iasdr 2021 [WITH DESIGN]

PolyU Design
School of Design
The Hong Kong Polytechnic University
5–9 December 2021

A Design-Led Exploration of Material Interactions between Machine Learning and Digital Portraiture

Green, David Philip*[a]; Lindley, Joseph[a]; Mason, Zach[a]; Coulton, Paul[a]

[a] Imagination Lancaster, Lancaster University, United Kingdom, Lancaster, UK
* d.p.green@lancaster.ac.uk

Design materials are defined as a combination of what they *are*, what they *do,* and the ways they interact with other materials. In this pictorial, we explore the interactions between machine learning (as a design material) and another design material—the human face—in the form of digital portraiture. Employing an exploratory *Research through Design* approach we consider how machine learning simultaneously enriches and subverts the materiality of the human face. Through a combination of images and text, we offer some considerations and provocations for further research.

Keywords: design materials; materiality; machine learning; digital portraiture; faces; art

1 Introduction

1.1 What is a design material?

Design materials are manifold and multiform; they include metals, woods and plastics, as well as *digital* materials like metadata (Odom et al, 2018) and sensors (Zhuying et al, 2020). All design materials have *properties* (toughness or malleability, for instance) as well as *uses* (red ink to flag mistakes, squid ink to colour pasta); this forms a combination of what it *is*, and what it *does* (Ingold, 2013). Moreover, design materials can be characterised by *interactions* with other materials; cement and sand interact to make concrete, for instance. However, “materiality is not always synonymous with physicality” (Patterson and Garner, 2007:274). In this respect, interactions are particularly important in characterising *digital* design materials since they are often invisible but for their interactions with other (non-digital) materials (Fuchsberger et al, 2013). For example, Hypertext Mark-up Language (HTML) interacts with browser software and hardware (e.g., screens) to generate websites. In this pictorial, we reflect on machine learning as a digital design material and consider it primarily through its interactions with another design material—the human face.

1.2 Machine Learning as a Design Material

Machine Learning (ML) has been around since the 1950s. However, there was a surge of interest in 2016 when Deepmind's *AlphaGo* system defeated Lee Sedol, a master of the abstract strategy game *Go*. Sedol later remarked, "*Surely, AlphaGo is creative*". Since then, ML has been tentatively considered a 'design material' (Leahu, 2016; Dove et al, 2017). Recent work has clarified its value for design as pertaining to its hallmark property of '*uncertainty*' and proposed reflecting upon uncertainty–as a material quality of ML–to be a potentially fruitful direction for future design research (Benjamin et al, 2021). This work is a response to this invitation. Our starting point is the observation that we know relatively little about how ML interacts with other design materials, or even how to characterise those interactions. Yet, given that we know the potential impact of "algorithmic bias" (Garcia, 2016) via problematic feedback loops (Noble, 2018) enabled by what Benjamin et al call "*pattern leakage*" (2021), understanding these interactions is a crucial design consideration. Can ML disrupt other design materials in ways we can't completely understand? How does ML influence material expertise? What are the implications of introducing uncertainty into design materials, through their interactions with ML?

1.3 The Human Face as a Design Material

Let's consider a design material that interacts with ML in an interesting way. The human face is one of the most recognisable forms in nature. Pareidolia (seeing human faces in non-human forms, e.g., door handles) is a reflection of our highly-tuned ability to perceive and *read* faces and this is a foundation of both our sociality and—in the Lacanian mirror—our sense of self. But is it a design material? When Picasso asked "*who sees the human face correctly?"* he alludes to its material properties. The face is a reflection of the self, but it is also changeable (e.g., as we age) and what it *does* depends upon how it is *used.* It can be *used* in many different ways, especially in digital forms.

Fig 1 – *Walden Kirsch (1957).*

The first digital image (176px x 176px).

This image was created by Russell A. Kirsch of the US National Bureau of Standards, who developed the first image scanner. Kirsch produced this image of his 3-month-old son from a scanned analogue photograph. The shades of grey were achieved by compositing several black and white images. It is remarkable that the first digital image was a human face. Note how the production process leaves a material trace in the form of artefacts that – particularly when compared to digital portraits today – create an 'impressionistic' – even 'uncanny' – appearance.

