Early Machine Learning and Artificial Animal Intelligences

James E. Dobson

Anthropomorphism is pervasive in contemporary discourse about artificial intelligence. It is now common for researchers and the corporations developing and marketing generative AI (GenAI) and AI-enhanced products to describe their models as if they possess the cognitive abilities of a human at a particular developmental stage. A company might claim, for instance, that the latest version of its model or chatbot has the intelligence of a secondary-school graduate. The next version, right on the horizon, we are promised, will possess PhD-level intelligence. While these markers are comparable, at least in form, to industry or research benchmarks that track performance on standardized tasks, the attempt to signal the intellectual "age" of a model or application falls short because these supposedly determinative markers (progress through formal education or demarcating age ranges) are themselves uneven, vague, and ill-defined. While similar to algorithmically determined assessment systems such as the well-known Flesch–Kincaid readability test, which scores texts via key parameters including sentence, word, and syllable counts to assign reading levels from U.S. fifth grade to professional (Kincaid), AI "grades" also advertise complexity but without the well-defined, if still problematic, formula used by traditional weighted metrics.

The labeling of machine learning models and artificial intelligence applications with such graduated comparisons to human performance is not new. These descriptive categories and human education-level comparisons have been employed since almost the beginning of machine learning to measure performance, attract funding agencies and users, and structure the protocols of experiments. Perhaps surprisingly, comparisons were also made to animal perceptual performance during the critical launching point of machine learning, in the mid-1950s and early 1960s, when researchers designed and implemented neural networks based on research derived from experimentation with various animal species. These networks, which persist in a more complex form in the present, were often explicitly referred to as artificial animal intelligence by their creators. In order to understand how this pivotal moment in the history of AI casts its influence on the present and our future, it is essential to reconstruct and critique the relays that were created through the early empirical research that linked together animal, child, and machine performance. This original work continues to impact the function and evolution of artificial neural networks (ANNs) and has also shaped the way we assess the efficacy of modern evaluation metrics.

Examining the prevalence of biological references in computing, social scientist Nigel Thrift writes that "[a]lmost since its inception, the biological and informational have been intertwined in software. Right at the birth of the modern computer, the new machines were framed in biological terms" (463). On his way toward linking the presence of these metaphors to the advent of animal-like companion or pet robots—what he calls "electric animals"— Thrift briefly mentions neural networks in the context of connectionism. While Thrift's larger claim, that there might be ethical stakes to the understanding of reconfigured forms of dominance and agency deployed in interactions with electronic companion animals, seems only more compelling and urgent in light of the new forms of human-chatbot relations produced through GenAI, the animal referent serves as much more than metaphor in the genealogy of the neural networks powering this technology. The animal provided the model because it was understood to be a simplification of the more complex target, the human. The animal thus exists somewhere between what is possible in contemporary artificial intelligence and the fantasy goal of artificial general intelligence, represented by human or beyond-human intelligence. As model and metric, in early neural network research the animal shaped and determined what professor N. Katherine Hayles calls the "cognitive assemblages" (8) of this field. This leads to a further destabilization of the human supremacy and exceptionalism that is the target of Hayles's critique: "Even when humans appear to be in control, their assumptions have been formed and mediated by prior decisions and interpretations made by computational media, so that they can scarcely be considered autonomous or self-determining" (3). ANNs, as computational media, are yet another site in which, as Thrift writes of Sony's robot dog AIBO, "theoretical models of biology come back to haunt the surfaces that define us" (461).

Neural networks are no longer solely the province of tech-industry conferences and computer-science research; they are now core components of the large and complex deep learning systems that power a wide array of consumer-facing applications. As a result, neural networks and the discourse surrounding the meaning and capabilities of artificial intelligence have become an increasingly visible part of everyday life. Despite other possible genealogies, approaches, and intellectual contexts, the "artificial" in artificial intelligence has essentially come to mean ANNs, and they are called "artificial neural networks" because the general design and specific architectures take as their inspiration formal descriptions of highly simplified simulations of biological neurons. This origin story makes sense in part because ANNs were originally developed as models of brains, to help psychologists and physiologists understand human perception. Despite this design goal—of building research instruments that could model cognitive processes—ANNs were originally the result of insights drawn from research conducted on animals, not humans. The advances made in deep learning as a result of complicated new architectures and training procedures and the impressive abilities of contemporary generative AI are all indebted to these early neural networks, based on animal-perception models developed in the 1950s and 1960s. But to what degree does this history reside in today's neural networks?

That ANNs are simulations of biological processes derived from what was known about animal perception in the mid-twentieth century does not mean that these models "see" the world as an animal would see them. This is to say, theories about the possibility of contemporary deep learning networks reasoning or even displaying so-called general artificial intelligence are not anchored to research of animals' (let alone humans') actual consciousness or visual systems. Animal models, however, did inspire and influence the architecture and design of neural networks in the form of suggestive organizational strategies for the components. Furthermore, animal research provided baselines for the evaluation of these models. In shedding light on the previously overlooked cross-functional relay between two field-defining mid-twentieth-century researchers at Cornell University, Frank Rosenblatt and Eleanor J. Gibson, we can revise our understanding of the development of AI and AI imaginaries through the zoomorphic and anthropomorphic figurations of machine intelligence. Gibson was a psychologist studying perceptual learning in animals and children, who became famous for her "visual cliff" studies, in which she studied depth perception via an apparatus that produced the illusion of sudden drop as a subject crosses a transparent surface (Robles-Anderson and Ferguson). Rosenblatt, also trained as a psychologist, was the creator of the best-known prototype of neural networks for pattern matching and machine learning, the Perceptron (Dobson, "Objective Vision"). Rosenblatt's extended research program, which was cut short due to his premature death and a loss of funding from the Office of Naval Research, provides the locus for this rich intersection of animal, child, and machine perception.

