


MEAN IMAGES

A WHILE AGO, science-fiction writer Ted Chiang described ChatGPT's text output as a 'blurry JPEG of all the text in the web'—or: as a semantic 'poor image'.¹ But the blurry output generated by machine-learning networks has an additional historical dimension: statistics. Visuals created by ML tools are statistical renderings, rather than images of actually existing objects. They shift the focus from photographic indexicality to stochastic discrimination. They no longer refer to facticity, let alone truth, but to probability. The shock of sudden photographic illumination is replaced by the drag of Bell curves, loss functions and long tails, spun up by a relentless bureaucracy.

These renderings represent averaged versions of mass online booty, hijacked by dragnets, in the style of Francis Galton's blurred eugenicist composites, 8k, Unreal engine. As data visualizations, they do not require any indexical reference to their object. They are not dependent on the actual impact of photons on a sensor, or on emulsion. They converge around the average, the median; hallucinated mediocrity. They represent the norm by signalling the mean. They replace likenesses with likelinesses. They may be 'poor images' in terms of resolution, but in style and substance they are: mean images.

An example of how a set of more traditional photographs is converted into a statistical render: the search engine, 'Have I been trained?'—a very helpful tool developed by the artists Mat Dryhurst and Holly Herndon—allows the user to browse the massive LAION-5B dataset used to train Stable Diffusion, one of the most popular deep-learning text-to-image generators. These pictures of mine (Figure 1) show up inside this training data. What does Stable Diffusion make of them? Ask the model to render 'an image of hito steyerl', and this (Figure 2) is the result.

FIGURE 1: 'Images of Hito Steyerl' Found on LAION-5B



Source: havebeentrained.com

FIGURE 2: Stable Diffusion's 'Image of Hito Steyerl'



Source: stablediffusionweb.com

So, how did Stable Diffusion get from A to B? It is not the most flattering ‘before and after’ juxtaposition, for sure; I would not recommend the treatment. It looks rather mean, or even demeaning; but this is precisely the point. The question is, what mean? Whose mean? Which one? Stable Diffusion renders this portrait of me in a state of frozen age range, produced by internal, unknown processes, spuriously related to the training data. It is not a ‘black box’ algorithm that is to blame, as Stable Diffusion’s actual code is known. Instead, we might call it a white box algorithm, or a social filter. This is an approximation of how society, through a filter of average internet garbage, sees me. All it takes is to remove the noise of reality from my photos and extract the social signal instead; the result is a ‘mean image’, a rendition of correlated averages—or: different shades of mean.

The English word ‘mean’ has several meanings, all of which apply here. ‘Mean’ may refer to minor or shabby origins, to the norm, to the stingy or to nastiness. It is connected to meaning as signifying, to ideas of the common, but also to financial or instrumental means. The term itself is a composite, which blurs and superimposes seemingly incompatible layers of signification. It bakes moral, statistical, financial and aesthetic values as well as common and lower-class positions into one dimly compressed setting. Mean images are far from random hallucinations. They are predictable products of data populism. They pick up on latent social patterns that encode conflicting significations as vector coordinates. They visualize real existing social attitudes that align the common with lower-class status, mediocrity and nasty behaviour. They are after-images, burnt into screens and retinas long after their source has been erased. They perform a psychoanalysis without either psyche or analysis for an age of automation in which production is augmented by wholesale fabrication. Mean images are social dreams without sleep, processing society’s irrational functions to their logical conclusions. They are documentary expressions of society’s views of itself, seized through the chaotic capture and large-scale kidnapping of data. And they rely on vast infrastructures of polluting hardware and menial and disenfranchised labour, exploiting political conflict as a resource.

