DOWNLOAD PDF AI THE VERY IDEA

Chapter 1 : The very idea of AI is courageous. And exciting. And scary. â€“ Emerging Tech Japan

The very idea of AI may or may not be successful: we can only find out by trying. Haugeland's explanation of why this is so is usually informative, even if sometimes cursory: there are many details which beg for further analysis.

Additional Information In lieu of an abstract, here is a brief excerpt of the content: The usual picture of Ives as the Transcendentalist composer par excellence simplycannotaccountforwhatis most basic to his music, nor for the origins of his complex aesthetic. That Ives sought to depict Transcendental figuresandsomethingofTranscendental thinkinginmusicisobvious. Buthedid not develop his musical method or conceive his musical aims through or becauseofTranscendentalistphilosophy p. Perhaps these will be given in the promised sequel. Simon and Schuster, New York, It is an ambitious program with therapeutic encores. In it, there is something for everyone. Exercisesareprovided-more than twenty of them-to help you to implement the lessons, learning by doing. Before you begin to listen, take a momenttorelax. Breathedeeply,finda comfortableposition for listening,and tell yourself that you will let your mind openly respond to the music. Close your eyes and let yourself experience whatever your mind is creating in response to the music, whether it be recounting an experience from work, reliving a childhood memory, or visiting a happy yet forgottenplace Occasionally,thetopics touch aesthetic concepts and philosophical names are dropped, but not often. Because the authors are so evidently devotedtothevalueof musicalexperiences in human life, they present a mixed approach to their topic. On the onehand, they are intent on citing research others have done, giving the impression, the aura of scientificauthority and detached objectivity. On the other hand, they wax enthusiastic over the silliest highly speculative ideas: You know music from your very first stirrings of life. Within fourweeks of conception, your own heart has begun beating,announcingyourveryexistence to the world. The authors quote with apparent approval a nonsense statement by John Cageastheir theme: It has little to teach musicians or a sophisticated audience. Dreyfus and Stuart B. Free Press,New York, Many artists engage new scientific and technological ideas more in the spirit of dilettantes than as serious practitioners. While this approach allows them to You are not currently authenticated. View freely available titles:

Chapter 2 : Artificial Intelligence: The Very Idea - John Haugeland - Google Books

Artificial intelligence, Haugeland notes, is based on a very good idea, which might well be right, and just as well might not. That idea, the idea that human thinking and machine computing are "radically the same," provides the central theme for his illuminating and provocative book about this exciting new field.

Artificial intelligence could be our saviour, according to the CEO of Google This is where we come to the question of how we are going to spend our time. Most people still rely on selling their time to have enough income to sustain themselves and their families. We can only hope that this opportunity will enable people to find meaning in non-labour activities, such as caring for their families, engaging with their communities and learning new ways to contribute to human society. If we succeed with the transition, one day we might look back and think that it was barbaric that human beings were required to sell the majority of their waking time just to be able to live. How do we distribute the wealth created by machines? Our economic system is based on compensation for contribution to the economy, often assessed using an hourly wage. The majority of companies are still dependent on hourly work when it comes to products and services. But by using artificial intelligence, a company can drastically cut down on relying on the human workforce, and this means that revenues will go to fewer people. Consequently, individuals who have ownership in AI-driven companies will make all the money. We are already seeing a widening wealth gap, where start-up founders take home a large portion of the economic surplus they create. In , roughly the same revenues were generated by the three biggest companies in Detroit and the three biggest companies in Silicon Valley How do machines affect our behaviour and interaction? Artificially intelligent bots are becoming better and better at modelling human conversation and relationships. In , a bot named Eugene Goostman won the Turing Challenge for the first time. In this challenge, human raters used text input to chat with an unknown entity, then guessed whether they had been chatting with a human or a machine. Eugene Goostman fooled more than half of the human raters into thinking they had been talking to a human being. This milestone is only the start of an age where we will frequently interact with machines as if they are humans; whether in customer service or sales. While humans are limited in the attention and kindness that they can expend on another person, artificial bots can channel virtually unlimited resources into building relationships. Even though not many of us are aware of this, we are already witnesses to how machines can trigger the reward centres in the human brain. Just look at click-bait headlines and video games. This and other methods are used to make numerous video and mobile games become addictive. Tech addiction is the new frontier of human dependency. On the other hand, maybe we can think of a different use for software, which has already become effective at directing human attention and triggering certain actions. When used right, this could evolve into an opportunity to nudge society towards more beneficial behavior. However, in the wrong hands it could prove detrimental. How can we guard against mistakes? Systems usually have a training phase in which they "learn" to detect the right patterns and act according to their input. Once a system is fully trained, it can then go into test phase, where it is hit with more examples and we see how it performs. Obviously, the training phase cannot cover all possible examples that a system may deal with in the real world. How do we eliminate AI bias? But it can go wrong, such as when a camera missed the mark on racial sensitivity, or when a software used to predict future criminals showed bias against black people. Once again, if used right, or if used by those who strive for social progress, artificial intelligence can become a catalyst for positive change. How do we keep AI safe from adversaries? The more powerful a technology becomes, the more can it be used for nefarious reasons as well as good. This applies not only to robots produced to replace human soldiers, or autonomous weapons, but to AI systems that can cause damage if used maliciously. How do we protect against unintended consequences? What if artificial intelligence itself turned against us? Rather, we can imagine an advanced AI system as a "genie in a bottle" that can fulfill wishes, but with terrible unforeseen consequences. In the case of a machine, there is unlikely to be malice at play, only a lack of understanding of the full context in which the wish was made. Imagine an AI system that is asked to eradicate cancer in the world. After a lot of computing, it spits out a formula that does, in fact, bring about the end of cancer â€” by killing everyone on the planet. The computer would have

