Search Space Features Underlying the Performance of Stochastic Local Search Algorithms for MAX-SAT

Holger H. Hoos[1]*, Kevin Smyth[1], and Thomas Stützle[2]

[1] Department of Computer Science, University of British Columbia
Vancouver, BC, V6T 1Z4, CANADA, {hoos,ksmyth}@cs.ubc.ca
[2] Fachbereich Informatik, Technische Universität Darmstadt
64283 Darmstadt, GERMANY, stuetzle@informatik.tu-darmstadt.de

Abstract. MAX-SAT is a well-known optimisation problem that can be seen as a generalisation of the propositional satisfiability problem. In this study, we investigate how the performance of stochastic local search (SLS) algorithms — a large and prominent class of algorithms that includes, for example, Tabu Search, Evolutionary Algorithms and Simulated Annealing — depends on features of the underlying search space. We show that two well-known measures of search space structure, the autocorrelation length of random walks and the so-called fitness distance correlation, reflect complementary factors underlying instance hardness for high-performance SLS algorithms. While the autocorrelation measure is computationally cheap, the fitness distance correlation serves mainly as an *a posteriori* measure for explaining performance. We also study the dependence of SLS performance on features of the distribution of clause weights for individual instances and show that, depending on the variance of the clause weight distribution, different search strategies seem to be suited best for dealing with the structure of the respective search spaces.

1 Introduction and Background

The satisfiability problem in propositional logic (SAT) is the task to decide for a given propositional formula whether it has a model. More formally, given a set of m clauses $\{C_1, \ldots, C_m\}$ involving n Boolean variables $x_1, \ldots, x_n$, the SAT problem is to decide whether an assignment of truth values to variables exists such that all clauses are simultaneously satisfied. This problem plays a prominent role in various areas of computer science, mathematical logic and artificial intelligence, as well as in many applications [4, 10, 2, 7].

MAX-SAT is the following optimisation variant of SAT: Given a propositional formula in conjunctive normal form (CNF), the (unweighted) MAX-SAT problem is to find a variable assignment that maximises the number of satisfied clauses. In *weighted MAX-SAT*, each clause C_i has an associated weight w_i, and the goal becomes to maximise the total weight of the satisfied clauses. Solving a given instance of a unweighted or weighted MAX-SAT instance corresponds to finding a global optimum of the objective function that maps each variable assignment a to the number or total weight of the clauses satisfied under a.

[*] to whom correspondence should be addressed

Both SAT and MAX-SAT are $\mathcal{NP}$-hard combinatorial problems which are mainly solved using heuristic search methods. Stochastic local search (SLS) algorithms for MAX-SAT are based on the idea of iteratively maximising the number of satisfied clauses. SLS algorithms are among the state-of-the-art methods for solving SAT, and by far the most effective methods for finding optimal or close-to-optimal solutions for large and hard instances of MAX-SAT (see, *e.g.*, [17]).

SAT and unweighted MAX-SAT can be seen as special cases of weighted MAX-SAT, and in principle, the same SLS methods can be applied to all these closely related problems. It is well-known that the performance of SLS algorithms on a given problem instance critically depends on the structure of the respective search space.[3] This raises the question whether there are substantial differences in the structure of the search spaces for SAT, unweighted MAX-SAT, and weighted MAX-SAT instances which might require different SLS methods to be used in order to find optimal solutions with minimal search cost.

In this work, we investigate the search space structure of various types of MAX-SAT instances and its impact on the performance of ILS-HSS, a high-performance SLS algorithm for MAX-SAT, whose performance we show to be highly correlated with that of state-of-the-art algorithms for the problem instances studied here. We use two well-known summary measures of search space structure: the autocorrelation length (ACL) [20, 18, 19] of random walks and the fitness distance correlation (FDC) for local optima [9]. Both measures are widely used for examining the search spaces of optimisation problems and are known to correlate with the hardness of problem instances for SLS algorithms [1, 11, 12]. While the autocorrelation measure is computationally cheap, the fitness distance correlation serves mainly as an *a posteriori* measure for explaining performance. We analyse (i) how these measures depend on the variance and the granularity of the clause weights, (ii) how these measures correlate with ILS-HSS performance between and within our test-sets and (iii) whether ACL and FDC reflect the same search space features. Our results show that ACL reflects well the differences of the instance hardness between test-sets with varying distributions of clause weights, while FDC seems to be better for explaining hardness differences within test-sets from a same instance distribution.

