

# On the Differences between Human and Machine Intelligence

Roman V. Yampolskiy

Computer Science and Engineering, University of Louisville  
roman.yampolskiy@louisville.edu

## Abstract

Terms *Artificial General Intelligence* (AGI) and *Human-Level Artificial Intelligence* (HLAI) have been used interchangeably to refer to the Holy Grail of Artificial Intelligence (AI) research, creation of a machine capable of achieving goals in a wide range of environments. However, widespread implicit assumption of equivalence between capabilities of AGI and HLAI appears to be unjustified, as humans are not general intelligences. In this paper, we will prove this distinction.

## 1 Introduction<sup>1</sup>

Imagine that tomorrow a prominent technology company announces that they have successfully created an Artificial Intelligence (AI) and offers for you to test it out. You decide to start by testing developed AI for some very basic abilities such as multiplying 317 by 913, and memorizing your phone number. To your surprise, the system fails on both tasks. When you question the system’s creators, you are told that their AI is human-level artificial intelligence (HLAI) and as most people cannot perform those tasks neither can their AI. In fact, you are told, many people can’t even compute  $13 \times 17$ , or remember name of a person they just met, or recognize their coworker outside of the office, or name what they had for breakfast last Tuesday<sup>2</sup>. The list of such limitations is quite significant and is the subject of study in the field of Artificial Stupidity [Trazzi and Yampolskiy, 2018; Trazzi and Yampolskiy, 2020].

Terms *Artificial General Intelligence* (AGI) [Goertzel *et al.*, 2015] and *Human-Level Artificial Intelligence* (HLAI) [Baum *et al.*, 2011] have been used interchangeably (see [Barrat, 2013], or “(AGI) is the hypothetical intelligence of a machine that has the capacity to understand or learn any intellectual task that a human being can.” [Anonymous, Retrieved July 3, 2020]) to refer to the Holy Grail of Artificial Intelligence (AI) research, creation of a machine capable of: achieving goals in a wide range of environments

[Legg and Hutter, 2007a]. However, widespread implicit assumption of equivalence between capabilities of AGI and HLAI appears to be unjustified, as humans are not general intelligences. In this paper, we will prove this distinction.

Others use slightly different nomenclature with respect to general intelligence, but arrive at similar conclusions. “**Local generalization, or “robustness”**: ... “adaptation to known unknowns within a single task or well-defined set of tasks”. ... **Broad generalization, or “flexibility”**: “adaptation to unknown unknowns across a broad category of related tasks”. ... **Extreme generalization**: human-centric extreme generalization, which is the specific case where the scope considered is the space of tasks and domains that fit within the human experience. We ... refer to “human-centric extreme generalization” as “generality”. Importantly, as we deliberately define generality here by using human cognition as a reference frame ..., it is only “general” in a limited sense. ... To this list, we could, theoretically, add one more entry: “universality”, which would extend “generality” beyond the scope of task domains relevant to humans, to any task that could be practically tackled within our universe (note that this is different from “any task at all” as understood in the assumptions of the No Free Lunch theorem [Wolpert and Macready, 1997; Wolpert, 2012]).” [Chollet, 2019].

## 2 Prior work

We call some problems ‘easy’, because they come naturally to us like understanding speech or walking and we call other problems ‘hard’ like playing Go or violin, because those are not human universals and require a lot of talent and effort [Yampolskiy, 2012]. We ignore ‘impossible’ for humans to master domains, since we mostly don’t even know about them or see them as important. As LeCun puts it: “[W]e can’t imagine tasks that are outside of our comprehension, right, so we think, we think we are general, because we’re general of all the things that we can apprehend, but there is a huge world out there of things that we have no idea” [LeCun, August 31, 2019]. Others, agree: “we might not even be

<sup>1</sup> Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

<sup>2</sup>Some people could do that and more, for example 100,000 digits of  $\pi$  have been memorized using special mnemonics.

aware of the type of cognitive abilities we score poorly on.” [Barnett, December 23, 2019].

