Research statement

Emanuele Viola

05/06/20

I work in theoretical computer science. This is the study of the reach of efficient computation, carried with mathematical tools. Theoretical computer science provides – or aims to provide – the foundations for the computer technology we enjoy every day, ranging from algorithmic efficiency to security of electronic transactions. More generally, it gives a computational viewpoint that is increasingly relevant to all of science, as articulated for example in Avi Wigderson's recent book [Wig19] titled "Mathematics and Computation: A Theory Revolutionizing Technology and Science."

In theoretical computer science I have worked in a wide variety of areas, including pseudorandomness (see Section 1), cryptography, communication complexity, group theory (these last three are discussed in Section 1.1), data structures (Section 2), and "Complexity Lower Bounds for Distributions" (Section 3), an area I initiated. My research has influenced not only computer scientists, but also leading researchers in mathematics and finance. Recently I received three long-term invitations to the Simons Institute for computing, and accepted two.

I describe a few of the research directions I pursued in the following sections, and I finish with a brief discussion of future plans.

1 Pseudorandomness

Background. A *pseudorandom generator* is an efficient procedure that stretches a short input seed into a much longer sequence that "looks random." These fascinating objects challenge our intuition of randomness and have a striking variety of applications. For example, ever since researchers in the 1940's have been employing the Monte Carlo method to obtain efficient randomized algorithms [MU49], it has been important to produce very long sequences that are "random enough" for those algorithms to work correctly. This need has spurred the development of pseudorandom generators, whose use has evolved over time and has also led to algorithmic breakthroughs such as "Undirected Connectivity in Logarithmic Space" [Rei08] or more recently the solution to long-standing open problems in coding theory [Ta-17].

Obtaining *unconditional* generators, i.e. generators that do not rely on any unproven assumption, is one of the ultimate goals of complexity theory, closely related to proving

P $\neq$ NP. The current ability in this direction is confined to restricted computational models. However, even generators for restricted models have a wide variety of applications. For example, the breakthroughs mentioned above are related to such generators.

My work. I have done extensive work on unconditional pseudorandom generators. I obtained in [Vio09b] a new generator for the important class of tests given by polynomials of low degree. Generators for the special case of linear polynomials had been known since the 90's [NN93], and are one of the most useful tools in theoretical computer science, with applications ranging from algorithm design to expander graphs and probabilistically checkable proofs. However, an extension to higher-degree polynomials had been an open problem ever since. My work gives the first such generators, and was awarded the *Best Paper Award* at the Computational Complexity Conference in 2008. After more than ten years and considerable effort by the community, my generator for low-degree polynomials [Vio09b] remains the state of the art. This result has been taught in classes offered at institutions such as UT Austin [Zuc17], and is devoted a section in the textbook [O'D14].

Moreover, this line of work – started with collaborator Andrej Bogdanov [BV10] – sparked a fruitful interaction between computer scientists and mathematicians such as Terence Tao, contributing to the solution of an outstanding conjecture in combinatorics [GT07, LMS08]. This conjecture is related to a certain norm introduced by Gowers in [Gow98, Gow01] and independently by Alon et al. in [AKK$^+$03]. Our work had multiple impacts. First, it formulated and promoted the study of a special case of the conjecture which was then shown to be false via a counterexample (thus disproving the more general conjecture). Second, it simplified both the proof of the conjecture over large domains (the counterexample only holds over small domains) and the proof of the counterexample, see [GT09].

For another example, a pseudorandom generator I have constructed [DGJ$^+$10] with a team of researchers gives as a byproduct a result of independent interest in probability theory. Recall that the central limit theorem asserts that the sum of independent random variables behaves approximately like a normal random variable. Our result implies the same conclusion starting from the weaker assumption that the variables enjoy only limited independence. Over the years, this result has found a number of applications, including an application to extractors which is mentioned below in Section 3.

I have also designed pseudorandom generators for various types of circuits [Vio07, FSUV13]; for the work [Vio07] I received the 2006 *SIAM Student Paper Prize.*

In the last few years, together with my Ph.D. student Chin Ho Lee (now a postdoc at Columbia University) and postdoc Elad Haramaty (subsequently a postdoc at Harvard University) we have published a series of papers [LV17, HLV18, LV, Lee19] on a paradigm in pseudorandomness which we called "bounded independence plus noise." In an equivalent form, this paradigm goes back to the influential works [AW89, GMR$^+$12]; our works gives it a different rendering. My team and I have in particular obtained new generators for several models, including small-space algorithms, and read-once polynomials (of any degree), a model which had been highlighted by researchers as a bottleneck for further progress [Tre10]. Our results have already been used by several other researchers in follow-ups, including

[MRT18] and [FK18]. Very recently, I have used these results to give a new pseudorandom generator for Turing machines [Vio19c], improving on a classical result [INW94].