The remainder of this article is structured as follows. We first give a brief introduction to SLS algorithms for MAX-SAT and present ILS-HSS; next, we report the results of our ACL and FDC analyses. Finally, we discuss the implications of our findings, and we end with some concluding remarks and directions for further work.

2 SLS Algorithms for MAX-SAT

A large number of different SLS algorithms for MAX-SAT are known from the literature. Among these are SAMD (an early form of Tabu Search) and various forms of Simulated Annealing [5]; GSAT with Random Walk, a randomised iterative improvement algorithm [16]; GRASP [15]; the Discrete Lagrangian Method [21]; Guided Local Search [13]; variants of WalkSAT, one of the most successful SLS algorithms for SAT

[3] Following common practice, we use the the term 'search space structure' synonymously for 'search landscape structure'.

(see, *e.g.,* [8]); Ant Colony Optimisation [22]; Iterated Robust Tabu Search (IRoTS) [17]; and many others. Of these, only IRoTS, GLS and a variant of WalkSAT called Novelty$^+$ appear to reach state of the art performance on a diverse range of MAX-SAT instances [7].

In this study we use an iterated local search (ILS) algorithm for MAX-SAT that achieves excellent performance on the problem instances we considered. In general, iterated local search algorithms can be seen as performing a biased random walk in the space of the local optima encountered by an underlying local search algorithm. This walk is built by iteratively perturbing a locally optimal solution, then applying local search to obtain another locally optimal solution, and finally using an acceptance criterion for deciding from which of these solutions to continue the search. The underlying local search algorithm used in our ILS algorithm for MAX-SAT, dubbed ILS-HSS, is a tabu search algorithm with fixed tabu list length $0.12 \cdot n$. The solution perturbation is implemented by randomly flipping each variable with a probability of 0.4. If applying perturbation and subsequent local search to a solution s results in an inferior solution s', s' is accepted with a probability of 0.1; otherwise the next perturbation starts from s. In case of a tie between s and s', each is chosen with probability 0.5. The algorithm terminates when a given solution quality (*i.e.,* number or total weight of unsatisfied clauses has been reached), or a user-specified run-time limit has been exceeded, returning the best variable assignment encountered during the search. (ILS-HSS can be seen as a variant of IRoTS [17].)

3 Benchmark Sets and SLS Performance

For the computational experiments conducted in this study, we used test-sets of various types of randomly generated MAX-SAT instances. These test-sets consist of both weighted and unweighted Random 3-SAT instances with 100 instances in each set. The first of these, `rnd-u`, is a set of unweighted MAX-SAT instances that were sampled from the well-studied Uniform Random 3-SAT distribution [14] for 100 variables and 500 clauses, corresponding to the over-constrained region of Uniform Random-3-SAT. The other five test-sets were obtained from `rnd-u` by adding clause weights that were randomly generated according to truncated and discretised normal distributions. In all cases, the mean of the clause weight distribution was chosen as $\mu = 500$, and the distribution was symmetrically truncated such that all clause weights are restricted to lie in the interval $[1, 1000]$. Symmetric truncation guarantees that the actual mean is close to μ. Within this class of distributions, standard deviations σ' (before truncation) of $100, 200$ and 500 were used for generating our test-sets. For $\sigma' = 100$, test-sets for three levels of clause weight granularity were generated by discretising all weights to a grid g of size 1,10, and 100, respectively. For $\sigma' = 200$ and $\sigma' = 500$, we used test-sets with grid size one (high granularity) only. The resulting test-sets are summarised in Table 1.

For evaluating the relative performance of SLS algorithms on our test-sets, we measured run-length distributions (RLDs) for ILS-HSS on each individual problem instance [6]. This was done by running ILS-HSS $1\,000$ times on every problem instance from a given test-set, each time until a putative optimal solution for the respective instance was found. (Since for most of the MAX-SAT instances used in this study,

Name	μ	σ'	g
`rnd-u`	(unweighted)		
`rnd5-d100`	500	100	100
`rnd5-d10`	500	100	10
`rnd5-v100`	500	100	1
`rnd5-v200`	500	200	1
`rnd5-v500`	500	500	1

Table 1. Test-sets of MAX-SAT instances with different standard deviation σ' and granularity g of clause weight distributions for this study. Each test-set contains 100 instances.

determining provably optimal solutions using state-of-the-art complete MAX-SAT algorithms was found to be practically impossible due to prohibitively long run-times, putatively optimal solutions were determined as described in [17].) We refer to the mean of the RLD, that is, the average number of steps over the respective 1 000 runs, as the 'local search cost' (lsc) for the given MAX-SAT instance.

