OXFORD INTERNET INSTITUTE

UNIVERSITY OF OXFORD

DEPARTMENT OF ENGINEERING SCIENCE

AI and the Arts

How Machine Learning Is Changing Artistic Work

Anne Ploin, Rebecca Eynon, Isis Hjorth, Michael A. Osborne

About the Authors

Anne Ploin is a DPhil candidate at the Oxford Internet Institute, University of Oxford.

Rebecca Eynon is Professor of Education, the Internet and Society at the University of Oxford.

Isis Hjorth is a Research Associate at the Oxford Internet Institute, University of Oxford.

Michael A. Osborne is Professor of Machine Learning at the University of Oxford.

Principal investigators Prof. Michael A. Osborne and Dr. Isis Hjorth designed and wrote the project proposal, secured funding, and managed the project. Co-investigator Prof. Rebecca Eynon oversaw the case study research in this report. Project researcher Anne Ploin developed the arts case study, carried out and analysed the interviews, and wrote the report.

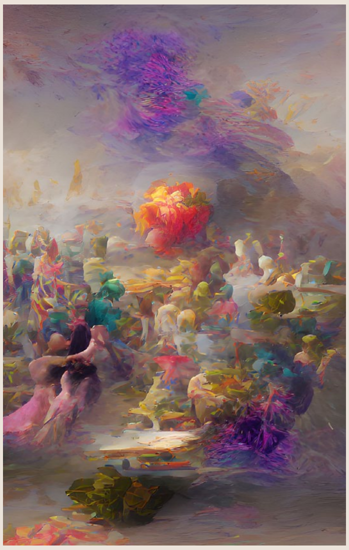 

About the Illustrator

Alexandra Francis is an illustrator and designer based in Manchester, UK. Her work can be found here: https://alexandrafrancis.com/.

About the Cover

Our report's cover came out of a conversation between analogue and AI methods used by our illustrator, Alex––and between Alex and the authors.

The image on the top left was produced by prompting an AI-backed online text-to-image generator with the title of the report ("AI and the Arts: How machine learning is changing artistic work"). The generator used was WOMBO Dream (https://app.wombo.art/).

The image on the top right is the report's cover, produced by Alex using the generated image as an artistic prompt.

Since this process expanded on the conversations about AI and art contained in the report itself, we include a short reflection with Alex here.

Q: What was it like doing the cover this way?

It was great. One of the hardest parts for me is getting the inspiration for a piece, so having the generated image was a great resource. I had this blurred image I could focus in on, where I could see different things I could take forward in my own piece. In the generated image we chose, I could see a mystical landscape, and immediately something clicked. It was really helpful for getting through that first stage of deciding what to do with the composition and the colours. It's a really good starting point.

Q: So it acted as a prompt more than anything else.

Yes. Sometimes, starting a piece of work, you've got an idea in your head, but later realise that you're rehashing something you've seen. You think, oh no, I'm ripping someone off, I've seen too many things on the internet and I'm taking somebody else's idea. But a generated image is completely unique. It's a nice way to make something original––especially in my field, where people draw inspiration from a lot of the same things.

Q: Would you use this process in your own practice in the future?

After doing this cover, I'm definitely keen to use more AI tools. I already keep a library of my own images, to try and avoid looking at what everybody else looks at, and to get inspiration for the initial stages. I would use it in the same way––if I needed a composition idea, as a prompt. To spark something.

Acknowledgements

First and foremost, our gratitude goes to our participants: Jason Bailey, Robbie Barrat, Nicolas Boillot, Sofia Crespo, Luba Elliott, Jake Elwes, Lauren Lee McCarthy, Sarah Meyohas, Manfred Mohr, Frieder Nake, Anna Ridler, Helena Sarin, David Young, and Joanna Zylinska.

Our thanks also go to Alexandra Francis for the beautiful illustration and design work she produced for this report. For coming up with the process which resulted in our report's cover art, we also thank Lise du Buisson.

Our thanks also to Paul Duckworth, who worked on a different strand of this project, and to Marie von Heyl, who came to share an artist's perspective as part of our project's Creativity x AI panel at the Rhodes AI Lab's Annual Conference on 8 June 2019, where we presented initial findings.

We thank David Sutcliffe and Jennifer Sin for their careful editing.

Finally, we are very grateful to the World Templeton Foundation for supporting this project.

All images are courtesy of the artists, except the Frieder Nake work which is courtesy of the V&A.

Table of Contents

Summary

This report accounts for the findings of the "Creative Algorithmic Intelligence: Capabilities and Complementarity" project, which ran between 2019 and 2021 as a collaboration between the University of Oxford's Department of Engineering and Oxford Internet Institute. In this report, we investigate the scope of human/AI creative complementarity through an interview-based case study of the use of current AI techniques in artistic work. By exploring practicing artists' engagement with ML technology, the report provides a fuller understanding of the spectrum of implications of AI––from automation to complementarity––in a domain at the heart of human experience: creativity.

Having identified the many communities working at the intersection of AI research and the art world, we focused on media and fine artists using machine learning (ML) as part of their practice. In the new space of "ML art", we found that artists self-identified in a range of ways which did not always centre AI technology––although technical competence is a highly valued skill. We found that new activities required by using ML models involved both continuity with previous creative processes and rupture from past practices. Major changes concerned the reorganisation of creative workflows around the generative process, the evolving ways ML outputs were conceptualised, and artists' embodied experiences of their practice.

We also found that artists highlighted a difference in scope between human and machine creativity. While ML models could help produce surprising variations of existing images, practitioners felt that the artist remained irreplaceable in giving these images artistic context and intention––i.e., in making artworks. They highlighted that the creativity involved in artmaking is about making creative choices, a practice outside of the capabilities of current ML technology.

Artists found many similarities between contemporary ML art and other periods in art history: for instance, the code-based and computer arts of the 1960s and 1970s and the harnessing of randomness by much experimental art. Many also found the generative capabilities of ML models to be a "step change" departure from past tools. Ultimately, most agreed that despite the increased affordances of ML technologies, the relationship between artists and their media remained essentially unchanged, as artists ultimately work to address human––rather than technical–– questions.

We concluded that human/ML complementarity in the arts is a rich and ongoing process through which artists refract technological capabilities to produce artworks. Although ML-based processes raise challenges around skills, resources, a common language, and inclusion, it is clear that the future of ML arts will belong to those with both technical and artistic skills.

Introduction

Recent advances in artificial intelligence (AI) have caused an explosion of interest in "creative AI". While we know AI technology is changing many areas of work, its potential impact on creative tasks and creative work is still unclear.

In this report, we take on this question.

Through a case study of the use of current AI techniques in artistic work, we investigate the scope of AI-enhanced creativity and whether human/algorithm synergies may help unlock human creative potential.

This deep dive into experiences of using AI techniques in creative work seeks to answer the following questions.

1. **How does using generative algorithms alter the creative processes and embodied experiences of artists?**
2. **How do artists sense and reflect upon the relationship between human and machine creative intelligence?**
3. **What is the nature of human/algorithmic creative complementarity?**

Towards these aims, we interviewed media and fine artists whose work centred around generative machine learning (ML) techniques, as well as curators and researchers in this field.

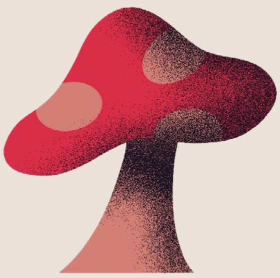

By exploring artists' engagement with ML technology, the report provides a fuller understanding of the spectrum of implications of AI—ranging from automation to complementarity––in a domain at the heart of human experience: creativity.

Cited Participants

Artists

Robbie Barrat	https://robbiebarrat.github.io
Nicolas Boillot	https://www.fluate.net
Sofia Crespo	http://www.sofiacrespo.com
Jake Elwes	https://www.jakeelwes.com
Lauren Lee McCarthy	http://lauren-mccarthy.com
Sarah Meyohas	http://www.sarahmeyohas.com
Anna Ridler	http://annaridler.com
Helena Sarin	https://twitter.com/NeuralBricolage
David Young	http://www.triplecode.com

Curators

Luba Elliott	AI Art Curator http://elluba.com
Jason Bailey	Founder and Curator of Artnome.com https://www.artnome.com/about-artnome

Researchers

Prof. Frieder Nake	Algorithmic artist, art theorist, and computer scientist (University of Bremen), https://compart.uni-bremen.de
Prof. Joanna Zylinska	Artist and Professor of New Media and Communications (Goldsmiths, University of London), http://www.joannazylinska.net

ve
nunities
Creati
Comm

Creative Communities

Many different communities work at the intersection of machine learning research and the arts. Focusing on fine artists and media artists who use machine learning in their work, we take another look at the field through their eyes.

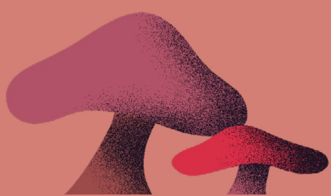

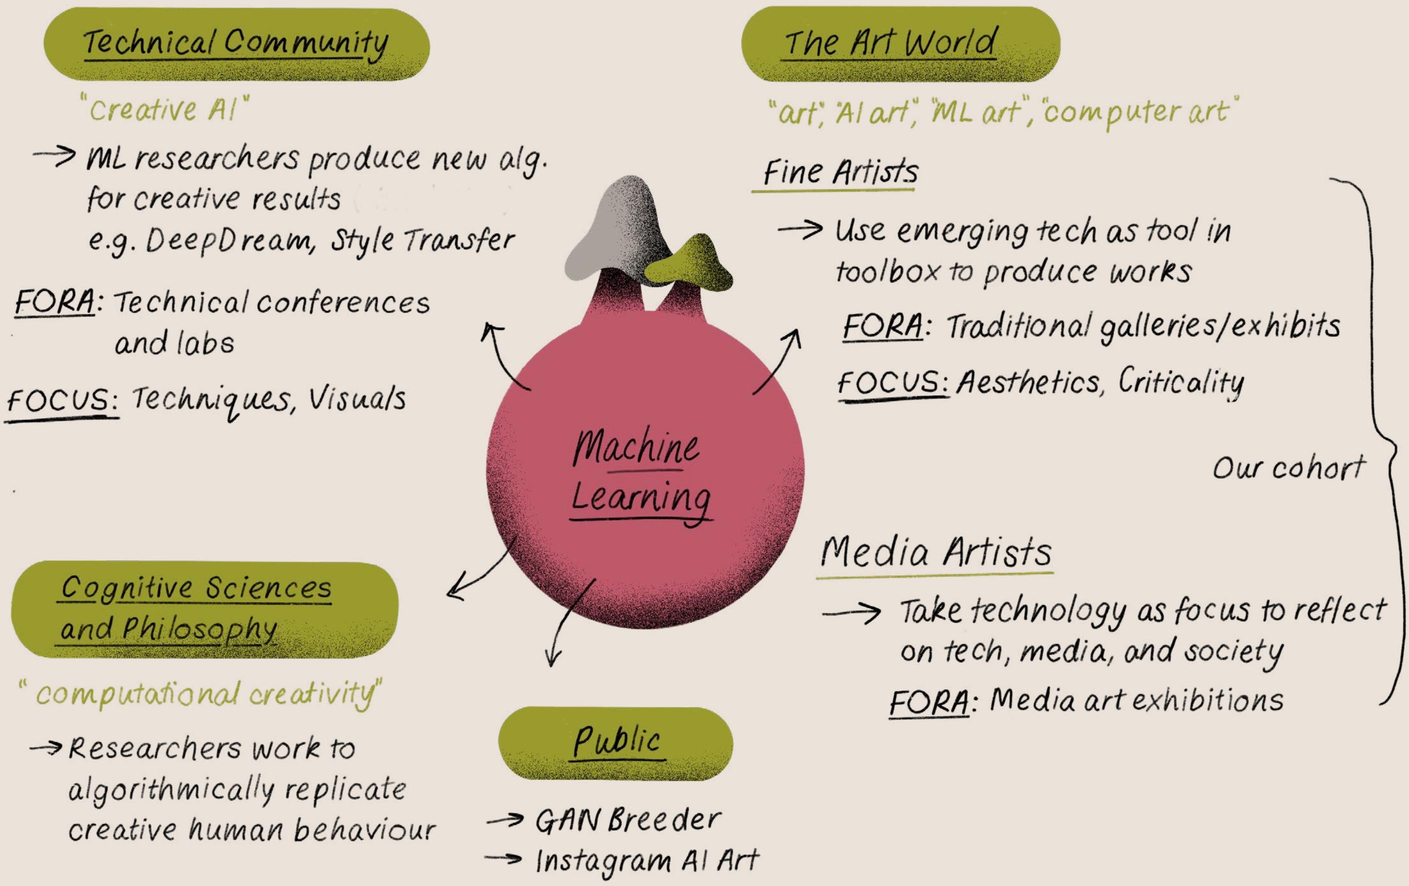

Figure 1. The Creative Communities Working with Machine Learning

1. A Varied Field

Several communities work at the intersection of machine learning (ML) research and the arts. Each is organised around their own goals, outputs, specialist language, forums, and audiences (see **Figure 1**).

In examining the shared use of ML techniques for creative purposes, our study is particularly concerned with the overlap between two of these communities, namely, the technical community and the art world[1].

In the **technical community,** ML researchers in computer vision work on models capable of producing interesting aesthetic results––with two notable examples being the DeepDream and StyleTransfer algorithms.

[1] Other communities also work at this intersection. First, computational neuroscience and cognitive science seek to model, and eventually emulate, human creativity as it occurs in the human brain. The associated term for this project is "computational creativity". Second, members of the public without technical or artistic training also make creative use of ML-based online platforms such as RunwayML or GANBreeder, sharing outputs on social media (e.g. using the #AIArt hashtag on Instagram). Debates about the "democratising" effects of ML-enhanced creativity often cite this second development.

Organised around the term "**creative AI**", these efforts are evaluated according to **technical** and **visual** merit, and results generally presented at technical conferences and circulated between research groups. While they are certainly motivated by producing visual results, most in the technical community would not self-define as working artists. As AI Art curator Luba Elliott explains,

> The researchers who develop the models, who came up with DeepDream or Style Transfer and so on, have a very different understanding of what art is. They focus on the aesthetics of an image. They're much more interested in trying to replicate the styles of paintings of the past than in participating in current developments in contemporary art.

This technical community is often mistakenly conflated with those working on more philosophical or neuroscientific questions around **computational creativity**, which focuses on meaningfully emulating creativity with software. For insiders, however, as Elliott explains,

> This is a known distinction. The field of creative AI is mainly interested in applications, not in philosophical questions around whether machines can be creative. The creative AI world is about doing cool things with AI tools. Most people in the field are less interested in this computational creativity question.

The second community sitting at this intersection is the art world. Responding to contemporary developments in technology and their impacts on society, artists pick up on emerging technologies in different ways. **Fine artists** adopt ML techniques to produce works within their own traditions, and often exhibit in traditional or mixed-media art spaces. **Media artists**, on the other hand, use the technology itself––often in the form of installations–– to stage a conversation around technology and society. Both streams—art using tech and art about tech— have specialised platforms for exhibiting works, as well as more general public art spaces.[2] Both use a range of field-specific terms for their work, including "AI art", "ML art", "digital art", "media art", "conceptual art", "algorithmic art", "computer art"––or, simply, "art". The outputs of both streams, in important contrast to the technical community, are judged both by **aesthetic** and **art-critical** criteria.

The open release of ML models and systems forms the central bridge between these groups. However, the different frames of reference described above influence possible interactions between communities, the circulation of skills and processes across subfields, and the stances different communities tend to take on human–technology relations.

[2] The separation drawn here between media arts and the fine arts is only a heuristic, and is not so clear-cut in practice. Indeed, efforts to categorise art into distinct fields is complicated, e.g. by a number of cross-cutting strands of practice (e.g., conceptual art), or by the ways in which members of different fields identify in different ways while doing similar kinds of work (or vice versa). Knowing this, our aim here is to draw only as many boundaries as are needed to have a conversation about high-level dynamics.