achieved its goal of "no more cancer" very efficiently, but not in the way humans intended it. How do we stay in control of a complex intelligent system? The reason humans are on top of the food chain is not down to sharp teeth or strong muscles. Human dominance is almost entirely due to our ingenuity and intelligence. We can get the better of bigger, faster, stronger animals because we can create and use tools to control them: This poses a serious question about artificial intelligence: How do we define the humane treatment of AI? While neuroscientists are still working on unlocking the secrets of conscious experience, we understand more about the basic mechanisms of reward and aversion. We share these mechanisms with even simple animals. In a way, we are building similar mechanisms of reward and aversion in systems of artificial intelligence. For example, reinforcement learning is similar to training a dog: Right now, these systems are fairly superficial, but they are becoming more complex and life-like. Could we consider a system to be suffering when its reward functions give it negative input? This happens over many generations and is a way of improving a system. The unsuccessful instances are deleted. At what point might we consider genetic algorithms a form of mass murder? Should they be treated like animals of comparable intelligence? Will we consider the suffering of "feeling" machines? Some ethical questions are about mitigating suffering, some about risking negative outcomes. While we consider these risks, we should also keep in mind that, on the whole, this technological progress means better lives for everyone. Artificial intelligence has vast potential, and its responsible implementation is up to us.

Chapter 3 : History of artificial intelligence - Wikipedia

I'm a fan of that classic parody of parliamentary life Yes, Minister. Particularly the acerbic adviser Sir Humphrey whose response to his minister's ambitious vision and bold plans was inevitably "courageous decision Minister".