This is most obvious in how we test for intelligence. For example, Turing Test [Turing, 1950], by definition, doesn’t test for universal general intelligence, only for human-level intelligence in human domains of expertise. Like a drunkard searching for his keys under the light because there it is easier to find them, we fall for the Streetlight effect observation bias only searching for intelligence in domains we can easily comprehend [Yampolskiy, 2019]. “The *g* factor, by definition, represents the single cognitive ability common to success across all intelligence tests, emerging from applying factor analysis to test results across a diversity of tests and individuals. But intelligence tests, by construction, only encompass tasks that humans can perform – tasks that are immediately recognizable and understandable by humans (anthropocentric bias), since including tasks that humans couldn’t perform would be pointless. Further, psychometrics establishes measurement validity by demonstrating predictiveness with regard to activities that humans value (e.g. scholastic success): the very idea of a “valid” measure of intelligence only makes sense within the frame of reference of human values.” [Chollet, 2019].

Moravec further elaborates the difference between future machines and humans: “Computers are universal machines, their potential extends uniformly over a boundless expanse of tasks. Human potentials, on the other hand, are strong in areas long important for survival, but weak in things far removed. Imagine a “landscape of human competence,” having lowlands with labels like “arithmetic” and “rote memorization,” foothills like “theorem proving” and “chess playing,” and high mountain peaks labeled “locomotion,” “hand-eye coordination” and “social interaction.” Advancing computer performance is like water slowly flooding the landscape. A half century ago it began to drown the lowlands, driving out human calculators and record clerks, but leaving most of us dry. Now the flood has reached the foothills, and our outposts there are contemplating retreat. We feel safe on our peaks, but, at the present rate, those too will be submerged within another half century.” [Moravec, 1998].

Chollet writes: “How general is human intelligence? The No Free Lunch theorem [Wolpert and Macready, 1997; Wolpert, 2012] teaches us that any two optimization algorithms (including human intelligence) are equivalent when their performance is averaged across every possible problem, i.e. algorithms should be tailored to their target problem in order to achieve better-than-random performance. However, what is meant in this context by “every possible problem” refers to a uniform distribution over problem space; the distribution of tasks that would be practically relevant to our universe (which, due to its choice of laws of physics, is a specialized environment) would not fit this definition. Thus we may ask: is the human *g* factor universal? Would it generalize to every possible task in the universe? ... [T]his question is highly relevant when it comes to AI: if there is

such a thing as universal intelligence, and if human intelligence is an implementation of it, then this algorithm of universal intelligence should be the end goal of our field, and reverse-engineering the human brain could be the shortest path to reach it. It would make our field close-ended: a riddle to be solved. If, on the other hand, human intelligence is a broad but ad-hoc cognitive ability that generalizes to human-relevant tasks but not much else, this implies that AI is an open-ended, fundamentally anthropocentric pursuit, tied to a specific scope of applicability.” [Chollet, 2019].

Humans have general capability only in those human accessible domains and likewise artificial neural networks inspired by human brain architecture do unreasonably well in the same domains. Recent work by Tegmark et al. shows that deep neural networks would not perform as well in randomly generated domains as they do in those domains humans consider important, as they map well to physical properties of our universe. “We have shown that the success of deep and cheap (low-parameter-count) learning depends not only on mathematics but also on physics, which favors certain classes of exceptionally simple probability distributions that deep learning is uniquely suited to model. We argued that the success of shallow neural networks hinges on symmetry, locality, and polynomial log-probability in data from or inspired by the natural world, which favors sparse low-order polynomial Hamiltonians that can be efficiently approximated.” [Lin *et al.*, 2017].