In another direction, recently my collaborators and I have analyzed an important method to construct pseudorandom generators, based on "amplifying" a function which high circuit complexity to a pseudorandom generator in a generic way. We have shown in [GSV18, Viob] that, given the state of circuit complexity, such methods have essentially been exhausted, and cannot be used to obtain new pseudorandom generators from the circuit complexity results that we have. These results close a line of research which I initiated in my Ph.D. thesis [Vio06], and improve on a number of incomparable previous works.

1.1 From cryptography to group theory via communication complexity

Motivated by successful attacks on cryptographic hardware, an exciting line of work known as *leakage-resistant cryptography* considers models in which the adversary obtains more information from cryptographic algorithms than just their input/output behavior. A general goal is to compile any circuit into a new "shielded" circuit such that any attack exploiting the internals of the circuit during its computation can in fact be carried out just using input/output access (and hence does not succeed under standard hardness assumptions).

Together with my first Ph.D. student, Eric Miles, subsequently a postdoc at UCLA, I gave a new construction of shielded circuits [MV13] based on *finite groups.* This construction resists stronger computational attacks than previous works, such as those computable in stronger circuit classes [Mil14]. Also, the construction was candidate to simultaneously resisting another type of attacks considered in the literature, known as "only computation leaks" attacks [MR04].

The proof of correctness hinged however on a communication complexity conjecture. The setting of the conjecture is that several collaborating parties wish to compute the product of a tuple of elements from a group. The twist is that each party doesn't know a few of the elements (the so-called "number-on-the-forehead" model). The missed elements are arranged in a specific, *interleaved* fashion. The conjecture was that computing this product requires the parties to exchange a lot of communication. The validity of the conjecture critically relies on the underlying group: for example the conjecture is completely false for abelian groups, for which constant communication suffices.

Together with Timothy Gowers, I have made progress [GV19, GV15] on understanding interleaved group products. In particular, we have proved the conjecture in [MV13], thus enabling the motivating application in cryptography. These works by Gowers and myself have brought a new set of mathematical tools to bear on communication complexity and cryptography, including the "trace method," simple groups, and algebraic geometry. I presented these works at invited talks for the FOCS 2014 workshop on higher-order Fourier analysis and the Harvard 2017 workshop on additive combinatorics. Some of this material also appeared in a SIGACT survey I wrote upon invitation [Vio19b], and which was based on lectures from an advanced class I taught at Northeastern [Vio17], and posts on my

Wordpress blog [Vio16a, Vio16b]. Finally, our work has also sparked the interest of other mathematicians [Sha16].

2 Succinct data structures

Background. For many systems collecting massive amount of data, it is important to store it with only a negligible *redundancy* (i.e., overhead) in memory. Even a small factor redundancy in memory is a significant disadvantage when dealing with gigabytes of data. At the same time, we would like to retrieve the data or even answer various queries about it as fast as possible. These two goals are in tension, and the aim of *succinct data structures* is to achieve both simultaneously.

Researchers have recently developed new data structures with surprisingly good tradeoffs between redundancy and retrieval [LY08, Păt08, DPT10]. To illustrate, consider the common problem of storing a tuple of elements from the *ternary* alphabet $\{0, 1, 2\}$. The obvious difficulty is that computer memory is organized as *bits*, but we have to store "trits." Via so-called arithmetic coding, you can store the trits using the minimum number of bits, but then to retrieve a trit you need to read all the bits. Or you can use two distinct bits per trit, but then you waste a linear amount of memory.

