



UNIVERSITY OF
CAMBRIDGE

Institute of Continuing Education

MSt in AI Ethics and Society
Indicative Course Guide
2021-2023

Contents

Welcome.....	5
The Institute of Continuing Education.....	5
The Leverhulme Centre for the Future of Intelligence	5
The Institute of Continuing Education MSt Programme.....	5
Academic Course Leads and Administration.....	5
Colleges and College Tutors.....	6
Welcome to the MSt in AI Ethics and Society	6
Welcome from the Course Team.....	7
Cognate Faculties	7
MSt Degree Sub-Committees	8
Cognate Faculty Libraries.....	8
Your MSt Course Team	9
Course Information	9
Course Overview.....	9
Course Aims and Objectives.....	10
Course Outcomes	10
Knowledge and Understanding.....	10
Skills and Other Attributes	10
Advice on approaching the course.....	11
Pre-course Preparation	11
Learning and Teaching Approaches for Modules of the MSt.....	11
Self-study requirements	12
Module 1: The Nature and History of AI	12
Module 2: Ethical and Societal Challenges	14
Module 3: Theories and Methods <i>Exploring critical and analytical frameworks</i>	15
Module 4: Governing AI <i>Exploring AI Leadership, Practice and Governance</i>	16
Module 5.....	17
Further information regarding Assessment.....	18
Essays	18
Dissertation	19
Identifying Your Topic.....	19
Supervisions and Your Supervisor	19
Word counts.....	19
Referencing.....	19
Assessment Criteria.....	19

Course Calendar	21
Special Circumstances: Dissertation Submission Deadline Deferral	21
Reading Lists	21
General Readings	22
Useful Contacts	23
Provisional module teaching format.....	26

Welcome

Welcome to the University of Cambridge. We are delighted that you have chosen to join our Master of Studies (MSt) programme at the University's Institute of Continuing Education (ICE). We look forward to working with you.

The Institute of Continuing Education

Continuing education has been part of the University of Cambridge's work since the 1870s and the University remains committed to the principle of providing opportunities for learning throughout life. The Institute of Continuing Education (ICE) has students studying part-time courses that range from day-schools to full MSt degrees.

ICE's administrative headquarters, teaching facilities and residential centre are at Madingley Hall, an elegant country house built in the 16th century and set in gardens of about seven acres, designed in the 18th century by Capability Brown. The Hall is located in the village of Madingley, three miles west of Cambridge with easy access from the M11, the A14, Cambridge rail station and with good links to London airports. See www.ice.cam.ac.uk and www.madingleyhall.co.uk for further information.

This Handbook provides you with key information that will support you through your period of study. **Please ensure you read it in its entirety and continue to refer to it throughout your time with us.** This Handbook includes information on the University, ICE and Madingley Hall as well as details of your course, the syllabus and your assessments.

The Leverhulme Centre for the Future of Intelligence

Our mission at the Leverhulme Centre for the Future of Intelligence (CFI) is to build a new interdisciplinary community of researchers, with strong links to technologists and the policy world, and a clear practical goal: to work together to ensure that we humans make the best of the opportunities of artificial intelligence as it develops over coming decades.

The Institute of Continuing Education MSt Programme

A University of Cambridge MSt degree is an academically rigorous postgraduate programme undertaken on a part-time basis; it contains a significant taught element and a dissertation. Our MSt programmes are not simply full-time programmes studied on a part-time basis; they are specifically designed as part-time programmes, to be flexible and accessible and to accommodate the needs of students with a range of commitments. The modular structure of consolidated teaching blocks enables students living away from Cambridge, including international students based overseas, to study and develop their professional and academic interests for career progression or personal development.

Depending on your course, you will be a member of either a Department or a Faculty and the content, structure and supervisory systems of your course will be overseen by one of the University's Degree Committees. The Degree Committee is responsible for the structure and content of your course, for monitoring your progress and for final approval of your award.

The regulations of the MSt programme are set out in the University's Statutes and Ordinances here www.admin.cam.ac.uk/univ/so/2019/chapter07-section17.html

Academic Course Leads and Administration

You will have three Course Leads who are responsible for the academic standard of your course and who will be supporting your overall progress. Your Course Leads will ensure you

are informed of the course structure, including details of mandatory and optional modules (if applicable) and timetables for lectures and assessments; assessment; academic conventions and feedback. Your Course Leads, along with the wider course team, will also advise you on personal and skills development and will be a source of advice and support.

You will also be supported by a Course Administration Team who will be in regular contact with you throughout your study. In addition to supporting students, the Course Administration Team support the Course Leads in arranging lectures, modules and assessments.

Colleges and College Tutors

There are thirty-one independent Colleges affiliated with the University, each with its own unique character. As an MSt student you will also be a member of a College. Your College is an important part of your life at Cambridge and will allow you to mix with individuals from many different disciplines. The University provides teaching and research supervision for graduate students, and the Colleges provide pastoral support within a stimulating academic community.

