
Potholes on the Royal Road *

Theodore C. Belding

Center for the Study of Complex Systems

University of Michigan

Ann Arbor, MI 48109-1120 USA

Ted.Belding@umich.edu

<http://www-personal.umich.edu/~streak/>

Abstract

It is still unclear how an evolutionary algorithm (EA) searches a fitness landscape, and on what fitness landscapes a particular EA will do well. The validity of the building-block hypothesis, a major tenet of traditional genetic algorithm theory, remains controversial despite its continued use to justify claims about EAs. This paper outlines a research program to begin to answer some of these open questions, by extending the work done in the royal road project. The short-term goal is to find a simple class of functions which the simple genetic algorithm optimizes better than other optimization methods, such as hillclimbers. A dialectical heuristic for searching for such a class is introduced. As an example of using the heuristic, the simple genetic algorithm is compared with a set of hillclimbers on a simple subset of the hyperplane-defined functions, the pothole functions.

1 BACKGROUND

Evolutionary algorithms (EAs) are computational search methods based on biological evolution. Some common EAs are *genetic algorithms* (GAs) [12, 9, 18], *evolutionary programming* (EP) [6], *evolution strategies* (ESs) [24], *genetic programming* (GP) [17], and *classifier systems* [14]. The study of EAs is called *evolutionary computation* (EC).

EAs are increasingly important in such areas as function optimization, machine learning, and modeling. However, as Mitchell et al. emphasized in the royal road (RR) papers [19, 7, 20], it is still unclear how an EA

searches a fitness landscape, or even what an EA *does*. It is also unknown what types of problem are easy or hard for a particular EA, how various landscape features affect problem difficulty for an EA, or under what circumstances an EA will outperform another search method. Even less work has been done to classify the features of real-world problems that may be relevant to EA performance. Moreover, the selection of EA parameters such as mutation rate or population size is still largely a black art, despite some promising research in this area. This lack of theory makes the selection and configuration of an EA for a given problem difficult.

A major open question in EC is the function and importance of the *crossover* operator, which recombines two individuals. Holland [12] has argued that crossover is central to an EA's efficacy. The theoretical basis for this is the *building-block hypothesis* (BBH) [9], which states that an EA uses crossover to repeatedly combine compact subsolutions with high observed fitness from different individuals, forming more complete subsolutions with even higher observed fitness. Such subsolutions are called *building blocks*. When an EA uses crossover on symbol strings from the set A^ℓ , where A is the set of possible symbols and ℓ is the length of a string, the building blocks are short *schemata* with high observed fitness. Schemata are members of the set $(A \cup \{*\})^\ell$, where $*$ is a wildcard symbol. They are hyperplanes in the search space. Holland's schema theorem [12] implies that short schemata with consistently above-average observed fitness tend to increase exponentially in frequency over several generations. (If operators other than reproduction are neglected, this is true for all partitions of the search space, not just for a partition into schemata. However, applying crossover to symbol strings induces a natural partition of the space into schemata [27]. Furthermore, short schemata are preserved by crossover. This makes schemata particularly relevant when studying EAs that use crossover on strings.) Applying crossover to individuals with high fitness is a plausible heuristic for generating offspring that will also be highly fit. The chance that this heuris-

*To appear in Lee Spector et al., editors, GECCO 2001: Proceedings of the Genetic and Evolutionary Computation Conference. Morgan Kaufmann, San Francisco, 2001.

tic succeeds can be quantified using Price’s covariance and selection theorem [1]. Implicit in the BBH is also the hypothesis that there are many real-world problems amenable to solution by this process.

A common misconception is that a schema has a unique, well-defined fitness, which is the average fitness of all of its possible instances, and that observed fitnesses are estimates of these “actual” values. In general, no such unique schema fitness exists, and the schema theorem makes no such assumption [11]: The observed fitness of a schema is the average fitness of its instances in the current population. This value depends on the distribution of schemata in the current population, which is biased by the EA over time. A uniform distribution is only seen immediately after generating an initial random population, if ever. Hence, there is no justification for arbitrarily defining a schema’s fitness to be the average over a uniform distribution. A schema may be a highly-fit building block in one population but not in another, even under the same fitness function. Grefenstette [10] made essentially the same point when he criticized the “static building-block hypothesis”.