The Janus problem

When a text-to-3D tool called Dreamfusion was trialed in autumn 2022, users began to spot an interesting flaw. The AI-generated 3D models

¹ Ted Chiang, ‘ChatGPT is a Blurry JPEG of the Web’, *New Yorker*, 9 February 2023.

often had multiple faces, pointing in different directions (Figure 3). This glitch was dubbed the Janus problem.² What was its cause? One possible answer is that there is an over-emphasis on faces in machine-learning image recognition and analysis; training data has relatively more faces than other body parts. The two faces of Janus, Roman god of beginnings and endings, face towards the past and towards the future; he is also the god of war and peace, of transition from one social state to another.

The machine-learning Janus problem touches on a crucial issue—the relation between the individual and the multitude. How to portray the crowd as one? Or conversely the one as crowd, as collective, group, class

FIGURE 3: *3D ML-Generated Model of a Squirrel with Three Faces*



Source: <https://twitter.com/poolio/status/1578045212236034048>

² The Janus problem was initially identified by Ben Poole, a research scientist at the experimental Google Brain lab, and posted on his Twitter feed.

or Leviathan? What is the relation between the individual and the group, between private and common interests (and property), especially in an era in which statistical renders are averaged group compositions?

Likelihoods

Here (Figure 4) is another statistical composite in which my face is entangled.

FIGURE 4: *Examples and Composites from Racial Faces in the Wild Database*



The ghostly gendered ‘racial’ blurs on the right could be called a vertical group photo, in which people are not located side by side but one on top of another. How did they come into being?

In 2016 my name popped up in a compendium called MS-Celeb-1M, a Microsoft database composed of 10 million images of 100,000 people who happened to be on the internet. The information was released by the artists Adam Harvey and Jules Laplace as part of their database research project, Megapixels.³ If your name appeared on this list,

³ See the Exposing.ai website. A small sample of the names on the list included Ai Weiwei, Aram Bartholl, Astra Taylor, Bruce Schneier, Cory Doctorow, danah boyd,

Microsoft would encourage researchers to download photos of your face from the internet to build a biometric profile. I had become part of an early training dataset for facial-recognition algorithms. But what was it used for and by whom?

As it turned out, MS-Celeb-1M was of interest to several groups and institutions. It was, for example, used to optimize racial classification by the developers of another dataset, Racial Faces in the Wild. They lamented the fact that facial-recognition technology underperformed on non-white people. So they set out to 'fix' this problem. They uploaded images from the MS-Celeb-1M dataset to the recognition interface Face++ and used the inferred racial scores to segregate people into four groups: Caucasian, Asian, Indian and African. The explicit reason was to reduce bias in facial-recognition software and diversify training data.⁴ The results were ghostlike apparitions of racialized phenotypes, or a quasi-platonic idea of discrimination as such.

If these ghostly renderings evoke the famous photographic composites created by Francis Galton in the 1880s, it is no coincidence. A pioneer social scientist, statistician and eugenicist, Galton developed a method of photographic superposition in order to create portraits of so-called types, such as 'Jewish', 'tubercular' and 'criminal'.⁵ Eugenicists were believers in 'racial improvement' and 'planned breeding', some advocating methods such as sterilization, segregation and even outright extermination to rid society of types they considered 'unfit'. The ghosts often turned out to be mugshots of categories that were supposed to disappear.

Much has been written on the eugenicist backgrounds of the pioneers of statistics, including for example Ronald Fisher. But statistics as a science

Edward Felten, Evgeny Morozov, Glenn Greenwald, Hito Steyerl, James Risen, Jeremy Scahill, Jill Magid, Jillian York, Jonathan Zittrain, Julie Brill, Kim Zetter, Laura Poitras, Luke DuBois, Michael Anti, Manal al-Sharif, Shoshanna Zuboff and Trevor Paglen. As Harvey and Laplace write, Microsoft's definition of 'celebrity' extended to journalists, activists and artists, many of whom were 'vocal critics of the very technology Microsoft is using their name and biometric information to build'.

⁴ Mei Wang, Weihong Deng, Jianli Hu, Xunqiang Tao, Yaohai Huang, 'Racial Faces in the Wild: Reducing Racial Bias by Information Maximization Adaptation Network', Computer Vision Foundation research paper, 2019.

