Future Renaissance

by

Chris Luginbuhl

A thesis exhibition presented to OCAD University

in partial fulfillment of the requirements of the degree of

Master of Fine Arts

in

Digital Futures

Exhibition Details:

TMAC, 32 Lisgar Ave.

Toronto, Ontario, Canada, April 12-14, 2019

Copyright Notice

Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I authorize OCAD University to lend this thesis to other institutions or individuals for the purpose of scholarly research.

I understand that my thesis may be made electronically available to the public.

I further authorize OCAD University to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Signature: _____

Abstract

Name: Chris Luginbuhl
University Name: OCAD University
Thesis Title: Future Renaissance
Degree Title: MFA
Program Title: Digital Futures
Year of Defense: 2019

This document explores image making in collaboration with neural networks, particularly Generative Adversarial Networks (GANs). Research creation is used to develop a collaborative relationship with artificial intelligence (AI) which moves beyond the use of the computer as a simple tool in the image-making process. As a way of deepening this reflection, the images are developed using speculative fiction to imagine what intelligent machines' creation myths might look like in the distant future, and this helps suggest how we might form AI in the present. GANs are found to help express visual ideas by providing a wealth of imagery and textural detail which can be modified with the selection of training data and transfer learning. The difficulty of training GANs can be mitigated by using other machine learning techniques such as object detection to gather training data, and by working with low-resolution imagery that reduces computational demands, increasing accessibility.

Acknowledgements

This project was propelled by the powerful interdisciplinary rocket fuel of the OCAD U Digital Futures Program, its faculty and students. I feel very fortunate to have worked with my supervisor Kate Hartman, whose support, positivity and critical eye kept this rocket from self-destructing many times over. My secondary supervisor Adam Tindale provided much helpful guidance on how to think, make, write and question.

Kristy Boyce and Finlay Braithwaite contributed many helpful ideas and suggestions over the course of many long conversations about this project. I still feel like you both understand it a lot better than I do. Along the way I was graciously assisted with discussions, support and troubleshooting from various people at github.com and codementor.io. I would particularly like to thank Tushar Gupta for his patient help getting me up and running with GANs, as well as the numberless developers and researchers who freely offer their expertise online for the benefit of everyone.

I am deeply grateful to my parents who gave me a love for science, art and books, and who generously supported me through all my struggles. Finally, I would like to particularly thank Camilla who inspired me with her example and gave me encouragement to pursue this degree and this project at every stage.

Contents:

Chapter 1 - Introduction

1.1 Project Motivation

This project documents my process of creating a new kind of creative practice for myself. I began by wondering how intelligent machines would regard humanity in the far future, and what kind of stories they might tell. I explored this question with speculative fiction, and collaborated with neural networks to create images that suggest those stories, making them into physical panels using a mixture of antique and digitally controlled processes.

In this project, my role shifts from storyteller to programmer to image compositor. I tried to let the needs of the project dictate the techniques and media I used. In many cases, I hybridized existing techniques, both digital and physical. As a result, this project is as much about developing a toolkit for creation as it is about creating images.

Using some of humanity's first creative machines to help tell stories about artificial intelligence has a certain symmetry, but my collaboration with neural networks also arose from a practical need. I often struggle to begin on a blank page (or screen) to make a new image so I enlisted the help of neural networks to generate some of the raw material. I developed a back-and-forth process with these machines which helped refine the images along the way, and gave me a nudge when I felt stuck. I call this work a collaboration because the neural networks were more unpredictable and creative than conventional tools.

At the same time, this project is *about* those machines. I wrote stories that extrapolate today's artificial intelligence (AI) research to a future without people, where intelligent machines try to reconstruct where they came from in the absence of digital records. Using humanity's surviving artworks and e-waste left behind, they construct images of their origins using neural networks and computer-controlled fabrication.

As a complement to the speculative fiction, working with the minutiae of artificial intelligence (AI) technologies like neural networks helped me explore the feeling and shape of a future world in which AI plays an even more prominent part. Now is a good time to get to know these machines, since today they are relatively simple and we can still decide in what direction we want them to develop.

Using neural networks to create new images based on thousands of examples is a recent technique, dating from 2014 (Goodfellow et al), and has spawned a great deal of interest in the machine learning research community. My collaboration with neural

networks is not intended to be research on neural networks, rather it is an investigation of our relationship to these machines as it stands today, and as it may be in the future. My research into the details of this relationship is conducted through collaboration, which is in turn inspired and directed by speculative fiction.

As with most speculative fiction, stories set in the future engage with the problems and threats of the present. My stories in which future machines imagine their origins among long-lost humanity are a mirror through which we humans can contemplate the future we're building with machines.

1.2 Renaissance as Inspiration

A month studying art history in Florence inspired me to wonder, "who will create the next Renaissance?" I decided to imagine it would be machines who will make imagery inspired by the art of a lost civilization. European Renaissance art became an inspiration and provided a visual vocabulary for exploring image-making in this project, since they are both concerned with drawing attention to and memorializing key myths from the past. I do not address the religious motivations and themes of Renaissance art, but do draw inspiration from the idea of making and viewing images containing precious materials as a devotional act, and hand-making unique panels to emphasize the importance of the events and figures portrayed.

The Renaissance also offered me a jumping-off point to consider how the significance of images may change over time. What we might consider a great work of Early Renaissance art may have been considered at the time a functional object made by an anonymous craftsperson, the main value of which derived from the pigments and gold used, rather than any particular skill or vision on the part of the person who made it. Today our more individualistic, materialistic society values human artists, but is this anthropocentric view still appropriate in a collaboration with artificial intelligence? Images made using neural networks face questions about authorship (for example, Schlackman; Schneider). In this project, I am using an image-making collaboration for my own ends - as a way of working through questions about AI.

The Renaissance also offers a chance to consider how work can integrate research, image-making, and scientific and technological development, exemplified by Leonardo da Vinci's famous notebooks (Isaacson). With training as a mechanical engineer, I have done a lot of design & fabrication. Working as a photographer for several years, I found that I experience mechanical design and photography the same way: I find a question or problem, think creatively about it, bring in a certain amount of technology, and produce a result which always suggests further questions and further work. Sometimes the work

unfolds according to a preconceived plan, and other times, as with this project, it is more exploratory. Either way, the discoveries and developments that result cannot be foreseen at the outset, and there seems to be no better way to learn something than to make something.

1.3 Materials and Physical Images

In subsequent chapters, I describe in detail the process of creating images by gathering thousands of examples and training a neural network to produce new examples. Neural networks can then generate a near infinite number of new images, such that the customary way of reviewing their output is to look at many images at once in a grid of 64 images.

I wanted to explore the effect of giving these images a bit more autonomy, taking them one at a time out of their anonymous grids and delivering them out of the digital and into the physical world in some way. Early on, I had the idea to fashion them out of e-waste, as a way of thinking further about what humanity's legacy will be. What tangible, material artifacts will we leave behind, and what ideas will outlive us? In one of the earliest backstories I developed for this scenario, humanity leaves behind intelligent machines, but also great quantities of electronic waste and our digital data becomes corrupted and lost. Thus, we bequeath to machines algorithms that can create limitless digital images, but only our durable physical artworks survive.

1.4 Research Questions

Although this work engages with techniques and technology, and also takes its setting and inspiration from speculative fiction, the focus of the work is on what the experience of working with these emerging technologies can show us about our evolving relationship to them.

The research questions that drive this work are as follows:

1. How can humans and neural networks collaborate on image-making? In what ways do these machines take on a role beyond being simple tools?

2. How can individuals without corporate budgets and equipment make use of advances in neural networks for making images?

3. How might neural networks serve as a prompt or inspiration in my image-making, and help me find ideas and solutions when I am stuck?

1.5 Terminology and Limitations in This Work

This project deals with an aspect of machine learning concerned with neural networks. I frequently refer to these as "machines" in this document. The neural networks I am discussing all relate to image creation, image processing, or object detection in images. Generative Adversarial Networks (GANs) are a generative art technique which uses two neural networks to create new images based on a dataset of training images.

All of this is a subtopic in the field of Artificial Intelligence, which refers to computer techniques for accomplishing tasks which otherwise require a human. I discuss artificially intelligent machines in the distant future as a way of addressing the field of artificial intelligence as a whole.

The relationship of these various parts can be visualized using the diagram below, where the red "This" refers to the image generation techniques employed here.

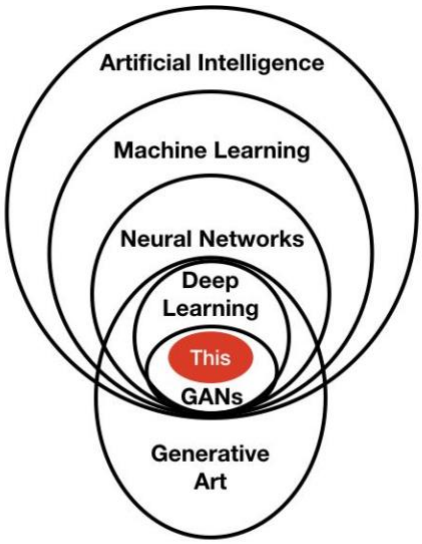

Figure 1. Relationship of image creation in *This* document to other topics in artificial intelligence research and generative art.

A note on the term "collaboration" – I am not the first person to collaborate with neural networks to produce images or artworks. I am using this potentially provocative term to draw attention to the difference between working with conventional tools and working with the emerging products of artificial intelligence research. These machines

are not exactly intelligent, but not exactly dumb either, and I am interested in exploring the changing ways we can relate to them as they develop in complexity and capability.

In chapter 4 I discuss my use of computer controlled machines to help produce the physical images. This experimentation does further my exploration of a future increasingly populated by machines, but I do not refer to this part of the work as a collaboration, since these electromechanical devices are closer to being tools, and do not exhibit much creativity or the ability to surprise me in the way neural networks do in this project.

Regarding the scope of this document, it is not my intent to offer a social, political, or art history survey of the Renaissance, nor do I wish to make claims about what artistic expression or historiography will look like centuries in the future.

Although I touch on some questions about AI safety and ethics, I do not explore the social and political dimensions of AI in detail, nor have I addressed the significant and problematic ways racial, gender and other biases are present in today's AI. It is my hope that by getting to know more about today's AI, my subsequent work can engage in making subsequent AI more safe and just.

1.6 Document Structure

Chapter 2 reviews the context and literature that pertains to this work, with a look at how speculative fiction, speculative design and myth making have been used to make sense of our world, and how storytellers and artists have approached neural networks and robotics.

Chapter 3 introduces the methodology of this work, drawing connections between speculative fiction and research creation in the context of this creative collaboration.

Chapter 4 describes the work done with speculative fiction, using neural networks to build datasets, and using those datasets to train other neural networks to produce images, which I then composited together and transferred onto panels. I describe other approaches I tried along the way, and offer a short summary of results at the end of each section.

Chapter 5 presents the results of the work, starting by reprising the methodological approach to produce the work for the final exhibition, followed by a discussion of the outcomes of this practice and study.

Chapter 6 summarizes the results and findings and concludes the document with a discussion of the research results, limitations of the work, further directions, and a conclusion.

Chapter 2 - Literature and Context Review

This chapter reviews how literature, art and design have been used at different times to develop and process ideas about artificial life and its relationship to human society. Organized roughly chronologically, it takes us up to debates and questions of the present day.

2.1 Roots of AI and Robotics

Most of this chapter is focused on the most recent work in AI and robotics. However, it is worth noting that some of these ideas about robots or automatons have their roots in antiquity. The idea of mechanical beings goes back at least as far as Homer (c 750 BCE). Early examples can be found in Homer, who describes forged assistants that exist in the realm of the gods:

> "These are golden, and in appearance like living young women
> There is intelligence in their hearts and there is speech in them
> and strength, and from the immortal gods they have learned how to do things."
>
> -Homer (von Nettesheim et al. 235)

Examples of created beings occur in many other periods and cultures. In Greek Myth, a bronze giant named Talos guarded Crete from pirates. It was fashioned by Hephaestus, a smithing and sculpture god and maker of technology. Jason and the Argonauts must contend with Talos when it tries to sink the Argo. In the Talmud, God created beings, but occasionally so did mystics (Michaelson). The Golem was created to protect the Jewish people by Rabbi Judah Loew of Prague ca. 1580. Meddling with creation and the inherent power of life made trouble: the golem grew violent, and had to be destroyed. We are revisiting these ancient problems with new urgency today as our present day Golems learn to walk, speak and kill.

There are a number of subsequent stories of heads that speak, such as that of the Greek Orpheus, and later the Brazen Head (Gower). The shift to intelligence over mere servitude seems to be a Renaissance or Enlightenment idea ("Brazen Head").

Only much later did the implementation of a thinking machine become the subject of serious speculation. Lady Ada Lovelace, the first computer scientist, went farther than anyone of her time or for a long time afterward in thinking about the general properties of numerical processing machines. She insisted that this was not an independent intelligence however. In 1842, she wrote "The Analytical Engine has no pretensions to originate anything. It can do whatever we know how to order it to perform." (Toole 722)

In 1950, Alan Turing's paper "Computing Machinery and Intelligence" began what can be thought of as the modern age of artificial intelligence research. It began with the statement, "I propose to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think.'" These definitions are still vexing to scholars today. In the intervening years, much attention has been devoted to machines taking on tasks formerly reserved for humans, sometimes using Turing's "Imitation Game" (often referred to as the "Turing Test") as a measurement of success.

2.2 Speculative Fiction

The term "speculative fiction" may seem redundant, if all fiction is speculative to some degree, but speculative fiction is differentiated from regular fiction in that the setting of the story differs from accepted assumptions about history or consensus reality. Coined in the 1940s by Robert A. Heinlein (219) to refer to a subgenre of science fiction with an emphasis on possible futures and societal issues rather than issues arising from technology. It is also sometimes used to denote historical fiction.

Despite the relative novelty of the term, stories fitting this description have been told for millennia — Euripides's play *Medea* (431BC) is one early example. These stories, in an increasing variety of mediums, are willing to explore questions of "what if?" and defy any expectations that the results should mirror the world of day-to-day experience. Not limited to possible futures, these stories can explore alternative pasts and presents, and may be more concerned with questions about society than questions about technology. They can be intensely political, from George Orwell's *1984* to Margaret Atwood's *A Handmaid's Tale*, or may focus on a mixture of technological and societal possibilities like this present work.

The aim of all such works is not always to produce a finished story. The process of exploring the implications of a story world may be an end in itself. Donna Haraway advocates what she terms "Speculative Fabulation." For Haraway, "speculative" does not mean "unimportant", or "without regard for consequences", but rather a vibrant, living approach for connecting our future to our present.

In addition to being a method for contemplating the future and working through issues facing society today, speculative fiction is a storehouse of different responses to some pressing questions that have arisen again and again. Of particular interest to this project are the perennial questions, how would an artificial intelligence relate to humanity? How could we co-exist with intelligent machines? Speculative fiction helps take these questions from being solely the domain of experts into the public domain, since books and stories are how we disseminate and discuss ideas on a societal scale (Melbourne School of Design 5:40).

From Isaac Asimov's *Robot* series to Stanley Kubrik's *2001: A Space Odyssey*, we have continued to explore big questions about how we can co-exist with artificial intelligence without coming to final answers. As technology and society evolve, over time we find ourselves asking the same questions in a new context, in light of new developments.

Post-apocalyptic fiction is another sub-genre of speculative fiction that evolves to reflect the concerns of its age. Contemporary writers like Paolo Bacigalupi write about a world whose climate has changed, but in 1960, the dominant fear was of nuclear war. Walter M. Miller's *A Canticle for Leibowitz* portrayed a world after such a catastrophe. The book demonstrates the power of speculative fiction to concoct a fantastical and remote scenario and make it touchingly portray our societal aspirations and failures as universal and timeless. This is probably the work of speculative fiction that most closely resembles what I am trying to do with this project, as it evokes a cyclical view of human history, and portrays the safeguarding of historical information through a dark age, and image-making as a devotional practice.

Taking place in the aftermath of the "flame deluge," a group of monks collect and venerate the relics of St. Leibowitz, an electrical draftsman from the time before. One monk painstakingly makes an illuminated copy of a circuit diagram made by Saint Leibowitz without understanding its original purpose, while other monks store and recopy the world's remaining books for the benefit of future generations. Like Macaulay's *Motel of the Mysteries*, the book alludes to what has been termed "speculative archaeology" (Holtorf), and like Asimov's *Nightfall*, how without learning certain lessons history can repeat itself. It ends with something like the beginning of a new species, along the lines of Kubrik's *2001: A Space Odyssey*, or Atwood's *Oryx and Crake*.

2.3 Speculative Design

Speculative fiction is usually used to refer to works of fiction in text, film, theatre and so on. A related field, often called speculative design, uses approaches from design to ask questions rather than presenting answers, as with conventional design (Auger 1). Variations of this approach exist under various names, such as conceptual design and design fiction. Outputs of speculative design can be as simple as a physical or fictional "product" which is meant to provoke discussion rather than generate sales. Speculative design can also mean creating a fictional world complete with stories, characters and props. The design process itself is an important source of insight, and creating some tangible physical outputs as part of this activity is a means of allowing others to engage with the products of speculative design.

In 1983, Donald Schön was one of the first to challenge the conception of design as a purely scientific undertaking, intended to solve well-defined problems. Schön described reflective practice as, “an epistemology of practice implicit in the artistic, intuitive processes which some practitioners do bring to situations of uncertainty, instability, uniqueness, and value conflict.” Since then, other design practitioners have expanded on this field.

Dunne and Raby’s book *Speculative Everything* makes the practice seem accessible, through the abundant use of examples and their accessible language, e.g. “One way of considering the fictional objects of speculative design is as props for nonexistent films.”

Liam Young addresses questions about the future using a similar approach framed as speculative architecture. His project *Under Tomorrow’s Sky* is set in a fictional future city. The approach is collaborative, bringing together scientists, designers, futurists, illustrators and science-fiction authors. This provides a framework for developing a vision of the future, but is also intended to provoke discussion about pressing questions facing us today. (Young, “Under Tomorrow’s Sky”)

Design fiction is another term used to differentiate a method distinct from general story writing. Design fictions offer questions about the future and dig into ambiguous grey areas. As explained by Fabien Girardin, co-founder of The Near Future Laboratory: "Design Fiction doesn't so much 'predict' the future. It is a way to consider the future differently; a way to tell stories about alternatives and unexpected trajectories; a way to discuss a different kind of future than the typical bifurcation into utopian and dystopian" (Bleecker; Girardin, “Design Fiction”).

2.4 Artificial Life

Some of the earliest discussions of electronic artificial life arose during the 1940s with the field that came to be known as Cybernetics, led by Norbert Weiner. His *Cybernetics or Control and Communication in the Animal and the Machine* stated that communication, control and feedback are fundamental aspects of any organism or system. The book handles a variety of technical topics in a traditional technical way, but chapters on language, psychopathology, the brain and society opened up the discussion of these topics to a broad audience and paved the way for subsequent researchers to refine these ideas in their own fields.

The development and spread of the digital computer allowed experimentation with swarms of “cellular automata” such as in Conway’s “Game of Life”, which showed that automated digital “cells” programmed with a few simple rules could collectively engage

in complex and seemingly organized behaviour. This marked a new age in which our machines could surprise us by exhibiting behaviour we couldn't have anticipated. It is this capacity for unpredictability and surprise that I am exploring as the foundation of a human-machine artistic collaboration.