The BBH is often used to explain how EAs work and to justify the importance of crossover. However, there is no theory that specifies in detail the conditions necessary for the BBH to be valid and thus for crossover to be beneficial. While there is empirical evidence in favor of the BBH [13], its validity in general for the SGA and for other EAs using crossover remains controversial [6], and the schema theorem’s relevance to EA theory has been questioned as well [27, 28]. In particular, uniform crossover, which is more likely to break up short building blocks than traditional crossover, is very effective on some problems [26], and it may be that on some problems crossover acts as a macromutation operator, rather than as an operator that recombines building blocks [16]. More generally, it is not clear how to formulate a BBH that is valid for an arbitrary EA operating on an arbitrary representation of solutions [22].

2 RESEARCH PROGRAM

This paper presents a research program to extend the RR papers [19, 7, 20] in testing the validity of the BBH, focusing specifically on the *simple genetic algorithm* (SGA) [9]. The SGA is a GA that uses fitness-proportionate selection, single-point crossover, and point mutation to evolve a single panmictic population of bit strings, with each generation completely replacing the previous one. I focus on the SGA because it is a relatively simple EA with a large theoretical literature and because many EAs descend directly or indirectly from it. A theory developed for the SGA has a relatively good chance of being applicable to other EAs.

As in the RR papers, I use function optimization to compare the SGA with other search algorithms. I do this because it is an increasingly important application for EAs, with relatively clear performance criteria, and because a simple fitness function is easy to design and implement. Also, function optimization can be viewed as search, so theories developed for it may be relevant to other applications of search, for instance artificial intelligence [21] and evolutionary biology [30].

As De Jong [2] pointed out, the SGA is not a function optimizer, per se. But if the BBH is valid, the SGA should work particularly well as an optimizer on functions rich in building blocks that can be recombined to reach the optimum. Hence, to determine the validity of the BBH, it would be useful to know the class of functions on which the SGA outperforms other optimizers. Examining what makes this class of functions particularly easy for the SGA will also help us to predict which functions the SGA will perform well on. One step towards this goal is to find a simple class of functions on which the SGA outperforms hillclimbers. This was the goal of the RR papers and is also the immediate goal here. To meet this goal, the SGA should consistently perform extremely well on the functions and outperform hillclimbers by a wide margin, for a reasonable set of performance criteria. (Note that this is a different question from that addressed by Wolpert and Macready’s [29] no free lunch (NFL) theorem.)

This paper takes a different approach from that of the RR papers, although the ultimate goal remains the same. In those papers, Forrest and Mitchell [7] determined that *hitchhiking* was a major factor limiting SGA performance on the RR functions, causing it to perform worse than the random mutation hillclimber (RMHC). Hitchhiking occurs when detrimental or neutral alleles increase in frequency due to the presence of nearby beneficial alleles on the same chromosome [8]. This can cause beneficial alleles at the same loci as the hitchhiking alleles to die out in the population, preventing the SGA from finding any highly-fit individuals that contain those alleles. (The fact that schemata do not have unique, well defined fitnesses is a necessary precondition for hitchhiking.) After identifying this problem, Mitchell et al. [20] investigated how to make the SGA perform more like an *idealized genetic algorithm* that was unaffected by hitchhiking. They developed the RR function *R4*, which reduced hitchhiking by lowering the fitness jump from one level of the function to the next. In effect, they made the SGA outperform the RMHC by making the functions easier for the SGA. In contrast, I attempt to make them harder for hillclimbers. The RR functions are easy for hillclimbers like the RMHC because they are convex: An algorithm never needs to go downhill in order to reach the global optimum. To make

these functions hard for hillclimbers, I add *potholes* to them: valleys in the fitness landscape that block a hillclimber’s path to the optimum [19]. This produces the *pothole functions*, described in Section 3. (Holland proposed a class of RR functions, described by Jones [15], that also contained potholes.)

