2

A Deep Blue Grasshopper: Playing Games with Artificial Intelligence

ANDY MIAH

> But perhaps this has nothing to do with chess, but is a game to the death between two forms of intelligence—the bounded field of postmodern intelligence and the unbounded vectors and virtualities of digital intelligence. A game of symbolic exchange in which Kasparov as the last chess player needs not only to be defeated, but humiliated by his own imaginary catastrophe. And he was.
>
> —David Cook and Arthur Kroker,
> "The Digital Prince and the Last Chess Player"

Can a machine think as you and I do? Can a machine have consciousness, feelings, or knowledge, rather than just simulate these capacities? These are the interests of artificial intelligence (AI) research and these questions have been asked for over half a century. The philosophy of AI is far reaching and continues to be fuelled by the early ideas of Alan Turing (1950) and his contemporary John Searle (1980). The critical questions remain the same as they always have been: Are the kinds of capabilities demonstrated by machines an indication of intelligence and, if so, is this kind of intelligence similar to that of sentient beings, such as humans?

Various authors have questioned this approach to philosophizing AI, asking whether the comparison between "wetware" and "hardware" is appropriate. Does it matter if machines have the same form of intelligence to humans? Does this manner of framing the question limit our conceptualization of intelligence? Indeed, if one models AI on sentient beings, does this constitute an unexplained speciesistic move? These questions about intelligent life are integral to broader philosophical inquiries into the nature of being. Modeling intelligence on biological life raises doubts about the existence of the soul or the spiritual, since the ends of this research have the potential to explain away these concepts, reducing being human to "mere" physics.

Various defenses of wetware's superiority consist of claims about the some spiritual, soulful, or simply intangible biologic quality that cannot be operationalized and replicated by computing (Haldane 1997). Others consider this discourse to be speciesistic in that it unreasonably measures nonbiological matter in terms of the specific biological capacities of humans. On this view, the response might be that, it is not that the human being has a unique quality, but merely that it has not yet been shown how human cognition is entirely a mechanistic process.

Traditionally, the method through which AI has been tested is game theory, or, more accurately, through the creation of closed-system games, which do not require participants to encounter any of the more morally problematic choices that face people in real-life situations (Luger and Subblefield 1998). As Lee explains:

> Games such as chess, go, and drafts (or checkers) are ideal games for computerization. This is partly because they are games of perfect information as all the pieces are visible to both players and there is no element of chance or probability. This means that it should be possible to calculate the ideal move for any given situation, and hence, in theory at least, a computer should be able to play perfect chess. (1989, 108–9)

Yet, while a considerable amount of work has sought to propose games that test for intelligence, few take into account what it means to play a game. I will argue that the game playing of AI theory functions as mere algorithm rather than game playing and that further thought is needed. Moreover, I will argue that, contrary to Lee (1989), the specific game of chess is not the best game we might utilize to test for the comparative intelligence capacity of humans and machines.

Within the philosophy of AI and the related cultural discourse surrounding it, chess occupies a privileged position as a measure for understanding how close machines have come to reaching the intelligence capabilities of humans. A significant amount of media attention and research has been focused on creating computers that can beat Grandmasters of chess. In this respect, chess has been afforded a rather privileged game status, for the considerable intellectual capability that is required of its players.

Opinions vary about what kind of intelligence is required for being good at chess. From one perspective, Grandmaster status is an indication of intellectual superiority just as capability in chess is a measure of one's mind. Such a view would conclude that the great human minds can be found sitting across a checkered board, busily thinking through move possibilities. Chess is a game that involves creative thought, imagination, the possibility of abstraction, and thinking ahead. For others, chess is merely a computational form of knowledge, a specific kind of knowledge that allows one to consider a number of complex options at the same time, as well as to think laterally. From this perspective, being good at chess is certainly some measure of intellect, but it does not approximate the richer and wider construct of human intelligence, as it is simply not measuring other kinds of knowledge.