Precursors[edit] The dream of artificial intelligence was first thought of in Indian philosophies like those of Charvaka , a famous philosophy tradition dating back to BCE and surviving documents around BCE. The faithful believed that craftsman had imbued these figures with very real minds, capable of wisdom and emotionâ€" Hermes Trismegistus wrote that "by discovering the true nature of the gods, man has been able to reproduce it. The study of mechanicalâ€"or "formal"â€"reasoning has a long history. Chinese , Indian and Greek philosophers all developed structured methods of formal deduction in the first millennium BCE. For it would suffice to take their pencils in hand, down to their slates, and to say each other with a friend as witness, if they liked: In the 20th century, the study of mathematical logic provided the essential breakthrough that made artificial intelligence seem plausible. This photo has been artificially darkened, obscuring details such as the women who were present and the IBM equipment in use. First, they proved that there were, in fact, limits to what mathematical logic could accomplish. But second and more important for AI their work suggested that, within these limits, any form of mathematical reasoning could be mechanized. The Church-Turing thesis implied that a mechanical device, shuffling symbols as simple as 0 and 1, could imitate any conceivable process of mathematical deduction. The key insight was the Turing machine â€"a simple theoretical construct that captured the essence of abstract symbol manipulation. This invention would inspire a handful of scientists to begin discussing the possibility of thinking machines. In the early 19th century, Charles Babbage designed a programmable computer the Analytical Engine , although it was never built. Ada Lovelace speculated that the machine "might compose elaborate and scientific pieces of music of any degree of complexity or extent". The latter two of these machines were based on the theoretical foundation laid by Alan Turing [24] and developed by John von Neumann. In the s and 50s, a handful of scientists from a variety of fields mathematics, psychology, engineering, economics and political science began to discuss the possibility of creating an artificial brain. The field of artificial intelligence research was founded as an academic discipline in Cybernetics and early neural networks[edit] The earliest research into thinking machines was inspired by a confluence of ideas that became prevalent in the late s, s, and early s. Recent research in neurology had shown that the brain was an electrical network of neurons that fired in all-or-nothing pulses. The close relationship between these ideas suggested that it might be possible to construct an electronic brain. These machines did not use computers, digital electronics or symbolic reasoning; they were controlled entirely by analog circuitry. They were the first to describe what later researchers would call a neural network. If a machine could carry on a conversation over a teleprinter that was indistinguishable from a conversation with a human being, then it was reasonable to say that the machine was "thinking". This simplified version of the problem allowed Turing to argue convincingly that a "thinking machine" was at least plausible and the paper answered all the most common objections to the proposition. Symbolic reasoning and the Logic Theorist[edit] When access to digital computers became possible in the middle fifties, a few scientists instinctively recognized that a machine that could manipulate numbers could also manipulate symbols and that the manipulation of symbols could well be the essence of human thought. This was a new approach to creating thinking machines. Simon created the " Logic Theorist " with help from J. The proposal for the conference included this assertion: Simon , all of whom would create important programs during the first decades of AI research. The programs that were developed during this time were, to most people, simply "astonishing": Few at the time would have believed that such "intelligent" behavior by machines was possible at all. Among the most influential were these: Reasoning as search[edit] Many early AI programs used the same basic algorithm. To achieve some goal like winning a game or proving a theorem , they proceeded step by step towards it by making a move or a deduction as if searching through a maze, backtracking whenever they reached a dead end. This paradigm was called " reasoning as search ". Researchers would reduce the search space by using heuristics or " rules of thumb " that would eliminate those paths that were unlikely to lead to a solution. She simply gave a canned