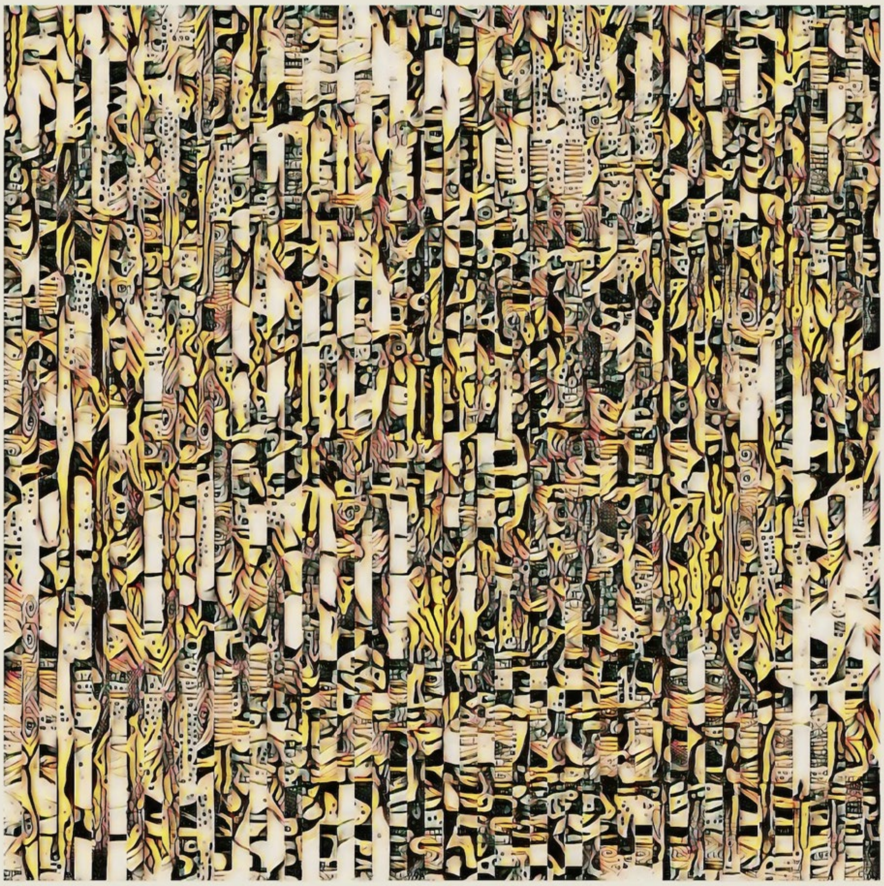

2. How Artists Conceptualise the Field

Our cohort of artists is diverse in disciplinary backgrounds, approaches, and directions; they also vary in how they self-define. Our respondents' backgrounds include fine arts (Elwes, Ridler), analogue techniques (Sarin, Crespo), digital techniques (Barrat, Young, and McCarthy), or mixed (Boillot and Meyohas).

"I'm not from a computer science background", Elwes explains. "I'm from a pure art background, which is quite unusual in this field." They consider their work to lean on the more conceptual side: "the whole system is what I think is interesting, and interesting to investigate. I see myself as a conceptual artist working with these tools." In contrast, Crespo recounts her transition to code-based art: "I was doing digital drawing, or digital collage, then a year and a half ago I discovered this kind of hands-on work. I started playing with machine learning and discovered the world of creative coding through it." Young, for his part, explains: "I definitely will say I am an artist working with AI."

spaghetti alla chitarra GANcommedia Erudita, volume II (2021) by Helena Sarin

These identifications are both reflective and strategic, and are in conversation with the evolving levels of societal interest in AI technologies. As Sarin explains:

> These days, for lack of a better term, I call myself an "AI artist". When you introduce yourself to somebody now, you don't need to explain what the "AI" means. You can just say "AI artist" and that's it. Because most people don't know what a GAN is, so you can't say "GAN artist". But these days pretty much everybody is in the loop with AI.

Learning Nature (b38,4106,16) (2018) by David Young

Our respondents also account for the **differences between technical and artworld communities** and qualify their relations. They describe a complex field, with **soft lines between technical and non-technical expertise and practice**, and with members of the community engaging to different extents with ML. As McCarthy comments:

> AI art is an interesting space because there's a broad range of people in it. There are researchers and engineers working with AI, then there are engineers with some sort of artistic background or inclination that are working primarily as researchers, but maybe also exploring artistic outputs. And then, there are artists, media artists, who are coming with some technical background also, but their primary medium is art. I've seen artists, for example, going back to school to get a PhD to study artificial intelligence, who have realised that there's a need to really engage in the technical realm as well. Then there are lots of artists engaging with the key questions and themes of AI, but in a less research-oriented or technical mode of working. But I'm hesitant to draw hard lines between technical and non-technical. The artists that are working to exploit social implications of these technologies are doing work that is crucial for technical development as well. Then, there're a lot of people who might not be engineers or artists, but who work in this space as sociologists, writers, activists, exploring AI in different ways.

While the boundaries are porous, **the AI Art field does tend to differentiate itself from those who are highly technically orientated but who have little knowledge or understanding of the art world**. As Elwes noted:

> I don't want to put down engineers working in art at all––there are people who succeed doing that––but some people don't necessarily come at it from quite the right angle. I see lots of parallels with the early days of photography, when many photographers were trying to validate this new medium by mimicking paintings. These were the engineers, the people creating the chemical development processes, and they would take photos of Renaissance paintings and say, "this is a valid art form".

For Elwes, imitative work adds nothing new to the conversation:

> That is so uninteresting, in my opinion, and so derivative. If it's making a work in the style of a painting to ask a question that provokes people into having a really interesting discussion about something, then that's a different question. But that is the human artist's intention, and their skill in framing the work to make people have that debate.

A background in art is thus felt to be necessary to ensure an adequate depth of conversation. Without that understanding, as Boillot noted, a creator of ML art "does not know where to go, where to put depth, how to make people think about something that's happening."

Drawing from her own practice, McCarthy reflects on distinctions within the field and her commitment to centring the artwork:

> Now primarily, I think of myself as an artist, though some of that is my capacity to work as a programmer. Programming is part of my artistic needs. But when it comes to thinking about what projects I work on, and how I work on them, I remind myself that I'm an artist. My projects often involve some amount of technical innovation, and I often have moments in projects where I realise how far I could take things technically to implement the idea that I have. Remembering that I'm acting as an artist helps me find that balance. Some of my projects could become research projects, but that's not really the goal. The goal is to think about the audience of this art piece and to create the experience that I'm envisioning for the artwork.

Photograph from *SOMEONE* (2019), an installation by Lauren Lee McCarthy

This exploration of the boundary between art and technology is nothing new––and certainly predates the invention of machine learning. Indeed, as Young describes, finding this line is a central question for the technology-based arts:

> There's a surprising diversity of artists working with technology. My whole career has always been about working with the emerging technology of the time, whether it was AI in the '80s or interactive media in the '90s or later the web and mobile devices. It's always been about asking, 'how do I find what's interesting in that threshold of the technology that's emerging, and in what we do with it?'

For our respondents, an important feature of working in this field was the **lack of a common language**––and the need to develop one. About one particular collaboration, Elwes recounts:

> I was brought in as a translator of sorts, which is an interesting role for an artist. Most of my artist friends can't talk to programmers about what they're doing, because they don't have the entry-level knowledge to be able to talk about machine learning algorithms. I understand enough of both sides that I can try to bridge the gap, and that's an exciting place to be. But it goes both ways: when an artist tries to talk about their more philosophical or critical ideas, quite often, computer scientists get quite lost. There is a lot of talking past each other.

Elwes stresses this is not a one-way conversation: for the field to develop, a two-way relationship is needed.

> Once, I was on a panel discussing what AI can offer the arts. But it's also worth asking, 'what can the arts offer to AI?' I think the artist can offer fresh perspectives to researchers, as well as the other way around. Artists can find limitations, find hacks, point out biases––but using art as a frame, and getting more people to enter the conversation from that different perspective.

In artworld terms, the field of ML Art is still **very new**––its inception often traced back to Alexander Mordvintsev's release of DeepDream in 2015––and **still emerging**. As Ridler explains, "there's still not a consensus as to whether this is a new movement, or whether machine learning is a tool that people are using in their practice that fits into kind of existing strands of art history. But there's no reason that it can't be both."

Consensus or no, this emerging field has already produced manifestos (e.g., Helena Sarin's Neural Bricolage and Mario Klingemann's Neurography) and dedicated exhibitions: Gradient Descent at Nature Morte in New Delhi (2018), Artificially Intelligent at the Victoria & Albert Museum in London (2018); AI: More Than Human at the Barbican Centre in London (2019); Entangled Realities: Living With AI at the House of Electronic Arts in Basel (2019); BARRAT/BARROT: Infinite Skulls at L'Avant Galerie Vossen in Paris (2019). The field also enjoys dedicated fora such as Artnome.com on the arts side and the NeurIPS Creativity & Design Workshop on the technical side––a combination which usually outlines the shape of a new movement in the arts.

In terms of its reception, ML art is still on the margins of the art world, not unlike previous technology-based communities in the arts. As Mehoyas describes it, ML art is "still the cousin of the art world, still on the fringes". Yet even within this niche intersection, there is a mainstream and an edge. As Sarin puts it:

> I see a lot of mainstream work, like the Obvious piece that auctioned for half a million. If you try the #AIart hashtag on Instagram, you would immediately see what I mean. It's mainstream and derivative, but that's what the public probably thinks about when they think about AI art. But now, a lot of writing is also trying to show a different side to AI art, like Jason Bailey's articles in Artnome which showcase Anna Ridler, myself, Mario Klingemann, Robbie Barrat–– how each of us works in a different direction and in a different way. Everybody is going in a different direction, and that's exciting.

{organic_resonance} from Neural Zoo (2018–2022) by Sofia Crespo

"Hype" is an important component in responding to technological developments in the arts. As Boillot explains, "I first wanted to use machine learning because there was hype. There was huge hype. You want to be the first to use it. Every time a new tool is created, artists engage with it in their practice". Sarin concurs. "Technology drives the artistic process––because novelty is the thing. You need to invent or find breakthroughs, so you look around. You see a new medium and try to use it for your own purposes." Though hype occurs with every technology, Crespo notes that interest eventually wanes. "If wanting to create something new is something that drives you", she says, "then once a technology becomes very naturalized in art, it stops being exciting. Then we look for the next thing." Despite an uncertain post-hype future, the fact that the ML art space is new strongly appeals to its pioneers and early adopters. As Elwes comments:

> That's why this is a fascinating time. Saying 'this is a new medium' or 'art movement' or 'genre'––as some people are saying––seems very grandiose, but it's also really exciting. The rulebook hasn't been written yet, and that's such an exciting place to be. We don't have the answers to these things. We are constantly repositioning ourselves and working out our views.

Film still from *CUSP* (2019) by Jake Elwes

Partly because the genre is new, partly because it is technology-based, responses from the traditional mainstream are often critical. Indeed, critical responses to code-based art are a deep-seated stance in the fine art world. As Young explains, "Christie's and Sotheby's have hosted panels and conversations about AI art—the problem is that if they're reviewed by somebody coming to them from a more traditional background, they're going to be quite critical and cynical about it." Elwes concurs, recounting an experience while training at the Slade, a British fine arts academy:

> I was down in the basement coding. None of the professors really got what I was doing. They gave me a really hard time. I'd gone to Chicago to really get into coding and generative art. I'd met some of the establishing figures in the community: Casey Reas, who invented Processing, Christopher Baker, who was one of the main contributors to OpenFrameWorks, the main languages that artists used, and others. But when I came back to London, my professors said, 'What is this? It's all spectacle, all demo. There's no criticality or message'.

Opinions vary on whether this stance will ever shift. Some, like Meyohas, believe that technology-based art will always occupy a marginal position in relation to the fine art, because of the lack of what she calls "the physical index". Others, like Sarin, believe divergence is due to the field's novelty, and that lines will blur once artists no longer need to spend so much time writing their own code. This is a possibility Meyohas also considers: "ML is too difficult to use for now, but it will become a tool. There's no doubt about that." With a lower technical barrier to entry, ML art may go mainstream, "but you need people who are in the art world to be adopting it for it to change. If Damien Hirst starts making AI butterflies, that becomes part of the tool set." For now, our cited artists speak from the edge, and take the tensions of their position in the artworld as a challenge.

Machine Learning and the Creative Process

Generative ML models clearly have an important impact on the creative process of many artists, with activities required by using ML models involving both continuity and rupture with previous creative processes. These shifts are most visible in the reorganisation of creative workflows around the generative process, conceptual shifts around the nature of ML outputs, and evolutions in artists' embodied experience of their practice.

1. New Activities in Artistic Practice

Overall, we found five new activities associated with the use of ML models in artistic practice: (1) technical research, (2) selecting or building models, (3) building datasets, (4) training models, and (5) curating outputs.

1. Technical Research

Technical research involves understanding what machine learning models are available, how they work, and how to leverage them for artistic purposes. Especially for artists without a computer science or engineering background, intensive research into technical ML literatures is essential to understand the behaviour of different models. This can be a time-consuming and challenging part of the process As Elwes describes,

> Using machine learning is such a steep learning curve for me. I understand enough of the technology to use it and hack it, but I'm not writing algorithms myself, so it often takes months of research to work out how to use a model and get it to do what I want it to do. To be able to see some of my artistic voice coming through a black box or a ready-made, and then find an interesting way of subverting it. It's a long process, not something you can just play with lightly. I can't just make a bad painting, I have to have the concepts and then do the research, much like what people do in academia.

Vectoglyph (from emoji) (2019) by Nicolas Boillot

2. Using and Building Machine Learning Models

Once the research process yields a promising model, either specific **algorithms** are then **written from scratch** or existing ones **procured**, and potentially **modified**, to suit a desired outcome. Artists vary in the extent to which they modify their algorithms, but some degree of coding is always required––amounting, in the end, to a kind of artistic signature. As Sarin describes,

> Right now, everybody who works seriously with GANs writes their own code, to different degrees. I don't change my algorithms much. Mario Klingemann does a lot of tweaking of his algorithms, as does Robbie Barrat. Anna Ridler doesn't, and neither do I. I don't write my own frameworks or algorithms––I use CycleGAN But I still write a lot of code for pre-processing, for post-processing, for changing high-level parameters to automate some parts of the pipeline. That's what adds uniqueness to the art.

Outputs from a GAN trained on ink drawings from *Fall of the House of Usher* I (2017) by Anna Ridler

3. Using and Building Datasets, Training Models

The models––ranging from off-the-shelf and pre-trained to completely DIY––**must then be trained** on **existing, curated, or custom-made visual datasets**. The training process opens up a new kind of artistic working space: "latent space", the statistical topography of the aggregated features of the training images, mapped onto a high-dimensional mathematical space––and an additional set of choices for artistic outputs. As Bailey describes, "the unique point of latent space is that it's accessible, almost as a physical space. It's what DeepDream did first, access something that was there in a cool way." Elwes, echoing others, also highlights the role of latent space as working material: "You can move through this multidimensional space, go through this journey through what this network has learned, through the high-level and low-level features."

The **landscape of latent space** depends on what kind of image dataset was used to train the model. These can vary according to two axes: (1) custom-made, curated, or ready-made images, and (2) large-scale or small-scale datasets.

On the **first axis (dataset content)**, custom-made datasets are the most labour-intensive, involving individually generating training images (e.g., photographs or drawings), cleaning, and finally sorting the data. Curated datasets involve the comparatively lighter process of curating existing material into a meaningful set of images. Artists training models on their own data tend to amplify their own style and work, entering into a game of mirrors with their artistic voice. By contrast, artists using existing data, either ready-mades or pre-curated datasets, tend to turn outwards to reflect or amplify aspects of the social world. For example, Elwes' process for Zizi (2019) involved selecting a network pre-trained on celebrity faces and fine-tuning it with a hand-curated dataset: "I gathered a few thousand high-resolution images of drag performers and used that to sort of corrupt, or dirty, their dataset". By doing so, they staged a conversation around queer erasure and normative visual grammars by engaging with, and correcting for, absences in existing data.

Outputs from a GAN trained on drag performer portraits from *Zizi* (2019) by Jake Elwes

On the **second axis (dataset scale)**, artists can either choose to use very large datasets or favour smaller, human-scale data. From a technical standpoint, ML models function best with large amounts of data. However, as artists' goals are not naturally aligned with practical ML applications, many of our respondents choose to work the model against the grain. Young describes this stance in his own practice, which he calls "Little AI":

> The work that I'm doing with "Little AI" contrasts with what we tend to think AI is. We think of AI as something that requires massive amounts of data. But we can also approach AI in a more intimate way, a more personal way, that allows us, as individuals, to have a different intuition for what AI is. Rather than feeding the machine learning system thousands or hundreds of thousands of images, my work is about feeding it a handful, almost "breaking" the system by giving it so little. I'm curious to see if that gives a different feeling for what the technology is, as well as change the kind of outputs that the system produces. I don't find AI systems trained on massive amounts of photos, aiming for hyper-realism or photorealism, particularly interesting. That is about bigness and quantity––but if you're feeding it small stuff, you suddenly get a little bit more understanding of what is happening within the machine, and how the code that's driving it works. Part of it is also, again, in contrast to the notion of AI being used for business, for scale, and for efficiency. I wanted to find something different.

In this way dataset scale is a **proxy for aesthetics**: large datasets yield **photorealistic results**––as well as, incidentally, require a proportionally larger investment of time into cleaning and organising data at scale––while tiny ones yield **appealingly imperfect** ones.

