



Stockholm Resilience Centre
Sustainability Science for Biosphere Stewardship



Human-Machine-Ecology: A Workshop on the Emerging Risks, Opportunities, and Governance of Artificial Intelligence

Princeton University
Friday-Saturday, January 11-12
Survey Responses
(As of January 8, 2019)

Workshop Participants

Chloé Bakalar	Temple University – Assistant Professor, Political Science; Princeton University – Visiting Research Collaborator, CITP
Solon Barocas	Cornell University – Assistant Professor, Information Science
Seth Baum	Global Catastrophic Risk Institute – Executive Director
Irina Brass	University College London – Lecturer in Regulation, Innovation and Public Policy, Department of Science, Technology, Engineering and Public Policy (STeAPP)
Sarah Brayne	University of Texas at Austin – Assistant Professor, Sociology
Joanna Bryson	University of Bath – Reader, Computer Science
Peter Callahan	Princeton University – Research Assistant, PIIRS Global Systemic Risks research community
Amar Causevic	Global Economic Dynamics and the Biosphere Programme – Research Assistant
Miguel Centeno	Princeton University – Director, PIIRS Global Systemic Risks research community
Darryl Farber	Penn State University – Assistant Teaching Professor, School of Engineering Design
Edward Felten	Princeton University – Directory, CITP
Danit Gal	Keio University, Tokyo – Project Assistant Professor, Cyber Civilization Research Center
Victor Galaz	Stockholm University – Stockholm Resilience Centre
David Garcia	Complexity Science Hub Vienna – Faculty
Daniel Jiménez R	International Center for Tropical Agriculture (CIAT)

Brian King	CGIAR – Coordinator of the Platform for Big Data in Agriculture (Co-led by CIAT and the International Food Policy Research Institute)
Paul Larcey	University of Cambridge – Research Fellow, Engineering
Karen Levy	Cornell University – Assistant Professor, Information Science
Arka Mukherjee	Global IDs Princeton University – PIIRS Global Systemic Risks research community
Thayer Patterson	Princeton University – Research Fellow, PIIRS Global Systemic Risks research community
Jennifer Rexford	Princeton University – Professor, Computer Science,
Jonathan Tannen	Facebook – Research Scientist
Michael Veale	University College London – PhD candidate, Department of Science, Technology, Engineering & Public Policy University of Birmingham – Honorary Research Fellow, Birmingham Law School
Ben Zevenbergen	Princeton University – CITP

Survey Responses

Participant	Chloé Bakalar
Affiliation	Temple University – Assistant Professor, Political Science; Princeton University – Visiting Research Collaborator, CITP
Email	cbakalar@princeton.edu
Primary Field	Politics
Brief Bio	Chloé Bakalar is Assistant Professor of Political Science. She is also a Visiting Research Collaborator at Princeton University's Center for Information Technology Policy (CITP). She is a political and legal philosopher with a background in empirical social science. Professor Bakalar received her Ph.D. from the University of Pennsylvania and her B.A. from New York University. Before joining Temple, she was a Senior Research Specialist at CITP and completed the Values and Public Policy Postdoc at Princeton's University Center for Human Values (UCHV) and the Center for the Study of Democratic Politics (CSDP).
Description of Research	Professor Bakalar is a member of the Steering Committee for the Princeton Dialogues on AI and Ethics. In this capacity, she has authored and co-authored a series of educational use case studies in the emerging field of AI. Broadly, her research focuses on philosophical and legal questions surrounding freedom of speech, especially in terms of the relationship between information technologies and liberal democratic citizenship. Additional research interests include: ethics and public policy (esp. technology ethics); normative ethics; democratic theory; American political thought; modern political thought; contemporary political thought; the history of ideas; constitutional law; and civil rights/civil liberties (esp. First Amendment). She is currently completing a book, "Small Talk: The Impact of Social Speech on Liberal Democratic Citizenship," which examines neglected concerns for the everyday in political theory and Anglo-American public law, and introduces a framework for understanding how everyday talk (i.e., “social speech”), especially online, influences liberal democratic citizenship and political outcomes in both positive and negative ways.
Website	- https://aiethics.princeton.edu - https://liberalarts.temple.edu/academics/faculty/bakalar-chlo
Recommended General Readings	Frischmann and Selinger, Re-Engineering Humanity; Gillespie, Custodians of the Internet; Nissenbaum, Privacy in Context
Participant	Solon Barocas
Affiliation	Cornell University – Assistant Professor, Information Science
Email	sbarocas@cornell.edu
Primary Field	Information Science
Brief Bio	Solon Barocas is an assistant professor in the Department of Information Science at Cornell University. He is also a faculty associate at the Berkman Klein Center for Internet & Society at Harvard University. He co-founded the annual workshop on Fairness, Accountability, and Transparency in Machine Learning (FAT/ML) and later established the ACM conference on Fairness, Accountability, and Transparency (FAT*). Previously, he was a postdoctoral researcher at Microsoft Research, working with the Fairness, Accountability, Transparency, and Ethics in Artificial Intelligence Group, as well as a postdoctoral researcher at the Center for Information Technology Policy at Princeton University.
Description of Research	Solon's current research explores ethical and policy issues in artificial intelligence, particularly fairness in machine learning, methods for bringing accountability to automated decision-making, and the privacy implications of inference.

Website	http://solon.barocas.org/
Case Studies	Artificial intelligence has begun to play an important role in the long-standing movement toward so-called "precision agriculture," a largely data-driven approach to farming which stands to reshape the everyday practice and business of agriculture. Artificial intelligence has begun to inform how farmers farm, how they manage their farms, and how they engage with other actors in the larger economy. But these new technologies have been met with skepticism and concern about their likely benefits for farmers.
Case Study Readings	David Schimmelpfennig, "Farm Profits and Adoption of Precision Agriculture," https://www.ers.usda.gov/webdocs/publications/80326/err-217.pdf?v=0
Benefits of these technologies	Far more efficient farming, leading to reduced environmental impact, sustainability, and expenses (both for producers and consumers).
Risks of these technologies	An uneven distribution of the economic benefits of these new technologies.
Recommended General Readings	-Megan Stubbs. "Big Data in U.S. Agriculture," https://fas.org/sgp/crs/misc/R44331.pdf -Government Accountability Office, "Artificial Intelligence: Emerging Opportunities, Challenges, and Implications," https://www.gao.gov/assets/700/690910.pdf

