# Chapter 10: Beyond Games

Please ...

1. Read "Chapter 10: Beyond Games" of Melanie Mitchell's "Artificial Intelligence: A Guide for Thinking Humans" book.

2. With respect to the questions presented for the "Chapter 10: Beyond Games" reading, construct a document containing question/answer pairs, one pair of each question, where the answers, with perhaps just a small number of exceptions, are more or less lifted from Melanie Mitchell's text. Save your document as a **pdf** file.

3. Post your question/answer document to your web worksite.

4. Do your best to internalize your question/answer pairs in some sort of semantic sense, so that the answers are likely to come back to you when prompted by the questions.

5. Come to class for the discussion of "Chapter 10: Beyond Games," when the time rolls around, prepared to participate in the discussion.

6. Please do all of this within one week of the "distribution" of this assignment.

# Chapter 10: Beyond Games - Questions

1. TRUE/FALSE - Over the past decade, reinforcement learning has transformed from a relatively obscure branch of AI to one of the field's most exciting (and heavily funded) approaches.

2. TRUE/FALSE - Developing superhuman game-playing programs is, for most AI researchers, an end in and of itself.

3. MM quotes him as saying: "Games are just our development platform .... It's the fastest way to develop these AI algorithms and test them, but ultimately we want to use them so they apply to real-world problems and have a huge impact on things like healthcare and science. The whole point is that it's general AI–it's learning how to do things based on its own experience and its own data." Who is he?

4. TRUE/FALSE - In stark contrast with humans, most "learning" in current-day AI is not transferable between related tasks. In this regard, the field is still far from what Hassabis calls "general AI." While the topic of transfer learning is one of the most active areas of research for machine-learning practitioners, progress on this front is still nascent.

5. TRUE/FALSE - The results of AlphaGo Zero comprehensively demonstrate that a pure reinforcement learning approach is fully feasible, even in the most challenging of domains: it is possible to train to superhuman level, without human examples or guidance, given no knowledge of the domain beyond basic rules.

6. In just a sentence or two, defend your answer to the previous question.

7. TRUE/FALSE - Many things we humans consider quite easy are extremely challenging for computers. Conversely, many things we humans would find terrifically challenging, computers can do in a split second with a one-line program.

8. TRUE/FALSE - A corollary to the proposition presented in the previous question is that the question of what is considered to be a "most challenging domain" is not universally well-defined.

9. What example did psychologist and AI researcher Gary Marcus propose as a human game that would prove to be significantly more challenging for an AI than the game of Go, because it requires sophisticated visual, linguistic, and social understanding far beyond the abilities of any current AI system.

10. TRUE/FALSE - Deep Q-learning systems have achieved superhuman performance in some narrow domains, but they are lacking something absolutely fundamental to human intelligence. Whether it is called abstraction, domain generalization, or transfer learning, imbuing systems with this ability is still one of AI's most important open problems.

11. Provide two lines of evidence that support the conclusion that deep Q-learning systems (sophisticated reinforcement learning systems), like supervised-learning systems, do not learn humanlike concepts or come to understand their domains in the ways that humans do.

12. TRUE/FALSE - For humans, a crucial part of intelligence is, rather than being able to learn any particular skill, being able to *learn to think* and to then apply our thinking flexibly to whatever situations or challenges we encounter.

13. Other than differences in machine/human architecture or machine/human substrate, what do you think is the most essential difference between how deep Q-learning systems learn to play games like checkers, chess, and Go, and the way that humans learn to play games like checkers, chess, and Go.

14. TRUE/FALSE - Deep reinforcement learning was named one of 2017's "10 Breakthrough Technologies" by MIT's *Technology Review* magazine.

15. What does MM have to say about Demis Hassabis's statement that the ultimate goal of DeepMind's work on reinforcement learning AI system is to "use them so they apply to real-world problems and have a huge impact on things like healthcare and science."?

16. The need for transfer learning is one obstacle to achieving Hassabis's goal of applying the AI methods championed by DeepMind to real world problems. What is another significant impediment?

17. What chore does MM use to illustrate the problems involved in using deep Q-learning to train a robot to do a real world task?

18. What does Andrej Karpathy, Tesla's director of AI, have to say about applying deep Q-learning to real-world tasks?