

## A STUDY OF PARAMETER RELEVANCE IN EVOLUTIONARY ALGORITHMS

### Abstract

We present an empirical study on the impact of different design choices on the performance of an evolutionary algorithm (EA). Four EA components are considered—parent selection, survivor selection, recombination and mutation—and for each component we study the impact of choosing the right operator and of tuning its free parameter(s). We tune 174 different combinations of EA operators to 4 different classes of fitness landscapes and measure the cost of tuning. We find that components differ greatly in importance. Typically the choice of operator for parent selection has the greatest impact, and mutation needs the most tuning. Regarding individual EAs however, the impact of design choices for one component depends on the choices for other components, as well as on the available amount of resources for tuning.

### 3.1 Introduction

Evolutionary Algorithms (EA) form a class of search methods that work by incrementally improving the quality of a set of candidate solutions by variation and selection (Eiben and Smith, 2003). The most important *components* of EAs are thus recombination and mutation (umbrella term: variation), parent selection, and survivor selection. To obtain a working EA, each component needs to be instantiated by a specific *operator*, e.g., the one-point crossover operator for the recombination component. Furthermore, an EA has *parameters* that need to be instantiated by a specific *parameter value*, e.g., 0.5

---

This chapter is an extension of Nannen et al. (2008).

for the crossover rate. In this paper we maintain the distinction between components and parameters and say that the instantiation of EA components by concrete operators specifies a particular EA, e.g., uniform crossover, bit-flip mutation, random parent selection and  $k$ -tournament survivor selection. Further details regarding the parameters do not lead to a different EA, only to variants of the one defined by the operators.<sup>1</sup> A complete EA design includes the definition of an EA (operators for its components) and the specification of a particular variant of it (values for its parameters).

Setting EA parameters is commonly divided into two cases, parameter tuning and parameter control (Eiben et al., 1999, 2007). In case of parameter control the parameter values are changing during an EA run. This requires initial parameter values and suitable control strategies, which in turn can be deterministic, adaptive or self-adaptive. The problem of parameter tuning is hard because for any given application there is a large number of options, but only little knowledge about the effect of EA parameters on EA performance. EA users mostly rely on conventions (mutation rate should be low), ad hoc choices (why not use uniform crossover), and experimental comparisons on a limited scale (testing combinations of three different crossover rates and three different mutation rates). Here we address the problem of parameter tuning. Our main research questions are:

1. How does the choice of operator for each component contribute to EA performance? To this end we compare the absolute performance achieved with different combinations of operators.
2. The parameters of which EA component need the most tuning? For this question we measure the amount of information needed to tune the free parameter(s) of each operator (e.g., crossover rate or tournament size).

For a systematic exploration of the space of EA configurations we use exhaustive search for the combination of operators and Relevance Estimation and Value Calibration (REVAC) to tune the free (numeric) parameters. REVAC is an Estimation of Distribution Algorithm (Mühlenbein and Höns, 2005) that tunes an EA by optimizing marginal probability distributions over the free parameters, see Section 2. Starting from a set of uniform distributions and an initial drawing of 100 vectors of random parameter values, REVAC iteratively generates new marginal distributions of increasing expected EA performance by *drawing* a new vector of parameter values from the current distributions, *evaluating* the vector by measuring the performance of the EA with these values, *updating* all marginal distributions based on this evaluation, and *smoothing* the updated distributions. Smoothing is a unique feature of REVAC that forces all marginal distributions to approach the maximum Shannon entropy distribution for a given EA performance. This maximized Shannon entropy is independent from the computational cost of any particular tuning method and can be used as a general estimator of the minimum amount of information required to reach a certain level of EA performance. Hence, it can be regarded as a general indicator of how difficult it is to tune a certain EA parameter, and how relevant it is to overall EA performance.

---

<sup>1</sup>Alternatively, components & operators could also be called symbolic parameters & values, and we could say these values only define different EA variants.

Related work includes the general discussion of EA design (Czarn et al., 2004) and parameter setting (Lobo et al., 2007), in particular within parameter tuning as defined by Eiben et al. (1999), Birattari (2004), and Eiben et al. (2007). Throughout the relevant literature we find that the cost of tuning parameters is largely ignored. Notable exceptions are the theoretical considerations of Oliver et al. (1987) and Goldberg (1989), as well as the systematic parameter sweeps of Jong (1975), Schaffer et al. (1989), and Samples et al. (2005) and the statistical analysis of parameters by François and Lavergne (2001). In the general field of experimental design, a paradigm shift that emphasizes a low cost of tuning over the performance of optimal parameter values was due to Taguchi and Wu (1980). In our field, Freisleben and Hartfelder (1993) propose a meta-GA approach in which both EA components and EA parameters are tuned and shows the importance of the right choice for the GA operators. Samples et al. (2005) show how parameter sweeps can be used for robustness and correlation analysis.

