

15 FEBRUARY -
22 MARCH 2019

**GROUND
ZERO
E@RTH**

OLIVIA DOMINGOS
BOB BICKNELL-KNIGHT
DANIEL SEAN KELLY
DAVID LISSER
JILLIAN MAYER

#groundzeroearth

Exhibition continues on upstairs floors

INTRODUCTION

Existential Risks are risks which threaten the existence of the human species. The threat of extinction has always been with us (seventy thousand years ago the Toba volcanic eruption may have reduced the total human population to between three and ten thousand individuals), but from the 20th Century onwards, increasingly powerful technologies have made it a very real possibility that we will bring about our own destruction. Many existential risks are low probability, but very high impact. That is, unlikely as their occurrence may be, their capacity to threaten all of us makes understanding them urgent. Researchers including Cambridge's Centre for the Study of Existential Risk (CSER) work to identify these risks, their likelihoods, and how to mitigate them. This raises questions concerning human value: why, exactly, would human extinction be a bad thing? And raises questions about justice: since we all stand to pay the price of technological existential crisis, shouldn't we all stand to benefit from these technologies?

This exhibition brings together five artists who question, provoke and explore our relationships with technology, our environments and the future, and how these relate to our humanity. In a world of increasing technological power, and multiplying existential risks, art's ability to help understand those relationships, and provoke dialogue about them, could turn out to be a critical component of our toolkit for survival.

Our increasingly powerful technologies are a double-edged sword. On the one hand, risks in the face of which we were previously helpless—earthquakes, tsunamis, volcanoes, and so on—are increasingly well understood, predictable and (perhaps one day) controllable. On the other hand, for the first time ever *Homo sapiens* have the technological prowess required to affect the Earth (and ourselves) at an enormous—potentially disastrous—scale. Our rising technological power has not yet been accompanied by a parallel

increase in the wisdom required to understand the consequences of use, and to design the safe-guards, social organization, and so forth, required to avoid unintended disasters (or, ominously, malicious use).

Founded in 2012, CSER is a multidisciplinary research centre within CRASSH at the University of Cambridge, dedicated to the study and mitigation of existential risks. Part of CSER's work involves asking after existential risks themselves: what are they? How probable are they? Other work involves examining calamities in the past and trying to learn from them. CSER also strives to bring together the people with the relevant expertise and influence to make a meaningful difference to our chances of surviving the next century. This exhibition explores the themes and questions arising from CSER (and other existential risk researchers') work.

The philosopher Nick Bostrom defines existential risk as 'an event that could cause human extinction or permanently and drastically curtail humanity's potential.' An existential risk threatens either to wipe us all out, or to damn us to mediocrity. Global catastrophic risks are threats to human wellbeing on a global scale, although not necessarily rising to the level of existential risks. We can bounce back from a global catastrophe, while existential catastrophes are permanent and pan-generational. Existential and catastrophic risks include increasingly familiar disasters like those associated with global climate change, as well as less familiar worries, such as the rise of artificial super-intelligence.

Existential risk researchers cover a lot of ground: philosophical, as well as scientific, legal, political and regulatory. In discussing these challenges, Martin Rees compares the scientific ingenuity of Einstein's early work, and the humanistic reflection of Einstein's later life. Both, he thinks, will be required:

The science done by the young Einstein will continue as long as our civilization, but for civilization to survive, we'll need the wisdom of the old Einstein – humane, global and farseeing. And whatever happens in this uniquely crucial century will resonate into the remote future and perhaps far beyond the Earth, far beyond the Earth.

GROUND ZERO

Threats to human extinction come in many shapes and sizes. Some emerge from natural events, often recorded in earth's history: extra-terrestrial impacts, super-volcanic eruptions, extreme solar flares, and oceanic acidification have led to extinctions (even mass extinctions) in the deep past. Others are more anthropogenic: nuclear war and the consequences of extreme climate change are more familiar, but existential risk researchers also worry about risks from bio- and nano-technology, automated weaponry, and artificial intelligence. As technology gets more powerful and cheaper, the consequences of mistakes and the potential capacities of malicious agents increase as well. We might also worry about 'perfect storms' of smaller risks, or less dramatic threats, which might nonetheless have devastating cumulative effects.