It is for this reason that the present chapter wishes to inquire into the suitability of chess as an indication of machine intelligence. At this juncture in the developing philosophy of information, it seems pertinent to reconsider this enduring debate. Specifically, this chapter will argue that other kinds of games are more suitable than chess to serve as a measure of machine intelligence compared with human intelligence. This is reflected in the title, *A Deep Blue Grasshopper,* which elaborates on the notion of game playing and endeavors to consider whether Deep Blue (and its subsequent incarnations) responds to such an idea. It is suggested that the privileged status of chess has been largely unwarranted, in part deriving from the cultural significance of chess as a game for intellectuals. Instead of wondering what the implications of a machine playing chess are, a more interesting question to ask would be "can a machine play soccer?" The example of soccer is relatively arbitrary and serves simply to illustrate that research into the philosophy of AI has appropriated a relatively biased notion of knowledge, which is premised wholly upon discrete and *relatively* uncomplicated logic.

In contrast to the rationale of using chess playing as a measure of intelligence, I suggest that our criteria must encompass broader characteristics of human knowledge than is demonstrated solely through the linear tests of Turing. A more suitable measure of humanness will be to reveal whether machines are able to exhibit the characteristics of play activity, which are present in a number of physical games.[i] Importantly, the reverse of this is not true—human play is not merely the replication of patterned behaviors. Nevertheless, it is first necessary to articulate the artificial-intelligence position and its limits.

Contextualizing AI

The work of Alan Turing became widely known through his famous Turing test and Turing machines. The former entails a proposition for considering how one might conclude that computers have intelligence comparable to that of humans. Through his imitation game, Turing argued that if it is possible for a computer to fool a human being into believing that it is human (through conversing with one another on a computer screen), then it can be concluded that the computer has the same degree of intelligence as a human. The implications of this hypothesis were far reaching, since the possibility of the two being indistinguishable aspired to collapse the special status of humans as defined by a Cartesian view. The possibility that humans are not defined or, at least, not unique because of their mental capabilities had threatened significant scientific research that endeavored to speak the contrary. If a computer could be passed off as a human, then, Turing argued, for all intents and purposes, a computer should be regarded as human. After all, our only way of interacting with other humans is through language, which is precisely what the Turing test measures.

In contrast, Searle (1980) considers that Turing's test is misleading and, instead, analogizes it to his famed "Chinese room argument." Here, Searle argues that even if the computer were able to fool another human into believing that it was talking to another human, this would not connote its having intelligence. Searle's Chinese room and his extended version, named the Chinese gym, describes how a computer program might be able to demonstrate its syntactic grasp and thus, provide the appearance of intelligence through a comprehensible conversation. However, Searle contends that it would lack any comprehension of the semantic quality of words, including those ascribed to the emotions, which are considered as constitutive of that which defines human intelligence.

Searle asks that we imagine a room with a person in it who receives Chinese symbols through a slot in a door. The person has a manual that allows him or her to match up the Chinese symbols with the English counterpart and thus, the person can output the translation. To anyone outside of the room, it would appear that the person has translated the Chinese words and thus understands what they refer to. In actual fact, the person has demonstrated merely a capability for pattern matching and has no real understanding of the meaning of the Chinese symbols (Searle 1990).

The example reveals a disparity within research about AI that is premised upon its artificiality standing against some, supposedly, natural intelligence. Furthermore, such natural intelligence tends to be directed towards that exhibited by a human being. Thus, it implies that if anything should be considered intelligent, then at least we must grant that humans are). Consequently, it is possible to identify two assumptions and thus distinct pursuits of AI research:

> **Strong AI**: That intelligence is entirely algorithmic and can be formalized through the kinds of computer programs that we currently use. The aims of such research are to determine what degree of intelligence can be held by a machine and to what degree such intelligence can be used.
>
> **Weak AI**: That intelligence is comprised partly, if not largely, by the semantic, nonalgorithmic characteristics of information that are not accessible to the syntactically defined computer programs. The aims of such research are to determine whether this degree of intelligence is comparable to the intelligence of a human being and could, in any way, exceed the capabilities of human intelligence.