A College Tutor is generally responsible for the wellbeing and pastoral care of all students and the coordination of tutorial activities. College Tutors will routinely communicate information about health and welfare services, travel grants and bursaries as well as more specialised needs like financial assistance and disability resources.

Colleges arrange your matriculation (enrolment) at the start of your course and, on the successful completion of your studies, the College Praelector presents you for the award of the MSt degree at the Senate House. The College fee payment, together with the composition (University tuition) fee, is a requirement of formal matriculation (the declaration that formally marks your admission to membership of the University).

If you are ill or experience other problems that may, in any way, affect your attendance, performance or submission of any assignments or dissertation you *must immediately* contact your Academic/Course Director, your Course Administration Team and your College Tutor.

Welcome to the MSt in AI Ethics and Society

Welcome to Cambridge. We are pleased that you have chosen to study the Master of Studies in AI Ethics and Society which is managed by the Institute of Continuing Education (ICE) with oversight from the University's Faculty of Philosophy. We look forward to working with you and meeting you during the modules, online activities and social events we have planned.

Welcome from the Course Team



Dr Stephen Cave is Executive Director of the Leverhulme Centre for the Future of Intelligence, Senior Research Associate in the Faculty of Philosophy, and Fellow of Hughes Hall, all at the University of Cambridge. Stephen earned a PhD in philosophy from Cambridge, then joined the British Foreign Office, where he served as a policy advisor and diplomat. He has subsequently written and spoken on a wide range of philosophical and scientific subjects, including in the *New York Times*, *The Atlantic*, and on television and radio around the world. His research interests currently focus on the nature, portrayal and governance of AI.



Dr Henry Shevlin (PhD, CUNY Graduate Center, 2016; BPhil, Oxford, 2009) is a Senior Research Fellow with the [Kinds of Intelligence](#) programme. His work focuses on issues at the intersection of philosophy of mind, cognitive science, and animal cognition, with a particular emphasis on perception, memory, and desire. Since 2015, he has been serving as a student committee member of the Association for the Scientific Study of Consciousness.



Dr Jonnie Penn, FRSA, is a historian of information technology, broadcaster, and public speaker. He is an Affiliate at the Berkman Klein Center at Harvard Law School, a Research Fellow at St. Edmunds College at the University of Cambridge, a *New York Times* bestselling author, and a fellow of the Royal Society of the Arts. He has held prior fellowships at the MIT Media Lab, Google, and the British National Academy of Writing. He writes and speaks widely about the future of work, data governance, youth and worker empowerment, and sustainable digital technologies.



Maya Indira Ganesh is a technology researcher and writer whose work investigates the social, cultural, and political implications of the 'becoming-human' of machines, and vice versa. Maya spent 15 years working at the intersection of gender justice, technology, and human rights with Indian and international NGOs.

She is completing a DPhil that will be awarded by the Doctoral College on Knowledge Cultures and Digital Media at Leuphana University, Lüneburg, Germany. Her work examines AI and autonomous technologies as ontologically multiple, and fragmented, and the implications of this for the shaping of 'ethics' in data-centric worlds.

Cognate Faculties

This MSt degree is developed and taught by the University's [Leverhulme Centre for the Future of Intelligence](#) (LCFI), a global research centre at the forefront of AI Ethics and impact research, in partnership with the Institute of Continuing Education. Because the programme is run by a specialist research centre, rather than by a department, the curriculum is uniquely multidisciplinary, informed by up-to-the-minute research developments, and incorporates experts from diverse areas, including philosophy, machine learning, computer science, policy, law, and more. Visit <http://lcfi.ac.uk/master-ai-ethics/> for more information on the Leverhulme Centre for the Future of Intelligence.

The MSt is offered in association with the Faculty of Philosophy. The administration and teaching of the MSt is co-ordinated by the ICE and its academic standards monitored and assured by a sub-committee of the Degree Committee of the Faculty of Philosophy. This sub-committee also includes representatives from the Faculty of Philosophy.

The Faculty of Philosophy, University of Cambridge is one of the world's leading institutions for teaching and research in philosophy and is part of the University's School of Arts and Humanities. The website of the Faculty of Philosophy can be found at: www.phil.cam.ac.uk

MSt Degree Sub-Committees

Your MSt Degree Sub-Committee is responsible for overseeing the delivery of your course. Its responsibilities include (but are not limited to): the initial selection of applicants; appointment of dissertation supervisors; approving dissertation titles, design of the syllabuses; issues of quality assurance; and oversight of assessment and moderation procedures.

The Degree Sub-Committee has delegated responsibilities from each cognate Degree Committee and acts as the formal conduit for communication between your respective Faculty and ICE in all matters relating to the MSt.

