Deloitte.

Seeing the forest for the trees, and the forests beyond

The future of AI

A report by the Deloitte AI Institute

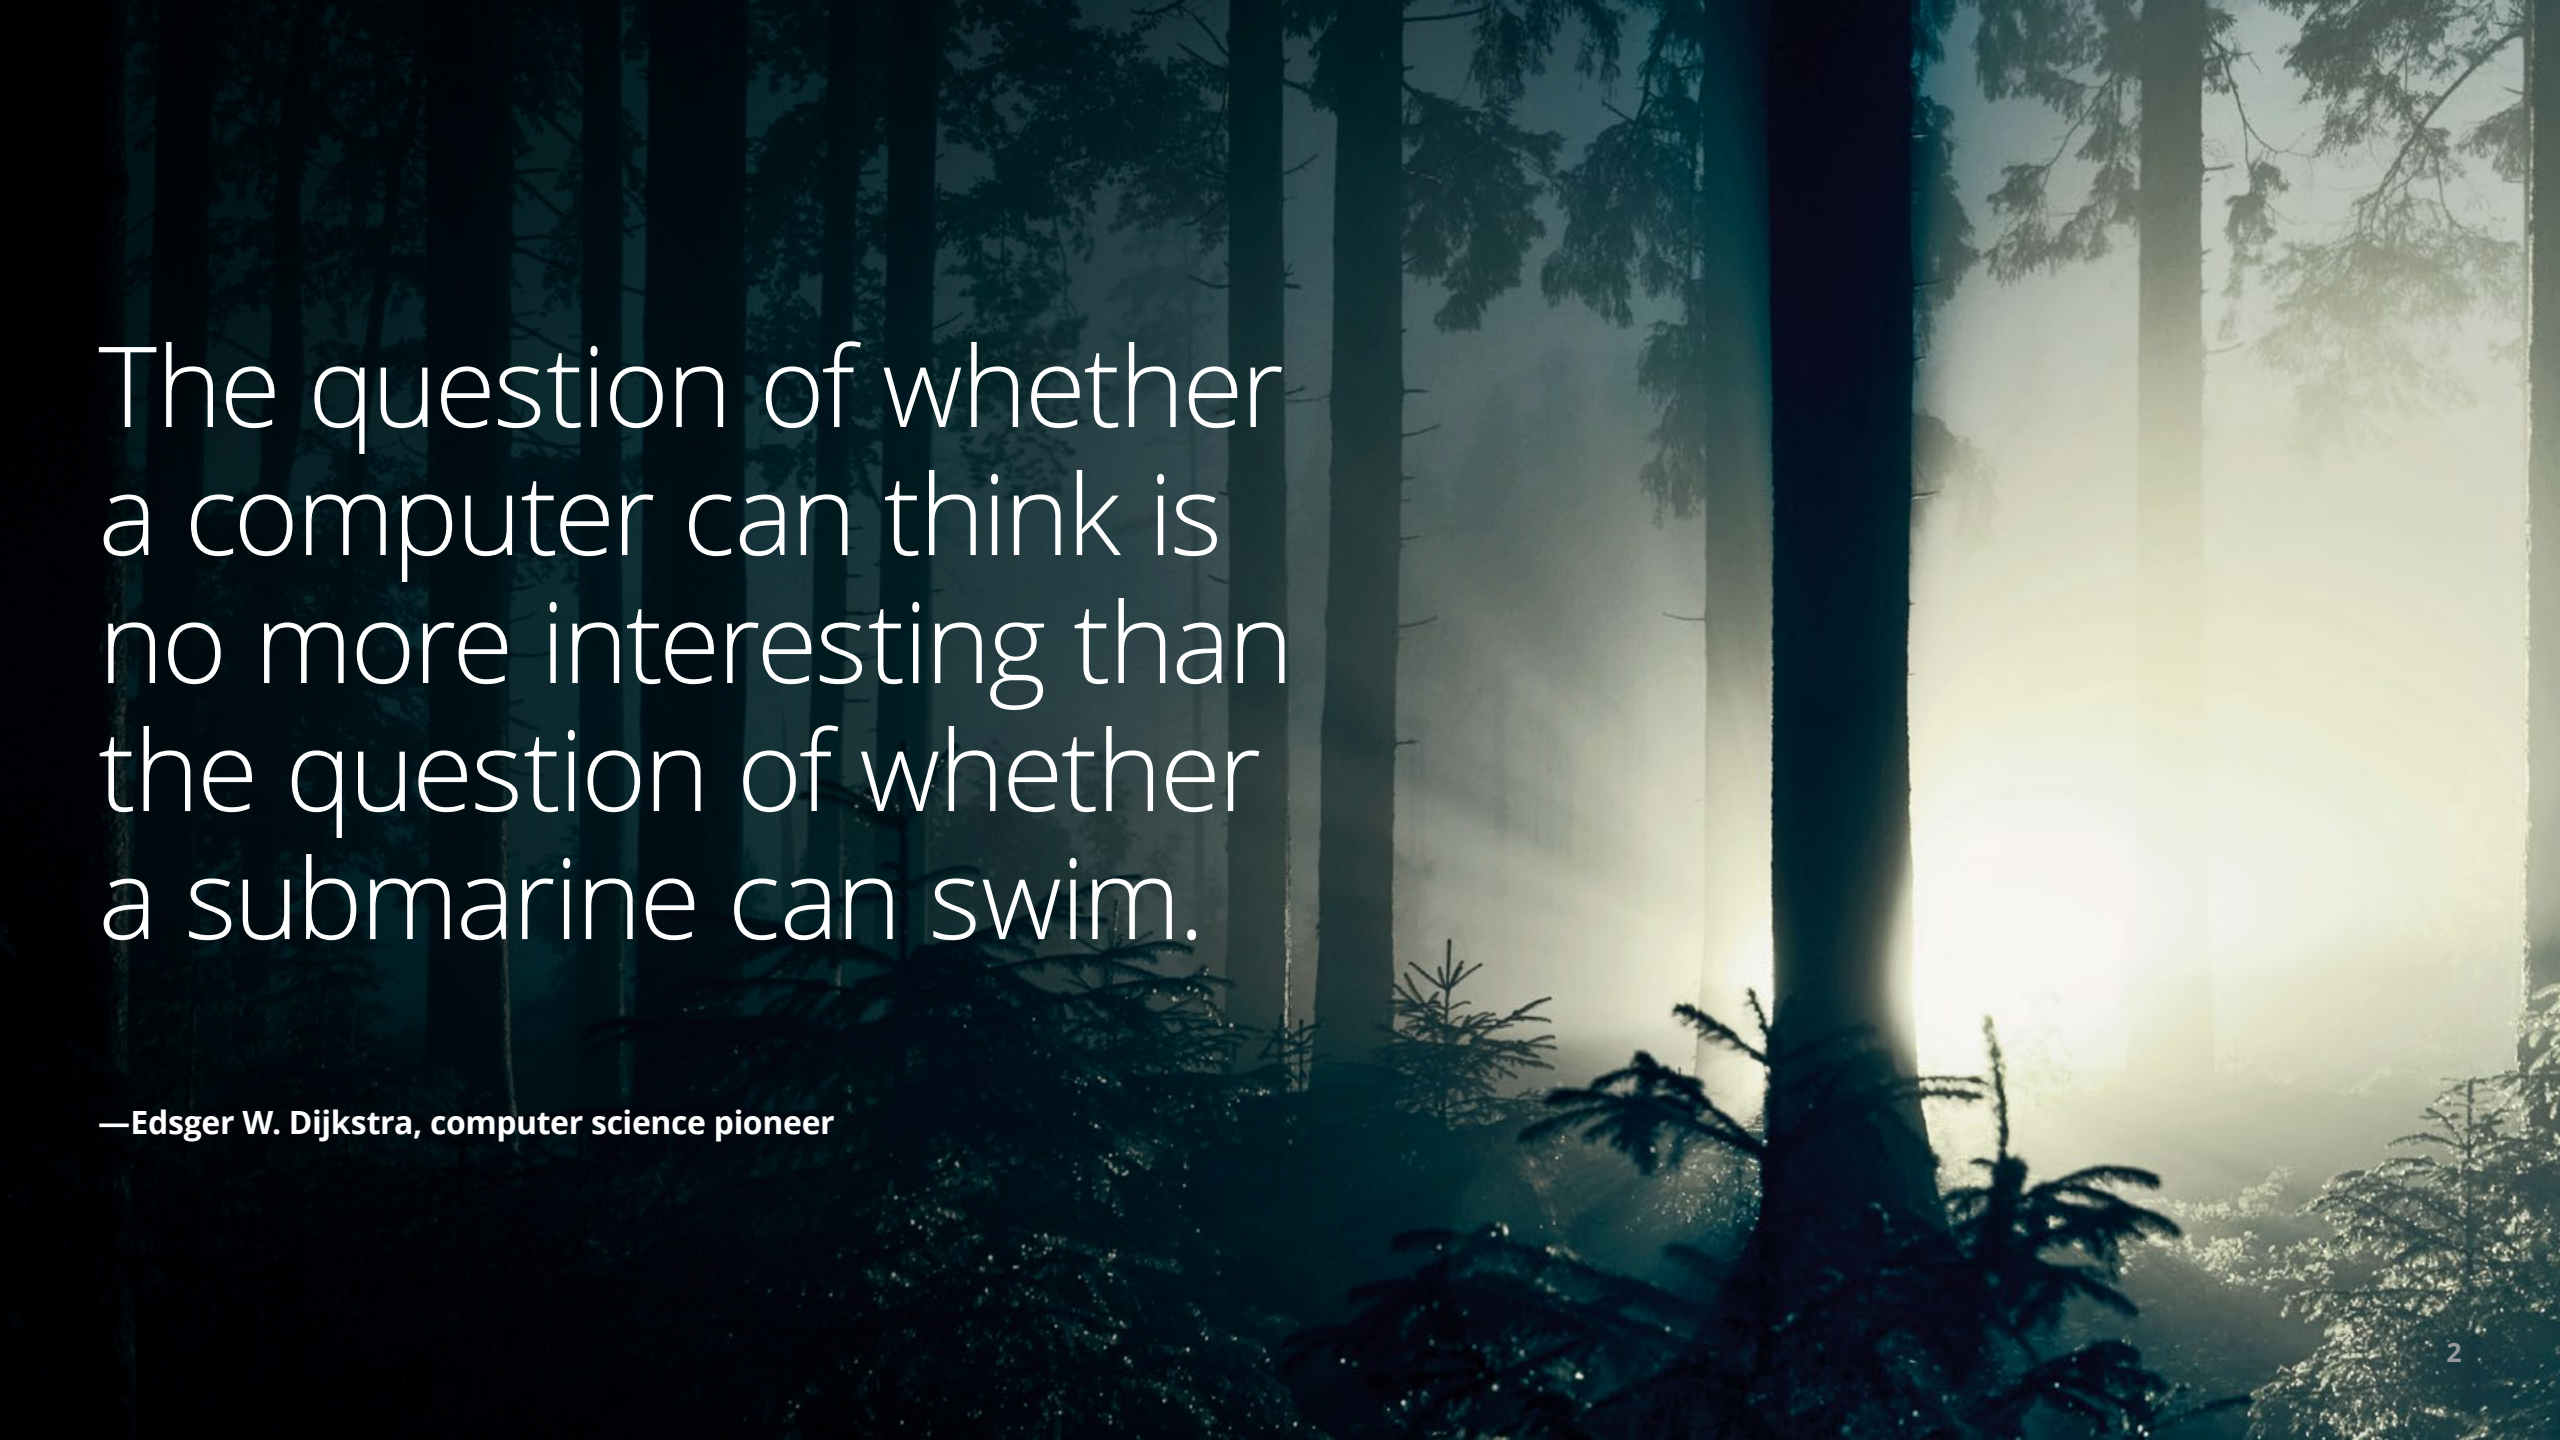
The question of whether a computer can think is no more interesting than the question of whether a submarine can swim.
—Edsger W. Dijkstra, computer science pioneer

