
**Property Testing:
A Learning Theory
Perspective**

Property Testing: A Learning Theory Perspective

Dana Ron

Tel-Aviv University

Ramat-Aviv

Tel-Aviv 69978

Israel

danar@eng.tau.ac.il

now

the essence of **knowledge**

Boston – Delft

Foundations and Trends[®] in Machine Learning

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
USA
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is Dana Ron, Property Testing: A Learning Theory Perspective, *Foundations and Trends[®] in Machine Learning*, vol 1, no 3, pp 307–402, 2008

ISBN: 978-1-60198-182-0

© 2008 Dana Ron

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc. for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1-781-871-0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

**Foundations and Trends[®] in
Machine Learning**
Volume 1 Issue 3, 2008
Editorial Board

Editor-in-Chief:

Michael Jordan

*Computer Science Division
University of California, Berkeley
Berkeley, CA 94720-1776
USA*

Editors

Peter Bartlett (UC Berkeley)	Michael Littman (Rutgers University)
Yoshua Bengio (Université de Montréal)	Gabor Lugosi (Pompeu Fabra University)
Avrim Blum (Carnegie Mellon University)	David Madigan (Columbia University)
Craig Boutilier (University of Toronto)	Pascal Massart (Université de Paris-Sud)
Stephen Boyd (Stanford University)	Andrew McCallum (University of Massachusetts Amherst)
Carla Brodley (Tufts University)	Marina Meila (University of Washington)
Inderjit Dhillon (University of Texas at Austin)	Andrew Moore (Carnegie Mellon University)
Jerome Friedman (Stanford University)	John Platt (Microsoft Research)
Kenji Fukumizu (Institute of Statistical Mathematics)	Luc de Raedt (Albert-Ludwigs Universitaet Freiburg)
Zoubin Ghahramani (Cambridge University)	Christian Robert (Université Paris-Dauphine)
David Heckerman (Microsoft Research)	Sunita Sarawagi (IIT Bombay)
Tom Heskes (Radboud University Nijmegen)	Robert Schapire (Princeton University)
Geoffrey Hinton (University of Toronto)	Bernhard Schoelkopf (Max Planck Institute)
Aapo Hyvarinen (Helsinki Institute for Information Technology)	Richard Sutton (University of Alberta)
Leslie Pack Kaelbling (MIT)	Larry Wasserman (Carnegie Mellon University)
Michael Kearns (University of Pennsylvania)	Bin Yu (UC Berkeley)
Daphne Koller (Stanford University)	
John Lafferty (Carnegie Mellon University)	

Editorial Scope

Foundations and Trends[®] in Machine Learning publishes survey and tutorial articles in the following topics:

- Adaptive control and signal processing
- Applications and case studies
- Behavioral, cognitive and neural learning
- Bayesian learning
- Classification and prediction
- Clustering
- Data mining
- Dimensionality reduction
- Evaluation
- Game theoretic learning
- Graphical models
- Independent component analysis
- Inductive logic programming
- Kernel methods
- Markov chain Monte Carlo
- Model choice
- Nonparametric methods
- Online learning
- Optimization
- Reinforcement learning
- Relational learning
- Robustness
- Spectral methods
- Statistical learning theory
- Variational inference
- Visualization

Information for Librarians

Foundations and Trends[®] in Machine Learning, 2008, Volume 1, 4 issues. ISSN paper version 1935-8237. ISSN online version 1935-8245. Also available as a combined paper and online subscription.

Foundations and Trends[®] in
Machine Learning
Vol. 1, No. 3 (2008) 307–402
© 2008 Dana Ron
DOI: 10.1561/2200000004



Property Testing: A Learning Theory Perspective

Dana Ron

*Department of Electrical Engineering — Systems, Tel-Aviv University,
Ramat-Aviv, Tel-Aviv 69978, Israel, danar@eng.tau.ac.il*

Abstract

Property testing deals with tasks where the goal is to distinguish between the case that an object (e.g., function or graph) has a pre-specified property (e.g., the function is linear or the graph is bipartite) and the case that it differs significantly from any such object. The task should be performed by observing only a very small part of the object, in particular by querying the object, and the algorithm is allowed a small failure probability.

One view of property testing is as a relaxation of learning the object (obtaining an approximate representation of the object). Thus property testing algorithms can serve as a preliminary step to learning. That is, they can be applied in order to select, very efficiently, what hypothesis class to use for learning. This survey takes the learning-theory point of view and focuses on results for testing properties of functions that are of interest to the learning theory community. In particular, we cover results for testing algebraic properties of functions such as linearity, testing properties defined by concise representations, such as having a small DNF representation, and more.

