

**ETHICS IN MATHEMATICS
ESSAY TITLES
WINTER 2019-20**

MAURICE CHIODO

For students who attended the Ethics in Mathematics lecture course in Michaelmas 2019 and wish to prepare an essay on a related topic or case study.

CONTENTS

0. Introduction	2
0.1. Aims	2
0.2. Format and length	3
0.3. Submission	3
0.4. Feedback	3
1. Essay 1: Algorithmic fairness	4
2. Essay 2: Mathematicians and cryptography	5
3. Essay 3: A topic of your choosing!	6

0. INTRODUCTION

0.1. Aims.

These winter essays serve many purposes, and there are many new skills that you can learn, and knowledge that you can gain. Ultimately, what you get out of this is entirely up to you and what you choose to put in; there are no set expectations. As a guide, here are some things that you *might* learn or gain from doing the essay:

(1) Letter of reference.

The key objective of this exercise is for me to see your work and the ways in which you have understood concepts around Ethics in Mathematics. To that extent, I will write a letter of reference outlining this, which covers your participation in lectures, and the work that you submit. A letter of reference where I have seen your written work counts a great deal more, and is substantially more useful to you, than one where I have not. There is no pass or fail notion here, just the idea that I can comment on your understanding and ability more accurately.

(2) Learning to use L^AT_EX.

Yes, this is important! You're mathematicians, or are in the process of becoming mathematicians. Our community communicates in documents written almost exclusively in L^AT_EX. This is excellent practice, as you'll get some feedback on your L^AT_EX formatting at the end of the essay. Better to do it now in a non-examinable relaxed setting, than later in a CATAM write-up / Part III essay / PhD thesis! For example, when you first try and do quotations marks like "this", you'll probably end up doing "this".

(3) Looking up research papers.

It is actually a non-trivial task to find research papers and preprints online. Various publications are placed behind paywalls, and although you all have access to most of these (in theory) via the university, getting to them is never very easy. You'll also learn about arXiv, which is where mathematicians deposit preprints of papers before they go through the (long and arduous) refereeing and publication process. It is not uncommon for an arXiv preprint to go up online in year N , yet for the published paper to appear in year $N + 3$.

(4) Writing prose.

When was the last time you wrote an entire paragraph of text that didn't include any numbers, variables or subscripts? And what proportion of the sentences that you've written in the last 3 years have begun with words other than "If, then, else"? If we want to communicate mathematically-based ideas to non-mathematicians, then we need to learn how to write prose! This is important: we can't send politicians or upper management lines of C++ code or complicated formulae and expect them to extract the same meaning that we do. We need to *write* for them, and they simply won't read code or formulae.

(5) Looking in to something with technical and non-technical aspects.

Some of this will involve looking at the technical details of the work of various mathematicians; you're probably already quite competent at that. But you'll also need to do things like make judgement calls, interpret the actions, and analyse and predict things in a qualitative way ("The effects were harmful", without having a precise definition of "harmful".) You'll need to discuss what the mathematicians did, and why it may have been a bad idea! You'll need to make an *argument*, not a *proof*. This is a different skill to what is taught in tripsos, and now is as good a time as any to practise it. Remember, mathematics lecturers say "Here are true things", while lecturers in all other faculties simply say "Here are things".

(6) Doing research on Ethics in Mathematics.

This is **hard**. There are practically no well-defined resources on this; I've literally created a research area from almost nothing. So you may find it somewhat unproductive to simply Google "Ethics in Mathematics". Even with the various essay titles I've set, these topics aren't exactly phrased or discussed in the way that I've been discussing them. You'll need to find and read the existing works, and then analyse them in your own mind (think of this as rotating them to get a different/better/more useful point of view), and then discuss them from that particular point of view. There are no right or wrong answers here. The challenge is that you need to find something sensible to study, and then figure out something sensible to say about it.