⁵ See also Allan Sekula's seminal text on Galton's composites, 'The Body and the Archive', *October*, vol. 39, Winter 1986, p. 19.

has evolved since their day.⁶ As Justin Joque explains, statistical methods were fine-tuned over the course of the twentieth century to include market-based mechanisms and parameters, such as contracts, costs and affordances, and to register the economic risks of false-positive or false-negative results. The upshot was to integrate the mathematics of a well-calibrated casino into statistical science.⁷ Using data, Bayesian methods could reverse Fisher's procedure of proving or disproving a so-called null hypothesis. The new approach worked the other way around: start from the data and compute the probability of a hypothesis. A given answer can be reverse engineered to match the most likely corresponding question. Over time, methods for the calculation of probability have been optimized for profitability, adding market mechanisms to selectionist ones.

Statistical renderings add a quasi-magical visual effect to this procedure. As categories seem to emerge from the data themselves, they acquire the authority of an immediate manifestation or apparition. Data are no longer presented via the traditional media of graphs, clusters, curves, diagrams or other scientific abstractions. Instead, they are visualized in the shape of the thing from which they are supposed to abstract.

⁶ Wendy Hui Kyun Chun, *Discriminating Data: Correlation, Neighbourhoods, and the New Politics of Recognition*, Boston MA 2021, p. 59; Lila Lee-Morrison, 'Francis Galton and the Composite Portrait', in *Portraits of Automated Facial Recognition*, Bielefeld 2019, pp. 85–100. Ronald Fisher, a founding figure in statistics, argued in 'The Genetical Theory of Natural Selection' (1930) that civilizations were at risk because people of 'low genetic value' were more fertile than people with so-called 'high genetic value' and recommended that caps be put on the fertility of lower classes. Fisher created the well-known Iris dataset, publishing the results in the *Annals of Eugenics* in 1936, to prove that one can discriminate between different species of irises using superficial measurements. Fisher's flowers were proxies for a more sinister idea: if one could discriminate different floral species from superficial measurements, so too could one prove the existence of different races by measuring skulls. The Iris data set is still taught to computer science students as a kind of foundational example.

⁷ The Neyman-Pearson method, for example, encourages the construction of two hypotheses between which a statistical test can select, depending on the circumstances of the experiment. Like a profitable casino, 'if the costs and risk are appropriately calculated, the house is bound to lose some hands, but over time it will ultimately win more'. Justin Joque, *Revolutionary Mathematics: Artificial Intelligence, Statistics and the Logic of Capitalism*, London and New York 2022, pp. 124–5.

They skip mediation to gesture towards fake immanence. Processes of abstraction and alienation are replaced with confusing back-propagation processes—or, more simply: social filters. If Joque reaches for Sohn-Rethel's concept of real abstraction to describe statistical modes of representation, 'mean images' could be described as 'authenticist abstractions'.⁸ This paradox reflects a foundational incoherence at the heart of this mode of rendition. Even though those renders are based on correlated means, they converge towards extreme and unattainable outliers—for example, anorexic body ideals. An unrealistic and probably unsurvivable result is prescribed as a norm; a mode of human social conditioning that long predates machine learning.

In 'mean images', statistics are integrated directly into an object's likeness, through determinations of likelinesses. If Galton performed this trick for faces, statistical renderings further expand his method into the realm of actions, relations and objects; the world at large. The seemingly spontaneous apparition of these distributions obscures the operations within the 'hidden layers' of neural networks that bend existing social relations to converge towards a highly ideological 'optimum' through all sorts of market-related weights and parameters. Markets were already seen as hyper-efficient computers by von Mises and Hayek. Indeed, within liberal economic mythologies, markets play the part of artificial general intelligences, or AGIs—superior, supposedly omniscient structures that should not be disturbed or regulated. Neural networks are thus seen to mimic a market logic, in which reality is permanently at auction.

