

‘One of the most surprising and important stories of our time’

Ashlee Vance, bestselling author of *Elon Musk*

Genius Makers



The Mavericks Who Brought AI
to Google, Facebook and the World

CADE METZ

GENIUS MAKERS

GENIUS MAKERS

THE MAVERICKS WHO BROUGHT AI
TO GOOGLE, FACEBOOK AND THE WORLD

CADE METZ



1 3 5 7 9 10 8 6 4 2

Random House Business
20 Vauxhall Bridge Road
London SW1V 2SA

Random House Business is part of the Penguin Random House group of companies whose addresses can be found at global.penguinrandomhouse.com.



Copyright © Cade Metz 2021

Cade Metz has asserted his right to be identified as the author of this Work in accordance with the Copyright, Designs and Patents Act 1988.

First published in the United States by Dutton in 2021
First published in the UK by Random House Business in 2021

www.penguin.co.uk

A CIP catalogue record for this book is available from the British Library.

ISBN 9781847942135 (Hardback)
ISBN 9781847942142 (Trade paperback)

Printed and bound in Great Britain by Clays Ltd, Elcograf S.p.A.

The authorised representative in the EEA is Penguin Random House Ireland,
Morrison Chambers, 32 Nassau Street, Dublin D02 YH68.

Penguin Random House is committed to a sustainable future for our business, our readers and our planet. This book is made from Forest Stewardship Council® certified paper.



Copyrighted Material

*In memory of Walt Metz,
who believed in truth, goodness, and beauty*

It's the best possible time to be alive, when almost everything you thought you knew is wrong.

—Tom Stoppard, *Arcadia*, Act One, Scene IV

When we have found all the mysteries and lost all the meaning, we will be alone, on an empty shore.

—Act Two, Scene VII

CONTENTS

PROLOGUE

THE MAN WHO DIDN'T SIT DOWN 1
December 2012

PART ONE

A NEW KIND OF MACHINE

- 1 GENESIS 15
“Frankenstein Monster Designed by Navy That Thinks.”
- 2 PROMISE 28
“Old ideas are new.”
- 3 REJECTION 46
“I was definitely thinking I was right the whole time.”
- 4 BREAKTHROUGH 66
“Do what you want at Google—not what Google wants you to do.”
- 5 TESTAMENT 80
“The speed of light in a vacuum used to be about 35 mph.
Then Jeff Dean spent a weekend optimizing physics.”
- 6 AMBITION 99
“Let’s really go big.”

CONTENTS

PART TWO

WHO OWNS INTELLIGENCE?

- 7 RIVALRY** 119
“Hello, this is Mark, from Facebook.”
- 8 HYPE** 133
“Success is guaranteed.”
- 9 ANTI-HYPE** 152
“He could produce something evil by accident.”
- 10 EXPLOSION** 167
“He ran AlphaGo like Oppenheimer ran the Manhattan Project.”
- 11 EXPANSION** 179
“George wiped out the whole field without even knowing its name.”
- 12 DREAMLAND** 189
“It is not that people at Google drink different waters.”

PART THREE

TURMOIL

- 13 DECEIT** 203
“Oh, you can actually make photo-realistic faces.”
- 14 HUBRIS** 214
“I knew when I gave the speech that the Chinese were coming.”
- 15 BIGOTRY** 228
“Google Photos, y’all fucked up. My friend is not a gorilla.”
- 16 WEAPONIZATION** 239
“You probably heard Elon Musk and his comment about AI causing WW3.”
- 17 IMPOTENCE** 251
“There are people in Russia whose job it is to try to exploit our systems. So this is an arms race, right?”

CONTENTS

PART FOUR

HUMANS ARE UNDERRATED

- 18 DEBATE** 263
“However long the rapid progress continues, Gary will still argue it is about to end.”
- 19 AUTOMATION** 276
“If the room looked like things had gone crazy, we were on the right track.”
- 20 RELIGION** 286
“My goal is to successfully create broadly beneficial AGI. I also understand this sounds ridiculous.”
- 21 X FACTOR** 303
“History is going to repeat itself, I think.”
- ACKNOWLEDGMENTS 313
- TIMELINE 317
- THE PLAYERS 321
- NOTES 327
- INDEX 357

PROLOGUE

THE MAN WHO DIDN'T SIT DOWN

DECEMBER 2012

By the time he stepped onto the bus in downtown Toronto, bound for Lake Tahoe, Geoff Hinton hadn't sat down for seven years. "I last sat down in 2005," he often said, "and it was a mistake." He first injured his back as a teenager, while lifting a space heater for his mother. As he reached his late fifties, he couldn't sit down without risking a slipped disk, and if it slipped, the pain could put him in bed for weeks. So he stopped sitting down. He used a standing desk inside his office at the University of Toronto. When eating meals, he put a small foam pad on the floor and knelt at the table, poised like a monk at the altar. He lay down when riding in cars, stretching across the back seat. And when traveling longer distances, he took the train. He couldn't fly, at least not with the commercial airlines, because they made him sit during takeoff and landing. "It got to the point where I thought I might be crippled—that I wouldn't be able to make it through the day—so I

took it seriously,” he says. “If you let it completely control your life, it doesn’t give you any problems.”

