

Publications of Maximilian Katzmann

This document lists all peer-reviewed publications of Maximilian Katzmann, Chair for Algorithm Engineering, Hasso Plattner Institute, Potsdam, Germany. This listing was automatically generated on February 22, 2024. An up-to-date version is available online at hpi.de/friedrich/docs/publist/katzmann.pdf.

Journal articles

- [1] Bläsius, T., Fischbeck, P., Friedrich, T., Katzmann, M., Solving Vertex Cover in Polynomial Time on Hyperbolic Random Graphs. In: *Theory of Computing Systems* 67, pp. 28–51, 2023.

The computational complexity of the VertexCover problem has been studied extensively. Most notably, it is NP-complete to find an optimal solution and typically NP-hard to find an approximation with reasonable factors. In contrast, recent experiments suggest that on many real-world networks the run time to solve VertexCover is way smaller than even the best known FPT-approaches can explain. We link these observations to two properties that are observed in many real-world networks, namely a heterogeneous degree distribution and high clustering. To formalize these properties and explain the observed behavior, we analyze how a branch-and-reduce algorithm performs on hyperbolic random graphs, which have become increasingly popular for modeling real-world networks. In fact, we are able to show that the VertexCover problem on hyperbolic random graphs can be solved in polynomial time, with high probability. The proof relies on interesting structural properties of hyperbolic random graphs. Since these predictions of the model are interesting in their own right, we conducted experiments on real-world networks showing that these properties are also observed in practice.

- [2] Bläsius, T., Freiberger, C., Friedrich, T., Katzmann, M., Montenegro-Retana, F., Thieffry, M., Efficient Shortest Paths in Scale-Free Networks with Underlying Hyperbolic Geometry. In: *ACM Transactions on Algorithms*, pp. 1–32, 2022.

A standard approach to accelerating shortest path algorithms on networks is the bidirectional search, which explores the graph from the start and the destination, simultaneously. In practice this strategy performs particularly well on scale-free real-world networks. Such networks typically have a heterogeneous degree distribution (e.g., a power-law distribution) and high clustering (i.e., vertices with a common neighbor are likely to be connected themselves). These two properties can be obtained by assuming an underlying hyperbolic geometry. To explain the observed behavior of the bidirectional search, we analyze its running time on hyperbolic random graphs and prove that it is $\tilde{O}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ with high probability, where $\alpha \in (0.5, 1)$ controls the power-law exponent of the degree distribution, and δ_{\max} is the maximum degree. This bound is sublinear, improving the obvious worst-case linear bound. Although our analysis depends on the underlying geometry, the algorithm itself is oblivious to it.

- [3] Bläsius, T., Friedrich, T., Katzmann, M., Krohmer, A., Hyperbolic Embeddings for Near-Optimal Greedy Routing. In: *Journal of Experimental Algorithmics (JEA)* 25, pp. 1–18, 2020.

Greedy routing computes paths between nodes in a network by successively moving to the neighbor closest to the target with respect to coordinates given by an embedding into some metric space. Its advantage is that only local information is used for routing decisions. We present different algorithms for generating graph embeddings into the hyperbolic plane that are well suited for greedy routing. In particular, our embeddings guarantee that greedy routing always succeeds in reaching the target, and we try to minimize the lengths of the resulting greedy paths. We evaluate our algorithm on multiple generated and real-world networks. For networks that are generally assumed to have a hidden underlying hyperbolic geometry, such as the Internet graph, we achieve near-optimal results (i.e., the resulting greedy paths are only slightly longer than the corresponding shortest paths). In the case of the Internet graph, they are only 6% longer when using our best algorithm, which greatly improves upon the previous best known embedding, whose creation required substantial manual intervention. In addition to measuring the stretch, we empirically evaluate our algorithms regarding the size of the coordinates of the resulting embeddings and observe how it impacts the success rate when coordinates are not represented with very high precision. Since numerical difficulties are a major issue when performing computations in the hyperbolic plane, we consider variations of our algorithm that improve the success rate when using coordinates with lower precision.

