21st century philosophy

An introduction to information ethics

Artificial intelligence and information technology are part of our everyday lives; from the way we carry out our work, to the ways in which we spend our leisure time or communicate with others. It is becoming increasingly important then that we understand not just the way that this technology works and how it is developed, but also how our individual and collective thoughts might be influenced by and even manipulated through it. In this article, Dr Brian Ball of New College of the Humanities, London explains how he and his colleagues are tackling this from a philosophical perspective, and describes his aim to build capacity in information ethics and the philosophy of information technology (IT) and artificial intelligence (AI) through both research and teaching.

Information and information technologies are increasingly essential to our social, economic, and political interactions. Given this, serious reflection on information ethics is imperative.

Head of Faculty in Philosophy at New College of the Humanities, London, Brian Ball works primarily in the philosophy of mind and language, and the theory of knowledge.

I n 2021 we are living in the 'Age of Information'; information technology (IT) and artificial intelligence (AI) have become an integral part of society, economics, politics and culture. This rise of IT and AI has led to a new branch of philosophy: information ethics.

Research Outreach caught up with Dr Brian Ball of New College of the Humanities (NCH), London, who believes that a good understanding of information ethics is critical in order for us to navigate our new technologically-driven world, and that a poor understanding, individually, collectively, and institutionally could be a threat to democracy.

Please tell us about yourself. What sparked your interest in philosophical research?
When I went to university as an undergraduate my plan was to study the natural sciences, biology specifically, but I also took an introduction to philosophy course in the first term. What I really liked about that course, and what compelled me to switch to a philosophy major, was the sense that philosophy was a subject that was both relevant and broad-ranging (it's a humanities discipline) and

yet intellectually rigorous and precise (like the sciences).

Your current research interests, on the face of it, seem quite diverse, but there is an underlying theme there of 'intentionality'. Can you tell us more?
I conduct research in the philosophy of language, the philosophy of mind, and the theory of knowledge (a.k.a. epistemology). I've written about speech act theory, defending the view that knowledge is the norm of assertion, that one ought not to assert something unless one knows it. I've also written about the mental analogue of assertion, namely judgement. I've written about the format of representation in certain cognitive processes, particularly around numerical cognition, and about belief, and talk about belief - so-called 'attitude ascriptions'.

All these topics are concerned with aspects of the world that are themselves about something. If I intend to bring about some situation, then my intention is about that future situation, but the medieval term that philosophers use for this, 'intentionality', extends beyond just intentions. All of these cases - assertion, judgement, cognition, belief, knowledge - they all exhibit this 'intentionality' or

'aboutness'. If I know that Paris is the capital of France, my knowledge is about Paris. If I tell you that Paris is the capital of France, my assertion is about Paris. And similarly for judgment. So, in each of these cases what we have are some acts, processes or states, all of which are about something else. Philosophers like to use this medieval jargon, but in contemporary terms, we might talk about 'information' - one kind of thing carries information about another. Indeed, the philosophy of information emerged just this century - and I'm interested in this, and in related issues in the philosophy of artificial intelligence, which involves the computational manipulation of information.

Moving to your more recent work, your 'Defeating Fake News' paper, you put forward the argument that fake news poses a threat to democracy. Can you please tell us more about this research?
As I said, I work in the philosophy of language and in the theory of knowledge. One area where they intersect is the epistemology of testimony - basically, how knowledge and other information can be transmitted, maybe even generated, through people telling each other

things. When I was last on research leave in 2018, the Cambridge Analytica scandal was breaking. It seemed to me that there are two elements to that scandal. One is that there is a massive privacy violation. The other is the manipulation of voting behaviour through the use of social media for the dissemination of information, some of which might be misinformation and psychological manipulation as well. I found myself asking what philosophers might have to say about this.

In the paper you mention, I was trying to understand the phenomenon through the lens of the theory of knowledge, and using that understanding to explain how and why there's a danger to democratic society. The central thought is that fake news and other forms of misinformation online undermine the ability of democratic societies to be informed, so as to take appropriate action through our elected representatives. In particular, they undermine the ability of journalistic outlets to transmit knowledge on matters that are pertinent to members of democratic societies. I thought, we need to defeat fake news because of the threat it poses to democracy, but also because, if we

People talk about the Fourth Industrial Revolution; AI is everywhere now - in our professional lives, in our homes, our cities, and so on.

Brian Ball aims to establish an interdisciplinary research centre at NCH, engaging with issues of broad intellectual interest surrounding the idea of information.

Jozsef Bagota/Shutterstock.com

don't, it will defeat our knowledge, and in particular the transmission of knowledge by journalistic channels. Effectively, the presence of fake news on the internet and its dissemination on social media provides grounds for consumers of information to doubt genuine reports of actual events, and if you doubt something, that's the opposite of believing it - and you can't know anything unless you at least believe it.

