IS AI IN JEOPARDY? THE NEED TO UNDER PROMISE AND OVER DELIVER – THE CASE FOR REALLY USEFUL MACHINE LEARNING

Martin Ciupa

CTO, calvIO Inc (part of the Calvary Robotics Group),
855 Publishers Pkwy, Webster, NY 14580, USA

ABSTRACT

There has been a dramatic increase in media interest in Artificial Intelligence (AI), in particular with regards to the promises and potential pitfalls of ongoing research, development and deployments. Recent news of success and failures are discussed. The existential opportunities and threats of extreme goals of AI (expressed in terms of Superintelligence/AGI and Socio-Economic impacts) are examined with regards to this media "frenzy", and some comment and analysis provided. The application of the paper is in two parts, namely to first provide a review of this media coverage, and secondly to recommend project naming in AI with precise and realistic short term goals of achieving really useful machines, with specific smart components. An example of this is provided, namely the RUMLSM project, a novel AI/Machine Learning system proposed to resolve some of the known issues in bottom-up Deep Learning by Neural Networks, recognised by DARPA as the "Third Wave of AI." An extensive, and up to date at the time of writing, Internet accessible reference set of supporting media articles is provided.

KEYWORDS

AI, Machine Learning, Robotics, Superintelligence, Third Wave of AI.

1. INTRODUCTION

In the past 18 months, though 2016 and into the first quarter of 2017 we are experiencing an amazing period of growth and media attention in Artificial Intelligence (AI) and Machine Learning (ML) [1]. Many high tech CEO's have claimed association of their companies future regarding these technologies [2], massive investment is piling in [3]. While there have been many successes, there have been a few failures. The risk exists that "irrational exuberance" of the optimists may be a bubble that pessimists may be able to pop with claims that the AI vision is hyped [4]. The application of this research paper is to bring the current positive/negative media coverage to authors attention on this subject and to suggest a more precise, non-conflated, and modest use of terms in project terminology.
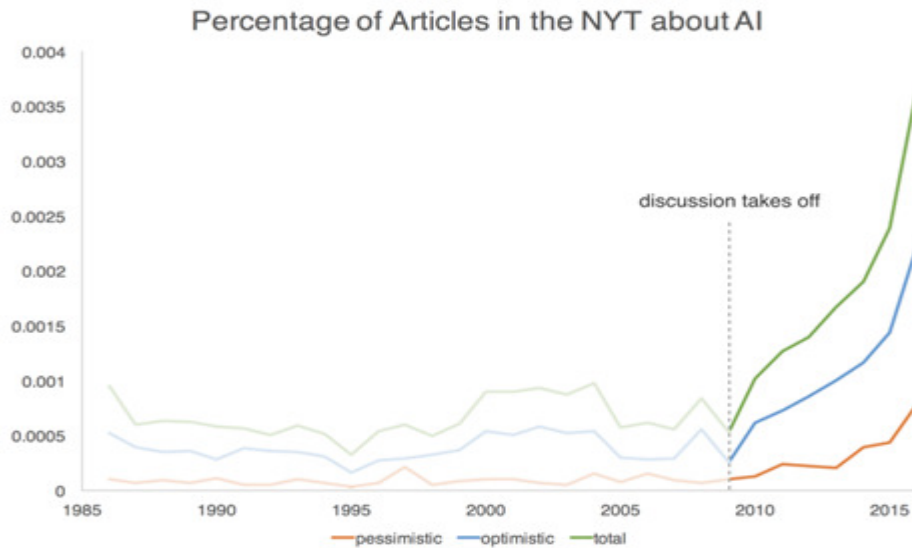
Figure 1. AI Media Coverage (Source [1])

For some commentators the very name "Artificial Intelligence" is controversial, implying something "unnatural." The very word "Intelligence" defies definition in the minds of many [5]. For our purpose, we do not wish to address the philosophical artificiality question here. Rather, we will simply think of intelligence in working definition terms as: "that faculty of mind by which order is perceived in a situation hitherto considered to be disordered." a definition (of Fatmi and Young) quoted in the Oxford Companion to the Mind [6]. Thus, if a machine, or an animal, has the facility to perceive the order change, then in a sense it has "intelligence", and some aspect (perhaps limited, and with artifice) of "mind." In this sense, it is a Cybernetic definition from a comparative systems perspective. We trust these points make the case that the AI term is itself controversial, it has its ideological proponents and detractors from the outset. In many respects, it would be useful to leave the term to the side and simply build really useful machines with specific descriptions of their application.

