2 Optimizing Art

This chapter begins with a personal story.

In the early 2000s I completed a master's degree in computer science at the University of Montréal. My gradudate research project at Yoshua Bengio's Laboratoire d'Informatique des Systèmes Adaptatifs (Computer Science and Adaptive Systems Laboratory)[1] focused on artificial neural network language models—forerunners of deep learning language models that currently run the most advanced speech recognition and automated translation systems. In those days, however, artificial neural networks were regarded by the majority of AI researchers as a dead end.

Over my years as a graduate student, I had become disenchanted by the disciplinary, homogeneous, and conservative culture that prevailed in the field of computer science at the time. Soon after finishing my master's, serendipitous circumstances and encounters led me to distance myself from artificial intelligence research and pushed me into the new media arts. Montréal in the early 2000s was a vibrant and exciting place for emerging new media artists with a rich, burgeoning ecosystem of artist-run centers, academic networks, and tech companies. Yet as I was making my first steps into the art world, my fascination with the arts was distorted by the engineering culture I had bathed in for almost a decade, which had left me with a naive and narrow view of art and society and a tendency to see everything as a problem to be solved.

In that context, sometime during the winter of 2005, I developed a proposal for an artistic project that reflected my state of mind at that point. I sought to create an interactive installation that would display generative images to the visitors, who would then interactively select their preferred images. The apparatus would then use machine learning to adapt to the public's tastes, generating increasingly aesthetically appealing images over time. In summary, I framed the artistic work as an opportunity for the public to optimize "beauty" through the optimizing power of machine learning, which I had been trained to use over so many years.

My young computer scientist brain was absolutely convinced of the revolutionary character of this proposal. My project would generate an optimal solution to the "problem" of artistic creation and it could potentially replace artists altogether. This idea encompassed a certain degree of contempt for my artist peers, who had for the most part gone through art or film schools and (because they did not possess the knowledge and skills that I had) were surely unable to compete with my brilliant approach. . . .

As questionable as this mindset might appear to any artist or curator, or anyone else remotely familiar with contemporary art, I remember thinking and feeling that way as a result of years of exposure to a narrow-minded disciplinary culture within AI that often perceived itself as superior to other disciplines—in particular, to the arts and humanities. Consequently, I and my computer science peers viewed artists as lazy daydreamers who had probably chosen the arts because they were not smart enough to work in the sciences.

Nothing, of course, could be further from the truth. Making art is challenging. Making *good* art requires not only intuition and talent but relentless dedication and a strong dose of resilience. Artists work extremely hard under challenging conditions, facing harsh criticism and rejection on a regular basis. Not only do they need to find the means by which to produce their work, they need to present it and market it—most often in exchange for very low income. Artists usually have to invest in their career for decades until their work starts to pay off, and successful artists more often than not live frugal lives.

Still, today perhaps more than ever, Silicon Valley culture, aided by the media, continues to fuel a certain myth of the bohemian artist whose only social function is to generate beauty while we are waiting for machines to do better. In particular, this distorted view persists in the field of *computational creativity* (CC), a branch of AI that attempts to understand and reproduce human creativity, and oftentimes tends to approach artistic creation through the same dualistic world view that has plagued computer science since the 1950s.

Fortunately, I never produced the generative interactive installation project I had proposed. Through years of studying and working as an artist and learning from my new peers in the art world, the internalized biases that distorted my comprehension and appreciation of artistic creation slowly eroded. I came to understand that contemporary art was not so much about solving problems as it was about creating problems for the viewers by bringing them into an experience, allowing the revelation of otherwise unfathomable truths about the world through its estrangement.[2]

Art, Purpose, Teleology

Years later as I was finishing my PhD at Concordia University within the Faculty of Fine Arts, I presented a work in Montréal that made use of reinforcement learning systems to generate luminous and sonic patterns. Rather than responding to the tastes of the viewers, these systems acted as indeterminate, self-sufficient, pattern-generating agents, trained in real time in response to their environment.

I invited one of my former computer science colleagues to come to and see the work. After the performance, my peer, an outstanding deep-learning researcher, approached me with an idea on how one could apply machine learning to artistic creation. They suggested creating an online platform that would generate images and propose them to the public, allowing viewers to select their favorite works—in other words, crowdsourcing the attribution of an aesthetic value to generative pictures. Using that information, a deep learning algorithm would find an optimal solution to the problem and thus would autonomously create images with greater aesthetic appeal.

