'The raw emotion of that encounter in New York bursts out of the pages of Kasparov's gripping story. What is striking, and reassuring, is that far from raging against the machine, Kasparov marvels at the capabilities of computers and is excited by the possibilities for future collaboration. This reads at times like a fast-paced psychological thriller. Chess fans will be engrossed by Kasparov's tale but the book deserves a far broader readership' John Thornhill, *Financial Times*

'An absorbing, page-turning thriller that weaves a personal account of intellectual combat with the wider picture of what it's like to come up against a powerful corporation that is determined to do whatever it takes to crush opposition' *Observer*

'As Kasparov recounts in arresting detail what it felt like to compete cognitively with a machine, he extrapolates his experience into an optimistic perspective on how computerized intelligence can enhance rather than overwhelm human brainpower, and instead of eliminating jobs and opportunities, can actually generate them' *Booklist*

'Intelligent, absorbing . . . thoughtful reading for anyone interested in human and machine cognition and a must for chess fans' *Kirkus*

'Garry Kasparov's perspectives on artificial intelligence are borne of personal experience – and despite that, are optimistic, wise and compelling. It's one thing for the giants of Silicon Valley to tell us our future is bright; it is another thing to hear it from the man who squared off with the world's most powerful computer, with the whole world watching, and his very identity at stake'
Charles Duhigg, author of *Smarter Faster Better*

'The great Garry Kasparov takes on the key economic issue of our time: how we can thrive as humans in a world of thinking machines. This important and optimistic book explains what we as humans are uniquely qualified to do. Instead or wringing our hands about robots, we should all read this book and embrace the future' Walter Isaacson, author of *The Innovators*

'An absorbing, often brilliant book which no chess-lover should miss' Edward Winter, chesshistory.com

DEEP THINKING

where machine intelligence ends and human creativity begins

GARRY KASPAROV

WITH MIG GREENGARD

JOHN MURRAY

a friendly exhibition, or a sort of science experiment. This wouldn't be the case in the coming years, however, as the machines got stronger and began to appear in serious tournaments where money and prestige were at stake, not merely the future of humanity.

CHAPTER 7

THE DEEP END

I AM A SORE LOSER.

I want to clear that up right at the start. I hate losing. I hated losing bad games and I hated losing good ones. I hated losing to weak players and I hated losing to world champions.

I have had sleepless nights after losses. I have had angry outbursts at award ceremonies after a bad defeat. I have been annoyed to discover that I missed a good move in a game I lost twenty years ago when analyzing it for this book.

I hate to lose, and not just at chess. I hate to lose at trivia games. I hate to lose at card games. (My complete lack of a poker face is why I rarely play them.)

Being a sore loser is not the attribute I'm most proud of, nor am I ashamed of it. To be the best in any competitive endeavor you have to hate losing more than you are afraid of it. The thrill of victory is wonderful, although I think any elite sportsman gets used to that feeling at a very young age. Everyone has different methods for finding motivation, especially over a long career. But no matter how much you love the game, you have to have to hate to lose if you want to stay on top. You have to care, and care deeply.

A database can bring up a list of practically every serious game of chess I've played since I was twelve years old, over twenty-four hundred games. Of those, I lost roughly 170 times. Counting only the tournament and match games over the course of my twenty-five-year professional career, starting when I was seventeen, the number of losses drops to around half that. If I was never a good loser, it was partly because I never had the chance to get good at it. In 1990, English

Grandmaster Raymond Keene wrote a book called *How to Beat Gary Kasparov* that collected all of my defeats up to that point. The book's introduction begins: "Beating Gary Kasparov at chess is considerably more difficult than climbing Mount Everest or becoming a dollar billionaire. . . . I learnt that it was six times easier to reach the peak of Everest . . . five times easier to acquire more than $1,000,000,000." Those few who defeated me might wonder if they should have gone into a different line of work.

I want to get all this out of the way because my attitude about losing inevitably comes up in any discussion of my match with the IBM supercomputer Deep Blue. To be more precise, my rematch with Deep Blue in 1997.

I am resigned to the fact that almost no one remembers I beat Deep Blue in our first match in 1996. "This Day in History" calendars don't have entries for all the failed attempts to fly across the Atlantic before Charles Lindbergh succeeded in 1927. When the 1996 match is remembered at all, it's because my loss in game one was the first time a machine had beaten the world champion in a classical time control game. Prior to that, I had played quite a few games against machines at faster time controls and lost a number of them. What we call "rapid" games allow between fifteen and thirty minutes per player for the whole game. Faster still is "blitz" chess, in which the players have five minutes or even less on the clock to begin the game. There is even "bullet" chess of just one or two minutes, which almost turns chess into an aerobic activity.

At least since the 1970s, the faster the game, the greater the advantage for a computer against a human. Grandmasters may play largely by intuition, but chess is a concrete game in the end. Without the time to calculate properly against a machine that is checking millions of positions per second, a blitz game can quickly become a bloodbath. The slight inaccuracies and tactical oversights humans make routinely against each other at fast time controls are instantly punished by the machines, and they never return the favor.

After beating Deep Thought in 1989, a few years went by before I played another machine opponent in a public match. Partly this was

because there was no market for me beating up on computers when they clearly had some work to do before really challenging me, and my time was valuable. I narrowly won my fifth world championship match against Karpov in 1990 while dealing with the sudden collapse of my home country. Along with thousands of others, my family and I had been forced to flee Baku in the face of Armenian pogroms as the Soviet Union unraveled.

But I was keeping an eye on the machines' progress. I had the latest programs installed on my personal computer, using them for analysis on occasion and playing against them for fun now and then. They didn't play good chess, but programs with names like Genius and Fritz were already tactically very dangerous even on an average home PC or laptop. One moment of human inattention in a quick game and pow, it was over.

I also crossed paths with Deep Thought again, in 1991 at a computer exhibition in Hanover, Germany. The machine's team had lost and gained a few members in its transition to becoming a major IBM project. Feng-hsiung Hsu and Murray Campbell were still the team leaders and they were both in Hanover, where Deep Thought had been invited to participate in the strongest tournament yet to include a machine. It was a closed event, with six German Grandmasters and a strong International Master, with an average rating of 2514.

Now with the formidable resources of IBM behind him, Hsu was still working on his upgraded dream machine with a thousand VLSI chips, but it wasn't yet ready. Deep Thought was still the strongest machine in the world and was expected to be a contender in Hanover, based on its past performances. It was a little surprising that it finished next to last with 2.5/7, winning two, drawing one, and losing four. The team blamed two of the losses on mistakes in the opening book (another reoccurring theme), although looking at its Hanover games now, it also just didn't play very good chess.

Of more interest was a little test for me, proposed by my friend Frederic Friedel, who was one of the Hanover event's organizers. I was shown the games from the first five rounds of the tournament to see if I could figure out which player was Deep Thought. It was a chess twist

on the Turing test, to see if a computer could pass for a Grandmaster. I managed to pick out two correctly and narrowed down another round to two games before choosing the wrong one, so three of the computer's five games passed the test. To me, this was a better indicator of computer chess progress than its score in the tournament. Some of its games followed the old patterns of terrible strategic play and unseemly greed balanced out by startling tactics. But other games just looked like chess, if still far from the world championship level.

I also thought this was interesting because I could imagine one day turning the tables. In ten years' time, roughly my guess for when a computer would be strong enough to beat me, would a super-strong machine be able to analyze human games insightfully? I spent a lot of time scrutinizing the tendencies and weaknesses of my opponents, but I was aware that this analysis was colored by my own tendencies and weaknesses. Machines, on the other hand, were objective. Chess engines were already proving to be useful for assisting in analysis, if mostly only for a tactical "blunder check." But once they were strong enough, I thought, maybe they would be able to detect patterns and habits in human games, both in my opponents' and my own.

This idea never really got off the ground, partly because the potential market for it was so small. There are only a few hundred players in the world who play the same opponents regularly enough to need to prepare for them specifically on a regular basis. ChessBase did eventually add some useful features like automatically building data-based player profiles, including their favorite openings and selected games. These were more time-savers than analytical tools, however. There was no advanced tendency breakdown like "often makes mistakes when his king is under attack" or "likes to trade queens when playing with black." The thought of such in-depth profiling also made some players a little uncomfortable, even though the data was all publicly available—their own games. I would love to know what a machine would say about me and my games.

I'm also very interested in what data-driven computer analysis of human behavior can do for fields like psychology, or in my realm of decision making. No one reading this would want to hand over all their

texts, email, social media posts, search history, shopping history, and the rest of the long digital trail we create hourly, at least, not to a human. But different apps and services already have all that information, for better or worse, and I'm sure that enough data and enough crunching would find many fascinating correlations, perhaps even diagnose things like depression or the early signs of dementia.

