Introduction: Three Questions

Computational limitations are simultaneously frustrating and useful; we are dismayed by the apparent difficulty of learning, but we need hard problems for secure cryptography. I study *computational complexity theory* to achieve useful outcomes — such as learning algorithms and provably correct pseudorandom number generators — by understanding and exploiting the limits of efficient computation. I seek formal answers to three basic questions:

(Q1) How are efficient algorithms and complexity limits related?
(Q2) Which natural phenomena can be efficiently and convincingly simulated by computation?
(Q3) What rich properties (e.g., privacy, fairness, transparency) can algorithm designers enforce?

My general approach is linguistic. By analyzing the languages used to describe algorithms and write proofs, we can study and manipulate broad classes of *arguments* and *algorithms* simultaneously. This "bulk" treatment of formal objects was key to my results in structural complexity [8], pseudorandomness [6], and computational learning theory [5, 4]. For each question above, I will provide background, sketch my accomplishments so far, and describe future directions.

Background & Past Work

(Q1) Duality between Meta-Computation & Complexity Lower Bounds

Background: Complexity theorists seek the limits of efficient computation. They try to *lower bound* the amount of computational resources (such as space, time, or random bits) needed to solve important and natural problems. Algorithm designers try to develop efficient algorithms for problems, giving *upper bounds* on the computational resources required for these tasks. These goals appear opposed. Yet a complexity lower bound also contains information about efficient computation; it must somehow analyze *all* efficient computations of some form to prove a limitation on them. When can these analyses themselves be made algorithmic? What algorithmic problems would they solve?

Meta-computational problems, where the inputs may be computations themselves, have deep connections to complexity lower bounds. To define some meta-computations, fix a *complexity class* $\mathfrak{C}$: a set of Boolean functions that can be computed within certain resource constraints. We often define complexity classes by restricting the structure and size of *logical circuits* — devices composed of small "gates," each computing a simple function like AND, OR, NOT. Two key meta-computational problems are:

- *Learning* ($\mathfrak{C}$-LEARN): given only the ability to query a Boolean function f from $\mathfrak{C}$, output a small circuit approximating f.
- *Minimum Circuit Size Problem*, ($\mathfrak{C}$-MCSP): given the complete truth table of a Boolean function f and a number s, is there a $\mathfrak{C}$-circuit of size less than s that computes f?

Intuitively, there are connections between meta-computation and lower bounds because both endeavors require a deep understanding of *all* $\mathfrak{C}$-functions: either to automatically analyze them, or to prove that some "hard" function is not in $\mathfrak{C}$. My work has produced new learning algorithms from novel reductions between $\mathfrak{C}$-LEARN and $\mathfrak{C}$-MCSP.

Learning Algorithms from Complexity Lower Bounds. Most complexity lower bounds fit the *Natural Proofs* framework [29]. Natural Proofs against a complexity class $\mathfrak{C}$ implicitly contain an algorithm that distinguishes between truth-tables of functions computable by $\mathfrak{C}$-circuits and random strings; that is, they solve an approximate version of the $\mathfrak{C}$-MCSP problem. Razborov and Rudich showed (under widely-believed cryptographic assumptions) that Natural Proofs cannot separate P from NP. Thus, the ability of the Natural Proofs framework to capture most known complexity lower bounds is generally cited as a barrier to separating P from NP.

My joint work [5] shows that the Natural Proofs framework is also a powerful tool for meta-computational algorithm design. We gave a generic construction of $\mathfrak{C}$-LEARN algorithms from Natural Proofs against a class $\mathfrak{C}$, for any $\mathfrak{C}$ satisfying a mild technical condition. Our construction immediately yielded the first non-trivial (quasi-polynomial time) learning algorithms for $\mathsf{AC}^0[p]$, the class of functions computed by constant-depth circuits of polynomial size with AND, OR, NOT and counting modulo p gates (where p is a prime). Obtaining any kind of learning algorithm for $\mathsf{AC}^0[p]$ was open for 23 years. Previous work gave a learning algorithm for AC^0 (the same class without counting modulo p) in 1993 [23]. Our work received the 31st Annual

Computational Complexity Conference Best Paper award. Our main construction has been applied to make progress on questions about learning, complexity, and cryptography [27, 16, 28, 14].