- [4] Friedrich, T., Katzmann, M., Krohmer, A., Unbounded Discrepancy of Deterministic Random Walks on Grids. In: *SIAM Journal on Discrete Mathematics* 32, pp. 2441–2452, 2018.

Random walks are frequently used in randomized algorithms. We study a derandomized variant of a random walk on graphs, called rotor-router model. In this model, instead of distributing tokens randomly, each vertex serves its neighbors in a fixed deterministic order. For most setups, both processes behave remarkably similar: Starting with the same initial configuration, the number of tokens in the rotor-router model deviates only slightly from the expected number of tokens on the corresponding vertex in the random walk model. The maximal difference over all vertices and all times is called single vertex discrepancy. Cooper and Spencer (2006) showed that on \mathbb{Z}^d the single vertex discrepancy is only a constant c_d . Other authors also determined the precise value of c_d for $d = 1, 2$. All these results, however, assume that initially all tokens are only placed on one partition of the bipartite graph \mathbb{Z}^d . We show that this assumption is crucial by proving that otherwise the single vertex discrepancy can become arbitrarily large. For all dimensions $d \geq 1$ and arbitrary discrepancies $\ell \geq 0$, we construct configurations that reach a discrepancy of at least ℓ .

Conference papers

- [5] Bläsius, T., Friedrich, T., Katzmann, M., Ruff, J., Zeif, Z., [On the Giant Component of Geometric Inhomogeneous Random Graphs](#). In: *European Symposium on Algorithms (ESA)*, pp. 20:1–20:13, 2023.

In this paper we study the threshold model of geometric inhomogeneous random graphs (GIRGs); a generative random graph model that is closely related to hyperbolic random graphs (HRGs). These models have been observed to capture complex real-world networks well with respect to the structural and algorithmic properties. Following comprehensive studies regarding their connectivity, i.e., which parts of the graphs are connected, we have a good understanding under which circumstances a giant component (containing a constant fraction of the graph) emerges. While previous results are rather technical and challenging to work with, the goal of this paper is to provide more accessible proofs. At the same time we significantly improve the previously known probabilistic guarantees, showing that GIRGs contain a giant component with probability $1 - \exp(-\Omega(n^{(3-\tau)/2}))$ for graph size n and a degree distribution with power-law exponent $\tau \in (2, 3)$. Based on that we additionally derive insights about the connectivity of certain induced subgraphs of GIRGs.

- [6] Friedrich, T., Göbel, A., Katzmann, M., Schiller, L., [Cliques in High-Dimensional Geometric Inhomogeneous Random Graphs](#). In: *International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 62:1–62:13, 2023.

A recent trend in the context of graph theory is to bring theoretical analyses closer to empirical observations, by focusing the studies on random graph models that are used to represent practical instances. There, it was observed that geometric inhomogeneous random graphs (GIRGs) yield good representations of complex real-world networks, by expressing edge probabilities as a function that depends on (heterogeneous) vertex weights and distances in some underlying geometric space that the vertices are distributed in. While most of the parameters of the model are understood well, it was unclear how the dimensionality of the ground space affects the structure of the graphs. In this paper, we complement existing research into the dimension of geometric random graph models and the ongoing study of determining the dimensionality of real-world networks, by studying how the structure of GIRGs changes as the number of dimensions increases. We prove that, in the limit, GIRGs approach non-geometric inhomogeneous random graphs and present insights on how quickly the decay of the geometry impacts important graph structures. In particular, we study the expected number of cliques of a given size as well as the clique number and characterize phase transitions at which their behavior changes fundamentally. Finally, our insights help in better understanding previous results about the impact of the dimensionality on geometric random graphs.

- [7] Bläsius, T., Friedrich, T., Katzmann, M., Stephan, D., [Strongly Hyperbolic Unit Disk Graphs](#). In: *Symposium Theoretical Aspects of Computer Science (STACS)*, pp. 13:1–13:17, 2023.