Facebook has suicide prevention tools that allow friends to flag posts for staff review and possible referral, but this requires human intervention. Fitness trackers are already monitoring everything from sleep habits to heart rates to calories burned. Google, Facebook, and Amazon probably know more about you than you know about yourself already, but people would be unnerved by seeing that analysis reflected back at them, perhaps revealing uncomfortable truths.

There are countless privacy issues to be negotiated anytime such data is accessed, of course, and that trade-off will continue to be one of the main battlefields of the AI revolution. I would want to know what a machine says about my chess or my mental and physical health, but would I want anyone else to know? You might want your family and your doctor to have all this information, but what about your insurance company or your employer? Social media reviews are already part of the hiring process in some companies. Anti-discrimination laws in the United States make it illegal to ask applicants about age, gender, race, and health, but algorithmic social media analysis can identify those in a split second, as well as make very accurate guesses at things like sexual preference, political leanings, and income level.

History tells us that eventually the desire for services wins out over a vague desire for privacy. We like sharing personal information on social media. We like to have books and music recommended to us by the algorithms of Netflix and Amazon. We won't give up GPS maps and directions even though using them means dozens of private companies know where we are practically every minute of the day—information that can also be accessed by governments and courts. When Gmail introduced ads based on scanning the content of people's email there was a collective shock, but it didn't last long. It's only an algorithm and,

if you're going to see an ad, wouldn't you rather see one you're interested in than one you aren't?

This is not an argument for surrendering to Big Brother. Coming from the country on which George Orwell based his novel *1984*, I am particularly sensitive to any encroachment on individual freedom. Surveillance can be an instrument of security or of repression, especially with the sophisticated tools available now. All the wonderful communication technology we depend on today is agnostic, neither good nor evil. Assuming that the Internet would magically set everyone free, as some appeared to believe, was foolish. Modern dictatorships and other political cliques are tech savvy and have learned how to limit and exploit these powerful new mediums. I'm glad privacy advocates are on the job, especially regarding the powers of the government. I just think they are fighting a losing battle because the tech will continue to improve and because the people they are trying to protect won't defend themselves. The barrage of privacy notices has become like all the disregarded warnings about the dangers of trans fats and corn syrup. We want to be healthy, but we like doughnuts more. The greatest security problem we have will always be human nature.

Technology will continue to make the benefits of sharing our data practically irresistible. Digital assistants like Amazon's Echo and Google Home listen to every word and sound in the home and people are buying them by the millions. Utility always wins. Even more invasive tech, like microsensors in our plumbing or implanted in our food or bodies, will likely be deployed first in countries with weak privacy regimes, especially in the developing world. When the results come back and show that the economic and health benefits are tremendous, the floodgates will open everywhere.

Our lives are being converted into data. This trend will accelerate as the tools become vastly more powerful and it will happen both voluntarily in exchange for services and due to the increasing public and private demand for security. This cannot be stopped, so what matters more than ever is watching the watchers. The amount of data we produce will continue to expand, and largely to our benefit, but we must

monitor where it goes and how it is used. Privacy is dying, so transparency must increase.

WITH ALL the attention going to massive parallel-processing beasts with specialized hardware and custom-designed chips, there was also a PC chess revolution going on. Thanks to a growing programming community being able to share ideas on the Internet and to the ever-faster CPUs coming out from Intel and AMD, personal computers running MS-DOS and Windows were becoming very strong. By 1992, they were surpassing the strength of most of the popular stand-alone chess machines, the all-in-one chess computers built into a board made by companies like Saitek and Fidelity, with names like Mephisto and even the Kasparov Advanced Trainer.

An endorsement message from me that accompanied some models in the late 1980s said, "I wish you enjoyment and satisfaction from your Kasparov chess computer—and who knows, maybe we'll meet in combat across the chessboard in the future!" I played long enough for this to come true, and more than one young player I have faced in exhibitions has brought a Kasparov chess computer with him for me to sign.

For those too young to remember, the capabilities of personal computers in the early 1990s were never enough for what you needed to do. Even spending $5,000 on a top-of-the-line machine soon left you scrambling for more RAM, more storage, and a faster CPU. And nothing taxes processing power like a chess engine. It will happily use 100 percent of the processor, and all four, ten, or twenty of a modern CPU's cores. After fifteen minutes of running an engine my old laptop would get hot enough to double as a toaster. Even today's superpowered machines can be slowed to a crawl by a chess engine grabbing every available CPU cycle for its search.

PC programs were, and are still, far slower than the specialized hardware machines like Deep Blue, often by several factors. They compensated by being much smarter, and by using optimized programming techniques to extend the search far deeper than it could get by simple exhaustive search. They are all still Type A brute force programs,

but a great deal of finesse has been added to the brute over the years. Using a multipurpose CPU allowed for more programming creativity and adaptability, and the commercial chess engines were competing constantly and tuning their evaluations, often with the help of Grandmasters. Meanwhile, although it had controller hardware that could be adjusted, Deep Thought's special chess chips were set in stone once fabricated, even though that stone was silicon.

Hardware speed depends greatly on circuit simplicity, as the Deep Thought/Deep Blue team wrote in an article about their machine in 1990. "Sacrifices in the knowledge content of the evaluation function were deemed justifiable if they simplified the circuit design significantly." They also acknowledged that "the best commercial chess programs appear to have measurably better evaluation than the research ones at this point in time." This sounds bad, but it actually gave them reason to expect greater improvements down the line when they had the opportunity to make the next generation of chips and to improve Deep Thought's evaluation function.

In 1992, I played a long casual blitz match against one of this new generation of PC programs, one that would go on to become nearly synonymous with PC chess engines. Fritz was published by ChessBase, which explains the sardonic German nickname. Its creator was a Dutchman, Frans Morsch, who had also written programs for tabletop chess machines like Mephisto. As such, he was used to having to cram tightly optimized code into very limited resources. He also helped pioneer several of the search enhancements that allowed chess machines to keep improving despite the increasing branching factor that was supposed to slow them down.

One of these is worth a brief technical detour because it's an interesting example of how machine intelligence has been augmented in ways that have nothing to do with the workings of the human mind. Called the "null move" technique, it tells the engine to "pass" for one side. That is, to evaluate a position as if one player could make two moves in a row. If the position has not improved even after moving twice, then it can be assumed that the first move is a dud and can be quickly discarded from the search tree, reducing its size and making

the search more efficient. Null moves were used in some of the earliest chess programs, including the Soviet Kaissa. It's elegant and a little ironic that algorithms designed on the principle of exhaustive search are augmented by being less exhaustive.

Humans use a very different heuristic when making plans. Strategic thinking requires setting long-term goals and establishing milestones along the way, leaving aside for the moment how your opponent, or business or political rivals, might respond. I can look at a position and think, "Wouldn't it be great if I could get my bishop over there, my pawn up there, and then work my queen around to join the attack." There are no calculations involved yet, only a type of strategic wish list. Only then do I begin to work out whether it's actually possible and what my opponent might do to counter it.

Programmers who worked on human-style or "selective search" Type B chess programs had visions of teaching machines to do this sort of goal setting. Instead of only working through the tree of possible moves, the program would also look at related hypothetical positions and evaluate those. If they were good, it would raise the values of elements in those positions in its search. It improved the quality of the evaluation in many cases, but it rendered the search so slow that results suffered, the sad tale of Type B programs in general.

More success was had with another method for allowing machines to extend their thinking into the hypothetical outside of the direct search tree. Monte Carlo tree search simulates entire games played out from positions in the search and records them as wins, draws, or losses. It stores the results and uses them to decide which positions to play out next, over and over. Playing out millions of "games within the game" like this was not particularly effective or necessary for chess, but it turned out to be essential in Go and other games where accurate evaluation is very difficult for machines. The Monte Carlo method doesn't require evaluation knowledge or hand-crafted rules; it just keeps track of the numbers and moves toward the better ones.

With so many interesting ideas to improve the output of intelligent machines, you can understand why tackling things like how the human mind works and the secrets of consciousness could fall to the

side. What matters most, the process or the results? It's always results that people want, whether it's in investing, security, or chess. As many of the programmers themselves lamented, this attitude was good for making strong chess machines and bad for anything to do with science and artificial intelligence. A chess machine that thinks like a human and loses to the world champion isn't going to make the news. And when a chess machine beats the world champion, nobody cares how it thinks.

And lose to a machine I finally did, to Fritz 3 at a blitz tournament in Munich in May 1994. The tournament was sponsored by Intel Europe, which had thrown its considerable weight behind the new Professional Chess Association (PCA) I had launched with my colleague and world championship challenger Nigel Short the year before. Along with many of the best players in the world, the event also included Fritz 3 running on a new Pentium chip. This was just the sort of promotion and sponsorship for chess that I had dreamed of when I saw all the publicity my 1989 match with Deep Thought received.