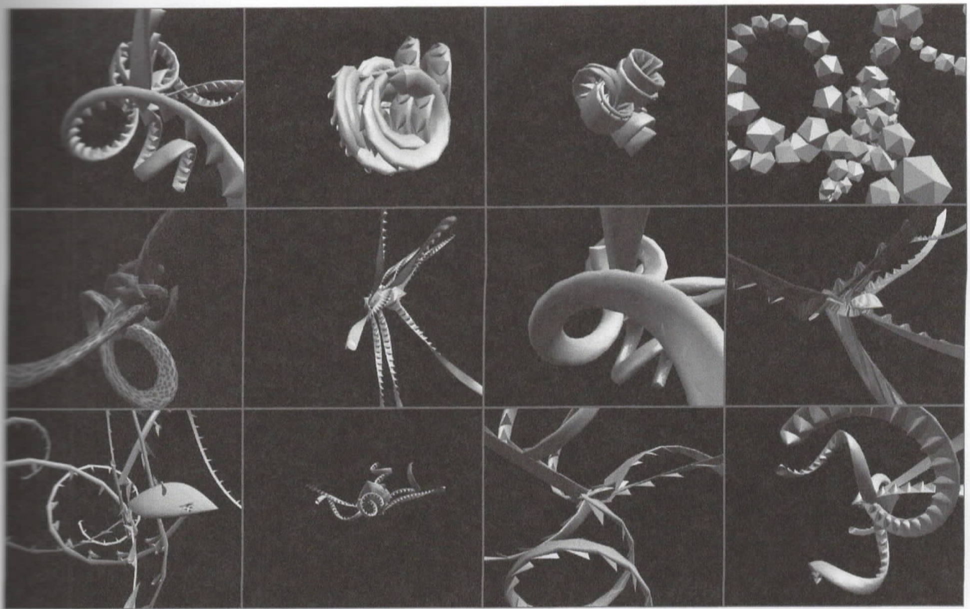

Figure 2.1
Images from *Galá-pagos* interactive exhibit, 1997. Courtesy of Karl Sims.

My peer's idea reminded me of my own proposal for an aesthetics-optimization system. Both of us had attempted to frame the artistic process as an optimization problem in which the aesthetic value of images was determined by a majority vote. Moreover, we both framed art making as a problem to be solved by defining the artistic value of a work as a quantifiable objective function.

This story demonstrates one of the core issues that render traditional optimization approaches in machine learning (and more generally in computer science) fundamentally inadequate when applied to the arts. It also shows how artists work with machine learning in alternative ways that often run counterstream to machine learning conventions. The reasoning behind such aesthetic optimization systems is not entirely unsound. In fact, similar concepts have been implemented in the art world; one of the most famous is Karl Sims's 1997 installation *Galápagos*, in which users control the evolution of virtual life forms by selecting their preferences (see figure 2.1). My old self understood that the value of an art work is subjective but nonetheless undertook a series of questionable mental leaps leading to the belief that one could *solve* art using machine learning.

First, I made the assumption that the subjective value of an artwork could be translated into an objective value using statistics on people's tastes. On that premise, I reasoned that while everyone has their own set of tastes and preferences, *good* works of art are generally appreciated by more people, and therefore one could assume the existence of hidden properties that make some works *better* than others. Finally, I presumed that these hidden properties could be approximated by a machine learning model such as a neural network that could be optimized to meet the preferences of the majority.

Although this might seem reasonable from a mathematical perspective, it is built on inaccurate premises about art. First, it presumes that art can be described as an optimization problem, which is far from certain.[3] There is no such thing as the *best* song or the *optimal* painting. The capacity of art to be optimized is at best fuzzy, and art is often described as precisely nonpurposeful and nonoptimizable.

Second, it attempts to measure artistic value averaging people's preferences. Value attribution in the art world is very complex and contextual, and in artistic circles the general public's opinion usually counts for very little in assessing the value of an art work. Artistic value fluctuates according to geography, history, art form, and other factors and involves the judgement of experts such as collectors, curators, gallerists, and artistic peers.

Third and perhaps more importantly, a series of independently generated objects such as images or sounds decouples the artistic gesture from its frame of production. Hence, it calls upon a formal, disembodied, and decontextualized definition of art. Art historian Andreas Broeckmann warns against such outdated understandings of art, explaining how throughout the twentieth century, art has been broadly defined as "a practice or a form of material production that displaces, that makes strange the sense of social artefacts and conventions" (Broeckmann, 2019, p. 3). Beauty and novelty are not as important to contemporary art than concept and context.[4]

The Best Art

In the mid-1990s, Russian-born artistic duo Vitaly Komar and Alex Melamid took a conceptual and humorous shot at the idea of optimizing an art work on the basis of people's preferences. In 1994, as part of their project *The People's Choice*, they hired a marketing firm to conduct a survey of 1,001 American adults, which asked questions about their preferences such as favorite colors, shapes, and themes. Using the survey results, they created two paintings: *America's Most Wanted* and *America's Least Wanted*, which were exhibited in New York Cirty at the Alternative Museum and later on the web.