In preliminary experiments, we determined that for the test-sets considered here, the local search cost for ILS-HSS and other high-performance SLS algorithms for MAX-SAT, such as Robust Tabu Search, Iterated Robust Tabu Search, and Novelty$^+$/wcs+we [17, 7], is relatively highly correlated (see Figure 1). Similar correlations in the performance of these algorithms are observed between our various test-sets. Nevertheless, especially for weighted MAX-SAT instances with highly variable clause weights, substantial performance differences are observed in some cases.

These observations suggest that to a large extent, the performance differences between the different test-sets as well as between different instances from the same test-set are due to search space features that affect a wide range of different SLS algorithms for MAX-SAT similarly; at the same time, there appear to be other features that have a more algorithm-specific impact. (Intuitively, this is what one would expect, given that all algorithms use the same objective function, but different search heuristics.) In the following, we focus on the former type of search space features, and investigate the fundamental question to which extent these features depend on the nature of the clause weight distributions underlying a given set of MAX-SAT instances.

4 Autocorrelation Length (ACL) Analysis

The autocorrelation length (ACL) of a random walk $y_1, \ldots, y_k$ is defined based on the autocorrelation series $\hat{\rho}(h) := c_h/c_0$, where $c_i := 1/k \cdot \sum_{j=1}^{k-i}[(y_j - \bar{y}) \cdot (y_{j+i} - \bar{y})]$ is the autocovariance function with lag i [20], and $\bar{y}$ the mean over $y_1, \ldots, y_k$. Note that $-1 \leq \hat{\rho}(i) \leq 1$. ACL is defined as the lag h for which $\hat{\rho}(h)$ falls below $1/e$. (Slightly different definitions can be found in the literature, *e.g.*, ACL := $-1/ln(\hat{\rho}(1))$, but for these we have observed very similar results to those reported in the following.)

Intuitively, the ACL of a random walk through the search spaces of the given problem instances indicates how many random steps are required before the correlation between the objective function value at the current search position and the objective function value at the starting position of the walk becomes insignificant. High ACL values

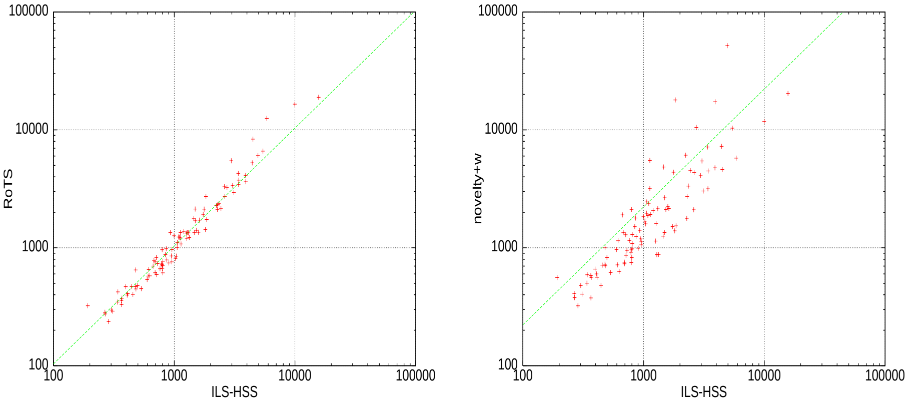

Fig. 1. Local search cost correlation between ILS-HSS and other SLS algorithms for MAX-SAT on the `rnd-u` (unweighted) test-set. Left: ILS-HSS *vs* RoTS. Right: ILS-HSS *vs* Novelty^{+}/wcs+we. Along the diagonal lines, the two given algorithms require equal CPU time (compensating for differences in CPU time per search step between the algorithms).

are an indication that search positions which are 'Hamming-close' tend to have roughly similar objective function values; this is indicative of a relatively 'smooth' search space. Conversely, low ACL values indicate that the search space is rather 'rugged' with many uncorrelated local features. The latter type of search space structure should intuitively be more difficult to deal with for SLS algorithms based on randomised improvement strategies, such as the ILS algorithm considered here.