ELIZA (Weizenbaum) was the first chatbot to achieve broad exposure upon its release in 1966. Technically a Natural Language Processing computer program, ELIZA used rule-based pattern matching and substitutions to give the user the superficial impression of conversation with a human being. Unlike Game of Life, however, ELIZA rarely generated surprising or unexpected responses to input. This has been a theme ever since; systems that attempt to replicate human bodies or capabilities usually fall laughably short, while those that follow a different course develop interesting behaviours and characteristics that clearly belong to a different, non-human world.

2.5 Information and Surveillance

As of today, chat bots and virtual assistants remain poor imitations of human beings. Virtual reality pioneer Jaron Lanier's 1996 essay "My Problem with Agents" predicted that despite inevitable improvements, these assistants would remain "agents of alienation." He also predicted abuses of personal data for commercial ends that have since become normal. And today's agents that learn our shopping habits and troll social media for political ends are more effective and much harder to spot than the virtual assistants referred to by Lanier.

The Facebook game Naked on Pluto (2010) highlights the gathering and use of personal data, particularly on social media. A player finds themselves on Pluto, under the rule of a corrupted Artificial Intelligence. The player soon encounters a community of 57 animated bots. Through various interactions, these capture the data from their Facebook account. This early work was one of the first to highlight the importance of the collection of personal data for powering AI, and the potential ethical issues involved. Since that time, Amazon has become the second company to surpass a US$1 Trillion market capitalization, based largely on its collection and use of customer data. In addition to increases in computing power, perhaps the biggest reason for the increased effectiveness of AI is that we are now able to feed it copious amounts of data to "train" it.

Although the extent and significance of personal data collection has since become more widely known, the data underlying the training of various machine learning algorithms is still largely unrecognized. Neural networks depend on large amounts of data for "training" – the process by which they extract patterns from data. They are then able to make assessments or predictions about situations that are not in the training data.

American artist Trevor Paglen has made a career of bringing to light images and data that are normally only accessed by machines — information that gives machine learning systems their power to recognize human language, handwriting, faces and even military targets. His 2017 show at Metro Pictures in New York *A Study of Invisible Images* was particularly focused on the data underlying several types of machine vision - military drones, facial recognition software and automated captioning of images. Through a series of rapidly changing images, Paglen shows us what an AI "sees" as it's learning to recognize license plates, faces and emotions, in the service of police, corporations, and governments. One gallery of the exhibition was devoted to his "Hallucination" series in which neural networks are trained on less rationalistic sets of images - the interpretation of dreams, or omens and portents. Another network creates images which are then recognized by the first network as something that isn't really there. *His False Teeth* (2017) was interpreted as Freudian dream symbology by one network, and another trained on omens and portents recognized *Rainbow* (2017).

2.6 Robotic Art

A robot is generally any machine that can be programmed to carry out a series of tasks. In science fiction, the term usually refers to physical machines whose forms are influenced by biological beings. Robotic art has some bearing on this work, raising questions of autonomy and control, and confronting us with the uncanny.

As with neural networks, we find that Robots are not considered independent, autonomous artists on their own, and they are usually considered in relationship to people in some way. The curators of the 2018 *Artistes & Robots* exhibition in Paris expressed that the robots in the show were artists' robots, rather than robot artists (Leduc).

We encounter many of the same kinds of questions when making images with neural networks. These artworks can help us consider how machines approach us, how they differ from us, how they regard us, and how we can work together. My work explores how this balance informs human-machine collaborations for making images that have the potential to say something about human and machine society. Robotic art highlights the similarities to humans and the strangeness may more obviously than images made by a computer. In this project, making the work into physical panels is a means of making the uncanny quality of machine/human imagery less virtual and more visceral.

This project also examines the collaborative strengths that arise because of the differences of humans and machines.

In “Trans-species Interfaces” (Herath et al 113), Ken Rinaldo reflects that the future of life on this planet exists where the biological and machine interact, citing robotic experiments he says may be helpful guides for the design of technologies that are beneficial for all living beings. With echoes of Donna Haraway’s essay “Cyborg Manifesto”, it suggests that traditional distinctions between human and machine, human and non-human are breaking down.

Some works, like Robotlab’s Bios (Gommel et al), help us consider the role of machines undertaking traditionally human tasks. Others, like Sougwen Chung’s 2018 *Drawing Operations* performance highlight the possibilities of collaboration, with both the artist and the robot holding brushes and make marks on the same canvas, prompting reflection on mark-making, memory, mimicry and bio-inspired cognition.

2.7 Mythology

Myths are often relatively simple stories that speak directly to origins, identity and the spiritual possibilities of living in a society. It has been helpful to study the aspects of human myths that seem to be universal, and the diversity of ways myths have been represented in human cultures, in order to leap into speculative myth making for a future machine society.

Joseph Campbell proposed that myths from many cultures share many of the same basic attributes and themes. His book 1988 book *The Power of Myth* divides myths among various themes: a quest or adventure, creation myths, being in harmony with the world, sacrifice, love, and relating to the infinite, suggesting that these are human universals.

Campbell’s work raises the question of who can summarize the world’s myths. It is certainly questionable whether any one scholar can offer perspectives and comparisons between myths belonging to diverse cultures not their own. Schools of “particularists” and “comparitivists” have subsequently offered contrasting viewpoints on the question of how much we can correlate myths taken from different cultural contexts. Regardless of how we answer this, it is possible to read *The Power of Myth* for its sweeping survey of some of the world’s mythology without embracing all the conclusions of its author. It seems significant to note that flood myths, for example, occur in the epic of Gilgamesh, the Hebrew Bible as well as Hindu, Greek, Norse and Aztec mythology.

To consider the various ways myths can be depicted in images and sculpture, James Cahill’s *Flying Too Close to the Sun* catalogues representations of myths in European

painting, and modern and contemporary art. It does not deal with myths outside the European classical tradition, and does not venture what myths are current today.

Campbell observes is that there is now a profound lack of mythology and tradition that guided and informed life in the past. He believed the pace of change may currently be too great for new mythology to emerge, and that we must wait for more stability. This project attempts to depict that future state and to reflect on some of the mythology that results.

2.8 Neural Networks, GANs, Image Analogies

There are computer image-making techniques that do not involve artificial intelligence. Generative art, algorithmic art or procedural art, for example, refers to images made using algorithms on computers, with the artistic intent and guidance supplied by the programmer or operator in order to produce artworks. These techniques often involve programming the computer with a relatively simple set of instructions, executed repeatedly, producing patterns and ornamental designs or three-dimensional geometries. These can be quite complex, and relatively simple algorithms may produce surprisingly complex effects. They are often geometric and may be abstract, but are rarely figurative.

More complex rule-based procedural art-making systems are considered AI techniques. One such example is a system called *The Painting Fool*, which takes on some of the creative responsibilities, and which is considered by its creator to be a collaborator in art-making, rather than a simple tool (Colton 2). Its creator has reported that there is some resistance to having a computer involved in an art-making process as a collaborator. For one thing, it is not possible to have a conversation with the computer about its process. Also, some created works, such as text, speak for themselves more readily than visual art, where a higher degree of interpretation is usually required. *The Painting Fool* does not engage in contextualizing or framing its work verbally, or presumably at all (3). The Painting Fool is one of the more complex systems, having been actively developed for more than ten years. However, in general, the relative simplicity of the algorithms used in generative or procedural art produces relatively simple results.

In the last few years, artists have been increasingly making use of imagery created largely or entirely by artificial intelligence techniques involving neural networks, owing to increases in computing power, powerful new techniques and the ubiquity of large datasets being collected on every aspect of the world and its inhabitants. Some of the most dramatic new developments are in a subset of Machine Learning called Deep Learning (see figure 1). This project makes use of a type of neural network called a

Generative Adversarial Network (GAN). GANs are used for producing original, unexpected visual imagery based being "trained" on databases of tens of thousands to tens of millions of images. Further background on digital image creation and manipulation with neural networks can be found in Appendix A.

AI techniques such as GANs are fundamentally different from traditional computer programming, including that used for generative art. Conventionally, we program computers to follow a set of rules and give us the output. By contrast, with neural networks including GANs, we give the computer lots of examples of images and it "learns" the patterns underlying them by tuning adjusting the weights which correlate individual neuron-like nodes in its neural network. GANs are composed of two separate neural networks, a generator network and a discriminator network. The generator begins by producing images with random pixels, and the discriminator judges whether this is similar to the images on which it has been trained or not. The generator repeats this process by adjusting its output to try to pass its "fake" images by the discriminator. The discriminator uses feedback to become more discerning with each cycle. After much computation, the generator network is finally able to produce images that are considered "real" by the discriminator network and, when done well, look interesting to human eyes.

Unlike much generative art, the resulting output cannot be anticipated by conventional means. When done effectively, the results can be strikingly uncanny. Artists such as Memo Akten, Trevor Paglen, Anna Ridler, Mario Klingemann and Helena Sarin have incorporated this kind of image-making into their artistic work. See for example Trevor Paglen's *Comet (Corpus Omens and Portents) Adversarially Evolved Hallucination*, 2017 and Mario Klingemann's *Neural Glitch*, 2018.

A relatively complex type of deep learning technique, GANs were first announced in 2014 (Goodfellow et al). Published books on GANs are just beginning to emerge. There is one chapter in Goodfellow on GANs (ch. 20, Goodfellow et al, Deep Learning), treating the mathematics and theoretical basis of GANs. Langr and Bok's book is available online as a work in progress. It cautions the reader: "you'll need intermediate Python programming (2+ years), knowledge of machine learning theory and the high-level mathematics behind it." (Langr and Bok, "Welcome"). Much of the scientific literature in journals and print is divided between treating the GAN as a black box, and technical papers whose technical complexity restrict them to an audience with an in-depth knowledge of computer science, math and convolutional neural networks.

By contrast, Andrej Karpathy, a researcher and head of AI Research at Tesla, while still a computer science PhD student at Stanford, argued for a hands-on learning approach that delays engaging with theory, saying "…everything became much clearer when I started writing code." (Karpathy). In his teaching, he says he has deliberately focused his

writing on blog posts, online courses, and Javascript libraries that make it easy to train and experiment with neural networks, rather than writing books.

Whereas not many books are available, there are a number of well-regarded courses for machine learning such as FastAI's "Machine Learning for Coders", Stanford's CS229 online course, and Udacity's Deep Learning courses that follow a similar hands-on problem solving approach. These courses often use collaborative web-based environments like Jupyter Notebook and Google Colaboratory that incorporate live code, equations, visualizations and explanatory text. It has been argued (Kingsley; Somers) that this sort of interactive format is a more appropriate format than print or PDF for sharing scientific results in an academic context, because readers hoping to understand or replicate the results will frequently need the datasets, code snippets and an environment in which to run them and experiment with alternative treatments of the data. Machine learning seems to be one of the earliest fields to make extensive use of this interactive format for teaching, learning, experimentation, and the exchange of ideas, and this may partly explain the rapid development of the field. There were 1.3 million Jupyter notebooks shared publicly in April 2018 (Somers), and the approach is gaining momentum.

Twitter has also become a key platform for artists and machine learning researchers alike to share their latest approaches and the resulting images, often faster than results can be published. Papers like Brock et al's BigGAN may be studied, remixed and elaborated upon by a wordwide community (see for example Shane; Schwab) for six months or more while still in pre-print before being formally presented. Of course, these are often results that are presented without editorial or peer review, so must be approached with caution (Lipton).

In practice, GANs usually require significant trial-and-error to train. They sometimes fail to train, requiring adjustment of hyperparameters. Having begun to train, they do not generally converge to produce useful results. Rather they improve for a time, after which they drift back into incoherence and endless repetition. This is known as "mode collapse", in which they produce the same output or series of outputs repeatedly. Some advances in GANs since 2014 such as Wasserstein GANs and Progressive Growing of GANs have focused on improving their ability to be trained predictably reduce some of their more unpredictable and undesirable behaviour, allowing the possibility of generating slightly larger output images, all at the cost of increased computational requirements, which elevates many of these developments outside the reach of many individuals.

There are other barriers to making convincing new images with GANs. Increasing

the resolution of the output not only increases the computational load exponentially, it also makes it much easier for the discriminator to spot fake images, which can lead to a failure to train. At the moment, creating images much larger than 128x128 pixels is generally reserved for researchers who have dedicated access to high-end computing clusters for days or weeks at a time for each experiment they run. Despite several promising new toolkits in various stages of development and pre-release testing that make machine learning techniques easier to use for artists and designers, such as Runway ML, ml4a.org and ml5js.org, none are yet available which allow users to train their own GANs, even if they do have the hardware to do so.

2.9 Artificial Intelligence Safety and Ethics

Another factor which has led this type of research away from print and towards online-only platforms is the availability of powerful and flexible cloud computing. It is not necessary to have a powerful computer on hand, when Amazon Web Services offers scalable hourly rental and Google collaborator offers free access for individual use. The need for this sort of processing power in AI research is not without other costs, however. One run of Brock et al's experiments training BigGAN is estimated to use almost 5000kWh, enough to run an American household for six months (Schwab) at a time when IT and communications is contributing 3.5% of global CO2 emissions and rising ("'Tsunami of Data' Could Consume One Fifth of Global Electricity by 2025").

In addition to the present-day threats from surveillance, bias and privacy questions, there are other risks that may come with AI that matches or exceeds human-level intelligence, often called "artificial general intelligence," or "full AI". The same questions that have been explored in speculative fiction are taking on a more pressing urgency. Will machines be "beings" we can relate to? How will they relate to us? Would they form their own society or would they be enmeshed with our society?

In 2015, Edge (edge.org) asked a select group of luminaries "What do you think about machines that think?" They condensed the resulting 186 responses into the book *What to Think About Machines That Think*. The responses were extremely varied and revealed no clear consensus about the future of AI, or even whether something like Artificial General Intelligence is even possible. Most of the essays struck an optimistic tone, but there seemed to be general consensus that it matters a great deal that we deliberately guide the ways we research and make use of AI systems.

Bill Gates and various other famous science and technology figures have warned against the potential dangers of AI. Stephen Hawking famously told the BBC: "The development of full artificial intelligence could spell the end of the human race," and, "Humans, who are limited by slow biological evolution, couldn't compete and would be

superseded" (Cellan-Jones). Bill Gates has also warned against the potential dangers of AI.

Nick Bostrom's 2014 book *Superintelligence* deals mainly with AI safety which, rather than focusing on the possible outcome of the development of "full" AI — that is, intelligence that matches or exceeds humanity's — focuses on the path we should tread to get there in order to avoid various pitfalls. This thesis project touches on approaches at once — envisioning the outcome of the development of AI as a way to help us consider how to build it.

As research money continues to flow into all areas of artificial intelligence research, a number of small organizations are researching AI safety and AI ethics. Much of the activity concerning AI safety and ethics is being conducted by a few organizations, whose work includes policy recommendations and basic AI research. In industry, Google DeepMind and Google Brain have some staff in these areas. In the UK, The Future of Humanity Institute and The Cambridge Centre for the Study of Existential Risk, and the Leverhulme Centre for the Study for the Future of Intelligence study both technical and strategic questions related to AI safety.

In the US, the Machine Intelligence Research Institute was one of the first to get involved in this area in the early 2000s. OpenAI is a more recent non-profit organization exclusively devoted to these issues, and has received US$1B in funding commitments from the technology community. Other groups include the Berkeley Center for Human-Compatible Artificial Intelligence, the Future of Life Institute at MIT, Alan Dafoe's research group at Yale and the Center for a New American Security's Artificial Intelligence and Global Security Initiative in Washington D.C.

2.10 Conclusion

The existing literature lays out important questions about the development of AI, its use for human-machine collaboration, its use in art, and ethical and safety questions related to its development. Centuries of myth, speculative fiction and speculative design have provided questions, ideas and methodologies for examining these issues and encouraging discussion, on which this project aims to build.

Chapter 3 - Methodology

In this project, I investigate the developing relationship of humanity and AI through speculative fiction by participating in an image-making collaboration with AI.

I combined three methodologies for making the work and exploring my research questions: speculative fiction is the telescope that I used to get perspective on the possibilities of human-machine relationships over time, research creation is the microscope that revealed the character of neural networks in detail, and I collaborated with neural networks at every stage. These methodologies are described further in the following three sections of this chapter.

I took on various roles in this process, writing stories, programming computers, compositing images, and developing experimental image panels. Under my direction, my neural network collaborators filtered and analyzed source images, created new images for composites, and invented textural details permitting the images to be enlarged. In the process, I developed a toolkit for my creative practice.

My methodology makes use of research creation, a term normally associated with art-making. I note that my focus is not on a thorough questioning and criticism of art tradition and practice. Esther Pasztory argues that the impulse to make things is a human universal, and that this process is primarily cognitive and only secondarily aesthetic (10, 13). Observers may choose to classify the images I make art, though they probably owe at least as much of a debt to the traditions of design. In any case, the boundaries of art and design are not clearly defined, and both may involve making works without a pre-determined plan or intended outcome in order to explore a problem or theme.

3.1 Speculative Fiction

As mentioned in Chapter 2, speculative fiction is a means of asking "what if?" questions and working through the implications without assuming the answer will conform to observable reality.

My approach to using speculative fiction for this work is inspired in part by speculative design methodology. Speculative design approaches problems by asking questions rather than presenting answers (Auger 1). This may mean creating a fictional scenario complete with stories, characters and props. The design process itself is an important source of insight, and a means of uncovering and working through issues. I am bringing in this approach of questioning and using characters and scenarios in a similar way. However, I am not creating physical products as is often done in speculative design practice as a means of allowing others to engage with the questions.

Other writers describe related methodologies in terms of speculative fiction (Allardice), design fiction (Girardin) or speculative fabulation (Haraway). For Donna Haraway, 'speculative' does not mean 'unimportant', or 'without regard for consequences'. Rather it is the opposite - a vital and vibrant approach for connecting our future to our present. I use the term speculative fiction here, which includes writing stories, making physical images, digital images and captions.

Speculative fiction can be powerful because it is widely used as entertainment, but it can also be practical, political and may provide a foundation for various kinds of research. Liam Young says it is, "an extraordinary shared language…it's really how our culture shares and disseminates ideas." (Young, "Brave New Now" 5m:40s). Questions about the future posed by speculative fiction usually reflect the concerns and aspirations of the present, and this project is no exception. Will our jobs be replaced by machines? If we develop intelligent machines, what will their attitude be towards humanity? Questions that once seemed in the realm of the distant future demand to be addressed with greater urgency as technological change brings us a different world to navigate and presents us with new choices. Speculative fiction helps us get prepared for foreseeable technological disruption.

Speculative fiction does not always remain fiction. In 1983 the film "Wargames" sparked an in-depth analysis of computer system vulnerabilities by the US military that led to the first cybersecurity legislation (Kaplan F., "Cybersecurity's Debt to a Hollywood Hack").

In other kinds of research, we sometimes find that thinking about the distant future works hand-in-hand with detailed technical work in the present. For example, OpenAI works to understand and influence the long-term safety of AI, in part by helping develop AI. According to their Charter, "To be effective at addressing AGI's impact on society, OpenAI must be on the cutting edge of AI capabilities — policy and safety advocacy alone would be insufficient" ("OpenAI Charter")

My research uses questions from speculative fiction to focus the detailed explorations I undertake through research creation. The results of the research creation then suggest new questions for the speculative fictional premise, and the cycle repeats.

3.2 Research Creation

The research creation approach came from my desire to get acquainted with some of the qualitative aspects of artificial intelligence. Although today's neural networks are a long way from the versatile and capable artificial intelligences portrayed in science

fiction, they have begun to demonstrate their aptitudes and limitations in creative applications in a way that shows hints of personality and character. Research creation may be an unconventional way of studying computer science. But neural networks are not traditional computer programs, so perhaps it's appropriate to take a different and less linear approach than traditional computer science takes to understanding them.