### 3 Humans are not AGI

An agent is general (universal [Hutter, 2004]) if it can learn anything another agent can learn. We can think of a true AGI agent as a superset of all possible NAIs (including capacity to solve AI-Complete problems [Yampolskiy, 2013]). Some agents have *limited domain generality*, meaning they are general, but not in all possible domains. The number of domains in which they are general may still be Dedekind-infinite, but it is a strict subset of domains in which AGI is capable of learning. For an AGI it’s domain of performance is any efficiently learnable capability, while humans have a smaller subset of competence. Non-human animals in turn may have an even smaller repertoire of capabilities, but are nonetheless general in that subset. This means that humans can do things animals cannot and AGI will be able to do something no human can. If an AGI is restricted only to domains and capacity of human expertise, it is the same as HLAI.

Humans are also not all in the same set, as some are capable of greater generality (*G* factor [Jensen, 1998]) and can succeed in domains, in which others cannot. For example, only a tiny subset of all people is able to conduct cutting-edge research in quantum physics, implying differences in our general capabilities between theory and practice. While theoretical definition of general intelligence is easy to understand, its practical implementation remains uncertain.

“LeCun argues that even self-supervised learning and learnings from neurobiology won’t be enough to achieve artificial general intelligence (AGI), or the hypothetical intelligence of a machine with the capacity to understand or learn from any task. That’s because intelligence — even human intelligence — is very specialized, he says. “AGI does not exist — there is no such thing as general intelligence,” said LeCun. “We can talk about rat-level intelligence, cat-level intelligence, dog-level intelligence, or human-level intelligence, but not artificial general intelligence.”” [Wiggers, May 2, 2020].

An agent is not an AGI equivalent if it could not learn something another agent could learn. Hence, we can divide all possible tasks into human learnable and those, which no human can learn, establishing that humans are not AGI equivalent. We already described ‘easy’ and ‘hard’ for humans problems, the third category of ‘impossible’ is what we would classify as abilities impossible for humans to learn efficiently [Valiant, 2013]. Computer-unaided humans [Blum and Vempala, 2020] do not possess capabilities in this category, to any degree, and are unlikely to be able to learn them. If performed by a human, they would be considered magical, but as Arthur Clarke has famously stated: “Any sufficiently advanced technology is indistinguishable from magic.”

Some current examples include: estimating face from speech [Oh *et al.*, 2019], DNA [Sero *et al.*, 2019] or ear [Yaman *et al.*, 2020], extracting passwords from typing sounds [Zhuang *et al.*, 2009; Shumailov *et al.*, 2019], using lightbulbs [Nassi *et al.*, 2020] and hard drives [Kwong *et al.*, 2019] as microphones, communicating via heat emissions [Guri *et al.*, 2015b], or memory-write-generated electromagnetic signals [Guri *et al.*, 2015a], and predicting gender, age and smoking status from images of retinal fundus [Poplin *et al.*, 2018]. This is what is already possible with Narrow AI (NAI) today, AGI will be able to see patterns where humans see nothing but noise, invent technologies we never considered possible and discover laws of physics far above our understanding. Capabilities, we humans will never possess, because we are not general intelligences. Even humans armed with simple calculators are no match for such problems.

LeCun gives an example of one task no human could learn: “So let me take a very specific example, it’s not an example it’s more like a quasi-mathematical demonstration, so you have about 1 million fibers coming out of one of your eyes, okay two million total, but let’s talk about just one of them. It’s 1 million nerve fibers in your optical nerve, let’s imagine that they are binary so they can be active or inactive, so the input to your visual cortex is 1 million bits. Now, they connected to your brain in a particular way and your brain has connections that are kind of a little bit like a convolution net they are kind of local, you know, in the space and things like this. Now imagine I play a trick on you, it’s a pretty nasty trick I admit, I cut your optical nerve and I put a device that makes a random permutation of all the nerve

fibers. So now what comes to your, to your brain, is a fixed but random permutation of all the pixels, there’s no way in hell that your visual cortex, even if I do this to you in infancy, will actually learn vision to the same level of quality that you can.” [LeCun, August 31, 2019].