Pothole functions are a very simple subset of Holland’s hyperplane-defined functions (HDFs) [13]; potholes are examples of HDF *refinements*. A long-term goal is to design a class of pothole functions with parameters that can be varied to select the landscape features present in a function, as well as its overall difficulty. An arbitrary number of functions could then be generated with the desired characteristics by using those parameters to define probability distributions, which in turn could be used to choose the schemata that contribute to an individual’s fitness, along with their fitness contributions. (In this paper, these will be called a function’s *significant schemata*.) Such a class would allow a researcher to use statistical methods to calculate the certainty of statements about an algorithm’s performance across the entire class, while being easier to understand than more general classes of HDFs. However, it is not yet clear what parameters or distributions should be used, if the goal is to describe a class of functions that is easy for the SGA yet hard for hillclimbers. The immediate goal is to use hand-designed pothole functions as testbeds to determine what regions of distribution space should be used for randomly-generated functions.

I base my work on the RR functions because they were explicitly designed to investigate the validity of the BBH by studying the SGA’s performance on functions rich in building blocks. The significant schemata in a RR function are not building blocks in every population, since their fitness depends on the current population. However, the functions are defined so that they are building blocks in all contexts except in the occurrence of hitchhiking, since they make only positive fitness contributions. The functions are “rich in building blocks” in this sense; the pothole function p_1 described in Section 3 has the same characteristic. In this paper, a schema that makes a positive fitness contribution (ignoring the fitness contribution of other schemata) will be called a *beneficial schema*. (It is possible to define a building block as any beneficial schema, in contrast to the definition given earlier. This definition is related to Fisher’s [5] *average excess*, but it makes the relationship between the BBH and the schema theorem less clear [J. H. Holland, personal communication].) Like the RR functions, the pothole functions are not meant to be realistic. Since the fitness contribution of every schema is specified in advance, schemata can be used as tracers: They can be related to individual landscape features, and their frequency in the EA population can

1. Create a function that is easy for the SGA, for some performance criteria.
2. Use domain-specific knowledge to design a simple algorithm that is able to optimize that function better than the SGA. If no such algorithm can be found, or if all such algorithms incorporate unreasonable amounts of domain-specific knowledge, go to Step 4.
3. Modify the function so that it is hard for the simple algorithm yet still easy for the SGA. If no such function can be found, go back to Step 1 and start over. Otherwise, go back to Step 2.
4. Stop — a candidate function has been found.

Algorithm 1: A dialectical heuristic for finding a simple function that is easy for the SGA but hard for other optimizers. (Note that this heuristic may never succeed.)

be tracked over time [19]. Hypotheses about the effects of various landscape features on EA behavior can then be formulated and tested. This knowledge can then be applied to the study of real-world functions.

Given enough domain-specific knowledge, it is plausible that a specialized optimization method can be designed to outperform the SGA on *any* sufficiently restricted class of functions. (The NFL theorem does not hold if the subset of functions being considered has measure 0 in distribution; this is true for many subsets of interest, in particular all countable subsets [J. H. Holland, personal communication].) Therefore, the issue is not whether the SGA will outperform *all* other algorithms on a given class. Rather, it is: How much domain-specific knowledge is it reasonable to incorporate into an algorithm before it becomes over-specialized or too expensive to design and implement, outweighing any performance advantage over the SGA? A related question is: How broad must a class of functions be before the SGA outperforms a specialized optimizer on it? More generally, the RR papers suggest a *dialectical heuristic* to search for a simple function that is easy for the SGA but hard for hillclimbers (Algorithm 1). (Note that “dialectic” here simply denotes “the existence or working of opposing forces” [25].) While this heuristic is very straightforward, it has never been articulated explicitly. Since it is a heuristic, it may never succeed; however it is a plausible way to search for the desired class of functions.

The remainder of this paper provides a concrete example of using this heuristic. The pothole function p_1 is introduced and shown to be difficult for the RMHC but easy for the SGA. Then a variant of the RMHC, the lax random-mutation hillclimber (LRMHC), is defined,

which knows the depth of the potholes in p_1 and is able to cross over them to reach the optimum. This hillclimber is shown to outperform the SGA on p_1 . The paper concludes with a discussion of the results.

