FOUNDATIONS OF PROBABILISTIC PROGRAMMING

What does a probabilistic program actually compute? How can one formally reason about such probabilistic programs? This valuable guide covers such elementary questions and more. It provides a state-of-the-art overview of the theoretical underpinnings of modern probabilistic programming and their applications in machine learning, security, and other domains, at a level suitable for graduate students and non-experts in the field. In addition, the book treats the connection between probabilistic programs and mathematical logic, security (what is the probability that software leaks confidential information?), and presents three programming languages for different applications: Excel tables, program testing, and approximate computing. This title is also available as Open Access on Cambridge Core.

GILLES BARTHE is Scientific Director at the Max Planck Institute for Security and Privacy and Research Professor at the IMDEA Software Institute, Madrid. His recent research develops programming language techniques and verification methods for probabilistic languages, with a focus on cryptographic and differentially private computations.

JOOST-PIETER KATOEN is Professor at RWTH Aachen University and University of Twente. His research interests include formal verification, formal semantics, concurrency theory, and probabilistic computation. He co-authored the book *Principles of Model Checking* (2008). He received an honorary doctorate from Aalborg University, is member of the Academia Europaea, and is an ERC Advanced Grant holder.

ALEXANDRA SILVA is Professor of Algebra, Semantics, and Computation at University College London. A theoretical computer scientist with contributions in the areas of semantics of programming languages, concurrency theory, and probabilistic network verification, her work has been recognized by multiple awards, including the Needham Award 2018, the Presburger Award 2017, the Leverhulme Prize 2016, and an ERC Starting Grant in 2015.

FOUNDATIONS OF PROBABILISTIC PROGRAMMING

Edited by

GILLES BARTHE
Max-Planck-Institut für Cybersicherheit und Schutz der Privatsphäre, Bochum, Germany

JOOST–PIETER KATOEN
Rheinisch–Westfälische Technische Hochschule, Aachen, Germany

ALEXANDRA SILVA
University College London

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

79 Anson Road, #06–04/06, Singapore 079906

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781108488518
DOI: 10.1017/9781108770750

When citing this work, please include a reference to the DOI 10.1017/9781108770750

First published 2021

Printed in the United Kingdom by TJ Books Limited, Padstow Cornwall

A catalogue record for this publication is available from the British Library.

ISBN 978-1-108-48851-8 Hardback

Contents

Contributors

Giorgio Bacci *Department of Computer Science, Aalborg University, Selma Lagerlöfs Vej 300, DK-9220 Aalborg, Denmark*

Gilles Barthe *Max Planck Institute for Cybersecurity and Privacy, Exzenterhaus, Universitätsstr. 60, 44789 Bochum, Germany*

Johannes Borgström *Department of Information Technology, Uppsala University, 752 37 Uppsala, Sweden*

Jose Manuel Calderón Trilla *Galois, Inc., Arlington, VA 22203, USA*

Michael Carbin *MIT CSAIL, 77 Massachusetts Ave, 32-G782 Cambridge, MA 02139, USA*

Krishnendu Chatterjee *Am Campus 1, IST Austria, A-3400 Klosterneuburg, Austria*

Fredrik Dahlqvist *University College London, Department of Computer Science, Gower Street, London WC1E 6BT, UK*

Ugo Dal Lago *Dipartimento di Informatica – Scienza e Ingegneria Università degli Studi di Bologna, Mura Anteo Zamboni, 7, 40127 Bologna, Italy*

Hongfei Fu *John Hopcroft Center for Computer Science, Shanghai Jiao Tong University, 800 Dongchuan Road, Minhang District, Shanghai 200240, China*

Diane Gallois-Wong *Inria de Paris 2, rue Simone Iff, CS 42112, 75589 Paris CEDEX 12, France*

Jeremy Gibbons *University of Oxford, Department of Computer Science, Parks Road, Oxford OX1 3QD, UK*

Andrew D. Gordon *Microsoft Research Ltd, 21 Station Road, Cambridge CB1 2FB, UK*

Thore Graepel *University College London, Department of Computer Science, Gower Street, London WC1E 6BT, UK*

Michael Hicks *Department of Computer Science, University of Maryland, College Park, MD 20742, USA*

Cătălin Hriţcu *Inria de Paris 2, rue Simone Iff, CS 42112, 75589 Paris CEDEX 12, France*

John Hughes *Department of Computer Science and Engineering, SE-412 96, Gothenburg, Sweden*

Justin Hsu *School of Computer, Data and Information Sciences, 1210 W. Dayton Street Madison, WI 53706-1613, USA*

Daniel Huang *EECS, University of California, Berkeley, CA 94720-1770, USA*

Benjamin Lucien Kaminski *Software Modeling and Verification Group, RWTH Aachen University D-52056 Aachen, Germany*

Joost-Pieter Katoen *Software Modeling and Verification Group, RWTH Aachen University D-52056 Aachen, Germany*

Dexter Kozen *Computer Science Department, 436 Gates Hall, Cornell University, Ithaca, New York 14853–7501, USA*