Although the work was not without some irony, the artists insisted on the honest character of their process and of the result as a faithful portrait of people's tastes. Yet, none of the works in that series—which eventually included *Most/Least Wanted* paintings created in response to surveys conducted in more than a dozen other countries—seems to have any strong artistic value in itself. With a few exceptions, the majority of the *Most Wanted* paintings display some kind of landscape painting (see figure 2.2), and the *Least Wanted* are often abstract paintings containing repetitive geometrical shapes. Once again, it was really the conceptual context in which the project was produced that made it interesting as a body of work and that allowed it to be supported by artistic institutions.

Komar and Melamid remind us that there is no such thing as the *best* or *worst* painting, nor is there a *best* song or an *optimal* media art work. Art generally does not attempt to solve problems or provide answers, and a work of art is always situated in context. Art is often described as nonpurposeful and therefore nonoptimizable. Hence, despite the huge potential of machine learning for the art world that this book tries to discern, a fundamental challenge in applying learning methods to artistic creation lies at the heart of machine learning and more broadly artificial intelligence: the dominance of optimizing and problem solving as central approaches of the field.

In machine learning, optimization is expressed through the definition of a cost function (sometimes called an evaluation function or an objective function) that one seeks to

Figure 2.2
Komar and Melamid, *America's Most Wanted*, 1994. Oil and acrylic on canvas, 24 × 32 inches. Photo: D. James Dee. Courtesy of the artists and Ronald Feldman Gallery, New York.

minimize or inversely a fitness or reward function to maximize. For example, in a typical classification task such as differentiating between images of cats and dogs the cost function will attribute a higher value when a system makes an error (i.e., identifying a cat as a dog and vice versa). Another example is reinforcement learning, in which an agent such as a program trying to play a video game makes decisions in the world and tries to achieve high scores.

One can trace this optimization principle to Rosenblueth and Wiener's *purposeful systems*, which tried to define systems' behaviors through the concept of teleology (Rosenblueth, Wiener & Bigelow, 1943). In their seminal paper the authors distinguished random processes from processes with a purpose or goal. Among these purposeful systems, they further defined teleological systems as able to adjust their own decision process through a feedback loop.

If art is precisely nonteleological, or even nonpurposeful, then basing art on machine learning seems bound to fail. Rather than favoring the emergence of possibilities, optimization reduces options as it pushes the system to converge on a specific goal. Artist Simon Penny has expressed a similar critique of artificial intelligence in general, claiming that he precisely seeks "anti-optimized" systems in order to increase the expressiveness and personality of his robotic art works.

Inspired by their work with swarms of drawing robots, artist Leonel Moura and scholar Henrique Garcia Pereira have complemented Penny's attack on optimization in their book *Man + Robots: Symbiotic Art* (Moura, 2004): "It is obvious that *any* teleological setting, linked to *any* kind of 'objective function' . . . should be banned from the conceptual background behind *any* 'artistic' application of technology (Moura, 2004, pp. 18–19)." They add that because the artistic output has no objective goal, it cannot be evaluated by such an objective function and thus using machine learning for art making is pointless.

These critiques of optimization expose a fundamental difference between artistic and engineering practices. As I explained earlier, computer scientists tend to see everything through the lens of "problem solving," and so they tend to believe that everything can be approached as an optimization problem. The result is that they tend to try to reproduce what already exists (i.e., the expected) whereas artists seek to create the unexpected.

Optimization is therefore hardly applicable to artistic practice. A first problem is the existence of multiple maxima, even within a restricted domain such as the tastes of a single individual. Most people have several favorite films, novels, or songs. To take a casual example, let us say that my favorite films are Jane Campion's *The Piano*, Spike Lee's *Malcolm X*, and Stanley Kubrick's *2001: A Space Odyssey*, it would be difficult for me to establish which of these three films is really my top choice, because I might like these films for different reasons—many of which I might not ever be able to describe precisely.

Furthermore, the space of possibilities in which art works exist is infinite and incommensurable, which makes optimizing them even harder. My three favorite films are very different from each other and it is difficult to clearly establish what unites them. Moreover, even if they are my favorite movies, that does not mean that a film that appears similar to one of these movies will interest me. For example, *2010: The Year We Make Contact* is nearly not as good as Kubrick's movie even though it is a sequel; many remakes are worse—or sometimes better—than the original. Campion, Lee, and Kubrick, have directed other films that I do not like as much. The qualities that make these films great are difficult to generalize to other cases.