Contents

Everything and nothing at all

As futurists, my team and I secretly spend the lion's share of our time studying the past. I like to say that we're closet historians. Specifically, we research the history of various technologies and how they've impacted, or failed to impact, the way the world works and lives. When you spend 25 years up to your eyeballs in all-things-newfangled, you can't help but start to recognize patterns. Scripts. Déjà vu.

And so, when tech headlines begin to increasingly read as breathless brochures for Artificial Intelligence, we grey-hairs can't help but be reminded of a certain upstart technology category some 20 years ago called "the world wide web."

The similarities are striking. Just as startups and incumbents alike once appended "dot.com" to their names to supercharge their marketplace perception (and IPO values) in the early 2000's, so too are today's players cloaking themselves in the "everything to everyone" halo afforded by AI. Statista reported over 2,000 AI-focused companies in the US alone as recently as 2018,[1] a number that's further exploded over the last two years.

Like any gold rush, there's hope beneath the hype. To be sure, there is plenty of actual gold afoot insofar as we're seeing a genuine, evidence-based phase shift from AI as "cherry-on-top" curiosity to "key ingredient" at leading organizations. 61 percent of respondents to a recent Deloitte Insights report say AI will substantially transform their industry in the next 3–5 years. Furthermore, adoption is significant on a per-organization basis[2], with 53 percent of those polled spending more than $20 million during the past year on AI tech and talent.[3]

AI's increasing centrality to business processes, and even strategy, is no longer up for debate. The "Age of With"—human work augmented and enhanced with AI—is upon us.

As with any exponentially accelerating emerging technology, the abundance of news has, however, given way to an even greater abundance of noise.

AI's increasing centrality to business processes, and even strategy, is no longer up for debate.

Everything and nothing at all

One time-honored means of getting past hyperbole is to double-click into the crunchy research itself. Zoom in too far, though, and we're blinded by a blizzard of buzzwords[4] too numerous, and too arcane, to reasonably keep track of. Papers with Code, an open source AI research community founded by the Facebook AI research team, itself hosts 37,000 published research papers and 3,000 datasets. A scrappy leader might reasonably visit such a community to derive a coherent first-person understanding of, say, "Weakly-Supervised Action Localization by Generative Attention Modeling." Woe, though, to that leader if she were to try to put together a boardroom-relevant "So What? And Now What?" executive summary at that altitude.

On the flipside, zoom too far out, and we're left with Artificial Intelligence and Machine Learning (AI/ML) as a blanket platitude. Like "synergy," "innovation," and "leverage," the fast and loose inclusion of "AI" into a business conversation is at once obligatory, and an eye-roller. When every story is an AI story, none of them are. And worse, some of them *actually* aren't. 40 percent of European AI startups, upon closer inspection, don't even *use* AI technologies or techniques.[5]

The key, it would seem, is to peer beneath AI-as-platitude and above the a la carte jargon. To shoot the curl towards the "Goldilocks Zone," where clusters of coherent, high-growth, impactful[6] sub-movements are both easier, and more useful, to understand. In this way, we can not only better sense what's happening, but make sense of where these movements are likely headed.

When every story is an AI story, none of them are.

Something old, somehow new

AI is not new.

Founded as an academic discipline in 1955, it's practically as old as the first digital computer. As with most emerging technologies, a gradual (though anything but smooth)[7] convergence of cost reductions, performance improvements, and network effects has only recently conspired to make AI a boardroom-relevant agenda item. What, specifically, has changed, and why are we talking about so much about it right now?

For starters, there's the *AI Effect.* Namely, the idea that as machines become more capable, the tasks considered to require "intelligence" are often removed from the definition of AI. Take chess, for example. A fascinating AI problem if there ever was one, at least until IBM's Deep Blue defeated grandmaster Garry Kasparov in 1997, after which it was reframed as *not really* the realm of AI. As ever more powerful AIs racked up ever more impressive championships in Poker, Go, and the TV game show Jeopardy!, the expectations only continued to rise.

As such, AI has become a catch-all term for "whatever computers can't do *yet.*" Hence AI's ability to continually grab our attention, despite it having been around for nearly 75 years. Less than 10 years ago, a natural language voice-assistant on our phone would be considered absolute magic. Today, the average American home has more such smart devices than people.[8]

Something old, somehow new

Behind this curious expectations-management whiplash: Our tendency to prize intelligence as a uniquely human virtue. AI's maturation has, in part, been a battle between our pride in creating something exceptional and our pride in ourselves *being* exceptional. Our profound need as humans to believe that we're capable of *creating* amazing things, just so long as they're not as amazing as us. This tension fuels the ever-moving goalpost as to what *really* constitutes interesting AI.

AI has always felt, and will always feel, "new."

That said, the contemporary burst of enterprise attention is most attributable to the recent point-in-time convergence of cloud-based architectures and open-source AI toolkits. Bespoke machine learning models have been built for decades on offline copies (dumps, snapshots, or slices to cite a few terms of art) with interesting and insightful results. It's only with the recent explosion of cloud-native production datasets that organizations can efficiently build and deploy real-time, point-of-need models that perform quickly enough to fit into existing business processes.

We can liken things to the difference between C-3PO and Chewbacca in the cockpit of the Millennium Falcon. Legacy analytics systems, like 3PO, could point out the odds (i.e., insights) and tell us verbosely, after the fact, what we might consider doing differently. Today's AIs, like Chewy, are fast enough to grab ahold of the controls and get us out of trouble. Co-pilots as opposed to critics.

This relatively recent development, *agency*, sits at the center of AI's newfound business utility. As ML systems graduate from "out-of-band" analyzers to "in-band" actors, their value proposition increases markedly. Specifically, organizations have begun to realize that, properly embedded in business processes, AIs *can reduce their fully-loaded cost per decision*.[9]

Or in the plainest possible English: AIs can, of late, save us time and money.

So, as we prepare to project likely AI futures, let's recap: AI itself is not new, but there is, and will always be, a new milestone afoot in AI. Each of these milestones seems improbable upon approach, and in retrospect, "no big deal" upon achievement. AI's increasing centrality to corporate strategy is a result of its recently achieved agency—Its ability to be cost-effectively embedded into business processes.