These two perspectives about intelligence are united in their wanting to understand whether it is possible to replicate the intelligence of a human being—or at least to replicate the affective mental processes of animals. The former claims that such an end goal for AI research is possible, even though the computing power does not yet exist. Conversely, the latter is more interested in harnessing the ways in which artificial intelligence can be used. Weak AI acknowledges that the limitations of AI do not necessarily reduce its significance but is cautious about claiming that such findings represent intelligence, since it identifies intelligence as including semantic qualities that cannot be formalized.

Turing's work set about to assert a strong AI thesis, where the machine was able to appear convincingly human. In contrast, Searle argues for weak AI and rejects the possibility that a machine could possess nonalgorithmic aspects of intelligence. For this reason, Searle's position has provoked a broader interest in researching AI that acknowledges that quantifiable measures of intelligence are not sufficient.

Deep Blue

In 1997, IBM's famous machine Deep Blue successfully conquered Grandmaster Garry Kasparov, after a arduous battle. As Moravec (1998) describes, "Kasparov won a long first game against Deep Blue, but lost next day to masterly moves by the machine. Then came three grueling draws, and a final game, in which a visibly shaken and angry Kasparov resigned early, with a weak position. It was the first competition he had ever lost" (1998). The discourse surrounding the capabilities of Deep Blue must thus be seen from these two different perspectives. For proponents of weak AI, Deep Blue reflects only a measure of one aspect of mind functioning, which does not lead to the conclusion that a machine is comparable to a human mind. Indeed, this is the view of its creators, who claim no level of intelligence within the machine (Moravec 1998). For those who consider chess-playing capacity as reflective of some special intellectual capacity, the symbolic significance of Deep Blue beating Garry Kasparov is the rise of machines and their capacity to surpass the capabilities of the self-described superiority of the human species. As Cook and Kroker (1997) identify, the role of chess in culture means that the contest between Kasparov and Deep Blue was not simply a matter of experimentation. Indeed, it defined the struggle between human and automaton, a theme that underpins many dystopian visions of the future and which is framed by a presumption that machines will make people redundant. Thus, chess represents the last stand for humans against machines. Its appeal functions on a number of levels. The claims are impressive and startling, though, I suggest, contingent upon how one understands intelligence and the special cultural value of chess.

I argue on behalf of the latter of these possible claims over artificial intelligence, though the basis for rejecting the possibility of strong AI is based less on the impossibility of a machine's capacity to adequately replicate a human mind. Rather, it involves the inadequacy of the tests that have been used to make such conclusions. In short, on the current, traditional model of AI research, the kind of AI that can be built does not permit the realization of strong AI. In accordance with Bringsjord (1998), I believe that "chess is too easy," so even if we succeed in beating the human, the claim does not allow us to justify claims that this constitutes human intelligence. Admittedly, this seems a particularly unforgiving reaction to the very impressive capabilities of Deep Blue and the research that has led to its success as a chess player. Yet, to infer from this megamachine's success that it is now comparable to a mind is not accurate. However, there is something crucial about the role of chess in society and what expertise in chess conveys philosophically about intelligence that should not be neglected. I will propose that chess is the wrong kind of game and that strong AI seeks a more complex and sophisticated but closed game system.

To elaborate on this, it is necessary to consider further the development of chess-playing machines and how they work. Deep Blue operates in a sequential manner and, when faced with its move in chess, it must analyze every possible move before being able to determine which is the better. This constant "trial and error" strategy contrasts with the human method that appears to involve learning which available moves are worthy of consideration and choosing from competing moves of significant value. Thus, humans are said to possess an intuition that computers can never gain. Such systems as Deep Blue and its offspring represent tentative efforts at creating a much broader, autonomous intelligence that is of interest in Artificial Life (AL) research. Evidence of these crucial differences is found in the stories surrounding the Kasparov versus Deep Blue contests. As Moravec (1998) explains, "The event was notable for many reasons, but one especially is of interest here. Several times during both matches, Kasparov reported signs of mind in the machine. At times in the second tournament, he worried there might be humans behind the scenes, feeding Deep Blue strategic insights!" It is interesting that this event should lead the world's greatest chess player into an accidental Turing test admission. Yet, even if we discount Kasparov's response as a mere psychology of impressions, as opposed to something more profound about what took place, these ways of talking about the computer's actions allude to the requirements of intentionality.