I had played quite a few games against Fritz's predecessor in an informal blitz match in Cologne in December 1992. Frederic Friedel says I played thirty-seven games against his beloved pet, as I poked and prodded it like a lab animal, pointing out when it made a particularly good move or chose a weak plan. It was far from the savage beast it would become, but it wasn't tame either. I lost nine times with a couple of draws, winning around thirty of the games.

Munich was another story. It was a serious tournament despite being blitz and one I fully expected to win, with or without the presence of a machine player. After a slow start, I scored eight straight wins, but Fritz 3 was right there with me, setting up our head-to-head encounter. I played aggressively in the opening and after just a dozen moves was rewarded with a crushing position. Then began the script that would become all too familiar to human players facing machines for the next decade. I played one lazy move and it counterattacked. Annoyed at my slip, I decided to sacrifice material, a rook for a bishop, in order to keep my initiative. The position was roughly equal, but in a blitz game I could not summon the accuracy to make good on my chances.

Despite a mutual blunder toward the end, where the machine and I both missed a chance for me to draw, Fritz 3 held on to win.

It was only blitz, with five minutes per side, but it was still the first victory over the world chess champion in a serious game by a machine. If not the moon landing, it was at least the launch of a small rocket. Fritz 3 and I ended up on top at the end, an impressive result for the machine. It was also an excellent silver lining, as I would get to meet it in a match for the title and exact my revenge. Here I managed to focus better and completely crush it, taking the play-off with three wins and two draws. I was even completely winning in one of the draws, but had no time on my clock to prosecute an easily winning position with a queen versus a rook.

Things did not work out so well for me a few months later when I met another PC program in an Intel PCA tournament, Chess Genius by Richard Lang. The London event was a rapid chess knockout event, with twenty-five minutes per side. I was paired against Genius in the very first round, which of course drew a lot of attention. It still wasn't a classical time control game, but the stakes were high. Whoever lost the two-game mini-match would be knocked out of the tournament, which was part of a Grand Prix series, so every point counted.

I got an excellent position with the white pieces in the first game, but missed a move that allowed the machine to equalize the position. It was then I committed another cardinal sin when playing against a computer: pushing too hard. Instead of acquiescing to the logical draw and moving on to the next game, I tried to keep the simplified position alive and immediately had cause to regret it. A surprising series of queen maneuvers by Genius left my king and knight in an awkward position and I ended up losing a pawn, and then the game. It was a brutal turn of events, and you can see my shock if you look up the clips from the game on YouTube.

Despite my blunder, I had every expectation of coming back and beating it with black in the next game and then to win the tie-break and move on in the tournament. I again got a very good position and this time won a pawn to enter another queen plus knight endgame. But Genius found a long series of improbable queen maneuvers that

prevented me from advancing my pawns. Head in my hands, I had to agree to a draw. I was out. It was rapid chess, yes, but a serious event and the machine had played quite well in parts. Still no moon landing, but low Earth orbit had been achieved.

Both games with Genius reflected the unique nature of computer chess, especially the second game. Chess players have the most trouble visualizing the moves of knights because their move is unlike anything else in the game, an L-shaped hop instead of a predictable straight line like the other pieces. Computers, of course, don't visualize anything at all, and so manage every piece with equal skill. I believe it was Bent Larsen, the first GM victim of a computer in tournament play, who stated that computers dropped a few hundred rating points if you eliminated their knights. This is an exaggeration, but it certainly seemed that way sometimes. There is a similar effect with the queen, by far the most powerful piece. On an open board, that is, one mostly uncluttered by pawns, the queen can reach nearly every square in just a move or two. This raises the level of complexity dramatically, something computers manage far better than humans. Facing a computer with a queen and knight in an open position near your king is a horror fit for a Stephen King novel.

For all of chess history, even the greatest players had been sheltered from the sort of incredibly complex tactical play that computers handled almost trivially by 1993. You knew that your human opponents had roughly the same limitations as you did when it came to dealing with whatever arose on the board. In my case, I always felt that I had the advantage in calculation over anyone except the Indian star Viswanathan Anand, who was justly famous for his speedy tactical play. Generally, I always knew that if I couldn't be completely sure of what the consequences of my move were going to be, my opponent couldn't be sure either. That perceived equilibrium went out the window when you were facing a strong computer. It played chess well, but also differently.

The psychological asymmetry and physical factors I've already mentioned were an issue, but the new sensation of always wondering if your opponent might be seeing something you could scarcely imagine

was very disturbing. It created a terrible tension in complex positions, a sense of dread that at any moment a shot could ring out in the dark. In response, you double- and triple-checked your calculations instead of trusting your instincts the way you would against a human opponent. All of this extra calculation cost you time on the clock and also made the games more physically taxing.

After a lifetime at the chessboard, you have no choice but to become a creature of habit, and those habits were all disrupted when playing against a machine. I didn't like it, but I also wanted to prove I could overcome these handicaps, and to prove that I was still the best chess player in the world, human or machine.

PC PROGRAMS were making impressive progress, but Deep Thought was not off my radar. I had had another close encounter with the IBM group in Copenhagen in February 1993, when the machine took on a Danish team that included Bent Larsen. IBM Denmark was eager to put their new employee to work. At this point the machine was Deep Thought II, but the IBM PR team had decided to call it Nordic Deep Blue in Copenhagen, apparently to distinguish it from the further-upgraded version they were building to challenge me at some point in the future. But I believe I'll avoid more confusion than I'll cause by simply referring to it as Deep Blue from now on.

Whatever they called it, the machine they brought to Denmark did not impress me. We used it to analyze one of my games for the audience, curious to see what suggestions it might have. Its evaluations of the game were poor, consistently underestimating my attacking chances, and it would only slowly realize that its proposed improvements didn't work. Still, it did capably well against Larsen and the other Danes for a performance rating of nearly 2600, and I was made aware that great improvements were in the IBM pipeline. The founding teammates Feng-hsiung Hsu and Murray Campbell had added Joe Hoane as a programmer, not to mention a sizable team and resources back at IBM, where the Deep Blue team would soon be moved to the company's premier research facility in Yorktown Heights, New York.

IBM had a new CEO, Lou Gerstner, who had come in during a very low point for the eighty-year-old company. IBM's stock had plummeted as the company struggled to keep up with a plethora of nimble new competitors. Among other things, Gerstner put a stop to a plan to dismember IBM into separate companies, which might have put an end to the chess project altogether.

In May 1995, I was able to get revenge on Chess Genius in a rapid match on German TV back in Cologne. I suppose it's a little silly to talk about revenge against a piece of software that may as well be counting grains of sand, but it felt good nonetheless. The first game should have ended in a draw, but Genius caught a case of that old chess machine disease, excessive greed, and came under a decisive attack against its king after it grabbed a distant pawn. The second game I drew with black without somersaults. In the interview afterward I confessed that I had been practicing at home against a version of the program in order to be as prepared as possible.

At the end of the year, I played another mini-match, this time against Fritz 4 in London. The constantly increasing version numbers were beginning to be a little intimidating, honestly. Perhaps I should have insisted on being called "Kasparov 6.0" after I won my sixth world championship match. It's not too far-fetched, considering that a PC program called "Kasparov's Gambit" came out in 1993, published by the American software giant Electronic Arts. It had a strong engine, colorful graphics, and occasionally a little video of me would pop up with basic advice on how the game was going. "Watch your pawn!" or "You're not on the right track now." It felt very cutting-edge at the time, but I'd probably laugh if I could find a working version now.

One of the interesting things about following the evolution of the PC programs from one version to the next was that I could always detect the programs' DNA, as it were. There would be new code added, new search algorithms and optimizations for a new generation of processors, but for the lack of a better word the damn things had style. I joke about the programmers treating their machines like children, or at least pets, but there is no doubt their creations take after them in some ways, and that these characteristics are passed down from one version

to the next like green eyes or red hair. The traits weaken over time as well, as you would expect in any hereditary system.

For example, Fritz was infamously materialistic, always keen to grab a pawn and hold on to it for dear life no matter how ugly its position got. This isn't to slight its programmer, Morsch, at all, but the soft-spoken Dutchman would be the first to admit his program was never one of the most aggressive on the market. Then you had the program Junior, a winner of many championships and created by the Israeli duo Shay Bushinsky and Amir Ban. It was revolutionarily aggressive, readily giving up material for open lines and attacking chances in a way that could only be described as completely uncomputerlike at the time. Is it going too far to wonder if the stolid Dutch-German program and the fiery Israeli engine had absorbed some of their stereotypical national characteristics? Well, probably so, but a program taking on the personality of its programmer is quite natural, especially if the programmer is a strong enough chess player to appreciate the stylistic qualities of his creation.