That fall, before lying down at the back of the bus for the trip to New York, taking the train all the way to Truckee, California, at the crest of the Sierra Nevadas, and then stretching across the back seat of a taxi for the thirty-minute drive up the mountain to Lake Tahoe, he created a new company. It included only two other people, both young graduate students in his lab at the university. It made no products. It had no plans to make a product. And its website offered nothing but a name, DNNresearch, which was even less appealing than the website. The sixty-four-year-old Hinton—who seemed so at home in academia, with his tousled gray hair, wool sweaters, and two-steps-ahead-of-you sense of humor—wasn’t even sure he wanted to start a company until his two students talked him into it. But as he arrived in Lake Tahoe, one of the largest companies in China had already offered \$12 million for his newborn start-up, and soon three other companies would join the bidding, including two of the largest in the United States.

He was headed for Harrah’s and Harvey’s, the two towering casinos at the foot of the ski mountains on the south side of the lake. Rising up over the Nevada pines, these twin slabs of glass, steel, and stone also served as convention centers, offering hundreds of hotel rooms, dozens of meeting spaces, and a wide variety of (second-rate) restaurants. That December, they hosted an annual gathering of computer scientists called NIPS. Short for Neural Information Processing Systems—a name that looked deep into the future of computing—NIPS was a conference dedicated to artificial intelligence. A London-born academic who had explored the frontiers of AI at universities in Britain, the United States, and Canada since the early 1970s, Hinton made the trip to NIPS nearly every year. But this was different. Although Chinese interest in his

company was already locked in, he knew that others were interested, too, and NIPS seemed like the ideal venue for an auction.

Two months earlier, Hinton and his students had changed the way machines saw the world. They had built what was called a *neural network*, a mathematical system modeled on the web of neurons in the brain, and it could identify common objects—like flowers, dogs, and cars—with an accuracy that had previously seemed impossible. As Hinton and his students showed, a neural network could learn this very human skill by analyzing vast amounts of data. He called this “deep learning,” and its potential was enormous. It promised to transform not just computer vision but everything from talking digital assistants to driverless cars to drug discovery.

The idea of a neural network dated back to the 1950s, but the early pioneers had never gotten it working as well as they had hoped. By the new millennium, most researchers had given up on the idea, convinced it was a technological dead end and bewildered by the fifty-year-old conceit that these mathematical systems somehow mimicked the human brain. When submitting research papers to academic journals, those who still explored the technology would often disguise it as something else, replacing the words “neural network” with language less likely to offend their fellow scientists. Hinton remained one of the few who believed it would one day fulfill its promise, delivering machines that could not only recognize objects but identify spoken words, understand natural language, carry on a conversation, and maybe even solve problems humans couldn't solve on their own, providing new and more incisive ways of exploring the mysteries of biology, medicine, geology, and other sciences. It was an eccentric stance even inside his own university, which spent years denying his standing request to hire another professor who could work alongside him in this long and

winding struggle to build machines that learned on their own. “One crazy person working on this was enough,” he says. But in the spring and summer of 2012, Hinton and his two students made a breakthrough: They showed that a neural network could recognize common objects with an accuracy beyond any other technology. With the nine-page paper they unveiled that fall, they announced to the world that this idea was as powerful as Hinton had long claimed it would be.

Days later, Hinton received an email from a fellow AI researcher named Kai Yu, who worked for Baidu, the Chinese tech giant. On the surface, Hinton and Yu had little in common. Born in postwar Britain to a family of monumental scientists whose influence was matched only by their eccentricity, Hinton had studied at Cambridge, earned a PhD in artificial intelligence from the University of Edinburgh, and spent the next thirty years as a professor of computer science. Born thirty years after Hinton, Yu grew up in Communist China, the son of an automobile engineer, and studied in Nanjing and then Munich before moving to Silicon Valley for a job in a corporate research lab. The two were separated by class, age, culture, language, and geography, but they shared an unusual interest: neural networks. They had originally met in Canada at an academic workshop, part of a grassroots effort to revive this nearly dormant area of research across the scientific community and re-brand the idea as “deep learning.” Yu was among those who helped spread the new gospel. Returning to China, he took the idea to Baidu, where his research caught the eye of the company’s CEO. When that nine-page paper emerged from the University of Toronto, Yu told the Baidu brain trust they should hire Hinton as quickly as possible. With his email, he introduced Hinton to a Baidu vice president, who offered \$12 million for just a few years of work.

