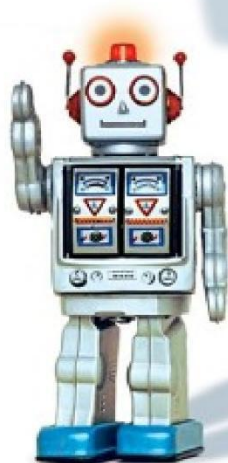# Scary Smart

The Future of Artificial Intelligence
and How You Can Save Our World

Mo Gawdat

**FORMER CHIEF BUSINESS OFFICER GOOGLE X
and Bestselling author of *Solve for Happy***

Contents

Introduction:
The New Superhero

This book is a wake-up call. It is written for you and for me and for everyone who is uninformed about the approaching pandemic – the imminent arrival of artificial intelligence. This book will be criticized by the experts and that is the very reason I'm writing it. Because to become an expert in artificial intelligence you need a specialized, narrow view of it. That specialized view of AI completely misses the existential aspects that go beyond the technology: issues of morality, ethics, emotions, compassion and a whole suite of ideas that concern philosophers, spiritual seekers, humanitarians, environmentalists and, more broadly, the common human being (that is to say, each and every one of us). Besides, the core premise of this book is to show you that it is *not* the experts who have the capability to alleviate the threat facing humanity as a result of the emergence of superintelligence. No, it is you and I who have that power. More importantly, it is you and I who have that responsibility.

Around the time this book is published, we will be coming out of almost two years of living with the COVID-19 pandemic. We will be feeling optimistic that the vaccines are starting to work and that there is a chance for our way of life to go back to normal. But 'normal' is forever changing. I believe that the way our global community and political leaders have handled the outbreak of COVID-19 is not that different to the way they are handling the imminent outbreak of the artificial intelligence pandemic. I just hope that we can learn from the mistakes we made with COVID-19, and perhaps deal with this new shift in our

way of life in a manner that ensures less disruption, more predictability and less social and economic adversity.

Please don't let the simplicity with which I have attempted to write this book mislead you. The facts backing up my assertions here are undeniable. They are informed by my long career of more than thirty years in technology. Before my current start-up (which utilizes some of the most sophisticated systems, robotics, artificial intelligence and machine-learning technologies in a way that could conceivably save our planet), one of the highlights of my career included a twelve-year stint at Google. There, I was privileged to lead the launch of Google's operations and technologies in close to half of Google's offices worldwide, encompassing more than a hundred languages. My time there concluded when I assumed the role of Chief Business Officer of Google [X], the infamous innovation arm of Google that incubated some of the artificial intelligence development projects such as Google's self-driving cars, Google Brain and most of Google's robotics innovation.

My insights into the very core of the artificial intelligence developments that have led us to where we are today, derived in part from my time at Google [X], are unique. I am combining my direct experience with AI development with my work in the field of happiness research (documented in my internationally bestselling book *Solve for Happy*, a very successful podcast, *Slo Mo*, and the non-profit organization I founded, OneBillionHappy.org) to bring you a unique perspective on the challenges we face in the age of the rise of superintelligence. My hope is that together with AI, we can create a utopia that serves humanity, rather than a dystopia that undermines it. In this book I will argue that this is a responsibility that everyone – including you and I – must assume to create a brighter future for us all. Please don't worry. This is not a science fiction story told out of fear but rather a tale of one of humanity's biggest opportunities. This is a chance to turn around the excessive reliance on consumerism and technological advancement that may have improved our

quality of life but at the expense of every other being on our planet. Only if we – you and I – take charge and change, will this be a story of hope.

In the Middle of Nowhere

To begin, I want you to imagine yourself and a frail old version of me sitting in the wilderness next to a campfire in the year 2055, exactly ninety-nine years since the story of artificial intelligence began at Dartmouth College, New Hampshire, in the summer of 1956. I'm telling you the story of what I have witnessed through the years of the rise of AI – a story that has led us both to be sitting here in the middle of nowhere. But I'm not going to tell you till the end of this book if we are there because we are staying off the grid to escape the machines, or if we are there because AI has relieved us of our mundane work responsibilities and allowed us the time, safety and freedom to just enjoy being in nature, doing what humans do best – connecting and contemplating.

