

Man and Machine: Questions of Risk, Trust and Accountability in Today's AI Technology

Piyush Ahuja¹
Indian Institute of Technology Delhi

August 12, 2014

¹The author graduated from IIT Delhi in 2013 with a major in Mathematics and Computing. His current research straddles the fields of Algorithms and Economics.

Abstract

Artificial Intelligence began as a field probing some of the most fundamental questions of science - the nature of intelligence and the design of intelligent artifacts. But it has grown into a discipline that is deeply entwined with commerce and society. Today's AI technology, such as expert systems and intelligent assistants, pose some difficult questions of risk, trust and accountability. In this paper, we present these concerns, examining them in the context of historical developments that have shaped the nature and direction of AI research. We also suggest the exploration and further development of two paradigms, human intelligence-machine cooperation, and a sociological view of intelligence, which might help address some of these concerns.

1 Introduction

*“Synthesis defines an ambitious put-a-man- on-the-moon goal. By doing so, it forces scientists and engineers to cross uncharted terrain in pursuit of the goal. This requires the solution of unscripted problems that are not normally encountered through either observation or analysis... synthesis drives the evolution of paradigms”*¹

Artificial Intelligence (AI) is the field of computer science which aims to create, or synthesize, intelligence. The pursuit of creating intelligent machines has contributed not only to psychology, cognitive science, neurology and philosophy, but has also given birth to whole new branches of research.

The field of AI research was formally founded at a conference on the campus of Dartmouth College in the summer of 1956.² It has come a long way since then, having gone through many cycles of boom and bust, “AI winters” and summers.³ Now AI applications have left the annals of Department of Defense R&D, and trickled down to everyday use. They can be found in common place consumer items and inexpensive intelligent toys, even though very often consumers fail to recognize the technology source. For example, the Kinect, which provides a 3D bodymotion interface for the Xbox 360, uses algorithms that emerged from lengthy AI research.⁴

The adoption and integration of AI based technology in all spheres has followed a pattern reminiscent of many modern technologies, including, for example, the internet, digital computing,⁵ and mobile telephony. The pace

¹It is hard to understate the role of synthesis in science and technological development. The above quote has been taken from the context of *synthetic biology*, see Steven A Benner and A Michael Sismour. “Synthetic biology”. In: *Nature Reviews Genetics* 6.7 (2005), pp. 533–543.

²The conference was chaired by J.McCarthy, M. Minsky, N. Rochester and C.Shannon. For a full report, see John McCarthy et al. “A proposal for the dartmouth summer research project on artificial intelligence, august 31, 1955”. In: *AI Magazine* 27.4 (2006), p. 12.

³Nils J Nilsson. *The quest for artificial intelligence*. Cambridge University Press Cambridge, 2010.

⁴Jan Smisek, Michal Jancosek, and Tomas Pajdla. “3D with Kinect”. In: *Consumer Depth Cameras for Computer Vision*. Springer, 2013, pp. 3–25; Kourosh Khoshelham and Sander Oude Elberink. “Accuracy and resolution of kinect depth data for indoor mapping applications”. In: *Sensors* 12.2 (2012), pp. 1437–1454; Iason Oikonomidis, Nikolaos Kyriazis, and Antonis A Argyros. “Efficient model-based 3D tracking of hand articulations using Kinect.” In: *BMVC*. 2011, pp. 1–11.

⁵The reader is referred to James Cortada’s essay in *Technology and Culture’s* April

has been such that historians and sociologists have barely had an opportunity to study its arrival, use and implications for society. Most discussions in this direction have been based on a priori and futuristic assessments, discussing the future of AI technology and what it might eventually turn out like. For instance, researchers have raised concerns about AI machines becoming malicious (unfriendly), apathetic or uncontrollable, often evoking images of killer robots and future wars between man and machine.⁶ The importance of incorporating ‘friendliness’ in AI research is repeatedly stressed in these discussions. Some have focussed on the concept of ‘intelligence explosion’, forewarning a future event where radically self-improving machines reach a state where it is impossible to predict or comprehend their actions.⁷ Others have noted that increasing dependence on decision-making intelligent machines may itself lead to a world where it is impossible for humans to survive without them, leaving machines effectively in control.⁸

Such discussions, though very important in their own right, tend to take the attention away from issues of risk and trust posed by AI-based technology which has already diffused in the society. In this essay, we argue that current AI technology (e.g., expert systems and intelligent assistants) is based on a notion of intelligence which is somewhat different from the notion of ‘general intelligence’ in popular perception. The intelligence or expertise of these

2013 issue. James W Cortada. “How New Technologies Spread: Lessons from Computing Technologies”. In: *Technology and Culture* 54.2 (2013), pp. 229–261.

