

Link to the Latest Version • You can always find the latest version online at http://people.mpi-inf.mpg.de/~doerr/gecco17_tutorial_theory.pdf Benjamin Doerr: Theory for Non-Theoreticians

Instructor: Benjamin Doerr

- Benjamin Doerr is a full professor at the French École Polytechnique.
- He received his diploma (1998), PhD (2000) and habilitation (2005) in mathematics from Kiel University. His research area is the theory both of problem-specific algorithms and of randomized search heuristics like evolutionary algorithms. Major contributions to the latter include runtime analyses for evolutionary algorithms and ant colony optimizers, as well as the further development of the drift analysis method, in particular, multiplicative and adaptive drift. In the young area of black-box complexity, he proved several of the current best bounds.
- Together with Frank Neumann and Ingo Wegener, Benjamin Doerr founded the theory track at GECCO and served as its co-chair 2007-2009 and 2014. He is a member of the editorial boards of several journals, among them Artificial Intelligence, Evolutionary Computation, Natural Computing, and Theoretical Computer Science. Together with Anne Auger, he edited the book Theory of Randomized Search Heuristics.

Benjamin Doerr: Theory for Non-Theoreticians

This Tutorial: A Real Introduction to Theory

- GECCO, CEC, PPSN always had a good number of theory tutorials
- They did a great job in educating the theory community
- However, not much was offered for those attendees which
 - have little experience with theory
 - but want to understand what the theory people are doing (and why)
- This is the target audience of this tutorial. We try to answer those questions which come before the theory tutorials

Benjamin Doerr: Theory for Non-Theoreticians





Focus: EAs with Discrete Search Spaces

- We try to answer these questions independent of a particular subarea of theory
- However, to not overload you with definitions and notation, we focus on evolutionary algorithms on discrete search spaces
- Hence we intentionally omit examples from
 - genetic programming, estimation of distribution algorithms, ant colony optimizers, swarm intelligence, ...
 - all subareas of continuous optimization
- As said, this is for teaching purposes only. There is strong theory research in all these areas. All answers this tutorial give are equally valid for these areas

Benjamin Doerr: Theory for Non-Theoreticians

A Final Word Before We Start

- If I'm saying things you don't understand or if you want to know more than what I had planned to discuss, don't be shy to ask questions at any time!
 - This is "your" tutorial and I want it to be as useful for you as possible
- This is still a young tutorial. To further improve it, your feedback (positive and negative) is greatly appreciated!
 - → So talk to me after the tutorial, during the coffee breaks, social event, late-night beer drinking, ... or send me an email

Benjamin Doerr: Theory for Non-Theoreticians



What Do We Mean With Theory?

- Definition (for this tutorial): By theory, we mean results proven with mathematical rigor
- Mathematical rigor:
 - make precise the evolutionary algorithm (EA) you regard
 - make precise the problem you try to solve with the EA
 - make precise a statement on the performance of the EA solving this problem
 - prove this statement
- Example:

<u>Theorem</u>: The (1+1) EA finds the optimum of the OneMax test function $f: \{0,1\}^n \to \mathbb{R}; x \mapsto \sum_{i=1}^n x_i$ in an expected number of at most $en \ln(n)$ iterations. <u>Proof</u>: blah, blah, ...

Benjamin Doerr: Theory for Non-Theoreticians

Part I:

What We Mean by "Theory of EC"

Benjamin Doerr: Theory for Non-Theoreticians

Other Notions of Theory Theory: Mathematically proven results Experimentally guided theory: Set up an artificial experiment to experimentally analyze a particular question example: add a neutrality bit to two classic test functions, run a GA on these, and derive insight from the outcomes of the experiments Descriptive theory: Try to describe/measure/quantify observations example: some parts of landscape analysis

- <u>"Theories":</u> Unproven claims that (mis-)guide our thinking
 - example: building block hypothesis

Benjamin Doerr: Theory for Non-Theoreticians

12

10



Why Do Theory? Because of Results Absolute guarantee that the result is correct your can be sure reviewers can check truly the correctness of results readers can trust reviewers or, with moderate maths skills, check the correctness themselves Many results can only be obtained by theory; e.g., because you make a statement on a very large or even infinite set all bit-strings of length n, all TSP instances on n vertices, all input sizes n ∈ N, all possible algorithms for a problem

Why Do Theory? Because of the Approach A proof (automatically) gives insight in how things work (→ working principles of EC) why the result is as it is Self-correcting/self-guiding effect of proving: when proving a result, you are automatically pointed to the questions that need more thought Trigger for new ideas clarifying nature of mathematics playful nature of mathematicians