Participant	Seth Baum
Affiliation	Global Catastrophic Risk Institute – Executive Director
Email	seth@gcrinstitute.org
Primary Field	risk analysis
Brief Bio	Dr. Seth Baum is Co-Founder and Executive Director of the Global Catastrophic Risk Institute, a nonprofit and nonpartisan think tank. He leads an interdisciplinary research agenda of risk and policy analysis of catastrophic risks, focusing primarily on artificial intelligence and nuclear war. Baum received degrees in Optics, Applied Mathematics, and Electrical Engineering before completing a Ph.D. in Geography from Pennsylvania State University and a post-doctoral fellowship at Columbia University. He is based in New York City.
Description of Research	Dr. Baum’s overall research focus is on developing effective solutions for reducing the risk of global catastrophe. This work has two specific thrusts. One is social science to understand the social processes that influence the global catastrophic risk. Current research in this thrust include (1) government application of AI for national security purposes and (2) corporate governance of AI, including tensions between social risks and profit motives. The other thrust is quantitative risk and decision analysis for evaluating risks and risk-reduction solutions. This includes attention to the considerable uncertainty inherent to the global catastrophic risks. This research agenda is coupled with outreach to relevant stakeholders and decision-makers in government, industry, and NGOs.
Website	Global Catastrophic Risk Institute: http://gcrinstitute.org Personal website: http://sethbaum.com
Case Studies	While this is not my specific focus, two examples come to mind. One is the increased energy consumption from the computer usage underlying the widespread application of AI. The other is the prospect of advanced future AI gaining enough capability to be a major force on the global environment, for better or for worse.
Case Study Readings	I briefly discussed the environmental implications of advanced future AI in the paper "The great downside dilemma for risky emerging technologies", Physica Scripta, vol. 89, no. 12 (December), article 128004, DOI 10.1088/0031-8949/89/12/128004.

Benefits of these technologies	If designed and applied well, the technologies can help identify means of reducing environmental harms and addressing other societal challenges.
Risks of these technologies	Increased energy consumption; the potential for harm from the AI not behaving in ways we expect or want it to; the potential for AI to be used as a tool by malevolent or antisocial actors
Recommended General Readings	A major theme of my current work on AI is to apply insights from the (much more substantial) study of environmental issues to the study of issues in AI. In particular: -Seth D. Baum, 2017. On the promotion of safe and socially beneficial artificial intelligence. <i>AI & Society</i> , vol. 32, no. 4 (November), pages 543-551, DOI 10.1007/s00146-016-0677-0. -Seth D. Baum, 2018. Superintelligence skepticism as a political tool. <i>Information</i> , vol. 9, no. 9 (August), article 209, DOI 10.3390/info9090209. -Seth D. Baum, 2018. Countering superintelligence misinformation. <i>Information</i> , vol. 9, no. 10 (September), article 244, DOI 10.3390/info9100244.

Participant	Irina Brass
Affiliation	University College London – Lecturer in Regulation, Innovation and Public Policy, Department of Science, Technology, Engineering and Public Policy (STeAPP)
Email	i.brass@ucl.ac.uk
Primary Field	Regulation of emerging technologies; public administration; public policy
Brief Bio	Irina Brass is a Lecturer in Regulation, Innovation and Public Policy at UCL STeAPP. She teaches Public Administration, Risk Assessment and Governance, and is the co-lead of a new MPA in Digital Technologies and Public Policy. Dr. Brass is also Co-Investigator of the Standards, Governance and Policy Stream of the PETRAS Internet of Things Research Hub, as well as a Co-Investigator of the Future Targeted Healthcare Manufacturing Hub looking at new regulatory approaches for targeted healthcare. Dr. Brass has a background in public policy, regulation and standardization, with research that spans across radio spectrum management, intellectual property and competition rules, cybersecurity risk and safety regulations. She holds a PhD in Government from the London School of Economics and Political Science (LSE).
Description of Research	Irina Brass' research focuses on the economic, social and risk regulation of disruptive technologies, especially digital technologies, and the establishment of more dynamic and adaptive modes of governing these technologies. She has published several studies on the misalignment between data protection, cybersecurity and safety regulations for the Internet of Things, the role of technical transnational epistemic communities in governing risks associated with emerging digital technologies, and the implications that digital technologies such as IoT or ML raise for public sector transparency and accountability. She works closely with the UK government and the standards-making community to develop best practices for tackling the risks associated with emerging technologies.
Website	Main web page: http://www.ucl.ac.uk/steapp/people/brass Publications: https://iris.ucl.ac.uk/iris/browse/profile?upi=IIORD26
Case Studies	Agriculture: detecting pests; monitoring soil composition; predicting yield; predicting and adapting to abrupt shifts in weather conditions; parametric insurance. Note that most of these applications are reliant on a combination of IoT (for data gathering and monitoring) and ML/AI (for data analytics and actuation). It is, thus, crucial to understand that we are relying on more than a single digital technology here, but an ecosystem of digital technological solutions. Environment: protecting biodiversity; maritime protection (acidification, biodiversity); clean air (connected autonomous vehicles/ transport); water security and distribution; energy management and diversification.

Case Study Readings	<p>-J. M. V. Nostrand and J. G. Nevius, “Parametric Insurance: Using Objective Measures to Address the Impacts of Natural Disasters and Climate Change,” <i>Environmental Claims Journal</i>, vol. 23, no. 3–4, pp. 227–237, Jul. 2011.</p> <p>-M. U. Scherer, “Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies,” <i>SSRN Electronic Journal</i>, 2015 [Online]. Available: http://www.ssrn.com/abstract=2609777</p> <p>-Nesta UK, “10 principles for public sector use of algorithmic decision making,” <i>nesta</i>, 2018. [Online]. Available: https://www.nesta.org.uk/blog/10-principles-for-public-sector-use-of-algorithmic-decision-making/</p> <p>-L. Tanczer, I. Steenmans, I. Brass, and M. Carr, “Networked World: Risks and Opportunities in the Internet of Things,” <i>Lloyd’s of London and UCL</i>, 2018 [Online]. Available: https://www.lloyds.com/news-and-risk-insight/risk-reports/library/technology/networked-world%20?utm_source=UCL_website&utm_medium=referral&utm_campaign=emergingrisks_networkedworld</p>
Benefits of these technologies	Close to real-time intelligence; tailored solutions (e.g. parametric insurance); distributed governance of resources.
Risks of these technologies	Low transparency and accountability of ML algorithms/ decision models; human bias in data used to train algorithms; easily compromised systems for data collection, monitoring and actuation due to cybersecurity vulnerabilities; data protection issues especially for applications in energy diversification and management, transport, agriculture; lack of standards for assessing and governing organizational risks.
Recommended General Readings	<p>-AI Now Report 2018 https://ainowinstitute.org/AI_Now_2018_Report.pdf</p> <p>-Coglianese, C. and D. Lehr. 2017. ‘Regulating by Robot: Administrative Decision-Making in the Machine-Learning Era’. <i>Georgetown Law Journal</i> 105, pp. 1147. https://papers.ssrn.com/sol3/Papers.cfm?abstract_id=2928293</p> <p>-Royal Society. 2017. <i>Machine Learning: The Power and Promise of Computers that Learn by Example</i>. https://royalsociety.org/~media/policy/projects/machine-learning/publications/machine-learning-report.pdf</p>