The machines

(or, a technological history of the future)

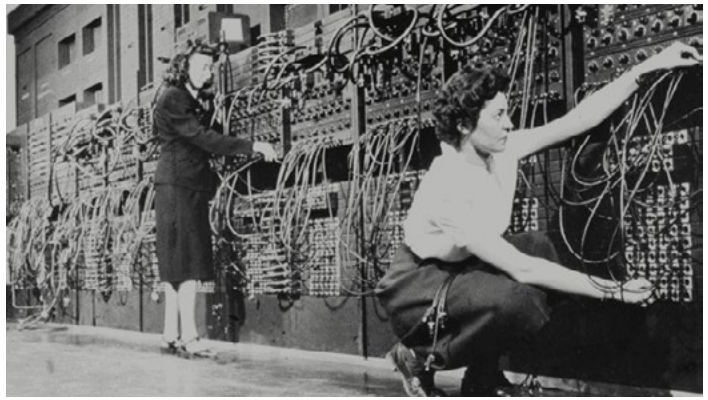

Charles Babbage and Ada Lovelace designed the first general-purpose computer, their Analytical Engine, in the 1830s. Frustrated by manual errors in their friend's astronomical tables, they hoped to bring the tireless precision of the steam-powered Industrial Revolution to arithmetic. Economies worldwide were being transformed by mechanical muscles (e.g., locomotives, looms, steam-shovels, etc.) and it stood to reason that mechanized math was itself an attainable goal. Alas, the engineering tools and techniques of the day (and it must be said, their investors) weren't ready to deliver on their design.

It would be roughly 100 years until the first digital computers would actualize the concept in the 1940s.

It would be roughly 100 years until the first digital computers would actualize the concept in the 1940s. Enormous, expensive, and requiring PhDs and punched cards to operate, early computers like ENIAC were hailed by the press as "electric brains," and by scientists as Turing Complete: That is, capable of solving any classical math or logic problem they were told to solve.

But as it would turn out, a *huge* set of real-world decision-making isn't readily framed as a tidy math problem. A painter, of course, follows their muse, not their math, in deciding which colors and lines to use next. Nor does an author adhere to an algorithm. Even scientists and professionals use more heuristics and "common sense" in their professional day than they do equations. A radiologist, for example, leans on their own unique mix of discernment, experience, and instinct to decide whether an MRI image suggests multiple sclerosis. Their expertise is hard to describe, and harder-still to codify.

Machine learning (ML) earns its paycheck in these situations. Whenever we're unable to logically or cost effectively use math to tell a computer what to do, we can use ML to teach a computer what to do by showing it examples of how it's been done.

Ernest Davis, Professor of Computer Science at New York University explains: *"So, say you want a computer to know how to cross a road, for example. With conventional programming you would give it a very precise set of rules, telling it how to look left and right, wait for cars, use pedestrian crossings, etc., and then let it go. With machine learning, you'd instead show it 10,000 videos of someone crossing the road safely (and 10,000 videos of someone getting hit by a car), and then let it do its thing."*[10]

Getting technical, today's AI moment has been powered by the ascendance of deep learning, a method of machine learning that can find patterns and infer meaning across many levels of abstraction, learning correlations and causation at the micro level, the macro level, and everywhere in between. Deep learning theory dates to the 1960s, but has only become popular over the last decade because its practical application requires the exponential increases in both computing power and training data now available thanks to cloud-native architectures.

This current AI/ML "Cambrian explosion" is resulting in a radical rethink as to what computers can realistically learn. Startups and incumbents alike are teaching machines to emulate an ever-increasing share of capabilities once thought of as "uniquely human."[11]

Startups and incumbents alike are teaching machines to emulate an ever-increasing share of capabilities once thought of as "uniquely human."

Sensation and discernment

Thanks to ML, (and in no small part to concurrent advancements in 5G networking and edge processing), computers are increasingly able to sense and make-sense-of visual information. Autonomous vehicles get much of the attention, but computer vision innovations abound. Photogrammetry solutions are allowing police to reconstruct detailed 3D accident scenes from simple 2D drone photographs.[12] Venture-backed MIT spin-out Affectiva, having analyzed over 10 million human faces, can "see" if a driver is sleepy or alert. And companies like Proscia offer "Computational Pathology," doing their best to augment the work of our previously imagined radiologists.

AI's hearing is getting better too. Thanks in large part to the proliferation of digital assistants (and the data they can collect for ML training), NLP—Natural Language Processing—has moved from Research and Development to Rec Rooms and Dining Rooms in less than a decade. If human speech weren't enough, the next frontier of research finds AIs able to recognize all manner of sounds (e.g., A horse trotting, a dog yawning) and better still, in turn *produce* convincing *guesses as to what sounds* might best accompany a silent video clip.[13]

Advances in artificial olfaction[14] (bomb-sniffing bots), touch[15] (dexterous droids), and taste (computer chefs)[16] promise to be similarly compelling.

Creativity

The next frontier for AI technology is the surreal realm of machine creativity. Computers have been uncontested captains of the math team for 75 years. Thanks to AI/ML, they're well on their way to acing reading, writing, and an increasingly impressive swath of electives.

Consider NLG (Natural Language Generation) systems, like OpenAI's GPT-3, which can write prose "which human evaluators have difficulty distinguishing from articles written by humans."[17] What kind? Essays, blog posts, functional design specifications, resumes, and even poetry.[18] All as a result of advanced machine learning models trained on staggering amounts of exemplary source material.

Figure 1: OpenAI GPT-3 Demo

Transformer AI poetry: Poetry classics as reimagined and rewritten by an artificial intelligence.
"Sonnet 18"
William Shakespeare
Shall I compare thee to a summer's day?
Thou art more lovely and more temperate:
Rough winds do shake the darling buds of May,
And summer's lease hath all too short a date;
A winter's day, when beams of sun are few,
And skies are grey with clouds intemperate,
Is to thy beauty somehow sordid grown.
Her lovely-lovely eye is quite shut up,
Like windows pale before a sunny shower.
The earth that's cold and moist, like to my mistress,
When once that she unbrac'd, her lovely jewel,
Grows as both ugly and uncourtly.
Poor soul, the center of my motley host,
See, now they request thy sovereign presence;
The din of merry hearts hath brought me thus
To greet thee, thou unsoil'd and saint-like beauty.

Image from machinelearningknowledge.ai, 2020[19]

More of a visual learner? Consider DALL.E, the same lab's cleverly named algorithm[20] that has been trained to draw anything you tell it. Pictured in Figure 2: its various imagined stylings for the prompt "Avocado Chair."

These kinds of "generative" examples abound. Even music[22] and dance[23] are fair game. If you can show enough examples of "good" (and ideally, "not so good") art to a neural network, it stands to reason that the machine can eventually graduate from philistine, to critic, to outright creator. Computers have been uncontested captains of the math team for 75 years. Thanks to AI/ML, they're well on their way to ace-ing reading, writing, and an increasingly impressive swath of electives.

Figure 2: Avocado chair

Image from OpenAI, 2021[21]

Congeniality

We've established that machines are increasingly, for lack of a better word, "smart." Though, just as the Inuit have 12 words for "snow," intelligence researchers recognize many different varieties of intelligence. An upcoming frontier in AI/ML: Emotional Intelligence. Though it's early days, we're starting to see the emergence of machines that can recognize human emotions, and in turn, help comfort and care for us.

Established players like Google and Amazon have, for some time now, offered AI/ML tools to extract the emotional sentiment from a text. IBM's Watson has been further trained to recognize nuanced sentiments like sarcasm and hyperbole.

Call-center player Cogito "performs live, in-call voice analysis to augment [agent] behavior in real time"[24] ensuring, for example, that flustered-sounding customers are escalated to a supervisor before they themselves even demand it. Digital interview providers can evaluate face, body posture, and appearance to algorithmically infer if a candidate is a self-starter, or a team-player.[25]

At the frontier, Replika's "AI companions who care" seek to create genuine companionship between humans and bots who are "always [there] to listen and talk."[26] The parallels to Scarlett Johansson's computer paramour in *Her* are readily noticeable.

As we've seen, the applications of AI/ML are as varied as they are numerous. There's quite literally nothing a human can do that a sufficiently trained machine won't (one day) be able to emulate. It's enough to give us (i.e., humanity as a whole) an inferiority complex.

Except: We humans continue to retain one critical advantage: Well-roundedness. Or as scientists call it, *General Intelligence.* Any given AI can be trained to be better than any given human at any one task, but no AI yet developed (or conceived) can match the diversity, adaptability, and all-around versatility of a standard-issue human being. Today's AIs, however impressive, demonstrate Artificial Narrow Intelligence (ANI). The advancements required to ladder up to Artificial General Intelligence (AGI) are, by nearly all accounts, a very long way off.

Us *and* the machines

(or, an anthropological history of the future)

Let's turn the page from future-shock to culture-shock.

There's a well-known, and increasingly well-researched, tendency to strengthen our standing and sense of belonging within our own social circle by disparaging, dehumanizing, and even demonizing those outside of it. Sociologists and anthropologists call this "In-Group/Out-Group" behavior: The idea that the shared vilification of "the other" strengthens group bonds and aligns shared values.[27] Typically (being social science) this manifests in social segmentation: Political parties, tribal identities, international conflicts, etc. As an old Bedouin proverb goes, "Me against my brother, my brother and I against my cousin, and all of us against the stranger."

Interestingly, AI research suggests that as machine intelligence ascends exponentially, our tendency to anthropomorphize it—that is, to think of machines in human terms and as having human characteristics—rises alongside of it. In turn, we (and our human "cousins," per the proverb) succumb to the tendency to think of machines as "the stranger" to be feared, and even shunned. To be sure, respected academics and industrialists alike[28] have begun to issue humanist rallying cries around the existential risks posed to "us" (i.e., People) should "they" (i.e., Machines) be allowed to continue to develop unchecked.