This integration of statistics is clear in the case of Dreamfusion's 3D models. The most common statistical analogy is the famous coin that, if flipped, has a 50 per cent probability of landing on either heads or tails, if it is fair and even. But in the Janus problem, the probability of encountering a head rather than a tail is far higher than 50 per cent. In fact, there may be no tails at all. Developers note the basic problem of deriving 3D renderings from 2D images. In addition, as noted above, the data may be skewed, the algorithm could be flawed, or missing something—or the experiment as such, and the tools used for it, neither fair nor even. Whichever it is, Dreamfusion has produced its own fork of probability theory: instead of heads or tails, the likelihood is heads and heads.

⁸ Joque, *Revolutionary Mathematics*, p. 179.

How does all this apply to the multiheaded Racial Faces in the Wild composites with which I became entangled? In the liberal logic of digital extraction, exploitation and inequality are not questioned; they are, at most, diversified. In this vein, RFW's authors tried to reduce racial bias within facial recognition software. The results were easily repackaged to identify minorities more accurately by machine-vision algorithms. Police departments have been waiting and hoping for facial recognition to be optimized for non-Caucasian faces. This is exactly what seems to have happened with the research generated from MS-Celeb-1M.

A company called SenseTime picked up on it too.⁹ SenseTime is an artificial intelligence firm that, until April 2019, provided surveillance software to Chinese authorities that was used to monitor and track Uighurs; it had been flagged numerous times as having potential links to human-rights violations.¹⁰ It seems the combination of my name and face was not only used to optimize machine vision for racial classification, but that this optimization was swiftly put into practice to identify and track members of an ethnic minority in China. The fact of my existence on the internet was enough to turn my face into a tool of literal discrimination wielded by an actually existing digital autocracy. By now, the majority of the faces that have appeared on the internet have probably been included in such operations.

The means of mean production

There is another, more pertinent reason why the supposed elimination of bias within datasets creates more problems than it solves. The process limits changes to parts of the output, making these more palatable for Western liberal consumers, while leaving the structure of the industry and its modes of production intact. But the problem is not only with the (social) mean but the overall means of production. Who owns them? Who are the producers? Where does production take place and how does it work?

Creating filters to get rid of harmful and biased network outputs is a task increasingly outsourced to underprivileged actors, so-called microworkers, or ghostworkers. Microworkers identify and tag violent, biased and illegal material within datasets. They perform this duty in the form of

⁹ See: Megapixels, as above.

¹⁰ Anna Swanson and Paul Mozur, 'US Blacklists 28 Chinese Entities over Abuses in Xinjiang', *New York Times*, 17 October 2019.

underpaid ‘microtasks’ that turn digital pipelines into conveyor belts. As *Time* magazine reported in January 2023, underpaid workers in Kenya were asked to feed a network ‘with labeled examples of violence, hate speech and sexual abuse.’¹¹ This detector is now used within OpenAI’s ChatGPT systems. In Western metropolises, microworkers are often recruited from constituencies that are barred from the official labour market by refugee or migrant legislation, as this anonymized interview with a digital worker in a large German city describes:

Digital worker: We were all kind of in the same situation, in a very vulnerable situation. We were new to the city and to the country, trying to integrate, and we desperately needed a job. All the staff on my floor had at least a Master’s degree, I was not the only one. One of my colleagues was a biologist who specialized in butterfly research and had to work on the exact same tasks as I did. Because it was too difficult to find a real job, with a connection to her specialism, people just took this kind of part-time job. They were highly qualified people, with different language backgrounds.

Interviewer: They were all foreigners?

Digital worker: All of them.

Interviewer: What was the work like?

Digital worker: Terrible. And it was the same for everyone I met there. During training you are told that you’re going to see paedophilia, graphic content, sexually explicit language. And then when you actually start working, you sit at your desk and you see things which are unbelievable. Is this really true? The long-term effects of this work are pretty nasty. There was no one in my group who didn’t have problems afterwards. Problems like sleep disorders, loss of appetite, phobias, social phobia. Some even had to go to therapy. The first month you go through a very, very intense training. We had to learn how to recognize content that was too drastic. Because the AI or machine learning mechanisms were not able to decide the delicate cases. The machine has no feeling, it was not accurate enough.