Through our modern lens, the earliest neural-network models, long claimed to have been inspired by animate nervous systems, might not bear much resemblance to the biology they purport to model. An oft-cited origin for ANNs is Warren S. McCulloch and Walter Pitts's 1943 "A Logical Calculus of the Ideas Immanent in Nervous Activity," published in the *Bulletin of Mathematical Biophysics*. In their article McCulloch and Pitts propose a theory of "nervous nets" and describe, in logical terms, the operation of their highly simplified neurons (100). Philosopher of science Matteo Pasquinelli turns received accounts of the McCulloch–Pitts neuron on their head by arguing that the pair "saw, in the first instance, biological neurons as technological artefacts" and that they "implicitly envisioned brain physiology as homologous with the communication technology of the age, comprised of electromechanical relays, feedback mechanisms, television scanners, and notably, telegraph networks" (Pasquinelli 136–37). To a degree, such was understood at the time. Walter Maurice Elsasser, a physicist and biologist, gives a comparable account of the brain-computer analogy in his application of insights from computing in *The Physical Foundation of Biology* (1958):

> When the histologist looks at the brain he sees something which is very reminiscent of large electronic computers. He sees a small number of basic components repeated over and over again. All the complexity lies in the innumerable interconnections, not in the variety of basic components. So far as we know, the brain consists exclusively of neurons. Again, so far as we know, a neuron does nothing but conduct electrochemical pulses from its head end to its tail end. (138)

Orit Halpern, a historian of science, further argues that artificial neural networks like Rosenblatt's

> both captured the idea that cognition and intelligence derive from the rule-based interactions of neurons and suggested that cognition could be modelled by means of a set of algorithms. These rules execute themselves, however, through the small actions of individual decision-making units that are networked together. This assemblage can produce more sophisticated decisions, much like the Hayekian market can coordinate many small decisions to self-organize and act in more complex ways. Intelligence here is reformulated as networked and capable of evolution through population level coordination of data. (350)

There were and continue to be divergent readings of the brain-computer analogy from both sides, which is to say accounts of the brain as a computer and the computer as a brain. Despite the use of these metaphors, animal research was undeniably important as an influence on the development of early neural networks. And while debates about the architectures used for these networks were informed by a range of research goals—from reaching engineering-focused requirements to achieving better performance on specific tasks, to help psychologists and physiologists study brain function and human perception—the strategies for assessing performance and evaluating various architectures, no matter those goals, were, like today, linked to comparisons with empirical research on humans and animals.

Building Animal Models

In *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms* (1962), Frank Rosenblatt's book-length treatment of his neural network research program, he notes what he sees as

> a failure to comprehend the difference in motivation between the perceptron program and the various engineering projects concerned with automatic pattern recognition, "artificial intelligence", and advanced computers. For this writer, the perceptron program is not primarily concerned with the invention of devices for "artificial intelligence", but rather with investigating the physical structures and neurodynamic principles which underlie "natural intelligence". A perceptron is first and foremost a brain model, not an invention for pattern recognition. As a brain model, its utility is in enabling us to determine the physical conditions for the emergence of various psychological properties. It is by no means a "complete" model, and we are fully aware of the simplifications which have been made from biological systems; but it is, at least, an analyzable model. (v–vi)

While this presentation of his Perceptron research did match a number of his statements made elsewhere and appears supported by his interactions with other researchers, it was primarily defensive. The phrase "artificial intelligence" had been coined by John McCarthy a handful of years prior for what would become known as the Dartmouth Summer Research Project on Artificial Intelligence, a six-week workshop held in the summer of 1956. Rosenblatt was keen to distance his work from a number of the early struggling projects that were identified with artificial intelligence. He also wanted to position his efforts within and against certain aspects of the developing fields of pattern recognition and computer vision. Later in *Principles of Neurodynamics*, Rosenblatt doubles down on his description of his research program in machine learning as dedicated to the basic scientific goal of simulating, in a mostly knowable and interpretable platform, biological systems for understanding complex psychological and physiological processes:

> Perceptrons are not intended to serve as detailed copies of any actual nervous system. They are simplified networks, designed to permit the study of lawful relationships between the organization of a nerve net, the organization of its environment, and the "psychological" performances of which the network is capable. Perceptrons might actually correspond to parts of more extended networks in biological systems; in this case, the results obtained will be directly applicable. More likely, they represent extreme simplifications of the central nervous system, in which some properties are exaggerated, others suppressed. In this case, successive perturbations and refinements of the system may yield a closer approximation. (26)

While he makes considerable effort in *Principles of Neurodynamics* to establish the Perceptron and neural networks as basic rather than applied research, his program was initially funded to create physical machines—essentially alternatives to the cumbersome general-purpose digital computers of the 1950s and 1960s—that would be usable by U.S. Department of Defense agencies for a variety of military applications. Nonetheless, the architectures developed by Rosenblatt were grounded in theories and experimental research concerning animal and human perception, and the resulting networks, as simplified simulations of perceptual systems, had thus been mapped through research on animal models.

Before the Perceptron was a physical machine, though, it was a simulation of perceptual networks in animals, and its ultimate fate (and the fate of machine learning in the present) would be to remain such a simulation. Rosenblatt himself preferred this broader context, noting that the term "perceptron" was "originally intended as a generic name for a variety of theoretical nerve nets," further explaining that it had the "unfortunate tendency to suggest a specific piece of hardware, and it is only with difficulty that its well-meaning popularizers can be persuaded to suppress their natural urge to capitalize the initial P" (*Neurodynamics*, x). Nonetheless, "Perceptron" will mostly be capitalized in this essay in order to gesture toward the importance of the materiality of the series of systems, both mechanical and simulated in software, developed by Rosenblatt.

While Warren S. McCulloch and Walter Pitts's description of a simplified neural network was one important source available to Rosenblatt, it was not the only or even the primary theoretical account that he applied in designing his early simulations of neural networks. Rosenblatt cited McCulloch and Pitts, alongside Nicolas Rashevsky, Pitts's advisor at the University of Chicago, but it was the work of Canadian Gestalt psychologist Donald O. Hebb that would cast a much longer shadow over Rosenblatt's research. Rosenblatt would criticize ANNs based on the standard McCulloch–Pitts archetype because they were "improbable as biological models . . . [and] highly inefficient in the size of the network which is required" ("Comparison of Several Perceptron Models" 468). Hebb's *The Organization of Behavior* (1949) was principally a theoretical work, although he did have findings from his research with animal models, including testing rodents' intelligence and memory with carefully constructed mazes (now known as a Hebb–Williams maze). Grounding his Gestalt theories in empirical research, Hebb writes that "animal experiments and the human clinical data alike indicate that the perception of simple diagrams as distinctive wholes is not immediately given but slowly acquired through learning" (35). These were all behavioral and observational experiments, however, and not based on findings from more invasive neurological research.