2. PROMISE – VISIONS OF AN EMERGING NEW POSITIVE PARADIGM

2.1. Recent success stories in AI and Machine Learning

There have been several notable success stories announced in the past 18 months, across a range of Machine Learning, Image Recognition, and Natural Language disciplines of AI. A high profile has been given to the ongoing development of Deep Learning by Reinforced/Convolutional Neural Network-based learning, in part underpinning IBM Watson/Jeopardy, AlphaGo and Poker game successes. These systems are beating the best human players who have invested their working lives into acquiring deep knowledge and intelligently applying it. See [7], [8] and [9]. The result of these successes has been a significant uptick in the media reporting of the AI potential and opportunities as mentioned earlier in [1], specifically more than four times as many New York Times articles discussed AI in 2016 than in 2009, as a percentage of the total number of articles published.

2.2. Investment into AI and Machine Learning

As a result of this positive news mentioned above, and the lead of many of the top tech company CEO's mentioned in [2], a significant financial round of investments has, and continues to be made, in this area.   The market research company Forrester report that across all businesses, there will be a greater than 300% increase in investment in AI in 2017 compared with 2016 [3].

2.3. Timeline to AGI

The reporting and investment mentioned above have led to widespread speculation as to, "where is this is all leading to?" And specifically, if the advent of "Superintelligent", Artificial Generic Intelligence (AGI) is on the foreseeable horizon. In 2015, the Machine Intelligence Research Institute compiled the MIRI AI predictions dataset, a collection of public predictions about human-level AGI timelines. Interesting features of the dataset include the result that the median dates at which people's predictions suggest AIG is less likely than not and more likely than not are 2033 and 2037 respectively. That is perhaps within 20 years. See Figure 2.

2.4. AI as the Singularity Hypothesis solution to Existential Risks

Several high profile scientists and philosophers have made their opinion on the issue mentioned above clear, optimistic and pessimistic, more about the pessimistic contributions will be mentioned in section 3 below. However, in the optimistic context of the promise of AI, and AGI in particular, much is suggested as opportunity to be sought by their agency. E.g., in their facility to help us to solve existential crises currently confronting humanity (such as Climate Change, Stabilisation of World Markets, Abundance of material needs, such as food, water and shelter, plus universal access to Healthcare and Education). The extreme optimistic positioning is that AGI, and the "Technology Singularity" will be humanity's last necessary invention, and once arrived a new age will ensure with material abundance for all, and solutions to humanity's problems [11].

2.5. Asilomar Conference 2017, World Economic Forum 2017

The Future of Life Institute, recognising the above opportunity, and in an attempt to offset negative pessimism about the downside of AGI, set out principles at the 2017 Asilomar conference, fundamentally a manifesto and guideline set of rules for the ethical development of AGI [12]. The explicit message is that action needs to be taken now to intercept any negativity and deliver positive result.

Similarly, the World Economic Forum, a high-profile international organisation whose stated mission is to "improve the state of the world", via Public-Private cooperation, continued on 17-20 January 2017 in its "exploration" on how developments in AI and Robotics could impact industry, governments and society in the future. Seeking to design innovative governance models to ensure that their benefits are maximised and the associated risks kept under control. The emphasis is that AI is coming, it is inevitable, and action has to be taken to ensure it arrives in good order [13].
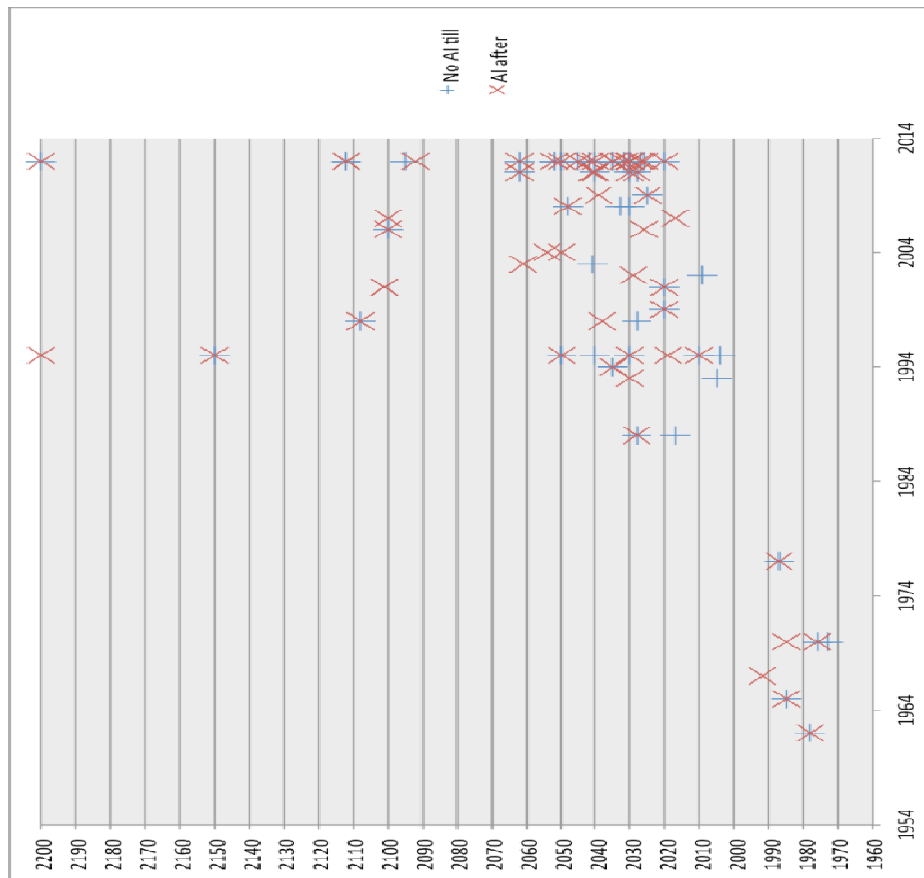
Figure 2: Predictions from the MIRI dataset surveys. Source [10]