The genetic fingerprints of the different engines were also a practical matter for me and the other Grandmasters who were battling them at the board for a decade or so of competition. You couldn't expect to practice against the exact engine you were going to face in a tournament or match, but even having an older version, or at least as many of its past games as possible, made a big difference in preparing for them. As the machines accumulated a track record over years of human-machine and machine-machine play, we could prepare for them much like we would for our Grandmaster peers. There was always the problem that they could adopt completely new openings or even a new "personality" between events, or even between games, but rarely did they change completely, although they did keep getting stronger.

The two London rapid games against Fritz 4 were only memorable because of another unique aspect of playing against a computer. On my seventh move with the black pieces I played my bishop two squares, from c8 to a6, to use standard algebraic board notation. But the human operator of Fritz wasn't paying close attention and thought I had placed it one square short, on b7, and entered that move instead.

Incredibly, the game progressed for four more moves before the operator noticed his mistake. Even more incredibly, the game was actually playable when the bishop was then placed on the correct square in the computer, although of course it would have played quite differently. I won that game and then drew the second to take the match, though it was hardly satisfying after the bizarre blunder. At least Fritz wasn't capable of being annoyed at its human handler for getting it into trouble.

IN EARLY 1995, at long last, there were inquiries by David Levy and Monty Newborn about the possibility of a match against Deep Blue, probably the next year, and I told my agent Andrew Page to keep an eye on it. When I had met their team in Denmark two years earlier, I had joked that they had to hurry up and get it ready because I wanted to face it while I was still young and strong, since I was about to turn thirty at the time. And as confident as I always was in my immortality, I wouldn't be the world champion forever. IBM wanted the match and so did I; the question was whether Deep Blue would be ready.

Hsu's compulsive perfectionism with his chess chips kept pushing back deadlines, although as a fellow obsessive I can only sympathize with him. If any one small category of people did more to build the American century than any other, it was gifted engineers who had big dreams and followed them, come hell or high water. But the parts of the machine that were working were always having problems. Reading Hsu's and the many other accounts of Deep Blue's development and play during 1994–95 quickly starts to sound like the diary of someone from the Geek Squad repair company. Bugs, crashes, disconnected phone lines, interrupted Internet connections, opening book errors, more bugs, loose circuits—everything but a virus. Meanwhile, IBM still wanted the machine on the road playing in tournaments and exhibitions for its PR value.

One of these events was the 1995 World Computer Chess Championship in Hong Kong. Deep Blue Prototype—as it was called this time, though apparently still the same basic machine as Deep Thought II

since the new hardware still wasn't ready—was the big favorite. It hadn't lost to another machine in a tournament in years and, according to Hsu, it beat the top commercial programs by a three-to-one margin in their tests. (It was a big advantage that they could test against many of their competitors simply by buying a copy of the engine, while no one else could test against them.)

But as the saying goes, upsets can happen and that's why the games are played. Deep Blue drew its fourth game against a PC program named WChess and would play Fritz 3 in the fifth and final round. Deep Blue had a half-point lead and, again according to Hsu, "It won about nine out of ten games against Fritz in our pre-tournament tests back at IBM," and would have the advantage of the white pieces. Fritz played a sharp line of the Sicilian and got a fine position when apparently Deep Blue got tricked by a transposition of moves and was out of its own opening book thinking for itself.

If Deep Blue were really so much stronger than Fritz, that shouldn't be considered much of a problem. To be fair, however, the opening was indeed a difficult one that even a modern computer could have trouble navigating out of book. Deep Blue was like the junior players I criticize in my coaching sessions for blindly following opening theory and then having no understanding of the position once their memorized lines end. Still, looking at the game, it wasn't as bad as all that. A player with an estimated 200-point rating advantage shouldn't have too much trouble holding such a position.

But out comes the Geek Squad again! Deep Blue's connection between Hong Kong and New York was lost and the entire machine had to be restarted and reconnected. According to Hsu, this "cold" restart set its thinking back and it chose a different move than the one it had been considering before the disconnection occurred.

Before moving on to the exciting conclusion of this little machine versus machine drama, I want to draw attention to what just happened because it is relevant to my own encounters with Deep Blue. In nearly every Deep Blue game description I can find from this period there are resets, crashes, reboots, and disconnects. It had to forfeit a game in a Harvard competition due to a power failure. It resigned a game against

women's world champion Xie Jun in Beijing due to crashes. But this is the nature of experimental technology being assembled in haste, and rules are usually in place to handle such eventualities.

The crashes themselves don't overly concern me, but two other things about this incident do. The first is that operator intervention is required to get the machine back into the game. This wasn't just reconnecting the phone line or waiting for the Internet connection to come back online. Input was required—"We had to restart Deep Thought II," writes Hsu—and, I assume, the entire game had to be entered into the machine before telling it to start playing again. As a logical consequence, Deep Thought made a different move than the one it had favored before the crash. Hsu again: "According to Joe [Hoane], who was watching the game from our lab in Hawthorne, Deep Thought II did switch to an alternative move. But the new move never showed up on our screen in Hong Kong before the line drop, and we did not know about it until after the game."

For the sake of argument, let the assumption stand that the move Deep Blue was considering before the crash was superior to the one it played. (Looking at the game now I can say that, yes, its post-disconnect thirteenth move was indeed unlucky.) Unfortunate, of course, but what if the new move had turned out to be stronger instead of weaker? The vagaries of computer chess thinking being what they are, it's perfectly plausible that the machine might have taken more time after a reboot and found an improvement, or simply made a different move quickly that turned out to be better; who can say? The implications of this are alarming, even if you want to be charitable.

The game continued with a big advantage for Fritz. In an unfortunate attempt to defend Deep Blue's honor, Hsu's book provides commentary on the rest of the game that is complete nonsense. I may not know much about the "0.8-micron CMOS process" or the other things that make Deep Blue tick, but I still know a few things about chess. He writes about "muddling along" and "not busted yet" as if the game was competitive. In reality, although it apparently wasn't aware of it at the time, Deep Blue was completely lost after making two more terrible moves soon after the disconnection. The first blunder, its very

next move, actually went unpunished when Fritz missed a crushing response. Two moves later, already losing badly, Deep Blue committed suicide by overlooking the power of black's kingside attack. It was over. Both the 3000-rated engine on my PC and the 2800-rated engine in my skull can see at a glance that white is dead meat after Fritz's sixteenth move. With nothing to lose, Deep Blue played on down a huge amount of material before finally resigning on move thirty-nine. In a huge upset, the little German David had downed the IBM Goliath and went on to win the world championship.

I was happy for Frederic and my friends at ChessBase, but this result had the potential to be a little awkward for any match against me down the line, since Deep Blue wasn't the world computer chess champion and the next championship was likely years away. In the end, this turned out to be nothing. There wasn't really any doubt that Deep Blue was still the strongest chess machine around, especially since the version I would face in Philadelphia nine months later would finally be the upgraded one that was far stronger than the machine that lost to Fritz in Hong Kong.

Meanwhile, there was the little matter of making sure that *I* was still a world champion. My 1995 title defense was a twenty-game match against India's Viswanathan Anand. We played in New York City, on the 107th floor of the South Tower of the World Trade Center. The ceremonial first move of game one was made by Mayor Rudy Giuliani, and the date was 9/11.

I'LL SHARE some details on that human versus human event later, and how a machine helped me retain my title, but one opponent at a time. February 10, 1996, would become yet another addition to my dubious collection of This Day in History dates. Previous to sitting down against Deep Blue in Philadelphia for that first game of our six-game match, I had been the first world champion to lose a blitz game to a computer and the first world champion to lose a rapid match to a computer. The trend was clear. By the time I sat down across from Hsu for game one, I understood that eventually, if I kept my title long enough, I would also

become the first world champion to lose a classical game and match against a computer. But I was not ready for it to be today.

The match was sponsored and hosted by the Association for Computing Machinery (ACM), which had a long involvement in computer chess. They were celebrating the fiftieth anniversary of the first digital computer, the ENIAC, at their annual Computing Week event in Philadelphia. Monty Newborn, a chess programmer himself, had used his post at ACM to become an impressive promoter of human versus machine chess. As an intermediary between the parties, he helped work out the rules of the Philadelphia match, heralded as the ACM Chess Challenge. The International Computer Chess Association (ICCA) was the sanctioning body for the match and ICCA vice president David Levy assisted with the negotiations and organization. The prize fund was $500,000, with $400,000 going to the winner. The 4–1 split was a compromise after my counterproposal of a winner-take-all match instead of the 3–2 in the original proposal. I was very confident, and, after over six years of waiting since I defeated Deep Thought in 1989, it was fair to think that they needed me more than I needed them.