Cognate Faculty Libraries

As a student on the MSt in AI Ethics and Society you are encouraged to join the Faculty of Philosophy's Library, (in addition to the University Library and other University libraries), which is located at the Philosophy Faculty on Sidgwick Avenue. This library holds a large collection of books and journals which may assist you with your assignments and dissertation. See <https://www.phil.cam.ac.uk/library/finding-us> for further information.

Your MSt Course Team

Your course team is made up of a number of dedicated academic and administrative staff who will support you throughout your time at ICE. You will be able to discuss any issues relating to your course with either your teaching module tutors, your Academic Director/Course Director or your Course Administration Team.

Academic Staff	Title	Email
Dr Stephen Cave	Academic Director for AI Ethics and Society	sjc53@cam.ac.uk
Dr Henry Shevlin	Course Co-Lead of the MSt in AI Ethics and Society	hfs35@cam.ac.uk
Dr Jonnie Penn	Course Co-Lead of the MSt in AI Ethics and Society	inp28@cam.ac.uk
Dr Maya Indira Ganesh	Course Co-Lead of the MSt in AI Ethics and Society	mi373@cam.ac.uk
Dorian Peters	Research Programmes Coordinator	dp605@cam.ac.uk
Course Administration		
Zara Kückelhaus	Head of Academic Centre Administration	zara.kuckelhaus@ice.cam.ac.uk
Susannah Lewis	Course Co-ordinator	susannah.lewis@ice.cam.ac.uk
AI Course Team	Course Administrators	ai@ice.cam.ac.uk

Your Course Leads will receive reports on your performance from the course tutors and will discuss your progress with your dissertation supervisor and/or co-supervisor as necessary.

Course Information

Course Overview

The Master of Studies (MSt) in AI Ethics and Society is a two-year, part-time degree.

The content of the MSt aims to equip students with the skills and knowledge necessary to address the ethical and societal challenges arising from the use of Artificial Intelligence. It provides a professionally relevant teaching and learning environment informed by the forefront of academic research on AI and its ethical and societal impacts.

The MSt takes place over two academic years. During the first year you will be required to undertake three of the four taught elements of the MSt. These modules of the MSt are delivered using a blended learning approach. The final taught element is during the first term of the second year of the course which then continues through individual supervision.

The first-year modules normally consist of about eight taught classes per term, of between 60 – 90 minutes each, and eight accompanying Discussion Seminars of between 30-60 minutes each. An indicative timetable is supplied at the back of this handbook; finalised timetables for the modules will be circulated to students a few weeks before each module and the VLE area for that module will open at the same time. All students are assigned an individual tutor. As well as the classroom teaching, students will have a 30-minute tutorial for

each module. For the first module this is likely to be in a group of about six students; tutorials for the other three modules will be one-to-one.

Course Aims and Objectives

The overall aims of the MSt in AI Ethics and Society are to:

- Ensure that engagement with the ethical and societal challenges of AI is thoroughly informed by the knowledge, theories and methods of relevant established academic disciplines
- Provide a critical overview of the state of current interdisciplinary research on the challenges of AI
- Develop the critical research skills and analytical abilities needed to identify and address ethical challenges as they arise in practice from applications of AI
- Instil thorough knowledge of the role that different governance approaches can play in navigating the challenges of AI, and the ability to critically analyse those approaches
- Develop experts in the ethical and societal implications of AI with the ability and initiative to identify and address the challenges of AI across sectors and society
- Create a network for such experts to collaborate and continue learning as leaders in the field of AI ethics and governance

Course Outcomes

Knowledge and Understanding

By the end of this programme, students will be able to:

- Discuss the ethical and societal challenges of AI with a full understanding of its history and its relationship to other disciplines and technologies
- Identify the capabilities of current AI systems, their key applications and the potential ethical and societal challenges of those applications
- Evaluate key ethical and societal challenges arising from the use of AI and the existing critical literature
- Apply theories and methodologies from a range of established disciplines and demonstrate their use in addressing ethical and societal challenges raised by AI
- Analyse the strengths and weaknesses of current governance approaches for addressing the challenges posed by AI
- Apply the systematic understanding of AI ethics and governance to develop new insights

Skills and Other Attributes

By the end of this programme, students will be able to:

- Synthesise and analyse research and advanced scholarship across disciplines

- Put theoretical and academic knowledge into practice
- Structure extended pieces of written work and present arguments clearly and effectively
- Plan and implement an independent research project
- Deal with complex issues both systematically and creatively, and show originality in tackling and solving problems

Advice on approaching the course

The transition from professional life back into academia can be challenging. This is especially true in a field like AI Ethics and Society, the substance of which remains relatively new and thus open for debate. The Course Leaders recommend that all students set adequate time to absorb their weekly readings so that they come to weekly discussion sessions prepared to ask questions of the instructors and of each other. Given students' knowledge of applied settings, each will be expected to contribute their own experiences and to problem solve together throughout the duration of the course.