The Price for All This All this has a certain a price... Possible drawbacks of theory results include: • Restricted scope: So far, mostly simple algorithms could be analyzed for simple optimization problems • Less precise results: Constants are not tight, or not explicit as in " $O(n^2)$ " = "less than cn^2 for some unspecified constant c" Less specific results: You get a weaker guarantee for all problem instances instead of a stronger one for the instances that show up in your real-world application Theory results can be very difficult to obtain: The proof might be short and easy to read, but finding it took long hours Usually, there is no generic way to the solution, but you need a completely new, clever idea Benjamin Doerr: Theory for Non-Theoreticians 16

Theory and Experiments: Complementary Results		
THEORY	EXPERIMENTS	
 cover all problem instances of arbitrary sizes → guarantee! 	 only a finite number of instances of bounded size → have to see how representative this is 	
 proof tells you the reason 	 only tells you numbers 	
 only models for real-world instances (realistic?) 	 real-world instances 	
 limited scope, e.g., (1+1) EA limited precision, e.g., 0(n²) 	everything you can implementexact numbers	
 implementation independent 	 depends on implementation 	
 finding proofs can be difficult 	• can be cheap (well, depends)	
→ Ideal: Combine theory and expension and good experimental people to the second second s	riments. Difficulty: Get good theory people talk to each other	
Benjamin Doerr: Theory for Non-Theoreticians	17	

Part II:

A Guided Walk Through a Famous Theory Result

Benjamin Doerr: Theory for Non-Theoreticians

18

Outline for This Part • We use a simple but famous theory result • as an example for a non-trivial result • to show how to read a theory result • to explain the meaning of such a theoretical statement • to discuss typical shortcomings of theory results In an expected Reference: [DJW02] S. Di evolutionary al 2002. -- famous pape -- famous pape -- famous probuseful method Benjamin Doerr: Theory for Non-Theoreticians

A Famous Result

Theorem: The (1+1) evolutionary algorithm finds the maximum of any linear function

$$f: \{0,1\}^n \to \mathbb{R}, (x_1, \dots, x_n) \mapsto \sum_{i=1}^n w_i x_i, \qquad w_1, \dots, w_n \in \mathbb{R},$$

in an expected number of $O(n \log n)$ iterations.

[DJW02] S. Droste, T. Jansen, and I. Wegener. On the analysis of the (1+1) evolutionary algorithm. Theoretical Computer Science, 276(1–2):51–81, 2002.

- -- famous paper (500+ citations, maybe the most-cited pure EA theory paper)
- -- famous problem (20+ papers working on exactly this problem, many very useful methods were developed in trying to solve this problem) Benjamin Doerr: Theory for Non-Theoreticians 20













"Drift analysis":

Translate *expected*

progress into

expected (run-)time





- An unrealistically simple EA: the (1+1) EA
- Linear functions are artificial test function only
- Not a precise result, but only 0 (n log n) in [DJW02] or a most likely significantly too large constant in the [DJW12] result just shown
- \rightarrow We discuss these points on the following slides

Theorem: The (1+1) evolutionary algorithm finds the maximum of any linear function

 $f: \{0,1\}^n \to \mathbb{R}, (x_1, \dots, x_n) \mapsto \sum_{i=1}^n w_i x_i, \qquad w_1, \dots, w_n \in \mathbb{R},$

in an expected number of $O(n \log n)$ iterations.

Benjamin Doerr: Theory for Non-Theoreticians



Maybe Uncool: Only Linear Functions Again, this was the starting point. Today, we know how the (1+1) EA (and some other algorithms) compute Eulerian cycles [Neu04,DHN06,DKS07,DJ07] shortest paths [STW04,DHK07,BBD+09] minimum spanning trees [NW07,DJ10,Wit14] and many other "easy" optimization problems • We also have some results on approximate solutions for NP-complete problems like partition [Wit05], vertex cover [FHH+09,OHY09], maximum cliques [Sto06] We are optimistic that we will enlarge the set of problems we understand. However, like in many fields, it is also clear that "theory will always be behind"; that is, it will take guite some time until theoretical analyses become available for typical algorithms used in practice and realistic realworld problems Benjamin Doerr: Theory for Non-Theoreticians 31





- We have seen one of the most influential theory results: The (1+1) EA optimizes any linear function in $O(n \log n)$ iterations
- We have seen how to read and understand such a result
- We have seen why this result is important
 - non-trivial and surprising
 - gives insights in how EAs work
 - spurred the development of many important tools (e.g., drift analysis)
- · We have discussed strengths and limitations of theory results

Benjamin Doerr: Theory for Non-Theoreticians

Part III:

How Theory Can Contribute to a **Better Understanding** of EAs

Benjamin Doerr: Theory for Non-Theoreticians

34

Outline for This Part 3 ways how theory can help understanding and improving EAs 1. Debunk misconceptions 2. Help choosing the right parameters, representations, operators, and algorithms 3. Invent new representations, operators, and algorithms Benjamin Doerr: Theory for Non-Theoreticians