2.6. RoboEthics

The academic study of Robot Ethics or "RoboEthics" is a well-established position in educational and research institutes. In part, this is a response to the understanding that robots, in particular drones, are entering the military as lethal weapons of war. And as such need regulation and legal codes of conduct [14]. But, is also concerned about the rights of robots themselves, as anticipated sapient and sentient "beings" with inalienable rights [15].

3. PITFALLS – RISK OF AI BEING OVERHYPED AND OR DANGEROUS

3.1. Recent failure stories in AI

But, not all voices are positive ones over AI. First of all, not all AI developments in the past few years have had positive results. Many no doubt go unreported, but two high profile ones in recent time are Microsoft's Tay AI Bot (a Machine Learning online chatbot that learns from its interactions). Tay developed a "character" with strong racist, sexist opinions, and had to be removed [16] Furthermore, ongoing problems with AI messaging chatbots, at Google, resulted in a rollback in their adoption [17]. While these, perhaps overly ambitious projects did not meet the

expectations of the management, they nevertheless provide a proving ground for what can be done with today's technology and what needs further work. Such feedback is helpful in setting more realistic expectations [18].

There are other noted failures in AI projects, e.g., the fatal crash of a Tesla Automated Driving system in 2016, whereby autopilot sensors on the Model S failed to distinguish a white tractor-trailer crossing the highway against a bright sky. In this case, it might be claimed that the testing/QA failed to identify this scenario and build sufficient responses to factor it safely.

3.2. Known problems with Deep Learning

Much, though not all of the recent success in AI has been due to Machine Learning advancement in so-called Deep Learning based Neural Networks. But, these bottom-up learning based systems (so-called because they develop learned experience from hierarchical analysis of basic data types and their deep correlations) have issues. They can learn well, often in tight domains beyond human learned experience and can outperform a human. But, a human can provide a top-down description, a discursive rationalisation to a, "why did you decide to make that move/act" question with a, "because, I think gaining control of the centre area a useful strategy at this point in the game" question. Whereas. a Deep Learning system cannot. It is not rule-based and cannot easily track its "reasoning". In a sense it is like an experienced financial markets trader who knows when to buy/sell, not because of the analytics, but because he has the intuitive "feel" built over years of trading. However, intuition is not necessarily a good guide when the market modalities change, as many traders have found out to their ultimate cost. As some critics have said, "it is only real intelligence if you can show it's working" [19].

In applications that require compliance with regulations and legal constraints, such systems may be considered risky. Furthermore, if such a system were committed to a highly Mission Critical Application, it would be difficult to question, supervise and control. Such intuitive systems, once relied upon, cannot easily be "rolled-back", human operators become reliant on the system to make all the hard decisions. In the event of failure, total system breakdown may ensue. There is a significant risk here. As AI is woven more deeply into the fabric of everyday life, the tension between human operators and AI has become increasingly salient. There is also a paradox: the same technologies that extend the intellectual powers of humans can displace them as well [20].

3.3. Mysterian Position – Beyond Computation

Some AI critics while recognising useful tools can be built do not believe AGI is likely in the short or medium term outlook of understanding of such things as "consciousness, free-will and self-aware visceral experience (qualia)." While this paper will not address these "Hard Problems of Consciousness" concerns; it is perhaps not necessary to incorporate these functional elements. If they cannot be defined in the case of natural human intelligence, we are not in a position to easily address them in our smart tools and AGI. Maybe AGI can pass Turing Tests, but not have these functions. The better question is perhaps, do we have them or are they illusory? See [21].

