

COMP4920 - First Essay Questions

Term 1, 2025 - UNSW

Word length - 800w-1000w in total (not including bibliography or footnotes).

Due date - Week 3 - Friday March 7, 17:00.

Value - 20% of final mark.

Please submit your essay via Moodle/Turnitin.

Answer (1) of the following questions

Do not answer more than one (1) of the following questions:

1. Assume that computer science is being established as a licensed profession (along the lines of law and medicine). Assume also that you have been tasked with the formulation of the ethical guidelines for the governing association of this new profession. How might the use of act versus rule utilitarianism to motivate the ethical guidelines for this professional association of computer science result in different professional guidelines? Which of act and rule utilitarianism do you prefer for this purpose? Why? Justify - that is give reasons for - your answer. Incorporate at least one leadership challenge and explain how it is that your response to this challenge turns on your choice of rule or act utilitarianism for the guidelines that govern this response.
2. Expound and assess rule-based/Kantian ethics. Analyse the extent to which such an ethics might be used to design an automated ethics as per the readings in section 1.2 below. What do you think that the risks and opportunities of such an automated ethics might be? Why? Justify your answer with explicit, detailed, expository reference to at least one of the suggested readings in section 1.2 below.
3. Is kicking a robot dog morally wrong, or morally permissible? In your answer, make explicit detailed reference to virtue ethics. Which answer or answers might virtue ethics give us? Are any of these answers correct? Why? Justify your answer with explicit, detailed, expository reference to the debate between Coeckelbergh and Sparrow in the readings noted in section 1.3 below.

This is the end of your questions. There are no more questions.

Please read on for essay tips, referencing guidelines, UNSW's plagiarism protocols, an important rule with regard to Generative AI, and details of Competency Grading.

1 Essay writing tips

1.1 For Question 1 above

This is a fun question. Question 1 is applied. It really is a question - well actually it is a bunch of questions - and this means that your essay should be an answer to it. Then, and this is super important, you need to give reasons to your reader for why it is that you think your answer(s) is the correct one. Your goal is to convince your reader that you are right! So, let's move through Question 1 carefully, one sub-question at a time. The first thing that Question 1 asks you to do is to make a pair of assumptions - so you do not need to justify or give reasons for these. Firstly that computer science is about to become a licensed profession. This means that one could become barred from practicing computer science legally. Secondly that your job is to draw up guidelines for the association that governs this profession. The second thing that Question 1 asks you to do is to explain how you think that act and rule utilitarianism might cause one to design different guidelines for such an association, should either of these utilitarianisms be used to constrain the design of such guidelines. So, this part of the question is of a "compare and contrast" nature. The third thing that Question 1 asks you is this - do you prefer either act or rule utilitarianism for the purpose above? A sensible start here would be to explain both act and rule utilitarianism. It is open for you to say that you think that both act and rule utilitarianism are so terrible for the task that you dislike each of them equally. The important thing is that you state why it is that you think this, and give detailed reasons for your belief. Similarly, it is open for you to say that, on balance, you think that both act and rule utilitarianism are good, and equally good at the task. Here too, the important thing is that you provide em reasons. Note again again - a good essay will engage with contrary points of view, and provide responses to them.

1.2 For Question 2 above

The first thing that Question 2 is asking you to do is to expound Kantian ethics. "Expound" is just a fancy way of asking you describe and explain something in detail.

The second thing that Question 2 is asking you to do is to assess Kantian ethics. So, this is a great opportunity for you to state what it is that you think about Kantian ethics, why you think it, and what reasons you have for your reader to agree with you that you are right! As always, a very useful thing to do here is to anticipate and articulate possible objections to your point of view, and to then respond to them in detail.

The third thing that Question 2 ask you is just how useful you think a Kantian ethics might be for automated ethics. Another name for automated ethics is “machine ethics”. Automated ethics is the automation of ethical/moral decision making. In other words, it is the attempt to make moral decision making computable. This is just to say that it is that attempt to subsume moral decision making within a finite mechanical process.

The fourth and fifth things that Question 2 asks you respectively are to give reasons for your answer to the third thing, and to then argue for and justify these reasons. I am sure that you can see a pattern here by now :)

Again as always, a very useful thing to do here is to anticipate and articulate possible objections to your point of view, and to then respond to them in detail. Here are some readings that will help you with your answer to Question 2. Note that all of them can be found on Google Scholar (you will need to be inside the UNSW network to access some of them, so VPN is essential if you are not on campus):

Manna, R. and Nath, R. (2021): Kantian Moral Agency and the Ethics of Artificial Intelligence, *Problemos* vol. 100, pp. 139–151.

Nath, R., and Sahu, V. (2021): The problem of machine ethics in artificial intelligence, *AI & Society*, 35:103–111.

Tonkens, R. (2009): A Challenge for Machine Ethics, *Minds & Machines* (2009) 19:421–438.

Singh, L. (2022): Automated Kantian Ethics: A Faithful Implementation. Online at <https://github.com/lsingh123/automatedkantianethics>

1.3 For Question 3 above

The first thing that Question 3 asks you is what you think about the moral status is of kicking a robot dog. A good answer here would give some space to the considerations that motivate your answer.

The second thing that Question 3 ask you is to engage in detail with virtue ethics! This is fun, as there is a lot of great literature out there on this very same topic.

Think about it now for just a moment. Neither utilitarianism nor Kantian ethics give as any obvious reason for why it is that kicking a robot dog is the wrong thing to do. In spite of this, some people have felt that there is /something/ wrong with kicking a robot dog. There is some great literature on this topic. The videos of the Boston Dynamics robot dog “Spot” being kicked led to a very interesting debate, and virtue ethics played a central role in it.