Russell A. Kirsch / Public domain.

Representations of human faces have a very long history. In art, from cave paintings to modern deepfakes, faces (in combination with other materials from oil paint to compression algorithms) produce powerful messages that stimulate a wide range of emotional responses. To the makeup artist, the plastic surgeon and the character designer, the face is a three-dimensional design material. For live performers and animators, time adds a fourth dimension. Here, we focus on a simpler form; the two-dimensional representation of the face in a still image – the portrait (*Fig 1*).

Fig 2 *–* ***Left****: A portrait of one of the authors as a boy, digitised from an analogue print, in this case using a smartphone camera.* ***Right****: The same image altered using Adobe Photoshop to simulate prosopagnosia (also known as 'face blindness') a cognitive disorder in which the ability to recognize face—including one's own—is impaired. [https://www.nhs.uk/conditions/face-blindness/]*

We frame our exploration of interactions with machine learning around *digital portraiture* for two main reasons. Firstly, we can already find evidence in society for how problematic this interaction can be. Deepfakes have caused a moral panic on account of their misuse and potential for further abuse. Algorithmic bias in face detection systems has been shown to perpetuate antisocial biases such as racism. This is an important research space that will benefit from varied approaches. Secondly, with the rare exemption of those with *prosopagnosia* (*Fig 2*), faces are universally recognisable. In this respect, we can leverage our innate familiarity with faces to highlight the ways machine learning uncertainty renders this material *unfamiliar*. Through this, we aim to reflect on the materiality of ML through its interactions with digital portraiture.

2 Method

We—the authors—have expertise in art and design, but no significant experience or expertise in ML. For this reason, we have adopted an exploratory *Research through Design* approach (Gaver, 2012), using our own digital (self-) portraiture as a raw material. Through experimentation with an accessible online tool (*RunwayML* - https://app.runwayml.com), we generated a large quantity of visual materials, which forms the basis of this pictorial. Applying an interpretive framework anchored in theories of media, art and design, we offer reflections on the interactions between ML and our digital (self-)portraits. The work is inspired by and informed by artwork by one of the authors (*Fig 3*) that also used *RunwayML* to explore how ML translates digital self-portraits into outputs including 3D printed models and virtual environments.

(a) *(b)* *(c) top | (d) bottom*

Fig 3 *– Machine Learnt Landscapes (Zach Mason, 2020) – This project explores the embodiments we can apply to outputs of ML, as well as the way this affects perceptions of the outputs. The dataset was trained on images of the artist's face taken between 2016-2020. Two different sets were created, one used cropped images of the face (a), the other was fed the entire (uncropped) image to decipher (b). From the two sets, the latter provided more diverse outputs, revealing sections of the face that the Generative Adversarial Network (the type of ML used) recognised as eyes, hair and glasses. The second set was used for two distinct outputs: a 3D printed plastic head (c), based on a specific image (b) and, the final output, which better represented the nature of the training and outputs; a 3D modelled landscape. This was compiled from a thousand randomised outputs from the trained, un-cropped GAN, using colour data to create depth geometry as a displacement map (d). This not only conveyed the scale of the ML process, but had the potential to visualise the unique patterns across the resulting images.*

RunwayML is an online platform that aims to be "*the most accessible machine learning platform*". Via a web interface, users can create and publish images, text and videos with pre-trained, open-source ML models, including *StyleGAN2*, a generative adversarial network (GAN) by Nvidia (2020) that can be used to generate photorealistic portraits. Based on our aim of exploring digital portraiture, we used the StyleGAN2 model to generate new 'portraits' based on models trained with our own datasets (images of ourselves). The number of training steps (from 500-25,000) determined how much the algorithm learns. We then used the updated StyleGAN2 model to generate new portraits. One parameter—a truncation slider—improved *variety* (i.e., the visible differences between multiple outputs) at the expense of *fidelity* (i.e., how much each output looked like a real person) or vice versa.

