

CV Johannes Lengler

Personal Information

Johannes Lengler
Department of Computer Science, ETH Zürich
Universitätsstrasse 6, 8092 Zürich, Switzerland
Phone +41-44-632-70-53
[Email](#), [Webpage](#), [Google Scholar](#)



Education and Career

2012 – present	Researcher at ETH Zürich, since 2019: permanent faculty
2010 – 2012	PostDoc at ETH Zürich
2009 – 2010	PostDoc at Saarland University, Germany
2009	PhD thesis in Mathematics, summa cum laude (Saarland University, Advisor: Ernst-Ulrich Gekeler)

Teaching Experience

I took part in developing the very successful **bonus system** for the first year of Bachelor at D-IFNK, especially in the 2016/17 lectures “Algorithmen und Datenstrukturen” und “Algorithmen und Wahrscheinlichkeit”. Since then, I have helped shaping and refining the system as Head TA (several times for both lectures) and as lecturer (A&W).

I have given the following lectures at ETH:

Bachelor D-IFNK:

- Algorithmen und Wahrscheinlichkeit (1st year, once)
- Graphs and Algorithms (3rd year, four times)

Master D-IFNK (and advanced BSc students):

- Graphs and Algorithms: Advanced Topic (once)
- Algorithms, Probability, and Computing: Honors Course (twice)
- Randomized Algorithms and Probabilistic Method: Advanced Topics (once)

Service Lectures:

- Algorithmen und Komplexität (D-MATH, BSc 2nd year, four times)
- Diskrete Mathematik (D-ITET, BSc 2nd year, once)

Other:

- Interdisciplinary Algorithms Lab (D-IFNK MSc, three times)

Apart from that, I have served as Head TA (main organizer) for 8 lectures. I have supervised 22 BSc/MSc theses and was co-supervisor for 7 PhD students.

Algorithm Consulting

From 2018-2022, I have run the **Algorithm Consulting Service** of the D-IFNK. This service **has made the knowledge of the algorithm groups at ETH available to members of other departments.**

Recent Committee Activities (since 2018)

I have been in the **program committees** of GECCO (2018-2022), SSCI (2018), FOGA (2018-2022) and PPSN (2018-2022). I was **track chair** for the theory track at GECCO in 2019 and 2020, and I have been in the **organizing committee** of the 2022 **Dagstuhl seminar** “Theory of Randomised Optimization Heuristics”. I am currently **editor** of two special issues of *Algorithmica*.

Current Grants

I am currently supported by the [SNF Research Project 192079](#) DynaGIRG “Dynamic Processes in Geometric Inhomogeneous Random Graphs”, which finances two PhD students.

Academic Awards

- **Best Paper Award** at Genetic and Evolutionary Computation Conference (GECCO) 2016
- **Best Paper Nomination** at Parallel Problem Solving from Nature (PPSN) 2018
- **Hojjat Adeli Award** for Outstanding Contributions in Neural Systems, 2018
- **Best Paper Award** at the IEEE Symposium Series on Computation Intelligence (SSCI) 2018
- **Best Paper Nomination** at Foundations of Genetic Algorithms (FOGA) 2019
- **Best Paper Nomination** at Parallel Problem Solving from Nature (PPSN) 2022

Research (Summary)

My general aim as a scientist is to **bridge the gap between different fields of science**. I am proud to contribute actively to three different academic communities:

- **Mathematics** (e.g, publications in *Journal of Algebra* and *Journal of Number Theory* during my PhD; in *Annals of Applied Probability, Combinatorics, Probability & Computing, Annales de l'Institut Henri Poincaré* for my recent research)
- **Computer Science** (e.g., conferences *Principles of Distributed Computing, Genetic and Evolutionary Computation Conference*, journals *Journal of Computer and System Sciences, Algorithmica, Theoretical Computer Science*)
- **Neuroscience** (e.g., in *Frontiers in Computational Neuroscience, Biological Cybernetics, Hippocampus*)