A few other factors made that not entirely true, however. Intel was dropping its support of my fledgling Professional Chess Association and its Grand Prix tournaments and I was hoping to establish a similar partnership with IBM. My dramatic and ill-advised breakaway from the World Chess Federation (FIDE) in 1993 had made me even more of a lightning rod in the chess world, but by bringing in new sponsors with the PCA we were organizing great events and putting money into the pockets of many players. But Intel Europe informed us that they weren't renewing the agreement. One of the reasons I played the Philadelphia match and the New York City rematch for less than the million dollars I thought they were worth was in the hope of establishing a long-term sponsorship arrangement for the PCA with IBM.

Predictions around the long-awaited match were very much in my favor. David Levy boldly predicted a 6–0 sweep for me. IBM's team leader C. J. Tan and I both predicted a 4–2 victory—him for Deep Blue and me for myself. I was confident, but worried about the lack of information available about this new version's capabilities—not the

technical specifications, which were useless to me, but what mattered to a Grandmaster's preparation: games. The version I was facing had never played publicly before, so I really had no idea what it was capable of.

Certainly the numbers were impressive. The previous model, the last officially to be called Deep Thought, searched between three and five million positions per second. This new one, with its 216 new chess chips connected to an IBM RS/6000 SP supercomputer, reached one hundred million. I knew that twenty times faster didn't mean it was twenty times better, but it was still a black box to me and that is never a pleasant feeling. According to the experts, the "speed to depth to strength" formula that had held steady in machine chess for decades might put this new version at over 2700. A better opening book and more chess knowledge might add another 50 or 100 points, approaching my 2800+ level. But this was all theoretical. And who knew what other tricks it might have up its sleeve?

Along with all the hardware and software improvements, Deep Blue had acquired an important new teammate, American Grandmaster Joel Benjamin. The debacle with the opening book in Hong Kong had convinced the IBM team that they needed professional help, so they hired a GM to prepare the opening book and to perform as Deep Blue's second during the match in case any book adjustments were needed. Benjamin would also play a role as the machine's sparring partner and in tuning its evaluation function. Even the fastest chess machine in the world needed a little human chess knowledge.

I was also taking the match seriously. I flew into Philadelphia from Rio de Janeiro, where I had just beaten a strong Brazilian team in a simultaneous exhibition. I arrived with my own second, my trainer Yuri Dokhoian. My mother, Klara, also attended, making sure all the conditions were correct in the playing hall, and was always seated in the front row. Frederic Friedel was there to serve as my unofficial computer chess advisor. Ken Thompson, the creator of Belle and who was still very involved in computer chess, agreed to be a sort of neutral overseer of the computer. Compared to the circus that the rematch would become in New York City a year later, this first match seemed almost

quaint. A considerable media presence built up as the match drew more attention, with journalists from most major print publications and even regular coverage on CNN. But it still felt relatively casual and open there in the giant convention hall. With ACM and the ICCA running the show, IBM had a relatively discreet presence, usually through the team leader, C. J. Tan. It all felt very much like any other top-level chess match, right up until the moment I sat down at the board with Deep Blue for the first time.

I'VE HAD twenty years to come up with a good way to describe what it's like for a world champion chess player to play against a world-champion-level chess machine. I'm still not sure I've succeeded. Directly competing against a computer at the highest level of a human discipline is a unique experience. It's not a video game against a computer AI or a metaphorical competition in the job market, the "race against" or "race with" the machines explained so capably by MIT's Erik Brynjolfsson and Andrew McAfee in their books.

John Henry competed with a steel-driving steam engine before a crowd of witnesses, his muscle and bone versus the implacable iron beast. Jesse Owens's races against cars and motorcycles also boasted that same tragicomic asymmetry; it was exploitation and entertainment, not serious competition. If a person wins a footrace against a car, it's funny. If he loses, what else could anyone expect?

The other difference was apparent from the popular news coverage, which echoed centuries of romantic notions about chess and intelligence, and misconceptions about artificial intelligence and Deep Blue. "The Brain's Last Stand," "Kasparov Defends Humanity," "The machines are entering the last human refuge, intelligence." Even the jokes about the match on shows like Jay Leno and David Letterman had a nervous, slightly apocalyptic feel to them. "Kasparov looks pretty nervous. You may think this is no big deal, but wait until that thing comes for *your* job!" "He's playing chess against a supercomputer and I still can't program my VCR!" "In a related story, earlier today the New York Mets were defeated by a microwave oven."

For the most part, the more flattering narratives were indulged by the organizers—and by the participants, I admit. Who was I to say that chess was not the "pinnacle of human intellectual activity"? Or that I wasn't a "living Mount Everest" or potentially a "chess champ guilty of letting down the whole human race"? And IBM had no incentive to disagree with any assumptions about its machine's "creativity" or "potential to revolutionize entire industries." ACM's Monty Newborn was in his element. He is a natural raconteur who never lets his background in computer science and chess get in the way of a P. T. Barnum sensibility. I didn't have much time for such things at the time, but now even I am almost inspired by listening to Newborn talk in interviews at the match about "what it means to be human" and likening a potential Deep Blue win to the moon landing.

Finally, all the hype and mythologizing could be put aside and the first game could begin. At least, it could after another bug was squashed by the operator. To my amazement, Deep Blue wasn't running yet when the arbiter started its clock, and it took Hsu, the operator that day, a few minutes to get it going. It may sound petty to speak of such things as distractions, but of course they are. It's difficult enough to summon your usual focus under such strange conditions, especially when you know your opponent has no such concerns. The scrum of photographers around the table doesn't annoy a computer. There is no looking into your opponent's eyes to read his mood, or seeing if his hand hesitates a little above the clock, indicating a lack of confidence in his choice. As a believer in chess as a form of psychological, not just intellectual, warfare, playing against something with no psyche was troubling from the start.

After a few moments Deep Blue was running and Hsu made its move, 1.e4. That's moving the pawn in front of the king forward two squares, and I answered with 1..c5, my favorite Sicilian Defense, a sharp counterattacking opening. Don't worry, I'm not planning on presenting the whole game! It's one of the most famous games in chess history and there is plenty of analysis of it available if you are interested. Unfortunately, it's not a very good game, as I rediscovered when I went over these matches again. To help me maintain objectivity, a couple of

strong players in Moscow also went over them with the best chess engines currently available. I did not play very well in Philadelphia, even if I did play well enough.

Deep Blue declined my challenge of an open game, somewhat surprisingly since computers thrive in the complex tactical positions that the Open Sicilian is known for. The IBM team was worried about running into an opening surprise I might have cooked up and clearly didn't feel like it was wise to match Joel Benjamin's preparation against mine in a risky variation. Instead it played the same second move it had used against me with white in our 1989 match, although of course they wouldn't expect me to repeat that game despite the result. Trying to repeat your past victories without having improvements prepared in your own play is a very good way of walking into a preparation land mine. They instead made a solid choice and the machine was well prepared, reaching the ninth move still in its opening book.

I was also prepared, and varied from a previous game of mine on the tenth move with an improvement. I wasn't going to go into a defensive crouch. I wanted to see what this thing could do. This wasn't a blitz or rapid game; we had hours on our clocks, not just a few minutes. This gave me enough time to think, so I wasn't afraid of entering sharp complications. Deep Blue played well in the early stages, gaining a slight advantage typical of the white pieces. After I made an inaccuracy it played several more strong moves to create real threats for the first time. I glanced up at Hsu, a habit rendered pointless in this match. My position was deteriorating. This thing was strong. This was different.

READING THROUGH some of the dozens of books and hundreds of articles written about this match, and this game in particular, you would think the authors had all attended different events and analyzed a different game. Disagreements in analysis are normal, of course, and healthy. When someday chess has been completely solved by some technology we cannot even imagine now, we will be able to speak of objective truth on the board. Until then, we will have regular

disagreements about the quality of some moves. Different Grandmasters and different machines will prefer different ideas that may be equally strong, which is what makes chess interesting.

This isn't to say that some moves aren't simply blunders or inaccuracies, or that often there isn't a clearly best move on the board. In many positions the right move is obvious and will be made by any reasonably strong player. Perhaps 10 to 15 percent of positions require a master's experience or calculation skills to find a sophisticated plan or complicated tactic. Then there is that last 1 or 2 percent of moves, the very difficult ones that even strong Grandmasters might miss. Under such conditions, with the stress of competition and the pressure of the clock, it's remarkable that humans play chess as well as we do. In fact, I discovered that it is often the case that we perform better under pressure, not worse.

During my work on the *My Great Predecessors* book series, I gained not just a deeper respect for the achievements of the past world champions I was studying, but a greater admiration for chess players in general. Few activities are as taxing to the human faculties as a game of professional chess. Rapid calculation is essential, adrenaline is surging, and the outcome hangs on every move. This goes on for hour after hour, day after day, often with the whole world watching. It is the ideal scenario for a mental and physical meltdown.