The development of AL (in contrast to AI) offers a perspective on intelligence that can enrich this suggested "presence" of a mind that was discussed in the Deep Blue stories. AL is inspired more by biological traditions than the linearity of Turing's imitation game (Mainzer 1998). It is this approach that speaks to the present critique of AI research. Aspiring to a different model for replicating biological systems avoids the assumed importance of algorithmic knowledge. However, I do not wish to pursue any further the recent achievements of AL. Rather, I want to suggest that one important dimension of AL's foundation must involve the pursuit of game-playing more broadly. Thus, it is important not to neglect what has been achieved in the context of chess, but to further extend what is relevant about this kind of test.

A similar idea is found in Bringsjord (1998), who suggests that a "storytelling" test rather than a Turing test should be used to measure computational intelligence. Bringsjord considers that stories constitute a significant (and even dominant) aspect of human communications. People make sense of the world by creating stories, which entail the creation of characters and identifying with them to assimilate their perspective. While Bringsjord's strategy is novel and plausible, it omits an important aspect of the mind/machine test that is an integral part of chess—that of game playing. The importance of chess being a gamelike activity seems to have been taken for granted. Yet, it does not seem that chess aspires to anything gamelike at all and, thus, its claims to creative, spontaneous, or strategic play do not ring true. Deep Blue cannot be strategic since it does not know it is playing a game. Thus, to extend the relevant aspects of the chess-intelligence test, it is necessary to pursue further the construction of game playing.

Grasshoppers

AI research has neglected to undertake a critical consideration of the meaning of game playing, which is surprising since the complex human interaction that we might typically identify as the richness of intelligence closely relies on playfulness. Thus, AI developers must seek a test that is more broadly reflective of humanness if the aspiration is to claim that machines exhibit humanlike intelligence. The importance of game playing is reflected in Johan Huizinga's classic text *Homo Ludens*, which, coincidentally, was published in 1950, the same year that Alan

Turing delivered his seminal paper *Computing Machinery and Intelligence* in the renowned journal of psychology and philosophy *Mind*. The salience of Huizinga's work is brought alive in Bernard Suits's tale of *The Grasshopper* (1978). Within his manuscript, Suits develops a unitary notion of game playing through the story of the Grasshopper, a character who seeks to spend his entire life playing games convinced that only activity with intrinsic value—games—can be a good life.

Suits suggests that the distinct value of being human involves the propensity to play games and that we find the richest conceptualization of humanness through this form of activity. From this, I argue that game playing embodies qualities that reflect what might be termed *broad intelligence* or *artificial general intelligence* (AGI). Suits's Grasshopper knew that intelligent animals spent as much time as possible playing games, rather than working or engaging with any form of instrumental activity. The characteristics of game-playing behavior thus include such examples as disinterestedness and a feeling of being distinct from ordinary life. Additionally, Suits stipulates:

> To play a game is to attempt to achieve a specific state of affairs [prelusory goal], using only means permitted by rules [lusory means], where the rules prohibit use of more efficient in favor of less efficient means [constitutive rules], and where the rules are accepted just because they make possible such activity [lusory attitude]. I also offer the following simpler and, so to speak, more portable version of the above: playing a game is the voluntary attempt to overcome unnecessary obstacles. (Suits 2005)

From *The Grasshopper* and his more formal models of games, Suits notes that game playing entails formal and informal elements that encompass both aspects of the game's structural limits *and* a game-playing attitude that must be observed by the game player (Suits 1969, 1967, 1977). Yet, within Suits's work, he emphasizes concepts that we would not typically identify as machinic, such as agency. To this extent, one might consider that my claim to pursue Suits's game-playing characteristics as tests for human intelligence are, again, speciesistic. Yet, the kind of agency associated with other types of games is, nevertheless, still a form of closed decision-making. To illustrate, we might consider the game of soccer where a player must undertake a series of closed decisions that can be easily transformed into algorithms. For instance, the set of options available to a player when a goal is scored, or when a whistle is blown. Yet, players also exhibit what appear to be "open" system decisions. For instance, the split-second decision-making that constitutes the way a player kicks the ball—the angle of the foot, the timing, and so on. In sport, it is common to describe such intuitions as talents, which are possible to teach only up to a point.

