CYBERNETICS AND THE CONSTRUCTED ENVIRONMENT

Design Between Nature and Technology

Zihao Zhang

First published 2025

ISBN: 9781032341743 (hbk)
ISBN: 9781032341750 (pbk)
ISBN: 9781003320852 (ebk)

Chapter 8

CO-PRODUCTIVE INTELLIGENCE

DOI: 10.4324/9781003320852-11

The Bernard and Anne Spitzer School of Architecture

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

8

CO-PRODUCTIVE INTELLIGENCE

Although posthumanist scholars have given sufficient account of nonhuman objects and assemblages, occasionally including infrastructures and computer software, their investigations leave the concept of intelligence unexplored. Many posthumanist ideas have focused solely on agency, not on the intelligence intrinsic to originating a sense of agency; we assume an object that acts is intelligent. Recent AI and machine-learning developments have provided empirical material to consider intelligence with a posthumanist reframing.

Over the past decade, AI has become pervasive in contemporary culture. Media coverage on AlphaGO and its predecessors, and other AI systems, led to further rounds of excitement and fear for the possibilities of artificially intelligent machines. Since mid-2022, generative AI systems like ChatGPT and image generators like Midjourney and Stable Diffusion have flooded social media; AI is no longer a farfetched concept from science fiction but applications in your browsers. This development propels one to speculate how to integrate AI in environmental systems, and poses a question: Will the exponential growth of AI applications shift the role of machines—other than as tools of optimization and control—in how the environment is conceived and constructed?

To answer, we must understand not only the technical capacity but also the mode of thinking that undergirds mainstream AI research. I will analyze cases in AI research with a posthumanist lens to evoke their inherent presumptions in understanding machines as tools of optimization that extend imagined human agency in managing the environment. Even though the boundary between nature and technology dissolves in the field of heterogeneous assemblages (see Part I), agency's individuation process leads to tenacious objects such as machines, plants, animals, and humans. To consider

DOI: 10.4324/9781003320852-11

intelligent machines in the environment risks falling into the normalized "adaptive management" that uses advanced technologies to automate and optimize natural processes; machines are regarded as a layer of control strategies through which humans manage natural processes.

Instead, I argue that intelligence cuts across different assemblages and reveals shared abilities or tendencies of the assemblages to exploit each other's surplus and produce effects. Most importantly, unlike agency that emphasizes individuation, intelligence involves interactions between individuals. However, this argument faces three challenges: *anthropocentrism*, *intelligence as a measurement*, and *individualism*. By deconstructing these three obstacles, I will restructure a posthumanist concept of co-productive intelligence.

Anthropocentrism and Universal Intelligence

The Merriam-Webster online dictionary defines intelligence as

> the ability to learn or understand or to deal with new or trying situations […] the ability to apply knowledge to manipulate one's environment or to think abstractly as measured by objective criteria (such as tests) […] the act of understanding.[1]

Though unspecified, these definitions point to human cognitive capacity, including learning and applying knowledge. In other words, to discuss and define intelligence, we maintain a human image in our minds and then venture to consider cognition in nonhuman entities, such as animals and machines. Anthropocentrism and human exceptionalism are intrinsic in the conceptualization of intelligence. However, a posthumanist framework is interested in how different assemblages, human and nonhuman, living and nonliving, coproduce each other, and this presents anthropocentrism and human exceptionalism with the first obstacles to overcome. In recent years, the research on machine intelligence has shed light on this issue and may help to map a posthumanist understanding of intelligence.

Machine intelligence has long been an important frontier in intelligence research. Can machines think? In 1950, Alan Turing proposed this question and the famous Turing test to argue that machines *can* think. The Turing test involves three players: a human discriminator in one room interviewing a human player and a machine in a separate room. The discriminator will try to discern between the machine and the human by asking questions without seeing them. The machine will try to imitate the human player to deceive the discriminator into believing that the machine is a human. Another concept Turing proposed was "child-machine," a theoretical strategy for machines to learn through random mutations and natural selection.[2] This strategy underpins many of today's most advanced machine-learning techniques.

The Turing test and the "child-machine" are evidence that, from its earliest stages, artificial or machine intelligence has been envisioned based on human intelligence. The Turing test essentially proposes to evaluate machine intelligence based on its resemblance to human behavior. We have long compared machines to humans; the first thing people do to test a powerful machine is to match it with the best of its human competitors. Over the years, many cases have been reported and discussed: Deep Blue (1997) defeated Garry Kasparov, the best human chess player; AlphaGo and its successors like AlphaZero and MuZero (2015–2019) trounced many of the best GO players and mastered Atari video games; AlphaStar (2019) mastered StarCraft, a real-time strategy video game notorious for decision-making based on incomplete knowledge; and OpenAI Five (2016–2019), a team of five separate AI agents, worked together and outperformed a human team in the game of DOTA2, known for its real-time collaborative strategies and corporations, as well as the ability to understand another player's intentions and act accordingly.

Underpinning our obsession with pitting humans against machines is, ironically, a sense of human exceptionalism—the unstated belief that humans are the most intelligent entities on Earth. If a machine outperforms, or at least matches, humans in one aspect, that is seen as a breakthrough. It is an ambivalent state of mind. The notion that another entity could replace humans at the top of the intelligence pyramid evokes terror, yet we secretly enjoy this sense of unease through the reassurance that we are the ultimate creators of intelligent machines. Our obsession with human-machine intelligence reveals another level of technological sublimity—the mixed sense of terror and joy. This state of mind was fruitful for popular culture products over past decades, making technological dystopia popular in video games, films, and television. Many televised franchises imagined a version of rogue AI attempting to end humanity. However, accompanying our terror is a sense of sublimity because these are, after all, fantasies that pose no real threat. Most importantly, these shows inevitably end with humanity regaining control over the rogue AI, reassuring the audience with a sense of human exceptionalism.

Other forms of intelligence are also measured against human intelligence and cognitive abilities. For example, when evaluating animal cognition, the common criteria used include teaching, short-term memory, causal reasoning, planning, deception, transitive inference, theory of mind, and language. Based on these criteria, a significant gap has been conceptualized between human and animal cognition.[3] Nevertheless, these evaluative criteria are modeled in favor of the human. Anecdotes often compare dogs to two-year-old human toddlers and dolphins to three-year-olds. Yet, this type of comparison essentially treats animals as diminished versions of us, reflecting deep human exceptionalism in framing the concept of intelligence. If there were a test of intelligence based on olfactory senses, many animals, such as dogs, would be far more intelligent than humans. Moreover, plant intelligence has emerged as a topic. In one study, researchers demonstrated that plants possess learning and memory, too, through systems unique to them. For example, *plasmodesmata*, a type of intercellular organelle found only in plant and algal cells, are crucial for plants to transmit information.[4] Yet research still attempts to locate human cognitive faculties in plants rather than discussing plant intelligence *for* plants. These examples speak to an intrinsic bias in our examination of nonhuman intelligence. By comparing other entities with the human image, we deny the possibility of noting and embracing how nonhuman entities interact with their environments in entirely different ways.