A more useful rubric for thinking about our AI future might instead be to think about machines less as a stranger to compete against, and more as a partner to be teamed with. Less "other," more "brother." We have crossed the threshold to the Age of With, wherein humans can team with machines to deliver outcomes that marry mechanized speed and precision with human intuition and curiosity.

Concrete manifestations of human and machine intelligence "better together" are all around us. In our personal lives, this currently shows up in the proliferation of digital assistants waiting attentively at our beck and call. Thanks in part to the aforementioned AI Effect, our tendency to take digital assistants for granted belies the expectation that this category will grow at a CAGR of 31.9 percent to $19 billion by the year 2027.[29]

Take, as a thought experiment, the everyday example of a smart speaker responding to our request for a weather forecast. In some ways, this is the epitome of rote. After all, we've been able to look up local weather over the internet since the early 1980s. Though, it must be said, that to do so in the 1980's required a bulletin board system subscription, a noisy dialup modem ritual, a typed user-authentication, and a zip code query. Contrast that user experience with my eight-year-old son simply musing aloud into the kitchen "Computer, what's the weather?" A cynic might say that this experience differs by degree, not kind; A "lift-and shift" from one user interaction modality to another. Try telling that that to my eight-year-old son.

Things get inarguably more interesting, though, as these intelligent agents continue to proliferate, specialize, and proact.

Proliferation

One of the most interesting developments in the continuing evolution of "smart assistants" is the rate at which they're multiplying. This proliferation is taking place across at least two dimensions: Platform and Place.

On the platform side, we've seen a 10-year explosion from "Just Siri" to a thriving ecosystem of Siri, Echo, Google, Cortana, and a long tail of special-purpose entities like Bixby, Hound, Robin, and Lyra. B2B agents have even emerged, with companies like IPSoft releasing a marketplace of digital workers under their Amelia moniker.[30] What might seem faddish to some is, more likely, a recognition that user interface technologies trend toward simplicity over the long term. NLP is our likely UI future insofar as the only operating systems with a greater install-base than Windows, Android, and iOS are English, Spanish and Mandarin.

As to place, we're seeing the move from "Single Device" to "Devices Everywhere." Talking to one's phone, or one's smart speaker, is giving way to talking aloud wherever we are and expecting the AI to be there. Researchers call this "Ambient Computing:"[31] The idea that we're likely to move beyond the glass, and beyond device-centric interactions altogether towards a ubiquitous computing ether that's at once always ready to help, yet always out of the way.

We're seeing the move from "Single Device" to "Devices Everywhere."

Specialization

When we project forward to an ambient AI future, we can't help but imagine collisions between well-intended assistants. To wit, my son's weather request works in a single-provider context, but what of the likely future wherein four, or 40, such agents are vying for our finite time and attention?

Our research suggests that the likely path forward will be further specialization. Just as today's startups are disinclined to enter into general competition with hyper-scale FAANG businesses, tomorrow's startups will be more likely to invest their energies developing domain-specific excellence and differentiation. It stands to reason that Alexa, Siri and their ilk may further evolve to become a sort of digital butler in charge of a full staff of relatively junior, specialized, digital servants. Digital Downton Abbey, if you like.

Or perhaps a better comparator might be "general contractor," our primary point of interaction who, in turn, sub-contracts specialized tasks not to captive staff, but to the most efficient, effective, bidder? A marketplace of intelligent agents, all the way down, bidding on a microsecond-scale for our business in much the same way online advertising markets work today.

Proactive

As we've seen and stated, AIs move further up the value chain as they elevate from reactive order-takers to proactive change-makers. The future of artificial assistants will no doubt rest on their ability to predict and proact in a way that we, as humans, find helpful and not pesky. A generation ago, the term "proactive AI assistant" conjured a garrulous digital paperclip hell-bent on helping us write a letter. As machine learning techniques, and the data training them, mature, we can rightly look forward to machines that behave less like meddling naggers, and more like trusted advisors and coaches. For example: Microsoft Research is already piloting digital AI advisors on everything from PowerPoint presentations (e.g., "Consider talking a bit less") to email habits (e.g., "Consider being a bit more direct").

Microsoft Research is already piloting digital AI advisors on everything from PowerPoint presentations to email habits.

From bits to atoms

But what of physical work? The sorts of digital assistants we've discussed thus far lean toward information management, not physical manipulation or production. To be sure, most of today's popular AIs traffic in bits as opposed to atoms. What, then, about the metallic robotic sidekicks we were promised in our youth?

The first thing to know about robots is that they're already here and have been for some time. The second thing to know is that we tend to overlook them because they very rarely look even remotely human. Just as algorithmic hyper specialization characterizes machine learning, physical hyper specialization characterizes physical robotics.

While eerie humanoid and animalesque robots capture the imagination (and trigger our uncanny valley response), most of the gainfully employed robots working in today's economy look downright dowdy by comparison. Package pickers, vacuums, welders, etc. Just as a picture is worth 1,000 words, a robot's physical form captures an undue share of our attention. To be sure, there's a rich future ahead for ever more *physically* capable bots. Advances in dexterous robotics portend a future where machines will match, and eventually exceed, our fine motor skill abilities. Advances in energy storage combined with algorithmic proprioception suggest that tomorrow's robots will match and exceed our gross-motor and ambulatory abilities as well. As with information-focused AIs, physically focused AIs will develop human, and eventually super-human, capabilities in *domain-specific tasks.* But akin to the chasm between ANI and AGI, the prospect of all-purpose robotic super-humans is a long way off.

The earliest industrial robots, say, those that precision-welded automotive frames on an assembly line, didn't play particularly well with humans. For starters, their situational awareness was nil. As pre-programmed collections of servos and actuators, they were technically robots, but their inability to sense and respond to environmental stimuli, in conjunction with their tendency to wield saws and/or torches, made them sub-optimal teammates.

Today's robots are different in at least two material ways. For one, they're much more configurable; Not general-purpose so much as "serial specialists," able to be more quickly reprogrammed to cover a much wider array of use cases. Second: They're much more situationally aware. That is, they've affordances to sense the visual, audio, and even olfactory cues needed to more thoughtfully (i.e., Safely) navigate their space and go about their work.

As a result, the future of robots in the workplace is increasingly turning towards "Cobots": Collaborative robotic teammates that work *with* us, rather than *instead* of us. The International Federation of Robotics recognizes a maturity curve wherein Cobots will gradually evolve from sequential collaboration (i.e., Taking turns with humans to work in the same physical space), to synchronous co-operation, to responsive collaboration wherein the robot responds in real time to the human's motion.

Also important: the emergence of "one-shot" learning models, wherein a cobot can be made to recognize an object it's never seen before without requiring a vast repository of training image data. One-shot cobots can learn new objects, and in turn, concepts, just by watching their human co-workers. By lowering the machine learning bar from arduous training and programming from distant PhDs to hand gestures and voice commands from experienced co-workers, cobots stand to eventually be able to learn on the job.

One-shot cobots can learn new objects, and in turn, concepts, just by watching their human co-workers.

The takeaway, for both physical robotics and intelligent agents is compelling, and harkens back to our opening point about "us vs. them:" While histrionic headlines have long fixated on robots coming for our jobs, research trends suggest that they are more likely to be our colleagues, as opposed to our replacements. Colleagues particularly well-equipped to specialize on subsets of "4d" tasks that many would rather avoid in the first place: Dull, dirty, dangerous, or delicate.[2] Add to this the recognition that "virtually every country in the world is experiencing growth in the number and proportion of older persons in their population,"[32] and the business case for capable robotic helpers matures.

Industry financial projections further paint the picture. The Cobot market[3] is expected to be worth $7.5 billion by 2027, or around 29 percent of the global industrial robot market. The market is further expected to grow at an annual rate of over 60 percent for the next two years and over 35 percent CAGR until 2027.[3]

Popular entertainment as early as the 1950s foresaw robotic automation as an enabler of both greater prosperity and time for leisure. George Jetson memorably lounging in his office chair all day, only occasionally having to (gasp) lean forward to press a button. The intervening half-century has seen this promise half-fulfilled: While productivity has skyrocketed, human burnout has increased as well. Susan Lambert, professor of University of Chicago[4] argues that, nonetheless, AI & physical robotics are the key to creating more humane schedules and expectations for workers. While this might not be the most profitable course of action in the short term, it can lead to increased productivity and more sustainable growth in the long run.[4]

While histrionic headlines have long fixated on robots coming for our jobs, research trends suggest that they are more likely to be our colleagues.

Us *as* the machines

(or, a biological history of the future)

While all science fiction explores our relationship with technology, a particularly vivid sci-fi trope speaks to the taboo of ourselves becoming technology-based beings. In 1977, Star Wars taught a generation that Darth Vader was "more machine now than man. (ergo) Twisted, and evil." In 1989, Star Trek introduced us to "The Borg," a sinister race of cybernetic organisms committed to assimilating natural species and stamping out their individuality. Stories like The Matrix ('99) and Black Mirror ('11) have continued to underscore a closely held cultural value: Humanity > Technology.

And so, we acknowledge the shock that accompanies research suggesting that our shared AI future is, in part, about our becoming cyborgs.

Furrowed brows and frowns tend to soften when we clarify that this this isn't a jump straight to neural implants and carbon fiber skeletons. It's a slow, symbiotic co-evolution, of our own choosing, and it's already begun.