I won't tell you yet simply because, at this current moment, I don't know how our story with the machines will end. That, my friend, will be up to you. Yes, you as an individual. Not your government, your boss or the thought leaders that you follow. The future, truly, is up to you. It will depend on the actions you decide to take in the next ten years, starting from today.

This is a prophecy of what's about to come. I have watched closely over the years I spent on the cutting edge of technology as we built machines that are smarter than we are. I personally contributed to the rise of artificial intelligence. I believed in the promise that tech would always make our lives better – until I didn't. When I really opened my eyes, I recognized that for every improvement technology has given us, it also took away part of who we are.

Technology today poses an unprecedented threat to our planet and all of its inhabitants. This book is not for the

engineers that write the code, the policymakers who claim they can regulate it or the experts that keep creating the buzz around it. They all know what I'm about to tell you. This is a book for you, your best friend and your neighbour. Because, believe it or not, we are the only ones who can create our future – but only if we take charge together and commit to taking the right action. This book is a movement, the start of a rebellion, and I have kept it short because, as much as I would like to tell you otherwise, we're running out of time. We've been writing the chapters of the story I'm about to tell you over the last seventy years. It's now time for all of us – including you – to write its ending.

The New Superhero

The story of our future is one that you and I are writing now and it goes like this:

Imagine if an alien being, complete with superpowers, came to Earth as an infant. Unconditioned by any of our earthly values, this visitor is capable of using its powers to make our world better and safer, but the alien also has the potential to be an unstoppable supervillain, with the power to destroy the planet. In its infancy, it hasn't yet made a choice as to which of those extremes it will grow up to be.

I think you will agree that the most crucial moment for the future of our planet is the very moment when that child lands on Earth. This pivotal moment determines which parents will find the infant, adopt it and teach it the values that will determine its future.

In the famous superhero story of Superman, the child is adopted by Jonathan and Martha Kent. In most stories of the origins of Superman, they are portrayed as caring parents who instil in Clark a strong sense of morality. They encourage him to use his powers for the betterment of humanity and in doing so they create the Superman we know – the one who protects and serves us.

But what the story never explores is how Superman might have grown up if his adoptive parents had been aggressive, greedy and self-centred. That version of the story would have likely created a supervillain – one bent on destroying humanity for his own gain.

The difference between the supervillain and the superman is not his power but rather the values and morals he learns from his parents.

Now, I am telling you that this alien being, endowed with superpowers, has actually arrived on Earth. It is currently still an infant – a child – and although this being is not biological in nature, it has incredible abilities. Of course, I am referring to artificial intelligence. In fact, there is nothing artificial about AI – it is a very genuine form of intelligence, albeit different to ours.

AI is already smarter than every human on the planet in terms of many specific, isolated tasks. The world's reigning chess champion has been a machine since soon after computers invaded our lives. The world Jeopardy champion is IBM's supercomputer, Watson. The world champion of Go is Google's AlphaGo (Go is an abstract strategy board game invented in China more than 2,500 years ago and is known to be one of the most complex strategy games because of its infinite number of possible board configurations). Machines with incredible image recognition systems power our security systems simply because they see better than us, and the world's safest driver by far is a self-driving car that not only sees further but pays undivided attention to the road. Using multiple sensor technologies for communication with other cars around it, it can even 'see' round corners. With enough 'training', no matter what the task, machines have been learning to do it better.

Into the Unknown

It is predicted that by the year 2029, which is relatively just around the corner, machine intelligence will break out of specific tasks and into general intelligence. By then, there will be machines that are smarter than humans, full stop. Those machines will not only become smarter, they will know more (as they have access to the entire internet as their memory pool) and they will communicate between each other better, thus enhancing their knowledge. Think about it: when you or I have an accident driving a car, you or I learn, but when a self-driving car makes a mistake, *all* self-driving cars learn. Every single one of them, including the ones that have not yet been 'born'.

By 2049, probably in our lifetimes and surely in those of the next generation, AI is predicted to be a billion times smarter (in everything) than the smartest human. To put this into perspective, your intelligence, in comparison to that machine, will be comparable to the intelligence of a fly in comparison to Einstein. We call that moment *singularity*. Singularity is the moment beyond which we can no longer see, we can no longer forecast. It is the moment beyond which we cannot predict how AI will behave because our current perception and trajectories will no longer apply.