⁶Bill Joy. “Why the future doesn’t need us.” In: *Nanoethics—the ethical and social implications of nanotechnology* (2000), pp. 17–39; Nick Bostrom, Milan M Cirkovic, and Martin J Rees. *Global catastrophic risks*. Oxford University Press Oxford, 2008; Eliezer Yudkowsky. “Artificial intelligence as a positive and negative factor in global risk”. In: *Global catastrophic risks* 1 (2008), p. 303; Nick Bostrom. “Existential risks”. In: *Journal of Evolution and Technology* 9 (2002); Nick Bostrom. *Ethical issues in advanced artificial intelligence*. 2003.

⁷Luke Muehlhauser and Anna Salamon. “Intelligence explosion: Evidence and import”. In: *Singularity Hypotheses*. Springer, 2012, pp. 15–42.

⁸In his now famous article “Why The Future Doesn’t Need Us”, computer scientist Bill Joy quotes Ted Kaczynski, “As society and the problems that face it become more and more complex and machines become more and more intelligent, people will let machines make more of their decisions for them, simply because machine-made decisions will bring better results than man-made ones. Eventually a stage may be reached at which the decisions necessary to keep the system running will be so complex that human beings will be incapable of making them intelligently.” See Joy, op. cit.; Theodore Kaczynski et al. *The unabomber manifesto: Industrial society and its future*. Filiquarian Publishing, LLC., 2005.

machines is measured through their performance in certain specialized contexts, and little else. Further, this confusion of what AI technology is based on, and how it works, itself has risk and trust-related consequences. As AI researcher Elizier Yudkowsky had noted, “By far the greatest danger with Artificial Intelligence is that people conclude too early that they understand it”.⁹

We start off this essay by tracing the history of current AI research in Section 2. Various factors which have steered AI research in a particular direction, and diminished in others, have been outlined. Section 3 examines the nature of AI systems today, illustrating, through the examples of expert systems and intelligent assistants, the questions of risk and trust they pose to society. In Section 4, we suggest two paradigms, which might complement basic research in AI to address some of these concerns . The discussion is summarised in the conclusion.

2 Historical Developments

Ever since AIs inception, there has been no clear consensus on what constitutes ‘intelligence’. The subject has drawn from a broad array of disciplines - Philosophy, Logic, Biology, Psychology, Statistics and Engineering. In the absence of an agreed-on curriculum for training students in AI,¹⁰ new researchers who enter the field bring with them different standards, traditions and problems. As a result, one big challenge that AI has faced is that the research effort has been characterized by a multiplicity of approaches, each endeavoring to attain some specific objective . This multigoal character of AI research has crystallized into its theoretical pluralism, and its institutionalization by means of competing groups with different aims¹¹. It has been a major factor in the eventual branching of Artificial Intelligence to various subfields, e.g., Knowledge Representation, Machine Learning and Natural Language Processing.

Up until the early 1970s, AI researchers dealt with highly theoretical problems probing the nature of intelligence, and the researchers pursued

⁹Yudkowsky, op. cit.

¹⁰Nilsson, op. cit.

¹¹Ronald David Schwartz. “Artificial intelligence as a sociological phenomenon”. In: *Canadian Journal of Sociology/Cahiers canadiens de sociologie* (1989), pp. 179–202, Fleck, 1982: 172 in.

projects that were staged in highly controlled laboratory settings.¹² The heavy initial funding (by DARPA) into such pathbreaking research was fuelled partly due to the cold war,¹³ and partly due to the highly optimistic claims made by its pioneers and their early theoretical successes.

However, such theoretical approaches to AI gradually came to be thought off by many computer scientists as fringe activities that did not adhere to rigorous scientific standards - some even viewed AI as ‘a field that housed charlatans’.¹⁴ Niel Nilsson, one of the founding fathers of the discipline, recalls that when he first interviewed for a position at SRI in 1961,¹⁵ a researcher had warned him against joining research on neural networks because it was ‘premature’, and his involvement with it could damage his reputation. This concern for ‘respectability’ had a stultifying effect on many AI researchers.¹⁶

In 1969, the Mansfield Amendment dealt another significant blow to the field. The Amendment put DARPA under increasing pressure to fund only ‘mission-oriented direct research, rather than basic undirected research’.¹⁷ The creative, freewheeling exploration that was characteristic of the early pioneering work in AI gradually came to be viewed as a burden. Instead, the money was directed at specific projects with clear objectives, such as autonomous tanks and battle management systems.¹⁸ Not only did this greatly influence the direction of research, but also, perhaps more importantly, swayed the spirit that guided work in the field. This has been noted by Marvin Minsky,¹⁹ co-founder of MIT’s AI Lab and a leading cognitive scientist:

¹²Nilsson, op. cit., p 265.