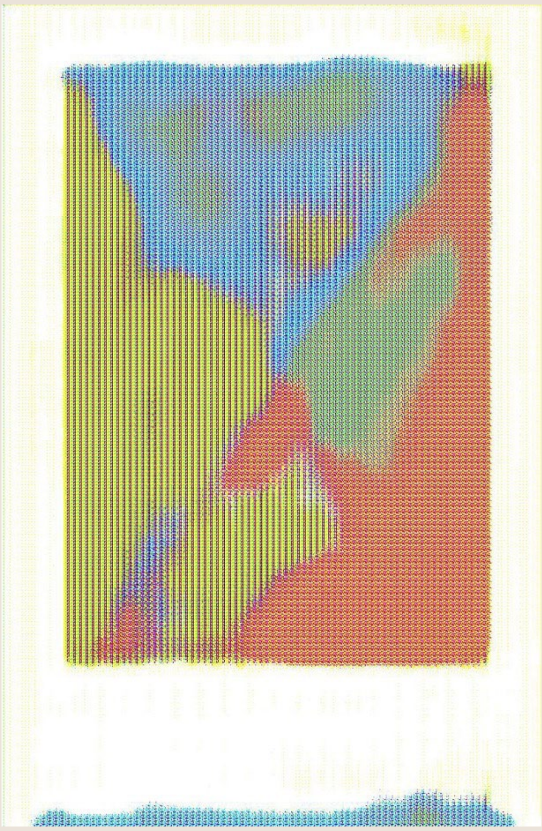

Tabula Rasa (b62x2,715) (2019)
by David Young

4. Combining Models

Different ways of **combining models** are available to artists to increase the granularity of their control over their desired output. "**GAN-chaining**" is one such technique. As Sarin explains, GAN-chaining can happen in one of two ways. The first concerns generating the content of an image:

> I usually start from noise. I make small images, and then I pass them through several GANs. That's called GAN-chaining. Then I upscale them. The first output of the GAN is basically completely unexpected. Let's say I give it a bunch of my flower photos, and it generates hopefully something that looks like a flower, but because they are small and I start upscaling, the next GANs can take me to some images that, semi-abstracted, might have nothing to do with flowers. Then I decide whether I can do something with them, or abandon the idea and start something different.

Drawings before (top) and after (bottom) application of a trained CycleGAN 'filter' by Helena Sarin

GAN-chaining can also be used as a controlled intervention to modify the texture of an existing image, a strategy Sarin calls "**personal filters**":

> If I want to change the colour of an image, we enter the realm of what I call my "personal filters", a bunch of pre-trained models that can impart certain textures. I can take images, pass them through this other GAN, and anticipate, to some degree, the output. This is a more controlled process.

5. Selecting Outputs

Alongside the curation processes at work around datasets, **curation** is also involved when **selecting outputs from latent space**. As Ridler explains,

> The issue of latent space is also a question of artistic agency. Some people will really edit what comes out of latent space, others just allow latent space to run as part of their work. There's the curatorial aspect of the dataset and then there's the curatorial aspect of latent space. Both are about human agency at different ends of the spectrum.

Ridler, for her part, prefers to edit: "when I make work, I never want to display all of latent space. I want to carefully control and edit the latent space to tell a story." This can prove a challenging step, as Sarin describes:

> The most challenging part is doing the curation because I want to post them all, to print them all. Making the selection is the difficult part. That's why it stays fresh, because with every step you get more and more interesting new material to work with.

Fall of the House of Usher II (2017) by Anna Ridler

Further, Barrat highlights the importance of timing in the process of curating from latent space. He describes this through his experimentation with landscapes:

> I watched the GAN learn, and I chose when to cut it off from learning. That sounds a bit arbitrary, but how long it learns really does affect the output. After the networks had already learned how to make landscape things well, I noticed that every few days of training it would alternate between making really bright, psychedelic landscape paintings, and really dark, moody ones. So I cut it off when it was making the darker paintings because I enjoyed those more.

A GAN-generate landscape (2018) by Robbie Barrat

Individual artists vary widely in how they carry out these new activities, including in the kinds of techniques used, the role and scope of ML techniques in the artistic process, the degree of tool customisation (or "**hacking**"), the extent of intervention in datasets, aesthetics (linked to dataset scale), exhibition choices, and the artist's stance on the nature of their artwork. Ridler calls these variations "a spectrum of working" (see **Figure 2**):

> Trying to generalize about machine learning art sometimes flattens out the fact that different artists are using different approaches, trying to explore very different things, and working in very different ways. There's a spectrum of working. The first way is where you are taking very readily available things, like ImageNet, and exploring what that does and how it works. That can be a really interesting way of working, exploring the limitations of big tech companies and things like that. The second is doing work like Mario Klingemann does. He'll scrape images from the internet, but it's his technical skill that is contributing something really innovative. The third is more the way that I work, where I build everything myself, the data set, and modify the algorithms. These are three very different ways of working and I think they very often get collapsed.

Many of our respondents also highlight the flattening effect of generalisation in this field. As Young expands,

> There are so many different artists working with technology. Some of them are working around how technology relates to the body or to gender or to other themes; a lot of artists work with blockchain, with augmented reality and virtual reality, etc. What's interesting is people's confusion, from the outside, about how these things are different. At the same time, there's also confusion internally––for example, some of the images that I'm making are available for purchase on blockchain-based websites and not as prints.

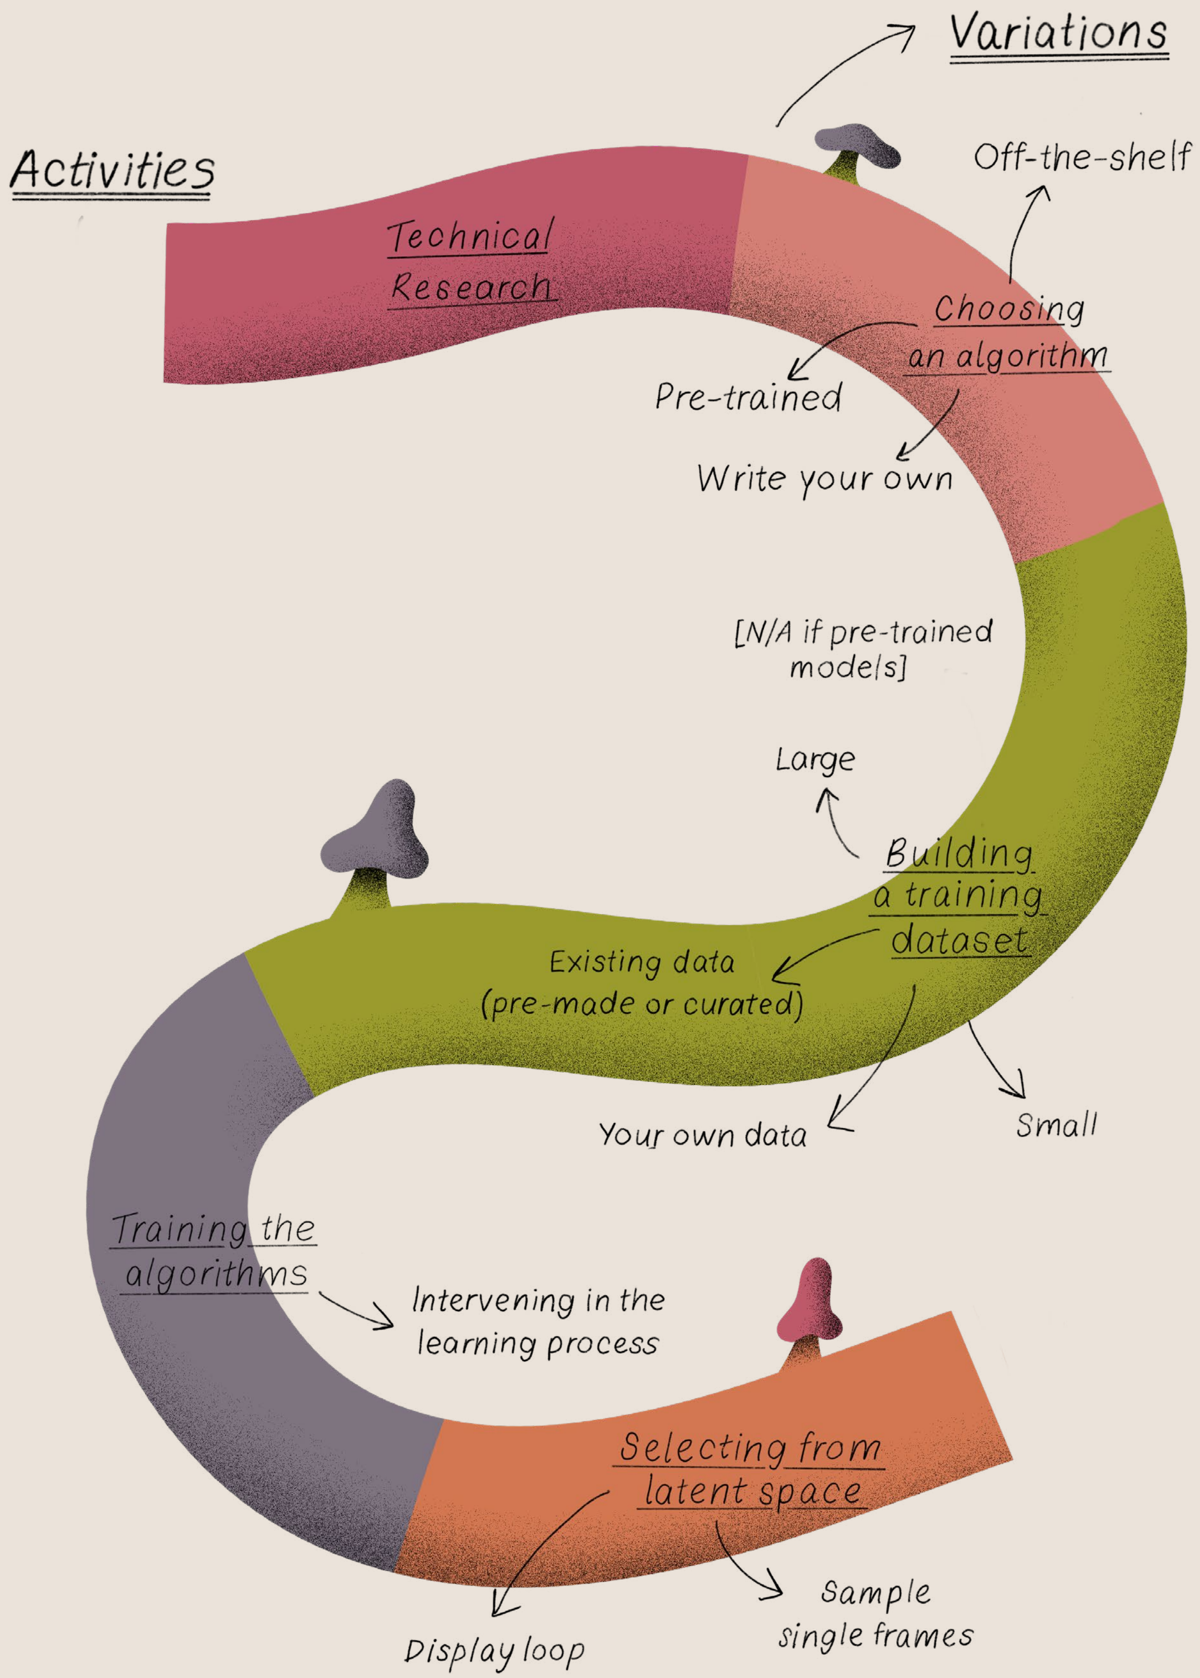

Figure 2: Five new activities involved in ML-based artistic practices, and their variations. These constitute a "spectrum of working".

2. Reorganisation of Creative Workflows

Artists describe **ML-based workflows** as having a particular **structure centred around the generative capabilities** of the models, and consisting of distinct parts: "a research part", "a creative part", and "a mechanical part". These phases, as Ridler explains, are later followed by an out-of-studio "dissemination phase"––with each phase providing a space to explore technology as material:

> My practice as an artist goes into three phases for each project: a research phase, a making phase, and then a dissemination phase–– when I'm installing, exhibiting, speaking, or writing about something. For each of those phases, I'm always really interested in what technology can do, how can it push what I'm trying to do in a way that I wouldn't be able to achieve otherwise.

Importantly, **an ML-based process is never an ML-only process**: there is often a dialogue between analogue and algorithmic techniques. Barrat describes how he integrated techniques used by his collaborator, French painter Ronan Barrot, into his own algorithmic work:

> I've been doing these works called 'corrections', inspired by what I saw in Ronan's studio. Ronan would cover up a little piece of a painting that he was making that he didn't like, for example with bright orange paint, and then fill it back in to correct it. I've been teaching neural networks to do a very similar thing. One of the favourite pieces that I've made recently is a correction of Peter Paul Rubens' '*Saturn Devouring His Son*'.

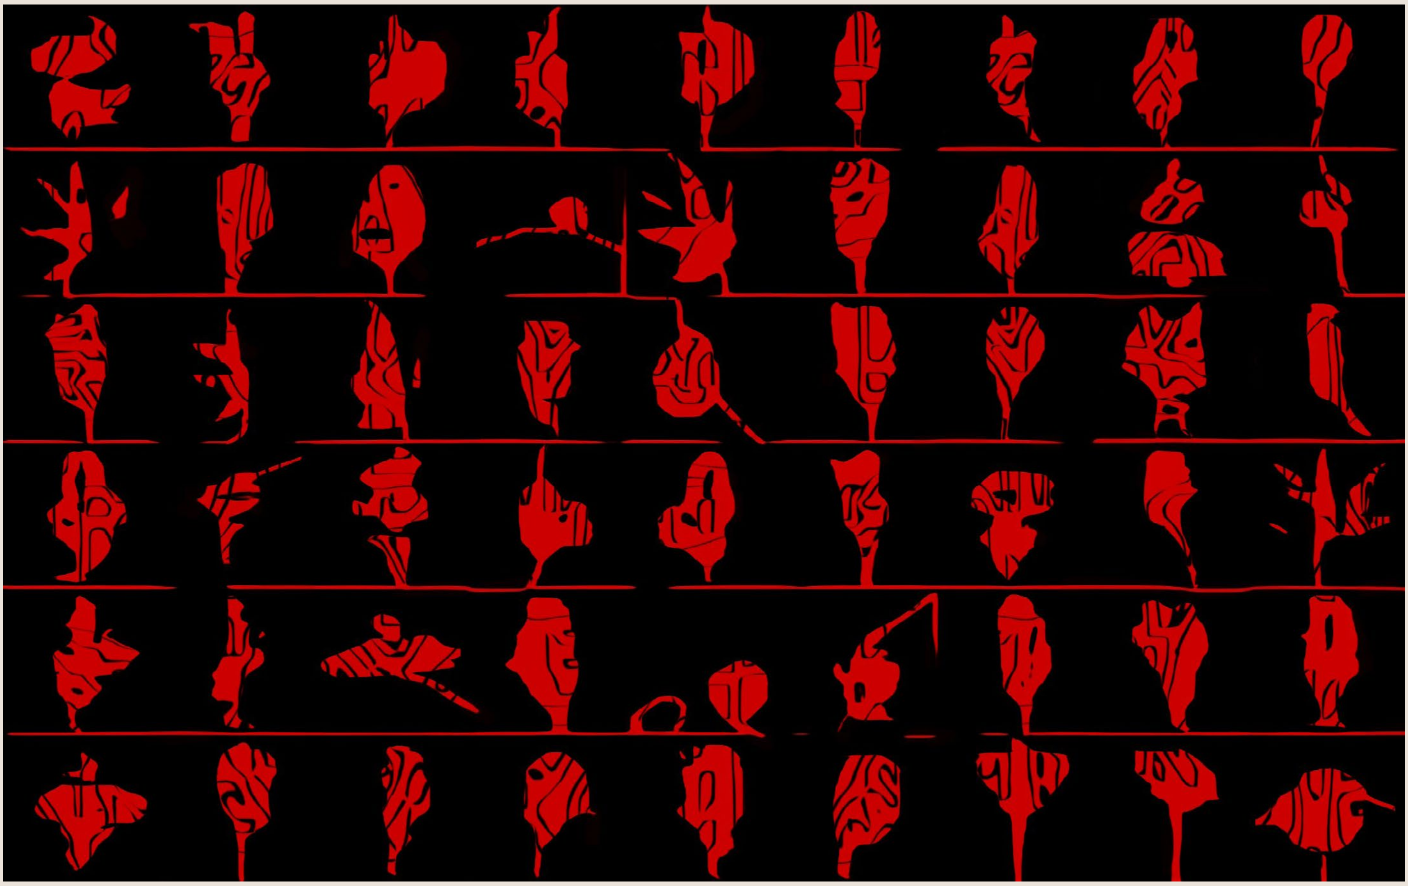

Sarin describes a similar process, inspired by a past teacher:

> Say you made a painting––not like oil painting, which needs a lot of advance preparation, but something lighter like pastels or watercolours. In this teacher's workshop, the mantra was, 'There is something in this picture that works. Cut it out. That's your final work.' And with GANs, I sometimes do the same. I cut some piece of it that's interesting and work it into something else.

54 characters in search of a platform (2021) by Helena Sarin

Artists differ with regards to the relative importance of ML models in their practice: some build a full digital practice around algorithms' generative capabilities, while others translate back-and-forth between digital and analogue methods at various points of the process.