Participant	Sarah Brayne
Affiliation	University of Texas at Austin – Assistant Professor, Sociology
Email	sbrayne@utexas.edu
Primary Field	Sociology
Brief Bio	Sarah Brayne is an Assistant Professor of Sociology at The University of Texas at Austin. Prior to joining the faculty at UT-Austin, she was a Postdoctoral Researcher at Microsoft Research. She received her Ph.D. in Sociology and Social Policy from Princeton University. She is a member of the MacArthur Foundation Pretrial Risk Management Project, and the founder of the Texas Prison Education Initiative.
Description of Research	In her research, Brayne uses qualitative and quantitative methods to study how data is used for surveillance within bureaucratic organizations. Her substantive focus is on the criminal justice system. She is currently completing a book on the use of predictive analytics and new surveillance technologies in the Los Angeles Police Department. In previous research, she studied the sociolegal implications of the use of predictive algorithms and surveillant technologies in policing and the courts, and the relationship between individuals' contact with the criminal justice system and their involvement in medical, financial, labor market and educational institutions.
Website	https://liberalarts.utexas.edu/sociology/faculty/sb49337

Case Studies	Modeling and forecasting urbanization, land change, energy use and emissions, biodiversity, and food systems.
Benefits of these technologies	To detect previously unobserved patterns and allocate resources more efficiently.
Risks of these technologies	Bias in the training data can lead to misdiagnosing problems, inaccurate forecasts, misspecified interventions, and biased assessments of treatment effects.
Recommended General Readings	- https://ainowinstitute.org/AI_Now_2018_Report.pdf - https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3257004 - https://journals.sagepub.com/doi/abs/10.1177/0003122417725865?journalCode=asra

Participant	Joanna Bryson
Affiliation	University of Bath – Reader (tenured Associate Professor), Computer Science
Email	jjb@alum.mit.edu
Primary Field	Artificial Intelligence (Psychology)
Brief Bio	Joanna Bryson is a Reader at the University of Bath. She has broad academic interests in the structure and utility of intelligence, both natural and artificial. She is best known for her work in systems AI and AI ethics, both of which she began during her PhD in the 1990s, but she and her colleagues publish broadly, in biology, anthropology, sociology, philosophy, cognitive science, and politics. Current projects include “The Limits of Transparency for Humanoid Robotics” funded by AXA Research, and “Public Goods and Artificial Intelligence” (with Alin Coman of Princeton University’s Department of Psychology and Mark Riedl of Georgia Tech) funded by Princeton’s University Center for Human Values. Other current research includes understanding the causality behind the correlation between wealth inequality and political polarization, generating transparency for AI systems, and research on machine prejudice deriving from human semantics. She holds degrees in Psychology from Chicago and Edinburgh, and in Artificial Intelligence from Edinburgh and MIT.
Description of Research	Joanna Bryson’s natural science research interests concern distributed intelligence: interacting learning systems and their benefits to both individuals & collectives. She began studying modularity in individual intelligence—differences between brain regions in architectural structure, memory and processing—and how animals regulate attention and expression across different goals and actions. This has extended through social learning and culture (the distribution and retention of intelligence) to Dr. Bryson’s current interests in individual & population variation in cognitive & cooperative strategies, including punishment, polarization, and public goods investment. With her PhD students Dr. Bryson actively researches Systems AI—methods for promoting the ease and safety of designing intelligence. Their work includes transparency, action selection, system design, architecture, and integration for real-time systems—robots, smart buildings, virtual reality, game characters—and scientific simulations. Systems AI makes autonomy easier to build, understand, and control. Contributions include the Behavior Oriented Design development methodology, the ABODE real-time AI IDE, POSH action selection, and several systems of synthetic emotions.
Website	http://www.cs.bath.ac.uk/~jjb/
Benefits of these technologies	AI is just a means by which humans and our institutions extend our agency, we can both know and do more as a consequence of technology in general, and AI in particular. The specific potential benefits of this with sustainability and agriculture are obvious: that we may detect and avoid the unsustainable, and that we may make the sustainable and beneficial more cost-effective to pursue. With respect to agriculture, the best examples I know of the latter is the more heterogeneous, dense, proficient, and sustainable use of agricultural land

because equipment is able to shift planting strategies rapidly and reliably in response to microlocal contexts.

Risks of these technologies The main problem is that ICT in general, and AI as a part of that, benefit transnational aggregation of business and other special interests, but governance mostly happens at the national and local level, so technology may facilitate lawlessness, corruption, and exploitation. This is an ongoing and heated battle, the outcome is uncertain. Ultimately AI should allow us to track responsibility, and will hopefully help us improve regulation and undercut the interests funding populist politicians that are compromising the rule of law that is essential to environmental protection. But this outcome is not certain, and already considerable damage is being done because that outcome is not yet achieved.