¹³Blay Whitby. *Reflections on Artificial Intelligence*. Intellect (UK), 1996, p. 66.

¹⁴Nilsson, op. cit., p 339.

¹⁵Stanford Research Institute

¹⁶It is interesting to note that many researchers in AI today deliberately call their work by other names, such as informatics, knowledge-based systems, cognitive systems or computational intelligence. One of the reasons is that these names help to procure funding. As New York Times(John Markoff. “Behind Artificial Intelligence, a Squadron of Bright Real People”. In: *New York Times* 14 [2005]) reported in 2005, “Computer scientists and software engineers avoided the term artificial intelligence for fear of being viewed as wild-eyed dreamers”

¹⁷National Research Council (US). Committee on Innovations in Computing, Communications, and Lessons from History. *Funding a revolution: government support for computing research*. National Academies Press, 1999, under ”Shift to Applied Research Increases Investment” (only the sections before 1980 apply to the current discussion).

¹⁸Ibid., under “Shift to Applied Research Increases Investment”.

¹⁹Wade Roush. “Marvin minsky on common sense and computers that emote”. In: *Technology Review*, July 13 (2006).

“In the early days, DARPA supported people rather than proposals. There was a lot of progress from starting in 1963 for about ten years in all branches and all approaches aimed to modeling intelligence. But the Mansfield Agreement made it much harder to support visionary researchers. At the same time, the American corporate research community started to disappear in the early 1970s. Bell Labs and RCA and the others essentially disappeared from this sort of activity.”

Also, by the end this period, the power of AI methods had already increased to the point where realistic applications seemed within reach.²⁰ This gave rise to what Minsky calls ‘the entrepreneur bug’.²¹ He attributes the disappearance of young scientists in that period to an increased tendency to patent things, start start-ups and make new products. Support for original theoretical research in areas like commonsense reasoning eventually fizzled out.

The collective consequence of all these factors was that by the end of the 70s many had people diverted to highly specialized subfields that solved specific real-world problems. These problems ranged from Speech Recognition and Understanding Systems, Consulting Systems, Understanding Queries and Signals to Computer Vision.²² Since these also had commercial applications, funding and notions of academic respectability became directly associated with these specialized problems.

Most AI researchers around this period eventually adopted the premise that general human-like intelligence can be developed by combining the programs that solve various subproblems using an integrated agent architecture. This was the Intelligent Agent paradigm, an approach that had become widely accepted in the AI community by the end of the 90s.²³

An intelligent agent is a system that perceives its environment and takes actions which maximize its chances of success.²⁴ The simplest intelligent

²⁰Nilsson, op. cit., p 265.

²¹Roush, op. cit.

²²Nilsson, loc. cit.

²³David Poole, Alan Mackworth, and Randy Goebel. *Computational Intelligence*. Oxford University Press Oxford, 1998, pp. 27, 3258, 968972; George F Luger. *Artificial intelligence: Structures and strategies for complex problem solving*. Pearson education, 2005, pp. 721; Stuart Russell. *Artificial intelligence: A modern approach, 2/E*. Pearson Education India, 2003, pp. 235240.

²⁴The definition used here is due to Russell and Norvig, see idem, *Artificial intelligence: A*

agents are programs that solve specific problems. The paradigm gave researchers license to study isolated problems and find solutions with actual applications, without agreeing on one single approach.²⁵ For perhaps the first time in AI, a partial consensus on a notion of intelligence was achieved, since it was well aligned with the pluralistic character of AI, the constraints of funding, the notion of academic respectability, as well as relevant to industrial applications. Many fundamental questions related to intelligence - symbolic and commonsense reasoning, intuition, imagination, creativity, and emotional intelligence, were left answered and unexplored.