Furthermore, speaking of human intelligence unwittingly evokes an image of a perfect human being. This gives rise to a series of problematic interpretations similar to those that fueled scientific racism and discrimination narratives through the twentieth century. The concept of intelligence has long been intertwined with the intelligence quotient (IQ) test. French psychologist Alfred Binet invented the first practical IQ test, which was later translated into English and revised in 1916 as the Stanford–Binet Intelligence Scales. Intelligence tests were then quickly adopted as tools to explain racial, class, and sex differences—however narrowly "intelligence" was defined by the tests—justifying all sorts of disturbing narratives of colonialism, slavery, social Darwinism, sexism, and racist eugenics.[5]

The problem of intelligence is not the concept itself, but using a white European male image of humanity as a reference for measurement. It ignores the fact that our cognitive functions result from the long-term co-evolution of human and other nonhuman assemblages, including other species, languages, tools, and other distributive systems.

Taking human-machine intelligence as an example, we will see that what we thought was human intelligence may instead be reframed as an outcome of co-production between humans, transistors, and circuit boards. The central processing unit (CPU) in any computer is a conglomeration of transistors wired in specific ways to compute logical and arithmetic statements through the binary language of ones and zeros. However, several early electronic computers were not binary computers; they were ternary, with three states, or even quinary, with five. Problematically, the more intermediate states, the harder it became to keep them separate, because disturbances such as power surges, low voltage, or electromagnetic interference would cause mixed signals. The binary was simple to track since it gave distinct signals of "on" and "off." In addition, before the twentieth century, rules and operations to manipulate ones and zeros were already reasoned because an entire branch of mathematics that dealt exclusively with true and false values, Boolean algebra, already existed. Thus, many saw Boolean algebra as the foundation of modern computer science theory. Self-taught English mathematician George Boole developed Boolean algebra in *The Mathematical Analysis of Logic* (1847), in which he presented the truth as systematically and formally represented through logic equations. Unlike the algebra we are familiar with, in Boolean algebra, instead of numbers, the values are true and false; instead of addition, subtraction, multiplication, and division, the Boolean operations are AND, NOT, and OR. For example, if A is true, and B is true, then "A *AND* B" is true.

Conveniently, transistors can be wired in ways to build different types of logic gates and perform these Boolean operations. The assemblage of these logic gates is called the arithmetic logic unit (ALU), which is central to any modern computer. Building computing machines has thus long been viewed as a way to formalize and represent human logical reasoning through material assemblages. Although Boolean algebra laid the foundation for modern computers, the material reality of transistors and electronic circuits eliminates other possibilities regarding how human logic may be systematically and materially formalized. In this regard, quantum computing sheds new light on other aspects of human logic. Quantum computing has gained increased attention over the past 30 years because a qubit (quantum bit) possesses three states; it can be in quantum states, which means a state between 1 and 0 with certain probabilities; but, when measured, a qubit is in the superposition of either 1 or 0 states. Thus, if classical computing uses Boolean logic to erase uncertainty by choosing between 1 and 0, then quantum computing harnesses the power of uncertainty and embraces the ability to be either 1 or 0. For this reason, a quantum computer presents different material assemblages, with the potential to rethink alternative aspects of human logic.

The moral here is that what we perceive as human intelligence is not a sole effort by *human*, but instead is enabled and co-produced by the material world around us. Using human intelligence as a measure for other types of intelligence is thus great hubris, which denigrates the efforts of nonhuman assemblages in the process of co-producing human intelligence. However, if *humans* should no longer be the measure, is there another way to frame intelligence? Recent AI research may provide insights into a non-anthropocentric definition of intelligence.

Generally, there are two types of approaches to building an AI system. The first may be described as an expert system, for which human experts write a rule book for decision-making. For example, if we want to build a plant identification system, we require a group of botanists to list all plant identification rules. This resembles searching for a plant in an encyclopedia. We address multiple questions: is the plant woody or herbaceous? What is the shape of the leaves? What is the shape of the bud? When does it flower? One may create many other questions related to identifying a plant. Many traditional online plant identifier websites are expert systems.

The second approach may be understood as machine learning, for which we need only provide the machine with a large amount of data; the machine itself will reconstruct potential rules for decision-making. Building the same plant identification system with machine learning, we first collect a large number of images of plants labeled with their names. Then, we build an artificial neural network (ANN) and expose this ANN to the plant dataset until it makes predictions with a high success rate. This process is called training. To evaluate how well the model performs, scientists commonly track the prediction error rate using a loss function, which can be optimized through model training. The loss function reflects the AI system's ultimate goal. In a way, all machine-learning problems may be understood as optimization problems; an underlying loss function will always be optimized in order to train a usable model.

Though the concepts of machine learning and AI have existed since the mid-twentieth century, machine-learning techniques did not find currency until the early days of the twenty-first century, due to limited computational power, as well as the amount of data available. With the rise of "big data," as well as increased computing powers, machine learning has produced fruitful results. In 2012, the AI community experienced a major breakthrough in the ImageNet Large Scale Visual Recognition Challenge, a benchmark for computer vision. Computer scientists were challenged to design machine-learning algorithms to perform image recognition; the error rate of the recognition results became a measurement. A team led by Alex Krizhevsky from the University of Toronto designed a model called SuperVision, later known as AlexNet; based on a convolutional neural networks (CNNs) approach, it won the 2012 competition.[6] AlexNet presented a significant drop in the error rate, from 25.8% in 2011 and 28.2% in 2010 to 16.4% in 2012.[7] A mere four years later, in 2016, computer scientists lowered the error rate to only 3%, based on CNNs, while the average error rate for humans doing such tasks was 5%.[8]

Advances in AI research also provide transformative cases for scholars to ask more profound questions about intelligence. Sean Legg and Marcus Hutter, both AI scientists, have proposed the term *universal intelligence* and formalized it through mathematics as a framework with which to compare different forms of intelligence. After reviewing a collection of definitions of intelligence from research groups and organizations, psychologists, and AI researchers,[9] Legg and Hutter ascertained a common thread in these definitions, which involved the interaction of an agent, human or nonhuman, with its environment. Based on this observation and their goal of measuring machine intelligence, they proposed a working definition for universal intelligence: *intelligence measures an agent's ability to achieve goals in a wide range of environments*.[10]

Legg co-founded DeepMind Technologies, later acquired by Google, and was behind the development of AlphaGo and its predecessors; Hutter is a professor and a senior scientist at DeepMind. Their definition of intelligence is deeply rooted in their extensive experience in machine learning, particularly in reinforcement learning (RL), which is the foundation of AI systems like AlphaGo. Essentially, their concept of universal intelligence mirrors an RL framework. In RL, an agent observes the environment's state and takes actions. The effectiveness of these actions is measured by a reward mechanism, typically expressed through a loss function. For instance, in training an RL agent to play Tetris, the agent initially performs random actions such as arbitrarily moving and placing blocks. Certain actions, clearing a line of blocks, for example, result in rewards (i.e., scoring). After thousands of rounds, the agent begins to favor actions that yield higher scores, eventually mastering advanced techniques such as the "Tetris Clear"—clearing four lines simultaneously.[11] Again, any machine-learning technique can be seen as an optimization problem. In RL, the agent strives to optimize its reward function to maximize rewards or minimize punishments. Thus, universal intelligence is a measure of an agent's ability to optimize these reward functions.

