

## The Final Big Class Discussion on whether advanced robotics will end life as we know it

Wednesday, Dec. 13, 10:30 a.m.-12:20 p.m.

Our final exam will be a long discussion, and you will be graded on your participation — not just that words come from your mouth, but that your words are true, relevant, and helpful in furthering the discussion.

We will use a talking ball. If you want to say something, raise your hand, and whoever has the ball, when they are done talking, will throw you the ball. (Or maybe just toss the ball; make sure the other person can catch the ball, OK?)

This is not a debate, but a discussion. There are no teams, and no winners or losers. There's just sharing, listening, thinking, evaluating, possibly changing your mind, or at least shifting how you understand something.

### On grading

At the very **minimum**, you must make this many comments to receive the following grades:

**10 (A), 8 (A-), 7 (B+), 6 (B), 5 (B-), 4 (C+), 3 (C), 2 (C-), 1 (D+).**

To receive a point, each comment must be true, relevant, and it must contribute to the discussion.



## Study Guide for FYS (Minds, Robots, and the End of Humanity) final exam

Here is a list of study questions for the final exam discussion. You should gather your reading materials and notes, and write down relevant facts and arguments relevant to each of the following discussion questions. **Feel free to make use of your class notes and readings during the discussion.** We will begin with topics that you are most interested in discussing, and go from there.

- (1) What are strong and weak AI, and how do they fit in with ANI, AGI, and ASI? What are examples of each of these?
- (2) What are the main obstacles to creating AGI and ASI? What are the implications of these developments?
- (3) Should we be worried about developing AGI and ASI?
- (4) What is “the technological singularity” mentioned by Vinge (and many others since him)?
- (5) What is the difference between exponential and linear growth. What sort of growth do we find in computer technology and what are the implications of this growth?
- (6) What are some measures that humans could take to prevent AI from self-improving itself into a threat to human existence?
- (7) What is a Turing machine and what relevance is it to the topics discussed in this class?
- (8) What is at stake in Daniel Dennett’s thought experiment (where his brain is removed from his body, with the neural connections maintained by radio transmitters/receivers)? Where is Dennett, and why is he there?
- (9) What would prevent your mind from being uploaded into a computer? If such an upload could be done, would there now be two of you? If your body and brain were then destroyed, would you still exist?
- (10) Can machines become conscious? (Are there conscious machines now?) Why or why not? How do we know or determine whether something is conscious?
- (11) What is the difference between a physicalist and a dualist understanding of the mind?
- (12) What is Cartesian dualism?
- (13) What was the point of the defecating duck?
- (14) What is the difference between a person and a human being? Is this difference important?
- (15) What is panpsychism? Why does Koch relate it to “integrated information theory” (IIT)?
- (16) Should we encourage the development and use of social robots for taking care of the sick and elderly? What are the advantages of using so-called “emotional robots” like Pepper?
- (17) To what extent, if any should autonomous machines be used in warfare and/or policing?
- (18) Can we program morality into these intelligent machines?
- (19) Who were the Luddites? What are the implications of the growth of AI on the economy? On our jobs?
- (20) At what point in the development of AI will a machine with AI develop moral rights (or protections) against human beings?