3 POTHOLE FUNCTIONS

Following the dialectical heuristic described in Section 2, I modified the RR functions to make them harder for simple hillclimbers by adding *potholes*. Potholes are detrimental schemata that contain beneficial schemata, and which, in turn, are necessary to reach beneficial schemata with higher fitness contributions [19]. This produces the class of *pothole functions*. All experiments in this paper were performed on the pothole function p_1 , which is defined in Table 1. That table lists all of the schemata that contribute to an individual’s fitness, along with their fitness contributions; these schemata are called the function’s *significant schemata*.

The fitness $p_1(x)$ of a string $x \in \{0, 1\}^{64}$ is given by

$$p_1(x) = \max \left\{ 0, 100 + \sum_{s \in S | x \in s} \mu(s) \right\}, \quad (1)$$

where S is the set of significant schemata for p_1 , s is a schema in S , and $\mu(s)$ is the fitness contribution of s . The notation $x \in s$ stands for “the string x is an instance of the schema s ”. Individuals have a base fitness of 100, so that in other pothole functions they may be less fit than the base fitness without having a negative fitness; the fitness is forced to be equal or greater than 0 so that fitness-proportionate selection may be used. The global optimum is a string of 64 1s, which has a fitness of 115.

The function consists of 4 *levels*, defined in Table 1. The first level consists of *elementary schemata*, each of which is a block of 8 1s. Each higher level consists of *compound schemata* composed of schemata from the previous one. The elementary and compound schemata are all beneficial schemata. An algorithm is said to *reach* a level when it finds an individual that is an instance of at least one significant schema from that level.

If p_1 consisted only of schemata s_0 – s_{14} , it would be a RR function, similar to $R2$ [19]. The additional schemata s_{15} – s_{31} are potholes. The potholes s_{15} and s_{16} prevent a single-mutation hillclimber, such as the RMHC [7], that has reached the first-level schema s_0 from reaching the second-level schema s_8 . This is because every sequence of single-bit mutations that leads from s_0 to s_8 would force the hillclimber to go downhill in fitness through one of these potholes (assuming neither of them is present to begin with), which it cannot do. Similarly, the potholes s_{17} and s_{18} prevent it from reaching s_8 if it has reached s_1 . The remaining potholes block the path to the other second-level schemata.

4 EXPERIMENTS ON p_1

I first compared the SGA against a variety of hillclimbers on p_1 , to verify that it was more difficult than the RR function $R2$ for hillclimbers such as RMHC.

4.1 SIMPLE GENETIC ALGORITHM (SGA)

The SGA used a population of 512 individuals. Two offspring were produced for each pair of parents, and the entire population was replaced in each new generation. Standard one-point crossover was used with a probability of 0.7 per mating pair. Point mutation was applied to each offspring with a probability of 0.005 per allele (mutations simply flipped the allele from 0 to 1 or vice-versa). Fitness-proportionate, or “roulette wheel”, selection was used, with σ -truncation scaling [9]:

$$f' = \max\{\min[f - (\bar{f} - c\sigma), 1.5], 0\}. \quad (2)$$

Here f' is an individual’s scaled fitness, f is its unscaled fitness, \bar{f} is the population average unscaled fitness, σ is the standard deviation of unscaled fitness in the population, and the scaling constant $c = 2$. The maximum and minimum possible scaled fitnesses are 1.5 and 0, respectively. The scaled fitnesses were then used to select the parents of the next generation. If $\sigma < 0.0001$, the unscaled fitnesses were used instead. Scaling appears to be necessary for the SGA to do well on these functions.

These parameters were chosen rather arbitrarily, since the goal is to find a class of functions that the SGA can optimize easily, without being sensitive to the exact parameter settings.

4.2 HILLCLIMBERS

When comparing the SGA with hillclimbers, it is important to report results from a variety of hillclimbers. In these experiments, I used the next-ascent hillclimber (NAHC), steepest-ascent hillclimber (SAHC), and random-mutation hillclimber (RMHC) described by Forrest and Mitchell [7], and Jones’s crossover hillclimber (XOHC) [16].