Based on this simple agent-environment framework, DeepMind has trained many RL-based AI systems, such as AlphaGo. Only one year after AlphaGo defeated its human competitors, two newer AI systems based on self-play—AlphaGo Master and AlphaGo Zero—overwhelmed the original algorithm. In 2020, MuZero mastered not only one but also four games.[12] The self-play technique implies that the AI system has been playing against itself without any human knowledge input related to the game of Go, except for the basic rules. Not only did these self-play AI systems outperform humans, but they also devised novel strategies that human players had never attempted. To a certain extent, they developed a machine understanding of the game.

In 2019, again using the self-play method, DeepMind developed an AI system called AlphaStar, which attained grandmaster level (the highest rank reachable by competing with other players) in the real-time, strategy video game StarCraft.[13] The AI community regards this experiment as a leap forward because real-time strategy games such as StarCraft are infamous for their infinite combinations of actions based on imperfect information. After watching or playing against AlphaStar, many professional players reported that it had devised new strategies from which the humans could learn; AlphaStar provided new ways to play the game itself. One commentator reported that observing the AI playing the game was akin to watching a drunken kung fu master performing martial arts: awkward, but outrageously effective.[14] Another professional player commented,

> AlphaStar is an intriguing and unorthodox player—one with the reflexes and speed of the best pros but strategies and a style that are entirely its own. The way AlphaStar was trained, with agents competing against each other in a league, has resulted in gameplay that's unimaginably unusual; it really makes you question how much of StarCraft's diverse possibilities pro players have really explored.[15]

These examples could respond to the AI criticism that machines cannot apply knowledge creatively. As we have seen, not only do machines create novel strategies, but importantly, they also develop strategies with a “machine flavor.” The implications of universal intelligence extend past training advanced AI systems such as AlphaGo and AlphaStar.

In a way, universal intelligence is a non-anthropocentric definition in which *human* no longer provides the measure; instead, human intelligence becomes an instance of universal intelligence, manifested in the human assemblage. However, it is not yet a posthumanist definition of intelligence because it faces two other challenges: *intelligence as a measurement*, and *individualism*.

Autopoiesis and Artificial Neural Networks

From IQ testing to universal intelligence, measurement has been a primary instrument for questions about intelligence. Thus, the problems associated with “intelligence,” such as racism, sexism, classism, and even speciesism, have been products of human exceptionalism in consort with human intelligence as a measurement. For example, critics claim that IQ testing measures nothing but the subject’s test-taking skills at the moment of testing; the result, therefore, has little to do with the individual’s intelligence. The concept of universal intelligence is confronted with similar critiques, which tie intelligence to metaphysical concepts such as conciseness, soul, and free will, all of which lie beyond measurement.[16] However, Legg and Hutter defend their thesis through a tautology of measurement:

> Our goal is to build powerful and flexible machines and thus these somewhat vague properties are only relevant to our goal to the extent to which they have some measurable effect on performance in some well-defined environment. If no such measurable effect exists, then they are not relevant to our objective.[17]

These critiques miss the point; they realized that the problem lies in measurement, but tying intelligence to metaphysical concepts beyond measurement ironically justifies Legg’s and Hutter’s approach as a pragmatic choice. Rather than relying on consciousness or soul as *a priori* for intelligence, we need to critique mathematism and formalism embedded in universal intelligence. We must confront what this measured “intelligence” really is. Perhaps formalizing intelligence helps to build powerful machines, but there is no need to rely on this formalized idea to define intelligence. Indeed, if we ponder the video game player’s comment above on AlphaStar, the reason the professional players believe this AI to be intelligent is not based on measurement or evaluation. Instead, it is a general impression, a belief that AlphaStar possesses certain potentials that allow it to change and adapt.

What precisely are these abilities that we recognize as intelligence? To answer, we need to detach the concept of intelligence from the urge to test and measure. Only then can universal intelligence become a posthuman concept. Universal intelligence implies agents' abilities to achieve *goals*. However, human scientists define these goals, thus imposing a human standard as to whether a behavior counts as an intelligent move. Ultimately, what universal intelligence measures remains perceived effectiveness from a human perspective, even though the concept itself seems to remove humans from the scale of measurement. It is an intelligence *for* humans rather than intelligence *for* the agent itself.

I will employ the concept of autopoiesis to elucidate the meaning of "goal-directedness" in the context of measuring intelligence. In the influential paper "What the Frog's Eye Tells the Frog's Brain," Chilean biologist and philosopher Humberto Maturana, along with his colleagues, revealed that a frog's eye does not simply capture and transmit an exact image to its brain for analysis. Instead, the process involves four distinct groups of nerve fibers, each responsible for a specific type of operation on the visual data. These fibers first process the image and then relay this organized information to the brain. This finding contests the prevalent belief that the nervous system functions primarily as a tool for gathering information to construct an internal representation of the external environment. On the contrary, the environment triggers a set of operations in the nervous system that produce reality inside the organism. Autopoiesis turns the environment outside-in. Based on this biological observation, Maturana and Francisco Varela, also a Chilean biologist and philosopher, argue that:

> [w]e as observers have access both to the nervous system and to the structure of the environment. We can thus describe the behavior of an organism as though it arose from the operation of its nervous system with representations of the environment or as an expression of some goal-oriented process. These descriptions, however, do not reflect the operation of the nervous system itself. *They are good only for the purpose of communication among ourselves as observers*. They are inadequate for a scientific explanation.[18]

The environment merely *triggers* the system to operate and construct a reality of its own. When folding the environment and reality outside-in, concepts such as consciousness, soul, and free will become epiphenomena produced through human system operations. This could be understood as a version of radical constructivism. Goal-directedness becomes an epiphenomenon, too, for humans to describe causal relationships. To use the logic of goal-directedness to describe how and why other systems operate is the privilege of the human observer. We explain system operations in a manner that fits within the descriptive category of goal-directedness, which now becomes a tactic of observation. We choose to believe that goals exist, then observe phenomena that fit within the framework of goal-oriented behaviorism.

However, as Maturana and Varela note, evaluating and measuring goal-directed phenomena has nothing to do with system operations themselves. I will analyze past examples of AI biases with autopoiesis theory to further illustrate what this means. One notable case involved an image-recognition model that erroneously identified husky dogs as wolves. This error was traced back to the training dataset, where images of wolves often possessed snowy backgrounds. The model, therefore, learned to associate snow with wolves, leading to confusion when presented with images of huskies against snowy backgrounds.[19] Other systems are gender-biased. In one example, an image-recognition AI system tended to label human figures at the forefront of a kitchen background as female.[20] In Midjourney, a generative AI system that produces images, if a user gave the prompt "a person in a garden," Midjourney would generate a white woman or a feminine figure in front of a garden scene (Figure 8.1).

FIGURE 8.1 Gender-biased images generated with Midjourney. Images generated by the author through Midjourney.

In these examples, there existed inherent biases in the training data because the images used for training were merely culled from the internet, and thus reflected society's persistent gender stereotypes. If one searches "people gardening," Google's image search will return many images featuring white women gardening. Moreover, these images are likely to be used to train the above AI systems. Further examples abound, and have been used to challenge scientists to create fair and "unbiased" AI.[21] Yet, in early 2024, Google's Gemini created some historically inaccurate images, such as Nazi-era German soldiers as people of color, possibly due to an overcorrection of AI racial bias.[22] Training "unbiased AI" is apparently more complex than we think.