The class of Euclidean unit disk graphs is one of the most fundamental and well-studied graph classes with underlying geometry. In this paper, we identify this class as a special case in the broader class of hyperbolic unit disk graphs and introduce strongly hyperbolic unit disk graphs as a natural counterpart to the Euclidean variant. In contrast to the grid-like structures exhibited by Euclidean unit disk graphs, strongly hyperbolic networks feature hierarchical structures, which are also observed in complex real-world networks. We investigate basic properties of strongly hyperbolic unit disk graphs, including adjacencies and the formation of cliques, and utilize the derived insights to demonstrate that the class is useful for the development and analysis of graph algorithms. Specifically, we develop a simple greedy routing scheme and analyze its performance on strongly hyperbolic unit disk graphs in order to prove that routing can be performed more efficiently on such networks than in general.

- [8] Angrick, S., Bals, B., Hastrich, N., Kleissl, M., Schmidt, J., Doskoč, V., Katzmann, M., Molitor, L., Friedrich, T., [Towards Explainable Real Estate Valuation via Evolutionary Algorithms](#). In: *Genetic and Evolutionary Computation Conference (GECCO)*, pp. 1130–1138, 2022.

Human lives are increasingly influenced by algorithms, which therefore need to meet higher standards not only in accuracy but also with respect to explainability. This is especially true for high-stakes areas such as real estate valuation. Unfortunately, the methods applied there often exhibit a trade-off between accuracy and explainability. One explainable approach is case-based reasoning (CBR), where each decision is supported by specific previous cases. However, such methods can be wanting in accuracy. The unexplainable machine learning approaches are often observed to provide higher accuracy but are not scrutable in their decision-making. In this paper, we apply evolutionary algorithms (EAs) to CBR predictors in order to improve their performance. In particular, we deploy EAs to the similarity functions (used in CBR to find comparable cases), which are fitted to the data set at hand. As a consequence, we achieve higher accuracy than state-of-the-art deep neural networks (DNNs), while keeping interpretability and explainability. These results stem from our empirical evaluation on a large data set of real estate offers where we compare known similarity functions, their EA-improved counterparts, and DNNs. Surprisingly, DNNs are only on par with standard CBR techniques. However, using EA-learned similarity functions does yield an improved performance.

- [9] Friedrich, T., Göbel, A., Katzmann, M., Krejca, M. S., Pappik, M., [Algorithms for hard-constraint point processes via discretization](#). In: *International Computing and Combinatorics Conference (COCOON)*, pp. 242–254, 2022.

We study the algorithmic applications of a natural discretization for the hard-sphere model and the Widom–Rowlinson model in a region of d -dimensional Euclidean space $V \subset \mathbf{R}^d$. These continuous models are frequently used in statistical physics to describe mixtures of one or multiple particle types subjected to hard-core interactions. For each type, particles are distributed according to a Poisson point process with a type-specific activity parameter, called fugacity. The Gibbs distribution over all possible system states is characterized by the mixture of these point processes conditioned that no two particles are closer than some type-dependent distance threshold. A key part in better understanding the Gibbs distribution is its normalizing constant, called partition function. Our main algorithmic result is the first deterministic approximation algorithm for the partition function of the hard-sphere model and the Widom–Rowlinson model in box-shaped regions of Euclidean space. Our algorithms have quasi-polynomial running time in the volume of the region $\nu(V)$ if the fugacity is below a certain threshold. For the d -dimensional hard-sphere model with particles of unit volume, this threshold is $e/2^d$. As the number of dimensions d increases, this bound asymptotically matches the best known results for randomized approximation of the hard-sphere partition function. We prove similar bounds for the Widom–Rowlinson model. To the best of our knowledge, this is the first rigorous algorithmic result for this model.

- [10] Bläsius, T., Friedrich, T., Katzmann, M., [Efficiently Approximating Vertex Cover on Scale-Free Networks with Underlying Hyperbolic Geometry](#). In: *European Symposium on Algorithms (ESA)*, pp. 20:1–20:15, 2021.

