



Course Proposal Form

Please see Page 2 for instructions on which parts of this form to complete, whom to consult with to avoid unnecessary effort, and where to send the completed form.

Proposer(s): Nadin Kokciyan

Date: 11 November 2020

Cover page: Basic permanent course information

Unless otherwise noted, items in this section are entered into EUCLID and **cannot** be changed without creating an entirely new course.

Course Name	Case Studies in AI Ethics
Course Acronym <i>(used by the School only, e.g., for the Sortable Course List)</i>	CSAI
Course Level If the course is only available to MSc students, then it must be classed as Postgraduate. All other courses, regardless of level, are Undergraduate.	<input type="checkbox"/> Undergraduate <input checked="" type="checkbox"/> Postgraduate
Normal Year Taken	<input type="checkbox"/> UG1 <input type="checkbox"/> UG2 <input type="checkbox"/> UG3 <input type="checkbox"/> UG4 <input type="checkbox"/> UG5 <input checked="" type="checkbox"/> MSc
Also available in years <i>[This can be changed later if need be.]</i>	<input type="checkbox"/> UG1 <input type="checkbox"/> UG2 <input type="checkbox"/> UG3 <input checked="" type="checkbox"/> UG4 <input checked="" type="checkbox"/> UG5 <input type="checkbox"/> MSc
SCQF Credit Level Level 8 should normally be used for pre-honours courses. Level 10 should normally be used for optional UG3 courses (so UG4 students may also take them) and for courses aimed mainly at UG4 students. Level 11 should be used for courses aimed mainly at MSc students, whether or not UG4 students can also take them.	<input type="checkbox"/> 7 <input type="checkbox"/> 8 <input type="checkbox"/> 9 <input type="checkbox"/> 10 <input checked="" type="checkbox"/> 11
SCQF Credit Points	<input checked="" type="checkbox"/> 10 <input type="checkbox"/> 20 <input type="checkbox"/> 40 <input type="checkbox"/> 60 <input type="checkbox"/> 80 <input type="checkbox"/> Other:
Delivery Location	<input checked="" type="checkbox"/> Campus <input type="checkbox"/> On-line Distance Learning
Course Type	<input checked="" type="checkbox"/> Standard (default) <input type="checkbox"/> Dissertation <input type="checkbox"/> Online Distance Learning <input type="checkbox"/> Other (specify: Placement, Student Led Individually Created Course, Year Abroad)
Marking Scheme By default, courses use a numerical marking scheme. If you wish to use a grade-only marking scheme, your course proposal below should justify this.	<input checked="" type="checkbox"/> Standard (numerical) <input type="checkbox"/> Letter grade only

Guidance for remaining sections:

For an initial course proposal, please complete the **cover page and Section 1 (Case for Support)**, which asks you to describe the need for this course and to provide an overview of the course design, including the learning outcomes. **Please discuss your plans as early as possible with the head of Curriculum Review to avoid unnecessary effort.**

Send the form with these sections completed to the BoS Academic Secretary and head of Curriculum Review (listed on the BoS page) to obtain their comments before filling out the remainder of the form.

If a full proposal is invited, please complete the remaining sections and send to iss-bos@inf.ed.ac.uk.

2. Student-facing course description and additional feedback and assessment information.

This section provides most of the information students see in the DRPS entry for this course, as well as related details for BoS consideration.

3. Further information for BoS consideration: sample materials.

4. Additional Course Details required for DRPS. *[Administrative information such as delivery timing and prerequisites.]*

5. Placement in degree programme tables. *[Required for all level 9-11 courses; used to determine where the course will be added to existing degree programme tables.]*

6. Comments from colleagues. *[All course proposal should be sent to relevant colleagues in the area as well as to the appropriate year organizer and BoS Academic Secretary for comment in good time before the BoS meeting. Use this section to indicate what feedback has been solicited and received.]*

Colour coding and item-by-item guidance:

Guidance is provided in italics for each item. Please also refer to the guidance for new course proposals at <http://www.inf.ed.ac.uk/student-services/committees/board-of-studies/course-proposal-guidelines>. Examples of previous course proposal submissions are available on the past meetings page <http://web.inf.ed.ac.uk/infweb/admin/committees/bos/meetings-directory> but note that the proposal form was updated in Jan 2019.