## 3.2 Experimental setup

For a clear discussion we distinguish three different layers in the analysis of an EA: the problem/application (here: fitness landscapes created by a generator), the problem solver (here: an EA), and the method for tuning the problem solver (here: REVAC). For an unbiased study we use independent software implementations for each layer and combine them through simple interfaces. For the problem layer we use a generator of real-valued fitness landscapes that are formed by the max-set of Gaussian curves in high dimensional Cartesian spaces (Gallagher and Yuan, 2006). Where a Gaussian mixture model takes the average of several Gaussians, a max-set takes their enveloping maximum, giving full control over the location and height of all maxima. For the implementation we followed Rudolph (1992) on rotated high dimensional Gaussians, and used 10 dimensions, 100 Gaussians, and the same distributions over height, location, and rotation of these Gaussians as specified in the exemplary problem sets 1–4 of Gallagher and Yuan (2006). These sets offer an increasing amount of exploitable structure to the EA. Set 1 has the least structure, with peaks of different height scattered at random, while set 4 is the most structured, with peaks that get higher the closer they get to the origin. For each set, different landscapes are created by passing a different random seed to the generator. Initialization of all EA populations is uniform random in the domain of the fitness landscapes. The optimal fitness value is 1 on each problem instance and the condition for successful termination is defined as “fitness > 0.9999 or 10,000 fitness evaluations”.

For the EAs we use the Evolutionary Computation toolkit in Java (ECJ) (Luke), which allows the specification of a fully implemented EA through a simple parameter file, including the choice of operator for each component and the values for the free parameters. The ECJ offers several operators for each EA component, cf. Table 3.1. For any given EA, the population size parameter is always present. Most operators have zero or one free parameter. One operator has 2 free parameters—Gaussian( $\sigma$ ,  $p$ ) with parameters  $\sigma$  for step size and  $p$  for mutation probability, which takes the value 1 in case of Gaussian( $\sigma$ , 1). Due to technical details of the ECJ, not all combinations of operators are possible. For example, neither fitness proportional nor best selection as parent selection can be combined with tournament selection as survivor selection. In total we have

Table 3.1: EA components, operators, and parameters used in this study

Component	Operator	Parameter(s)
		population size $\mu$
parent selection	tournament	parent tournament size
	best selection	number $n$ of best
	fitness proportional	-
	random	-
survivor selection	generational	-
	tournament	survivor tournament size
	$(\mu, \lambda)$	$\lambda$
	$(\mu + \lambda)$	$\lambda$
recombination	random	-
	none	-
	one-point	crossover probability
	two-point	crossover probability
mutation	uniform	crossover probability
	reset (random uniform)	mutation probability
	Gaussian( $\sigma, 1$ )	step size
	Gaussian( $\sigma, p$ )	step size, mutation probability

*Notes.* We follow the naming convention of the ECJ. Arguably,  $(\mu, \lambda)$  and  $(\mu + \lambda)$  define both parent *and* survivor selection. Here we classify them under survivor selection because that is what the parameter  $\lambda$  influences.

174 combinations of operators, of which 6 with 2, 33 with 3, 65 with 4, 55 with 5, and 15 with 6 free parameters.

The performance of an EA with a given set of parameter values is measured in three different ways: SR (Success Rate, percentage of runs with fitness  $> 0.9999$ ), MBF (Mean Best Fitness of all runs), and AES (Average number of Evaluations to Solution of successful runs; undefined when SR = 0). Each EA is tuned 5 times on each of the 4 problem sets. During each tuning session on a given set REVAC generates 1,000 different vectors of parameter values. Each vector of values is written to the ECJ configuration file, together with the specification of the operators and the problem generator. The resulting EA is evaluated on 10 different instances of the problem set, generated by different random seeds.

For each REVAC tuning session and each EA, the best performance after  $n$  evaluations is the best performance measured after evaluating  $n$  vectors of parameter values. The average best performance after  $n$  evaluations is averaged over multiple tuning sessions on the same EA. We define *near best performance* as the average best performance after 1,000 evaluations minus 5%<sup>2</sup>. If  $n$  is the lowest number of vectors for which the average best performance after  $n$  evaluations exceeds this value, then we say that REVAC

<sup>2</sup>In case of MBF this is calculated after subtracting a performance level of .5, a level that is reached by any reasonably sized population upon random uniform initialization. The maximum possible near best MBF is therefore .975

needs  $n$  evaluations to tune the EA to near best performance. Section 3.3 uses this to study the impact of choosing an operator for each component.

In Section 3.4 we analyze the cost and benefits of tuning per EA component. REVAC continuously maximizes the Shannon entropy of the marginal distributions that it optimizes during a tuning session. This maximized Shannon entropy provides a generic information-theoretic measure of the minimum amount of information needed per parameter to reach a given performance level. The differential Shannon entropy  $h$  of a probability density function  $\mathcal{D}$  over the continuous interval  $[a, b]$  is commonly defined as

$$h(\mathcal{D}_{[a,b]}) = - \int_a^b \mathcal{D}(x) \log_2 \mathcal{D}(x) dx. \quad (3.1)$$

The sharper the peaks of a probability density function, the lower its Shannon entropy. In order to compare the entropy of distributions that are defined over different parameter intervals in a meaningful way, we normalize all parameter intervals to the interval  $[0, 1]$  before calculating the Shannon entropy. In this way the initial uniform distribution has a Shannon entropy of zero, and any other distribution has a negative Shannon entropy  $h(\mathcal{D}_{[0,1]}) < 0$ .

