

RESEARCH STATEMENT

Neil Lutz

For studying the intricate, irregular patterns exhibited by real-world systems, both natural and computational, fractal geometry and dynamical system theory constitute a powerful framework that has yielded advances throughout the sciences. Fractal geometry provides tools for analyzing sets that have non-trivial structure at every scale, for which “standard” geometry is inadequate. Dynamical system theory considers the possible trajectories of a system’s state over time. Even apparently simple systems can induce complex dynamics that map out highly irregular, fragmented sets—i.e., fractal sets. My research uses ideas from theoretical computer science to examine both fractal geometry and dynamic behavior.

Algorithmic Fractal Geometry

The goal of my work on fractals is to develop and apply connections between classical fractal geometry and the theory of algorithmic information. The approach is to view the points in fractal sets as infinite data objects with varying degrees of compressibility. So far, this work has been primarily focused on two notions of dimension and the relationship between them.

The first is *Hausdorff dimension*, which is the most standard notion of fractal dimension. Since fractal sets of interest typically have the cardinality of the continuum but zero volume, their “size” cannot be adequately quantified by counting or measure. Hausdorff dimension addresses this, providing a meaningful quantitative classification of fractal sets. As a general rule, Hausdorff dimension is easy to approximate from above, but it is often difficult to prove lower bounds.

The second notion is *algorithmic dimension*, which was developed by analogy to Hausdorff dimension in order to quantitatively classify individual data objects according to their level of unpredictability and later characterized using Kolmogorov complexity, the standard measure of algorithmic information. In the latter formulation, the algorithmic dimension of an individual point x represents the *incompressibility* of x ’s binary expansion.

In [14], we established a *point-to-set principle* that uses relativization to make the relationship between these two dimension concepts precise. In particular, this principle characterizes the Hausdorff dimension of any set in \mathbb{R}^n terms of the algorithmic dimensions of its individual points, relative to an optimal oracle. Together with other new algorithmic dimensional tools we developed, this principle gives rise to a new, pointwise approach for bounding Hausdorff dimensions: To prove a lower bound on the Hausdorff dimension of a set, it is sufficient to find a single point that is not too compressible.

Some of the early triumphs of this method have been in removing structural hypotheses from prominent classical theorems. For example, the following two statements about Hausdorff dimension, \dim_{H} , have long been known to hold for every Borel set $E \subseteq \mathbb{R}^2$. The first is known as *Marstrand's slicing theorem*, and the second is a well-known characterization of *packing dimension*, \dim_{P} , another extensively studied notion of fractal dimension.

- Almost every vertical slice of E has Hausdorff dimension $\leq \max\{0, \dim_{\text{H}} E - 1\}$ [20].
- $\dim_{\text{P}} E = \max_{F \subseteq \mathbb{R}^2} \dim_{\text{H}}(E \times F) - \dim_{\text{H}} F$ [3].

Using a pointwise, algorithmic information theoretic approach, I have shown that both of these statements hold for arbitrary $E \subseteq \mathbb{R}^2$, removing the assumption that E is Borel [16]. These results and others, such as the projection theorems we proved in [18], have demonstrated that the pointwise approach is well-suited to handle pathological sets, circumventing the kind of global set properties—measurability, compactness, etc.—that are often required under traditional analytic techniques.

Working with others, I have also used the pointwise approach to study a prominent class of problems in geometric measure theory concerning the minimum size of sets that contain large subsets of lines in many directions. A *Keakeya set*, for instance, is any set that contains length-1 segments in all directions. Keakeya sets in \mathbb{R}^n may have measure zero [1, 2], and the famous *Keakeya conjecture* asserts that they must have Hausdorff dimension n . This has been shown to hold for $n = 2$ [4] and is open for all $n > 2$.

Using the point-to-set principle, problems of this type can be posed in terms of algorithmic dimensions. For example, the existence of measure-zero Keakeya sets follows from the existence, for every slope $a \in \mathbb{R}$, of an intercept $b \in \mathbb{R}$ such that the line $y = ax + b$ contains no algorithmically random point, as we showed in [13]. Similarly, the fact that Keakeya sets in \mathbb{R}^2 must have dimension 2 is implied by our result in [14] that for all algorithmically random $a \in \mathbb{R}$ and all $b \in \mathbb{R}$, almost every point on the line $y = ax + b$ has effective dimension 2.

A more general problem in the same class concerns sets in \mathbb{R}^2 that contain, for some parameters $\alpha, \beta \in (0, 1]$, α -dimensional intersections with lines in a β -dimensional set of directions. These sets are called *generalized Furstenberg sets*, and their minimum Hausdorff dimension is not known [21]. In [19], we improved on the known lower bound by proving a general bound on the effective dimension of any point $(x, ax+b)$ in terms of the compressibility of the individual components x , a , and b . One of my current research objectives is to further refine the new proof techniques we introduced in that work to strengthen the main result.