'Ground zero' is a term used to describe the nearest point on the earth's surface to detonation or disaster: a nuclear explosion, the centre of an earthquake's destruction, or an epidemic's point of origin. Existential risk researchers are concerned that we are fast approaching ground zero on climate change and runaway technology. Concerning catastrophic risks, ground zero is the point of no return: the point at which it's too late to avoid horrifying loss of life. For existential risks, ground zero is game over.

'Ground zero' also highlights how we conceive of calamitous risks: as powerful, sudden, unexpected 'one-shots' like nuclear war, asteroid impacts, and so forth. But focusing too much on ground zero—on sudden disasters—might lead us to miss slower and more subtle dangers. Creeping poverty, increasing populations, or pollution could slowly, unbeknownst to us, build to catastrophe. These are what Liu, Lauta & Maas have called 'boring apocalypses', or Karen Kuhleman calls 'unsexy' risks: these are distributed, diffuse and gradual risks without identifiable ground zeros.

CSER also studies 'cascading effects': when the impact of a physical event generates a sequence of events in human subsystems that result in physical, social, or economic disruption. The trigger or domino effect potentially causes significant downstream disruptions. Cascading effects and boring apocalypses call to mind T.S Eliot's oft-quoted conclusion to *The Hollow Men*:

*This is the way the world ends
Not with a bang but a whimper*

Avoiding both bangs and whimpers requires understanding existential risks. Identifying risks is the first step of mitigation. Following this, we must understand the likelihood of occurrence, its likely impact, and so forth. CSER often focuses on low-probability but high-impact events: it might be unlikely that, say, our powerful technologies carry unforeseen risks at the existential level, but the consequences, should such events occur, are so extreme—human extinction—that their low probability doesn't prevent us from caring about them deeply.

Thus far, we've avoided existential catastrophe, but new technologies mean our choices have effects, detrimental or otherwise, on much larger scales. They cannot be limited to traditional national borders, and cannot be managed in isolation from each other. Collaborative approaches and responsibility, as well as a significantly longer-term perspective, are required to understand and mitigate existential risk.

Studying existential risk doesn't require being anti-technology, or prophesizing doom. Rather, with our new technological powers come new responsibilities: if we are to reap the benefits of these technologies, we also need to understand the dangers they bring. At least sometimes, studying existential risk is an optimistic pursuit: if we are as wise as we are smart, and are careful to manage the risks of our new technologies, we may be able to harness their power in positive, transformative ways.

The study of existential and catastrophic risk should not be deemed depressing or negative: we believe that if we act quickly and carefully enough, threats to humanity can be minimized. CSER and other researchers of catastrophic and existential risk do not see their work as inherently pessimistic. Rather, the central aim is to enable us to use the potentially transformative—and transformative for the better—capacities of our new technologies, while minimizing exposure to the disastrous tail of the probability bell-curve. CSER's co-director Seán Ó hÉigearthaigh describes the existential risk community as an 'insurance policy for a society developing more and more powerful technologies'. As he told the *The National*:

I sometimes joke that if we do our job correctly, you'll never know we did anything because what we will have done is reduce the possibility of something from 5 per cent to 0.005 per cent...

PERSPECTIVE


Art can communicate, challenge received ideas, and confront the past and future. It also facilitates understanding of the here-and-now. For existential risk, this last capacity is perhaps the most important. Because much of our new technology is unprecedented and transformative, our ability (whether scientific or artistic) to predict or imagine the future is extremely limited. Media tends to explore dystopias and utopias: exaggerated, extreme future societies. When we consider the fictional dystopias and utopias produced in the past, we find that these are more a reflection of the hopes and fears of the time than any real prescience on what the future will be like. But understanding and mitigating existential risks requires understanding ourselves, and here is where art's capacity to be reflexive, critical, and revelatory becomes valuable. Facing up to the destructive potential of our new technology doesn't simply require predicting our future, but understanding our present. These together might provide a rough map of our pathway forwards, or at least the wisdom required to avoid some of the traps.