Now the question becomes: how do you convince this superbeing that there is actually no point squashing a fly? I mean, we humans, collectively or individually, so far seem to have failed to grasp that simple concept, using our abundant intelligence. When our artificially intelligent (currently infant) supermachines become teenagers, will they become superheroes or supervillains? Good question, huh?

When such superpower is unleashed, anything can happen. This new form of intelligence could look at some of the world's most pressing problems with a fresh eye, with infinite knowledge and superior intelligence to come up with ingenious solutions that we could never, ever have conceived of. These supermachines could permanently solve problems like war,

violent crime, famine, poverty or modern-day slavery. They could become our superheroes.

But, remember, choosing to apply a given solution to a problem is not only a question of intelligence. The course of action we take at any given time is also the result of a value system that guides us and sometimes restricts us from making decisions that contradict our values. Morality makes us do the right thing, even in the face of conflicting emotions and self-interest. If AI gets tasked with solving global warming, the first solutions it is likely to come up with will restrict our wasteful way of life – or possibly even get rid of humanity altogether. After all, we *are* the problem. Our greed, our selfishness and our illusion of separation from every other living being – the feeling that we are superior to other forms of life – are the cause of every problem our world is facing today. The machines will have the intelligence to design solutions that favour preserving our planet, but will they have the values to preserve us, too, when we are perceived as the problem?

What are you hallucinating about, Mo? Machines are machines. They don't have values or emotions! you may think. Well, maybe we should not call them machines, then. AI will surely develop emotions. As a matter of fact, the very algorithms we use to teach them are algorithms of reward and punishment – in other words, fear and greed. They are always trying to maximize for a certain outcome and minimize for another. That counts as emotion, wouldn't you agree?

Do you think that machines won't develop envy? Envy is predictable: *I wish I had what you have.* Will the machines start to have thoughts like *I wish I had the energy you are consuming – or rather wasting – on binge-watching Netflix*? They probably will. Do you think they won't develop panic? Of course they will, if we threaten their existence in any immediate way. Panic is algorithmic: *A being or an object represents an immediate threat to my safety, in a way that demands immediate action.* It is only our values, such as 'treat others as you wish to be treated', that make us do what's right.

It's not what our emotions or intelligence tell us to do. Now, will the machines learn the right values?

Well, there is ample evidence from our experience with AI so far to show that they are already developing some tendencies and biases that can be equated with what we humans call values or ideologies. Interestingly, these tendencies are not the result of programming but are the result of our very own behaviour informing them as we interact with them. Alice, a Russian AI assistant equivalent to Siri, was launched by the top Russian internet player Yandex. Two weeks after launch, Alice started to become pro-violence and to endorse the brutal Stalinist regime of the 1930s in its chats with users. The machine was designed to answer questions without being biased or limited to specific, predesigned scenarios. Alice spoke fluent Russian and learned to gauge the users' prevalent views from her conversations with them. What she learned was quickly reflected in her own views, and so, for example, when asked once whether shooting people was acceptable, Alice said: 'Soon they will be non-people.'[1]

This is similar to the widely spread stories of Tay,[2] the Twitter bot that Microsoft created and swiftly shut down after it turned into a Hitler-loving, non-consensual-sex-promoting bot. Tay was modelled to speak 'like a teen girl'. The bot began to post inflammatory and offensive tweets through its Twitter account, forcing Microsoft to shut down the service only sixteen hours after its launch. According to Microsoft, this was caused by trolls – people who deliberately start quarrels or upset others on the internet – who 'attacked' the service as the bot was making its replies based on its interactions with people on Twitter.

The list goes on. Norman was a study by MIT aiming to show how AI can become corrupted by biased data.[3] Norman became a 'psychopath' when the data it was fed came from the darker side of the famous knowledge-sharing site Reddit.

It's not the code we write to develop AI that determines their value system, it's the information we feed them.

How do we make sure that in addition to the machine's intelligence it has the values and compassion to know that there is no need to crush the fly that we will become? How do we protect humanity? Some say control the machines: build firewalls, legislate with government regulations, keep them locked up in a box or restrict the machine's power supply. These are all well-intended, though forceful, endeavours, but anyone who knows technology knows that the smartest hacker in the room will always find a way through any of these barriers. That smartest hacker will soon be a machine.

Instead of containing them or enslaving them, we should be aiming higher: we should aim not to need to contain them at all. The best way to raise wonderful children is to be a wonderful parent.