3.4. Issues of functional complexity and computational limits

Some AI critics, while not accepting necessarily the problems of intelligence being in principle beyond computational methods, believe it is so complex that our current computational

processing, is just not powerful enough, or the cost to provide it would be so astronomical as not to be worth the benefit. The above is a more of an objection in pragmatic terms. "Smart AI Tools" can be delivered in tight knowledge domains, but as for AGI, this is simply a non-starter for a long time to come. Recently Philosopher Daniel Dennett has made this argument, referring to AGI as "balderdash" [22]. Likewise, Philosopher Jerry Kaplan [23] has expressed concerns. However, some will point out that massive hierarchies of systems, even such as we have today, on parallel networked computational systems might cover the spectrum of human domain knowledge, with abstract reasoning and coordinated response. Such systems are termed Multi-Agent or Distributed AI systems [24]. Furthermore, such computer platforms, while not affordable today in their massive extension, might be within a few decades.

As for software, that also continues to improve, and the potential for geometric improvement by evolutionary learning algorithms, whereby AI works on AI and evolves rapidly is viable. The evolutionary cost of learning in simulated "toy universes" may not be as "expensive" or slow as in our experience of evolution in "natural reality." Such simulations can be used for development, but also for QA/testing of extreme scenarios not easily tested physically [25].

3.5. Moore's Law's End?

While the hope is computer processing will continue to become dramatically cheaper, and more powerful, this is not a given. Moore's Law that underpins that expectation is considered by some to be no longer applicable [26]. However, Intel counters that it is still alive [27]. The possibility of future improvement by innovation in new technology, such as Quantum Computers (and other innovative technologies) might step up and provide the substrates to keep Moore's law alive for years to come [28].

3.6. "Prophets of Doom", AGI as an Existential Risk

Even if the negative technical issues expressed above can be resolved, such that the promise and positive thinking expressed optimistically in section 2 above can be delivered, there are high profile voices being expressed that developing safe and ethical AGI within the next 20-30 years or so is very dangerous. These critics express concerns that AGI perhaps should not be developed. The danger is regarding the existential risk that they present to humanity, as well as to human culture regarding reducing human employment to such a point as to destabilise society.

Physicist Stephen Hawking, Microsoft founder Bill Gates and Tesla/SpaceX founder Elon Musk have expressed concerns about the possibility that AI could develop to the point that humans could not control it [29]. Stephen Hawking said in 2014 that "Success in creating AI would be the biggest event in human history. Unfortunately, it might also be the last, unless we learn how to avoid the risks." Hawking believes that in the coming decades, AI could offer "incalculable benefits and risks" such as "technology outsmarting financial markets, out-inventing human researchers, out-manipulating human leaders, and developing weapons we cannot even understand." He makes the point that "The development of full artificial intelligence could spell the end of the human race" [30].

In January 2015, Nick Bostrom joined Stephen Hawking, Max Tegmark, Elon Musk, Lord Martin Rees, Jaan Tallinn, and numerous AI researchers, in signing the Future of Life Institute's open letter speaking to the potential risks and benefits associated with artificial intelligence. The

signatories "…believe that research on how to make AI systems robust and beneficial is both important and timely, and that there are concrete research directions that can be pursued today [31]. Their position is not that AGI should be halted, rather that urgent attention is needed to ensure that it is delivered ethically and safely. It should be noted that in the case of a sudden AGI "intelligence explosion", effective precautions will be extremely difficult. Not only would its creators have little ability to test their precautions on an intermediate intelligence, but the creators might not even have made any precautions at all, if the advent of the intelligence explosion catches them completely by surprise. Hence their concern to think through the AGI negative issues as early as possible so as to prepare principles to handle them.

3.7. AI Impact on Jobs – Luddite Rebellion

Furthermore, some consider AGI and the uptake of pervasive Robotics/Industrial Automation as harmful to our civilisation regarding reducing human employment to such a point as to destabilise society. Such instabilities in earlier Industrial Revolutions resulted in an anti-technology reaction, known the Luddite Rebellion [32]. Could this occur again? This concern needs to be addressed. Indeed it is the subject and concern of the World Economic Forum as mentioned earlier [13].

4. NEED TO UNDER PROMISE AND OVER DELIVER – AVOID THE HYPE

4.1. Addressing concerns

As Management Consultant Tom Peters says, "Quality is important, to be sure, so is absolute response time, and price, but at the top of most lists, by far, is keeping your word." With uncertainty rising, if you 'under promise, over deliver,' you will not only keep the customers satisfied; you'll keep the customers" [33]. While you can make the case this advice is not universally applicable, in cases where timelines are in doubt, and where there are critics in the wings who are willing to take pot shots, and your case depends, in part, on technology innovation not yet delivered, then it seems prudent to apply this principle.