33



Misconception: Monotonic Functions are Easy to Optimize for EAs

- A function *f*: {0,1}ⁿ → ℝ is monotonically strictly increasing if the fitness increases whenever you flip a 0-bit to 1
 - special case of "no local optima" where each neighbor with more ones is better
- Misconception: Such functions are easy to optimize for standard EAs...
 - because already a simple hill-climber flipping single bits (randomized local search) does the job in time O(n log n)
- [DJS+13]: There is a monotonically strictly increasing function such that with high probability the (1+1) EA with mutation probability 16/n needs exponential time to find the optimum
 - very different from linear functions with positive weights: $O(n \log n)$ time

Benjamin Doerr: Theory for Non-Theoreticians

<list-item> Summary Misconceptions Intuitive reasoning or experimental observations can lead to wrong beliefs. It is hard to falsify them experimentally, because counter-examples may be rare (so random search does not find them). counter-examples may have an unexpected structure There is nothing wrong with keeping these beliefs as "rules of thumb", but is is important to distinguish between what is a rule of thumb and what is a roven fact. Theory is the right tool for this!

Contribution 2: Help Designing EAs When designing an EA, you have to decide between a huge number of design choices: the basic algorithm, the operators and representations, and the parameter settings. Theory can help you with deep and reliable analyses of scenarios similar to yours The question "what is a similar scenario" remains, but you have the same difficulty when looking for advice from experimental research 2 examples: fitness-proportionate selection edge-based representations for optimization problems in graphs

Benjamin Doerr: Theory for Non-Theoreticians



- Fitness-proportionate selection has been criticized (e.g., because it is not invariant under re-scaling the fitness), but it is still used a lot.
- Theorem [OW15]: If you use
 - the Simple GA as proposed by Goldberg [Gol89] (fitness-proportionate selection, comma selection)
 - to optimize the OneMax test function $f: \{0,1\}^n \to \mathbb{R}; x \mapsto x_1 + \dots + x_n$
 - with a population size *n*^{0.2499} or less

then with high probability the GA in a polynomial number of iterations does not create any individual that is 1% better than a random individual

- Interpretation: Most likely, fitness-proportionate selection and comma selection together make sense only in rare circumstances
 - more difficulties with fitness-proportionate selection: [HJKN08, NOW09

```
Benjamin Doerr: Theory for Non-Theoreticians
```

Designing EAs: Representations Several theoretical works on shortest path problems [STW04, DHK07, BBD+09], all use a vertex-based representation: each vertex points to its predecessor in the path mutation: rewire a random vertex to a random neighbor typical theorydriven curiosity [DJ10]: How about an edge-based representation? individuals are set of edges (forming reasonable paths) mutation: add a random edge (and delete the one made obsolete) • **<u>Result</u>**: All previous algorithms become faster by a factor of $\approx \frac{|V|^2}{|E|}$ [JOZ13]: edge-based representation also preferable for vertex cover Interpretation: While there is no guarantee for success, it may be useful to think of an edge-based representation for graph-algorithmic problems Benjamin Doerr: Theory for Non-Theoreticians 42

Contribution 3: Invent New Operators and Algorithms

- Theory can also, both via the deep understanding gained from proofs and by "theory-driven curiosity" invent new operators and algorithms.
- Example 1: What is the right way to do mutation?
 - A thorough analysis how EAs optimize jump functions suggests that we should use mutation operators such that the Hamming distance between parent and offspring follows a heavy-tailed distribution (and not a binomial one)
 - \rightarrow this GECCO, best-paper nominee in the GA track
- Example 2: The (1 + (λ, λ)) GA
 - Invent an algorithm that truly profits also from inferior search points

Benjamin Doerr: Theory for Non-Theoreticians

Example 1: Invent A New Mutation Operator · Short storyline: The recommendation to flip bits independently with probability 1/n might be overfitted to ONEMAX or other unimodal functions Longer storvline of this (longer) part: 4 young researchers ask themselves what is the right mutation rate to optimize jump functions (which are not unimodal) • surprise: for jump size m, the right mutation rate is m/n and this speeds-up things by a factor of $(m/e)^m$ • but: missing this optimal mutation rate by a factor of $(1 \pm \varepsilon)$ increases the runtime again by a factor of at least $\frac{1}{2}e^{m\varepsilon^2/5}$ reason: With standard-bit mutation, the Hamming distance between parent and offspring is strongly concentrated solution: design a mutation operator where this Hamming distance follows a power-law (not strongly concentrated) Benjamin Doerr: Theory for Non-Theoreticians 44

41













Solution: Heavy-tailed Mutation What do we need? Larger numbers of bits flip with reasonable probability No strong concentration 1-bit flips occur with constant probability (we don't want a massive slow-down for easy problems) NEW: Heavy-tailed mutation (with parameter β > 1): choose α ∈ {1, 2, ..., n/2} randomly with Pr[α] ~ α^{-β} [power-law distrib.] perform standard-bit mutation with mutation rate α/n Some maths: The probability to flip k bits is Θ(k^{-β}) → no exponential tails around the mean © Pr[1 bit flips] = Θ(1), e.g., ≈32% when β = 1.5 (≈37% for classic mut.)