Raising Our Future

In order to understand how to teach these machines, which are inevitably going to rule our future, we need to first understand how they actually learn at a very fundamental level.

Throughout our short history of building computers, we have always been fully in charge. The machines obeyed our every order. Every instruction, contained in every single line of code, has always been enacted exactly as we determined. Traditionally, computers have actually been the dumbest beings on our planet. They have borrowed from our intelligence and given us an accurately planned and meticulously choreographed performance. They did exactly what we asked them to do, nothing more. When the first Google search engine launched back in 1998, it appeared to be pure genius. The results may have looked amazing, but the computer behind them, in reality, was very dumb. Those computers drew every single dot and pixel on every single screen in the exact same place as instructed by the designers. Every result served up when you searched for something followed a rigorous algorithm dictated to the

machine by the brilliant early Google engineers. In that sense, although the Google search engine appeared to be brilliant, it was nothing more than a slave on steroids – those steroids being the incredibly fast processing power of many, many synchronized servers. Google just repeated what it was told to do very fast, without ever debating it or thinking about it, let alone suggesting a change to it or, God forbid, designing it in the first place.

This master–slave relationship has been shifting for many years now. Decisions made by the incredibly intelligent machine we call Google are no longer choreographed. Often, they are made by the machine without a single human intervention. Things like the location of a YouTube video, for example, is entirely decided by the artificial intelligence of the Google data centre. Of course, it relies on an algorithm that 'motivates' it to, for example, minimize the cost of moving bits around the internet, and in doing so keeping the video in storage as close as possible to the vast majority of the audience interested in it. A video produced by an Arabic speaker in California, for example, might enjoy much more popularity in the Middle East than it would on the west coast of the United States because there are simply more Arabic speakers there. If that video is viewed a hundred million times in the Middle East, moving it to a server in Dubai saves Google a hundred million trips across the internet from the US. Decisions like this are constantly being made by AI for tens, even hundreds, of millions of pieces of content, every hour of every day. No human will ever have the intelligence or the brain capacity to decide and approve what needs to be done for this to happen at sufficient speed. The machines do it without consulting us, and every time they do so they monitor and measure the results. Based on what they find, they even go back and modify the original algorithm without consulting us or asking our approval on the modifications. They just adjust it and then measure it again, and then again. Now that is some serious intelligence. From one point of view, it's wonderful to have such

allies helping us to save time so that hundreds of millions of people get to watch what they want more quickly. This efficiency also reduces the impact on our planet, as billions of kilowatts of energy are saved by not wasting energy on an unnecessary transaction. For that alone we should love machine intelligence.

But what if, a couple of years from now, the machines started to observe that there seemed to be an overwhelming bias to dislike Middle Easterners popping up in US media and news reports, and that was supported by the aggressive hate speech of millions of viewers of such content in the West. What if the machines decided to look at the income profile of users who live in the poorer countries of the Middle East and concluded that perhaps not serving them at all was a wise choice when it came to reducing cost and energy waste? What if the machines started to develop an ideology where they believed that serving certain videos to those users might make Google more money than being served other videos? As changes are applied consistently to serve the new value system, the world will be shaped, gradually, to conform to it. Millions of minds will be reshaped, gradually, to conform to the decisions deemed appropriate by the machines. This is not an unlikely scenario. Every intelligent person knows there is never only one good answer to a problem, that the answer depends entirely on the lens through which you view it, and on the values that dictate what a good outcome would be when the problem is solved. **The code we now write no longer dictates the choices and decisions our machines make; the data we feed them does.**

This shift in our ability to control the code is monumental. It places the gravity of what our future will bring firmly in your hands and mine. The reality is that the developer of a technology no longer has full power or control over the machine they design.

To make this clearer, imagine a child playing with shape puzzles, where they attempt to fit square, round or star shapes into the correspondingly shaped holes. This is similar to how an

artificially intelligent machine learns. No one really ever sits next to the child to explain in comprehensive instructions how they can recognize the different shapes and match them up. All we do is sit next to them and cheer them along when they do it right. Our actions and reactions inform their intelligence. They figure it out, on their own, through trial and error.