More broadly, I plan to advance this research program in three ways. First, I am working to build a more complete understanding of the structure of algorithmic information in \mathbb{R}^n , or at least in \mathbb{R}^2 . This includes an ongoing collaboration on the dimension spectra of lines [17] and studying the dimensions of points on other simple geometric figures like circles and polynomial curves. Second, I will continue using the pointwise approach to generalize classical

results in fractal geometry. And finally, I will prove point-to-set principles for other metric spaces and notions of dimension; already this year, my co-authors and I have established point-to-set principles for gauged dimensions in arbitrary separable metric spaces [12], expanding the scope of problems susceptible to point-to-set techniques. My work in this area has appeared in recent survey articles by Downey and Hirschfeldt in the *Bulletin of the American Mathematical Society* and in *Communications of the ACM* [6, 7]. This work is also one of the main topics of a recent workshop on algorithmic randomness at the American Institute of Mathematics (<https://aimath.org/workshops/upcoming/algorandom/>).

Distributed Strategic Dynamics

The disciplines of game theory and distributed computing are both concerned with describing possible outcomes of interaction between multiple agents, but they differ in focus. In game theory, the agents have individual economic preferences, and each is expected to act strategically according to its own self-interest, although these interests may conflict with one another. By contrast, distributed computing research is focused on agents trying to coordinate to accomplish some computational task in a way that is robust to difficulties like unreliable communication channels, asynchrony, or memory failures.

For many real-world systems, both of these perspectives are relevant. A motivating example is network routing protocols, in which each server in a computer network is assumed to have preferred outcomes and to act strategically, but the interaction is highly decentralized and asynchronous. This dual character is also exhibited by systems in large-scale markets, social networks, and multi-processor computer architectures.

My research draws on both game theory and distributed computing to investigate the long-term global behavior—i.e., the dynamics—of such systems when the agents repeatedly interact. Of particular interest is the possibility of convergence to a single global configuration, which captures the notions of both game theoretic equilibrium and termination of a distributed computation. The broad goal of my work in this area is to describe the computational and information resources necessary so that this convergence can be guaranteed.

My coauthors and I have taken major steps in this direction. In [9], we considered dynamics that are *uncoupled*, meaning that agents' behaviors are determined by decentralized private inputs [8]. For game dynamics, these inputs represent each player's private preferences, of which the other players are initially ignorant. We characterized the robustness of these dynamics to memory failures by bounding the number of past configurations that each agent must be able to recall in order to guarantee convergence to a game-theoretic equilibrium. One of my short-term research objectives is to show that our recall bound for deterministic uncoupled dynamics is tight.

In [10], we studied asynchronous dynamics, in which interaction is governed by a schedule that determines, in each discrete time step, which of the agents will be activated and permitted

to take action. Each agent may be forced to wait arbitrarily (but finitely) many time steps in between activations. We presented a broadly applicable non-convergence theorem describing a large and natural class of asynchronous systems in which convergence cannot be guaranteed. This theorem implies that many routing protocols, asynchronous circuit systems, and social dynamics are fundamentally unstable, in that they are prone to indefinite oscillations of the global configuration.

An important special case of this non-convergence result, which is sufficient for many of these applications, is that oscillations can occur whenever each agent’s behavior is *historyless*, meaning that their decisions depend only on the system’s current configuration. In [5], my coauthors and I introduced a more general notion of *stateless computation*, in which decisions may additionally depend on messages sent by neighboring agents. This change preserves many robustness properties that are desirable for distributed applications, but we showed that it adds significant computational power, even when message length is severely restricted and the communication network is sparse.

We also showed in that work that asynchronous stateless computations can fail to converge, even when the wait times imposed by the schedule are bounded, which is a more limited and “realistic” form of asynchrony than we studied in [10]. This distinction reflects the fact that asynchrony is not purely adversarial in practice, and that robustness to asynchrony is not an all-or-nothing proposition. Even among systems that converge with probability 1 under uniformly random schedules, we can ask how aberrant the schedules that lead to non-convergence are. Accordingly, I plan to use fractal dimension to develop a more quantitative approach to robustness to asynchrony.

