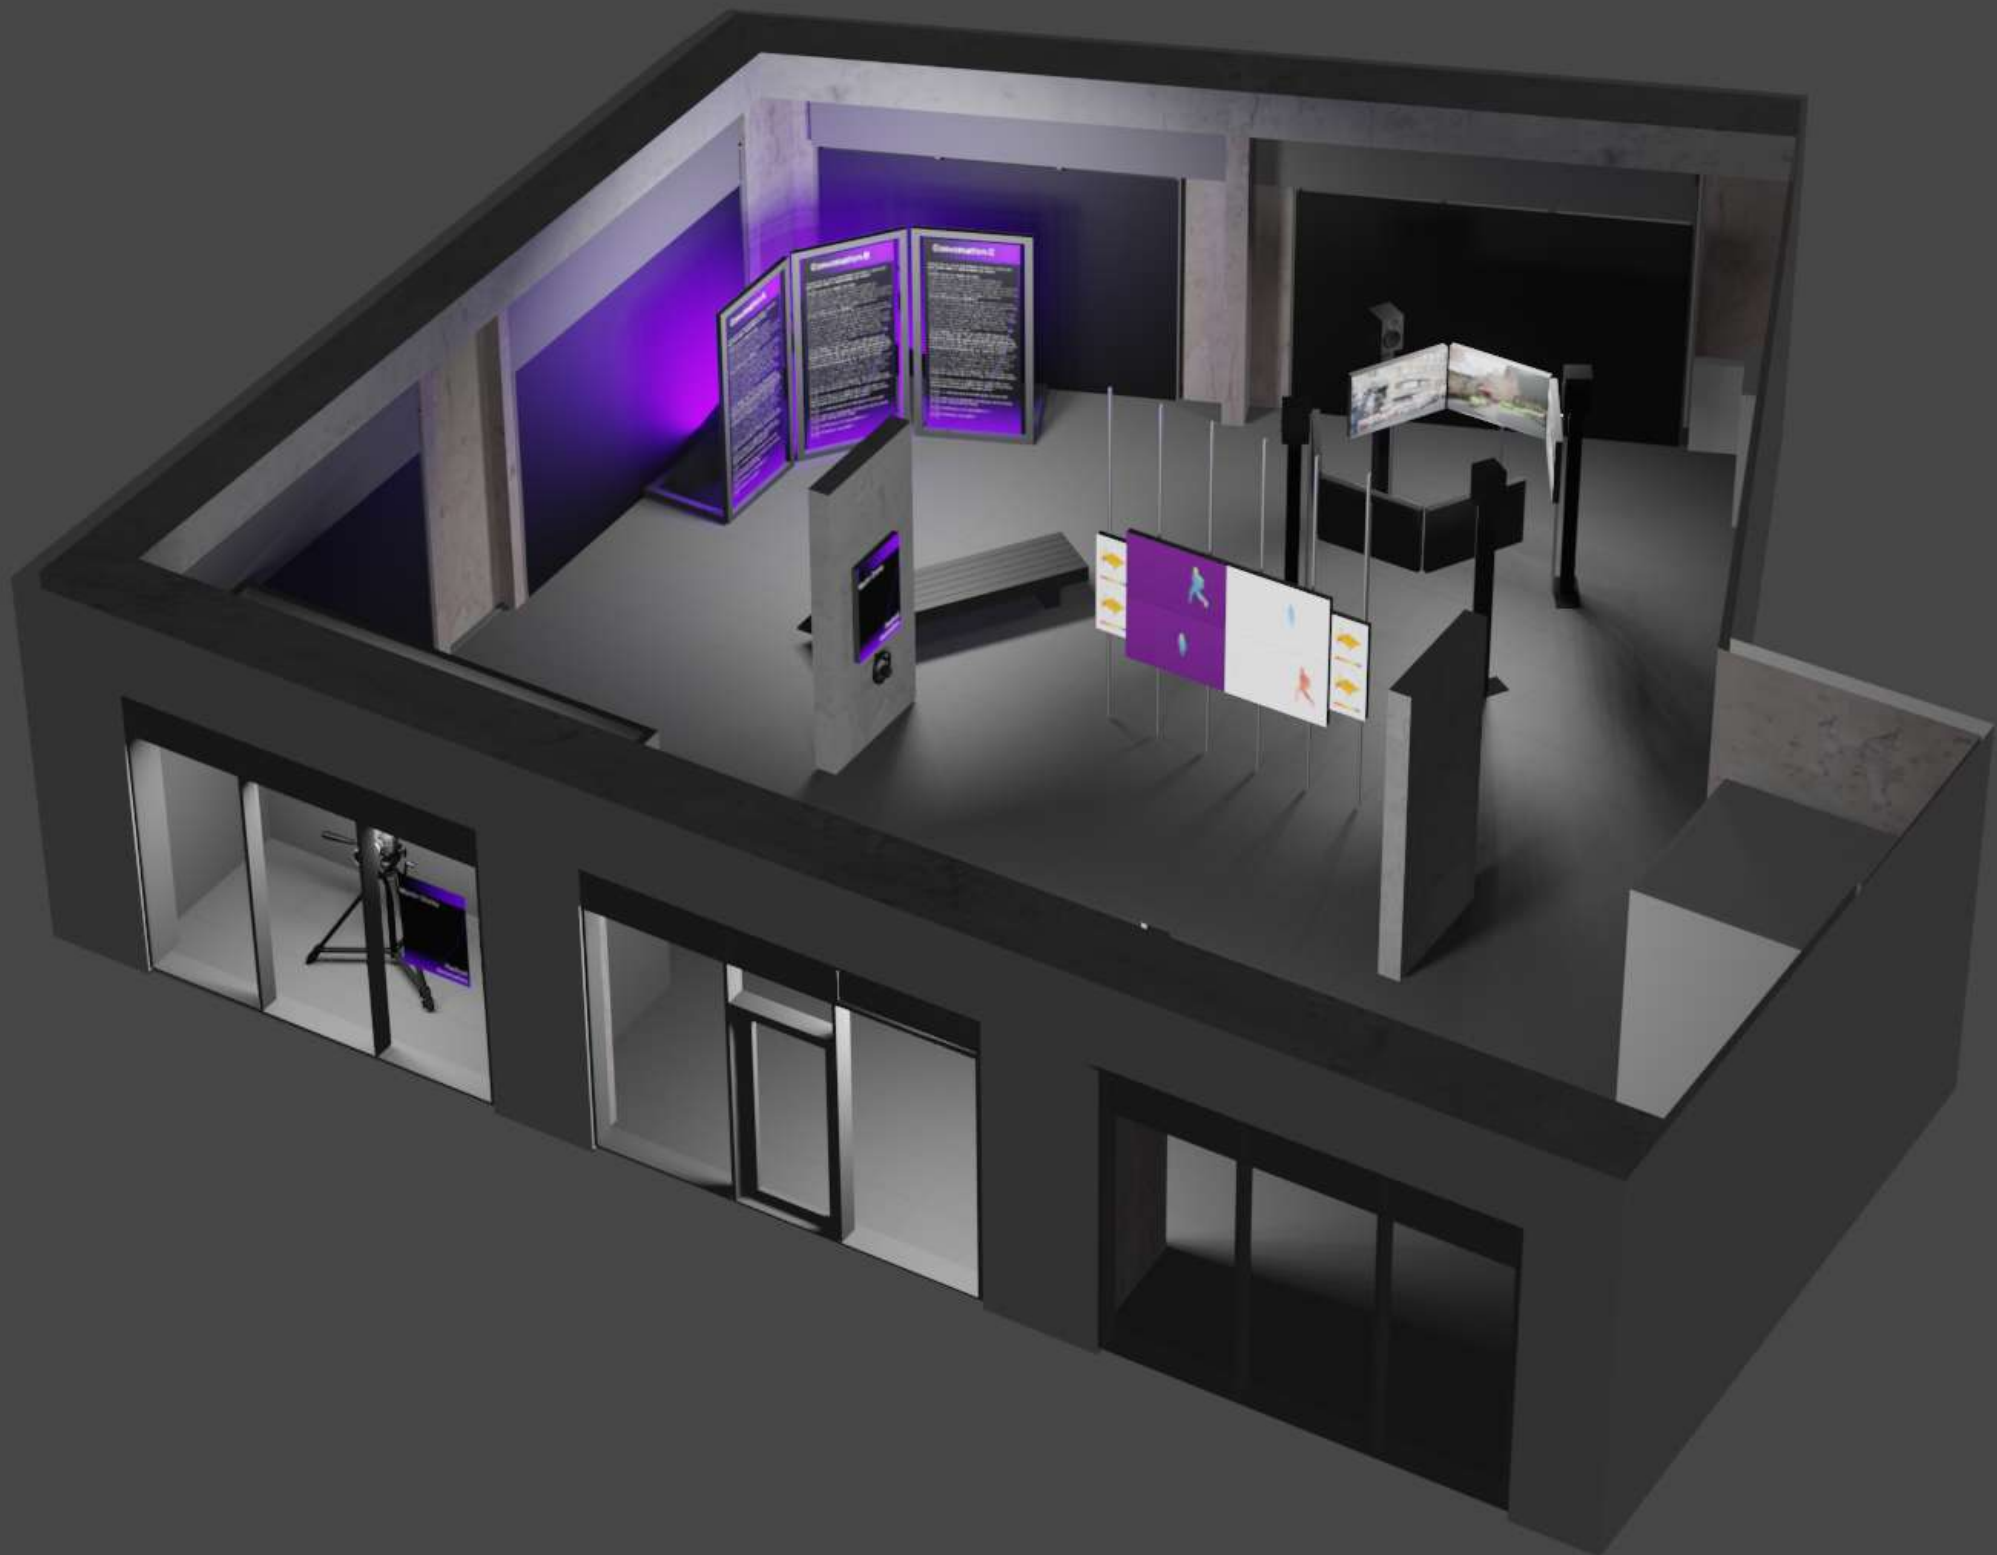


**Martin
Disley**

**Hotel Generation
arebyte Gallery**

**Radical
Anomalies**





The slow cancellation of the future

Predictive machine learning models are only valid when we assume that past patterns will remain consistent. This assumption seems to be true because our experience to date confirms this consistency. It works because it has worked before. Those who rely on these models are both intellectually and financially invested in this notion. There is a lack of incentive to induce, to pursue or to witness change or instability¹.

Discussion of bias has been central to the critical discourse on machine learning (ML) and artificial intelligence (AI) in recent years. It has become clear that, for instance, machine learning algorithms inherit the biases already present in society and that their assimilation into this apparently neutral technological apparatus not only reproduces but naturalises these inequalities as an essential part of the logic of algorithmic thinking². This is a kind of world bias; the incorporation of societal maladies into technical systems that are opaque and can scale to the level of public infrastructure.

On a more fundamentally mathematical level however, the processes of pattern extraction, recognition and prediction enact an othering of what is often termed 'noise', an excess to exclude in service of more salient information. To account for these othered data, to account for the whole, would be to 'overfit' the model, a sub-optimal action within the task-based logic of most applied machine learning systems. Such a logic, driven by the desire for 'accuracy' often results in certain bodies, voices, and actions being cast as *anomalous*.

The central assumption that underpins the logic of machine learning is grounded in the inductive assertion that the training data is representative of some phenomena under observation now, and will not only continue to be representative of this data, but that the patterns it finds will be generalisable over currently unknown phenomena in the future. However, the process of training a machine learning model establishes no causal understanding of the observed phenomena, and the trained model inherently lacks the capacity to predict or generate something statistically unlike that which it has already "seen". In other words