There has been literature, philosophy and art about the future since there has been literature, philosophy and art. The 4,000 year old Indian Vedas conceive of time cyclically, a tradition carrying through to contemporary Tibetan Buddhism's *Kalachakra* or "wheel of time". Ancient Greek philosophers debated time's existence: in the 5th Century BCE, Antiphon argued that time was a human-imposed measure, and Parmenides believed that past and future were an illusion. Heraclitus disagreed, arguing that time truly existed and flowed. Again in contrast, Epicurean philosophers argued that the future was nothing to fear, as it didn't exist. Medieval Christians were often obsessed with apocalypse, inspired in part by the Jewish prophets of the Old Testament and the New Testament's dramatic closing act of John's Revelation.

Partly prompted by booms in economic growth and technological advancement, and the unprecedented transformations these inspired (in wealthier societies at any rate), the 20th century was—perhaps more so than previous eras—particularly concerned with visions of the future, from the utopian visions of 1950s and '60s United States of America and their Soviet counterparts, to the dystopian warnings of environmental calamity inspired by, among others, Rachel Carson's *Silent Spring*.

It may be that dystopias and utopias tell us as much (if not more) about the people who produced them, their fears and hopes, than they tell us about their potential futures. Flying cars haven't eventuated, but the idea that technology will free individuals from mundane chores tells us a lot about how the West

valued technology in the 1950s and '60s (to say nothing of the deeply ingrained gender roles). No wonder these utopian ideas became increasingly dark and dystopian in light of the deep anxiety of the cold war. Looking back from 2019's vantage point, were previous prophecies and warnings over-exaggerated, merely reflecting the anxieties of the time, or are we now fulfilling those projections?



YOUR PERSONAL "FLYING CARPET" Step into it, press a button, and off you go to market, to a friend's home, or to your job. Take off and land anywhere, no parking problems. Plug in to any electric outlet for recharging. They're working on it!

MORE POWER TO YOU!

America's independent light and power companies build for your new electric living

Tomorrow's higher standard of living will put electricity to work for you in ways still unheard of!

The time isn't too far off, the experts say, when you'll wash your dishes without soap or water—ultrasonic waves will do the job. Your beds will be made at the touch of a button. The kid's homework

will be made interesting and even exciting when they are able to dial a library book, a lecture or a classroom demonstration right into your home—with sound. (Some of this is happening already.)

To enjoy all this, you'll want a lot more electric power, and the independent electric companies of America are already building

new plants and facilities to provide it. Right now these companies are building at the rate of \$5,000,000,000 a year, and planning to double the nation's supply of electricity in less than 10 years.

America has always had the best electric power service in the world. The electric companies are resolved to keep it that way.

AMERICA'S INDEPENDENT ELECTRIC LIGHT AND POWER COMPANIES

Art and media often respond to contemporary life and society, sometimes employing stereotypes and exaggerations to highlight and communicate themes and ideas. However, knowledge gaps between the sciences and the arts, as well as the aesthetic and communicative constraints of various media, may lead to the misrepresentation of particular concepts. In projecting an image of present-day or future crises, how likely is such misinterpretation? And how might certain artistic approaches (be they mass media or more local artistic projects) contribute to this misinterpretation and distortion on the one hand, or resolve it on the other? Such biases in communication can have consequences. The philosopher Adrian Currie has warned that a science of existential risk must be speculative and deal with uncertainty, and there are dangers in how such subject matter is communicated: a balance must be struck between urgency and importance, the low probability of occurrence, and the spectacular, apocalyptic nature of the subject matter.

On the one hand, the public or policy makers might take the science too seriously, and act rashly in the light of that. But on the other hand, repeated potential 'failures' could lead to a loss of faith in the science.

This raises the question: what is art about the future created for? 'Communication' alone is far too narrow, and 'prediction' seems way too difficult. Given that art cannot help but be influenced by and reflect the contemporary world and tends to exaggerate and focus on extremes for effect (art is full of dystopia and utopia, but dystopias and utopias seem so rarely to eventuate), is it a good source of ideas about the future? Perhaps it is to some extent: art doesn't simply exaggerate and represent, it transforms,

distorts and questions. More critically, perhaps, art is a means of conversation, of raising questions, challenging perspectives and helping us understand ourselves. A major part of understanding existential and catastrophic risk is understanding ourselves, our own limitations and foibles. Insofar as art helps us to understand ourselves—and to question our own assumptions and values—it is critical for approaching the deep future and the risks it might bring.