response or repeated back what was said to her, rephrasing her response with a few grammar rules. ELIZA was the first chatterbot. They pointed out that in successful sciences like physics, basic principles were often best understood using simplified models like frictionless planes or perfectly rigid bodies. Much of the research focused on a " blocks world ," which consists of colored blocks of various shapes and sizes arrayed on a flat surface. At the same time, Minsky and Papert built a robot arm that could stack blocks, bringing the blocks world to life. It could communicate in ordinary English sentences, plan operations and execute them. Simon and Allen Newell: DARPA continued to provide three million dollars a year until the 70s. Licklider , then the director of ARPA , believed that his organization should "fund people, not projects! Its limb control system allowed it to walk with the lower limbs, and to grip and transport objects with hands, using tactile sensors. Its vision system allowed it to measure distances and directions to objects using external receptors, artificial eyes and ears. And its conversation system allowed it to communicate with a person in Japanese, with an artificial mouth. AI researchers had failed to appreciate the difficulty of the problems they faced. Their tremendous optimism had raised expectations impossibly high, and when the promised results failed to materialize, funding for AI disappeared. Even the most impressive could only handle trivial versions of the problems they were supposed to solve; all the programs were, in some sense, "toys". Although some of these limits would be conquered in later decades, others still stymie the field to this day. There was not enough memory or processing speed to accomplish anything truly useful. He suggested an analogy: Intractability and the combinatorial explosion. Finding optimal solutions to these problems requires unimaginable amounts of computer time except when the problems are trivial. This almost certainly meant that many of the "toy" solutions used by AI would probably never scale up into useful systems. Many important artificial intelligence applications like vision or natural language require simply enormous amounts of information about the world: This requires that the program know most of the same things about the world that a child does. Researchers soon discovered that this was a truly vast amount of information. No one in could build a database so large and no one knew how a program might learn so much information. Proving theorems and solving geometry problems is comparatively easy for computers, but a supposedly simple task like recognizing a face or crossing a room without bumping into anything is extremely difficult. This helps explain why research into vision and robotics had made so little progress by the middle s. AI researchers like John McCarthy who used logic discovered that they could not represent ordinary deductions that involved planning or default reasoning without making changes to the structure of logic itself. They developed new logics like non-monotonic logics and modal logics to try to solve the problems. The pattern began as early as when the ALPAC report appeared criticizing machine translation efforts. After spending 20 million dollars, the NRC ended all support. Hans Moravec blamed the crisis on the unrealistic predictions of his colleagues. Funding for the creative, freewheeling exploration that had gone on in the 60s would not come from DARPA. Instead, the money was directed at specific projects with clear objectives, such as autonomous tanks and battle management systems. Philosophy of artificial intelligence Several philosophers had strong objections to the claims being made by AI researchers. If the symbols have no meaning for the machine, Searle argued, then the machine can not be described as "thinking". Problems like intractability and commonsense knowledge seemed much more immediate and serious. It was unclear what difference " know how " or " intentionality " made to an actual computer program. Minsky said of Dreyfus and Searle "they misunderstand, and should be ignored. A feud began, and the situation was not helped when Colby did not credit Weizenbaum for his contribution to the program. In , Weizenbaum published Computer Power and Human Reason which argued that the misuse of artificial intelligence has the potential to devalue human life. Like most AI researchers, he was optimistic about their power, predicting that "perceptron may eventually be able to learn, make decisions, and translate languages. The effect of the book was devastating: Eventually, a new generation of researchers would revive the field and thereafter it would become a vital and useful part of artificial intelligence. Rosenblatt would not live to see this, as he died in a boating accident shortly after the book was published. Alan Robinson had discovered a simple method to implement deduction on computers, the resolution and unification algorithm. However, straightforward implementations, like those attempted by McCarthy and his students in the late s, were especially intractable: Simon that would lead to Soar and their unified theories of cognition. He argued

that what is really needed are machines that can solve problemsâ€”not machines that think as people do. Marvin Minsky , Seymour Papert and Roger Schank were trying to solve problems like "story understanding" and "object recognition" that required a machine to think like a person. In order to use ordinary concepts like "chair" or "restaurant" they had to make all the same illogical assumptions that people normally made. Unfortunately, imprecise concepts like these are hard to represent in logic. For example, if we use the concept of a bird, there is a constellation of facts that immediately come to mind: We know these facts are not always true and that deductions using these facts will not be "logical", but these structured sets of assumptions are part of the context of everything we say and think. He called these structures " frames ". Schank used a version of frames he called " scripts " to successfully answer questions about short stories in English. Boom â€”[edit] In the s a form of AI program called " expert systems " was adopted by corporations around the world and knowledge became the focus of mainstream AI research. In those same years, the Japanese government aggressively funded AI with its.

Chapter 4 : A Brief History of AI

The idea that human thinking and machine computing are "radically the same" provides the central theme for this marvelously lucid and witty book on.

But can we agree, at least for now, that having an AI determine your guilt or innocence in a court of law is a step too far? Worryingly, it seems this may already be happening. Loomis, who has a criminal history and was sentenced for having fled the police in a stolen car , now asserts that his right to due process was violated as neither he nor his representatives were able to scrutinise or challenge the algorithm behind the recommendation. The report was produced by a software product called Compas, which is marketed and sold by Nortpointe Inc to courts. The program is one incarnation of a new trend within AI research: The Wisconsin Supreme Court convicted Loomis, adding that the Compas report brought valuable information to their decision, but qualified it by saying he would have received the same sentence without it. But how can we know that for sure? Other courts are free to do the same. We also know that there are a number of AI startups that are competing to build similar systems. A similar trial is underway from the police in Antwerp, Belgium. However, past AI and deep learning projects involving massive data sets have been have been problematic. Benefits for the few? Technology has brought many benefits to the court room, ranging from photocopiers to DNA fingerprinting and sophisticated surveillance techniques. Algorithms can be racist, too. Why do we go through the trouble of selecting juries composed of our peers? The standard in law has never been one of perfection, but rather the best that our abilities as mere humans allow us. We make mistakes but, over time, and with practice, we accumulate knowledge on how not to make them again â€” constantly refining the system. This must be resisted forcefully. Legal systems depend on continuity of information, transparency and ability to review. What we do not want as a society is a justice system that encourages a race to the bottom for AI startups to deliver products as quickly, cheaply and exclusively as possible. An open source, reviewable version of Compas would be an improvement. However, we must ensure that we first raise standards in the justice system before we begin offloading responsibility to algorithims. AI should not just be an excuse not to invest. While there is a lot of money to be made in AI, there is also a lot of genuine opportunity. I have no perfect solutions for all these problems right now. But I do know that when it comes to the role of AI in law we must ask in which context they are being used, for what purposes and with what meaningful oversight. Until those questions can be answered with certainty, be very very sceptical. Or at the very least know some very good lawyers.