In this section, we examine how ACL is affected by changing the variance and granularity in the clause weights, and investigate the correlation between ACL and hardness within each test-set.

ACL measurements. To generate the ACL data, we generated 100 conflict-directed random walk trajectories for each instance (that is, during each local search step, an unsatisfied clause was first selected, and then a random variable from the clause was flipped), with a length of 10 000 steps each. The ACL of each of these trajectories was calculated, and then all of the ACLs were averaged to obtain the ACL value for each instance. Table 2 shows the mean ACL over all instances in each test-set. The ACL data was also analysed for a correlation with lsc (*cf.* Table 3).

Influence of variability and granularity of clause weights on ACL. As shown in Table 2, ACL depends strongly on the variability and granularity of clause weights: average ACL decreases with increasing variability and granularity, indicating that the (weighted) objective function values over the respective search spaces tend to be most rugged for instances with high variability and fine granularity in the clause weights. Variability in clause weights appears to have a more pronounced effect on average ACL than granularity.

Correlation of ACL and local search cost over different test-sets. As can be seen from Table 2, the variation in average ACL (where the average is over different random walks for the same problem instance) between different problem instances from the

Test-Set	Mean ACL	$\sigma_x/\bar{x}$	Mean lsc (ILS-HSS)
rnd-u	41.76	0.09	1 652
rnd5-d100	40.89	0.09	3 380
rnd5-d10	40.85	0.09	4 099
rnd5-v100	40.69	0.09	4 162
rnd5-v200	38.46	0.09	8 576
rnd5-v500	36.09	0.10	13 772

Table 2. Mean ACL over all instances in each test-set, compared with the mean lsc (*i.e.*, the mean number of steps required by each algorithm to find an optimal solution for a given problem instance). The column labelled $\sigma_x/\bar{x}$ shows the coefficient of variation within each test-set.

Test-Set	a	b	r
rnd-u	-1.26	50.60	**-0.278**
rnd5-d100	-1.15	49.90	**-0.228**
rnd5-d10	-1.05	49.29	-0.189
rnd5-v100	-0.98	48.56	-0.181
rnd5-v200	-1.45	51.07	**-0.344**
rnd5-v500	-0.41	39.84	-0.103

Table 3. Correlation analysis of ACL *vs* lsc; a and b are the slope and intercept, respectively, of the regression line, r is the correlation coefficient (statistically significant if shown in bold).

same test-set is very low. This is in contrast with the high variability in local search cost observed within these test-sets (see Figure 1). However, our ACL data suggests a inverse relationship between average ACL and average local search cost. This is consistent with the intuition that SLS algorithms based on (randomised) greedy descent strategies should have more difficulties to find optimal or high-quality solutions in rugged search spaces full of local features, most of which do not provide useful guidance [1].

Correlation of ACL and local search cost within test-sets. To investigate whether the observed correlation is solely due to the differences in the syntactical features of the respective instances (i.e., to due to the differences in their clause weights), or whether it reflects a deeper dependency between ACL and local search cost, we analysed the same correlation for the instances within each individual test-set. The results shown in Table 3 indicate a weak negative correlation between average ACL and average local search cost per instance within each test-set. Although the correlation is generally too weak to be statistically significant, it is present for all test-sets and is consistent with the negative correlation of local search cost and ACL observed over multiple test-sets. [4]

5 Fitness Distance Correlation (FDC) Analysis

FDC is defined as the correlation between the 'fitness' of local minima states, and the 'distance' from those local minima states to the closest global optimum [9]. In this

[4] Correlations were measured using the Pearson correlation coefficient r, and their significance was assessed based on a standard test statistic for r (which, for the sample sizes used here, is approx. t-distributed); in our tables, correlations that are statistically significant at $\alpha = 0.05$ are shown in boldface.