In 1957, when Rosenblatt, then a researcher at Cornell Aeronautical Laboratory, first proposed building a machine, the Mark I Perceptron, to the Office of Naval Research as part of Project PARA (Perceiving and Recognizing Automaton), there were theories about the operation of neural networks but no major experimental studies making use of single-cell recordings. In implementing a neural network based on Hebb's cell assemblies—and also referencing the work of McCulloch and Pitts as providing proof of the ability of these machines to implement Boolean logic—Nathaniel Rochester and his colleagues at IBM cast their research as a mechanical simulation in the absence of an ability to test Hebb's theories in an animal model: "In the experimental study of the brain it is not yet possible to observe well the electrical interconnections among neurons. No one has yet been able to simultaneously record input and output signals of a single neuron in the brain. For this reason it has not yet been possible to test certain theories about how the brain works by experimentation on animals" (Rochester et al. 80). These *in silico* experiments were simulations of theories about perceptual functions and neural responses conducted by engineers and early computer scientists on the boundaries of the brain-computer metaphor. Rosenblatt thus shared with the IBM group an interest in testing theories about perception using mechanical simulations. In the absence of single-cell recordings of neuronal activity, artificial neural networks would be the mechanism by which researchers could evaluate theoretical work.

As Rosenblatt developed his neural networks, adding complexity and reconfiguring the architectures, he would increasingly draw on findings from new research. Two major models of perception derived from animals—frogs first and then cats—informed the designs and strategies of Rosenblatt and other machine learning researchers. The research on frog vision bridged the theory and practice of understanding perceptual systems through neural networks; McCulloch and Pitts worked with Jerome Y. Lettvin and Humberto R. Maturana to publish the groundbreaking 1959 article "What the Frog's Eye Tells the Frog's Brain." This paper reported on research that examined recordings from nerve fibers in order to map the receptive field of the frog retina and identify four major "detectors" that are specialized for different visual features (Dobson, *Birth* 71). Lettvin and his colleagues were not operating in isolation from emergent research on neural networks; they were motivated by and responding to experiments with machine learning conducted by artificial intelligence researcher Oliver G. Selfridge. Their findings served as a major intervention in neural network research, providing indications for adjusting existing paradigms and implementing a different sort of architecture for neural networks, one that was not based on simple presence-absence detectors of light from the lowest layer of the network that resembled and was thus named the retina by Rosenblatt.

The image concept in machine learning, as input data and as increasingly higher-level abstract representations within neural networks, was deeply influenced by the Lettvin et al. frog-vision model, but it would be another, tangentially related, description of animal perception that confirmed Rosenblatt's prior hunches about improving the Perceptron, ultimately altering the trajectory of his neural network research. That transition would also bring Rosenblatt into closer contact with animal research, eventually leading him to conduct his own experiments on memory and intelligence with animals. David Hubel and Torsten Wiesel's 1962 article "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex" would become an extremely pivotal publication in the history of artificial intelligence, machine learning, and computer vision. These fields were, and remain, intimately connected, although the nature of that connection has shifted over time. In the early 1960s, AI researchers believed that computers were capable of thought but that thought was disconnected from the environment. Adding environmental awareness through the modeling of perception, especially visual perception, was considered key to making machines more intelligent. Oliver G. Selfridge and Ulric Neisser wrote of this imperative to solve the perceptual problem in the pages of *Scientific American*:

> Can a machine think? The answer to this old chestnut is certainly yes: Computers have been made to play chess and checkers, to prove theorems, to solve intricate problems of strategy. Yet the intelligence implied by such activities has an elusive, unnatural quality. It is not based on any orderly development of cognitive skills. In particular, the machines are not well equipped to select from their environment things, or the relations, they are going to think about. (60)

Hubel and Wiesel may not have solved the thinking problem, but their findings changed the way researchers designed neural networks and represented image data, eventually leading to the development of deep learning networks such as the convolutional neural networks used for image recognition and the vision transformers implemented in contemporary large multi-modal models.

Similar to Lettvin's methods of recording single fibers in the optic nerves of frogs, Hubel and Wiesel attached probes to the neurons of cats in order to observe nerve activity while the cats were shown visual stimuli. In the methods section of their paper, Hubel and Wiesel described their experiment in detail, noting how they prepared, positioned, and displayed stimuli while recording activity in visual cortex neurons:

> Recordings were made from forty acutely prepared cats, anaesthetized with thiopental sodium, and maintained in light sleep with additional doses by observing the electrocorticogram. Animals were paralysed with succinylcholine to stabilize the eyes. Pupils were dilated with atropine. . . . The animal faced a wide tangent screen at a distance of 1–5 m, and various patterns of white light were shone on the screen by a tungsten-filament projector. All recordings were made in the light-adapted state. . . . For each cell receptive fields were mapped out separately for the two eyes on sheets of paper, and these were kept as permanent records. (107)

Hubel and Wiesel's study provided clear evidence that the prior perceptual model that had structured the earliest ANNs was quite limited and not aligned with how visual perception worked in mammals. This discovery motivated what would eventually be an almost field-wide transformation from the use of pixel-wise data, in which the initial processing elements were mapped to specific and static locations in image space, to higher-level visual features derived from the sort that Hubel and Wiesel identified as key to feline vision. They would continue their invasive experiments with cats as the title of a 1963 article makes clear, "Single-Cell Responses in Striate Cortex of Kittens Deprived of Vision in One Eye," and would eventually apply a similar experimental apparatus to monkeys. In modifying components of his networks to include these complex visual features, Rosenblatt returned to his initial experiments with multi-layered neural networks and introduced additional complexity to the Perceptron architecture, in the form of specialized layers.