Recommended General Readings -Joanna J. Bryson, Mihailis E. Diamantis, and Thomas D. Grant, Of, For, and By the People: The Legal Lacuna of Synthetic Persons. *Artificial Intelligence and Law* 25(3):273–291 [Sep 2017]. Two professors of law and I argue that it would be a terrible, terrible idea to make something strictly AI (in contrast to an organisation also containing humans) a legal person. In fact, the only good thing about this is that it gives us a chance to think about where legal personhood has already been overextended (we give examples).
<https://link.springer.com/article/10.1007/s10506-017-9214-9>
 -This is still in review / a draft, but it might be of interest: Alexander J. Stewart, Nolan McCarty, Joanna J. Bryson, Explaining Parochialism: A Causal Account for Political Polarization in Changing Economic Environments, <https://arxiv.org/abs/1807.11477>
 -This is even further from the topic but may be of interest: Miles Brundage, Shahar Avin, Jack Clark, Helen Toner, Peter Eckersley, Ben Garfinkel, Allan Dafoe, Paul Scharre, Thomas Zeitzoff, Bobby Filar, Hyrum Anderson, Heather Roff, Gregory C. Allen, Jacob Steinhardt, Carrick Flynn, Seán Ó hÉigeartaigh, Simon Beard, Haydn Belfield, Sebastian Farquhar, Clare Lyle, Rebecca Crootof, Owain Evans, Michael Page, Joanna Bryson, Roman Yampolskiy, Dario Amodei, The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation , a technical report apparently published by all seven of the Future of Humanity Institute, University of Oxford, Centre for the Study of Existential Risk, University of Cambridge, Center for a New American Security, Electronic Frontier Foundation, and OpenAI. <https://arxiv.org/abs/1802.07228>

Participant	Peter Callahan
Affiliation	Princeton University – Research Assistant, PIIRS Global Systemic Risk research community
Email	pwcallah@princeton.edu
Primary Field	Natural Resource Management, Renewable Energy Technology & Policy
Brief Bio	Peter Callahan is a research assistant for the PIIRS research community on Global Systemic Risk. A graduate from Princeton University, Peter went on to earn his M.S. in Geography and Environmental Studies from the University of New Mexico.
Description of Research	Peter’s research has focused on the study of systemic risk, resilience in socio-ecological systems, natural resource management, renewable energy technology, and environmental policy.
Website	http://risk.princeton.edu

Participant	Amar Causevic
Affiliation	Global Economic Dynamics and the Biosphere Programme – Research Assistant
Email	amar.causevic@kva.se
Primary Field	Climate finance, climate change and security, emerging technologies and the biosphere
Brief Bio	Amar Causevic is a research assistant at the Global Economic Dynamics and the Biosphere Programme at the Royal Swedish Academy of Sciences. His research focuses on examining

interactions between financial markets and the biosphere. Prior to his current position, he worked on finance, energy and climate change-related issues at the World Bank, Carbon War Room, and the United States Agency for International Development (USAID). Amar holds a Masters of Arts in International Economics and Energy, Resources & Environment from the Paul H. Nitze School of Advanced International Studies (SAIS) at the Johns Hopkins University.

Description of Research	Amar is currently examining interactions between financial markets and the biosphere at Global Economic Dynamics and the Biosphere Programme. Also, he works on Beijer Institute's Complexity, Technology and Governance looking at environmental, technological and socio-economic changes and their impact on the behavior of complex systems.
Website	http://www.gedb.se/
Case Studies	I conducted research that examined the use of artificial intelligence and algorithms in the precision agriculture realm. The agro-tech start-up ecosystem in emerging markets (China, India, Brazil, etc.) has been a good example where new technologies are developing at very fast rate while at the same time receiving huge amounts of funding from different investors.
Case Study Readings	<ul style="list-style-type: none"> - Digital Agriculture Feeding the future document (http://breakthrough.unglobalcompact.org/disruptive-technologies/digital-agriculture/) - Radical ocean futures-scenario development using science fiction prototyping (https://www.sciencedirect.com/science/article/pii/S0016328716301914) - Secretary-General's Strategy on New Technologies (http://www.un.org/en/newtechnologies/index.shtml?lrs=5de0a812-84d6-45d4-90d1-f66c050ac985)
Benefits of these technologies	The biggest benefits I currently see is the optimization in the use of resources (e.g. water in crop production) as well as minimizing the use of fertilizers (e.g. reducing the amount of fertilizer used per m ²). Furthermore, farmers can obtain better information about soil characteristics and gather other important data for the sake of better understanding not only of their production, but ecological footprint as well.
Risks of these technologies	I see several issues, those are: disturbing labor markets (i.e. loss of manual employment), high energy dependency of precision agriculture devices, technical failures and resultant damages, problems related to hacking of these devices, environment-related concerns (i.e. increase in e-waste), legal status of the machines operating independently, the overall cost of high tech equipment for average farmer.
Recommended General Readings	<ul style="list-style-type: none"> - Machines learning development (https://rethink.earth/machines-learning-development/) - 'New Wilderness' Requires Algorithmic Transparency: A Response to Cantrell et al. (https://www.cell.com/trends/ecology-evolution/fulltext/S0169-5347(17)30161-1) - The Biosphere Code Manifesto (http://thebiospherecode.com/index.php/manifesto)

Participant	Miguel Centeno
Affiliation	Princeton University—PIIRS Global Systemic Risk research community
Email	cenmiga@princeton.edu
Primary Field	Sociology
Brief Bio	Miguel Centeno is Musgrave Professor of Sociology and International Affairs, and Director of PIIRS Global Systemic Risk at Princeton University. Professor Centeno has published many articles, chapters, and books including <i>War and Society</i> (Polity 2016), <i>State and Nation Making in the Iberian World</i> (Cambridge 2013) and <i>State Making in the Developing World</i> (Cambridge 2016). He is the founder of the Research Community on Global Systemic Risk funded by PIIRS from 2013-20 and recently published "The Emergence of Global Systemic Risk" in the <i>Annual Review of Sociology</i> (2015). From 2012-2017, he served as Chair of the Sociology Department.

Description of Research	<p>As the director of the PIIRS research community on Global Systemic Risk at Princeton University, Miguel provides structure for a core group of scholars across a wide variety of disciplines in Princeton as well as partner institutions across the world to establish a common dialogue in this emerging field. Working as a group, the depth and breadth of their interests will help to establish a comprehensive and cohesive framework for the study of risk and thus move the field forward. Through external dialogue and collaboration, the PIIRS Global Systemic Risk group seeks to help develop the field of systemic risk into a cohesive field of academic study with the goal of working with practitioners and policymakers to create more robust and resilient global systems.</p> <p>Independently, Professor Centeno is primarily interested in how states work. Much of his research has focused on Latin American cases, including the development of Latin American states in the 19th century. He is also interested in globalization and is working on a project comparing contemporary state capacity, with an emphasis on Brazil, India, and South Africa.</p>
Website	http://risk.princeton.edu
Recommended General Readings	<p>-Perrow, Charles, <i>Normal Accidents</i> (1984)</p> <p>-Guillen, Mauro, <i>The Architecture of Collapse</i> (2016)</p> <p>-Mitchell, Melanie, <i>Complexity: A Guided Tour</i> (2011)</p> <p>-Centeno, MA, et al., The Emergence of Global Systemic Risk, <i>Annu. Rev. Sociol.</i> 2015. 41:65–85</p>

