



How much intelligence is there in artificial intelligence? A 2020 update

Han L.J. van der Maas^{*}, Lukas Snoek, Claire E. Stevenson

University of Amsterdam, The Netherlands

ARTICLE INFO

Keywords:

Artificial Intelligence
Deep learning
Individual differences
Intelligence tests
Reinforcement

ABSTRACT

Schank (1980) wrote an editorial for *Intelligence* on “How much intelligence is there in artificial intelligence?”. In this paper, we revisit this question. We start with a short overview of modern AI and showcase some of the AI breakthroughs in the four decades since Schank’s paper. We follow with a description of the main techniques these AI breakthroughs were based upon, such as deep learning and reinforcement learning; two techniques that have deep roots in psychology. Next, we discuss how psychologically plausible AI is and could become given the modern breakthroughs in AI’s ability to learn. We then access the main question of how intelligent AI systems actually are. For example, are there AI systems that can solve human intelligence tests? We conclude that Shank’s observation, that intelligence is all about generalization and that AI is not particularly good at this, has, so far, withstood the test of time. Finally, we consider what AI insights could mean for the study of individual differences in intelligence. We close with how AI can further *Intelligence* research and vice versa, and look forward to fruitful interactions in the future.

In 1980 Roger C. Schank wrote an editorial for *Intelligence* entitled “How much intelligence is there in artificial intelligence (AI)?”. His first observation was the lack of any interaction between the fields of intelligence research and artificial intelligence research. Since then limited interactions have taken place. His second observation was that AI is relevant for intelligence research. This was based on the state of the art of research in both fields, at that time, when AI was still in its infancy. Given the breathtaking developments in modern AI, it is worthwhile to ask Schank’s question again. We contend that the question’s relevance has increased over time. Schank’s third observation was that real intelligence is all about generalization, which at that time was a weak point in AI, but perhaps not anymore.

Schank was modest in claiming to speak for AI, given the multitude and diversity of AI approaches at that time. The current explosion of developments in AI research makes a systematic overview even more impossible. But we think it is useful to present a short overview here and discuss the relations between AI and intelligence research as it is reported in *Intelligence*.

Perhaps it is helpful to start discussing possible reasons why AI may be irrelevant to intelligence research. The first well known argument is that AI may be building intelligent machines, but that machine intelligence is not similar to human intelligence or cognition. Dennett’s (2006) cognitive wheel idea illustrates this. The wheel is a beautiful invention, but it is not used in nature as a method of movement and therefore of no

interest to biologists. This may apply to many AI inventions, e.g., cognitive wheels. The second reason is that AI, even if it has psychological relevance, can only inform us about cognition and not intelligence, as the latter primarily concerns individual differences in cognitive functioning, and not cognitive functioning itself. This relates to Cronbach’s famous distinction between the psychology of processes or mechanisms versus the psychology of individual differences, where both apparently exist independently of each other (Cronbach, 1957). We will argue against both arguments, but first discuss current developments in AI.

1. A short history of modern AI

One constant in the short but rich history of AI is the tension between two main goals. The first goal is to understand (human) intelligence and the second is to let computers do information processing tasks. Chess is an exemplary example. The earliest attempts to build computer chess programs were inspired by studies on human chess thinking (de Groot, 1946). However, the grand steps were only made when approaches were developed that made use of the unique strengths of computers. Deep Blue, the program that beat Gary Kasparov, the world champion in 1996, was a typical cognitive wheel. It applied advanced tree search algorithms that allowed deep search of positions based on a rather simple evaluation of resulting positions. Deep Blue mainly told us how

^{*} Corresponding author at: Department of Psychology, University of Amsterdam, Nieuwe Achtergracht 129-B, 1018WS Amsterdam, The Netherlands.
E-mail address: h.l.j.vandermaas@uva.nl (H.L.J. van der Maas).

humans do not play chess, as it depended on a brute force computing approach, evaluating 200 million positions per second, that humans are particularly bad at. Humans are much better at evaluating positions (Gobet & Chassy, 2009). How chess players think is still largely a mystery (Gobet & Charness, 2018), but chess involves both high-level conceptual processing and low-level recognition of familiar patterns (Lane & Chang, 2018; Van Harrevel, Wagenmakers, & Van Der Maas, 2007; see Blanch (2020) for an overview of research on individual differences in chess).

The renewed interest in AI among psychologists is based on novel techniques that have also revolutionized computer chess in the last 5 years. AlphaZero chess (Silver et al., 2018) combines a deep learning network that values positions and predicts next moves, with a reinforcement learning system (explained in more detail below). The amazing thing is that AlphaZero initially performs badly, but the underlying systems help each other to learn and improve during millions of self-played games. This approach radically departs from earlier, non-learning, programs where brute force search and built-in indexes of openings and endgames were the key to success. The open-source variant of AlphaZero chess, Leela Chess Zero, currently dominates chess rating lists and is praised for her human-like play (including occasional incomprehensible blunders). So, it could be said that the initial divergence between the two goals of AI (cognition vs application) has diminished given the latest spectacular developments in computer chess.

Chess is not the only impressive example of recent successes in AI. In the last 10 years we have seen breakthroughs in areas such as natural language (NLP) processing, computer vision, robotics, and more. In the field of NLP, machine translation, AI assistants (like Alexa, Google and Siri), and automated closed captions in videos have become commonplace in our lives. These have evolved from AI systems that obtained historical milestones such as IBM's Watson question answering system defeating human champions at the popular trivia TV show Jeopardy in 2011 and Microsoft's speech recognition software successfully transcribing conversational telephone speech on the widely benchmarked Switchboard task in 2017. More recently, instead of transcribing or interpreting human text or speech, NLP models have been developed to generate original text and speech. A particularly impressive example is the recent GPT-3 model, which was able to write a complete and coherent article in the Guardian (GPT-3, 2020) arguing that robots will not destroy humanity, a fear held by several well-known people in science and industry.

The advances in computer vision are similarly impressive. Perhaps the most well-known examples are in the domain of object and face recognition, where Google Lens (available on most Android smartphones) is an excellent example of object recognition and Facebook's automatic face tagging system uses the power of face recognition. AI powered vision systems have also been applied in other, more practical domains, such as medicine, in which they are used to assist medical doctors in disease diagnosis and pharmacologists with treatment development. For example, AI models are used to detect cancer from medical scans and to assist the development of novel medical drugs (Kaul, Enslin, & Gross, 2020).