Summary Heavy-tailed Mutation

- Key working principle:
 - reduce the probability of a 1-bit flip slightly (say from 37% to 32%)
 - distribute this free probability mass in a power-law fashion on all other k-bit flips
 - → increases the prob. for a *k*-bit flip from roughly $\frac{1}{e \cdot k!}$ to roughly $k^{-\beta}$ → reduces the waiting time for a *k*-bit flip from $e \cdot k!$ to k^{β}
 - this redistribution of probability mass is a good deal, because we usually spend much more time on finding a good multi-bit flip
 - $JUMP_{m,n}$: spend $\Theta(n \log n)$ time on 1-bit flips, but $\binom{n}{m}$ time to find the one necessary *m*-bit flip
- There is good reason to believe that heavy-tailed mutation is useful for a wide range of multi-modal problems.

```
Benjamin Doerr: Theory for Non-Theoreticians
```



New Algorithms (2/3)

- A simple idea to exploit inferior search points (in a (1+1) fashion):
 - 1. create λ mutation offspring from the parent by flipping λ random bits
 - 2. select the best mutation offspring ("mutation winner")
 - 3. create λ crossover offspring via a biased uniform crossover of mutation winner and parent, taking bits from mutation winner with probability $1/\lambda$ only
 - 4. select the best crossover offspring ("crossover winner")
 - 5. elitist selection: crossover winner replaces parent if not worse
- Underlying idea:
 - If λ is larger than one, then the mutation offspring will often be much worse than the parent (large mutation rates are destructive)
 - However, the best of the mutation offspring may have made some good progress (besides all destruction)
 - Crossover with parent repairs the destruction, but keeps the progress 57

Benjamin Doerr: Theory for Non-Theoreticians

New Algorithms (3/3) • Performance of the new algorithm, called $(1+(\lambda,\lambda))$ GA: • solves OneMax in time (=number of fitness evaluations) $O\left(\frac{n \log n}{n} + \frac{1}{n}\right)$ λn), which is $O(n \sqrt{\log n})$ for $\lambda = \sqrt{\log n}$ • the parameter λ can be chosen dynamically imitating the 1/5th rule, this gives an O(n) runtime experiments: • these improvements are visible already for small values of λ and small problem sizes nGP14]: good results for satisfiability problems Interpretation: Theoretical considerations can suggest new algorithmic ideas. Of course, much experimental work and fine-tuning is necessary to see how such ideas work best for real-world problems.

Summary Part 3

Theory has contributed to the understanding and use of EAs by

- debunking misbeliefs (drawing a clear line between rules of thumb and proven fact)
 - e.g., "no local optima" does not mean "easy"
- giving hints how to choose parameters, representations, operators, and algorithms
 - e.g., how useful is crossover when we hardly find an example where is provably improves things?
- inventing new representations, operators, and algorithms; this is fueled by the deep understanding gained in theoretical analyses and "theorydriven curiosity"

59

Part IV:

Current Topics of Interest in Theory of EC

Benjamin Doerr: Theory for Non-Theoreticians

Benjamin Doerr: Theory for Non-Theoreticians





Dynamic Parameter Choices Instead of fixing a parameter (mutation rate, population size, ...) once and forever (static parameter choice), it might be preferable to use parameter choices that change depending on time depending on the current state of the population depending on the performance in the past Hope: different parameter settings may be optimal early and late in the optimization process with self-adjusting parameters, we do not need to know the optimal parameters beforehand, but the EA finds them itself Experimental work suggests that dynamic parameter choices often outperform static ones (for surveys see [EHM99,KHE15]) Benjamin Doerr: Theory for Non-Theoreticians 63

Theory for Dynamic Parameter Choices: Deterministic Schedules

- Deterministic variation schedule for the mutation rate [JW00, JW06]:
 - Toggle through the mutation rates $\frac{1}{n}, \frac{2}{n}, \frac{4}{n}, \dots, \approx \frac{1}{2}$
 - Result: There is a function where this dynamic EA takes time $O(n^2 \log n)$, but any static EA takes exponential time
 - For most functions, the dynamic EA is slower by a factor of log n

Benjamin Doerr: Theory for Non-Theoreticians



Theory for Dynamic Parameter Choices: Success-based Dynamics

- Success-based choice of island number: You can reduce of the parallel runtime (but not the total work) of an island model when choosing the number of islands dynamically [LS11]:
 - double the number of islands after each iteration without fitness gain
 - half the number of islands after each improving iteration
- A success-based choice (1/5-th rule) of λ in the (1+(λ,λ)) GA automatically finds the optimal mutation strength [DD15a]
 - $\lambda := \sqrt[4]{F} \lambda$ after each iteration without fitness gain, F > 1 a constant
 - $\lambda := \lambda/F$ after each improving iteration
 - Important that *F* is not too large and that the fourth root is taken $(\rightarrow 1/5$ -th rule). The doubling scheme of [LS11] would not have worked