Further, the dialogue between analogue and algorithmic is not only linear, but **synthetic**, as ML-based processes bring together different activities into one coherent practice. In Ridler's words:

> Machine learning really did revolutionise my practice. Because of the way it's structured, with the dataset and the algorithm, I found that I'm able to condense lots of different interests into one process of working. Before, if I was feeling interested in drawing, in narrative, or in sound, I would have to pick just one or two of those to work with. But with machine learning, I'm suddenly able to take all of those interests and use them in the same process. I'm not just clicking a button, or scraping stuff from the internet and hitting enter in the command line. It's a very iterative process. There's work around building a dataset, and questions around archiving, labelling, categorization, language, and memory. You can include and play with all of those things in the dataset. Then you move towards working with an algorithm, and there's a whole other history of rule-based art there, a whole different set of issues around chance and control.

Crespo, for her part, compares this deliberate synthesis of processes to constructing a symbiotic organism:

> This jellyfish, the Portuguese Man o' War, is a combination of different organisms living together, forming one. I find this concept of collecting different pieces of life together to create a single one really interesting. I explored this idea with another artist during a workshop: how do we combine different machine learning techniques to do one thing? Automate one thing, then automate another thing, and another thing, and put them together. It's like one organism made from a lot of different little automations. This somehow feels natural, as it allows me to focus on different things––a research process, a process of combining things in a deliberate way, and then seeing what happens.

Meyohas, for her part, accounts for the process as a meta-curatorial one:

> ML makes the creative process really different. I can be the curator putting all the different elements together, and that makes machine learning really exciting. Working with technology means things never work at the beginning. It requires patience. It does require new skills. It requires being open-minded to technology, and also being able to think really holistically. I try to go into a project that is not just about AI, that's using AI as a tool for something bigger.

Ridler describes another kind of evolution in her own drawing style––not aimed towards a particular output from the GAN, but as a way of incorporating her experience of working with GANs into her hand-drawn style:

> With 'Fall of the House of Usher', it took my drawings and made them kind of wilder and freer and much more interesting than they were before. I spend up to a year making a single project. That immersion changes the way that I construct imagery. It changes the way that I engage with the world and my work. My drawing style has shifted dramatically from doing the Fall of the House of Usher project––even to this day, I draw artifacts into my work, and things like that. All these very unexpected traces of each project became embedded in my practice.

GANs thus function as **adversarial counterparts** of a kind––filtering choices, providing feedback, and opening up new spaces for artists to take their practice forward.

This synthetic aspect of ML-based processes thus gives artists the space for multiple creative interventions at different moments and levels of their practice. Further, as Ridler describes, this is a **very iterative** way of working:

> When you get to the whole suite of making and constructing datasets, building algorithms, working with the output, it never works going one, two, three. It always works going, one, two, one. One, two, two, three. Three, two, one. You're always going back and forth between all of those different parts.

Further, the iterative nature of ML processes allows the generation of **feedback loops**––explicit or otherwise–– in practice. Artists, for instance, find traces of their hand-made datasets in the outputs, or modify their hand-drawn styles to influence the model's generation. Sarin, for instance, describes shaping her practice in response to her GANs' outputs:

> GANs give me feedback into my analogue work. When I worked making drawings to train GANs, I started to notice that things like dirty shades negatively affected my GAN's outputs. Now when I draw, I keep this in mind to make sure I get a clean output from the GAN.

As with the discussion of activities above, these overarching themes play out differently in artists' workflows. Examples from four of the artists we interviewed are shown in **Figure 3**.

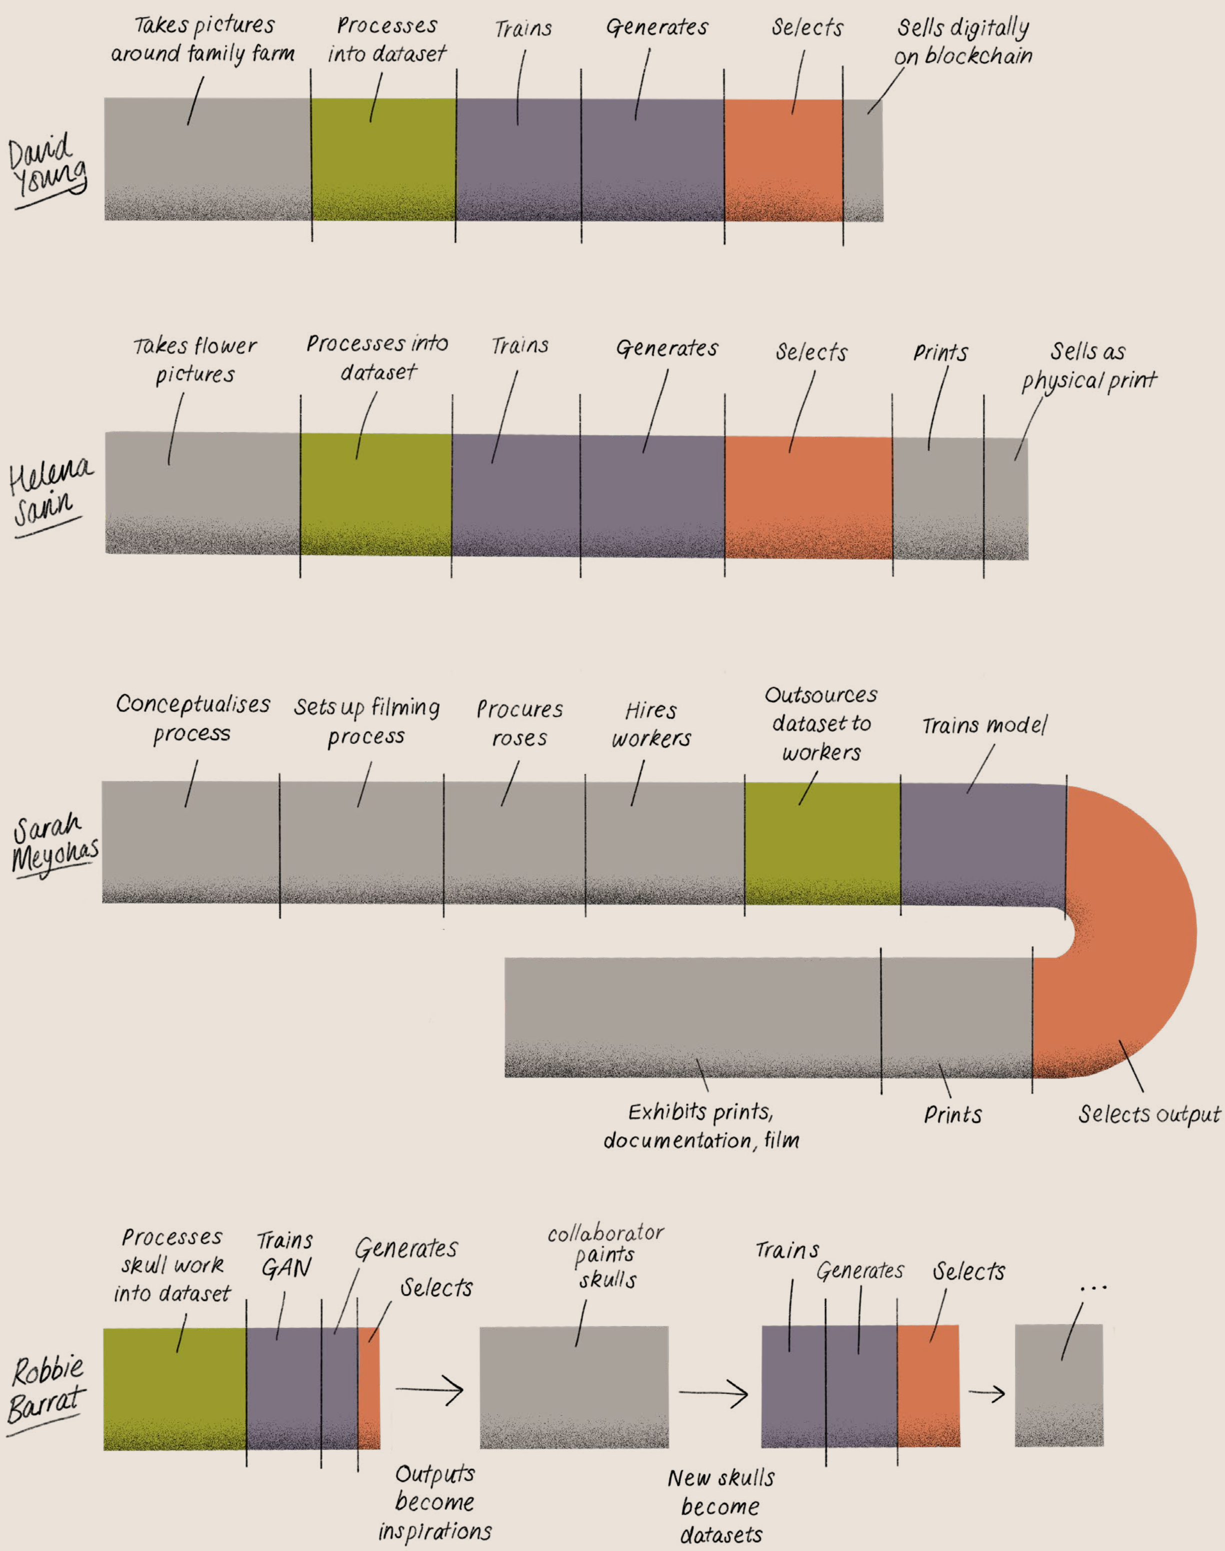

Figure 3: Four artists' workflows, with the ML-based process in bold

3. Artistic Affordances

A **medium** in the arts refers to both the method with which artworks are made and the material they are made from––more generous than the term tool, and more suitable in this context than the technical term agent. We find that artists use ML media either **with** or **against the grain**: in other words, artists are either interested in the models' **generative capabilities**, or aim to explore the artistic potential of machine failure. While the former produces smooth, often photorealistic results, the latter––favoured by our cohort––exploits 'glitches' and machine 'misunderstandings'.

The main affordance of ML models lies with their ability to learn. Instead of specifying rules for the model to execute, artists, through curating datasets, get the model to infer them––a paradigm shift from the previous wave in algorithmic art. As Barrat explains:

> With old computer art, the key difference lies in rules. ML art is a bit more exciting, because in traditional computer art, a programmer is feeding the machine rules through programming. If you wanted to make a generated landscape in the '70s, you would have to program it into the computer, draw ground on the bottom, draw trees in the middle, and draw sky on the top of the painting. You would have to hard-code rules into the computer, and the computer would execute those perfectly. But with neural networks, what happens is you're still feeding rules to the machine, but it happens through the dataset. You choose a dataset that conveys the rules that you want to get across, then you allow the machine to try and pick up those rules.

Ridler concurs:

> I do think it is a different paradigm because I think just the possibility for it to just kind of do something totally unexpected is much higher. The possibilities of it doing something unexpected within algorithmic art are always within certain bounds. The type of unexpectedness that you get from machine learning, it's just so wild, and so crazy, and just amazing. I think it is, fundamentally, a shift.

Aesthetically, ML models also allow for photorealistic outputs, sometimes leaving audiences confused as to whether the work is generative or non-generative digital art. As Crespo notes, "a lot of people ask me whether the images are done using Photoshop". Though models can produce sophisticated visuals, they are currently still in what Sarin calls "**the textured regime**", producing wabi-sabi patterns but lacking any semantic understanding.

Respondents unanimously report that the most interesting dimensions of practising with ML models relate to **mislearning**. Indeed, the models' inferences stage a confrontation between artists' expectations and machine perception: what did the model actually learn? Did it learn what the artist thought it would? This process generates surprise and is often conceptually or aesthetically productive. As Barrat continues,

> The misinterpretation part, I think, is the most exciting part, because you can get really surprising outputs, that you, as the artist, were not expecting. You could do that with traditional computer art, but it would be because you made an error within your parenthesis, or you messed up the logic on a line, so you get some glitchy output. The glitches or misunderstandings that happen with neural network art are a lot more meaningful. The sort of errors that happen with neural networks are more high-level, not like misplacing the position of a line. With my nude portraits, the network totally misunderstood the high-level organization of a nude portrait. It's a semantic glitch.

From the *Corrections* series (2019–)
by Robbie Barrat

The history of art shows that **glitches** are often **artistically desirable**. ML art is no exception to this: while the capabilities of ML models are valued by our respondents, most were particularly interested in their edges: the artistic potential of machine failure. As Barrat explains,

> I really wanted to introduce some sort of misinterpretation. I thought that the landscapes were a bit boring because the network got it right. We have so many landscapes. It just seemed boring for a neural network to produce more plain old landscapes.

In this way, that models are not technically perfect keeps them interesting. As Sarin explains, if models "got too good", she would have to find other tools:

> GANs are not perfect and that's why I work with them. I use two types of GANs, one is maybe three years old, and the other over a year old. I intentionally don't upgrade to new advances because they push the more photorealistic stuff, which is exactly against my process. If they start going for completely photorealistic, I will have to find a new branch in this area.

In the end, often our respondents deliberately avoid working with the grain of machine capability in order to keep, as Sarin calls it, to the "human side of AI".

Film still from Zizi – Queering the Dataset (2019) by Jake Elwes

4. Translated Outputs

ML art has particular features which require ongoing translations between digital and material. First, since generative models are **probabilistic** in nature, the outputs artists generate are, strictly speaking, **not single images, but classes of images**. As Nake explains, in algorithmic art:

> The entire set is the work of art, so to speak. But there'll never be all of them because it would take way too long. Therefore, the individual piece now is only an instance of the class that it belongs to. The works that people like me, if we do static works that are being put on a wall, those works are no longer interesting. Of course, I do hope that they are nice. But I as a theoretician, I must, particularly nowadays, I must suffer, if you like, these discrepancies. That I'm doing only trivial things, parts of an infinity of works.

This complicates the status of the print as the output of a visual art process. No longer an output by default, a print is a translation of the output into a static visual work. The next step, as both Nake and Boillot explain, is finding ways to exhibit the entire class of works in non-static installations. As Boillot suggests: "For example, you know Mario Klingemann? He sold a computer to Sotheby's: that's the next step".

Some, including Boillot and Barrat, push this reasoning further to argue that the **true artwork** is, in fact, **the algorithm itself**. As Boillot explains:

> What most artists said, and what I think also, is that the artistic process is in the writing. It's not in what's produced. It's more conceptual. But it's very difficult for the non-expert to think like this, because AI is like a magical box.

However, these curious features of algorithmic art remain mostly of interest to theorists and conceptual artists, as the public and artists themselves overall prefer the materiality of prints rather than digital displays or conceptual gestures. As Meyohas comments, "there are definitely motions towards screen-based experiences as a more prevalent form of art, but from my perspective, that's still very narrow". She continues:

> What I've always seen is that photography always sells at a huge discount to painting because it doesn't have the physical touch. This is fetishism, but people really love the fact that you touched it. It's an index of your movement at a certain point in time. My VR piece that runs all the time is not the same. My hand isn't in it in the same way. There's no physical index. I'm in conversation with a museum in Shanghai and they want to show my videos, which is great, but what they really want to show are the pressed petals and the sculptures, because they want something material. When people figure out how to use AI as a tool to create compositions or forms that they wouldn't have been able to create otherwise, but then turn that into physical pieces, that's where the money will be.

Sarin agrees, describing the way her audience gravitates towards prints even when she exhibits her work in different formats:

> I recently did this private event where I had a big screen projecting my images non-stop, and also a small exhibit with a dozen images. My images are usually small, A4 or A3. People still spent all their time looking at the prints, studying them. Hardly anybody looked at the screen. I mean, you sit at a screen the entire day People aren't interested.

A few decades prior, Nake describes going through the same cycle with his pen plotter:

> The computer I used produced a paper tape with holes, which then went into the drawing machine to control it. Once, somebody asked, "Oh, why don't you just look at these holes?" They were right. The code stands for the drawing. But I want to see it. I myself used to think "think the image, don't make it". But I don't want to think the image, I want to see it. And the seeing comes from the action making.

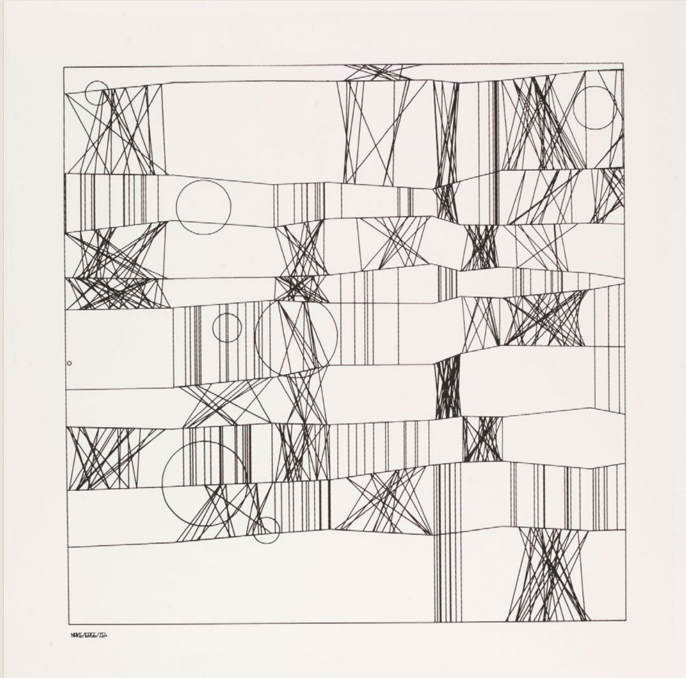

Hommage à Paul Klee, 13/9/65 Nr. 2 by Frieder Nake © Victoria and Albert Museum, London

He further argues that the field of algorithmic art is moving through a second necessary translation, that between the static and the dynamic image.