In addition to the above, students are strongly encouraged to engage the many services offered by ICE to help learn and/or re-learn aspects of academic writing such as grammar, citations, and basic argumentation. These skills will prove invaluable when the time comes for your written materials to be graded. It would be unwise to wait until the first essay is due to 'stretch your legs' in regards to academic writing. To make the most of this course, and of your time at the University of Cambridge, we ask students to take these responsibilities seriously.

Pre-course Preparation

You will be given a reading list before the programme starts and subject-specific reading before each module.

Learning and Teaching Approaches for Modules of the MSt

We recognise that studying for a part-time Master's degree is a serious undertaking particularly when students are balancing work and family commitments. To ensure you have the best possible chance of completing the course, we take a blended learning approach which involves a combination of online learning and face-to-face teaching for each module of the MSt. The teaching and learning approaches adopted for your course are as follows:

Lectures: are likely to be in person and online during [term time](#). Where appropriate, a lecture will be accompanied by handouts containing key readings and other relevant information.

Discussion Seminars: will be held online each week during term time. In a seminar, your tutor will lead the session by posing the key questions and chairing the discussion. All students will be expected to participate in group discussion.

Online tasks: On occasion, assigned and recommended tasks will be completed through the VLE. These tasks might involve working through a series of questions, watching a video clip, listening to a podcast or reading a document. You will be required to upload examples of work to the VLE, and to comment on the work of others.

Independent reading: there is a considerable quantity of independent reading expected of you and you will need to take appropriate steps to ensure you can access the material you need in good time.

Tutorials: you will have the opportunity to discuss work with your tutor prior to the final submission of work.

Self-study requirements

You will be expected to read all of the pieces identified by your Course Leads as essential and to select appropriate readings from the additional reading list. In addition, you will need to complete the necessary online tasks set through the VLE. In general, the essential and additional reading lists are supposed to be starting points, and you are expected to read widely. In addition to writing essays and completing the tasks on the VLE, you are also required to participate in online discussions with other students, upload resources and comment on the work of others.

For each module, you are expected to:

1. complete all written work agreed with the tutor punctually and to a high standard;
2. act upon feedback from the tutor in order to develop the work submitted;
3. complete all required online tasks on the VLE;
4. upload any requested work onto the VLE to share with other students;
5. comment on the work of other students in a supportive and critical fashion;
6. access the required written materials through the VLE, online or in person e.g. the University Library.

Module 1: The Nature and History of AI

Residential dates: 27 September - 1 October 2021

Aims: To provide students with theoretical, academic and practical understanding of how artificial intelligence has been developed, used and understood historically across different traditions, and how it is being applied in society today.

Key areas:

- The technical foundations of AI and the current capabilities and status of the technology
- Current applications of AI across a range of domains and sectors
- The history of AI and its relationship to other disciplines and technologies, including the history of computing and administration
- The nature and measurement of intelligence, and comparisons between human, animal and artificial intelligence

Structure:

- Approximately 60 minutes worth of lecture material per week during Michaelmas Term. Available for asynchronous study. (Tues 6 Oct – Fri 4 Dec)
- One to two discussion Seminars each week. Students are not expected to attend both but are strongly encouraged to attend at least one. If two Seminars are offered, it is to accommodate different students' working schedules.

Assignment One

2,000-word essay (8% of final grade).

In this essay, students are asked to consider how historical forces influence contemporary understandings of AI ethics and its implications for society.

Deadline for submission of assignment one:

Wednesday 8 December 2021, 12:00 noon GMT

Module 2: Ethical and Societal Challenges

Residential dates: 10 - 14 January 2022

Aims: To provide students with a comprehensive understanding of key ethical and societal challenges raised by AI, through engagement with the contemporary critical literature and case studies.

Key areas:

- Critical discussion of the following themes:
 - Privacy
 - Fairness and equality
 - Safety
 - Accountability
 - Human dignity and autonomy
- The relationship between the near- and long-term challenges of AI
- Comparison of different global perspectives

Assignment Two

4,500-word essay (14% of final grade).

Students are asked to apply theories and frameworks from normative ethics, meta-ethics, and/or political theory to assess the moral dimensions of an issue in artificial intelligence research. Students are encouraged both to offer context-specific detail about the problem and its technological dimensions while also locating it within broader ethical and political debates.

Deadline for submission of assignment two:

Wednesday 23 March 2022, 12:00 noon GMT

Module 3: Theories and Methods

Exploring critical and analytical frameworks

Residential dates: 20 - 24 June 2022

Aims: To increase rigour and depth in understanding and analysing the ethical and societal challenges of AI by introducing students to foundational knowledge, theories and methods in established academic disciplines.

Key areas:

- Theories and methods from the following disciplines:
 - Philosophical ethics
 - The history and philosophy of science
 - Literary and cultural studies
 - Social and behavioural sciences
 - Futures studies and foresight methods
 - Critical design studies

Assignment Three

4,500-word essay (14% of final grade)

Students are asked to assess an aspect of AI's social implications using one of the various methodologies surveyed during Module 3. Where possible, students are encouraged to compare taught methods with their own complex lived experience of AI ethics in use in industry, government, medicine, and other areas of civil society.