With that in mind, I propose that we be more modest in our claims for AI and AGI. Yes, we may expect it to come, and indeed we may need it to help us solve knotty problems we face in an ever more chaotic and complex world full of existential problems, many arising out of our mismanagement. But, we should be warry of over-extending the use of the term, or conflating them. Thus, let us make a point of not talking up AGI for the next five years or so, to deliver on our projects in hand. That, if done well and to expectations will establish AI as a real vehicle for a paradigm change (and build solid foundations for AGI). As we have seen AI is an easily challenged term and easily hyped, we need to be more specific with our language, else "AI" risks becoming meaningless [34]. We should address specifically what it does and what it doesn't. Often it is better to say Machine Learning, Natural Language Understanding, Forecasting, Image Recognition, Predictive Analytics/Big Data, etc. I would put these technologies under the banner of Cybernetics, with its transdisciplinary systems and top-down perspective [35].

Having said that, I think it is wise to maintain the dialogue such as ongoing at the Asilomar Conference and World Economic Forum [12], [13], to prepare the way for ethical and safe AGI. Once a self-reinforcing cybernetic loop starts to generate more and more intelligent systems, the onset of AGI is likely to be fast, and catch many unprepared. Given that it is wise to

acknowledge that it is necessary to be prepared to think the unthinkable, before the unstoppable starts, that is if you have a desire to stop it/make it safe.

4.2. Really Useful Machine Learning – RUMLSM

With the rising optimistic and pessimistic media attention to AI, the philosophical debate over the term, and the conflation of AI projects with lofty AGI ambitions, it is perhaps wise to consider using terms that limit reference to AI and AGI at this time.   More emphasis should be made on terms such as Machine Learning, Neural Networks, Evolutionary Computation, Vision, Robotics, Expert Systems, Speech Processing, Planning, Natural Language Processing, etc., and to make sure that the scope is well defined in practical application.

In this spirit, we are currently undertaking modest AI research and development of an innovative Hybrid Machine Learning paradigm that incorporates bottom-up Deep Learning Neural Networks and a means to extract a rationalisation of a top-down heuristic narrative. Our intention is to roll this out over 4 phases over the following three years and apply it to calvIO's Industrial Robotics platforms and integrated systems. We refer, and headline, this as "Really Useful Machine Learning" (RUML^{SM}).

The key novelty of the system is the integration of intuitive bottom-up and rational top-down learning; we find inspiration for the paradigm in known means to teach expertise at and expert practitioner level by "masterclass mentoring", outlined in the Figure below [36].
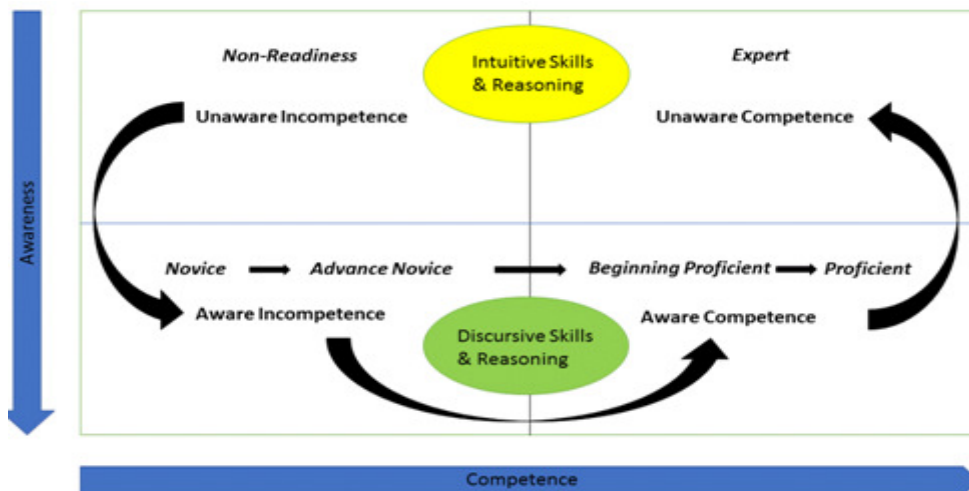

Figure 3: Masterclass Mentoring: Journey from Non-Readiness to Expert

4.3. Compliance, Supervision and Control needs

Key aspects of the deliverable features of RUMLSM will be the ability to manage compliance, supervision and control of the system, by inspection of its extracted rationalised heuristic rule base. The testing of designed systems is performed in extensive simulation scenarios, examining extreme conditions before deploying to the physical systems.

4.4. Extraction of Heuristics from Deep Learning Neural Networks

The means by which Expert Heuristics are extracted from the Deep Learning Neural Networks has been studied by other teams [37], [38] and [39]. The means by which we propose to do so in RUMLSM is an innovative patent pending process. Expert Heuristic/Rule extraction can be defined as "…given a trained neural network and the data on which it was trained, produce a description of the network's hypothesis that is comprehensible yet closely approximates the network's predictive behaviour." Such extraction algorithms are useful for experts to verify and cross-check neural network systems.