Participant	Darryl Farber
Affiliation	Penn State University – Assistant Teaching Professor, School of Engineering Design, Technology, and Professional Programs
Email	dlf112@psu.edu
Primary Field	Sociotechnical systems research
Brief Bio	<p>Darryl Farber is an Assistant Teaching Professor of Engineering Design in the School of Engineering Design, Technology, and Professional Programs at Penn State University. His scholarship focuses on strategic foresight and analysis of emerging technologies and critical infrastructure, sociotechnical systems. He has worked across several infrastructures that include agricultural biotechnology, drinking water systems, nuclear engineering systems, and transportation. Dr. Farber earned a B.S. (honors) in Plant Sciences from Cornell University, a M.S. in Agricultural Economics, and a Ph.D. in Technology Management and Policy from Penn State University. He served as a postdoctoral fellow at the Belfer Center for Science and International Affairs at Harvard University, and as a visiting lecturer-project associate at the Public Policy Programme, (now the Lee Kuan Yew School), National University of Singapore. He serves on the executive committee of the Council of Engineering Systems Universities and as an associate editor of the journal <i>Systems Engineering</i>.</p>
Description of Research	<p>The focus of Dr. Farber’s research is to understand the impact of emerging technologies on critical infrastructure with applications to strategic and global security. To do this, he works to bridge model-based systems engineering methods with scenario analysis and planning. The idea is to generate a whole systems perspective. The messy reality is that there are multiple stakeholders with sometimes competing interests and positions that make alignment for critical infrastructure functioning a challenging design problem. However, it is not a design problem in a common-sense use, but rather a problem of designing an evolving system. A pioneer in the field is the late Herbert A. Simon and is reflected in his book <i>The Sciences of the Artificial</i>. The problem may be thought of as a communication systems problem where interpretation, meaning, and context are central to effective decision making and the functioning of a system.</p>
Website	https://www.sedtapp.psu.edu/department/directory-detail-g.aspx?q=DLF112

- Recommended General Readings**
- de Weck, Olivier L., Daniel Roos, and Christopher L. Magee. 2011. Engineering Systems: Meeting Human Needs in a Complex Technological World. Cambridge, MA: The MIT Press.
 - Mitchell, Melanie. 2018. Artificial Intelligence Hits the Barrier of Meaning. New York Times, November 5.
 - Simon, Herbert A. 1996. The Sciences of the Artificial, 3rd edition, chapters 5, 6, 7, and 8. Cambridge, MA: The MIT Press.

Participant	Edward Felten
Affiliation	Princeton University – Directory, CITP
Email	felten@cs.princeton.edu
Primary Field	
Brief Bio	Edward Felten is a professor of computer science and public affairs and director of both the Center for Information Technology Policy and the Program in Information Technology and Society. He was appointed the chief technologist for the US Federal Trade Commission (FTC) in 2010 and currently serves as a technology consultant to the FTC.
Description of Research	Dr. Felten’s research interests include computer security and privacy, and public policy issues relating to information technology. Specific topics include software security, Internet security, electronic voting, cybersecurity policy, technology for government transparency, network neutrality and Internet policy.
Website	https://citp.princeton.edu

Participant	Danit Gal
Affiliation	Keio University – Project Assistant Professor, Cyber Civilization Research Center
Email	gal.danit@gmail.com
Primary Field	Technology Policy
Brief Bio	Danit Gal is a Project Assistant Professor at the Cyber Civilizations Research Center at the Keio University Global Research Institute in Tokyo, Japan. She chairs the IEEE P7009 standard on the Fail-Safe Design of Autonomous and Semi-Autonomous Systems and serves on various committees of The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems. She is an Associate Fellow at the Leverhulme Centre for the Future of Intelligence at the University of Cambridge, and an Affiliate at the Center for Information Technology Policy at Princeton University.
Description of Research	Technology geopolitics and governance, safety and security, diversity and inclusion, and maximizing shared social benefit.
Website	https://danitgal.com
Recommended General Readings	<ul style="list-style-type: none"> - https://maliciousaireport.com - https://www.theguardian.com/environment/2018/dec/02/world-verge-climate-catastrophe - https://www.bloomberg.com/graphics/2017-feeding-china/

Participant	Victor Galaz
Affiliation	Stockholm University – Deputy Director, Stockholm Resilience Centre
Email	victor.galaz@su.se
Primary Field	global environmental change and politics

Brief Bio	Victor Galaz is associate professor in political science and Deputy Director at the Stockholm Resilience Centre (Stockholm University). Galaz is also Program Director at the Beijer Institute for Ecological Economics at the Royal Swedish Academy of Sciences.
Description of Research	Victor Galaz explores institutional and political dimensions of global environmental change, especially the notion of non-linear changes in the Earth system (e.g. planetary boundaries, tipping elements in the climate system). H also has a strong interest in exploring how technological change interplays with environmental change, especially for information technological innovations.
Website	www.stockholmresilience.org
Case Studies	Digital farming, climate change monitoring and modeling, extraction of marine resources
Case Study Readings	- https://rethink.earth/machines-learning-development/ - https://www.theguardian.com/science/political-science/2015/oct/05/a-manifesto-for-algorithms-in-the-environment - Cantrell, Bradley, Laura J. Martin, and Erle C. Ellis. "Designing autonomy: Opportunities for new wildness in the Anthropocene." Trends in ecology & evolution 32.3 (2017): 156-166.
Benefits of these technologies	Increased efficiency, adaptability to changing circumstances, early warnings and response to ecological surprise
Risks of these technologies	limited understanding of systemic risks, hidden data gaps and biases, socio-political implications unknown
Recommended General Readings	-Pasquale, Frank. The black box society: The secret algorithms that control money and information. Harvard University Press, 2015. -Kelly, Kevin. What technology wants. Penguin, 2010. -Arthur, W. Brian. The nature of technology: What it is and how it evolves. Simon and Schuster, 2009.