Deadline for submission of assignment three:

Wednesday 6 July 2022, 12:00 noon BST (GMT+1)

Module 4: Governing AI

Exploring AI Leadership, Practice and Governance

Residential dates: 12 - 16 September 2022 (tbc)

Aims: To critically engage with a range of practical approaches to navigating the ethical and societal challenges of AI, including those found in policy, regulation, law, ethics principles, and social action.

Key areas:

- Comparison and critical analysis of current AI policy initiatives worldwide
- Overview and critical discussion of different codes of practice and principles for AI ethics, and their implementation
- Critical discussion of methods for ethical impact assessment
- Critical discussion of methods for ethical design
- The role of activism and civil society

Assignment Four

4,500-word essay (14% of final grade):

Indicative assignment: Students are asked to contrast and analyse a political-economic, systemic, or institutional aspect of AI in the world and to assess how best to develop actionable policy and practices in varied contexts and workplaces. As part of this they will develop a framework on what might constitute an actionable AI governance model for a select work or public context.

Deadline for submission of assignment four:

Wednesday 7 December 2022, 12:00 noon GMT

Module 5

Module dates: September 2022 – July 2023

Assignment Five

Aims: To enable students to apply and develop their learning from Modules 1-4 through an innovative, independent research project in an area relevant to the course, topic and scope to be agreed with the supervisor.

In the second year (module 5), students will write a 15,000-word dissertation which accounts for 50% of the final grade.

Please see page 30 for a fuller description of the dissertation requirements.

Deadline for submission of the dissertation:

Wednesday 14th June 2023, 12.00 noon, BST (GMT+1)

Further information regarding Assessment

Assignments on the MSt AI Ethics and Society are divided into two components: the essays, taken as a group, and the dissertation.

Essays

The modules are assessed as follows:

- Module 1: 2,000-word essay (8% of final grade)
- Modules 2, 3 and 4: 4,500-word essay each (14% each of final grade)

When marking coursework, examiners will be asking:

1. What is the main achievement of this work? Is there an original contribution? If so, what is it?
2. Does the candidate show a good understanding of relevant material? Is the content of the work informative and insightful?
3. Does the candidate advance effective arguments contributing towards well-articulated conclusions?
4. Has the candidate used a sufficient number and range of appropriate sources? Are they effectively used and properly credited and cited?
5. Does the work have a clear and effective structure? Is the writing clear, grammatical, and free of typographical and other errors? Is the style of the references and footnotes clear and consistent?

(This work will receive a formal mark, which will count towards the student's final mark.)

All summative assessment is compulsory. Students will receive continual formative feedback throughout the course using a variety of strategies and techniques, including evidence of regular reflection.

When your assignments are returned to you (expected to be within 15 working days of the assessment submission deadline) they will have been 'double-marked'. This means that each tutor gives feedback on your work independently and awards a provisional mark (subject to moderation). Only when each tutor has done this separately do they confer and arrive at a reconciled mark (which takes account of both their individual marks).

Dissertation

The MSt in AI Ethics and Society Dissertation is the broad title given to the final work you will produce and submit in your second year. It will consist of a portfolio of work of 15,000 words (or equivalent, please see below).

Identifying Your Topic

You must decide upon your chosen form for your dissertation by the end of the final module in Year 1 of the course. Once a supervisor has been identified for your project, you will work with him or her over the course of the second year to produce your dissertation.

Supervisions and Your Supervisor

You will be assigned a supervisor for year 2 of your course who has expertise in your research area and who will offer advice on aspects of research and preparation of your dissertation. Your supervisor will help you clarify your own ideas and guide your own independent efforts. They will advise on the scope and structure of your dissertation and will help you plan your work. You might be allocated a co-supervisor as well. If so, this decision has been taken to reflect the combination of expertise required to help you complete your research and dissertation.

Further information on the relationship with your supervisor can be found here www.student-registry.admin.cam.ac.uk/information-supervisors.

Please speak to your Course Administration Team if you would like to see copies of previous dissertations.

Word counts

All word counts are inclusive of footnotes and appendices but exclusive of bibliographies. You should aim to submit work of a length close to that specified. If you submit significantly less you will naturally be at a disadvantage, and if you submit significantly more you may be penalised.

Referencing

When writing essays and dissertations, correct and consistent referencing is an important way to avoid accusations of plagiarism and ensure that people reading your work can locate the item.

This [link](#) provides information and guidance to help you with your references and bibliographies. If you are in any doubt about the style you should adopt, talk to your supervisor for further guidance or consult the Chicago Manual of Style 17th Edition or later.

Assessment Criteria

The assessment criteria adopted on the MSt in AI Ethics and Society are reproduced below for your guidance.