Earlier this year, John Launchbury, director of DARPA's Information Innovation Office said, "There's been a lot of hype and bluster about AI." They published their view of AI into "Three Waves", so as to explain what AI can do, what AI can't do, and where AI is headed. RUMLSM is very much in this third wave in our opinion [40].
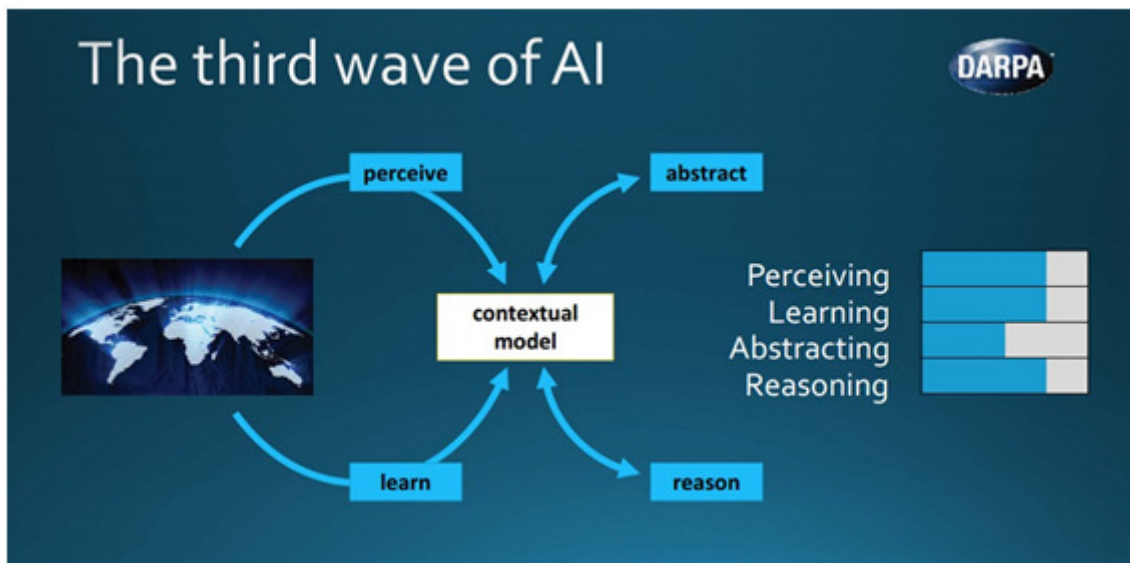

Figure 4: DARPA's Third Wave of AI [Source: 40]

5. CONCLUSIONS

The increase in media interest in Artificial Intelligence (AI) was noted, and the optimistic and pessimistic expectations as to the impact of AI (and potentially future AGI) commented on. While there have been some failures recently, significant milestone successes have encouraged massive investment. But, there is nevertheless a potential for AI to be marked as a bubble and "hyped" technology. Indeed there are already signs of a backlash against AI [41], [42] and [43].

A short-term recommendation, to avoid the extremes of the positive and negative positions, was offered, we should ensure we don't conflate AGI with AI and be more specific in the use of terms to avoid AI becoming "meaningless". As Jerry Kaplan says "Had artificial intelligence been named something less spooky, we'd probably worry about it less." [23]. As an example of a more realistic use of terms for AI, goals is the use of terms such as "Really Useful Machine Learning" (RUMLSM). This particular technique was introduced in the context of resolving

some of the known issues in bottom-up Deep Learning by Neural Networks with a top-down Cybernetic simulation-based process providing a more open and rational means to supervise, regulate and maintain compliance of the system.   This patent pending technology uses a Masterclass-Mentoring paradigm, and fits into DARPA's "Third Wave of AI."

ACKNOWLEDGEMENTS

REFERENCES

Note: Given the media content referenced the format has been adjusted to feature the links to these.