Test-Set	Mean FDC	$\sigma_x/\bar{x}$	Mean lsc (ILS-HSS)
`rnd-u`	0.416	0.44	1 652
`rnd5-d100`	0.343	0.51	3 380
`rnd5-d10`	0.354	0.47	4 099
`rnd5-v100`	0.358	0.49	4 162
`rnd5-v200`	0.362	0.47	8 576
`rnd5-v500`	0.379	0.45	13 772

Table 4. Mean FDC over all instances in each testset, compared with the mean lsc.

study, we define the fitness of a variable assignment search position as the difference between its objective function value and the optimal objective function value for the respective problem instance. We define the distance between two states as the Hamming distance between the states. The FDC coefficient is then defined as

$$\rho_{fdc}(f,d) := \frac{Cov(f,d)}{\sigma(f)\cdot\sigma(d)} = \frac{\langle f(s)\cdot d(s)\rangle - \langle f(s)\rangle\cdot\langle d(s)\rangle}{\sqrt{\langle f^2(s)\rangle - \langle f(s)\rangle^2}\sqrt{\langle d^2(s)\rangle - \langle d(s)\rangle^2}},$$

where $Cov(f,d)$ denotes the covariance of the fitness-distance pairs $(f(s),d(s))$ over all variable assignments s; $\sigma(f)$ and $\sigma(d)$ are the respective standard deviations of the fitness and the distance values for all s; and $\langle f(s)\rangle$, $\langle f^2(s)\rangle$, $\langle f(s)\cdot d(s)\rangle$ denote the averages of $f(s)$, $f^2(s)$, and $f(s)\cdot d(s)$, respectively, over all assignments s. By definition, we have that $-1 \leq \rho_{fdc}(f,d) \leq 1$. A significantly positive FDC coefficient indicates that, on average, with increasing solution quality the search is also getting closer to an optimal solution, and therefore the objective function can be expected to effectively guide the local search.

FDC measurements. FDC coefficients for each problem instance in our test-sets were determined as follows. First, we constructed a set S of putatively optimal solutions.[5] Second, we ran ILS-HSS on the given instance, recording every local minimum (LMIN) state encountered during the search along with the fitness of the state, until a putatively optimal solution was encountered. (LMIN states correspond to assignments which cannot be improved by a single variable flip.) This process was repeated until a minimum of 50 000 LMIN states had been found. We then computed the Hamming distance between each of these LMIN states and its Hamming-closest element of S, and output the respective fitness-distance pair. Finally, the FDC coefficient was calculated as the correlation coefficient over these sets of fitness-distance pairs.

Influence of variability and granularity of clause weights on FDC. The results reported in Table 4 indicate that, with the exception of our set of unweighted instances, there are small, but systematic differences in FDC depending on the variablity and granularity of clause weights: Mean FDC values increase monotonically with granularity and variability of clause weights. Compared to the respective ACL results, the variation of FDC values within each test-set is considerably higher (as reflected in the higher

[5] Since measuring S exactly was computationally intractable for the instance types and sizes used here, we determined putatively optimal solution qualities as described in [17] and approximated S for each instance using the set of unique solutions of that quality obtained from many runs of a high-performance SLS algorithm for MAX-SAT.

Test-Set	FDC *vs* lsc			ACL *vs* FDC		
	a	b	r	a	b	r
rnd-u	-0.12	1.31	**-0.593**	0.02	-0.22	**0.319**
rnd5-d100	-0.13	1.38	**-0.553**	0.01	-0.15	**0.254**
rnd5-d10	-0.12	1.23	**-0.486**	< 0.01	0.20	0.085
rnd5-v100	-0.13	1.39	**-0.506**	0.01	0.13	0.118
rnd5-v200	-0.08	1.07	**-0.405**	< 0.01	0.20	0.087
rnd5-v500	-0.07	0.97	**-0.334**	< 0.01	0.21	0.091

Table 5. Correlation analysis of FDC *vs* lsc (left side) and of ACL *vs* FDC (right side).

coefficient of variation). These results indicate a positive correlation between FDC values and mean local search cost — a counterintuitive finding, since we would expect that strong fitness-distance correlations (*i.e.*, high FDC values) should make a problem instance easier for an SLS algorithm such as ILS-HSS, which is essentially based on a (randomised) iterative improvement strategy. One possible explanation for this somewhat surprising result is that ruggedness effects (as captured by our ACL analysis) might dominate the impact of FDC for these test-sets.