Sections in gold are for student view and are required before a course can be entered into DRPS. You must complete these sections even if your course has already been approved based on other documentation.
Sections in orange are for School use but are still required for all courses (even those that have already been approved based on other documentation).
Section in gray are for consideration by the Board of Studies. They are normally required for all new course proposals but may be omitted in some circumstances (e.g., for invited course proposals) if you obtain permission in advance.

1. Case for support

This section is for consideration by the Board of Studies. The final two boxes (Learning Outcomes, Graduate Attributes) will also go into the student-facing course description.

Overall contribution to teaching portfolio and relation to existing curriculum

Please explain (a) what motivates the course proposal (e.g. a previous course having become outdated/inappropriate, an emergent or maturing research area or new research activity in the School, offerings of our competitors) and (b) how it relates to existing courses and degree programmes (including any prerequisite courses). Every new course should make an important contribution to the delivery of our [Degree Programmes](#).

Artificial intelligence (AI) technology is being deployed in real-world settings more than before. In recent years, social and technical issues arise exponentially in deployed AI systems. As the School of Informatics, we should increase awareness around these topics by including such a course within our existing curriculum. The aim will be to emphasize the importance of using AI for Social Good.

In our current offerings, we have a gap in discussing social and technical ethical issues that we are currently facing in AI technology. Students should have a better understanding of this topic, so that they can consider such issues while designing or developing AI technology. Such a course would be interesting for most of our Informatics students. It would be beneficial to have such a standalone course complementing add-on lectures within some Informatics courses we are currently offering.

Earlier this year, I sent an email to the teaching-staff to learn more about how we cover *data ethics* within our courses. I have shared an Excel spreadsheet with all the teaching-staff to share the responses that I have received [note that the list of the courses is not complete]. My findings can be summarized as follows: 1) NLP-based courses cover ethics more than other courses (e.g. ANLP, NLU+). 2) We do not talk about ethical issues much; in some courses there are dedicated lectures for one/two weeks. 3) There is a growing interest in adding more lectures about ethics, some lecturers contacted me regarding this.

The link to the Excel spreadsheet is here: https://uoemy.sharepoint.com/:x/g/person/nkokciya_ed_ac_uk/EdKOEwYCNQRBPcv3r3gzZUsBySWitrEBoZ1ui7YZNsGfhg?e=BMoBSA

I was also part of some of the discussions about changing the course structure in Professional Issues. But such changes will affect the UG curriculum mostly.

Target audience and expected demand

Describe the type of student the course would appeal to in terms of background, level of ability, and interests, and the expected class size for the course based on anticipated demand. A good justification would include some evidence, e.g. by referring to projects in an area, class sizes in similar courses, employer demand for the skills taught in the course, etc

I would expect to see 50-80 students interested in this course; since students will not require meeting any pre-requisites. The course is aimed at MSc students on all Informatics degrees. This may even become a compulsory course in some programmes.

The course would be of interest for students in most of our MSc programmes (e.g. Informatics, AI, Cybersecurity, Privacy and Trust, Data Science); CDT programmes (e.g. Biomedical AI, NLP, Data Science).

There is a course running by PPLS that is called *Ethics of Artificial Intelligence (PHIL10167)*. This course investigates the ethics questions from a philosophical perspective. The average number of students taking this class is around 30-35 students since there is a quota specified. I had a conversation with Mark Sprevak who is the course organiser for several years now. He also mentioned that there is much interest from Informatics students as well.

Anticipated Resource Requirements

Estimate how much lecturing, tutoring, exam preparation and marking effort will be needed in steady state, and any additional resources needed to set the course up initially. Provide estimates relative to class size where applicable and discuss how support staff will be recruited and supervised, if the class is likely to be very large. Please mention any scaling limits due to equipment or space. If equipment is required, say how it will be procured and maintained.]

Written exam: 80%

Coursework: 20%

Tutorials: 2

Each tutorial group is expected to have 10-15 students; a tutor will be required for each group to guide the students. There will be also a need for markers for the coursework and the written exam (depending on the size of the class). The coursework will have a limit in length (1500-2000 words).