Chollet elaborates on the subject of human unlearnable tasks: “[H]uman intellect is not adapted for the large majority of conceivable tasks. This includes obvious categories of problems such as those requiring long-term planning beyond a few years, or requiring large working memory (e.g. multiplying 10-digit numbers). This also includes problems for which our innate cognitive priors are unadapted; ... For instance, in the [Traveling Salesperson Problem] TSP, human performance degrades severely when inverting the goal from “finding the shortest path” to “finding the longest path” [MacGregor and Ormerod, 1996] – humans perform even worse in this case than one of the simplest possible heuristic: farthest neighbor construction. A particularly marked human bias is dimensional bias: humans ... are effectively unable to handle 4D and higher. ... Thus, ... “general intelligence” is not a binary property which a system either possesses or lacks. It is a spectrum,” [Chollet, 2019]. “Human physical capabilities can thus be said to be “general”, but only in a limited sense; when taking a broader view, humans reveal themselves to be extremely specialized, which is to be expected given the process through which they evolved.” [Chollet, 2019]. “[W]e are born with priors about ourselves, about the world, and about how to learn, which determine what categories of skills we can acquire and what categories of problems we can solve.” [Chollet, 2019].

If such tasks are in fact impossible for any human to perform, that proves that humans are not AGI equivalent. But, how do we know what a highly intelligent agent is capable of or more interestingly incapable of learning? How do we know what human’s can’t learn [Ziesche and Yampolskiy, 2020]? One trick we can use, is to estimate the processing speed [Roberts and Stankov, 1999] for an average human on a particular learning task and to show that even 120 years, a very optimistic longevity estimate for people, is not sufficient to complete learning that particular task, while much faster computer can do so in seconds.

Generality can be domain limited or unlimited. Different animals, such as dolphins, elephants, mice, etc. and humans are all general in overlapping but not identical sets of domains. Humans are not a superset of all animal intelligences. There are some things animals can do that humans cannot and vice versa. For example, humans can’t learn to speak animal “languages” and animals can’t learn to play chess [Yampolskiy, 2018b]. Richard Hamming made this point in his famous paper - “The Unreasonable Effectiveness of Mathematics”: “Just as there are odors that dogs can smell and we cannot, as well as sounds that dogs can hear and we cannot, so too there are wavelengths of light we cannot see and flavors we cannot taste. Why then, given our brains wired the way they are, does the remark, “Perhaps there are thoughts we cannot think,” surprise you? Evolution, so far,

may possibly have blocked us from being able to think in some directions; there could be unthinkable thoughts." [Wigner, 1990].

Only AGI is universal/general intelligence over all learnable domains. AGI is not just capable of anything a human can do; it is capable of learning anything that could be learned. It is a Superset of all NAIs and is equal in capability to Superintelligence.

## 4 Conclusions

There is no shortage of definitions of intelligence [Legg and Hutter, 2007a; Legg and Hutter, 2007b; Hernández-Orallo, 2017; Wang, 2019; Yampolskiy, 2020a], but we felt it was important to clarify that humans are neither fully general nor terminal point in the space of the possible minds [Yampolskiy, 2015]. As Chollet says: "We may even build systems with higher generalization power (as there is no a priori reason to assume human cognitive efficiency is an upper bound), or systems with a broader scope of application. Such systems would feature intelligence beyond that of humans." [Chollet, 2019]. Humans only have a subset of capabilities an AGI will have and the capability difference between us and AGI is far greater than capability difference between AGI and superintelligence (SAI). Bostrom describes three forms of superintelligence (p. 53-57) [Bostrom, 2014]: Speed SAI (like a faster human), Collective SAI (like a group of humans), and Quality SAI (does what humans can't). All three can be accomplished by an AGI, so there is no difference between AGI and SAI, they are the same ( $HLAI \leq AGI = SAI$ ) and the common takeoff-speed debate [Yudkowsky and Hanson, 2008] resolves to hard takeoff, from definitions. This implies even stronger limitations [Yampolskiy, 2017; Yampolskiy, 2019; Yampolskiy, 2020b] on our capability to control AI and a more immediate faceoff. We are already having many problems with Ignorance Explosion [Lukasiewicz, 1974; Lukasiewicz, 1994], an Intelligence Explosion [Loosemore and Goertzel, 2012; Muehlhauser and Salamon, 2012] will be well beyond our capabilities to control.