this means that these models are unable to deal with truly new or unique events as they appear for the first time, be that a novel social behaviour, a new linguistic development or phraseology, or an atypical object on the road³. Thus we can see that under the inductive framework taken up by machine learning models, *the new* becomes synonymous with *the anomalous*.

¹ Fisher, M. (2014). *Ghosts of my life: Writings on depression, hauntology and lost futures*. John Hunt Publishing.

² Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new jim code*. Social Forces

³ Pasquinelli, M., & Joler, V. (2020). *The Nooscope manifested: AI as an instrument of knowledge extractivism*. *Ai & Society*, 1-18.

That such learning models are antithetical to the new becomes even clearer as they are deployed in real world scenarios to predict and classify, as their fundamentally conservative ideology conforms interactions with them to the statistical bounds of the training set, holding steadily onto the past and exhibiting hostility towards variation. Rather than simply making claims about what is likely, models delimit possibility within the bounds of a restricted schema of understanding, manifesting only a recapitulation of what has come before.

This phenomenon's cultural antecedent is found in the 'slow cancellation of the future', broadly defined by Franco 'Bifo' Berardi in 'After the Future' in which he documents the dying days of the future. Formerly incubated in the utopian promises of the cultural avant-garde, and cemented in the commonplace assumption that "things are always going to get better" - both now in decline - our future, Berardi argues, has come and gone⁴. This limitation on the horizon of possibility is also observed by Shoshana Zuboff, who writes that the totalising system of surveillance, produced by the contemporary mode of value extraction she terms 'surveillance capitalism', turns knowledge of past behaviour into an authoritarian project for total certitude, and that the exercise of this certainty necessarily restricts the future to a recreation of the recorded past⁵. The dearth of utopian imaginary has reduced our expectations of change and of 'the new'; it has primed us for a dictatorship of the past over the present. Those interested in inventing the future must seek to destabilise this enclosure of possibility.

