1. Name the two founders of the computer age who wrote programs to play chess before there were even computers that could run their code.

Alan Turing and Claude Shannon

2. He is the “face” of DeepMind Technologies, a company that has also been responsible for significant technical advancements in machine learning, having produced a number of award winning papers. In particular, the company has made significant advances in deep learning and reinforcement learning, and pioneered the field of deep reinforcement learning which combines these two methods. Who is he? ((Thanks to Wikipedia for helping me to frame the question.))

Demis Hassabis

3. The video game Breakout was the result of Atari’s effort to create a single-player version of its successful game Pong. To whom was the design and implementation of Breakout originally assigned in 1975?

Steve Jobs

4. TRUE/FALSE - In 2013, a group of Canadian AI researchers released a software platform called the Arcade Learning Environment that made it easy to test machine-learning systems on forty-nine Atari video games. This was the platform used by the DeepMind group in their work on reinforcement learning.

True

5. TRUE/FALSE - The DeepMind group combined reinforcement learning–in particular Q-learning– with deep neural networks to create a system that could learn to play Atari video games. The group called their approach deep Q-learning.

True

6. What game did Melanie Mitchell use as a running example to explain how deep Q-learning works?

Breakout

7. TRUE/FALSE - DeepMind first presented its work on computational methods for learning to play Atari video games in 2013 at an international machine-learning conference. The audience was dazzled. Less than a year later, Google announced that it was acquiring DeepMind for about 650 million dollars, presumably as a result of its work on the machine learning of Atari video games.

True

8. TRUE/FALSE - With a lot of money in their pockets and the resources of Google behind them, DeepMind–now called Google DeepMind– to on a bigger challenge, one that had in fact long been considered one of AI's "grand challenges": creating a program that learns to play the game Go better than any human.

True

9. In the very early years of the computer age, Arthur Samuel joined IBM's laboratory in Poughkeepsie, New York and immediately set about programming an early version of IBM's 701 computer to play checkers. In what language did he do this early work?

He had to write everything using the op codes and addresses.

10. TRUE/FALSE - Samuel's checkers-playing program was among the earliest machine-learning programs; indeed, it was Samuel who coined the term machine learning.

True

11. Samuel's checkers player was based on a method of searching a game tree, which is essentially the same method used by Deep Blue to play chess and AlphaGo to play Go. What is the name of this method of searching a game tree?

Monte Carlo Tree Search

12. TRUE/FALSE - What Samuel's program learned was which features of the board should be included in the static evaluation function at a given turn, as well as how to weight these different features when summing their points.

True

13. TRUE/FALSE - In the most interesting version of Samuel's program, the system learned while playing itself!

True

14. What are the three main concepts that Samuel's checkers player illustrates?

The game tree, the evaluation function and learning by self-play

15. TRUE/FALSE - In 1958, Alan Newell and Herb Simon wrote, “If one could devise a successful chess machine, one would seem to have penetrated to the core of human intellectual endeavor.”

True

16. TRUE/FALSE - In 1997, IBM had its second big game-playing triumph with Deep Blue, a chess-playing program that beat the world champion Garry Kasparov in a widely broadcast multi-game match.

True

17. Other than the difference of game specific knowledge (checkers vs chess), what were the major differences between Samuel’s chess machine and the Deep Blue chess machine?

Deep Blue had a deeper look-ahead in its game-tree, it had a more complex evaluation function, hand programmed chess knowledge and specialized parallel hardware to make it run very fast.

18. After Deep Blue defeated Kasparov, how did most people in AI view Newell and Simon’s famous 1958 statement about a successful chess-playing machine?

19. TRUE/FALSE - According to the great Go champion Lee Sedol, Go is incomparably more subtle and intellectual than chess.

True

20. TRUE/FALSE - Go is a game that has fairly simple rules but produces what you might call emergent complexity. At each turn, a player places a piece of his or her color (black or white) on a 19-by-19 square board, following rules for where pieces may be placed and how to capture one’s opponent’s pieces. Unlike chess, with its hierarchy of pawns, bishops, queens, and so on, pieces in Go (“stones”) are all equal. It’s the configuration of stones on the board that a player must quickly analyze to decide on a move.

True

21. TRUE/FALSE - In 2016, DeepMind’s AlphaGo program spectacularly defeated Lee Sedol in a very high profile 5 game match.

True

22. TRUE/FALSE - Demis Hassabis, following the AlphaGo-Sedol match, noted that “the thing that separates out top Go players is their intuition,” and suggested that “what we’ve done with AlphaGo is to introduce with neural networks this aspect of intuition.”

True

23. DeepMind built several different versions of AlphaGo, so to keep them straight, DeepMind started naming them after the human Go champions the programs had defeated–AlphaGo Fan and AlphaGo Lee. What image did this naming convention evoke in MM’s mind?

The image of the skulls of the vanquished enemies in the collection of a digital viking

24. A year after the Lee Sedol match, DeepMind developed a version of the Go machine, called AlphaGo Zero, that was both simpler than and superior to the previous versions. Why was the program given this name?

It started off with zero knowledge of Go besides the rules.

25. TRUE/FALSE - The word intuition has an aura of mystery, but AlphaGo’s “intuition” arises from its combination of deep Q-learning with a clever method called Monte Carlo tree search.

True

26. Describe “Monte Carlo tree search” in just one paragraph.

An assigned score is given to geach possible move from a given board position. A minuscule amount of the possible sequences of moves is looked at. The program counts how many wins and losses these sequences lead to and uses those counts to give a score to each of its moves. At first the choices are random but as time goes on the programs chooses moves that have led to wins.

27. TRUE/FALSE - The scientists at DeepMind invented Monte Carlo tree search.

False

28. TRUE/FALSE - Imagine that you are AlphaGo staring at a board position: before you start the Monte Carlo process of performing roll-outs from that position, The ConvNet is whispering in your ear which of the possible moves from your current position are probably the best ones.

True

29. TRUE/FALSE - Imagine yourself as AlphaGo after a Monte Carlo tree search. The results of your search are new probabilities assigned to all your possible moves, based on how many times those moves led to wins or losses during the roll-outs you performed. These new probabilities are now used to correct your ConvNet's output, via back-propagation.

True

30. TRUE/FALSE - Once AlphaGo "matures," its ConvNet, informed by a huge number of Monte Carlo tree search applications over a huge number of plays of the game against itself, will play the role of the program's "intuition."

True

31. TRUE/FALSE - With its AlphaGo project, DeepMind demonstrated that one of AI's longtime grand challenges could be conquered by an inventive combination of reinforcement learning, convolutional neural networks, and Monte Carlo tree search (and adding powerful modern computing hardware to the mix).

True