Finding a minimum vertex cover in a network is a fundamental NP-complete graph problem. One way to deal with its computational hardness, is to trade the qualitative performance of an algorithm (allowing non-optimal outputs) for an improved running time. For the vertex cover problem, there is a gap between theory and practice when it comes to understanding this tradeoff. On the one hand, it is known that it is NP-hard to approximate a minimum vertex cover within a factor of $\sqrt{2}$. On the other hand, a simple greedy algorithm yields close to optimal approximations in practice. A promising approach towards understanding this discrepancy is to recognize the differences between theoretical worst-case instances and real-world networks. Following this direction, we close the gap between theory and practice by providing an algorithm that efficiently computes nearly optimal vertex cover approximations on hyperbolic random graphs; a network model that closely resembles real-world networks in terms of degree distribution, clustering, and the small-world property. More precisely, our algorithm computes a $(1 + o(1))$ -approximation, asymptotically almost surely, and has a running time of $\mathcal{O}(m \log(n))$. The proposed algorithm is an adaption of the successful greedy approach, enhanced with a procedure that improves on parts of the graph where greedy is not optimal. This makes it possible to introduce a parameter that can be used to tune the tradeoff between approximation performance and running time. Our empirical evaluation on real-world networks shows that this allows for improving over the near-optimal results of the greedy approach.

- [11] Bläsius, T., Friedrich, T., Katzmann, M., [Force-Directed Embedding of Scale-Free Networks in the Hyperbolic Plane](#). In: *Symposium on Experimental Algorithms (SEA)*, pp. 22:1–22:18, 2021.

Force-directed drawing algorithms are the most commonly used approach to visualize networks. While they are usually very robust, the performance of Euclidean spring embedders decreases if the graph exhibits the high level of heterogeneity that typically occurs in scale-free real-world networks. As heterogeneity naturally emerges from hyperbolic geometry (in fact, scale-free networks are often perceived to have an underlying hyperbolic geometry), it is natural to embed them into the hyperbolic plane instead. Previous techniques that produce hyperbolic embeddings usually make assumptions about the given network, which (if not met) impairs the quality of the embedding. It is still an open problem to adapt force-directed embedding algorithms to make use of the heterogeneity of the hyperbolic plane, while also preserving their robustness. We identify fundamental differences between the behavior of spring embedders in Euclidean and hyperbolic space, and adapt the technique to take advantage of the heterogeneity of the hyperbolic plane.

- [12] Bläsius, T., Fischbeck, P., Friedrich, T., Katzmann, M., [Solving Vertex Cover in Polynomial Time on Hyperbolic Random Graphs](#). In: *Symposium on the Theoretical Aspects of Computer Science (STACS)*, pp. 25:1–25:14, 2020.

The VertexCover problem is proven to be computationally hard in different ways: It is NP-complete to find an optimal solution and even NP-hard to find an approximation with reasonable factors. In contrast, recent experiments suggest that on many real-world networks the run time to solve VertexCover is way smaller than even the best known FPT-approaches can explain. Similarly, greedy algorithms deliver very good approximations to the optimal solution in practice. We link these observations to two properties that are observed in many real-world networks, namely a heterogeneous degree distribution and high clustering. To formalize these properties and explain the observed behavior, we analyze how a branch-and-reduce algorithm performs on hyperbolic random graphs, which have become increasingly popular for modeling real-world networks. In fact, we are able to show that the VertexCover problem on hyperbolic random graphs can be solved in polynomial time, with high probability. The proof relies on interesting structural properties of hyperbolic random graphs. Since these predictions of the model are interesting in their own right, we conducted experiments on real-world networks showing that these properties are also observed in practice. When utilizing the same structural properties in an adaptive greedy algorithm, further experiments suggest that, on real instances, this leads to better approximations than the standard greedy approach within reasonable time. We link these observations to two properties that are observed in many real-world networks, namely a heterogeneous degree distribution and high clustering. To formalize these properties and explain the observed behavior, we analyze how a branch-and-reduce algorithm performs on hyperbolic random graphs, which have become increasingly popular for modeling real-world networks. In fact, we are able to show that the VertexCover problem on hyperbolic random graphs can be solved in polynomial time, with high probability. The proof relies on interesting structural properties of hyperbolic random graphs. Since these predictions of the model are interesting in their own right, we conducted experiments on real-world networks showing that these properties are also observed in practice. When utilizing the same structural properties in an adaptive greedy algorithm, further experiments suggest that this leads to even better approximations than the standard greedy approach on real instances.