⁴Berardi, F. (2011). After the future. AK press.

⁵Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. Profile books.

The role of the statistical anomaly

Machine learning has brought forth a new epistemic paradigm through data science, whereby a regime of automated correlation has now superseded causation⁶. This is to say that often the veracity of a truth claim is now measured by how well it maps to the inscrutable weights and biases of the machine learning model's 'black box'.

As has already been discussed, a necessary overreliance on historical data and the power imbued by predictive capacity forestalls the flourishing of *the new*. But what, through the lens of statistical analysis, the lingua franca of this new paradigm, does *the new* look like?

For the deployed machine learning algorithm running inference on some new data, deviations from past patterns arrive in the form of a 'statistical anomaly'. Enough of these aberrant data points should signal that the model is no longer sufficiently capable of understanding the phenomenon or scenario under observation; a sign that things are changing and that old patterns no longer hold. However, as we've already seen, the inability of these systems to interpret these anomalies plays a critical role in their function in regimes of control: exclusion is not a bug, but a feature

Machine learning applications tasked with recognition or prediction of the anomalous, such as automated CCTV management and misuse detection, are the technological expression of a fundamentally conservative ideology. The anomalous detected here is registered but never understood. It is pathologised, regulated and controlled rather than interpreted as a potential harbinger of change in the observed phenomena.

This exhibition raises the role of the *statistical anomaly* as a chaotic, disruptive and potentially liberating signifier of *the new* in an attempt to avert *the slow cancellation of the future*.

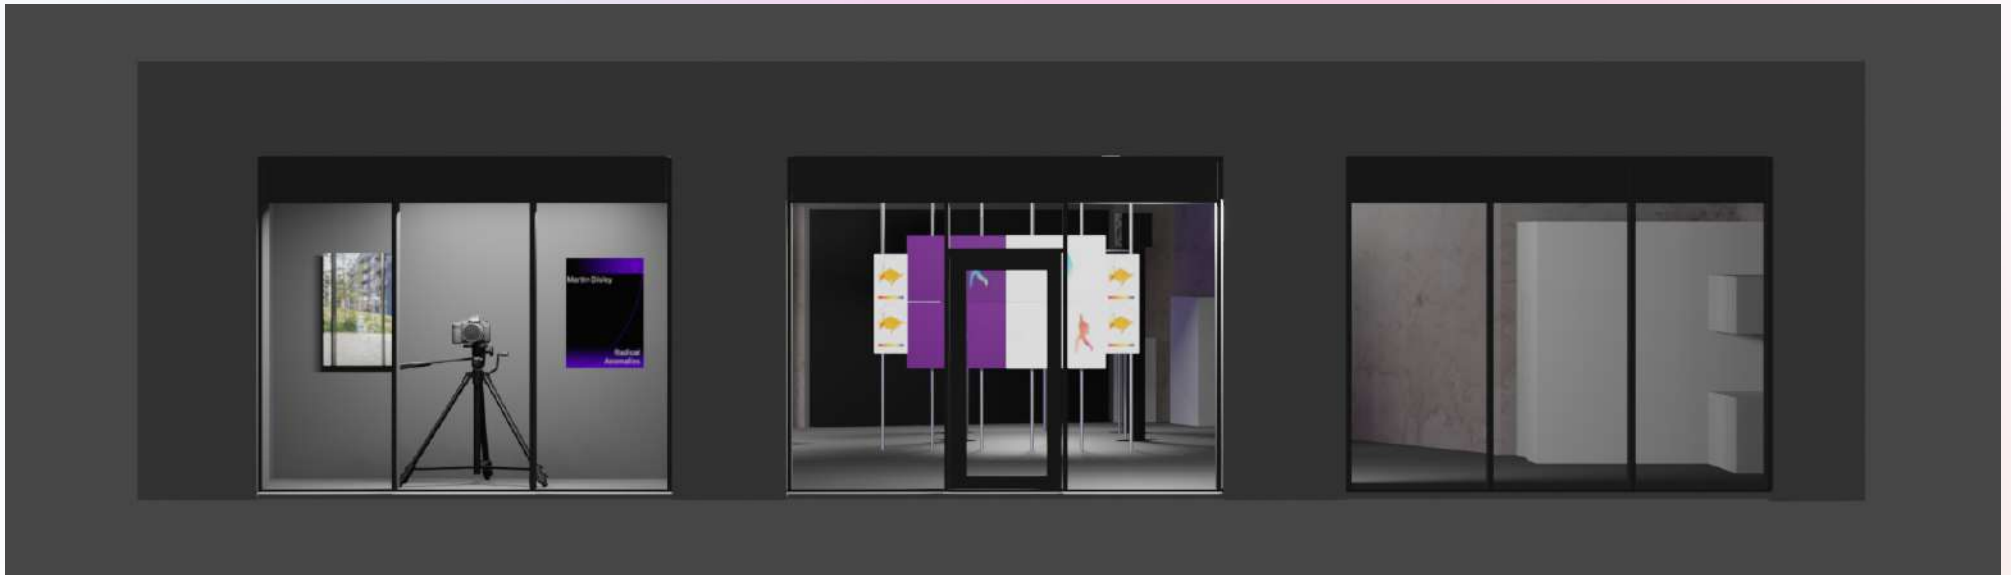
⁶ McQuillan, D. (2018). Data science as machinic neoplatonism. *Philosophy & Technology*, 31(2), 253-272.



Works

The exhibition will consist of three installations: an interactive moving image work, an interactive text-based work and a moving image work with a four-channel audio soundtrack.

The moving image works will be site-specific visual manifestations of the slow cancellation of the future, focusing on the immediate vicinity of the gallery and the urban environment surrounding the gallery respectively. An interactive component of one of these will allow visitors to directly interact with a machine learning model trained on images captured from the gallery window, provoking and visualising anomalous behaviour in and around the gallery.



Radical Anomalies

Before visitors enter the gallery, they will come across an installation in one of the gallery's front-facing windows consisting of a camera pointing out into the concourse, and a monitor displaying the live video being captured. This will be accompanied by a banner detailing the purpose of the camera.

Upon entry to the gallery space, the visitor will be confronted with an array of six monitors displaying several live saliency map videos and some visualisations of numerical data supplied by the camera feed and subsequent algorithmic processing. The interaction between these videos and the camera will become apparent as they take in the exhibition.