66

 Simple mechanisms to automatically find the current-best parameter setting (note: this is great even when the optimal parameter does not change over time)

Benjamin Doerr: Theory for Non-Theoreticians







- Strong Selection Weak Mutation (SSWM) algorithm [PPHST15], inspired by an inter-disciplinary project with populations-genetics:
 - worsening solutions are accepted with some positive probability
 - for improving offspring, acceptance rate depends on the fitness gain
 - Examples are given in [PPHST15] for which SSWM outperforms classic EAs
- Black-box complexity view: there are examples where any elitist algorithm is much worse than a non-elitist algorithm [DL15]
- State of the art: Not much real understanding apart from sporadic results. The fact that non-elitism is used a lot in EC practice asks for more work.

Benjamin Doerr: Theory for Non-Theoreticians

EAs are *black-box algorithms*: they learn about the problem at hand only by evaluating possible solutions
Algorithm

x, f(x)
(y, f(y))
f(z)

What is the price for such a problem-independent approach?

This is the main question in black-box complexity.

In short, the black-box complexity of a problem is the minimal number of function evaluations that are needed to solve it

= performance of the best-possible black-box algorithm

Benjamin Doerr: Theory for Non-Theoreticians

70



Black-Box Complexity vs. Games – Where EA Theory Meets Classic CS

- Black-box algorithms are strongly related to Mastermind-like guessing games:
 - algorithm guesses a search point
 - opponent reveals the fitness
- Such guessing games have a long history in classic computer science due to applications in security and privacy
- We have several (hidden) black-box complexity publications in classic CS venues (including a paper to appear in the *Journal of the ACM*)
 - EC theory meets classic theory
 - a chance to get the classic CS community interested in our field!

Benjamin Doerr: Theory for Non-Theoreticians



74



Summary

- Theoretical research gives deep insights in the working principles of EC, with results that are of a different nature than in experimental work
 - "very true" (=proven), but often apply to idealized settings only
 - for all instances and sizes, ..., but sometimes less precise
 - often only asymptotic results instead of absolute numbers
 - proofs tell us why certain facts are true
- The different nature of theoretical and experimental results implies that a real understanding is best obtained from a combination of both
- Theory-driven curiosity and the clarifying nature of mathematical proofs can lead to new ideas, insights and algorithms



Recent Books (Written for Theory People, But Not Too Hard to Read)



- Jansen (2013). Analyzing Evolutionary Algorithms, Springer
- Neumann/Witt (2010). Bioinspired Computation in Combinatorial Optimization, Springer

```
Benjamin Doerr: Theory for Non-Theoreticians
```

```
78
```





Appendix A

Glossary of Terms Used in This Tutorial

Benjamin Doerr: Theory for Non-Theoreticians

81

Discrete and Pseudo-Boolean Optimization

In this tutorial we are mostly interested in the optimization of problems of the type $f: \{0,1\}^n \to \mathbb{R}$

- Problems f: S → ℝ with finite search space S are called *discrete* optimization problems

 (in contrast to *continuous problems* f: ℝⁿ → ℝ or, more generally f: S → ℝ with continuous S)
- When $S = \{0,1\}^n$ and $f: \{0,1\}^n \to \mathbb{R}$, we call f a pseudo-Boolean function
- Please note: don't get fooled! Even if optimizing a function *f*: {0,1}ⁿ → ℝ
 may look harmless, a HUGE range of problems (even NP-hard ones like
 Max-SAT and many others!) can be expressed this way

Benjamin Doerr: Theory for Non-Theoreticians

82

What we Mean by "Optimization"

- Recall: we assume that we aim at optimizing a function $f: \{0,1\}^n \to \mathbb{R}$
- For this tutorial "optimization" = maximization, that is, we aim at finding a bit string x = (x₁,..., x_n) such that f(x) ≥ f(y) for all y ∈ {0,1}ⁿ
- Note in particular: we are not interested in this tutorial in identifying local optima, only the global best solution(s) are interesting for us



Expected Runtimes – Introduction

- All EAs are *randomized algorithms*, i.e., they use random decisions during the optimization process (for example, the *variation step*, i.e., the step in which new search points are generated, is often based on random decisions---we will discuss this in more detail below)
- Our object of interest, the *runtime of EAs*, is the number of function evaluations that an EA needs until it queries for the first time an optimal solution. Since EAs are randomized algorithms, their runtime is a random variable

Benjamin Doerr: Theory for Non-Theoreticians











References (1/6)