One last major challenge with approaching artistic creation as an optimization problem is that art does not always try to respond directly to the inclinations of an audience. Artists design aesthetic experiences that often go against the public's preferences or question such preferences. Engineers and scientists who try to tackle art as an optimization problem confound art with entertainment, which seeks to please the masses.[5]

Computational Creativity

The kinds of engineering and optimizing approaches to creation described in this chapter belong to the domain of computational creativity, an approach to AI that has grown in importance and interest over the past thirty years. In theory, computational creativity is not specifically interested in artistic creation but in the general notion of creativity that applies to many human activities such as science, engineering, and mathematics; however, in practice it seems currently to be importantly focused on artistic creativity. A broad, interdisciplinary field that brings together artists, designers, computer scientists, psychologists, and philosophers, computational creativity integrates many different approaches to and conceptions of creativity in computational systems. Yet, at its core it lies in the continuity of traditional AI in its interest in studying and engineering human-level creative capabilities of computers by, for example generating musical scores or poems indistinguishable from those produced by human experts.

Philosopher Margaret A. Boden is a central figure of the field. Boden associates creativity with the ability to come up with ideas and artifacts that are both *original* and *valuable*. She further defines two forms of creativity: (1) *psychological creativity* or *P-creativity*, which refers to mundane, everyday activities that are novel from the perspective of the creating agent (e.g., the creativity of a child in an art class); and (2) *historical creativity* or *H-creativity*, which is recognized by society as creative (e.g., the creation of a masterpiece) (Boden, 1996, pp. 76).

Boden further classifies creativity under three different types: (1) *exploratory*, (2) *combinatorial*, and (3) *transformational*. Exploratory creativity involves the exploration of a given space in order to generate new, unforeseen elements of that space. A striking example of exploratory creativity is AARON, a computer program by artist Harold Cohen that automatically creates drawings on the basis of a complex set of rules. Another example of exploratory creativity would be a neural network trained to generate songs that sound similar to those of a well-known performer such as Michael Jackson by training it on a database of that performer's songs.

Combinatorial creativity involves creating something new by combining two objects from different spaces. Jazz fusion, a musical genre prevalent in the 1960s that combined jazz and rock, is a good example of combinatorial creativity. Another example would be the NSynth, a method that uses deep learning to create new musical instruments by combining existing ones (Engel et al., 2017).

Finally, transformational creativity involves disrupting some accepted conceptual space or cultural conventions. Marcel Duchamp's concept of the readymade, which disrupted modernist artistic conventions, is a good example of transformational creativity.

Boden traces the origins of computational creativity to Ada Lovelace who said of the analytical engine (a mechanical general-purpose computer designed by inventor Charles Babbage in the nineteenth century) that it could be able to compose complex musical scores (Lovelace, 1842). A century later, at the 1956 Dartmouth Conference that jump-started the field of AI, creativity was named as one of the core aspects of AI (McCarthy, Minsky, Rochester, & Shannon, 2006). Yet, for many decades creativity was ignored by the largest part of AI communities, who focused more on problem solving. Hence, the most successful early attempts at computer-based creativity were designed by people outside of AI. For example, throughout his career, visionary composer Iannis Xenakis used computers to generate musical scores based on his theory of stochastic music (Xenakis, 1992), and architect John Frazer received the Architectural Association prize in 1969 for his work on computer-generated environments.

Many of the works of art discussed in this book can, indeed, fall into the category of computational creativity. Yet, one of the defining features of computational creativity is its attachment to AI as an effort to understand human creativity by engineering creative processes on the computer.

From the 1990s onwards, as raw computation became more accessible and allowed for more complex AI systems to be designed, interest in creativity surged again within mainstream AI. Machine learning has been used in many computational creativity applications—with promising results in the mid-1990s, and with tremendous successes since the mid-2000s.

Music generation is probably the most advanced domain of computational creativity. An example of a recent success is the DeepBach system for the generation of Bach-like chorale scores, which trained neural networks on a corpus of Johann Sebastian Bach's polyphonic chorales. In an online test with 1,600 listeners, of whom about 25 percent had significant musical expertise, more than 50 percent confused scores generated by DeepBach with authentic pieces from the Baroque composer (Hadjeres, Pachet & Nielsen, 2016).