When I began to analyze the games of my world champion forebears, I was therefore prepared to be a little forgiving. Not in my analysis, where I had to be as merciless as my teacher Botvinnik had instructed me, but in the tone I adopted toward their mistakes. Here I was in the twenty-first century, with databases of millions of games and gigahertz of chess engine processing power at my fingertips. "With these advantages and the benefit of hindsight I shouldn't judge my predecessors too harshly," I told myself.

An important part of the project was to collect all the relevant analysis that had been done on these games before, especially the published analyses of the players themselves and their contemporaries. My colleague Dmitry Plisetsky did a phenomenal job of tracking down sources in a half-dozen languages. One would assume that the analyst,

working in the calm of his study and with unlimited time to move the pieces and make notes, would have a much easier job than the players themselves. Hindsight is 20/20, is it not? But one of my first discoveries was that when it comes to chess analysis in the precomputer age, hindsight was badly in need of bifocals.

Paradoxically, when other top players wrote about games in magazines and newspaper columns they often made more mistakes in their commentary than the players had made at the board. Even when the players themselves published analyses of their own games they were often less accurate than when they were playing the game. Strong moves were called errors, weak moves were praised. It was not only a few cases of journalists who were lousy players failing to comprehend the genius of the champions, or everyone missing a spectacular move that I could easily find with the help of an engine, although that did happen regularly. The biggest problem was that even the players would fall into the trap of seeing each game of chess as a story, a coherent narrative with a beginning and a middle and a finish, with a few twists and turns along the way. And, of course, a moral at the end of the story.

I took two lessons away from this discovery. The first is that we often do our best thinking under pressure. Our senses are heightened and our intuition is activated in a way that is unique to stress and competition. I would still rather have fifteen minutes on my clock than fifteen seconds to make a critical move, but the fact remains that our minds can perform remarkable feats under duress. We often do not realize how powerful our intuitive abilities are until we have no choice but to rely on them.

The second lesson was that everyone loves a good story, even if it flies in the face of objective analysis. We love it when the most annoying character in the movie finally gets what he deserves. We root for underdogs, cringe at a hero's downfall, and sympathize with the unlucky victim of the Fates. All these tropes are in play in a chess game, just as they are in an election or the rise and fall of a business, and they feed the powerful cognitive fallacy of seeking a narrative where often none exists.

Computer analysis exploded this lazy tradition of analyzing chess games as if they were fairy tales. Engines don't care about story. They expose the reality that the only story in a chess game is each individual move, weak or strong. This isn't nearly as fun or interesting as the narrative method, but it's the truth, and not just in chess. The human need to understand things as a story instead of as a series of discrete events can lead to many flawed conclusions. We are easily drawn away from the data by a nice anecdote that fits our preconceived notions or that fulfills one of the popular tropes. This is how urban legends propagate so efficiently; the best ones tell us something we really want to believe is true. I'm certainly not immune to this tendency myself, and it's impossible to overcome all our intellectual biases. But becoming aware of them is a good first step, and one of the many benefits of human-machine collaboration is helping us overcome lazy cognitive habits.

With all that in mind, let's return to the board, where I was getting into real trouble in my first game against Deep Blue. The machine had played several surprising moves and I was accumulating weaknesses in my position. Looking over the analysis of others and listening to the commentary that was being provided live by several Grandmasters (and Fritz 4!), the tendency toward narrative has overwhelmed objectivity once again. The consensus seems to be that I made the fatal mistake of counterattacking a computer in an open position where its unmatched tactical abilities would be overwhelming instead of trying to consolidate and sit tight. Perhaps this is true, but it was not my intention to play to the computer's strength. I simply did not see a better choice.

After my 1989 victory over Deep Thought, I was interviewed by the *New York Times* for a lengthy magazine article. We were looking over the news coverage the match had received and one quote from Deep Thought team member Murray Campbell caught my eye. "Deep Thought didn't get a chance to show what it can do," he said. "That's exactly the point!" I shouted to the interviewer. "I didn't let it! The

highest art of the chess player lies in not allowing your opponent to show you what he can do."

Seven years later, Deep Blue was proving too strong to so easily dictate terms to, especially since it had the white pieces. And while my choice to attack its king can be criticized as ill advised against a machine, it was not a bad move, and certainly not the losing move. That would come two moves later when, ironically, I held up my attack to preserve a pawn. Had I continued to play as aggressively as all the commentators criticized me for doing, I might have saved the game. But that would go against the popular narrative, so the losing move is often overlooked.

What I overlooked, on the other hand, has been correctly diagnosed. Deep Blue grabbed a pawn far from the action in what appears to be a terrible loss of time with its king under attack. But, in the time-honored tradition of human-machine chess, it had calculated deeply enough to get away with it. Despite what I've said about the dangers of narrative, I cannot resist sharing this passage on the game from Charles Krauthammer's story on the match for *TIME* magazine. This sort of storytelling I completely endorse.

> Late in the game, Blue's king was under savage attack by Kasparov. Any human player under such assault by a world champion would be staring at his own king trying to figure out how to get away. Instead, Blue ignored the threat and quite nonchalantly went hunting for lowly pawns at the other end of the board. In fact, at the point of maximum peril, Blue expended two moves—many have died giving Kasparov even one—to snap one pawn. It was as if, at Gettysburg, General Meade had sent his soldiers out for a bit of apple picking moments before Pickett's charge because he had calculated that they could get back to their positions with a half-second to spare.
>
> In humans, that is called sangfroid. And if you don't have any sang, you can be very froid. But then again if Meade had known absolutely—by calculating the precise trajectories of all the bullets and all the bayonets and all the cannons in Pickett's division—the time of arrival of the enemy, he could indeed, without fear, have ordered his men to pick apples.

Which is exactly what Deep Blue did. It had calculated every possible combination of Kasparov's available moves and determined with absolute certainty that it could return from its pawn-picking expedition and destroy Kasparov exactly one move before Kasparov could destroy it. Which it did.

It takes more than nerves of steel to do that. It takes a silicon brain. No human can achieve absolute certainty because no human can be sure to have seen everything. Deep Blue can.

I held out my hand to resign on move thirty-seven and a computer had defeated the world chess champion in a classical game for the first time in history. I was a bit in shock, as were the spectators and commentators. Even Hsu, who would have been aware of Deep Blue's winning evaluation from his screen, looked a bit confused, almost apologetic on the moment of his great triumph. I honestly feel a little bad about that now, despite the bad blood that would arise out of the rematch a year later. I'm sure he wanted to jump up and down with his teammates, not answer my questions.

Still in a mild daze at how well the machine had played, I asked a reflexive question immediately after resigning, the way two Grandmasters might begin what we call the "postmortem" of a completed game. "Where did I go wrong?" I asked. But Hsu wasn't much of a chess player and, probably a bit dazed himself, he couldn't recall enough of Deep Blue's analysis on the screen to answer, so it was a slightly awkward moment for both of us.

A month after the match, I wrote in *TIME* that I felt I could sense "a new kind of intelligence across the table" that day, and in a way, it was true. I wasn't suggesting any metaphysical interpretation, but could sheer speed really produce such impressive chess? Several of its moves were almost as if it was saying, "I bet you didn't think a computer could make a move like this!" For example, at one point in the middlegame it sacrificed a pawn for activity, a very humanlike idea not at all in keeping with the usual machine materialism.

It was the best I had ever seen a machine play, against me or anyone else and, at least at the moment of my loss, I even considered the

possibility that it might be too strong to beat. Later that day, I wondered aloud to Frederic, "What if this thing is invincible?" I had known that day would arrive eventually; was it here already?

I didn't have to wait very long for the answer. In game two the next day, I played a slow, maneuvering opening with white. The idea was to not provide Deep Blue with any clear targets, knowing it couldn't formulate strategic plans the way a human could. At least, I hoped it couldn't. As usual, there were a few technical problems, although I was only aware of one of them at the time. Deep Blue played a poor move very early, on move six. According to Frederic, I was visibly delighted by what I could only assume was a major flaw in the machine's opening book. Not only wasn't it invincible, I was going to have an easy day of it. You can imagine my disappointment when the arbiter ran over to say that Hsu had accidentally made the wrong move on the board, capturing the wrong pawn, as had happened in my London match with Fritz. The rules allowed them to correct it and the game proceeded along normal lines. It all turned out fine, but it illustrated the danger of having a weak player making the moves and how distractions like this only affected the human player.

Hsu's book blames Murray Campbell for failing to properly upload the updated opening book file he and Benjamin had worked on after game one to the machine back in Yorktown Heights. This left it to rely on something he calls an "extended book," which had vague guidelines based on database statistics from Grandmaster games. Regardless, I was oblivious to this and Deep Blue played the opening just fine, following high-level Grandmaster theory until I introduced a new idea on move fourteen. Several books also mention an "evaluation bug" in Deep Blue that affected its play in this game, but honestly I tire of trying to figure out which bugs are "bugs" and which "bugs" are bugs, and which are just lousy evaluations.