Another field that has seen remarkable developments is robotics. While robots before the turn of the century were mostly limited to simple repetitive tasks, modern AI powered robots are able to perform surgery, have been used in nursing homes to assist with medical documentation and combat loneliness among their residents, and commercially available fully self-driving cars are expected to arrive at the market within 1 or 2 years (Siciliano & Khatib, 2016). The developments are spectacular and worrisome at the same time. The study of the societal consequences of AI is becoming a research area in itself (Dubhashi & Lappin, 2017; Jobin, Ienca, & Vayena, 2019; Rudin, 2019).

2. Modern AI techniques

At the heart of the progress in AI lies the development of novel techniques and algorithms commonly known as "artificial neural networks". In the late 20th century, ANNs started to become the model of choice in many AI domains and are currently used in most modern AI applications. Because of their importance in modern AI, we discuss their origin and subsequent development, which shows a long history of mutual interaction between and inspiration from psychology, neuroscience, and (artificial) intelligence research. Then, we discuss (artificial) reinforcement learning and other techniques that are common in modern AI.

2.1. Artificial neural networks

Nowadays, many applications using artificial intelligence are powered, at least partly, by (deep) artificial neural networks (ANNs). ANNs come in many different varieties, but they all are composed of a collection of largely identical processing "units" structured in a hierarchical fashion, not unlike the way biological neurons make up biological neural pathways. In fact, psychology and neuroscience had, and still have, a strong influence on the development of ANNs.

As their name suggests, the development of ANNs in the 1950s was originally inspired by the structure and function of biological neurons and networks (McCulloch & Pitts, 1943). The goal was to model cognition using models that stayed true to what, at the time, was known about brain processing. This approach, sometimes called "connectionism", is complementary to the more symbolic approaches to artificial intelligence advocated by, among others, the psychologists Alan Newell and Herbert Simon (e.g., Newell & Simon, 1956).

Initially, ANN models were composed of small arrays of artificial neurons with a single input layer and one output layer, such as the famous Perceptron model (Rosenblatt, 1958). These models were often used to demonstrate that they can learn simple logical functions (like the AND function) and embeddings (such as recoding decimals to binary numbers). Importantly, instead of explicitly programming the parameters (or "rules") of the models (as was common in symbolic artificial intelligence), the parameters (or "weights") of these networks were learned.

After a period of reduced interest in and general pessimism about the potential of artificial intelligence in general and ANNs in specific (known as the first "AI winter"), ANNs regained their popularity when datasets and computing power increased throughout the first decade of the 21st century. Massive datasets, such as ImageNet (Deng et al., 2009), and the technological developments that allowed networks to be trained on powerful graphical cards allowed researchers to develop more complex models, which often manifested as networks with many hidden layers (i.e., deep neural networks; LeCun, Bengio, & Hinton, 2015). Also, a particular variant of ANNs, convolutional neural networks (CNNs), were developed specifically to take advantage of locally structured data (such as images and audio) by rectangular sets of parameters, called filters, that were learned. This made it unnecessary to carefully hand-craft filters (as was common in computer vision at the time). CNNs proved to be extraordinarily powerful on visual tasks such as object recognition, which was convincingly demonstrated when a CNN beat the then state-of-the-art object recognition model during the 2012 ImageNet Large Scale Visual Recognition Challenge (Krizhevsky, Sutskever & Hinton, 2012).

Since 2012, research into ANNs has exploded, fueled both by a wider interest from the AI community as well as the adaptation of ANNs in different academic domains. The general trend has been to develop more complex architectures with more layers, often referred to as deep learning. Extreme examples are the winning network from the ImageNet competition in 2015, which featured over 150 hidden layers (He, Zhang, Ren, & Sun, 2016) and the aforementioned GPT-3 language model featuring 96 layers and 175 billion parameters. Apart from (computer)

vision, ANNs have been adopted by other scientific domains, including neuroscience (Marblestone, Wayne, & Kording, 2016), biology and medicine (Ching et al., 2018), genomics (Zou et al., 2019), and geoscience (Zhang, Zhang, & Du, 2016). Common to these applications is that (deep) neural networks are used to estimate a (very) nonlinear mapping from a set of inputs to a set of outputs.

The idea of stacking computations, as in the hierarchical layering of (deep) ANNs, has been integrated in other AI domains such as reinforcement learning and natural language processing (NLP). This has given rise to deep reinforcement learning, used in the previously mentioned AlphaZero model, and deep-learning-based NLP models, such as the aforementioned GPT-3 model. In these layered models, ANNs are used as a separate “perceptual module” (which may or may not be pre-trained) that embeds information, such as images or speech, in a way that a reinforcement learning model or a language model can use the ANN as a basis. Importantly, although reinforcement learning models may use ANNs as part of their model architecture, they are fundamentally different from ANNs in terms of goal. Because of this, and because artificial reinforcement learning also has roots in psychology and neuroscience, we now discuss this class of models in more detail.

2.2. Reinforcement learning

Artificial reinforcement learning models learn optimal behavior in a given context, through experience. In contrast to ANNs, which tend to excel at perceptual tasks, reinforcement learning is used to make an artificial agent learn to behave. Unsurprisingly, much of the progress in artificial reinforcement learning was inspired by and used concepts from animal and human learning (Sutton & Barto, 1998, see chapter 11). Early artificial reinforcement learning models, which did not feature ANN modules, directly adapted or extended models from animal learning. For example, the extended Rescorla-Wagner classical conditional model (Rescorla & Wagner, 1972), called temporal difference learning (Sutton & Barto, 1987), is extensively used in both psychology (e.g., Gershman, Pesaran, & Daw, 2009), neuroscience (Neftci & Averbeck, 2019) and AI research (e.g., deep Q learning; Mnih et al., 2015).

These early reinforcement learning models were often evaluated on relatively simple problems using a relatively simple environment (i.e., with few “states”) and a restricted and discrete set of possible behaviors (i.e., with few “actions”). Examples of such “toy problems” are navigating a maze (in animal learning) or learning to choose the most rewarding (virtual) slot machine (in artificial reinforcement learning). Driven by the desire to scale up reinforcement learning models to complex environments and behaviors, 21st century researchers started to add ANNs to their reinforcement learning models. In these models, ANNs function as a preprocessing step that embeds a complex (possibly continuous) environment as a more tractable set of representations that can be fed into the reinforcement learning model. By adding ANNs researchers could scale up reinforcement learning models that in turn led to the development of systems capable of (super) human performance on a range of tasks, including playing board games (Silver et al., 2018) and video games (Mnih et al., 2015), as well as autonomous driving (Kiran et al., 2020).

2.3. Other techniques

ANNs and reinforcement learning are arguably the dominant techniques used in modern AI models, but far from the only ones. Here, we briefly discuss two other important techniques: generative modelling and Bayesian models.