> In the world of computer art, the static image is okay. I would call this "the static phase", the "McLuhan phase" of algorithmic art. There, the medium is the message. Static images, as pieces on the wall, only say, look, I'm not very good, because the painters have done it much better and the graphic artists also, but I come from the computer. That's the message. A medium first tells us, look, I'm the old content in a new way, in a new medium. Either that new medium disappears because it doesn't have much potential, or it develops. In our case, the computer, it has developed tremendously. It reached its own mediality, if that is a word, in movement, in the dynamic image.

ML art is thus a space in which different and contested kinds of translations occur: the move from the single image to the class of image; the move from the material work of the print to the conceptual work of the algorithm; and the move from the static image to the dynamic image.

5. Embodied Experiences

Artists unanimously report valuing algorithmic models' ability to **generate fresh, novel, surprising outputs**––"images one couldn't have imagined" (Meyohas). As Ridler puts it:

> When I did the tulip project, the model made these weird, eerie, non-tulip tulips. It crossbred them to make things that didn't exist. A lot of it is kind of personal aesthetic choices as to when it becomes something that is interesting and when something becomes not interesting. That's a lot of kind of messing about and changing things, but I think there's this uncanniness to it that you can never program. You can never tell it to do the type of things that it will come out with. I think if you could, then it would all become a lot less interesting.

Documentation from *Myriad (Tulips)* (2018), the training set used to create *Mosaic Virus* (2019) by Anna Ridler

Working with ML models, however, involves balancing heightening surprise with the **frustration of having less control** over one's medium. As Barrat explains,

> A lot of my work is about misinterpretation, and that can be a double-edged sword. It's cool when the network misinterprets things and gives you surprising results, but it's also very frustrating when you cannot get the exact image that you want. I feel like I have a lot less control over the exact thing that I'm making. It's not like I'm a painter where I can just place a brushstroke wherever I please. I'm working with a larger system. If I had to describe working with algorithms, I'd have to say it's a split between frustration and nice surprise all the time.

Artists describe the need for patience and an openness to new technologies. Indeed, Sarin describes ML as a 'stubborn medium':

> Oftentimes, you know, working with GANs is like working with watercolour. It's a very stubborn medium that likes to do its own thing. GANs often do something that you don't expect. That's why I love them.

The ML process also requires **different modes of thinking**. For example, Sarin talked about the programming process as a meditative one:

> Being a programmer my whole life, I enjoy mechanical work. It's meditation. I enjoy every piece because I like the process, and I anticipate how it will look at the end. I think that's part of the artistic process, or any kind of creative process. You need to be able to enjoy the boring part or mechanical part.

Others, like Ridler and Crespo, talked about how using ML models fosters analytic thinking and 'algorithmic ways of working'. This was felt to be quite unique to generative art. As Crespo reflects:

> I was interested in exploring different mechanisms that automate certain tasks, or the algorithms have this automation thing to it. It just feels like working with generative art. It's different than, well, non-generative art. It really changed my way of thinking about when I'm creating.

Further, the **interplay between the cognitive and physical work** involved in the ML process has a particular structure. Ridler described this as "like feast and famine":

> The thing that I find most difficult, and the thing that did change when I included machine learning in the process, is the fact that it's like feast and famine with the different parts of it. There's lots of analogue work, like photographing or drawing, and then lots of work coding.

Ridler found this a different cognitive process to past experiences of using technology to create where the coding and analogue work were more integrated. The process of building the dataset and making the model encouraged her to flip back and forth "between two very different mindsets":

> Doing the dataset is quite a lot of manual work. You've got a lot of space to think because you tend to be doing quite repetitive things, like taking the same type of photograph. It's also physically exhausting that many photographs, thousands. Your body cramps up. Then when you do the coding part, it's like the total opposite. It's very sedentary, it's high cognitive load. The total opposite. So you have these two states of working. I flip a switch halfway through my work, going from one state to another state, and it's always really difficult.

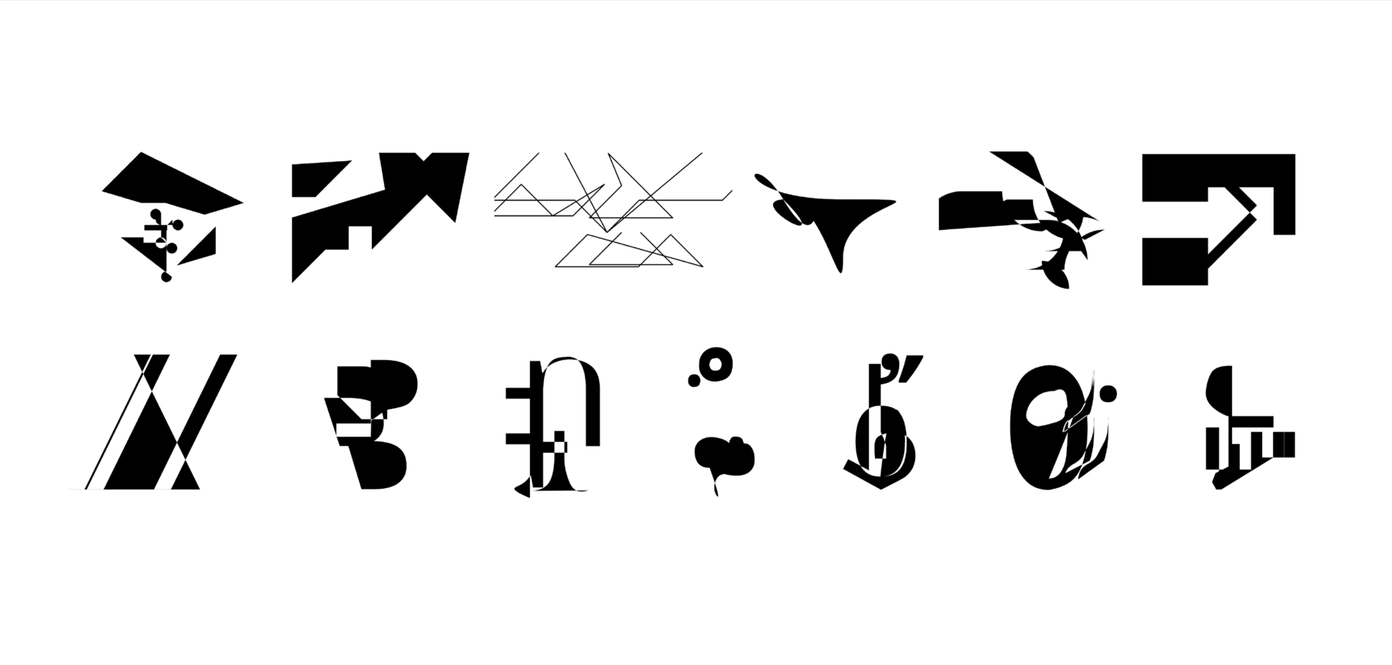

The **physical demands of working with ML** are often **overlooked**. Ridler's body 'cramped up' when taking the photos necessary for the data set. Sarin spoke about 'visual tiredness' after sorting through endless outputs. Respondents sometimes thought about these physical demands within the context of the contrast between the finite resources of the human and the infinite resources of the machine. Sarin talked about how looking at the endless outputs is 'physically draining.' McCarthy highlighted this **mismatch in 'embodied' resources** by noting that 'I get tired, the model doesn't':

> My piece '24h Host' was built off this utopian idea that I could build a software system to guide me through this party for a never-ending amount of time. And then the reality is that, as a human I break down. But also, just trying to build a system that actually does that effectively is quite difficult. And requires a lot of iteration. So I found that, despite the tools that I had built, in the end I was kind of left with myself. And I had to ultimately be the one that would translate the tool into something meaningful in terms of experience.

Far from being a 'tool of convenience', as much instrumental rhetoric around AI Art may suggest, artists sometimes found **ML a stressful landscape to navigate**. Elwes expresses feeling 'lost and suffocated' by the breadth of the technical landscape of ML:

> This is the problem with being a media artist, or more of a conceptual artist, as opposed to a painter. There you have a canvas, a physical limitation or constraint within which you can just experiment. You can do a bad painting. It's quite easy and fast to do a bad painting, and that's fine, you learn from that. I can't just make a bad painting. First, I have to have the concepts, and then do a lot of research. I'm slightly envious of my friends working with more analogue materials who can just play in their studio. I'm post-studio, apparently–– that's what it's called now.

The variety of embodied encounters were often challenging yet also enlightening, enabling artists to become more aware of data science and the important societal questions it raised.

Vectoglyph (from shapes and fonts) (2019) by Nicolas Boillot

Human and Machine Creative Intelligence

Recent successes in machine learning research have revived debates around machine intelligence, machine creativity, and machine autonomy. Providing new perspectives on claims about machine learning having crossed a threshold, artists share their perceptions of the scope of their ML models' intelligence and creative agency in the artmaking process.

1. "Machine Intelligence"

The strong consensus among our respondents was that their medium is an "**automation tool**". Sarin, for instance, is unequivocal about her models being tools—–even their valued capacity for surprise is often a limited one: "I expect it to do something that I will be surprised by, but in many cases it's more like my personal filters on Instagram".

This is not to say that the medium does not open up a space for new things. As Crespo says, "I'm still learning from it". Reversing the usual direction of the analogy, Crespo reflects that she sees herself as "working algorithmically" when doing analogue collages:

> I expect it to do something that I will be surprised by, but in many cases it's more like my personal filters on Instagram, or something like that.

This is not to say that the medium does not open up a space for new things. As Crespo says, "I'm still learning from it". Reversing the usual direction of the analogy, Crespo reflects that she sees herself as "working algorithmically" when doing analogue collages:

> I feel like they are statistical methods, right? You can explore data, extract information from it, and I see it as automated creation, which is really interesting. It doesn't mean that I don't learn from it. I really learn a lot of things, and it doesn't mean that we are still learning from it, that we are discovering things from it.

Our respondents, further, **question whether models are intelligent** in a meaningful sense—–and **suggest perception** may be a more **useful metaphor** to conceptualise machine capability. Boillot first questions the association of models with intelligence, suggesting that the scope of machine capability fails to qualify: "I'm not sure machine intelligence is intelligence. It's mostly training—it's not general." He continues: "For me, now, it's a tool. It's a tool, you can train a tool, and it will do new tricks, and sometimes tricks are good. Sometimes you can sell the tricks at Sotheby's." Ridler also highlights continuity, proposing that ML is "a tool that people are using in their practice that fits into existing strands of art history". However, as she notes, "it is understanding the world in a fundamentally different way to how a human understands the world. It's halfway between man and machine".

As Zylinska explains, the competition between the metaphors of intelligence and perception is a live issue in AI research and adjacent fields:

> The notion of artificial intelligence is still very much a cognitive one. I'm thinking about the role of perception here: not just human perception, but also how other organisms perceive. The research in machine vision still uses human vision as a model. It's not perception in an expanded sense, but human vision, seeing how can we get machines to see what humans see. At the end of the day, the goal is to get machines to see like humans do. The idea of the object which they are supposed to see is already predefined in the research.

Nake concurs, framing the perspectival differences between humans and computers in terms of the contrast between globality and locality:

> The computer is the machine of locality. We are the beings of globality. We recognise. If you look here, you immediately, in no time, take in everything. I could continue for hours giving you examples of how crazy it is to think for a second of intelligence in the computer. They are local machines that are fantastic in the cycles they run through, repeating what they are supposed to repeat. If we make a tiny change, they immediately run it. But they are very bad at pattern recognition. We, ourselves, we just look.

In addition to the limitations of the metaphor of intelligence for machine capability, respondents highlight that **the sense of whether or not a model is intelligent** heavily **depends on both context and technical understanding**. Indeed, McCarthy points out that perceptions of machine intelligence differ with prior expectations, a trend in the reception of her installation work:

> I think I realized that the expectations of machine intelligence vary wildly. People's experience of these works was very much shaped by their expectations and by their understanding of what AI could or should do.

She also recounts getting people to reflect on human/machine relations by confronting them to substitution in her 2017 piece LAUREN:

> One of the central ideas of the project was, how does it feel to swap in a human? And does that help us recontextualize what some of these systems are doing? And I think for most people that participated, that shift had a big effect on them. It made them reconsider issues of privacy, as we aren't really built to have an intuition for what it means to be giving up this much data or privacy. But, once you put a human in, instantly it's much easier to conceptualise.

Still from *LAUREN* (2019)
by Lauren Lee McCarthy

Overall, however, McCarthy distinguishes between technical and philosophical differences in intelligence:

> I think it comes down to control and knowledge. I don't personally think that there's a large difference philosophically between machine learning versus other types of algorithmic tools, but I recognize that technically, there's some things that are unlocked.

Crespo agrees:

> What's the difference between doing a creative coding project and working withan artificially intelligent algorithm to generate art? Both involve me sitting on my machine trying to understand something using code repositories. You're always stealing codes that somebody else wrote, remixing it, making it do something else. It's part of the practice. Why does one sound so charming, so mysterious, and the other so normal?

Importantly, we found that **perception of agency** on the part of the model **diminished** with artists' **technical understanding.** As Crespo explains,

> When I was doing Trauma Doll, I didn't feel like a person doing it. I felt like it was an algorithmic way to do it. I had an automated process, a very organized method for doing these collages. I began to think, "How can I automate this more and more?" Then the difference between an algorithm doing it or a person doing it became really interesting to me. When I started Neural Zoo, it felt like that had a collaborative aspect to it in the beginning, but then I realised that it was just the magic of seeing something that I wasn't expecting. Eventually that went away, as I began to better understand what was happening with the model, and learning more about the maths. Then I realised that it wasn't that collaborative. It's just like creative coding. I began asking myself, "Why is machine learning 'collaborative'?" If you're using Processing or OpenFrameworks, nobody talks about collaboration.

Pointing to dataset bias, Young flips the narrative of logical algorithmic intelligence:

> We think that it's logical and rational, but there's so much underneath it that remains enigmatic and mysterious. And so maybe even that in itself is interesting for people, to think AI isn't this sort of hyper-rational, hyper-logical thing. It begins to help us realize how much bias is fed into these systems, which creates systems that are biased and illogical.

rare-earth flower aphid from Artificial Remnants 2.1 (2020) by Entangled Others / Sofia Crespo

Lastly, our respondents highlight the fact that, together with the challenge of technicity, **interpretability problems keep the process feeling mysterious** and open. Illustratively, Elwes was in two minds about machine intelligence. On the one hand, "like Mario often says, you wouldn't say that a piano is making the art. It's just a tool. It's like the paintbrush, it's just a tool." On the other: "at the same time, you like to think about where this could go and what sort of philosophical questions it can pose". Suggesting a relation between the perception of intelligence and the lack of interpretability, they continue:

> Right now, we are teaching computers to see, and then, for artists, reverse engineering that to find hacks and poetry in the way that we are teaching computers to see the world. I just think there is something a bit beyond there than just using a random system, although it might just be how complex the algorithm's becoming. It's no longer interpretable. You can open up a neural network, you don't know how it's getting to those weights. Maybe that is the key difference.

Surprise is conditioned on our lack of knowledge about how models work: if we understood the optimisation space, we would know exactly what models would generate––but we don't. As Young summarises,

> Within these networks, there are these layers and weights, and the notion that these connections correspond to the way the model recognises the world. It's this massively dimensional space. Some of the dimensions we can sort of understand, like is this more blue or more green, is this a picture of a person wearing glasses or with facial hair. But there are other dimensions alongside which the model learns about the subject matter, which we don't understand. We can try to move along that dimension but it's indecipherable to us. So there is a strange magic of how this network has been constructed, and it may not be ever something that we can open up.

2. "Machine Creativity"

When it comes to creativity, our respondents' accounts highlight a **difference in scope between human and machine capability**—a sense of different creative strengths. As outlined in previous sections, the generative capability of models is highly valued. As Meyohas describes:

> In the process, the model was entirely more creative than a human. It created images of petals. I can't create those on Photoshop. Even the Markhov chain-created phrases I couldn't have composed on some piano. So, in terms of process, it's way more creative.