I became depressed. I had to go to therapy. I was prescribed medication. When I started there, my main job was to sieve through posts with sexually explicit content and so-called high priority cases, which usually had to do with suicide or self-harm. There were a lot of pictures of cutting. I had to analyse which were self-harm and which suicidal. The second month, I asked my team leader to put me in a different content workflow because I was feeling bad.

¹¹ Billy Perrigo, ‘OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic’, *Time*, 18 January 2023.

There were a lot of rules for the desk: no phones, no watches, nothing to take pictures with. No paper, no pens, nothing to take notes with. A few times, we saw drones flying outside our windows. Supposedly spies were trying to film what was going on in the company. Everyone was instructed to lower the curtains when that happened. Once, a journalist was outside the building. We were told not to leave the building and not to talk to this journalist. For the company, the journalists were like enemies.

Interviewer: Which form of AI was involved in your work?

Digital worker: I don't know much about the AI that was used where I worked.

I think they tried to hide it somehow. What we did know was that there was some sort of machine learning going on. Because basically they wanted to replace the humans with the AI software. I remember at one point they tried to use this software, but it was very inaccurate. So they stopped using it.

Interviewer: What kind of software was that, what exactly was its task?

Digital worker: I have no idea. This kind of knowledge was only passed on at the executive level of the company. They kept this information secret. And even though we heard rumours here and there, they hid this project from us. But the AI is a failure, because the algorithm is not accurate. People talk about artificial intelligence, but the technology behind it is still very, very mainstream, I would say. That's why they need us humans. Machines don't have feelings. They can't touch. The main goal was to make people like me work like robots.¹²

Another interview conducted as part of the same project described how Syrian digital workers in Germany had to filter and review images of their own home towns, destroyed by the recent earthquake in the region—and, in some cases, the ruins of their former homes.¹³ They were deemed too violent for social media consumers, but not for the region's inhabitants, who had been expelled by war and destruction and were forced to become ghostworkers in exile. Conveniently, military violence had provided digital corporations located in Germany with a new, supremely exploitable refugee workforce.

¹² *Die Zeit*, Feuilleton Ausgabe no. 14, 30 March 2023—a joint cultural supplement project on artificial intelligence by Marie Serah Ebcinoglu, Eike Kühn, Jan Lichte, Peter Neumann, Hanno Rauterberg, Malin Schulz, Hito Steyerl and Tobias Timm. This interview was conducted by Tobias Timm.

¹³ This interview was conducted by Marie Serah Ebcinoglu.

Tweaking technology to be more ‘inclusive’ can thus lead to improved minority identification while outsourcing traumatic and underpaid labour. It can optimize discrimination, superficially sanitizing commercial applications while creating blatantly exploitative class hierarchies in the process. Political and military conflict as well as racially motivated migration barriers are important tools in creating this disenfranchised labour force. Perhaps bias is not a bug, but an important feature of a mean system of production. Bias is not only productive on the level of representation, by visually demeaning people. Its supposed elimination is equally productive in helping to consolidate class hierarchies buttressed by wars, energy conflicts and racist border systems, and can be taken advantage of within a mean system of production.

