International Journal of Intelligence and CounterIntelligence, 36: 1337–1351, 2023

ISSN: 0885-0607 print/1521-0561 online
DOI: 10.1080/08850607.2022.2073542

PERSPECTIVE

JACK O'CONNOR

Undercover Algorithm: A Secret Chapter in the Early History of Artificial Intelligence and Satellite Imagery

Abstract: In the second decade of the twenty-first century, computer algorithms accurately and rapidly identify features of some objects on digital satellite imagery. These feature-recognition algorithms are expected to transform geospatial intelligence: to enable rapid retrospective searching of imagery archives and focus prospective analytic attention. This Perspective article establishes the beginning of the U.S. Intelligence Community's research and development on this capability. Archival research of declassified Central Intelligence Agency documents produced two discoveries: one identifies that the earliest feature-recognition initiative predated the creation of the National Photographic Interpretation Center (NPIC) in January 1961. The other discovery reveals that the earliest neural network software, Frank Rosenblatt's Mark I Perceptron, from which current feature-detection software descends, had been part of a previously secret four-year NPIC effort from 1963 through 1966 to develop this algorithm into a useful tool for photo-interpreters. The manager of that research effort, John Cain, defined the prospective utility of this software in 1963, and Cain's criteria, derived from NPIC's experience during the Cuban Missile Crisis, continues to shape the prospective geospatial intelligence uses of feature-recognition software.

Jack O'Connor directs and teaches in the M.S. in Geospatial Intelligence program at Johns Hopkins University. He served nearly all his 31-year government career in Geospatial Intelligence and Imagery Analysis. He is the author of NPIC: Seeing the Secrets, Growing the Leaders *and he is a recipient of the ODNI [Office of the Director of National Intelligence] Galileo Prize and the National Intelligence Medal of Achievement. The author can be reached at joconnor@jhu.edu.*

A FLAW IN A SUCCESSFUL ORGANIZATION

At the end of 1962, Arthur Lundahl reached his professional zenith as an intelligence officer. Over the previous six months, the organization he created and led for the previous eight years, beginning as the Central Intelligence Agency (CIA)'s Photographic Interpretation Division (PID), codenamed HTAUTOMAT for the U-2 program, and becoming the National Photographic Interpretation Center (NPIC), discovered Soviet strategic missiles in Cuba, tracked the daily events of the Cuban Missile Crisis, and monitored the first nuclear weapons withdrawal in history.[1]

At that time, along with its unique work on the Cuban Missile Crisis, NPIC had been exploiting film returned from the first space photo-intelligence missions of the CORONA satellite for more than two years. Due to its successes and the need to increase its workforce to exploit the increasing volume of photo-satellite film, NPIC moved over New Year's Day weekend in 1963. It left the decrepit Steuart Motor building and relocated to Building 213, a newly refurbished 400,000 square-foot building at 1st and M streets in the Washington Navy Yard Annex in Southeast Washington, DC. NPIC was in the early stages of its first hiring surge since 1956, when the U-2 program started. The nearly four hundred people at NPIC who moved from the Steuart building would soon grow to nearly a thousand.

Neither Lundahl nor NPIC had time to bask in the glory. A review of previously classified CIA documents indicates that Lundahl had been aware since 1960 that one key to NPIC's operational success was unsustainable. Since July 1956, Lundahl's organization had relied on overtime to interpret and analyze film rapidly after the receipt of U-2 and CORONA/KH-1–4 missions. During the Cuban Missile Crisis and the subsequent Soviet force withdrawal, NPIC continued to spend large amounts of overtime to exploit additional daily tactical low-altitude aircraft missions over Cuba. Yet Lundahl and NPIC's senior leadership knew that in 1963 the next generation of photo-satellites, the high-resolution[2] GAMBIT-1/KH-7, was scheduled to launch, and that each future CORONA/KH-4 mission was being modified to increase its resolution and film load. This awareness of future satellite developments with their additional film caused Lundahl to request technical assistance for the photo-interpreters.

From 1956 through 1962, PID, the Photographic Interpretation Center, and NPIC[3] photo-interpreters had worked prodigious amounts of overtime. To quote Lundahl, "By 1958 our numbers had grown from around 60 to about 225, and the overtime was going out at a mad rate. I think we spent a total of 200,000 hours of overtime in the St[e]uart Building alone."[4] From July 1956 through May 1960, the U.S. government had authorized 28 U-2 overflights of the Soviet Union. Every delivery of film from each of the 28 Soviet penetration missions commenced round-the-clock analytic efforts at the Steuart building. Between the U-2 missions, extended follow-on efforts continued to glean every bit of new information and intelligence. Less than two months after the first successful photo-satellite flight on 18 August 1960, CORONA/KH-1, mission 9009, Lundahl first documented the requirement for automated technical assistance for the photo-interpreters. His 5 October 1960 memo to the CIA's deputy director for intelligence[5] describes the new technical initiatives and states the initial requirement for computer-assisted feature-identification:

> The objectives of this program are as follows:
>
> a. To determine the feasibility of using optical electronic instrumentation to distinguish objects on photography above and below the normal range of human vision;
> b. To develop instrumentation which will demonstrate the feasibility of machines distinguishing between man-made and natural features on photography, and
> c. The ultimate development of equipment for locating specific shapes on photography as a means of target recognition.