BEING HUMAN

Understanding existential risks isn't simply a scientific and technological undertaking alone: it requires reflection on what it is to be human. First, exactly what value should we place on human lives? Our extinction wouldn't simply harm ourselves, but close off the possible existence of many future generations. Should we (and if so, how can we) take these not-yet-actual people into account? Second, new technologies (artificial intelligence and increasingly flexible, sophisticated automation most directly) threaten to outperform humans in many professional contexts. How should we understand legal responsibility if our health, for instance, is in the care of an artificial agent? Third, if we want to use our new technology in positive, transformative ways, what are those ways? That is, what does a flourishing life look like, and how could new technology help us to pursue it? Fourth and finally, if we are to take on existential risks, who should decide on actions, and through what processes? And how should the benefits of new technology be distributed?

Researching and mitigating existential risks requires enquiry into fundamental questions about human nature and value. Sometimes questions about existential risk are questions about what it means to be human. What would it take for our species to become extinct, and why would extinction be so bad? What value should we place on our current needs, as opposed to those of future generations? Given that existential risks are everyone's risks, is it fair if the benefits of new technologies are unevenly distributed? Given that the consequences of an existential calamity are so enormous, who should be responsible for making decisions concerning them? What risks are worth taking, and who is in a position to decide? Many of these questions play out in the rapidly developing field of artificial intelligence.

Artificial intelligence is the capability for machines to recognize and respond to their environment, as well as demonstrate human-like abilities in learning and problem solving. There is a lot of debate about the potential for artificial intelligence becoming human-like, human-level, or super-human. Although

machine-learning algorithms increasingly outperform us at explicitly defined tasks (playing chess and Go, for instance), it is not clear whether ramping up the capacity for task-oriented artificial intelligence is enough to achieve general artificial intelligence. There is a further question about why we would want something like this: what would we learn from engineering a general intelligence, what would it be able to do for us? Moreover, should we consider that intelligence to 'be like us'—with moral status and rights? Can we, and should we, form meaningful relationships with it? Finally, if AI is to have a general, or super-human intelligence, presumably that AI will need to have values: how should they decide what to protect, what to encourage and what to discourage? As Anca Dragan has put it:

Robots aren't going to try to revolt against humanity... they'll just try to optimize whatever we tell them to do. So we need to make sure to tell them to optimize for the world we actually want.

But what world do we want? And given that human values are diverse and complex, which 'we' are we talking about? Artificial intelligence isn't the only technology which challenges conceptions of humanity.

The transhumanist movement aims to use scientific knowledge and new technologies to expand human life-spans and augment our intelligence and physical capacities: perhaps super-intelligence won't be a revved-up super-computer, but a part-human, part-technology cyborg. And this raises a strange existential risk: if our technologically-augmented evolution means we leave our humanity behind, would that count as the extinction of our species? More generally, what is it that we value about our humanity, and how might we retain—or even increase—that value in an increasingly technologized future? Our new technology offers a great future, but what are we willing to give up for it?

Our new technologies, and the risks they raise, are challenging: they are transforming human institutions and how we organize ourselves at an alarming rate. Notions like privacy, ownership, and work are changing at speeds which outpace the capacity of society to safely adapt. And there is little sign that the

world is becoming less hierarchical: the gap between the haves and the have-nots is widening not closing. Facing up to existential risk requires facing up to the relationship between democratic justice and decision-making, and emerging technologies.

The pace of technological change is accelerating such that our usual approaches to governance and the organization of human worlds is struggling to keep up. The new 'industrial revolution' ushered by advancing AI affects jobs that have traditionally been filled by highly-trained, well-paid human labour: teachers, lawyers, accountants, doctors, and many others may become increasingly unnecessary; replaced not by automation, but by intelligence (just of an artificial kind). Meanwhile, increasingly powerful digital technology can be used to replicate human speech and modify videos, contributing to the break-down in trust between the public and the media they consumed, which characterised the 20th

Century, and underwrote the effectiveness of the '4th estate'. This misinformation is spread at super-virus-like speeds along the new channels of information flow represented by social media—new companies who have access to unprecedented amounts of data about our everyday lives, preferences and behaviours. Legal conceptions of privacy and free-speech are breaking down in the light of new technology to collect, manipulate and analyse data, and to share it.

