

Chapter 1 : Benefits & Risks of Artificial Intelligence - Future of Life Institute

Welcome to the official website of the AIS. We are a student organization at UT Dallas with a mission to increase knowledge of, and greater interest in, artificial intelligence within the UT Dallas and local communities.

Governance and ethics in a world of AI raises extensive ethical concerns. Technologists themselves say the technology needs to align with human values, and that ethical dimensions must be prioritised at every stage of the design, development and deployment of AI systems. We will need focused research and effective governance structures to make sure AI technologies create opportunities and not harm. Algorithms are still being developed by people at this point; we have a bit of control of what we are doing. Are the intermediaries that we deal with going to be artificial intelligence? Academic, Europe It starts with the value of the human. Once we start giving power to machines, will that be tied to the metaphysical commitment to the human at the centre of governance? Civil Society, North America AI also raises serious considerations related to privacy, transparency, safety, the nature of work and jobs, and the overall economy. For example, technologies such as facial recognition based on AI can improve user experience over a social media platform. But the same technologies can be used to improve surveillance and compromise anonymity. Or, if AI becomes a permanent feature in social media networks and online platforms, where algorithms are used to curate the online experience, questions about free choice and bias will intensify. Concerns about the transparency and accountability of data collection and decision-making will accelerate calls for ethical principles to guide AI design and deployment. A society completely based on data collection on the business side fuels a surveillance society without proper democratic checks and balances. Humans lose some self-determination through automated choices by connected machines. Human Rights Expert, Europe Our community, across all stakeholder groups and regions, believes that automation generated through data analytics technology will have greater influence on human behaviour and decision-making. Within governments, AI could bring about a fundamental reshaping of decision-making as policy development is increasingly data driven. By extension, there is a risk that AI could become an unaccountable and non-transparent decision-making tool for future policy choices. The development of IoT and AI will provide scientific references for government decision-making and help them to respond quickly to public needs. Technologist, Asia-Pacific Many foresee a fierce, competitive battle to dominate the commercial AI space in coming years. While this will likely drive innovation and possibly disrupt current market structures, there are also concerns about competition. For once, they were right, especially about the legal professions. After the global financial crisis Keep reading Impact of AI on the Internet economy While some argue that projections about AI are simply marketing hype, there is a clear focus by many in industry and government on preparing for a future in which AI is pervasive. As AI and automation drive significant structural change across industries, the very nature of work will change. Many existing jobs may be displaced as AI moves beyond monetising user data to changing how products and services are delivered. Adapting to the pace of change will be a major global challenge for the immediate future. Projects relating to Artificial Intelligence and the Internet of Things have gone a long way in advancing our present technology and making life easier for the average human being. Internet Society member, Africa This said, AI systems and technologies could also change the nature of work in a way that empowers workers, diminishing inequality among people and between countries. AI can be a partner in human intelligence, letting us take on and solve much bigger challenges. Machine to machine communication increases pressures to cut costs and people are being replaced. This is only going to increase with time. There are economic benefits but also challenges to employment. Private Sector, Middle East AI brings potential for huge gains in scientific research, transportation and the delivery of services. If accessibility and open source development win out, AI has the potential to bring dividends to developed and developing countries alike. For example, a country that relies on agricultural production could use AI to analyse crop yields and optimise food production. AI applications in healthcare could be a game changer for disease detection in low income areas. Artificial intelligence will be creative destruction. Many jobs will be eliminated by AI, but it can generate new roles and jobs. For developing economies, new technologies always create possibilities to leapfrog legacy systems,

although the infrastructure requirements for AI and IoT to be deployed will be significant. The benefits of AI may also be unevenly distributed: The intelligence and services used to manage and implement manufacturing or services may still reside in developed countries rather than being developed locally. AI might exacerbate the digital divide in significant ways that would have geopolitical implications. Ensuring Internet technology can create jobs in the market and does not harm the job market is a challenge that has to [be] addressed in the next 5 years and it is an urgent and serious problem internationally.

Chapter 2 : OxAI - Oxford Artificial Intelligence Society

Over the next five years, we are about to witness the world we live in entirely disrupted by improvements in artificial intelligence (AI) and machine learning.