At first, Hinton's suitors in Beijing felt they had reached an agreement. But Hinton wasn't so sure. In recent months, he'd cultivated relationships inside several other companies, both small and large, including two of Baidu's big American rivals, and they, too, were calling his office in Toronto, asking what it would take to hire both him and his students. Seeing a much wider opportunity, he asked Baidu if he could solicit other offers before accepting the \$12 million, and when Baidu agreed, he flipped the situation upside down. Spurred on by his students and realizing that Baidu and its rivals were much more likely to pay enormous sums of money to acquire a company than they were to shell out the same dollars for a few new hires from the world of academia, he created his tiny start-up. He called it DNNresearch in a nod to the "deep neural networks" they specialized in, and he asked a Toronto lawyer how he could maximize the price of a start-up with three employees, no products, and virtually no history. As the lawyer saw it, he had two options: He could hire a professional negotiator and risk angering the companies he hoped would acquire his tiny venture, or he could set up an auction. Hinton chose an auction. In the end, four names joined the bidding for his new company: Baidu, Google, Microsoft, and a two-year-old start-up most of the world had never heard of. This was DeepMind, a London company founded by a young neuroscientist named Demis Hassabis that would grow to become the most celebrated and influential AI lab of the decade.

The week of the auction, Alan Eustace, Google's head of engineering, flew his own twin-engine plane into the airport near the south shore of Lake Tahoe. He and Jeff Dean, Google's most revered engineer, had dinner with Hinton and his students in the restaurant on the top floor of Harrah's, a steak house decorated with a thousand empty wine bottles. It was Hinton's sixty-fifth birthday. As he stood at a bar table and the others sat on high

stools, they discussed Google's ambitions, the auction, and the latest research under way at his lab in Toronto. For the Googlers, the dinner was mostly a way of running the rule over Hinton's two young students, whom they had never met. Baidu, Microsoft, and DeepMind also sent representatives to Lake Tahoe for the conference, and others played their own roles in the auction. Kai Yu, the Baidu researcher who'd kicked off the race for Hinton and his students, held his own meeting with the Toronto researchers before the bidding began. But none of the bidders ever gathered in the same place at the same time. The auction played out over email, with most bids arriving via corporate executives elsewhere in the world, from California to London to Beijing. Hinton hid the identity of each bidder from all the rest.

He ran the auction from his hotel room, number 731 in the Harrah's tower, which looked out over the Nevada pines and onto the snowy peaks of the mountains. Each day, he set a time for the next round of bidding, and at the designated hour, he and his two students would gather in his room to watch the bids arrive on his laptop. The laptop sat on a trash can turned upside down on a table at the end of the room's two queen-sized beds, so that Hinton could type while standing up. The bids arrived via Gmail, the online email service operated by Google, just because this was where he kept an email account. But Microsoft didn't like the arrangement. In the days before the auction, the company complained that Google, its biggest rival, could eavesdrop on its private messages and somehow game the bids. Hinton had discussed the same possibility with his students, though for him this was less a serious concern than an arch comment on the vast and growing power of Google. Technically, Google could read any Gmail message. The terms of service said it wouldn't, but the reality was that if it ever violated those terms, no one was likely to know. In the end, both

Hinton and Microsoft set their concerns aside—"We were fairly confident Google wouldn't read our Gmail," he says—and though no one quite realized it at the time, it was a moment loaded with meaning.

The auction rules were simple: After each bid, the four companies had an hour to raise the buying price by at least a million dollars. This hour-long countdown started at the time stamped on the email holding the latest bid, and at the end of the hour, if no one lodged a new bid, the auction was over. DeepMind bid with company shares, not cash, but it couldn't compete with the giants of the tech world and soon dropped out. That left Baidu, Google, and Microsoft. As the bids continued to climb, first to \$15 million and then to \$20 million, Microsoft dropped out, too, but then it came back in. Each small moment seemed large and heavy with meaning as Hinton and his students debated which company they'd rather join. Late one afternoon, as they looked out the window at the peaks of the ski mountains, two airplanes flew past from opposite directions, leaving contrails that crossed in the sky like a giant *X*. Amid the atmosphere of excitement in the room, they wondered what this meant, before remembering that Google was headquartered in a place called Mountain View. "Does that mean we should join Google?" Hinton asked. "Or does it mean we shouldn't?"