In that paper, I argue that fake news and related communicative phenomena online and particularly on social media can lead to problems for democracy by undermining knowledge. This has two consequences. One is, we might just get worse outcomes; but then there's also a kind of meta-reflection - if people start to think that they're getting worse outcomes through democracy, they might disvalue democracy itself. I also wanted to stress that we might look to structural causes here rather than individual cognitive failings - we need to make sure that our informational environment is not polluted. In

particular, we need to think about what responsibilities might be imposed on social media platforms.

The 'Defeating Fake News' paper is part of your broader research initiative on information ethics. Could you tell us a bit more about that broader research initiative?
'Information ethics' is concerned with various permissions and obligations that we might have around the dissemination and use of information. There are a range of legal obligations and permissions that we might have (e.g. under the General Data Protection Regulation), but there are also moral ones (should I repeat this thing you told me in confidence?), and even epistemic ones - ought I to believe what you're telling me? You're presenting me with some information - should I accept it? The way I approach information ethics is as an aspect of applied, social epistemology, focusing on questions like these last ones, but in the context of larger social groups. How does

information flow in social networks? How does this affect our beliefs and our knowledge, both on an individual basis, and collectively? And related issues arise in specific ways in the digital sphere. What should we do about e.g. the way misinformation flows on social media networks? I'm working with some graduate student research assistants now, trying to address some of these issues.

Together with Ron Sandler, Director of the Ethics Institute at Northeastern University in Boston, you established a transatlantic centre devoted to investigating information ethics - despite the challenging conditions resulting from the coronavirus pandemic. What have you achieved with this?
Ron has been doing some really interesting work at the Ethics Institute – he and his colleague John Basl

published a report for Accenture on how companies can build AI and Data Ethics Committees that are fit for purpose; and he is now working on the ethics of content labelling with a great team of researchers. Ron and I successfully applied for some internal funding a little while ago, so I went to Boston with a couple of colleagues from NCH - Paula Boddington, who has been involved with AI4People, and written an important textbook on AI ethics; and Ioannis Votsis, a philosopher of science who has published on AI and the philosophy of information. We met a number of Ron's collaborators, and made valuable connections. We were going to put on an information ethics event in London, but unfortunately, we had to cancel because the pandemic broke out. What we've done instead this academic year is institute an online seminar series, The Information Ethics Roundtable, which is co-hosted between the Ethics Institute and myself at NCH.

With your colleagues at NCH you are building interdisciplinary programmes to enable students to research the issues that you yourself explore in

information ethics. Could you tell us more?
I think it's critical societally for people to be equipped to deal with these intellectual and practical problems in information ethics. People talk about the Fourth Industrial Revolution; AI is everywhere now - in our professional lives, in our homes, our cities, and so on. It's really important that we think carefully about its deployment - both collectively, as a democratic society, but also within companies that are developing these systems just to make sure they anticipate foreseeable issues and navigate around them.

So at NCH we've launched minor programmes in psychology and in data science - students can take these alongside majors in other disciplines, such as philosophy. We also launched

an MA in Philosophy and Artificial Intelligence. We have courses in AI and Data Ethics, and one on Minds and Machines, which explores issues around the nature of intelligence, both natural and artificial. Those are both in the Philosophy Faculty, but we offer the chance for students on the programmes to learn some programming and data science techniques as well so that they can see what the technology looks like and how ethical issues might arise in the process of producing an AI or IT system. In September, we're also launching an MSc in Artificial Intelligence with a Human Face. We feel it's quite important to make sure that computer science students, as well as humanists, get the right exposure to thinking about the ethical and theoretical issues that arise around AI.

You are passionate about building capacity. What are your future aspirations for capacity development at NCH?

There are two aspects of this: research capacity; and teaching capacity. I just laid out what some of our new programmes are that I've been pushing through and that I'm excited about. In general, having a cluster of people at all levels of academic pursuit focusing on these issues is going to be valuable. By building up a base of students, students we can work with on research projects (like I'm doing now in information ethics), we can make progress on these pressing issues for society. NCH is now part of Northeastern University's global network. Northeastern has a number of research institutes - for instance, the Ethics Institute that I mentioned, and the Network Science Institute - and we're looking to build up clusters of researchers here in London that can interact with people over in Boston

The central thought is that fake news and other forms of misinformation online undermine the ability of democratic societies to be informed...

and beyond, within the network and of course in other parts of academia here in the UK and Europe, and beyond academia as well. I would love to see a research centre focusing on the study of information and its ethics here in London at NCH.

New College of the Humanities