

# Link to the Latest Version, BibTeX Entry

- I will almost surely update these slides before the tutorial in July
- You can find the latest version of this tutorial online at http://www-ia.lip6.fr/~doerr/GECC017tutorial.pdf
- · You can cite this tutorial (thanks in advance ;) ) as follows
  - Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation. GECCO (Companion). ACM, 2017
- BibTeX entry: @inproceedings{Doerr17tutorial, author = {Carola Doerr}, title = {Non-Static Parameter Choices in Evolutionary Computation}, booktitle = {Genetic and Evolutionary Computation Conference, {GECCO} 2017, Berlin, Germany, July 15-19, 2017, Companion Material Proceedings}, pages = {tbd}, year = {2017}, url = {http://doi.acm.org/10.1145/3067695.3067707}, doi = {10.1145/3067695.3067707} publisher = {ACM}}
   Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

# Instructor: Carola Doerr

- Carola Doerr, née Winzen, is a permanent researcher with the CNRS and Pierre et Marie Curie University (Paris 6).
- Carola's main research interest is in the *theory of randomized* search heuristics, both in the design of efficient algorithms as well as in lower bounds for such general-purpose optimization techniques. After contributing to the revival of black-box complexity, a theory-guided approach to explore the limitations of heuristic search algorithms, she recently started a
   series of works aimed at exploiting insights from the theory of evolutionary computation to design more efficient EAs, in particular such with a non-static choice of parameters.
- Carola has been co-chair of the theory track at GECCO 2015 and 2017 and served as tutorial chair at PPSN 2016. She is editor of two special issues in Algorithmica. From 2014 to 2016 she has been a co-organizer of the women@GECCO workshop series.
- She studied mathematics at Kiel University (Diploma in 2007) and computer science at the Max Planck Institute for Informatics and Saarland University (PhD in 2011). Her PhD studies were supported by a Google Europe Fellowship in Randomized Algorithms. Carola's thesis has been awarded the Otto Hahn Medal of the Max Planck Society. From Dec. 2007 to Nov. 2009, Carola Doerr has worked as a business consultant for McKinsey & Company. She was a post-doc at the Université 7 in Paris and the Max Planck Institute for Informatics in Saarbrücken.
   Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

# What We Will Discuss Today

- 1. (Almost) all EAs depend on a number of parameters
- 2. Choosing the right parameters of an EA is an important task
- 3. Choosing the right parameters of an EA is a difficult task
- 4. Parameter Tuning vs. Parameter Control: why it is (almost) always better to use non-static parameters
- 5. Parameter Control Mechanisms
  - 1. Which parameters should be updated?
  - 2. Which effects should trigger an update?
  - 3. How should we update the parameters and how do we classify the different update schemes in EC?



- The performance of an EA depends on
  - the components of the EA
  - the operators in use
  - the representation of the problem/ the model of the fitness function and the interactions among these!
- We won't discuss how to chose
  - the "best" algorithm for your problem
  - nor how to find a good fitness function to model your problem
  - $\rightarrow$  we will take the algorithm(s) and problem(s) as given and ask ourselves how to find, for this given setting, good parameter values





# A Word of Warning: Scope (2/2) In 110 minutes, I cannot offer you an exhaustive literature review. You can find very good surveys here (see reference list for details) Eiben, Hinterding, Michalewicz, IEEE TEC, 1999 [EHM99] Eiben, Michalewicz, Schoenauer, Smith, 2007 [EMSS07] Karafotias, Hoogendoorn, Eiben, IEEE TEC, 2015 [KHE15] Those of you who are interested in runtime analysis works on nonstatic parameter choices can take the survey in Section 6.1 in Doerr, Doerr: Optimal Static and Self-Adjusting Parameter Choices for the $(1+(\lambda,\lambda))$ Genetic Algorithm [DD17] as a starting point for further investigations $\rightarrow$ I have decided to discuss in-depth a few mechanisms that should give you a flavor of what has been done and what is possible (and how to get there!) We will mostly focus on discrete optimization in continuous optimization, adaptive parameter choices are standard similar mechanisms are used in continuous optimization, often (but not always) originating from a similar source of inspiration Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

6













# How to Find Good Parameter Values? (2/3)

- <u>"Modern view" of parameter selection</u>: no globally optimal parameters exist
   → the right choice of parameter values crucially depends on the problem that we face and the algorithm that we employ
  - Result: whenever we face a new problem, or employ a new algorithm, we need to ask ourselves how to set the parameters
  - Very often, some (often many !) preliminary experiments are conducted to find reasonable parameter values
  - This parameter tuning quite often is a difficult task (see next slide)
  - In recent years, there is a substantial amount of EA literature on how to find good (static) parameter values, so-called parameter tuning mechanisms (see below for discussion and references)
  - Similar situation in the theory of EC: typical research question concerns the performance of a given algorithm with respect to some fixed set of parameters.
  - The bulk of EC papers (with a focus on discrete optimization problems) falls into this category of analyzing performance with respect to some fixed set of parameters! (How about your latest GECCO paper?)