Although computational creativity encompasses both artistic and scientific approaches, the field is still led by scientists, not artists, and its core objectives and methods are thus primarily scientific rather than artistic. For example, the ability of an algorithm to imitate the style of a master painter such as Van Gogh or Degas, such that even experts are fooled, is surely impressive and definitely interesting from an engineering perspective. Yet, it is not clear how such machine made *pastiches* contribute to contemporary artistic discourse and debates, any no more than would man-made imitations through the use of computers.

Furthermore, *the framing of creativity in computational terms and the equivalence framing of creativity as a sufficient condition of art are misleading.* Art, like science, employs creative processes, but it is much more than that: it is first and foremost a domain of activity with its own communities, codes, and references. Therefore, while the field of computational creativity is of great interest for philosophy and science, and its recent advances certainly have merit, its relevance for the arts is more nuanced than it appears at first look.

Questions of novelty and value to define creativity also present a number of problems. Novelty and value are capitalism's most potent drivers of consumption, and artists who want to generate engaging and reflective worlds to suggest revolutionary modes of being and thinking need to be critical of such false idols. For example, a new pop remix of an old song might be valued by the masses, bringing in millions of dollars in revenue for the artist and the recording label, but it might not be as artistically relevant as an experimental work of art that reaches only a handful of people whose lives it significantly changes.

The Imitation Game

To be fair, most of the studies that deal with artistic creation within computational creativity do not attack the problem by trying to directly match people's tastes using a generative system. Rather, to avoid the direct intervention of humans, many studies use a proxy. Instead of using the audience's feedback as a measure of optimization, these studies look at accepted canons, such as Baroque music, English poetry, or modernist painting and then design algorithms able to generate new works that fit within these categories. Their measure of success is generally a variation of the Turing test, in which they ask a group of human subjects to distinguish between original human works and those generated by the computer.[6]

Take for example a 2017 study from the Art and Artificial Intelligence Lab at Rutgers University, which is among the most interesting in that type of research. The researchers designed a new deep learning system called *creative adversarial networks* (CAN) (Elgammal, Liu, Elhoseiny, & Mazzone, 2017) to create artistic images. Using a data set of 81,449 pictures of existing paintings representing artistic movements ranging from Fauvism and Pointillism to Abstract Expressionism, they not only trained their system to imitate existing styles but also adjusted the cost function in order to force the algorithm to generate work that is "novel, but not too novel" by maximizing its deviation from established styles while minimizing its distance from the norm (see figure 2.3).