These changes are not necessarily good or bad—in a complex world consequences are almost always mixed—but it does challenge us to imagine new ways in which our society and institutions might be organized in order to retain and amplify what we value, protect what needs protecting, and lower the risks of the super bad, the catastrophic and existential outcomes.

CONCLUSION

Martin Rees has argued that this century—the 21st—will be critically important for our species.

This century is crucial because if you're very pessimistic, you can imagine that we will misuse powerful technology and snuff ourselves out or foreclose a bright, longer-term future. On the other hand, if we use technology wisely, then it allows us to perhaps jump-start an even more exciting kind of civilization here on Earth and far beyond. That's why even though the Earth has existed for 45 million centuries, and will go on existing for many million more centuries, this century is special.

Where are we heading? Is the power of collective change, corporate social responsibility and political influence enough to mitigate the immediate threats which lie ahead? Will we continue to change the landscape of our fragile earth beyond repair, or is there hope that we can ensure our survival?

Can we escape, to Mars perhaps? Such a courageous human endeavour may be possible in our lifetime and is being explored by those wealthy enough to fund such missions. However, it is worth remembering that considerations of existential risk are also considerations of human value. We might ask whether a species who sacrifices billions of its member in order to save a wealthy few is a species worth caring for. Or, for that matter, what would stop those colonists from inheriting the same unsustainable practices in their new worlds? Some argue that our resources could be better utilised in resolving the issues on our home planet before we start a new chapter on distant planets.

There is only so much we can achieve as individuals. Catastrophic and existential risks are global risks—they belong to all of us—and as such require collective, often regulatory and often international measures. As members of the human species (and the crew of spacecraft earth) it is our collective responsibility to ask what risks we're willing to take, what values we want to express and what comes at too high a cost.

PROGRAMME

14th February 2019, 5.30pm

Panel Discussion

Alison Richard Building (ARB) Large Seminar Room

13th March 2019, 4.30pm

Curator Tour

Alison Richard Building Foyer

14th March 2019, 6pm

Rise of the Machines, Mini-Film Screening with Q&A

Hosted by Yasmine Rix (CSER Research Affiliate, Curator) and Beth Singler (Faculty of Divinity, Junior Research Fellow in Artificial Intelligence) watch a short documentary series on the social, ethical, philosophical, and religious implications of advancements in artificial intelligence and robotics, with 'Pain in the Machine' winning 2017 AHRC Best Research Film of the Year award.

Alison Richard Building, Large Seminar Room

Please book your free ticket for all events on Eventbrite

ARTISTS

Olivia Domingos (b.1991) is an artist and illustrator based in London. Exploring installation and drawing, she has produced works with a focus on celebrity culture and challenging the public domain of wellbeing. Her rendering of specific events and their detail brings attention to falling victim to voyeurism of celebrities or news sensationalism.

Bob Bicknell-Knight (b. 1996, Suffolk) is a London-based artist and curator working in installation, sculpture, video and digital media. Using found objects and tools made readily available by the Internet, as well as drawing from a unique sensibility influenced by participation in online communities and virtual games, Bicknell-Knight's work explores the divergent methods by which consumer capitalist culture permeates both online and offline society.

Daniel Sean Kelly (b.1989, Leicester) is an artist and co-director of Two Queens artist led gallery and studios. Working largely in painting, printmaking and ceramics, his work seeks to create a speculative space for the imagining of other realities – a science fictional universe comprised only of objects existing in the world up to this point.

David Lisser (b. Wolverhampton 1987) is an artist based in Newcastle who investigates our relationship with food and emerging technologies, playfully creating artefacts excavated from an imagined past, documentation of protests that haven't yet materialised, and mechanisms for producing novelty meats.

Jillian Mayer (b. 1986) is an artist and filmmaker living in Miami, Florida. Through video, sculptures, online experiences, photography, performances and installations, she explores how technology affects our lives, bodies and identities. Mayer investigates the points of tension between our online and physical worlds and makes work that attempts to inhabit the increasingly porous boundary between the two.

Please make any enquiries about purchasing of art works to jw571@cam.ac.uk

Acknowledgement and thanks to the artists involved, staff members at the Centre for Study of Existential Risk, ARTatARB, CRASSH and our sponsors including Cambridge BID and Arts Council England.

Curated by Yasmine Rix



Supported using public funding by
**ARTS COUNCIL
ENGLAND**