At about \$22 million, Hinton temporarily suspended the auction for a discussion with one of the bidders, and half an hour later, Microsoft dropped out again. That left Baidu and Google, and as the hours passed, the two companies took the price still higher. Kai Yu handled the initial Baidu bids, but when the price reached \$24 million, a Baidu executive took over from Beijing. From time to time, Yu would stop by room 731, hoping to glean at least a small sense of where the auction was headed.

Though Yu wasn't even remotely aware of it, these visits were a

problem for Hinton. He was sixty-five years old, and he often got sick when he traveled to Lake Tahoe, where the air was cold, thin, and dry. He was worried he might get sick again, and he didn't want Yu, or anyone else, to see him that way. "I didn't want them thinking I was old and decrepit," he says. So he removed the mattress from the pullout couch against the wall, laid it on the floor between the two beds, stretched an ironing board and a few other sturdy objects across the gap between the beds, dampened several towels with water, laid them across the gap, too, and slept each night in the wet air under this makeshift canopy. This, Hinton thought, would keep the illness at bay. The trouble was that as the auction continued, Yu, a small, round-faced man with glasses, kept dropping in for a chat. Now Hinton didn't want Yu to see how determined he was not to get sick. So each time Yu stopped by, Hinton turned to his two students, the only other people in his three-person company, and asked them to hide the mattress and the ironing board and the wet towels. "This is what vice presidents do," he told them.

After one visit, Yu left the room without his backpack, and when Hinton and his students noticed it sitting on a chair, they wondered if they should open it to see if anything inside would tell them how high Baidu was willing to bid. But they didn't, feeling it just wasn't the right thing to do. In any case, they soon realized Baidu was willing to go much higher: \$25 million, \$30 million, \$35 million. Inevitably, the next bid wouldn't arrive until a minute or two before the top of the hour, extending the auction just as it was on the verge of ending.

The price climbed so high, Hinton shortened the bidding window from an hour to thirty minutes. The bids quickly climbed to \$40 million, \$41 million, \$42 million, \$43 million. "It feels like

we're in a movie," he said. One evening, close to midnight, as the price hit \$44 million, he suspended the bidding again. He needed some sleep.

The next day, about thirty minutes before the bidding was set to resume, he sent an email saying the start would be delayed. About an hour later, he sent another. The auction was over. At some point during the night, Hinton had decided to sell his company to Google—without pushing the price any higher. His email to Baidu said that any other messages the company sent would be forwarded to his new employer, though he didn't say who that was.

This, he later admitted, was what he had wanted all along. Even Kai Yu had guessed that Hinton would end up at Google, or at least another American company, because his back would keep him from traveling to China. As it was, Yu was happy just for Baidu to have taken its place among the bidders. By pushing its American rivals to the limit, he believed, the Baidu brain trust had come to realize how important deep learning would be in the years ahead.

Hinton stopped the auction because finding the right home for his research was ultimately more important to him than commanding the maximum price. When he told the bidders at Google he was stopping the auction at \$44 million, they thought he was joking—that he couldn't possibly give up the dollars that were still coming. He wasn't joking, and his students saw the situation much as he did. They were academics, not entrepreneurs, more loyal to their idea than to anything else.

But Hinton didn't realize how valuable their idea would prove to be. No one did. Alongside a small group of other scientists—spread across those same four companies, one more American Internet giant, and, eventually, a new upstart—Hinton and his students soon pushed this single idea into the heart of the tech industry. In

doing so, they suddenly and dramatically accelerated the progress of artificial intelligence, including talking digital assistants, driverless cars, robotics, automated healthcare, and—though this was never their intention—automated warfare and surveillance. “It changed the way that I looked at technology,” Alan Eustace says. “It changed the way many others looked at it, too.”

Some researchers, most notably Demis Hassabis, the young neuroscientist behind DeepMind, even believed they were on their way to building a machine that could do *anything* the human brain could do, only better, a possibility that has captured the imagination since the earliest days of the computer age. No one was quite sure when this machine would arrive. But even in the near term, with the rise of machines that were still a very long way from true intelligence, the social implications were far greater than anyone realized. Powerful technologies have always fascinated and frightened humanity, and humanity has gambled on them time and again. This time, the stakes were higher than the scientists behind the idea could know. The rise of deep learning marked a fundamental change in the way digital technology was built. Rather than carefully defining how a machine was supposed to behave, one rule at a time, one line of code at a time, engineers were beginning to build machines that could learn tasks through their own experiences, and these experiences spanned such enormous amounts of digital information, no human could ever wrap their head around it all. The result was a new breed of machine that was not only more powerful than before but also more mysterious and unpredictable.