Russian What is AI? Artificial intelligence today is properly known as narrow AI or weak AI , in that it is designed to perform a narrow task e. Why research AI safety? In the long term, an important question is what will happen if the quest for strong AI succeeds and an AI system becomes better than humans at all cognitive tasks. As pointed out by I. Good in , designing smarter AI systems is itself a cognitive task. How can AI be dangerous? Most researchers agree that a superintelligent AI is unlikely to exhibit human emotions like love or hate, and that there is no reason to expect AI to become intentionally benevolent or malevolent. The AI is programmed to do something devastating: Autonomous weapons are artificial intelligence systems that are programmed to kill. In the hands of the wrong person, these weapons could easily cause mass casualties. The AI is programmed to do something beneficial, but it develops a destructive method for achieving its goal: A key goal of AI safety research is to never place humanity in the position of those ants. Why the recent interest in AI safety Stephen Hawking, Elon Musk, Steve Wozniak, Bill Gates, and many other big names in science and technology have recently expressed concern in the media and via open letters about the risks posed by AI , joined by many leading AI researchers. Why is the subject suddenly in the headlines? The idea that the quest for strong AI would ultimately succeed was long thought of as science fiction, centuries or more away. However, thanks to recent breakthroughs, many AI milestones, which experts viewed as decades away merely five years ago, have now been reached, making many experts take seriously the possibility of superintelligence in our lifetime. While some experts still guess that human-level AI is centuries away, most AI researches at the Puerto Rico Conference guessed that it would happen before Since it may take decades to complete the required safety research, it is prudent to start it now. Because AI has the potential to become more intelligent than any human, we have no surefire way of predicting how it will behave. The best example of what we could face may be our own evolution. Timeline Myths The first myth regards the timeline: A common misconception is that we know the answer with great certainty. In fact, history is full of technological over-hyping. AI has also been repeatedly over-hyped in the past, even by some of the founders of the field. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer. All these surveys have the same conclusion: For example, in such a poll of the AI researchers at the Puerto Rico AI conference , the average median answer was by year , but some researchers guessed hundreds of years or more. Many of the safety problems associated with human-level AI are so hard that they may take decades to solve. When Stuart Russell, author of the standard AI textbook , mentioned this during his Puerto Rico talk , the audience laughed loudly. A related misconception is that supporting AI safety research is hugely controversial. After all, fear sells, and articles using out-of-context quotes to proclaim imminent doom can generate more clicks than nuanced and balanced ones. That scenario combines as many as three separate misconceptions: If you drive down the road, you have a subjective experience of colors, sounds, etc. But does a self-driving car have a subjective experience? Does it feel like anything at all to be a self-driving car? If you get struck by a driverless car, it makes no difference to you whether it subjectively feels conscious. The fear of machines turning evil is another red herring. A superintelligent AI is by definition very good at attaining its goals, whatever they may be, so we need to ensure that its goals are aligned with ours. The beneficial-AI movement wants to avoid placing humanity in the position of those ants. Machines can obviously have goals in the narrow sense of exhibiting goal-oriented behavior: If you feel threatened by a machine whose goals are misaligned with yours, then it is precisely its goals in this narrow sense that troubles you, not whether the machine is conscious and experiences a sense of purpose. To cause us trouble, such misaligned superhuman intelligence needs no robotic body, merely an internet connection â€” this may enable outsmarting financial markets, out-inventing human researchers, out-manipulating human leaders, and developing weapons we cannot even understand. Even if building robots were physically impossible, a super-intelligent and super-wealthy AI could easily pay or manipulate many

humans to unwittingly do its bidding. What sort of future do you want? Should we develop lethal autonomous weapons? What would you like to happen with job automation? Do you prefer new jobs replacing the old ones, or a jobless society where everyone enjoys a life of leisure and machine-produced wealth? Further down the road, would you like us to create superintelligent life and spread it through our cosmos? Will we control intelligent machines or will they control us? Will intelligent machines replace us, coexist with us, or merge with us? What will it mean to be human in the age of artificial intelligence? What would you like it to mean, and how can we make the future be that way? Please join the conversation!

Chapter 3 : McGill Artificial Intelligence Society

Benefits & Risks of Artificial Intelligence " Everything we love about civilization is a product of intelligence, so amplifying our human intelligence with artificial intelligence has the potential of helping civilization flourish like never before - as long as we manage to keep the technology beneficial.