### 3.3 How does the choice of operator per component contribute to performance?

Figure 3.1–3.8 on page 56–61 show the average near best fitness in AES and MBF and the tuning cost in number of evaluations of REVAC needed to reach this fitness. Results in AES are based on those EAs with  $SR > 0$  for which the AES could be calculated. Depending on the problem set, 60–70 EAs could be tuned to terminate with success. Results on the MBF are based on all 174 EAs. Each figure contains four scatter plots that show the performance in AES or MBF after tuning, and the cost of tuning, averaged over 5 tuning sessions per EA. The  $y$ -axes show the near best performance in AES. The  $x$ -axes show the number of REVAC evaluations needed to tune the EA to this performance. Each of the four plots of a figure shows the same EAs but labels them according to the operator choice for a different component. To read the full specification of an EA, one needs to look at the same location in all four plots. Under each AES figure a table (i.e., Table 3.2–3.5) shows the average performance on the same set that was achieved with each operator of each component. The table reports the average success rate, the near best performance, and the cost of finding it, for AES and for MBF respectively. Each value is averaged over all those EAs that use this operator and, in case of AES, terminated with success.

The choice of operator for the parent selection component has the strongest effect on EA performance. The 16 EAs that are clustered together in the lower left of each plot of Figure 3.1–3.4 display the best performance and the lowest number of evaluations needed to reach this performance. These EAs all use tournament selection for parent selection, either tournament selection or random selection for survivor selection, any recombination operator, and either Gaussian( $\sigma, p$ ) or Gaussian( $\sigma, 1$ ) for mutation. On the other hand, those EAs that have the never terminate with success and generally reach only a minimum MBF (see Figure 3.5–3.8) share one common feature, namely a lack of

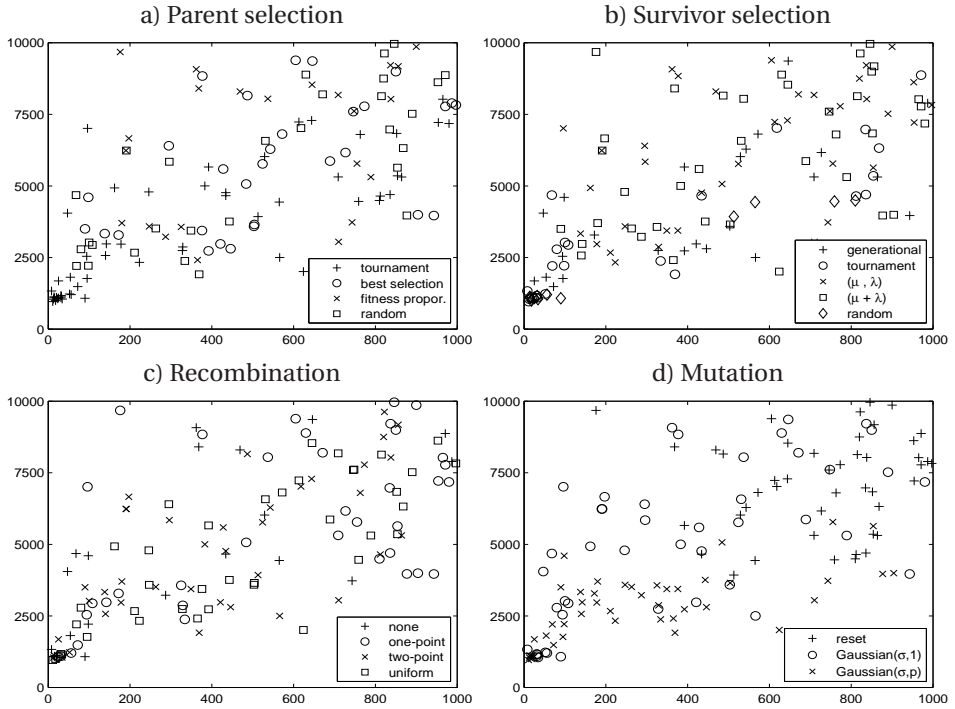


Figure 3.1: Near best performance in AES on set 1 against cost of tuning

Table 3.2: Average near best performance on set 1, per operator

Component	Operator	Success rate	AES	Cost	MBF	Cost
parent selection	tournament	1	3484	344	0.87	281
	best select.	1	5782	556	0.89	279
	fitness pro.	0.64	6582	532	0.87	219
	random	0.54	5611	517	0.77	388
survivor selection	generational	0.5	4224	426	0.8	367
	tournament	1	3320	318	0.87	324
	$(\mu, \lambda)$	0.97	6259	546	0.89	209
	$(\mu + \lambda)$	0.97	6072	550	0.9	210
	random	0.5	2175	244	0.74	444
recombination	none	0.67	4798	343	0.82	328
	one-point	0.79	5549	548	0.84	289
	two-point	0.81	4727	412	0.85	273
	uniform	0.81	4732	479	0.85	324
mutation	reset	0.78	7168	729	0.83	309
	Gauss. $(\sigma, 1)$	0.78	5001	368	0.85	306
	Gauss. $(\sigma, p)$	0.79	2785	285	0.85	287