Bart Jacobs *Interdisciplinary Hub for Security, Privacy, and Data Governance, Radboud University Nijmegen, Erasmusplein 1, 6525 HT Nijmegen, The Netherlands*

Lampropoulos Leonidas *University of Pennsylvania, Department of Computer and Information Science, 3330 Walnut Street, Philadelphia, PA 19104-6389, USA*

Stephen Magill *Galois, Inc., Arlington, VA 22203, USA*

Radu Mardare *Computer and Information Sciences University of Strathclyde, 26 Richmond Street, Glasgow G1 1XH, UK*

Piotr Mardziel *Electrical and Computer Engineering Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*

Christoph Matheja *Software Modeling and Verification Group, RWTH Aachen University D-52056 Aachen, Germany*

Annabelle McIver *Department of Computing, Macquarie University, Sydney, NSW 2109, Australia*

Sasa Misailovic *Department of Computer Science, University of Illinois, 4110 Siebel Center, Urbana, IL 61801, USA*

Carroll Morgan *Faculty of Engineering, UNSW, Sydney, NSW 2052, Australia*

Greg Morrisett *Cornell Tech, Cornell University, 2 West Loop Road, New York, New York 10044, USA*

Petr Novotný *Faculty of Informatics, Masaryk University, Botanickà 68a, 60200 Brno, Czech Republic*

Prakash Panangaden *School of Computer Science, McGill University, 3480 rue University, Montreal, Quebec H3A 0E9, Canada*

Benjamin C. Pierce *University of Pennsylvania, Department of Computer and Information Science, 3330 Walnut Street, Philadelphia, PA 19104-6389, USA*

Gordon Plotkin *Laboratory for Foundations of Computer Science, School of Informatics, Informatics Forum, 10 Crichton Street, Edinburgh EH8 9AB, UK*

Nicolas Rolland *University College London, Department of Computer Science, Gower Street, London WC1E 6BT, UK*

Claudio Russo *DFINITY, Stockerstrasse 47, 8002 Zürich, Switzerland*

Sriram Sankaranarayanan *Department of Computer Science, University of Colorado, Boulder CO 80309–0430, USA*

Tom Schrijvers *Department of Computer Science. KU Leuven. Celestijnenlaan 200A. 3001 Leuven. Belgium*

Alexandra Silva *University College London, Department of Computer Science, Gower Street, London WC1E 6BT, UK*

Bas Spitters *Department of Computer Science, Aarhus University, Nygaard-268 Aabogade 34, DK-8200 Aarhus N, Denmark*

Sam Staton *Department of Computer Science, University of Oxford Wolfson Building, Parks Road, Oxford OX1 3QD, UK*

Ian Sweet *Department of Computer Science, University of Maryland, College Park, MD 20742, USA*

Marcin Szymczak *Lehrstuhl für Informatik 2, RWTH Aachen University, 52056 Aachen, Germany*

Daniel Tarlow *Google Research, Brain Team, Montreal, Quebec H3B 2Y5, Canada*

Li-yao Xia *University of Pennsylvania, Department of Computer and Information Science, 3330 Walnut Street, Philadelphia, PA 19104-6389, USA*

Fabio Zanasi *University College London, Department of Computer Science, Gower Street, London WC1E 6BT, UK*

Preface

Probabilistic programs

Probabilistic programs describe recipes for inferring statistical conclusions from a complex mixture of uncertain data and real-world observations. They can represent probabilistic graphical models far beyond the capabilities of Bayesian networks and are expected to have a major impact on machine intelligence. Probabilistic programs are ubiquitous. They steer autonomous robots and self-driving cars, are key to describe security mechanisms, naturally code up randomised algorithms for solving NP-hard or even unsolvable problems, and are rapidly encroaching on AI. Probabilistic programming aims to make probabilistic modelling and machine learning accessible to the programmer.

What is this book all about?

Probabilistic programs, though typically relatively small in size, are hard to grasp, let alone check automatically. Elementary questions are notoriously hard – even the most elementary question "does a program halt with probability one? – is "more undecidable" than the halting problem. This book is about the theoretical foundations of probabilistic programming. It is primarily concerned with fundamental questions such as the following: What is Bayesian probability theory? What is the precise mathematical meaning of probabilistic programs? How show almost-sure termination? How determine the (possibly infinite) expected runtime of probabilistic programs? How can two similar programs be compared? It covers several analysis techniques on probabilistic programs such as abstract interpretation, algebraic reasoning and determining concentration measures.

These chapters are complemented with chapters on the formal definition of concrete probabilistic programming languages and some possible applications of the use of probabilistic programs.

How to read this volume?

The volume consists of five parts: semantics, verification, logic, security, and programming languages.

Semantics.
The first part on semantics consists of four chapters on different aspects of the formal semantics of probabilistic programming languages. Dahlqvist *et al.* start off in Chapter 1 by presenting an operational and denotational semantics of an imperative language with discrete and continuous distributions. The chapter by Staton presents a compositional measure-theoretic semantics for a first-order probabilistic language with arbitrary soft conditioning. Chapter 3 by Huang *et al.* studies semantics with a focus on computability. Motivated by the tension between the discrete and the continuous in probabilistic modelling, type-2 computable distributions are introduced as the elementary mathematical objects to give semantics to probabilistic languages. Dal Lago's Chapter 4 completes the semantics part by treating a probabilistic version of the λ-calculus, the backbone language for functional programming languages.