Mark 80+ First Class (Starred Distinction)	An outstanding and memorable performance in which all the qualities deemed to constitute first-class work are present in a remarkable degree. The work should be well researched and substantially original, bearing in mind that originality has many dimensions: it may reside, for instance, in the thesis defended; or in the way a known thesis is presented and defended. Such work might well form the basis for publication. Potential for outstanding PhD work.
Mark 75–79 First Class (Distinction) Mark 70–74 First Class (Boundary for PhD)	Work which is of high calibre both in the range and in the command of the material and in the argument and analysis that it brings to bear. The examiner would expect some elements of originality – which may consist in putting together material in novel ways – although originality alone would not guarantee marks in this range. Work in this class will generally meet the following criteria: the argument may be sophisticated, incisive or demonstrate flair; there may be a wealth of relevant information, showing exceptional knowledge and understanding of the issues involved; the approach may be unorthodox in the best sense, suggesting new and worthwhile ways of considering material. The submitted work may display evidence of extensive research imaginatively and convincingly deployed. Work which receives a mark of 70 to 74: A solid performance in which some of the criteria for first class work will clearly be present but not necessarily all. Shows potential for PhD work. Work which receives a mark of 75 to 79: A very strong performance demonstrating clear originality, and in which the qualities deemed to constitute first-class work are consistently well represented. Clear potential for good PhD work.
Mark 65–69 High Performance	Clearly proficient with a proper coverage of relevant material. Work may indicate broader range than the Pass category and should be reasonably well presented. Solid but on occasion unimaginative. Ambition of work clearly visible but not always carried through. The analysis and argument are generally good. Work at the upper end of this category shows evidence of a good and broad-based engagement with, and understanding of, the relevant material and organised in a clearly-argued, well-illustrated and relevant fashion. The essay or dissertation will usually contain material which displays evidence of high intelligence, and which is regularly, but not consistently, sophisticated in analysis, impressive in its display of relevant knowledge, and occasionally demonstrate flair.
Mark 60–64 Pass	Work which is basically competent, and, in the case of dissertations, reasonably independent. Interesting and provocative ideas may not be not carried through fully convincingly. The main thesis may be vague, too general, too unambitious or else over-ambitious. There may be gaps in the bibliography, deficiencies in the overall structure; and weaknesses of analysis and argument. A piece of work which is not always clearly written.
Mark 0–59 Fail	Work that, while it may show reasonable knowledge of the material, and serious effort, reveals deficiencies in understanding, organisation or breadth of reference. Work that is derivative or irrelevant, ignorant or extremely superficial. Work showing minimal understanding of material or serious deficiencies in argument.

Candidates are required to pass in each part of the examination separately (i.e. the essays and the dissertation), except in the following special circumstances:

- A candidate whose failure in the essays is **marginal** should be allowed to submit a dissertation, and a **high performance** in the dissertation may be taken into account by the Degree Committee in determining their recommendation to the Board of Graduate Studies. Students whose overall essay mark is a marginal fail will be warned by the MPhil Manager in May.

- Where a candidate's failure in the dissertation is **marginal**, a **high performance** in the essays may be taken into consideration by the Degree Committee in determining their recommendation to the Board of Graduate Studies.

At the end of the course, examiners may decide to hold an oral examination. Such an examination will in any case be necessary if the dissertation is judged to be a marginal failure or if there is a very marked discrepancy between the two examiners' independent reports on the dissertation.

Course Calendar

MSt in AI Ethics and Society Modules and Dissertation	Residential dates	Assessment Submission Dates
Year 1		
Module 1: The Nature and History of AI	27 September - 1 October 2021	Wednesday 8 December 2021 12.00 noon GMT*
Module 2: Ethical and Societal Challenges	10 - 14 January 2022	Wednesday 23 March 2022, 12.00 noon GMT*
Module 3: Theories and Methods	20 - 24 June 2022	Wednesday 6 July 2022 12.00 noon BST*
Year 2		
Module 4: Governing AI	12 - 16 September 2022 (tbc)	Wednesday 7 December 2022 12.00 noon GMT*
Module 5: Details of supervisions, dissertations, presentations etc		Dissertation submission : Wednesday 14 June 2023 12.00 noon BST*

Special Circumstances: Dissertation Submission Deadline Deferral

In *exceptional cases*, where you experience difficulties that are likely to mean you will be unable to complete the dissertation by the submission date, you should notify your Course Administration Team, Supervisor, College Tutor and MSt Academic/Course Director as soon as possible; your case will be reviewed by the cognate Degree Committee. If you are likely to need a break of a full academic year, then you should request permission to intermit. For further information, please see

www.cambridgestudents.cam.ac.uk/your-course/graduate-study/your-student-status/extending-your-end-registration-date.