The XOHC used differs from Jones’s in that it repeats if the maximum number of jumps is reached, until the maximum number of evaluations has been performed. Jones’s original algorithm also quit if no fitness increase was found within 10000 steps; the one used here does not.

4.3 PERFORMANCE CRITERIA

In order for one algorithm to outperform another in this study, it should do so over a wide range of reasonable performance metrics. I use the number of function evaluations needed to reach the optimum as the primary

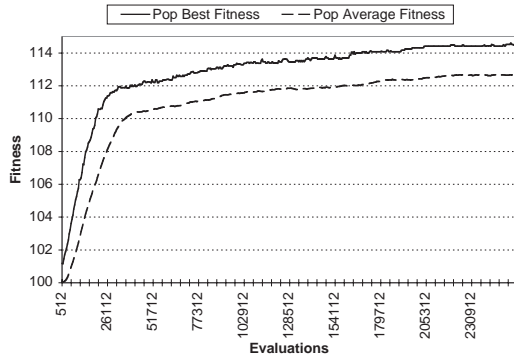


Figure 1: SGA population best and average fitness on p_1 , sampled every generation and averaged over 50 runs.




Figure 2: RMHC best and current fitness on p_1 , sampled every 512 function evaluations and averaged over 50 runs.




Figure 3: LRMHC best and current fitness on p_1 , sampled every 512 function evaluations and averaged over 50 runs. ($\epsilon = 0.1$).

Table 2: The SGA, RMHC, NAHC, SAHC, XOHC, and LRMHC on p_1 , 50 runs, 256000 function evaluations: Mean evaluations needed to reach each level (level 4 is the optimum). Here n is the number of times a level was reached, \bar{x} is the sample mean number of evaluations needed to reach each level, averaged over n , and s is the sample standard deviation of the number of evaluations.

		Level			
		1	2	3	4
SGA	n	50	50	50	48
	\bar{x}	36.9	4026.9	17838.2	65063.2
	s	32.6	1833.5	13546.4	50361.2
NAHC	n	50	48	0	0
	\bar{x}	224.6	72733.2	—	—
	s	219.1	62128.0	—	—
SAHC	n	50	48	0	0
	\bar{x}	213.3	71687.3	—	—
	s	183.1	57099.5	—	—
RMHC	n	50	2	0	0
	\bar{x}	333.7	3133.5	—	—
	s	289.6	1051.5	—	—
XOHC	n	50	50	34	0
	\bar{x}	392.0	10454.8	110519.0	—
	s	453.8	8675.6	63414.0	—
LRMHC ($\epsilon = 0.1$)	n	50	50	50	50
	\bar{x}	249.9	1342.5	3547.1	6244.0
	s	229.0	915.0	2025.4	3055.2

1.4, there is a foolproof method for determining whether a drop in fitness resulted from encountering a pothole, or from something else. If the drop is 0.1, then it is due to just a pothole and can be ignored. If it greater than or equal to 0.8, one or more beneficial schemata have been lost.

The lax random-mutation hillclimber (LRMHC) listed in Algorithm 2 resulted from incorporating this domain-specific knowledge into the RMHC. The LRMHC is exactly like the RMHC, except that it accepts any new string whose fitness is no more than ϵ below the fitness of the current string; in these experiments, ϵ is 0.1. (On p_1 , ϵ can be any value in the interval $[0.1, 0.8)$.) The algorithm is very similar to the constant threshold algorithm (CTA) developed independently by Quick et al. [23]. The only difference is that the CTA accepts a new string only if its fitness is strictly greater than the old string's fitness minus ϵ , rather than greater than or equal as in the LRMHC. (Due to a typo, the LRMHC algorithm published by Holland [13] also differs in this way from the algorithm listed here.) In turn, both algorithms are similar to the record-to-record travel algorithm and the great deluge algorithm [3], and to threshold accepting [4].

1. Initialize the current individual to a random string.
2. Mutate one randomly-chosen allele. If the new string has a fitness equal to or greater than the current individual's fitness minus ϵ , replace the current individual with the new individual.
3. If the number of fitness evaluations performed so far is less than the maximum, go to Step 2. Otherwise, stop.