Correlation between FDC and local search cost within test-sets. Our analysis of the correlation between FDC and mean local search cost per instance within each test-set revealed a different picture: Table 5 shows that there is a significant negative correlation, which is consistent with the previously discussed, expected impact of FDC on local search cost. This correlation seems to weaken with increasing granularity and variability of clause weights, which is probably due to the fact that for these types of instances, ILS-HSS finds significantly less high-quality local minimum states than for the much easier instances from the test-sets with low granularity and low clause weight variability. However, the correlations between FDC and mean local search cost within our test-sets are all statistically significant and much stronger than the correlations between ACL and mean local search cost within test-sets.

Correlation between ACL and FDC within test-sets. Our experimental results indicate that both ACL and FDC are correlated to some extent with mean local search cost. This raises the question whether both of these measures merely reflect the same factors underlying the hardness of a problem instance for ILS-HSS, or whether they are complementary. It seems that ACL better captures the significant differences in local search cost between test-sets, while within each test-set, search cost was much stronger correlated with FDC on a per-instance basis. This is further supported by the results of a correlation analysis between ACL and FDC within each test-set, shown in Table 5. The correlation between ACL and FDC is very weak and, except for the sets of unweighted and low granularity instances, statistically insignificant, though consistently positive for all test-sets. The latter observation is interesting, because it indicates that there appears to be a slight tendency (within the test-sets) for instances with more rugged search spaces to have a lower fitness-distance correlation. However, the weakness of this correlation and the qualitative differences in our observations regarding ACL and FDC suggests that both measures reflect different relevant aspects of search space structure.

6 Conclusions and Future Work

The results presented in this study show that autocorrelation length (ACL) and fitness-distance correlation (FDC) are useful for predicting and explaining the hardness of MAX-SAT instances for various stochastic local search algorithms. Our empirical results show that ACL and FDC capture different factors underlying local search cost. While ACL reflects the significant differences in instance hardness between test-sets with varying distributions of clause weights, FDC seems better suited for explaining the search cost differences within test-sets of instances sampled from the same distribution. ACL and FDC are also complementary in the sense that they both correlate with instance hardness, but the correlation between them is rather weak; hence, a combination of both measures has a higher predictive and explanatory power than either ACL or FDC alone. Furthermore, experimental results not included here indicate that FDC actually correlates better with local search cost than the number of optimal solutions for test-sets of instances with very high variance in the clause weights. This comes as a surprise considering the dominance of number of solutions on local search cost for unweighted SAT instances [3].

This work raises a number of interesting questions and suggests various avenues for further research. First and foremost, it would be very interesting to further explore the performance differences observed between different SLS algorithms for MAX-SAT, such as ILS-HSS, RoTS and Novelty^{+}/wcs+we, and to analyse these for possible correlations with the measures of search space structure investigated here. Secondly, it may be possible to use computationally cheap measures of search space structure, such as ACL, to predict the run-time of specific SLS algorithms for MAX-SAT, and to select the algorithm that is expected to solve a given instance most efficiently. Further directions for extending this research include an analysis of the impact of problem size (*i.e.*, number of variables and clauses) on the results reported here; an investigation of search space features of structured MAX-SAT instances (as obtained, *e.g.*, from encodings of other combinatorial problems into MAX-SAT) and their differences to those of the Random-3-SAT considered here; and a more detailed analysis of other search space characteristics (*e.g.*, the plateau and basin structure of a given search landscape) that explain some of the observed variation in local search cost that cannot be explained by global measures such as ACL and FDC.

Overall, we strongly believe that insights gained from a thorough understanding of the search space characteristics and their dependence on syntactic features of MAX-SAT instances as well as their impact on the performance of SLS algorithms for MAX-SAT will provide the key to further improving the efficiency of these algorithms, and hence the state-of-the-art in MAX-SAT solving.

References

1. E. Angel and V. Zissimopoulos. On the classification of NP-complete problems in terms of their correlation coefficient. *Discrete Applied Mathematics*, 99(1–3):261–277, 2000.
2. P. Asirelli, M. de Santis, and A. Martelli. Integrity constraints in logic databases. *J. of Logic Programming*, 3:221–232, 1985.