If we use Legg's definition of intelligence [Legg and Hutter, 2007a], and average performance across all possible problems, we can arrive at a somewhat controversial result that modern AI is already smarter than any human is. An individual human can only learn a small subset of domains and human capabilities can't be trivially transferred between different humans to create a union function of all human capabilities, but that is, at least theoretically, possible for AI. Likewise, humans can't emulate some computer algorithms, but computers can run any algorithm a human is using. Machines of 2020 can translate between hundreds of languages, win most games, generate art, write poetry and learn many tasks individual humans are not capable of learning. If we were to integrate all such abilities into a single AI agent it would on average outperform any person across all possible problem domains, but perhaps not humanity as a whole seen

as a single agent. This may have been true for a number of years now, and is becoming more definitive every year. As an AI agent can be a superset of many algorithms from which it can choose it would not be a subject to the No Free Lunch (NFL) theorems [Wolpert and Macready, 1997; Wolpert, 2012].

While AI dominates humans in most domains of human interest [Goodfellow *et al.*, 2014; Mnih *et al.*, 2015; Silver *et al.*, 2017; Devlin *et al.*, 2018; Clark *et al.*, 2019; Vinyals *et al.*, 2019], there are domains in which humans would not even be able to meaningfully participate. This is similar to the Unpredictability [Yampolskiy, 2020b] and Unexplainability/Incomprehensibility of AI [Yampolskiy, 2019] results, but at a meta-level. The implications for AI control and AI Safety and Security [Callaghan *et al.*, 2017; Yampolskiy, 2018a; Babcock *et al.*, 2019; Babcock *et al.*, July 16-19, 2016] are not encouraging. To be dangerous AI doesn't have to be general, it is sufficient for it to be superior to humans in a few strategic domains. If AI can learn a particular domain it will quickly go from Hypohuman to Hyperhuman performance [Hall, 2009]. Additionally, common proposal for merging of humanity with machines doesn't seem to work as adding HLAI to AGI adds nothing to AGI, meaning in a cyborg agent human will become a useless bottleneck as AI becomes more advanced and the human will be eventually removed, if not explicitly at least implicitly from control. What does this paper tell us? Like the dark matter of the physical universe, the space of all problems is mostly unknown unknowns, and most people don't know that and don't even know that they don't know it. To paraphrase the famous saying: "The more AI learns, the more I realize how much I don't know."

## References

- [Anonymous, Retrieved July 3, 2020] Anonymous. Artificial general intelligence. Wikipedia. Available at: [https://en.wikipedia.org/wiki/Artificial\\_general\\_intelligence](https://en.wikipedia.org/wiki/Artificial_general_intelligence).
- [Babcock *et al.*, July 16-19, 2016] James Babcock, Janos Kramar and Roman Yampolskiy. The AGI Containment Problem. The Ninth Conference on Artificial General Intelligence (AGI2015). NYC, USA.
- [Babcock *et al.*, 2019] James Babcock, János Kramár and Roman V Yampolskiy. Guidelines for artificial intelligence containment. Next-Generation Ethics: Engineering a Better Society. A. E. Abbas: 90-112.
- [Barnett, December 23, 2019] Matthew Barnett. Might humans not be the most intelligent animals? Available at: <https://www.lesswrong.com/posts/XjuT9vgBfwXPxsdfN/might-humans-not-be-the-most-intelligent-animals>.
- [Barrat, 2013] James Barrat. Our final invention: Artificial intelligence and the end of the human era, Macmillan.
- [Baum *et al.*, 2011] Seth D Baum, Ben Goertzel and Ted G Goertzel. "How long until human-level AI? Results from