Contents

1	Introduction	1
1.1	Property Testing as Relaxed Decision	2
1.2	Property Testing and Learning (Estimation)	3
1.3	Property Testing and Hypothesis Testing	4
1.4	Topics and Organization	5
2	Preliminaries	7
2.1	Definitions and Notations	7
2.2	A Basic Observation on the Relation Between Learning and Testing	9
3	Algebraic Properties	13
3.1	Linearity	13
3.2	Low-Degree Polynomials	20
4	Basic (Boolean) Function Classes	35
4.1	Singletons, Monomials, and Monotone DNF	35
4.2	Juntas	47
4.3	Testing by Implicit Learning: General DNF, Decision Trees and More	54
4.4	Testing Linear Threshold Functions	59

5 Other Models of Testing	61
5.1 Distribution-Free Testing	62
5.2 Testing From Random Examples	65
6 Other Results	75
6.1 Monotonicity	75
6.2 Clustering	77
6.3 Properties of Distributions	79
6.4 Testing Membership in Regular Languages	81
6.5 Testing Graph Properties	82
6.6 Tolerant Testing and Distance Approximation	87
Acknowledgments	89
A The (Multiplicative) Chernoff Bound	91
References	93

1

Introduction

Property testing [82, 128] is the study of the following class of problems.

Given the ability to perform (local) queries concerning a particular object the problem is to determine whether the object has a predetermined (global) property or differs significantly from any object that has the property. In the latter case we say it is far from (having) the property. The algorithm is allowed a small probability of failure, and typically it inspects only a small part of the whole object.

For example, the object may be a graph and the property is that it is bipartite, or the object may be a function and the property is that it is linear. It is usually assumed that the property testing algorithm is given query access to the object. When the object is a function f the queries are of the form: “what is $f(x)$?” while if the object is a graph then queries may be: “is there an edge between vertices u and v ” or: “what vertex is the i^{th} neighbor of v ?”. In order to determine what it means to be far from the property, we need a distance measure between objects. In the case of functions it is usually the weight according to the uniform

2 Introduction

distribution of the symmetric difference between the functions, while in the case of graphs it is usually the number of edge modifications divided by some upper bound on the number of edges. When dealing with other objects (e.g., the object may be a set of points and the property may be that the set of points can be clustered in a certain way) one has to define both the types of queries allowed and the distance measure.

1.1 Property Testing as Relaxed Decision

Property testing problems are often viewed as a relaxation of decision problems. Namely, instead of requiring that the algorithm decide whether the object has the property or does not have the property, the algorithm is required to decide whether the object has the property or is far from having the property. Given this view there are several scenarios in which property testing may be useful.

- If the object is very large, then it is infeasible to examine all of it and we must design algorithms that examine only a small part of the object and make an approximate decision based on what they view.
- Another scenario is when the object is not too large to fully examine, but the exact decision problem is \mathcal{NP} -hard. In such a case some form of approximate decision is necessary if one seeks an efficient algorithm and property testing suggest one such form. We note that in some cases the approximation essentially coincides with *standard* notions of approximation problems (e.g., Max-Cut [82]) while in others it is quite different (e.g., k -Colorability [82]).
- It may be the case that the object is not very large and there is an efficient (polynomial-time) algorithm for solving the problem exactly. However, we may be interested in a *very* efficient (sublinear-time) algorithm, and are willing to tolerate the approximation/error it introduces.
- Finally, there are cases in which typical no-instances of the problem (that is, objects that do not have the property) are actually relatively far from having the property. In such cases we may first run the testing algorithm. If it rejects the object

then we reject it and otherwise we run the exact decision procedure. Thus, we save time on the typical no-instances. This is in particular useful if the testing algorithm has one-sided error so that it never rejects yes-instances (that have the property).

In all the aforementioned scenarios we are interested in testing algorithms that are much more efficient than the corresponding decision algorithms, and in particular have complexity that is sublinear in the size of the object.

1.2 Property Testing and Learning (Estimation)

Here when we say *learning* we mean outputting a good estimate of the target object.¹ Thus, another view of property testing is as a relaxation of learning (with queries and under the uniform distribution).² Namely, instead of asking that the algorithm output a good estimate of the function (object), which is assumed to belong to a particular class of functions \mathcal{F} , we only require that the algorithm decide whether the function belongs to \mathcal{F} or is far from any function in \mathcal{F} . Given this view, a natural motivation for property testing is to serve as a preliminary step before learning (and in particular, agnostic learning (e.g., [107]) where no assumption is made about the target function but the hypothesis should belong to a particular class of functions): we can first run the testing algorithm to decide whether to use a particular class of functions as our hypothesis class.

Here too we are interested in testing algorithms that are more efficient than the corresponding learning algorithms. As observed in [82], property testing is no harder than *proper* learning (where the learning algorithm is required to output a hypothesis from the same class of functions as the target function). Namely, if we have a proper learning

¹One may argue that property testing is also a certain form of learning as we learn information about the object (i.e., whether it has a certain property or is far from having the property). However, we have chosen to adopt the notion of learning usually used in the computational learning theory community.

²Testing under non-uniform distributions (e.g., [1, 92]) and testing with random examples (e.g., [105]) have been considered (and are discussed in this survey), but most of the work in property testing deals with testing under the uniform distributions and with queries.

4 Introduction

algorithm for a class of functions F then we can use it as a subroutine to test the property: “does the function belong to F ” (see Section 2.2 for a formal statement and proof).

Choosing between the two viewpoints. The choice of which of the aforementioned views to take is typically determined by the type of objects and properties in question. Much of property testing deals with combinatorial objects and in particular graphs. For such objects it is usually more natural to view property testing as a relaxation of exact decision. Indeed, there are many combinatorial properties for which there are testing algorithms that are much more efficient than the corresponding (exact) decision algorithms. On the other hand, when the objects are functions, then it is usually natural to look at property testing from a learning theory perspective. In some cases, both viewpoints are appropriate. This survey focuses on the latter perspective.