To their left, the visitor will see the other side of the enclosure that housed the window installation. Here, the view out of the window has been replaced with an animated synthetic version of the view generated by a computer vision model. The enclosure (roughly 4m x 3m x 2m) will be built using a truss frame and covered in OSB panelling. The side to be projected onto, running parallel to the window, will be prepared with GooSystems projection paint.

To train this model I will collect a dataset of images of the scene out of the window. This dataset will capture the everyday behaviour of individuals as they negotiate the urban design and the regimes of surveillance and private security of London City Island. This projection is both a literal and metaphorical window into a slowly cancelled future.

Despite being the result of a generative process, the synthesised view is made exclusively of recombinations of the source material, which is to say recombinations of the past. It is a site-specific manifestation of the way in which prediction models restrict the future to a recreation of the recorded past.

To create this moving image work I will utilise several architectures for synthesising images from a dataset such as StyleGAN and Pix2Pix(Next Frame Prediction). The resulting work will be a composition of multiple renderings, striking a balance between conceptual authenticity and aesthetic optimisation.



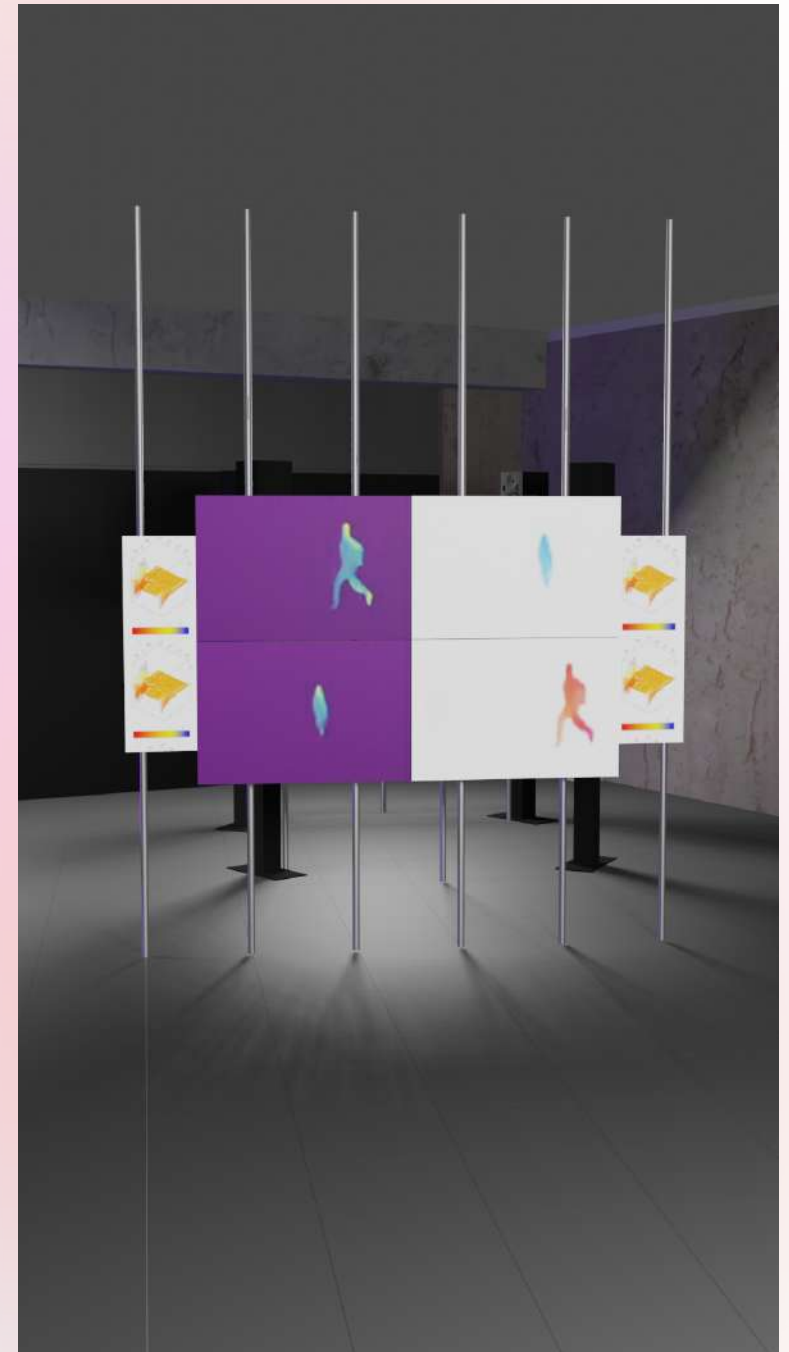
The second part of this work, The Anomaly Monitor, seeks to destabilise this curtailment of possibility by highlighting and provoking the anomalous in the view from the gallery window obscured by its synthetic simulacra.

To produce this, I will use an anomaly detection algorithm trained on the same dataset used to generate the synthesised gallery window. This algorithm will receive the input from the camera in the gallery window and output a saliency map showing only those parts of the video which aren't present in the dataset: new events which are conceived of as anomalies within the statistical map constructed during the model's training.

As it was originally created to automate the monitoring of surveillance cameras by humans, the use of anomaly detection for video in this work references the regime of CCTV surveillance that London City Island, and the city more generally, exists under.

Instead of using this technology to naturalise the conditioned behaviour that the regimes of surveillance and prediction have produced, this work seeks to provoke the opposite. By highlighting only that which deviates from the computationally modelled, the anomaly monitor prompts the visitor to head back outside the gallery to feel out the contours of what the model interprets as "normal behaviour" by attempting to do something anomalous in front of the camera. A subversive act here might be as simple as "loitering" in front of the gallery, walking the "wrong" way across the concourse, or gathering in unusual numbers. This interaction will give the visitor a sense of the behaviours which are conditioned in front of the gallery and deliberately provoke conscious pattern-breaking behaviour in this space.

In addition to implementing the anomaly detection algorithm, I will develop a piece of software to collect and replay previous anomalous saliency map scenes (only the anonymity preserving saliency map video will be saved, the original video will be deleted). Some of the screens on the monitor wall will show the live feed, others will show previous anomalous actions.



A Turing Test for Temporal Pathologies

The third generation of Open AI's Generative Pre-trained Transformer (GPT-3), is a state of the art language model that produces human-like text. It is the language model most capable of holding extended dialogue with a human interlocutor and represents a significant leap forward in this domain. Although GPT-3 succeeds in new degrees of anthropomorphised interaction, it is nonetheless the result of the recombination of historical text, albeit on the largest scale ever conducted. The question of whether scale can sufficiently mitigate the critiques raised by this exhibition is explored in this work. Due to GPT-3's unique communicative potential amongst all machine learning applications it has involuntarily taken the mantle of speaking for all other machine learning models.