As Google and other tech giants adopted the technology, no one quite realized it was learning the biases of the researchers who built it. These researchers were mostly white men, and they didn’t see the breadth of the problem until a new wave of researchers—both women and people of color—put a finger on it. As the technology

spread even further, into healthcare, government surveillance, and the military, so did the ways it might go wrong. Deep learning brought a power even its designers did not completely understand how to control, particularly as it was embraced by tech superpowers driven by an insatiable hunger for revenue and profit.

After Hinton's auction played out in Lake Tahoe and the NIPS conference came to an end, Kai Yu boarded a plane for Beijing. There, he ran into a China-born Microsoft researcher named Li Deng, who had his own history with Hinton and had played his own role in the auction. Yu and Deng knew each other from years of AI conferences and workshops, and they arranged for adjacent seats on the long flight to Asia. Because Hinton kept the bidders anonymous, neither was quite sure which companies were involved in the auction. They wanted to know, and Deng loved to talk. They spent hours standing in the back of the cabin, discussing the rise of deep learning. But they also felt bound by their employers not to reveal their own involvement in the auction. So they danced around the issue, trying to understand what the other knew without giving their own secrets away. Though they didn't say it, both knew that a new competition was on. Their companies would have to answer Google's big move. That was how the tech industry worked. It was the beginning of a global arms race, and this race would quickly escalate in ways that would have seemed absurd a few years before.

Meanwhile, Geoff Hinton took the train back to Toronto. He would eventually make his way to the Google headquarters in Mountain View, California, but even as he joined the company, he retained his professorship at the University of Toronto and held tight to his own aims and beliefs, setting an example for the many other academics who soon followed him into the world's largest tech companies. Years later, when asked to reveal the companies that bid for his start-up, he answered in his own way. "I signed

contracts saying I would never reveal who we talked to. I signed one with Microsoft and one with Baidu and one with Google,” he said. “Best not to go into that.” He did not mention DeepMind. But that was another story. In the wake of the auction in Lake Tahoe, Demis Hassabis, the founder of the London lab, imprinted his own views on the world. In some ways, they echoed Hinton’s. In other ways, they looked even further into the future. Soon, Hassabis, too, was caught up in the same global arms race.

This is the story of Hinton and Hassabis and the other scientists who sparked this race, a small but eclectic group of researchers from around the globe who nurtured an idea for decades, often in the face of unfettered skepticism, before it suddenly came of age and they were sucked into the machinery of some of the largest companies on Earth—and a world of turmoil none of them expected.

PART ONE

A NEW KIND OF MACHINE

1

GENESIS

“FRANKENSTEIN MONSTER DESIGNED
BY NAVY THAT THINKS.”

On July 7, 1958, several men gathered around a machine inside the offices of the United States Weather Bureau in Washington, D.C., about fifteen blocks west of the White House. As wide as a kitchen refrigerator, twice as deep, and nearly as tall, the machine was just one piece of a mainframe computer that fanned across the room like a multipiece furniture set. It was encased in silvery plastic, reflecting the light from above, and the front panel held row after row of small round lightbulbs, red square buttons, and thick plastic switches, some white and some gray. Normally, this \$2 million machine ran calculations for the Weather Bureau, the forerunner of the National Weather Service, but on this day, it was on loan to the U.S. Navy and a twenty-nine-year-old Cornell University professor named Frank Rosenblatt.

As a newspaper reporter looked on, Rosenblatt and his Navy cohorts fed two white cards into the machine, one marked with a

small square on the left, the other marked on the right. Initially, the machine couldn't tell them apart, but after it read another fifty cards, that changed. Almost every time, it correctly identified where the card was marked—left or right. As Rosenblatt explained it, the machine had learned this skill on its own, thanks to a mathematical system modeled on the human brain. He called it a Perceptron. In the future, he said, this system would learn to recognize printed letters, handwritten words, spoken commands, and even people's faces, before calling out their names. It would translate one language into another. And in theory, he added, it could clone itself on an assembly line, explore distant planets, and cross the line from computation into sentience.

"The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself, and be conscious of its existence," read the article that appeared the next morning in the *New York Times*. A second article, in the Sunday edition, said that Navy officials hesitated to call this a machine because it was "so much like a human being without life." Rosenblatt grew to resent the way the popular press covered the event, particularly a headline in Oklahoma ("Frankenstein Monster Designed by Navy That Thinks"). In later years, among colleagues and in his published writings, he described the project in more measured terms. He insisted it was not an attempt at artificial intelligence, and he acknowledged its limitations. Still, the idea slipped from his grasp.

The Perceptron was one of the first neural networks, an early incarnation of the technology Geoff Hinton would auction to the highest bidder more than fifty years later. But before it reached that \$44 million moment, let alone the extravagant future predicted across the pages of the *New York Times* in the summer of 1958, it descended into academic obscurity. By the early 1970s, after those

lavish predictions met the limitations of Rosenblatt's technology, the idea was all but dead.