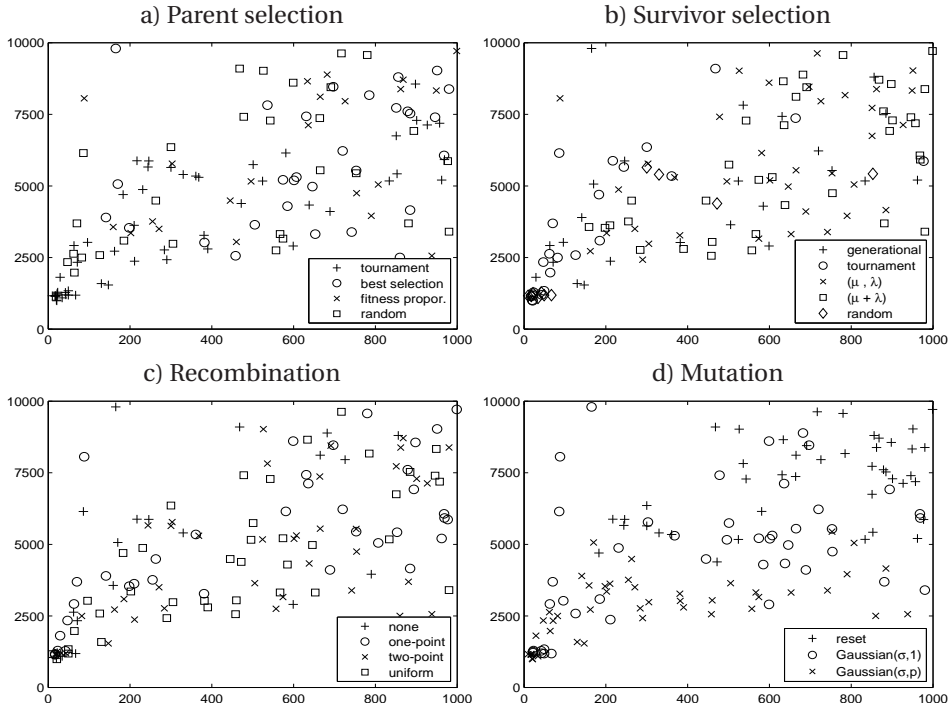


Figure 3.2: Near best performance in AES on set 2 against cost of tuning

Table 3.3: Average near best performance on set 2, per operator

Component	Operator	Success rate	AES	Cost	MBF	Cost
parent selection	tournament	1	3559	325	0.91	265
	best select.	1	5870	654	0.92	157
	fitness pro.	0.61	6042	590	0.91	162
	random	0.54	5357	479	0.8	324
survivor selection	generational	0.5	4708	427	0.83	248
	tournament	1	3524	181	0.92	255
	$(\mu, \lambda)$	0.92	5792	595	0.93	201
	$(\mu + \lambda)$	1	5643	663	0.93	157
	random	0.5	2529	185	0.77	431
recombination	none	0.67	5046	322	0.85	275
	one-point	0.77	5251	523	0.88	245
	two-point	0.81	4705	515	0.89	232
	uniform	0.81	4559	468	0.89	233
mutation	reset	0.76	7259	689	0.87	239
	Gauss. ( $\sigma, 1$ )	0.78	4620	430	0.88	232
	Gauss. ( $\sigma, p$ )	0.79	2810	313	0.88	260