Within artificial intelligence, generative modelling is a broad set of models that are in fact (often) a specific implementation of a neural network. The difference with “standard” neural networks is that generative neural networks try to learn the distribution of the data (X) given a particular feature (y), $p(X | y)$, in contrast to standard (or “discriminative”) neural networks that try to predict a particular feature

given the data, $p(y | X)$. For example, instead of predicting object category (y) from images (X), a corresponding generative model could try to learn what a particular image (X) of a particular object (y) would look like. For example, researchers trained a generative model on images of faces using a generative adversarial network (GAN; e.g., Karras, Aila, Laine, & Lehtinen, 2017) and were then able to generate photo-realistic images of faces of non-existent people (convincingly demonstrated on <https://thispersondoesnotexist.com>). Generative models have also been successfully applied in NLP, with the recent GPT-3 model as a striking example (Brown et al., 2020), and across domains, such as the text-to-image model DALL-E, which is able to create images from text descriptions (see Fig. 1).

Bayesian modelling (or Bayesian reasoning) is another technique worth mentioning. Although Bayesian techniques are used more in psychology than in AI, many have argued that Bayesian modelling is crucial to overcoming the current limitations of AI (Lake, Ullman, Tenenbaum, & Gershman, 2017; Marcus, 2018; Pearl, 2019). What’s more, the abstract problem domains that Bayesian models are often applied to are much closer to the problems presented in human intelligence tests than those presented to ANNs.

Bayesian models come in many forms, but their common denominator is that they aim to formalize interpretable and generalizable mechanisms and knowledge. In contrast to most ANNs, which are trained on a very specific task (e.g., object recognition) and often need massive amounts of data, Bayesian models are designed to learn generalizable mechanisms that need little data to train (Tenenbaum, Kemp, Griffiths, & Goodman, 2011). In a way, this approach resembles symbolic AI, with the major difference being that modern Bayesian techniques are able to handle more complex data and models. To achieve models that can better generalize across domains (“far transfer”), Bayesian models stress the importance of abstract knowledge, functioning as priors in the Bayesian framework, that constrain their structure and learning process. A well-known application of Bayesian modelling is in the realm of intuitive physics (e.g., Battaglia, Hamrick, & Tenenbaum, 2013). Here, Bayesian models are used to model how humans (learn to) understand the physical world and its physical laws in a causal way. While Bayesian modelling is still mostly confined to realm of psychology, many have advocated integrating “connectionism” (represented by ANN modelling) and “symbolic AI” (represented by symbolic/Bayesian models; Lake et al., 2017; Marcus, 2018), and initial efforts have produced some promising results (Mao, Gan, Kohli, Tenenbaum, & Wu, 2019; Yi et al., 2018). The field of intelligence also sees a renewed interest in cognitive modelling (e.g., Frischkorn & Schubert, 2018), which may provide new links between AI and intelligence research.

To summarize, modern AI relies on a set of techniques that are able to solve complex tasks. ANNs are extraordinarily effective in the context of perceptual tasks (such as object and speech recognition), but are increasingly used as modules in other applications as well, for example by combining them with reinforcement learning and NLP models. Whereas ANNs are particularly suited to model perceptual processes, reinforcement learning represents a set of techniques that are very effective in terms of learning adaptive behavior and can outperform humans at a variety of tasks (including board and video games). In addition, generative modelling complements techniques for prediction by generating data, and Bayesian reasoning may provide AI with models that alleviate AI of its current shortcomings such as limited generalizability and data-intensive training regimes.

3. Is modern AI (more) psychologically plausible?

Schank (1980) stated that early AI work [1950’s–1970’s] “did not have a very particular psychological flavor”, where questions about how “people do various tasks... were not of particular interest”, although AI researchers were becoming increasingly interested in cognition in recent years [late 70’s]. With the impressive development of AI techniques in



Fig. 1. Examples of generated images in response to the text prompt “an armchair in the shape of an avocado” by DALL-E. Adapted from <https://openai.com/blog/dall-e>.

the past decades, we believe that the interest in cognition has only grown. Although much can be said about the limitations of current AI methods and models, such as the amount of data required to train state-of-the-art ANNs and inhuman complexity of some models (such as GPT-3), we see many ways in which AI has become more psychologically plausible since the time of Schank’s, 1980 article. We explain this according to David Marr’s (1982) three levels of analysis.

In terms of the computational level, the problems modern AI solve are much more humanlike than the simple toy problems often used in early AI efforts: modern AI is often used, as witnessed by the applications introduced before, to solve complex tasks in real-world contexts and using real-world data. And these real-world modern AI models sometimes achieve (super)human performance. Although most AI models that achieve (super)human performance are constrained to a single, narrow task, the observation that these artificial systems perform on par with actual human behavior is amazing nonetheless. Interestingly, one direction some researchers take to improve the performance of their models is to make the task that needs to be solved more “human”, or in other words, to align the computational level between artificial systems and humans. One characteristic example of this is “curiosity-based learning”, a specific implementation of reinforcement learning models which use “curiosity” instead of external reward to motivate learning (Burda, Edwards, Storkey, & Klimov, 2018).

Equivalence at the computational level, however, does not necessarily mean equivalence at the algorithmic level (i.e., a cognitive wheel). In other words, although modern AI and human cognition aim to solve similar problems, the way AI and humans implement these solutions may differ fundamentally. There is, however, reason to be optimistic in this regard, given the abundance of studies showing that ANNs and

reinforcement learning models learn representations and contain mechanisms that have plausible neural correlates. Especially in the domain of perception, researchers have shown that the representations learned by ANNs are similar to stimulus representations in the sensory cortex (e.g., Yamins et al., 2014). And, importantly, that deeper layers in ANNs contain information corresponding to stimulus features represented in brain regions further down the sensory processing hierarchy (Güçlü & van Gerven, 2015). Indeed, at the moment, features from deep neural networks are better able to explain neural activity during naturalistic vision than features from classical computational models of the visual cortex (Khaligh-Razavi & Kriegeskorte, 2014; Schrimpf et al., 2018). In addition to ANNs, (deep) reinforcement learning models appear to have plausible neural correlations, such as the correspondence of parameters from the temporal difference learning model with midbrain dopamine activity during animal and human learning (Schultz, Dayan, & Montague, 1997; for a review on the correspondence between reinforcement learning models and the brain, see O’Doherty, Lee, & McNamee, 2015). This correspondence between artificial and biological representations is further supported by the fact that many AI models, once they are trained, can often be successfully applied to other tasks and (perceptual) domains with minimal or no retraining – a technique called transfer learning (Torrey & Shavlik, 2010). The success of transfer learning (Yosinski, Clune, Bengio, & Lipson, 2014), which includes the aforementioned reuse of ANNs in deep reinforcement learning (Botvinick, Wang, Dabney, Miller, & Kurth-Nelson, 2020) and deep language models (Lu, Grover, Abbeel, & Mordatch, 2021), thus shows that AI systems can learn generalizable representations, just like the brain.