- an expert assessment." *Technological Forecasting and Social Change* 78(1): 185-195.
- [Blum and Vempala, 2020] Manuel Blum and Santosh Vempala. "The complexity of human computation via a concrete model with an application to passwords." *Proceedings of the National Academy of Sciences* 117(17): 9208-9215.
- [Bostrom, 2014] Nick Bostrom. *Superintelligence: Paths, dangers, strategies*, Oxford University Press.
- [Callaghan *et al.*, 2017] Vic Callaghan, James Miller, Roman Yampolskiy and Stuart Armstrong. *The Technological Singularity: Managing the Journey*, Springer.
- [Chollet, 2019] François Chollet. "On the measure of intelligence." arXiv preprint arXiv:1911.01547.
- [Clark *et al.*, 2019] Peter Clark, Oren Etzioni, Tushar Khot, Bhavana Dalvi Mishra, Kyle Richardson, Ashish Sabharwal, Carissa Schoenick, Oyvind Tafjord, Niket Tandon and Sumithra Bhakthavatsalam. "From 'F' to 'A' on the NY Regents Science Exams: An Overview of the Aristo Project." arXiv preprint arXiv:1909.01958.
- [Devlin *et al.*, 2018] Jacob Devlin, Ming-Wei Chang, Kenton Lee and Kristina Toutanova. "Bert: Pre-training of deep bidirectional transformers for language understanding." arXiv preprint arXiv:1810.04805.
- [Goertzel *et al.*, 2015] Ben Goertzel, Laurent Orseau and Javier Snaider. "Artificial General Intelligence." *Scholarpedia* 10(11): 31847.
- [Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville and Yoshua Bengio. *Generative adversarial nets*. *Advances in neural information processing systems*.
- [Guri *et al.*, 2015a] Mordechai Guri, Assaf Kachlon, Ofer Hasson, Gabi Kedma, Yisroel Mirsky and Yuval Elovici. *GSMem: Data Exfiltration from Air-Gapped Computers over GSM Frequencies*. 24th USENIX Security Symposium (USENIX Security 15).
- [Guri *et al.*, 2015b] Mordechai Guri, Matan Monitz, Yisroel Mirski and Yuval Elovici. *Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations*. 2015 IEEE 28th Computer Security Foundations Symposium, IEEE.
- [Hall, 2009] J Storrs Hall. *Beyond AI: Creating the conscience of the machine*, Prometheus books.
- [Hernández-Orallo, 2017] José Hernández-Orallo. *The measure of all minds: evaluating natural and artificial intelligence*, Cambridge University Press.
- [Hutter, 2004] Marcus Hutter. *Universal artificial intelligence: Sequential decisions based on algorithmic probability*, Springer Science & Business Media.
- [Jensen, 1998] Arthur Robert Jensen. *The g factor: The science of mental ability*, Praeger Westport, CT.
- [Andrew *et al.*, 2019] Andrew Kwong, Wenyuan Xu and Kevin Fu. *Hard drive of hearing: Disks that eavesdrop with a synthesized microphone*. 2019 IEEE Symposium on Security and Privacy (SP), IEEE.
- [LeCun, August 31, 2019] Yann LeCun. *Yann LeCun: Deep Learning, Convolutional Neural Networks, and Self-Supervised Learning*. AI Podcast, Lex Fridman. Available at: <https://www.youtube.com/watch?v=SGSOCuByo24>.