Verification.
The second part on formal verification starts with a chapter by Barthe and Hsu on the use of couplings, a well-known concept in probability theory for verifying probabilistic programs. Kaminski *et al.* present a weakest pre-condition calculus in the style of Dijkstra for determining the expected run-time of a probabilistic program. This can be used to determine whether a program needs infinitely many steps to terminate with probability one. Chapter 7 by Chatterjee *et al.* presents various techniques based on supermartingales to decide the almost-sure termination of a probabilistic program, i.e., does a program terminate with probability one on all possible inputs? The verification part ends with a chapter by Sankaranarayanan about the quantitative analysis of probabilistic programs by means of concentration of measure inequalities. This includes Chernoff–Hoeffding and Bernstein inequalities.

Logic.
The third part focuses on logic and consists of two chapters. In Chapter 9, Jacobs and Zanasi present an insightful new perspective on fundamental aspects of probability theory which are relevant for probabilistic programming. Their chapter introduces a category-theoretic formalization of Bayesian reasoning, and in particular a string-diagram-friendly one. This contribution is complemented by the chapter by Bacci *et al.* which surveys some recent contributions on extending the classical

Birkhoff/Lawvere theory of equational reasoning to the quantitative setting. In that setting, compared entities are not necessarily equal but rather treated by a notion of distance.

Security.
Part four is concerned with security, an important application field in which probabilities are pivotal. Chapter 11 by Calderon *et al.* collects together results on probabilistic abstract interpretation and applies them to probabilistic programming in the context of security. Chapter 12 by Gibbons *et al.* presents an embedded domain-specific language in Haskell to compute hyper-distributions induced by programs. This is used to compute the amount of leakage of a program by measuring variations on post-distributions that include Shannon entropy and Bayes' risk (that is, the maximum information an attacker can learn in a single run).

Programming languages.
The final part of this volume is concerned with three concrete probabilistic programming languages: Luck, Tabular, and Rely. Chapter 13 describes Luck, proposed by Lampropoulos *et al.*, a language for test generation and a framework for property-based testing of functional programs. Luck combines local instantiation of unknown variables and global constraint solving to make test generation more efficient than existing approaches. Gordon *et al.* introduce Tabular, a domain-specific programming language designed to express probabilistic models and to perform probabilistic inference over relational data. Tabular can be used from Microsoft Excel or as stand-alone software. Chapter 14 presents the syntax, semantics and type system of Tabular and shows how it can be used to design probabilistic models and to perform probabilistic inference. Chapter 15, the last chapter of this volume, by Carbin and Misailovic, presents Rely, a programming language that enables reasoning about the probability that a program produces the correct result when executed on unreliable hardware.

How this volume emerged

This book consists of 15 contributed chapters and a preface. The idea for this volume emerged at the first summer school on Foundations of Programming and Software Systems held in Braga, Portugal, May–June 2017. This biennial school series is supported by EATCS (European Association for Theoretical Computer Science), ETAPS (European Conference on Theory and Practice of Software), ACM SIGPLAN (Special Interest Group on Programming Languages) and ACM SIGLOG (Special Interest Group on Logic and Computation). It was felt that there is no comprehensive book on the theoretical foundations of probabilistic programming

languages. We sincerely hope that this volume contributes to filling this gap. Enjoy reading!

Acknowledgements

Many people have helped us to make this volume possible. First and foremost, we like to thank all authors for their fine contributions and for their patience in this book process. All chapters have been subject to peer reviews. We thank the reviewers Alejandro Aguirre, Krishnendu Chatterjee, Ugo Dal Lago, Thomas Ehrhard, Claudia Faggian, Marco Gaboardi, Francesco Gavazzo, Jeremy Gibbons, Andrew D. Gordon, Ichiro Hasuo, Jan Hoffmann, Justin Hsu, Bart Jacobs, Benjamin Lucien Kaminski, Pasquale Malacaria, Radu Mardare, Christoph Matheja, Annabelle McIver, Sasa Misailovic, Petr Novotný, Prakash Panangaden, Corina Pasareanu, Alejandro Russo, Sriram Sankaranarayanan, Steffen Smolka, and Marcin Szymczak for their thorough reviewing work.

We thank Luis Barbosa, Catarina Fernandes and Renato Neves for their invaluable efforts in the local organisation of the FoPPS 2017 summer school. We thank Joshua Moerman for his efforts in the editing process. David Tranah and Anna Scriven from Cambridge University Press are thanked for their support during this process. We thank Hanna Schraffenberger for designing the cover of the book. Finally, we are grateful for the financial support from the ERC (AdG FRAPPANT and StG ProFoundNet) and the MPI on Security and Privacy which enabled to publish this volume under gold open access.