My strategy worked out quite well and Deep Blue was saddled with the sort of long-term structural weakness it had no idea how to defend. I realized that just avoiding wild tactical positions wasn't enough. I should aim for positions in which general principles would outweigh short-term calculations. Deep Blue did have evaluation functions, but

it was not very sophisticated and something I could exploit once I became aware of its hard-coded preferences. For example, if I noticed it had been set to keep queens on the board—generally a good idea for a machine against a human—I could play moves that offered it the choice of exchanging queens or making an inferior move.

This sort of human adaptation was one of the reasons some computer scientists thought that chess machines wouldn't defeat Grandmasters for far longer than turned out to be the case. Once a human figures out the rules and knowledge that govern a machine's play, they thought, they would figure out how to exploit them. But it turned out that with super-fast brute force, little exploitable knowledge was needed and most weaknesses were amply covered up by sheer depth of search.

Deep Blue hadn't achieved perfection just yet, however. In game two, I offered a pawn sacrifice that it couldn't resist and in compensation the light squares around its king were fatally weakened. It was close to reaching a draw, but the best lines were always a little too deep for its search and it didn't know the general principles of how to defend such positions. After hours of careful maneuvering I won one pawn and then another and Murray Campbell resigned for Deep Blue on move seventy-three. I had leveled the score and, more importantly, I knew it was only mortal.

Now that I knew that the "new kind of intelligence across the table" was only a much faster version of the computer programs I understood well, I relaxed a bit. It was very strong, yes, but it wasn't stronger than I was and it had clear deficiencies. As with a human opponent, if I could target its weaknesses and avoid its strengths, I would win the match.

Game three repeated the opening from the first game until Deep Blue deviated with a move inserted into its book that day by Benjamin. We continued along the line he planned until move eighteen, when Deep Blue noticed that the line Benjamin had intended, but lucky for him, not inserted into the book, actually lost a piece. This left me with a small advantage and a clear target to focus on, so I thought my chances were good for a second consecutive win. But Deep Blue started to do what machines are known for, impossibly tenacious defending, harder

to kill than a real bug, a cockroach. If there is only one move to save the position, they always find it. Much to my frustration, Deep Blue found a long sequence of resourceful moves to escape danger and draw the game.

Precision under fire is another aspect of human versus machine asymmetry. What we call a "sharp" position in chess is when there is high complexity and grave consequences for any error. Both players are balancing on a tightrope and the first slip can be fatal. For a computer, this actually makes it easier to find the right path because all the other moves return a very low score. Humans can never enjoy such confidence. What's more, only the human player is aware that there is a tightrope. I can sense danger in a position, feel the tree of variations growing exponentially. That's just another day at the beach for a machine, especially one that has, as Deep Blue did, special search extensions that added extra depth in consequential variations.

The match was level after three of the six games, but I had white in two of the final three and was feeling more comfortable. The media attention for the match had grown tremendously after Deep Blue's win in game one, but of course the machine didn't have to give interviews. I disappointed my computer expert by ignoring his advice and opening the position in game four. I didn't shy away from playing sharply with white. I pondered a piece sacrifice against Deep Blue's kingside on move thirteen for a while before deciding it was simply too risky. It's notable, however, that I would have played it against any other chess-playing entity on the planet, man or machine. I knew that if I made the slightest miscalculation in such a position I was dead and would be behind in the match with only two games to go. It was an important moment, in retrospect. I wasn't just playing chess, I was making specific adjustments to playing against a machine whose capabilities in certain areas far exceeded mine or anyone else's.

There was yet another technical snafu during game four, and it came exactly when I was preparing a dangerous attack. I had spent a long time on my previous move, planning to sacrifice a knight for two pawns and an attack. Before Deep Blue replied, it crashed and had to be restarted. I was furious, ripped out of my state of deep concentration

at a key moment in the game. It took twenty minutes to get it working again and when it came back it played a strong move that avoided my sacrifice. It was enough to make me wonder if something more than bugs was going on. (Subsequent analysis shows the sacrifice would likely have led to a roughly equal position.)

The position was now balanced, but sharp, and I was approaching time trouble. If I reached move forty, more time would be added to the clocks; the question was whether or not I would make it. After making several precise moves I got to the safe shore of the time control at move forty with a defensible position. I found a nice way to force a drawn position and the game was soon over. The score was still level with two games to go and I was exhausted. Attendance at the match had continued to climb and the media attention was approaching a frenzy. There were interviews and TV appearances for both teams, and IBM definitely noticed that their little chess project was getting more attention than just about anything else they had done in years.

Despite the rest day between games four and five, I had trouble mustering my energy. I avoided my usual Sicilian for the Russian Defense, a.k.a. the Petroff. This wasn't a display of patriotism; the Petroff is very solid, some would say boring. It often leads to many piece exchanges and symmetrical pawn structures that reduce the dynamism in the position, something I thought sounded ideal when tired and facing a supercomputer, even though it wasn't the sort of position I usually played. Deep Blue transposed instead into a Four Knights Opening, which was no more or less as dull as the Petroff.

After many exchanges took place I had the tiniest of advantages. Thinking of saving my energy for the final game with white the next day, I offered an early draw on move twenty-three. For those new to the chess world, the idea of offering a draw must sound very strange. Imagine two boxers simply agreeing to stop fighting in the second round, or a soccer match ending after fifteen minutes because the coaches decide a tie is a good result. Typically, until rules were put in place to discourage it, in chess either player can offer her opponent a draw after any move. The other player can then think about it and accept, or ignore the offer and make a move, and the game continues.

to kill than a real bug, a cockroach. If there is only one move to save the position, they always find it. Much to my frustration, Deep Blue found a long sequence of resourceful moves to escape danger and draw the game.

Precision under fire is another aspect of human versus machine asymmetry. What we call a "sharp" position in chess is when there is high complexity and grave consequences for any error. Both players are balancing on a tightrope and the first slip can be fatal. For a computer, this actually makes it easier to find the right path because all the other moves return a very low score. Humans can never enjoy such confidence. What's more, only the human player is aware that there is a tightrope. I can sense danger in a position, feel the tree of variations growing exponentially. That's just another day at the beach for a machine, especially one that has, as Deep Blue did, special search extensions that added extra depth in consequential variations.

The match was level after three of the six games, but I had white in two of the final three and was feeling more comfortable. The media attention for the match had grown tremendously after Deep Blue's win in game one, but of course the machine didn't have to give interviews. I disappointed my computer expert by ignoring his advice and opening the position in game four. I didn't shy away from playing sharply with white. I pondered a piece sacrifice against Deep Blue's kingside on move thirteen for a while before deciding it was simply too risky. It's notable, however, that I would have played it against any other chess-playing entity on the planet, man or machine. I knew that if I made the slightest miscalculation in such a position I was dead and would be behind in the match with only two games to go. It was an important moment, in retrospect. I wasn't just playing chess, I was making specific adjustments to playing against a machine whose capabilities in certain areas far exceeded mine or anyone else's.

There was yet another technical snafu during game four, and it came exactly when I was preparing a dangerous attack. I had spent a long time on my previous move, planning to sacrifice a knight for two pawns and an attack. Before Deep Blue replied, it crashed and had to be restarted. I was furious, ripped out of my state of deep concentration

at a key moment in the game. It took twenty minutes to get it working again and when it came back it played a strong move that avoided my sacrifice. It was enough to make me wonder if something more than bugs was going on. (Subsequent analysis shows the sacrifice would likely have led to a roughly equal position.)

The position was now balanced, but sharp, and I was approaching time trouble. If I reached move forty, more time would be added to the clocks; the question was whether or not I would make it. After making several precise moves I got to the safe shore of the time control at move forty with a defensible position. I found a nice way to force a drawn position and the game was soon over. The score was still level with two games to go and I was exhausted. Attendance at the match had continued to climb and the media attention was approaching a frenzy. There were interviews and TV appearances for both teams, and IBM definitely noticed that their little chess project was getting more attention than just about anything else they had done in years.

Despite the rest day between games four and five, I had trouble mustering my energy. I avoided my usual Sicilian for the Russian Defense, a.k.a. the Petroff. This wasn't a display of patriotism; the Petroff is very solid, some would say boring. It often leads to many piece exchanges and symmetrical pawn structures that reduce the dynamism in the position, something I thought sounded ideal when tired and facing a supercomputer, even though it wasn't the sort of position I usually played. Deep Blue transposed instead into a Four Knights Opening, which was no more or less as dull as the Petroff.

After many exchanges took place I had the tiniest of advantages. Thinking of saving my energy for the final game with white the next day, I offered an early draw on move twenty-three. For those new to the chess world, the idea of offering a draw must sound very strange. Imagine two boxers simply agreeing to stop fighting in the second round, or a soccer match ending after fifteen minutes because the coaches decide a tie is a good result. Typically, until rules were put in place to discourage it, in chess either player can offer her opponent a draw after any move. The other player can then think about it and accept, or ignore the offer and make a move, and the game continues.