So, what would it mean if a computer could play soccer effectively? What kind of claim about intelligence could be made? Chess seems to lack a number of the types of the decisions that face a soccer player and, for this reason, I suggest tests for human intelligence must develop into other kinds of games, which have the advantage of being mixed-space—both real and yet closed. Crucially, understanding the sociohistorical development of chess as a test for intelligence informs this proposal. Indeed, it is precisely the presence of a sociocultural context that allows chess to appear as a persuasive test for intelligence. However, it is necessary to develop a test that involves a broader range of decision-making actions. Thus, my thesis is that chess is the wrong kind of game to test for intelligence, rather than a criticism of game playing generally as our form of measurement. Chess is certainly the right kind of test since it is gamelike, but it is not the kind of game that can reveal whether a machine has humanlike intelligence. Other kinds of games, such as those that more closely approximate sportlike games, where creativity and spontaneity adopt a more complex variation, are more useful to study. Within such activities, opportunities for nonlinear decision-making and deviance from preconceived strategic patterns exist. The temporal element of such activities demands of its players a different kind of knowledge than does chess playing.

From this perspective, there is an expectation that machines must appropriate some level of facility for reflexivity that goes beyond following preprogrammed moves available in a chess game. The possibility for deviating from one's strategy or losing concentration in chess is what makes a human player different from a machine. Being able to play a game in the Suitsian sense, thus, would allow the fulfillment of such conditions. In short, one might describe the flaws in Deep Blue's intelligence as an inability to know when to break the rule. Yet, such capacities are not unreasonably demanding in the sense that expecting a computer to "enjoy" a game might be. Instead, a broader range of game-playing characteristics develops the achievements of AI research through such tests as playing chess, without asking too much of a machine's capabilities.

NOTES

1. This proposition does not neglect the playfulness of games like chess, but aims to establish how such an activity is distinguishable from the physical games I propose.

REFERENCES

Bringsjord, Selmer. 1998. "Chess Is Too Easy." *Technology Review* 101, no. 2.

Cook, David, and Arthur Kroker. 1997. "The Digital Prince and the Last Chess Player." CTheory. Event-scene e042, May 20 1997. www.ctheory.net/articles.aspx?id=175.

Haldane, J. 1997. "Could the Soul Be Software?" *Ends and Means: Journal of the University of Aberdeen Centre for Philosophy, Technology, and Society.*

Huizinga, Johan. 1950. *Homo Ludens*. London: Routledge and Kegan Paul.

Lee, Mark H. *Intelligent Robotics.* New York: Milton Keynes / Open University Press, 1989.

Luger, G. F., and W. A. Subblefield. 1998. *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*. Reading, MA: Addison-Wesley Longman.

Mainzer, Klaus. 1998. "Computer Technology and Evolution: From Artificial Intelligence to Artificial Life." *Techne: Society for Philosophy and Technology* 4, no. 1.

Moravec, Hans. "When Will Computer Hardware Match the Human Brain?" *Journal of Evolution and Technology* 1 (1998).

Searle, John R. 1980. "Minds, Brains, and Programs." *Behavioral and Brain Sciences* 3.

———. 1990. "Is the Brain's Mind a Computer Program?" *Scientific American* 262.

Suits, Bernard. 1967. "What Is a Game?" *Philosophy of Science* 34.

———. 1969. "Games and Paradox." *Philosophy of Science* 36, no. 3.

———.1977. "Words on Play." *Journal of the Philosophy of Sport* 4.

———. 2005. *The Grasshopper: Games, Life and Utopia.* Ontario: Broadview Press. Orig. pub. 1978.

Turing, Alan M. 1950. *Programmers' Handbook for the Manchester Electronic Computer*. Manchester, UK: University of Manchester Computing Library.