1.3 Property Testing and Hypothesis Testing

The notion of property testing is related to that of *hypothesis testing* (see e.g., [108, Chap. 8]) and indeed the distinction between estimation and testing is well known in the mathematical statistics literature. In this context, having the tested property (belonging to the corresponding class of objects) is called the *null hypothesis*, while being ϵ -far from the property (where ϵ is the distance parameter that the algorithm is given as input) is the *alternative hypothesis*. There are two major mathematical approaches to the study of testing in statistics (see, e.g., [136] and [113]). In the first, the alternative is taken to approach the null hypothesis at a certain rate as a function of the number of data points; when the correct rate is chosen the error probabilities stabilize at values strictly greater than zero and strictly less than one. In the second approach, the alternative is held fixed as the number of data points grows; in this case error probabilities go to zero and large deviation methods are used to assess the rate at which error probabilities go to zero. Aspects of both of these approaches can be found in the property testing literature.

While in many cases the particular problems studied in the property testing literature are somewhat different from those typically studied

in the mathematical statistics literature, the work on testing properties of distributions (which is discussed shortly in Section 6.3) deals with problems that are similar (or even the same) as those studied in mathematical statistics.

We also note that there are several works with a mathematical statistics flavor that are related to property testing and appeared in the computational learning literature (e.g., [33, 112, 137]).

1.4 Topics and Organization

We start with some preliminaries, which include basic definitions and notations. The preliminaries also include a precise statement and proof of the simple but important observation that testing is no harder than learning.

In Section 3, we consider the first type of properties that were studied in the context of property testing: algebraic properties. These include testing whether a function is (multi-)linear and more generally whether it is a polynomial of bounded degree. This work has implications to coding theory, and some of the results played an important role in the design of Probabilistically Check Proof (PCP) systems.

In Section 4, we turn to the study of function class that have a concise (propositional logic) representation such as singletons, monomials, and small DNF formula. This section includes a general result that applies to many classes of functions, where the underlying idea is that testing is performed by *implicit* learning.

The results in Sections 3 and 4 are in the *standard* model of testing. That is, the underlying distribution is uniform and the algorithm may perform queries to the function. In Section 5, we discuss distribution-free testing, and testing from random examples alone.

Finally, in Section 6, we give a more brief survey of other results in property testing. These include testing monotonicity, testing of clustering, testing properties of distributions, and more.

References

- [1] N. Ailon and B. Chazelle, “Information theory in property testing and monotonicity testing in higher dimensions,” *Information and Computation*, vol. 204, pp. 1704–1717, 2006.
- [2] N. Ailon, B. Chazelle, S. Comandur, and D. Liue, “Estimating the distance to a monotone function,” in *Proceedings of the Eight International Workshop on Randomization and Computation (RANDOM)*, pp. 229–236, 2004.
- [3] N. Alon, “Testing subgraphs of large graphs,” *Random Structures and Algorithms*, vol. 21, pp. 359–370, 2002.
- [4] N. Alon, A. Andoni, T. Kaufman, K. Matulef, R. Rubinfeld, and N. Xie, “Testing k -wise and almost k -wise independence,” in *Proceedings of the Thirty-Ninth Annual ACM Symposium on the Theory of Computing*, pp. 496–505, 2007.
- [5] N. Alon, S. Dar, M. Parnas, and D. Ron, “Testing of clustering,” *SIAM Journal on Discrete Math*, vol. 16, no. 3, pp. 393–417, 2003.
- [6] N. Alon, R. A. Duke, H. Lefmann, V. Rodl, and R. Yuster, “The algorithmic aspects of the regularity lemma,” *Journal of Algorithms*, vol. 16, pp. 80–109, 1994.
- [7] N. Alon, E. Fischer, M. Krivelevich, and M. Szegedy, “Efficient testing of large graphs,” *Combinatorica*, vol. 20, pp. 451–476, 2000.
- [8] N. Alon, E. Fischer, and I. Newman, “Testing of bipartite graph properties,” *SIAM Journal on Computing*, vol. 37, pp. 959–976, 2007.
- [9] N. Alon, E. Fischer, I. Newman, and A. Shapira, “A combinatorial characterization of the testable graph properties: It’s all about regularity,” in *Proceedings of the Thirty-Eighth Annual ACM Symposium on the Theory of Computing*, pp. 251–260, 2006.