Since this is a new course, I will need more TA support (~50 hours) to help me in developing course materials (e.g. case studies) and preparing the coursework. The TA and me will checking the coursework outlines to provide formative feedback to the students.

Quotas, special arrangements or unusual characteristics

Please specify if this course requires any special arrangements such as quotas or other registration arrangements; is a collaboration with another school or institution, or has other atypical characteristics that may affect finances or student registration. Further justification/information may be requested for such courses.

None

Narrative description of the course aims and structure

Please describe the main goals of the course and how the course design will allow students to achieve those goals. This section should be consistent with the student-facing information provided below, but should provide additional information to help colleagues at BoS understand the vision and structure of the course. This description may refer to the learning outcomes and graduate attributes (next two boxes) and should explain how activities such as tutorials, labs, or in-lecture activities will support them, and how the proposed assessments will assess them.

For courses that are important pre-requisites for other courses, this section may also provide content/syllabus information which is too detailed for the student-facing description, such as a lecture-by-lecture syllabus.

Artificial intelligence (AI) technology is being deployed in real-world settings more than before. In recent years, people are more concerned using AI technology as there are lots of ethical challenges to consider. Especially, fully automated AI systems started to make critical decisions such as who should

be employed or who is a criminal. This calls for a new course that could potentially increase awareness around this topic. The goal of the course is to introduce basic concepts around ethical issues (e.g. bias, fairness, privacy) and bring together different case studies from various contexts (e.g. design and programming, law, human rights).

Topics to be covered include (but are not limited to):

- Data Ethics
 - Deployed AI technologies
 - Ethical and social issues arising with data
 - Case Studies
- Fairness, Accountability and Transparency
 - Overview of the definitions
 - Types of bias
 - Explainability
 - Case Studies
- Privacy
 - Arising issues (e.g. surveillance, usability vs privacy trade-off)
 - State of the art: ML approaches, Agent-based approaches
 - Case Studies
- Towards implementing ethical tools
 - Implementing AI Ethics (e.g. agent-based perspectives, AI approaches)
 - Ethics guidelines for Trustworthy AI (e.g. European Commission), AI Auditing guidelines (e.g. ICO)
 - Applied Ethics (e.g. IEEE Ethics in Action, Markkula Centre's Ethics Toolkit)
 - Case Studies

Each week, there will be 2 lectures. Some lectures will include small case-studies where the students will be asked to discuss the ethical and social issues arising in the case study in small-groups. The TA and me will contact some of the course organizers to discuss possible case studies that we can collect internally. I will use online resources (e.g. Embedded EthiCS course running at Harvard University) or published papers to develop new case studies. Some case studies will be used as part of the lectures whereas some will be discussed in detailed during tutorials. I provide some example case studies in Section 3.

Summary of Intended Learning Outcomes (MAXIMUM OF 5)

*List the learning outcomes of the course. These must be assessable (i.e., observable), so must specify what the student should be able to do concretely, not simply what they should "understand". Use concrete verbs that indicate (a) what type of assessment would be appropriate, and (b) what level of knowledge/thinking is expected (from recall to analysis to novel creation). **Example verbs:** define, explain, implement, compare, justify. Assessments (described later) should be tied to the learning outcomes.*

Outcomes should typically focus more on the types of thinking/skills developed than on the detailed course content, and the level of thinking should be appropriate to the level of the course: outcomes for a Level 11 course should include more higher-level thinking skills than for a Level 8 course. Further guidance on writing learning outcomes can be found at <https://www.ncl.ac.uk/ltds/assets/documents/res-writinglearningoutcomes.pdf>

On completion of this course, the student will be able to

- 1) Understand data ethics and arising issues (e.g. bias, fairness, privacy) in AI systems
- 2) Explain and provide examples of how AI systems can play a critical role in decision-making
- 3) Analyse case studies to identify and mitigate potential risks considering legal, social, ethical or professional issues
- 4) Apply ethical methodologies in the design of responsible AI systems

Graduate Attributes, Personal & Professional Skills

List the personal attributes and generic transferrable skills this course will help develop. Examples include

Cognitive skills: *problem-solving, critical/analytical thinking, handling ambiguity*