Fig 4. *Visage. A (fictional) magazine cover, produced as a kind of ‘visual abstract’ for this work. It suggests some of the concepts explored within the pictorial. The headlines read, “uglyism / vectors / VGG-Face vs StyleGan2”, “HAIR! EYES! SKIN! ARTEFACTS!”, “your face or mine?”, “the materiality is the message”, “what have algorithms ever done for portraiture?”, “The new Picasso? Inside the amateur machine learning revolution” and, “From the Uncanny to the Grotesque”.*

3 Findings

The materials we used to 'seed' our experiments were our personal photo archives. *RunwayML* recommends using 500-5000 images to train a model. Unsurprisingly, collating these materials was time-consuming (even aided by automatic face detection and identification algorithms). For some of us, it was also the first time we had (ever) systematically reviewed our personal photo archive. Chronologically browsing the image archive and seeing old friends, younger versions of ourselves, and in some cases seeing images of people who had recently passed away was a surprisingly emotional experience. Subsequently, spending time cropping ourselves out of group photos (to train a model based on images of *just us*) was an exercise in self-isolation, unfortunately reminiscent of pandemic lockdown measures. Our early results did not lighten the mood.

3.1 A Gallery of Grotesquerie

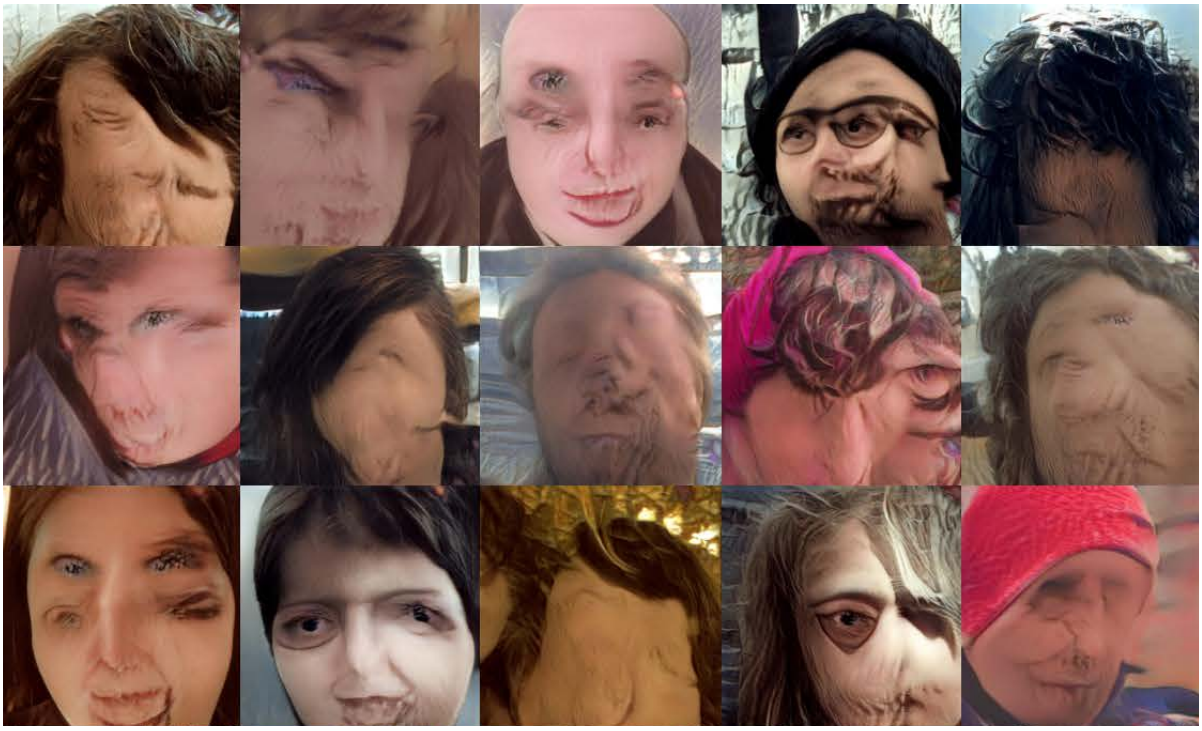

Fig 5 *– 'Grotesque' results from our early efforts to generate new faces based on our minimum-values models (i.e. after <1000). Note the duplicated, distorted and reorganised facial features. These data were collected*

The word 'grotesque' is defined as, "*comically or repulsively ugly or distorted.*" (OED). We have characterised the examples above (*Fig 5*) as *grotesque* on the basis that they are distorted, twisted, ugly and un-human. The relationship between prejudice and ML is well-documented, but subtle. *Cacophobia*—prejudice that discriminates against people based on their perceived attractiveness is a pernicious threat, like any form of prejudice. Since we know that "*pattern leakage*" creates feedback loops between machine learning and society; we may wish to consider what underlying biases underpin our visceral reactions to images like these and why.