However, though models have the ability to generate surprising outputs, they entirely **lack the capacity to anchor these outputs in the world**. Unlike the kind of creativity valued in humans both in and beyond the arts, ML has **little scope for contextualisation**. This, as Meyohas explains, can be turned into an artistic advantage: "the great thing about those piano phrases is that a lot of them can't be played by human hands. It's helping me create something that no human would actually be able to play. That's the exciting part!" But there are limits to this move. As Meyohas continues, there still needs to be a "human creator to contextualise the present moment":

> When I view creating an artwork, there needs to be soul in an artwork. This is different than design or other things. Artworks have to touch your soul, touch you emotionally as well as visually and intellectually. For that, AI runs into trouble. It doesn't know yet how to create things that have friction with the real world. Cloud of Petals happened at the space at Bell Labs. Petals, the stock performance––these things have friction because they're coming from the real world. This is difficult to do with just a beautiful graphic that the model is creating. The pictures of the petals are beautiful, and the interpolation is beautiful, but it's only made more meaningful by knowing where they come from. That's where I think you still need a human creator to contextualize and understand the present moment, because art is created for people in a specific cultural moment and that's really difficult to understand.

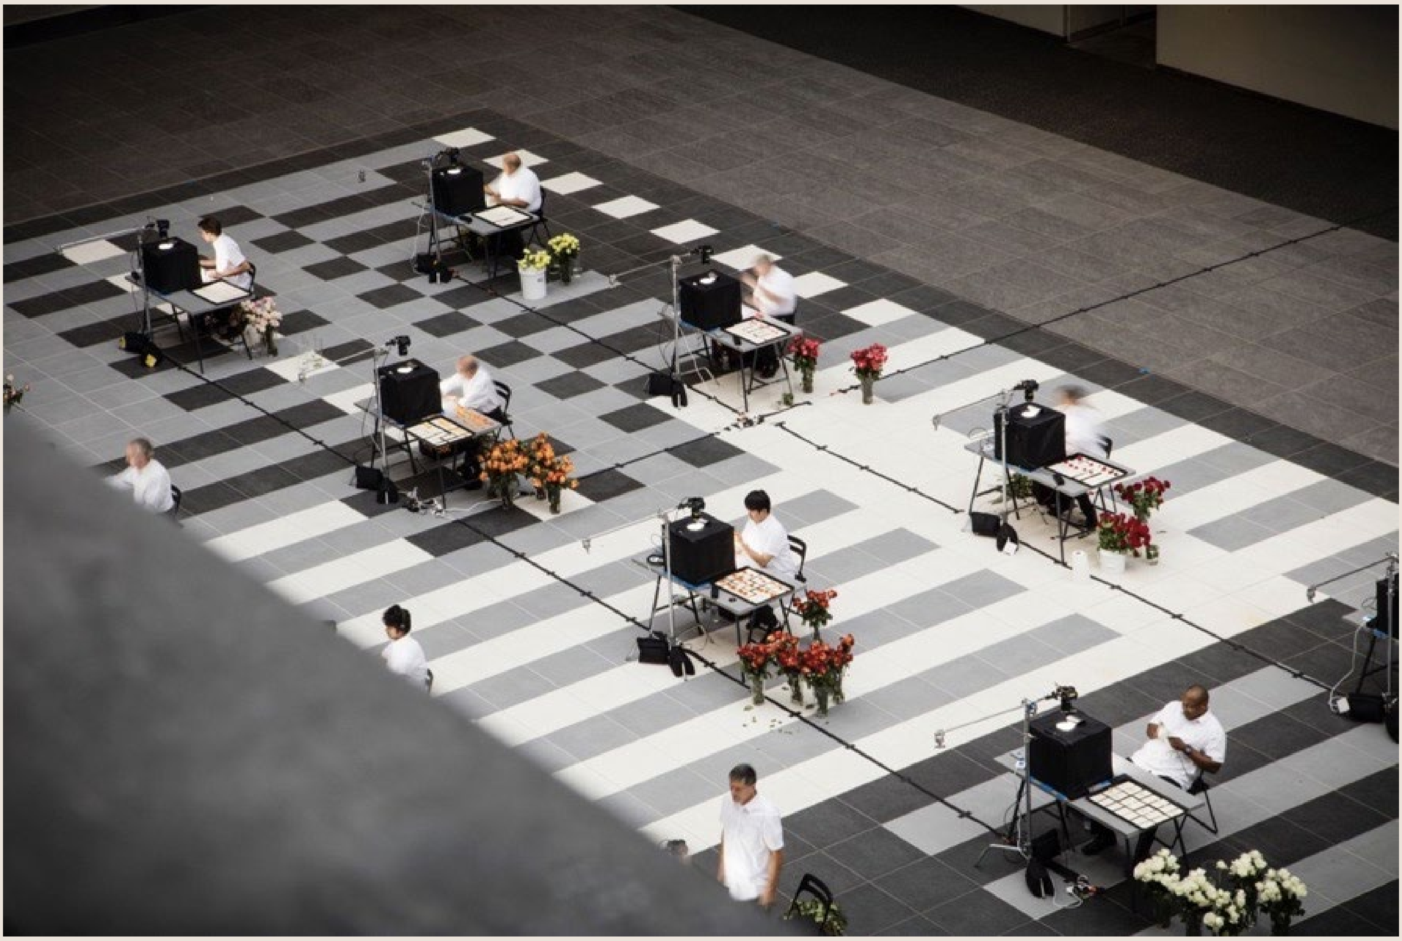

Elwes concurs. As they explain, the fine arts are a particularly challenging field for automation for this very reason:

> I do think taking fine art, especially as a case study, is interesting because it's the furthest out-there thing, perhaps aside from music. But you know, fine art, beyond painting, looking at modernism and post modernism, and conceptualism, it is such a complex thing. You can't even really envisage how a machine could even start to do something like put a urinal in a museum.

Photograph from *Cloud of Petals* (2017) by Sarah Meyohas

In addition to the challenge of context, respondents also point out the **failure of ML models to clear the intention criteria of artistic creativity**. Models may feel creative, they explain, but this is in appearance only: they are not making any choices. As Barrat explains:

> The network is really learning a very technical task. It's trying to pick up a set of rules from the data set that will allow it to generate images that look like images from the data set. There's really nothing artistic in what the network is doing. It's not trying to make pleasing images. It's not trying to make beautiful images. It's just trying to make images that reflect its understanding of the data set. I'm looking for interesting images, pleasing images, ways to mess the network up a little bit, make it misinterpret. I don't think that the network is learning anything creative at all. It's not really making any choices. It's just really assembling images. All of the creative choice to be on my end.

Indeed, as Elwes expands, the perception of creativity––rather than meaningfully creative activity itself––is what keeps the conversation about ML models open.

> That whole idea of creativity as being able to make mistakes, and then learn from those mistakes, and branch off and do something different and new, is by some standards a definition of creativity. In one of my older works, the model was constantly branching off itself, and misinterpreting itself and making sorts of quite unusual, erratic connections. It felt to me like I was watching these creative machines, but I also don't want to overstate that, because I don't want to add to the mystification. These things are stupid. They're limited by the datasets. They're not really doing anything special. But it is still provoking to ask, "what is our creativity?"

In this sense, **creativity is an easier target than art:**

> I believe that these algorithms are and can be creative. I think I believe that, you know, AlphaGo, does make creative moves that make us question what creativity is as an ability. But there is a big distinction between that and making art, or art that is interesting or valid. That requires a lot of intentionality.

And, further, as Elliott explains, **machine capability is of a much more limited scope that human artists' practices:**

> Art is just so broad these days. If you went to a contemporary art degree course and you told people to create an artwork based on a human face, you'll probably get anything from a performance to a sculpture to a painting, or documentation of some intervention somebody did. There could be a very broad spectrum of what humans could do when asked to create a work based on a human face. But if you currently asked a machine to do that then you'd first need to figure out which strand you want the machine to go down. Do you want it to create a work that's based on facial recognition or do you want it to use sculpture designs, or take a photograph and then make a painting like with DeepDream or StyleTransfer? A lot of this sort of creativity is just quite narrow.

What becomes clear overall is that **the artistically valued kind of creativity happens when artists understand the models.** As Ridler explains, once an artist has a fine grain understanding of the models, they can make creative choices:

> When you understand the way that different models behave and the way that different architectures behave, you can really start to pull out, through the way that these different architectures behave, stuff around collapse, stuff around memory, stuff around decay, stuff around compression and memory. Each of these different models will have different associations and contexts that you can then start to play with. Then you get the output and have these associations displayed in how you're choosing to emphasize different elements of the project, and what it's about.

Mapping the process is the first step to understand the difference between "what can be generated by anyone, what is genuinely novel, what is genuinely craft or skill by an artist in using these technologies, versus someone taking a model that's already on the internet, a data set that is very readily available" (Ridler). In so doing, this mapping highlights "where human agency is at different parts of the process" (Ridler).

This question of creative choice maps onto the difference between an artwork and a tech demo, the unavoidable art-critical question facing every new wave of technology-based art. As Sarin explains:

> 'Isn't it just pushing a button?' is the question I get every single time. The perception that you just push the button, like in Photoshop, right? You spend ten minutes on your image, then you push the button and publish it on Facebook or print it. That's people's main misunderstanding. But what does AI art-making involve, with GANS in particular? What does it actually mean? It's my goal to educate people about this part.

This is the landscape upon which the artworld norms are negotiated: a struggle with traditional artworld gatekeepers, stakeholders, and the public over aesthetic and practice conventions, as well as over the nature, meaning, message, and worth of new work.

As well as fielding iterations of the gatekeeping question facing any new development in the arts––"is it art or not?"; or in our case more specifically, "is it an artwork or is it a tech demo?" or "isn't it just pushing a button?"—ML artists are also always engaged in more precise self-theorising. Where is creative agency in this process? **What is the role of the artist?** As Ridler explains:

> There's a difference between computation and creativity and art practice. "AI art" is bandied about a lot, and I think a lot of those works are creative experiments, but they're not artworks. A lot of the stuff that is labelled "AI art" on Twitter or Instagram are tech demos, in a way. They're just taking technology and making something pretty. I think that for it to be an artwork, there does need to be a critical context––where you're considering what each of the parts of the material are doing and how they're doing it, and what that can bring out. The interesting GAN-generated work is trying to do that. It's treating GANs as a material. It has a message, a central concept that is being explored through the way that the GAN is working. Creativity and authorship and control are being considered, as opposed to just downloading something and playing with it without trying to push those ideas.

This necessarily remains an open question. As Elwes asks, "then where does the art lie? I don't have an answer for this. The whole system is what I think is interesting to investigate".

Practices in Context

Striking continuities exist between ML art and code-based art since the 1960s. Highlighting important precursors, we then turn to the future of ML art and consider its prospects.

1. Looking Back on a Long Tradition

Thinking about algorithmic techniques as forms of creative augmentation is grounded in a rich tradition. In the 1960s and 1970s, the pioneers of computer art (Frieder Nake, Vera Molnar, and Manfred Mohr, among others) explicitly harnessed early computers to pitch chance against control in order to produce unexpected results––what Friede Nake later described as drawing "with eyes wide shut". In this long tradition of using artificial constraints to productively narrow the space of possible outputs––and thus heighten the creativity necessary to do something interesting––**the role of machine learning techniques is up for debate.**

There was some **disagreement among our respondents** as to whether machine learning represents a paradigm shift from older code-based art. For example, Crespo voted no, Ridler voted yes; overall, most agreed that everybody is still, as Luba Elliott puts it, "making art with computers". As Bailey recounts from his curatorial work:

> We recently did this show about generative art, and I made the decision to include AI art under this umbrella. Some of the AI artists disagreed, but I'm comfortable with lumping them. It's all code-based art. The objections were around a lack of complete control in AI art. But the truth is, generative artists don't have control over everything either, even though the code executes the same way every time. There's a lot of trial and error and accidents in that work. At the moment, AI art is a similar process.

As well as looking back within the algorithmic tradition, artists also make wider, more **lateral connections**. Indeed, the **analogies** used by artists to describe their work with ML range well beyond computer art and illustrate the wide range of ways in which artists relate to their practice. For instance, Ridler finds reference for her practice in the land and environmental arts:

> I rarely think of myself as being in that lineage of algorithmic artists. I feel more closely aligned with land and environmental artists. In that tradition, artists will have an idea of what they want, build it, and then allow elements that they can't control, like the weather or dust or the desert, to act on what they created. It's that same tension, having a vision of what you want and then being open to something else acting on the things that you spent quite a lot of time setting up. It's the same way I spend all this time and effort constructing my dataset, and then allow something that I can never really control or predict to act on that. Sometimes, the label algorithmic art boxes the work in. Thinking about it in this wider way helps me think about it as a longer process.

Sarin, on the other hand, finds parallels for GANs' textured aesthetics in printmaking:

> I don't find the comparison between photography and GAN outputs useful because we are still in this textured regime. A much more useful analogy for me is printmaking. Printmaking is a perfect metaphor. You have some idea of what will happen, but then the medium will have its own way, images will come differently because of varying amounts of pressure. This is where I see GAN outputs as close to printmaking.

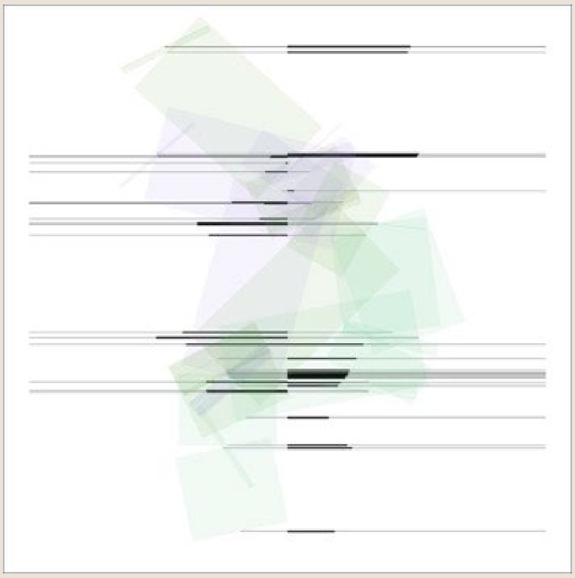

Crespo, for her part, is "really inspired by scientists", while Young looks back to the Hudson River School of painting and the long traditions of artistic work on nature. These comparisons bring out the balance of similarity and difference with other methods and with the past. They demonstrate that **working with ML is not a disavowal of analogue methods, but a dialogue with and through them.**

Looking back, though there are intra-tradition differences, Bailey argues that overall, ML art is made of the "**same flour, same eggs**" as older traditions in the arts, from algorithmic arts to classical painting. Drawing on his experience going from analogue to digital techniques, and then onto learning to code, he explains:

> When I swapped from painting and drawing to Photoshop or Illustrator or 3D software, there was a richness that was missing. I felt like when I was working an analogue medium, I was baking a cake from scratch, with flour and eggs and water and that level of control. And when I used tools like Photoshop or Illustrator, it felt like a store-bought cake. I didn't have that level of granularity. Then, once I learned how to program, I was back to working with ingredients, and all the accidents and exploration and nuance of analogue studio practice came back. When you can work with first principles, accessing the code and the hardware at that higher level, it opens up a much broader range and vocabulary for things you can do.

The apparent differences between analogue and digital, digital and generative, are erased through experience. "It's so hard for people who have done one and not the other, or neither", Bailey explains, to understand that "programming art codes is almost exactly like painting: same eggs, same flour".

From the *Algorithmic Modulations* series (2019) by Manfred Mohr

2. Looking Ahead to the Future of ML Art

Looking ahead, it seems the **jury's still out on ML art**. As Elliott explains, it is difficult to tell whether we are at **the height of an AI-art trend** or transitioning into something more lasting:

> There's been a lot of hype and overinflated interest in AI art, it's very much the height of the AI art summer. But I don't know if it will continue as its own movement. I expect some of the art, particularly generative art, to become part of the lineage and narrative of the computer arts movement. And some of the art will be absorbed into fine or contemporary art or media art, where AI artists will compete with all the other artists who are painting landscapes or creating critical works that look at society, and there the technique will become a bit more secondary than it is now.

Bailey, for his part, is more optimistic. To him, generative art is "**the most important artwork of our generation**", a currently undervalued movement which will come to define the early 21st century. Pushing back against what he calls "a massive bias against digital art and digital artists" in the art world, he explains:

> Not only is digital art important, it's the most important artwork of our generation. I think of generative art in particular as the history of our generation made visible. To me, everybody's missing the bus on that, which is also a good signal, because historically, we do a very lousy job of celebrating the most important art as it's happening. It's usually after artists die and we've moved on that we figure out what actually mattered. For me, it's clear that digital art is the art of our generation. It's undervalued now, but I don't think that'll last.

Practically, Bailey anticipates that the future of AI art will entail an expansion of the range of techniques used by generative artists:

> AI means a lot more than GANs. Machine learning means a lot more than GANs. We're seeing a lot of work based on GANs right now. I think that may play itself out, especially as that technology gets democratised. It'll wear on us. But that doesn't mean we'll have exhausted what we can do with machine learning and AI. The same folks who were the first to tap into this technology are going to be the first to expand the space, because they're technical enough to be able to do that. I think we'll see an expansion. We'll see artists using tech and AI in more interesting ways moving forward.