While prior work has also developed algorithms from complexity lower bounds, ours is a rare example of a generic construction. All known algorithms for Boolean meta-computation were developed by adapting proof techniques originally intended for complexity lower bounds [17, 23, 9]. Those results manually inspect the lower bound proof to uncover a concrete weakness of the target complexity class, and then exploit it for algorithm design. Our learning algorithm is almost completely generic: it uses the Natural Proof against $\mathfrak{C}$ as a black box.

This genericity in our work suggests a structural relationship between proving circuit lower bounds and constructing algorithms. I want to develop a new algorithm design paradigm tailored to meta-computation, that "automatically" exploits complexity lower bounds. Progress towards this end is ongoing: in subsequent work with the same co-authors, I showed that many Natural Proofs can actually distinguish between random strings and strings merely close to the truth-table of a $\mathfrak{C}$-function. This enables our learning algorithms (once suitably modified) to tolerate adversarially corrupted answers to queries, obtaining the first *agnostic* learning algorithms for $\mathsf{AC}^0[p]$ in [4].

My research program uses meta-computation as a "Rosetta Stone" to translate algorithmic questions into complexity lower bound questions, and vice-versa. Such translations have already resolved longstanding open problems in both complexity theory and algorithm design, in my previous work and elsewhere. I believe this is no accident, and that understanding the nature of the "Stone" itself will illuminate the mathematical foundations of efficient computation.

(Q2) Pseudorandomness: A Durable Illusion from Computational Limitations

Background: One of the richest areas in complexity theory is *pseudorandomness*: the study of how to simulate perfectly uniform random bits. These investigations established many "hardness versus randomness trade-offs," constructions that use hard functions to build pseudorandom sources [26]. Intuitively, hard functions allow us to show computationally bounded agents patterns that they cannot "understand." These patterns, to the limited observer, may as well be random noise. This is a striking example of the utility of complexity lower bounds. If we could prove strong enough circuit complexity lower bounds, we could transform any efficient randomized algorithm into an efficient deterministic algorithm [18].

This idea of "fooling" algorithms can be pushed further. *Average-case complexity* relaxes the definition of correctness: an algorithm A is correct "on average" if no efficient adversary can generate "bad" inputs that cause A to err. Average-case complexity formalizes the idea that, if efficient adversaries cannot tell the difference between a perfect algorithm A and mostly-correct A', we can substitute A' for A and get away with it under reasonable circumstances, parameterized by how carefully our enemies pay attention [22, 20].

The combination of average-case complexity with pseudorandomness is fruitful. The hardness to randomness results cited above require strong *circuit* complexity lower bounds. What if we started from strictly weaker complexity lower bounds about Turing Machines? It turns out that Turing Machine lower bounds imply *average-case* derandomization [19]! Since complexity separations for Turing Machines seem easier than circuit lower bounds, this brings us closer to the ultimate goal of total derandomization — or at least, the *convincing illusion* of total derandomization. If we believe reality is populated by computationally efficient phenomena, then this "pseudo-derandomization" is just as good. Starting from popular conjectures about the Randomized Turing Machine complexity of certain "key" problems, I developed dramatically more efficient average-case derandomizations than were previously known [6].