94 References

- [10] N. Alon, T. Kaufman, M. Krivelevich, and D. Ron, "Testing triangle freeness in general graphs," in *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 279–288, 2006.
- [11] N. Alon and M. Krivelevich, "Testing k -colorability," *SIAM Journal on Discrete Math*, vol. 15, no. 2, pp. 211–227, 2002.
- [12] N. Alon, M. Krivelevich, T. Kaufman, S. Litsyn, and D. Ron, "Testing Reed-Muller codes," *IEEE Transactions on Information Theory*, vol. 51, no. 11, pp. 4032–4038, 2005. (An extended abstract of this paper appeared under the title: Testing Low-Degree Polynomials over $\text{GF}(2)$, in the proceedings of RANDOM 2003).
- [13] N. Alon, M. Krivelevich, I. Newman, and M. Szegedy, "Regular languages are testable with a constant number of queries," *SIAM Journal on Computing*, pp. 1842–1862, 2001.
- [14] N. Alon and A. Shapira, "Testing satisfiability," *Journal of Algorithms*, vol. 47, pp. 87–103, 2003.
- [15] N. Alon and A. Shapira, "Testing subgraphs in directed graphs," *Journal of Computer and System Sciences*, vol. 69, pp. 354–482, 2004.
- [16] N. Alon and A. Shapira, "A characterization of easily testable induced subgraphs," *Combinatorics Probability and Computing*, vol. 15, pp. 791–805, 2005.
- [17] N. Alon and A. Shapira, "A characterization of the (natural) graph properties testable with one-sided error," in *Proceedings of the Forty-Sixth Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 429–438, 2005.
- [18] N. Alon and A. Shapira, "Every monotone graph property is testable," in *Proceedings of the Thirty-Seventh Annual ACM Symposium on the Theory of Computing (STOC)*, pp. 129–137, 2005. (To appear in SICOMP).
- [19] N. Alon and A. Shapira, "Linear equations, arithmetic progressions and hypergraph property testing," *Theory of Computing*, vol. 1, pp. 177–216, 2005.
- [20] D. Angluin, "Queries and concept learning," *Machine Learning*, vol. 2, pp. 319–342, 1988.
- [21] S. Arora and S. Safra, "Probabilistic checkable proofs: A new characterization of NP," *Journal of the ACM*, vol. 45, no. 1, pp. 70–122, 1998. A preliminary version appeared in Proc. 33rd FOCS, 1992.
- [22] S. Arora and M. Sudan, "Improved low-degree testing and its applications," in *Proceedings of the Thirty-Second Annual ACM Symposium on the Theory of Computing (STOC)*, pp. 485–495, 1997.
- [23] L. Babai, L. Fortnow, L. Levin, and M. Szegedy, "Checking computations in polylogarithmic time," in *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing (STOC)*, pp. 21–31, 1991.
- [24] L. Babai, L. Fortnow, and C. Lund, "Non-deterministic exponential time has two-prover interactive protocols," *Computational Complexity*, vol. 1, no. 1, pp. 3–40, 1991.
- [25] T. Batu, S. Dasgupta, R. Kumar, and R. Rubinfeld, "The complexity of approximating the entropy," *SIAM Journal on Computing*, vol. 35, no. 1, pp. 132–150, 2005.
- [26] T. Batu, F. Ergun, J. Kilian, A. Magen, S. Raskhodnikova, R. Rubinfeld, and R. Sami, "A sublinear algorithm for weakly approximating edit distance," in

- Proceedings of the Thirty-Fifth Annual ACM Symposium on the Theory of Computing (STOC)*, pp. 316–324, 2003.
- [27] T. Batu, E. Fischer, L. Fortnow, R. Kumar, and R. Rubinfeld, “Testing random variables for independence and identity,” in *Proceedings of the Forty-Second Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 442–451, 2001.
- [28] T. Batu, L. Fortnow, R. Rubinfeld, W. Smith, and P. White, “Testing that distributions are close,” in *Proceedings of the Forty-First Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 259–269, 2000.
- [29] T. Batu, R. Kumar, and R. Rubinfeld, “Sublinear algorithms for testing monotone and unimodal distributions,” in *Proceedings of the Thirty-Sixth Annual ACM Symposium on the Theory of Computing (STOC)*, pp. 381–390, 2004.
- [30] T. Batu, R. Rubinfeld, and P. White, “Fast approximate PCPs for multi-dimensional bin-packing problems,” *Information and Computation*, vol. 196, no. 1, pp. 42–56, 2005.
- [31] A. Beimel, F. Bergadano, N. Bshouty, E. Kushilevitz, and S. Varricchio, “Learning functions represented as multiplicity automata,” *Journal of the ACM*, vol. 47, no. 5, pp. 506–530, 2000.
- [32] M. Bellare, D. Coppersmith, J. Håstad, M. Kiwi, and M. Sudan, “Linearity testing over characteristic two,” *IEEE Transactions on Information Theory*, vol. 42, no. 6, pp. 1781–1795, 1996.
- [33] S. Ben-David, “Can finite samples detect singularities of real-valued functions?,” in *Proceedings of the Twenty-Fourth Annual ACM Symposium on the Theory of Computing (STOC)*, pp. 390–399, 1992.
- [34] I. Ben-Eliezer, T. Kaufman, M. Krivelevich, , and D. Ron, “Comparing the strength of query types in property testing: The case of testing k -colorability,” in *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2008.
- [35] M. Ben-Or and P. Tiwari, “A deterministic algorithm for sparse multivariate polynomial interpolation,” in *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing (STOC)*, pp. 301–309, 1988.
- [36] E. Ben-Sasson, P. Harsha, and S. Raskhodnikova, “3CNF properties are hard to test,” *SIAM Journal on Computing*, vol. 35, no. 1, pp. 1–21, 2005.
- [37] M. Bender and D. Ron, “Testing properties of directed graphs: Acyclicity and connectivity,” *Random Structures and Algorithms*, pp. 184–205, 2002.
- [38] L. Bisht, N. Bshouty, and H. Mazzawi, “An optimal learning algorithm for multiplicity automata,” in *Proceedings of the Nineteenth Annual ACM Conference on Computational Learning Theory (COLT)*, pp. 184–198, 2006.
- [39] A. Blum, C. Burch, and J. Langford, “On learning monotone Boolean functions,” in *Proceedings of the Thirty-Ninth Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 408–415, 1998.
- [40] A. Blum, L. Hellerstein, and N. Littlestone, “Learning in the presence of finitely or infinitely many irrelevant attributes,” *Journal of Computer and System Sciences*, vol. 50, no. 1, pp. 32–40, 1995.
- [41] A. Blum and M. Singh, “Learning functions of k terms,” in *Proceedings of the Third Annual Workshop on Computational Learning Theory (COLT)*, pp. 144–153, 1990.