Another development that points to AI becoming more human at the

algorithmic level is that researchers started to incorporate concepts and ideas from psychology and cognitive science in their AI systems. For example, researchers have shown that implementing an artificial version of “attention” in both language models (Vaswani et al., 2017) and vision models (Linsley, Shiebler, Eberhardt, & Serre, 2018; Xu et al., 2015) improves their performance (see for a comprehensive review Lindsay, 2020), which mirrors recent developments in intelligence research that highlight the role of attention and attentional control in human intelligence (e.g., Burgoyne & Engle, 2020).

A similar example in the domain of reinforcement learning is the implementation of “episodic memory” and meta-learning (“learning to learn”) in artificial reinforcement learning models, which has been shown to dramatically improve the speed and efficiency when training these models (Botvinick et al., 2019; Wang, 2021). Psychologically-infused reinforcement models also appear to have plausible neural correlations (Wang et al., 2018).

Finally, even though early ANN models were in fact abstractions from biological neurons and networks, modern AI models are much more human and brain-like at the implementational level. Like the brain, modern AI models often feature hierarchical processing streams that contain a large set of parameters that match the complexity of the human brain. These developments include hierarchical layers, convolutions, synaptic pruning (Blalock, Ortiz, Frankle, & Guttat, 2020), and recurrence — all of which have plausible equivalents in the human brain (Kriegeskorte, 2015).

4. Is AI intelligent?

Although these recent breakthroughs are very impressive and based on methods with parallels in the human brain, the question whether such systems display real intelligence remains relevant.

On this, opinions are divided. As said, AI already outperforms humans in several areas. In other areas, more computation power, a combination of existing techniques and new techniques that are under development now, may lead to similar levels of (super)human performance. According to Newell and Simon’s (2007) physical symbol system hypothesis, there is no principal reason why an extrapolation of recent advances in AI should not lead to true intelligence.

Yet, the general assessment of how intelligent AI really is, and how much it still differs from human intelligence, is contradictory, at best. On the one hand, cognitive psychologists and neuroscientists generally acknowledge the merits of modern AI methods (Naseleris et al., 2018), especially in the context of perception (VanRullen, 2017) and language (Henderson, 2020). On the other hand, others argue that previous critiques of early AI systems — including the famous arguments by Searle (1980), Dreyfus (1965), and Penrose (1994) — still hold for modern AI systems. Common to these critiques is that AI systems lack understanding.

Chess again helps to illustrate the problem. In certain so-called “anti-computer positions” (such as the impenetrable fortress of pieces shown in Fig. 2), human chess players quickly see the point of the position, while AI systems are only able to evaluate the position after considering long sequences of moves. For example, the Leela Chess Zero engine has a difficult time with this kind of position. Humans somehow have the ability to use both general reasoning and automatic pattern recognition when playing chess, which suggests that the critique about lack of understanding is still relevant.

In the case of chess, the “lack of understanding” limitation is perhaps temporary. Leela chess might, in time, learn to understand “anti-computer positions”, where perhaps a modular approach in which Leela chess is combined with a symbolic or Bayesian approach might work. But this clearly would not end the debate. According to Searle, Dreyfus and others, the problem lies deeper. AI systems lack semantics, feelings, a body and goals. Some of these limitations can perhaps be overcome by using robots that do have a body or real time interactions with the actual (rather than digital, simulated) world, which is often thought to be



Fig. 2. A typical anti-computer position. Human chess players quickly see that black, in spite of its material advantage, can not make progress. The best computer chess programs assess the position as much better for black.

necessary to grasp causality — an element that is sorely absent from most AI systems (Lake et al., 2017; Tenenbaum et al., 2011).

So how should we assess the “lack of understanding” issue from the intelligence researcher’s point of view? We will answer this with two definitions of intelligence in mind. The first is a practical definition: intelligence is what the intelligence test measures. This definition is not taken very seriously (however, see van der Maas, Kan, & Borsboom, 2014), but how well AI performs on human intelligence tests provides an interesting insight into the limitations of AI nonetheless (which is discussed in the next section).

In the most commonly accepted definitions of intelligence, we see two important requirements. One is the ability to deal with various forms of information and to solve all kinds of cognitive problems, and the other is the ability to both quickly and effectively learn how to deal with new situations (Legg & Hutter, 2007).

AI clearly meets the first requirement, as there is arguably nothing wrong with AI’s crystallized intelligence: AI programs can play games like go and chess, compose music, create art, solve complex math problems, provide medical diagnosis, etc. But things are different when we confront a trained AI system with completely novel situations. Schank (1980) already mentioned the problem of generalizability. Humans are able to transfer solutions from one problem domain to another, apply general, abstract concepts (such as presented in “anti-computer problems” in chess) in reasoning and to develop solutions for completely new problems. AI, on the other hand, often fails to learn such abstract rules and if they do, it requires an immense amount of training examples (Lake et al., 2017).

We think it is justified to say that generalization is still a weakness of AI systems. At the same time, we see many promising developments. We already mentioned transfer learning, which works by virtue of learning general representations that facilitate broad generalization. Multitask learning is an active area of research (Ruder, 2017). Ongoing attempts to integrate deep learning and symbolic AI or Bayesian reasoning may also be crucial for progress in this research area. New work on one-shot learning (Vinyals, Blundell, Lillicrap, Kavukcuoglu, & Wierstra, 2016) may also lead to breakthroughs.

Whether these future developments also make AI systems

‘understand’ things like we do is a topic of hot philosophical debate (for further discussion we refer to [Cole, 2020](#)).