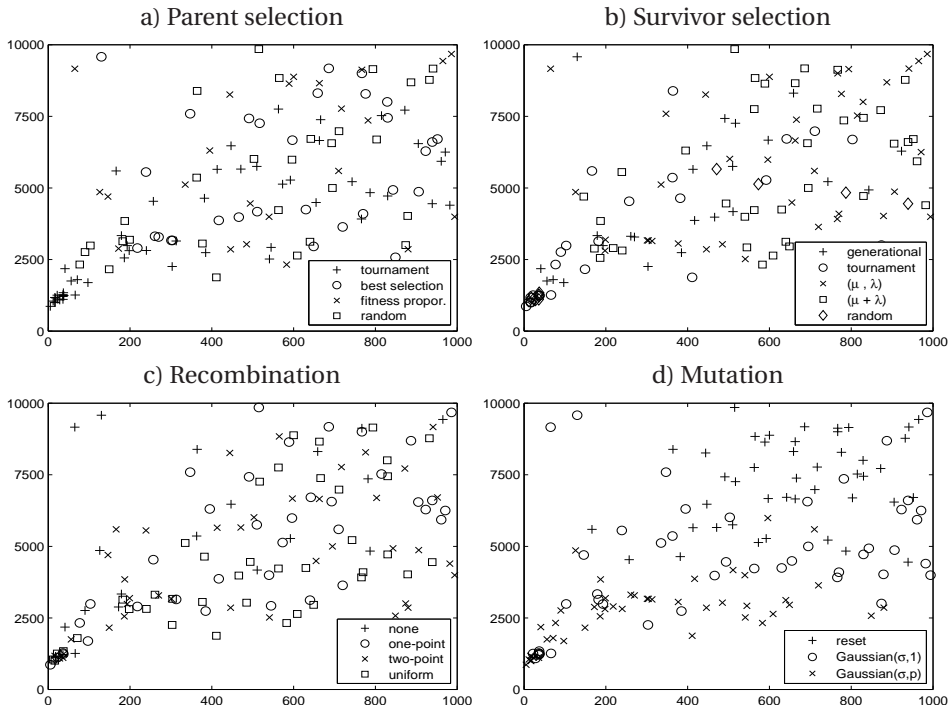


Figure 3.3: Near best performance in AES on set 3 against cost of tuning

Table 3.4: Average near best performance on set 3, per operator

Component	Operator	Success rate	AES	Cost	MBF	Cost
parent selection	tournament	1	3526	365	0.92	253
	best select.	1	5636	593	0.93	203
	fitness pro.	0.64	6085	560	0.91	192
	random	0.54	5329	529	0.8	325
survivor selection	generational	0.5	4622	397	0.84	263
	tournament	1	3395	266	0.93	242
	$(\mu, \lambda)$	0.95	5856	612	0.94	225
	$(\mu + \lambda)$	1	5525	621	0.94	176
	random	0.5	2451	250	0.77	424
recombination	none	0.7	5149	342	0.85	323
	one-point	0.77	5233	522	0.89	245
	two-point	0.81	4590	502	0.89	237
	uniform	0.81	4442	504	0.89	238
mutation	reset	0.76	7283	656	0.87	253
	Gauss. ( $\sigma, 1$ )	0.79	4532	493	0.89	260
	Gauss. ( $\sigma, p$ )	0.79	2719	309	0.89	250



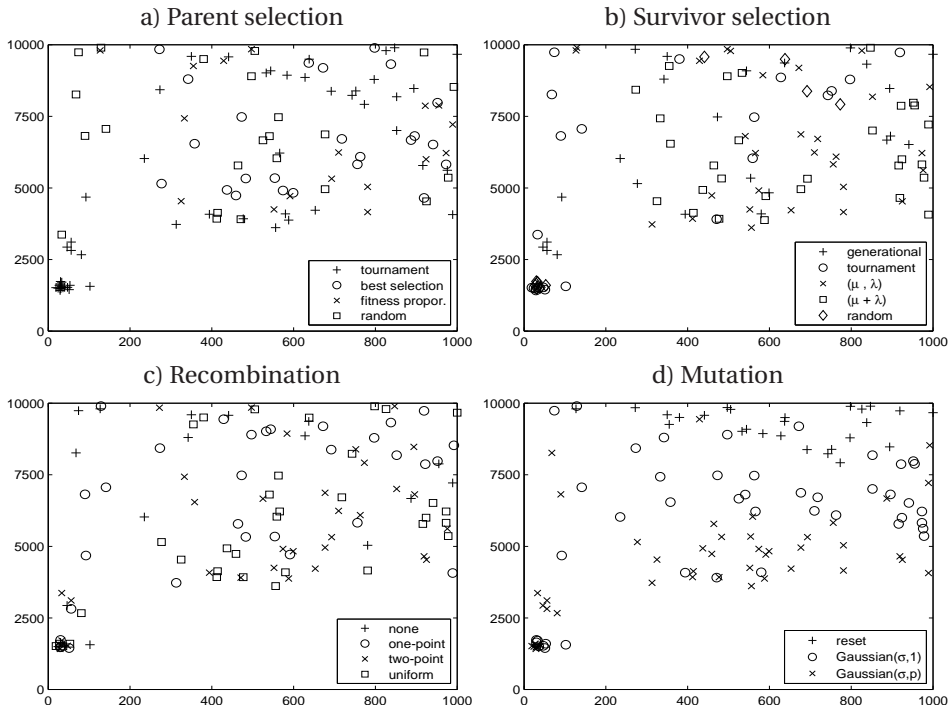


Figure 3.4: Near best performance in AES on set 4 against cost of tuning

Table 3.5: Average near best performance on set 4, per operator

Component	Operator	Success rate	AES	Cost	MBF	Cost
parent selection	tournament	0.96	5078	408	0.89	319
	best select.	0.8	6782	649	0.89	249
	fitness pro.	0.47	6778	643	0.88	266
	random	0.43	6872	480	0.77	398
survivor selection	generational	0.5	6554	493	0.8	353
	tournament	0.92	5370	298	0.9	313
	$(\mu, \lambda)$	0.74	6660	631	0.9	265
	$(\mu + \lambda)$	0.74	6321	663	0.9	271
	random	0.5	4019	235	0.73	434
recombination	none	0.6	6446	375	0.81	390
	one-point	0.65	6535	490	0.85	311
	two-point	0.69	5601	545	0.86	283
	uniform	0.71	5785	554	0.86	325
mutation	reset	0.45	9250	630	0.84	323
	Gauss. ( $\sigma, 1$ )	0.76	5890	514	0.86	301
	Gauss. ( $\sigma, p$ )	0.79	4358	430	0.86	338