[AD11] [BBD ⁺ 09]	Anne Auger and Benjamin Doerr. Theory of Randomized Search Heuristics. World Scientific, 2011. Surender Baswana, Somenath Biswas, Benjamin Doerr, Tobias Friedrich, Plyush P. Kurur, and Frank Neumann. Computing single source shortest paths using single-objective fitness functions. In Proc. of Foundations of Genetic Algorithms (FOGA), pages 59–66. ACM, 2009.		
[BDN10]	S. Böttcher, B. Doerr, and F. Neumann. Optimal fixed and adaptive mutation rates for the LeadingOnes problem. In Proc. of Parallel Problem Solving from Nature (PPSN), pages 1–10. Springer, 2010.		
[DD15a]	Benjamin Doerr and Carola Doerr. Optimal parameter choices through self-adjustment: Applying the 1/5 rule in discrete settings. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 1335–1342. ACM, 2015.	-th	
[DD15b]	Benjamin Doerr and Carola Doerr. A tight runtime analysis of the $(1+(\lambda, \lambda))$ genetic algorithm on OneM In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 1423–1430. ACM, 2015.	ax.	
[DDE15]	Benjamin Doerr, Carola Doerr, and Franziska Ebel. From black-box complexity to designing new genetic algorithms. <i>Theoretical Computer Science</i> , 567:87 – 104, 2015.		
[DDY16]	Benjamin Doerr, Carola Doerr, and Jing Yang. Optimal parameter choices via precise black-box analysis. Proc. of Genetic and Evolutionary Computation Conference (GECCO). ACM, 2016. To appear.	In	
[DHK07]	Benjamin Doerr, Edda Happ, and Christian Klein. A tight analysis of the (1+1)-EA for the single source shortest path problem. In <i>Proc. of Congress on Evolutionary Computation (CEC)</i> , pages 1890–1895. IEEE 2007.	£,	
[DHK08]	Benjamin Doerr, Edda Happ, and Christian Klein. Crossover can provably be useful in evolutionary computation. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 539–546. A 2008.	.CM,	
[DHN06]	Benjamin Doerr, Nils Hebbinghaus, and Frank Neumann. Speeding up evolutionary algorithms through restricted mutation operators. In Proc. of Parallel Problem Solving from Nature (PPSN), volume 4193 of Lecture Notes in Computer Science, pages 978–987. Springer, 2006.		
[DJ07]	Benjamin Doerr and Daniel Johannsen. Adjacency list matchings: an ideal genotype for cycle covers. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 1203–1210. ACM, 2007.		
Benjamir	Doerr: Theory for Non-Theoreticians	91	

References (2/6)

[DJ10] Benjamin Doerr and Daniel Johannsen. Edge-based representation beats vertex-based representation in shortest path problems. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 758-766, ACM, 2010. [DJK⁺11] Benjamin Doerr, Daniel Johannsen, Timo Kötzing, Per Kristian Lehre, Markus Wagner, and Carola Winzen. Faster black-box algorithms through higher arity operators. In Proc. of Foundations of Genetic Algorithms (FOGA), pages 163-172, ACM, 2011. [DJK⁺13] Benjamin Doerr, Daniel Johannsen, Timo Kötzing, Frank Neumann, and Madeleine Theile. More effective crossover operators for the all-pairs shortest path problem. Theoretical Computer Science, 471:12-26, 2013. [DJS⁺13] Benjamin Doerr, Thomas Jansen, Dirk Sudholt, Carola Winzen, and Christine Zarges. Mutation rate matters even when optimizing monotonic functions. Evolutionary Computation, 21:1-27, 2013. [DJW98] Stefan Droste, Thomas Jansen, and Ingo Wegener. A rigorous complexity analysis of the (1 + 1) evolutionary algorithm for separable functions with Boolean inputs. Evolutionary Computation, 6:185-196, 1998. [DJW02] Stefan Droste, Thomas Jansen, and Ingo Wegener. On the analysis of the (1+1) evolutionary algorithm. Theoretical Computer Science, 276:51-81, 2002. [DK15] Benjamin Doerr and Marvin Künnemann. Optimizing linear functions with the $(1+\lambda)$ evolutionary algorithm different asymptotic runtimes for different instances. Theoretical Computer Science, 561:3-23, 2015. [DKS07] Benjamin Doerr, Christian Klein, and Tobias Storch. Faster evolutionary algorithms by superior graph representation. In Proc. of Symposium on Foundations of Computational Intelligence (FOCI), pages 245-250. IEEE, 2007. [DL15] Carola Doerr and Johannes Lengler. Elitist black-box models: Analyzing the impact of elitist selection on the performance of evolutionary algorithms. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 839-846. ACM, 2015.