Crucially, this future will belong to **those who have both technical and artistic skills:**

> The same principles that made good artists good, regardless their medium, make artists good when it comes to AI. It's not just tech races and first-ism. The artists who will create this new canon are the ones that understand how to code, how to produce GANs, how to chain them together and use them in new ways. These artists are using these new tools and technologies to expand what we can do in the art-making process. Those will be the people who will help form the canon going forward.

Is ML then revolutionising current practice? Meyohas proposes that while **ML does not, in itself, unlock previously unreached levels of creativity**, using ML technology **opens an exciting space for exploration** and experimentation:

> I feel very motivated by discovering a tool thatdoes something, even if that thing is very mechanical. Even when it's a very mechanical tool, I'm thinking, "What could this do? Could I use this to create a new image or a new video?" It's the motivation to explore some new land. It feels like we're walking on a new territory. It doesn't mean that if you remove machine learning from my life, I wouldn't have any creativity. But playing with new technologies has this feel to it, like walking onto new territory.

McCarthy concurs. She explains that **while the technical affordances of ML technology are a step above previous tools, the relationship between artists and their media remains unchanged:**

> Every tool that we use really augments our thought and our creative process, AI or not. And part of what it means to be artistically exploring a medium is to not just see what the affordances of that medium are, but explore what it actually brings up for you. How does it change the way you think or see the world? What's exciting about AI is that it's a tool that is quite new. It's not just one step away from tools of the past. There's a different paradigm there. I don't see it as a significant shift from how artists worked with their tools and technologies in the past, but I do think that the tool itself is a significant technological advancement. It's opening up a lot of things, and I'm excited to see what that brings up for people in their individual practices.

Boillot concurs, proposing that ML will, instead of revolutionising creativity, simply become one of its modes, and may get old "like everything else". Barrat also aligns with peers in questioning whether the creative affordances of ML are specific to ML itself:

> This idea that AI brings new forms of creativity that weren't possible before sounds really interesting. But is this just because we're using a new tool? Is it necessarily because the tool has something to do with AI? You could probably say the same thing about advancements in pigments, a hundred years ago. If some advancement made pigments more accessible to more people, you could say that that was enabling creativity that wouldn't have been possible before, because now a wider amount of people can paint. So I question whether that kind of progress is unique to AI.

Elwes goes further to say that the proposition for machine creativity relies on very "low-level" thinking around what creativity is in the arts:

> Algorithms have become very creative design solutions. But I'm from quite a strict fine art background, where it's very much all about the intentionality and the message, and design and craft are secondary. From that point of view, the question becomes whether the algorithm can take the role of the human and do something truly innovative, or with intentionality, rather than something that's just mimicry. Engineers in this field like to say that their algorithms can be creative. But they're thinking about it on a very low level. They're saying that the algorithm can make a Van Gogh-style painting, therefore it's making art. But that's just mimicry, which is the death of art.

Further, Crespo highlights the recurring problems caused by public perception of machine autonomy:

> Art is still going to be around human expression and human emotions. That's the way that art works. I had a few people tell me that my work isn't really art. That it's made by an algorithm, and not by me, and therefore it's not art. But there's a real human doing this data stuff. It's me. I'm a real person feeling something. Why is that less art? In a way, I think having to ask ourselves these kinds of questions means that Ai is somehow changing the field.

Ridler, for her part, prefers the kind of ML art which isn't centrally about ML itself, but about how ML is used to explore critical questions:

> For me, it's when art-making splinters off into these niche fields that it becomes really exciting. When people are working with technology from a different angle, it becomes less about the technical things that can or can't be done and more about how it's being used and explored from a very strong critical standpoint. I'm more drawn to those types of projects.

Further, the more important impact of ML on creative work may lie outside of image generation and around tasks relating to curation and judgment. As Ridler continues:

> In a way, the GAN-generated imagery and the machine-generated music are a bit of a red herring when it comes to the impact that AI will have on the artworld and on creative fields. The biggest impact that AI has had in the creative field is not being able to produce DeepFakes, or anything like that. It's the Netflix recommendation algorithm. The way that choice is being curated through algorithms is having an impact on the type of media that people are consuming, and also how people are approaching curation and deciding what art is good.

Finally, the hope for the field is a **human-centred one**––focused on **artistic and cultural rather than technical questions.** As Barrat says:

> I really am still hopeful about AI art.
> Even though I'm not all on board with this idea that artists teaming up with AI will unlock new creative potentials, I'm hopeful that interesting work will happen. Work that will help us figure out where exactly AI fits in in the field of art right now, as an artist's tool, and what new technologies can offer. I'm more interested in those questions than in technical questions, and I think that they're important.

In the end, our respondents all agree, **it's about them as artists**––about artistic questions over technical ones (Barrat), about "hacks and poetic ways of using algorithms" (Elwes), and "not just about affordances, but about what it brings up for people" (McCarthy).

Neural Zoo ((analogs)) (2019-2020) by Sofia Crespo lenges

Chal

Challenges

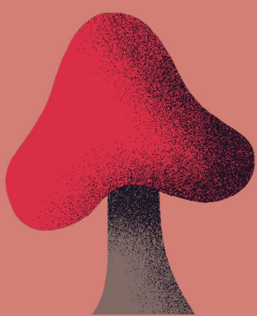

Key challenges emerge from the adoption of ML-based practices. These challenges include skilling up, issues around resources and finances, the absence of a shared language, barriers of entry to the field, and the environmental impact of machine learning.

1. Skilling Up

Throughout the interviews, artists reflected upon the array of technical and social skills needed to work within this domain. As Bailey summarised:

> It's one thing to understand the technology, and it's another thing to have a very nuanced and sensitive mind to sort out the way things are working in the world socially, and to be able to communicate this. Technical skill, communication skills, and then sort of a third skill that I don't fully know what to call it, but it's sort of this ability to appropriately synthesize and see into the future and look at what's happening now in a way that's artistically nuanced. When you put those three things together, I think you get the Rubens of AI art.

Artists described the need to **continuously upskill** to be able to use the latest ML models and keep up to date with developments in machine learning research. Related issues of hardware and software obsolescence, especially for installations, were also raised as common problems.

Artists also talked of the sometimes surprisingly **technical limitations they had encountered in using AI**. As Young reflected:

> I've learned that AI is really slow. We think of these models as super-fast, buttraining is slow. To train on a batch of 50 images can take three or four days. And not only is it slow, but you can't get very high-res images. So the images are surprisingly small, with a lot of challenges in making them large enough to be interesting to print. That was an interesting insight that I got fairly early on.

The need to **translate knowledge and practice across technical and artistic communities is also important**—and is an issue well recognised by artists. For Elliott, this process of 'building bridges' formed a central part of their work, primarily through talks and workshops within and across different technical and art communities. As Bailey described:

> In my art training, visual communication and verbal and written communication were really my strengths. And because I'm used to being around technical people, I don't get intimidated. I just keep asking questions until I can break it down and understand it and explain it to other people.

Due to the multiple processes ML art can involve, skill-building goes beyond the individual skill-set and often requires **collaboration**. As Meyohas explained:

> As a creator, my skillset doesn't lie in any one thing. I am not animating the birds, I worked with an animator. I am not figuring out which speakers we are using, I have a sound engineer who helps me with that. I play piano quite well but I'm not a composer. I'm not the specialist at any one thing.

Alongside these communication skills, artists were also required to be aware of the wider social context. This included the challenges of positioning themselves vis-à-vis aesthetic and critical developments, and the difficulties related to communicating their artistic identity in a hybrid field.

2. Finding a Common Language

As noted above, **language plays an important role in perceptions and misperceptions of machine intelligence**. In the interviews, artists stressed the need to be **careful about the language** used when talking about AI Art.

As McCarthy noted, these issues are compounded by the mystical and often anthropomorphising language typically used in discussions about AI:

> There's something that captures people's imagination because ML feels more mysterious, even though it's not that much different than other technical ways of solving problems. It feels a little bit more unknown, and I think the languages we use to talk about it––like artificial intelligence, or this algorithm that's learning, or headlines that say Google's new algorithm developed its own language to talk to itself–– evokes such a sci-fi feeling, so that's one of the reasons people get captivated by it.

For Crespo, some of the mythical status attached to AI is due to a lack of understanding:

> Why is machine learning so over-hyped? What is 'artificially intelligent'? What was perceived as artificial intelligence some years ago isn't what we perceive now. The technology's become really natural for us. When it becomes normal, then it doesn't feel magical anymore. It's this lack of understanding of what's happening with the technology that makes us feel 'wow'. Then it must be magic, or it must be a lie, because it looks alive.

The artists we spoke to often tried to **subvert the hyperbolic narratives of algorithmic agency**, and found ways of resisting the discourse. For example, Young tried to get away from the problem of language through using emotion as a metaphor for a different way of talking about uncertainty and capability:

> That's why I play with the idea of emotions. I don't think the model is really intelligent and I don't think there are really emotions in there, but it's language that helps us think about it in a way that's different from traditional computer science. Most computer science is a top-down process of encoding rules. AI is a much more bottom-up, organic kind of learning. We don't know what's happening inside of AI models, so maybe emotions are a good way to at least break away from the idea that they're so logical.

Edmond de Belamy (2018) by collective Obvious

Overall, artists considered **mystification and anthropomorphising** language to be irresponsible. They stressed that the wider group of actors engaged in AI art sometimes deliberately overplayed a model's contribution to an artwork for creative or commercial reasons—although there is a clear tension between these aims, and a responsibility to be clear in communicating what AI can actually do. As Elwes argued:

> There are conflicts and contradictions in this as well. One part of me wants to say, "Look at this crazy thing, it's making us think about our own consciousness. It's doing this incredible creative thing." And explain it to people that way. But I've realized that's irresponsible. The more I learned about it, the more I noticed artists who are more in the public eye, who don't fully understand the technology, writing curatorial statements and texts saying 'this is artificial consciousness', and mystifying it on purpose for the art audience.

As Barrat expanded:

> The narrative that GANs are creative or whatever is gaining in popularity. It's a strong thing to market. We saw this at Christie's. Christie's auctioned off that Obvious piece, and the whole narrative attached to it was "a robot made this" or "a computer made this". They're not looking for a dialogue between a fashion designer and an artist that works with algorithms. They're looking for a narrative of "computers can design clothes on their own". I don't know why you would want that. It's a lot less interesting than the real narrative of creatives using AI or machine learning to produce their works. I don't know why people are so attached to that first narrative, and want it so badly.

For Ridler, the problems related more to the high knowledge barriers in the field, in that often curators and galleries focus on the AI, instead of the content and the interplay between artist and machine:

> I think it's partly because the movement is so young, but also I think that there isn't a high level of understanding as to how these things work. A lot of the time with curators and galleries, though this is slowly changing, there isn't a high level of knowledge as to the real nitty gritty of how we're working as artists, which means that it just becomes that someone will produce an AI show, rather than kind of think about the different themes, the different strands of working.

Poor use of language, deliberate or otherwise, was felt to also lead to a **lack of understanding of the role of the artist.** In fact, because of the expectation of agency on the part of ML models, these exacerbate the usual problems of tech-based art-making. As Barratt explains:

> In some collaborations, I just cannot get people to understand that I'm an artist, and that the neural network is not making art. That it's not making any creative choices. That's all on my end. Just the way that people talk about it, you can tell that they think it's an automatic system. But it's not. It might be because of the hype, or just the word "AI" and what we associate it with, but people really want to assume that AI is making choices, or learning, or self-aware.

Developing a language that appropriately reflects the role and activities of artist and machine, as well as a better understanding of AI itself, is thus needed.

3. Recovering Voices

As with wider discussions of the social implications of AI, issues of bias within datasets and the ways that algorithms could amplify this effect were raised as a concern by artists. They also stressed the 'missing voices' in AI art—not just who is (mis)represented within datasets, but also in the people working on them. As McCarthy commented:

> There's been so much great work about thinking about bias in AI, in terms of the data sets, but it's not just the data sets, it's who's actually working with them and on them. And it's not only thinking about race and ethnicity, but also thinking about financial or social privilege.

There was a feeling by some that in the field of ML Art, there was "always the same crowd". It was felt that part of the reason for the missing voices, was because the community membership was shaped by the high financial, technical, and social barriers of entry, which overlap with barriers of entry into technology-based fields. McCarthy explains:

> Something that's hard is that a lot of the tools have quite a high technological barrier, or at least are framed to feel that way. So often it might feel like I need to have studied this in a formal context to be able to understand. Or, if you're working with models, it could be actually quite expensive in terms of just computer processing units. I think there's a lot of people doing great work around how the algorithms themselves are biased. But I'm also interested in how the processes of working these systems are biased or missing voices.

Potential strategies to address these issues included technical developments that made machine learning more accessible and open, and a more open definition of what could be included as AI art. As McCarthy suggested:

> I hope that there are more tools that keep opening up machine learning as a technique that people can work with. I also hope that there's more embrace of creative and technical work that is not necessarily like building models. That could be offering important critiques and viewpoints on these systems without having to run the model themselves. Or without that being the only acceptable way to be working as an artist in that space.

Others talked about the need for better connections between tech developers and the art community, as a lack of conversation has critical social implications. Elwes commented:

> Having a conversation between these people is so important, because often, these people in the Bay Area are so wrapped up, and they're making so much money and they're able to do these disruptive technologies on their lunch break, but they don't really take a step back and actually meet someone from a humanities background who is considering these things they might not have even been aware of.

McCarthy suggests that AI art may be the place to facilitate these conversations:

> There's still a big disconnect between researchers talking about the societal implications of AI and researchers looking at technical possibilities. Sometimes these camps feel really separate. Not that they're at odds with each other necessarily, but they're not engaged in a conversation. I think that art is the space where you see a little bit more crossover between these two ways of approaching the subject. There's a lot of possibility but also a lot of hype, so the challenge will be finding ways to engage these different parties in meaningful conversation.

4. Resource and Financial Pressures

Artists reported experiencing interrelating resource and technical pressures that had an impact on their work, and on the field of AI art more broadly. Some of these pressures related to financial and practical issues such as the challenge of selling hardware-based installations—"ML-based installations are unsellable because you need to sell the computer", as Boillot explained—and the exclusionary costs of computing.

Others reflected the challenges of making AI Art in a changing technical landscape, where open source and 'off the shelf' models had both positive and negative impacts, both encouraging a democratisation of practice but also the devaluing of artworks under a push-the-button misunderstanding of ML art.

These challenges collided in the problem Bailey called "**first-ism**":

> When you have a tool that's hungry for data, we've already seen, people are racing around looking for large datasets. It's inevitable that some people are going to pick the same dataset and have similar output. And then it becomes a race. Who can get to the cool new datasets first, which is not necessarily in the spirit of good art-making. It's just who can get the most computing power and shove it in the newest dataset before anyone else can get there. It becomes more of a tech race than a nuanced examination of art-making. But maybe that's just part of the new thing.

As ML-based works become commercially valuable, **community values also shifted**. As Bailey noted:

> Lots of people shared their code and people learned from open-source practices. No one ever worried too much if someone took credit for their code, or just tweaked a few colours, and then spread it, because there was nothing at stake other than kudos within the community, because no one was buying it or selling it, or writing about it for that matter.

However, with the field's recent successes, Bailey saw a **shift towards less open source practices**. This may not only be a problem for the development of the field, but also for the **decisions about credit and authorship:**

> It's not like a painting with a single author. When someone makes generative art, or AI art, they're using algorithms that are written by other people. So it's not always clear where the credit giving stack should stop. Should Obvious have credited not only Robbie Barrat, but also Ian Goodfellow? And maybe also the person who made the graphics card that made it possible and the person who made the MacBook that they used? I know it can get absurd, but where do you stop? And depending on who you ask, it changes.

Relatedly there were challenges for **programming-based artists who may wish to not make their code so straightforwardly open** as in the past. As Bailey went on to argue:

> It's the nature of algorithm-to-code that a work that looks unique can quickly be replicated by lots of people. Then it doesn't feel unique anymore or special anymore, so you need to protect it more. We're democratizing the tools, which is forcing the specialized forefront artists to be a little bit more careful with whether or not they open source their code.

The democratisation of ML tools also opens up **legal issues related to the "radical remixing"** (Boillot) of potentially copyrighted training materials. As Bailey argued:

> These tools that were once only available to a small number of people are now going to be available to everyone, and, to get them to do anything, we're all going to need data. None of us are going to worry about copyright and legal nuance. We're just going to grab whatever we can find on the internet and put it in there.

Despite the challenges, Bailey and others remained optimistic about the future, as there was likely to be a **positive interplay between democratisation, first-ism, and good art**.

> The ability to understand new tools sometimes is where the ideas come from. So it'll be convenient to say, "Oh, well, this person's going at new tools. This person's good at ideas. And this person is good at artistic execution." But the three sort of drive each other to varying degrees. Everybody has a balance of skills, not just one. And no one group will win. I am excited to see what happens when a bunch of people with an art background and training can play with these new accessible tools. As you expand the audience, there's a higher likelihood that more interesting work will come along.