Again as always, a very useful thing to do here is to anticipate and articulate possible objections to your point of view, and to then respond to them in detail. The other things that Question 3 asks you require you to engage with this literature. Here is the literature. Again, note that all of them can be found on Google Scholar (you will need to be inside the UNSW network to access some of them, so VPN is essential if you are not on campus):

Coeckelbergh, M. (2021): How to Use Virtue Ethics for Thinking About the Moral Standing of Social Robots: A Relational Interpretation in Terms of

Practices, Habits, and Performance, *International Journal of Social Robotics*, 13:31–40.

Sparrow, S. (2021): Virtue and Vice in Our Relationships with Robots: Is There an Asymmetry and How Might it be Explained? *International Journal of Social Robotics* (2021) 13:23–29.

Coeckelbergh, M. (2021): Does kindness towards robots lead to virtue? A reply to Sparrow’s asymmetry argument, *Ethics and Information Technology* (2021) 23:649–656.

Coeckelbergh, M. (2021): Should We Treat Teddy Bear 2.0 as a Kantian Dog? Four Arguments for the Indirect Moral Standing of Personal Social Robots, with Implications for Thinking About Animals and Humans, *Minds and Machines* (2021) 31:337–360.

2 Plagiarism rules

It is very, very important that you do not fall afoul of UNSW’s plagiarism rules. You need to read through and understand all of the following:

<https://www.student.unsw.edu.au/plagiarism>

I cannot over-emphasise just how serious this is.

3 Generative AI Policy

I have read through the recent document from UNSW’s Pro-Vice Chancellor on the issue of Generative AI. I find that this assessment in SENG/COMP4920 falls under the following rules for simple editing assistance: For this assessment task, you may use standard editing and referencing software, but not Generative AI. You are permitted to use the full capabilities of the standard software to answer the question (e.g. you may wish to specify particular software such as Microsoft Office suite, Grammarly, etc.).

If the use of generative AI such as ChatGPT is detected, it will be regarded as serious academic misconduct and subject to the standard penalties, which may include 00FL, suspension and exclusion. Here too, I really cannot over-emphasise just how serious this is.

4 Late Penalties

The School of Computer Science and Engineering at UNSW focuses on developing and promoting professional standards of practice and a professional approach to workloads and deadlines. As such if you are not able to submit your assessment by the due date and time above, you should apply for Special Consideration if you are eligible, and/or an adjustment to your ELP if you have one. In all other cases, **late submissions will not be accepted for marking**, and a grade of CN will be awarded for the assessment task.

5 Referencing Guidelines

All essays submitted for the First Essay and the Group Project Report **must** follow the referencing guidelines specified here:

<https://www.student.unsw.edu.au/harvard-referencing>

6 Marking Guidelines - Competency Grading

Hi everyone! Thank you in advance for all of your essays - I know that they will be fantastic. Below are the marking guidelines for **Competency Grading**. Please note that the essay questions denote open problems. This means that there will be many ways to satisfy the guidelines below. An open problem is one for which there are many good solutions. They are different in kind to closed problems - for which the correct solution is very tightly constrained, and for which the search space for solutions is very limited. Answers to open problems cannot be auto-marked. By contrast, the search space for open problems is big. Really big. For open problems, you cannot just follow a set of codified instructions or a flow chart or a check-list, and arrive at a good solution/answer. There is further general advice at the bottom of this document.

(CN) - **Non-Competent**. A grade of CN demonstrates a failure to grasp the elementary issues relevant to a competent answer. Reliable ways to receive a grade of CN are to not answer the question at all, to get even basic definitions very wrong, or for the essay to not actually make any sense. Other ways include failing to give sensible reasons for your claims, failing to have a logical structure, and so on.

(CO) - **Competent**. In order to achieve a grade of CO, you need to answer the question(s) and give sensible reasons for your answers. You must anticipate objections to your reasons, articulate them, and articulate your responses to these objections. You must demonstrate a proper grasp of the course material. In your essay, you must engage with the relevant course material at the relevant points in such a way for it to support the points that you were making.

To go even further, you might reconstruct complex points and arguments from the literature in your own words and with considerable clarity.

(CM) - **Competent with Merit**. All of the above qualities for a grade of CO, plus your essay contains convincing arguments not anticipated easily by a marker, and makes genuine, novel contributions to the issues. This might be achieved by having demonstrated an extraordinarily subtle analysis of the relevant issues, or by having read and understood a wide range of relevant literature that you discovered as a function of your own research, or by a combination of the two.

7 General advice:

If you have not started on your first draft already, then get going! Then revise the draft as many times as you can before the deadline. Essays left to the last minute are usually obvious (and not very good).

You may write in the first person if you wish.

You are very welcome to pursue further research into the literature on relevant areas. Please do be sure to use peer-reviewed academic sources (and not random magazine articles or videos etc.).

Give reasons for the claims that you make, and back up claims about other people's claims by referring to the relevant literature and citing it properly.

Write your introduction at the very end. By this I mean write the rest of your essay first, and then go back and write the introduction.

Read your drafts out loud to yourself. Trust me, you will spot a bunch of opportunities for improvement that you will not spot otherwise.

A short sentence is a good sentence. English does not like long sentences. A long sentence in English is a bad sentence.

Remember that your essay is a piece of academic writing.

END OF DOCUMENT