Popular Fine-Grained Hardness Conjectures Imply *Efficient* Derandomization. It was known that *weak* uniform hardness implies *non-trivial* derandomization: if *polynomial time* randomized computation (BPP) is less powerful than *exponential time* deterministic computation (EXP), then BPP can be simulated in deterministic sub-exponential time $O(2^{n^\epsilon})$ for every constant $\epsilon > 0$ [19, 31]. Note that any sub-exponential time simulation of BPP is non-trivial, and this does not follow easily from $\mathsf{BPP} \neq \mathsf{EXP}$. We know $\mathsf{BPP} \subseteq \mathsf{EXP}$ trivially, by brute-forcing the random coins of any probabilistic algorithm and explicitly computing the acceptance probability on a given input. Before [19], nothing was known about the implications of uniform separations like $\mathsf{EXP} \neq \mathsf{BPP}$ for non-trivial derandomization.

My joint work obtained *efficient* (polynomial-time) simulations of BPP from some popular and well-studied uniform hardness assumptions of "fine-grained" complexity theory, such as the k-Orthogonal-Vectors conjecture. These strong but widely believed hypotheses assert that solving simple families of combinatorial

problems inside P require brute-force search over a polynomial-sized domain [33]. Our work converted these problem-centric assumptions into structural conclusions about the power (or weakness) of randomness as a computational resource. I continue to work on these topics, aiming to study and characterize phenomena X for which efficient pseudo-X constructions fool computationally bounded entities.

(Q3) Algorithmic Desiderata Beyond Correctness & Efficiency

Background: Algorithms are now deployed in contexts where simple correctness and efficiency guarantees are not sufficient. For example, we want to learn from and release information about highly-sensitive data — such as medical records — without compromising the privacy of individuals who participate in these studies. Recently, several related desiderata have been studied intensively:

- Differential Privacy: Algorithms *should not* leak personal information about individuals whose data they operate on. [11]
- Fairness in Classification: Algorithms *should not* discriminate based on sensitive attributes such as race, gender, and ability status. [10]
- Transparency in Modeling: Algorithms *should* be explicable to those who use, regulate, or are affected by them. [24]

I approach enforcing these desiderata as problems of language design: we can hope to design constrained frameworks where any correct algorithm automatically satisfies useful properties such as those listed above. This transforms the task of analyzing algorithms for useful properties into one of language design, and introduces a "dual" problem: what computational primitives and laws of combination naturally give rise to especially "well-behaved" algorithms? Aside from theoretical interest, provable linguistic constraints could make it easier to enforce these desiderata in real-world systems. There is reason to be optimistic: *locally* private learning is equivalent to giving a learning algorithm in the *statistical query* model [21]. In the past, I worked on interpretable models for the Mars Curiosity mission. Currently, I am working on transparent models of student grades [7] and private learning algorithms for large-margin halfspaces.

Sparse & Explicable Models for Mars Curiosity. I know from personal experience that it is crucial to understand the *context* in which algorithms are used, and that human cognition and conversation is often the most expensive resource during deployment.

The ChemCam on Mars Curiosity is a Laser-Induced Breakdown Spectroscopy (LIBS) instrument: it fires a laser at samples to create plasma, and images that plasma with a telescoped spectrometer. Because each element has characteristic spectrochemical emissions, we can infer the elemental composition of remote samples using the data collected by a LIBS. This instrument is currently deployed on Mars. I worked with Professor Darby Dyar at Mount Holyoke, a participating scientist on the Curiosity mission, to develop and deploy statistical models for this composition-inference task [13, 12, 2].

I communicated intensively with scientists to understand the ChemCam instrument. We worked together to create accurate and transparent models whose predictions were explicable in terms of underlying physics. This process was ultimately successful, but time-consuming and difficult. I now study theoretical foundations for *automatically* enforcing transparency, privacy, and other rich desiderata in modeling tasks.

Work in Progress & Future Directions

(Q1) Circuit Synthesis vs Top-Down Complexity Lower Bounds

The Natural Proofs framework is a powerful tool because it gives an *algorithmic* characterization of many circuit lower bound techniques. However, Natural Proofs seem to capture only "bottom-up" approaches: ideas that would separate larger and larger sub-classes of functions computable by polynomial-sized circuits ($\mathsf{P/poly}$) from NP. We could instead approach $\mathsf{NP} \neq \mathsf{P/poly}$ "top-down:" start by separating a huge super-class of NP from $\mathsf{P/poly}$, and then separate smaller and smaller super-classes of NP from $\mathsf{P/poly}$ [32, 25].