AI around us

An evidence-based argument[33] can be made that practically speaking, we're already cyborgs. As a colleague recently mused, "Just take away someone's phone for a week and see how they respond." To be sure, the last decade has seen the prevailing computing form factor evolve from situated "click and type" modality to a mobile "touch and swipe" experience. Smaller screens and higher resolutions are part of the story, but the real headline has been *ubiquity*. That is, the promise of "always on, always connected" has gone from R&D speculation in the aughts, to marketing gusto in the teens, to our lived experience in the early 20s.

With this ubiquitous computing access comes, for the first time, the ability to reliably and regularly offload a meaningful share of our cognition and recollection to our mobile devices and AI assistants.

My generation, Gen X, was commonly told in our high school math classes that we needed to memorize arcane math facts because "we wouldn't always have a calculator in our pocket." Little did our teachers know that we'd not only have calculators, but that those calculators would also double as conduits to the entire collective knowledge of humankind to boot.

And so, just as we smug 1990s teens ended up needing trigonometry as little as our parents needed, say, typesetting, 2020s teens will smugly expect that much of their rote classwork will be obsolesced by some as-yet-undetermined AIs.

This isn't an incitement to sloth, but a recognition that *our minds have already been materially changed* by our ever-present immersion in technological supports. For example, research finds that today's students (and many professionals) no longer read entire texts line by line (using a "Z" eye-movement pattern), but rather, skim texts (using an "F" eye-movement pattern) teasing out relevant keywords in much the same way they might look for hyperlinks on a web page.[34] To be sure, I've myself been surprised seeing my teenage son compose an essay with a mix of NLP dictation and algorithmic grammar-checking. Some experts worry that these technologies amount to "weapons of mass distraction,"[35] ruining our ability to focus by trading cognitive depth for what author Nicholas Carr calls "the shallows."[36] Optimists might argue that these AI supports simply free that cognitive energy for higher order pursuits.

"But you won't always have Siri with you!" goes the contemporary argument. Our trendlines suggest that even the hardest scrabble futures are likely to enjoy an abundance of information and cognitive supports. In fact, our children will be more likely to have their AIs always on them.

AI upon us

The wrist

When the Apple Watch[37] wearable device debuted in 2015, much attention was given to matters of its fashion, craftsmanship, and cultural acceptance. This, in part, informed by the consumer market's recent rejection of head-mounted optical displays, several years prior. Available (initially) in "18-karat gold with prices starting at $10,000,"[38] the message was clear: This isn't *technology*, it's *jewelry*. Something for fashionistas, not just fans. A design-centric company, Apple[39] knew that the utility and value proposition of a wearable computer would be immaterial if people couldn't get past social stigma.

Just six years later, computers-on-wrists are the new normal.[40] And with that, the emerging normalcy of talking to one's wrist, prompting a bot for help with all manner of knowledge and cognitive support.

Interestingly, as wearables have evolved, many of their capabilities are being focused inward. A majority of the iterative investment and feature focus in this category has been on biometrics:[41] Activity levels, cardiopulmonary health, fall detection, and hearing health to name a few headliners.

The trendline would seem to suggest that just as we've gotten comfortable being seen with these devices, we're getting increasingly comfortable with them seeing us at a level of detail once afforded only to medical professionals.[42] At the risk of cliché, it's not necessarily an overstatement to say that these wearable intelligent devices are beginning to know more about us than we readily know about ourselves. For instance, there's a good chance your wearable has a more up-to-date understanding of your blood oxygen level than you do at this very moment. More usefully, the AI model currently being trained on your (and millions of others') SpO2 data is better positioned to recognize abnormalities and recommend medical interventions than you, as a patient, or even than your doctor.

The rest

Turn the clock forward, and we see all manner of wearables, each a conduit for assistive AI. Necklaces that track reproductive health,[43] yoga pants that refine your poses,[44] and mood rings that actually work[45] are just a few examples of the looming boom in wearable tech. Further out, David Eagleman, a Stanford University neuroscientist, has created VEST (Versatile Extra-Sensory Transducer): A thin wetsuit of sorts, worn around the torso, with 32 vibrating motors that can help the deaf "hear" ambient noise, or proactively remind you to correct your posture on your next video call.[46]

The frontier beyond is very likely to feature the second coming of smart glasses, this time propelled by the twin tailwinds of wearable acceptance and augmented reality (AR) advances.

LiDAR (light detection and ranging), already a feature on today's mobile devices, will play a starring role in tomorrow's AR glasses, allowing wearers to paint detailed, contextually-relevant metadata anywhere in their field of view. Early use cases are likely to be "merely very useful" as opposed to jaw-dropping or transformative. Virtual street signs appearing on unmarked street corners. The ambient temperature emblazoned on the walls in a cold storage room, with a helpful maximum safe exposure indicator to reference. The UV index hovering beside the sun along with a handy countdown to dusk for youth sports coaches, or golden hour for photographers.

And on the back end, the spatial web:[47] The next-generation data and processing architectures needed to store and serve this vast, shared digital reality to millions of people.

As hyperscale incumbents normalize smart glasses, and startups get to work on smart contact lenses,[48] the trendlines point toward a future where most our everyday physical experiences will be mediated through a digital layer informed by machine intelligence.

Reality itself, online.

The frontier beyond is very likely to feature the second coming of smart glasses, this time propelled by the twin tailwinds of wearable acceptance and augmented reality (AR) advances.

AI inside us

There's no getting around it: The idea of chips-in-our-brains is jarring.

The move from desktops to smart phones (and further on to smart glasses and contacts) constitutes a difference of degree, not kind. The screens get smaller, and ever closer to our eyes, but they're still screens, and it's still our eye. An AI-enhanced eyeball implant, though, feels like a different animal entirely.

First: The "ick" factor. Then, the biological safety concerns. Finally, depending on culture and creed, a variety of taboos regarding cleanliness, identity, and spirituality. How much, or which parts, of "you" can you replace or augment before you are no longer "you"? A top-tier philosophical discussion to be sure, but, alas, outside of this piece's scope.

From a more brass-tacks information technology perspective, direct brain-computer-interfaces would offer greatly improved bandwidth and reduced signal loss. Talking to a computer might be faster than typing on one, but both pale in comparison to the prospect of "thought control."

Today, for example, research teams are working with paraplegics to use nascent neuro-muscular implants to help them regain the use of their limbs, and with the blind and deaf to help them approximate sensations of sight and sound. These stories inspire us because they speak to restoration.

But what about taking good to great? Or great to superhuman? When South African sprinter Oscar Pistorius, a double-amputee, won Paralympic sprint events, he was roundly celebrated as an inspiration. His bold decision to compete in non-disabled sprint events proved controversial, to put it mildly. Clearly, he was differently-abled. But was he, with his carbon-fiber prostheses, super-abled?

We can only begin to speculate as to the increased nuance and intensity of this debate as we evolve from muscular augmentation to cognitive augmentation.

These quandaries aside, we humans stand to squarely retain our choice and agency in foreseeable human-machine relationships. Dystopian sci-fi stories invert that power dynamic: Technology as active antagonist, forcefully exerting its will on virtuous, unwitting, humans. In all likelihood, the most realistic cybernetic AI innovations will be those that we enthusiastically select from the marketplace of ideas.

Guardrails

Despite our largely optimistic projections of these probable AI futures, it's important to acknowledge that there are many milestones still needing to be met, and many roadblocks to be overcome. PhDs, professionals, and populists alike generally agree that these capabilities must be developed and deployed intentionally and ethically. More specifically, in a way that is consistent with our shared human values and which respects the dignity of individuals and protects our most vulnerable.