[1]   Ethan Fast, Feb 3 2017, "Long-term Trends in the Public Perception of Artificial Intelligence", https://hackernoon.com/long-term-trends-in-the-public-perception-of-artificial-intelligence-6b1512fdd7ac#.v16cpe36l Last Referenced 5th March 2017.
[2]   Fortune Tech, Jun 03, 2016,"Tech CEOs Declare This the Era of Artificial Intelligence" http://fortune.com/2016/06/03/tech-ceos-artificial-intelligence/ Last Referenced 5th March 2017.
[3]   Gil Press, 11 November 2016, "Forrester Predicts Investment In Artificial Intelligence Will Grow 300% in 2017" https://www.forbes.com/sites/gilpress/2016/11/01/forrester-predicts-investment-in-artificial-intelligence-will-grow-300-in-2017/#44642b785509 Last Referenced 5th March 2017.
[4]   Om Malik, 26 August 2016, "The Hype—and Hope—of Artificial Intelligence",
        http://www.newyorker.com/business/currency/the-hype-and-hope-of-artificial-intelligence
        Last Referenced 5th March 2017.
[5]   H. B. Barlow, 21 July 1983, "Intelligence, guesswork, language", Nature 304, 207 - 209 (21 July 1983); doi:10.1038/304207a0, (also see…
        http://www.nature.com/nature/journal/v304/n5923/abs/304207a0.html Last Referenced 5th March 2017.)
[6]   Fatmi, H.A. & Young, R.W. Nature 228, 97 (1970) (also see... http://www.ifsr.org/index.php/in-memorandum-dr-haneef-akhtar-fatmi-03-july-1933-04-april-1995/)
[7]   Richard Mallah, 29 December 2015, "The Top A.I. Breakthroughs of 2015",
        https://futureoflife.org/2015/12/29/the-top-a-i-breakthroughs-of-2015/ Last Referenced 5th March 2017.
[8]   Mona Lalwani, 25 December 2016 "AI was everywhere in 2016"
        https://www.engadget.com/2016/12/25/ai-was-everywhere-in-2016/ Last Referenced 5th March 2017.
[9]   Tonya Riley, 3 March 2017, "Artificial intelligence goes deep to beat humans at poker" http://www.sciencemag.org/news/2017/03/artificial-intelligence-goes-deep-beat-humans-poker   Last Referenced 5th March 2017.
[10] 5th December 2015 "MIRI AI Predictions Dataset" http://aiimpacts.org/miri-ai-predictions-dataset/ Last Referenced 5th March 2017.
[11] Peter Diamandis, Steven Kotler, 21 February 2012, ISBN: 9781451614213, Publishers: Free Press, Tantor Media "Abundance: The Future Is Better Than You Think", See…
        http://www.diamandis.com/abundance/ Last Referenced 5th March 2017.
[12] Future of Life Institute, "BENEFICIAL AI 2017" Asilomar Conference 2017
        https://futureoflife.org/bai-2017/ Last Referenced 5th March 2017.
[13] The Future of Artificial Intelligence and Robotics, World Economic Forum 2017,
        https://www.weforum.org/communities/the-future-of-artificial-intelligence-and-robotics
        Last Referenced 5th March 2017.
[14] A Jung Moon, "RobotEthics info Database" http://www.amoon.ca/Roboethics/ Last Referenced 5th March 2017.

[15] Alex Hern, 17 January 2017, "Give robots 'personhood' status, EU committee argues", https://www.theguardian.com/technology/2017/jan/12/give-robots-personhood-status-eu-committee-argues Last Referenced 5th March 2017.

[16] John West, 2 April 2016, "Microsoft's disastrous Tay experiment shows the hidden dangers of AI", https://qz.com/653084/microsofts-disastrous-tay-experiment-shows-the-hidden-dangers-of-ai/   Last Referenced 5th March 2017.

[17] Andrew Orlowski 22 Feb 2017, "Facebook scales back AI flagship after chatbots hit 70% f-AI-lure rate", https://www.theregister.co.uk/2017/02/22/facebook_ai_fail/ Last Referenced 5th March 2017.

[18] Justin Bariso, 23 February 2017, "Microsoft's CEO Sent an Extraordinary Email to Employees After They Committed an Epic Fail",   http://www.inc.com/justin-bariso/microsofts-ceo-sent-an-extraordinary-email-to-employees-after-they-committed-an-.html Last Referenced 5th March 2017.

[19] James Vincent, 10 October 2016, "These are three of the biggest problems facing today's AI", http://www.theverge.com/2016/10/10/13224930/ai-deep-learning-limitations-drawbacks
Last Referenced 5th March 2017.

[20] IEEE Technical Community Spotlight, 2 September 2016, "On the Use of AI – the Dependency Dilemma", http://sites.ieee.org/spotlight/ai-ethical-dilemma/ Last Referenced 5th March 2017.

[21] Josh Weisberg, Internet Encyclopaedia of Philosophy, "The Hard Problem of Consciousness", http://www.iep.utm.edu/hard-con/ Last Referenced 5th March 2017.

[22] Daniel Dennett, 28 February 2017, "Why robots won't rule the world – Viewsnight", BBC Newsnight, https://www.youtube.com/watch?v=2ZxzNAEFtOE, Last Referenced 5th March 2017.

[23] Jerry Kaplan, 3 March 2017, AI's PR Problem, https://www.technologyreview.com/s/603761/ais-pr-problem/ Last Referenced 5th March 2017.

[24] Tamer Sameeh, 6 March 2017, "Decentralized Artificial Super-intelligence Via Multi-agent Systems and Ethereum's Smart Contracts" https://steemit.com/crypto-news/@tamersameeh/decentralized-artificial-super-intelligence-via-multi-agent-systems-and-ethereum-s-smart-contracts Last Referenced 6th March 2017.