To fund this requirement, Lundahl reallocated some of his 1961 operating budget to fund an initial target recognition initiative. On 18 August 1960, the first successful CORONA mission had brought back more area coverage of the USSR, although at a much lower resolution, than all 28 prior U-2 missions.[6] By October 1960, even before the second successful CORONA mission, Lundahl understood that reliance on overtime would lead eventually to future mission failure. As satellite, camera, and film technology improved and the frequency of satellite missions increased, continued reliance on overtime to look quickly and thoroughly at all satellite photography would become operationally, mathematically, and humanly unsustainable.[7]

While technical problems caused many early CORONA missions to have aborted, limited, and unsuccessful missions, by March 1963 CORONA had successfully flown eighteen missions, each returning to Earth with greater amounts of film with better resolution. By then, the KH-4 camera and the KH-5 ARGON mapping camera had been on orbit.

Throughout his leadership of HTAUTOMAT and NPIC, Lundahl's decades of knowledge of research and development in the photogrammetry[8] community enabled him to provide the best technical assistance for the CIA photo-interpreters.[9] In the early 1950s, he had been the managing editor of *Photogrammetric Engineering*, the journal of the American Society of Photogrammetry. By the early 1960s his responsibilities as the NPIC director caused Lundahl to rely on John Cain, the NPIC research and development director, to keep up with current developments in photogrammetry.

MARK I EYEBALL AND THE MARK I PERCEPTRON

After a dialog with Lundahl in 1963, John Cain started an initiative that would develop through photo-interpretation, lie dormant during the era of imagery analysis, and reemerge in geospatial intelligence. Cain initiated the first effort to develop machine learning in the U.S. photographic Intelligence Community (IC). Academic research on feature recognition and machine learning had begun in the same year that U-2 missions began to change the thinking and foreign assessments of the U.S. IC.[10] In 1957, Frank Rosenblatt, a recent alumnus and Ph.D. psychology researcher,[11] was on the Cornell Aeronautical Laboratory faculty. He had been conducting experiments to validate his theory that computers could "learn" to make visual distinctions, and through repetition improve their capability to make correct distinctions. Rosenblatt's 1958 paper, "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain,"[12] reached a number of audiences. The *New York Times*[13] and *New Yorker*[14] reported his discovery to the general reading public. In an aviation and space periodical, the technical press wrote about his research.[15] The Naval Research Laboratory in Washington, DC, which funded Rosenblatt's research, provided the most informed and most critical audience.[16]

The challenge of creating the hardware to prove or disprove Rosenblatt's theory loomed over the project. Under the Navy contract, to test and evaluate his conception and software, Rosenblatt created an early version of a digital sensor that he connected to an IBM computer. Rosenblatt called his learning machine the Mark I Perceptron.[17] He envisioned and constructed a 20×20 representation of an artificial retina that could perceive differences in the gray scale. He processed the output from the 400 primitive sensors through 512 potentiometers, devices that measured the electrical resistance from a sensory response against a constant.[18] This hardware configuration would help Rosenblatt understand if his model of biological learning could be verified through replication (Figure 1).

Rosenblatt's Perceptron experiment succeeded wildly and failed miserably. His concept of a single layer of digital neurons that could improve its ability to distinguish objects proved valid, successful, and capable of further development. However, his concept was not without mathematical shortcomings. And the Perceptron turned out to be a slow learner. Three technologies, that would not exist for 40 to 50 years, would be required for it to become practical. The first technology was computer processing speed. Rosenblatt tested his concept on an IBM 704 mainframe that could process about 4,000 operations a second. Even with the gains provided by Moore's Law in computer processing speed, the necessary computer processing speed to validate later versions of Rosenblatt's Perceptron model would not exist until the twenty-first century.