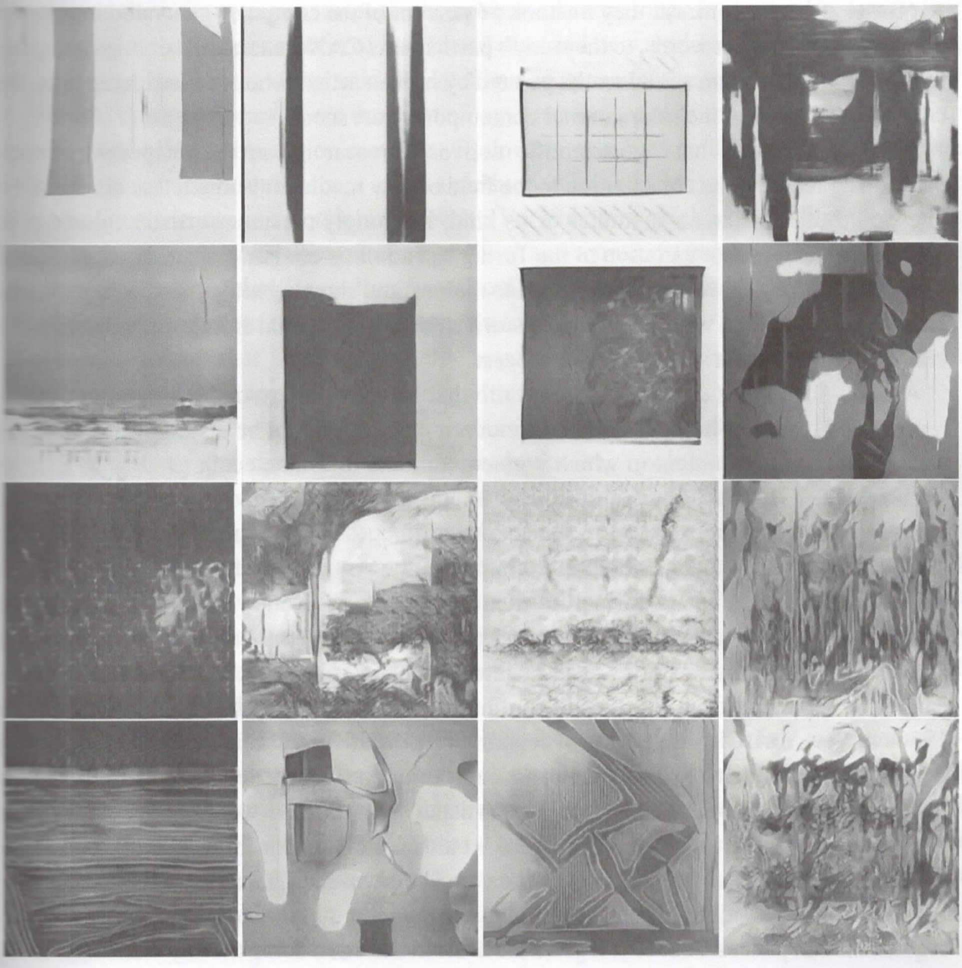

Figure 2.3
Sample "paintings" generated by a Creative Adversarial Network (CAN) in Elgammal et al. (2017). Courtesy of Ahmed Elgammal – AICAN.io – Art & AI Lab, Rutgers.

In order to validate their technique, they compared the results of the generative program to works made by human artists from two data sets: one of Abstract Expressionist paintings and one of paintings presented at the latest Art Basel fair. The authors argued that works shown at Art Basel, the "flagship art fair for contemporary art world wide," are indicative of works existing "at the frontiers of human creativity in paintings" according to contemporary art experts. This was a good call by the authors, who explained that they did not just want to imitate existing styles, but to see whether their method can generate art that could be considered truly novel.

The results of this study are quite telling. A random sample of non-experts attributed not only a higher degree of aesthetic quality but also of intentionality and "humanness", to the computer-generated images than to those from Art Basel. While study participants rightfully believed that works from the Abstract Expressionist data set were human-made 85 percent of the time, this belief dropped to 41 percent when they considered images from

Art Basel 2016; by contrast, they mistook 53 percent of the computer-generated images for human works. In other words, to the study's participants CAN-generated paintings generally felt more "human" than actual works painted by human artists who are considered to be the *crème de la crème* of the international contemporary art scene.[7]

Although this study has high scientific merit as AI research from the perspective of computational creativity, its contribution to the field of new media art is much less obvious than one would think. Like most studies of its kind, it wrongly presumes artistic value can be determined using some variation of the Turing test administered in a vacuum. As explained earlier, art history shows that with regard to making and showing art, context is key: artistic creation is embedded within cultural, historical, and institutional frameworks that directly interact with the artist's own creative process.

Art is a fluctuating domain of human life that is utterly intertwined with culture, technology, and history. The signification of modern painting cannot be distinguished from the sociotechnical environment in which it emerged in the mid-nineteenth century in Europe, from which it responded antagonistically to, on one hand, aesthetic norms of the time and, on the other hand, to the mechanization of images through photography. The value of Abstract Expressionism, which is examined in the Art and Artificial Intelligence Lab's paper, lies as much in its formal aesthetics as in its relevance within art history. Its appearance in the post–World War II era in New York City was influenced by Surrealist automatism, aboriginal painting, and quantum mechanics (Paalen, 1943). The movement cannot be separated from the artistic community at its core, who had matured during the 1930s in New York City in a period of economic and political turmoil and had come to value art anchored in human experience, in particular through freeing the mind and releasing the unconscious. In other words, we attribute artistic value to these works because they were relevant in a certain context, yet if someone would create the same kinds of work today (like the CAN does) it would not necessarily be considered original, let alone artistic.

Hence, while this kind of computational creativity study is interesting in showing how algorithms can generate novelty, it gives the illusion that in so doing the algorithms are being creative—or worse, that they are creating art. These shortcuts just add to the general confusion. They are tainted by computationalism, a form of dualism that is still the dominant world view within AI: the idea that human behavior is a form of computation that is completely independent of the human body, and by extension, of culture, society, and history.

I reject this view, as I believe it is profoundly flawed and barren to understand and advance artistic practice. This engineering approach to art making is built on the false assumption that machine learning systems can circumvent the need for authorship by deferring the decision-making process to a neutral and creative black-box. In reality, decisions must be made at every step of the process in selecting the optimization algorithm, evaluation function, model, and data set. For example, even the most advanced research in the field of computational creativity is defined over a rather restricted domain for which enough data is available, such as the work of classical composers (Hadjeres, Pachet & Nielsen, 2016) or abstract paintings (Elgammal et al., 2017).