96 References

- [42] M. Blum, M. Luby, and R. Rubinfeld, “Self-Testing/Correcting with applications to numerical problems,” *Journal of the ACM*, vol. 47, pp. 549–595, 1993.
- [43] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth, “Occam’s razor,” *Information Processing Letters*, vol. 24, no. 6, pp. 377–380, April 1987.
- [44] A. Bogdanov, K. Obata, and L. Trevisan, “A lower bound for testing 3-colorability in bounded-degree graphs,” in *Proceedings of the Forty-Third Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 93–102, 2002.
- [45] A. Bogdanov and L. Trevisan, “Lower bounds for testing bipartiteness in dense graphs,” in *Proceedings of the Nintheenth Computational Complexity Conference (CCC)*, 2004.
- [46] B. Bollig, “Property testing and the branching program size,” in *Proceedings of FCT, Lecture notes in Computer Science 3623*, pp. 258–269, 2005.
- [47] B. Bollig and I. Wegener, “Functions that have read-once branching programs of quadratic size are not necessarily testable,” *Information Processing Letters*, vol. 87, no. 1, pp. 25–29, 2003.
- [48] C. Borgs, J. Chayes, L. Lovász, V. T. Sós, B. Szegedy, and K. Vesztegombi, “Graph limits and parameter testing,” in *Proceedings of the Thirty-Eighth Annual ACM Symposium on the Theory of Computing*, pp. 261–270, 2006.
- [49] N. Bshouty, “On learning multivariate polynomials under the uniform distribution,” *Information Processing Letters*, vol. 61, pp. 303–309, 1996.
- [50] N. Bshouty, J. Jackson, and C. Tamon, “More efficient PAC-learning of DNF with membership queries under the uniform distribution,” in *Proceedings of the Twelfth Annual ACM Conference on Computational Learning Theory (COLT)*, pp. 286–295, 1999.
- [51] N. Bshouty and Y. Mansour, “Simple learning algorithms for decision trees and multivariate polynomials,” *SIAM Journal on Computing*, vol. 31, no. 6, pp. 1909–1925, 2002.
- [52] B. Chazelle, R. Rubinfeld, and L. Trevisan, “Approximating the minimum spanning tree weight in sublinear time,” in *Automata, Languages and Programming: Twenty-Eighth International Colloquium (ICALP)*, pp. 190–200, 2001.
- [53] H. Chernoff, “A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations,” *Annals of the Mathematical Statistics*, vol. 23, pp. 493–507, 1952.
- [54] H. Chockler and D. Gutfreund, “A lower bound for testing juntas,” *Information Processing Letters*, vol. 90, no. 6, pp. 301–305, 2006.
- [55] M. Clausen, A. Dress, J. Grabmeier, and M. Karpinski, “On zero-testing and interpolation of k -sparse multivariate polynomials over finite fields,” *Theoretical Computer Science*, vol. 84, no. 2, pp. 151–164, 1991.
- [56] A. Czumaj, A. Shapira, and C. Sohler, “Testing hereditary properties of non-expanding bounded-degree graphs,” Technical Report TR07-083, Electronic Colloquium on Computational Complexity (ECCC), 2007. Available from <http://www.eccc.uni-trier.de/eccc/>.

