

# 1st Symposium on Foundations of Responsible Computing

FORC 2020, June 1, 2020, Harvard University, Cambridge, MA,  
USA (virtual conference)

Edited by

Aaron Roth



*Editors*

**Aaron Roth**

University of Pennsylvania, Philadelphia, PA, USA  
aaroht@cis.upenn.edu

*ACM Classification 2012*

Theory of computation → Algorithmic game theory and mechanism design; Theory of computation → Theory of database privacy and security; Theory of computation

**ISBN 978-3-95977-142-9**

*Published online and open access by*

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at <https://www.dagstuhl.de/dagpub/978-3-95977-142-9>.

*Publication date*

May, 2020

*Bibliographic information published by the Deutsche Nationalbibliothek*

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at <https://portal.dnb.de>.

*License*

This work is licensed under a Creative Commons Attribution 3.0 Unported license (CC-BY 3.0):  
<https://creativecommons.org/licenses/by/3.0/legalcode>.



In brief, this license authorizes each and everybody to share (to copy, distribute and transmit) the work under the following conditions, without impairing or restricting the authors' moral rights:

- Attribution: The work must be attributed to its authors.

The copyright is retained by the corresponding authors.

Digital Object Identifier: 10.4230/LIPIcs.FORC.2020.0

ISBN 978-3-95977-142-9

ISSN 1868-8969

<https://www.dagstuhl.de/lipics>

## LIPICs – Leibniz International Proceedings in Informatics

LIPICs is a series of high-quality conference proceedings across all fields in informatics. LIPICs volumes are published according to the principle of Open Access, i.e., they are available online and free of charge.

### *Editorial Board*

- Luca Aceto (*Chair*, Gran Sasso Science Institute and Reykjavik University)
- Christel Baier (TU Dresden)
- Mikolaj Bojanczyk (University of Warsaw)
- Roberto Di Cosmo (INRIA and University Paris Diderot)
- Javier Esparza (TU München)
- Meena Mahajan (Institute of Mathematical Sciences)
- Dieter van Melkebeek (University of Wisconsin-Madison)
- Anca Muscholl (University Bordeaux)
- Luke Ong (University of Oxford)
- Catuscia Palamidessi (INRIA)
- Thomas Schwentick (TU Dortmund)
- Raimund Seidel (Saarland University and Schloss Dagstuhl – Leibniz-Zentrum für Informatik)

**ISSN 1868-8969**

**<https://www.dagstuhl.de/lipics>**



## ■ Contents

Preface	
<i>Aaron Roth</i> .....	0:vii
Efficient Candidate Screening Under Multiple Tests and Implications for Fairness	
<i>Lee Cohen, Zachary C. Lipton, and Yishay Mansour</i> .....	1:1–1:20
Metric Learning for Individual Fairness	
<i>Christina Ilvento</i> .....	2:1–2:11
Recovering from Biased Data: Can Fairness Constraints Improve Accuracy?	
<i>Avrim Blum and Kevin Stangl</i> .....	3:1–3:20
Can Two Walk Together: Privacy Enhancing Methods and Preventing Tracking of Users	
<i>Moni Naor and Neil Vexler</i> .....	4:1–4:20
Service in Your Neighborhood: Fairness in Center Location	
<i>Christopher Jung, Sampath Kannan, and Neil Lutz</i> .....	5:1–5:15
Bias In, Bias Out? Evaluating the Folk Wisdom	
<i>Ashesh Rambachan and Jonathan Roth</i> .....	6:1–6:15
Individual Fairness in Pipelines	
<i>Cynthia Dwork, Christina Ilvento, and Meena Jagadeesan</i> .....	7:1–7:22
Abstracting Fairness: Oracles, Metrics, and Interpretability	
<i>Cynthia Dwork, Christina Ilvento, Guy N. Rothblum, and Pragya Sur</i> .....	8:1–8:16
The Role of Randomness and Noise in Strategic Classification	
<i>Mark Braverman and Sumegha Garg</i> .....	9:1–9:20
Bounded-Leakage Differential Privacy	
<i>Katrina Ligett, Charlotte Peale, and Omer Reingold</i> .....	10:1–10:20





## ■ Preface

With the rise of the consumer internet, algorithmic decision making became personal. Beginning with relatively mundane things like targeted advertising, machine learning was brought to bear to make decisions *about people* (e.g. which ad to show them), and was trained on enormous datasets of personal information that we increasingly generated unknowingly, as part of our every-day “digital exhaust”. In the last several years, these technologies have been deployed in increasingly consequential domains. We no longer just use machine learning for targeting ads. We use it to inform criminal sentencing, to set credit limits and approve loans, and to inform hiring and compensation decisions. All of this means that it is increasingly urgent that our automated decision-making upholds social norms like “privacy” and “fairness” that we are accustomed to thinking about colloquially and informally, but are difficult to define precisely enough to encode as constraints on algorithms. It also means that we must grapple with strategic interactions, as changes in our algorithms lead to changes in the behavior of the users whose data the algorithms operate on.

It is exactly because the definitions are so difficult to get right that strong theoretical foundations are badly needed. We need definitions that have meaningful semantics, and we need to understand both the limits of our ability to design algorithms satisfying these definitions, and the tradeoffs involved in doing so. Foundations of Responsible Computing is a venue for developing this theory.

Our first program is a great example of the kind of work we aim to feature. We have 17 accepted papers, 10 of which appear in this proceedings (we allow authors to opt to instead have a one-page abstract appear on the website but not in the proceedings, to facilitate different publication cultures). The program includes formal proposals for how to reason about different kinds of privacy that fall short of differential privacy, but that we much reckon with because of legal or other practical realities. It includes work studying the implications of imposing fairness constraints in the presence of faulty data. It contains work aimed at making strong but to-date impractical fairness constraints more actionable. And it contains papers studying the strategic and game theoretic effects of deployed algorithms. This is all to say, our inaugural conference has much to say about the foundations of computation in the presence of pressing social concerns.

Finally, let me note that our program committee finished their work in the midst of a historic global pandemic, that has and continues to disrupt all of our lives. Despite this, they did a remarkable job. The ongoing pandemic will mean that we cannot meet in person for FORC 2020, but it will not lessen the impact of the work, now to be presented in a remote format.

Aaron Roth  
Philadelphia, PA  
April 12, 2020