A breakthrough work by Pătraşcu, later with Dodis and Thorup [Păt08, DPT10], gives a data structure whose redundancy decays *exponentially* with the number of bits read. This exponential decay gives a vast improvement over the parameters achieved by previous constructions, typically allowing for constant query time and negligible overhead. Similar succinct data structures have been obtained for a number of other problems. For another example consider the famous "rank/select" problem. Here we want to store succinctly a vector of bits so as to retrieve partial sums with few cell probes. This problem is the bread-and-butter of succinct data structures: It finds use in many other data structures (for representing dictionaries, trees, etc.). The trivial solution that allocates one cell for each partial sum results in constant retrieve time, but has large redundancy, superlinear in the input length. Pătraşcu in [Păt08] again shows how to reduce the redundancy exponentially in the retrieve time.

Given the surprising nature of these results, it is natural to ask if they can be further improved, or to establish a lower bound. However, on the lower-bound side the progress had been slower. Pătraşcu in [Păt08] notes that proving lower bounds for these problems *"seems beyond the scope of current techniques."*

My work. I have proved first and tight lower bounds for various fundamental data structure problems in [Vio12a, Vio09a] and, together with Pătraşcu, in [PV10]. Among these problems are both the ones mentioned above: the problem of storing ternary elements and the problem of storing partial sums. For both problems, my lower bounds match the upper bounds [DPT10, Păt08].

Recently I worked more on succinct data structures. I proved a first lower bound for the well-studied problem of storing permutations [Vio20]. I have also designed several new succinct data structures. In [Vio19a] I designed a new data structure for storing error-correcting codes. The parameters of the data structure essentially match a seminal lower bound in the area [GM07]. My data structure is obtained by drawing a new connection between data structures and circuits with few wires. I show that every problem which can be solved by a circuit with few wires admits an efficient data structure. Plugging in an efficient encoding circuit that I developed with coauthors [GHK$^+$13] gives the data structure. My connection with circuits has already been used in other works such as [KW19]. Finally, my coauthors and I have obtained a new succinct data structure for storing walks on graphs [VWY20]. The parameters obtained are similar to those in the breakthrough work mentioned above, in particular we can answer queries in constant time while still having very small redundancy. But our setting generalizes some previous work by allowing more complex dependencies among queries.

3 The complexity of distributions

In 2010 I initiated [Vio12b] the study of "Complexity Lower Bounds for Distributions" and I have been advocating it ever since. The object of this study is the computational complexity of sampling (also known as generating) probability distributions, given random numbers. This is in contrast with the more traditional study of the complexity of computing a function *on a given input.* While some related works existed in the literature, e.g. [ASTS$^+$03, GGN10], this area was largely uncharted. Since I published [Vio12b], a number of studies followed up building on this paper: [LV12, DW11, Vio14, BIL12, BCS14, Vio12c, Vio20], and in 2018 I gave a survey talk at the Simons Institute [Vioa].

Recently, I exhibited [Vio20] an explicit boolean function whose input-output distribution cannot be sampled by small-depth circuits significantly better than picking a uniform input and guessing the output bit at random. This result resolves a question that I asked in my earlier papers [Vio12b, Vio14], and is the strongest sampling lower bound to date for small-depth circuits.

My work has also laid out new connections between sampling lower bounds and several other areas, including data structures and randomness extractors. For example, I showed that if a problem admits an efficient data structure then it can also be sampled somewhat well. The data structure lower bound for permutations [Vio20] mentioned in Section 2 is obtained using this connection. Moving to the connection to extractors let us first recall that a randomness extractor is an algorithm that can convert distributions with high entropy into a nearly uniform distribution. Using the connection from my work I obtained the first extractor for high-entropy distributions that are sampled by small-depth circuits [Vio14].

My work has also had an impact on recent breakthrough constructions of *two-source extractors.* Specifically, in [Vio14] I introduced a new class of sources (some bits have bounded independence, the others are adversarially chosen), gave the first extractor for them (majority, analyzed using our previous work in pseudorandomness [DGJ$^+$10]), and finally asked if better extractors exist. Answering my question affirmatively is a main step in the breakthrough construction of two-source extractors for polylogarithmic entropy [CZ16]. Follow-up work [Li16] gives better yet extractors for the sources I introduced. Interestingly, subsequent papers leading to better and better two-source extractors [CS16, Coh16, BDT16] use instead the original extractor I gave in [Vio14].

4 Research goals

In the future, I will continue to work on questions about pseudorandom generators. For example, my generator for polynomials is proved correct only for polynomials of degree that is logarithmic (in the number of variables). However, it could work for higher degrees as well. Proving that this is indeed the case would have dramatic consequences, including solving major problems in circuit complexity (via [Raz87]). Another direction that I will pursue is to understand better the power of "bounded independence plus noise." Despite a surge of interest in recent years, basic questions remain, and the power of the method is unknown.