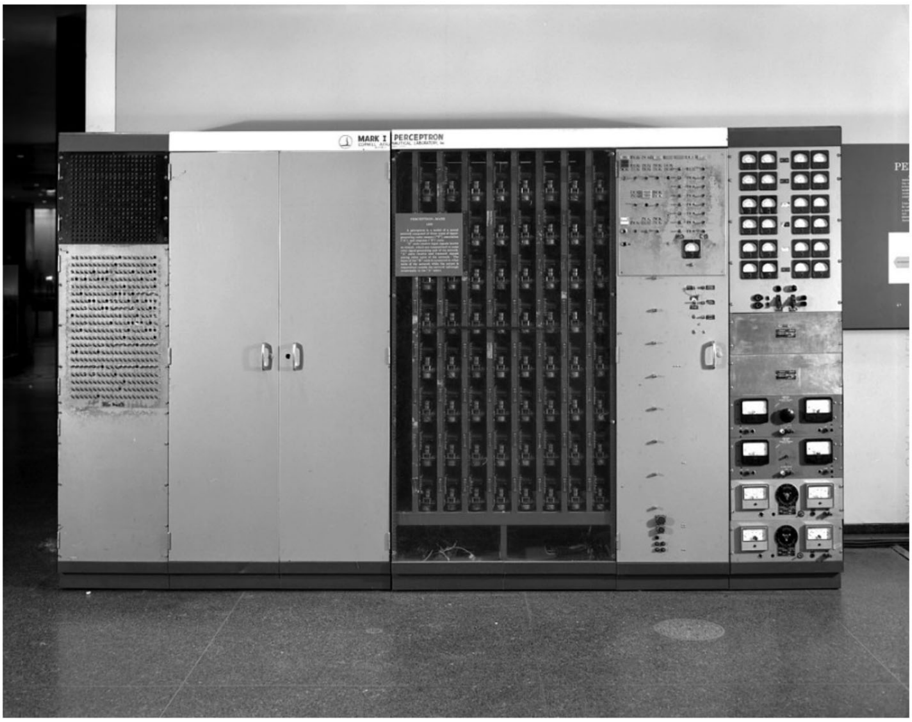

Figure 1. The Mark I Perceptron. Image used with permission of the Division of Medicine and Science, National Museum of American History, Smithsonian Institution.

The second shortcoming was the insufficient amount of visual data to "train" the algorithm. Rosenblatt worked with a handful of images. To improve the model's statistical validity, the necessary volume of spatial data would not be available until the first decade of the twenty-first century. Computer storage was the third shortcoming. This technology became available before the other two, but not until late in the twentieth century, many years after Frank Rosenblatt's untimely 1971 death at age 43 in a sailing accident on the Chesapeake Bay.[19]

PERCEPTRON DISAPPEARS

By mid-1961, the publicity and news about the Perceptron had ended except in technical and academic journals. But an article in *Photogrammetric Engineering* that had its genesis in a 1961 talk at the American Society for Photogrammetry may have helped the Perceptron disappear. Albert Murray, of the Cornell Aeronautical Laboratory, spoke about Perceptron Applications to Photo Interpretation.[20] Murray's abstract made the first comparison between a computer program and the human eye:

> Perceptron is a general name for a large family of perceiving and/or recognizing automata. Many of these, when given suitable sensory equipment, resembling an eye, and consisting of a lens and a retinal mosaic of light-sensitive elements, may be "taught" to reliably recognize or classify simple visual patterns.[21]

Murray's metaphor of the electronic retina likely caught the interest of Lundahl and Cain. Both men sought a way to manage the increased NPIC workload caused by the earliest photographic satellites as well as continuing U-2 missions over China and Cuba.

In a handwritten NPIC memo on 11 March 1963, alluding to the Navy withdrawing Perceptron funding, Lundahl asked Cain: "Is this device (Perceptron) something we want to support?[22]" Cain had already coordinated with the Navy Photographic Interpretation Center (the other NPIC) and scheduled on 19 March a briefing in Lundahl's NPIC to determine how much support it could provide to the Perceptron research and development. Cain and his staff knew about the Perceptron model[23] and its capabilities. On 23 April 1963, after a visit to the Rome Aeronautical Development Center at Griffiss Air Force Base, Cain noted that their automatic target recognition efforts had been developed from the Perceptron.[24] By June 1963, *The NPIC Technical Development Program*,[25] written by the Plans and Development staff under John Cain's signature, indicated the current state of progress:

> 5. Evaluation of Automatic Photographic Image Recognition Systems
>
> The increasing development of both the quantitative and qualitative aspects of photographic intelligence requirements and acquisition capabilities indicates rapidly increasing demands for improvements in exploitation technology. A particularly significant factor is the disproportionately large manpower requirement characteristics of the exploitation phase. It has become obvious that a large-scale, high-priority development program for automation of the exploitation phase is required.
>
> Developments in this realm may be divided into two basic categories.
>
> 1. Automation which assists the film analyst in film handling and viewing.
> 2. Automation which assists the analyst in scanning and evaluating the photo image.
>
> Although there are urgent requirements in both categories, it is apparent that a performance plateau has been reached. This plateau is basically defined by the available manpower and the limited efficiency of using highly-trained manpower for first-phase scanning of tremendous quantities of *negatively significant or highly redundant*