Reading Lists

It is expected that you will carry out an extensive programme of reading as part of your studies. The study of AI in this context is not best served by tying every student to the same narrow group of texts (needs and interests vary) but some guidance is obviously necessary. The lists related to each module are indicative rather than prescriptive: you should have read

a substantial number of these texts by the end of the course, and your reading should be apparent in your writing and in your seminar contributions; however, there's no need to read in any particular order – read what seems useful to you at the time.

For all modules, please ensure that you also check the VLE, where your tutors will have added other assigned materials, including online resources. **The readings given below are suggested but not assigned.** There will also be a reading suggestion forum where students can recommend texts to each other.

General Readings

1. Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim code*. Oxford, England: Polity.
2. Boden, M. A. (2016). *AI: Its nature and future*. Oxford University Press.
3. Coeckelbergh, M. (2020). *AI ethics*. MIT Press.
4. Costanza-Chock, S. (2020). *Design justice: Community-led practices to build the worlds we need*. The MIT Press.
5. Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life and appropriating it for capitalism*. Stanford University Press.
6. Crawford, K. (2021). *The Atlas of AI*. Yale University Press.
7. D'Ignazio, C., & Klein, L. F. (2020). *Data feminism*. MIT Press.
8. Eubanks, Virginia. (2017). *Automating inequality: how high-tech tools profile, police, and punish the poor*. St. Martin's Press.
9. Lee, K.-F. (2018). *AI superpowers: China, Silicon Valley, and the new world order*. Houghton Mifflin Harcourt.
10. Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Penguin.
11. Scharre, P. (2018). *Army of none: Autonomous weapons and the future of war*. WW Norton & Company.
12. Zerilli, J. et al. (2021). *A Citizen's Guide to Artificial Intelligence*. MIT Press.

Edited Collections:

1. Cave, S., Dihal, K., & Dillon, S. (Eds.). (2020). *AI Narratives: A history of imaginative thinking about intelligent machines*. Oxford University Press.
2. Dubber, M. D., Pasquale, F., & Das, S. (Eds.). (2020). *The Oxford handbook of ethics of AI*. Oxford University Press.
3. Liao, Matthew (Ed.). (2020). *Ethics of Artificial Intelligence*. Oxford University Press.

Useful Contacts

Cambridge Rape Crisis Centre	http://cambridgerapecrisis.org.uk
Careers Service	www.transkills.admin.cam.ac.uk/skills-portal or the University Careers Service at www.careers.cam.ac.uk/index.asp
Childcare Office	www.childcare.admin.cam.ac.uk
Colleges	www.graduate.study.cam.ac.uk/colleges
Counselling Service	www.counselling.cam.ac.uk
CUSU (Cambridge University Student Union)	www.cusu.co.uk
Degree Ceremony	www.cambridgestudents.cam.ac.uk/your-course/graduation-and-what-next/degree-ceremonies .
Disability Resource Centre	www.disability.admin.cam.ac.uk
Equality and Diversity	www.equality.admin.cam.ac.uk
Events	http://talks.cam.ac.uk/ and CRASSH (Centre for Research in the Arts, Social Sciences and Humanities) at www.crassh.cam.ac.uk/ .
General Data Protection Regulations (GDPR)	www.information-compliance.admin.cam.ac.uk/data-protection/general-data-protection-regulation and www.information-compliance.admin.cam.ac.uk/data-protection/research-participant-data/
idiscover	http://idiscover.lib.cam.ac.uk
ICE (Institute of Continuing Education)	www.ice.cam.ac.uk
ICE Records	ice.records@ice.cam.ac.uk
Intermission	www.cambridgestudents.cam.ac.uk/your-course/graduate-study/your-student-status/non-medical-intermission and www.cambridgestudents.cam.ac.uk/your-course/graduate-study/your-student-status/medical-intermission
Library (ICE)	www.ice.cam.ac.uk/ice-library
Madingley Hall	www.madingleyhall.co.uk

Mitigating Circumstances	www.ice.cam.ac.uk/info/mitigating-circumstances
MSt programme regulations in the University's Statutes and Ordinances	www.admin.cam.ac.uk/univ/so/2019/chapter07-section17.html
OSSCA (Office for Student Conduct, Complaints and Appeals)	www.studentcomplaints.admin.cam.ac.uk/home/who-ossca
Plagiarism and Academic Misconduct and Turnitin	www.plagiarism.admin.cam.ac.uk/ and www.academic.admin.cam.ac.uk/education/student-operations/office-student-conduct-complaints-and-appeals-ossca and www.plagiarism.admin.cam.ac.uk/turnitin-uk
Raven	https://raven.cam.ac.uk/
Research Ethics	www.research-operations.admin.cam.ac.uk/managing-research-projects/research-ethics and www.research-integrity.admin.cam.ac.uk/research-ethics and www.data.cam.ac.uk/fag
Samaritans	www.samaritans.org
Social media guidelines	www.ice.cam.ac.uk/files/downloads/social_media_guidelines_students.docx
Society of College, National and University Libraries (SCONUL) Vacation Access scheme	www.sconul.ac.uk/page/about-sconul
Students with Additional Requirements	www.ice.cam.ac.uk/info/students-additional-requirements-courses-madingley-hall
Student Registration	www.cambridgestudents.cam.ac.uk/new-students/student-registration .
Student Wellbeing	www.studentwellbeing.admin.cam.ac.uk and www.cambridgestudents.cam.ac.uk/your-course/graduate-study
Supervisors	www.student-registry.admin.cam.ac.uk/information-supervisors .
Travel for research/study	www.cambridgestudents.cam.ac.uk/your-course/graduate-study/your-student-status/work-away
University Centre	www.unicen.cam.ac.uk
University Information Service	www.uis.cam.ac.uk and training www.training.cam.ac.uk/ucs .
University Library	www.lib.cam.ac.uk and training events https://training.cam.ac.uk/cul/event-timetable
University Sports Centre	www.sport.cam.ac.uk