However, the analytical values of these cases were not discussed fruitfully by computer scientists and critics, especially within an autopoiesis framework. If these AI systems are understood as autopoietic systems, then their environments are constructions of their own system operations, which are entirely different from human constructions. If a human is presented with an image of a figure in a kitchen setting and asked to identify the figure's gender, they would ignore the background and turn their attention to the figure itself. However, the machine-learning architecture for most image-recognition algorithms is called convolutional neural networks (CNN). In ANN, artificial neurons are constructed in layers. In CNN, particularly, a convolutional layer acts as a "filter" that passes over the image, scanning several pixels at a time from top to bottom, and left to right, and creating a feature map. Then, a pooling layer checks this feature map and abstracts it into small edges that represent the object in the image. Finally, fully connected layers would make a prediction based on these small edges. The training process uses training data to fine-tune each neuron. A CNN is therefore a perfect example of an autopoietic system; after training, the CNN is hard-wired such that the environment only triggers a set of neural activations. A saying in the AI community notes, "Neurons that fire together, wire together." A CNN does not absorb an image as a whole but as pieces; it is as if a person "reading" a picture used a magnifying glass to search for the small edges that define objects. Using this magnifying glass analogy, a person born holding a magnifying glass would be exposed to a completely different "environment" than the rest of us. Just as Matuana's frog constructs a reality with four groups of fibers, a CNN constructs a reality through a combination of ANN layers.

Based on CNN architecture, scientists combine multiple and various layers for building hierarchies in feature complexity, thus constructing deep convolutional neural networks (DCNN). In DCNN, the first few layers detect edges, the middle layers detect portions of the object, and the final layers detect the object itself. Neuroscientists have found that the increasing feature complexity of DCNN resembles the increasing complexity that occurs in visual object recognition in humans.[23] Others found that the activation of DCNN was similar to that of gamma-band activity within the human visual cortex.[24] Many have duly argued that DCNN resembles the method by which humans see, and some regard DCNN as a biologically inspired approach. These findings seem to contradict autopoiesis theory by drawing similarities between machines and human system operations. On the contrary, these examples prove that DCNN does not resemble humans at all because they perform an operation that merely resembles the *human cortex*, which is only a part of the cognitive assemblage we call human.

We may explain the behavior of a CNN as a goal-directed behavior, but this explanation has nothing to do with system operations; AI systems will execute what we *say* rather than what we *mean*. After training, CNN will label an image, but this action does not equal image recognition. We have no idea what CNN maintains as goals for itself or if *goal* is even a concept within CNN's reality as constructed by its system operation.

The AI community has struggled with the disconnect between goal and system operation. Autopoiesis sheds new light on problems in AI safety research, such as the problem of reward hacking. In an amusing example, a simulated AI robotic arm was trained to flip a pancake; scientists set up a reward function such that the robot would receive a small reward for each second that the pancake did not hit the floor. This sounded logical since a human would continuously flip the pancake to keep it in the air as long as possible. However, the robot arm hurled the pancake into the air with as much force as possible, to maximize the reward.[25] Articulating reward goals and AI behaviors is a constant task in AI research because AI systems may exploit the reward function by performing unexpected behaviors.[26] Powerful AI systems, such as AlphaGo or MuZero, might be understood as gratifying accidents, where the articulated reward goal matches the system operation. This synergy between goals and system behaviors points us in another direction: we can frame intelligence with ideas of co-production and co-evolution.

To consider synergy between scientists and AI systems, we must practice what Niklas Luhman calls *second-order observation*, or the practice of observing others observe. New insights arise when we bring the observer into the equation. When intelligence is tied to measurement, it becomes a product of observation; thus, an observer plays an essential role in conceptualizing intelligence—how intelligent an entity appears to the observer at the moment of observation. In human–robot interactions, *perceived intelligence* is a key concept. Robot engineers know the underlying rules that govern how a robot would behave in a given circumstance; to them, the robot's behaviors are transparent and predictable, thus non-intelligent. However, for a non-expert without prior knowledge of the robot's underlying rules, its behaviors may appear intelligent. Perceived intelligence was used to mediate the knowledge gap between engineers and non-expert users. To make robot behaviors appear more intelligent, engineers would intentionally program randomness so that the behavior patterns were less predictable. However, after many interactions, users will eventually detect patterns in these limited random behaviors, and decide that the robot is not intelligent, after all. Thus, a popular approach in human–robot interaction is the so-called Wizard of Oz trick, in which a human hides behind the robot and conducts a conversation with another human to create the illusion of intelligence.[27]

What separates these two situations—the Wizard of Oz versus limited random actions? We may conclude that users find the robot unintelligent because its behaviors become predictable over time. However, we could draw a different conclusion if we absorb observers into the equation: the users become more intelligent over time, discerning the robot's deceptive behaviors. The robot is unintelligent because it cannot co-evolve with the users, whereas another human hiding behind the robot adapts to the users and advances the conversation.

Similarly, if we reconsider the Turing test, the seemingly innocent discriminator becomes a critical player in the imitation game because the game eventually becomes a test for all parties. Recursive observation and learning between observer and participants will force them to adapt to each other and become more intelligent in playing the imitation game. Measurement overlooks the intelligence emerging from dynamic feedback loops between the observer and the observed. From this perspective, intelligence points to a direction on the line of co-production and co-evolution within an assemblage framework. Or, to use Donna Haraway's term, *sympoiesis*—making with.[28]

Sympoiesis and the Emergence of Intelligence

With current AI advancement, we lack vocabularies and concepts for a non-individualistic view of intelligence, in both theory and practice, because AI was originally conceived as an autonomous machine. Collaborative intelligence is a concept proposed in AI research to challenge the rivalry between humanity and AI constructed over decades by popular culture and media. It posits that a synergy between people and AI is the basis for considering machine intelligence.[29] We should extend this argument to include all types of human and nonhuman agents, and construct a posthumanist definition of intelligence as the basis for AI research.

Computer scientists regard an agent as a black box with input and output; similarly, researchers regard intelligence as a person's capacity in IQ testing. Assemblage thinking exists at odds with the inherent individualism of how intelligence is conceived—a quality of a pre-existing agent. However, as I have explored in previous chapters, an agent is simply a product of observation to attribute perceived effectiveness distributed *within* and *around* the so-called agent. It is merely a carrier of intelligence that has emerged from dynamic co-evolution and co-production processes involving actors more than the agent itself.

One way to observe the emergence of intelligence is to construct an arbitrary diagram of an agent and its outside environment, and then analyze both inside and outside the agent. First, intelligence is an emergent property of the interactions between the components that give rise to the assemblage. For example, swarm intelligence points to the collective behavior of decentralized, self-organized systems, such as a colony of ants. While exploring their environment, ants leave pheromones to guide each other to resources. The behavior of an ant appears simple, but if each ant repeats the same actions, a colony of ants appears remarkably efficient in exploring its environment. This behavior inspired a famous instance in computer science, called the ant colony optimization algorithm, which has proven useful in pathfinding optimization problems.[30] This type of intelligence, also known as "hive mind," occurs where intelligence is not located within any of the individuals, but as a property that emerges from the interactions of each part.