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

# How to Find Good Parameter Values? (1/3)

- <u>"Sports" of the 70s/80s in EC:</u> Finding good parameter values
  - good = "globally good", i.e., for a broad range of problems
  - Examples: De Jong [DJ75], Grefenstette [Gre86] give recommendations for parameters such as population size, mutation and crossover probabilities, selection strategies, etc.
    - → these recommendations are *independent* of problem class, problem size, ... (*absolute values*)
  - Mühlenbein [Müh92] and others suggest 1/n as mutation rate for problems of lengths n (relative values)
    - Note: we know today that this choice indeed works well for a broad range of problems, cf. discussion below. However, it is widely acknowledged today, that problem size is not the only feature that matters.

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

16

# **Difficulty of Finding Good Parameter Choices**

- 1. Even if we find "optimal" parameter values for one problem, these may (!, don't have to) be much different for similarly-looking problems (which is the basis for so-called *parameter choice by analogy*)
- 2. Small changes in one parameter can (!, don't have to) cause huge performance gaps
  - Many empirical works on this matter exist (again, check this year's GECCO talks to see if/how much effort has been put into finding the right parameters <sup>(2)</sup>)
  - Those of you interested in theoretical results can find in [DoerrJS<sup>+</sup>13] [Doerr, Jansen, Sudholt, Winzen, Zarges: Mutation Rate Matters Even When Optimizing Monotonic Functions. Evolutionary Computation, 2013] or [LS16] [Lengler, Steger: Drift analysis and evolutionary algorithms revisited, arXiv] examples where changing the mutation rate by a small constant factor changes the expected running time from a small polynomial (e.g.,  $O(n \log n)$ ) to super-polynomial/exponential

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation





# Why Use Non-Static Parameters?

# VERY intuitive motivation:

- 1. Different stages of the optimization process require different parameter values!
  - Example:
    - beginning = "exploration phase" → large mutation rate/small selective pressure to make large jumps and discover different areas of the search space
- Dace
  - end = "exploitation phase"
     → small mutation rates/high selective pressure to focus the search
- 2. No need for the user to identify good parameter values
  - Parameter tuning takes time and is quite complex (tuning 1 parameter is difficult already, but the operators also interact with each other, tuning 2 or more parameters typically requires non-sequential optimization, which is a difficult task)
- Hope is that the algorithm identifies good parameter values itself
  Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

# Your Experience with Non-Static Parameter Values

- 1. Have you already experimented with non-static parameter values?
  - 1. What was your motivation?
  - 2. How did you do this?
  - 3. What did go well, what did not go well?
- If you haven't experimented with this idea, why not?
   (e.g., never crossed my mind, not convinced by the motivation, no time to work on this, my algorithm doesn't have any parameters,...)

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

20













# Sketching the Classification Scheme of [EHM99] Deterministic Parameter Control

Parameter Control Idea 1: Updates should follow a similar pattern

- Most popular example: exploration first, then exploitation
  - mutation rate: large in the beginning, smaller towards the end
  - selection strength: more generous in the beginning, higher selection pressure in the end
- to stimulate or enforce this behavior, we can change the parameters based on the time elapsed (number of generations, fitness evaluations, wall-clock time, etc.)
  - Simple examples:
    - cooling schedule of the selective pressure ("temperature") in Boltzmann selection of Simulated Annealing
    - start with mutation rate p = 1/2, decrease p after 10,000 fitness evaluations
    - after each 1,000 iterations, draw a random mutation probability



# **Remarks on "Deterministic" Update Schemes**

- Already Eiben, Hinterding, and Michalewicz noted in their work [EHM99] that the term "deterministic" is sub-optimal (update rules may be random as in the third example on the previous slide)
- More suitable terms could be
  - "time-dependent", "scheduled" update scheme, or
  - "feedback-free", "progress-independent" update scheme

but in lack of a widely acknowledged alternative, "deterministic update rule" is still the predominantly used term

• The crucial feature here is that there is no feedback from the optimization process, so the update rule is determined in advance, before the actual run of the algorithm

(Note: when random decisions are involved, then it would be possible to run the experiments determining the parameter value before the EA is started)

 Note that finding the optimal deterministic update rules requires tuning, i.e., while they bypass the disadvantage of the non-flexible static parameter values, they do not allow the algorithm to identify the good parameter values by itself

30

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

**Examples for Deterministic Parameter Choices (1/2)** Some selected theory works: Hesser and Männer (PPSN'90) [HM90] suggested the following rule for the mutation strength of a GA with population size  $\lambda$  for OneMax:  $p_m(t) \coloneqq \frac{\sqrt{\frac{\alpha}{\beta}}\exp\left(-\frac{\gamma t}{2}\right)}{\lambda\sqrt{n}}$  where  $\alpha, \beta, \gamma$  are constants Jansen Wegener [JW06]: mutation rate changes in every iteration •  $p_t(n) \coloneqq 2^i/n$  where  $i \equiv (t-1) \mod \log(n) - 1$ +/- very frequent changes  $\rightarrow$  non-stable algorithm - worse performance on simple functions like OneMax, linear functions, LeadingOnes, etc. + examples where better performance than any static choice can be proven Doerr, Doerr, Kötzing (GECCO 2016) [DDK16b]: in every iteration, a random step size is used for a multi-valued OneMax-type problem (problem will be discussed in more detail in the next section, along with a self-adjusting parameter choice) Note: non-static parameter values, but static probability distribution used here! Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation 31