Responsibility, autonomy, effectiveness: *independent learning, self-awareness and reflection, creativity, decision-making, leadership, organization and time management, flexibility and change management, ethical/social/professional awareness and responsibility, entrepreneurship*

Communication: *interpersonal/teamwork skills, verbal and/or written communication, cross-cultural or cross-disciplinary communication*

Cognitive skills: problem-solving (via tutorials, coursework), critical thinking (via lectures/tutorials/coursework), handling ambiguity (via in-class discussions)

Responsibility, autonomy, effectiveness: independent learning (via readings, videos), self-awareness and reflection (via tutorials, coursework, lectures), leadership (via discussions about case-studies in small groups), time management (via coursework, discussions during classes), ethical/social/professional awareness and responsibility (via tutorials, lectures, coursework, readings, videos)

Communication: written communication (via coursework), verbal communication (via in class-discussions)

2. Student-facing course description and additional feedback and assessment information

Except where noted, all fields are required and will go into the DRPS entry for the course (for use by students). **Important:** any text in DRPS is effectively a contract with students, so should not include details that are likely to change from year to year.

<p>Summary Description <i>Provide a brief official description of the course, around 100 words. This should be worded in a student-friendly way, it is the part of the descriptor a student is most likely to read. If this course replaces another course, please say so in this summary.</i></p>	<p>Artificial intelligence (AI) is being deployed in real-world settings more than before. Especially, fully automated AI systems started to make critical decisions such as who should be employed or who is a criminal. In this course, the students will increase their understanding of data ethics. The course gives an overview of the ethical issues (e.g. bias, fairness, privacy) and bring together different case studies from various contexts. The students will analyse case studies to identify and mitigate potential risks considering legal, social, ethical or professional issues.</p>
<p>Keywords <i>Give a list of searchable keywords.</i></p>	<p>Artificial intelligence, ethics, machine learning, data science</p>
<p>Course Description <i>A more detailed student-facing description of the course, which should normally include (a) a more in-depth academic description of the learning aims, nature and context of the course, (b) a rough outline of the content or syllabus, often as bullet points, and (c) a description of how the course will be taught, how students are expected to engage with their learning and how they will be expected to evidence and demonstrate their achievement of the intended learning outcomes.]</i></p>	<p>In this course, we will discuss the following topics:</p> <ul style="list-style-type: none"> • Data Ethics <ul style="list-style-type: none"> ○ Deployed AI technologies ○ Ethical and social issues arising with data • Fairness, Accountability and Transparency <ul style="list-style-type: none"> ○ Overview of the definitions ○ Types of bias ○ Explainability • Privacy <ul style="list-style-type: none"> ○ Arising issues (e.g. surveillance, usability vs privacy trade-off) ○ State of the art: ML approaches, Agent-based approaches • Towards implementing ethical tools <ul style="list-style-type: none"> ○ Implementing AI Ethics ○ Ethics guidelines for Trustworthy AI (e.g. European Commission), AI Auditing guidelines (e.g. ICO) ○ Applied Ethics (e.g. IEEE Ethics in Action, Markkula Centre’s Ethics Toolkit) <p>The students will be expected to prepare for the lectures by reading papers, news; or watching videos. Some lectures will include case studies where students will discuss the ethical issues in small discussion groups for 15 minutes; and report back their findings.</p>