One remarkable aspect of these images is their resemblance to self-portraiture by people who have experienced physical and mental illnesses. Self-portraiture by contemporary artists such as Adrienne Gantenberg (Fig 6) exude similar asymmetries, exaggerated, incomplete and augmented features. The celebrated Russian biologist Eugen Gabritschevsky (1893 - 1979) produced a series of drawings with similar visual qualities following his diagnosis with schizophrenia. What accounts for these similarities? Is the uncertainty of machine learning analogous to the mysteries of mental illness? Can one shed light upon the other?

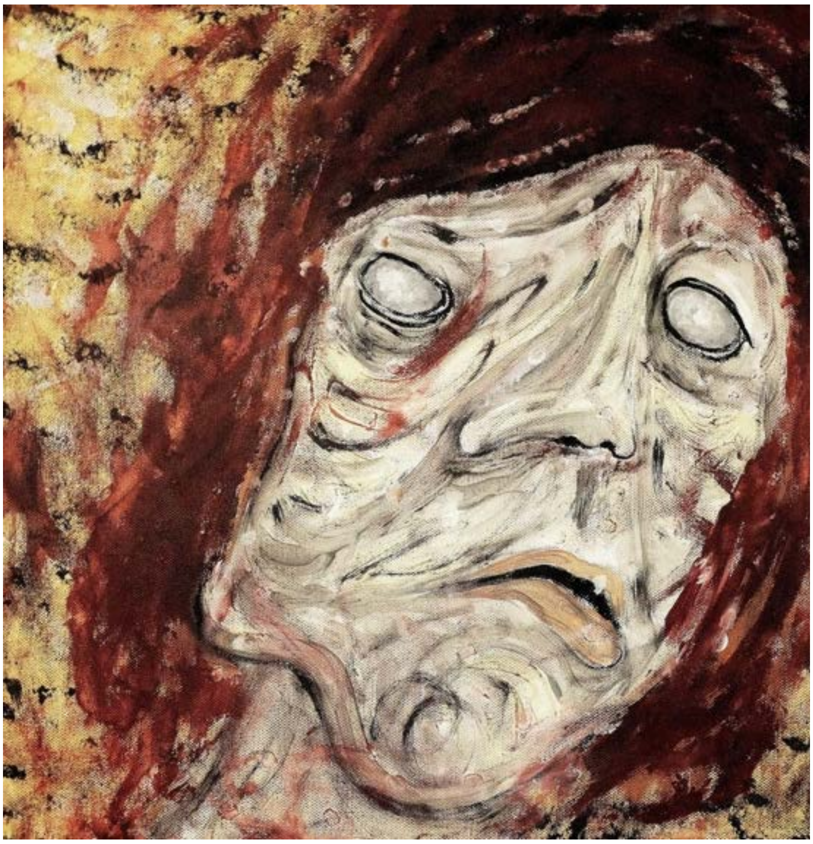

Fig 6 *- Shame (2019)*
by Adrienne Gantenberg
https://www.adriennegantenberg.com/

"This painting is called, "Shame". It was painted in 2018 when I was going through a severe depressive episode. Making art during episodes of mental illness is not always possible. I am grateful and proud that I was still able to produce this gnarly piece! You should buy it!"

Adrienne Gantenberg (Etsy.com, 2021)

Image reproduced with permission from the artist.

3.2 A Beautiful Batch

Some of the images we generated were striking in ways that align with more traditional aesthetic values, albeit not necessarily those associated with portraiture. While some of the more visually striking images include hints of a human presence – even hints of a face – many tended to be impressionistic, surrealist, or abstract in nature; often exuding psychedelic or dreamlike qualities.