> *photographic images* [italics added]. It is thus implied that a definite limit in the capacity for first-phase[26] readout is being approached and that this limit will not be significantly changed by any developments in the first category.
>
> For this reason members of the Plans and Development staff have started a comprehensive, accelerated program for searching out and evaluating all the automatic imagery detection and recognition systems presently being proposed, developed, or produced.
>
> Initial results of this investigation indicate that important strides are being made in this field. One of the most promising areas of potential development for such automation is in the realm of "biological computers." This terminology is used to describe computer systems which generally consist of a sensory matrix, connected by a means of statistically significant coupling to the input of a digital computer, which is programmed in a fashion simulating brain mechanisms. Such systems are capable of being "taught" to recognize, with a high degree of reliability and discrimination a specific stimulus falling on the sensory matrix.
>
> The PERCEPTRON which was developed by the Cornell Aeronautical Laboratory in 1958 is the generally accepted forerunner of these systems. Since that time many related systems have been developed by other organizations, and continued improvements have been made in the Perceptron. The common denominator of limitation in all these systems appears to be the need for prenormalization, or standardization of scale and orientation of the image before it is presented to the sensory matrix. The complexity of the "biological computer" system precludes its utilization on the mass of redundancy required to identify a single image in all its possible variations of scale and attitude. However, the potential capability of these systems for fine discrimination of photo images appears to be accepted. Therefore, most of the current development programs related to these systems are in the realm of prenormalization of the image.

Cain's assessment of the performance plateau for photo-interpreters documented for the first time the operational shortfall in photo-interpretation and the first attempt to combine human and electronic vision to analyze photographic images. Cain wished to automate two very different analytic tasks: the exploitation of negatively significant images and the exploitation of highly redundant images. The first task—negatively significant images—addresses a time-consuming task that most analysts found easy, identifying that nothing had changed at a location since the last observation. Negative intelligence—a report that nothing had changed—can be essential.

The other analytic challenge for which Cain sought help in 1963 was how to accelerate the exploitation of highly redundant images. The age of highly redundant photographic images had not yet become routine, but from 23 October until mid-November 1962, NPIC experienced it. In response to the discovery of Soviet nuclear missiles in Cuba on U-2 imagery, the Defense Department commenced Operation BLUE MOON. Initially, the U.S. Navy and later the U.S. Air Force flew 168 low-level reconnaissance missions over known targets in Cuba in fewer than 25 days before the operation ended on 16 November 1962.[27] These missions focused on targets discovered on earlier U-2 collection in October. At least half of the missions, those flown by the Navy, were flown by pairs of reconnaissance aircraft. Exploiting the BLUE MOON missions would have been the NPIC interpreters' first experience with highly redundant first-phase collection.

While the BLUE MOON missions carried much less film than a U-2 mission,[28] the number of daily missions, averaging between six and seven, meant that the amount of redundant collection became an exploitation burden for the NPIC interpreters. As NPIC analyzed no other collection with a similar frequency of observation of individual targets before 1963, these missions likely provided the foundation of Cain's judgment in the June 1963 report. At that time, twenty successful photo-satellite missions had been flown (eighteen CORONA missions and two successful KH-5 ARGON[29] missions). The first KH-7 GAMBIT would not be launched for another month. But the incremental successes of the KH-4 CORONA missions and their number of images steadily increased. Cain's judgment about a large number of negatively significant or highly redundant photographic images had its basis in the ARGON and CORONA missions as well as the recent NPIC BLUE MOON experience from mid-October through mid-November 1962.[30] Cain's distinction remains essential.

In the imagery and geospatial world, negative intelligence means no reportable or accountable change since the last observation or the most recent baseline report about known facilities. At the time of Cain's report, the attention cost for an individual target was much lower than it is today, but the growing numbers of satellites and aircraft missions, even in 1963, started to make NPIC senior managers aware of analytic attention as an increasingly scarce resource. Cain's comment about highly redundant images is also significant.

At that time of comparatively scarce photographic or imagery collection, redundant collection meant that a target was a priority. But even priority targets have intervals when nothing happens. Imagery and geospatial analysts and photo-interpreters always pay attention to priority targets, yet some part of this attention is not rewarded. Cain's insight—that automated assistance might quickly identify where no change was detected— would become increasingly more important as the number of imaging satellites and the volume of photographs and images increased over the years. If computer software could be trusted to tell human analysts rapidly what to ignore, the software would help focus human attention on what changed or partially changed in the image and questions that humans could answer better than computers. By the twenty-first century, Cain's 1963 insight about the applicability of AI was used increasingly to inform decisions about where to focus analytic attention.