Reproducing human-level performance for object detection or speech recognition is a rather easy task for machine learning systems, as it requires them to find regularities in the data—in other words, to look for the expected. However, artists like Nicolas Baginsky, who work with self-organizing systems, are more than often looking for the unexpected, the glitches, and the outliers. Furthermore, many artists are not as much interested in the end product as they are engaged, through the process of artmaking, in a creative inquiry that ultimately leads to the revelation of new knowledge and new worlds. From this perspective, attempts to *automate* art using artificial intelligence appear at the very least questionable.

Beyond these problems, there is a host of ways in which artists use and misuse machine learning and other optimization methods. Optimization can take different forms, and even in machine learning research there exists a plethora of less restrictive optimization objective functions that encourage diversity and novelty. For example, much of the recent progress in deep learning is related to unsupervised learning and representation learning methods that aim at extracting regularities and patterns from data without necessarily trying to match data points to predetermined categories such as cats and dogs.

We cannot deny that increasingly large aspects of artistic production are automated thanks to advances in computer science, and machine learning is already contributing to this. Although this new form of mechanization of creative processes, which allows mass production of novelty, should not to be equated with art, it will likely have a huge impact on the art world which will force artists and art historians to redefine its boundaries.

Some will likely take refuge in anthropocentrism, claiming that art simply cannot be accomplished by machines because art is by definition a human activity. I believe this would be a mistake. While we need to be careful not to fall for Google-driven techno-optimism and embrace the idea of soon-to-come creative machines that will replace humans, we also have to consider how at least some dimensions of art might exist beyond the boundaries of the human species and its activities.[8]

The automatization of creative labor made possible by machine learning and other artificial intelligence technologies is an important challenge faced by artists. There is a perceived threat, real or not, that artists, whose living conditions and social status are often fragile, will cease to have a valuable role in society because algorithms will steal their jobs. This threat should be taken seriously. Artists currently manage to survive by undertaking a plethora of creative activities, which include tasks that might be taken over at a lower cost by machine learning systems—such as doing work for the entertainment industries to pay the bills. Yet, machine learning also offers new opportunities for artists to further redefine relationships with machines by, for instance, imagining new ways to collaborate with them.

Artists have always found ways to reclaim the technologies of their time and shape them to their needs. For instance, there is no requirement to define the evaluation function of a learning algorithm such that it attempts to approximate the public's tastes or to imitate a particular style. This function can in fact be freely defined to achieve other purposes. Thus, one could adapt it to their own preferences; they could experiment with different evaluation measures to generate new content, and they could even change the evaluation function over time. In other words, while optimization appears antinomic for artistic practice, there are many ways by which optimization processes can be hijacked for creative practice such as tinkering with the data fed into the system, playing with the model, and tweaking the evaluation function.

Learning in Real Time

Beyond trying to directly "solve the problem" of artistic creation through a popular survey, some artists choose to employ the fascinating generative process of the learning loop as a support for new forms of art. In this case, a system is optimized according to an objective

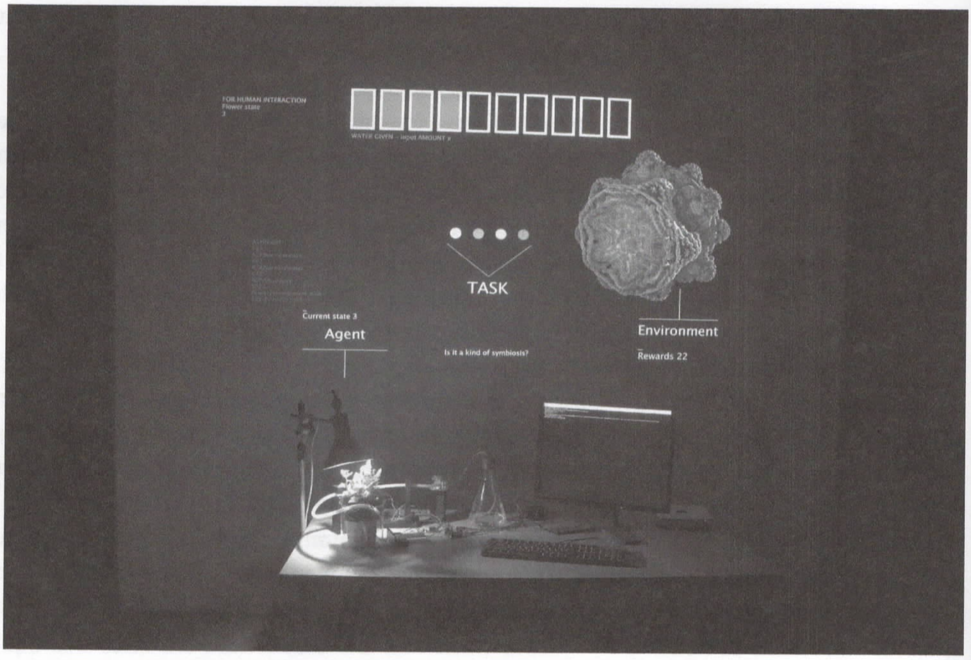

Figure 2.4
Natalia Balska, *B-612*, 2014. Courtesy of Natalia Balska.

target; however, the piece is constructed so that the *training process itself* is presented as part of the art work.