In a way, all forms of intelligence—human, machine, or animal—should be understood as emergent phenomena. Human and animal intelligence are nothing but electrochemical interactions between neurons and a distributed nervous system. Machines are the same; as we saw in the above CNN example, in ANNs, AI is achieved through the interactions of layers of interconnected artificial neurons, or functions between input and output. An AI system is thus a conglomeration of connected functions.

Second, outside the boundary of the agent, the term *mediation* may help us conceptualize that an agent's intelligent behaviors are constantly mediated by other assemblages. The mediation may be understood through Andy Clark's "extended mind" theory, which reminds us that the effectiveness of human intelligence is achieved by off-loading human cognitive functions onto other objects in the environment. Human intelligence becomes a continually changing phenomenon, depending on the temporary and ephemeral assemblage at the moment of observation. A paper-based IQ test evaluates the effectiveness of finishing the questions by the temporary assemblage of human, pen, paper, and perhaps corrective lenses; all factors contribute to how well the subject performs at the test-taking moment. We must also account for the breakfast eaten that morning, as well as the diverse bacteria and enzymes in the test-taker's digestive system, which help process food into energy for the electrochemical nerve impulses that give rise to "intelligent behaviors."

On the other hand, we may add a time scale to the mediation process and see its effects as co-evolution and co-production. In the story of transistors and Boolean algebra, the users of the humanoid robots and the discriminator in the Turing test are all instances where the human and the machine become each other's media and co-evolve over time. We need to reframe intelligence as a collection of observed phenomena that is, in fact, a result of random interactions in co-production and co-evolution among assemblages.

To understand intelligence based on the inside and outside of the agent is arbitrary. In fact, it is crucial to recognize that reality is a continuous mesh of heterogeneous bodies within the assemblage framework. Random interactions and feedback loops exist between different agents, and if certain interactions become useful for a temporary assemblage, then synergies and loose couplings would happen among these bodies. Such thinking is found in many contemporary feminist materialists' writing. As Jane Bennett noted, "bodies enhance their power in or as a heterogeneous assemblage."[31] This process is also essential to Donna Haraway's notion of *sympoiesis* (making with); nothing is truly autopoietic and self-organizing, and things are always "making with" others.[32] Sympoiesis produces intelligence.

In a posthumanist, sympoietic framework, intelligence becomes a shared ability among different assemblages to exchange effects through recursive observation and learning, thus forming synergies and couplings to gain power. Intelligence speaks to the process of attunement and symbiosis between assemblages, as in the moment when computer scientists' goals accidentally match operations in an AI system. We can frame this phenomenon as "intelligence of co-production" or "co-productive intelligence." What we perceive as intelligence—whether human, machine, or animal—becomes a specific instance of observed local manifestation of co-productive intelligence distributed across the heterogeneous landscape of human and nonhuman, biotic and abiotic assemblages.

Co-Productive Intelligence

To summarize this part, I propose three types of co-productive intelligence as an attempt to depart from how intelligence is currently discussed: (1) intelligence in adversarial relations, (2) intelligence in symbiosis (symmetrical and asymmetrical), and (3) intelligence in loose coupling.

Adversarial intelligence gestures to competition in long-term evolution, in which complex behaviors and responsive strategies emerge. In examples like AlphaGo and AlphaStar, humans and AI agents participated in competitive environments, and new strategies emerged; together, humans and machines expanded the possibilities of a game.

In recent years, the AI community has exploited adversarial relations and achieved promising results. Many machine-learning approaches involve two adversarial parts training together. One example is that of generative adversarial networks (GANs). Before diffusion models (e.g., DALL-E.2 and Midjourney) took over image generation in mid-2022, GANs were the go-to solution for computer scientists. There, GANs might be demonstrated with a forgery example. Imagine two neural networks, one generator (forger) and one discriminator. The forger's goal is to produce photorealistic images to deceive the discriminator, and the goal of the discriminator is to ascertain whether it is a real image. The two networks are trained simultaneously until the discriminator cannot distinguish a fake image from a real image. Researchers have used this technique to train AI systems to render photorealistic images and stylized paintings.[33] In 2018, GANs was named one of the "10 breakthrough technologies" by the Massachusetts Institute of Technology (MIT); a commenter noted that GAN "gives machines something akin to a sense of imagination, which may help them become less reliant on humans."[34]

Even though diffusion-based systems, including DALL-E.2 and Midjourney, quickly superseded GAN-based image generators, the concepts embedded in GANs are not completely obsolete. For example, Variational Autoencoder (VAE), a key component of the neural network model used in Stable Diffusion to improve the quality of AI-generated images, essentially assumed two neural networks (encoder and decoder) in an adversarial relationship. Similarly, in RL, Deep Deterministic Policy Gradient (DDPG) is a technique that establishes an actor-critic architecture within the learning agent, in which the actor decides which action to take in an environment, and the critic evaluates the action and informs the actor how effective the action was and how it could be improved.

Moreover, computer scientists have found that competitive multi-agent techniques in RL not only train models faster but they also give rise to behaviors far more complex than those allowed by the environment alone.[35] In an example developed by OpenAI, two agents were asked to play hide-and-seek in a simulated environment with walls, boxes, and boards. After training, both agents (seeker and hider) learned to use these tools to their advantage and began to develop methods to exploit the environment in which they were trained. One agent learned to stand on a box and apply force to the agent itself; it then "surfed" the boxes and jumped over walls.[36]

Intelligence in symbiosis or structural coupling is the second type of co-productive intelligence. The evolution from prokaryotic cells to eukaryotic cells has been widely discussed among posthumanism scholars.[37] It is now believed that mitochondria were once a type of bacteria before being co-opted as permanent organelles of cells in which they were originally parasites. It seems that adversarial relationships may transform into symbiotic relationships where two or more competing assemblages self-synchronize in terms of their inputs and outputs so that they can exchange effects. We may find the relationship of symbiosis in many plant communities. For example, root nodules are primarily found on the roots of legumes, or the pea family (*Fabaceae*), including peas and soybeans, as well as on trees such as the black locust (*Robinia pseudoacacia*). These root nodules are the result of legumes forming a symbiosis with nitrogen-fixing bacteria that helps convert dinitrogen (N_2) from the atmosphere into ammonia (NH_3), which can then be used by plants. In ANNs, neurons (non-linear functions) are structurally coupled in layers and networks. One neuron (function) cannot do much, but structurally coupled neural networks become intelligent.

Furthermore, OOO and other posthumanist ideas expand symbiosis beyond biology as a metaphor for how living entities form symbiotic relations with other assemblages, including non-material organizations, institutions, and historical objects.[38] Symbiotic relationships involve *structural coupling*, first discussed by Maturana and Varela when introducing autopoiesis theory. Levi Bryant provides a more accessible explanation of structural coupling by describing it as “a relation in which one or two entities are dependent for stimuli or flows from one another in order to engage in their own operations and becoming.”[39] Bryant divides structural coupling into bidirectional and unidirectional; this is similar to Harman’s symmetrical and asymmetrical symbiosis. Symmetrical symbiosis, or bidirectional coupling, occurs when both assemblages or systems require flows from the other to engage in a joint operation, such as mitochondria acting as permanent organelles in living cells. In asymmetrical symbiosis, or unidirectional coupling, one system depends entirely on the other; an ecosystem type may be said to be unidirectionally coupled with the climate in its region. Asymmetrical symbiosis or unidirectional coupling as a metaphor speaks to adaptive strategies which a system develops to take advantage of another system. In symbiosis or structural coupling, intelligence manifests as the ability to link one system to another so that its own system operation becomes assimilated into another system’s operation; both systems benefit from the assimilation of system operations, and both gain power.