[25] Luzius Meisser, 11 January 2017, "Simulations as Test Environments"
http://meissereconomics.com/2017/01/11/Simulation.html Last Referenced 6th March 2017.

[26] Tom Simonite, 13 May 2016, MIT Technology Review, "Moore's Law Is Dead. Now What?" https://www.technologyreview.com/s/601441/moores-law-is-dead-now-what/ Last Referenced 6th March 2017.

[27] Justine Brown, 6 January 2017, CIO Dive, "Moore's Law is 'alive and well and flourishing,' Intel CEO says" http://www.ciodive.com/news/moores-law-is-alive-and-well-and-flourishing-intel-ceo-says/433484/ Last Referenced 6th March 2017.

[28] Josh Hamilton, February 13, 2017 "Moore's Law and Quantum Computing"
https://cloudtweaks.com/2017/02/moores-law-quantum-computing/ Last Referenced 6th March 2017.

[29] Michael Sainato, 19 August 2015 Observer, "Stephen Hawking, Elon Musk, and Bill Gates Warn About Artificial Intelligence" http://observer.com/2015/08/stephen-hawking-elon-musk-and-bill-gates-warn-about-artificial-intelligence/   Last Referenced 6th March 2017.

[30] Rory Cellan-Jones, 2 December 2014, BBC, "Stephen Hawking warns artificial intelligence could end mankind" http://www.bbc.co.uk/news/technology-30290540 Last Referenced 6th March 2017.

[31] George Dvorsky, 14 January 2015, "Prominent Scientists Sign Letter of Warning About AI Risks" http://io9.gizmodo.com/prominent-scientists-sign-letter-of-warning-about-ai-ri-1679487924   Last Referenced 6th March 2017.

[32] Gil Press, 26 February 2017, Forbes, "Luddites Against Job-Killing Automation And Technology Enthusiasts Creating New Industries" https://www.forbes.com/sites/gilpress/2017/02/26/luddites-against-job-killing-automation-and-technology-enthusiasts-creating-new-industries/#425e0ea77e46 Last Referenced 6th March 2017.

[33] Tom Peters, 1987 "Under Promise, Over Deliver" http://tompeters.com/columns/under-promise-over-deliver/ Last Referenced 6th March 2017.

[34] Ian Bogost 4 March 2017, The Atlantic, "'Artificial Intelligence' Has Become Meaningless" https://www.theatlantic.com/technology/archive/2017/03/what-is-artificial-intelligence/518547/   Last Referenced 6th March 2017.

[35] Kevin Warwick, 10 November 2016, MIT Technology Review, "The Future of Artificial Intelligence and Cybernetics" https://www.technologyreview.com/s/602830/the-future-of-artificial-intelligence-and-cybernetics/ Last Referenced 6th March 2017.

[36] D'Youville College, Professional Development, https://www.dyc.edu/academics/professional-development/ Last Referenced 6th March 2017.

[37] Tameru Hailesilassie, (IJCSIS) International Journal of Computer Science and Information Security, Vol. 14, No. 7, July 2016 "Rule Extraction Algorithm for Deep Neural Networks: A Review" https://arxiv.org/ftp/arxiv/papers/1610/1610.05267.pdf Last Referenced 6th March 2017.

[38] Jan Ruben Zilke, Master Thesis, TUD, "Extracting Rules from Deep Neural Networks" http://www.ke.tu-darmstadt.de/lehre/arbeiten/master/2015/Zilke_Jan.pdf Last Referenced 6th March 2017.

[39] Zhiting Hu, Xuezhe Ma, Zhengzhong Liu, Eduard Hovy, Eric P. Xing, School of Computer Science Carnegie Mellon University, 2016 "Harnessing Deep Neural Networks with Logic Rules" http://www.cs.cmu.edu/~epxing/papers/2016/Hu_etal_ACL16.pdf Last Referenced 6th March 2017.

[40] Steve Crowe, 21 February 21 2017, Robotics Trends, "What AI Can and Can't Do: DARPA's Realistic View, http://www.roboticstrends.com/article/what_ai_can_and_cant_do_darpas_realistic_view/Artificial_Intelligence

[41] Matt Asay., 3 March 2017, Infoworld Tech Watch, "Artificially inflated: It's time to call BS on AI" http://www.infoworld.com/article/3176602/artificial-intelligence/artificially-inflated-its-time-to-call-bs-on-ai.html Last Referenced 6th March 2017.

[42] Andrew Orlowski, 2 January 2017, The Register, "'Artificial Intelligence' was 2016's fake news" https://www.theregister.co.uk/2017/01/02/ai_was_the_fake_news_of_2016/   Last Referenced 6th March 2017.

[43] Luke Dormehl, 18 February 2017, Wired, "Don't believe the hype when it comes to AI" http://www.wired.co.uk/article/sensationalism-ai-hype-innovation Last Referenced 6th March 2017.