Participant	David Garcia
Affiliation	Complexity Science Hub Vienna – Faculty
Email	garcia@csh.ac.at
Primary Field	Computational Social Science
Brief Bio	David Garcia has been a group leader at the Complexity Science Hub Vienna since September 2017. He leads a research group funded by WWTF (Vienna Research Groups for Young Investigators Call). He holds computer science degrees from Universidad Autonoma de Madrid (Spain) and ETH Zurich (Switzerland). David did a PhD and Postdoc at ETH Zurich, working at the chair of systems design.
Description of Research	David’s research focuses on computational social science, designing models and analyzing human behavior through digital traces. His main work revolves around the topics of emotions, cultures, and political polarization, combining statistical analyses of large datasets of online interaction with agent-based modeling of individual behavior. His thesis “Modeling collective emotions in online communities” provides an example of how agent-based modeling can be used to construct an integrated approach to collective emotional phenomena in cyberspace. His postdoctoral research on opinion polarization and online privacy has been funded by the Swiss National Science Foundation and the ETH Risk Centre. David’s work lies at the intersection of various scientific disciplines, combining methods from network science, computer science, and statistical physics to answer questions from psychology, economics, and political science.
Website	www.dgarcia.eu

- Recommended General Readings**
- Exposure to opposing views on social media can increase political polarization. Bail et al., PNAS, 2018 <https://www.pnas.org/content/115/37/9216.short>
 - Machine Behavior Needs to Be an Academic Discipline. Rahwan & Cebrian, Nautilus, 2018 <http://nautil.us/issue/58/self/machine-behavior-needs-to-be-an-academic-discipline>
 - Locally noisy autonomous agents improve global human coordination in network experiments. Shirado & Christakis, Nature, 2017 <https://www.nature.com/articles/nature22332>

Participant	Daniel Jiménez R
Affiliation	International Center for Tropical Agriculture (CIAT)
Email	d.jimenez@cgiar.org
Primary Field	Agriculture
Brief Bio	Daniel Jiménez is a scientist at the International Center for Tropical Agriculture (CIAT). He is an agronomist and has been a pioneer of using data-mining techniques for agricultural research in developing countries. Expertise in research project management, fundraising, ensuring science impact, strategic positioning, and establishment of international scientific collaborations. Currently, he coordinates the Data-Driven Agronomy Community of Practice of the CGIAR Platform for Big Data in Agriculture. His work has received recognition from the World Bank Group (2015) and the United Nations (2014 and 2017), and the team he leads at CIAT took the top prize at the Syngenta 2018 Crop Challenge in Analytics.
Description of Research	Agriculture, Data Mining, Research for Development
Website	<ul style="list-style-type: none"> - https://www.usaid.gov/digitalag/ciat-case-study - https://www.ideaconnection.com/syngenta-crop-challenge/finalists.php - https://ccafs.cgiar.org/research/results/cracking-patterns-big-data-saves-colombian-rice-farmers-huge-losses - http://ciat.cgiar.org/outcome/farmers-associations-across-colombia-institutionalized-climate-site-specific-management/ - http://blog.ciat.cgiar.org/big-data-the-team-that-seldom-rests/ - http://blog.ciat.cgiar.org/es/los-gremios-colombianos-entran-a-la-era-de-los-datos/ - http://ciat.cgiar.org/outcome/ciat-ccafs-science-activates-climate-adaptation-in-colombia/
Case Studies	<ul style="list-style-type: none"> -Putting data at the service of agriculture (CIAT) -Digitizing the science of discovery and the science of delivery (ICRISAT)
Case Study Readings	<ul style="list-style-type: none"> - https://www.usaid.gov/digitalag/ciat-case-study - https://www.usaid.gov/digital-development/machine-learning/AI-ML-in-development - https://www.usaid.gov/sites/default/files/documents/15396/Data_Driven_Agriculture_Web.pdf - https://www.weforum.org/reports/innovation-with-a-purpose-the-role-of-technology-innovation-in-accelerating-food-systems-transformation
Benefits of these technologies	Benefits: Sustainable agriculture (economically viable, environmentally friendly, socially fair)
Risks of these technologies	Risks: Automation of tasks (a farm that doesn't smell like a farm is not a farm anymore)
Risks of these technologies	Making decisions based on misinterpretation of the results
Recommended General Readings	<ul style="list-style-type: none"> - https://www.usaid.gov/digitalag -Reflecting the Past, Shaping the Future: Making AI Work for International Development (https://www.usaid.gov/digital-development/machine-learning/AI-ML-in-development) -Big Data for Resilience (https://www.iisd.org/library/big-data-resilience-storybook)

Participant	Brian King
Affiliation	CGIAR – Coordinator of the Platform for Big Data in Agriculture (Co-led by CIAT and the International Food Policy Research Institute)
Email	b.king@cgiar.org
Primary Field	Digital Development
Brief Bio	Brian King leads CGIAR’s Big Data Platform, a cross cutting program to mobilize CGIAR data to accelerate food security research and impact. Previously he served as a Senior Technical Advisor for Digital Agriculture, Team Lead for Digital Financial Services at the Global Development Lab, and Senior Digital Development Advisor to Bureau for Africa at the United States Agency for International Development (USAID). He holds a BA from UC Davis, an MBA from the University of Maryland and an MA from University of California, Berkeley.
Description of Research	How can we leverage semantic web and knowledge graph/symbolic approaches to make machine learning more "human"? How can this be turned into machine-driven organizational capabilities to: mobilize legacy data, glean new insights from unstructured data, manage the complexity of data interoperability and queryability across multiple research domains, and put this to use for building more resilient food systems and ecosystems?
Website	https://bigdata.cgiar.org/
Case Studies	The interaction of intelligent agents and digital communications channels, at scale is an important theme. One publication demonstrates the opportunity nicely: https://www.ictworks.org/wp-content/uploads/2019/01/farmer-chatbot.pdf
Case Study Readings	More of a position paper from Cornell and Deep Mind: https://arxiv.org/pdf/1806.01261.pdf
Benefits of these technologies	Industrialized agricultural contexts tend to have very rich historical data and as a result a lot of the innovation in using digital technologies for agriculture come from these contexts. These technologies offer concrete opportunities to fill the data gap and also support decision-making with messy or incomplete data. This is a critical capability to leverage for building resilient agro-ecologies on a global scale.
Risks of these technologies	We can be wrong! The underlying data can be messy, the predictions or correlations themselves spurious. Homogenization--we can use these tools to fill data gaps, but in the process, we can classify ecologies or people with too broad a brush and this can have implication on equity as well as sustainability.