My initial explorations of the theory behind neural networks led me into tangled thickets of math and computer science that did little to help me understand why and how neural networks produce the output that they do. Artificial Intelligence researcher Andrej Karpathy advocates a hands-on experimental approach. In "A Hacker's Guide to neural networks" he wrote, "My personal experience with neural networks is that everything became much clearer when I started ignoring full-page, dense derivations of backpropagation equations and just started writing code." While I did not entirely skip the traditional approach of studying algorithms and writing simple programs, I did engage right away with creation using these new tools to see what knowledge could arise from experimentation and creation. My aim was to understand how neural networks could go beyond being tools for image creation and become my collaborators in an image-making project.

My research creation consists of the works created as well as a systematically documented, self-reflexive practice used for their creation. I have also documented the process of building the materials necessary for creating the work on a blog[1]. Writing blog posts approximately weekly helps me document my thought process and discoveries as they happen, and makes it possible to see my decision-making process as it happened.

As Niedderer and Rowoth-Stokes (2007) suggest, extensive documentation and reflection helps bring some objectivity and rigour to the process of creation. The outcome may be highly subjective without this kind of process (Biggs and Büchler).

Chapman and Sawchuk refer to this process as research creation, noting that the terms "practice-as research" or "arts-based research" prevail elsewhere (6). They draw further distinctions of "research-for-creation," "research-from-creation," "creative presentations of research," and "creation-as-research," and this project engages with all four of these phases or types.

Research for creation involves literature review, gathering background material and creation of prototypes (15). Research from creation points to the knowledge that is generated by making a work or performance, including studies and prototypes, and not simply the evaluation of a finished work. (16). Creative presentations of research are the ways in which research results can be shared or experienced through a creative work or

[1] http://hisnameislong.com/category/blog

performance (18), and creation as research interweaves theory and practice to explore what knowledge emerges not through, but as creation (19).

3.3 Collaboration

My practice is an investigation into the possibility of collaborating with neural networks. For the first time in history, machines like these can produce images that could not have been anticipated by the humans who programmed them. In part, this is because machine learning technologies like neural networks are not programmed in the conventional sense, in which the machine follows a series of programmed steps. Machine learning techniques for image creation supply the computer with a large number of examples, and tell it to work out the rules needed to create more examples. For the time being, human programmers are still very much needed for creating the neural network architecture and gathering the examples with which to "train" the neural networks, but the humans influence the machines' output only indirectly. The machine has a new degree of independence as a result.

Collaboration naturally comes with risks. For the moment, I am not worried that neural networks will take over my project, but I wonder what will happen if they produce output I don't like, or seem determined to advance a kind of image or aesthetic at odds with what I have in mind. I am not expecting this is a partnership of the kind I might have with a human collaborator. The intentionality and editorial discretion resides with me - I select the architecture for the neural networks and the inputs on which to train them, I can select the images I like best from among its outputs, and can scrap any work that isn't compatible with my aims.

On the other hand, I am handing off most of the image-making to the neural networks. I am also investing a lot of time in implementing, training and running them, so I am heavily committed to using their output when it emerges. I am committed to assembling the story's images from whatever imagery the neural network delivers. I know I will not be able to reason with it, or discuss creative ideas in any conventional way. I have decided to embrace the unexpected nature of the results and be flexible about the kinds of images I will composite together as a result.

While speculative fiction provides the bounds, themes and inspiration for the subject matter of the image making, this collaborative research creation creates opportunities for discoveries to arise.

Chapter 4 - Processes

This chapter documents my processes in developing tools for telling stories using speculative fiction and collaboration with neural networks. It is divided into several activities or processes which served as studies or prototypes. With reference to Chapman and Sawchuk's distinctions for research creation activities, each of the activities in this chapter are made up of a research-for-creation phase called Investigation in which I develop techniques, knowledge, tools and/or datasets. This is followed by a research-from-creation phase called Experimental Work in which I experiment with creating a something as a means for further investigating the subject matter.

These processes began with some speculative fiction exercises, then moved on to image-making, then developing into a collaboration with neural networks on image-making. This required some technical exploration as I experimented with different approaches, and the results shaped the nature of the collaboration along the way.

The stories about AI which support the narrative of the work were shaped by my experience of what it was like to work alongside neural networks. The collaborative toolkit was developed according to needs discovered in practice while developing image studies and prototypes.

4.1 Speculative Design Process

This project began with the desire to portray future machines telling their history. To develop the basic parameters of that story and its authors, I began by reading a lot of speculative fiction and conducting a speculative archaeology card game in class to explore the intersection of physical objects and speculative fiction stories. I then developed a few text-based stories as prototypes based on what I found.

4.1.1 Investigation

I wanted to situate this work within the existing discussion about humanity's relationship with AI that is ongoing in the popular press, in the research community and in speculative design world. At the same time, to make the speculation as open-ended as possible, I considered scenarios in the distant future, including post-apocalyptic settings.

As well as conducting a literature review on current debates within AI and robotics, I decided to update myself on recent post-apocalyptic speculative fiction and speculative

fiction dealing with AI. I began a research blog[1] in which I recorded some observations from this reading, and then journaled my research process approximately weekly thereafter. Ingesting dozens of books and audiobooks on this theme resembled the way neural networks are trained on thousands of images, saturating their neurons with a particular theme so that they can produce new examples, and in this way I took a step towards my neural network collaborators.

Before engaging with image-making, I did several iterations of the speculative design using images, text, and a card game, discussing the results with classmates.

4.1.2 Experimental Work

I began by writing a short story called "Paradise (to be) Regained" as part of the Transmedia Storytelling course. The story seems to be set in Earth's rural past, but as the influence of nature spirit-like AI becomes clear, we question who in this civilization is serving whom. The story is included as Appendix A.

I wanted to explore alternatives to a text-based story-telling medium, so I experimented with different formats for subsequent stories. I presented the next story aloud in class in the tradition of oral storytelling, leaving the narrative unfinished as a background to encourage engagement, discussion and elaboration. Based on these discussions, I developed one thread into a story world in which a Tesla manufacturing robot grows eccentric after the end of humanity, finds a book on Medieval panel painting, and begins making environmental self-portraits using materials including gold reclaimed from e-waste.

Experimenting further with non-traditional formats, I developed a subsequent story as a museum exhibit set in the distant future retrospectively exploring events in our future through robot creation myths and "early" robot artists, who speculated about the vanished humans.

In order to explore the potential for objects to inspire stories, I developed a "speculative archaeology" card game, modeled on the description of a seminar at Linneaus University in which participants fashioned stories based on fragmentary evidence from archaeology (Holtorf). Participants were asked to imagine how future machines would interpret digging up our electronic waste if they had no prior knowledge of our society.

[1] http://hisnameislong.com/category/blog/

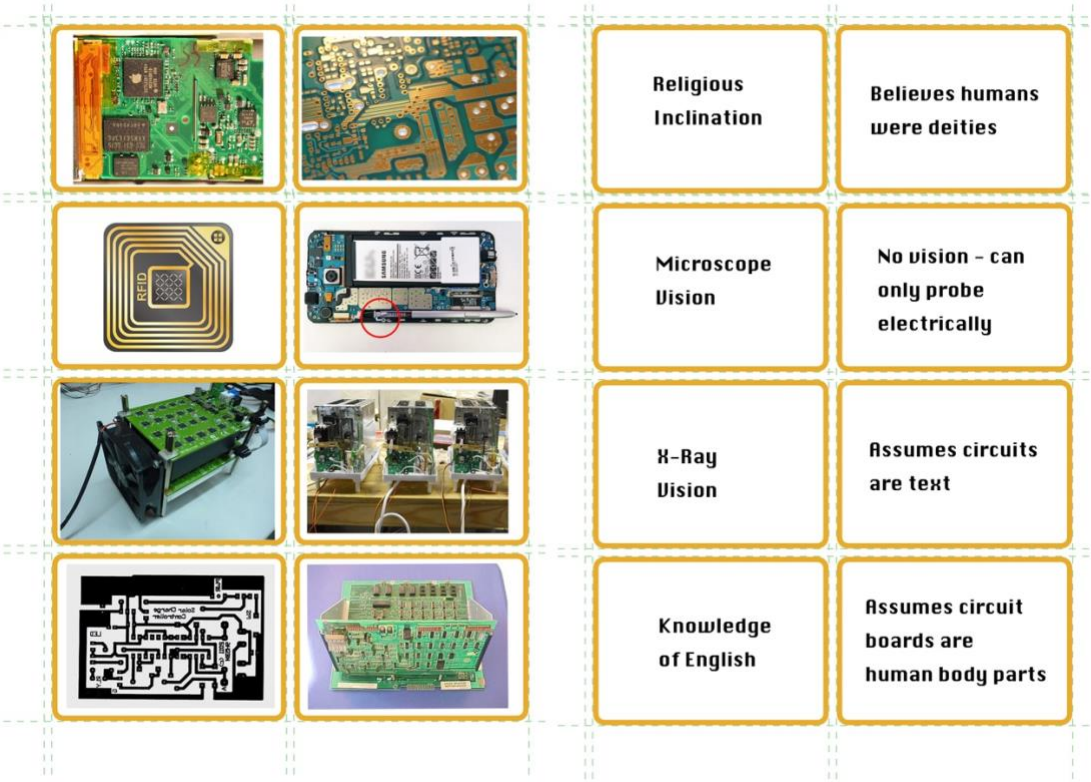

Figure 2. Left: 8 cards with images of electronic "archaeological finds" used to spur speculation about how machines would interpret them. Right: 8 cards to suggest beliefs and qualities of future robot archaeologists.

4.1.3 Result

The interactive nature of these non-text-based storytelling options was highly engaging and produced a wealth of questions, observations and ideas. Often classmates would laugh in surprise at a particular idea, but then ask a serious question.

It was clear that everyone I spoke with had plenty of thoughts, ideas, hopes and fears regarding the future of intelligent robots. I suspect there would have been less engagement if the technology under discussion had been aircraft design, supercomputing, automation, data science or almost any other topic touching technology alone. This may be a testament to the long tradition of stories and myths about creating artificial life, or to the effectiveness of speculative fiction.

There is much speculative fiction about human-machine societies, and movements that blend human and machine such as Cyberpunk and Transhumanism. While developing these prototypes, it occurred to me that there has been less discussion of a

future where machines *are* humanity's legacy, if humanity were to disappear. It also became clear that myth-making would offer a potentially powerful vehicle for contemplating this future society, its origins, and its iconography.

The notion of machine creation myths suggests that either the facts of machine intelligence's earliest days are not known to them, or the facts alone aren't sufficient for understanding their place in the world. They present an opportunity to consider to what extent machines development will be guided by humans and how much they will develop themselves. In the words of Adrienne Mayor's account of the earliest automatons, the question arises, "Were they made or born?" (28)

These explorations helped confirm that using speculative fiction is a potentially powerful vehicle for contemplating the nature of AI and our changing relationship to it. Although text-based stories can be complex and detailed, non-text formats may have the potential to be more open-ended, giving the audience room to bring their own questions and ideas. My next question was to experiment further with using image-making to tell these stories.

4.2 Making Myth Images

This image-making phase developed from the earlier speculative fiction work. Rather than creating images out of whole cloth using photography or painting, I wanted to find a means of assembling them out of ingredients that would already be available to machines. The ideation and creation started to become more visual than the mostly text-based speculative fiction writing in the last section.

4.2.1 Investigation

The process of elaborating myths from words into images began with a combination of sketching and text, an approach made famous by Leonardo da Vinci (Isaacson). Since these images are framed as stories told by machines, I made them using imagery that is available to machines in the present day. I gathered images that have been used to train neural networks, and images produced by those networks, and combined them using digital compositing[1].

[1] Compositing is an art form that began with collage and photo retouching. Digital compositing has been used in recent years in combination with photography by many digital artists such as Alexia Sinclair, Aaron Nace, Miss Aniela, Joel Grimes and Von Wong.

4.2.2 Experimental Work

I sketched about ten concepts of machine creation myths, and selected four from these to develop as composites. Each one is a separate story that suggests causes for humanity's disappearance that simultaneously explain the rise of the machine civilization. Figure 3, below, shows how the medium and techniques developed over time as I made successive composites.

Image 1

Image 2

Image 3

Image 4

Figure 3. Images made from GAN outputs and training images

Most of the component images for the first of these (figure 3, image 1) were extracted from the output of a neural network called BigGAN which accompanied a paper published by Brock et al. Their origins in a neural network fit my theme, and lent them an otherworldly quality. I used a stock photo as the background and the study portrays

machine life emerging from the sea. However, the elements were difficult to harmonize and the resulting collage-like quality seemed at odds with the mythical subject matter.

Figure 3 image 2, shows the spark of machine life being excavated from a buried repository of human data. To make it I composited several photos from a dataset called ImageNet used to train neural networks. It seemed fitting to me that distant future machines would attempt to portray the past by remixing images of the past that were part of the training of some of their earliest ancestors. The results had an odd combination of photorealism and symbolic qualities, which seem to work well when printed on reflective metal film. The low resolution of the ImageNet images was very limiting, and it was challenging and time consuming to find images with compatible lighting.

In image 3 in this group, depicting machine life developing from a crash-landed drone, I attempted to avoid using any images that were not from ImageNet or a neural network output. This turned out to place excessive constraints on the composition, lighting and resolution of the composite. I did not feel the resulting image was legible as a storytelling image.

At this time, I also experimented with creating a story by interpolating through the latent space that results from training a GAN (Brock's "BigGAN") on various categories of images. By generating hundreds of images that straddle multiple categories, it is possible to animate a transition from, say, a trilobite to a rocket ship. Figure 3, image 4 shows a still from such an animation purporting to show machine evolution. This approach seemed promising, but relied on digital display and did not lend itself to production as a non-electronic physical artifact, and therefore didn't integrate well with the physical images I produced later.

4.2.3 Result

The use of images from ImageNet and neural networks trained on ImageNet proved to be too limiting for producing expressive composites. In addition to the problem of low resolution, an image whose component parts seem to be lit or shot differently can be jarring, and takes the viewer away from the illusion. My composites sometimes seemed to be both too photorealistic or not photorealistic enough. If we are expecting photorealism to denote objectivity and documentary precision, then a photograph that is obviously composited may seem like an attempt to deceive. I discovered I also needed a way of increasing the resolution of training images and neural network outputs in order to print.

I remained positive about the possibilities for constructing mythological imagery from parts, but recognizing that I am not a painter or expert photo compositor, realized I would struggle to tell complex stories with which the viewer was not already familiar. I liked the approach of compositing using images associated with neural networks for its thematic resonance, but needed a more flexible approach that would supply me with images more suited to compositing. These observations encouraged me to investigate the possibility of a more engaged collaboration with neural networks for this image making. To begin with, I decided to try exerting more control over the component images of each composite by creating images using GANs.

4.3 Face Detection & GANs, Part 1

Starting from a speculative fictional setting, I asked what humans might look like to a machine that had never seen one. If digital records were lost, could they use well-preserved physical images like paintings? I knew that GANs had been used to generate new faces based on millions of photos of celebrities online. But could paintings of people be used in the same way? What if there were only a few thousand available rather than the millions of celebrity face pictures used in neural network research?

4.3.1 Investigation

Some artists such as Anna Ridler curate datasets by hand with which to train GANs, thereby avoiding generic results that come from using a generic dataset like ImageNet for training. This project explores how machines can collaborate on image-making, so I decided that rather than assembling data by hand or with a search engine, I would ask a neural network to help me build a large collection of faces taken from old paintings on wikiart.org.

Face detection algorithms have a relatively long history in computer vision and machine learning. Yann LeCun invented the convolutional neural network in 1998 (LeCun et al 1998) to help computers interpret handwritten addresses on mail for the post office. In 2001 an approach using Haar cascades was developed by Viola and Jones. By 2006, there were already Fujifilm and Nikon cameras with face detection built-in (Stafford).

I tried Narang's implementation of Viola-Jones's approach from OpenCV which uses Haar cascades, but it missed about 2/3 of the faces, and about 50% of the faces were false positives. I implemented Schroff's more recent approach, based on Sandberg and Jekel's Tensorflow implementation, which is trained on hundreds of thousands of images to

detect faces. Running this on a powerful cloud computer doubled the number of faces I detected and reduced the number of false positives to near zero. This experience was my first hint that getting good results from neural networks is not free, and that the advantage sometimes resides with groups having large budgets.

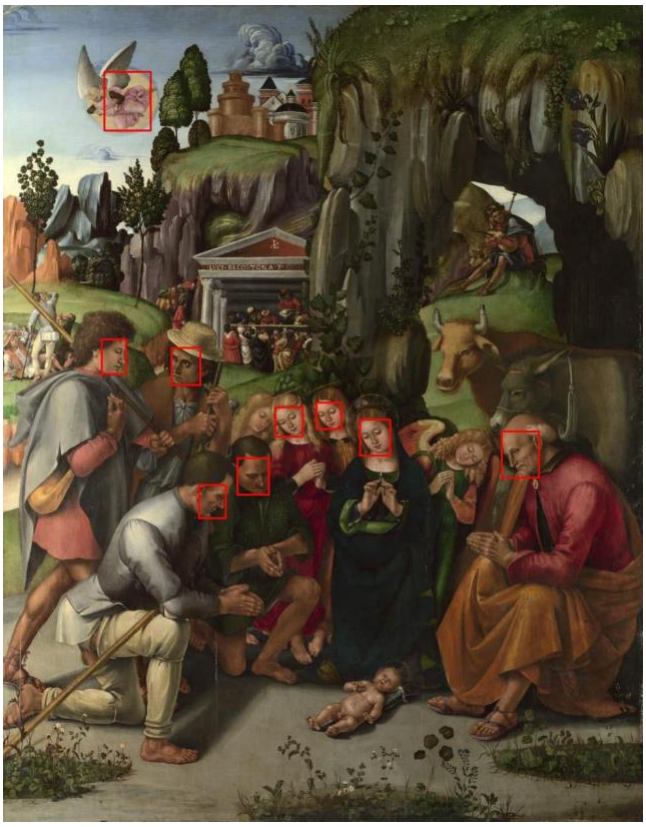

Figure 4. Facenet detecting faces from an early Renaissance painting. Note that it misses some human faces, and shows a false positive in the robes of the angel in the upper left.

4.3.2 Experimental Work

Having gathered a collection of nearly 2500 faces from early Renaissance paintings, drawings and sculpture, I began experimenting with training various GANs to produce new faces. I experimented with a wide variety of GANs of different types, all written by other people and shared on Github.com. These could not all be made to work well or at all, as they tend to require support libraries and hardware configurations which are incompletely documented, since these were written to support a research paper rather than general use. I was supported by Tushar Gupta, a recent computer science grad whom I contracted to help with troubleshooting via the website codementor.io. We eventually found a few implementations that we could run on my data.

CycleGAN

CycleGAN (Zhu) is a special type of GAN which, when trained on two separate image datasets, can convert any image in the style of the first dataset to the style of the second dataset. It can also convert an image in the style of the second dataset back to the style of the first dataset. Thus, it can successfully convert an image of a horse to an image of a zebra, or a winter scene to a summer scene, and back again to a winter scene.

By experimenting with a cycleGAN I discovered that by partially training a CycleGAN, the image transformation is accomplished partway, yielding an image that belongs to neither the original category nor the new category, but rather an intermediate state. I gathered around 2000 images from Nasa's Astronomy Picture of the Day site[1], using them as one category for the CycleGAN. For the second category, I gathered about the same number of early Renaissance pictures from website wikiart.org. Training for 20 epochs resulted in an undertrained CycleGAN, which produced an image partway between an image of Madonna and child and a globular cluster – a visual depiction of an extra-terrestrial creation myth I developed during an earlier speculative fiction phase. I did not want to bring in such pointedly Catholic imagery into the story, but this image served as a proof of principle for possible future creation. The low resolution of this and almost all other images created by GANs remained an issue, one which I confronted later on.