But, in 1965, Cain's Perceptron research initiative ran into some hard realities. Cain had engaged the Cornell Aeronautical Laboratory, and from an undated document, most likely written in 1963, it can be determined that Cornell had been assisting NPIC. The undated compilation includes documents written by members of the Cornell faculty, but not Frank Rosenblatt, who by this time had moved to a different area of study.[31]

Lundahl continued to fund research on recognition algorithms, and in January 1966, NPIC awarded another contract for an Automatic Target Recognition System.[32] Throughout NPIC's history, it would continue investing in research and development initiatives on computer-assisted target recognition. But the 1966 contract ended Perceptron research at NPIC.[33] The contract justification outlines other research efforts:

> It has been demonstrated by various electro/optical/digital techniques, several of which have been hardware implemented in crude prototype systems. Some of the more promising ATR developmental systems are; [*sic*] the Tactical Target Recognizer for the Army; the Natural Image Computer for GIMRADA [Geodesy, Intelligence, Mapping Research and Development Agency, U.S. Army]; the Automatic Target Recognition Device for RADC [Rome Air Development Center, U.S. Air Force]; and the Minos II by Stanford for Ft. Monmouth. All of the systems developed to date satisfy some of the required objectives, but all have serious deficiencies when related to the ultimate operation requirements.

The January 1966 contract, which extended through 1967,[34] was NPIC's first research effort at automatic target recognition that excluded the Perceptron. Consequently, the Mark I Perceptron was donated to the Smithsonian in 1967, where it still remains, although not on exhibit. The Smithsonian's records indicate that it was shipped from Cornell in that year under a government transfer administered by the Office of Naval Research.[35] Until now, no record of the Mark I Perceptron being evaluated by the U.S. IC has been published.

While Rosenblatt's Mark I Perceptron hardware went to the Smithsonian, his Perceptron software went into academic limbo at the same time. Two prominent MIT researchers, Seymour Papert and Marvin Minsky, identified a mathematical limitation of the single-layer Perceptron model and attacked

Rosenblatt's ideas. Their book caused academic research and government funding on perceptrons to go dormant for many years.[36]

PERCEPTRON REBORN

Until late in the twentieth and the early twenty-first century, the academic study, experimentation, and technology development that revived the Perceptron model did not occur. Part of the reason for this hiatus had been the mathematical validity of Papert's and Minsky's recognition of the limitations of single-layer perceptrons. Another reason was Rosenblatt's untimely death. But the most significant reason for the delay in developing Perceptron-type algorithms was that three required engineering, scientific, and mathematical developments had to be invented before Perceptron-based neural networks could demonstrate their utility at recognizing features in photographic images.

The first discovery was the computer engineering that led to advanced computational speed, greater and cheaper computer storage, and faster chips, particularly the graphics chips called graphical user interfaces.[37] The second discovery grew from Geoff Hinton's mathematical work. Hinton played a significant part in addressing and resolving the issues raised by Minsky and Papert about backpropagation and value assignments in the algorithms. Hinton continued to work with a succession of discoveries. The parallel-distributed-processing group included David Rumelhart, who, like Frank Rosenblatt, had a degree in psychology. Rumelhart's work provided a way to advance the Perceptron model from single-layer to multilayer neural-net algorithms. Hinton next worked with Terry Sejnowski, and their partnership led to Hinton's work on backpropagation as a way of improving the Perceptron model. The research driven by Hinton, Rumelhart, and Sejnowski in the late twentieth century brought academic attention to the neural-net algorithms. Still, the utility of their ideas did not emerge until the computer hardware improved, and the data sets grew. The next generation of AI researchers, such as Alex Krizhevsky, Ian Goodfellow, Ilya Sutskever, and Andrew Ng, continued their work.[38] Rumelhart's parallel distributed processing enabled multilayer perceptrons to function. They have improved their detection rates to the current state where hundreds of layers "look" at millions[39] of images through billions of permutations with generative adversarial networks.[40]

Michael Wooldridge's history of AI captures the intermittent development of this software.[41] His analysis of the stages of neural-net development from Rumelhart's development of Parallel Distributed Processing to neural-net algorithms such as AlexNet, and finally Deep Learning algorithms is helpful, as it characterizes three kinds of depth in Deep Learning. One measure of depth is increasing the number of layers of neurons that operate in the neural-net algorithms. Another measure of depth was the number of neurons that would interrelate in each layer of the model. While Rosenblatt's Mark I Perceptron had one layer of neurons, current algorithms have hundreds of layers and a network that incorporates different algorithms to improve the overall performance of the neural network.