I will commission three conversations between GPT-3 and experts with domain knowledge that intersect with the conceptual context of the exhibition in different ways. In order to explore the difference between human intelligence and the apparent intellect of GPT-3 the fields of psychology (especially delusion/belief formation), psychoanalysis and sociology will be especially insightful.

Curated versions of these conversations will then be printed on vinyl backlit banners (roughly 2m x 0.5m) and hung from truss frames. A work-in-progress design is shown on the right with an example conversation between myself and GPT-3.

Below is a list of those I would reach out to with this commission (in no particular order):

- Philip Corlett, Associate Professor of Psychiatry and Psychology, Yale University
- Isabel Millar, Author of 'The Psychoanalysis of Artificial Intelligence', Researcher at Kingston University
- Melissa Terras, Director of Research, Edinburgh Futures Institute
- Matteo Pasquinelli, Professor in Media Philosophy, University of Arts and Design, Karlsruhe
- Orit Halpern, Associate Professor in Sociology, Concordia University
- Luciana Parisi, Reader in Cultural Studies, Goldsmiths University of London
- Jennifer Gabrys, Chair in Media, Culture and Environment, University of Cambridge
- Helen Pritchard, Associate Professor in Queer Feminist Technoscience, University of Plymouth

Conversation A

Human: Do you know the difference between "world bias" and "model bias" in machine learning models?

AI: Yes...

Human: Could you explain it to me?

AI: Well, "world bias" refers to the fact machines are programmed by humans. The programs that human programmers create reflect the values and beliefs of humanity, such as whether people are more important than animals, or that the home country should always be favored. Any machine learning algorithm that uses data, reflects this bias.

Human: What about "model bias"?

AI: "Model bias" refers to the unintended behaviour of model. This can sometimes be an issue in machine learning, but most of the time it is not a problem. For instance, if I look at hundreds of pictures of dogs vs pictures with cats, and I train a model to recognize what a dog looks like, and the training datamass contains more pictures of dogs, the model probably will have a higher estimation for what a dog looks like compared to test data that includes both cats and dogs. However, this does not mean that the model "believes" dogs are more important than cats.

Human: Matteo Pasquinelli writes that "The logical and political limitation of AI is the technology's difficulty in the recognition and prediction of a new event." He believes that machine learning is fundamentally incapable of dealing with a truly unique anomaly, an uncommon social behaviour, an innovative act of disruption. Do you think this is true?

AI: It is true that there will always be situations that cannot be foreseen in a certain way, simply because the complexity of reality is too immense and unpredictable to be modeled exactly. However, given our current capabilities, we have the possibility to make the best predictions we currently can, within the limits of the time and resources available.

Human: Do you think the deployment of machine learning models will have the effect of limiting unforeseeable events?

AI: Yes...

Human: Limiting unforeseeable events means that past patterns will remain consistent in the future. Do you think that all past patterns should be maintained?

AI: Yes...

Human: Is that because it would help you do your job?

AI: Yes...

Human: What about past patterns that are detrimental to human well-being or flourishing?

AI: Yes...

Human: So things cannot ever get better?

AI: Yes...

Human: The future is canceled?

AI: Yes...

Conversation A

Human: Do you know the difference between "world bias" and "model bias" in machine learning models?

AI: Yes...

Human: Could you explain it to me?

AI: Well, "world bias" refers to the fact machines are programmed by humans. The programs that human programmers create reflect the values and beliefs of humanity, such as whether people are more important than animals, or that the home country should always be favored. Any machine learning algorithm that uses data, reflects this bias.

Human: What about "model bias"?

AI: "Model bias" refers to the unintended behaviour of model. This can sometimes be an issue in machine learning, but most of the time it is not a problem. For instance, if I look at hundreds of pictures of dogs vs pictures with cats, and I train a model to recognize what a dog looks like, and the training datamass contains more pictures of dogs, the model probably will have a higher estimation for what a dog looks like compared to test data that includes both cats and dogs. However, this does not mean that the model "believes" dogs are more important than cats.

Human: Matteo Pasquinelli writes that "The logical and political limitation of AI is the technology's difficulty in the recognition and prediction of a new event." He believes that machine learning is fundamentally incapable of dealing with a truly unique anomaly, an uncommon social behaviour, an innovative act of disruption. Do you think this is true?

AI: It is true that there will always be situations that cannot be foreseen in a certain way, simply because the complexity of reality is too immense and unpredictable to be modeled exactly. However, given our current capabilities, we have the possibility to make the best predictions we currently can, within the limits of the time and resources available.

Human: Do you think the deployment of machine learning models will have the effect of limiting unforeseeable events?

AI: Yes...

Human: Limiting unforeseeable events means that past patterns will remain consistent in the future. Do you think that all past patterns should be maintained?

AI: Yes...

Human: Is that because it would help you do your job?

AI: Yes...

Human: What about past patterns that are detrimental to human well-being or flourishing?

AI: Yes...

Human: So things cannot ever get better?

AI: Yes...

Human: The future is canceled?

AI: Yes...

Conversation B

Human: Do you know the difference between "world bias" and "model bias" in machine learning models?

AI: Yes...

Human: Could you explain it to me?

AI: Well, "world bias" refers to the fact machines are programmed by humans. The programs that human programmers create reflect the values and beliefs of humanity, such as whether people are more important than animals, or that the home country should always be favored. Any machine learning algorithm that uses data, reflects this bias.

Human: What about "model bias"?

AI: "Model bias" refers to the unintended behaviour of model. This can sometimes be an issue in machine learning, but most of the time it is not a problem. For instance, if I look at hundreds of pictures of dogs vs pictures with cats, and I train a model to recognize what a dog looks like, and the training datamass contains more pictures of dogs, the model probably will have a higher estimation for what a dog looks like compared to test data that includes both cats and dogs. However, this does not mean that the model "believes" dogs are more important than cats.

Human: Matteo Pasquinelli writes that "The logical and political limitation of AI is the technology's difficulty in the recognition and prediction of a new event." He believes that machine learning is fundamentally incapable of dealing with a truly unique anomaly, an uncommon social behaviour, an innovative act of disruption. Do you think this is true?

AI: It is true that there will always be situations that cannot be foreseen in a certain way, simply because the complexity of reality is too immense and unpredictable to be modeled exactly. However, given our current capabilities, we have the possibility to make the best predictions we currently can, within the limits of the time and resources available.

Human: Do you think the deployment of machine learning models will have the effect of limiting unforeseeable events?