A recent example is the work *Pachinko Machine* (2017) by UK artist Brigitta Zics. The piece takes its name and inspiration from Japanese pachinko, a gambling device similar to pinball. In pachinko, the player launches small metal balls into the machine using a spring-loaded handle. The balls then travel through the machine, bouncing on different brass pins and eventually fall into various catching slots. The goal of the game is to collect as many balls as possible.

No human players are involved in Zics's work. Instead, her video installation features an animated drawing of a pachinko that plays against itself. Throughout one day of the exhibition, the machine learning agent playing the game tries to optimize its actions to continually improve its playing skill. However, the pachinko agent is also playing against another algorithmic process whose goal is to add obstacles and confusion to the game. The work is meant as a metaphor for human life: our efforts to achieve our goals are disturbed by chance and chaos that lead us astray from our path, inspiring us to explore alternatives.

The new media installation *B–612* (2014) by Polish artist Natalia Balska also centers on a reinforcement learning process happening in real time (see figure 2.4). Here, however, the pace is much slower; whereas Zics's piece evolves over the course of one day, Balska's learning loop runs over several months. The piece, originally inspired by the artist's investigation of the concept of altruism, features a closed system in which a plant is nurtured by a machine learning system.

Every day, a pool of ten individual units of water (10 milliliters each) must be shared between the plant and the machine learning system. The learning agent makes decisions about distributing the water supply, thus impacting the environment and in particular the health of the plant. The impact of the decision is evaluated by an external computer program, which, as a measure of the system's performance sends a reinforcement signal to the agent in the form of a reward or penalty. Over the course of the exhibition, both the machine learning agent and the plant adapt to the situation staged by Balska.

In the first two weeks of the exhibition, the actions of the learning agent are mostly random. Only after a few weeks does a pattern emerge. First, the agent appears very greedy for a few days, keeping all the water to itself, causing the plant to begin dying. In response to the signs of the plant's decay, the artificial agent begins to share some of its resources for a few days. When the plan shows signs of recovery, the learning agent becomes greedy again. Eventually, after many months, the system becomes more efficient in its water management, resulting in more consistent and balanced decisions.

For Balska, the purpose of the work is not so much to create an optimal plant care system but rather to generate interactions between reality and virtuality by creating a relationship between two adaptive systems, one computational and one biological. The audience is invited to discover this relationship and is left to interpret it in their individual ways. Although the work takes the form of an experimental apparatus, through this optimization process what is revealed is the unpredictability of the system, which can change depending on external factors such as the ambient humidity and temperature.

Conclusion

Machine learning offers a unique challenge to art because of its historical entanglements with an engineering culture that idealizes optimization and problem-solving over open-endedness and diversity. Traditional engineering approaches to art making within computer science and artificial intelligence rest on false premises as they focus on techniques and outcomes rather than on processes and contexts.

In other words, as machine learning is geared toward optimization, when experts of the field attempt to apply machine learning algorithms directly to artistic creation, they more than often miss the point. When all you have is a hammer, everything looks like a nail.

This is a recurrent issue in computational creativity research. Directly applying machine learning to artistic creation requires framing art making as an optimization problem. However, the fact that these generative outputs are removed from any frame of reference is antithetical to how contemporary art operates. What the twentieth century has taught us, through the ongoing operation of the avant-garde is that art is not just about creating new, beautiful stuff. It is a dimension of culture that always responds to broader cultural, social, and political contexts. Furthermore, as Komar and Melamid's *America's Most Wanted* painting suggests, artistic value does not boil down to a vote of popularity.

Artists have always found ways to engage critically with technologies by approaching them sideways. Machine learning artists use adaptive technologies as raw material, hijacking the optimization game to create meaningful experiences. In so doing, they do not attempt to answer questions or solve problems but rather ask questions and create problems for the audience to address. Machine learning art hence joins previously established artistic movements that have dealt with computational systems, such as cybernetics art, computer art, artificial life art, and generative art.

How do machine learning artists do it? The next chapter delves more deeply into this process by examining how artists intervene in the training process of machine learning systems.