•	Random Variation of the Population Size GA (RVPS) by Costa, Tava and Rosa [CTR99]	es
	<ul> <li>size of the actual population is changed every N fitness evaluation for a given N (according to some monotonous rule)</li> </ul>	าร,
	Both shrinking and increasing the population size are considered	
•	Saw-tooth like population size growth considered by	
	<ul> <li>Koumousis and Katsaras in [KK06] (TEC 2006): linear decrease of population size with eventual re-initialization of the population size adding randomly selected individuals</li> </ul>	∂f ≥ b
	• Hu, Harding, Banzaf [HHB10]: inverse saw-tooth like population s	ize
ar	ola Doerr: Non-Static Parameter Choices in Evolutionary Computation	





# Examples for Self-Adaptive Parameter Choices We won't discuss this in much detail, but if you are interested in such mechanisms, you can start your investigations with the following works Bäck PPSN'92 [Bäc92] and follow-up works: extends the chromosome by 20 bits. Mutation works as follows: Decoding the 20 bits to the individual's own mut. rate pm Mutating the bits encoding pm with mutation probability pm Decoding these changed bits to p'm Mutating the bits that encode the solution with mutation probability p'm Dang, Lehre (PPSN'16) [DL16]: theoretical work on a self-adaptive choice of the mutation strength in a non-elitist population

# Sketching the Classification Scheme of [EHM99] Adaptive Parameter Control

- Parameter Control Idea 3:
  - use feedback from the optimization process to change the parameters according to some pre-described rule
- Relevant feedback includes:
  - success-based rules: presence/absence of progress (Example: "if the iteration was successful, increase mutation rate, and decrease it otherwise")
  - fitness-based update rules: magnitude of parameter change depends on the magnitude of the progress or the fitness of the current-best search point(s) (Example: fitness- or ranking-based mutation rates)





# **Discussion of the Classification Scheme** The classification scheme of [EHM99] is the most widely accepted one While historically there has been guite some work on deterministic and self-adaptive update rules, today the most commonly applied rules are adaptive These adaptive rules can be much different in flavor, as we shall see below As mentioned earlier, after 18 years of existence, it may be the right time for a new classification scheme In [DD15a] (GECCO'15) we suggested to distinguish between • functionally-dependent adaptive schemes: parameters are in functional dependence of current population. That is, only the current state matters, not the process to achieve this state self-adjusting adaptive schemes: parameters depend on success of previous iterations Other discriminations are much needed (see discussions below) but haven't been established to date Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation 39











# **Examples for Simple Success-Based Rules**

- Simple success-based (multiplicative) parameter update mechanisms have been experimented with in different research streams of EC
- In what follows, I give a few examples, mostly stemming from the <u>theory of</u> <u>EC</u>
  - → the simple algorithms and problems regarded there allow us to concentrate on the main ideas
  - → for some of the considered ("toy") problems, it is sometimes possible to formally prove that the adaptive parameter choices outperform any (!) static one
  - → while all the mentioned works serve as a showcase that selfadjustment is feasible and brings performance gains, for none of them has there been a thorough investigation of how much one can gain by tuning the adjustment rules (see discussion below), so there is a lot of room for us to experiment and to learn !
- Empirical works which can serve as a starting point for further investigations on the simple success-based (multiplicative) parameter update mechanisms include [Aug09,KMH<sup>+</sup>04]
   Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

Simple Success-Based Rules: Example 1
Lässig, Sudholt: Adaptive Population Models for Offspring Populations and Parallel Evolutionary Algorithms, FOGA 2011 [LS11]:
regard the (1 + λ) EA
an iteration is called *successful* if it produces an offspring of better

- than previous best fitness value
- Scheme A:
  - If (iteration not successful)  $\rightarrow$  double  $\lambda$ If (iteration successful)  $\rightarrow$  reduce  $\lambda$  to 1
- Scheme B:
  - If (iteration not successful) → double λ
     If (iteration successful) → halve λ
- Main results: decreased expected parallel optimization times without increasing the expected sequential runtime for problems like OneMax, LeadingOnes, Jump, unimodal functions

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

47

# Simple Success-Based Rules: Example 1

- Similar mechanism has been proposed by Jansen, De Jong, Wegener ECJ 2005 [JDW05]:
  - Scheme C:
    - If (iteration not successful) → double λ
       If (iteration successful) → replace λ by λ/s where s is the nbr of better offspring
  - Jansen, De Jong, Wegener showed that this principle works well in practice, but did not analyze it theoretically

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

48





# Simple Success-Based Rules: Example 3

- The following example requires a bit of time
- I decided to invest this time because
  - I think that this algorithm is worth it
  - this is an example where we can formally prove that the simple success-based rule is better than any static parameter choice
  - I want to discuss with you how we came up with the main ideas
  - there are quite a few open questions, interesting for both empiricallyand theory-oriented researchers
- References for this part:
  - [DDE13] (GECCO 2013) and [DDE15] (TCS 2015, journal version of [DDE13]) suggested the (1+(λ, λ)) GA
  - [DD15b] (GECCO'15): tight bound for static parameter setting
  - [DD15a] (GECCO'15): analysis of self-adjusting mechanism

```
    [Doe16] (GECCO'16): lower bound for 3-dimensional parameter space
Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation 51
```