Virtual Learning Environment	www.ice.cam.ac.uk/about-us/elearningsupport
Visas and Immigration	www.ice.cam.ac.uk/info/visas-international-students and www.internationalstudents.cam.ac.uk/ and www.iso.admin.cam.ac.uk/about-team/about-international-student-team .

Provisional module teaching format

Confirmed timetable will be provided in the VLE approximately four weeks before start date

	<u>Module 1</u> Michaelmas 2021 The Nature and History of AI	<u>Module 2</u> Lent 2022 Ethics and Societal Challenges	<u>Module 3</u> Easter 2022 Theories and Methods	<u>Module 4</u> Michaelmas 2022 AI Governance
Lead:	Dr Jonnie Penn	Dr Henry Shevlin	Combined leads	Dr Maya Indira Ganesh
Week 0	Residential 1 Introductions, foundations, and frameworks for AI ethics & Society	Residential 2. Debates, dilemmas, and controversies in the ethics of artificial intelligence	-	Residential 4: Situating AI in the world; developing practical applications and tools.
Week 1	Responsible AI: Introducing attempts to conceptualise 'AI'	Foundations of AI Ethics 1: Introducing methods and theories	Quantitative and Technical Analysis: Use and Abuse of Statistics	Principles in practice: AI and ethics principles and frameworks in practice
Week 2	Responsible AI: Prehistories of AI and measures of progress	Foundations of AI Ethics 2: First steps in applying ethical thinking to AI systems	Qualitative and Cultural Analysis	AI's Impact: Practical tech impact assessment tools: limits and opportunities.
Week 3	AI Histories 1: Datafication Logics - Biopolitics, Data Epistemologies, and Data Doubles	Foundations of AI Ethics 3: Moral disagreement and cross-cultural perspectives.	Intersectional approaches 1: engaging power, gender, race, and coloniality in making AI	Values in action: Participatory approaches to negotiating values in design and decision-making: Labour perspectives.
Week 4	AI Histories 2: Administrative Logics - Bureaucracy as Information System	AI Ethics Challenges 1: Privacy, censorship, and free speech	Intersectional approaches 2: engaging power, gender, race, and coloniality in making AI	Governance through infrastructures: Developing standards and benchmarking for AI technologies.
Week 5	AI Histories 3: Surveillance Logics - AI, Race, and Necropolitics	AI Ethics Challenges 2: Well-being, happiness, and autonomy	Accessibility 1 Disability, value sensitive design, HCI and digital anthropology	Working with government and public bodies: National and international approaches to engaging non-specialist institutions.

	<u>Module 1</u> Michaelmas 2021 The Nature and History of AI	<u>Module 2</u> Lent 2022 Ethics and Societal Challenges	<u>Module 3</u> Easter 2022 Theories and Methods	<u>Module 4</u> Michaelmas 2022 AI Governance
Week 6	AI Histories 4: Consumption Logics - The Hard Ecological Limits on an AI Future	AI Ethics Challenges 3: Nudging, targeting, autonomy	Accessibility 2 Data Justice and participatory action Research	AI governance from the ground-up: Civil society perspectives, activism, and re-imagining AI.
Week 7	'Cutting Edge AI': Case Studies 1 - Autonomous Vehicles	AI Ethics Challenges 4: Societal and algorithmic biases	Methods in Psychology and cognitive science in artificial intelligence	Whose future(s)? AI, future speculation, development and coloniality.
Week 8	'Cutting Edge AI': Case Studies 2 - Synthetic Text and Image Generation (e.g. GPT-3 et. al.)	AI Ethics Challenges 5: Justice, equality, and the digital divide	Methods in Legal and Policy Studies	Transnational governance of AI
Week 9	-	-	Residential 3 Focus: Dissertation Preparation & Case Studies	Communicating AI: Journalism, and public engagement on AI
Core outcome	Understanding and applying historical frameworks	Understanding and applying philosophical ethics	Establishing familiarity with varied frameworks and methods	Developing actionable policy and practices in varied workplaces