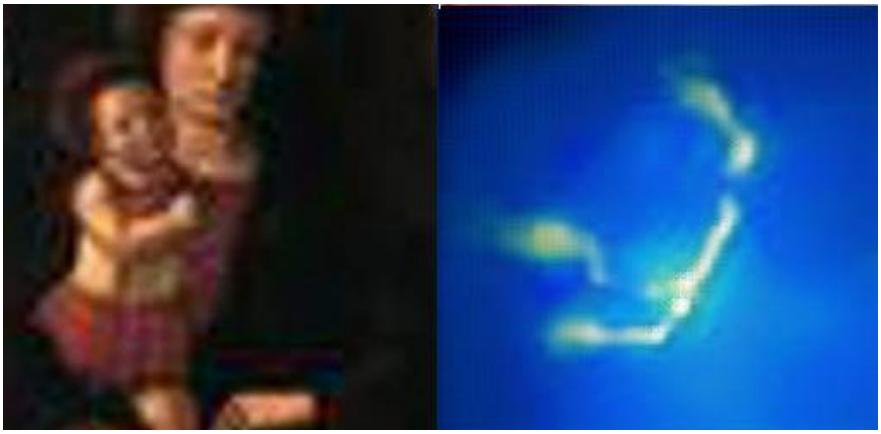

Figure 5. Madonna and Child from Early Renaissance training dataset (left), Globular Cluster from NASA APOD training dataset (right).

[1] https://apod.nasa.gov

Figure 6. Images produced by CycleGAN: Madonna and Child partially converted to astronomical images (left), Globular cluster partially converted to an early Renaissance image (right).

DCGAN for Renaissance Faces

Training a GAN on a selection of undifferentiated early Renaissance paintings resulted in incoherent blobs in colours appropriate to the pigments known at the time. Using the face detection neural network to grow a larger dataset of segmented faces, combined with using a different implementation of Deep Convolutional GAN (DCGAN) generated the first new faces for my myth images.

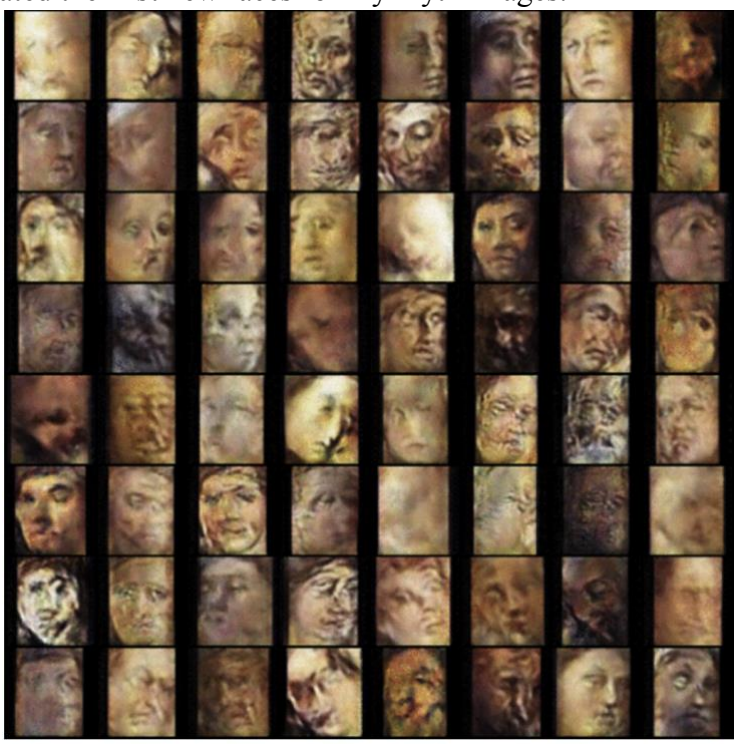

Figure 7. Grid of 8x8 images generated by DCGAN based on 2500 faces from paintings

4.3.3 Result

These two results suggested my collaboration was working. Lacking drawing and painting expertise, I have no way of producing new images in a genre of my choosing by any other means. The images did not yet tell a story, and were too low in resolution to print, but they were recognizable examples of the type of images I intended, and were not direct derivatives of the training data.

The practice of studying the theory and using these neural networks for face recognition and for image generation helped me understand what a neural network "sees" when it processes an image. The layers within the neural network are finely tuned to respond to the distinctive pattern of light and shadow that signifies a face. They compress this information down into fewer neurons than there are pixels in the image, suggesting that they have extracted the essence of the images without retaining all the details.

Without having a background in computer science or machine learning, I had struggled to get this far, and wondered if it would have been quicker to learn to paint. On the other hand, having established this collaboration, I could train it with more faces or other kinds of images, and could receive limitless new images as output in return. My machine collaborator was beginning to exhibit some of the same virtuosity and temperamental character of a human collaborator, though with a much more arcane interface.

In order to coax higher quality, higher resolution outputs, I would need to source a larger dataset and possibly a more sophisticated GAN. I tried various experiments to increase the resolutions of the outputs, from adding a layer to the output, to implementing Progressive Growing of GANs (Karras et al.), Wasserstein GANs and many others, without success. Some of the more sophisticated GANs require computational power greater by an order of magnitude than the already high-end machines I was renting by the hour. I reflected that my aim differed from those of machine learning researchers whose work I was studying. While they were extending the bounds of detail, sharpness and resolution, I was looking for an evocative face or pose that was expressive but which left something to the imagination. At first it seemed unhelpfully limiting to be restricted to 64x64 or 128x128 pixel images, but after a while working with these images, I came to appreciate their aesthetic more.

As a result of the various successful and unsuccessful GAN experiments, I decided to embrace the opportunities afforded by working at lower resolutions. At the same time, I discovered that artist Helena Sarin has resisted the trend of pushing GAN image resolution (Bailey), advocating for the increased expressiveness and freedom that comes from "small GAN". In a similar vein, artist Anna Ridler avoids the generic aesthetic that can result from everyone using the same large datasets to train GANs by hand-curating

her own training datasets. In a creative collaboration like my own, I found that working with smaller images increased the speed with which

I found that the generated faces didn't need to be extremely detailed to allow us connect with a pose and facial expression. At the same time, we want to see detail around the eyes and mouth. As with Leonardo da Vinci's *sfumato* technique perfected in the Mona Lisa's mysterious smile, when these details are unclear, faces may take on an intriguing ambiguity. In the case of these faces generated by a GAN, often cues about gender or about emotion are ambiguous, which adds an open-ended interest to their interpretation.

4.4 Increasing Expressiveness in GAN Images

Distinct from the issue of image resolution was the readability and expressiveness of the image's subject. I found that images meant to evoke mythological figures depend on having a compelling figure as their subject, but the faces produced by my GAN collaboration so far were excessively blurry and distorted, with relatively few interesting characters among many unreadable faces. My next series of experiments sought to increase the clarity, range and expressiveness of the figures. I found three techniques in particular that helped with this: transfer learning, inventing image detail with SRGAN, compositing outputs from multiple GANs, and these are detailed in the sections below.

Additional techniques that showed promise, but which I did not incorporate in the finished work, can be found in Appendix D.

4.4.1 Transfer Learning

It is sometimes possible to use a large dataset to pre-train a neural network, and then switch the training data to a smaller, more specific dataset for the last portion of the training. This technique, called transfer learning, can produce images that are a hybrid of the two datasets. I wanted to see if it could help produces images in the style of my Renaissance painting dataset, but with a stronger foundation in the variety, detail and depth of a larger dataset.

To this end, I used Robbie Barrat's "Art DCGAN", which is based on Soumith Chintala's DCGAN written in Torch (Chintala; Radford). This was the GAN that art collective Obvious used to produce a print that Christie's sold at auction for US$432,500 in 2018. Obvious was criticized for not crediting the researchers who created the GAN, or Barrat who trained it, which fueled further discussion of authorship in art produced with

GANs[1]. This implementation made it possible to change the training data at the end of training. I used this to supply the GAN trained on more general portrait paintings with my Renaissance faces dataset. Significant trial and error is needed to produce images that balance the old training data with the new, but in the end the results were encouraging. The resulting images adopted a palette representative of Renaissance painting, and some of the same facial characteristics (see figure 8, below).

Figure 8. This image, produced using transfer learning, takes on some qualities of a Renaissance silverpoint drawing after two epochs of training Barat's portrait network on the Renaissance faces.

The use of transfer learning made the results more unpredictable and often unusable, but when successful, the technique did provide me with more thematically complex and detailed images than I obtained using GANs trained on a single dataset. These results still lacked a degree of detail, especially around the eyes, which I felt would be needed to give the subject more presence and the appearance of life, particularly in larger physical prints.

4.4.2 Inventing Image Detail with SRGAN

Simply scaling up low resolution images produced by a GAN using traditional techniques such as Adobe Photoshop's "bicubic" resize resulted in a blocky, pixelated image. But GANs have also been used to increase image resolution, by guessing at appropriate fine detail to add to images that have been down-sampled to lower resolution (Ledig). I wondered if I could add interest and depth to an image starting with a low-

[1] Barrat's implementation was based heavily on Chintala's, which in turn was based on Goodfellow's paper. How inclusive should the attribution be? Questions of attribution are discussed further in the next chapter.

resolution output from a GAN and adding detail using this Super-resolution GAN (SRGAN) technique. Despite being trained for adding detail to photographs, it often produced good results with images made by GANs.

Figure 9. 128x128 image produced by Art DCGAN resized to 512x512 pixels using Adobe Photoshop Bicubic resize (left) and SRGAN (right)

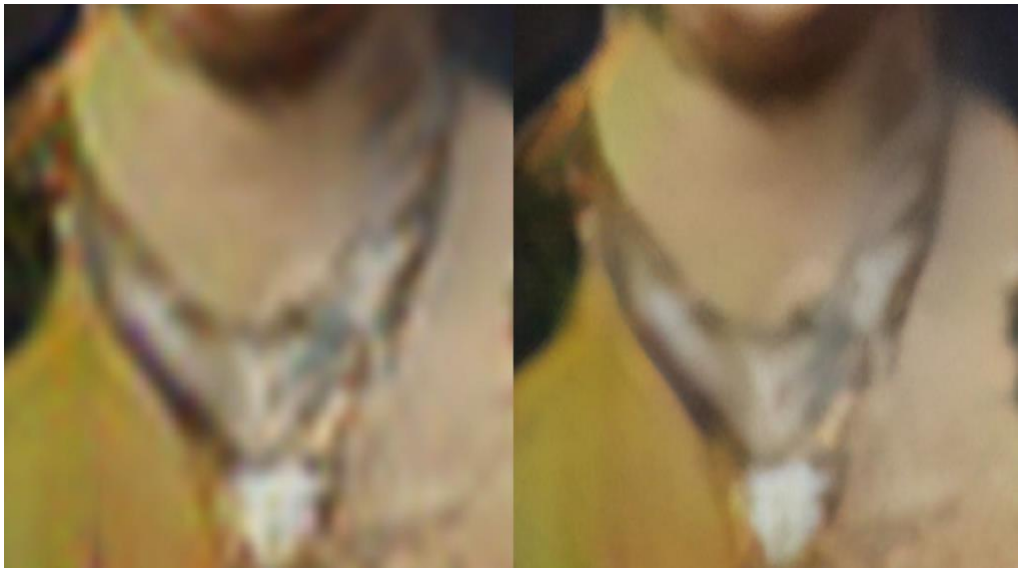

Figure 10. Detail of image above showing detail created by bicubic resize (left) and SRGAN (right). Both images were subsequently resized using Adobe Photoshop "bicubic sharper" algorithm for print.

Image uniformity, smoothness and sharpness are better with the GAN-resized image. The technique sometimes introduces artifacts such as a checkerboard pattern, for example in 64 x 64 pixel images. These artifacts could be interesting as the SRGAN tried to invent textures for ambiguous parts of the image. This GAN supplied an important function that

made it possible to produce prints of a modest size (e.g. 13 x 17cm) without obvious pixellation.

As a tool for collaboration, SRGAN was less interesting and interactive than the DCGANs that produced new figures and faces. I did not have a way of producing both high and low-resolution GAN images with which to train one specific to my application, so I was limited to adjusting basic settings. As a result, this felt more like a simple tool than a collaborator in this project, despite its capacity to sometimes produce inventive and interesting results.

4.4.3 Compositing

The transfer learning-trained GAN described in 4.4.1 was producing interesting characters (e.g. figure 8), but the faces were often too blurry and distorted to present clear emotion or presence, even with the detail added using and SRGAN described in the previous section. I experimented with combining outputs from multiple GANs to produce the desired effect, using one GAN to produce figures, and another to produce faces, combining them manually with Adobe Photoshop.

An example of one such composited image is shown in figure 11, below.

Figure 11. Portrait composited from GAN outputs.

I wanted to incorporate the expressive faces generated earlier to produce an image of a possible human forebear of intelligent machines. The human figure in figure 11 remains fluid as to time period and gender. I prepared a gilded copper circuit board panel, printed the figure on a metal film and attached it to the panel. I then used a computer controlled machine to inscribe a geometric motif of circuit traces to the gold, in a mechanized version of the technique used to draw haloes, names, and ornamentation around figures portrayed in Medieval and early Renaissance panel paintings. Showing this to classmates, I called the piece *Corporation* as a reference to corporeality and a machine's notion of the identity of its creator.

Upon showing this panel, the question of authorship came up again – "what part did you do, and what part was done by the machine?" The diagram below shows the various processes and stages that resulted in the final image, which incorporated a gilded background inscribed with a circuit trace pattern by computer controlled machine.

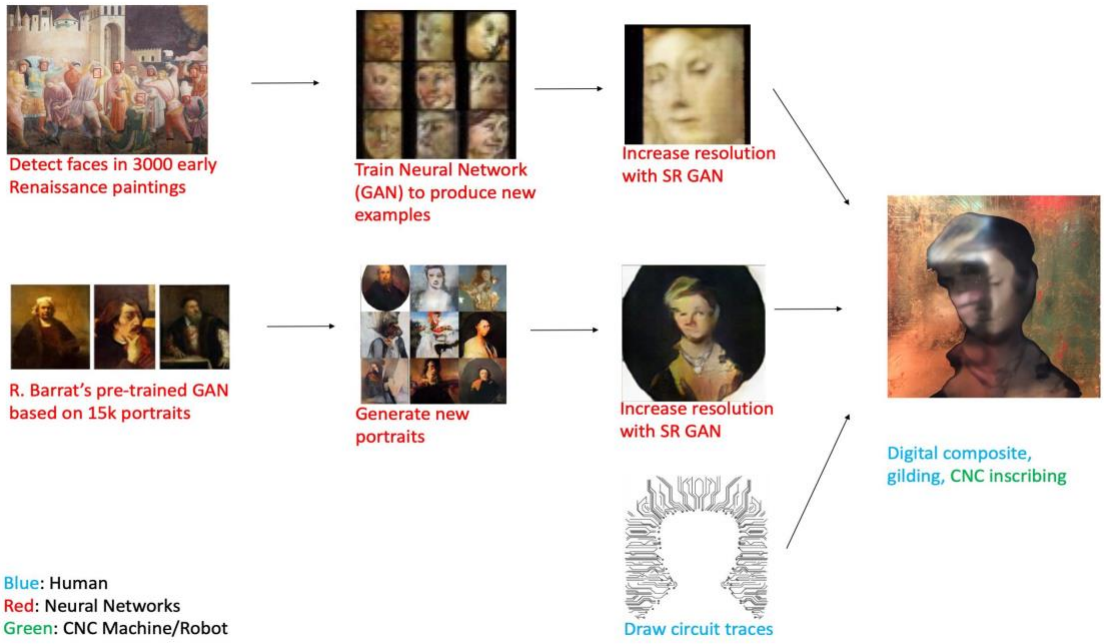

Figure 12. Diagram showing process of producing image shown in figure 11.

4.4.4 Result

It was eventually possible to make images that suggest an origin myth using neural networks. Several stages of my experimental image-making practice were made possible

by GANs, including various activities that would have been either very time-consuming to do manually, such as extracting thousands of faces from old paintings, or impossible for me, such as generating new faces.

At times, the experience was similarity to photography, in which the creator must attend to the requirements of the technology, and then tries to arrange the elements of the composition, while acknowledging that these are not completely under control. In film photography, that interplay culminates in seeing an image slowly appear on paper in the developer, prompting pleasant surprise or disappointment, and often another pass through the enlarger with different settings. In this work, the difficult and often time-consuming work of setting up the neural networks was offset by the delight and surprise I experienced when they finally produced images, which became the input for the next stage of creative practice.

The experience also differed from the usual experience of using technology to make images. My activities interleaved with the machines' activities with many instances of interaction and intervention from me, and influence and inspiration from the neural networks. The initial premise of the image-making was directed by the speculative fiction backstory, but I also found myself adjusting the story in response to compelling characters suggested by the GANs. The GANs provided limitations to the possibilities for image composition – I would not be making large narrative panels – but together we could produce engaging character portraits around which to hang story details. I found that despite their limitations, the GANs continued to produce interesting and unexpected imagery which energized the collaboration and propelled it forward.

When asked about authorship or attribution, I struggled to find a helpful way to express which of us did which part, and whom to credit for the various ideas and techniques. Upon reflection, dividing the work into phases as in figure 12 and attributing the different phases to different actors probably obscures more about the process than it illuminates. It is hard to describe with one-way arrows and boxes all the false starts, chance discoveries and surprise inspirations that characterized every stage of this collaborative image making process. I was also aware of the debt I owed to all the artists and researchers whose work made my output possible at every stage, and didn't know how to properly acknowledge their role. I began looking for a better way to communicate the authorship of this work.

4.5 Making the Digital Images Physical - CNC Gilding

Part of the premise of this image making activity is that future machines wanted to memorialize their myths in physical objects. Metal seemed like a suitably timeless and

evocative medium for machines to use to immortalize an image. Gold has the properties of being impervious to corrosion and malleable for inscribing, and has been used variously for its value, symbolism and utilitarian properties for human and machine purposes throughout history. If humanity were to disappear from the earth, after concrete and iron dissolved in rain the most durable parts of our buried waste would be the last artifacts to survive us - the glass fibre and gold used in circuit boards being among the most durable of those.

My use of programming computer controlled machines (also known as Computer Numeric Control or CNC) for scribing and punching images in gilded panels was partly inspired by my desire to enable machines to create durable images that could outlast our digital data, and was also inspired by a need for a collaborative partner in image making that had the patience, precision and attention to detail for physical image making that I lack.

4.5.1 Research for Creation

In the Middle Ages and early Renaissance in Europe, panel paintings such as altarpieces combined painted figures with gilded backgrounds before these fell out of favour as the aims of image-making changed (for example, Alberti and Sinisgalli). This laborious process, known as water gilding, involves the preparation of panels using many layers of animal-based glue, canvas, gesso and clay bole before gold and paint are applied. While still uncured, the gold leaf can be burnished to a shine with an agate stone (or preferably a wolf's tooth), and patterns such as halos, text and decoration are added using styli and punches (Frinta; Martin).

Water gilding is a traditional process for making golden backgrounds using precious materials suitable for devotional images. Unlike other forms of guiding (oil gilding, transfer leaf) and gold paint, water gilded surfaces can be burnished to a reflective shine, and can be decorated with punched and scribed decoration. Producing a panel takes a minimum of 3 days.