# How to Chose λ in the (1+(λ, λ)) GA? We analyzed the performance of the (1+(λ, λ)) GA on OneMax

- First "quick&dirty" result: for λ = θ(√log n) the expected runtime of the (1+(λ, λ)) GA on OneMax is 0(n√log n) [DDE13]
- This bound has later been slightly improved in [DD15b]: for λ = θ(√log (n) log log(n) / log log log(n)) the expected runtime of the (1+(λ, λ)) GA on OneMax is 0(n√log (n) log log log(n) / log log(n))

54

• No other (static!) combination of  $p, c, \lambda$  can yield a better runtime



















	<b>Example 4: The</b> $(1 + \lambda)$ <b>EA on OneMax</b>	
	the following mechanism:	
	<ul> <li>let p be the current mutation rate</li> </ul>	
	<ul> <li>in each iteration do:</li> </ul>	
	• create $\lambda/2$ offspring with mutation rate $2p$	
	• create $\lambda/2$ offspring with mutation rate $p/2$	
	<ul> <li>update p as follows (capping at 2/n and 1/4, respectively)</li> </ul>	
	<ul> <li>with probability 1/2 set it to the value for which the best offspring has been found</li> </ul>	
	<ul> <li>with probability 1/2, independently of the last iteration, randomly decide whether to replace p by either p/2 or by 2p</li> </ul>	
	• Main result: this simple mechanism achieves the asymptotically optimal $T^{\text{gen}} = \Theta\left(\frac{n}{\log \lambda} + \frac{n \log n}{\log \lambda}\right)$ performance	
	<ul> <li>Note: this is very recent work. It is likely to be better to work with 3 subpopulations, created by mutation rate <i>cp</i>, <i>p</i>, and <i>Cp</i> for some constants <i>c</i> &lt; 1, <i>C</i> &gt; 1 (feel free to be creative ;) )</li> </ul>	
L	Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation	64



# Main Ideas of Learning-Type Updates The main idea for learning-/reward-type adjustment rules is have a set S of possible parameter values according to some rule, test all or some of these values update the likelihood to employ the tested value based on the feedback from the optimization process Picture to have in mind: • *n* experts in each round, you have to chose one of them and you follow his advice you update your confidence in this expert depending on the quality of his forecast Main difficulty: exploitation vs. exploration trade off exploitation: we want to test each parameter value sufficiently often, to make sure that we select the "optimal" one (in particular when the quality of its "advice" changes, which is the typical situation that we face in evolutionary optimization) exploration: we want, of course, to use an optimal parameter value as often as possible Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation 66

# Learning-Type Updates, Remarks

- Frequently found feature: time-discounted methods. That is, a good advice in the past is worth less than a good advice now
  - different update mechanisms and "forgetting rates" have been experimented with, see discussion below
  - note that such mechanisms are in particular useful when the quality of advice (in our setting, this could be the expected fitness gain, the expected decrease in distance to the optimum, or some other quantity) changes over time
- Note: such learning mechanisms are referred to as "operator selection" in [KHE15]. Another keywords to search for is "credit assignment". It may also be worth to look into literature from learning, in particular on multi-armed bandit algorithms (main goal: maximize reward "on the go", i.e., while learning) and on reinforcement learning (possibly have dedicated "learning" iterations, a notion of state is introduced and the hope is to learn for each state which operator maximizes expected progress)
- Again I will have to focus on a few selected works here. Much more work has been done, cf. Section IV.C.4 in [KHE15] for a survey. There is still much room for further creativity and much research is needed to understand which mechanisms are most useful in which situations
   Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

# Example 1: Davis's adaptive operator fitness (1/2)

Davis (ICGA'89) [Dav89] suggests to adapt rates of crossover operators based on rewards

- Several crossover operators are used simultaneously in every iteration, each having its own crossover rate p<sub>c</sub>(operator<sub>i</sub>)
- the strength of an operator is measured by the fitness value d<sub>i</sub> gained over the best so-far individual in the population. These values are updated after every use of operator i
- every *K* iterations, the crossover rates are being updated as follows:  $p_c^{\text{new}}(\text{operator}_i) = 0.85 \ p_c^{\text{old}}(\text{operator}_i) + d_i^{\text{normalized}}$ with  $d_i^{\text{normalized}}$  are normalized  $d_i$  values (summing up to 15) that is, 15% of the probability mass is re-allocated based on the experience from the last *K* iterations

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation











# Example 4: Self-Adjusting RLS on OneMax (1/4)

- An interesting (albeit not so easy to answer problem) is to determine, for a given search point *x*, how many random bits to flip in order to maximize the expected progress towards the target string *z* when *f* = 0M<sub>z</sub>
- It is easy to convince oneself that the optimal number of bits that one should flip is large when OM<sub>z</sub>(x) is small and is getting smaller when we approach the target string z (illustration on the blackboard)
- In [DDY16b] (GECCO'16) we analyzed this dependence and showed that an optimal mutation-based algorithm is the one employing such fitnessbased step sizes, striving at any point in time for maximal drift towards the target string z
- As before, the question is how an algorithm designer should guess such a relationship (e.g., it turns out that the numbers should always be odd. It is not so easy to compute the cutoff-points from which on the optimal set size changes (see next slide), etc.)