Participant	Paul Larcey
Affiliation	University of Cambridge – Research Fellow, Engineering
Email	pl381@cam.ac.uk
Primary Field	Infrastructure
Brief Bio	Paul Larcey studied engineering and materials science for both his undergraduate & master’s degrees (Universities of Oxford & Cambridge respectively), and received an MBA in Finance (Imperial College Business School). He worked initially in a corporate research environment, followed by venture capital, before moving into global industrial sectors at board and senior levels closely involved in funding strategies for major projects (public and private) primarily in the infrastructure/engineering sectors, M&A, Private Equity and risk analysis in challenging environments. He has also worked in strategic management consulting with a focus on technology and sustainability. His current commercial work is in industrial / technology ventures.
Description of Research	Attached to the University of Cambridge and involved in the risk perception in finance decision making for infrastructure, the understanding of resilience in systems and the

fragility of new technologies especially incorporation machine learning. He is also seeking to develop methodologies, frameworks and cross-disciplinary understanding of systemic risk from academic and practitioner perspectives.

Website	http://www.eng.cam.ac.uk
Case Studies	<ul style="list-style-type: none"> -Incorporating AI systems in robotics used in extreme environments enabling enhanced resource harvesting -Using machine learning to rapidly identify environmentally deleterious changes in water quality -Crop yield prediction and mapping
Case Study Readings	<ul style="list-style-type: none"> -The grand challenges of Science Robotics Yang et al., Sci.Robot, 3, eaar 7650 (2018) -Artificial intelligence technology could help protect water technologies (2018) Phys.org, https://phys.org/news/2018-07-artificial-intelligence-technology.html -Machine Learning in Agriculture: A Review Konstantinos G. Liakos , Patrizia Busato , Dimitrios Moshou , Simon Pearson. Sensors 2018, 18, 2674; doi:10.3390/s18082674
Benefits of these technologies	<ul style="list-style-type: none"> -Lower risk to humans in challenging environments -Dramatically improved production capabilities in construction and enhanced infrastructure performance -Enhanced crop and livestock yields
Risks of these technologies	<ul style="list-style-type: none"> -Loss of human agency in the complex working environments -Potential cyber malfeasance of integrated complex infrastructure -Skills shortage in gathering and analyzing AI big data output in developing economies
Recommended General Readings	<ul style="list-style-type: none"> -An AI Race for Strategic Advantage: Rhetoric and Risks ÓhÉigeartaigh & Cave (2018) Association for the Advancement of Artificial Intelligence (www.aaai.org) -Is Artificial Intelligence Stuck in A Rut? W Knight (2018) MIT Technology Review https://www.technologyreview.com/s/603945/is-artificial-intelligence-stuck-in-a-rut/ -A good life for all within planetary boundaries Daniel W. O’Neill , Andrew L. Fanning , William F. Lamb and Julia K. Steinberger https://doi.org/10.1038/s41893-018-0021-4 -The Impact of Changing Technology on the Evolution of Complex Informational Networks, Bettencourt, L.M.A (2014) Santa Fe Working paper 2014-11-043

Participant	Karen Levy
Affiliation	Cornell University – Assistant Professor, Information Science
Email	karen.levy@cornell.edu
Primary Field	Information Science, Law, Sociology
Brief Bio	Karen Levy is a faculty member in Information Science at Cornell University, where she is also affiliated with Cornell Law School and the departments of Sociology and Science and Technology Studies. Her training is in law (JD from Indiana University) and sociology (PhD from Princeton). Prior to joining Cornell, she was a postdoc at NYU’s Information Law Institute and the Data & Society Research Institute.
Description of Research	Karen Levy studies the social, legal, and ethical implications of emerging technologies, with a particular focus on privacy, control, and information flow. Much of her research has focused on work and on intimate social relationships, and how information technologies interact with these contexts. She is working on a project with Solon Barocas about privacy and economic redistribution with respect to agricultural technologies.

Website	www.karen-levy.net
Case Studies	Emergence of precision agriculture and forecasting technologies in U.S. (and global) agricultural production; associated economic issues (e.g. mergers between ag-tech conglomerates) and privacy-protective responses (e.g., development of Ag Data Transparency Evaluator by American Farm Bureau Federation)
Case Study Readings	[Solon Barocas and I have a paper in draft, to be completed soon!]
Benefits of these technologies	Opportunities for more sustainable and efficient practices, particularly around reduced application of nitrogen and pesticides; increased food security
Risks of these technologies	Farmer privacy issues; lack of clear economic benefit to farmers, price discrimination, predominant distribution of economic benefits of ag tech to conglomerates and large corporate farms
Recommended General Readings	-Woodard et al., The Power of Agricultural Data http://science.sciencemag.org/content/362/6413/410 -Fraser, Land grab/data grab: precision agriculture and its new horizons https://www.tandfonline.com/doi/abs/10.1080/03066150.2017.1415887?journalCode=fjps20

Participant	Arka Mukherjee
Affiliation	Global IDs Princeton University – PIIRS Global Systemic Risks research community
Email	arka.mukherjee@globalids.com
Primary Field	Computational Chemistry
Brief Bio	Arka Mukherjee is the CEO of Global IDs, a data management software company with headquarters in Princeton, NJ. Dr. Mukherjee has led a global team of software developers in creating an innovative platform for understanding Data Ecosystems. The software is currently deployed in many organizations that have globally distributed data environments. The software helps organizations understand their data assets, and create value from data assets through cost reduction, revenue generation and risk mitigation initiatives. He holds a Ph.D. in Computational Chemistry from Princeton University (Princeton, NJ). He received his M.Sc. degree from the Indian Institute of Technology (Kanpur, India) and his B.Sc. degree from Presidency College (Calcutta, India).
Description of Research	Dr. Mukherjee has researched Complex Data Ecosystems for the last 20 years, building a systematic understanding of the data owned by large corporations. The goal of his research team has been to build software agents that can obtain a statistical and semantic understanding of complex data landscapes that are found inside most corporations. With this foundation, organizations can improve the efficiency, quality and security of their data environments and perform advanced analytics on large data sets. The complexity of large data ecosystems has often led to a poor understanding of systemic risk. The intent of our research is to surface hidden risk using computational approaches that construct knowledge graphs.
Website	www.globalids.com

Participant	Thayer Patterson
Affiliation	Princeton University – Research Fellow, PIIRS Global Systemic Risks research community
Email	tspatter@princeton.edu
Primary Field	Risk, Finance

