CALL FOR PAPERS
AI ETHICS: THE PRIVACY CHALLENGE

The Future of Privacy Forum and the Brussels Privacy Hub of the Vrije Universiteit Brussel are partnering with IEEE Security & Privacy in a call for papers on **AI Ethics: The Privacy Challenge**.

Abstracts due to guest editors: July 15 2017
Manuscripts due in ScholarOne: 1 October
Publication date: May/June 2018

Researchers are encouraged to submit interdisciplinary works in law and policy, computer science and engineering, social studies, and economics for publication in a special issue of IEEE Security & Privacy. Authors of selected submissions will be invited to present their work at the Brussels Privacy Symposium, which will be hosted by the VUB on 6 November 2017, in Brussels, Belgium.

Sixty-five years after Alan Turing published "Computing Machinery and Intelligence," (*Mind*, vol. 59, no. 236, pp. 433–460) thinking machines have matured from scientific theory into practical reality. Developers deploy artificial intelligence in domains ranging from social networks, autonomous vehicles, and drones to speech and image recognition, universal translators, precision medicine, criminal justice, and ad targeting. Enhancing efficiency, increasing safety, improving accuracy, and reducing negative externalities are just some of AI's key benefits. However, AI also presents risks of opaque decision making, biased algorithms, security and safety vulnerabilities, and upending labor markets. In particular, AI and machine learning challenge traditional notions of privacy and data protection, including individual control, transparency, access, and data minimization. On content and social platforms, it can lead to narrowcasting, discrimination, and filter bubbles.

A group of industry leaders recently established a partnership to study and formulate best practices on AI technologies. Last year, the White House issued a report titled Preparing for the Future of Artificial Intelligence and announced a National Artificial Intelligence Research and Development Strategic Plan, laying out a strategic vision for federally funded AI research and development. These efforts seek to reconcile the tremendous opportunities that machine learning, human–machine teaming, automation, and algorithmic decision making promise in enhanced safety, efficiency gains, and improvements in quality of life, with the legal and ethical issues that these new capabilities present for democratic institutions, human autonomy, and the very fabric of our society.

Scientists and reporters have shown the vulnerability of automated and autonomous systems and deep learning processes to hacking, security breaches, inaccuracies, and propagation of societal biases. Critics argue that unsupervised, deep machine learning can sever the link between accountable agents and consequential decisions, reducing responsibility and preventing individuals from learning why they've been treated a certain way. Nature suggested, "There are many sources of bias in algorithms. One is the hard-coding of rules and use of data sets that already reflect

common societal spin. Put bias in and get bias out. Spurious or dubious correlations are another pitfall." ("More Accountability for Big-Data Algorithms," Nature, vol. 537, no. 7621, 2016).

The new European data protection framework, the General Data Protection Regulation, tightens restrictions on profiling and automated decision making, casting doubt on the viability of technological breakthroughs in the field. Last October, the 38th International Conference of Data Protection and Privacy Commissioners devoted part of its closed session to AI and robotics. In its final communique, the commissioners' conference expressed concern about the challenges these technologies present to the consent model of data collection as well as to the predictability, transparency, and accountability of algorithmic decision making. The French privacy regulator, CNIL, launched a public debate to address fundamental questions AI raises with regard to ethics, morality, and values.

At the same time, researchers hope to harness the power of AI and machine learning to better protect individuals' privacy and security, helping consumers to navigate complex sociotechnical architectures in smart cities and homes, transportation systems, financial transactions, and content platforms, in accordance with their preferences, and to implement their privacy policies and choices. For example, payment networks already use AI to allow financial institutions to increase the accuracy of real-time approvals of genuine transactions and reduce declined transactions. Social networking platforms use AI to flag hate speech, cyberbullying, and harassment and even to automatically identify intentionally misleading content—in real time and with minimal false positives.

Successful submissions will address the following issues:

• Privacy values in design
• Algorithmic due process and accountability
• Fairness and equity in automated decision making
• Accountable machines
• Formalizing definitions of privacy fairness and equity
• Societal implications of autonomous experimentation
• Deploying machine learning and AI to enhance privacy
• Cybersafety and privacy

Submissions should adhere to the following schedule:

15 July 2017: deadline for abstracts
31 July: workshop participants announced
1 October: first draft of paper due
6 November: workshop in Brussels
1 December: final draft of paper due

Direct any questions and submit abstracts (by July 15 2017) to the Guest Editors via sp3-2018@computer.org:

- Jules Polonetsky, The Future of Privacy Forum, julespol@fpf.org
- Omer Tene, International Association of Privacy Professionals, otene@iapp.org
- Ahmad-Reza Sadeghi, EIC, eic.sp@trust.cased.de
- Christopher Kuner, VUB

Submission Guidelines

- Submissions will be subject to the IEEE Computer Society's peer-review process.
- Abstracts should be at least 800 words long.
- Articles should be at most 7,200 words, with a maximum of 15 references, and should be understandable to a broad audience of people interested in security, privacy, and dependability.
- If portions of the paper have been previously published, at least 30% of the final paper must constitute new content.
- The writing style should be down to earth, practical, and original. Authors should not assume that the audience will have specialized experience in a particular subfield.
- All accepted articles will be edited according to the IEEE Computer Society style guide.
- Submit draft papers to ScholarOne **by October 1, 2017** at https://mc.manuscriptcentral.com/cs-ieee