- [13] Bläsius, T., Friedrich, T., Katzmann, M., Meyer, U., Penschuck, M., Weyand, C., [Efficiently Generating Geometric Inhomogeneous and Hyperbolic Random Graphs](#). In: *European Symposium on Algorithms (ESA)*, pp. 21:2–21:14, 2019. **EATCS Best Paper Award**.

Hyperbolic random graphs (HRG) and geometric inhomogeneous random graphs (GIRG) are two similar generative network models that were designed to resemble complex real world networks. In particular, they have a power-law degree distribution with controllable exponent β , and high clustering that can be controlled via the temperature T . We present the first implementation of an efficient GIRG generator running in expected linear time. Besides varying temperatures, it also supports underlying geometries of higher dimensions. It is capable of generating graphs with ten million edges in under a second on commodity hardware. The algorithm can be adapted to HRGs. Our resulting implementation is the fastest sequential HRG generator, despite the fact that we support non-zero temperatures. Though non-zero temperatures are crucial for many applications, most existing generators are restricted to $T = 0$. We also support parallelization, although this is not the focus of this paper. Moreover, we note that our generators draw from the correct probability distribution, i.e., they involve no approximation. Besides the generators themselves, we also provide an efficient algorithm to determine the non-trivial dependency between the average degree of the resulting graph and the input parameters of the GIRG model. This makes it possible to specify \bar{d} as input and obtain a graph with expected average degree \bar{d} . Moreover, we investigate the differences between HRGs and GIRGs, shedding new light on the nature of the relation between the two models. Although HRGs represent, in a certain sense, a special case of the GIRG model, we find that a straight-forward inclusion does not hold in practice. However, the difference is negligible for most use cases.

- [14] Bläsius, T., Friedrich, T., Katzmann, M., Krohmer, A., [Hyperbolic Embeddings for Near-Optimal Greedy Routing](#). In: *Algorithm Engineering and Experiments (ALENEX)*, pp. 199–208, 2018.

Greedy routing computes paths between nodes in a network by successively moving to the neighbor closest to the target with respect to coordinates given by an embedding into some metric space. Its advantage is that only local information is used for routing decisions. We present different algorithms for generating graph embeddings into the hyperbolic plane that are well suited for greedy routing. In particular our embeddings guarantee that greedy routing always succeeds in reaching the target and we try to minimize the lengths of the resulting greedy paths. We evaluate our algorithm on multiple generated and real world networks. For networks that are generally assumed to have a hidden underlying hyperbolic geometry, such as the Internet graph, we achieve near-optimal results, i.e., the resulting greedy paths are only slightly longer than the corresponding shortest paths. In the case of the Internet graph, they are only 6% longer when using our best algorithm, which greatly improves upon the previous best known embedding, whose creation required substantial manual intervention.

- [15] Bläsius, T., Freiberger, C., Friedrich, T., Katzmann, M., Montenegro-Retana, F., Thieffry, M., [Efficient Shortest Paths in Scale-Free Networks with Underlying Hyperbolic Geometry](#). In: *International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 20:1–20:14, 2018.