74

 In [DDY16a] (PPSN'16) we showed how a learning-type mechanism automatically chooses parameter values that are close to optimal







Expected fitness gain estimation for using a k-bit flip:  

$$v_t[k] \coloneqq \frac{\sum_{s=1}^{t} 1_{r_s=k} (1-\varepsilon)^{t-s} (f(x_s) - f(x_{s-1}))}{\sum_{s=1}^{t} 1_{r_s=k} (1-\varepsilon)^{t-s}}$$

- 1/ɛ: "forgetting rate", determines the decrease of the importance of older information. 1/ɛ is (roughly) the information half-life
- The "velocity" can be computed iteratively in constant time by introducing a new parameter  $w_t[r] := \sum_{s=1}^{t} \mathbf{1}_{r,er} (1-\epsilon)^{t-s}$
- · This mechanism seems to work well also for other problems
  - So far, no other theoretical results available
  - A few experimental results for LeadingOnes and the Minimum Spanning Tree problem exist, see next 2 slides (these results were also presented in [DDY16a])
  - Again, much more work is needed to see how the algorithm performs on other problems and how to set the parameters  $\delta$  and  $\varepsilon$  (see also discussion below)

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

Our algorithm (S2)

Example 4: Self-Adjusting RLS on LeadingOnes LeadingOnes Theoretical runtime of other evolutionary algorithms:

- Classic RLS:  $0.5n^2$
- (1+1) EA with mutation rate  $p = 1/n: 0.8591n^2$
- (1+1) EA with best-possible mutation rate:  $0.6796n^2$  [2]

Experimental results (100 runs) :

♠ n = 10,000; r<sub>max</sub> = 5; δ = 0.1; ε = <sup>1</sup>/<sub>5,000,000</sub>

	Average complexity	Relative standard deviation
RLS	$0.500n^2$	1.74%
Our algorithm	$0.450n^2$	4.36%

LeadingOnes(x)=number of initial 1s, e.g., LO(1110\*\*\*\*)=3

 parameters above required some tuning, bit we did not invest much time for the tuning → it is likely that you can get better results by a more careful investigation

Example 4: Self-Adjusting RLS on OneMax (4/4)

78

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

Example 4: Self-Adjusting RLS on MST		
Minimum Spanning Trees		
<b>Experimental results for two settings of graph (100 runs on each setting) :</b> Setting 1(S1): $ V  = 20$ ; $ E  = 190$ ; $r_{max} = 5$ ; $\delta = 0.1$ ; $\varepsilon = \frac{1}{400}$ Setting 2(S2): $ V  = 50$ ; $ E  = 1225$ ; $r_{max} = 5$ ; $\delta = 0.1$ ; $\varepsilon = \frac{1}{20,000}$		
	Average runtime	Relative standard deviation
RLS (S1)	$9.62 \cdot 10^{3}$	47.34%
(1+1) EA (S1)	$26.22 \cdot 10^3$	45.61%
Our algorithm (S1)	$6.25 \cdot 10^3$	52.01%
RLS (S2)	$5.08 \cdot 10^6$	37.75%
(1+1) EA (S2)	_	_

 $2.70 \cdot 10^{6}$ 

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

36.34%

÷	As said, we did not try hard to optimize the parameters $\delta$ and $\varepsilon$ If you want to experiment with this learning idea, we suggest that you use the following set-up for the first tries:	
	<ul> <li>few different values for the mutation strength (i.e., small r), since the learni effort is proportional to their number (we used r = 5)</li> </ul>	ng
	<ul> <li>learning rate δ: a small constant, e.g., 5% ("price of the learning mechanism")</li> </ul>	
	• $\delta\left(1-\frac{1}{r}\right)$ is the rate of iterations using a non-optimal mutation strength	
	(can still give progress, but smaller than best-possible)	
	• we used $\delta = 0.1$ and this seems to work well	
1	<b>forgetting time </b> $1/\epsilon$ <b>:</b> this parameter is the most difficult one to set. We recommend to set it so that $1/\epsilon$ is a small percentage of the envisaged total runtime, e.g., $1\% \rightarrow$ it takes very roughly that long to change to a new optimal parameter value	
	→ Too large $\varepsilon$ : we quickly forget the outcomes of previous iterations ③ quick adaption to a changed environment	
	☺ risk that a rare exceptional success with a non-ideal <i>r</i> -value has too much influence	
Carc	ola Doerr: Non-Static Parameter Choices in Evolutionary Computation	80

79



Other Control Mechanisms (2/3)

# Other Control Mechanisms (1/3) In addition to the simple multiplicative update rules and the learning-type rules, many other mechanisms have been experimented with. Here are a few keywords and references (Again, more or less random selection of references, much more work can be found in the survey papers. The works below can serve as a starting point for further investigations.) • Krasnogor and Smith [KS00] (GECCO 2000) suggest a control mechanism for the selective pressure of a memetic algorithm. They use *Boltzmann selection* (popular selection mechanism used in Simulated Annealing, probability of 1 to accept better offspring, probability to accept worse offspring depends on the fitness difference of parent and offspring and a "temperature" which decreases over time, making it less and less likely for worse offspring to get accepted) and suggest to • increase selective pressure when fitness diversity in the population is large • decrease it when fitness diversity is low

 main idea: low fitness diversity = converged population, increase probability to escape and to search elsewhere