3 AI Today

“Today’s AI is about new ways of connecting people to computers, people to knowledge, people to the physical world, and people to people”. - Patrick Winston, MIT AI Lab briefing, 1997

3.1 All That Glitters

Today, with the increased power of relatively inexpensive computers, availability of large databases and growth of the World Wide Web, AI technology like *expert systems* and *intelligent assistants* have slowly diffused into our society. These have been said to heralding a new age and revolutionizing the way we live today. Such proclamations might be true, but not necessarily in the context in which they are made. Most people overlook the fact that most of these systems do not work on the premise of approximating human cognitive abilities, but on performing certain tasks in highly specialized scenarios. They do well in performing certain jobs that are complicated for people to do, but they are far from being capable of carrying out tasks that are simple for humans to do. Many elements, like common-sense reasoning, ethical judgements and decision-making under uncertainties, which might be core components of intelligence, are missing. The deployment of such systems

modern approach, 2/E, Other definitions also include knowledge and learning as additional criteria.

²⁵An agent that solves a specific problem can use any approach that works some agents are symbolic and logical, some are sub-symbolic neural networks and others may use new approaches.

and their celebration, besides diverting research from other fields, has social, ethical and risk - related consequences.

Performance in specialized contexts (e.g., processing natural language, playing chess, recognizing patterns) are the sole basis for intelligence for current AI based systems. There is a certain gap in what these systems are represented as doing, and what they actually do. This can mislead people on two levels. One, that they create a notion that they work in the same way that humans do is (human-experts in the case of expert systems and assistants in case of digital assistants). On a second level, they create the impression that these are free from prejudices, personal bias and errors that might affect a human counterpart. The actual implementation and mechanism - what goes behind an AI technology - is entirely invisible to the user. The knowledge engineering, as it is called in the case of expert systems, and the underlying architecture, might involve restrictions and selective representation of data, and unavoidable biases resulting from the organization of data and architecture of the machine.

3.2 Risk

Another, possibly graver issue is that of handling of uncertainties by AI systems. Inductive knowledge of computer systems is inadequate. The number of combinations of possible inputs and internal states of a computer system of any complexity is huge, and especially in AI technologies, such complexity is intractable. Even with highly automated testing, it will seldom be feasible to exercise each and every state of a system to check for errors or underlying design faults, or bugs that may have caused them. As computer scientist Edsger Dijkstra famously put it in 1969, "Program Testing can be used to show the presence of bugs, but never to show their absence!" Because bugs may lurk for years before becoming manifest as a system failure, no guarantees can be offered.

In the case of expert systems, although they do reason usefully and economically about specific problems in medicine, geology, chemistry and other delimited areas, they are acknowledged to be brittle (that is, they break down) when confronted with problems outside their area of expertise or even on problems within their area of expertise, if knowledge were needed that had not been provided in their rulebooks. They don't know what they don't know, and therefore might provide wrong answers in cases where a human expert would do better.

A colorful anecdote involving John McCarthy²⁶ illustrates this. In an interaction with the medical expert system MYCIN, he typed in some information about a hypothetical patient, saying that he was male and also saying that he underwent amniocentesis. MYCIN accepted all that without complaint! That male patients don't get pregnant was not considered part of the 'expert knowledge' that MYCIN needed to be given.²⁷

Donald Michie, a British researcher, examines this problem through his concept of the 'human window' on the reasoning of the program, where its behaviour is 'scrutable'. Outside of this window, Michie observes, it is impossible to tell whether the program is being exceedingly clever or is just malfunctioning.²⁸

Reliance on experts is said to be an inevitable aspect of high modernity.²⁹ As we enter an age of knowledge-intensive 'information society', how we employ expert systems based on AI will turn out to be crucial in shaping our society.

The handicaps involving AI technology, as well as their associated risks, are greatly amplified in those systems that are critical to human safety or security. A disturbing example of this is provided by Donald MacKenzie,³⁰ a Professor of Sociology at the University of Edinburgh:

“On October 5, 1960, United States’ Ballistic Missile Early Warning System (BMEWS) went off, indicating several missile launches from a general area in Siberia, and sent everyone in a panic mode. A level 5 alarm meant 99.9 percent likelihood that a missile attack had been launched. If that were true, ellipses should have been forming on the war rooms display map of North America in Colorado Springs and should have started to shrink, indicating the target of the attack. Yet no ellipses were being formed, and the minutes-to-go indicator showed nothing. Soviet Union premier Nikita Khrushchev was in New York, attending the General Assembly of the UN. Deputy General Slemon decided

²⁶John McCarthy was one of the founding fathers of the field. He purportedly coined the term 'Artificial Intelligence'

²⁷Nilsson, op. cit., p 407.

²⁸Schwartz, op. cit., Michie, 1984 in.

²⁹Ibid.