- [Legg and Hutter, 2007a] Shane Legg and Marcus Hutter. "A collection of definitions of intelligence." *Frontiers in Artificial Intelligence and applications* 157: 17.
- [Legg and Hutter, 2007b] Shane Legg and Marcus Hutter. "Universal intelligence: A definition of machine intelligence." *Minds and Machines* 17(4): 391-444.
- [Lin *et al.*, 2017] Henry W Lin, Max Tegmark and David Rolnick. "Why does deep and cheap learning work so well?" *Journal of Statistical Physics* 168(6): 1223-1247.
- [Loosemore and Goertzel, 2012] Richard Loosemore and Ben Goertzel. *Why an intelligence explosion is probable. Singularity hypotheses*, Springer: 83-98.
- [Lukasiewicz, 1974] Julius Lukasiewicz. "The ignorance explosion." *Leonardo* 7(2): 159-163.
- [Lukasiewicz, 1994] Julius Lukasiewicz. *The ignorance explosion: Understanding industrial civilization*, McGill-Queen's Press-MQUP.
- [MacGregor and Ormerod, 1996] James N MacGregor and Tom Ormerod. "Human performance on the traveling salesman problem." *Perception & psychophysics* 58(4): 527-539.
- [Mnih *et al.*, 2015] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland and Georg Ostrovski. "Human-level control through deep reinforcement learning." *Nature* 518(7540): 529-533.
- [Moravec, 1998] Hans Moravec. "When will computer hardware match the human brain." *Journal of evolution and technology* 1(1): 10.
- [Muehlhauser and Salamon, 2012] Luke Muehlhauser and Anna Salamon. *Intelligence explosion: Evidence and import. Singularity hypotheses*, Springer: 15-42.
- [Nassi *et al.*, 2020] Ben Nassi, Yaron Pirutin, Adi Shamir, Yuval Elovici and Boris Zadov. *Lamphone: Real-Time Passive Sound Recovery from Light Bulb Vibrations*. *Cryptology ePrint Archive*. Available at: <https://eprint.iacr.org/2020/708>.
- [Oh *et al.*, 2019] Tae-Hyun Oh, Tali Dekel, Changil Kim, Inbar Mosseri, William T Freeman, Michael Rubinstein and Wojciech Matusik. *Speech2face: Learning the face behind a voice*. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- [Poplin *et al.*, 2018] Ryan Poplin, Avinash V. Varadarajan, Katy Blumer, Yun Liu, Michael V. McConnell, Greg S. Corrado, Lily Peng and Dale R. Webster. "Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning." *Nature Biomedical Engineering* 2(3): 158-164.
- [Roberts and Stankov, 1999] Richard D Roberts and Lazar Stankov. "Individual differences in speed of mental processing and human cognitive abilities: Toward a taxonomic model." *Learning and Individual Differences* 11(1): 1-120.
- [Sero *et al.*, 2019] Dzemila Sero, Arslan Zaidi, Jiarui Li, Julie D White, Tomás B González Zarzar, Mary L Marazita, Seth M Weinberg, Paul Suetens, Dirk Vandermeulen and Jennifer K Wagner. "Facial recognition from DNA using face-to-DNA classifiers." *Nature communications* 10(1): 2557.