0.2. Format and length.

Obviously, this should be formatted in L^AT_EX. Which style file you use is up to you (if you don't know what that means, then look it up). As for length: I'm not looking for a phone book here. Try and aim for 2000–3000 words; as an absolute upper-bound, I will *literally* stop reading after 4000 words (~8–10 pages, depending on your font size and margins). I have time constraints, sorry.

Don't feel pressured to complete everything that is listed in an essay description; you can chop and choose as you like. So long as you present *something* resembling a complete argument (so if an essay has 3–4 parts, perhaps you cover 1–2 of them). If you're feeling creative, you can add/modify part of an essay, provided it makes sense to do so.

Try and lay out your essay as you might lay out a Part III essay, with a title, abstract, table of contents, introduction, conclusions, and proper (full) referencing. There are places where it may be necessary to insert some formatted mathematics; that's fine, and I leave you to judge how much detail to put in. A couple of lines is ok, but a page of algebra should be relegated to an appendix (or removed entirely), as this is not a mathematics essay!

0.3. Submission.

So as to ensure that this does not take up any of your precious time in term, I'm setting a hard deadline of

5pm on Friday 17 January 2020

Please email the compiled pdf to me at mcc56@cam.ac.uk by this date. There is no point printing a hard copy.

0.4. Feedback.

If you would like feedback on your essay, please let me know and I can arrange an informal meeting (not a supervision!) during term.

1. ESSAY 1: ALGORITHMIC FAIRNESS

With the rapid development and deployment of algorithmic processes in society, many people have raised concerns about potential bias in these algorithms, and the ways in which that bias may affect people’s lives. Some mathematicians seem to have “come to the rescue”, creating systems that identify and measure bias in algorithms, but this has spawned even more (technical and non-technical) debate about how reasonable the bias-checking processes are! The issue is complicated.

We start by looking at a particular example of algorithmic decision making: the prison sentencing algorithm known as COMPAS, developed by Northpointe. You can find the “Practitioners Guide to COMPAS” at: http://www.northpointeinc.com/files/technical_documents/FieldGuide2_081412.pdf , which gives a lot of details of the algorithm.

1. Summarise the COMPAS algorithm does, how it works, how it’s derived, and how it’s applied.

The work of ProPublica was one of the first pieces to give a thorough analysis and testing of the COMPAS algorithm. It can be downloaded from:

<https://www.propublica.org/datastore/dataset/compas-recidivism-risk-score-data-and-analysis>

2. Download the COMPAS data set and explore it. Which variables seem to be ethically relevant and which variables were not collected but could have given insight into ethical issues?

You could then read the original ProPublica articles:

<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>

<https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>

3. Compare the insights of ProPublica with you own perspective on the data set and your own ideas on algorithmic fairness.

There are several technical, and non-technical, papers which deal with fairness in algorithms. Here are a few that you can read over:

- Cynthia Dwork et al., “Fairness Through Awareness”. Technical paper giving a mathematical way of measuring if an algorithm is fair.
- Cynthia Dwork et al., “Fairness Under Composition”. Technical paper showing how composition of fair algorithms need not be fair.
- Sorelle Friedler et al., “On the (im)possibility of Fairness”. Technical paper showing that various mathematical definitions of “fair” are inconsistent.
- Michael Skirpan, “The Authority of ‘Fair’ in Machine Learning”. Non-technical paper discussing the fairness of *trying* to use an algorithm for a particular decision-making task.
- Alexandra Chouldechova et al., “The Frontiers of Fairness in Machine Learning”. A semi-technical paper surveying existing works which discuss algorithmic bias and fairness.
- Reuben Binns, “Fairness in Machine Learning: Lessons from Political Philosophy”. A non-technical paper discussing how mathematicians are trying to re-invent the wheel when it comes to describing fairness, and that philosophers have already thought about this.

4. Try and make a summary of the algorithmic fairness debate, or at least certain parts of it.

Remember, you are free to pick and choose whichever aspects of the essay interest you the most.