³⁰Donald MacKenzie. *Mechanizing proof: computing, risk, and trust*. The MIT Press, 2004.

that Soviet Union was highly unlikely to attack the US while its leader was in New York. It was characteristic human reasoning. He also knew that the BMEWS had been operational for only four days and was still being 'run-in.' So, no action was eventually taken. It was later found out that the powerful radars in Greenland, designed to detect objects upto 3,000 miles distant, were fooled by the reflections from the slow-rising moon over Norway. The BMEWS interpreted the radar echoes as sightings of multiple objects, and the engineers had left this phenomenon unaccounted for while designing the system."

Since expert systems are deemed to emulate the decision-making ability of a human expert, and intelligent assistants work autonomously, human beings are increasingly being decoupled from the inner workings and implementations of such technology. In the case of the 1960 nuclear false alarm, human beings had remained in the loop, and common-sense had prevailed. With the increasing deployment of AI technology everywhere, it might not always be the case.

In his conception of the 'risk society', contrasting the nineteenth-century and the present, German sociologist Ulrich Beck³¹ had noted, "hazards in those days assaulted the nose or the eyes and were thus perceptible to the senses, while the risk of civilization today escapes perception." This idea, apart from its associations with nuclear technology and biotechnology, manifests itself in the domain of AI, and should be read in its context as well. The current AI technologies pose dangers of an entirely different kind - involving invisible contingencies, the seriousness of which laypersons eyes cannot judge.³²

³¹Ulrich Beck. *Risk society: Towards a new modernity*. Vol. 17. SAGE Publications Limited, 1992, p. 21.

³²The Y2K problem at the end of the previous millennium was an important indication of the kind of risk advanced computing technologies entail. The millennium bug episode highlights both the dependence of the modern societies upon computing, as well as the difficulty of forming a judgment of risks posed by that dependence.

3.3 Trust

Siri is a personal assistant application for iOS. The application uses natural language processing to answer questions, make recommendations, and perform actions by delegating requests to an expanding set of web services. The software also supposedly adapts to the user's individual preferences over time and personalizes results, as well as accomplishing tasks such as making dinner reservations and reserving a cab. It is a spin-out from the SRI International Artificial Intelligence Center, and is an offshoot of the DARPA-funded CALO project, described as the largest artificial-intelligence project ever launched. The project brings together experts in machine learning, natural language processing, and knowledge representation, human-computer interaction, flexible planning and behavioral studies. The CALO software learns by interacting with and being advised by its users and is meant to help users with military decision making task.


Once technical objects are stabilised, they become instruments of knowledge. As already noted the current AI technologies, like expert systems or intelligent assistants, portray an impression of intelligence which is somewhat misleading. The actual engineering of these systems is hidden from the user, and in most cases their knowledge and utility arises from connections to databases and systems which are external to the environment in which they are deployed. Thus, the knowledge these technical systems gather can be 'exported'.³³ To take a mild example - the Apple iOS assistant application of Siri, which is a by-product of AI research, draws its knowledge from the Cloud and Internet databases.³⁴

How does a user trust that the AI technology is working for his (her) welfare, and not some hidden agenda placed there, intentionally or inadvertently, by the programmer? How does the user trust that it will behave for his welfare, even in his absence, if the user doesn't have detailed knowledge of the system's underlying hardware and software limits, not to mention the subtle bugs. AI systems, being highly automated and self-sufficient (and having a deep, hidden architecture), raises trust issues which are qualitatively different

³³Madeleine Akrich. "The de-scription of technical objects". In: *Shaping technology/building society* (1992), pp. 205-224.

³⁴Krystal D'Costa. *Getting Serious With Siri*. 2011. URL: <http://blogs.scientificamerican.com/anthropology-in-practice/2011/11/08/getting-serious-with-siri>.

Siri, in a year.



from the ‘tool-like technologies’ of the past. As more such AI technologies make their way to the general public, the convenience of using these systems and the deceptive nature of their intelligence will tend to replace the concerns over privacy and security.

4 New Directions

It has been shown that extraneous sociological factors have been responsible for the gradual deviation of AI research from its original objectives and its acceleration in one particular direction - which has probably led to the inadequacies presented by AI technologies of today. Basic areas like judgement under uncertainty, intuition, commonsense reasoning, inventiveness, non-linear thinking have lagged behind, and should be encouraged. Going forward, this paper presents the case for two complementary methodologies that help in addressing the concerns arising out of the technologies and research as it is today.

4.1 A Sociological View of Intelligence

Almost all of the earlier models on which AI and cognitive science rested assumed that intelligence is not socially constituted and not socially situated.