<p>Assessment Weightings:</p> <p><i>These should correspond approximately to the proportion of learning outcomes that each component assesses. More than 30% coursework requires specific justification.</i></p> <p><i>The expectation for a 10pt course is 20% coursework with the equivalent of one 15-20hr assessed assignment (but possibly split into smaller pieces). See 'components of assessment' below.</i></p>	<p>Written Exam 80%</p> <p>Practical Exam 0% (for courses with programming exams)</p> <p>Coursework 20%</p>
<p>Further Assessment Information</p> <p><i>Provide any further information that should go on DRPS for students. E.g., if the assessment includes required group work or if students must pass some individual component of assessment as well as the course overall.</i></p>	-
<p>Components of assessment and time spent on assignments (for BoS only)</p> <p><i>If not already included in the course narrative description, please describe the type of assessments (oral presentation, report, programming, etc) and how each component of assessment will assess the intended learning outcomes. Where coursework involves group work, it is important to remember that every student has to be assessed individually for their contribution to any jointly produced piece of work.</i></p> <p><i>Also estimate how many hours students will spend on assignments. Please see the School policy on Workload and Assessment, which states that students should not be expected to spend more than 6-7 hrs/wk per 10 credits, including contact hours.</i></p> <p><i>Note that it often desirable to include formative assignments which are not formally assessed but submitted for feedback, often in combination with peer assessment.</i></p>	<p>There will be one piece of individual coursework. It will involve the student working on a specific case study where the student will be expected to apply an ethical toolkit to solve the case study in hand. I anticipate that it will take around 20 hours. The length limit will be around 1500-2000 words.</p> <p>I anticipate 2 lectures per week, and 2 tutorials in the semester. The student will need to spend about 30 minutes-2 hours preparing for lectures each week. The student will read some materials or watch videos prior to lectures.</p> <p>The exam will assess the learning outcomes (LO1-LO3). The individual coursework will assess LO3 and LO4 specifically.</p>
<p>Feedback Information</p> <p><i>Provide a high-level description of how and what type of feedback will be provided to students, for inclusion in DRPS.</i></p>	<p>During the tutorials, the students will receive formative feedback from the tutors.</p> <p>The individual coursework will be a written essay where the students will analyse a case study in depth, and they will prepare a design document to address ethical concerns in the chosen case study. A formative feedback</p>

	<p>will also be provided when the students submit an outline for their coursework.</p> <p>The exam will be pen and paper, and raw marks will be given.</p>																				
<p>Additional Feedback Information (for BoS use only)</p> <p><i>If not already included in the course narrative, provide further details on planned feedback arrangements. This includes how course feedback is solicited from the class and responded to, as well as what feedback students will get (either on work that contributes to their final mark, or not).</i></p> <p><i>The University is committed to a baseline of principles regarding feedback that we have to implement at every level, and the School encourages submission of at least one piece of written work for formative feedback.</i></p> <p><i>In general, formative feedback:</i></p> <ul style="list-style-type: none"> • Should say how students can improve. • Need not be on individual work (e.g., consider a lecture or document summarizing common issues.) • Can include oral feedback during labs/tutorials • Can include feedback from peers • Clickers/TopHat/equivalents can provide in-class feedback for both students and lecturer. • Is returned in time for other forms of assessment to which it relates, to allow feedforward. 																					
<p>Breakdown of Learning and Teaching Activities</p> <p><i>State how many hours students spend on each part of the course. The total should be 10 x course credits, but please also see the School policy on Workload and Assessment, which states that students should not be expected to spend more than 6-7 hrs/wk per 10 credits, including contact hours.</i></p> <p><i>Assume 10 weeks of lectures slots and 10 weeks of tutorials, but these need not all be used. As a guideline, a 10-pt course typically has 18-20 lecture hours, but should have only around 15 lectures of examinable material; the rest should be used for guest lectures, revision sessions, introductions to assignments, etc.</i></p>	<p>Contact hours</p> <table border="1"> <thead> <tr> <th>Hours</th> <th>Type</th> </tr> </thead> <tbody> <tr> <td>18</td> <td>Lecture Hours</td> </tr> <tr> <td>2</td> <td>Seminar/Tutorial Hours</td> </tr> <tr> <td>0</td> <td>Dissertation Project Supervision Hours</td> </tr> <tr> <td>0</td> <td>Supervised practical/Workshop/Studio hours</td> </tr> <tr> <td>2</td> <td>Feedback/Feedforward hours</td> </tr> <tr> <td>2</td> <td>Summative assessment hours</td> </tr> <tr> <td></td> <td>Revision Session Hours</td> </tr> </tbody> </table> <p>Non-contact hours</p> <table border="1"> <thead> <tr> <th>Hours</th> <th>Type</th> </tr> </thead> <tbody> <tr> <td>76</td> <td>Directed Learning & Independent Learning hours</td> </tr> </tbody> </table> <p>Total hours: 100</p>	Hours	Type	18	Lecture Hours	2	Seminar/Tutorial Hours	0	Dissertation Project Supervision Hours	0	Supervised practical/Workshop/Studio hours	2	Feedback/Feedforward hours	2	Summative assessment hours		Revision Session Hours	Hours	Type	76	Directed Learning & Independent Learning hours
Hours	Type																				
18	Lecture Hours																				
2	Seminar/Tutorial Hours																				
0	Dissertation Project Supervision Hours																				
0	Supervised practical/Workshop/Studio hours																				
2	Feedback/Feedforward hours																				
2	Summative assessment hours																				
	Revision Session Hours																				
Hours	Type																				
76	Directed Learning & Independent Learning hours																				