Brief Bio	Thayer Patterson is a research fellow and coordinator of the PIIRS research community on Global Systemic Risk at Princeton University. Subsequent to receiving a Master in Finance from Princeton's Bendheim Center for Finance, his research has focused on the causes and consequences of catastrophic systemic risk.
Description of Research	Thayer's research has focused on the causes and consequences of catastrophic systemic risk.
Website	http://risk.princeton.edu

Participant	Jennifer Rexford
Affiliation	Princeton University – Professor, Computer Science
Email	jrex@princeton.edu
Primary Field	computer networking
Brief Bio	Jennifer Rexford is the Gordon Y.S. Wu Professor of Engineering and the Chair of Computer Science at Princeton University. Before joining Princeton in 2005, she worked for eight years at AT&T Labs--Research. Her research focuses on how to use programmable networking devices to enable innovation in the design and operation of computer networks.
Description of Research	Computer networking, with an emphasis on how to design and operate large networks (e.g., data-center, enterprise, and backbone networks), including Internet routing, measurement, and management. Recent work focuses on collaboration with researchers in programming languages and streaming algorithms, to leverage programmable network devices so that future computer networks can manage themselves.
Website	http://www.cs.princeton.edu/~jrex/
Recommended General Readings	https://researchportal.bath.ac.uk/en/publications/semantics-derived-automatically-from-language-corpora-necessarily

Participant	Jonathan Tannen
Affiliation	Facebook – Research Scientist
Email	jtannen@fb.com
Primary Field	Algorithmic Fairness
Brief Bio	A Philadelphia native, Tannen studied segregation and processes of neighborhood change at the Woodrow Wilson School. In doing so, he became interested in quantitative methods in social sciences. At Facebook, Tannen studies Algorithmic Fairness, focusing on their “Jobs” case study.
Description of Research	Jonathan Tannen studies systems to ensure algorithmic fairness broadly. He also has a personal passion for ways that neighborhoods cement or can combat inequality.
Case Studies	I personally study the algorithmic fairness of Facebook's Jobs Recommendation system. How can we build systems to prevent our algorithms from learning and embedding discrimination?
Case Study Readings	-A video about our project is at 46:46 here: https://developers.facebook.com/videos/f8-2018/f8-2018-day-2-keynote/ -Among generally useful readings on Algorithmic Fairness, my favorites include: Corbett-Davies and Goel, The measure and mis-measure of fairness. https://arxiv.org/pdf/1808.00023.pdf -Kilbertus et al, Avoiding discrimination through causal reasoning. https://arxiv.org/pdf/1706.02744.pdf

Benefits of these technologies	Making good matches for employment could be a huge boon to job-seekers and the economy broadly. The job search problem is a core societal challenge that Facebook is well positioned to solve.
Risks of these technologies	Naively training models introduces the risk that AI will learn the biases that are embedded in our current job market.
Recommended General Readings	-Corbett-Davies and Goel, The measure and mis-measure of fairness. https://arxiv.org/pdf/1808.00023.pdf -Kilbertus et al, Avoiding discrimination through causal reasoning. https://arxiv.org/pdf/1706.02744.pdf

Participant	Michael Veale
Affiliation	University College London – PhD candidate, Department of Science, Technology, Engineering & Public Policy University of Birmingham – Honorary Research Fellow, Birmingham Law School
Email	m.veale@ucl.ac.uk
Primary Field	Technology policy
Brief Bio	Michael Veale is a researcher in responsible public sector machine learning at University College London, specializing in the fairness and accountability of data-driven tools in the public sector, the interplay between advanced technologies, data protection law, and human-computer interaction. His research has been cited by national and international governments and regulators, discussed in the media, as well as debated in Parliament. Michael has acted as expert consultant on machine learning and society for the World Bank, United Nations, European Commission, the Royal Society and the British Academy, and a range of national governments. Michael is a Fellow at the Centre for Public Impact, an Honorary Research Fellow at Birmingham Law School, University of Birmingham, a Visiting Researcher at the BBC DataLab, and a member of the Advisory Council for the Open Rights Group. He holds a MS in Sustainability Science and Policy from U Maastricht, and a BSc from LSE. He tweets at @mikarv.
Description of Research	Michael looks at areas including: - responsible practices use of machine learning in consequential, high-stakes situation - the potential and limitations of the current legal frameworks concerning machine learning, in particular, European data protection law - human-computer interaction issues around individuals making use of their data rights - the governance of privacy-enhancing technologies
Website	http://michae.lv
Case Studies	Much discussion about precision agriculture, much of the 'AI' use is prospective. Systemic problems more likely to occur around 'smart' grid management at the moment. Algorithmic systems around smart meters and privacy are pertinent to environmental issues more broadly from a consumer point of view
Case Study Readings	On anthropology and algorithmic systems in water management: -Lansing, J. (1991). Priests and Programmers: Technologies of Power in the Engineered Landscape of Bali. Princeton University Press. -"Smart Meters" in Veale, Michael. Data Management and Use: Case Studies of Technologies and Governance. London: The Royal Society and the British Academy, 2017. https://royalsociety.org/media/policy/projects/data-governance/data-governance-case-studies.pdf -On the externalities from optimizations, see https://arxiv.org/abs/1806.02711
Benefits of these technologies	Some incremental improvement/better management capacity for certain natural resources.

Risks of these technologies Externalities from optimizations are not considered <https://arxiv.org/abs/1806.02711>

Recommended General Readings Charles Perrow, 'Normal Accidents'

Participant	Ben Zevenbergen
Affiliation	Princeton University – CITP
Email	benzevenbergen@princeton.edu
Primary Field	Law, ethics
Brief Bio	Ben Zevenbergen is a visiting professional specialist at the Center for Information Technology Policy. Previously a lawyer, policy advisor, he is finishing a PhD on the ethics of computer science research.
Description of Research	On the intersection of Internet and AI technology and society, specifically law and ethics.
Website	https://aiethics.princeton.edu
Benefits of these technologies	Efficiency, increased productivity, assistance for humans
Risks of these technologies	Lower employment, reliance on data systems and its predictions, surveillance,
Recommended General Readings	- https://www.foreignaffairs.com/reviews/review-essay/2018-11-16/beyond-ai-arms-race - https://www.vox.com/future-perfect/2018/12/21/18126576/ai-artificial-intelligence-machine-learning-safety-alignment - https://ec.europa.eu/digital-single-market/en/news/coordinated-plan-artificial-intelligence