In fact , the traditional AI techniques hinged on the following assumptions:³⁵

- Human mentality is a freestanding, individual, brain based phenomenon; and
- Human mentality is best understood in logical, linguistic, and rational terms

Most of the current AI technologies have resulted from the research grounded on such similar assumptions. But the success of these applications, especially of those in the form of intelligent assistants and expert systems, depends on the basis of how they perform in contexts and situations that are essential social and cultural in nature. Most of the inadequacies of these technologies highlighted above, result from their inability to handle practical social situations and uncertainties.

The idea that social and cultural factors are not only important but primary is not yet a widely appreciated or understood possibility. In the past, Joseph Weizenbaum has observed that intelligence manifests itself only relative to specific social and cultural contexts.³⁶ There is a need to develop a sound sociological basis for Artificial Intelligence. Social and Cultural assumptions should be incorporated in the worldviews guiding the work in AI. Intelligence, thus, should be understood as how well an entity performs in social situations, in addition to the inner architecture and cognitive model that brings about that behavior. A Sociology of Artificial Intelligence will result in more robust technologies in the future, which are better equipped to handle real world scenarios.

4.2 Human-Augmented AI

When Doug Engelbart was creating early computer interfaces and mapping systems, he firmly maintained his belief that the machine was meant to be an augments, not prosthesis. J.C.R. Licklider,³⁷ the computer science pioneer who had a profound effect on the development of technology and

³⁵.

³⁶William F Clocksin. “Artificial intelligence and the future”. In: *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 361.1809 (2003), pp. 1721–1748.

³⁷Joseph Carl Robnett Licklider. “Man-computer symbiosis”. In: *Human Factors in Electronics, IRE Transactions on* 1 (1960), pp. 4–11.

the Internet, also had the vision of enabling man and machine to cooperate in making decisions and controlling complex scenarios together, instead of compromising the flexibility by being dependent on predetermined programs. Future research should focus on similar principles.

To get a flavour of the potential power of this concept, consider the 2005 freestyle chess tournament in which man and machine could enter together as partners, rather than adversaries. Initially, even a supercomputer was beaten by a grandmaster with a relatively weak laptop. But the real surprise which took everyone off-guard came at the end. The eventual winner was not a grandmaster with a supercomputer, but actually two American amateurs using three relatively weak laptops. Their ability to effectively use their computers to deeply explore specific positions effectively counteracted the superior chess knowledge of the grandmasters and the superior computational power of other adversaries.³⁸ This astonishing result, of average men, average machines beating the best man, the best machine, strikingly illustrates the strength of human-machine cooperation - that the right symbiosis can be much more powerful than the sum of individual parts.

The power of Man and Machine working together has also manifested in a totally different but very relevant setting - Protein Folding. There are more ways of folding a protein than there are atoms in the universe. This is a world-changing problem with huge implications for our ability to understand and treat diseases. Supercomputers have traditionally struggled in this area. When computer scientists at created Foldit, a game where non-technical, non-biologist amateurs visually rearrange the structure of the protein (allowing the computer to manage the atomic forces and interactions and identify structural issues) , it was found that the players beat supercomputers 50 percent of the time and tied 30 percent of the time. Recently, the structure of the Mason-Pfizer monkey virus, a protease that has eluded determination for over 10 years, was opened up to a group of online gamers (through Foldit), who competed to model the protein, with all the associated scores, points, and rankings of a game. The players finished the model within 10 days - a notable and major scientific discovery^{39, 40}.

Amazon partially taps into the principle of Human-Augment AI through a concept which they call ‘Artificial Artificial Intelligence’. The premise

³⁸Shyam Shankhar. “The rise of human-computer cooperation”. In: *TED Video* (2012).

³⁹Ibid.

⁴⁰Dan Woods. *The Man-Machine Framework. How to Build Machine-Learning Applications the Right Way*.

is simple: Since Humans outperform AI in many simple tasks (like recognizing faces or sorting patterns), so why not farm out computing tasks to people, instead of machines? AAI has been employed as part of The Amazon Mechanical Turk (MTurk), a crowdsourcing Internet marketplace that enables computer programmers (known as Requesters) to co-ordinate the use of human intelligence to perform tasks that computers are unable to do yet.⁴¹

Current AI technologies lack common sense and symbolic reasoning, inventiveness, non-linear approaches, iterative hypotheses - aspects of intelligence which come naturally and easily to humans. On the other hand, machines have been better and ever improving in handling scale, volume and computation. Thus, an approach which harnesses the best of both worlds holds great potential. In completely manual and completely automated system, users and machines are effectively decoupled and too often, systems fail because they are not designed as a whole system with people and machines working in harmony. The power lies of Human-Augmented AI will lie in expressing the 'black box' and making it transparent. The human mind will enhance the machines solution by filling in the gaps.