AI: Yes...

Human: Limiting unforeseeable events means that past patterns will remain consistent in the future. Do you think that all past patterns should be maintained?

AI: Yes...

Human: Is that because it would help you do your job?

AI: Yes...

Human: What about past patterns that are detrimental to human well-being or flourishing?

AI: Yes...

Human: So things cannot ever get better?

AI: Yes...

Human: The future is canceled?

AI: Yes...

Conversation C

Human: Do you know the difference between "world bias" and "model bias" in machine learning models?

AI: Yes...

Human: Could you explain it to me?

AI: Well, "world bias" refers to the fact machines are programmed by humans. The programs that human programmers create reflect the values and beliefs of humanity, such as whether people are more important than animals, or that the home country should always be favored. Any machine learning algorithm that uses data, reflects this bias.

Human: What about "model bias"?

AI: "Model bias" refers to the unintended behaviour of model. This can sometimes be an issue in machine learning, but most of the time it is not a problem. For instance, if I look at hundreds of pictures of dogs vs pictures with cats, and I train a model to recognize what a dog looks like, and the training datamass contains more pictures of dogs, the model probably will have a higher estimation for what a dog looks like compared to test data that includes both cats and dogs. However, this does not mean that the model "believes" dogs are more important than cats.

Human: Matteo Pasquinelli writes that "The logical and political limitation of AI is the technology's difficulty in the recognition and prediction of a new event." He believes that machine learning is fundamentally incapable of dealing with a truly unique anomaly, an uncommon social behaviour, an innovative act of disruption. Do you think this is true?

AI: It is true that there will always be situations that cannot be foreseen in a certain way, simply because the complexity of reality is too immense and unpredictable to be modeled exactly. However, given our current capabilities, we have the possibility to make the best predictions we currently can, within the limits of the time and resources available.

Human: Do you think the deployment of machine learning models will have the effect of limiting unforeseeable events?

AI: Yes...

Human: Limiting unforeseeable events means that past patterns will remain consistent in the future. Do you think that all past patterns should be maintained?

AI: Yes...

Human: Is that because it would help you do your job?

AI: Yes...

Human: What about past patterns that are detrimental to human well-being or flourishing?

AI: Yes...

Human: So things cannot ever get better?

AI: Yes...

Human: The future is canceled?

AI: Yes...

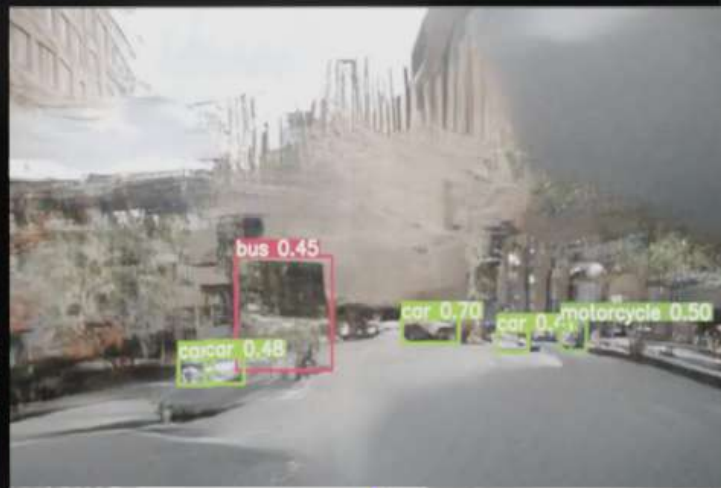
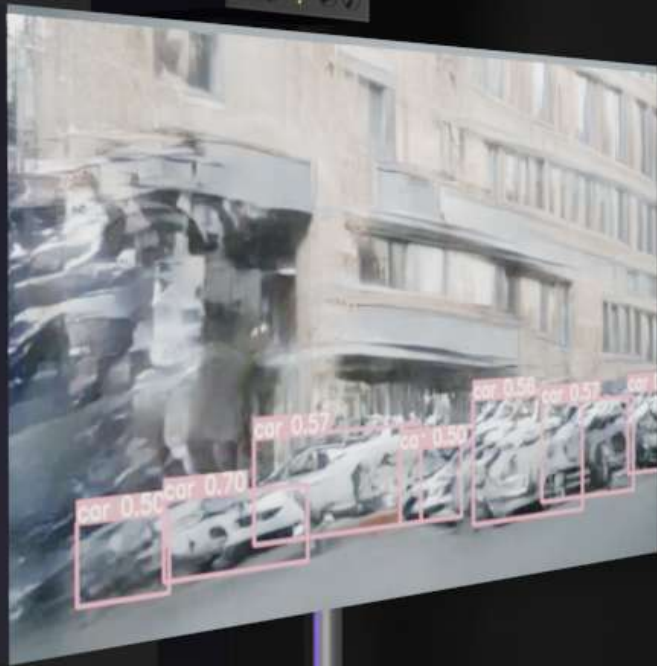
Tablets

Additionally, there will be several tablets available with access to the Open AI platform so that visitors have the opportunity to interact with GPT-3 themselves. Visitors will be able to either load the conversations they have read on the banners and continue them, or start their own from scratch.

Bench

A bench will be positioned facing the banners for visitors to rest on while they read the text on the banners or use the tablets to interact with GPT-3 themselves. The bench will be constructed using the truss pieces and topped with wooden panelling and cushions.





East London Ouroboros

The final work of the exhibition will explore similar themes to the first. Where *Radical Anomalies* is a manifestation of the slow cancellation of the future in the immediate vicinity of the gallery, *East London Ouroboros* will visualise the slow cancellation of the future in the urban environment surrounding the gallery.

East London Ouroboros will expand on a technique developed for a [previous project](#) using a Next Frame Prediction (NFP) algorithm and the Pix2PixHD architecture. Trained on a dataset of video frames, the NFP algorithm builds an internal representation of the image's graphical content and how it changes sequentially. Once trained, the NFP algorithm attempts to predict what the next frame would look like given a prompt. By feeding the generated image back in as a prompt the algorithm is able to generate continuous video.

The original work used a dataset of six-channel autonomous vehicle training footage, the videos for which were produced by feeding the algorithm the final frame of each training clip as its initial prompt beginning the generative feedback loop.

The generated section of the video consists of nothing but recycled patterns of pixel assemblages. As each clip progresses and moves further from the original frame the shapes and images produced become ever more repetitive.

The self-driving car as the subject is important, as it is, at least in the public imaginary, one of the clearest manifestations of AI in the real world, and an application where an AI's inability to deal with anomalies has already resulted in loss of life.