Benjamin Doerr: Theory for Non-Theoreticians

References (3/6)

- [Doe16] Benjamin Doerr. Optimal parameter settings for the (1 + (λ, λ)) genetic algorithm. In Proc. of Genetic and Evolutionary Computation Conference (GECCO). ACM, 2016. To appear. (dPdLDD15) Axel de Perthuis de Laillevault, Benjamin Doerr, and Carola Doerr. Money for nothing: Speeding up
- evolutionary algorithms through better initialization. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 815–822. ACM, 2015.
 [Dro02] Stefan Droste. Analysis of the (1+1) EA for a dynamically changing OneMax-variant. In Proc. of Congress on
- [Droos] Stefan Droste, Analysis of the (1+1) EA for a dynamically changing Onemax-variant. In Proc. of Congress on Evolutionary Computation (CEC), pages 55–60. IEEE, 2002.
 [Droos] Stefan Droste, Analysis of the (1+1) EA for a dynamically bitwise changing OneMax. In Proc. of Genetic and
- [Proof] Steam Diotect Amayaso of the (1+1) Error or opamineary powers changing include: in Proc. of Ochever and Evolutionary Computation Conference (GECCO), volume 2723 of Lecture Notes in Computer Science, pages 909–921. Springer, 2003.
- [DSW13] Benjamin Doerr, Dirk Sudholt, and Carsten Witt. When do evolutionary algorithms optimize separable functions in parallel? In Proc. of Foundations of Genetic Algorithms (FOGA), pages 51–64. ACM, 2013.
- [DT09] Benjamin Doerr and Madeleine Theile. Improved analysis methods for crossover-based algorithms. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 247–254. ACM, 2009.
 [EHM99] Agoston Endre Eiben, Robert Hinterding, and Zbigniew Michalewicz. Parameter control in evolutionary
- [Eritys] Agoston Endre Enden, Robert Hinderding, and Zoginew Michaewicz. Farameter Control in evolutionary algorithms. IEEE Transactions on Evolutionary Computation, 3:124–141, 1999.
 [FHH⁺09] Tobias Friedrich, Jun He, Nils Hebbinghaus, Frank Neumann, and Carsten Witt. Analyses of simple hybrid
- [FHH⁺09] Tobias Friedrich, Jun He, Nils Hebbinghaus, Frank Neumann, and Carsten Witt. Analyses of simple hybrid algorithms for the vertex cover problem. Evolutionary Computation, 17:3–19, 2009.
 [FM92] Stenhanie Forrest and Melanie Mitchell, Relative building-block fluxes and the building block hypothesis. In
- [FM92] Stephanie Forrest and Melanie Mitchell. Relative building-block fitness and the building block hypothesis. In Proc. of Foundations of Genetic Algorithms (FOGA), pages 109–126. Morgan Kaufmann, 1992.
 [FW04] Simon Fischer and Ingo Wegener. The ising model on the ring: Mutation versus recombination. In Proc. of
- [PW04] Sinon Fischer and ingo wegener. The sing model on the ring. Autoation vessus recombination. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), volume 3102 of Lecture Notes in Computer Science, pages 1113–1124. Springer, 2004.
- [Gol89] David E. Goldberg. Genetic Algorithms in Search, Optimization and Machine Learning. Addison-Wesley Longman Publishing Co., Inc., 1989.
- [GP14] Brian W. Goldman and William F. Punch. Parameter-less population pyramid. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 785–792. ACM, 2014.

Benjamin Doerr: Theory for Non-Theoreticians

93

References (4/6)

- [GW15] Christian Gießen and Carsten Witt. Population size vs. mutation strength for the (1+λ) EA on OneMax. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 1439–1446. ACM, 2015.
- [HGD94] Jeff Horn, David Goldberg, and Kalyan Deb. Long path problems. In Proc. of Parallel Problem Solving from Nature (PPSN), LNCS 866, pages 149–158. Springer, 1994.
- [HJKN08] Edda Happ, Daniel Johannsen, Christian Klein, and Frank Neumann. Rigorous analyses of fitness-proportional selection for optimizing linear functions. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 953–960. ACM, 2008.
- [Hol75] J.H. Holland. Adaptation in Natural and Artificial Systems. University of Michigan Press, 1975.
- [Jan13] Thomas Jansen. Analyzing Evolutionary Algorithms—The Computer Science Perspective. Springer, 2013.
- [JOZ13] Thomas Jansen, Pietro Simone Oliveto, and Christine Zarges. Approximating vertex cover using edge-based representations. In Proc. of Foundations of Genetic Algorithms (FOGA), pages 87–96. ACM, 2013.
- [JS05] Thomas Jansen and Ulf Schellbach. Theoretical analysis of a mutation-based evolutionary algorithm for a tracking problem in the lattice. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 841–848. ACM, 2005.
- [JW99] Thomas Jansen and Ingo Wegener. On the analysis of evolutionary algorithms A proof that crossover really can help. In Proc. of European Symposium of Algorithms (ESA), volume 1643 of Lecture Notes in Computer Science, pages 184–193, Springer, 1999.
- [JW05] Thomas Jansen and Ingo Wegener. Real royal road functions-where crossover provably is essential. Discrete Applied Mathematics, 149(1-3):111–125, 2005.
- [JW06] Thomas Jansen and Ingo Wegener. On the analysis of a dynamic evolutionary algorithm. Journal of Discrete Algorithms, 4:181–199, 2006.
- [KHE15] G. Karafotias, M. Hoogendoorn, and A.E. Eiben. Parameter control in evolutionary algorithms: Trends and challenges. IEEE Transactions on Evolutionary Computation, 19:167–187, 2015.
- [KM12] Timo Kötzing and Hendrik Molter. ACO beats EA on a dynamic pseudo-boolean function. In Proc. of Parallel Problem Solving from Nature (PPSN), volume 7491 of Lecture Notes in Computer Science, pages 113–122. Springer, 2012.