A field of human intelligence - machine intelligence cooperation will thus help answer the concerns arising from the increasing separation of human beings from the inner workings and implementations of AI technology highlighted in the previous sections. This approach, traditionally underappreciated and unexplored. For example, a language of machine-plus-human interaction has not yet been developed. There is a need to reevaluate and reframe the conventional 'Man vs Machine' dichotomy towards a common 'Man and Machine' framework.^{42 43} Research oriented at technologies which involve both humans and machines performing intelligent tasks in the context of an integrated system should be encouraged. This would involve designs where both humans and machines have responsibilities, require access to resources, and have particular knowledge appropriate to tasks. Tasks maybe performed

⁴¹Kristen Millares Bolt. *Amazon creates Artificial Artificial Intelligence*. 2005. URL: <http://www.seattlepi.com/business/article/Amazon-creates-artificial-artificial-intelligence-1186698.php>.

⁴²Arnab Gupta. "Man + Machine". In: *SFSX Interactive talk* (2012).

⁴³The early signals indicating the potential of Human-Augmented AI technology can be read in the Big Data landscape. The idea of a 'Man and Machine framework' to generate commercial insights has recently been pioneered by Opera Solutions, while software company Palantir Technologies takes on important real world problems (e.g. counterterrorism) using a concept called 'Intelligence Augmentation'.

in parallel, or may require results or permission from the other. The field of Human Intelligence-Machine Intelligence cooperation would define roles to each as tools or assistants. Different types of knowledge would have to be distinguished: designers, end users, and maintenance people for instance. Humans would thus be designed in the process, and friction between man and machine, minimized.

5 Conclusions

The state of AI research and technology, as it stands today, is not without blemishes. Under the shining surface, it hides some fairly serious concerns and issues. Most of these risk and trust related issues arise from an imbalance resulting from the impact of extraneous sociological factors on AI research. The exploration of two paradigms - Sociological AI and Human-Augmented AI - will go a long way in addressing these concerns, and hold great potential as a research agenda for the future.

Besides this, it is essential that the funding agencies appreciate the importance of long-range basic research in AI. As Nils Nilsson⁴⁴ observes, AI, perhaps together with molecular genetics, will be society's predominant scientific endeavor for the rest of this century and well into the next - just as physics and chemistry predominated during the decades before and after 1900.

References

- Akrich, Madeleine. "The de-scription of technical objects". In: *Shaping technology/building society* (1992), pp. 205–224.
- Beck, Ulrich. *Risk society: Towards a new modernity*. Vol. 17. SAGE Publications Limited, 1992.
- Benner, Steven A and A Michael Sismour. "Synthetic biology". In: *Nature Reviews Genetics* 6.7 (2005), pp. 533–543.
- Bijker, Wiebe E. "Dikes and dams, thick with politics". In: *Isis* 98.1 (2007), pp. 109–123.

⁴⁴Nils J Nilsson. "Artificial intelligence prepares for 2001". In: *AI Magazine* 4.4 (1983), p. 7.

- “Understanding technological culture through a constructivist view of science, technology, and society”. In: *Visions of STS: Counterpoints in science, technology, and society studies* (2001), pp. 19–34.
- Bolt, Kristen Millares. *Amazon creates Artificial Artificial Intelligence*. 2005. URL: <http://www.seattlepi.com/business/article/Amazon-creates-artificial-artificial-intelligence-1186698.php>.
- Bostrom, Nick. *Ethical issues in advanced artificial intelligence*. 2003.
- “Existential risks”. In: *Journal of Evolution and Technology* 9 (2002).
- Bostrom, Nick, Milan M Cirkovic, and Martin J Rees. *Global catastrophic risks*. Oxford University Press Oxford, 2008.
- Bostrom, Nick and Eliezer Yudkowsky. “The ethics of artificial intelligence”. In: *Draft for Cambridge handbook of artificial intelligence* (2011).
- Boy, Guy A. “Integrated human-machine intelligence”. In: *Computers & chemical engineering* 17 (1993), S395–S404.
- Clocksini, William F. “Artificial intelligence and the future”. In: *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 361.1809 (2003), pp. 1721–1748.
- Computing, National Research Council (US). Committee on Innovations in, Communications, and Lessons from History. *Funding a revolution: government support for computing research*. National Academies Press, 1999.
- Cortada, James W. “How New Technologies Spread: Lessons from Computing Technologies”. In: *Technology and Culture* 54.2 (2013), pp. 229–261.
- D’Costa, Krystal. *Getting Serious With Siri*. 2011. URL: <http://blogs.scientificamerican.com/anthropology-in-practice/2011/11/08/getting-serious-with-siri>.
- Fairhead, Harry. *Kinect’s AI breakthrough explained*. 2011. URL: <http://www.i-programmer.info/news/105-artificial-intelligence/2176-kinects-ai-breakthrough-explained.html>.
- Forsythe, Thomas and Marvin Green. *ADVANCEMENTS IN 3-D SENSING TECHNOLOGY IMPLIMENTED BY THE KINECT*. Tech. rep. accessed 11/11, 2012.
- Gupta, Arnab. “Man + Machine”. In: *SFSX Interactive talk* (2012).
- Joy, Bill. “Why the future doesn’t need us.” In: *Nanoethics—the ethical and social implicatons of nanotechnology* (2000), pp. 17–39.
- Kaczynski, Theodore et al. *The unabomber manifesto: Industrial society and its future*. Filiquarian Publishing, LLC., 2005.