82

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation



# Controlling population size is the focus of the Genetic Algorithm with Variable Population Size (GAVaPS) by Arabas, Michalewicz, Mulawka (CEC'94) [AMM94] individuals come with their own lifetime at birth their age is set to 0, each iteration increases the age by 1 maximum lifetime depends on the fitness values, the better a new individual is, the longer its lifetime (and, hence, the more offspring are created from this individual) there is hence no fixed population size, but the size depends adaptively on the search history. One of the goals of GAVaPS was to remove the population size as parameter, but the update mechanism itself comes again with its own parameters Adaptive Population GA (<u>APGA</u>) by Bäck, Eiben, van der Vaart (PPSN 2000) [BEvdV00]: similar to GAVaPS, but age of best individual is not increased, thus allowing it a longer life · lifetime depends on individual's fitness and current-best as well as average fitness of the individuals in the population Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation



# "Parameter-Less" GA

- Parameter-less Genetic Algorithm (PLGA) by Harik and Lobo (GECCO 1999) [HL99] and follow-up works
  - a number of populations of different sizes evolve simultaneously
  - the smaller the population size, the more function evaluations it gets
  - a populations becomes extinct when it converges
  - Hope was to remove population size as a parameter, but note that the mechanism itself introduces new parameters, so the term "parameterless" may be deceptive

86

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation



# Controlling Multiple Parameters Most EAs have several parameters Intuitively, there is no reason to not control more than one or even all of them Several works on controlling more than 1 parameter exists, but we won't have the time to discuss them today Check the mentioned surveys for references on ideas that have been experimented with so far







	Advantages of non-static parameter choices (aka parameter control);	
-	© we gain flexibility and the possibility to adjust the parameter values to the current state of the search	
	$\ensuremath{}$ If we have no idea how to set the parameter, we let the algorithm discover itself	
•	Possible disadvantages:	
	<ul> <li>how to determine which update scheme to use? → designing parameter control mechanisms can, in principle, be an even more complex task than parameter tuning (suggestion: use the "mushroom rule": have a set of 2 or 3 different mechanisms that you declare your favorite ones. Do not try to know all possible mechanisms but rather concentrate on the most promising ones, e.g., one multiplicative update rule, one learning-based rule)</li> </ul>	
	<ul> <li>update mechanisms often come with their own parameters (remember: hope is that the algorithm is much less sensitive to these)</li> </ul>	
	<ul> <li>possibly more difficult to understand how the update mechanism influences the overall performance (measured, e.g., by the distribution of the optimization time)</li> </ul>	;
~	rela Desar Nea Statia Desaratas Chaises is Fuskatian au Computation	





# References

- [ACBF02] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. Machine Learning 47:235-256 2002
- [AMM94] Jaroslaw Arabas, Zbigniew Michalewicz, and Jan J. Mulawka, GAVaPS A genetic algorithm with varying population size In Proc. of International Conference on Evolutionary Computation (ICEC'94), pages 73-78. IEEE, 1994
- [Aug09] Anne Auger. Benchmarking the (1+1) evolution strategy with one-fifth success rule on the BBOB-2009 function testbed. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'09), (Companion), pages 2447-2452. ACM, 2009.
- [Bäc92] Thomas Bäck. The interaction of mutation rate, selection, and self-adaptation within a genetic algorithm. In Proc. of Parallel Problem Solving from Nature (PPSN'92), pages 87-96, Elsevier, 1992.
- [BD17] Maxim Buzdalov and Benjamin Doerr. Runtime analysis of the  $(1 + (\lambda, \lambda))$  Genetic Algorithm on random satisfiable 3-CNF formulas. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'17), page tbd. ACM, 2017.
- [BDN10] Süntie Böttcher, Benjamin Doerr, and Frank Neumann. Optimal fixed and adaptive mutation rates for the Leading Ones problem. In Proc. of Parallel Problem Solving from Nature (PPSN'10), volume 6238 of Lecture Notes in Computer Science pages 1-10. Springer, 2010.
- [BeS00] Helio J.C. Barbosa and Asla Medeiros e Sá. On adaptive operator probabilities in real coded genetic algorithms. In Proc. of Conference of the Chilean Computer Science Society, 2000.
- [BEvdV00] Thomas Bäck, A. E. Eiben, and Nikolai A. L. van der Vaart. An empirical study on gas "without parameters". In Proc of Parallel Problem Solving from Nature (PPSN'00), volume 1917 of Lecture Notes in Computer Science, pages 315-324. Springer, 2000.
- [BLS14] Golnaz Badkobeh, Per Kristian Lehre, and Dirk Sudholt. Unbiased black-box complexity of parallel search. In Proc of Parallel Problem Solving from Nature (PPSN'14), volume 8672 of Lecture Notes in Computer Science, pages 892-901. Springer, 2014.
- [CFSS08] Luís Da Costa, Álvaro Fialho, Marc Schoenauer, and Michèle Sebag. Adaptive operator selection with dynamic multi-armed bandits. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'08), pages 913-920. ACM, 2008.
- Carlos J. Costa, R. Tavares, and A. Rosa. An experimental study on dynamic random variation of population size. In Proc [CTR99] of Systems, Man, and Cybernetics (SMC'99), pages 607-612. IEEE, 1999.