The early Perceptron, the architecture outlined in his 1957 proposal to the Office of Naval Research, and the mechanical implementation found in Cornell Aeronautical Laboratory's Mark I Perceptron were based on a relatively straightforward three-layer network, which Rosenblatt called a "*simple perceptron*" (". . . Several Perceptron Models" 465). The lowest and first layer, as noted, was defined as the retina, or, in the more technical descriptions and documentation, the sensory, or "S," system. This layer established a matrix of input pixel values, frequently a square of twenty-by-twenty values. Rosenblatt labeled the second layer (the middle layer) the association, or "A," system. This variously had either static multiple or random connections to the layers on either side. In the contemporary nomenclature, these "A" layers would be termed hidden layers. The final layer was the response, or "R," system. In the early models, this system's function was to sum the inputs from the "A" system and output a response, either a 1 (positive) or -1 (negative) as the predicted class, which is to say binary category, of the input data supplied to the "S" layer. However, Hubel and Wiesel's findings from their single-cell recordings inspired Rosenblatt to once again modify the architecture of his Perceptron. He noted that while a four-layer network had previously been developed, the erroneous assessment that "the constraints assumed were unlikely to be found in biological systems" had meant that "this model was relatively neglected until recently. The recent discoveries of Hubel and Wiesel on the visual system of the cat . . . completely alter this evaluation, however, and the model can now be shown to be one of particular biological interest" (". . . Several Perceptron Models" 468–69). Rosenblatt's changes, however, included more than just the addition of a second A layer, as he also modified his neural network architecture to implement higher-level feature detection in both A layers.

Like the three-layer precursor, these newer multi-layer Perceptron networks (four or more layers) were initially simulated in software, as a program written for and executed on a general-purpose computer, an IBM 7090 housed at New York University. In a paper that he delivered at the 1962 Self-Organizing Systems conference in Chicago, Illinois, Rosenblatt explicitly links the design of these networks to the account of the cat's visual cortex offered by Hubel and Wiesel, describing one layer of the A units as "organized as line and edge detectors with a distribution of origin configurations roughly approximating the distribution found by Hubel and Wiesel" (". . . Several Perceptron Models" 477). In reporting on experiments comparing these networks to his *simple perceptron*, Rosenblatt refers to the newer designs as a "cat model." In this paper, he alternates between quoting and not quoting "cat" and sometimes drops the reference to "model" completely: "Here the cat does only slightly better than the simple perceptron" (477). While recognizing that this is only a highly simplified simulation of inferences drawn from a single study, Rosenblatt does begin to think in terms that we could consider zoomorphic; he proposes to evaluate the model in relation to feline behavioral studies: "It would be interesting to determine whether this discrimination is equally difficult for flesh-and-blood cats" (477). This idea prompted him to establish a fresh line of research that he would continue to find fascinating and would later pursue from a slightly different angle.

Rosenblatt's early simulation experiments of the four and five-layer cat models gave him and his colleagues the confidence to implement more advanced versions in hardware. Developing a mechanical system held relevance for Rosenblatt and the funding agencies supporting his research. It was easier and less expensive to simulate a neural network as a program on a general-purpose computer, but these simulations were slower than a dedicated machine and required sharing access to the systems with others. While this approach made sense at the time, the rapid increase in the speed of digital computers would, across the longer history of scientific computing, later leapfrog specialized computing hardware. Building input devices directly into the systems, in the form of the S layer, was also understood to enable more rapid digitization of a variety of stimuli. The Office of Naval Research funded the construction of the first mechanical learning machine designed by Rosenblatt, the Mark I Perceptron. That system was constructed at the Cornell Aeronautical Laboratory (CAL), in Buffalo, New York, and remained there. It was used by CAL researchers for military reconnaissance photointerpretation, among other tasks (Murray 1961). By 1959, Rosenblatt had left CAL and moved back to the Ithaca campus of Cornell University (where he had earned his undergraduate and graduate degrees) to start the Cognitive Systems Research Program, or CSRP. His new system would thus be housed at Cornell University, and while still funded by the Office of Naval Research, along with support from the National Science Foundation, the research conducted with this machine would be much more varied than the projects executed with the Mark I Perceptron.

This advanced system featured a novel memory storage device, designed by George Nagy, a graduate student working with Rosenblatt (Nagy, "A Survey"). The architecture of this system would become the subject of

Nagy's 1963 dissertation. While still implementing the cat model, this new machine was dedicated to sound perception—it was intended to be what Rosenblatt termed a phonoperceptron, in contrast to the photoperceptron that was the earlier CAL Mark I machine. Instead of a retina, the S layer for this system accepted, as inputs, digitized audio signals from a microphone in a connected recording booth. Rosenblatt had more freedom in choosing the name for this second mechanical implementation of his Perceptron. As it was based on a simulation of a cat's visual cortex and designed primarily to process audio input, Rosenblatt selected the name "Tobermory," which he extracted from his reading and personal life. More imaginative than the sedate and militarized Mark I (or the obvious second in the sequence, Mark II), "Tobermory" was fitting for the machine's intended purpose, as it was inspired by a short story concerning a talking cat, written by British writer Hector Hugo Munro (better known by his pen name, "Saki") and published in 1911. Rosenblatt was especially captivated by fictional stories involving animals. He owned a large farmhouse in Brooktondale, New York, a town neighboring Ithaca, which he shared with an ever-changing commune-like gathering of students and others. He presided as something of a father figure over this collection of housemates and, as appropriate to that role, would read aloud from his favorite books. Rod Miller, a Cornell undergraduate, close friend, and scientific collaborator with Rosenblatt, referenced these reading sessions in his comments at Rosenblatt's memorial, shortly after Rosenblatt's untimely death:

> I lived at his house and he used to read to us after dinner. We read *Canticle for Liebowitz* [sic], *Alice in Wonderland*, *Through the Looking Glass*, the Ring Trilogy, *The Once and Future King*, and many others. *The Wind in the Willows*. One chapter Frank came to and he said, "This chapter's about me." And it's about Mr. Toad, who had wanted to get a bright shiny red motor car, and he was obsessed with this idea, and Frank identified with Mr. Toad, and he had many shiny red motor cars: projects. (*Congressional Record* H27717)

Saki's "Tobermory" was most likely another favorite as it served the source of a name for both the Tobermory machine and for one of Rosenblatt's own cats.