Chapter 5 : Go Moment Press â€" â€¦the very idea of AI, in a customer experience

The usual picture of Ives as the Transcendentalist composer par excel- lence simply cannot account for what is most basic to his music, nor for the origins of his complex aesthetic.

Fight of the 20th Century You need to have JavaScript enabled to view this clip. Supporters of top-down AI still had their champions: The IBM-built machine was, on paper, far superior to Kasparov - capable of evaluating up to million positions a second. But could it think strategically? The answer was a resounding yes. Some hailed this as the moment that AI came of age. But for others, this simply showed brute force at work on a highly specialised problem with clear rules. But Roomba was a big achievement. Despite relatively simple sensors and minimal processing power, the device had enough intelligence to reliably and efficiently clean a home. Roomba ushered in a new era of autonomous robots, focused on specific tasks. Having seen their dreams of AI in the Cold War come to nothing, the US military was now getting back on board with this new approach. They began to invest in autonomous robots. BigDog, made by Boston Dynamics, was one of the first. Built to serve as a robotic pack animal in terrain too rough for conventional vehicles, it has never actually seen active service. Their bomb disposal robot, PackBot, marries user control with intelligent capabilities such as explosives sniffing. Over PackBots have been deployed in Iraq and Afghanistan. In November , a small feature appeared on the new Apple iPhone â€" a Google app with speech recognition. But this heralded a major breakthrough. Google pioneered a new approach: Artificial intelligence would be the ultimate version of Google. It would understand exactly what you wanted, and it would give you the right thing. Find out how close we are to enabling robots to learn with mathematician Marcus Du Sautoy. At the same time as massive mainframes were changing the way AI was done, new technology meant smaller computers could also pack a bigger punch. These new computers enabled humanoid robots, like the NAO robot, which could do things predecessors like Shakey had found almost impossible. NAO robots used lots of the technology pioneered over the previous decade, such as learning enabled by neural networks. Fight of the 21st Century Watson is now used in medicine. This was a far greater challenge for the machine than chess. Watson had to answer riddles and complex questions. Its makers used a myriad of AI techniques, including neural networks, and trained the machine for more than three years to recognise patterns in questions and answers. Watson trounced its opposition â€" the two best performers of all time on the show. The victory went viral and was hailed as a triumph for AI. Across four states in America it is legal for driverless cars to take to the road. Sixty-four years after Turing published his idea of a test that would prove machine intelligence, a chatbot called Eugene Goostman finally passed. But very few AI experts saw this a watershed moment. It was other developments in that really showed how far AI had come in 70 years.

Chapter 6 : Why using AI to sentence criminals is a dangerous idea

In Artificial Intelligence: The Very idea, John Haugeland introduces four common responses people have to the very idea of artificial intelligence (AI).. Note: By "AI," Haugeland means GOFAI (Good Old Fashioned Artificial Intelligence).

Chapter 7 : Top 9 ethical issues in artificial intelligence | World Economic Forum

the very idea of AI, in a customer experience context, is to make the customer feel empowered.

Chapter 8 : 5 Big Predictions for Artificial Intelligence in - MIT Technology Review

Why using AI to sentence criminals is a dangerous idea May 16, am EDT. Until those questions can be answered with certainty, be very very sceptical. Or at the very least know some.

Chapter 9 : BBC - iWonder - AI: 15 key moments in the story of artificial intelligence

DOWNLOAD PDF AI THE VERY IDEA

At the very least, one hopes, the ideas and concepts being explored in these winning AI grants might help prevent some of the "unintended" and "disastrous" consequences hinted at by the.