Draws have always been a part of chess, at least in the modern history of the game. There are many positions that cannot be won by either side, including stalemate, in which the side to move has no legal moves and so the game is drawn. Draws are worth half a point for both players, so it's definitely better to draw than to lose and get nothing. The draw offer was created as a courtesy so strong players would not have to wear themselves out playing tedious and obviously equal positions all the way down to nothing. It was a way of saying, "I know you know how to draw this and you know that I know, so let's shake hands and retire to the smoking room." It may have disappointed some spectators to end the game early, but there weren't usually very many spectators to worry about. Additionally, back in the nineteenth century, the level of play was relatively low and almost all the games ended decisively.

The problem began when masters began to exploit the draw offer strategically, or even tactically. If a draw suited you in the tournament standings, why not see if your opponent would also like a short day's work and offer an early draw? Or if you felt that your position was deteriorating, perhaps offer a draw and see what your opponent thought about it? Soon enough, it became something close to a plague, with perfunctory games as short as a few minutes and a dozen moves, even between strong Grandmasters. The habit was contagious and today it's not unusual to see short draws even at the weak amateur level.

Eventually, organizers of top tournaments decided they no longer wanted to support such behavior and instituted rules like move minimums. Now it's fairly standard to have events where it is not permitted to offer a draw before move thirty or forty, although little can be done about draws by repetition of position. With players becoming stronger and more accurate decade by decade, the number of draws has increased at the top level, with roughly half the games at elite events finishing drawn. I don't see this as a problem as long as they are fighting games—a draw is a fair result. But there are regular pushes to introduce more rule changes to encourage more aggressive play and produce more decisive games, such as awarding three points for a win and one for a draw, as is used in many professional soccer and hockey leagues.

In match play, short draws can be strategically useful. I was still feeling exhausted in game five and also felt that there wasn't a great deal to play for in the position when I offered the early draw. It would have been a disappointment for the seven hundred or so spectators that day, however, so it was their good fortune that the Deep Blue team declined my offer and decided to play on. As an aside, this is another unique aspect of machine play, when to offer or accept draws. Should the decision be left up to the machine somehow? For example, if its evaluation is at zero or worse, should it automatically accept? But what if it is in a must-win situation? As with opening books, it's a case where there isn't a very good solution to what amounts to human intervention.

Deep Blue thought it was a little worse at the time of my draw offer. The team huddled and eventually followed Benjamin's recommendation that it was too early to end the game, especially since they would have black in the final game. This turned out to be my good fortune as well, as Deep Blue's next move was a serious mistake. Unable to see the long-term consequences, it stepped right into a pin that would tie up its pieces for a long time as I advanced my pawns. With no active plan, and not understanding that its only hope was to lash out, Deep Blue shuffled around for several moves. By the time the danger was close enough to reach its search horizon, it was too late to save itself. I won in forty-five moves to take the lead in the match for the first time, and was guaranteed at least a tie in the match going into the final game the next day.

I was feeling good headed into game six despite my tiredness. I had outplayed the machine in game five and felt like I was getting to know its weaknesses. This was probably an overestimation on my part after only five games, but I knew a lot more than I had a week earlier, and it would all come together in game six. We repeated the first few moves from my first two whites until Deep Blue varied. Being behind in the match, their team had the task of trying to find a way to play for a win with black, and it wasn't going to be easy. I could go entire calendar years without losing a game with the white pieces despite my aggressive style, and here I only needed a draw to win the match, and the $400,000 winner's check, so I wasn't going to take any unnecessary risks.

After my transposition of moves got Deep Blue out of its opening book it began to play weakly, and it fell into a passive position. Without its book, it didn't know, as a Grandmaster would, that certain pieces just belong on certain squares in certain openings. This is exactly the sort of generalized, analogous thinking that humans use all the time. Without it, Deep Blue had to rely on its search to keep it out of trouble, but its options were dwindling. I shoved my queenside pawns forward, driving its pieces back. It was exactly the sort of control game I had dreamed of: closed instead of open, strategic instead of tactical. I could smell blood, or whatever it had.

At move twenty-two I considered a tempting piece sacrifice against its king that looked winning. But could I be sure? Ninety percent sure, yes. Ninety-five percent, maybe. But against Deep Blue, and needing only a draw to win the match, I would have to be 100 percent sure. Analysis later showed that it was indeed a winning blow, although there is no way to guarantee I would have played it perfectly. And there was no reason for me to take any risks, since I was crushing it already. Black had no counterplay and my pawns were still on the march. The audience got quite excited when they understood what was happening. Deep Blue was being suffocated, its bishop and rook trapped on its first rank. At the end, black's pieces were so tied up that I didn't even have to break through. The machine was out of moves that didn't lose material and the Deep Blue team decided it was time to resign.

I had won the match 4–2, exactly the score I had predicted, but it had also been far tougher than I imagined it would be. I praised the Deep Blue team for their achievement. Beyond the score, it could occasionally play chess of a quality I never believed a computer could play. I adapted my strategy and won the last two games quite easily, which may not have been good for my mindset going into the rematch. I concluded my match article in *TIME*:

> In the end, that may have been my biggest advantage: I could figure out its priorities and adjust my play. It couldn't do the same to me. So although I think I did see some signs of intelligence, it's a weird kind, an inefficient, inflexible kind that makes me think I have a few years left.

In fact, I had exactly 450 days, until the end of the rematch on May 11, 1997. Looking back, I was the last world champion to win a match against a computer. Why don't those This Day in History calendars have a page for that!?

DESPITE BEGINNING with very little publicity, the first Deep Blue match became the largest Internet event in history at the time. IBM had to assign a supercomputer like the one that ran Deep Blue to handle the load on the website—and this was in 1996 when most people were on dial-up connections. It became an early example of the power of the new communications network, showing how the Internet might one day compete with television and radio. Imagine.

The Deep Blue team obviously weren't happy with the result of the match or the way the last game went in particular, but they said they were satisfied. They had beaten the world champion and made me sweat quite a bit in the first four games. Meanwhile, IBM was even happier than I was. The winner's check they gave me was nothing compared to what the match publicity had done for IBM's stock price and the company's image. Suddenly, stodgy old IBM was cool, on the cutting edge of artificial intelligence and supercomputing, battling for supremacy against the human mind. At least that's how it looked, and the stock market seemed to agree.

According to Monty Newborn's book on the match, IBM's stock rose an equivalent of $3,310 million in little more than a week, a week that the rest of the Dow Jones went down significantly. I should have demanded stock options instead of a 4–1 prize split! Deep Blue's name was everywhere in the media and the IBM team and the IBM brand went with it. It was good for me too, of course, especially in America where chess champions were hardly household names. I was getting more US media attention for beating Deep Blue in Philadelphia than I had for beating Anand in a world championship match in New York City. It turned out that even world champions are outranked by defenders of humanity.

The PR bonanza virtually guaranteed a rematch; the question was when. There was no way the Deep Blue team would want to play again until substantial improvements could be made. How long would it take for them to get a new version ready that was strong enough to be more of a threat? Because, as the negotiations went on, one thing became very clear: if there was a rematch, it wasn't going to be because the Deep Blue team wanted to improve, or because Garry Kasparov wanted another paycheck. It would because IBM wanted to win.

CHAPTER 8

DEEPE

contra their it'

KEN THOMPSON designed the revolutionary chess machine Belle, whose chips Deep Blue's were based on, while working at Bell Laboratories in New Jersey, the famous "Idea Factory" that did pioneering work on breakthroughs in everything from solar cells and lasers to transistors and cell phones. Thompson was also the principal inventor of the ubiquitous Unix operating system while there, which is the basis for what runs Apple Macs, Google Android, and the billions of devices and servers running Linux.

As with the early years of ARPA, the concept at Bell Labs was to describe big problems and then work on creating the technology to solve them, instead of starting with a specific product in mind. I heard similar stories when I was invited to speak at General Electric's new Innovation Center near Detroit in 2010. My hosts were eager to stimulate the sort of "blue sky" thinking that had gone out of fashion after decades of industry consolidation and acquisitions. During my seminar, someone pointed out that too often giant companies assume that even if they aren't innovating, somebody is somewhere, and when something good comes along they will simply buy it. You can see how it eventually becomes a problem when everybody thinks somebody else will innovate.

I was reminded of this particular seminar in the context of chess machines because of a slide I used with a quote from Alan Perlis, a computer science pioneer and the first recipient of the Turing Award, in 1966, awarded by ACM. In a famous list of epigrams about programming that he published in 1982, Perlis wrote, "Optimization hinders evolution." This jumped out at me because it at first sounds