Figure 3.5: Near best performance in AES on set 1 against cost of tuning




Figure 3.6: Near best performance in AES on set 2 against cost of tuning




Figure 3.7: Near best performance in AES on set 3 against cost of tuning

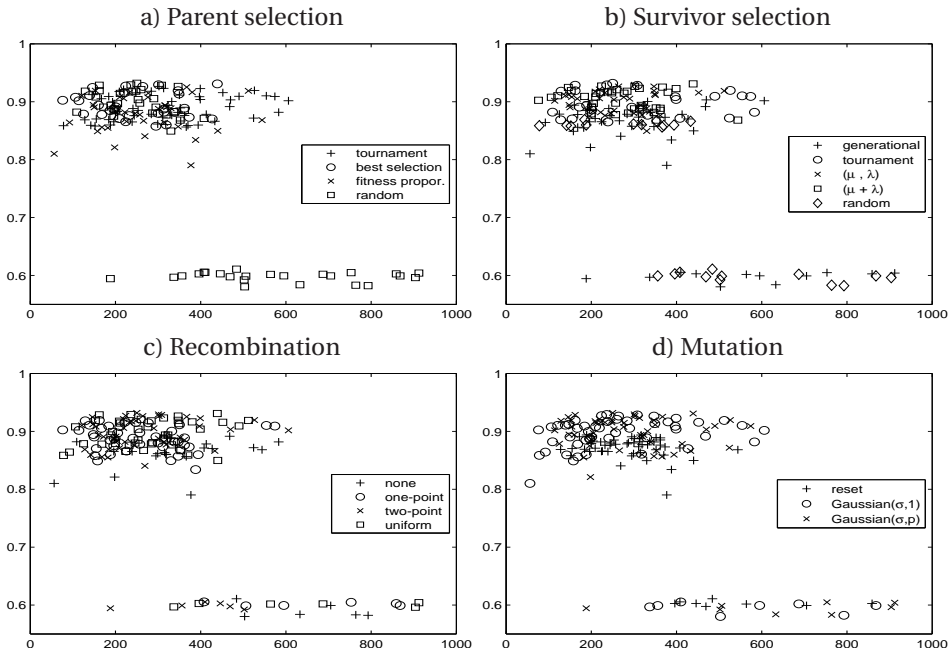


Figure 3.8: Near best performance in AES on set 4 against cost of tuning




Figure 3.9: Impact of recombination operators on AES and cost of tuning

selection pressure. In particular, EAs with random or fitness proportional selection for parent selection do not terminate with success unless combined with strong survivor selection pressure.

Of the two variation components, the choice of mutation operator has the stronger effect on EA performance, as can be seen from the differences in Table 3.2–3.5. On all our problem sets reset mutation is the worst mutation operator, and non-standard Gaussian( $\sigma, p$ ) mutation is superior to Gaussian( $\sigma, 1$ ) both in terms of performance and in terms of cost of tuning. The latter may come as a surprise, since the additional free parameter for mutation probability increases the parameter search space. We conclude that the tuning cost of different operators is not additive, and that the tuning cost of an operator can only be evaluated in the context of the overall EA composition.

While choosing the recombination operator has the least effect on EA performance, it demonstrates how the choice of operator can depend on the available resources for tuning. Figure 3.9 enlarges the lower left corner of Figure 3.4c, overlaid by four graphs that show the evolution of the average performance of 4 EAs with tournament selection for both parent and survivor selection, Gaussian( $\sigma, p$ ) mutation, and four different recombination operators. 20 tuning sessions were used for each graph. While the tunable recombination operator eventually outperform no recombination, an EA with no recombination consistently outperforms EAs with tunable recombination after about 30–40 parameter vectors have been evaluated, and it has at least average performance for anything under 100 evaluated parameter vectors. We observed this phenomenon over a wide range of operator choices for the other components and over all 4 problem sets. All in all, for recombination, the choice of operator can clearly depend on the amount of effort that can be invested in tuning.

### 3.4 Which EA component needs the most tuning?

The previous section related the performance of the near best parameter vector to the number of REVAC evaluations needed to find this vector and to achieve this performance. This section takes a rather unconventional approach based on the expected performance when parameter values are drawn from a probability distribution, namely those created by REVAC after 500 evaluations. To calculate the performance gain achieved by tuning, this expected performance is compared to the expected EA performance




Figure 3.10: Correlation between Shannon entropy and standard deviation of tuned parameter values




Figure 3.11: Correlation between performance gain and the sum of the Shannon entropy of all marginal distributions

when parameter values are drawn from the uniform distribution. All results are averaged over 5 REVAC tuning sessions of an EA on each of the 4 problem sets, 20 tuning sessions per EA. In order to extend our analysis to all 174 EAs, we use the Mean Best Fitness that an EA achieves at termination (successful or not), rather than the AES.