4.1. How well do AI systems perform on intelligence tests?

[Newell \(1973\)](#) argued that if a single AI system could solve a diverse set of intelligence test problems, then we could consider it to be intelligent. Since the 1960’s there have been waves of efforts to have AI systems solve intelligence test problems. This started with [Evan’s \(1964\)](#) ANALOGY program that could solve geometric analogies from the WAIS and was intended to be able to generalize to other reasoning tasks (but not pursued). In [Hernández-Orallo, Martínez-Plumed, Schmid, Siebers, and Dowe’s \(2016\)](#) review we see that AI has tackled a broad range of intelligence subtests, ranging from letter and number series to block design and exclusion tasks, and without a doubt excels at what are considered tests of crystallized intelligence (e.g., math, vocabulary) and processing capacity (e.g., memory, speed). The one test that has continually popped up in the literature over the years is Raven’s Progressive Matrices (RPM). AI programs have been used to study how people solve the RPM ([Carpenter, Just, & Shell, 1990](#); [Lovett & Forbus, 2012](#)) to fully independent programs that take raw images as input and induce rules and select the correct option ([Correa, Prade, & Richard, 2012](#); [Kunda, McGreggor, & Goel, 2013](#)) or even generate it ([Pekar, Benny, & Wolf, 2020](#)), in many cases reaching human level performance.

An especially interesting case is that of Bongard problems ([Bongard, 1967](#) in [Hofstadter, 1979](#)). In [Fig. 3](#) you see a simple Bongard problem where the test-taker has to induce which pattern the images on the left follow and those on the right do not. [Hofstadter \(1979\)](#) considered these types of reasoning tasks to lie at the core of human intelligence. His PhD student [Foundalis \(2006\)](#) created Phaeaco, a (by today’s standards) rudimentary AI system, to try to solve Bongard problems. Phaeaco contains a cognitive architecture with modules to process the images, match visual patterns, learn, and to store and retrieve learned patterns. Phaeaco was relatively successful, although the program did not reach human level performance. More recently, [Nie et al. \(2020\)](#) tried to solve Bongard problems using state-of-the-art deep learning methods. Their best model achieved 66% accuracy, and yet performance still falls behind that of humans who are able to solve about 90% of their Bongard

benchmark problems.

So, although in many cases intelligence tests can be solved by AI programs, most of these are specialized programs that can only solve one particular task. Some attempts have been made to create more general AI systems. For example, in [Nie et al.’s \(2020\)](#) study one of the models they used to solve Bongard problems was built upon one of the most successful Raven-like item solvers ([Barrett, Hill, Santoro, Morcos, & Lillicrap, 2018](#)), alas this model performed only around the level of chance. So currently, we cannot say that AI is “intelligent” by the psychometric definition in the sense of being able to achieve human level performance on all subtests of an intelligence test using one general AI system.

But, what if we approach this question from the other side: How well do humans perform on tests of intelligence for AI? Over the AI lifespan, an interesting field has emerged that takes the psychometric approach to assessing general intelligence in AI, but then applied to problems generated specifically for AI research (e.g., [Bringsjord & Schimanski, 2003](#); [Chollet, 2019](#); [Evans, 1964](#); [Hernández-Orallo et al., 2016](#)). This approach is characterized by hundreds or thousands of (often computer generated) items fit for a non-verbal test of intelligence ([Barrett et al., 2018](#); [Chollet, 2019](#); [Liu et al., 2019](#); [Nie et al., 2020](#); [Zhang, Gao, Jia, Zhu, & Zhu, 2019](#)). These item banks are referred to as benchmarks and used to objectively compare the performance of different AI systems, and, in some cases, also compare AI performance to that of humans. There is a clear role for Intelligence researchers in this field given our expertise in test development, sampling from the general population and in assessing human performance. There is also a role for AI in improving existing measures of human intelligence (e.g., automatic item generation; [Gierl & Haladyna, 2012](#)). For AI, currently more intelligence test benchmarks are needed (e.g., for problem analogies, [Ichien, Lu, & Holyoak, 2020](#)). Also, these benchmarks only sometimes include humans among their comparison samples, which appear limited to adults ([Nie et al., 2020](#)) or only a few experts ([Chollet, 2019](#)). Unsurprisingly, at this point in time humans generally outperform AI systems on these types of AI “fluid” intelligence tests (e.g., [Chollet, 2019](#); [Nie et al., 2020](#)). Given the activity in the field, perhaps many strides will be made in the coming years to achieve human-like performance on these benchmarks, but also psychology can inspire innovation in AI by providing a developmental account of if and how humans learn to solve

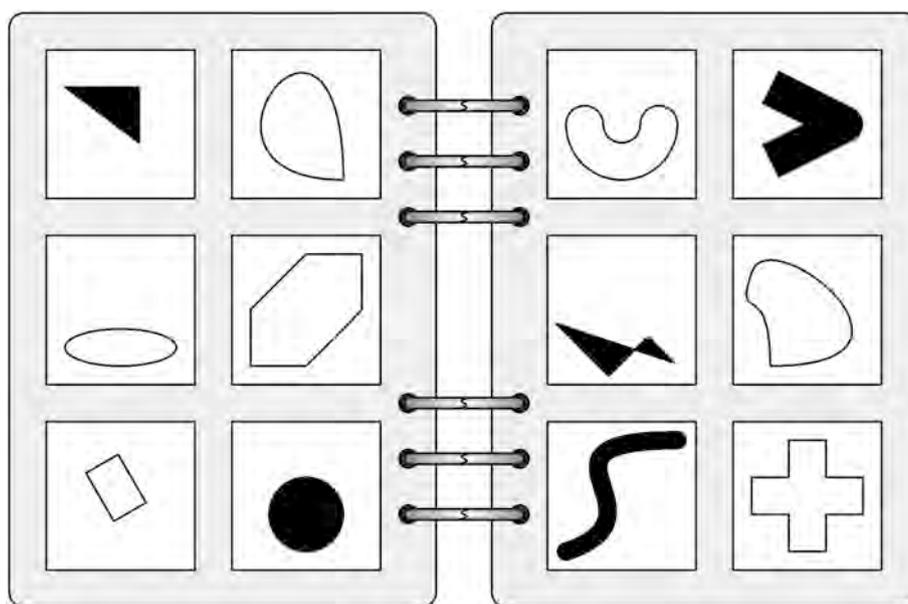


Fig. 3. Bongard visual pattern recognition problem, where the test-taker has to induce which pattern the images on the left follow and those on the right do not. In this example, the images on the left are convex and those on the right are concave. From https://upload.wikimedia.org/wikipedia/commons/thumb/e/e9/Bongard_problem_convex_polygons.svg.

these AI benchmarks.

4.2. Individual differences in intelligence

One question posed at the start of the paper, that has not yet been addressed, concerns the issue of individual differences. The intelligence researcher traditionally focuses on individual differences in cognitive functioning, and not so much on the cognitive architecture of the mind. Is AI then relevant for individual differences research? We can now answer this question with a definitive yes.