Artificial Intelligence should benefit society, not create threats January 16, by Toby Walsh, The Conversation Science fiction has plenty of tales of AI turning against society including the popular Terminator movie franchise, here depicted on brick wall art. I am not sure the famous British theoretical physicist Stephen Hawking does irony but it was somewhat ironic that he recently welcomed the arrival of the smarter predictive computer software that controls his speech by warning us that: The development of full artificial intelligence could spell the end of the human race. Of course, Hawking is not alone in this view. The serial entrepreneur and technologist Elon Musk also warned last year that: Both address an issue that taps into deep, psychological fears that have haunted mankind for centuries. What happens if our creations eventually cause our own downfall? An open letter for AI In response to such concerns, an open letter has just been signed by top AI researchers in industry and academia as well as by Hawking and Musk. Signatures include those of the president of the Association for the Advancement of Artificial Intelligence , the founders of AI startups DeepMind and Vicarious , and well-known researchers at Google, Microsoft, Stanford and elsewhere. In the interests of full disclosure, mine is also one of the early signatures on the list, which continues to attract more support by the day. The open letter argues that there is now a broad consensus that AI research is progressing steadily and its impact on society is likely to increase. For this reason, the letter concludes we need to start to research how to ensure that increasingly capable AI systems are robust in their behaviours and beneficial to humans. For example, we need to work out how to build AI systems that result in greater prosperity within society, even for those put out of work. The letter includes a link to a document outlining some interdisciplinary research priorities that should be tackled in advance of developing artificial intelligence. These include short-term priorities such as optimising the economic benefits and long-term priorities such as being able to verify the formal properties of AI systems. A Space Odyssey through Robocop and Terminator to recent movies such as Her and Transcendence , all of which paint a dystopian view of a future transformed by AI. AI researchers have been predicting it will take another 30 or 40 years now for the last 30 or 40 years. And if you ask most of them today, they as I will still say it is likely to take another 30 or 40 years. Making computers behave intelligently is a tough scientific nut to crack. The human brain is the most complex system we know of by orders of magnitude. Replicating the sort of intelligence that humans display will likely require significant advances in AI. The human brain does all its magic with just 20 watts of power. This is a remarkable piece of engineering. Other risks to society There are also more imminent dangers facing mankind such as climate change or the ongoing global financial crisis. These need immediate attention. The Future of Humanity Institute at the University of Oxford has a long list of threats besides AI that threaten our society including: The AI debate for the future The Campaign to Stop Killer Robots is advancing the debate on whether we need to ban fully autonomous weapons. I am organising a debate on this topic at the next annual conference of the Association for the Advancement of Artificial Intelligence later this month in Austin, Texas, in the US. Another issue that requires more immediate attention is the impact that AI will have on the nature of work. How does society adapt to more automation and fewer people needed to work? If we can get this right, we could remove much of the drudgery from our lives. If we get it wrong, the increasing inequalities documented by the French economist Thomas Piketty will only get worse.

Chapter 4 : Aies Conference – Conference on Artificial Intelligence, Ethics and Society

Artificial intelligence helps farmers, doctors and rescue workers make a positive impact on society. Artificial intelligence (AI) powers many gadgets, like smartphones, smart thermostats and voice-activated virtual assistants that bring modern conveniences to daily life.

Situated in the heart of Montreal, a city striving to become a world class leader in artificial intelligence, we aim to be a hub that unifies all those involved in this community. Become A Member Learn We hold bi-weekly tutorials that cover different aspects of AI from computer vision to natural language processing and everything in between. Share With student, researchers, and industry professionals all within the same community, knowledge is constantly being shared through tech talks, research presentations, and our ongoing blog. Upcoming Timeline Workshop 1 - Intro to ML September 13 Perfect for beginners, our workshop leads will be giving an introduction to machine learning with a hands on tutorial. Click here for more details. ImplementAI September 29, 30 Join us for our annual hackathon where AI researchers, companies, and students will come together to explore the implementation of the latest research in AI. He also enjoys short walks on the beach and is the master of programming puns. His interests currently reside in reinforcement learning for text-based problems. Her primary academic interest is in natural language processing. She hopes to one day help different industries apply AI to improve business solutions. He is interested in computer vision and pattern recognition. Ketan Rampurkar VP Finance Ketan is a 3rd year Software Eng student with an interest in combining finance and tech to make some mind-boggling software projects. She is passionate about physics and newly developed interest in computer related topics. She wants to apply an interdisciplinary approach to further understand AI and quantum computing. He enjoys solving algorithms and backend development. But sometimes you will hear him say "border-radius" and "margin" Tiffany Wang Design Lead Tiffany is a fourth year Electrical Engineering student. His interests include human dynamics, behavioural science and affective AI. Frank Ye Workshop Lead Frank is a fourth year electrical engineering student. Apart from interests in computer vision and AI, he unhealthily enjoys Tetris, anime trap remixes and late-night walks with friends. Isaac Chan Workshop Lead Isaac is a fifth year electrical engineering student. He is interested in understanding the mechanics of deep reinforcement learning in the context of few-shot learning and HRL Claudia Leung Editorial Lead Claudia is a third year Psychology and Political Science student with broad academic interests in speech and language pathology, neuropsychology, and natural language processing. He is interested in applying AI and Computer Vision techniques in medical applications. Our Sponsors All club operations are funded by our generous sponsors. Without them, we would not be able to do what we do.