Intelligence in loose coupling is the third type of co-productive intelligence. Computer scientists and system engineers use coupling to describe the interdependence between different programming objects or systems. In programming, there are two types of coupling—tight and loose—with the latter preferred for greater flexibility so that altering one programming object or system is unlikely to affect another. Similarly, sociologist Niklas Luhmann, who used structural coupling to describe social systems, argued that “[l]oosely coupled systems are more stable than tightly coupled ones. ‘Tight coupling’ is a very improbable arrangement. It is not to be found in nature.”[40] Many scholars disagree with Luhmann’s claim regarding loosely coupled systems.[41] Symmetrical symbiosis may be understood as a type of tight coupling, and symmetrical symbiosis is found nearly everywhere in “nature.” Certainly, in symmetrical symbiosis, if one of the tightly coupled systems (for example, the soybean) is destroyed, then the other tightly coupled system (here, nitrogen-fixing bacteria) can no longer survive. Viewing loose coupling as beneficial is based on its resistance to external disturbance, overlooking the intensive exchange of flows and functions in tightly coupled systems that do not exist in loosely coupled systems. Instead of placing a high value on loose coupling, it can be understood as another type of relationship from which intelligence emerges.

Loose coupling is a type of relationship we commonly encounter, such as the relationship between computer scientists and machine-learning algorithms. The synergies between these assemblages could be seen as "happy accidents" when scientists' goals match the observed system operation in an AI algorithm, and thus effectiveness comes into being.

These three types of co-productive intelligence attempt to create a set of vocabularies with which one could consider the role of intelligent machines within a posthumanist, sympoietic framework. I will illustrate this point further with Sougwen Chung's artwork "Drawing Operations," a series of performances in which the artist drew alongside robots (Figure 8.2).[42] In the first version, the artist drew alongside a robot arm called D.O.U.G_1 (Drawing Operations Unit: Generation_1). D.O.U.G_1 mimicked the artist's movements by analyzing her drawing gestures through an overhead camera and reproducing them. The final artifact is a co-produced drawing through collaboration. However, the robot's movements were not perfect reproductions of the artist's gestures. Although the algorithm tracked the artist's linework in the digital simulation, the movements were dramatically altered when translated to a robotic arm.

Imperfection was unavoidable in this loosely coupled system. The computationally heavy, live, computer-vision analysis created an inevitable system

FIGURE 8.2 Top: "Drawing Operations" (2015); bottom: "Drawing Operations" (2017). Courtesy of Sougwen Chung.

latency, creating a lag between the artist's movements and the robot's reproduction of them. The linework quivered as if the robotic hand was unsteady. In real-time, the artist was forced to adapt her movement to the robot's. Her new gestures fed back to the robot, and the robot produced a new set of movements to which the artist was again compelled to adapt. The artist and the robot formed a positive feedback loop; together, they became a loosely coupled system with new styles and techniques manifested in the co-produced artifact, gradually synchronizing and attuning to each other in a territory alien to both.

Later, the D.O.U.G_2 was designed around the notion of memory. Chung and her team deployed machine-learning techniques to teach the artist's drawing style to an AI system. They fed an algorithm with decades of the artist's work so that the AI would attempt to reproduce Chung's drawing patterns. The machine subsequently developed its own understanding of the artist's style, and expressed a machine interpretation in their later drawing collaborations.

Chung's art installations raise questions about perceived errors and glitches in systems, reframing the notion of failure as an important factor in art production and design. As the artist noted, "The robot mimics the artist like a partner in an improvisational performance. It is an AI that embraces every glitch, bug, and error."[43] From a posthumanist perspective, these so-called glitches, bugs, and errors are fundamental facets of operation defined by an algorithm and the physical armature of a robotic arm. To human eyes, they may appear to be errors, but they are the ways in which the robotic systems operate within the environment. In theory, we might minimize system latency by using a faster processing unit and by accounting for the robotic arm's physical limitations. However, if one repairs all the glitches, bugs, and errors so that the robot perfectly replicates every detail of the artist's gestures, the art does not exist. The results become anticipated and predictive, while Chung's art lies in the unexpected outcomes of human–robot collaboration. When D.O.U.G_2 learned the artist's style through her works, there were no right or wrong answers for interpretation because they were merely machine interpretations. If the machine replicates the artist's drawing style perfectly through learning, then the machine hardly differs from a photocopier. Thus, adaptation between the artist and the robot entails exploiting each other's limitations and errors. Or, to use Chung's words, her artwork is about "embracing the imperfections and recognizing the fallibility of both human and machine in order to expand the potential of both."[44] As discussed earlier, the foundation for the emergence of intelligence is co-production between assemblages. In Chung's work, intelligence manifests in the loose coupling and attunement between humans and machines. In this attuning process, the potentials and possibilities of the human-machine assemblage expand.

In the cybernetic environment, our relationship with machines must not be limited to the co-dependent relationship between users and tools. Instead, artistic explorations suggest that potential may be found in interdependent relationships, where new questions arise, new understandings form, and new strategies emerge. With the conviction that machines are not tools for automation but actors that have been involved in the co-production of the environment, we may build an alternative way to understand failures and errors when collaborating with intelligent machines. When machines produce unexpected strategies, those are not necessarily errors. They are, instead, opportunities for us to form different understandings of the questions posed, just as AlphaGo and AlphaStar have expanded players' understanding of their games.

A Posthumanist Perspective on AI Safety and Bias

I will end this chapter by providing notes on AI safety and bias from a posthumanist perspective because I have not read such analyses in my research. Viewing machine intelligence as a threat to humanity is a popular sentiment, especially after the twentieth century's two world wars, when society witnessed the potential destructive power of technology. Popular culture also aided the growth of this sentiment, since AI often assumes a supervillain role in dystopian novels and movies. Imagining rogue AI is a form of the technological sublime, in which we entertain a potential threat from a safe position. We comprehend the threat with human reasoning and rationality, which, according to Kant, is the basic structure of sublimity. However, this view is deeply rooted in an anthropocentric understanding of human intelligence and agency. Recognizing the sublime quality of machines extends the anthropocentric understanding of human agency. In turn, it adds to a sense of human hubris, since entertaining the potential threat of machines ironically reflects a sense of human control and mastery over machines.

What we understand as human agency has always been distributed in a heterogeneous landscape of human and nonhuman components. These components are biological and technological, including tools, machines, technological systems, models, algorithms, and ANNs. We have always been cyborgs, and we have always explored the environment in company with other actors. Machines have helped us expand our understanding of the environment, and the rejection of their perspectives of the environment is rebuffing what it means to be human. Thus, embracing machine perspectives is to accept the cyborg condition of being human.

I argue that adopting a posthumanist understanding of intelligence could greatly expand the discourse of research on AI safety and bias. Over the years, algorithms trained by AI research clusters, such as Google and Meta, have resulted in "biased" AI systems that uncannily resemble humanity's biases. In 2015, a Google image-recognition algorithm placed the label "gorilla" on Black people. In 2016, searching female names, such as "Stephanie Williams," on LinkedIn triggered search suggestions asking if the user sought similar male names, such as "Stephen Williams." That same year, Tay, a Microsoft chatbot, spent a day learning from Twitter, after which it began to tweet racist and sexually-charged messages. Other examples include gender-biased image-recognition and -generation algorithms, as we have seen earlier in this chapter. As we now recognize, AI biases are reflected from everyday gender and racial stereotypes.