No framework similar to Natural Proofs is known for top-down arguments. However, top-down arguments often embed efficient algorithms for circuit *synthesis*: the problem of printing small circuits for a "key" language L, if such circuits exist. I hope to use this pattern to give a comprehensive and algorithmic characterization of top-down circuit lower bounds. Two directions would follow:

1. *Barriers:* if the embedded algorithms contradict longstanding conjectures (as in the case of Natural Proofs) we would know to search for new and "unnatural" top-down techniques.

2. *Progress:* if the embedded algorithms are plausible, we can direct intensive efforts towards their construction and hopefully achieve unconditional progress on circuit lower bounds.

(Q2) Human Resources for Teaching (and Beyond)

As Computer Science courses scale up to teach more and more students, it is vital to maintain consistent grading and high-quality feedback. This is especially important when students work rich problems whose responses interleave prose with formalism; these problems cannot be effectively auto-graded because they are too complex. So, I work to optimize the cognitive, *human* resources spent on grading.

As a Teaching Assistant at UCSD, I co-developed an adaptive rubric creation method for such rich responses [7]. Briefly, the TA executes a two-pass "streaming algorithm" over randomly-ordered responses: the first pass explores a sub-sample to discover common patterns, and the second pass assigns grades and feedback by recognizing those patterns. The algorithm aims to minimize the number of times the TA switches tasks (because task-switching is know to impair human cognitive performance) while helping the TA to consistently justify scores across many (hundreds) of distinct student responses.

By what objective standards could we possibly assess such an algorithm? I propose an *indistinguishability criterion,* inspired by pseudo-randomness. In small classes, students would receive "truly personalized" ad-hoc feedback on their assignments. Our algorithm tries to produce *pseudo-personalized* feedback that is just as useful, by leveraging common patterns in student work. Can students tell the difference? This is an empirical question which my future work will seek to answer using randomized controlled trials.

The "human resources" framework used to develop our grading system can be immediately generalized. Any iterative process where humans must issue justified decisions about rich inputs can be assessed using our distinguishability criterion. Such tasks include: social media moderation, natural language dataset labeling, and resume screening. In future work, I hope to establish a firm theoretical foundation for the design and assessment of algorithms that use human "components" to efficiently address rich problems.

(Q3) Privacy (and Beyond) in Learning via Boosting

Boosting is a fundamental technique in both the theory and practice of machine learning for converting weak learning algorithms into strong ones. A (randomized) learning algorithm is differentially private if the distribution on hypotheses it produces does not depend too much on any single input sample. It is natural to design "private boosting" algorithms, to transform differentially private weak learners into differentially private strong learners.

In preliminary joint work, I designed a framework for privatizing boosting algorithms [3]. Our framework is expressive enough to privatize many smooth boosting algorithms [30, 15, 1]. We applied our framework to develop a noise-tolerant and private learner for large-margin halfspaces, whose sample complexity does not depend on the dimension of the domain. Previous dimension-free learning for halfspaces could not tolerate noise, and used a completely different algorithmic approach.

Boosting is an *ensemble* method: it produces a committee of weak learners whose collective opinion is accurate. A deep theory has been developed to explain the success of boosting. This foundation could give us sufficient leverage to enforce rich desiderata beyond privacy. For example, by analyzing *how* boosting creates a committee, I hope to both justify the predictions of the committee in terms of sensible features of the data and ensure adaptive validity of analyses. Since Boosting has also been enormously successful in practice, this work offers many opportunities for collaboration and experimentation.