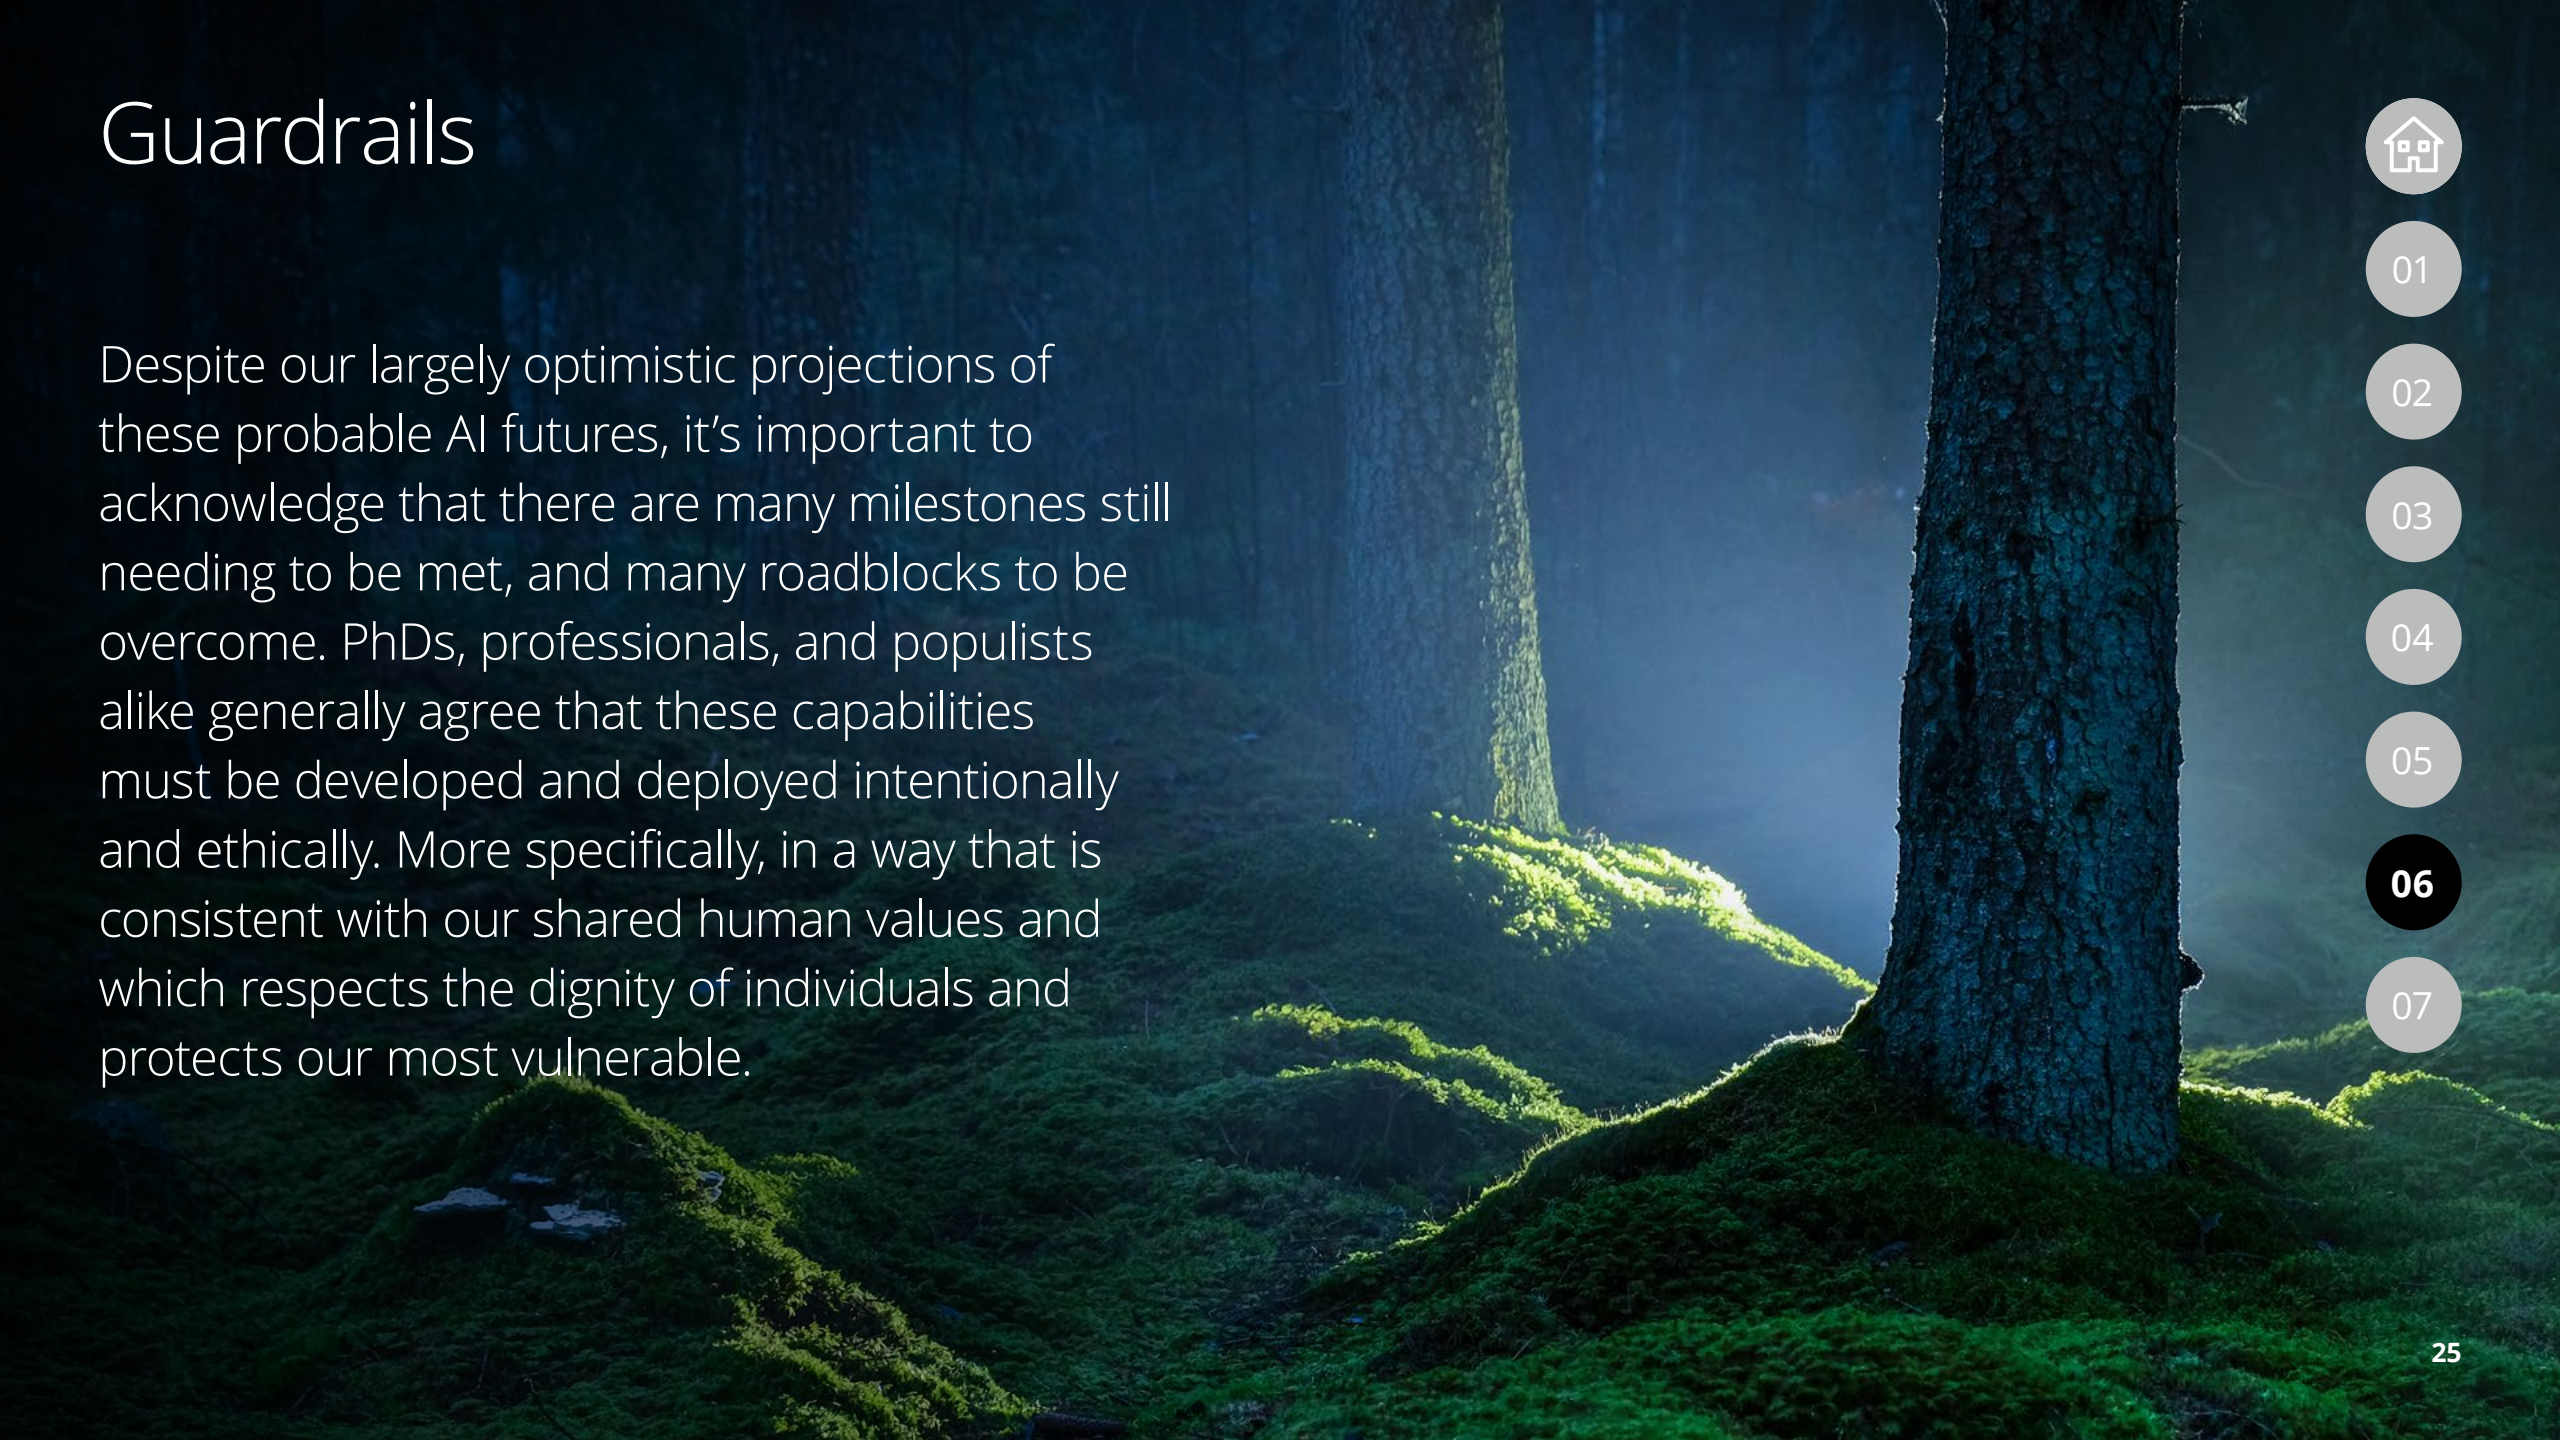

Our research suggests that the following four lenses will, taken together, guardrail our shared AI future.