The elimination of bias is not the only task of microworkers. They also tag street photographs for self-driving cars and categorize images of objects and people to help machine-learning networks distinguish between them. As many writers have noted, human ghostworkers are the engine of automation—self-driving cars could not function without them. Automation runs on the averaged microjudgments of whole detachments of underpaid humans, not some supersmart computer. In some cases, this leads to people impersonating AIs, even where no ML application exists. A researcher reports:

We interviewed K., a Parisian entrepreneur and start-up founder who blamed his competitors for their claim to do AI while instead, they outsource all work to humans recruited through platforms overseas. He went as far as to claim that ‘Madagascar is the leader in French artificial intelligence’. Even more upset was S., a student who did an internship in an AI start-up that offered personalized luxury travel recommendations to the better-off. His company’s communication strategy emphasized automation, with a recommender system allegedly based on users’ preferences extracted from social media. But behind the scenes, it outsourced all its processes to micro-providers in Madagascar. It did no machine learning, and the intern could not gain the high-tech skills he dreamt of.¹⁴

The hidden layers of neural networks also hide the reality of human labour, as well as the absurdity of the tasks performed. The seemingly

¹⁴ Paolo Tubaro, Antonio A. Casilli, and Marion Coville, ‘The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence’, *Big Data & Society*, vol. 7, no. 1, June 2020, p. 7.

unmediated magic of images spontaneously emanating from a heap of data rests on massive exploitation and expropriation at the level of production. Perhaps the seemingly ghostly emanations of faces in statistical renderings are in fact portraits of the hidden microlabourers, haunting and pervading mean images.

Hidden labour is also crucial for the datasets used to train prompt generators. The 5.8 billion images and captions scraped from the internet and collected on LAION-5B, the open-source dataset on which Stable Diffusion was trained, are all products of unpaid human labour, ‘from people coding and designing websites to users uploading and posting images on them.’¹⁵ It goes without saying that none of these people were offered remuneration, or a stake in the data pool or the products and models built from it. Private property rights, within digital capitalism and beyond, are relevant only when it comes to rich proprietors. Anyone else can be routinely stolen from.

From the mean to the common?

It now becomes clearer why the Janus-headed 3D models show far more heads than tails: the ‘coin’ is rigged. Whether it lands on head or ‘head’, on automation or—in Astra Taylor’s term—fauxtimation, the house always wins. But the question of labour conditions may allow for a more general observation about the relation of statistics and reality, or the question of correlation and causation. Many writers, including myself, have interpreted the shift from a science based on causality towards assumptions based on correlation as an example of magical thinking, or a slide into alchemy. But what if this slide also captures an important aspect of reality? A reality which, rather than being governed by logic or causality, is in fact becoming structured much more like a casino?

A blog post about depression and video games provides an illuminating example. The writer describes playing video games in moments of depression, especially enjoying small and repetitive tasks that lead to some form of constructive result—a crop being planted, a house being built:

What makes video game labour so exciting is the chance to fully enjoy a reward for your efforts. You put in what you get out of it, in a very direct

¹⁵ Chloe Xiang, ‘AI Isn’t Artificial or Intelligent’, *Vice Motherboard*, 6 Dec 2022.

sense. The most satisfying games are actually fantasy simulations of living out a basic Marxist value: labour is entitled to all it creates.¹⁶

The writer describes a causal relation between input and output, labour and reward. The striking conclusion is that such causality within actually existing capitalism is rare, especially in precarious work. Whatever effort you put in is not going to create an adequate output, in a linear fashion—a living wage, or an adequate form of compensation. Precarious and highly speculative forms of labour do not yield linear returns; cause and effect are disconnected. This also introduces a class aspect into the real-life distribution of causality versus correlation. Jobs paid per hour retain a higher degree of cause-and-effect than those that happen in the deregulated realm of chance, at both ends of the wage scale. It is thus more ‘rational’ for many people to treat daily existence as a casino and to hope for speculative windfalls. If all you can hope for within a causal paradigm is zero income, then buying a lottery ticket becomes an eminently rational choice. Work starts to resemble gambling. Phil Jones describes microwork along similar lines:

The worker thus operates increasingly in a quasi-magical economy of gambling and lottery. Microwork represents the grim apex of this trajectory, where the possibility of the next task being paid tempts workers time and again to return for more. Intricate reward schedules and contestable pricing gamify tasks and effectively repackage superfluity and precarity as new, exciting forms of work-cum-leisure.¹⁷

When the ‘wage transforms into a wager’, probability is not just an assessment of a real outcome but becomes a part of that outcome itself.¹⁸ Statistical renderings account for this. Once social causality is partially replaced by correlation, labour relations are pushed back into the times of Victorian sweatshops, images converge towards bets and the factory becomes a gaming house. Indexical photography was, at least partly, based on a cause-effect relation. But in statistical renders, causality is adrift within a mess of quasi-nonlinear processes, which are not contingent but opaquely tampered with.