FRANK Rosenblatt was born on July 11, 1928, in New Rochelle, New York, just north of the Bronx. He attended Bronx Science, the elite public high school that eventually produced eight Nobel laureates, six Pulitzer Prize winners, eight National Medal of Science winners, and three recipients of the Turing Award, the world's top computer science prize. A small, thin man with fleshy jowls and short, dark, wavy hair who wore standard-issue black-rimmed glasses, Rosenblatt was trained in psychology, but his interests were much wider. In 1953, the *New York Times* published a small story describing an early computer he used to crunch data for his PhD thesis. Called EPAC—short for “electronic profile-analyzing computer”—it analyzed the psychological profiles of his patients. As the years passed, he came to believe that machines could provide an even greater understanding of the mind. After finishing his PhD, he joined the Cornell Aeronautical Laboratory in Buffalo, about a hundred and fifty miles from the university's main campus in Ithaca, New York. Donated to Cornell by a company that designed aircraft during the Second World War, this flight research center morphed into a more eclectic lab in the postwar years, operating with little oversight from the administration back in Ithaca. It was here that Rosenblatt designed the Perceptron, backed by funding from the Office of Naval Research.

Rosenblatt saw the project as a window into the inner workings of the brain. If he could re-create the brain as a machine, he believed, he could plumb the mysteries of what he called “natural intelligence.” Drawing on ideas initially proposed by two researchers at the University of Chicago a decade earlier, the Perceptron analyzed objects and looked for patterns that could identify these

objects (for example, whether a card had a mark on the left-hand or right-hand side). It did this using a series of mathematical calculations that operated (in a very broad sense) like the network of neurons in the brain. As the Perceptron examined and tried to identify each object, it got some right and others wrong. But it could learn from its mistakes, methodically adjusting each of those mathematical calculations until the mistakes were few and far between. Much like a neuron in the brain, each calculation was nearly meaningless on its own—just an input for a larger algorithm. But the larger algorithm—a kind of mathematical recipe—could actually do something useful. Or at least that was the hope. At the Weather Bureau in the summer of 1958, Rosenblatt showed off the beginnings of this idea—a simulation of the Perceptron that ran on the bureau's IBM 704, the leading commercial computer of the day. Then, back at the lab in Buffalo, working alongside a team of engineers, he began building an entirely new machine around the same idea. He called it the Mark I. Unlike other machines of the day, it was designed to see the world around it. "For the first time, a non-biological system will achieve an organization of its external environment in a meaningful way," he told a reporter later that year, during another trip to meet with his backers in Washington.

His primary collaborator at the Office of Naval Research did not view the Perceptron in the same extravagant terms. But Rosenblatt was unmoved. "My colleague disapproves of all the loose talk one hears nowadays about mechanical brains," he told the reporter, over a cup of coffee. "But that is exactly what it is." A small, silver pitcher of cream sat on the table in front of him, and he picked it up. Though this was the first time he'd laid eyes on the pitcher, Rosenblatt said, he could still recognize it as a pitcher. The Perceptron, he explained, could do much the same. It could draw the conclusions needed to distinguish, say, a dog from a cat. He admit-

ted that the technology was a long way from having practical uses: It lacked depth perception and “the refinements of judgment.” But he was confident of its potential. One day, he said, the Perceptron would travel into space and send its observations back to Earth. When the reporter asked if there was anything the Perceptron was *not* capable of, Rosenblatt threw up his hands. “Love. Hope. Despair. Human nature, in short,” he said. “If we don’t understand the human sex drive, how should we expect a machine to?”

That December, the *New Yorker* hailed Rosenblatt’s creation as the first serious rival to the brain. Previously, the magazine had marveled that the IBM 704 could play a game of chess. Now it described the Perceptron as an even more remarkable machine, a computer that could achieve “what amounts to human thought.” Though scientists claimed that only biological systems could see, feel, and think, the magazine said, the Perceptron behaved “*as if* it saw, felt, and thought.” Rosenblatt had not yet built the machine, but this was dismissed as a minor obstacle. “It is only a question of time (and money) before it comes into existence,” the magazine said.