Shannon entropy measures the amount of information that a probability distribution provides on its random values. By definition, the lower the Shannon entropy of the maximum entropy distribution that achieves a given expected EA performance, the finer the parameter value has to be tuned in order to achieve that expected performance. This is demonstrated in Figure 3.10. The scatter plot shows the correlation between the Shannon entropy of the marginal distribution over the mutation probability and the standard deviation of the best found parameter values. The  $x$ -axis shows the Shannon entropy as estimated by REVAC. The  $y$ -axis shows the average (over 4 sets) of the standard deviation of the 5 best values found during the 5 REVAC tuning sessions on each set. The correlation coefficient is 0.9 and the  $p$ -value (the probability to observe this or a stronger correlation when the true coefficient is zero) virtually zero. The point here is that if the maximum entropy distribution has a higher Shannon entropy, there is less certainty on the precise parameter value, something that can otherwise be expensive to assess.

Figure 3.11 shows a clear correlation between a gain in expected MBF and the Shannon entropy of the maximum entropy distributions that REVAC has estimated after 500 evaluations. The  $x$ -axis shows the average performance gain in percent. The  $y$ -axis

Table 3.6: Entropy per EA component, averaged over all EAs

Component & population size		Correlation with MBF gain		Shannon entropy		
		correlation	p-value	max	mean	min
1)	population size	-0.4	0	0	-0.9	-1.7
2)	parent selection	0.2	0.1	-0.1	-0.6	-3.3
3)	survivor selection	-0.1	0.3	0	-0.7	-1.4
4)	recombination	-0.4	0	0	-0.2	-1.2
5)	mutation	-0.6	0	-0.1	-1.2	-4.5
Entire EA		-0.8	0	-0.2	-2.9	-5.1




Figure 3.12: Mean Shannon entropy per component &amp; population size

shows the Shannon entropy of the estimated distributions, summed over all tuneable parameters of the EA. Note that no EA lies above the main diagonal, which shows that there is a minimum information cost for every percent point of gain in expected performance, regardless of the EA specifications. Of those EAs that lie significantly below the diagonal, most use tournament selection for both parent and survivor selection. By 500 REVAC evaluations, their MBF had long been maximized. Further tuning only improved their AES, distorting their performance gain to entropy ratio.

Does the strong correlation between total Shannon entropy and the gain in expected performance carry over to individual EA components? The first two numeric columns of Table 3.6 show the correlation coefficient for each component and its p-value, i.e., the probability to observe this or a stronger correlation coefficient if the true coefficient is zero. Only EAs with a tunable operator were considered for the respective component. The correlation is generally weak, in particular for selection. In other words, the question which component needs tuning in order to improve the performance of a particular EA depends much on the EA in question.

With respect to the average Shannon entropy per component, we see that not all components require the same amount of tuning. The right numeric columns in Table 3.6 show the maximum, mean, and minimum Shannon entropy that we observed for each component (and the population size) when instantiated with an operator that needs tuning. The bar diagram of Figure 3.12 allows a visual comparison of this average mean Shannon entropy. Such a skewed distribution of a need for tuning is commonly known as *sparsity of effects*.

Typically, mutation requires the highest amount of tuning, and recombination the least. This rule has many exceptions, as can be concluded from the low correlation co-

efficient. While the relative order of Shannon entropy per component depends much on the EA in question, consistent patterns can be detected for small groups of EAs. Take for example the two EAs with tournament selection for both parent and survivor selection, Gaussian( $\sigma, 1$ ) mutation and either one-point, two-point or uniform crossover. We find that the Shannon entropy for mutation has the unusually high Shannon entropy of around  $-2$ , while the parent selection operator has a low Shannon entropy below  $-3$ . When combining the same selection operators with other recombination or mutation operators, we find that the Shannon entropy for parent selection is back to normal levels, while it is still comparatively high for mutation. Another example is recombination, which only exhibits a low Shannon entropy for uniform crossover in combination with either  $(\mu + \lambda)$ , or  $(\mu, \lambda)$ . Such irregular patterns are consistent over different problem sets and seem to be inherent to specific combinations of EA components.

### 3.5 Conclusions

This chapter introduces a novel approach to EA design that emphasizes the cost of tuning. To understand how this cost depends on the choice of operator per EA component, we combined exhaustive search over operators with REVAC for tuning their parameters. Our experiments revealed a number of notable insights.

Our tests confirmed the common wisdom that the choice of operator for one EA component depends on the choice of operator for the other components. Of all components, the choice of operator for parent selection has the biggest impact on EA performance. Furthermore, EAs differ greatly in the amount of tuning needed to reach a given performance, and this tuning cost depends on the overall setup of the EA, rather than the number of free parameters. With regard to recombination, we found that the best EA setup depends on the time and effort one can permit to tune the EA.