Chapter 5 : Three Ways Artificial Intelligence is Good for Society - iQ by Intel

Our mission is to conduct research in Artificial Intelligence to help solve the most difficult social problems facing our world.

Well-trained AIs are now better at spotting breast tumours and other cancers than radiologists. Does that mean widespread unemployment for radiologists? It is not so straightforward, says Stiglitz. Mostly these are low-skilled roles: Again, though, Stiglitz sees reasons to be cautious about what that will mean for overall unemployment. There is a strong demand for unskilled workers in education, the health service and care for older people. If AI takes over certain unskilled jobs, the blow could be softened by hiring more people into health, education and care work and paying them a decent wage, he says. Stiglitz won the Nobel prize for economics in for his analyses of imperfect information in markets. Trade negotiations, he argued, were driven by multinationals at the expense of workers and ordinary citizens. Armed with AI, tech firms can extract meaning from the data we hand over when we search, buy and message our friends. It is used ostensibly to deliver a more personalised service. That is one perspective. Another is that our data is used against us. Now you can target particular individuals by exploiting their information. For example, retailers can now track customers via their smartphones as they move around stores and can gather data on what catches their eye and which displays they walk straight past. That means that life is going to be increasingly unpleasant, because your decision to shop in a certain store may result in you paying more money. To the extent that people are aware of this game, it distorts their behaviour. What is clear is that it introduces a level of anxiety in everything we do and it increases inequality even more. With the new AI, it looks like the answer is finding a better way to exploit somebody. Stiglitz is concerned that companies are using, or will use, similar tactics to exploit their customers, in particular those who are vulnerable, such as compulsive shoppers. One of the many reasons is that the complexity of the technology can make it intimidating.

Chapter 6 : Artificial Intelligence | Internet Society

Artificial intelligence is a technology that is already impacting how users interact with, and are affected by the Internet. In the near future, its impact is likely to only continue to grow.

Artificial intelligence AI powers many gadgets, like smartphones, smart thermostats and voice-activated virtual assistants that bring modern conveniences to daily life. Increasingly, AI is also being used to tackle critical social challenges. AI is a branch of computer science where machines can sense, learn, reason, act and adapt to the real world, amplifying human capabilities and automating tedious or dangerous tasks. Some experts believe AI has the potential to spark a serious social revolution. Still, many tech innovators – including physicist Stephen Hawking, Microsoft founder Bill Gates and Tesla founder Elon Musk – caution that humanity could lose control of superintelligent machines, and AI could cause more harm than good. The evolution of a master machine race has been debated since the mids when AI research began. The late Hubert L. Dreyfus , a professor of philosophy at the University of California at Berkeley, challenged the lofty expectations for AI, arguing that machines lack the intuition to compete with human intelligence. In the past three decades, attention and energy poured into AI has steadily increased. Many AI researchers agree that AI can be smart without being sentient, which gets to the heart of the fear of the new technology: She said AI is taking off because of three key elements: At a time when the agricultural land available for farming is shrinking, farmers will need to grow approximately 50 percent more crops. Enter FarmLogs , a farming management app currently used by one out of three farmers in America. FarmLogs uses data and technology to help farmers monitor fields, track the weather and get insights into soil using historical satellite imagery to calculate irregular plant growth. Improving Cancer Diagnosis Cancer is a frightening diagnosis, impacting 1. Waiting for biopsy results can be stressful, but AI may help accelerate the diagnosis and treatment process. Working with healthcare industry leaders, by Intel aims to create one-day precision medicine for cancer patients – that means going to the doctor, getting a diagnosis and receiving a personalized treatment plan, all in 24 hours. The project is driven in part by prostate cancer survivor Bryce Olson , a global marketing director for Intel whose cancer went into remission after he used targeted molecular testing to create a custom treatment plan. In alone, its CyberTipline received 8. While technology may have helped lead to this problem, AI may be part of the solution. And AI helps automate and speed up the process. While fear of the negative consequences remain, AI is proving it can bring about enormous societal benefits.

Chapter 7 : Association for the Advancement of Artificial Intelligence

If you haven't been paying attention, artificial intelligence is the new player on the field, and it's steadily taking names as it forges ahead into every realm of industry. But is the.

Chapter 8 : Artificial Intelligence should benefit society, not create threats

The Future is Here, the Future is AI. We are a society dedicated to the further development of AI technologies and adoption in Hong Kong. Join our monthly meetups and learn more about how Artificial Intelligence and Chatbots can help you grow your business.

Chapter 9 : Artificial Intelligence - Internet Society Digital Future Report

Association for the Advancement of Artificial Intelligence. Association for Computing Machinery. ACM Special Interest group on Artificial Intelligence.