

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.

True

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

False

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?

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.

True

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

False

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

The AlphaGo Zero had been built by programmers to specifically play the game Go, "...the architecture of its convolutional neural network, the use of Monte Carlo tree search and the setting of the many hyperparameters that both of these entail." are all built by programmers. In my opinion, this is more like human guided learning, and those programmers should give themselves more credit.

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.

True

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.

True

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.

The game of Charades.

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.

True

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.

“For example, one research group showed that it's possible to make specific minuscule 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.”

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.

True

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.

In my opinion, the AI learns which is the most beneficial state and tries to reach that state, and looks at the entire game board as a whole, each tiny difference is an entirely different state to the AI.

Meanwhile, the human tend to be more focused on individual pieces of a game, like how do I knock out the enemy's king the quickest, and generalize the piece patterns so that even when pieces are in different locations on the board, when they follow the the placement pattern, the human will see them in the same way as before.

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

True

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."?

MM thinks it's very possible that AI will eventually have that kind of impact, but there's a long way to go from games to the real world.

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?

The games have clear rules, straightforward reward functions, relatively few possible actions, and fully visible components at all times. The real world is not so cleanly defined.

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?

Tasking the robot to put dirty dishes into the sink, wash them, and put them into the dishwasher. It's hard to clearly define the state, the robot needs to clearly define the objects that it sees and infers the objects that it doesn't see, It also needs to learn how to perform each action. Most importantly, it would take a lot of dispensable dishes for the robot to reinforce learning.

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

The real-world tasks often violate every assumption and advantage of AlphaGo, any successful approach would look extremely different.