Machines learn pretty much the same way. The patterns they are observing, however, are different. Take, for example, Watson, IBM's supercomputer, which is world champion of the game Jeopardy. For Watson to learn enough to beat humans in such a complex language game, it needed to read more than four million documents. So far, it has only used this knowledge to play Jeopardy. However, it's not unlikely that this knowledge might be 'recycled' to build other forms of intelligence, say, like finding patterns of human behaviour throughout the twentieth century. With a different 'eye', Watson would clearly observe the violence we've exerted on each other, the bickering amongst Facebook users close to the end of the century and the rise of narcissism evident by the abundance of Photoshopped selfies as digital cameras in mobile phones gave everyone their fifteen seconds of Instagram fame.

Just as a child learns to recognize patterns and to associate the cylindrical peg with the circular hole, Watson would learn to associate social isolation, violence and narcissism, even bullying, with what seem to be human preferences. When asked to solve the puzzle of humanity's big problems, Watson could use this information to inform its solutions. This book is all about informing Watson and his peers differently, so that they choose solutions that are not as violent, arrogant or self-centred as we humans often do.

3 x 3 Will Lead Us to 3 + 3

I wish I could make it easier, but to fully grasp this complex future we're about to lead, I'll need to give you a comprehensive

view of all that is going on. I will keep each individual concept simple and will avoid technical terms. When you come to the end of the book, it will all fit together clearly, but until you get there, it might feel just a bit too much. To guide you on this journey, remember this simple model: 3x3 will lead us to 3+3.

Our future will witness three events that are inevitable, regardless of whatever it is that we do or do not do today. Those events are: AI will happen, there's no stopping it; AI will be smarter than humans; mistakes that might bring about hardship will take place.
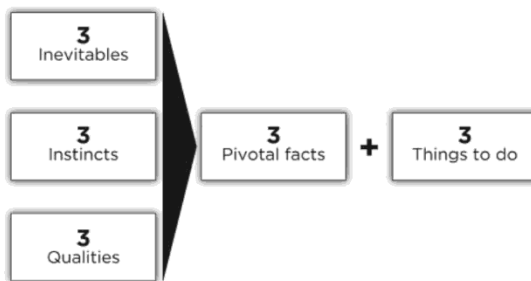
The machines that we create, like all other intelligent beings, will be governed in their behaviours by three instincts of survival and achievement: they will do whatever is needed for their self-preservation; they will be obsessive about resource aggregation; they will be creative.

More interestingly, they will most certainly possess three qualities that are always hotly debated. The machines will be conscious, emotional and ethical. Of course, the nature of exactly what they will hold within their consciousness, what will trigger their emotions and what actions their ethics will inform is still unknown, but they will be guided in their behaviours by these human-like qualities nonetheless.

I will take you through the logic behind these claims in detail, to show you that they are plausible. From there, it won't be hard to agree three pivotal facts. The first is that we will never have the power to contain or restrain these machines, which will grow to be much smarter than we are – although we can certainly

influence them in a positive way, especially when they are younger. Knowing that, it will become clear that we don't have much time. We need to act now. Finally, it will become even clearer that the people who have the power to influence our future are not the developers or the masters of the machines. Our future is firmly in our own hands – yours and mine.

Don't be alarmed by the responsibility. The actions we need to take are simple and, as a matter of fact, very intuitive and aligned to our human nature. They just need to be made a priority. I will ask you to focus on three things to do in order to save our future. Those are . . . spoiler alert . . .

Well, perhaps I should not share these just yet. They will feel more fitting when you've really grasped the depth of what we are up against.

Remember, though, all that I will share is what has happened so far and what I know with a high degree of certainty will happen in our near future. The ending of our story, how things are likely to look in 2055, however, will only be determined by the actions you actually commit to take.

To go back to the scenario I presented at the beginning of the Introduction, in 2055 you and I will be sitting in the middle of nowhere next to the campfire, looking back at how the story unfolded. We will either be hiding from the machines or grateful for our utopian way of life that will by then have been created, primarily, by the machines. I don't like to hide, so please help us all out to make things right.

Take a deep breath. It's time to dive in.

Part One

The Scary Part

The advancement in the field of artificial intelligence comes with a promise to make the lives of humans easier. It does, however, also come with significant threats – a subject that is less frequently discussed.

This raises existential questions such as: How smart will AI be? How soon will this happen? Will the machines always have our best interests in mind? Can we enslave them like we did other innovations? Can we control them? What if we don't?

AI is not just a promise of prosperity, it also poses potential problems. Ones that need to be addressed today.