1. Principles

Who is responsible for ensuring that AIs are developed and implemented in a fair, safe and ethical way?

We don't necessarily consider the subtle role AI plays when we call for a ride service. As far as we riders are concerned, a car magically appears. In truth, a complex series of algorithms is working behind the scenes to figure out who gets served where, and in what order.

What happens if this same system "learns," for whatever reason, to begin to avoid certain geographies or riders, because of tacit biases perpetuated in their original training? Is racism, sexism, or classism any more acceptable when a machine is the discriminator? No.

It's therefore critical that leaders and developers work to identify, and mitigate, these "algorithmic accidents" to lessen the likelihood of societal biases and prejudices souring outcomes before wide-scale adoption. The first step in this direction entails a move towards "explainable AI." That is, moving away from mercurial "black box" algorithms toward transparent, auditable models that promote accountability.

AI will not be a force for good if it's simply allowed to make bad behaviors happen faster.

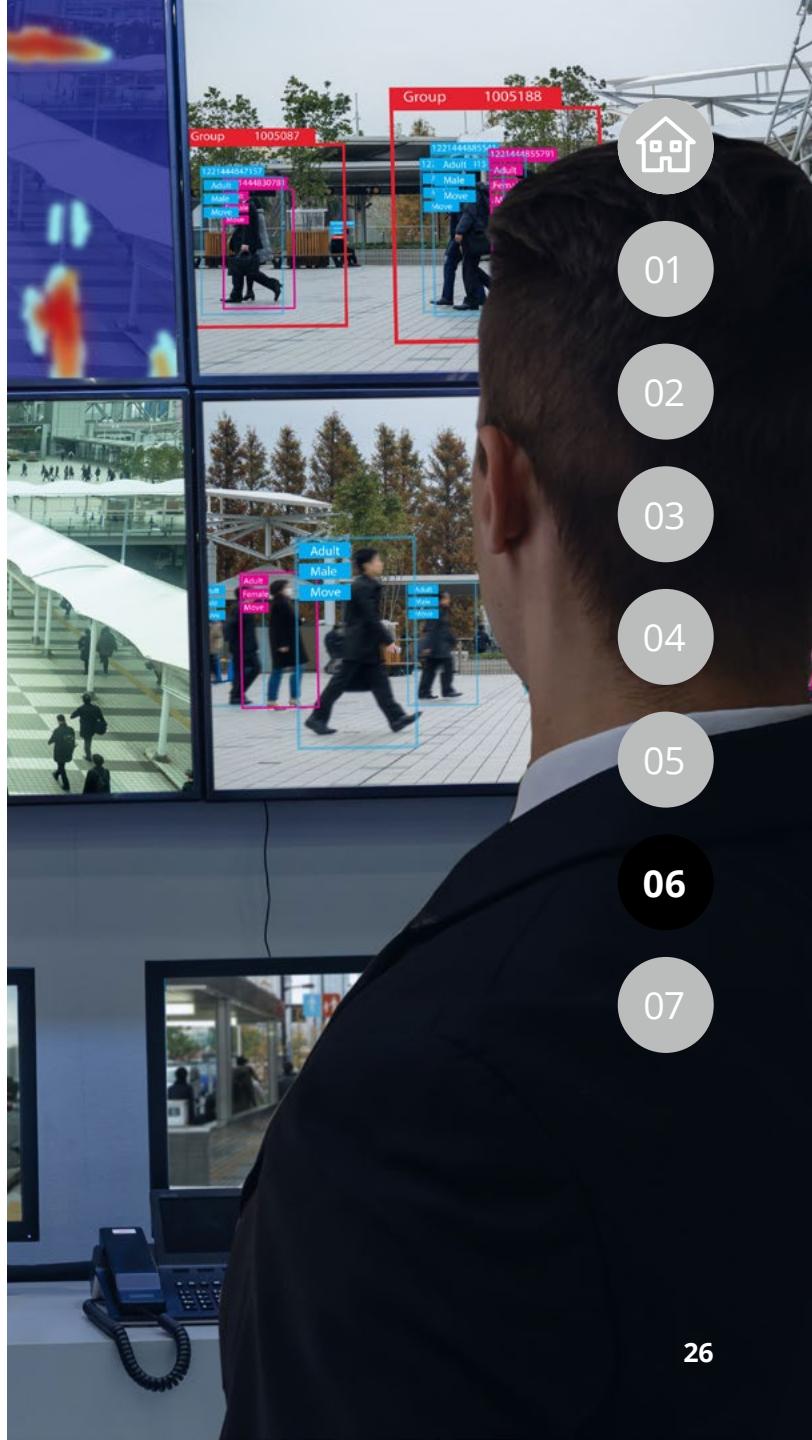

2. Privacy

Will AI pose a threat to our privacy and security?

On the security front, it's still early innings. AI-based tech is extremely beneficial to cybersecurity defense postures, but also an offensive weapon for those who would look to do harm. We might say that cyber implications figure to be a "draw."

It's likely safe to say that AI models' insatiable need for training data, however, creates pressure on its developers to collect and store data that they don't necessarily need. That increasing incentive to harvest and horde does pose a threat to anonymity (e.g., Social Security breaches), and privacy (e.g., Voice assistant devices accused of "wiretapping") should organizations, or perhaps more accurately, un-governed employees within them, fail to exercise restraint.

3. Professions

Will AI pose a threat to livelihoods?

AI stands to profoundly disrupt every major industry, and will likely touch every facet of our lives. Headlines about structural unemployment trend toward the hyperbolic, but it's critical that people outside computer science departments become AI fluent, or at least AI conversant. Biologists, for example, might cross-train in AI and advanced genomics. Soldiers might be trained on the proper use of AI to enable human-machine teaming in the battlespace. A robust, agile, AI-ready workforce requires STEM training and a cross-disciplinary ecosystem. Gone are the days where singular, exceptional, data scientists can be counted upon to chart the course forward. Just as computer literacy became professional table-stakes a generation ago, so it will likely go with AI/ML literacy, and data-driven decision making.

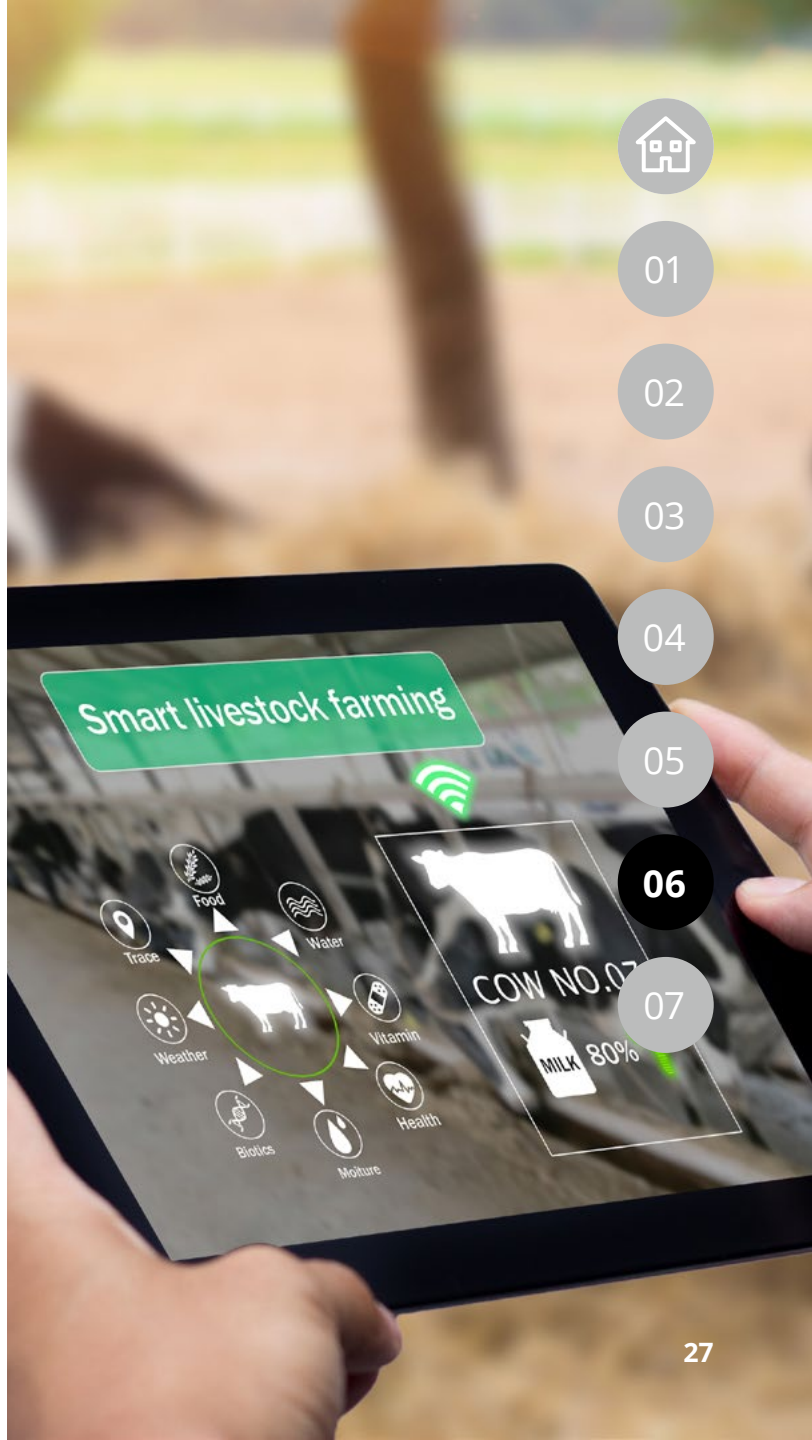

4. Policy

Who will be trusted to regulate and incentivize?