Rosenblatt completed the Mark I in 1960. It spanned six racks of electrical equipment, each the size of a kitchen fridge, and it plugged into what seemed to be a camera. This *was* a camera, though engineers had removed the film loader, swapping in a small square device covered with four hundred black dots. These were photocells that responded to changes in light. Rosenblatt and his engineers would print block letters on squares of cardboard—*A*, *B*, *C*, *D*, etc.—and when they placed these squares on an easel in front of the camera, the photocells could read the black lines of the letters against the white space of the cardboard. As they did, the Mark I learned to recognize the letters, much like the IBM mainframe inside the Weather Bureau learned to recognize the marked cards. This required a little help from the humans in the room: As

it worked to identify the letters, a technician would tell the machine whether it was right or wrong. But eventually, the Mark I would learn from its own hits and misses, pinpointing the patterns that identified the slanted line of an *A* or the double curve of a *B*. When demonstrating the machine, Rosenblatt had a way of proving this was learned behavior. He would reach into the racks of electrical equipment and pull out a few wires, breaking the connections between the motors that acted as faux-neurons. When he reconnected the wires, the machine once again struggled to recognize the letters, but then, after examining more cards and relearning the same skill, it worked as it did before.

This electrical contraption worked well enough to attract interest beyond the Navy. Over the next several years, the Stanford Research Institute, or SRI, a lab in Northern California, began exploring the same ideas, and Rosenblatt's own lab won contracts with both the U.S. Postal Service and the Air Force. The Postal Service needed a way of reading addresses on envelopes, and the Air Force hoped to identify targets in aerial photos. But all that was still in the future. Rosenblatt's system was only marginally effective when reading printed letters, a relatively simple task. As the system analyzed cards printed with the letter *A*, each photocell examined a particular spot on the card—say, an area near the lower right-hand corner. If the spot was black more often than white, the Mark I assigned it a high “weight,” meaning it would play a more important role in the mathematical calculation that ultimately determined what was an *A* and what wasn't. When reading a new card, the machine could recognize an *A* if most of the highly weighted spots were colored black. That was about it. The technology wasn't nearly nimble enough to read the irregularities of handwritten digits.

Despite the system's obvious deficiencies, Rosenblatt remained optimistic about its future. Others, too, believed the technology

would improve in the years to come, learning more complex tasks in more complex ways. But it faced a significant obstacle: Marvin Minsky.

FRANK Rosenblatt and Marvin Minsky had been contemporaries at Bronx Science. In 1945, Minsky's parents moved him to Phillips Andover, the model American prep school, and after the war, he enrolled at Harvard. But he complained that neither could match his experience at Science, where the coursework had been more challenging and the students more ambitious—"people you could discuss your most elaborate ideas with and no one would be condescending," he said. After Rosenblatt died, Minsky pointed to his old schoolmate as the kind of creative thinker who walked the halls of Science. Like Rosenblatt, Minsky was a pioneer in the field of artificial intelligence, but he viewed the field through a different lens.

As an undergraduate at Harvard, using over three thousand vacuum tubes and a few parts from an old B-52 bomber, Minsky built what may have been the first neural network, a machine he called SNARC. Then, as a graduate student in the early '50s, he continued to explore the mathematical concepts that eventually gave rise to the Perceptron. But he came to see artificial intelligence as a larger endeavor. He was among the small group of scientists who crystalized AI as its own field of study during a gathering at Dartmouth College in the summer of 1956. A Dartmouth professor named John McCarthy had urged the wider academic community to explore an area of research he called "automata studies," but that didn't mean much to anyone else. So he recast it as "artificial intelligence," and that summer, he organized a conference alongside several like-minded academics and other researchers. The agenda at the Dartmouth Summer Research Conference on Artificial

Intelligence included “neuron nets,” but also “automatic computers,” “abstractions,” and “self-improvement.” Those who attended the conference would lead the movement into the 1960s, most notably McCarthy, who eventually took his research to Stanford University on the West Coast; Herbert Simon and Alan Newell, who built a lab at Carnegie Mellon in Pittsburgh; and Minsky, who settled at the Massachusetts Institute of Technology in New England. They aimed to re-create human intelligence using whatever technology could get them there, and they were sure it wouldn’t take very long, some arguing that a machine would beat the world chess champion and discover its own mathematical theorem within a decade. Bald from a young age, with wide ears and an impish grin, Minsky became an AI evangelist, but his evangelism didn’t extend to neural networks. A neural network was just one way of building artificial intelligence, and Minsky, like many of his colleagues, began exploring other avenues. By the mid-’60s, as other techniques grabbed his attention, he questioned whether neural networks could handle anything beyond the simple tasks Rosenblatt demonstrated at his lab in upstate New York.

Minsky was part of a larger backlash against Rosenblatt’s ideas. As Rosenblatt himself wrote in his 1962 book *Principles of Neurodynamics*, the Perceptron was a controversial concept among academics, and he placed much of the blame on the press. The reporters who wrote about his work in the late 1950s, Rosenblatt said, “fell to the task with all of the exuberance and sense of discretion of a pack of happy bloodhounds.” He lamented, in particular, headlines like the one in Oklahoma, saying they were a long way from inspiring confidence in his work as a serious scientific pursuit. Four years after the event in Washington, pulling back on his own early claims, he insisted the Perceptron was not an attempt at artificial intelligence—at least, not as researchers like Minsky saw AI. “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,” he wrote. “Its utility is in enabling us to determine the physical conditions for the emergence of various psychological properties.” In other words, he wanted to understand how the human brain worked, rather than send a new brain out into the world. Because the brain was a mystery, he couldn’t re-create it. But he believed he could use machines to explore this mystery, and maybe even solve it.