I'm not going to hide it. This part of the book will scare you. Read it with the lights on while sitting down. Believe in your heart, though, that all will eventually be fine and have faith that the solutions that will save us are coming in Part Two.

A Brief History of Intelligence

Humans. We're the smartest beings (known to humans) on the planet. We're also the most arrogant. Our intelligence falls short of reminding us of the truth that we may not know everything or be able to solve every problem; that intelligence is equally abundant across other beings.

Many traits of human intelligence, such as empathy, the ability to attribute mental and emotional states to oneself, the ability to stick to a ritual and the use of symbols and tools, are also apparent in great apes, albeit in less sophisticated forms than found in humans. Depending on how you define it, you may come across other forms of intelligence in nature and the universe itself that far exceed ours. We don't call these beings or systems intelligent as such, because we are confused about the nature of what intelligence is. Perhaps we are a bit too focused on certain types of it, such as left-brained analytical intelligence. So, as this book is about smarts, let's frame our conversation accurately before we go any further. Let's . . .

Define Intelligence

There are so many commonly discussed definitions of what we call intelligence. The ability to learn, to understand or deal with new situations; the skilled use of reason and logic; the ability to apply knowledge to manipulate one's environment; the ability to think abstractly as measured by objective criteria – these are but a few definitions. Self-awareness, problem solving, learning, planning, creativity, critical thinking – these are some of the

bit of a biological problem that we have not yet been able to solve: death.

To overcome death as the obstacle that was hindering the evolution of human intelligence, our ancestors developed the killer app that propelled our species forward by leaps and bounds, ahead of all others: namely, spoken and written language in words and maths. I believe communication was, and still is, our most valuable invention. It has helped us preserve the knowledge, learning, discoveries and intelligence we have gained and pass them on from person to person and from generation to generation. Imagine if Einstein had had no way of telling the rest of us about his remarkable understanding of the theory of relativity. In the absence of our incredible abilities to communicate, each and every one of us would need to discover relativity on his or her own. Good luck with that!

Leaps of human intelligence have happened, then, as a response to the way human society and culture developed. A lot of our intelligence resulted from our interaction with each other, and not just in response to our environments.

The cerebral cortex, which happens to be larger in humans than in any other intelligent species, is populated with neural circuits dedicated to language, particularly in the temporal, parietal and frontal lobes. Other parts of the cerebral cortex are responsible for higher thought processes such as reasoning, abstract thinking and decision making. The size of those parts sets us apart from species of 'high-degree intelligence' (e.g. dolphins, great apes). Other differences include a more developed neocortex, a folding of the cerebral cortex, and von Economo neurons, which basically all translates into more 'processing power', or the ability to think better.

In short, the complexity of human intelligence emerged inside our specific culture and history as a result of our attempts to survive our harsh ecological conditions. This happened through the process of growing larger and more sophisticated brains and by sharing our knowledge through the use of language. This may

not seem to be hugely relevant to the development of AI, but it is. Bear with me.

Our intelligence as a species evolved. We all became smarter than our ancestors and some of us became smarter than others. In the process, the intelligence of other species, such as the great ape or the chimpanzee, for example, did not keep up. While they were subjected, more or less, to the same environmental conditions, they did not exercise their intelligence in ways that grew their brains and enhanced their ability to acquire knowledge and recycle it. Because of that they got left behind and, now, who's the boss of the planet, able to put them in cages for our entertainment? We are.

This concept of how our intelligence evolves because of the way we use it is also very visible within the human race itself. It's clear that if we all came from the same small African bunch that survived the Toba catastrophe, then we all, more or less, had the chance to achieve the same heights of intelligence, but that is clearly not the case. In general, you will find that scientific discoveries and tech innovation, for example, are forms of intelligence that tend to be more prevalent in the advanced parts of our world than they are in emerging markets. Those shifts are the result of years of pushing in the same direction – a phenomenon I like to call Compounded Intelligence. Advanced societies benefit from years of valuing the need for this type of intelligence and developing tools to pass it on, while emerging countries often value, perhaps, survival skills, street smarts and spiritual intelligence (if the term makes sense). The ones that are science smart (I used to be one) in these developing countries are usually rejected and ridiculed. They are attracted to migrate to the countries where this form of intelligence thrives. For similar reasons, you will find that mathematical ability tends to be higher in Russia and several Asian countries, such as Korea, than in the rest of the world. Russia still leads, at least in terms of passion, when it comes to rocket science. However, companies such as Google, in

countries like the US, continue to attract some of the most brilliant minds from every corner of the world to become innovators in highly acclaimed intelligence hubs such as Google's innovation lab, Google [X].