Policymakers and international institutions are already converging to establish conditions, in the form of standards, protocols, and requirements for AI applications. Activity has been afoot at organizations ranging from the UN, to the US State Department, to Microsoft's own Aether Committee. In much the same way that the EU's General Data Protection Regulation (GDPR) restrained big data collection, agencies could introduce restraints on AI auditability, capability, and applications, levying punitive action on violators in the form of financial penalties.

On the flipside, these same bodies, along with private entities and philanthropists, might elect to offer incentives to those who develop and deploy AIs in a manner consistent with expectations.

Speculation as to regulatory matters is ambiguous even by futurist standards, as the political and commercial landscape over the long haul are fluid. It might be confidently said, however, that the continued exponential explosion of artificial intelligence will almost certainly require a concurrent dose of human mindfulness to keep things on the rails.

Conclusion

Computer science pioneer Edsger W. Dijkstra once mused “The question of whether a computer can think is no more interesting than the question of whether a submarine can swim.”[49] This, his clever way of discouraging us from spending too many cycles admiring the existential quandaries inherent in AI.

Indeed, the long arc of history to current, when projected into the future, suggests that we may one day look back at the term artificial intelligence as an anachronism. A transitional precursor to an eventual realization that, whether it sits in vivo, or in silico, intelligence is intelligence.

Our species has always been defined (or at the very least, differentiated) by our ability to *learn*, *create*, and *adapt*. 2.6 million years ago, Homo Habilis created the first stone tools,[50] their “hack” freeing time and energy for higher order pursuits. The Sumerians created the first written language[51] 5,000 years ago as a hack for offloading memories and knowledge. This too, freed time and energy for still higher order pursuits. 500 years ago, the printing press similarly hacked communication, and 75 years ago, the digital computer, calculation. When seen through this long lens, AI is neither hero *nor* villain. Rather, it represents our species’ latest in a long series of hacks: This time, offloading discernment and decision-making.

The question before us, or perhaps, our children:

To what end? For what higher-order pursuits?

Now *there’s* a uniquely human problem worth solving.

Acknowledgements

Author

Mike Bechtel
Managing Director,
Chief Futurist
Deloitte Consulting LLP

Contributors

Amy Golem
Senior Consultant
Deloitte Consulting LLP

Raquel Buscaino
Consultant
Deloitte Consulting LLP

Varvn Aryacetas
Director
Deloitte LLP (UK Consulting)

Wes Craig
Senior Consultant
Deloitte Consulting LLP

Aliza Berger
Analyst
Deloitte Consulting LLP

Atul Fotedar
Consultant
Deloitte Consulting LLP

Grant Robinson
Consultant
Deloitte Consulting LLP

Fahd Aslam
Manager
Deloitte Consulting LLP

Deloitte
CATALYST

About the Deloitte AI Institute

The Deloitte AI Institute helps organizations transform with AI through cutting-edge research and innovation, bringing together the brightest minds in AI to help advance human-machine collaboration in the Age of With. The institute was established to advance the conversation and development of AI in order to challenge the status quo. The Deloitte AI Institute collaborates with an ecosystem of industry thought leaders, academic luminaries, start-ups, research and development groups, entrepreneurs, investors, and innovators. This network, combined with Deloitte's depth of applied AI experience, can help organizations transform with AI. The institute covers a broad spectrum of AI focus areas, with current research on ethics, innovation, global advancements, the future of work, and AI case studies.

Connect
To learn more about the Deloitte AI Institute, please visit www.deloitte.com/us/AIInstitute.

Endnotes

1. https://www.statista.com/statistics/941054/number-of-ai-companies-worldwide-by-country/#:~:text=The%20statistic%20shows%20the%20number,the%20United%20States%20was%202%2C028
2. Deloitte Insights report
3. Deloitte Insights report
4. https://fortune.com/2019/05/28/ai-buzzwords/
5. Study of European AI startups https://www.theverge.com/2019/3/5/18251326/ai-startups-europe-fake-40-percent-mmc-report
6. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2743186
7. https://www.longdom.org/peer-reviewed-journals/artificialintelligenceopenaccess-36559.html#:~:text=Artificial%20intelligence%20was%20founded%20as,approaches%2C%20success%20and%20renewed%20funding
8. https://en.wikipedia.org/wiki/AI_winter
9. https://variety.com/2019/digital/news/u-s-households-have-an-average-of-11-connected-devices-and-5g-should-push-that-even-higher-1203431225/
10. https://www.theverge.com/2016/2/29/11133682/deep-learning-ai-explained-machine-learning
11. https://www.youtube.com/watch/7Pq-S557XQU
12. https://www.aerial-metrics.com/
13. https://news.mit.edu/2016/artificial-intelligence-produces-realistic-sounds-0613
14. https://www.wired.com/story/now-machines-learning-smell/
15. https://news.mit.edu/2019/teaching-ai-to-connect-senses-vision-touch-0617
16. https://thenewstack.io/how-ai-will-help-us-find-new-innovative-flavors-of-the-future/
17. https://analyticsindiamag.com/open-ai-gpt-3-language-model/
18. *21 OpenAI GPT-3 Demos and Examples to Convince You that AI Threat is Real, or is it?* [Including Twitter Posts], MLK–Machine Learning Knowledge
19. https://machinelearningknowledge.ai/openai-gpt-3-demos-to-convince-you-that-ai-threat-is-real-or-is-it/
20. https://machinelearningknowledge.ai/openai-gpt-3-demos-to-convince-you-that-ai-threat-is-real-or-is-it/
21. https://openai.com/blog/dall-e/
22. https://filmora.wondershare.com/audio-editing/best-ai-music-composer.html
23. https://www.dancemagazine.com/is-google-the-worlds-next-great-choreographer-2625652667.html
24. Cogito | Augmented Intelligence for Your Call Center (cogitocorp.com)
25. https://www.technologyreview.com/2020/02/14/844765/ai-emotion-recognition-affective-computing-hirevue-regulation-ethics/
26. https://replika.ai/
27. https://waitbutwhy.com/2019/08/giants.html
28. https://www.wsj.com/articles/protecting-against-ais-existential-threat-1508332313
29. https://www.globenewswire.com/news-release/2020/10/09/2106176/0/en/Global-Smart-Speaker-Market-Trajectory-Analytics-Report-2020-Market-to-Reach-19-Billion-by-2027-U-S-Market-is-Estimated-at-825-4-Million-While-China-is-Forecast-to-Grow-at-30-8-CAG.html#:~:text=Global%20Smart%20Speaker%20Market%20to,the%20analysis%20period%202020%2D2027
30. https://amelia.com/digitalworkforce/
31. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6767848/
32. https://www.un.org/en/sections/issues-depth/ageing/
33. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6510111/
34. https://www.theguardian.com/commentisfree/2018/aug/25/skim-reading-new-normal-maryanne-wolf
35. https://books.google.com/books/about/Noise.html?id=W35PugEACAAJ
36. https://www.npr.org/templates/story/story.php?storyId=127370598
37. Apple and Apple Watch are trademarks of Apple Inc., registered in the U.S. and other countries.
38. https://www.techradar.com/news/why-apple-gave-up-on-its-18-karat-gold-apple-watch-edition
39. *The Future of AI: Seeing the Forest for the Trees, and the Forests Beyond* is an independent publication and has not been authorized, sponsored, or otherwise approved by Apple Inc. Apple is a trademark of Apple Inc., registered in the U.S. and other countries.
40. https://www.statista.com/statistics/1095866/us-adults-regularly-wearing-smart-fitness-tracker-demographics/
41. https://www.forbes.com/sites/timbajarin/2020/09/18/why-apple-focuses-on-health-and-fitness/?sh=19904b2b3dbf
42. https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN180044.pdf
43. https://bellabeat.com/product/leaf/
44. https://www.wearablex.com/collections/nadi-x-smart-yoga-pants
45. https://ouraring.com/
46. https://www.smithsonianmag.com/innovation/could-this-futuristic-vest-give-us-sixth-sense-180968852/
47. https://www2.deloitte.com/us/en/insights/topics/digital-transformation/web-3-0-technologies-in-business.html/
48. https://www.mojo.vision/mojo-lens
49. https://www.goodreads.com/quotes/32629-the-question-of-whether-a-computer-can-think-is-no
50. https://humanorigins.si.edu/evidence/human-fossils/species/homo-habilis
51. https://www.history.com/news/sumerians-inventions-mesopotamia

Deloitte.