To measure the need for tuning per component we use the Shannon entropy of maximum entropy distributions as estimated by REVAC, which expresses the minimum amount of information that is needed to achieve a given expected EA performance. It is a generic information-theoretic measure that is independent of any particular tuning algorithm. Inspired by theoretical considerations, it was validated by a strong correlation with the standard deviation of best solutions found during multiple tuning sessions. Based on this measure we observed that the need for tuning follows a skewed distribution, and that while total Shannon entropy is strongly correlated with performance gain, the correlation per component is weak. The question which component needs the most tuning depends on the precise composition of an EA and can not be answered on a general level. It needs to be addressed by the operational analysis of individual EAs. We recommend that a scientific discussion of individual operators addresses their effect on the overall tunability of an EA and on the need for tuning per component.

Regarding the scope of our results, an empirical study can only use a limited set of test problems, and strictly speaking our findings are only proven for our test problems. However, we consider it unlikely that the complex picture that has emerged here is an artefact of the test problems. What remains to be studied is whether the way in which the need for tuning per component depends on the choice of operator for other components is different on other complex fitness functions.

## REFERENCES

- Birattari, M., 2004. The problem of tuning metaheuristics as seen from a machine learning perspective. Ph.D. thesis, Université Libre de Bruxelles.
- Czarn, A., MacNish, C., Vijayan, K., Turlach, B., Gupta, R., 2004. Statistical Exploratory Analysis of Genetic Algorithms. *IEEE Trans. Evol. Comput.* 8 (4), 405–421.
- Eiben, A. E., Hinterding, R., Michalewicz, Z., 1999. Parameter control in evolutionary algorithms. *IEEE Trans. Evol. Comput.* 3 (2), 124–141.
- Eiben, A. E., Michalewicz, Z., Schoenauer, M., Smith, J. E., 2007. Parameter control in evolutionary algorithms. In: Lobo, F. G., Lima, C. F., Michalewicz, Z. (Eds.), *Parameter Setting in Evolutionary Algorithms. Studies in Computational Intelligence*. Springer, Berlin / Heidelberg, pp. 19–46.
- Eiben, A. E., Smith, J. E., 2003. *Introduction to Evolutionary Computing*. Springer, Berlin / Heidelberg.
- François, O., Lavergne, C., 2001. Design of evolutionary algorithms—a statistical perspective. *IEEE Trans. Evol. Comput.* 5 (2), 129–148.
- Freisleben, B., Hartfelder, M., 1993. Optimization of genetic algorithms by genetic algorithms. In: Albrecht, R. F., Reeves, C. R., Steele, N. C. (Eds.), *Artificial Neural Networks and Genetic Algorithms*. Springer, Berlin / Heidelberg, pp. 392–399.
- Gallagher, M., Yuan, B., 2006. A general-purpose tunable landscape generator. *IEEE Trans. Evol. Comput.* 10 (5), 590–603.
- Goldberg, D. E., 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- Jong, K. A. D., 1975. An analysis of the behaviour of a class of genetic adaptive systems. Ph.D. thesis, University of Michigan.
- Lobo, F. G., Lima, C. F., Michalewicz, Z. (Eds.), 2007. *Parameter Setting in Evolutionary Algorithms. Studies in Computational Intelligence*. Springer, Berlin / Heidelberg.
- Luke, S., 2007. The ECJ Evolutionary Computation System. Tech. rep., Department of Computer Science George Mason University.



- Mühlenbein, H., Höns, R., 2005. The estimation of distributions and the minimum relative entropy principle. *Evol. Comput.* 13 (1), 1–27.
- Nannen, V., Smit, S. K., Eiben, A. E., 2008. Costs and Benefits of Tuning Parameters of Evolutionary Algorithms. In: *Parallel Probl. Solving Nat., PPSN X*. Springer, Dortmund, Germany, pp. 528–538.
- Oliver, I. M., Smith, D. J., Holland, J. R. C., 1987. A study of permutation crossover operators on the traveling salesman problem. In: *Proc. 2nd Int. Conf. Genet. Algorithms Genet. algorithms appl.* Cambridge, Ma, pp. 224–230.
- Rudolph, G., 1992. On correlated mutations in evolution strategies. In: *Parallel Probl. Solving Nat., PPSN II*. Springer, Brussels, Belgium, pp. 107–116.
- Samples, M. E., Daida, J. M., Byom, M., Pizzimenti, M., 2005. Parameter sweeps for exploring GP parameters. In: *Proc. Genet. Evol. Comput. Conf., GECCO'05*. ACM, Washington, DC, pp. 1791–1792.
- Schaffer, J., Caruana, R., Eshelman, L., Das, R., 1989. A study of control parameters affecting online performance of genetic algorithms for function optimization. In: *Proceedings of the third international conference on genetic algorithms*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 51–60.
- Taguchi, G., Wu, Y., 1980. *Introduction to Off-Line Quality Control*. Central Japan Quality Control Association, Nagoya, Japan.