3. D. Clark, J. Frank, I. Gent, E. MacIntyre, N. Tomov, and T. Walsh. Local search and the number of solutions. In E. Freuder, editor, *Proc. of CP'96*, volume 1118 of *LNCS*, pages 119–133. Springer Verlag, Germany, 1996.
4. J. Gu and R. Puri. Asynchronous circuit synthesis with boolean satisfiability. *IEEE Transactions on Computer–Aided Design of Integrated Circuits*, 14(8):961–973, 1995.
5. P. Hansen and B. Jaumard. Algorithms for the maximum satisfiability problem. *Computing*, 44:279–303, 1990.
6. H.H. Hoos and T. Stützle. Characterising the behaviour of stochastic local search. *Artificial Intelligence*, 112:213–232, 1999.
7. H.H. Hoos and T. Stützle. *Stochastic Local Search—Foundations and Applications*. Morgan Kaufmann Publishers, USA, 2004 (to appear).
8. Y. Jiang, H. Kautz, and B. Selman. Solving problems with hard and soft constraints using a stochastic algorithm for MAX-SAT. In *Proceedings of the 1st International Joint Workshop on Artificial Intelligence and Operations Research*, 1995.
9. T. Jones and S. Forrest. Fitness Distance Correlation as a Measure of Problem Difficulty for Genetic Algorithms. In *Proceedings of ICGA'95*, pages 184–192. Morgan Kauffman, 1995.
10. A.P. Kamath, N.K. Karmarkar, K.G. Ramakrishnan, and M.G.C. Resende. A continuous approach to inductive inference. *Mathematical Programming*, 57:215–238, 1992.
11. P. Merz and B. Freisleben. Fitness landscapes and memetic algorithm design. In D. Corne, M. Dorigo, and F. Glover, editors, *New Ideas in Optimization*. McGraw-Hill, 1999.
12. P. Merz and B. Freisleben. Fitness landscape analysis and memetic algorithms for the quadratic assignment problem. *IEEE Transact. on Evol. Comput.*, 4(4):337–352, 2000.
13. P. Mills and E. Tsang. Guided local search for solving SAT and weighted MAX-SAT problems. In I.P. Gent, H. van Maaren, and T. Walsh, editors, *SAT2000 — Highlights of Satisfiability Research in the Year 2000*, pages 89–106. IOS Press, 2000.
14. D. Mitchell, B. Selman, and H. Levesque. Hard and easy distributions of SAT problems. In *Proc. of AAAI'92*, pages 459–465. MIT Press, 1992.
15. M.G.C. Resende, L.S. Pitsoulis, and P.M. Pardalos. Approximate solution of weighted MAX-SAT problems using GRASP. In D. Du, J. Gu, and P.M. Pardalos, editors, *Satisfiability problem: Theory and Applications*, volume 35, pages 393–405. AMS, 1997.
16. B. Selman, H.A. Kautz, and B. Cohen. Noise strategies for improving local search. In *Proc. of the 12th Natl. Conf. on Artificial Intelligence*, pages 337–343. MIT Press, 1994.
17. K. Smyth, H. H. Hoos, and T. Stützle. Iterated robust tabu search for MAX-SAT. In Y. Xiang and B. Chaib-draa, editors, *Advances in Artificial Intelligence*, volume 2671 of *LNCS*, pages 129–144. Springer Verlag, Germany, 2003.
18. P. F. Stadler. Towards a theory of landscapes. In R. Lopéz-Peña et al., editor, *Complex Systems and Binary Networks*, volume 461 of *Lecture Notes in Physics*, pages 77–163. Springer Verlag, Berlin, Germany, 1995.
19. P.F. Stadler and R. Happel. Random field models for fitness landscapes. *Journal of Mathematical Biology*, 38(5):435–478, 1999.
20. E.D. Weinberger. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biol. Cyber.*, 63:325–336, 1990.
21. Z. Wu and B.W. Wah. Trap escaping strategies in discrete lagrangian methods for solving hard satisfiability and maximum satisfiability problems. In *Proc. of AAAI'99*, pages 673–678. MIT Press, 1999.
22. M. Zlochin and M. Dorigo. Model-based search for combinatorial optimization: A comparative study. In J. J. Merelo et al., editors, *Proceedings of PPSN-VII*, volume 2439 of *LNCS*, pages 651–661. Springer Verlag, Germany, 2002.