By the twenty-first century, neural-net algorithms had the necessary components for development—very large data sets and very fast computers. These components led to the development of the 2010 Stanford Image-net contest where teams of developers competed in having their algorithms identify a dataset of 14 million images accurately.[42] Yet even modern feature-recognition algorithms continue to demonstrate limits that Cain foresaw.

Algorithms identify only the known. This makes them very effective at the two analytic challenges first described by John Cain—repeated exploitation at a target where change seldom occurs and exploitation of oversampled targets, or targets where the frequency of collection exceeds the rate of observed changes. In the past decade a third very useful application for these algorithms has been developed for a sensor type that was only in the early stages of research and development when Cain started to sponsor Perceptron research. Synthetic-aperture-radar (SAR) has been developed to a point where commercial imagery providers and geospatial analytic organizations have created feature-recognition algorithms that have automated the answering of specific questions about certain classes of targets, and have scaled the application so that it can address a global market issue every day. The energy sector and the global petroleum storage market rely on this algorithmic analysis to monitor global petroleum storage. This algorithm has trained itself to measure a geometrically simple technology—the floating-roof storage tank, a ubiquitous global feature of petroleum storage. Ursa has pioneered this analysis and provides daily updates on a global scale.[43]

Algorithms are useful with SAR for multiple reasons. SAR orbits are circular so the geometry is easier. SAR is an all-weather sensor so weather and climate do not affect its coverage frequency. And the increase in the number of commercial SAR providers enables more redundancy and greater accuracy. The number of algorithmic applications will continue to grow, but even the best of these algorithms has an attribute that may preclude analytic customers from trusting them entirely.

The depth of deep-learning algorithms is not always beneficial. Much like an ocean, the deeper the algorithm, the more opaque its processes, and the more difficult for humans to see. A deep-learning algorithm cannot explain its decisions, which makes the application of these algorithms problematic for some intelligence analyses. Consequently, for all the improvements since the Perceptron, algorithmic target recognition remains untrusted in many communities that have developed its capabilities. Their continued development will likely result in their increased use, as no other technologies under development are as likely to help geospatial intelligence analysts with their vastly increased burden of "negatively significant or highly redundant" images. So, the descendants of Rosenblatt's Mark I Perceptron will, for the foreseeable future, always be used in combination with a human "the Mark I eyeball."[44]

ACKNOWLEDGMENT

I would like to recognize the assistance of the following in the preparation of this article: Dr. Mark Stout, Ms. Ann Dee Reiter, Dr. Marty Smyth, and Dr. Mark C. Lowell.

REFERENCES

1 For his work at NPIC, Lundahl was awarded the National Career Civil Service Award from the National Civil Service League on 14 February1963 (*Civil Service Journal*, Vol 3, No. 4 [1963], p. 1); Jack O'Connor, *NPIC: Seeing the Secrets, Growing the Leaders* (Alexandria, VA: Acumensa Publishing, 2015), pp. 49–59; Dino A. Brugioni, *Eyeball to Eyeball*: *The Inside Story of the Cuban Missile Crisis* (New York: Random House, 1992), *passim*.

2 Resolution is the capability of an aerial camera to distinguish objects of various sizes. In digital photography, it is the amount of area covered by one pixel. In chemical photography it is the number of lines per millimeter that can be individually distinguished and measured against known targets by reticles in a magnifier.

3 The CIA Photographic Intelligence Division became the Photographic Interpretation Center (PIC) in August 1958, and PIC became the NPIC in January 1961, in the last week of the Eisenhower administration, when the president signed NSCID 8. O'Connor, *NPIC*, p. xxii.

4 Arthur C. Lundahl, "The NPIC and Its Work," Unpublished manuscript for a proposed *Studies in Intelligence* article, adapted from a lecture delivered at the DIA Intelligence School, p. 6. (Declassified in Part-Sanitized Copy Approved for Release 2013/09/26: CIA-RDP78B05167A000800060003-8.)

5 PIC/D-114/60, Memorandum for Deputy Director (Intelligence) Subject: Development of Automatad. 5 October 1960. SECRET, Declassified in Part: Sanitized Copy Approved for Release 2012/11/01: CIA RDP78B05702A000100070054-3.

6 Kenneth E. Greer, "CORONA," in *Studies in Intelligence*, Vol. 17 (Suppl.) (1973), pp. 1–37; Kevin C. Ruffner (ed.), *CORONA: America's First Satellite Program* (Washington, DC: Center for the Study of Intelligence, 1995).