I will also continue to work on several problems at the intersection of communication complexity and group theory. Some such problems are candidate for solving long-standing open problems in communication complexity. For example in the work [GV19] we conjecture that a certain problem is hard even for communication protocols with more than a logarithmic number of parties, a well-known barrier in the area.

Finally, I also intend to continue to work on data structures and to develop the study of the complexity of distributions and its ramifications to other areas.

References

[AKK$^+$03] Noga Alon, Tali Kaufman, Michael Krivelevich, Simon Litsyn, and Dana Ron. Testing low-degree polynomials over GF(2). In *7th Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*, volume 2764 of *Lecture Notes in Computer Science*, pages 188–199. Springer, 2003.

[ASTS$^+$03] Andris Ambainis, Leonard J. Schulman, Amnon Ta-Shma, Umesh V. Vazirani, and Avi Wigderson. The quantum communication complexity of sampling. *SIAM J. on Computing*, 32(6):1570–1585, 2003.

[AW89] Miklos Ajtai and Avi Wigderson. Deterministic simulation of probabilistic constant-depth circuits. *Advances in Computing Research - Randomness and Computation*, 5:199–223, 1989.

[BCS14] Itai Benjamini, Gil Cohen, and Igor Shinkar. Bi-lipschitz bijection between the boolean cube and the hamming ball. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2014.

[BDT16] Avraham Ben-Aroya, Dean Doron, and Amnon Ta-Shma. Explicit two-source extractors for near-logarithmic min-entropy. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:88, 2016.

[BIL12] Chris Beck, Russell Impagliazzo, and Shachar Lovett. Large deviation bounds for decision trees and sampling lower bounds for AC0-circuits. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:42, 2012.

[BV10] Andrej Bogdanov and Emanuele Viola. Pseudorandom bits for polynomials. *SIAM J. on Computing*, 39(6):2464–2486, 2010.

[Coh16] Gil Cohen. Making the most of advice: New correlation breakers and their applications. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 188–196, 2016.

[CS16] Gil Cohen and Leonard J. Schulman. Extractors for near logarithmic min-entropy. *Electronic Colloquium on Computational Complexity (ECCC)*, 23:14, 2016.

[CZ16] Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. In *ACM Symp. on the Theory of Computing (STOC)*, pages 670–683, 2016.

[DGJ$^+$10] Ilias Diakonikolas, Parikshit Gopalan, Ragesh Jaiswal, Rocco A. Servedio, and Emanuele Viola. Bounded independence fools halfspaces. *SIAM J. on Computing*, 39(8):3441–3462, 2010.

[DPT10] Yevgeniy Dodis, Mihai Pătraşcu, and Mikkel Thorup. Changing base without losing space. In *42nd ACM Symp. on the Theory of Computing (STOC)*, pages 593–602. ACM, 2010.

[DW11] Anindya De and Thomas Watson. Extractors and lower bounds for locally samplable sources. In *Workshop on Randomization and Computation (RANDOM)*, 2011.

[FK18] Michael A. Forbes and Zander Kelley. Pseudorandom generators for read-once branching programs, in any order. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2018.

[FSUV13] Bill Fefferman, Ronen Shaltiel, Christopher Umans, and Emanuele Viola. On beating the hybrid argument. *Theory of Computing*, 9:809–843, 2013.

[GGN10] Oded Goldreich, Shafi Goldwasser, and Asaf Nussboim. On the implementation of huge random objects. *SIAM J. Comput.*, 39(7):2761–2822, 2010.

[GHK$^+$13] Anna Gál, Kristoffer Arnsfelt Hansen, Michal Koucký, Pavel Pudlák, and Emanuele Viola. Tight bounds on computing error-correcting codes by bounded-depth circuits with arbitrary gates. *IEEE Transactions on Information Theory*, 59(10):6611–6627, 2013.

[GM07] Anna Gál and Peter Bro Miltersen. The cell probe complexity of succinct data structures. *Theoretical Computer Science*, 379(3):405–417, 2007.

[GMR$^+$12] Parikshit Gopalan, Raghu Meka, Omer Reingold, Luca Trevisan, and Salil Vadhan. Better pseudorandom generators from milder pseudorandom restrictions. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2012.