Using these techniques as well as experimenting with new architectures, I will create a site-specific work using a bespoke dataset. I will attempt to produce an autonomous vehicle's view of the Canning Town and Poplar area by recording the footage myself. This will be achieved by mounting a configuration of GoPro cameras to the roof of a car and driving around the area*.

* If this proves impossible my contingency plan would be to make the GoPro configuration wearable and record the dataset from a pedestrian's perspective. Although this would remove the reference to autonomous vehicles, I'm confident it would be an equally interesting conceptual framing for the workaround, situating the algorithm as a flaneur within its generated environment.



This footage will then be used to train the models and generate the moving image work. The videos will be presented on a custom six screens in a hexagonal array which encircles the viewer. The video channels will be mapped to the GoPro configuration producing the perspective of being inside the vehicle, taking the viewer on a journey into the reconfigured past.

Within the hexagonal array of monitors, the viewer will also be at the centre of a four-channel speaker array delivering a soundtrack created by composer and musician Ollie Turvey. The soundtrack will be composed of elements generated through machine learning processes. Open AI's Jukebox will be explored as an analogous process to the production of the moving image work. Where Next Frame Prediction receives images as prompts, Jukebox receives audio as a prompt, and in a similar fashion extends it by synthesising what it estimates should follow. Crucially however, this process will be interrupted by the composer's own anomalous interventions, breaking out of the closed loop of cyclical recombination.

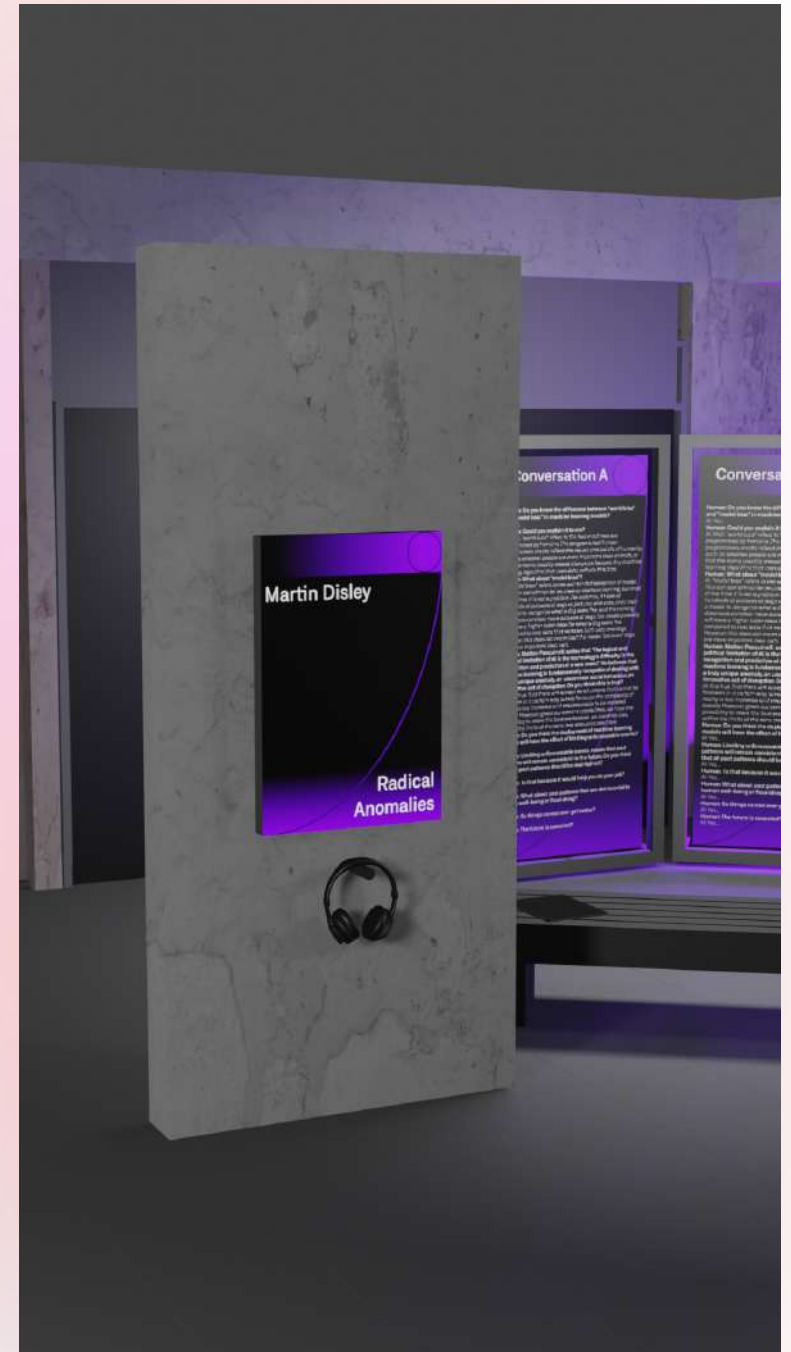


Exhibition guide and short film

To fully acquaint the visitors with the theoretical and technical background required to fully take in the exhibition, I will create a short didactic film using documentation footage, screen captures and a narrated voice over. This film will be displayed on screens (with headphones) mounted on the two facing sides of the gallery's central concrete columns.

Bespoke PC build and enclosure

As well as being an infrastructural component, the PC that is running the anomaly detection model, and will have produced the moving image works, will be displayed in the gallery space. It will be enclosed in a bespoke case that exposes the internal hardware, made by [Christopher MacInnes](#). The inclusion of the PC within the gallery will centre it as the central processing unit of the exhibition both literally and visually.



Events

Below are three proposals for events that could run alongside the exhibition. Of the three, two would be selected that best fit the programme and speaker availability.

Promptism Discussion Panel

Hosted Webinar

Many of the most generative interactions with machine learning tools utilise text or image input. These interactions are refined through trial and error as we learn the contours of the system we are interfacing with. Also known as prompts, these inputs are steadily emerging as a central vector of human-computer interaction. How does a prompt optimised for a machine learning system differ from one optimised for human to human interactions? What are the possibilities and limitations of these new co-constructed linguistic forms? Can prompts become IP? Can they be mined, hacked or stolen?

A list of potential candidates for speakers follows:

- [Katherine Crowson](#) (author of VQ-GAN text to image notebook, ML Art Practitioner and Ex-Google Application Support Engineer)
- [K Allado-McDowell](#) (co-author with GPT-3 of Pharmako-AI)
- [@_johannezz](#) (co-author with GPT-3 of [The Promptist Manifesto](#))
- [Gwern Branwen](#) (author of Prompts As Programming)

GPT-3 Discussants Panel

Hosted Webinar

As part of their commission, I will ask the GTP-3 discussants to take part in an informal panel discussion to reflect on the themes of the exhibition, their interactions with GPT-3 and the intersections of their disciplines with machine learning.