7 Lundahl and his Navy colleague, Allen Kramer, won a medal during World War II for applying their mathematical knowledge. The two men, serving on the Aleutian Islands off Alaska, were studying a Navy mathematical model for antisubmarine searches. They found a mathematical error when the model was applied at high latitudes, and persuaded an air commander to modify his previously unsuccessful searching technique. When Lundahl's and Kramer's

corrections were applied to the search plan, the air patrols found and sank a Japanese submarine that had previously eluded them. O'Connor, *NPIC*, p. 17.

8 Photogrammetry is the science of measuring objects on the ground from photographic or digital images taken at altitude.

9 In outfitting HTAUTOMAT for the initial U-2 missions, Lundahl wrote the justifications for non-U.S.-manufactured equipment when he knew it to be better than domestically produced technology. So, the Stereo-comparator was a Nistri, made in Italy; the Stereoscopes were Wild, made in Switzerland; and the CIA's first computer, installed in the Steuart building, was a Swedish-built ALWAC-2. O'Connor, *NPIC*, p. 25.

10 Jack O'Connor, "The Genome of GEOINT: Jam Session." Paper delivered at the North American Society for Intelligence History Inaugural Conference, 20–21 October 2019 (in press).

11 "Frank Rosenblatt," *Wikipedia*, https://en.wikipedia.org/wiki/Frank_Rosenblatt (accessed 23 April 2021).

12 Frank Rosenblatt, "The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain," *Psychological Review*, Vol 65, No. 6 (1958), pp. 386–408.

13 "Electronic Brain Teaches Itself," *New York Times*, 13 July 1958, Section E, p. 9.

14 *The New Yorker*, 6 December 1958, p. 44.

15 *Aviation Week*, 24 April 1961.

16 Contract Nonr2381(00), cited in Frank Rosenblatt, "The Perceptron, A Probabilistic Model."

17 The name of Rosenblatt's software and hardware may have been influenced by U.S. Navy slang. "Contemporary Geodesy" (Proceedings of a Conference Held at the Harvard College Observatory—Smithsonian Astrophysical Observatory, Cambridge, Massachusetts, 1–2 December 1958) says: "Now the first type of optical tracking, the most elementary, is that using merely the naked eye—as I heard a Navy man say the other day, 'Mark I eyeball.'"

18 A potentiometer measures the amount of resistance against a constant. If the constant of the potentiometer is 10Ω, and the potentiometer is set at 2, then the resistance is 8.

19 Cade Metz, *Genius Makers: The Mavericks who Brought AI to Google, Facebook, and the World* (New York: Dutton, 2021), pp. 26–27. Annually, the Institute of Electrical and Electronics Engineers presents an award in Frank Rosenblatt's name for outstanding contributions to biologically and linguistically motivated computational paradigms and systems.

20 Albert E. Murray, "Perceptron Applications to Photo Interpretation." Presented at the Society's 27th Annual Meeting, The Shoreham Hotel, Washington, DC, 19–22 March 1961. Published in *Photogrammetric Engineering,* Vol. XXVII, No. 4 (September 1961), pp. 627–637.

21 *Ibid*., p. 627.

22 Memorandum to Director, NPIC, From: Assistant for Plans & Development, *Subject PERCEPTRON*, 12 March 1963. Approved for Release 2005/05/02: CIA-RTP78B04770A002300030008-7.

[23] The compilation document begins with Design of a Photo Interpretation Automaton (22 pages, 14 figures). The citation of Murray's article in *Photogrammetric Engineering* means that the article had been written after September 1961. The subsequent articles are titled: "Illustrations to Accompany Perceptron on Investigation of Perceptron Application to Photo-Interpretation (Project PICS) Washington, D. C. December 20, 1962" (15 pages); PAPER PERCEPTRON, 25 June 1962 (4 pages text, six pages of illustrations); INPUT/OUTPUT EQUIPMENT FOR RESEARCH APPLICATIONS, reprinted from *Proceedings of the NEC*, 1962, Vol 18, pp 509-517; TWO-DIMENSIONAL SPATIAL FILTERING AND COMPUTERS, W. O. Fryer and E. E. Richmond, Cornell Aeronautical Laboratory, Inc (6 pages); SYNTHESES OF AN OPTIMAL SET OF RADAR TRACK-WHILE-SCAN-SMOOTHING EQUATIONS, reprinted from IRE TRANSACTIONS ON AUTOMATIC CONTROL, Vol AC-7, Number 1, July 1962 (6 pages); APPLICATIONS OF PERCEPTRONS TO PHOTO-INTERPRETATION, Summary Report, Report Number VE-1446-6-3, August 1963 (39 pages).