- [57] A. Czumaj and C. Sohler, “Sublinear-time approximation for clustering via random sampling,” in *Automata, Languages and Programming: Thirty-First International Colloquium (ICALP)*, pp. 396–407, 2004.
- [58] A. Czumaj and C. Sohler, “Abstract combinatorial programs and efficient property testers,” *SIAM Journal on Computing*, vol. 34, no. 3, pp. 580–615, 2005.
- [59] A. Czumaj and C. Sohler, “Testing hypergraph colorability,” *Random Structures and Algorithms*, vol. 331, no. 1, pp. 37–52, 2005.
- [60] A. Czumaj and C. Sohler, “Testing expansion in bounded-degree graphs,” in *Proceedings of the Forty-Eighth Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 570–578, 2007.
- [61] I. Diakonikolas, H. K. Lee, K. Matulef, K. Onak, R. Rubinfeld, R. A. Servedio, and A. Wan, “Testing for concise representations,” in *Proceedings of the Forty-Eighth Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 549–557, 2007.
- [62] I. Diakonikolas, H. K. Lee, K. Matulef, R. A. Servedio, and A. Wan, “Efficient testing sparse GF(2) polynomials,” in *Automata, Languages and Programming: Thirty-Fifth International Colloquium (ICALP)*, pp. 502–514, 2008.
- [63] Y. Dodis, O. Goldreich, E. Lehman, S. Raskhodnikova, D. Ron, and A. Samorodnitsky, “Improved testing algorithms for monotonicity,” in *Proceedings of the Third International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*, pp. 97–108, 1999.
- [64] F. Ergun, S. Kannan, S. R. Kumar, R. Rubinfeld, and M. Viswanathan, “Spot-checkers,” *Journal of Computer and System Sciences*, vol. 60, no. 3, pp. 717–751, 2000.
- [65] S. Fattal and D. Ron, “Approximating the distance to monotonicity in high dimensions,” Manuscript, 2007.
- [66] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy, “Approximating Clique is almost NP-complete,” *Journal of the ACM*, pp. 268–292, 1996.
- [67] E. Fischer, “On the strength of comparisons in property testing,” *Information and Computation*, vol. 189, no. 1, pp. 107–116, 2004.
- [68] E. Fischer, “The difficulty of testing for isomorphism against a graph that is given in advance,” *SIAM Journal on Computing*, vol. 34, pp. 1147–1158, 2005.
- [69] E. Fischer, “Testing graphs for colorability properties,” *Random Structures and Algorithms*, vol. 26, no. 3, pp. 289–309, 2005.
- [70] E. Fischer and L. Fortnow, “Tolerant versus intolerant testing for Boolean properties,” *Theory of Computing*, vol. 2, pp. 173–183, 2006.
- [71] E. Fischer, G. Kindler, D. Ron, S. Safra, and S. Samorodnitsky, “Testing juntas,” *Journal of Computer and System Sciences*, vol. 68, no. 4, pp. 753–787, 2004.
- [72] E. Fischer, E. Lehman, I. Newman, S. Raskhodnikova, R. Rubinfeld, and A. Samorodnitsky, “Monotonicity testing over general poset domains,” in *Proceedings of the Thirty-Fourth Annual ACM Symposium on the Theory of Computing (STOC)*, pp. 474–483, 2002.
- [73] E. Fischer, A. Matsliah, and A. Shapira, “Approximate hypergraph partitioning and applications,” in *Proceedings of the Forty-Eighth Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 579–588, 2007.

98 References

- [74] E. Fischer and I. Newman, “Testing versus estimation of graph properties,” *SIAM Journal on Computing*, vol. 37, no. 2, pp. 482–501, 2007.
- [75] E. Fischer, I. Newman, and J. Sgall, “Functions that have read-twice constant width branching programs are not necessarily testable,” *Random Structures and Algorithms*, vol. 24, no. 2, pp. 175–193, 2004.
- [76] P. Fischer and H. U. Simon, “On learning ring-sum expansions,” *SIAM Journal on Computing*, vol. 21, no. 1, pp. 181–192, 1992.
- [77] K. Friedl and M. Sudan, “Some improvements to total degree tests,” in *Proceedings of the 3rd Annual Israel Symposium on Theory of Computing and Systems*, pp. 190–198, 1995. Corrected version available online at <http://theory.lcs.mit.edu/~madhu/papers/friedl.ps>.
- [78] P. Gemmell, R. Lipton, R. Rubinfeld, M. Sudan, and A. Wigderson, “Self testing/correcting for polynomials and for approximate functions,” in *Proceedings of the Thirty-Second Annual ACM Symposium on the Theory of Computing (STOC)*, pp. 32–42, 1991.
- [79] D. Glassner and R. A. Servedio, “Distribution-free testing lower bounds for basic Boolean functions,” in *Proceedings of the Eleventh International Workshop on Randomization and Computation (RANDOM)*, pp. 494–508, 2007.
- [80] O. Goldreich, “Short locally testable codes and proofs (a survey),” Technical Report TR05-014, Electronic Colloquium on Computational Complexity (ECCC), 2005. Available from <http://www.eccc.uni-trier.de/eccc/>.
- [81] O. Goldreich, S. Goldwasser, E. Lehman, D. Ron, and A. Samordinsky, “Testing monotonicity,” *Combinatorica*, vol. 20, no. 3, pp. 301–337, 2000.
- [82] O. Goldreich, S. Goldwasser, and D. Ron, “Property testing and its connection to learning and approximation,” *Journal of the ACM*, vol. 45, no. 4, pp. 653–750, 1998.
- [83] O. Goldreich and D. Ron, “A sublinear bipartite tester for bounded degree graphs,” *Combinatorica*, vol. 19, no. 3, pp. 335–373, 1999.
- [84] O. Goldreich and D. Ron, “On testing expansion in bounded-degree graphs,” *Electronic Colloquium on Computational Complexity*, vol. 7, no. 20, 2000.
- [85] O. Goldreich and D. Ron, “Property testing in bounded degree graphs,” *Algorithmica*, pp. 302–343, 2002.
- [86] O. Goldreich and L. Trevisan, “Three theorems regarding testing graph properties,” *Random Structures and Algorithms*, vol. 23, no. 1, pp. 23–57, 2003.
- [87] S. Goldwasser and S. Micali, “Probabilistic encryption,” *Journal of Computer and System Sciences*, vol. 28, no. 2, pp. 270–299, 1984.
- [88] M. Gonen and D. Ron, “On the benefits of adaptivity in property testing of dense graphs,” in *Proceedings of the Eleventh International Workshop on Randomization and Computation (RANDOM)*, pp. 525–537, 2007.
- [89] D. Y. Grigoriev, M. Karpinski, and M. F. Singer, “Fast parallel algorithms for sparse multivariate polynomial interpolation over finite fields,” *SIAM Journal on Computing*, vol. 19, no. 6, pp. 1059–1063, 1990.
- [90] D. Guijarro, J. Tarui, and T. Tsukiji, “Finding relevant variables in PAC model with membership queries,” in *Proceedings of Algorithmic Learning Theory, 10th International Conference*, pp. 313–322, 1999.