¹⁶ Tabitha Arnold, ‘Depression and Videogames’, Labour Intensive Art Substack, 6 January 2023.

¹⁷ Phil Jones, *Work Without the Worker: Labour in the Age of Platform Capitalism*, London and New York 2021, p. 50ff.

¹⁸ *Ibid.*, p. 50.

Precarious labour in machine-learning industries, and the repetitive processes of conditioning and training it requires, raises the question: who or what is being trained? Clearly, it is not only the machines, or more precisely the neural networks. People are being trained, too—as microworkers, but also as general users. To come back to Herndon and Dryhurst’s excellent question: ‘Have I been trained?’ The answer is: yes. Not just my images but I, myself. Prompt-based image generators like DALL-E do of course rely on training machine-learning models. But much more importantly, they train users how to use them and, in doing so, integrate them into new production pipelines, soft- and hardware stacks that align with proprietary machine-learning applications.

They normalize a siloed production environment in which users constantly have to pay rent to some cloud system, not only to be able to perform but even to access the tools and results of their own labour. One example is Azure, according to Microsoft the ‘only global public cloud that offers AI supercomputers with massive scale-up and scale-out capabilities.’ Azure rents out machine-learning computational applications and computing power, while Microsoft has established a top-down proprietary structure which includes ML-ready hard- and software, browsers, access to models, application interfaces (APIs) and so forth. Adobe, the world’s most hated, convoluted and extractivist quasi-monopoly for image workers, is moving swiftly in the same direction. Dwayne Monroe calls this kind of quasi-monopoly a ‘super rentier structure’, in which digital corporations privatize users’ data and sell the products back to them: ‘The tech industry has hijacked a variety of commons and then rents us access to what should be open.’¹⁹

Many digital and administrative white-collar labourers are threatened by ML-based automation; among them, in no particular order, programmers, PR professionals, web designers and bookkeepers. But it is more likely that many of them will be forced to ‘upgrade’ by renting services built on their own stolen labour, in order to remain ‘competitive’, than that they will be replaced fully by ML automation. Today, training for these professionals means getting them used to their disappearance over the medium term and conditioning them to dependency on monopoly stacks, in order to be able to continue to work and access the results of their own labour.

¹⁹ Dwayne Monroe, ‘ChatGPT: Super Rentier’, *Computational Impacts*, 20 January 2023.

This draws attention to yet another aspect of ‘mean images’. Like NFTs, statistical renderings are onboarding tools into specific technological environments. In the case of NFTs, this is a crypto-environment run through tools such as wallets, exchanges or ledgers. In the case of machine learning, the infrastructure consists of massive, energy-hungry, top-down cloud architectures, based on cheap click labour performed by people in conflict regions, or refugees and migrants in metropolitan centres. Users are being integrated into a gigantic system of extraction and exploitation, which creates a massive carbon footprint.

To take the Janus problem seriously therefore would mean to untrain oneself from a system of multiple extortion and extraction. A first step would be to activate the other aspect of the Janus head, the one which looks forward to transitions, endings as beginnings, rather than back onto a past made of stolen data. Why not shift the perspective to another future—a period of resilient small tech using minimum viable configurations, powered by renewable energy, which does not require theft, exploitation and monopoly regimes over digital means of production? This would mean untraining our selves from an idea of the future dominated by some kind of digital-oligarch pyramid scheme, run on the labour of hidden microworkers, in which causal effect is replaced by rigged correlations. If one Janus head looks out towards the mean, the other cranes towards the commons.