Conclusion: Three More Questions

Linguistic expressibility is simultaneously useful and frustrating; we need rich languages to reason about complex problem domains, but are dismayed by the apparent difficulty of analyzing and managing these powerful tools. I study *the expressive power of languages* to achieve positive outcomes by enforcing "simplicity" in algorithms and proofs. I seek formal answers to three basic questions, derived from those above:

(Q1') What theorems can we prove using "simple" reasoning systems?
(Q2') Which phenomena can be succinctly described to computationally-bounded agents?
(Q3') How can we build languages that correspond to natural paradigms of algorithm design?

These questions are particularly relevant in the context of *computational complexity theory*; by attending to issues of expressibility, I have made progress in both upper and lower bounds on computational resources. My future work will test and develop this methodology against the widest possible range of problems.

[1] Boaz Barak, Moritz Hardt, and Satyen Kale. "The uniform hardcore lemma via approximate Bregman projections". In: *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*. Ed. by Claire Mathieu. SIAM, 2009, pp. 1193–1200. URL: http://dl.acm.org/citation.cfm?id=1496770.1496899.

[2] Thomas F. Boucher, Marie V. Ozanne, Marco L. Carmosino, M. Darby Dyar, Sridhar Mahadevan, Elly A. Breves, Kate H. Lepore, and Samuel M. Clegg. "A study of machine learning regression methods for major elemental analysis of rocks using laser-induced breakdown spectroscopy". In: *Spectrochimica Acta Part B: Atomic Spectroscopy* 107 (2015), pp. 1–10. ISSN: 0584-8547. DOI: https://doi.org/10.1016/j.sab.2015.02.003. URL: http://www.sciencedirect.com/science/article/pii/S0584854715000518.

[3] Mark Bun, Marco L. Carmosino, and Jessica Sorrell. *Private Boosting and Noise-tolerant Learning*. presented by Mark Bun at Simons Institute. 2019. URL: https://simons.berkeley.edu/talks/private-boosting-and-noise-tolerant-learning.

[4] Marco L. Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. "Agnostic Learning from Tolerant Natural Proofs". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2017, August 16-18, 2017, Berkeley, CA, USA*. Ed. by Klaus Jansen, José D. P. Rolim, David Williamson, and Santosh Srinivas Vempala. Vol. 81. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017, 35:1–35:19. DOI: 10.4230/LIPIcs.APPROX-RANDOM.2017.35. URL: https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2017.35.

[5] Marco L. Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. "Learning Algorithms from Natural Proofs". In: *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*. Ed. by Ran Raz. Vol. 50. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016, 10:1–10:24. DOI: 10.4230/LIPIcs.CCC.2016.10. URL: https://doi.org/10.4230/LIPIcs.CCC.2016.10.

[6] Marco L. Carmosino, Russell Impagliazzo, and Manuel Sabin. "Fine-Grained Derandomization: From Problem-Centric to Resource-Centric Complexity". In: *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*. Ed. by Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella. Vol. 107. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018, 27:1–27:16. DOI: 10.4230/LIPIcs.ICALP.2018.27. URL: https://doi.org/10.4230/LIPIcs.ICALP.2018.27.

[7] Marco L. Carmosino and Mia Minnes. "Adaptive Rubrics". In: *SIGCSE 2020 – TO APPEAR*. ACM, 2020.

[8] Marco Carmosino, Russell Impagliazzo, Valentine Kabanets, and Antonina Kolokolova. "Tighter Connections between Derandomization and Circuit Lower Bounds". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*. Ed. by Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim. Vol. 40. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015, pp. 645–658. DOI: 10.4230/LIPIcs.APPROX-RANDOM.2015.645. URL: https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2015.645.

[9] Ruiwen Chen, Valentine Kabanets, Antonina Kolokolova, Ronen Shaltiel, and David Zuckerman. "Mining Circuit Lower Bound Proofs for Meta-Algorithms". In: *Computational Complexity* 24.2 (2015), pp. 333–392. DOI: 10.1007/s00037-015-0100-0. URL: http://dx.doi.org/10.1007/s00037-015-0100-0.