<p>Reading List/Learning Resources <i>You are encouraged to create resource lists using</i> LEGANTO</p>	<p>Required readings will be primarily from open access papers listed on the course website.</p> <p>A representative reading list is as follows:</p> <ul style="list-style-type: none">- Lin, P., Abney, K. and Jenkins, R. (2019) Robot Ethics 2.0, Oxford University Press- Wallach, W., Allen, C. (2009) Moral Machines, Oxford University Press- Dignum, Virginia. "Responsible artificial intelligence: designing AI for human values." (2017).- Boddington, Paula. Towards a code of ethics for artificial intelligence. Cham: Springer, 2017.
--	---

3. Further information for BoS consideration: sample materials

A full proposal for a new course must include examples of exercises and assessment. Please provide these below, along with publicity information if the course is to be advertised outwith the School.

<p>Course information and publicity <i>The course web page (typically the Learn landing page) will be linked from the Sortable Course List, and information such as timetables and assignment deadlines must be made available prior to the start of the academic year. Please specify here if any additional info/publicity is needed for your course: typically only if it is aimed largely at non-Sol students.</i></p>	<p>-</p>
<p>Sample tutorial/lab sheet questions <i>Provide a list of tutorial questions and answers and/or samples of lab sheets. These need not be fully fleshed out but should indicate what sort of exercises will be provided to help students learn the material.</i></p>	<p>Tutorials will be based on case-studies that will be developed to discuss in small groups. The focus will be on particular ethical challenges that exist in the case studies. Some example short case-studies that could be developed are provided below:</p> <p>[1] Facebook Emotional Contagion Study In 2012, Facebook researchers conduct an experiment with 700.000 Facebook users. For a week, some users are exposed to happy and positive content; while the remaining users are shown sadder content. At the end of one week, the researchers find that users were more likely to share positive (negative) status messages if they were part of the group seeing positive (negative) content. Hence, emotional contagion was possible by only seeing friends' shared content.</p> <p>[2] Google - Project Maven In 2018, Google had an agreement with the United States Pentagon to support the development of to human-identifying drone technology. For this, Pentagon wanted to have access to data sets collected by Google products to train their AI models. This is one use case where the values of employees conflicted with the values of the organisation.</p> <p>[1] Kramer, A. D., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. <i>Proceedings of the National Academy of Sciences</i>, 111(24), 8788-8790.</p>

	<p>[2] https://www.theverge.com/2018/6/1/17418406/google-maven-drone-imagery-ai-contract-expire</p>
<p>Sample assessment materials</p> <p><i>If the course is primarily assessed by exam, provide a sample exam question with model answers. Any non-standard exam format must be justified. The online list of past exam papers gives an idea of typical and alternative exam formats:</i></p> <p>http://www.inf.ed.ac.uk/teaching/exam_papers/</p> <p><i>If the course is largely or primarily assessed by coursework, provide a sketch of a possible assignment with an estimate of effort against each sub-task and a description of marking criteria.</i></p>	<p>Company X decided to make an internal audit for one of its client projects. Happy-Go-Lucky, Inc., an imagined photo service company looking for a smile detection algorithm to automatically trigger the cameras in their installed physical photo booths [1].</p> <p>Company X has designated five AI principles: "Transparency", "Justice, Fairness & Non-Discrimination", "Safety & Non-Maleficence", "Responsibility & Accountability" and "Privacy".</p> <p>[1] Use case is adapted from the following paper: Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. In <i>Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)</i>. Association for Computing Machinery, New York, NY, USA, 33–44. DOI:https://doi.org/10.1145/3351095.3372873</p> <ol style="list-style-type: none"> In class, we discussed various AI principles that an organization may want to follow. Consider Company X's AI principles, and answer the questions. <ol style="list-style-type: none"> State the definition of one of the AI principles. Name one AI principle that Happy-Go-Lucky could violate. Justify your response. Name one solution to mitigate this problem. Use Failure Modes and Effects Analysis (FMEA), methodical and systematic risk management approach to examine Happy-Go-Lucky company. <ol style="list-style-type: none"> Describe what FMEA is. Choose one specific feature with low risk priority and apply the methodology. Choose one specific feature with high risk priority and apply the methodology. As Company X continues its internal audit, it gets access to the dataset that Happy-Go-Lucky is using to train their smile detection algorithm.