Fig 7 *– Some of the images developed during the Machine Learnt Landscapes project.*

Fig 8 *- This model was trained on randomly cropped images, hence, while each image contains a face, the pixels chosen for the model to train on may or may not include that face. The result is images which rarely look human, but, nonetheless, have human elements in them. Elements of hair and shapes which almost resemble a face are common here. In addition the model's outputs evoke memories of familiar scenes, but which are almost impossible to place. The feeling is similar to that of a dream which one has just forgotten.*

3.3 Almost Lifelike

Some outputs appeared at first glance to be photorealistic portraits but, on closer inspection, artefacts, distortions and aberrations lent the images an 'uncanny' appearance. Often, the source of the 'malfunction' was clear—a stray bit of beard here; an unfinished glasses frame there—but not always. This uncanniness is a direct material consequence of the influence of ML uncertainty. The "uncanny valley" is often evoked in relation to generative imagery; could it be leveraged as an educational tool for teaching people about the material consequences of ML uncertainty?

***Fig 9**. 12 portraits (**left**), generated from a training dataset (**right**) of 971 images (2160px x 2160px) - still frames taken at 10 frame intervals from a 4K video. In the self-filmed video, I made various facial expressions while reciting the alphabet, meanwhile sitting quite still so as to achieve consistent framing. A consequence of this approach is that the generated eyes are quite small (a reflection of the training images being framed wider than the base images). Furthermore, in combination with the base images, the eyes resolve to the same part of the image—irrespective of the shape and size of the original face. Flicking through the images results in the uncanny effect of eyes remaining static while other facial features move around them. Scanning the images, I catch glimpses of myself that dissolve on closer inspection; this uncanny quality is one of the more interesting material consequences that differentiate these portraits from authentic human likenesses. An earlier training dataset had more varied compositions and did not exhibit the same effect.*

***Fig 10**. I see this model's outputs as infinite mirrors for lives I could have had, or could still have. After being trained on around 900 images depicting a range of facial expressions, the model now outputs faces reminiscent of my own, but which seem to have lived lives which aren't mine. The majority of the images are clearly me, but occasionally the contortions are so extreme that it seems to be somebody else looking back at me, albeit with a similar smile, shape of nose, or facial hair pattern. The sequences cause me to wonder how recognisable I am to others through these images. Would a familiar member who had been out of touch for some time be able to spot a fake? Consistent training data and a higher number of training steps yields close variations, with only subtle hints of the original images remaining.*

3.4 Spatial Collage and 'Vectors'

RunwayML presents generated images in a grid-like spatial collage evocative of the photographic call sheet (or a very busy Zoom call!) These 'Vectors' facilitate exploration of generative datasets along vertical and horizontal axes (*Fig 11*), dynamically updating as we scroll; invisible parameters updating to create qualitative shifts in the appearances of the faces, as if navigating gene expressions. The generative outcomes share various characteristics, but the vectors demonstrate how subtle shifts compound to create markedly different manifestations—from identical source materials. Why? How? Here, we can clearly see the *influence* of ML, but we cannot observe its workings.