The first somewhat trivial reply is that a deeper understanding of how the human mind operates informs us about possible sources of individual differences in intelligence, but modern AI offers more interesting and explicit insights into individual differences in (artificial) intelligence.

Chess, and in particular the previously discussed Leela Zero Chess models, again provides an interesting example. What has to be done before Leela Chess Zero can play superhuman chess? First, a network architecture, defined by the number of layers and filters of the neural network, is chosen. The choice depends on the hardware and computation time available, where the largest networks are trained for months on an extensive and powerful array of computer hardware. Second, choices about the training methods are made. Official variants use reinforcement learning and self-play and learn largely from the ground up, but other variants incorporate supervised learning and rich knowledge bases (e.g., on openings) to improve performance. According to the Leela chess website the LCZero team has produced hundreds of thousands of AI chess playing systems.

One could thus say that the development of all of these Leela chess variants mirrors the case of people in intelligence research. Humans and Leela chess both have information processing systems of very high complexity, to such a degree that we do not really understand their mechanisms. We could say that both nature (the network architecture and the hardware used to train the system) and nurture (training regime) are involved. And we have individual differences: Which variants of Leela chess are actually the best?

The strength of chess engines is measured by the Elo rating system, which is mathematically almost identical to the Rasch model, the root model of item response theory that forms the basis of the best measurement models in intelligence research. The Elo rating system suffices when deciding which chess system is best. However, Elo ratings provide a one-dimensional measurement of strength, while the many Leela variants probably differ in other dimensions too. Some may excel in speed chess, some may be exceptionally good at end games or tactical positions. Some Leelas might be quite adventurous and others rather conservative. Insight into these individual differences may be best achieved by devising tests with subtests consisting of items meant to measure these sub-abilities in chess, akin to what has been done in human chess research (van der Maas & Wagenmakers, 2005).

The factor analytic study of these artificial individual differences is of interest too, not only to understand Leela zero chess better, but to also see what the factor analytic approach actually brings us in such a case. Suppose, and we think this is quite probable, we again find a general factor, what then does this mean? Is it a formative or a reflective factor (van der Maas, Kan, Marsman, & Stevenson, 2017)? Does this g-factor represent one underlying source of chess strength (computing speed, for instance), or does it summarize the collective force of the many building blocks of Leela? Note that a part of the success of current AI lies in the combination of techniques such as deep learning, reinforcement learning and powerful hardware. Having such a mirror system creates new possibilities to study intelligence and also raises interesting questions regarding individual differences research.

Similarly, as modern AI systems are systems that learn, it is interesting to rethink questions regarding the development of intelligence. One example might be the study of the third source hypothesis (Kan, Ploeger, Raijmakers, Dolan, & Van Der Maas, 2010), that states that

nonlinear epigenetic processes cause variation that cannot be explained by genetic and environmental sources, nor by their interactions. This would occur when identical copies of Leela zero chess with the same learning regime diverge during development. We again used Leela as an example, but many other modern AI systems could be used instead (Scholte, 2018).

One could even imagine that AI could be used to produce artificial research participants for experiments in intelligence research that would be unethical to perform on humans. Such experiments could consist of very rigid training regimes or drastic changes in the network structure of neural networks (e.g., adding lesions or simulating brain disease) after they have been trained.

5. Discussion

AI has seen multiple cycles of enthusiasm and disappointment, but the current wave seems to be of a different order. As we stated in the introduction of this paper, one of the original goals of AI was to learn more about human intelligence. This endeavor could be misguided as AI may only produce “cognitive wheels”, techniques that have no equivalent in human cognition. In this paper we argued that this might have been true for some older approaches (e.g., brute force search techniques), but is less the case for much of current AI. The progress made in recent years is certainly technologically driven, but inspired by biological and psychological knowledge about human information processing and learning.

We expect that the recent progress in AI will change the way we think about intelligence. AI forces us to rethink the definition of intelligence. Definitions that center on just information processing and problem solving are perhaps insufficient. Shank’s observation that intelligence is all about generalization has, so far, withstood the test of time. Many information processing problems, from processing speech to playing chess, appear to be less difficult than perhaps expected. The really hard problem is to deal with completely novel cases. One requirement for solving this hard problem is the ability to learn invariant and thus generalizable patterns. And especially with regard to learning, the progress in AI has been spectacular. The main difference between AI systems of the past, such as expert systems, and modern AI is the fact that they learn. That deep learning and reinforcement learning, the core techniques in current AI, have deep roots in psychology is remarkable and promising for studying how artificial and human intelligence are related.

AI is relevant to intelligence research because it enhances our understanding of the core mechanisms of human cognition. How the immense neural systems in our brain are able to process extremely complicated information such as speech and produce logical thinking is an extremely difficult question. Having an artificial system that performs such tasks using the same basic principles is extremely useful. Classic questions regarding the modularity of the mind, the origin of creativity, and the organization of long-term memory spring to mind. In addition, we argued that the psychological relevance of AI extends to unexpected areas such as the understanding of individual differences and the development of cognition. It is relatively easy to create a population of AI systems with minor differences in architecture and training regime. Modern AI provides us with a new playing field for individual differences research.

On a practical level we expect fruitful interactions regarding the measurement of natural and artificial intelligence. As modern AI systems are incredibly complex, our experience in examining such systems may be relevant for AI. Vice versa, insights from AI may lead to new developments in (adaptive) intelligence testing and educational interventions.

We attempted to shed light on the future of intelligence research from the point of view of AI. Our overview is necessarily limited and probably quickly outdated, but hopefully we have given intelligence researchers some insights in the rapid developments in AI and the

possible consequences for our field.

Acknowledgements

This research was partly supported by the Jacobs Foundation Fellowship 2019-2021 awarded to Claire Stevenson (2018 1288 12).