[24] TRIP Report, 23 April 1963, Rome Air Development Center, *Purpose: To Ascertain the Scope of Investigation of Automatic Image Recognition Techniques being Supported by RADC and to Obtain an Educated Opinion Regarding the Advisability of NPIC Supporting the (redacted) (PERCEPTRON) Effort in this Realm*. SECRET (Approved for Release 2005/05/02: CIA-RTP78B04770A002300030028-5.

[25] *National Photographic Interpretation Center Technical Development Program*. June 1963. SECRET NPIC Internal Use Only. Approved for Release 1999/09/07: CIA-RDP66R00546R000200030003-8, pp. 22–24.

[26] First-phase readout of aircraft missions meant a reporting deadline of two hours after the analyst received a copy of the film.

[27] All the literature about these missions mentions the interservice rivalry between U.S. Navy and U.S. Air Force tactical reconnaissance units; Michael Dobbs, "The Photographs that Prevented World War III," *Smithsonian Magazine*, October 2012 ; Brugioni, *Eyeball to Eyeball*, pp. 441–444; and William B. Ecker and Kenneth B. Jack, *BLUE MOON Over Cuba: Aerial Reconnaissance During the Cuban Missile Crisis* (Oxford, UK: Osprey, 2012), pp. 117–122, p. 223.

[28] In *Eyeball to Eyeball*, Brugioni mentions that the U-2 carried two 5,000-foot canisters of film, p. 185. Each of the BLUE MOON mission cameras carried 250 feet of film, but each mission used six aircraft for a total of 1,500 feet of film per mission. Ecker and Jack, *BLUE MOON over Cuba*, p. 70.

[29] KH-5 ARGON, https://en.wikipedia.org/wiki/KH-5_Argon (accessed 23 April 2021).

[30] O'Connor, *NPIC,* p. 55.

[31] Rosenblatt did publish a book on his research: *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms* (New York: Spartan Books, 1962).

[32] *Automatic Target Recognition Program.* R&D Catalog Form. 26 January 1966. SECRET, Approved for Release 2005/05/20: CIA-RDP78B04770A0011700010012-1.

33 As Cade Metz's recent history of AI recounts, neural network algorithms could only succeed with sufficient processing speed and sufficient data to train the algorithms multiple millions of times. This technology did not exist until the first decade of the twenty-first century, and it was not reliable until the second decade.

34 *AUTOMATIC TARGET RECOGNITION (ATR): TASK NO.6: NPIC OPERATIONS*, TOP SECRET. January 1967. Approved for release 2003/05/15: CIA-RDP99T01396R0003004400001-2, p. 2.

35 E-mail exchange between the author and Alana Staiti, curator of the History of Computers and Information Sciences, National Museum of American History, Smithsonian Institution, 6 April 2021.

36 Seymour Papert and Marvin Minsky, *Perceptrons: An Introduction to Computational Geography* (2nd ed. (Cambridge, MA: MIT Press, 1972); Woolridge, *A Brief History of Artificial Intelligence*, pp. 117–118.

37 Metz, *Genius Makers*, pp. 72–74.

38 *Ibid*., pp. 91–97, 86–88.

39 The Stanford Imagenet contest (https://en.wikipedia.org/wiki/ImageNet) uses a set of 14 million images.

40 Metz, *Genius Makers*, pp. 203–206.

41 Michael Wooldridge, *A Brief History of Artificial Intelligence: What It Is, Where We Are, and Where We Are Going* (New York: Flatiron Books, 2020), p. 119.

42 The AlexNet algorithm in 2012 raised the average number of successful image identifications in the Stanford ImageNet annual conference from 74% to 87%. Azeem Azhar, *The Exponential Age: How Accelerating Technology Is Transforming Business, Politics, and Society* (New York: Diversion Books, 2021), p. 20. To provide a sense of the growth in processing complexity, consider putting the mathematics in a temporal scale. If each permutation for Rosenblatt's 1960 Mark I perceptron took one second, it would take 56 hours, or 2.3 days, for all the permutations to take place. For the 2012 AlexNet Algorithm that won the ImageNet Contest at Stanford in 2012, with its 60 million parameters, it would take 694.4 days, or a year and nearly eleven months. And for the 2020 GPT-3 neural net, with its 175 billion parameters (Azhar, *The Exponential Age*, pp. 20–21) it would take 5,549.25 years to go through all its permutations. Said differently, it took 52 years, or from 1960 to 2012, to increase the neural-net processing by two orders of magnitude. It has increased by another two orders of magnitude in the past eight years.

43 Ursa Space, "An Inside Look at SAR-Based Measurements," 9 October 2020, https://ursaspace.com/blog/an-inside-look-at-sar-based-measurements/

44 Gilman Louie used this reference in his 2021 talk at the September GEOINT Symposium in St. Louis, Missouri.