- [91] V. Guruswami and A. Rudra, "Tolerant locally testable codes," in *Proceedings of the Ninth International Workshop on Randomization and Computation (RANDOM)*, pp. 306–317, 2005.
- [92] S. Halevy and E. Kushilevitz, "Distribution-free property testing," in *Proceedings of the Seventh International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*, pp. 341–353, 2003.
- [93] S. Halevy and E. Kushilevitz, "Distribution-free connectivity testing," in *Proceedings of the Eight International Workshop on Randomization and Computation (RANDOM)*, pp. 393–404, 2004.
- [94] S. Halevy and E. Kushilevitz, "A lower bound for distribution-free monotonicity testing," in *Proceedings of the Ninth International Workshop on Randomization and Computation (RANDOM)*, pp. 330–341, 2005.
- [95] S. Halevy and E. Kushilevitz, "Distribution-free property testing," *SIAM Journal on Computing*, vol. 37, no. 4, pp. 1107–1138, 2007.
- [96] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13–30, March 1963.
- [97] P. Indyk, "A sublinear-time approximation scheme for clustering in metric spaces," in *Proceedings of the Fortieth Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 154–159, 1999.
- [98] J. Jackson, "An efficient membership-query algorithm for learning DNF with respect to the uniform distribution," *Journal of Computer and System Sciences*, vol. 55, pp. 414–440, 1997.
- [99] C. S. Jutla, A. C. Patthak, A. Rudra, and D. Zuckerman, "Testing low-degree polynomials over prime fields," in *Proceedings of the Forty-Fifth Annual Symposium on Foundations of Computer Science (FOCS)*, 2004.
- [100] S. Kale and C. Seshadhri, "Testing expansion in bounded degree graphs," Technical Report TR07-076, Electronic Colloquium on Computational Complexity (ECCC), 2007. Available from <http://www.eccc.uni-trier.de/eccc/>.
- [101] T. Kaufman, M. Krivelevich, and D. Ron, "Tight bounds for testing bipartiteness in general graphs," *SIAM Journal on Computing*, vol. 33, no. 6, pp. 1441–1483, 2004.
- [102] T. Kaufman, S. Litsyn, and N. Xie, "Breaking the ϵ -soundness bound of the linearity test over $\text{GF}(2)$," Technical Report TR07-098, Electronic Colloquium on Computational Complexity (ECCC), 2007. Available from <http://www.eccc.uni-trier.de/eccc/>.
- [103] T. Kaufman and D. Ron, "Testing polynomials over general fields," *SIAM Journal on Computing*, vol. 35, no. 3, pp. 779–802, 2006.
- [104] M. Kearns, M. Li, and L. Valiant, "Learning Boolean formulae," *Journal of the ACM*, vol. 41, no. 6, pp. 1298–1328, 1995.
- [105] M. Kearns and D. Ron, "Testing problems with sub-learning sample complexity," *Journal of Computer and System Sciences*, vol. 61, no. 3, pp. 428–456, 2000.
- [106] M. Kearns and L. Valiant, "Cryptographic limitations on learning boolean formulae and finite automata," *Journal of the ACM*, vol. 41, no. 1, pp. 67–95, 1994.