Saki's story has interpretive possibilities for Rosenblatt's conceptualization of the machine, its biological inspiration, and its possible uses. "Tobermory" provides a satire of late-nineteenth and early-twentieth-century social norms and a send-up of the pseudoscience of that period. The story takes place during a party at the aristocratic home of Lady Blemley in the southwest of England, on the Bristol Channel. Joining the guests for an extended stay at the house is an outsider, a mysterious man named Cornelius Appin. Shortly after his arrival, Appin announces a recent scientific success. The narrator recalls, "now he was claiming to have launched on the world a discovery beside which the invention of gunpowder, of the printing-press, and of steam locomotion were inconsiderable trifles. Science had made bewildering strides in many directions during recent decades, but this thing seemed to belong to the domain of miracle rather than to scientific achievement" (Saki 27). Appin, we learn, has spent almost two decades experimenting with different animals and has just had his first breakthrough: he has been able to teach a cat to speak. Cats, we learn, are special animals, "wonderful creatures which have assimilated themselves so marvelously with our civilization while retaining all their highly developed feral instincts" (27). Upon hearing this unbelievable news, the guests go off in search of the Tobermory, a "Beyond-cat" that Appin has discovered at Lady Blemley's estate.

Once Tobermory has been located, questions are put to the cat to test its intelligence. Tobermory's responses to these questions evidence not only his listening capacities and memory but also his wit:

> "What do you think of human intelligence?" asked Mavis Pellington lamely.
>
> "Of whose intelligence in particular?" asked Tobermory coldly.
>
> "Oh, well, mine for instance," said Mavis, with a feeble laugh.
>
> "You put me in an embarrassing position," said Tobermory, whose tone and attitude certainly did not suggest a shred of embarrassment. "When your inclusion in this house-party was suggested Sir Wilfrid protested that you were the most brainless woman of his acquaintance, and that there was a wide distinction between hospitality and the care of the feeble-minded. Lady Blemley replied that your lack of brain-power was the precise

> quality which had earned you your invitation, as you were the only person she could think of who might be idiotic enough to buy their old car. You know, the one they call 'The Envy of Sisyphus,' because it goes quite nicely up-hill if you push it." (29)

These questions also become the basis for the guests' fear of Tobermory, not only of what he knows now but what knowledge he will soon have with his ability to listen, unobserved, to conversations among them. The guests worry about what to do with Tobermory, now that he has the power of speech. Lady Blemley, the cat's owner, says decisively, "My husband and I are very fond of Tobermory—at least, we were before this horrible accomplishment was infused into him; but now, of course, the only thing is to have him destroyed as soon as possible" (32). But the cat, in the meantime, has gone missing. With their anxiety raised, Clovis, one of guests, says: "He won't turn up tonight. He's probably in the local newspaper office at the present moment, dictating the first instalment of his reminiscences" (34). Fortunately for the guests, Tobermory is shortly found dead, the result of a run-in with another cat, "big Tom from the Rectory" (35).

One can see the humor and fascination that this pithy short story had for Rosenblatt and his colleagues, especially as they constructed a machine that they hoped would be able to learn to recognize spoken words. Machine learning has a standard pedagogical model. It positions the researcher as a teacher and the network as a student. What is known as reinforcement learning, in certain of these experiments, implemented punishment and reward procedures to train the networks. That these machines were typically created in university or university-owned laboratories only fortified an instructional relation that was already firmly in place. "Teaching" a Perceptron to "read" an image—especially in the case of stimuli that were single alphabetic letters, a favorite object of early researchers—or to "listen" to sound, placed the researchers in the role of Cornelius Appin, the mysterious scientist and teacher of Saki's story. When asked about his method, how he has managed to do the impossible and teach a cat to listen and speak, Appin responds, continuing to offend the other guests: "one teaches little children and savages and backward adults in that piecemeal fashion; when one has once solved the problem of making a beginning with an animal of highly developed intelligence one has no need for those halting methods. Tobermory can speak our language with perfect correctness" (27). One might speculate that Rosenblatt understood the "simple perceptron" as learning through halting methods and that he approached the Tobermory system, and other networks using the cat model inspired by Hubel and Wiesel's research, with the hopes that these would be more like the "Beyond-cat" of Saki's story.

As metaphor or model, accounts of perception drawn from single-cell recordings of cats were deeply important to Rosenblatt's multilayer Perceptron networks and his second-generation mechanical system. But a key question remains: while Saki's "Tobermory" was an enjoyable story and a humorous naming opportunity for Rosenblatt and Nagy, to what degree did the narrative content of this science-fiction story resonate with their thinking about the Tobermory machine and its possible applications? That the Tobermory had an important role to play in ongoing mid-century U.S. surveillance projects seems clear; it was designed to decode digitized audio recordings of speech, much like the original hardware implementation of the Perceptron in CAL's Mark I was designed to decode digitized images and text. This project, like all CSRP research, was executed under the backing and oversight of the Office of Naval Research, Rosenblatt's chief funding vehicle. It is clear that the surveillance and eavesdropping element found in the short story parallels the possibility of the device Tobermory being used to process and recognize surreptitiously recorded audio data.

Despite Rosenblatt's insistence that the Perceptron was a brain model and not an engineering project that would lead to artificial intelligence or improved pattern-matching machines, much of his research on neural networks was in fact directed toward applied problems, and these motivated not only the trajectory of his research but also the overdetermining logic implemented in the virtual and mechanical assemblages of his neural networks. Rosenblatt had specified some of these possible applications in his initial 1957 proposal to the Office of Naval Research, to fund CAL's Project PARA and build the Mark I Perceptron: "Devices of this sort are expected ultimately to be capable of concept formation, language translation, collation of military intelligence, and the solution of problems through inductive logic" ("The Perceptron" summary). While he slowly shifted away from work directed to military applications—in part connected to his growing opposition to the war in Vietnam and his move from CAL back to Cornell University—Rosenblatt remained ready with possible applications for his neural network research. As late as 1963, when presenting on recent progress with the Tobermory project, at Northwestern University, in Evanston, Illinois, Rosenblatt continued to provide predictions of near-future applications of neural-network based learning machines. Ronald Kotulak, a reporter for the *Chicago Tribune*, writes of Rosenblatt's imagined uses for the Tobermory technology: "Within the next five years it will be possible to build a machine that will perform all the functions of an office secretary and receptionist, he predicted" (13). Rosenblatt described other possible applications, including everyday use as something like a contemporary chatbot, "'It will not be long before these machines can be made to hold a conversation,' Rosenblatt said. 'For example, if you mention the name Polonius, the machine will launch into a discussion of the finer points of Hamlet.'" And he continues, "'Altho [sic] I am skeptical that a machine can ever completely duplicate a man's behavior, I'm sure that they can be built so that it would be almost impossible to distinguish many of their actions from those of a man,' Rosenblatt said" (13).