	<p>(a) Name one artifact that Company X could use to learn more about the dataset. Give a definition of this artifact. We will call this artifact A.</p> <p>The artifact A includes the following demographic details: (58.1% female, 42% male), (77.8% aged 0-45, 22.1% aged over 46 and 14.2% lighter-skinned, 85.8% darker-skinned).</p> <p>(b) By looking at A, what is the AI principle Happy-Go-Lucky violates? Justify your response.</p> <p>(c) What mitigation strategy could be recommended to address the problem in 3.(b)? How would you ensure that this strategy is good?</p>
<p>Any other relevant materials <i>Include anything else that is relevant, possibly in the form of links. If you do not want to specify a set of concrete readings for the official course descriptor, please list examples here.</i></p>	

4. Additional Course Details for DRPS

Except where otherwise noted, these fields are required for entry into EUCLID and will be visible to students in the DRPS entry.

<p>Planned Academic Year of Delivery <i>(The first year you anticipate the course running, e.g. AY 2019-20)</i></p>	AY 2021-22
<p>Course Organiser <i>(By default, the course proposer)</i></p>	Nadin Kokciyan
<p>Intended Delivery Period</p>	<input type="checkbox"/> Semester 1 <input checked="" type="checkbox"/> Semester 2 <input type="checkbox"/> Full Year <input type="checkbox"/> Summer <input type="checkbox"/> Other (please specify):
<p>Timetable considerations/conflicts <i>For School use. Please specify any constraints to be considered (e.g. overlap of popular combinations, other specialism courses, external courses etc). Include whether the semester delivery is constrained or could be flexible.</i></p>	
<p>Is this course available to visiting students?</p>	<input checked="" type="checkbox"/> Yes (default) <input type="checkbox"/> No If no, please provide a justification here:
<p>Required pre-requisite courses <i>Use sparingly: these are enforced in PATH and can only be waived by approval from the School's Curriculum Approval Officer. Note that cross-year required pre-requisites may prevent MSc students from registering; consider using recommended pre-requisites or "other requirements" instead.</i></p>	<input checked="" type="checkbox"/> No <input type="checkbox"/> Yes (please specify full course name(s) and code(s)):
<p>Recommended pre-requisite courses</p>	<input checked="" type="checkbox"/> No <input type="checkbox"/> Yes (please specify full course name(s) and code(s)):
<p>Required co-requisite courses <i>Specify any courses that must be taken in parallel with the existing course. Note that this leads to a timetabling constraint that should be mentioned elsewhere in the proposal.</i></p>	<input checked="" type="checkbox"/> No <input type="checkbox"/> Yes (please specify full course name(s) and code(s)):

<p>Prohibited Combinations Specify any courses that may not be taken in combination with the proposed course].</p>	<p><input type="checkbox"/> No <input checked="" type="checkbox"/> Yes (please specify full course name(s) and code(s)): <i>INFR10022 Professional Issues¹, PHIL10167 Ethics of Artificial Intelligence</i></p> <p>¹ PI is a prohibited combination because the updated version will also use ethical case studies (possibly some of the same ones), and to keep enrolment manageable initially. We can revisit this decision later if need be.</p>
<p>Other Requirements/Additional Information This information is often used by MSc students and students from other Schools to see if they have appropriate background without having done our School's courses. So please avoid course titles, instead list specific knowledge and skills (such as mathematical concepts, programming ability or specific languages, etc).</p> <p>Also list any other constraints on registration, for example: "Only available to 4th Year Informatics students including those on joint degrees." or "This course is open to all Informatics students including those on joint degrees, and to students in the School of Mathematics. Other external students whose DPT does not list this course should seek permission from the course organiser."</p>	<p><input checked="" type="checkbox"/> No <input type="checkbox"/> Yes (please specify):</p>
<p>Visiting Student Pre-requisites</p>	<p><input checked="" type="checkbox"/> Same as "other requirements" <input type="checkbox"/> Different than "other requirements" (please specify):</p>

5. Placement in degree programme tables: for level 9-11 courses only

This section is for consideration by the Board of Studies and will be used later by ITO to determine where the course will be added to existing degree programme tables.

