Chapter 10: Beyond Games

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?

The quoted individual is Demis Hassabis.

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.

The infrastructure of AlphaGo Zero does not solely rely on "pure" reinforcement learning in its infrastructure. Human intervention had to be used to build its infrastructure (convolutional neural network and Monte Carlo tree search) and fine-tune its hyperparameters (167).

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.

Gary Marcus proposed that charades would be far more difficult than Go for AI (168-169).

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.
    a. Mitchell writes, "...one research group showed that it's possible to make specific miniscule changes to the pixels in an Atari game-playing program's input–changes that are imperceptible to humans but that significantly damage the program's ability to play the game," (171).
    b. Mitchell writes, "As far as I know, none of the abilities it has learned are general in any way; none can be transferred to any other task…it would have to start essentially from scratch in learning a new skill," (171).
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.

The ability to generalize is incredibly important in how humans learn to play games. Much of human behavior is informed by past experiences, and the result of this can be good or bad depending on the information at hand. Machines training from scratch lack that intuition that underlies human intelligence.

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."?

Mitchell writes, " I think it's very possible that DeepMind's work on reinforcement learning may eventually have the kinds of impacts Hassabis is aiming for. But there's a long way to go from games to the real world," (172).

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?

Another significant impediment is the difficulty to parameterize the real world into a finite amount of decipherable states (172).

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?

Mitchell illustrates these problems using the chore of doing the dishes (172-173).

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

Andrej Karpathy says, "'basically every single assumption that Go satisfies and that AlphaGo takes advantage of are violated, and any successful approach would look extremely different,'" (173).