Five Selected Publications

- [Drift Analysis](#).
(Book chapter in: *Theory of Evolutionary Computation*, Springer, 2020)
- [Geometric Inhomogeneous Random Graphs](#)
(*Theor. Comp. Sci.*, 2019, joint work with K. Bringmann, R. Keusch)
- [A General Dichotomy of Evolutionary Algorithms on Monotone Functions](#)
(*IEEE Trans. Evol. Comp.*, 2019)
- [Penalising transmission to hubs in scale-free spatial random graphs](#)
(*Ann. Henri Poincare*, 2021, joint work with J. Komjáthy, J. Lapinskas)
- [Greedy Routing and the Algorithmic Small-World Phenomenon](#)
(*J Comput Syst Sci*, 2022, joint work with K. Bringmann, R. Keusch, Y. Maus, A. Molla)

Appendix A: Scientific Achievements

In the following, I will describe some of my main scientific contributions, in three unrelated fields.

Geometric Inhomogeneous Random Graphs (GIRGs)

GIRGs are a random network model that I have developed together with Bringmann and my PhD student Keusch [Bringmann, Keusch & Lengler: ESA 2017; Theor. Comp. Sc. 2018]. The model has been picked up by other research groups in mathematics and has entered several textbooks [van der Hofstad 2016, Heydenreich, van der Hofstad 2017]. It has also been used in applications, for example for estimating the effectiveness of different interventions related to the CoViD19 pandemic [Jorritsma, Hulshof, Komjáthy 2020; Goldberg, Jorritsma, Komjáthy, Lapinskas 2021]. In my own work, I have investigated several dynamic processes on these networks, of which I describe two in detail.

Together with Koch, we have studied bootstrap percolation on GIRGs [Koch & Lengler, ICALP 2016; Internet Mathematics 2021]. In this process, for some initial geometric region B every node is infected with some probability p . Then in each round every uninfected node with at least r infected neighbors becomes infected, while infected nodes stay infected forever. We showed that this process (in the supercritical case) proceeds by a non-trivial mixture of infecting nodes of large degrees (super-spreaders) and geometric constraints. Qualitatively speaking, the infection spreads locally, but super-spreaders create new infection sources in different regions. In this way, the first nodes that are infected in a community are the local hubs of this region, and from these hubs a new local infection cascade starts.

The results in the paper are quantitative and remarkably precise. For each node v , the paper derives a round t_v depending on the expected degree of v and its distance to the source of infection, such that v is infected in a time interval $[(1-o(1))t_v, (1+o(1))t_v]$, asymptotically almost surely. Due to the precision of the results, we could derive a containment strategy that restricts the infection to a small portion of the network by removing (vaccinating) only a small number of vertices, even if the set of infected individuals is never exactly known. Remarkably, this containment strategy only works if the fraction of weak ties stays below a certain threshold.

This study exemplifies the type of result that I am aiming for as a researcher. It gives a deep (and quantitative!) understanding of the interplay between community structure and heavy-tail degree distribution, it derives policies for action in case of such infections, and it makes the role of weak ties transparent and explicit.

Together with Bringmann, Keusch, Maus, and Molla, we investigated whether *local routing* is possible in GIRGs [Bringmann et al., PODC 2017; J. Comp. Sys. Sci. 2022]. The task is to route a message from node s to node t , but each node has only local information about the network. I.e., each node only knows the weight and geometric position of its direct neighbors. In addition, s knows the geometric position of t . Then a routing protocol should find a walk through the network that reaches t , and that is not much longer than the shortest path from s to t .

Following the protocol of the famous small-world experiment by psychologist Stanley Milgram for the friendship graph, who studied routing of letters that were only allowed to be sent to direct friends, we started with the following basic scheme. Node s computes whether some of its neighbors is more likely to be adjacent to t than s itself. If so, then s forwards the message to the neighbor that maximizes this probability. Otherwise, the protocol fails. This repeats until t is reached, or until failure. Even this simple protocol has a constant probability to reach the target if s and t are chosen randomly. More importantly, the authors showed that this simple protocol is shockingly efficient. The ratio between the path of the protocol and a shortest path in the graph (the *stretch*) is only $1+o(1)$. This may seem surprising, but matches remarkably well the Milgram experiments, who also found very short paths.