From the beginning, the lines that separated artificial intelligence from computer science, psychology, and neuroscience were blurred as various academic camps sprung up around this new breed of technology, each mapping the landscape in its own way. Some psychologists, neuroscientists, and even computer scientists came to see machines in the way Rosenblatt did: as a reflection of the brain. Others looked on this grandiose idea with scorn, arguing that computers operated nothing like the brain, and that if they were ever going to mimic intelligence, they would have to do it in their own way. No one, however, was anywhere close to building what could rightly be called “artificial intelligence.” Though the field’s founding fathers thought the path to re-creating the brain would be a short one, it turned out to be very long. Their original sin was that they called their field “artificial intelligence.” This gave decades of onlookers the impression that scientists were on the verge of re-creating the powers of the brain when, in reality, they were not.

In 1966, a few dozen researchers traveled to Puerto Rico, gathering at the Hilton hotel in San Juan. They met to discuss the latest advances in what was then called “pattern recognition”—technology that could identify patterns in images and other data. Whereas Rosenblatt viewed the Perceptron as a model of the brain, others

saw it as a means of pattern recognition. In later years, some commentators imagined Rosenblatt and Minsky locking horns at academic conferences like the one in San Juan, openly debating the future of the Perceptron, but their rivalry was implicit. Rosenblatt didn't even travel to Puerto Rico. Inside the Hilton, the tension emerged when a young scientist named John Munson addressed the conference. Munson worked at SRI, the Northern California lab that embraced Rosenblatt's ideas after the arrival of the Mark I. There, alongside a larger team of researchers, he was trying to build a neural network that could read handwritten characters, not just printed letters, and with his presentation at the conference, he aimed to show the progress of this research. But when Munson finished the lecture and took questions from the floor, Minsky made himself heard. "How can an intelligent young man like you," he asked, "waste your time with something like this?"

Sitting in the audience, Ron Swonger, an engineer from the Cornell Aeronautical Lab, the birthplace of the Mark I, was appalled. He bristled at Minsky's language, and he questioned whether the attack had anything to do with the presentation delivered from the front of the room. Minsky wasn't concerned with recognizing handwritten characters. He was attacking the very idea of the Perceptron. "This is an idea with no future," he said. Across the room, Richard Duda, who was part of the team trying to build the system for handwritten characters, was stung by the laughter from the audience as Minsky made light of claims that the Perceptron mirrored the network of neurons in the brain. The performance was typical of Minsky, who enjoyed stirring public controversy. He once greeted a roomful of physicists by saying that the field of artificial intelligence had made more progress in just a few years than physics had made in centuries. But Duda also felt the MIT professor had practical reasons for attacking the work at places like SRI

and Cornell: MIT was competing with these labs for the same government research dollars. Later in the conference, when another researcher presented a new system designed to create computer graphics, Minsky praised its ingenuity—and took another swipe at Rosenblatt’s ideas. “Can a Perceptron do that?” he said.

In the wake of the conference Minsky and an MIT colleague named Seymour Papert published a book on neural networks, which they titled *Perceptrons*. Many felt it shut the door on Rosenblatt’s ideas for the next fifteen years. Minsky and Papert described the Perceptron in elegant detail, exceeding, in many respects, the way Rosenblatt described it himself. They understood what it could do, but they also understood its flaws. The Perceptron, they showed, couldn’t deal with what mathematicians called “exclusive-or,” an esoteric concept that carried much larger implications. When presented with two spots on a cardboard square, the Perceptron could tell you if both were colored black. And it could tell you if both were white. But it couldn’t answer the straightforward question: “Are they *two different colors*?” This showed that in some cases, the Perceptron couldn’t recognize simple patterns, let alone the enormously complex patterns that characterized aerial photos or spoken words. Some researchers, Rosenblatt among them, were already exploring a new kind of Perceptron that aimed to fix this flaw. Still, in the wake of Minsky’s book, the government dollars moved into other technologies, and Rosenblatt’s ideas faded from view. Following Minsky’s lead, most researchers embraced what was called “symbolic AI.”

Frank Rosenblatt aimed to build a system that learned behavior on its own in the same way the brain did. In later years, scientists called this “connectionism,” because, like the brain, it relied on a vast array of interconnected calculations. But Rosenblatt’s system was much simpler than the brain, and it learned only in small ways.