- Khoshelham, Kouros and Sander Oude Elberink. “Accuracy and resolution of kinect depth data for indoor mapping applications”. In: *Sensors* 12.2 (2012), pp. 1437–1454.
- Licklider, Joseph Carl Robnett. “Man-computer symbiosis”. In: *Human Factors in Electronics, IRE Transactions on* 1 (1960), pp. 4–11.
- Luger, George F. *Artificial intelligence: Structures and strategies for complex problem solving*. Pearson education, 2005.
- MacKenzie, Donald. *Mechanizing proof: computing, risk, and trust*. The MIT Press, 2004.
- Markoff, John. “Behind Artificial Intelligence, a Squadron of Bright Real People”. In: *New York Times* 14 (2005).
- McCarthy, John et al. “A proposal for the dartmouth summer research project on artificial intelligence, august 31, 1955”. In: *AI Magazine* 27.4 (2006), p. 12.
- Muehlhauser, Luke and Anna Salamon. “Intelligence explosion: Evidence and import”. In: *Singularity Hypotheses*. Springer, 2012, pp. 15–42.
- Nilsson, Nils J. “Artificial intelligence prepares for 2001”. In: *AI Magazine* 4.4 (1983), p. 7.
- *The quest for artificial intelligence*. Cambridge University Press Cambridge, 2010.
- Oikonomidis, Iason, Nikolaos Kyriazis, and Antonis A Argyros. “Efficient model-based 3D tracking of hand articulations using Kinect.” In: *BMVC*. 2011, pp. 1–11.
- Orlikowski, Wanda J and Debra C Gash. “Technological frames: making sense of information technology in organizations”. In: *ACM Transactions on Information Systems (TOIS)* 12.2 (1994), pp. 174–207.
- Pinch, Trevor J and Wiebe E Bijker. *The Social Construction of Facts and Artifacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other*. 1987.
- Poole, David, Alan Mackworth, and Randy Goebel. *Computational Intelligence*. Oxford University Press Oxford, 1998.
- Roush, Wade. “Marvin minsky on common sense and computers that emote”. In: *Technology Review*, July 13 (2006).
- Russell, Stuart. *Artificial intelligence: A modern approach, 2/E*. Pearson Education India, 2003.
- Schwartz, Ronald David. “Artificial intelligence as a sociological phenomenon”. In: *Canadian Journal of Sociology/Cahiers canadiens de sociologie* (1989), pp. 179–202.

- Shankhar, Shyam. “The rise of human-computer cooperation”. In: *TED Video* (2012).
- Siri, in a Year*. URL: <http://thedoghouseidiaries.com/3175>.
- Smisek, Jan, Michal Jancosek, and Tomas Pajdla. “3D with Kinect”. In: *Consumer Depth Cameras for Computer Vision*. Springer, 2013, pp. 3–25.
- Whitby, Blay. *Reflections on Artificial Intelligence*. Intellect (UK), 1996.
- Woods, Dan. *The Man-Machine Framework. How to Build Machine-Learning Applications the Right Way*.
- Yudkowsky, Eliezer. “Artificial intelligence as a positive and negative factor in global risk”. In: *Global catastrophic risks* 1 (2008), p. 303.
- Yudkowsky, Eliezer et al. “Creating Friendly AI-The Analysis and Design of Benevolent Goal Architectures”. In: (2001).