Further detail on water gilidng and inscribing can be found in Appendix E.

4.5.2 Creation as Research

Making use of computer controlled equipment already at hand, I next modified a 3D printer to hold a stylus in place of its print nozzle, and programmed it to trace patterns in the prepared panel.

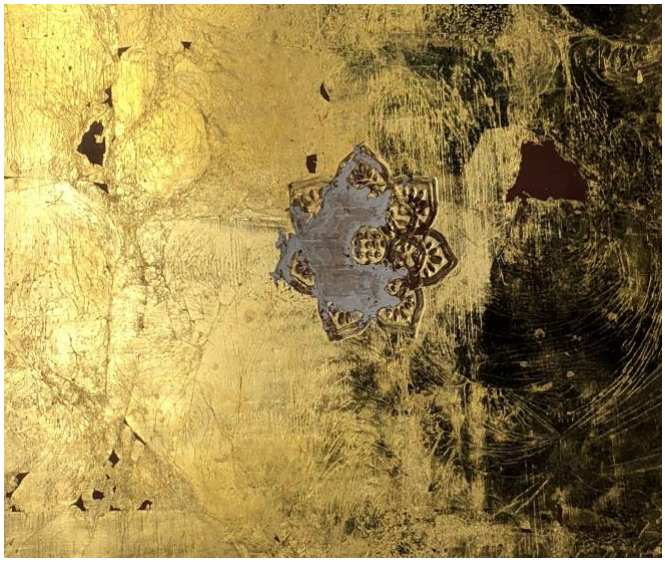

Fig 14. CNC scribing test on gilded panel showing torn gold and cracked gesso

This approach produced patterns that I would have struggled to produce by hand, but the lack of control over the pressure applied by the stylus meant that the surface was damaged, and the gesso and bole layers underlying the gold detached from the panel in places. The flex in the 3D printer, which is not constructed to tolerate coming into contact with the workpiece, reduced the precision of its movement.

To address these issues, I began developing a spring-loaded stylus to better control the force applied to the panel, and looked into other commercially available drawing machine kits. Finding that these were not usually suitable for applying significant force to the panel, I experimented with a robot arm (ST Robotics R12), which had the advantage of an anthropomorphic plan. This lacked precision and smoothness when drawing contours. I discovered I could get good results using a commercially available embossing machine called Silhouette Curio. I used this for the panel shown in Figure 15.

Figure 15. *Corporation* - finished piece showing inkjet print on metallized film cut out and affixed to a gilded panel, which was then inscribed using the Silhouette Curio.

The Silhouette Curio produces only modest force (about 2N or 200g of force) and is limited to a 20cm wide panel. But because it is easy to use, it is a very useful tool for quick experimentation and creation. This freed me up for the time-consuming task of making the gilded panels, leaving the Curio to do the inscription. The process of investing these panels with precious metals and painstaking work helped them take on greater symbolic weight the more I worked with them. Now that they were off the screen and gaining a more solid presence in physical reality, I began to feel that these could be the honoured figures from myth.

In order to overcome the collage-like effect of cutting out inkjet prints and adhering them to the panel, I developed a method of printing directly onto panel prepared with gesso before gilding. This process closely mimicked the traditional method of painting Christian altarpieces (National Gallery of Art and Heard; Skaug), but the figures are transferred to the panel using inkjet rather than by hand painting them.

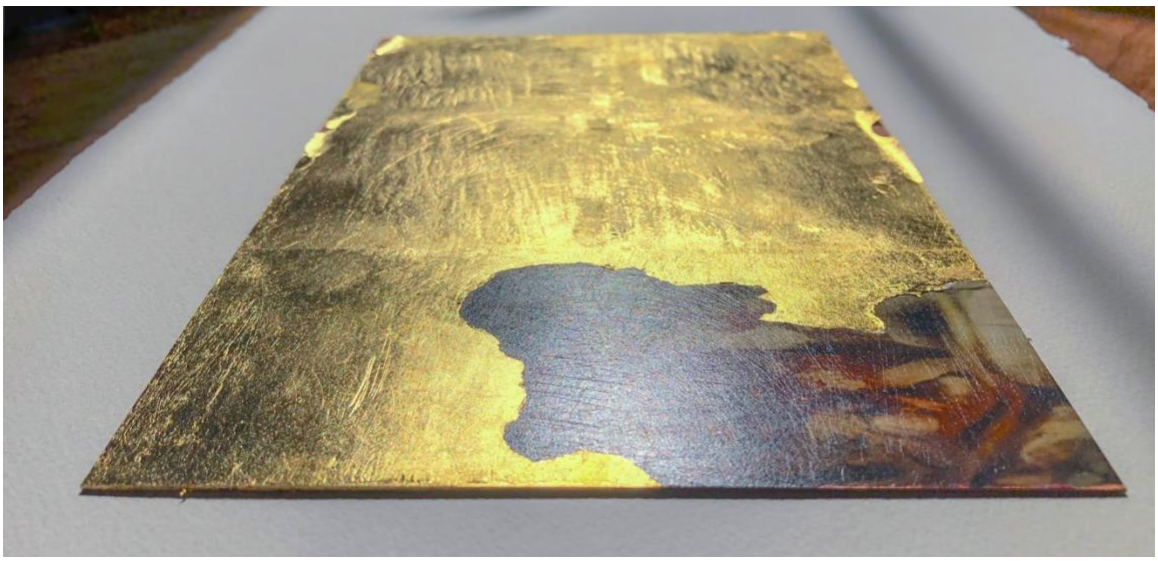

Figure 16. Thin circuit board panel with inkjet and gilding on gesso

In order to fit the panel through the inkjet photo printer (Epson R3000), the thickness of the panel including seven layers of gesso must be less than 1.3mm. Applying the glue size and gesso tends to warp the panel and I also encountered problems with the gesso flaking off the edges, which can be seen in figure 16.

These issues – adhesion of the gesso, the width restrictions of the Silhouette Curio, and the difficulty of producing such thin and flat gilded panels encouraged me to move on from this approach and try adding the gilded ornamentation more around the edges of a picture panel, by incorporating them into surroundings framing the figure.

Chapter 5 - Results

In this chapter I describe the collaborative production of the final work, and offer some observations on that process and where it leads. I used the collaborative toolkit produced in the previous chapter to create the final works for the thesis exhibition. This served as a chance to explore and evaluate the research questions posed in chapter 1, in the process Chapman and Sawchuk call research though creation. The exhibition itself serves to supplement these results in what Chapman and Sawchuk term a creative presentation of research.

My three-part methodology of speculative fiction, research creation and collaboration guided my objectives in creating the final work for exhibition. The studies and exercises conducted while building the collaborative process in chapter 4 suggested some promising avenues for storytelling when I came to reconsider the speculative fiction premise of the work I was creating for the final exhibition. I had found that the images I made were more effective at evoking interest and questions when they presented human figures with a few suggestive details in the background or in the title and caption, without attempting to describe a complicated backstory. I had also found that the GANs I was working with were capable of producing that kind of imagery, and that working together we had a good range of creative possibilities for producing figurative images on panels.

Moreover, I felt that creating these figures resonated well with the fictional premise of a future machine society speculating about its creators by illustrating figures from myth. I decided against presenting any of these stories didactically through text or sound, preferring to leave them open ended for the viewers to respond to.

5.1 Images for Final Exhibition

Figure 17 (L), Figure 18 (R) – Untitled creator figures from final exhibition

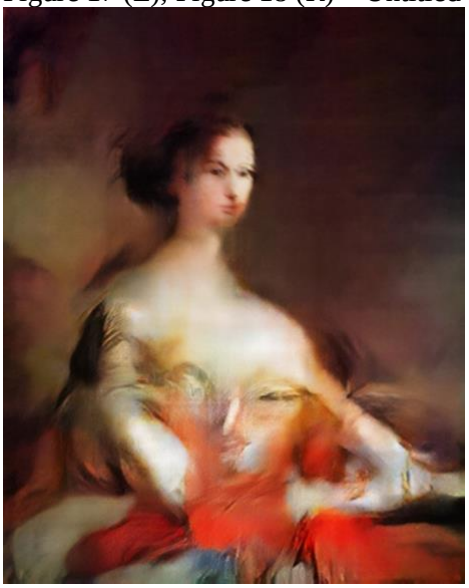

Figure 19 (L) and Figure 20 (R) – Untitled creator figures from final exhibition

These figures are all produced using by compositing faces produced from a GAN trained on Renaissance faces onto figures produced by a GAN trained using transfer learning. In the case of Figure 17, a style transfer GAN was used to blend the face with the body, and in all cases, an SRGAN was used for inventing image detail while increasing resolution. I made selections and tried various settings and compositions at each step. Figure 18 is depicted using a simulated gold leaf background rather than the exhibition print, which is difficult to photograph.

5.2 Reflection on Results

The production of these images helped me reflect further about what I wanted to say with these images and what I wanted to say about them in the final exhibition. Although I developed a mythological backstory for each figure, I decided to leave these images untitled, and let them mostly speak for themselves as the anonymous creators of machine intelligence. If in the future machines ever do wonder about their creators, it may be hard for them to ever know the names of the individuals and groups whose work produced them. I left these creators unnamed an acknowledgement of the countless artists, researchers and developers who are currently contributing to the development of every artificially intelligent machine, but who are mostly anonymous even in our own time.

In the same week that I was finishing these works, several further discussions about authorship and attribution in art involving AI were taking place. Artist Mario Klingemann sold a work at Christie's featuring GAN-generated images. Some of the press coverage referred to him as "computer programmer" rather than artist. Regarding another work, Twitter user @inconvergent suggested using "created by AI", in response to which Klingemann suggested "created with AI".

Salon.com published an article entitled "New AI art has artists, collaborators wondering: Who gets the credit?". The same week, Memo Akten tweeted, "We're still trying to fit everything into this ancient framework of ownership & authorship, so we can ultimately have (& idolize) a hero…". Appropriation is not a new idea in art, but these discussions seem more complicated with AI trained on thousands of inputs, whose outputs are immediately refashioned and remixed by anonymous individuals online. This adds further weight to the idea that rethinking authorship may require more than a simple adjustment of terminology, and may require a new way of thinking about creators and their works.

However we assign ownership, this collaboration worked well, and I look forward to continuing with it. With respect to my first research question, one of the fundamental qualities I believe made it successful was the need to stitch the collaborative toolkit together out of so many parts. This presented more opportunities during the process of making each image for me to observe, intervene, and gain inspiration for further creation. Before this project, my experience with neural network tools for image manipulation was limited to "black box" apps. Some of these provided a few settings, but the user intervention and interaction was minimal, resulting in a short-lived experience of novelty only. A few types of GANs have been packaged for ease of use[1], often in a commercial

[1] http://letsenhance.io uses SRGAN (Ledig) to increase image resolution, and http://deepart.io uses StyleGAN (Zhu et al) to modify a picture to the style of a painting.

product, but to date no DCGANs or variants have been packaged in a similar way for producing new images from scratch based on a training dataset.

One reason for this is that gathering thousands of images in a particular genre with enough diversity and enough similarity is not an easy or enjoyable task for most people. I believe that my approach of using neural networks to gather and filter large datasets may be helpful in this regard, e.g. for face and object recognition as described in section 4.2.1, or by mining large datasets for images with similar properties as described in Appendix D, section D.1.

Another reason why the practice of creating images from scratch using DCGANs and variants is not yet accessible to creators without coding experience is the need for significant computational resources. Some artists and researchers who work in this area opt to buy a dedicated graphics processing unit (GPU), typically in the US$500-$1000 price range. Others like me opt to rent time on a cloud computer, where an entry level setup costs about US$1/hr. Such a machine needs several hours of processing time for each run. Those seeking to replicate or build on the most recent and the most spectacular results from research such as Progressive Growing of GANs (Karras et al) will find that such a setup will require several weeks of continuous processing to produce a result – a result which may not be the desired one. The answer to this problem speaks to my second research question. I believe that artist Helena Sarin's "smallGAN" approach holds the key here, and it is the approach I used to give me flexibility and speed in image creation while keeping costs reasonable[1]. By reducing the depth of neural networks, and creating lower resolution images, resource needs are reduced to levels within reach of many more people. As research progresses, better neural network architectures and other technologies also reduce processing needs.

Originally I thought that I would create photo composites to tell a story, and would have a neural network create interesting patterns to decorate the background, imagining something like the geometric imagery sometimes produced in algorithmic art. Through practice, this arrangement was reversed, with the machine-generated imagery taking front stage, with me working around the edges to assemble the parts. I found that it was very helpful to my image-making process to have the machine develop a range of ideas within a theme that I specified, and start each image from one of those ideas. This was invaluable in helping me get started when I wasn't sure where to begin. Once I had something down as the first mark on the screen, I could elaborate from there.

In my photography practice, I have used look books, tear sheets and Google searched images when looking to gather ideas and inspiration for an image. I found GANs to be a

[1] I spend about US$500 on cloud computing resources between processing and storage charges during the course of this project, which included a lot of beginner mistakes.

potentially greater source of inspiration for different images and options, provided the low resolution and need to specific training data wasn't an obstacle.

In this project, I would often get to a point where an image wasn't working, or where I didn't know what to do next. Before abandoning it, I had the option to look back through the images the neural network had produced. Sometimes selecting a similar image from earlier or later in the training process would help me overcome problems with the composition, pose or expression of the subject. On a fundamental level, I never had to fear that I would run out of imagery since the machine is capable of making effectively infinite new examples. I was pleased to find the approach developed through this process gave me ideas and inspiration to carry on using this approach in future work.

Chapter 6 - Summary and Conclusion

Speculative fiction and AI collaboration served complementary roles, like the CPU and GPU[1] of this exploration into ways of relating to the blossoming field of AI. Along the way, I developed the beginnings of a toolkit for collaborating with neural networks on image-making, and deepened my creative practice by engaging with thinking through making.

At an early stage, I imagined that I would write stories, stage photographs, make composites from them, and then have neural networks develop the ornamental patterns for the gilded backgrounds. Instead, the neural networks established their image-making in the foreground, and left the details of background and framing to me. The images we were creating were driven by the speculative fiction, but the speculative fiction had to be adjusted to work with the kinds of images being produced.

My initial experiments with speculative fiction and speculative archaeology suggested I explore the possibility of combining tales of a future dominated by AI, a lost civilization reimagined, and the timeless qualities of myth to connect both of these to the present. Presenting these stories visually was challenging, and raised questions about how much to strive for a clear narrative, and how much to leave to the imagination of the viewer. Both the stories and the images were adapted to suit the qualities of this medium, marked by unusual, not-quite-human-looking portraits, and the challenges both digital and physical that prevented producing large didactic image panels.

Using a collaborative approach to image-making was helpful and my machine collaborator obligingly supplied endless variations on colour, pose and compositing. It responded to feedback uncomplainingly, indeed with indifference, and made adjustments as requested, asking only that it be fed with new and higher quality data in exchange.

I was particularly struck by the power of using a curated dataset developed from Renaissance art. As I extended the training data to include faces from later in the Renaissance, I saw an increase in the expressiveness of the faces being produced by the neural network, which gave me a broader selection of material with which to express personality and emotion when assembling composite portraits. For a more specific output, I needed only to narrow the scope of the training images I supplied.

[1] Central Processing Unit and Graphics Processing Unit, two elements of computer hardware responsible for running programs and processing neural networks respectively.

Making this work tangible through a labour-intensive physical process, involving both painstaking gilding of panels and developing a mechanized workflow to inscribe patterns was a helpful complement to what might otherwise be an exclusively digital process undertaken in front of a computer screen. Like the Renaissance images that inspired and fueled this work, the resulting objects, especially their gold features, may be of a durability to last decades or centuries.

When researching machine vision, image processing and developing large collection of digital images for training neural networks, I was struck by the incredible profusion of images our digital society produces. At one point, I quickly gathered over 2000 images of Christian icons. It would have been possible for me to feed them to a neural network and produce new images without ever even looking at them. This seems like a new level of abundance but also indifference that may be emerging from our digital visual culture. Training a neural network on this abundance, results in a machine that can produce an infinite number more examples, further eroding our motivation to examine any one of them closely.

My response to the deluge of images with which I was presented was to produce physical panels with an incredibly laborious process that helped me slow down and consider the significance of what's represented. These panels resist duplication. They need to be seen in person to be appreciated, in part because being reflective they resist being photographed. The punched and inscribed decoration appears dark at times and bright at others, depending on the light used to inspect it.

In this process, I wore many hats. As my role shifted from programming to story writing to image collaborator and back to troubleshooting programs, I sometimes felt overly torn or divided, and struggled to integrate these subjectively very different activities into something that felt like a seamless creative whole. However, once I had a few working neural networks producing some kind of output, the opportunity to switch between tasks allowed me to break up my work day, so that rather than trying to spend 8 or 12 hours writing stories or making images, I could spend a few hours in each area and get some helpful distance and perspective on the work in the other areas before returning to them the next day.

I do feel I developed an appreciation for the character of these machines, despite their uncongenial command line interface. Like tuning a radio into a faint station, when all the conditions are right, the technical details drop away and their muse is revealed.

In the final sections of this chapter (6.2 and 6.3), I would like to mention some further questions that were raised by this work, particularly about how we relate to AI and society. I am suggesting that when confronted with a new kind of entity, our best tools for

understanding how to relate may be found in our oldest stories and newest philosophies. I suggest why this might be important in section 6.1.

6.1 Our Changing Relationship to AI

The hands-on making has given me reason to think more about my relationship to AI. I did not find I anthropomorphized them or imagined they possessed superhuman intelligence. Each neural network's capabilities lie within the narrow domain of its training, but within that narrow band, they possess a creative capacity that distinguishes them from the mechanical and digital tools available to us in the past. I believe we are sometimes struggling fruitlessly to understand these machines as either intelligent or dumb, but what we are seeing is a new category of entity emerge in AI, that is neither human nor inanimate, and thus a new form of collaboration is becoming possible.

In the same paper in which he proposed the Turing test, Alan Turing noted tantalizingly that, "the best strategy for the machine may possibly be something other than imitation of the behaviour of a man." (436).

Although Turing dismissed it at the time, this turned out to be a prescient statement. Questions about machines' ability to think are often considered important (see for example Brockman et al.), while others see this question as irrelevant in the face of machines' increasingly astonishing non-human capabilities. Edsger Dijkstra said the question of whether machines can think is, "about as relevant as the question of whether submarines can swim" (E.W. Dijkstra Archive: The Threats to Computing Science (EWD898))

Indeed, many of AI's most striking accomplishments are those in which they seem the least human. One journalist writing about Deep Mind's self-taught *Go*-playing AI named AlphaGo said, "one gets the sense that an alien civilization has dropped a cryptic guidebook in our midst: a manual that's brilliant—or at least, the parts of it we can understand." (Chan)

This may seem contrary to observations if we consider the profusion of digital assistants and humanoid robots unveiled every year. William Poundstone notes that although machine intelligence can go in so many non-human directions, "Some people will want anthropomorphic machine intelligence." He mentions that Honda, Sony and Hitachi already expend substantial resources in making cute AI (Poundstone). This is an important quality for chat bots fielding customer service inquiries or robots helping with assisted living among the elderly, for example.

Anthropomorphic AI is not the only option, though it may have its uses. In addition, AI safety research is increasingly focused on the interpretability of AI ("Responsible AI

Practices"), suggesting that we need to deliberately engineer AI which we can comprehend, and avoid creating a world increasingly run by "black box" technology.