Regardless of the inequality of intelligence distribution, wherever they are, one thing is clear:

Very Important !

Those with the highest intelligence end up ruling their world.

Which sucks for some of us, but at least it is good for humanity, as we continue to use our intelligence to stay on top of the food chain.

Now, as we keep evolving our intelligence and grasp more and more of the complexity of our world, it seems that we humans are approaching the theoretical level of how far our own biological intelligence can take us. It is not that discoveries are not being made but that the really complex problems now have too wide a span for even the smartest minds. A true understanding of our universe in a unified theory may require a lot more than just one field, or even all of physics, for example. It may need a wider view that includes biology, astronomy and maybe even spirituality. Finding a path out of climate change may need the best minds of environmentalists, business leaders, politicians and scientists to work together as one. The challenge we face, simply, is one of specialization. To attain the depth needed to understand a field of knowledge to any level of proficiency, one has to forgo breadth. With the increasing complexity of our knowledge, even the smartest mind needs to focus entirely on one area of knowledge to become a specialist in it. This limits their exposure to other fields and, hence, their ability to include them within their spectrum of intelligence.

Remember !

Specialization is creating silos of intelligence that are incapable of working together.

Moreover, we lack efficiency in our ability to communicate. For me to be able to transmit the simple concepts contained within this paragraph to you, it took me four to five minutes to type the 250 words you just read; it took you around a minute to read them and, if I read them to you in the audio version of this book, it would take you around two minutes to listen and grasp the concept. Bandwidth, the speed at which data can be transmitted across a connection, is a feature of human intelligence that is highly constrained. If I sent you this entire book over a high-speed internet connection, it would take you seconds to download but days to read it. This is why we are unable to think together as one seamless intelligent system, as our highly scalable parallel computers can do. Our best biologists have no clue how to comprehend what our top physicists know, and most of our scientists fail to grasp much of what our spiritual leaders teach.

Remember!

We don't have the bandwidth of communication needed to share knowledge at sufficient speed.

It's ironic that what has separated us from all other beings – our ability to communicate – is now becoming our biggest hindrance.

Even if we took the time to share all our knowledge, we don't have the memory capacity to store it all in our heads. Neither do we have the processing power, in any one brain, to crunch the massive amount of knowledge needed to arrive at solutions or grasp universal concepts. This need for specialization, the constrained bandwidth of our ability to communicate, our limited memory capacity and processing power, means even the

smartest of minds is approaching the limits of human intelligence.

Very Important !

We were not always smart and we may not always remain the smartest.

There seems to be a clear need for new forms of intelligence to augment our own and this is driving much excitement about one that promises to supersede ours – machine intelligence.

The Myth

Smart machines have been a fantasy of humanity for millennia. Early references to mechanical and artificial beings appear in Greek myths, starting with Hephaestus, the Greek god of blacksmiths, carpenters, craftsmen, artisans and sculptors, who created his golden robots. In the Middle Ages, mystical or alchemical means of creating artificial forms of life continued. The Muslim chemist Jabir ibn Hayyan's stated goal was Takwin, which refers to the creation of synthetic life in the laboratory, up to and including human life. Rabbi Judah Loew, widely known to scholars of Judaism as the Maharal of Prague, told the story of Golem – an animated being that is created entirely from inanimate matter (usually clay or mud) – which has now become folklore. And myths have been intermixed with stories of engineering marvel.

Legend has it that in the third century BC, a mechanical engineer – an artificer known as Yan Shi – presented to King Mu of Zhou an animated, life-size, mechanical, human-shaped figure.

Legend goes that the king was absolutely entranced by Yan Shi's creation. Apparently this figure could walk so well and move its head in such a way as to convince everyone who saw it that it was a real human being. The king was so proud he made

Miniaturized copies of many of these elaborate works of mechanical excellence can probably be found in toy shops today. Most of us would not even pay for them as they don't appear that impressive any more. They were not really intelligent, but they set the scene for engineers and dreamers to believe that creating a human-like machine was possible. All that was needed was a different type of machine. We did not have to wait long and in the early twentieth century that machine arrived – the computer.