[10] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. "Fairness through awareness". In: *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*. Ed. by Shafi Goldwasser. ACM, 2012, pp. 214–226. DOI: 10.1145/2090236.2090255. URL: https://doi.org/10.1145/2090236.2090255.

[11] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. "Calibrating Noise to Sensitivity in Private Data Analysis". In: *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*. Ed. by Shai Halevi and Tal Rabin. Vol. 3876. Lecture Notes in Computer Science. Springer, 2006, pp. 265–284. DOI: 10.1007/11681878_14. URL: https://doi.org/10.1007/11681878%5C_14.

[12] M.D. Dyar, M.L. Carmosino, E.A. Breves, M.V. Ozanne, S.M. Clegg, and R.C. Wiens. "Comparison of partial least squares and lasso regression techniques as applied to laser-induced breakdown spectroscopy of geological samples". In: *Spectrochimica Acta Part B: Atomic Spectroscopy* 70 (2012), pp. 51–67. ISSN: 0584-8547. DOI: https://doi.org/10.1016/j.sab.2012.04.011. URL: http://www.sciencedirect.com/science/article/pii/S058485471200095X.

[13] M.D. Dyar, M.L. Carmosino, J.M. Tucker, E.A. Brown, S.M. Clegg, R.C. Wiens, J.E. Barefield, J.S. Delaney, G.M. Ashley, and S.G. Driese. "Remote laser-induced breakdown spectroscopy analysis of East African Rift sedimentary samples under Mars conditions". In: *Chemical Geology* 294-295 (2012), pp. 135–151. ISSN: 0009-2541. DOI: https://doi.org/10.1016/j.chemgeo.2011.11.019. URL: http://www.sciencedirect.com/science/article/pii/S0009254111004554.

[14] Shuichi Hirahara. "Non-black-box Worst-case to Average-case Reductions within NP". In: *Electronic Colloquium on Computational Complexity (ECCC)* 25 (2018), p. 138. URL: https://eccc.weizmann.ac.il/report/2018/138.

[15] Thomas Holenstein. "Key agreement from weak bit agreement". In: *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*. Ed. by Harold N. Gabow and Ronald Fagin. ACM, 2005, pp. 664–673. DOI: 10.1145/1060590.1060689. URL: https://doi.org/10.1145/1060590.1060689.

[16] Russell Impagliazzo, Valentine Kabanets, and Ilya Volkovich. "The Power of Natural Properties as Oracles". In: *33rd Computational Complexity Conference, CCC 2018, June 22-24, 2018, San Diego, CA, USA*. Ed. by Rocco A. Servedio. Vol. 102. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018, 7:1–7:20. DOI: 10.4230/LIPIcs.CCC.2018.7. URL: https://doi.org/10.4230/LIPIcs.CCC.2018.7.

[17] Russell Impagliazzo, William Matthews, and Ramamohan Paturi. "A satisfiability algorithm for AC^0". In: *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*. 2012, pp. 961–972. URL: http://portal.acm.org/citation.cfm?id=2095193&CFID=63838676&CFTOKEN=79617016.

[18] Russell Impagliazzo and Avi Wigderson. "$P = BPP$ if E Requires Exponential Circuits: Derandomizing the XOR Lemma". In: *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*. 1997, pp. 220–229. DOI: 10.1145/258533.258590. URL: http://doi.acm.org/10.1145/258533.258590.

[19] Russell Impagliazzo and Avi Wigderson. "Randomness vs Time: Derandomization under a Uniform Assumption". In: *J. Comput. Syst. Sci.* 63.4 (2001), pp. 672–688. DOI: 10.1006/jcss.2001.1780. URL: http://dx.doi.org/10.1006/jcss.2001.1780.

[20] Valentine Kabanets. "Easiness Assumptions and Hardness Tests: Trading Time for Zero Error". In: *Proceedings of the 15th Annual IEEE Conference on Computational Complexity, Florence, Italy, July 4-7, 2000*. IEEE Computer Society, 2000, pp. 150–157. DOI: 10.1109/CCC.2000.856746. URL: http://dx.doi.org/10.1109/CCC.2000.856746.