[Gow98] Timothy Gowers. A new proof of Szemerédi's theorem for arithmetic progressions of length four. *Geometric and Functional Analysis*, 8(3):529–551, 1998.

[Gow01] Timothy Gowers. A new proof of Szemerédi's theorem. *Geometric and Functional Analysis*, 11(3):465–588, 2001.

[GSV18] Aryeh Grinberg, Ronen Shaltiel, and Emanuele Viola. Indistinguishability by adaptive procedures with advice, and lower bounds on hardness amplification proofs. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2018. Available at http://www.ccs.neu.edu/home/viola/.

[GT07] Ben Green and Terence Tao. The distribution of polynomials over finite fields, with applications to the Gowers norms, 2007. arXiv:0711.3191v1.

[GT09] Ben Green and Terence Tao. The distribution of polynomials over finite fields, with applications to the Gowers norms. *Contributions to Discrete Mathematics*, 4(2):1–36, 2009.

[GV15] W. T. Gowers and Emanuele Viola. The communication complexity of interleaved group products. In *ACM Symp. on the Theory of Computing (STOC)*, 2015.

[GV19] W. T. Gowers and Emanuele Viola. Interleaved group products. *SIAM J. on Computing*, 48(3):554–580, 2019. Special issue of FOCS 2016.

[HLV18] Elad Haramaty, Chin Ho Lee, and Emanuele Viola. Bounded independence plus noise fools products. *SIAM J. on Computing*, 47(2):295–615, 2018.

[INW94] Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *26th ACM Symp. on the Theory of Computing (STOC)*, pages 356–364, 1994.

[KW19] Young Kun Ko and Omri Weinstein. An adaptive step toward the multiphase conjecture, 2019.

[Lee19] Chin Ho Lee. Fourier bounds and pseudorandom generators for product tests. In Amir Shpilka, editor, *34th Computational Complexity Conference, CCC 2019, July 18-20, 2019, New Brunswick, NJ, USA*, volume 137 of *LIPIcs*, pages 7:1–7:25. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[Li16] Xin Li. Improved two-source extractors, and affine extractors for polylogarithmic entropy. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2016.

[LMS08] Shachar Lovett, Roy Meshulam, and Alex Samorodnitsky. Inverse conjecture for the Gowers norm is false. In *40th ACM Symp. on the Theory of Computing (STOC)*, pages 547–556, 2008.

[LV] Chin Ho Lee and Emanuele Viola. More on bounded independence plus noise: Pseudorandom generators for read-once polynomials. *Theory of Computing*. Available at http://www.ccs.neu.edu/home/viola/.

[LV12] Shachar Lovett and Emanuele Viola. Bounded-depth circuits cannot sample good codes. *Computational Complexity*, 21(2):245–266, 2012.

[LV17] Chin Ho Lee and Emanuele Viola. Some limitations of the sum of small-bias distributions. *Theory of Computing*, 13, 2017.

[LY08] Hsueh-I Lu and Chia-Chi Yeh. Balanced parentheses strike back. *ACM Trans. on Algorithms*, 4:28:1–28:13, July 2008.

[Mil14] Eric Miles. Iterated group products and leakage resilience against NC^1. In *ACM Innovations in Theoretical Computer Science conf. (ITCS)*, 2014.

[MR04] Silvio Micali and Leonid Reyzin. Physically observable cryptography. In *Theory* *of Cryptography Conf. (TCC)*, pages 278–296, 2004.

[MRT18] Raghu Meka, Omer Reingold, and Avishay Tal. Pseudorandom generators for width-3 branching programs. *Electronic Colloquium on Computational Complexity (ECCC)*, 25:112, 2018.

[MU49] Nicholas Metropolis and S. Ulam. The monte carlo method. *Journal of the American Statistical Association*, 44(247):335–341, 1949.

[MV13] Eric Miles and Emanuele Viola. Shielding circuits with groups. In *ACM Symp. on the Theory of Computing (STOC)*, 2013.

[NN93] Joseph Naor and Moni Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM J. on Computing*, 22(4):838–856, 1993.

[O'D14] Ryan O'Donnell. *Analysis of Boolean Functions.* Cambridge University Press, 2014.

[Păt08] Mihai Pătraşcu. Succincter. In *49th IEEE Symp. on Foundations of Computer Science (FOCS).* IEEE, 2008.