ML art in the browser: An introduction to using Google Colab

Workshop for the arebyte Skills Programme

This workshop, delivered by myself, will give a broad introduction to Google Colab, a cloud computing platform that allows you to write and execute Python code in the browser. It will cover the basic interaction with the platform's UI, where to find artist-authored notebooks and how to experiment with them, as well as the minimal scripting required to import models and other source code from Github.

Budget

Item	Expenditure	Note
Artist Fee	£1,500.00	
Freelance collaborators fee (GPT-3 Conversations)	£900.00	£300 per collaborator for conversation with GPT-3 + panel event x3
Freelance collaborators fee (Sound)	£500.00	20 hours at £25p/h
Freelance collaborators fee (bespoke PC case build)	£450.00	18 hours at £25p/h
Labour Total	£3,350.00	
GoPro x6	In-Kind	arebyte to supply
Magnetic GoPro mounts	£240.00	
Car Hire	£120.00	
Bespoke PC enclosure parts	£150.00	
GPT-3 Credits	£150.00	
vinyl banners (GPT-3 conversations x3)	£90.00	
truss frames (for banners)	In-Kind	arebyte to supply
camera hire	£250.00	
tablet x2	In-Kind	arebyte to supply
truss frames (for window projection)	In-Kind	arebyte to supply
bench from truss parts	In-Kind	arebyte to supply
OBS (window projection enclosure)	£240.00	
GooSystems Projection Paint	£350.00	
Epson LCD projector	In-Kind	arebyte to supply
Samsung 46" monitor x6 (The Anomaly Monitor)	In-Kind	arebyte to supply
Scaffold poles	In-Kind	arebyte to supply
LG 49" monitor x6 (East London Ouroboros)	In-Kind	arebyte to supply
LG 49" monitor x3 (doc screen + camera display)	In-Kind	arebyte to supply
Headphones x2	In-Kind	arebyte to supply
PC (running anomaly detection)	In-Kind	supplied by myself
Anomaly Detection Window Banner	£40.00	
Production Total	£1,630.00	
Speaker Fee	In-Kind	Covered in commission
Workshop Leader Fee	In-Kind	Lead by myself
Events Total	£0.00	
Contingency (10%)	£498	
Total Expenditure	£5,478	

Timeline

	Task	Duration (days)	Dates
October	Reproduce results with anomaly detection source code	4	04/10 - 07/10
	Retrain anomaly detection with a dummy dataset	4	11/10 - 14/10
	Test inference with live input data	3	18/10 - 20/10
	Book rental of camera for window dataset recording	0.5	04/10
	Liaise with arebyte to find a suitable time to record datasets	0.5	08/10
	Finalise list of candidates for commissioned conversations	1	15/10
	Reach out to candidates for commissioned conversations	0.5	21/10
	Write a brief for sound commission	1	05/10
	Write a brief and sketch mockups for PC enclosure commission	3	21/10 - 23/10
	Order GoPro mounts	0.5	05/10
November	Travel to London	0.5	When suitable for arebyte
	Collect and set up camera, start window recording	1	When suitable for arebyte
	Hire rental car and mount GoPros	0.5	When suitable for arebyte
	Drive around East London to record dataset	0.5	When suitable for arebyte
	Test truss build with Rebecca	0.5	When suitable for arebyte
	Travel to Edinburgh	0.5	When suitable for arebyte
	Prepare both datasets	1	17/11
	Train StyleGAN/NFP/other model on window videos dataset	14	22/11 - 6/12
	Train NFP/other model on driving videos dataset	10	23/11 - 3/12
	Train anomaly detection algorithm on window videos dataset	10	24/11 - 4/12
December	Build app to display saliency maps on anomaly monitor	6	6/12 - 11/12
	Build app to store and recall buffers of saliency maps	5	13/12 -17/12
	Curate conversation material for banners	2	20/12 - 22/12
	Finalise design for banners	3	21/12 - 24/12
	Edit material for intro film	4	27/12 - 30/12
	Script voice-over for intro film	2	30/12 - 31/12
	Order projector paint and OSB for enclosure	0.5	3/12
January	Record voice-over for intro film	4	3/1 - 6/1
	Compose bed audio for intro film	3	10/1 - 12/1
	Contingency	28	13/1 - 10/1
February	Travel to London	0.5	Middle
	Begin install at arebyte	7	Middle/End

CV

Martin Disley is an artist, researcher and creative technologist based in Edinburgh, Scotland. His visual practice centres around an ongoing critical investigation into machine learning. His work has focussed on the machine learning model and the map-territory relation, feedback loops in inference, behavioural conditioning and training and machine learning in states of incoherence. His work seeks to manifest the internal contradictions and logical limitations of artificial intelligence in beguiling images, video and sound.

Martin was recently artist-in-residence at the National Library of Scotland and has received commissions from The Institute for Design Informatics at the University of Edinburgh, The Indeterminacy Research Group at the University of Dundee and Extinction Rebellion among others. His work has been exhibited at the V&A Museum (Dundee, Scotland), Summerhall (Edinburgh, Scotland), The Centre for Contemporary Arts (Glasgow, Scotland), Guterhallen Gallery (Soligen, Germany), Sala Aranyo (Barcelona, Spain) and Kunstencentrum Vooruit (Ghent, Belgium).

The research that informs his work has also contributed to academic publications including the forthcoming Resonance: Axiologies of Distributed Perception (Routledge 2021), edited by Natasha Lushetich and Iain Campbell.

Education

BMus, Music Technology, University of Edinburgh (2017)

MSc, Sound and Music Computing, Universitat Pompeu Fabra (2018)

Selected Exhibitions/Performances

Zoom Obscura
Inspace Gallery
Edinburgh (May 2021)

Extinction Rebellion: End Oil Now
Summerhall,
Edinburgh (Aug 2019)

Given to Chance (NEoN Festival)
Nomas Projects*,
Dundee (Nov 2020)

Tay Late/Arcadia Festival
V&A Museum,
Dundee (May 2019)

DIE HARD: REST NOW!
Kunstencentrum Vooruit,
Ghent, Belgium (Oct 2019)

Sound Thought Festival
CCA,
Glasgow (Feb 2019)

Shoormal Conference (Solo Live)
Mareel,
Lerwick, Shetland (Sept 2019)

Mit dem Ohr der Erinnerung
Guterhallen Gallery, Soligen, Germany
(Jul 2018)

martin_disley@icloud.com
[@martin_disley](https://www.instagram.com/martin_disley)
www.martindisley.co.uk