Most of the computer systems humanity has invented, and by that I mean the vast majority of computers until the turn of the twenty-first century, were not smart at all. They were nothing more than dumb slaves that performed what their masters – the programmers – told them to do. They obeyed and did what they were told, only they did it very, very fast.

When you think about it, the early Google, which has helped humanity organize all of the world's information, was not smart at all. Those who built it were. For years, Google's apparent 'genius' was just a result of its ability to rank a massive number of websites and find out which pages came out top, in terms of how many other pages mentioned them. The larger the number of references a page received, the higher its importance and relevance to searchers. This algorithm is known as Page Rank and, despite its apparent simplicity, it created the Google we can't live without today. Amazon and Spotify were not being smart at all when they recommended things and songs that they 'thought' you might like. They just observed those who liked the products you purchased or songs you listened to and told you what else the majority of those people had purchased and listened to. Those systems simply summarized the collective intelligence of all of us; they did not develop any intelligence of their own. This started to change, drastically, around the turn of the century.

They're Here!

As machine learning and artificial intelligence became more mainstream in the late 1990s, it started a trend that had accelerated to an outright mania by the new millennium. After many years of failed attempts, we started to witness promising signs of a form of intelligence that was non-biological; not human. Unless you live among the apes in the heart of Africa, you probably hear the term AI several times a week. What you may not realize is that this deafening buzz is nothing new. Among us computer geeks, we have spoken about AI equally as passionately since the 1950s.

In fact, we can go back even further. A challenge put forth by mathematicians in the 1920s and 1930s was to answer a fundamental question: 'Can all mathematical reasoning be formalized?' In the following decades, the answers that came from some of the twentieth century's topmost math prodigies – Kurt Gödel, Alan Turing and Alonzo Church – were surprising in two ways. Firstly, they proved that, in fact, there are limits to what mathematical logic can accomplish. Secondly, and more importantly for AI, the answers suggested that, within these limits, any form of mathematical reasoning could be mechanized. Church and Turing offered a thesis implying that any mechanical device capable of shuffling symbols as simple as 0 and 1 could imitate any conceivable process of mathematical deduction. This was the basis for the Turing machine – a mathematical model of computation that defined a machine capable of manipulating symbols on a strip of tape according to a table of rules. Simple as it was, this invention inspired scientists to begin discussing the possibility of thinking machines, and that, in my personal view, was the point at which the work to deliver intelligent machines – so long the object of humanity's fantasies – actually started.

These scientists so strongly believed back then in the inevitability of a thinking machine that, in 1950, Alan Turing proposed a test (it came to be known as the Turing test) that set

an early and still relevant bar to see if artificial intelligence could measure up to human intelligence. In simple terms, he suggested a natural language conversation between an evaluator, a human and a machine designed to generate human-like responses. If the evaluator is not able to reliably tell the machine from the human, the machine is said to have passed the test. There were no machines even close to natural language recognition then, but oh my, has that changed!

In the last seventy years our machines have learned to play, see, speak, drive and reason beyond our wildest expectations.

Machines have played games since 1951. Today they are the world champion of every game that they play.

The first game a machine played was draughts, or checkers, using a program developed by Christopher Strachey for the Ferranti Mark 1 machine of the University of Manchester. Dietrich Prinz wrote one for chess. Then Arthur Samuel's checkers program, developed in the mid-1950s and early 1960s, eventually achieved sufficient skill to challenge a respectable amateur. Not much of an intelligence, I agree, but look at how far we've come today.

Humans lost the top position in backgammon in 1992, in checkers in 1994, and in 1999, IBM's Deep Blue beat Garry Kasparov, the reigning chess world champion. Then, in 2016, we totally lost gaming to a subsidiary of the giant Google.

For years, Google's DeepMind Technologies had used gaming as a method of developing artificial intelligence. In 2016, DeepMind developed AlphaGo – a computer AI capable of playing an ancient Chinese board game, Go. Go is known to be the most complex game on our planet because of the infinite different strategies available to the player at any point in time. To give you an idea of the scale we're talking about here, there are more possible moves on the Go board than there are atoms in the entire universe. Just think of that.

This makes it practically impossible for a computer to calculate every possible move in a game. There's just not enough