Benjamin Doerr: Theory for Non-Theoreticians

94

References (5/6) [LS11] Jörg Lässig and Dirk Sudholt. Adaptive population models for offspring populations and parallel evolutionary algorithms. In Proc. of Foundations of Genetic Algorithms(FOGA), pages 181-192. ACM, 2011. [LW12] Per Kristian Lehre and Carsten Witt. Black-box search by unbiased variation. Algorithmica, 64:623-642, 2012. [LW14] Andrei Lissovoi and Carsten Witt. MMAS vs. population-based EA on a family of dynamic fitness functions. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 1399–1406. ACM, 2014. [LW15] Andrei Lissovoi and Carsten Witt. Runtime analysis of ant colony optimization on dynamic shortest path problems Theoretical Computer Science, 561:73-85, 2015. F. Neumann. Expected runtimes of evolutionary algorithms for the eulerian cycle problem. In Proc. of [Neu04] Congress on Evolutionary Computation (CEC), pages 904–910. IEEE, 2004. Frank Neumann, Pietro Simone Oliveto, and Carsten Witt. Theoretical analysis of fitness-proportional [NOW09] selection: landscapes and efficiency. In Proc. of Genetic and Evolutionary Computation Conference (GECCO). pages 835-842 ACM 2009 [NW07] Frank Neumann and Ingo Wegener. Randomized local search, evolutionary algorithms, and the minimum spanning tree problem. Theoretical Computer Science, 378:32-40, 2007. Frank Neumann and Carsten Witt. Bioinspired Computation in Combinatorial Optimization - Algorithms and [NW10] Their Computational Complexity. Springer, 2010. [OHY09] Pietro Simone Oliveto, Jun He, and Xin Yao. Analysis of the (1+1) -ea for finding approximate solutions to vertex cover problems. IEEE Transactions on Evolutionary Computation, 13:1006-1029, 2009. [OW15] Pietro Simone Oliveto and Carsten Witt. Improved time complexity analysis of the simple genetic algorithm. Theoretical Computer Science, 605:21-41, 2015 [OZ15] Pietro Simone Oliveto and Christine Zarges. Analysis of diversity mechanisms for optimisation in dynamic environments with low frequencies of change. Theoretical Computer Science, 561:37-56, 2015. 95 Benjamin Doerr: Theory for Non-Theoreticians

References (6/6)

- [PPHST15] Tiago Paixao, Jorge Pérez Heredia, Dirk Sudholt, and Barbora Trubenova. First steps towards a runtime comparison of natural and artificial evolution. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 1455–1462. ACM, 2015.
 [Rud97] Günter Rudolph. Convergence Properties of Evolutionary Algorithms. Kovac, 1997.
- [Sto06] Tobias Storch. How randomized search heuristics find maximum cliques in planar graphs. In Proc. of Genetic
- and Evolutionary Computation Conference (GECCO), pages 567–574. ACM, 2006. [STW04] Jens Scharnow, Karsten Tinnefeld, and Ingo Wegener. The analysis of evolutionary algorithms on sorting and
- shortest paths problems. Journal of Mathematical Modelling and Algorithms, 3:349–366, 2004.
 [Sud05] Dirk Sudholt. Crossover is provably essential for the ising model on trees. In Proc. of Genetic and Evolutionary Computation Conference (GECCO). ACM, 2005.
- [SW04] Tobias Storch and Ingo Wegener. Real royal road functions for constant population size. Theoretical Computer Science. 320(1):123–134, 2004.
- [Wit05] Carsten Witt. Worst-case and average-case approximations by simple randomized search heuristics. In Proc. of Symposium on Theoretical Aspects of Computer Science (STACS), pages 44–56. Springer, 2005.
- [Wit13] Carsten Witt. Tight bounds on the optimization time of a randomized search heuristic on linear functions. Combinatorics, Probability & Computing, 22:294–318, 2013.
- [Wit14] Carsten Witt. Revised analysis of the (1+1) EA for the minimum spanning tree problem. In Proc. of Genetic and Evolutionary Computation Conference (GECCO), pages 509–516. ACM, 2014.

Benjamin Doerr: Theory for Non-Theoreticians