References

- Barrett, D., Hill, F., Santoro, A., Morcos, A., & Lillicrap, T. (2018, July). Measuring abstract reasoning in neural networks. In *International conference on machine learning* (pp. 511–520). PMLR.
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, *110* (45), 18327–18332.
- Blalock, D., Ortiz, J. J. G., Frankle, J., & Gutttag, J. (2020). What is the state of neural network pruning? *arXiv preprint arXiv:2003.03033*.
- Blanch, A. (2020). *Chess and individual differences*. Cambridge University Press.
- Bongard, M. (1967). *The problems of knowledge formation*.
- Botvinick, M., Ritter, S., Wang, J. X., Kurth-Nelson, Z., Blundell, C., & Hassabis, D. (2019). Reinforcement learning, fast and slow. *Trends in Cognitive Sciences*, *23*(5), 408–422.
- Botvinick, M., Wang, J. X., Dabney, W., Miller, K. J., & Kurth-Nelson, Z. (2020). Deep reinforcement learning and its neuroscientific implications. *Neuron*, *107*(4), 603–616.
- Bringsjord, S., & Schimanski, B. (2003, August). What is artificial intelligence? Psychometric AI as an answer. In *International joint conference on artificial intelligence* (pp. 887–893).
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... Amodei, D. (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Burda, Y., Edwards, H., Storkey, A., & Klimov, O. (2018). Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*.
- Burgoyne, A. P., & Engle, R. W. (2020). Attention control: A cornerstone of higher-order cognition. *Current Directions in Psychological Science*, *29*(6), 624–630.
- Carpenter, P. A., Just, M. A., & Shell, P. (1990). What one intelligence test measures: A theoretical account of the processing in the Raven Progressive Matrices Test. *Psychological Review*, *97*(3), 404.
- Ching, T., Himmelstein, D. S., Beaulieu-Jones, B. K., Kalinin, A. A., Do, B. T., Way, G. P., ... Greene, C. S. (2018). Opportunities and obstacles for deep learning in biology and medicine. *Journal of the Royal Society Interface*, *15*(141), 20170387.
- Chollet, F. (2019). On the measure of intelligence. *arXiv preprint arXiv:1911.01547*.
- Cole, D. (2020). The Chinese room argument. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Winter 2020 ed.).
- Correa, W. F., Prade, H., & Richard, G. (2012, August). When intelligence is just a matter of copying. In , vol. 12. *European conference on artificial intelligence* (pp. 276–281).
- Cronbach, L. J. (1957). The two disciplines of scientific psychology. *American Psychologist*, *12*(11), 671.
- de Groot, A. D. (1946). *Het denken van den schaker: een experimenteel-psychologische studie* (p. 315). Amsterdam: Noord-Hollandsche Uitgevers Maatschappij.
- Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009, June). Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition* (pp. 248–255). IEEE.
- Dennett, D. C. (2006). The frame problem of AI. *Philosophy of Psychology: Contemporary Readings*, *433*, 67–83.
- Dreyfus, H. (1965). *Alchemy and AI*. Santa Monica, CA: RAND Corporation.
- Dubhashi, D., & Lappin, S. (2017). AI dangers: Imagined and real. *Communications of the ACM*, *60*(2), 43–45.
- Evans, T. G. (1964). *A program for the solution of a class of geometric-analogy intelligence-test questions*. Bedford, MS: Air Force Cambridge Research Labs.
- Foundalis, H. E. (2006). *Phaeaco: A cognitive architecture inspired by bongard's problems*. Doctoral dissertation. Indiana University https://www.foundalis.com/res/diss_rese_arch.html.
- Frischkorn, G. T., & Schubert, A. L. (2018). Cognitive models in intelligence research: Advantages and recommendations for their application. *Journal of Intelligence*, *6*(3), 34.
- Gershman, S. J., Pesaran, B., & Daw, N. D. (2009). Human reinforcement learning subdivides structured action spaces by learning effector-specific values. *Journal of Neuroscience*, *29*(43), 13524–13531.
- Gierl, M. J., & Haladyna, T. M. (2012). *Automatic item generation: Theory and practice*. Routledge.
- Gobet, F., & Charness, N. (2018). In K. A. Ericsson, R. R. Hoffman, A. Kozbelt, & A. M. Williams (Eds.), *Expertise in chess*. Cambridge handbooks in psychology.
- Gobet, F., & Chassy, P. (2009). Expertise and intuition: A tale of three theories. *Minds and Machines*, *19*(2), 151–180.
- GPT-3. (2020, September 8). A robot wrote this entire article. Are you scared yet, human? *The Guardian*.
- Güçlü, U., & van Gerven, M. A. (2015). Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *Journal of Neuroscience*, *35*(27), 10005–10014.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 770–778).
- Henderson, J. (2020). The unstoppable rise of computational linguistics in deep learning. *arXiv preprint arXiv:2005.06420*.
- Hernández-Orallo, J., Martínez-Plumed, F., Schmid, U., Siebers, M., & Dowe, D. L. (2016). Computer models solving intelligence test problems: Progress and implications. *Artificial Intelligence*, *230*, 74–107.
- Hofstadter, D. (1979). In *Gödel, Escher, Bach: an eternal golden braid*. New York, NY: Basic Books.
- Ichien, N., Lu, H., & Holyoak, K. J. (2020). Verbal analogy problem sets: An inventory of testing materials. *Behavior Research Methods*, *52*, 1803–1816.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, *1*(9), 389–399.
- Kan, K. J., Ploeger, A., Raijmakers, M. E., Dolan, C. V., & Van Der Maas, H. L. (2010). Nonlinear epigenetic variance: Review and simulations. *Developmental Science*, *13* (1), 11–27.
- Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Kaul, V., Enslin, S., & Gross, S. A. (2020). The history of artificial intelligence in medicine. *Gastrointestinal Endoscopy*, *92*(4), 807–812.
- Khaligh-Razavi, S. M., & Kriegeskorte, N. (2014). Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS Computational Biology*, *10*(11), Article e1003915.
- Kiran, B. R., Sobh, A., Talpaert, V., Mannion, P., Sallab, A. A., Yogamani, S., & Pérez, P. (2020). Deep reinforcement learning for autonomous driving: A survey. *arXiv preprint arXiv:2002.00444*.
- Kriegeskorte, N. (2015). Deep neural networks: A new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science*, *1*, 417–446.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, *25*, 1097–1105.
- Kunda, M., McGreggor, K., & Goel, A. K. (2013). A computational model for solving problems from the Raven's Progressive Matrices intelligence test using iconic visual representations. *Cognitive Systems Research*, *22*, 47–66.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, *40*.
- Lane, D. M., & Chang, Y. H. A. (2018). Chess knowledge predicts chess memory even after controlling for chess experience: Evidence for the role of high-level processes. *Memory & Cognition*, *46*(3), 337–348.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444.
- Legg, S., & Hutter, M. (2007). A collection of definitions of intelligence. *Frontiers in Artificial Intelligence and Applications*, *157*, 17.
- Lindsay, G. W. (2020). Attention in psychology, neuroscience, and machine learning. *Frontiers in Computational Neuroscience*, *14*, 29.
- Linsley, D., Shiebler, D., Eberhardt, S., & Serre, T. (2018). Learning what and where to attend. *arXiv preprint arXiv:1805.08819*.
- Liu, Y., He, F., Zhang, H., Rao, G., Feng, Z., & Zhou, Y. (2019). How well do machines perform on IQ tests: A comparison study on a large-scale dataset. In *International joint conference on artificial intelligence* (pp. 6110–6116).
- Lovett, A., & Forbus, K. (2012). Modeling multiple strategies for solving geometric analogy problems. In , vol. 34. *Proceedings of the annual meeting of the cognitive science society*. No. 34.
- Lu, K., Grover, A., Abbeel, P., & Mordatch, I. (2021). Pretrained transformers as universal computation engines. *arXiv preprint arXiv:2103.05247*.
- van der Maas, H. L., Kan, K. J., & Borsboom, D. (2014). Intelligence is what the intelligence test measures. Seriously. *Journal of Intelligence*, *2*(1), 12–15.
- van der Maas, H. L., Kan, K. J., Marsman, M., & Stevenson, C. E. (2017). Network models for cognitive development and intelligence. *Journal of Intelligence*, *5*(2), 16.
- van der Maas, H. L., & Wagenmakers, E. J. (2005). A psychometric analysis of chess expertise. *The American Journal of Psychology*, *29*–60.
- Mao, J., Gan, C., Kohli, P., Tenenbaum, J. B., & Wu, J. (2019). The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. *arXiv preprint arXiv:1904.12584*.
- Marblestone, A. H., Wayne, G., & Kording, K. P. (2016). Toward an integration of deep learning and neuroscience. *Frontiers in Computational Neuroscience*, *10*, 94.
- Marcus, G. (2018). Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*.
- Marr, D. (1982). *Vision: A computational approach*. Freeman & Co.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, *5*(4), 115–133.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533.
- Naselaris, T., Basset, D. S., Fletcher, A. K., Kording, K., Kriegeskorte, N., Nienborg, H., ... Kay, K. (2018). Cognitive computational neuroscience: A new conference for an emerging discipline. *Trends in Cognitive Sciences*, *22*(5), 365–367.
- Neftci, E. O., & Averbach, B. B. (2019). Reinforcement learning in artificial and biological systems. *Nature Machine Intelligence*, *1*(3), 133–143.
- Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing*. New York, NY: Academic Press.
- Newell, A., & Simon, H. (1956). The logic theory machine – a complex information processing system. *IRE Transactions on Information Theory*, *2*(3), 61–79.
- Newell, A., & Simon, H. A. (2007). Computer science as empirical inquiry: Symbols and search. In *ACM turing award lectures* (p. 1975).
- Nie, W., Yu, Z., Mao, L., Patel, A. B., Zhu, Y., & Anandkumar, A. (2020). BONGARD-LOGO: A new benchmark for human-level concept learning and reasoning. *arXiv preprint arXiv:2010.00763*.