According to their own accounts, Google and Facebook quickly "fixed" these "biased" algorithms after the public launched scathing criticisms. Researchers wanted to create "unbiased" AI algorithms. However, a deeper irony lurks behind these stories. From a posthumanist perspective, these AI algorithms are not biased at all, because they hold true to the datasets on which they were trained by cruelly reflecting an ugly aspect of humanity which many of us dare not recognize. In a way, to teach an algorithm to be "racial-neutral" equates to teaching someone to see no skin color. Yet "seeing no color" is a problem in itself, one of society's ignorance of the reality of systemic racism.

At the center of this irony is that we seek to treat AI algorithms as automation machines, instead of as yet more voices on issues we thought to address solely as humans. If we focus on the bigger picture, these blunt and "biased" AI systems played an important role in recognizing our biases in online spaces such as Twitter and LinkedIn. There is a saying in the machine-learning community: "Garbage in, garbage out." Yet, a garbage AI reveals how the internet has become a hoarding place for humanity's mental garbage.

A similar argument can be made concerning image generators that have disrupted creative professions. The work "Théâtre D'opéra Spatial" dropped a bomb onto the art professions that brokered a war between artists and AI image generators. Created using Midjourney, Jason Allen submitted this work to the Colorado State Fair's annual art competition using the name "Jason M. Allen via Midjourney." He won the blue ribbon for emerging digital artists. While Allen unapologetically boasted of his achievement, the news set off a backlash that eventually led to a mass protest of artists against AI-generated artwork. Because AI researchers scraped the internet for training data, including images by these protesting artists, the trained AI system would generate similar visual styles. Although rules and regulations concerning training data in AI research must be encouraged, a posthumanist twist is needed to bypass the inherent human-machine rivalry narratives. To some degree, the boycotting human artists must have felt insecure when many non-artists exclaimed that AI image generators would one day replace them. Such a comparison is based on an anthropocentric view, in which human artists become the measure of creativity and imagination. From a posthumanist perspective, the goal of building powerful machines should never be about replacing or replicating human capacity, but about providing alternative, non-anthropocentric perspectives to expand our all-too-human pathways of thinking and making.

Notes

1 Merriam-Webster, "Definition of INTELLIGENCE." https://www.merriam-webster.com/dictionary/intelligence
2 Turing, "Computing Machinery and Intelligence."
3 Premack, "Human and Animal Cognition."
4 Trewavas, "Aspects of Plant Intelligence."
5 Belkhir, "Race, Sex, Class & 'Intelligence' Scientific Racism, Sexism & Classism"; Dennis, "Social Darwinism, Scientific Racism, and the Metaphysics of Race"; McNally, "Scientific Racism and The Politics of Looking."
6 Krizhevsky, Sutskever, and Hinton, "Imagenet Classification with Deep Convolutional Neural Networks."
7 Russakovsky et al., "ImageNet Large Scale Visual Recognition Challenge."
8 Langlotz et al., "A Roadmap for Foundational Research on Artificial Intelligence in Medical Imaging."
9 Legg and Hutter, "A Collection of Definitions of Intelligence."
10 Legg and Hutter, "Universal Intelligence."
11 See https://www.youtube.com/watch?v=iI6TQOQ_Ccc for an example of a RL-based agent playing Tetris. Last accessed: April 29, 2024.
12 Schrittwieser et al., "Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model."
13 Vinyals et al., "Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement Learning."
14 Two Minute Papers, "DeepMind's AlphaStar: A Grandmaster Level StarCraft 2 AI - YouTube." https://www.youtube.com/watch?v=jtlrWblOyP4. Last accessed: April 29, 2024.
15 The AlphaStar Team, "AlphaStar." https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii. Last accessed: April 29, 2024.
16 Legg and Hutter, "Universal Intelligence."
17 Legg and Hutter, 42.
18 Maturana and Varela, *The Tree of Knowledge*, 129. Emphasize added.
19 Ribeiro, Singh, and Guestrin, ""Why Should I Trust You?"
20 Zhao et al., "Men Also like Shopping."
21 Computer scientists tend to avoid difficult social science questions regarding gender equality and justice; most of the time, unbiased AI means reflecting whatever the training data represents, excluding issues of gender equality outside the scope of the research. However, this "unbiased" attitude limits the potential for an AI system to be used creatively, for activism.
22 See a media report here: https://www.theverge.com/2024/2/21/24079371/google-ai-gemini-generative-inaccurate-historical. Last accessed: April 29, 2024.
23 Kriegeskorte, "Deep Neural Networks."
24 Kuzovkin et al., "Activations of Deep Convolutional Neural Networks Are Aligned with Gamma Band Activity of Human Visual Cortex."
25 See a video here:https://dzamqefpotdvf.cloudfront.net/p/images/2cb2425b-a4de-4aae-9766-c95a96b1f25c_PancakeToss.gif._gif_.mp4. Last accessed: April 26, 2024.
26 A list of unexpected AI behaviors https://docs.google.com/spreadsheets/d/e/2PACX-1vRPiprOaC3HsCf5Tuum8bRfzYUiKLRqJmbOoC-32JorNdfyT-

iRRsR7Ea5eWtvsWzuxo8bjOxCG84dAg/pubhtml. Last accessed: April 26, 2024.
27 Bartneck et al., "Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots."
28 Haraway, *Staying with the Trouble*.
29 Epstein, "Wanted: Collaborative Intelligence."
30 Dorigo and Di Caro, "Ant Colony Optimization."
31 Bennett, *Vibrant Matter: A Political Ecology of Things*, 23.
32 Haraway, *Staying with the Trouble*.
33 Goodfellow et al., "Generative Adversarial Networks."
34 MIT Technology Review, "2018 MIT Technology Review." https://www.technologyreview.com/10-breakthrough-technologies/2018/. Last accessed: April 29, 2024.
35 Bansal et al., "Emergent Complexity via Multi-Agent Competition."
36 https://www.youtube.com/watch?v=kopoLzvh5jY&ab_channel=OpenAI. Last accessed: April 29, 2024.
37 Authors such as Manual DeLanda, Graham Harman, and Levi Bryant all use this example to demonstrate the process of forming assemblages.
38 Harman, *Object-Oriented Ontology*, 112.
39 Bryant, *Onto-Cartography: An Ontology of Machines and Media*, 153.
40 Luhmann, *Introduction to Systems Theory*, 123.
41 Probert, "Book Review: Introduction to Systems Theory."
42 Drawing Operations, 2015; Drawing Operations, 2017
43 Chung, *Drawing Operations*, 2015.
44 Watch her full TED Talk here: https://www.youtube.com/watch?v=q-GXV4Fd1oA. Last accessed: April 29, 2024.

Bibliography

Bansal, Trapit, Jakub Pachocki, Szymon Sidor, Ilya Sutskever, and Igor Mordatch. "Emergent Complexity via Multi-Agent Competition." *arXiv:1710.03748 [Cs]*, March 14, 2018. http://arxiv.org/abs/1710.03748.