A common way to accelerate shortest path algorithms on graphs is the use of a bidirectional search, which simultaneously explores the graph from the start and the destination. It has been observed recently that this strategy performs particularly well on scale-free real-world networks. Such networks typically have a heterogeneous degree distribution (e.g., a power-law distribution) and high clustering (i.e., vertices with a common neighbor are likely to be connected themselves). These two properties can be obtained by assuming an underlying hyperbolic geometry. To explain the observed behavior of the bidirectional search, we analyze its running time on hyperbolic random graphs and prove that it is $\tilde{O}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ with high probability, where $\alpha \in (0.5, 1)$ controls the power-law exponent of the degree distribution, and δ_{\max} is the maximum degree. This bound is sublinear, improving the obvious worst-case linear bound. Although our analysis depends on the underlying geometry, the algorithm itself is oblivious to it.

- [16] Bläsius, T., Friedrich, T., Katzmann, M., Krohmer, A., Striebel, J., [Towards a Systematic Evaluation of Generative Network Models](#). In: *Workshop on Algorithms and Models for the Web Graph (WAW)*, pp. 99–114, 2018.

Generative graph models play an important role in network science. Unlike real-world networks, they are accessible for mathematical analysis and the number of available networks is not limited. The explanatory power of results on generative models, however, heavily depends on how realistic they are. We present a framework that allows for a systematic evaluation of generative network models. It is based on the question whether real-world networks can be distinguished from generated graphs with respect to certain graph parameters. As a proof of concept, we apply our framework to four popular random graph models (Erdős-Rényi, Barabási-Albert, Chung-Lu, and hyperbolic random graphs). Our experiments for example show that all four models are bad representations for Facebook’s social networks, while Chung-Lu and hyperbolic random graphs are good representations for other networks, with different strengths and weaknesses.

- [17] Katzmann, M., Komusiewicz, C., [Systematic Exploration of Larger Local Search Neighborhoods for the Minimum Vertex Cover Problem](#). In: *Conference on Artificial Intelligence (AAAI)*, pp. 846–852, 2017.

We investigate the potential of exhaustively exploring larger neighborhoods in local search algorithms for Minimum Vertex Cover. More precisely, we study whether, for moderate values of k , it is feasible and worthwhile to determine, given a graph G with vertex cover C , if there is a k -swap S such that $(C \setminus S) \cup (S \setminus C)$ is a smaller vertex cover of G . First, we describe an algorithm running in $\Delta^{O(k)} \cdot n$ time for searching the k -swap neighborhood on n -vertex graphs with maximum degree Δ . Then, we demonstrate that, by devising additional pruning rules that decrease the size of the search space, this algorithm can be implemented so that it solves the problem quickly for $k \approx 20$. Finally, we show that it is worthwhile to consider moderately-sized k -swap neighborhoods. For our benchmark data set, we show that when combining our algorithm with a hill-climbing approach, the solution quality improves quickly with the radius k of the local search neighborhood and that in most cases optimal solutions can be found by setting $k = 21$.

- [18] Friedrich, T., Katzmann, M., Krohmer, A., [Unbounded Discrepancy of Deterministic Random Walks on Grids](#). In: *International Symposium on Algorithms and Computation (ISAAC)*, pp. 212–222, 2015.

Random walks are frequently used in randomized algorithms. We study a derandomized variant of a random walk on graphs, called rotor-router model. In this model, instead of distributing tokens randomly, each vertex serves its neighbors in a fixed deterministic order. For most setups, both processes behave remarkably similar: Starting with the same initial configuration, the number of tokens in the rotor-router model deviates only slightly from the expected number of tokens on the corresponding vertex in the random walk model. The maximal difference over all vertices and all times is called single vertex discrepancy. Cooper and Spencer (2006) showed that on \mathbb{Z}^d the single vertex discrepancy is only a constant c_d . Other authors also determined the precise value of c_d for $d = 1, 2$. All these results, however, assume that initially all tokens are only placed on one partition of the bipartite graph \mathbb{Z}^d . We show that this assumption is crucial by proving that otherwise the single vertex discrepancy can become arbitrarily large. For all dimensions $d \geq 1$ and arbitrary discrepancies $\ell \geq 0$, we construct configurations that reach a discrepancy of at least ℓ .