- O'Doherty, J. P., Lee, S. W., & McNamee, D. (2015). The structure of reinforcement-learning mechanisms in the human brain. *Current Opinion in Behavioral Sciences*, 1, 94–100.
- Pearl, J. (2019). The seven tools of causal inference, with reflections on machine learning. *Communications of the ACM*, 62(3), 54–60.
- Penrose, R. (1994). *Shadows of the Mind*. Oxford: Oxford University Press.
- Pekar, N., Benny, Y., & Wolf, L. (2020). Generating correct answers for progressive matrices intelligence tests. *arXiv*. preprint arXiv:2011.00496.
- Rescorla, R. A., & Wagner, A. R. (1972). In A. H. Black, & W. F. Prokasy (Eds.), *A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement, classical conditioning II* (pp. 64–99). Appleton-Century-Crofts.
- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386.
- Ruder, S. (2017). An overview of multi-task learning in deep neural networks. *arXiv*. preprint arXiv:1706.05098.
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215.
- Schank, R. C. (1980). How much intelligence is there in artificial intelligence? *Intelligence*, 4, 1–14.
- Scholte, H. S. (2018). Fantastic DNimals and where to find them. *Neuroimage*, 180, 112–113.
- Schrimpf, M., Kubilius, J., Hong, H., Majaj, N. J., Rajalingham, R., Issa, E. B., ... DiCarlo, J. J. (2018). Brain-score: Which artificial neural network for object recognition is most brain-like? *BioRxiv*, 407007.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- Searle, J. R. (1980). Minds, brains, and programs. *Verbal Behaviour as the Hallmark of Intelligence. The Turing Test*, 201–224.
- Siciliano, B., & Khatib, O. (Eds.). (2016). *Springer handbook of robotics*. Springer.
- Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., ... Hassabis, D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419), 1140–1144.
- Sutton, R. S., & Barto, A. G. (1987, July). A temporal-difference model of classical conditioning. In *Proceedings of the ninth annual conference of the cognitive science society* (pp. 355–378).
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning*, 135. Cambridge: MIT press.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285.
- Torrey, L., & Shavlik, J. (2010). Transfer learning. In E. Soria, J. Martin, R. Magdalena, M. Martinez, & A. Serrano (Eds.), *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques* (pp. 242–264). IGI global.
- Van Harreveld, F., Wagenmakers, E. J., & Van Der Maas, H. L. (2007). The effects of time pressure on chess skill: An investigation into fast and slow processes underlying expert performance. *Psychological Research*, 71(5), 591–597.
- VanRullen, R. (2017). Perception science in the age of deep neural networks. *Frontiers in Psychology*, 8, 142.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. *arXiv*. preprint arXiv:1706.03762.
- Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., & Wierstra, D. (2016). Matching networks for one shot learning. *arXiv*. preprint arXiv:1606.04080.
- Wang, J. X. (2021). Meta-learning in natural and artificial intelligence. *Current Opinion in Behavioral Sciences*, 38, 90–95.
- Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., ... Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6), 860–868.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., ... Bengio, Y. (2015, June). Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning* (pp. 2048–2057). PMLR.
- Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., & DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23), 8619–8624.
- Yi, K., Wu, J., Gan, C., Torralba, A., Kohli, P., & Tenenbaum, J. B. (2018). Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. *arXiv*. preprint arXiv:1810.02338.
- Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks? *arXiv*. preprint arXiv:1411.1792.
- Zhang, C., Gao, F., Jia, B., Zhu, Y., & Zhu, S. C. (2019). Raven: A dataset for relational and analogical visual reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 5317–5327).
- Zhang, L., Zhang, L., & Du, B. (2016). Deep learning for remote sensing data: A technical tutorial on the state of the art. *IEEE Geoscience and Remote Sensing Magazine*, 4(2), 22–40.
- Zou, J., Huss, M., Abid, A., Mohammadi, P., Torkamani, A., & Telenti, A. (2019). A primer on deep learning in genomics. *Nature Genetics*, 51(1), 12–18.