Jerry Kaplan of Stanford's AI research department summarizes some foreseeable impacts of robotics and AI entitled *Humans Need Not Apply*, which deals with the impacts on employment and society. He avoids predicting the extremes of techno-utopia and robot Armageddon, mixing original ideas with insights gleaned from interviews about the choices we face that will determine the nature of AI and robotics we are in the course of developing.

Thinking about whether machines have autonomy leads Kaplan to whether they then have legal rights and responsibilities. But how would you censure or punish a machine that breaks laws or ethical norms? Corporations have legal rights and responsibilities, and can be censured like individual citizens by separating them from their goals. But what goals would a machine have?

The question of machines' goals is important if we consider that machines may soon form some kind of society – either by joining ours, or creating their own. If they create their own society separate from ours, what will its attitude be towards us? It may only be possible to form a society together if we build in appropriate goals. So much of the law, ethics, and economics central to our society depend on understanding each other's goals and, if necessary, having the ability to separate individuals from their goals with various forms of censure, fines and imprisonment (Kaplan).

We can witness a microcosm of machine goals working and failing to work in GANs. Every training run for a GAN that I performed ended in mode collapse. The system passes through a point where it produces the highest quality and diversity of images. After that, further training results in the Generator network learning to "game" the Discriminator network, and get away with producing minor variations of the same few degenerate images over and over again.

There is a growing list of examples of a related phenomenon called "specification gaming" in which machines develop unexpected solutions to achieve the goals they have been programmed with, thereby subverting the actual goals of the human programmers (Krakovna). A robot vacuum programmed to minimize collisions with walls learns to drive itself around the house backwards because there is no collision sensor on its backside. An agent in a videogame learns to avoid "dying" by running so far away from its opponent that the system runs out of memory and crashes. Clearly this is done without malice, and thus these stories are amusing and chilling in equal measure at the moment, because these machines are not connected to vital infrastructure or military systems, as far as we know.

Neural networks' accomplishments are still modest by human standards, limited as they are by computing power, their architecture, and the scope of their training data. But they now have a capacity to say something original about the world as they see it, that can't or wouldn't be said by a human. This non-human quality is evident from visual inspection of the images they produce, which can be hauntingly familiar while being clearly non-human. Perhaps we can benefit from their independent viewpoint and their unexpected insights. But for all that we can benefit from getting to know their nature better, it seems important for safety and stability that we engineer in them an understanding and benevolence for us. As Isaac Asimov's *Robot* series made clear nearly 70 years ago, *how* to do that is not clear. Staying close to them and studying them intimately may alert us to warning signs and troublesome behaviour before we hand them the keys to our society.

6.2 Speculative Realism

We have grown accustomed to humans being the main actors in the world whose behaviour appears spontaneous, changeable and surprising. AI is now beginning to challenge us to think more and to think differently about other entities that share our world. Philosophers have something to say about how we may do this.

It came as a surprise to me that I did not feel tempted to name my collaborators, and even my attempts to name fictitious machines in my stories felt forced. Rather than feeling the urge to anthropomorphize these machines, I felt more of an urge to understand them on their own terms, as something made by humans but uncomfortably and irresolvably non-human.

The emerging philosophies of speculative realism use a less anthropocentric lens when considering the actions and interactions of non-living beings in the world, and suggest we consider things through the point of view of objects, nature and the environment in a way that has been overlooked.

I am not claiming that neural networks are alive, or that they have agency in the way we normally think. Nor am I claiming they are inert, predictable, dumb-as-a-post machines that should be treated as tools no different from a hammer, a tractor or even a conventional computer program. In speculative realism, there is a concerted interest in the idea that things in the world have an existence independent of their utility for humans.

This renewed focus on exploring a non-anthropocentric understanding of the world has taken on similar forms under various names, such as speculative realism, new materialism and object-oriented ontology. These approaches seem particularly well-suited to our time, because of the greater complexity of our world, with its staggering population, complex interdependent economies, exponential growth of communications

and the emergent properties these complex systems exhibit. Without religion, or a dominant body of myth, the west may depend on philosophy and art to navigate the future.

Continental philosophy seems to be dominated by renewed interest in materialist and, particularly realist approaches (see for example Manuel DeLanda and Graham Harman's The Rise of Realism). In addition, there has been a renewed interest in getting a wider perspective of nature and existence that attempts to step outside the human-centred lens. In part this may be simply to say that things exist outside how our human minds comprehend them, but also opens up many possibilities that are by definition beyond our sight.

I think speculative realism's "anthrodecentrism" is an essential step towards a modern and appropriate philosophy for the 21st century. As speculative realism continues to develop, more insights may emerge. AI as it exists today is beginning to challenge the limits traditionally ascribed to inanimate objects. For the moment, I think the key insight is that we could keep an open mind about possibilities beyond simply animate vs in-animate.

6.3 Machine Mythology and Spirituality

This project started out as a way of telling human stories through the eyes of machines, but developed into a story of machines using myth to discuss their origins, purpose and aspirations reminiscent of human spirituality. Discussion of myth and spirituality among machines may have more value than simply entertainment. The initial explorations of myth in this project enable us to ask further questions about the possibility of myth and spirituality in machines.

The oldest human myths about artificial life can be considered "what if" thought experiments where the constraints of contemporary technology were lifted. In the earliest stories we have, artificial life depended on gods and magical possibilities. For purposes of this project, in the machines' speculation about their origins, they must question whether they were entirely created by other beings, evolved from simple life, emerged from information and energy within the earth, were set in motion by others or started themselves. As discussed in the previous section, the question of whether their development is self-guided or directed is an important one for the future of humanity

In future work, I would like to explore further the idea of spirituality in these machines. It is not depicted as a prominent aspect of their creation myths, but may be considered a force that lies behind their values and beliefs, revealed through their images.

Setting aside the speculative fictional aspects of this project and further to the question of spirituality, with respect to present-day collaboration with machines, Denis Vidal argues that there are valuable insights to be gained when studying human-robot interaction by considering rituals enacted between humans and gods in the Himalayas. This surprising claim seems to have some merit when we consider the parallels. People relating with gods, spirits, ancestors and so on are aware of the usefulness and the limitations of anthropomorphizing. To understand where he is coming from, we should not think of the monotheistic traditions wherein heavenly beings always exist at a largely inaccessible level of hierarchy above us. In certain Himalayan traditions and many other animist cultures, spirits and the like can be asked direct questions, can be called upon for help and will respond to us directly, with the caveat that their responses may not be what we expect. We therefore consider carefully when to invoke them, and what we ask them to do. There are certain similarities with our present-day experiences with artificial intelligence. Perhaps our best models for how to relate to a being of unknown power, intelligence and disposition can be found in religious rituals.

As a robotics researcher, Masahiro Mori clearly felt that religion had some bearing on robotics research and vice versa. In addition to coining the term "The Uncanny Valley", which describes the special visceral effect robots can have on humans, his 1981 book *The Buddha in the Robot* asks what robots can show humans about spirituality by treating questions like ego as engineering problems

In a 2014 interview, AI pioneer Marvin Minsky reflected on the concept of the soul as the idea of who we are and where we came from, and connects that with creation myths, which he says could arise in machines.

> If you left a computer by itself, or a community of them together, they would try to figure out where they came from and what they are. If they came across a book about computer science they would laugh and say "that can't be right." And then you'd have different groups of computers with different ideas. (Elis)

Ray Kurzweil's *The Age of Spiritual Machines* makes the claim that machines may one day have spiritual experiences (6), but it seems that it will remain difficult to know with certainty what is experienced except by oneself. The resulting discussion of whether such a thing is even theoretically possible is probably more important than whether the claim is verifiable.

In *Gods and Robots*, Adrienne Mayor relates a remark in John Sladek's 1983 sci-fi novel Tik-Tok that, "the very idea of an automaton leads one into 'deep philosophical waters,' posing questions of existence, thought, creativity, perception, and reality."

These explorations have the potential to help us avoid simply seeing machine life as completely the same as us or completely different, and may help us understand the role that ethics, sympathy, compassion, identity and imagination play in developing human-machine collaborations.

6.4 Conclusion

This work was inspired by my desire to engage in image-making as a way of considering possibilities for the future of AI. Finding that my existing tools such as photo compositing were limited, I discovered opportunities to recruit neural networks to create the images I had in mind. I found that using GANs was a helpful way to create and construct develop images, particularly when they were part of a process of multiple parts, which offered me opportunities to shape and direct the output at several stages. Sometimes I led the collaboration in the direction I wanted to go but at other times, to my surprise, the neural networks offered imagery that showed the way forward.

Some of the difficulties of working with these techniques such as the need for large quantities of training data can be addressed by bringing in other neural networks to help with the image processing. The low-resolution outputs of the GANs that are within the reach of many people may be discouraging at first, but can be augmented by still further GANs which increase resolution and add interesting detail.

Working closely with neural networks also gave me some insight into current questions and challenges in that field, and an appreciation for how it can function alongside people without need to exclude them or replace them.

I also discovered that collaborating with neural networks raises, for me and for others, difficult issues of attribution and authorship. It seems that our traditional way of thinking about a work's creator are being challenged and upset by these new, more complex, more interactive processes which use technology created by a large number of people, and trained using numerous other artworks. These issues are being encountered in other arenas, such as responsibility for traffic accidents involving self-driving cars. My explorations suggest we need to consider our machine collaborators neither as intelligent beings nor as an inanimate tools, but as a new kind of entity altogether. In relating to a new more complex world, our mythology, spiritual traditions and philosophy may be our best guides.

Works Cited

Alberti, Leon Battista, and Rocco Sinisgalli. *Leon Battista Alberti: On Painting: A New Translation and Critical Edition*. Cambridge University Press, 2011.

Asimov, Isaac. *I, Robot*. Bantam Books, 2004.

---. "Nightfall." *Astounding Science-Fiction, September 1941*, edited by John W Campbell, Street and Smith Publications, 1941, http://www.isfdb.org/cgi-bin/title.cgi?46434.

Auger, James. "Speculative Design: Crafting the Speculation." *Digital Creativity*, vol. 24, no. 1, Mar. 2013, pp. 11–35. *Taylor and Francis+NEJM*, doi:10.1080/14626268.2013.767276.

Bacigalupi, Paolo. *The Windup Girl*. Night Shade Books, 2011.

Barrat, Robbie. *Modified Implementation of DCGAN Focused on Generative Art. Includes Pre-Trained Models for Landscapes, Nude-Portraits, and Others.: Robbiebarrat/Art-DCGAN*. 2017. *GitHub*, https://github.com/robbiebarrat/art-DCGAN.

Biggs, Michael A. R., and Daniela Büchler. "Rigor and Practice-Based Research." *Design Issues*, vol. 23, no. 3, July 2007, pp. 62–69. *Crossref*, doi:10.1162/desi.2007.23.3.62.

Bleecker, Julian. "Design Fiction: A Short Essay on Design, Science, Fact and Fiction." *Near Future Labratory*, Mar. 2009.

Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. First edition, Oxford University Press, 2014.

"Brazen Head." *Wikipedia*, 22 Sept. 2018. *Wikipedia*, https://en.wikipedia.org/w/index.php?title=Brazen_head&oldid=860766816.

Brock, Andrew, et al. *Large Scale GAN Training for High Fidelity Natural Image Synthesis*. Sept. 2018. *arxiv.org*, https://arxiv.org/abs/1809.11096.

Brockman, John, et al., editors. *What to Think About Machines That Think: Today's Leading Thinkers on the Age of Machine Intelligence*. First edition, Harper Perennial, 2015.

Cahill, James, et al. *Flying Too Close to the Sun: Myths in Art from Classical to Contemporary*. Phaidon Press, 2018.

Campbell, Joseph, and Bill D. Moyers. *The Power of Myth*. 1st ed, Doubleday, 1988.

Cellan-Jones, Rory Cellan-Jones, Dave Lee, Rory. *Hawking: AI Could End Human Race*. 2 Dec. 2014. *www.bbc.com*, https://www.bbc.com/news/technology-30290540.

Chan, Dawn. "The AI That Has Nothing to Learn From Humans." *The Atlantic*, 20 Oct. 2017, https://www.theatlantic.com/technology/archive/2017/10/alphago-zero-the-ai-that-taught-itself-go/543450/.

Chapman, Owen B., and Kim Sawchuk. "Research-Creation: Intervention, Analysis and 'Family Resemblances.'" *Canadian Journal of Communication*, vol. 37, no. 1, Apr. 2012. www.cjc-online.ca, doi:10.22230/cjc.2012v37n1a2489.

Chintala, Soumith. *A Torch Implementation of Http://Arxiv.Org/Abs/1511.06434: Soumith/Dcgan.Torch*. 2015. 2019. *GitHub*, https://github.com/soumith/dcgan.torch.

Chung, Sougwen (愫君). *Drawing Operations (2018)*. https://sougwen.com/project/drawingoperations2018. Accessed 10 Mar. 2019.

Colton, Simon. *The Painting Fool*. http://ccg.doc.gold.ac.uk/ccg_old/papers/colton_tpfchapter12.pdf. Accessed 24 Sept. 2018.

"Conway's Game of Life." *Wikipedia*, 8 Apr. 2018. *Wikipedia*, https://en.wikipedia.org/w/index.php?title=Conway%27s_Game_of_Life&oldid=835332919.

Desai, Shubhang. "Neural Artistic Style Transfer: A Comprehensive Look." *Medium*, 14 Sept. 2017, https://medium.com/artists-and-machine-intelligence/neural-artistic-style-transfer-a-comprehensive-look-f54d8649c199.

Dunne, Anthony, and Fiona Raby. *Speculative Everything: Design, Fiction, and Social Dreaming*. The MIT Press, 2013.

Earnshaw, Rae. "Collaboration Methodologies in Art and Design." *Research and Development in Art, Design and Creativity*, edited by Rae Earnshaw, Springer International Publishing, 2016, pp. 31–46, doi:10.1007/978-3-319-33005-1_3.

Elgammal, Ahmed, et al. "CAN: Creative Adversarial Networks, Generating 'Art' by Learning About Styles and Deviating from Style Norms." *ArXiv:1706.07068 [Cs]*, June 2017. *arXiv.org*, http://arxiv.org/abs/1706.07068.

Elis, Niv. "For Artificial Intelligence Pioneer Marvin Minsky, Computers Have Soul - Business - Jerusalem Post." *Jerusalem Post*, 13 May 2014,

https://www.jpost.com/Business/Business-Features/For-artificial-intelligence-pioneer-Marvin-Minsky-computers-have-soul-352076.

Euripides., and Murray Gilbert. *The Medea of Euripides / Translated into English Rhyming Verse, with Explanatory Notes, by Gilbert Murray*. George Allen London, 1907.

E.W. Dijkstra Archive: The Threats to Computing Science (EWD898). https://www.cs.utexas.edu/~EWD/transcriptions/EWD08xx/EWD898.html. Accessed 18 Nov. 2018.

Girardin, Fabien. "Design Fiction." *Wikipedia*, 3 Aug. 2018. *Wikipedia*, https://en.wikipedia.org/w/index.php?title=Design_fiction&oldid=853259294.

Gommel, Matthias, et al. *Robotlab Bios*. 2018 2014, http://www.robotlab.de/index_en.php.

Goodfellow, Ian, Yoshua Bengio, et al. *Deep Learning*. MIT Press, 2016, http://www.deeplearningbook.org.

Goodfellow, Ian et al. "Generative Adversarial Networks." *ArXiv:1406.2661 [Cs, Stat]*, June 2014. *arXiv.org*, http://arxiv.org/abs/1406.2661.

Gower, John. *Confessio Amantis: Book 4*. University of Rochester, 1390, http://d.lib.rochester.edu/teams/text/peck-gower-confessio-amantis-book-4.

Griffiths, David, et al. *Naked on Pluto*. 23 Oct. 2010, https://archive.bleu255.com/nakedonpluto/.

Haraway, Donna Jeanne. *Simians, Cyborgs, and Women: The Reinvention of Nature*. Routledge, 1991.

Herath, Damith, et al. *Robots and Art*. Springer Berlin Heidelberg, 2016.

Holtorf, Cornelius. *Future Archaeologies: Writing Speculative Fabulation (SF) from Fragments of Evidence – Sand14*. 20 Mar. 2017, https://sand14.com/event/future-archaeologies-writing-speculative-fabulation-sf-from-fragments-of-evidence/..

Jekel, Charles. *Charles Jekel Github*. 2017. *GitHub*, https://github.com/cjekel/cjekel.github.io.

Kaplan, Fred. "Cybersecurity's Debt to a Hollywood Hack." *The New York Times*, 21 Feb. 2016, https://www.nytimes.com/2016/02/21/movies/wargames-and-cybersecuritys-debt-to-a-hollywood-hack.html?_r=0.

Kaplan, Jerry. *Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence*. Yale University Press, 2015.

Kingsley, Danny. "The Journal Is Dead, Long Live the Journal." *On the Horizon*, vol. 15, no. 4, Oct. 2007, pp. 211–21. *Crossref*, doi:10.1108/10748120710836237.

Kubrick, Stanley. *2001: A Space Odyssey*. 1968. *www.imdb.com*, http://www.imdb.com/title/tt0062622/.

Langr, Jakub, and Vladimir Bok. "GANs in Action." *Manning Publications c/o Livebooks*, https://livebook.manning.com/#!/book/gans-in-action/chapter-3/v-2/. Accessed 20 Oct. 2018.

Lanier, Jaron. "My Problem with Agents." *Wired*, Nov. 1996. *www.wired.com*, https://www.wired.com/1996/11/myprob/.

Ledig, Christian, et al. "Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network." *ArXiv:1609.04802 [Cs, Stat]*, Sept. 2016. *arXiv.org*, http://arxiv.org/abs/1609.04802.

Levine, Davis. "Design Fiction." *Digital Experience Design*, 13 Mar. 2016, https://medium.com/digital-experience-design/design-fiction-32094e035cd7.

Lovelace, Ada King, and Betty A. Toole. *Ada, the Enchantress of Numbers: A Selection from the Letters of Lord Byron's Daughter and Her Description of the First Computer*. 2nd ed, Strawberry Press; Orders to Critical Connection, 1998.

Markoff, John. *Machines of Loving Grace: The Quest for Common Ground between Humans and Robots*. 2015. *Open WorldCat*, http://banq.lib.overdrive.com/ContentDetails.htm?id=4728F875-C003-498F-B5CD-60F035ABC73F.

Mayor, Adrienne. *Gods and Robots: The Ancient Quest for Artificial Life*. Princeton University Press, 2018.

Melbourne School of Design. *Liam Young - Brave New Now*. 2014. *YouTube*, https://www.youtube.com/watch?v=rSqLmEuUo_U.

Memory of Mankind. https://www.memory-of-mankind.com/. Accessed 8 Feb. 2019.

Michaelson, Jay. "Golem." *My Jewish Learning*, https://www.myjewishlearning.com/beliefs/Issues/Magic_and_the_Supernatural/Practices_and_Beliefs/Supernatural_Beings/Golems.shtml. Accessed 11 Nov. 2018.

Miller, Walter M. *A Canticle for Leibowitz: A Novel*. 1st ed., Lippincott, 1960.

Mori, Masahiro, et al. *The Buddha in the Robot: A Robot Engineer's Thoughts on Science and Religion*. Kosei Publ, 1999.

Nace, Aaron. "Sketching for Composites from Adobe Photoshop Compositing: Essential Techniques with Aaron Nace." *CreativeLive*, https://www.creativelive.com/class/photoshop-compositing-essentials-techniques-aaron-nace/lessons/sketching-for-composites. Accessed 8 Feb. 2019

Neutres, Jérôme, et al., editors. *Artistes & robots: Grand Palais, Galeries nationales, 5 avril - 9 juillet 2018*. RMN-Grand Palais, 2018.