Algorithm 2: The lax random-mutation hillclimber (LRMHC) algorithm. ϵ is set to 0.1 in this paper.

In effect, the LRMHC assumes that any function it encounters has potholes that all have a depth of no more than ϵ , and that it can ignore them since building blocks have an observed fitness contribution higher than this value. It might seem unreasonable to incorporate this knowledge into the LRMHC. But the RMHC can be viewed as incorporating just as much knowledge; it merely assumes that the pothole depth is always 0. The issue here is not whether the LRMHC is a useful general-purpose optimizer for real functions. Rather, it is: How much domain-specific knowledge must be built into a hillclimber so that it outperforms the SGA on p_1 ?

5.1 LRMHC EXPERIMENTAL RESULTS

Results for LRMHC on p_1 are shown in Table 2 and Figure 2. It outperforms the SGA and all of the other algorithms on p_1 , always finding the optimum, and finding it much faster than the SGA. The timeseries for LRMHC also shows a rapid increase in fitness over the SGA and RMHC. These results demonstrate that there is a very simple algorithm that outperforms the SGA by a wide margin on p_1 . Similarly, Quick et al. [23] showed that the CTA outperformed a GA variant on a class of RR functions proposed by Holland and described by Jones [15], which also contained potholes. (However, the SGA outperforms this class of hillclimber on the RR function $R4$ [20].)

6 CONCLUSIONS

This paper has presented a research program to investigate the validity of the BBH, as well as some preliminary results. A dialectical heuristic for finding a simple class of functions on which the SGA outperformed other simple search algorithms was presented. A class of pothole functions was designed by adding potholes to the RR functions, in order to make them harder for simple hillclimbers. The pothole function p_1 was shown to be hard for hillclimbers such as the RMHC but easy for the SGA. Then a new hillclimber, the LRMHC, was designed by

incorporating domain-specific knowledge about p_1 into the RMHC. While LRMHC is not useful as a general-purpose optimizer, this simple hillclimber outperformed the SGA on p_1 , demonstrating that simple pothole functions such as p_1 are still too easy for hillclimbers. This result does not by itself invalidate the BBH. However, it reinforces the finding of the RR papers that simple assumptions about what functions are especially easy for the SGA, relative to other optimizers, are often unjustified. Simple hillclimbers can be surprisingly effective at optimizing simple functions.

The next step, following the dialectical heuristic from Section 2, is to modify p_1 so that it becomes hard for the LRMHC, while remaining easy for the SGA. First, however, the LRMHC's behavior on p_1 must be investigated, in order to predict what kinds of functions it will find difficult.

Acknowledgments

I thank the UM Royal Road Group, the UM Complex Systems Reading Group, and the anonymous reviewers for their comments and suggestions.

References

- [1] Lee Altenberg. The schema theorem and Price's theorem. In L. Darrell Whitley and Michael D. Vose, editors, *Foundations of Genetic Algorithms 3*, pages 23–49, San Francisco, 1995. Morgan Kaufmann.
- [2] Kenneth A. De Jong. Genetic algorithms are NOT function optimizers. In L. Darrell Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 5–17, San Mateo, CA, USA, 1993. Morgan Kaufmann.
- [3] Gunter Dueck. New optimization heuristics : The great deluge algorithm and the record-to-record travel. *Journal of Computational Physics*, 104:86–92, 1993.
- [4] Gunter Dueck and Tobias Scheuer. Threshold accepting: A general purpose optimization algorithm appearing superior to simulated annealing. *Journal of Computational Physics*, 90:161–175, 1990.
- [5] Ronald A. Fisher. *The Genetical Theory of Natural Selection*. Dover, New York, second revised edition, 1958.
- [6] David B. Fogel. *Evolutionary Computation: Toward a New Philosophy of Machine Learning*. IEEE Press, New York, 1995.
- [7] Stephanie Forrest and Melanie Mitchell. Relative building-block fitness and the building-block hypothesis. In L. Darrell Whitley, editor, *Foundations*