2. ESSAY 2: MATHEMATICIANS AND CRYPTOGRAPHY

In recent years the work of mathematicians in state surveillance and cryptography has come under some degree of public scrutiny. This essay aims to explore some of that scrutiny, and look at the arguments that mathematicians, along with computer scientists, are making that might be changing the way mathematicians view their own involvement in cryptographic work.

We take as our starting point the Snowden revelations of 2013. There it was revealed that the NSA, in collaboration with many other surveillance agencies collectively known as the “Five eyes” network, had engaged in systematic mass digital surveillance of practically the entire internet-connected world. This is a very very large topic, and not one that can be covered in depth over a matter of weeks. However, the technical details are still of interest. You can find a list of news articles here: <https://github.com/iamcryptoki/snowden-archive>, and the complete document archive here: <https://search.edwardsnowden.com>. Thus, the first part of the essay is:

1. *Discuss some of the technical capabilities uncovered in the Snowden revelations, give technical details of how these were carried out, and discuss the sorts of technically-trained people who may have done this work.*

From this treasure trove of information spawned many discussions, and many pieces of serious academic work. An early discussion piece was written by Tom Leinster in 2014 titled “Maths spying: The quandary of working for the spooks”, where he begins to discuss the merits and drawbacks of the involvement of mathematicians in such work. So the second part of the essay is:

2. *Analyse the points that Leinster raises in his piece, look at who else is making these points, and see if you can document a “community” of vocal objectors and their specific objections or arguments.*

A much more detailed academic paper written on this was the (now very well-known) work by Phillip Rogaway titled “The Moral Character of Cryptographic Work” in late 2015. In it, he gives a history and overview of the role of cryptography and cryptographers in society; past, present, and future. This is a long piece, and covers many sub-points and case studies. So the third part of the essay is:

3. *Read through the piece by Rogaway and summarise his main arguments, both technical and non-technical. Perhaps choose one or two particular case studies that he mentions, and investigate them further.*

If choosing a case study from the piece by Phillip seems daunting, here is one I have pre-selected that involves a very technical academic paper by David Adrian *et al.* from 2015 titled “Imperfect forward secrecy: how Diffie-Hellman fails in practice”.

4. *Discuss the technical aspects covered in the piece by Adrian *et al.*, as well as their practical consequences.*

And finally, see if you can find out who the following quote is from: “*The laws of mathematics are very commendable, but the only law that applies in Australia is the law of Australia.*”

Remember, you are free to pick and choose whichever aspects of the essay interest you the most.

3. ESSAY 3: A TOPIC OF YOUR CHOOSING!

Yes, that's right: you get to choose your own topic! This is a horribly double-edged sword, and not for the faint-hearted. You can choose any topic or point relating to Ethics in Mathematics that was covered in the Michaelmas 2019 lectures, or even think up your own topic to investigate. It really is up to you. Just make sure you run it past me *before* commencing; email mcc56@cam.ac.uk .

For those who wish to try this route, but are struggling to think up (or recall) relevant topics, please email me, and I'll send you a set of the pdf course notes from the Michaelmas 2019 lectures (approximately 4-10 pages of text per lecture).

There is a course given at MIT on "The Ethics and Governance of Artificial Intelligence", and they have a very comprehensive reading list which might give you some starting points. It can be found here:

<https://www.media.mit.edu/courses/the-ethics-and-governance-of-artificial-intelligence/>

There is also the now-popular (though somewhat devoid of technical detail) book by Cathy O'Neil titled "Weapons of math destruction", which you can read through and look for inspiration.

Whatever topic you choose, be sure to keep it to ethics in *mathematics*, as it can be very easy to wander in to other neighbouring disciplines such as computer science, engineering and physics. I cannot, however, stress the following strongly enough:

Be sure to run your proposal past me before commencing!