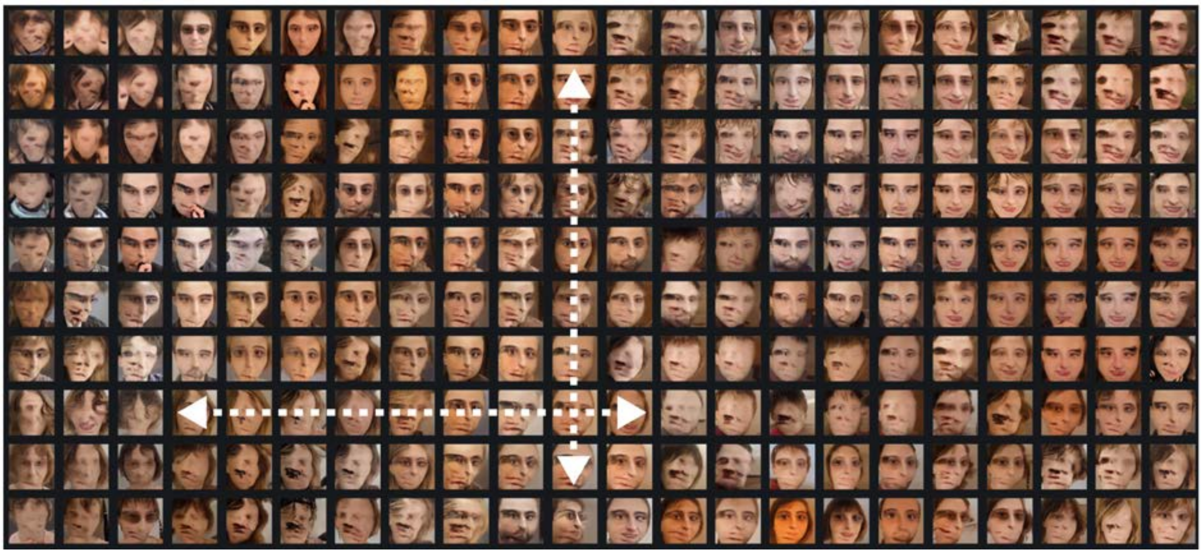

Fig 11 *– Vectors in RunwayML. Using our understanding of the human face, we can explore the ways in which machine learning augments the images, thereby gaining some insights into how the process works.*

3.5 Temporal Collage and 'Interpolations'

Fig 12 *– One way of seeing the impact of the training steps is through sequences of images that represent the interpolations between base images and new generated outputs. In these examples, we see how a training dataset can transform photorealistic human faces (produced by the untrained model) into grotesque portraits that become gradually less human.*

4 Discussion

The hallmark material quality of machine learning—*uncertainty*—led us to attribute the provenance of emergent traits in our digital portraits to its mysterious influence. To paraphrase Jackson Pollock, each image gained *a life of its own*. Yet, whereas Pollock's drip technique resulted in his iconic abstracts, the interaction between ML and the universal familiarity of the human face resulted in portraits that were variously *grotesque*, *beautiful*, *ugly*, *strange*, and *uncanny*. The same can be said of relatively photorealistic ML-generated digital portraiture, by artists such as Morphy_Me (@morphy_me). Our experiments were exploratory, the results were far-from-photorealistic, and *because of this*, they expose more clearly the otherwise subtle influence of ML upon portraiture.

We began with Ingold's distinction between what a design material *is* and what it *does*. Taking this point a step further, Ingold compares the work of the design practitioner to that of an *alchemist*. Unlike the *chemist*, who considers "*matter in terms of its atomic or molecular constitution*" (i.e., what it *is*), the alchemist considers a material in terms of what it *does*. ML-generated portraiture is an act of alchemy. It produces a real effect—it is *affecting*—because ML *did something*, but we don't know *what*. As a material interaction, this is intriguing, even inspiring, but also concerning. AI researcher Ali Rahimi evokes the alchemy metaphor to discuss concerns in the AI community about how ML is being widely applied without a clear of understanding of the fundamentals of how it works:

> "*There's a place for alchemy. Alchemists invented metallurgy, ways to make medication, dy[e]ing techniques for textiles, and our modern glass-making processes. Then again, alchemists also believed they could transmute base metals into gold*" (Rahimi, 2017).

Alchemists are both *innovators* and *charlatans*; do we trust them? ML has divided photorealistic portraiture into two camps. One represents real, individual persons (edited or not); the other represents merely *the archetype* of an individual person. We can no longer trust which camp any given portrait falls into. *Photoshop* and *CGI* undermined the authenticity of the portrait; ML has *obliterated* it. From a design perspective, this creates new possibilities for digital portraiture, but it has also fundamentally subverted and altered its traditional role in society and–by extension–design.

In a sense, this fundamental uncertainty is a continuation of the postmodern tradition in art and design, but a poststructuralist perspective is arguably more revealing of its potential impact. Let us consider this issue through the lens of Barthes' 'studium' and 'punctum' (1981). The studium relates to what we can infer from the formal elements of an image; "this is a photograph of an old man". The punctum relates to how we are *affected*; "this photograph makes me sad because it looks like my late grandfather". Barthes, who famously declared the "*death of the author,*" would be intrigued by the possibility of a studium isolated from intentional (human) authorship. ML materially alters the appearance of a portrait, but it also renders the very notion of a photographic subject uncertain. As a consequence of the mere *possibility* of ML; *only the punctum remains*.