- of *Genetic Algorithms 2*, pages 109–126, San Mateo, CA, USA, 1993. Morgan Kaufmann. Santa Fe Institute working paper 92-06-029.
- [8] Douglas J. Futuyma. *Evolutionary Biology*. Sinauer, Sunderland, MA, USA, third edition, 1998.
- [9] David E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA, USA, 1989.
- [10] John J. Grefenstette. Deception considered harmful. In L. Darrell Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 75–91, San Mateo, CA, USA, 1993. Morgan Kaufmann.
- [11] John H. Holland. Re: Building block hypothesis considered harmful. *Genetic Algorithms Digest*, 5(20), 1991. URL <http://www.aic.nrl.navy.mil/galist/digests/v5n20>.
- [12] John H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, MA, USA, second edition, 1992. First edition: University of Michigan Press, Ann Arbor, 1975.
- [13] John H. Holland. Building blocks, cohort genetic algorithms, and hyperplane-defined functions. *Evolutionary Computation*, 8(4):373–391, 2000.
- [14] John H. Holland, Keith J. Holyoak, Richard E. Nisbett, and Paul R. Thagard. *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, MA, USA, 1986.
- [15] Terry Jones. A description of Holland’s royal road function. *Evolutionary Computation*, 2(4):411–417, 1994.
- [16] Terry Jones. Crossover, macromutation, and population-based search. In Larry J. Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 73–80, San Francisco, 1995. Morgan Kaufmann.
- [17] John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA, 1992.
- [18] Melanie Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, USA, 1996.
- [19] Melanie Mitchell, Stephanie Forrest, and John H. Holland. The royal road for genetic algorithms: Fitness landscapes and GA performance. In Francisco J. Varela and Paul Bourguine, editors, *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, pages 245–254, Cambridge, MA, USA, 1992. MIT Press. Santa Fe Institute working paper 91-10-046.
- [20] Melanie Mitchell, John H. Holland, and Stephanie Forrest. When will a genetic algorithm outperform hill climbing? In Jack D. Cowan, Gerald Tesauro, and Joshua Alspector, editors, *Advances in Neural Information Processing Systems 6*, San Mateo, CA, USA, 1994. Morgan Kaufmann.
- [21] Allen Newell. *Unified Theories of Cognition*. Harvard University Press, Cambridge, MA, USA, 1990.
- [22] Una-May O’Reilly and Franz Oppacher. The troubling aspects of a building block hypothesis for genetic programming. In L. Darrell Whitley and Michael D. Vose, editors, *Foundations of Genetic Algorithms 3*, pages 73–88, San Francisco, 1995. Morgan Kaufmann.
- [23] Richard J. Quick, Victor J. Rayward-Smith, and George D. Smith. The royal road functions: Description, intent and experimentation. In Terence C. Fogarty, editor, *Evolutionary Computing: AISB Workshop, Brighton, U.K., April 1–2, 1996; Selected Papers*, volume 1143 of *Lecture Notes in Computer Science*, pages 223–235, Berlin, 1996. Springer.
- [24] Hans-Paul Schwefel. *Evolution and Optimum Seeking*. Wiley, New York, 1995.
- [25] J. A. Simpson and E. S. C. Weiner, editors. *The Oxford English Dictionary*. Oxford University Press, New York, second edition, 1989.
- [26] Gilbert Syswerda. Uniform crossover in genetic algorithms. In J. David Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 2–9, San Mateo, CA, USA, 1989. Morgan Kaufmann.
- [27] Michael D. Vose. A critical examination of the schema theorem. Technical Report UT-CS-93-212, University of Tennessee Computer Science Department, Knoxville, TN, USA, 1993. URL <http://www.cs.utk.edu/~library/TechReports/1993/ut-cs-93-212.ps.Z>.
- [28] Michael D. Vose. *The Simple Genetic Algorithm: Foundations and Theory*. MIT Press, Cambridge, MA, USA, 1999.
- [29] David H. Wolpert and William G. Macready. No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82, April 1997.
- [30] Sewall Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In *Proceedings of the Sixth International Congress of Genetics*, pages 356–366, 1932.