We did not stop at this point, but developed the protocol further, to obtain a more practical protocol. By minor modifications, we could achieve a protocol which *always* finds the target if s and t are in the same component, but which still achieves a stretch of $1+o(1)$. The results are robust in the sense that approximate knowledge of weights and position of the neighbors suffice, and that it is resilient against random failure of edges. The existence of such local protocols had been an open problem for a decade [Krioukov et al., SIGCOMM 2007].

Nature-Inspired Search Heuristics

I have intensively studied nature-inspired optimization heuristics in the last years, including genetic and evolutionary algorithms (GAs and EAs). These algorithms are optimization paradigms, which successively refine a population of solutions by mutation, crossover, and selective pressure. Here I just describe one of my many contributions.

Over the last years we have investigated failure modes of some standard algorithms in seemingly simple situations. The insights provide guidance on which search heuristics are promising for which type of optimization problems. The starting point was the discovery that a surprising number of mutation-based EAs fail badly to optimize monotone pseudo-Boolean functions, although they are generally considered easy [Lengler & Steger, Comb., Prob., & Comp. 2018; Lengler, PPSN 2018; Lengler & Schaller, SSCI 2018, Lengler, Trans. Ev. Comp. 2019]. More precisely, these algorithms fail if the mutation operator is too aggressive. Moreover, the situation becomes substantially worse for large (but standard) population sizes [Lengler & Zou, FOGA 2019; Theor. Comp. Sci. 2021]. On the other hand, if crossover is used with a sufficiently large population, then a repair mechanism emerges that leads to efficient optimization (in sufficiently hard fitness landscapes), no matter how aggressive mutation is. These results were awarded with a **best paper nominations** at PPSN 2018 and at FOGA 2019, and a **best paper award** at SSCI 2018.

In a recent line of research, we could show that many of the above insights transfer to very simple and basic optimization problems in dynamic environments [Lengler & Meier, PPSN 2020, Lengler & Riedi, EvoCop 2021, Janett & Lengler, **best paper nomination** at PPSN 2022]. For example, consider the case that the objective function is always a linear function with positive weights, but sometimes the weights of the first half of parameters are larger than the weights of the second half, and sometimes vice versa. Since all weights are positive, the trivial solution is to set all parameters to one. However, this solution is not found by some standard optimization heuristics, and we could substantially increase our understanding of which heuristics succeed or fail in which dynamic environments.

We have also studied failure modes of self-adaptation, which is an important parameter control mechanism. Specifically, we could show that for any (possibly) dynamic monotone environment, a standard way of parameter control is highly susceptible to the aggressiveness of parameter update, but not if the parameter landscape is hard [Kaufmann, Larcher, Lengler & Zou, PPSN 2022]. This leads to the paradox situation that some algorithms succeed in vicinity of the optimum (where improvements are hard) but fail to find this vicinity. Although the underlying mechanism is different, this resembles a situation that we also could observe in our other lines of research [Lengler & Zou, Theor. Comp. Sci. 2021], suggesting that this situation may be more common than assumed.

Computational Neuroscience

In neuroscience I have analyzed the dynamics of biological neuronal networks. An important theme is to understand the role of randomness. Randomness is ubiquitous in the brain: in its components, in spontaneous firing, or in failure of links (some synapses fail in 80% of the cases). This randomness is often regarded as noise, and an engineer might think of the brain as a computing device that is robust despite of the noise. In [Lengler et al., PloS One, 2013], we have challenged this view. We could show that signal propagation in neural networks does not work *despite of*, but *because of* the variance in its component, the link failure, and the background random noise. All these noise sources increased speed and robustness of signal propagation. In a nutshell, a lack of noise in the system lead to too strong synchronization, which made the system less receptive to its input.

Together with my PhD students, I have used this insight in subsequent work, when we studied how simple rules can turn an initially random network into a functional unit, either for a high-capacity memory system [Einarsson et al., Front. Comp. Neurosci., 2014], a normalization circuit [Einarsson et al., Front. Comp. Neurosci., 2014], or a reliable signal propagation network [Weissenberger et al., Int. J. Neur. Sys., 2017]. The latter work was recognized with the **Hojjat Adeli Award for Outstanding Contributions in Neural Systems**.