100 *References*

- [107] M. J. Kearns, R. E. Schapire, and L. M. Sellie, "Toward efficient agnostic learning," *Machine Learning*, vol. 17, nos. 2–3, pp. 115–141, 1994.
- [108] J. C. Kiefer, *Introduction to Statistical Inference*. Springer Verlag, 1987.
- [109] A. Klivans and R. A. Servedio, "Boosting and hard-core sets," in *Proceedings of the Fortieth Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 624–633, 1999.
- [110] Y. Kohayakawa, B. Nagle, and V. Rödl, "Efficient testing of hypergraphs," in *Automata, Languages and Programming: Twenty-Ninth International Colloquium (ICALP)*, pp. 1017–1028, 2002.
- [111] S. Kulkarni, S. Mitter, and J. Tsitsiklis, "Active learning using arbitrary binary valued queries," *Machine Learning*, vol. 11, pp. 23–35, 1993.
- [112] S. R. Kulkarni and O. Zeitouni, "On probably correct classification of concepts," in *Proceedings of the Sixth Annual ACM Conference on Computational Learning Theory (COLT)*, pp. 111–116, 1993.
- [113] E. K. Lehman and J. P. Romano, *Testing Statistical Hypotheses*. Springer Verlag, third edition, 2005.
- [114] N. Littlestone, "Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm," *Machine Learning*, vol. 2, no. 4, pp. 285–318, 1987.
- [115] Y. Mansour, "Randomized interpolation and approximation of sparse polynomials," *SIAM Journal on Computing*, vol. 24, no. 2, pp. 357–368, 1995.
- [116] S. Marko and D. Ron, "Distance approximation in bounded-degree and general sparse graphs," in *Proceedings of the Tenth International Workshop on Randomization and Computation (RANDOM)*, pp. 475–486, 2006. (To appear in *Transactions on Algorithms*).
- [117] K. Matulef, R. O'Donnell, R. Rubinfeld, and R. A. Servedio, "Testing Half-spaces," Report number 128 in the Electronic Colloquium on Computational Complexity (ECCC), 2007.
- [118] N. Mishra, D. Oblinger, and L. Pitt, "Sublinear time approximate clustering," in *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 439–447, 2001.
- [119] E. Mossel, R. O'Donnell, and R. A. Servedio, "Learning juntas," *Journal of Computer and System Sciences*, vol. 69, no. 3, pp. 421–434, 2004.
- [120] A. Nachmias and A. Shapira, "Testing the expansion of a graph," Technical Report TR07-118, Electronic Colloquium on Computational Complexity (ECCC), 2007. Available from <http://www.eccc.uni-trier.de/eccc/>.
- [121] I. Newman, "Testing membership in languages that have small width branching programs," *SIAM Journal on Computing*, vol. 31, no. 5, pp. 1557–1570, 2002.
- [122] M. Parnas and D. Ron, "Testing the diameter of graphs," *Random Structures and Algorithms*, vol. 20, no. 2, pp. 165–183, 2002.
- [123] M. Parnas, D. Ron, and R. Rubinfeld, "On testing convexity and submodularity," *SIAM Journal on Computing*, vol. 32, no. 5, pp. 1158–1184, 2003.
- [124] M. Parnas, D. Ron, and R. Rubinfeld, "Tolerant property testing and distance approximation," *Journal of Computer and System Sciences*, vol. 72, no. 6, pp. 1012–1042, 2006.

- [125] M. Parnas, D. Ron, and A. Samorodnitsky, “Testing basic Boolean formulae,” *SIAM Journal on Discrete Math*, vol. 16, no. 1, pp. 20–46, 2002.
- [126] S. Raskhodnikova, D. Ron, R. Rubinfeld, and A. Smith, “Strong lower bounds for approximating distributions support size and the distinct elements problem,” in *Proceedings of the Forty-Eighth Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 559–568, 2007.
- [127] R. Roth and G. Benedek, “Interpolation and approximation of sparse multivariate polynomials over $\text{GF}(2)$,” *SIAM Journal on Computing*, vol. 20, no. 2, pp. 291–314, 1991.
- [128] R. Rubinfeld and M. Sudan, “Robust characterization of polynomials with applications to program testing,” *SIAM Journal on Computing*, vol. 25, no. 2, pp. 252–271, 1996.
- [129] R. E. Schapire and L. M. Sellie, “Learning sparse multivariate polynomials over a field with queries and counterexamples,” *Journal of Computer and System Sciences*, vol. 52, no. 2, pp. 201–213, 1996.
- [130] A. Shpilka and A. Wigderson, “Derandomizing homomorphism testing in general groups,” in *Proceedings of the Thirty-Sixth Annual ACM Symposium on the Theory of Computing (STOC)*, pp. 427–435, 2004.
- [131] M. Sudan, “Efficient checking of polynomials and proofs and the hardness of approximation problems,” PhD thesis, UC Berkeley, 1992. Also appears as Lecture Notes in Computer Science, Vol. 1001, Springer, 1996.
- [132] E. Szemerédi, “Regular partitions of graphs,” in *Proceedings, Colloque Inter. CNRS*, pp. 399–401, 1978.
- [133] R. Uehara, K. Tsuchida, and I. Wegener, “Optimal attribute-efficient learning of disjunction, parity and threshold functions,” in *Proceedings of the 3rd European Conference on Computational Learning Theory*, pp. 171–184, 1997.
- [134] L. G. Valiant, “A theory of the learnable,” *CACM*, vol. 27, no. 11, pp. 1134–1142, November 1984.
- [135] P. Valiant, “Testing symmetric properties of distributions,” in *Proceedings of the Fourtieth Annual ACM Symposium on the Theory of Computing*, pp. 383–392, 2008.
- [136] A. Van der Vaart, *Asymptotic Statistics*. Cambridge University Press, 1998.
- [137] K. Yamanishi, “Probably almost discriminative learning,” *Machine Learning*, vol. 18, pp. 23–50, 1995.
- [138] R. Zippel, “Probabilistic algorithms for sparse polynomials,” in *Proceedings of the International Symposium on Symbolic and Algebraic Computation*, pp. 216–226, 1978.
- [139] R. Zippel, “Interpolating polynomials from their values,” *Journal of Symbolic Computation*, vol. 9, pp. 375–403, 1990.