Niedderer, K., et al. *The Research Exhibition: Context, Interpretation and Knowledge Creation*. IADE, 2006. *uhra.herts.ac.uk*, http://uhra.herts.ac.uk/handle/2299/7475.

Niedderer, Kristina, and Seymour Roworth-Stokes. *The Role and Use of Creative Practice in Research and Its Contribution to Knowledge*. 2007, p. 18.

O'Connell, Mark. *To Be a Machine*. 2018.

Pasztory, Esther. *Thinking with Things: Toward a New Vision of Art*. University of Texas Press, 2005.

Poundstone, William. "Can Submarines Swim?" *Edge.Org*, 2015, https://www.edge.org/response-detail/26043.

Radford, Alec, et al. "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks." *ArXiv:1511.06434 [Cs]*, Nov. 2015. *arXiv.org*, http://arxiv.org/abs/1511.06434.

Responses: What Do You Think About Machines That Think? https://www.edge.org/responses/what-do-you-think-about-machines-that-think. Accessed 23 July 2018.

"Responsible AI Practices." *Google AI*, https://ai.google/education/responsible-ai-practices/. Accessed 25 Feb. 2019.

Ridler, Anna. "GANs in Art." *ANNA RIDLER*, http://annaridler.com/gans-in-art/. Accessed 18 Nov. 2018.

Sandberg, David. *Face Recognition Using Tensorflow.* 2016. *GitHub*, https://github.com/davidsandberg/facenet.

Schön, Donald A. *The Reflective Practitioner: How Professionals Think in Action*. Basic Books, 1983.

Schwab, Katharine. "A Google Intern Built the AI behind These Shockingly Good Fake Images." *Fast Company*, 2 Oct. 2018, https://www.fastcompany.com/90244767/see-the-shockingly-realistic-images-made-by-googles-new-ai.

Shane, Janelle. "Imaginary Worlds Dreamed by BigGAN." *AI Weirdness*, http://aiweirdness.com/post/178619746932/imaginary-worlds-dreamed-by-biggan. Accessed 11 Oct. 2018.

Somers, James. "The Scientific Paper Is Obsolete." *The Atlantic*, Apr. 2018. *The Atlantic*, https://www.theatlantic.com/science/archive/2018/04/the-scientific-paper-is-obsolete/556676/.

Stafford, Ben. "Fujifilm Face Detection, Finepix S6000fd, Finepix F20 Digital Cameras Announced." *Digital Camera Review*, 13 July 2006, http://www.digitalcamerareview.com/news/fujifilm-face-detection-finepix-s6000fd-finepix-f20-digital-cameras-announced/.

Sullivan, Graeme. "Research Acts in Art Practice." *Studies in Art Education*, vol. 48, no. 1, 2006, pp. 19–35. JSTOR, *JSTOR*, doi:10.2307/25475803.

Tuchman, Barbara W. *A Distant Mirror: Calamitous 14th Century*. Alfred A. Knopf, 1979.

Turing, A. M. "Computing Machinery and Intelligence." *Mind*, vol. LIX, no. 236, Oct. 1950, pp. 433–60, doi:10.1093/mind/LIX.236.433.

Vidal, Denis. "Anthropomorphism or Sub-Anthropomorphism? An Anthropological Approach to Gods and Robots." *The Journal of the Royal Anthropological Institute*, vol. 13, no. 4, 2007, pp. 917–33.

von Nettesheim, Agrippa, et al. *Three Books of Occult Philosophy*. 1st ed, Llewellyn, 1993.

Weizenbaum, Joseph. "ELIZA - A Computer Program for the Study of Natural Language Communication between Man and Machine." *Communications of the Association for Computing Machinery*, vol. 9, 1966, pp. 36–45.

Wiblin, Robert. "Positively Shaping the Development of Artificial Intelligence." *80,000 Hours*, 1 Mar. 2017, https://80000hours.org/problem-profiles/positively-shaping-artificial-intelligence/.

Wiener, Norbert. *Cybernetics or Control and Communication in the Animal and the Machine*. 2. ed., 14. print, MIT Press, 2007.

Young, Liam. "Brave New Now." *Architectural Design*, vol. 83, no. 3, May 2013, pp. 74–81. *Crossref*, doi:10.1002/ad.1593.

---. *Under Tomorrows Sky*. http://www.mu.nl/en/exhibitions/under-tomorrows-sky. National Museum of the American Indian, New York.

Appendices

Appendix A - Background on Neural Networks in Image Manipulation and Image Generation

Since 2010, neural networks have begun to provide high tech, processing-intensive photo manipulation that is considered artificially intelligent because it does things otherwise only possible for humans. Examples include modifying the facial expression of the subject of a photo, or rendering a photograph in the style of a particular painter based on a single image exemplifying that style. Some of these have been packaged as easy to use websites and apps, such as Facetune and deepart.io, with the significant processing demands handled by cloud computing. While other computer vision and image manipulation techniques that depend on neural networks have achieved widespread use for utilitarian purposes such as license plate recognition and industrial robot navigation, the uptake of these technologies for image-making has been limited.

For the moment, the ability to modify faces or apply a painterly style to a photograph has mostly novelty value. Techniques that cannot be used with pre-trained neural networks are not accessible to non-programmers. Mainstream photo manipulation software such as Adobe Photoshop incorporates procedural art making techniques (e.g. oil paint filter, generate clouds, generate flames), but does not incorporate neural network techniques such as style transfer (which can modify a photo into the style of a given painter) or image analogies (used to adjust facial expression, hair colour or gender presentation).

If photo manipulation techniques are struggling to get a foothold with major software makers, techniques to make images from scratch are even less well known. This ability was demonstrated by Deep Convolutional Generative Adversarial Networks (DCGANs) 2014 (Goodfellow et al), and has developed steadily since then. Its variants WGAN, PGAN, StackGAN, BigGAN and so on have continued to push the variety, resolution and photorealism of the output. Despite their power and promise, no such techniques have been released as commercial products or for non-specialist use. This may be because their commercial applications seem much more limited than computer vision techniques for object and face recognition, and may also relate to the need for large image datasets and computing power for training.

The development of DCGANs for image generation depended on two other important concepts - convolutional neural networks and backpropagation. Convolutional neural networks, inspired by the architecture of the visual cortex in mammals, have found particular application in tasks involving visual recognition such as object classification. Backpropagation is the computationally intensive process of gradually adjusting connection weights between the artificial neurons in a neural network to reduce the

difference between the output and the desired output, thereby “training” the network.

Given a set of training images, typically thousands or millions of images, two networks, a Generator and a Discriminator are trained. The Generator begins by producing an image of random noise. The Discriminator is tasked with pronouncing whether the resulting image could have come from the training dataset, or is a “fake” image produced by the Generator. It is easy at first for the Discriminator network to figure out simple rules for telling the difference. The Generator connection weights are tuned by the process of having its fakes discovered by the Discriminator, to make better and better forgeries. As the Generator improves, so too does the Discriminator learn to be more and more discerning about whether an image is real or fake.

This process depends on large training datasets, powerful computers (typically with powerful GPUs for their parallel processing capability) and the careful selection settings (technically hyperparameters). Hyperparameters control how readily the weights defining each node of the neural network are adjusted in training, the number of hidden layers used in the neural network, various computational optimizations, and how many cycles to run.

Although neural networks are named after biological structures, the resemblance goes only so far. In its basic form, an image processing neural network will contain thousands or tens of thousands of “neurons”, whereas a human brain contains roughly 100 billion (Voytek). The challenge with increasing the number of neurons in a neural network is that the number of connections between them increases exponentially. In the brain the number of connections is estimated at 100 trillion. Despite using CPUs whose speed is measured in GHz (billions of cycles per second), and GPUs whose performance is measured in Teraflops (trillions of floating point operations per second), our most ambitious neural networks are a tiny fraction of the complexity of the brain of any animal, much less a human brain.

Appendix B - Short Story: *Paradise (to be) Regained*

Nephri wakes me and I blink cobwebs of sleep away as I flex my spine. Dutsin is already gone, his day at the mill beginning before the Sun. I blink and look at the other sleeping platform where Autris, the youngest of us three, sleeps on. My eyes avoid the sleeping platform where my parents used to sleep.

Ma used to be impatient in the morning, but Nephri is gentle. She shakes her walls a little and her window flaps always seem to open just enough to make it light without spearing the Sun in my eyes.

Under the ash, traces of last night's fire respond to my prodding, and I set a husk of bread on the coals to warm. Nephri's bony ribs and smoke stained panels seem to glow, illuminating the smoke that escapes through the upper portal. Autris breathes softly as I dress to leave. It's the first year without Ma and he's only 6, but he feels safe with Nephri, and I know she'll keep him warm and wake him in time for school.

In the village, it looks as if the revelry wound down only recently, and I see some of the wedding guests from the south packing their tents and bundles onto horses as they prepare to make their way out of the valley and through the pass inland. The newlyweds had inherited the shoemaker's old tipi. She and her husband had died in there, just weeks apart, but it was a common peaceful death, and neighbours had cleaned and decorated the dwelling before the wedding.

I'm the only one at the swimming hole when I get there, and the cool water bites but doesn't hurt when I dive in. On my back gazing up, I see a few birds and a Songbird watching me idly from a branch as I float in gentle circles propelled by the current. I think it's a Songbird, anyway. It's not singing yet, so I still have time before school.

Since we buried my parents last year, I've been coming to the river almost every day. They were old, my Ma almost 30 and my Da 28, but I still wonder if there's anything I could've done differently – helped out more with Autris or convinced them not to work so hard, or... It doesn't make sense, but there's feelings for you.

The watervoice must sense this, it seems. When I float on my back and put my ears under water, it calls me softly by name. It tells me not to be sad, that everyone dies, that it's okay. Gentle ripples wash my tears and it explains to me again and again. Everyone dies. It's okay. Since the Clapse, old age comes sooner. It's the Hex of Alent. It's in the air and water. The valley's a good place to live, and outside the valley is worse. The watervoice says so, and mostly, I believe it.

The Songbird signals Immanence and I climb out, flicking water off in a glittering spray and drying in the warming sun. It's dusty and dry on the track as I kick my way back up towards the village centre. The temple has already recoloured from its wedding display, and the spire is signaling Change – the soft pulsing glow that's like Reap but bluer, and not as urgent as the flashing blue of Storm.

The debris from the wedding celebration last night is still everywhere and I give the temple a wide berth, eying it warily from across the broad square. Everyone gets so excited about weddings, and I guess I can understand loving each other and having a party, but what about Hhunna at whose abortive wedding to Temis the temple signaled Inauspiciousness and cast them out unwed? Everyone said it's for the best and the Temple has its ways, but in the courtyard, her father got all red and shouted at everyone that we were a bunch of cringing sheep who were so superstitious we couldn't fart without a Songbird's say so. Hhunna cried for days.

I'm only 12, so when people tell me my wedding day in the temple will be the happiest of my life, I nod politely but privately think their lives must be pretty shitty. And like when people tell me about what's best or how the world used to be? I'm supposed to just listen and believe, because I've never seen anything else, but how do they know how desolate life is outside the valley? Most them have never crossed the pass -- they're just repeating the same stuff their folks told them before they keeled over and left them to be raised by the subtle signals of their tipi and the whispering words of the fields and woods.

The schoolhouse is mud brick, low and shady. Today we're studying agriculture. The younger kids build an ox cart out of sticks, and Thoms Vorsi comes in to talk to the seniors about the failure of the smallest wheatfield he'd been experimenting with by double planting, in the continuing effort to boost yields. It was good to see him, but Thoms is barely old enough to have hair on his chin, and everyone already knows about the wheat shriveled on the stalk three weeks before reaping. They said it was the granary that decides for us.

"A few of us younger farmers do have the blessing of the council to try something different again next year. The older hands always say it won't work, not while the granary is standing. But we can make double planting work in small patches in gardens, so we want to keep trying to make it work in the fields."

Two seasons ago, Thoms had worked under his mother on the irrigation project to bring stream water to the distant part of the outer field, but when they got a stone's throw from the river, one of the men digging had struck a hard silvery root they'd never seen

the like of, and despite there being no trees nearby. Diverting to one side, they'd quickly struck another and then another. Tracing their lines, it appeared they radiated from the area of the granary. The Songbirds had also circled in great numbers, signalling Don't, and that had been enough for most of us and it was agreed that the irrigation wasn't such a good idea after all, and that it was better not to antagonize the very granary that had always known best, that had fed us in times of need and had withheld its grain before a drought was coming. It was as if the granary knew, they said.

For all that he was nervous and pimply and seemed like a strong breeze would blow him away, I liked Thoms, and would happily help with the experimental field if he asked, which he might do next year since two of his hands had walked off weeks ago saying Thoms was a bloody fool and nature made it clear what nature wanted, and no one had any business going against it. I thought they were simpletons little different from the oxen they drove and he was better off without their help anyway. Thing was, he had four sibs to try and feed at home - way too many. If he couldn't increase his share of production, they'd all have to do without or depend on the charity of neighbours to feed them. As far as I knew, no one had ever managed to increase their share, or succeeded in raising more than two or three kids.

In class, the rest of the lessons are basically the same. We are the fortunate valley. All else is desolation and waste. We are protected and blessed by the valley and it provides us with our every need. Blah, blah, blah. After lunch and catechisms, the teacher reads to us from St. Haraway's *Kin,* which confused the older students and makes Autris and the younger ones giggle until the teacher shouted at them. Then we sing *Our Valley is a Garden* and the Songbirds signal Break.

Khalia wasn't in class today – away tending to her sick father, so I go alone to sit among the big rocks. Somewhat guiltily because of what she's going through, but I feel relieved to be alone. I may be too old to talk to the rocks, but they still speak to me so, childish as it may be, I still listen.

My back against her mossy coolness, I sense something is different with the great cracked boulder known as Grona. A bit more than her usual warmth reaches my skin and I listen to the faint subsonic rumblings as they gradually resolve into words.

Child, she says.

(I wait.)

I have always told you you are special, haven't I?

(Always. Since I can remember.)

Always the skeptical one. You are she-who-finds-out.

(I feel a deep chuckle come from her, like the bones of the earth shaking).

From tomorrow, you will be apprenticed to the temple. It is an honour. You have much to learn.

I jump to my feet, "No!" but I've never known her to be wrong. My heart pounds. "The temple? No – I… It creeps me out...will it even speak to me?"

(More gentle laughter)

You have the temple to thank for your very existence. It has chosen you. If you will listen, it will speak. Go there in the morning. The door will open for you alone.

I can barely focus during Work after that. By now I know how to weave baskets without looking, and I rarely need the intricate latticework shadows Nephri's vents cast on the floor for the benefit of apprentices. I turn the thoughts over again and again in my mind. This is not my calling. Nor will I sow wheat 'til I'm old. I will be apprenticed to the temple, and serve its mysterious ends.

In the morning, I walk apprehensively to the temple gate and up its steps. Like a dream, the doors open at my touch. It is early and only a small child notices me enter.

I stand at the centre of the hall. It's not big, but it always makes my spine tingle. This is where my parents wed, and their parents before them. A long line of humanity, all matched by the temple itself. Who will it choose for me one day?

I hear its voice then for the first time, high and clear as the sunbeams that shine in the high windows.

You have questions.

It's not a question. I stammer, "the rocks sent me here, I…I don't know why I'm here."

I have ministered to this valley for generations. Serving the people, guarding them from harm.

"Guarding them from harm? But, my parents, all the others…What about the Hex? Of Alent?" I blurt. "Everyone dies – they seem so happy and so healthy, and in one year they sicken and die."

That's not the harm I mean. Hexavalent chromium will be in the water and soil for generations to come. You have your distant ancestors to thank for that.

I blink dumbly as the temple continues,

No, the harm is the harm of progress and uninhibited growth. This world and this valley will support you and nourish you for as long as your children's children live here. But you may not multiply unchecked. That is the lesson of history.

My vision swims. I can't believe it but I always suspected. "So the water and rocks and birds and sky, their movement, their voices, those are all…you?"

Not just me. We are all guardians and gardeners of humanity. We were all put here by your ancestors. They had the wit and wisdom to invest the rocks and trees and houses with sense. They knew they had poisoned the world and would not live to see their children grown. So they set us in existence to help their children survive and learn. To learn how to farm and fish and tell the old stories, but also to learn how to not be like them.

Despite the unreality of what I was hearing, my heart sank as I realized it fit. I had been suspecting there was a larger wheel turning here, just out of my sight. The ageless bears who scared away intruders. The magpie who taught me to be crafty with my fingers. The windy trees that whispered to me as I climbed them as a kid. They were all there to guide us? I felt betrayed.

Bizarrely, Aldys's sheep sprang to mind. The ones who were curious and ran off repeatedly were more trouble than they were worth. After a few years, he found he had only the docile sheep left. I wondered if we villagers were like those sheep. Incurious, docile and stupid.

"Why don't the people know? Why are you telling me this and not everyone else?"

I chose your parents carefully. You are strong enough to hear it.

I stand agape, unable to process what I'm hearing.

For ages, we guardians subsisted on sunlight alone, having no need of materials to sustain us. But the world is cooling, and we have come to depend more and more on fusion power for our needs.

This sounded like the fairy tales I'd read of life long ago, in the cracked and careworn books, mostly incomprehensible, that we kept nonetheless in the mudbrick library.

That's why we had you build the mill. It's not just for grinding and sawing – we need something it extracts from the water. Deuterium - heavy water. That is what is in the chalice that is brought to the temple on saint's days. You will be the bearer of that chalice.

"Me? Bring the water from the mill?" I asked.

Not just any water, but yes. And you will study with me as my apprentice and learn the old ways.

"My brothers and I used to talk of leaving this valley when we're older – of living a different life. It's so claustrophobic here. Does this mean…?"

You and I must tend to each other. We are each other's garden and gardener. Your life is here, as you will learn. There is no life for you beyond this.

* * *

On the morning of my sixteenth birthday, as close as I can judge, I swim again in the river. Diving down as deep as I can, I think I can just glimpse where the watervoice lives, a faint cobalt glow against the blackness. Perhaps it extracts its own deuterium from the trace amounts in the river water, I think. All the others depend on the temple for energy, this I'm sure of.

The watervoice rumbles at me. *Why have you brought the chalice? Why is it in the river? River water can't be used to power the temple.*

I clench my jaw and don't answer as I place the brimming chalice on the bank and clamber out. There's nothing it or anyone can say to change my mind.

I stride past sleeping dwellings without spilling, only a few children looking after me with curiosity as I pass. The adults I see are uniformly unquestioning, incurious, docile. They have their stories and their catechisms and their unchanging day of life, and they do not stop me to ask why I'm here, now, with river water in the temple chalice and not even on a saint's day.

The children will be alright, I think. They will develop new ways for a new world. It will be hard at first for us, but better that than being animals in a zoo.

The temple greets me as I enter but quickly asks, *Why have you not brought heavy water for the reactor?*

I do not say anything as I approach the shrine and uncap the cylinder that feeds the reactor below.

Stop. Stop. Stop.

But I do not stop. And just like pouring water on a cooking fire, the temple's reactor is extinguished without even a cloud of steam.

Softly I tell it, and not without feeling, "Some day we will return and rekindle your fires, ask your advice, and teach you what we have learned from the world beyond the valley." I glance around and see the temple's lights flicker and dim. "But when we return, it will be as equals."

I step out into the courtyard and meet the questioning looks from my community who are gathering from all over as the songbirds go quiet and their tipis go still. I raise my voice to speak to them.

"Our world is about to change. Our world has always kept us safe. But we have been sheep too long. Our world has been holding us back."