Why does it matter whether a portrait was created using photography or ML? Uncertainty of provenance in digital portraiture is a disruptive shift for designers in various industries, notably the publishing industry. The network of human collaborators (e.g., models, make-up artists, photographers) and the wider supply chain of the entire industry will need to confront and adapt to this new reality. Its societal impact is also an ongoing concern (as deepfake forgeries attest), accelerated by, and accelerat*ing*, 'post-truth' culture. Marshall McLuhan's prescience is increasingly relevant—*mediation* (i.e. how a message is conveyed) matters. The deepfake video, "*In the Event of Moon Disaster*" (Panetta & Burgund, 2019)—a deepfake of a speech given by the US President Richard Nixon—strives not for photorealism but for *authenticity*. It achieves it through attention to material qualities, low resolution, inconsistent analogue signal artefacts, chromatic aberration, and soft focus. This simultaneously disguises the fact that the video is generative and makes it appear authentic to a historical moment (the late 1960s). It claims back what uncertainty has undermined. This works *because* of its materiality. However, as postmodernism continues and media converges (Jenkins, 2004), material anchors in time have been steadily replaced with endless simulacra, parody, and pastiche. We may want to consider how we will identify a digital portrait from the 2020s? This is an important reflexive consideration for design as it moves into the future, but also for futurism *now*. How can designers evoke the early 2020s through portraiture? Using artefact-laden images such as those here? Presently, images as distorted as these are culturally invisible—we are familiar with deepfakes already, but we have not been showing our workings. We *should* be showing images like these to expose how we got to the point we are at now, where photorealistic portraits are *not* photographs. We must invite critical consideration of the potential impacts and 'show the workings' of the new reality.

For digital portraiture, the horse has already bolted; photorealistic generative portraiture is already here. However, digital portraiture was low-hanging fruit for ML—simple, visual, recognisable, emotive, abundant. Interactions between ML and *other* design materials are relatively immature, so there are opportunities for deeper, earlier, more critical reflection and *doing* and *showing* this kind of working out. What happens when ML meets metadata, sensors, big data, HTML, digital audio, IoT, VR, AR? We need to interrogate these questions, not just in abstract terms, but materially through Design Research.

5 Conclusion

Through our experiments with ML and portraiture, we gained very little insight into *how* ML works, but we gained some intriguing insights into the *effects* of ML on digital portraiture. Future research into the interactions between ML and other design materials might draw upon the strength of this Research through Design approach, while bearing in mind the important limitation of ML remaining fundamentally mysterious.

The familiarity and formal simplicity of what the human face *is* masks the fact that what it *does* is complex and deeply intertwined with our complex human psychologies, societies, and cultures. Contrastingly, ML is *unknowable*—yet what it *does* is simple: It injects uncertainty into otherwise knowable materials. Here, it rendered the familiarity of the face *unfamiliar* in ways that were challenging, humorous, disturbing and surprising. It is tempting to look at the images we generated and propose that machine learning rendered the '*sacred geometry*' of the face *profane.* Yet our images are the product of inexperienced users wielding a nascent design material without skill and proficiency; like a scribble drawn in crayon. We are satisfied that these experiments raised as many questions as answers because these are important questions to consider, both from a technological perspective (how can we make ML better?) and from the perspective of design and materiality (how can we use / work with ML better?). However, given the mysterious "creative" power of machine learning, its strange influence over other design materials, and its potential to stir powerful emotions, perhaps we should be asking, *who are the alchemists* and what motivates them?

Acknowledgments. This work was funded by UK Research and Innovation (grant reference MR/T019220/1). Special thanks to Adrienne Gantenberg for permission to use her artwork, Shame (Fig 6).

References

Barthes, R. 1981., 'Camera Lucida: Reflections on Photography'. Macmillan, 1981.