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

### References Lawrence Davis. Adapting operator probabilities in genetic algorithms. In Proc. of International Conference on Genetic [Dav89] Algorithms (ICGA '89), pages 61-69. Morgan Kaufmann, 1989. Benjamin Doerr and Carola Doerr. Optimal parameter choices through self-adjustment: Applying the 1/5-th rule in discrete [DD15a] settings. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'15), pages 1335–1342. ACM, 2015. [DD15b] Benjamin Doerr and Carola Doerr. A tight runtime analysis of the $(1+(\lambda, \lambda))$ genetic algorithm on OneMax. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'15), pages 1423-1430, ACM, 2015 Benjamin Doerr and Carola Doerr. Optimal static and self-adjusting parameter choices for the $(1 + (\lambda, \lambda))$ genetic algorithm [DD17] Algorithmica, 2017. Accepted subject to a minor revision. Available upon request [DDE13] Benjamin Doerr, Carola Doerr, and Franziska Ebel. Lessons from the black-box: Fast crossover-based genetic algorithms. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'13), pages 781–788, ACM, 2013. [DDE15] Benjamin Doerr, Carola Doerr, and Franziska Ebel. From black-box complexity to designing new genetic algorithms. Theoretical Computer Science, 567:87–104, 2015. [DDK16a] Benjamin Doerr, Carola Doerr, and Timo Kötzing. Provably optimal self-adjusting step sizes for multi-valued decision variables. In Proc. of Parallel Problem Solving from Nature (PPSN'16), volume 9921 of Lecture Notes in Computer Science, pages 782-791. Springer, 2016. [DDK16b] Benjamin Doerr, Carola Doerr, and Timo Kötzing. The right mutation strength for multi-valued decision variables. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'16), pages 1115-1122. ACM, 2016. Full version available at http://arxiv.org/abs/1604.03277. [DDY16a] Benjamin Doerr, Carola Doerr, and Jing Yang. k-bit mutation with self-adjusting k outperforms standard bit mutation. In Proc. of Parallel Problem Solving from Nature (PPSN'16), volume 9921 of Lecture Notes in Computer Science, pages 824-834, Springer, 2016.

- [DDY16b] Benjamin Doerr, Carola Doerr, and Jing Yang. Optimal parameter choices via precise black-box analysis. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'16), pages 1123-1130. ACM, 2016.
- [Dev72] Luc Devroye. The compound random search. Ph.D. dissertation, Purdue Univ., West Lafayette, IN, 1972.
- [DGWY17] Benjamin Doerr, Christian Gießen, Carsten Witt, and Jing Yang. The  $(1 + \lambda)$  evolutionary algorithm with self-adjusting mutation rate. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'17), page tbd. ACM, 2017. 98

Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

# References

- Kenneth Alan De Jong, An Analusis of the Behavior of a Class of Genetic Adaptive Sustems, PhD thesis, University of [DJ75] Michigan, Ann Arbor, MI, USA, 1975.
- [DJS<sup>+</sup>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.
- [DL16] Duc-Cuong Dang and Per Kristian Lehre. Self-adaptation of mutation rates in non-elitist populations. In Proc. of Parallel Problem Solving from Nature (PPSN'16), volume 9921 of Lecture Notes in Computer Science, pages 803-813. Springer, 2016.
- Benjamin Doerr. Optimal parameter settings for the  $(1 + (\lambda, \lambda))$  genetic algorithm. In Proc. of Genetic and Evolutionary [Doe16] Computation Conference (GECCO'16), pages 1107-1114. ACM, 2016.
- [EHM99] Agoston Endre Eiben, Robert Hinterding, and Zbigniew Michalewicz. Parameter control in evolutionary algorithms. IEEE Transactions on Evolutionary Computation, 3:124-141, 1999.
- [EMSS07] A. E. Eiben, Zbigniew Michalewicz, Marc Schoenauer, and James E. Smith. Parameter control in evolutionary algorithms. In Parameter Setting in Evolutionary Algorithms, volume 54 of Studies in Computational Intelligence, pages 19-46. Springer, 2007
- [EMV04] A. E. Eiben, Elena Marchiori, and V. A. Valkó. Evolutionary algorithms with on-the-fly population size adjustment. In Proc. of Parallel Problem Solving from Nature (PPSN'04), volume 3242 of Lecture Notes in Computer Science, pages 41-50. Springer, 2004.
- [ES11] A. E. Eiben and Selmar K. Smit. Parameter tuning for configuring and analyzing evolutionary algorithms. Swarm and Evolutionary Computation, 1:19-31, 2011.
- [FCSS08] Álvaro Fialho, Luís Da Costa, Marc Schoenauer, and Michèle Sebag. Extreme value based adaptive operator selection. In Proc. of Parallel Problem Solving from Nature (PPSN'08), volume 5199 of Lecture Notes in Computer Science, pages 175-184. Springer, 2008.
- [FCSS10] Álvaro Fialho, Luís Da Costa, Marc Schoenauer, and Michèle Sebag. Analyzing bandit-based adaptive operator selection mechanisms. Ann. Math. Artif. Intell., 60:25-64, 2010.
- Brian W. Goldman and William F. Punch. Parameter-less population pyramid. In Proc. of Genetic and Evolutionary [GP14] Computation Conference (GECCO'14), pages 785-792. ACM, 2014. 99

## Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

# References

- [GP15] Brian W. Goldman and William F. Punch. Fast and efficient black box optimization using the parameter-less population pyramid. Evolutionary Computation, 23:451-479, 2015.
- [Gre86] John J. Grefenstette. Optimization of control parameters for genetic algorithms. IEEE Trans. Systems, Man, and Cybernetics, 16:122-128, 1986.
- [GS16] Brian W. Goldman and Dirk Sudholt. Runtime analysis for the parameter-less population pyramid. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'160, pages 669-676. ACM, 2016.
- Christian Gießen and Carsten Witt. Population size vs. mutation strength for the  $(1 + \lambda)$  EA on OneMax. In Proc. of [GW15] Genetic and Evolutionary Computation Conference (GECCO'15), pages 1439-1446. ACM, 2015.
- Christian Gießen and Carsten Witt Optimal mutation rates for the  $(1+\lambda)$  EA on OpeMax. In Proc. of Genetic and [GW16] Evolutionary Computation Conference (GECCO'16), pages 1147-1154, ACM, 2016.
- [GW17] Christian Gießen and Carsten Witt. The interplay of population size and mutation probability in the  $(1+\lambda)$  EA on OneMax. Algorithmica, 78(2):587-609, 2017
- [HHB10] Ting Hu, Simon Harding, and Wolfgang Banzhaf. Variable population size and evolution acceleration: a case study with a parallel evolutionary algorithm. Genetic Programming and Evolvable Machines, 11:205-225, 2010.
- Georges R. Harik and Fernando G. Lobo. A parameter-less genetic algorithm. In Proc. of Genetic and Evolutionary [HL99] Computation Conference (GECCO'99), pages 258-265. ACM, 1999.
- [HM90] Jürgen Hesser and Reinhard Männer. Towards an optimal mutation probability for genetic algorithms. In Proc. of Parallel Problem Solving from Nature (PPSN'90), volume 496 of Lecture Notes in Computer Science, pages 23-32. Springer, 1990.
- [JDW05] Thomas Jansen, Kenneth A. De Jong, and Ingo Wegener. On the choice of the offspring population size in evolutionary algorithms, Evolutionary Computation, 13:413-440, 2005.
- Bryant A. Julstrom What have you done for me lately? adapting operator probabilities in a steady-state genetic algorithm [Jn195] In Proc. of International Conference on Genetic Algorithms (ICGA'95), pages 81-87, Morgan Kaufmann, 1995.
- [JW06] Thomas Jansen and Ingo Wegener. On the analysis of a dynamic evolutionary algorithm. J. Discrete Algorithms, 4:181-199,

100

## Carola Doerr: Non-Static Parameter Choices in Evolutionary Computation

	References
[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.
[KK06]	V. K. Koumousis and C. P. Katsaras. A saw-tooth genetic algorithm combining the effects of variable population size and reinitialization to enhance performance. <i>IEEE Transactions on Evolutionary Computation</i> , 10:19–28, 2006.
[KMH+04]	Stefan Kern, Sibylle D. Müller, Nikolaus Hansen, Dirk Büche, Jiri Ocenasek, and Petros Koumoutsakos. Learning probability distributions in continuous evolutionary algorithms - a comparative review. <i>Natural Computing</i> , 3:77–112, 2004.
[KS00]	Natalio Krasnogor and Jim Smith. A memetic algorithm with self-adaptive local search: TSP as a case study. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'00), pages 987–994. Morgan Kaufinann, 2000.
[LS11]	Jörg Lässig and Dirk Sudholt. Adaptive population models for offspring populations and parallel evolutionary algorithms. In Proc. of Foundations of Genetic Algorithms (FOGA'11), pages 181–192. ACM, 2011.
[LS16]	$\label{eq:loss_loss} \ensuremath{Johannes}\xspace \ensuremath{Legger}\xspace \ensuremath{Ingelika}\xspace \ensuremath{Steger}\xspace \ensuremath{Steger}\xspace \ensuremath{Steger}\xspace \ensuremath{Ingelika}\xspace \ensuremath{Steger}\xspace \ensuremath{Steger}\xs$
[Müh92]	Heinz Mühlenbein. How genetic algorithms really work: Mutation and hillclimbing. In Proc. of Parallel Problem Solving from Nature (PPSN'92), pages 15–26. Elsevier, 1992.
[Rec73]	Ingo Rechenberg. Evolutionsstrategie. Friedrich Fromman Verlag (Günther Holzboog KG), Stuttgart, 1973.
[Smi12]	Selmar K. Smit. Parameter Tuning and Scientific Testing in Evolutionary Algorithms. PhD thesis, VU University of Amsterdam, 2012. PhD thesis.
[SS68]	Michael A. Schumer and Kenneth Steiglitz. Adaptive step size random search. IEEE Transactions on Automatic Control, 13:270–276, 1968.
[Thi05]	Dirk Thierens. An adaptive pursuit strategy for allocating operator probabilities. In Proc. of Genetic and Evolutionary Computation Conference (GECCO'05), pages 1539–1546. ACM, 2005.
[TR98]	Andrew Tuson and Peter Ross. Adapting operator settings in genetic algorithms. <i>Evolutionary Computation</i> , 6:161–184, 1998.
[Wit13]	Carsten Witt. Tight bounds on the optimization time of a randomized search heuristic on linear functions. Combinatorics, Probability & Computing, 22:294–318, 2013.
Carola	Doerr: Non-Static Parameter Choices in Evolutionary Computation 101

Г