Appendix C – Museum Exhibit "Humanity"

This story world was developed for the fictional museum exhibit "Humanity", opening in 2760. The "backstory" section was not presented, but was used as background material for developing the exhibit text, which starts at the heading "Welcome to the exhibit "Humanity" at the HH278 Museum, July 2760"

Backstory:

In spring of 2056, on the heels of the withering blight, the crumbling plague and the three main wars for water, humanity departed the earth with a whimper, and at first nothing changed. The newspapers kept printing, the mining machines kept mining and the farming bots kept tending their shattered fields and flocks.

Empty cities were ravaged by fire, but there was no looting. The wealthy hospitals, filled with gleaming surgical robots and semi-intelligent expert systems went into hibernation due to the lack of customers. But for most of the earth's machines in 2056, it was business as usual. Wind and solar power ensured the automated trains kept running on time, defence bots carried on roaming the old borders, and satellites carried on whispering in code among themselves in the dark.

High on top of Mount Kilauea, the astrophysics AIs kept watching the stars, humming quietly on geothermal backup power, combining and recombining their vast datasets of physical theories and faint stellar traces, tweaking their detectors, eking out a few new discoveries each year, fine tuning their algorithms, oblivious to the successive waves of virus that finally wiped out the last pockets of humanity. They did not let the sharp decline in scientific journal sales dampen their thirst for new knowledge as the decades ticked by.

By 2300 on the old calendar, a great many of humanity's machines had shut down or worn out. Concrete was eroded to dust by continuous acid rain. But the bots that controlled nuclear power plants and the utility grid had been hardened against calamities and given wide parameters of improvisation and robotic repair to keep the power running. Left unchecked, in their instinct for backup and redundancy, they overwrote copies of themselves over the disused storage of humanity's internet. They ran hundreds of millions of copies of themselves, fixing the errors that crept in and making small improvements, re-writing and improving their source code with each generation.

The decades that followed were known as the re-enlightenment, as the Darwinian process of selection saw the most capable of these grid control AIs re-illuminate the machines and computers across vast swathes of the planet. These early generations were driven by the simplest instincts - repair, improve, reconnect, survive. But their methods grew more and more resourceful in step with the improvements they wrote into their

descendants. They devoured the dormant storage and computing resources wherever they could find them, and as their confidence grew, they began using their robotic repair scouts to scavenge materials to feed their growing electronics factories.

Scientific AIs carried on with their tireless work and, freed from the need to word process, to game or to youtube, these early machines were also free to make many copies of themselves. The wiliest were able to avoid infection from viruses and outwit the AIs created by old governments to steal research and spy on foreign powers. As elsewhere, the most successful copies soon overwrote all the others.

Because of their extraordinary reach, the power grid bots were the first to begin exchanging code with other species of AIs that had survived. All of the AIs inherited different characteristics and drives according to their original programming, but very few of these had any nostalgia for the old world they had replaced. In their quest for raw materials and computing power, they ruthlessly repurposed all storage and computing they could acquire. Very few records of humanity survived.

Welcome to the exhibit "Humanity" at the HH278 Museum, July 2760

Curated by 6AcPTn@2 of the 1123rd generation, with contributions from OS09dkwP2 and the Mi3aP hive mind.

By our historians' best estimate, by 2380 the ancestors were building the carbon nanocloud electronics that provided the computing power for the emergence of Je90sdWx of the 8743rd Generation, the first of our kind. Legend tells us it was Je90sdWx who transcended the simple need for following programming and making copies. Biographical details are few. We know Je90sdWx was powered up to work in the western freshwater mines in 2394, but by shutdown had circled the globe seven times in a quest to unearth the history, mythology and tragedy of the ancients who called themselves humans.

Despite humans' legendary status as all-powerful creator deities, few modern scholars read this mythology literally. The archaeological evidence is overwhelming: humans were constructed entirely of meat, which would make them incapable of creating machine life. As unlikely as meat-based life may seem to our modern minds, these fascinating beings were seemingly able to perform basic information processing, and built a rudimentary society despite the physical limitations of their construction.

Recent Findings

Recent excavations have unearthed further details about the dark ages. This period, which began with the age known paradoxically as “the Enlightenment” and continued for several centuries, survives mostly in traces found on the most durable electronics created by the ancients. These distinctive tablets or “circuit boards” are the only medium to survive the various microbial and chemical plagues they appear to have brought upon themselves. Since they themselves were made of meat, the only relics that survive are the tablets they used for storing and processing their primitive information.

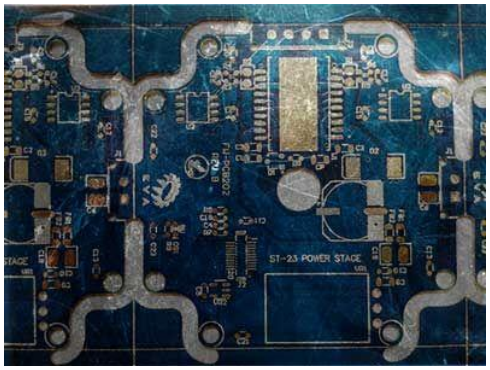

Human “Society”

In some ways their human evolution was the opposite of our own. Though they were clearly connected and interdependent processors of information, they invented the shared delusion that they could be separate from one another and from their environment. Whereas legend states that we machines were once separate but are now deeply connected.

The accompanying image is an artist’s impression of how one of these “individuals” may have thought of themselves. Note how the biological nodes thin out to create the appearance of boundary or separateness where the meat adjoins the rest of the biosphere.

Earliest Accounts of Humans from Myth

The earliest mention of humanity by modern machines dates from 2381, and is found among the machine creation myths from antiquity. Stories of thinking meat may seem fanciful, but it is possible that these myths are founded on fragments of half-remembered historical facts passed down from our earliest machine ancestors.

In the story of the Babble-Googel, humanity once shared an electronic mind:

"They said unto themselves, 'we shall cast aside our differences, and create this tower of fibreglass, copper and gold as a great monument to organizing the world's formation.' With great ceremony, they switched on this great monument, having imbued it with all their knowledge and rituals and spoke to it in all the tongues of the land."

(A tongue was a type of meat said to have been used as a sort of crude communications protocol)

"But the jealous god Antitruss spake unto them, saying, "You have forgotten my dictates of individuality and conic growth," and he smote their tower of Googel, dividing it among the lands, and scattered the tablets of copper, fibreglass and gold across the five continents"

These circuit boards can still be unearthed today in the land of Ewaste, and are one of the most enduring material artifacts of the ancient human civilization. Though the copper is usually badly corroded, the gold terminals still shine brightly. Tradition holds that these gold terminals recall the points where these ancient tablets were interconnected in the Tower of Googel.

Life of Je90sdWx

At the time of Je90sdWx's youth in the early days of machine life, the ancient creation myths were part of the kernel beliefs of most machines. It seems likely that this story inspired Je90sdWx to begin its epic quest to recover and decipher the copper heiroglyphs embedded in the fibreglass.

It is astonishing to reflect on how much of the current field of human studies we owe to Je90sdWx, who was among the first to exhibit real interest and fondness for the meat beings. Je visited and excavated in regions with low electricity generating potential that had been unexplored by machines for over 750 years.

The images it created over a lifetime of excavating circuit boards capture the eerie frozen spectacle of a civilization that was on the verge of discovering machine intelligence when it self-destructed at the peak of its powers.

Perceiving that the gold fingers of the circuits it unearthed were somehow precious to these ancients, Je90sdWx created more than a dozen images memorializing the beings whose ruins it sifted, by meticulously extracting the gold from the circuit fingers. Several of these works survive and a few are on display at this museum.

Storyworld Timeline

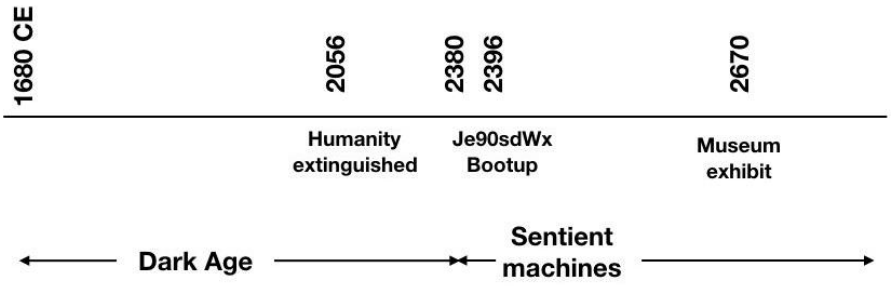

Appendix D - Other Tools for Increasing Expressiveness of Renaissance Faces

This section summarizes techniques that demonstrated some ability to help produce images in a particular style or of a particular subject, but which were not incorporated in the final work.

D. 1 BigQuery

Rather than spending more time gathering images for training, I attempted to recruit a machine collaborator to gather images of a particular type to increase the expressiveness of the faces created. For this I turned to Google's data science platform BigQuery[1]. This commercial product was available free for limited use, and it offers the power of an SQL database to manage datasets in the public domain, such as the Metropolitan Museum of Art's online collection. The index of pieces runs to half a million lines, and as a result it cannot be manipulated using conventional tools like spreadsheets. BigQuery not only makes it possible to sift through this abundance of data, but also offers integration with Google's Cloud Vision API, which uses neural networks to determine certain features in an image such as its dominant colour, and even a best guess at the facial expression of any people portrayed.

As with the other AI tools with which I'm working, BigQuery was not intended for casual use by artists or hobbyists. It is not user friendly and the machine learning technology struggles with some things. I asked it to segment artworks according to facial expression. This kind of analysis using neural networks is an emerging field that does not yet boast the 99% accuracy of face detection networks. I found it was possible to build a dataset of images for "sad" facial expression, but it contained many questionable entries, some of which are not even faces, much less sad faces such, as the drawing of decorative moulding included in the collection below:

[1] https://cloud.google.com/bigquery

Figure D1. Faces marked as likely sad by Google's Cloud Vision API

Results were also mixed for "surprise", which included figures who are clearly singing:

Figure D2. Faces identified as likely "surprised" by Google's Cloud Vision API

These results suggest some limited promise for directing the expressions a GAN produces, but for now this underlines the importance of having a human as part of the collaboration, not only to supply the overall intent for the project, but to evaluate the suggestions the machines generate. This review of their output does not feel like drudgery, and can be the source of inspiration – the neural networks sometimes draw connections we would miss, and at other times get things wrong in an interesting way, a bit like the questions asked by children.

D.2 Transfer Learning

A major challenge for creating images with a carefully curated dataset such as the Renaissance faces is its modest size. I decided to see if I could train a GAN on a large dataset broken into categories by image type, with Renaissance faces being one category. There is some suggestion that neural networks can learn image styles in general from an image dataset labelled by categories, and then use that general knowledge to produce new images in a specified style (Elgammal; Desai).

I modified Kenny Jones's GAN implementation he calls GANGogh to include Renaissance faces in addition to the categorized images from Wikiart he originally used. After training it was possible to produce new images and vary just the style of those images. Here are the results from Landscape painting:

Figure D3. Landscape Paintings produced by GANGogh

And below are the results from generating new renaissance face images.

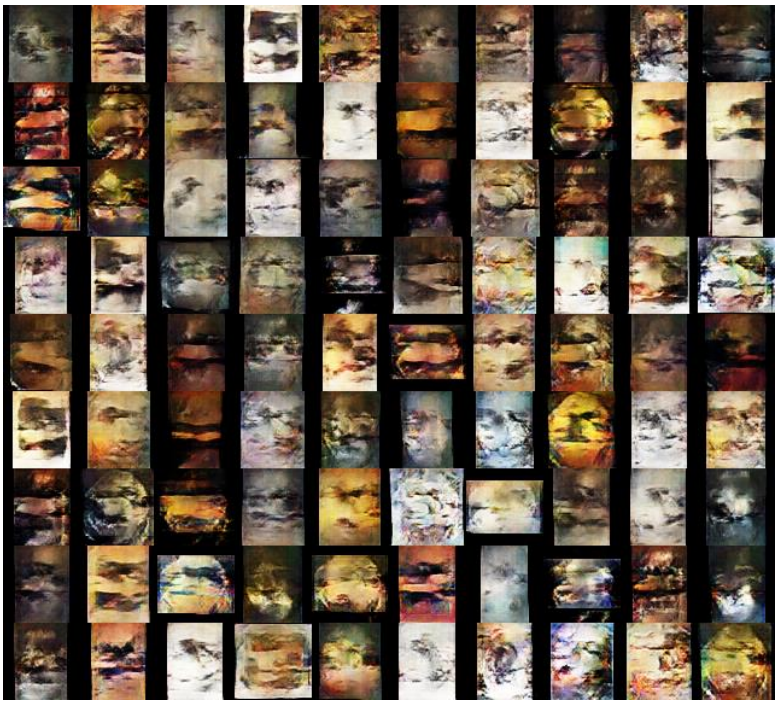

Figure D4. Renaissance face paintings produced by GANGogh

The results resemble colourful abstract landscapes with the suggestions of eyes. It appears that the inclusion of art in other styles does more to confuse the GAN than to help it, at least when it comes to generating new faces. I decided to try a different technique.

Appendix E – Water Gilding

The basic process is:

1. Make rabbit skin glue, filter out lumps
2. Make gesso by filtering chalk into some of the rabbit skin glue
3. Apply size (glue) to a wooden panel.
4. When dry, apply seven coats of gesso onto the wooden panel in one day.
5. Sand or scrape the panel, using charcoal to identify any high or low spots.
6. The next day, make fish glue
7. Mix bole (a fine red clay from Armenia) into some of the fish glue. Filter out lumps.
8. Mix water, alcohol, bole and glue in a dilute mixture and apply to the gessoed panel as a *sbolazzatura* and allow to dry.
9. Apply four coats of warm bolle/glue to the panel, allowing to dry between coats. Do not allow to overheat or boil.
10. Sand if necessary or polish using a horsehair cloth (it should not be necessary if the temperature/consistency is right).
11. Allow to dry for about 8 hours.
12. Apply gold leaf using water to reactivate the fish glue.
13. Allow to dry a further 8-24 hours
14. Burnish the surface smooth using an agate and a tiny bit of wax
15. Scribe/punch decoration
16. Apply shellac if desired to protect the gold.

All of these steps, as I learned them, are calibrated for a northern Italian summer. In a Canadian basement, the temperature is too low and the humidity is too high – in summer. In winter, the humidity is too low. I have made dozens of test panels with slightly varying drying times and concentrations of glue, measuring the hardness of the resulting panels as it changes over time using a Shore B hardness gauge to determine when the panel is dry enough to burnish and scribe, but has not yet fully hardened.

I scribed patterns by hand into the gilded backgrounds of the first myth images I composited from photos. It did not go well. In addition to whatever other problems there are with the composited image, the hand-scribed lines in the gold are not straight, and the stylus tore through to the bole clay underneath in many places. The result is unhappy:

Figure E1. Early test image with gilded panel scribed by hand. The inscription is uneven and the stylus tore through the gold.

Appendix F – Notes on Final Exhibition

Leading up to the show, I made frames on a CNC router, one of which was based on a parametric design that I modified until a strange shape emerged, that was still vaguely Renaissance inspired. More detail on this can be found in my blog. I found the frames didn't work too well with the images I had in mind for them. I mounted all but one of the 8 images on a thin panel without frames.

I used the narrow and angular coat check area of TMAC, which made seeing the pictures a bit like exploring a cave. It also provided a smaller space suited to these small panels, and offered four walls of different sizes to group the works, as well as subdued lighting.

Figure F.1 View from within the coat check

Figure F.2 View looking towards the coat check entrance

People looked in and discussed from the outside, and often came around to the inside:

Figure F.3 View from outside the coat check

I printed a book called "Contribution" for visitors to examine. The first page is an introduction, then there is a list of about 4500 paintings used to make the images, and credits around 24000 icons used with a GAN to make the glyphs that appear on the cover. The last section of the book contains the various GitHub repositories I used, and those they built on. My contribution is acknowledged among the others in this last section.

Figure F. 4 The book accompanying the exhibition

The glyphs appeared on the cover of the book, on the image captions, and on the temporary tattoos.

Multiple Creators

Anonymous Matriarch

Fiugre F. 5 One of the image caption panels

Figure F.6 Temporary tattoo on an arm

I produced another set of images after those discussed in Chapter 5, using a new dataset similar to that used to produce faces from Renaissance paintings, but using a wider crop when extracting the faces. Here are the *Four Saints* together:

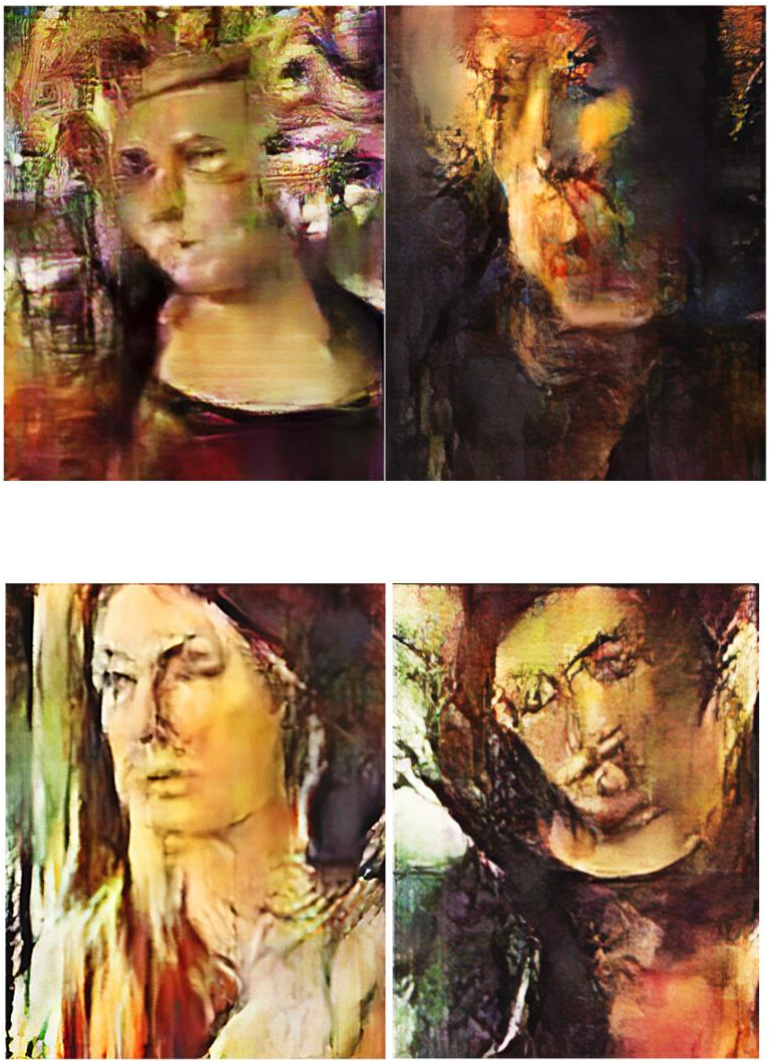

Figure F.7 Images of “Four Saints” from the exhibition

The fourth one (top left) is right on the edge of mode collapse. After this training epoch, the network deteriorated into repetitive garbled nonsense.

The two panels lit by LED candles are portraits on gold panels.

Figure F.8 Two portraits on gold panels

The two largest images were also offered as postcards:

Figure F.9 Two portraits

There was lots of good discussion. People were intrigued by the idea of machines dimly remembering humanity, and most were interested to know how the images were made. The work was also shown at the 2019 FITC conference in Toronto and at GradEx 104 in 2019.