- [Shumailov *et al.*, 2019] Ilia Shumailov, Laurent Simon, Jeff Yan and Ross Anderson. "Hearing your touch: A new acoustic side channel on smartphones." arXiv preprint arXiv:1903.11137.
- [Silver *et al.*, 2017] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai and Adrian Bolton. "Mastering the game of go without human knowledge." *Nature* 550(7676): 354.
- [Trazzi and Yampolskiy, 2018] Michaël Trazzi and Roman V Yampolskiy. "Building safer AGI by introducing artificial stupidity." arXiv preprint arXiv:1808.03644.
- [Trazzi and Yampolskiy, 2020] Michaël Trazzi and Roman V Yampolskiy. "Artificial Stupidity: Data We Need to Make Machines Our Equals." *Patterns* 1(2): 100021.
- [Turing, 1950] A. Turing. "Computing Machinery and Intelligence." *Mind* 59(236): 433-460.
- [Valiant, 2013] Leslie Valiant. *Probably Approximately Correct: Nature's Algorithms for Learning and Prospering in a Complex World*, Basic Books (AZ).
- [Vinyals *et al.*, 2019] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds and Petko Georgiev. "Grandmaster level in StarCraft II using multi-agent reinforcement learning." *Nature*: 1-5.
- [Wang, 2019] Pei Wang. "On Defining Artificial Intelligence." *Journal of Artificial General Intelligence* 10(2): 1-37.
- [Wiggers, May 2, 2020] Kyle Wiggers. Yann LeCun and Yoshua Bengio: Self-supervised learning is the key to human-level intelligence. Available at: <https://venturebeat.com/2020/05/02/yann-lecun-and-yoshua-bengio-self-supervised-learning-is-the-key-to-human-level-intelligence/>.
- [Wigner, 1990] Eugene P Wigner. The unreasonable effectiveness of mathematics in the natural sciences. *Mathematics and Science*, World Scientific: 291-306.
- [Wolpert, 2012] David H Wolpert. What the no free lunch theorems really mean; how to improve search algorithms. Santa Fe Institute.
- [Wolpert and Macready, 1997] David H Wolpert and William G Macready. "No free lunch theorems for optimization." *IEEE transactions on evolutionary computation* 1(1): 67-82.
- [Yaman *et al.*, 2020] Dogucan Yaman, Fevziye Irem Eyiokur and Hazım Kemal Ekenel. "Ear2Face: Deep Biometric Modality Mapping." arXiv preprint arXiv:2006.01943.
- [Yampolskiy, 2012] Roman V Yampolskiy. "AI-Complete, AI-Hard, or AI-Easy—Classification of Problems in AI." The 23rd Midwest Artificial Intelligence and Cognitive Science Conference, Cincinnati, OH, USA.
- [Yampolskiy, 2015] Roman V Yampolskiy. *The space of possible mind designs*. International Conference on Artificial General Intelligence, Springer.
- [Yampolskiy, 2017] Roman V Yampolskiy. "What are the ultimate limits to computational techniques: verifier theory and unverifiability." *Physica Scripta* 92(9): 093001.
- [Yampolskiy, 2018a] Roman V Yampolskiy. *Artificial Intelligence Safety and Security*, Chapman and Hall/CRC.
- [Yampolskiy, 2018b] Roman V Yampolskiy. "The singularity may be near." *Information* 9(8): 190.
- [Yampolskiy, 2019] Roman V Yampolskiy. "Unexplainability and incomprehensibility of artificial intelligence." arXiv preprint arXiv:1907.03869.
- [Yampolskiy, 2020a] Roman V Yampolskiy. "On Defining Differences Between Intelligence and Artificial Intelligence." *Journal of Artificial General Intelligence* 11(2): 68-70.
- [Yampolskiy, 2020b] Roman V Yampolskiy. "Unpredictability of AI: On the Impossibility of Accurately Predicting All Actions of a Smarter Agent." *Journal of Artificial Intelligence and Consciousness* 7(01): 109-118.
- [Yampolskiy, 2013] Roman V. Yampolskiy. Turing Test as a Defining Feature of AI-Completeness. *Artificial Intelligence, Evolutionary Computation and Metaheuristics - In the footsteps of Alan Turing*. Xin-She Yang (Ed.), Springer: 3-17.
- [Yudkowsky and Hanson, 2008] Eliezer Yudkowsky and Robin Hanson. The Hanson-Yudkowsky AI-foom debate. MIRI Technical Report. Available at: <http://intelligence.org/files/AIFoomDebate.pdf>.
- [Zhuang *et al.*, 2009] Li Zhuang, Feng Zhou and J Doug Tygar. "Keyboard acoustic emanations revisited." *ACM Transactions on Information and System Security (TISSEC)* 13(1): 1-26.
- [Ziesche and Yampolskiy, 2020] Soenke Ziesche and Roman V Yampolskiy. "Towards the Mathematics of Intelligence." *The Age of Artificial Intelligence: An Exploration*: 1.