Teaching Machines, Learning from Children

Rosenblatt's cat-derived neural network architectures were tested and evaluated through comparative studies of early childhood learning—specifically the process of learning to recognize characters and read. While there was already comparative work between children and animals in studies of visual perception, the introduction of neural networks presented novel opportunities for modeling and comparing learning processes and evaluation criteria. These came to Rosenblatt through his connections with Eleanor Gibson (Dobson, "Confusion Matrix"). Due to labor conditions and gender bias in the academic realm—for many years she was prevented from serving in the same department as her husband, due to prevailing university nepotism rules of the time (Rodkey)—Gibson took up animal studies of visual perception. Initially unable to carry out research on humans, she went on to conduct her studies on animals, off-campus, alongside Cornell professor Howard S. Liddell, at his Behavior Farm Laboratory. When Gibson later turned back to human-subject research and was finally appointed as a faculty member in Cornell's psychology department, her earlier animal research influenced her work on children and motivated her comparative studies between animals and humans. This research became crucial to the development of machine learning and artificial intelligence. These machines, as we have already seen, were based not only on animal models but also conceptualized as possessing cognitive capabilities akin to those of animals and young children. In the case of Rosenblatt's research, the addition of neural networks to the animal-human comparative framework was rooted in the geographies and histories of this branch of machine learning. Rosenblatt's first academic publication, as a graduate student, was co-authored with Eleanor's husband, James J. Gibson, and all three held appointments or affiliations with Cornell University's psychology department. Gibson's ecological theory of vision, developed in that publication, informed much of Rosenblatt's subsequent work (Irwin).

In an article co-authored by Rosenblatt and Carl Kesler, also of Cornell University, there is detailed description of how they came to use experimental data collected from early childhood studies of literacy education:

> The last experiment is only partially complete at this time, but the data are of sufficient interest to warrant presentation of some preliminary results. In a recent series of perceptual experiments with pre-literate nursery school children, E.J. Gibson has obtained a confusion matrix for discrimination of alphabetic characters. . . . It was decided that a comparison of her confusion matrix with one obtained from our cat model, using the identical alphabetic characters (in digitalized form) would be worth while. So far, only three discriminations have been tested. . . . The relative performance in these three cases agrees perfectly with the results shown in Gibson's confusion matrix. It is particularly striking that the discrimination of M and W, which look virtually identical to the "cat" model, was also the most difficult for Gibson's children. While it is obviously premature to place much weight on these results, they suggest the possibility that the analyzing mechanisms present in the cat model may be quite close, in performance, to those which are operating in human infants, before more sophisticated perceptual tests have

> been learned. It will be particularly interesting to see whether discriminations of curved letters continue to show the close correlation with human performance which is being found for the straight-line letters. (96)

In appropriating Gibson's confusion matrix, a concise visualization taking the form of a square chart of values showing the number of times one letter was misrecognized as another, Rosenblatt found both data and an instrument by which he could compare the performance of a neural network with human subjects. This matrix, along with the standardized rendering of alphabetic characters used by Gibson, enabled Rosenblatt to test different parameters and configurations via his neural networks. The cognitivist paradigm girding the studies undertaken by both Gibson and Rosenblatt enabled the comparison of these alphabetic figures and their representations, whether mental, in the case of the children, or numerical, in the case of Rosenblatt's artificial neural networks. In her critique of cognitivism in contemporary discussions of deep learning networks, anthropology expert Lucy Suchman argues that "the enduring commitment that informs the project of computational neural networks—and its embrace by computational neuroscience—is to cognitivism. The sense of cognitivism in this context is a theory of intelligence based in a correspondence between mental representations formed in the brain/mind and a world taken to stand outside of it" (Suchman 88). Cognitivist theory, as applied to machine learning, takes the particularly striking correspondence between representations of confusion, which is to say errors of classification, as evidence of similarity (at least in terms of effectiveness) of the human and machine "analyzing mechanisms."

In another article, a short conference paper appearing in a collection also published in 1963, Rosenblatt writes of his results from this series of experiments, now focused solely on the five-layer Perceptron network:

> a series of studies is now being initiated to determine the performance of cats in discriminating the letter pairs used for the perceptrons and children. The three-way comparison of the cats, children, and brain models will help in deciding whether residual discrepancies in performance are due to differences between

> perceptrons and biological brains in general, or between cats and children in particular. In addition, it is hoped that the cat data will provide a better pool of performance measures for the evaluation of future models of brain mechanisms, which must be based largely on physiological data derived from the cat as an experimental animal. (". . . Five-Layer Perceptron" 145)

Implicit in this three-way comparison was the understanding that these neural networks could be imagined as having similar cognitive capabilities as children. Based on Hubel and Wiesel's mapping of the cat visual cortex, these five-layer networks with simple feature detectors were also understood as partial simulations of cat brains and thus needed to be evaluated alongside the performance and capabilities of real cats. While the source of data from children and the procedures to obtain data from recognized and misrecognized characters was clear from his continued citation of Gibson's experiments, the methods that Rosenblatt intended to use to evaluate the literacy of "flesh-and-blood" cats were less apparent.