More recently, I have begun to explore game dynamics in two additional settings. The first is molecular programming: In [15], we developed a framework to quantify the robustness of molecular programs by modeling interactions between them as *catalytic games* between chemical reaction networks that can perturb each others’ relative clock speeds. The second is multi-armed bandit models, widely used in learning theory, in which an agent repeatedly chooses among several different options, then receives a reward based on its choice. In order to maximize its cumulative rewards, the agent typically balances *exploitation* of empirically lucrative options with *exploration* of less familiar options. This exploration can be costly in terms of *regret*, which is the difference between realized cumulative rewards and optimal cumulative rewards; in T rounds, an agent acting alone must incur regret on the order of $\log T$. In a system with multiple agents, though, a strategic agent might *free-ride* off of others’ exploration, unfairly avoiding that cost. In [11], we analyzed the feasibility of free-riding strategies, describing circumstances in which the free-rider’s regret can be bounded by a constant in the presence of *self-reliant* exploring agents. I plan to study the dynamics and equilibria that emerge when all players act strategically and to develop deterrence mechanisms for free-riding.

References

- [1] A. S. Besicovitch. Sur deux questions d'intégrabilité des fonctions. *Journal de la Société de physique et de mathématique de l'Université de Perm*, 2:105–123, 1919.
- [2] A. S. Besicovitch. On Kakeya's problem and a similar one. *Mathematische Zeitschrift*, 27:312–320, 1928.
- [3] C. J. Bishop and Y. Peres. Packing dimension and Cartesian products. *Transactions of the American Mathematical Society*, 348:4433–4445, 1996.
- [4] R. O. Davies. Some remarks on the Kakeya problem. *Proc. Cambridge Phil. Soc.*, 69:417–421, 1971.
- [5] D. Dolev, M. Erdmann, N. Lutz, M. Schapira, and A. Zair. Stateless computation. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, PODC '17, pages 419–421, New York, NY, USA, 2017.
- [6] R. Downey and D. R. Hirschfeldt. Algorithmic randomness. *Communications of the ACM*, 62(5):70–80, 2019.
- [7] R. Downey and D. R. Hirschfeldt. Computability and randomness. *Notices of the American Mathematical Society*, 66(7):1001–1012, 2019.
- [8] S. Hart and A. Mas-Colell. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5):1830–1836, 2003.
- [9] A. D. Jaggard, N. Lutz, M. Schapira, and R. N. Wright. Self-stabilizing uncoupled dynamics. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT '14)*, pages 74–85, 2014.
- [10] A. D. Jaggard, N. Lutz, M. Schapira, and R. N. Wright. Dynamics at the boundary of game theory and distributed computing. *ACM Trans. Econ. Comput.*, 5(3):15:1–15:20, Aug. 2017.
- [11] C. Jung, S. Kannan, and N. Lutz. Quantifying the burden of exploration and the unfairness of free riding. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2020)*, to appear.
- [12] J. Lutz, N. Lutz, and E. Mayordomo. Extending the reach of the point-to-set principle, 2020. URL: <https://arxiv.org/pdf/2004.07798.pdf>.
- [13] J. H. Lutz and N. Lutz. Lines missing every random point. *Computability*, 4(2):85–102, 2015.
- [14] J. H. Lutz and N. Lutz. Algorithmic information, plane Kakeya sets, and conditional dimension. *ACM Trans. Comput. Theory*, 10(2):7:1–7:22, May 2018.
- [15] J. H. Lutz, N. Lutz, R. R. Lutz, and M. R. Riley. Robustness and games against nature in molecular programming. In *Proceedings of the 41st International Conference on Software Engineering: New Ideas and Emerging Results, ICSE (NIER) 2019, Montreal, QC, Canada, May 29-31, 2019*, pages 65–68, 2019.
- [16] N. Lutz. Fractal intersections and products via algorithmic dimension. In *42nd International Symposium on Mathematical Foundations of Computer Science August 21–25, 2017, Aalborg, Denmark, Proceedings*, 2017. Extended version: <https://arxiv.org/abs/1612.01659>, 2019.
- [17] N. Lutz and D. M. Stull. Dimension spectra of lines. In *Unveiling Dynamics and Complexity - 13th Conference on Computability in Europe, CiE 2017, Turku, Finland, June 12-16, 2017, Proceedings*, pages 304–314, 2017.
- [18] N. Lutz and D. M. Stull. Projection theorems using effective dimension. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK*, pages 71:1–71:15, 2018.
- [19] N. Lutz and D. M. Stull. Bounding the dimension of points on a line. *Information and Computation*, 275, 2020.
- [20] J. M. Marstrand. Some fundamental geometrical properties of plane sets of fractional dimensions. *Proceedings of the London Mathematical Society*, 4(3):257–302, 1954.
- [21] U. Molter and E. Rela. Furstenberg sets for a fractal set of directions. *Proc. Amer. Math. Soc.*, 140:2753–2765, 2012.