[21] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. "What Can We Learn Privately?" In: *SIAM J. Comput.* 40.3 (2011), pp. 793–826. DOI: 10.1137/090756090. URL: https://doi.org/10.1137/090756090.

[22] Leonid A. Levin. "Average Case Complete Problems". In: *SIAM J. Comput.* 15.1 (1986), pp. 285–286. DOI: 10.1137/0215020. URL: https://doi.org/10.1137/0215020.

[23] Nathan Linial, Yishay Mansour, and Noam Nisan. "Constant Depth Circuits, Fourier Transform, and Learnability". In: *J. ACM* 40.3 (1993), pp. 607–620. DOI: 10.1145/174130.174138. URL: http://doi.acm.org/10.1145/174130.174138.

[24] Zachary C. Lipton. "The mythos of model interpretability". In: *Commun. ACM* 61.10 (2018), pp. 36–43. DOI: 10.1145/3233231. URL: https://doi.org/10.1145/3233231.

[25] Cody Murray and R. Ryan Williams. "Circuit lower bounds for nondeterministic quasi-polytime: an easy witness lemma for NP and NQP". In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*. Ed. by Ilias Diakonikolas, David Kempe, and Monika Henzinger. ACM, 2018, pp. 890–901. DOI: 10.1145/3188745.3188910. URL: http://doi.acm.org/10.1145/3188745.3188910.

[26] Noam Nisan and Avi Wigderson. "Hardness vs Randomness". In: *J. Comput. Syst. Sci.* 49.2 (1994), pp. 149–167. DOI: 10.1016/S0022-0000(05)80043-1. URL: http://dx.doi.org/10.1016/S0022-0000(05)80043-1.

[27] Igor Carboni Oliveira and Rahul Santhanam. "Conspiracies Between Learning Algorithms, Circuit Lower Bounds, and Pseudorandomness". In: *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia*. Ed. by Ryan O'Donnell. Vol. 79. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017, 18:1–18:49. DOI: 10.4230/LIPIcs.CCC.2017.18. URL: https://doi.org/10.4230/LIPIcs.CCC.2017.18.

[28] Igor Carboni Oliveira and Rahul Santhanam. "Pseudo-Derandomizing Learning and Approximation". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*. Ed. by Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer. Vol. 116. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018, 55:1–55:19. DOI: 10.4230/LIPIcs.APPROX-RANDOM.2018.55. URL: https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2018.55.

[29] Alexander A. Razborov and Steven Rudich. "Natural Proofs". In: *J. Comput. Syst. Sci.* 55.1 (1997), pp. 24–35. DOI: 10.1006/jcss.1997.1494. URL: http://dx.doi.org/10.1006/jcss.1997.1494.

[30] Rocco A. Servedio. "Smooth Boosting and Learning with Malicious Noise". In: *J. Mach. Learn. Res.* 4 (2003), pp. 633–648. URL: http://jmlr.org/papers/v4/servedio03a.html.

[31] Luca Trevisan and Salil P. Vadhan. "Pseudorandomness and Average-Case Complexity Via Uniform Reductions". In: *Computational Complexity* 16.4 (2007), pp. 331–364. DOI: 10.1007/s00037-007-0233-x. URL: http://dx.doi.org/10.1007/s00037-007-0233-x.

[32] Ryan Williams. "Improving Exhaustive Search Implies Superpolynomial Lower Bounds". In: *SIAM J. Comput.* 42.3 (2013), pp. 1218–1244. DOI: 10.1137/10080703X. URL: http://dx.doi.org/10.1137/10080703X.

[33] Virginia Vassilevska Williams. "On Some Fine-Grained Questions in Algorithms and Complexity". In: *Proceedings of the International Congress of Mathematicians* 3 (2018), pp. 3431–3472.