Convolving Perspectives

In the later years of his unfortunately short academic career, Rosenblatt made the transition into experimenting with real animals. Although different in methodology and scope, this research can be thought of as extending the modeling work done with Perceptron-based ANNs. In his preface to the program's technical papers, Rosenblatt would describe the larger project of the CSRP as "primarily concerned with the study of models of central nervous system functioning, and the testing and verification of such models by means of biological experiments" (iii). While he had initially focused on mechanical and computer simulation of neural networks, his later research turned to trials conducted with animals themselves. However, instead of following his earlier concept of utilizing cats, Rosenblatt launched these newer tests on rats. His initial effort was a collaborative study, conducted in 1963 (with findings published in 1964's "Serotonin Binding to Nerve-Ending Particles of the Rat Brain and its Inhibition by Lysergic Acid Diethylamide"), which involved giving rats LSD-25 in order to assess their neuroanatomy.

Following this collaborative experiment, Rosenblatt then launched his own investigations with rats. These studies involved experiments to "transfer" acquired knowledge and learning in animals through the rather crude injection of extracted brain matter. In the article "Behavioral Assay Procedures for Transfer of Learned Behavior by Brain Extracts," co-authored with Rod Miller, the pair describe this process, in which brain material from a rat that had completed a maze was then inserted into another that had not yet seen the maze. They used several methods of analyzing their data in their attempt to find signs of knowledge that had successfully migrated from one rat to another. In a 1969 article addressing these experiments, titled "Behavior Induction by Brain Extracts: A Comparison of Two Procedures," Rosenblatt reports on employing the well-known behavioral apparatus known as the Skinner box, with similar brain-material injection methods. Neither study has held up as correct, and the methods utilized were highly unlikely to result in the transference of training or learning. In his remarks at Rosenblatt's memorial, Richard O'Brien, the head of Rosenblatt's division at Cornell University and co-author on the LSD study, described the Perceptron research and the animal research as two distinct "research reputations." Yet it is clear, especially from historical distance, that these two programs were part of the same line of inquiry that had initially motivated Rosenblatt's research on the Perceptron and ANNs. While not explicitly linked to the construction of ANNs, these experiments were invested in another sort of virtualization of memory and knowledge, in the attempted transmission of tangible learning from one rat to another.

While Rosenblatt's research was frequently idiosyncratic and his collaborations enabled by his time and place, he was not alone in deriving neural network architectures from empirical research performed on animals (Deutsch) and explaining and evaluating these models in terms of the animals they emulated. Like Rosenblatt, Kunihiko Fukushima was inspired by Hubel and Wiesel's account of the cat visual perception system. He first experimented with this design in 1969 ("Visual Feature Extraction") and returned to it 1980 with his invention of the neocognitron, an early convolutional neural network (CNN). The neocognitron implemented feature detectors in the form of S, or simple, cells and C, or complex, cells. While Fukushima notes that his neural network is a highly simplified model, he also finds its referent model incomplete:

> The author does not advocate that the neocognitron is a complete model for the mechanism of pattern recognition in the brain, but he proposes it as a working hypothesis for some neural mechanisms of visual pattern recognition . . . the hierarchy model of the visual nervous system proposed by Hubel and Wiesel is not considered to be entirely correct. It is a future problem to modify the structure of the neocognitron lest it should be contradictory to the structure of the visual system which is now being revealed. ("Neocognitron" 201)

The CNN was later modified by Yann LeCun, representing a crucial step toward the development of deep learning (LeCun et al.). These biologically inspired networks remain in use today, especially for many computer vision tasks (Dobson, "Objective Vision"), and they have even been used in computational neuroscience to help understand biological vision (Lindsay).

Today, those deploying the numerous evaluation metrics, benchmarks, and popular assessments of neural networks and GenAI-based applications share, with individuals from the earlier historical moment examined here, this desire to concretize the brain-computer metaphor through comparisons between ANNs and elusive aspects of animate intelligence. Given the genealogy of contemporary ANNs, the anthropomorphism of AI technologies in the present cannot be separated from a residual zoomorphism. Taking a page from Amazon's use of the phrase *artificial artificial intelligence* for Mechanical Turk, a human-task crowdsourcing service used to train numerous AI and ML applications, it might be appropriate to call our modern artificial neural networks *artificial animal intelligence*. Thinking in such terms acknowledges the role played by early empirical animal research, including single-cell recordings from frogs and cats, and subsequent theories about the meaning of these experiments, which were crucial to the development of single- and multilayer neural networks in the 1950s and 1960s, as well as later architectures, such as convolutional neural networks. Although ANNs do not directly simulate actual animal (or human) perceptual systems, their foundational paradigms were profoundly shaped by theoretical and empirical animal research. To a degree, especially in the popular discourse of machine learning, these origins are now a lost biological referent in the history of neural networks. Yet these biological referents remain vital to our understanding of these systems—both their historical scope and their future in the evolution of machine learning. In their creation and training and through their crossing of the animal-human-machine boundary, ANNs are more than the key, enabling technology of artificial intelligence; they are posthuman cultural techniques (Siegert). Recognizing the relevance of their intertwined histories, as well as the new ontologies that they generate and enact, is crucial for the critical analysis of computational culture in the present and for reconceptualizing the human in relation to artificial intelligence.

Works Cited

Deutsch, Sid. "Conjectures on Mammalian Neuron Networks for Visual Pattern Recognition." *IEEE Transactions on Systems Science and Cybernetics*, vol. 2, no. 2, Dec. 1966, pp. 81–85.

Dobson, James E. *The Birth of Computer Vision*. U of Minnesota P, 2023.

———. "Objective Vision: Confusing the Subject of Computer Vision." *Social Text*, vol. 41, no. 3, Sept. 2023, pp. 35–55.

———. "On the Confusion Matrix." *Configurations*, vol. 32, no. 4, 2024, pp. 331–50.

Elsasser, Walter M. *The Physical Foundation of Biology*. Pergamon P, 1958.

Fukushima, Kunihiko. "Visual Feature Extraction by a Multilayered Network of Analog Threshold Elements." *IEEE Transactions on Systems Science and Cybernetics*, vol. 5, no. 4, Oct. 1969, pp. 322–33.

———. "Neocognitron: A Self-Organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position." *Biological Cybernetics*, vol. 36, no. 4, Apr. 1980, pp. 193–202.