<p>Is this course restricted to students on a specific degree? <i>E.g., some courses are only available to students on a specific CDT or MSc.</i></p>	<p><input checked="" type="checkbox"/> No <input type="checkbox"/> Yes (please specify and provide justification):</p>
<p>Is this course compulsory for students on any degree(s)?</p>	<p><input checked="" type="checkbox"/> No <input type="checkbox"/> Yes (please specify and provide justification):</p>
<p>Any issues for part-time students? <i>Normally, part-time students have access to the same courses as full-time students on the equivalent degree. If you anticipate any problems with this, please specify here.</i></p>	

For optional courses:

If this course is available but non-compulsory for students on various degrees (most courses), please fill in this section. The choices here determine where the course appears in degree programme tables (DPTs) and the 2-3 character tags are displayed in the Informatics sortable course list.

<p>Should this course be tagged as 'ML' (machine learning foundations and methods)? <i>Courses with the ML tag are typically very high-demand and most degrees limit the number of ML credits. If your course might appeal to a similar audience but draw off students from these large courses, please select 'no' and choose one of the tags below.</i></p>	<p><input checked="" type="checkbox"/> No <input type="checkbox"/> Yes</p>
<p>If you chose 'no', please choose at least one of the following tags... <i>Ideally, select exactly one, unless there is a good argument for more than one. These three are used in various combinations for many of our degrees.</i></p>	<p><input type="checkbox"/> FSS (CS foundations, systems, and software) <input checked="" type="checkbox"/> AIA (artificial intelligence applications and paradigms) <input type="checkbox"/> COG (cognitive science: including HCI and NLP courses, but not most other AI courses. Please restrict to courses most relevant to natural cognition.)</p>
<p>...and also tick if any of the following tags or categories apply. <i>Do not tick any of these if you selected 'ML' already.</i></p>	<p><input type="checkbox"/> NS (natural systems: e.g., computation by or about biological or social systems. Many COG courses are also NS. This tag is mainly relevant for MSc in Informatics.) <input type="checkbox"/> SE (software engineering: including courses that are highly relevant to SE degrees. All SE courses should also be FSS. This tag is mainly relevant for UG SE degrees.)</p>

	<ul style="list-style-type: none">___ Databases and data management systems (used for Data Science MSc and MSc(R))___ Unstructured data and applications (used for Data Science MSc and MSc(R))___ Level 11 Security courses (used for Security MSc)
If you are not sure which tags are most appropriate or have other questions about this section, please note any comments/issues here.	

6. Comments from colleagues

All course proposal should be sent to relevant colleagues in the area as well as to the appropriate year organizer and BoS Academic Secretary for comment in good time before the BoS meeting. Please indicate here what feedback has been solicited and received.

<p>Additional Comments <i>Summarise any comments received from relevant individuals prior to proposing the course. If you have not discussed this proposal with others please note this.</i></p>	<p>Sharon Goldwater reviewed this course proposal and provided important feedback. I revised it regarding her suggestions.</p>
<p>Year Organiser Comments <i>Year Organisers are responsible for maintaining the official Year Guides for every year of study, which, among other things, provide guidance on available course choices and specialist areas. The Year Organisers of all years for which the course will be offered should be consulted on the appropriateness and relevance on the course. Issues to consider here include balance of course offerings across semesters, subject areas, and credit levels, timetabling implications, fit into the administrative structures used in delivering that year.]</i></p>	
<p>BoS Academic Secretary Comments <i>Proposals must be checked by the Secretary of the Board of Studies prior to discussion at the actual Board meeting. This is a placeholder for their comments, mainly on the formal quality of the content provided above.</i></p>	