5. Opting In

Artists highlighted the challenges of working in a domain where the commercial sector prevails. While artists recognised the value of technical possibilities, they also had concerns about the links between industry and social issues.

A key theme was **privacy**. Artists keenly felt that privacy was the currency exchanged to access processing power and the affordances of AI. As Boillot commented, "you put everything on the cloud, you don't know what they do with it". The **environmental impact** of computing was also raised. As Boillot explained, "machine learning is very carbon-consuming. So I cheat a little by using Google Collab. You know about Google Collab? Google pollutes for you!" The conflict between individual flourishing through art and the behaviour of Big Tech companies was summed up by Barrat, who explained:

> Interesting things are going to happen between AI and art, depending on how people use it. But it's hard when you see things like the OpenAI language model, fake news, and what Facebook is doing. It's disheartening. I'm very conflicted between those two views about AI, where on the one hand it has a positive and radical impact in the arts, and on the other, seeing all these corporations doing pretty much exactly the opposite.

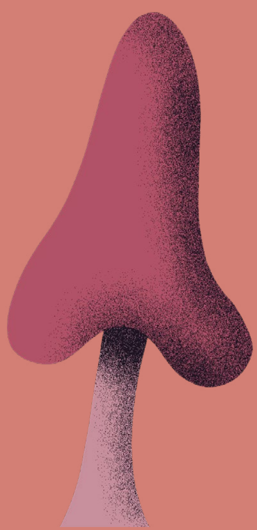

New Creatives?

Despite exciting technical developments, artists' creative agency remains the cornerstone of artistic work. To explore the affordances of machine learning in the arts and in creative work, we propose the terms 'statistical creativity' and 'human/machine complementarity' as productive alternatives to narratives of machine autonomy.

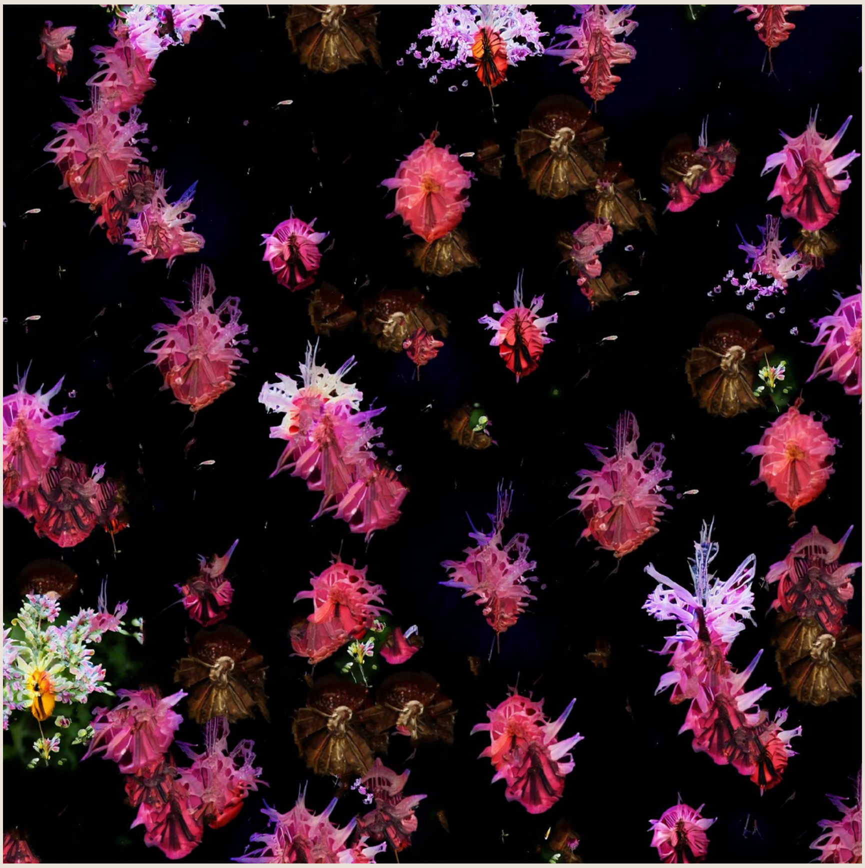

1. Statistical Creativity

Because of the particular limitations of creative algorithmic capabilities––namely the provision of context through the dataset; the difference between data and reality; and the reliance on statistical, not qualitative, proximity––we propose the more precise term statistical creativity for machine applications in creative domains. While statistical creativity, in itself, fails the bar of the full kind of creativity practitioners use in their work by the standards of their field, it can assist humans in their practice under structured conditions. Statistical creativity is meaningful not because of machine capability, but because human creativity occurs through encounter and using ML models provide such opportunities. New theories of creative complementarity must address the networked and interactional nature of the human creative process. While generative algorithms are bounded by human decisions, they may also expand creative cognition by sparking inspiration or by supporting the production of novel, surprising, and valuable outputs. Technology-augmented art-making may just be machine learning delivering on its full potential.

tiny agile swimmers (2019)
by Sofia Crespo

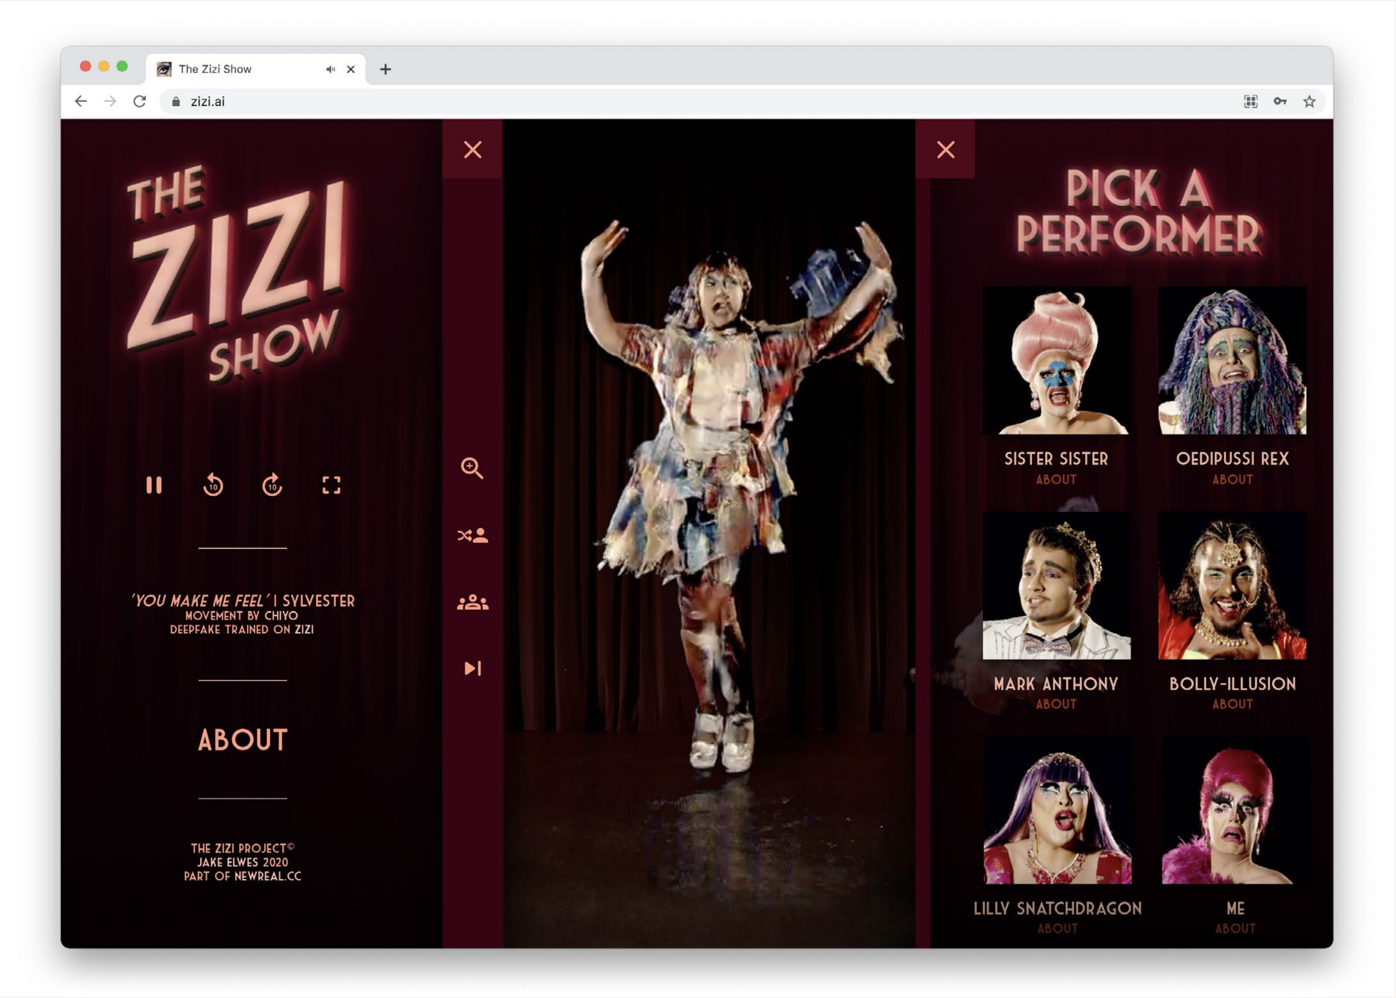

Still from *The Zizi Show* (2020)
by Jake Elwes

2. Continuum or Paradigm Shift

Uptake of machine learning tools in the arts do not, in themselves, unlock a step-change in human creativity, but instead build upon century-long trends in automation. Indeed, what we find is the progressive inclusion of the affordances of ML in the toolkit available to fine and media artists, accompanied by local changes in their creative process––the new activities around creating, curating, and cleaning data; as well as the particularities of selecting, post-processing, and exhibiting ML-based works.

New media or technologies in the art world integrate the mainstream in progressive waves. First, the new technology appears to threaten older media and artists themselves. Photography hailed the end of painting, video the end of photography, interactive NetArt the end of static artworks, and so on. The second wave is reflective: the new technology itself becomes the subject of the artwork, which reveals the technology's capabilities, limitations, and effects in the world. The third wave––which includes most of our respondents––looks for synergies. At this point, the new technology is no longer new, but has integrated into the artist's toolbox.

Photograph from *Cloud of Petals* (2017) by Sarah Meyohas

To create expected surprise has always been a strategy in art production––both in the algorithmic arts and in older traditions. However, **human agency is still needed to break the loop.** Humans are still fundamentally needed to generate outputs meaningful to other humans. **What we observe isn't automation, but complementarity**––not singularity (speculative) or self-sufficiency (untrue), but **ongoing conversation and embeddedness**, with a constant range of human interventions ensuring that works are not derivative, and that they "have friction in the real world" (Meyohas).

Image from *24h HOST* at
Brownie Project, Shanghai
(2021) by Lauren Lee McCarthy

Conclusion

In this new space of "ML art", we find that artists self-identify in a range of ways which do not always centre AI technology––although technical competence is a highly valued skill.

We find that new activities involved in the use of ML models involve both continuity with previous creative processes and rupture from the past: visible in the reorganisation of creative workflows around the generative process, conceptual shifts around the nature of ML outputs, and an evolution in artists' embodied experience of their practice.

When conceptualising human and machine creativity, we find that artists highlighted a difference in scope: while ML models could help produce surprising variations of existing images, the artist is irreplaceable in giving these images artistic context and intention. As artists highlighted, the creativity involved in art-making is about making creative choices––a practice beyond of the capabilities of current ML technology.

Reflecting on their ML-based practices, artists found many similarities with past periods in art history: the code-based and computer arts of the 1960s and 1970s and the harnessing of randomness by much experimental art. Many also found the generative capabilities of ML models to be a "step change" departure from past tools.

Ultimately, however, most agreed that despite the increased affordances of ML technologies, the relationship between artists and their media remains essentially unchanged––as artists ultimately work to address human, rather than technical, questions.

Although ML-based processes raises challenges around skills, a common language, resources, and inclusion, what is clear is that the future of ML arts will belong to those with both technical and artistic skills.

Human/ML complementarity in the arts is a thus rich and ongoing process through which artists refract technological capabilities to carry out their work. There is more to come.

Appendix: Methodology

1. Defining the Case Boundaries

The case study targets key players in the field of ML Art: the fine and media artists at the forefront of the exploration of human/machine creative complementarity.

Artists are a desirable subset of creative practitioners for several reasons. As a consequence of professional training and the necessities of producing accompanying material for exhibition spaces, catalogues, and the market, artists are particularly capable of articulating and reflecting upon their creative processes. As the production of a discourse on process is embedded in practice, artists are thus uniquely well-positioned to describe and reflect upon their interactions with intelligent systems, as well as being particularly sensitive—by position and necessity—to the different reactions that their peers, the public, and the markets may have to their work. They are thus well-placed to confirm or put pressure on the creative promises made on behalf of algorithmic techniques by different spheres.

Further, artists working with generative algorithms are situated at an additional intersection: the overlap between the field of the arts and that of engineering and industry, and thus are particularly able to reflect upon the dynamics of these very different spaces, and bring to light any overlaps and inconsistencies in discourse and practice.

Given the complexity of the ML Art field, and the lack of existing qualitative work exploring this phenomenon in depth, the case was designed to provide as rich as picture a possible and to be used to build theoretical ideas.

2. Immersion in the Field

To familiarise ourselves with the field, we attended key events in the field of algorithmic arts, including the V&A's "Chance & Control: Art in the Age of Computers" symposium (London, 2018); GROW.Paris, France's biggest creative coding conference (Paris, 2018); the "AI as Cultural Gesture" symposium (Stockholm, 2019), the vernissage of Learning Nature: A Machine's Exploration of Our World (Oxford, 2019), and AI Art curator Luba Elliott's Creative AI Meetups (London, regularly).

In order to gather and test preliminary data, we ran a data-gathering roundtable titled "Alternative Visions for the Future of Work" (Oxford Internet Institute, 15 November 2018), exploring the discourse around AI's impact on the future of work. In collaboration with creativity scholar Dan Holloway and the TORCH Futures Thinking Research Network, we also ran a workshop around creativity, socio-technical imaginaries, and AI futures (17 June 2019), and organised the AI x Creativity panel at the Rhodes AI Lab's Annual Conference (8 June 2019), where the research team and artist and guest speaker Marie von Heyl presented and gathered feedback on preliminary research findings.

3. Participant Interviews

We identified the following key exhibitions in the ML- using art world: Gradient Descent, Nature Morte, New Delhi (2018), Artificially Intelligent, Victoria & Albert Museum, London (2018), AI: More Than Human, Barbican, London (2019), Entangled Realities: Living With AI, House of Electronic Arts, Basel (2019), and BARRAT/BARROT: Infinite Skulls, L'Avant Galerie Vossen, Paris (2019). Starting in these spaces, we interviewed participating artists whose work centred around ML techniques, as well as curators and researchers in the same field to enrich our account of the broader dynamics the subfield of ML Art. Twenty interviews with artists, curators, and researchers in digital and algorithmic arts were conducted. Sampling was guided by immersion in the field, conducting preliminary interviews, and snowballing from participant recommendations. Sampling focused most heavily on contemporary artists using ML techniques as part of their practice, but strategic care was taken to include digital artists from earlier paradigms of computer arts in order to situate the features of ML's moment in the arts.

Interviews were conducted in person, on Skype, by telephone, or email (one interview), and lasted between one and four hours. The different modes of interviewing were guided by good practice in the field to ensure equivalent depth regardless of mode.

The interviews explored participants' reflections on the developing field of AI Art, in the context of their own experiences and practices. They were asked both to look back over the history of their own practices and to consider the impacts of AI on the future of their field. The interviews were semi-structured and deliberately open in style, to account for the emerging nature of the topic and the diversity of experiences that we wanted to capture within the study. The precise approach and questions varied according to the participants' backgrounds.

For artists, example questions included: could you tell me about a piece or project you used AI technology for? What kinds of techniques were you using? How has using AI altered your practice? What are your creative process(es) usually like? How would you describe your relationship with the AI technologies you use? Would you call it a collaboration or partnership? Do you think of AI as a tool? Why/why not? What role do you think AI plays or will play in your field? Do you feel like that direction of travel is the right one?

4. Data Analysis

Audio recordings from the interviews were transcribed and analysed through multiple rounds of thematic coding using software NVivo to examine each of the motivating research questions. This was achieved through an iterative coding process to refine the themes, visualize the data, and test alternative explanations. Well-established frameworks for ensuring quality in qualitative research were utilised during the analysis and writing up process.

5. Research Ethics

Research ethical approval for the study was granted by the University of Oxford (approval reference SSH_OII_CIA_19_052). All participants were provided with an information sheet detailing the aims of the case study, details of data collection and management, and their rights as participants. Participants then signed a consent form authorising the team to use the interviews as research material. All participants chose to be named in the research and gave permission to be directly quoted.