Geoghegan, Bernard Dionysius. *Code: From Information Theory to French Theory*. Duke UP, 2023.

Halpern, Orit. "The Future Will Not Be Calculated: Neural Nets, Neoliberalism, and Reactionary Politics." *Critical Inquiry*, vol. 48, no. 2, Winter 2022, pp. 334–59.

Hayles, N. Katherine. *Bacteria to AI: Human Futures with our Nonhuman Symbionts*. U of Chicago P, 2025.

Hebb, D. O. *The Organization of Behavior: A Neuropsychological Theory*. John Wiley and Sons, 1949.

Hubel, D. H., and T. N. Wiesel. "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex." *The Journal of Physiology*, vol. 160, no. 1, Jan. 1962, pp. 106–54.

Irwin, Julia A. "Artificial Worlds and Perceptronic Objects: The CIA's Mid-Century Automatic Target Recognition." *Grey Room*, no. 97, Fall 2024, pp. 6–35.

Kincaid, J. Peter, et al. *Derivation of New Readability Formulas (Automated Readability Index, Fog Count and Flesch Reading Ease Formula) for Navy Enlisted Personnel*. Naval Technical Training Command, Research Branch Report, Feb. 1975, pp. 8–75.

Kotulak, Ronald. "New Machine Will Type Out What It 'Hears': Can Fill Secretary's Job, Expert Says." *Chicago Tribune*, 18 June 1963, section 1, p. 13.

Kesler, C., and F. Rosenblatt. "Further Simulation Experiments on Series-Coupled Perceptrons." *Collected Technical Papers*, edited by Frank Rosenblatt, Cognitive Systems Research Program, vol. 2, report no. 5, Cornell U, 1963, pp. 73–98.

LeCun, Yann, et al. "Backpropagation Applied to Handwritten Zip Code Recognition." *Neural Computation*, vol. 1, no. 4, Dec. 1989, pp. 541–51.

Lindsay, Grace W. "Convolutional Neural Networks as a Model of the Visual System: Past, Present, and Future." *Journal of Cognitive Neuroscience*, vol. 33, no. 10, Sept. 2021, pp. 2017–31.

McCulloch, Warren S., and Walter Pitts. "A Logical Calculus of the Ideas Immanent in Nervous Activity." *Bulletin of Mathematical Biophysics*, vol. 5, Dec. 1943, pp. 115–33.

Murray, Albert E. "Perceptron Applications in Photo Interpretation." *Photogrammetric Engineering*, vol. 27, no. 4, Sept. 1961, pp. 627–37.

Nagy, George. "A Survey of Analog Memory Devices." *IEEE Transactions on Electronic Computers*, vol. EC-12, no. 4, Aug. 1963, pp. 388–93.

———. *System and Circuit Designs for the Tobermory Perceptron*. 1963. Cornell U, PhD dissertation.

Pasquinelli, Matteo. *The Eye of the Master: A Social History of Artificial Intelligence*. Verso Books, 2023.

Robles-Anderson, Erica, and Scott Ferguson. "The Visual Cliff: Eleanor Gibson and the Origins of Affordance." *Money on the Left*, 19 Apr. 2022.

Rochester, N., et al. "Tests on a Cell Assembly Theory of the Action of the Brain, Using a Large Digital Computer." *IEEE Transactions on Information Theory*, vol. 2, no. 3, Sept. 1956, pp. 80–93.

Rodkey, Elissa N. "The Woman Behind the Visual Cliff." *Monitor on Psychology*, vol. 42, no. 7, Jul./Aug. 2011, p. 30.

Rosenblatt, Frank. "Behavior Induction by Brain Extracts: A Comparison of Two Procedures." *Proceedings of the National Academy of Sciences*, vol. 64, no. 2, 1969, pp. 661–68.

———. "Comparison of a Five-Layer Perceptron with Human Visual Performance." *Natural Automata and Useful Simulations*, edited by H. H. Pattee et al., Spartan Books, 1966, pp. 139–48.

———. "A Comparison of Several Perceptron Models." *Self-Organizing Systems, 1962*, edited by Marshall C. Yovits et al., Spartan Books, 1962, pp. 463–84.

———. "The Perceptron: A Perceiving and Recognizing Automaton (Project PARA)." Report No. 85-460-1, Cornell Aeronautical Laboratory, Jan. 1957.

———. Preface. *Collected Technical Papers*, edited by Frank Rosenblatt, Cognitive Systems Research Program, vol. 2, report no. 4, Cornell U, 1963, pp. iii–iv.

———. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books, 1962.

Rosenblatt, Frank, and Rodman G. Miller. "Behavioral Assay Procedures for Transfer of Learned Behavior by Brain Extracts, II." *Proceedings of the National Academy of Sciences*, vol. 56, no. 6, 15 Dec. 1966, pp. 1683–88.

Saki [Hector Hugh Munro]. "Tobermory." *The Unrest-Cure and Other Stories*. New York Review Books, 2013.

Selfridge, Oliver G., and Ulric Neisser. "Pattern Recognition by Machine." *Scientific American*, vol. 203, no. 2, Aug. 1960, pp. 60–68.

Siegert, Bernhard. *Cultural Techniques: Grids, Filters, Doors, and Other Articulations of the Real*. Translated by Geoffrey Winthrop-Young, Fordham UP, 2015.

Suchman, Lucy. "The Neural Network at Its Limit." *Neural Networks*, edited by Ranjodh Singh Dhaliwal, Théo Lepage-Richer, and Lucy Suchman, U of Minnesota P, 2024.

Thrift, Nigel. "Electric Animals: New Models of Everyday Life?" *Cultural Studies*, vol. 18, no. 2–3, Jan. 2004, pp. 461–82.

United States, Congress, House. *Congressional Record*, vol. 117, part 21, 28 July 1971. p. 27717, https://www.congress.gov/92/crecb/1971/07/28/GPO-CRECB-1971-pt21-3-2.pdf.

Wiesel, Torsten N., and David H. Hubel. "Single-Cell Responses in Striate Cortex of Kittens Deprived of Vision in One Eye." *Journal of Neurophysiology*, vol. 26, no. 6, Nov. 1963, pp. 1003–17.