Benjamin, J.J., Berger, A., Merrill, N., Pierce, J., 2021. Machine Learning Uncertainty as a Design Material: A Post-Phenomenological Inquiry, in: *Proceedings of the 2021 CHI Conference on Human Factors in Compu.ng Systems, CHI '21*. Associalon for Compulng Machinery, New York, NY, USA, pp. 1–14. hJps://doi.org/ 10.1145/3411764.3445481

Budelman, M., and Kizelshteyn, M., 2019., 'Three Principles for Designing ML-Powered Products'. Online - https://spotify.design/article/three-principles-for-designing-ml-powered-products. Accessed Thursday 10th June, 2021.

Dove, G. Halskov, K., Forlizzi, J., Zimmerman, J. 2017. 'UX Design Innovation: Challenges for Working with Machine Learning as a Design Material'. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17). Association for Computing Machinery, New York, NY, USA, 278–288. DOI: https://doi.org/10.1145/3025453.3025739

Fuchsberger, V., Murer, M., & Tscheligi, M., 2013. Materials, materiality, and media. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13). Association for Computing Machinery, New York, NY, USA, 2853–2862. DOI: https://doi.org/10.1145/2470654.2481395

Fiore, Q. and McLuhan, M., 1967. *The medium is the massage*. New York: Random House.

Garcia, M., 2016. 'Racist in the machine: The disturbing implications of algorithmic bias'. *World Policy Journal*, *33*(4), pp.111-117.

Gaver, W., 2012. What should we expect from research through design? Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, 937–946.

Gellersen, H., 2018. 'Eye Movement as Design Material', In Chuang, Lewis, Andrew Duchowski, Pernilla Qvarfordt and Daniel Weiskopf, Report from Dagstuhl Seminar 18252, Ubiquitous Gaze Sensing and Interaction. June 18-21 2018.

Ingold, T., 2013. Making: Anthropology, archaeology, art and architecture. Routledge.

Jenkins, H., 2004. The cultural logic of media convergence. *International journal of cultural studies*, *7*(1), pp.33-43.

Leahu, L., 2016. Ontological Surprises: A Relational Perspective on Machine Learning. In Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16). Association for Computing Machinery, New York, NY, USA, 182–186.

Noble, S.U., 2018. Algorithms of oppression: How search engines reinforce racism. nyu Press.

Odom, W., and Duel, T., 2018. 'On the Design of OLO Radio: Investigating Metadata as a Design Material'. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18). Association for Computing Machinery, New York, NY, USA, Paper 104, 1–9. DOI: https://doi.org/10.1145/3173574.3173678

Paterson, G. and Garner, S., 2006. On Making: Issues Of 3d Representations In Design Education. In DS 38: Proceedings of E&DPE 2006, the 8th International Conference on Engineering and Product Design Education, Salzburg, Austria, 07.-08.09. 2006 (pp. 273-278).

Rahimi, A., and Recht, B. 2017, 'Reflections on Random Kitchen Sinks', online, http://www.argmin.net/2017/12/05/kitchen-sinks/ (accessed Thursday 10th June 2021)

Redström, J., 2005. 'On Technology as Material in Design'. Design Philosophy Papers. 3. 39-54. 10.2752/144871305X13966254124275.

Sundström, P., Taylor, A.S., & O'Hara, K., 2011. 'Sketching in software and hardware Bluetooth as a design material'. In Proceedings of the 13th International Conference on Human Computer Interaction with Mobile

Devices and Services (MobileHCI '11). Association for Computing Machinery, New York, NY, USA, 405–414. DOI: https://doi.org/10.1145/2037373.2037434

Sutton, S. J., Foulkes, P., Kirk, D., & Lawson, S., 2019. 'Voice as a Design Material: Sociophonetic Inspired Design Strategies in Human-Computer Interaction'. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Paper 603, 1–14. DOI: https://doi.org/10.1145/3290605.3300833

Li, Z., Wang, Y., Greuter, S., Mueller, F., 2020. 'Ingestible Sensors as Design Material for Bodily Play'. In Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA '20). Association for Computing Machinery, New York, NY, USA, 1–8. DOI: https://doi.org/10.1145/3334480.3382975