Bartneck, Christoph, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. "Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots." *International Journal of Social Robotics* 1, no. 1 (January 2009): 71–81. https://doi.org/10.1007/s12369-008-0001-3.

Belkhir, Jean. "Race, Sex, Class & 'Intelligence' Scientific Racism, Sexism & Classism." *Race, Sex & Class* 1, no. 2 (1994): 53–83.

Bennett, Jane. *Vibrant Matter: A Political Ecology of Things*. Durham, NC: Duke University Press, 2010.

Bryant, Levi R. *Onto-Cartography: An Ontology of Machines and Media*. Speculative Realism. Edinburgh: Edinburgh University Press, 2014.

Chung, Sougwen. *Drawing Operations*. 2015. Mixed media. https://sougwen.com/project/drawing-operations.

———. *Drawing Operations*. 2017. Mixed media. https://sougwen.com/project/drawingoperations-memory.

Dennis, Rutledge M. "Social Darwinism, Scientific Racism, and the Metaphysics of Race." *The Journal of Negro Education* 64, no. 3 (1995): 243–52. https://doi.org/10.2307/2967206.

Dorigo, M., and G. Di Caro. "Ant Colony Optimization: A New Meta-Heuristic." In *Proceedings of the 1999 Congress on Evolutionary Computation-CEC99 (Cat. No. 99TH8406)*, 2 (1999): 1470–77. https://doi.org/10.1109/CEC.1999.782657.

Epstein, Susan L. “Wanted: Collaborative Intelligence.” *Artificial Intelligence* 221 (April 2015): 36–45. https://doi.org/10.1016/j.artint.2014.12.006.

Goodfellow, Ian J., Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. “Generative Adversarial Networks.” *arXiv:1406.2661 [Cs, Stat]*, June 10, 2014. http://arxiv.org/abs/1406.2661.

Haraway, Donna J. *Staying with the Trouble: Making Kin in the Chthulucene*. Durham, NC: Duke University Press, 2016.

Harman, Graham. *Object-Oriented Ontology: A New Theory of Everything*. London: Penguin, 2018.

Kriegeskorte, Nikolaus. “Deep Neural Networks: A New Framework for Modeling Biological Vision and Brain Information Processing.” *Annual Review of Vision Science* 1, no. 1 (2015): 417–46. https://doi.org/10.1146/annurev-vision-082114-035447.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. “Imagenet Classification with Deep Convolutional Neural Networks.” In *Advances in Neural Information Processing Systems*, 1097–1105, 2012.

Kuzovkin, Ilya, Raul Vicente, Mathilde Petton, Jean-Philippe Lachaux, Monica Baciu, Philippe Kahane, Sylvain Rheims, Juan R. Vidal, and Jaan Aru. “Activations of Deep Convolutional Neural Networks Are Aligned with Gamma Band Activity of Human Visual Cortex.” *Communications Biology* 1, no. 1 (August 8, 2018): 1–12. https://doi.org/10.1038/s42003-018-0110-y.

Langlotz, Curtis P., Bibb Allen, Bradley J. Erickson, Jayashree Kalpathy-Cramer, Keith Bigelow, Tessa S. Cook, Adam E. Flanders, et al. “A Roadmap for Foundational Research on Artificial Intelligence in Medical Imaging: From the 2018 NIH/RSNA/ACR/The Academy Workshop.” *Radiology* 291, no. 3 (June 2019): 781–91. https://doi.org/10.1148/radiol.2019190613.

Legg, Shane, and Marcus Hutter. “A Collection of Definitions of Intelligence.” *arXiv:0706.3639 [Cs]*, June 25, 2007. http://arxiv.org/abs/0706.3639.

———. “Universal Intelligence: A Definition of Machine Intelligence.” *Minds and Machines* 17, no. 4 (December 1, 2007): 391–444. https://doi.org/10.1007/s11023-007-9079-x.

Luhmann, Niklas. *Introduction to Systems Theory*. 1 edition. Cambridge; Malden, MA: Polity, 2012.

Maturana, Humberto R., and Francisco J. Varela. *The Tree of Knowledge: The Biological Roots of Human Understanding*. The Tree of Knowledge: The Biological Roots of Human Understanding. Boston, MA,: New Science Library/Shambhala Publications, 1987.

McNally, Cayla. “Scientific Racism and The Politics of Looking.” In *Jordan Peele’s Get out: Political Horror*, edited by Dawn Keetley, 212–22. New Suns: Race, Gender, and Sexuality in the Speculative. Columbus: The Ohio State University Press, 2020.

Merriam-Webster. “Definition of INTELLIGENCE.” Accessed August 16, 2020. https://www.merriam-webster.com/dictionary/intelligence.

MIT Technology Review. “2018 MIT Technology Review.” MIT Technology Review, 2018. https://www.technologyreview.com/10-breakthrough-technologies/2018/.

Premack, David. “Human and Animal Cognition: Continuity and Discontinuity.” *Proceedings of the National Academy of Sciences* 104, no. 35 (August 28, 2007): 13861–67. https://doi.org/10.1073/pnas.0706147104.

Probert, Stephen K. "Book Review: Introduction to Systems Theory." *International Journal of Systems and Society* 1, no. 1 (2014): 55–57.

Ribeiro, Marco Tulio, Sameer Singh, and Carlos Guestrin. "'Why Should I Trust You?': Explaining the Predictions of Any Classifier." *arXiv:1602.04938 [Cs, Stat]*, August 9, 2016. http://arxiv.org/abs/1602.04938.

Russakovsky, Olga, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, et al. "ImageNet Large Scale Visual Recognition Challenge." *International Journal of Computer Vision* 115, no. 3 (December 1, 2015): 211–52. https://doi.org/10.1007/s11263-015-0816-y.

Schrittwieser, Julian, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, et al. "Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model." *Nature* 588, no. 7839 (December 2020): 604–9. https://doi.org/10.1038/s41586-020-03051-4.

The AlphaStar Team. "AlphaStar: Grandmaster level in StarCraft II using multi-agent reinforcement learning." Deepmind, October 30, 2019. https://deepmind.com/blog/article/alphastar-mastering-real-time-strategy-game-starcraft-ii.

Trewavas, Anthony. "Aspects of Plant Intelligence." *Annals of Botany* 92, no. 1 (July 1, 2003): 1–20. https://doi.org/10.1093/aob/mcg101.

Turing, A. M. "Computing Machinery and Intelligence." *Mind, New Series* 59, no. 236 (1950): 433–60.

Two Minute Papers. "DeepMind's AlphaStar: A Grandmaster Level StarCraft 2 AI - YouTube," 2019. https://www.youtube.com/watch?v=jtlrWblOyP4.

Vinyals, Oriol, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, et al. "Grandmaster Level in StarCraft II Using Multi-Agent Reinforcement Learning." *Nature* 575, no. 7782 (November 2019): 350–54. https://doi.org/10.1038/s41586-019-1724-z.

Why I Draw with Robots, 2020. https://www.youtube.com/watch?v=q-zGXV4Fd1oA&t=426s.

Zhao, Jieyu, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. "Men Also like Shopping: Reducing Gender Bias Amplification Using Corpus-Level Constraints." arXiv Preprint arXiv:1707.09457, 2017.