[PV10] Mihai Pătraşcu and Emanuele Viola. Cell-probe lower bounds for succinct partial sums. In *21th ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 117–122, 2010.

[Raz87] Alexander Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Akademiya Nauk SSSR. Matematicheskie Zametki*, 41(4):598–607, 1987. English translation in Mathematical Notes of the Academy of Sci. of the USSR, 41(4):333-338, 1987.

[Rei08] Omer Reingold. Undirected connectivity in log-space. *J. of the ACM*, 55(4), 2008.

[Sha16] Aner Shalev. Mixing, communication complexity and conjectures of Gowers and Viola. *Combinatorics, Probability and Computing*, pages 1–13, 6 2016. arXiv:1601.00795.

[Ta-17] Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In *ACM Symp. on the Theory of Computing (STOC)*, pages 238–251, 2017.

[Tre10] Luca Trevisan. Open problems in unconditional derandomization. Presentation at China Theory Week 2010, 2010. Available at http://conference.iiis.tsinghua.edu.cn/CTW2010/content/Slides/2.pdf.

[Vioa] Emanuele Viola. The Complexity of Distributions, Fall 2018 talk at the Simons Institute. https://www.youtube.com/watch?v=O78b085HE3w.

[Viob] Emanuele Viola. Constant-error pseudorandomness proofs from hardness require majority. *ACM Trans. Computation Theory.* Available at http://www.ccs.neu.edu/home/viola/.

[Vio06] Emanuele Viola. The complexity of hardness amplification and derandomization. *Ph.D. thesis, Harvard University*, 2006.

[Vio07] Emanuele Viola. Pseudorandom bits for constant-depth circuits with few arbitrary symmetric gates. *SIAM J. on Computing*, 36(5):1387–1403, 2007.

[Vio09a] Emanuele Viola. Cell-probe lower bounds for prefix sums, 2009. arXiv:0906.1370v1.

[Vio09b] Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d. *Computational Complexity*, 18(2):209–217, 2009.

[Vio12a] Emanuele Viola. Bit-probe lower bounds for succinct data structures. *SIAM J. on Computing*, 41(6):1593–1604, 2012.

[Vio12b] Emanuele Viola. The complexity of distributions. *SIAM J. on Computing*, 41(1):191–218, 2012.

[Vio12c] Emanuele Viola. Extractors for turing-machine sources. In *Workshop on Randomization and Computation (RANDOM)*, 2012.

[Vio14] Emanuele Viola. Extractors for circuit sources. *SIAM J. on Computing*, 43(2):355–972, 2014.

[Vio16a] Emanuele Viola. Thoughts: Mixing in groups, 2016. https://emanueleviola.wordpress.com/2016/10/21/mixing-in-groups/.

[Vio16b] Emanuele Viola. Thoughts: Mixing in groups ii, 2016. https://emanueleviola.wordpress.com/2016/11/15/mixing-in-groups-ii/.

[Vio17] Emanuele Viola. Special topics in complexity theory. Lecture notes of the class taught at Northeastern University. Available at http://www.ccs.neu.edu/home/viola/classes/spepf17.html, 2017.

[Vio19a] Emanuele Viola. Lower bounds for data structures with space close to maximum imply circuit lower bounds. *Theory of Computing*, 15:1–9, 2019. Available at http://www.ccs.neu.edu/home/viola/.

[Vio19b] Emanuele Viola. Non-abelian combinatorics and communication complexity. *SIGACT News, Complexity Theory Column*, 50(3), 2019.

[Vio19c] Emanuele Viola. Pseudorandom bits and lower bounds for randomized turing machines. Available at http://www.ccs.neu.edu/home/viola/, 2019.

[Vio20] Emanuele Viola. Sampling lower bounds: boolean average-case and permutations. *SIAM J. on Computing*, 49(1), 2020. Available at http://www.ccs.neu.edu/home/viola/.

[VWY20] Emanuele Viola, Omri Weinstein, and Huacheng Yu. How to store a random walk. In *ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2020. Available at http://www.ccs.neu.edu/home/viola/.

[Wig19] Avi Wigderson. *Mathematics and Computation: A Theory Revolutionizing Technology and Science*. Princeton University Press, 2019.

[Zuc17] David Zuckerman. Cs395t: Pseudorandomness (fall 2017). https://www.cs.utexas.edu/∼diz/395T/17/index.html, 2017.