Epilogue

In this thesis we have studied various aspects of stochastic local search (SLS), an algorithmic paradigm for solving hard combinatorial problems which has become increasingly popular in Artificial Intelligence and Operations Research over the past few years. Our work focussed on modelling SLS algorithms, empirically evaluating their performance, characterising and improving their behaviour, and understanding the factors which influence their efficiency. We studied these issues for the SAT problem in propositional logic as our primary application domain. SAT has the advantage of being conceptually very simple, which facilitates the design, implementation, and presentation of algorithms as well as their analysis. However, most of the methodology generalises easily to other combinatorial problems like CSP. The main contributions of this work can be summarised as follows.

- We introduced the *GLSM model*, a novel formalism for representing SLS algorithms which is based on the concept of a clear separation between search strategies and search control (Chapter 3). The GLSM model allows the adequate representation of many modern SLS algorithms for SAT (Chapter 4) and provides a basis for analysing and improving these (Chapter 5).

- We developed an adequate methodology for empirically analysing the behaviour of Las Vegas algorithms (a more general class of algorithms than SLS) which avoids weaknesses of established approaches (Chapter 2). We used this methodology for studying SLS algorithms for SAT and characterising their behaviour (Chapters 4 and 5).

- We studied the performance of a number of popular and very powerful SLS algorithms for SAT across a broad range of problem domains (Chapters 4 and 7). We found that the relative performance of modern SLS algorithms varies with the problem domain such that there is no single best algorithm. We could also show that there is a tight correlation between the performance of different SLS algorithms within the domains, which suggests that SLS performance generally depends on the same features of the underlying search spaces.

- We proved a number of results concerning the asymptotic behaviour (PAC property) of various modern SLS algorithms for arbitrary long runs of the corresponding pure strategies. We could show that some of the most powerful SLS algorithms for SAT suffer from essential incompleteness and demonstrated how to overcome this defect in practice (Chapter 5).

- We developed a functional characterisation of SLS behaviour, showing that for optimal and larger-than-optimal noise parameter settings, the RTDs for modern SLS algorithms, when applied to hard problem instances from different domains, can be well approximated using exponential distributions (EPAC behaviour, *cf.* Chapter 5). This result has a number of interesting consequences regarding the parameterisation and parallelisation of SLS algorithms, and it suggests a novel interpretation of SLS behaviour as random picking in a significantly smaller "virtual" search space.

- We analysed various aspects of search space structure and their influence on SLS performance (Chapter 6). We confirmed and refined earlier results regarding the dominant role of the number of solutions and investigated other search space features such as the variability of the objective function and the topology of local minima regions. Our results show that local minima regions resemble systems of narrow canyons rather than large basins; the branching of these structures is very low, especially for SAT-encoded problems from other domains. This explains the effectiveness of simple escape strategies like random walk.

- We investigated the influence of different SAT-encoding strategies on search space structure and SLS performance for test-sets of hard Constraint Satisfaction and Hamilton Circuit Problems (Chapter 7). Confirming earlier results for planning domains, our results show that compact encodings induce relatively small search spaces which are nevertheless very difficult for search for SLS algorithms. Similarly, eliminating symmetric solutions of the Hamilton Circuit Problem seems to have detrimental effects on SLS performance. We show how the techniques and methods developed earlier in this thesis can be used to analyse and identify and some of the search space features which are responsible for these effects. Furthermore, we show that for solving hard Random Binary CSP instances, using efficient SAT-encodings and state-of-the-art SLS algorithms for SAT is competitive to the best SLS algorithms for CSP.

Open Questions and Future Work

Many contributions of this thesis, including methodological aspects as well as the results of empirical and theoretical investigations, are interesting on their own. However, they also give rise to a number of questions, shed new light on some open problems, and suggest many routes for further research. In the following, we will briefly address some of these issues.

- Our empirical studies of SLS algorithms for SAT were mainly limited to algorithms which could be represented using simple instances of the basic GLSM model. Many of the advanced features and extensions of the GLSM model have not been used in this context; their potential for SAT and other problem domains has to be further investigated and assessed. In the light of our results for SAT, the heterogenous cooperative models, possibly including learning schemes and evolutionary population dynamics, appear to be particularly attractive.

- We have shown that the GLSM model can be used to improve some of the best currently known SLS algorithms for SAT. However, the corresponding GLSMs are still very simple; it remains to be seen, whether more complex combinations of pure SLS strategies show substantially improved performance and / or robustness. Further research on this issue should probably not be exclusively focussed on combinatorial decision problems like SAT, but also include classical optimisation problems like TSP or scheduling.

- While we could prove approximate completeness for GWSAT and essential incompleteness for WalkSAT+tabu, Novelty, and R-Novelty, the question whether WalkSAT is approximately complete remains open. Our empirical results give no indication of essential incompleness, therefore we suspect that WalkSAT might be approximately complete. Furthermore, we did not prove essential incompleteness for GSAT and GSAT+tabu; but we believe that the proofs are not difficult. As we have shown, essential incompleteness leads to poor performance and robustness of SLS algorithms in practice. Therefore, approximate completeness and the avoidance of stagnation behaviour in practice should be considered as important design objectives for future development of SLS algorithms.

- We developed functional characterisations of the run-time behaviour of a number of SLS algorithms for SAT when using optimal or larger-than-optimal noise settings. Since the same EPAC behaviour was found for all the algorithms and all hard problems studied here, we suspect that this type of behaviour is characteristic for SLS algorithms when applied to hard combinatorial problems. This hypothesis is backed up by preliminary results on modelling the behaviour of SLS algorithms for Constraint Satisfaction Problems. We did not model SLS behaviour for lower-than-optimal noise; but we believe that the corresponding RTDs can be approximated using combinations of the generalised exponential distribution family introduced in Chapter 5.

- As discussed in Chapter 5, the observed EPAC behaviour suggests a novel interpretation of SLS behaviour as random picking in a reduced "virtual" search space. The size of this search space is directly related to the average local search cost for solving the corresponding problem instance. While in Chapter 6, we found some search space features which are to some extent correlated to the average local search cost, we could not find any structural search space features which directly correspond to the virtual search space. Generally, from our perspective the attempt to find explanations for the observed SLS behaviour based on the simple mathematical models developed here seems to be one of the most exciting and promising issues for further research.

- While we studied search space structure and its influence on SLS performance as well as the effect of different SAT-encodings on search space structure, many questions in this context remain unanswered. Our results suggest that features different from the ones investigated here play an important role; further work should try to identify and characterise these. Furthermore, the analyses of the search space structural features induced by different encodings and their impact on SLS peformance presented in Chapter 7 should be extended to different problem sizes and other domains. We are convinced that further insights into these dependencies will be crucial for the further development of SLS algorithms, encoding strategies, and the general applicability of the generic problem solving approach discussed in Chapter 1.

- Finally, we are convinced that most of the methodology developed in the context of this thesis can equally well be used for studying SLS algorithms for optimisation problems. To some extent, the required generalisations have already been introduced in this work (*cf.* Chapter 2). Consequently, conducting investigations analogous to the ones presented here for optimisation problems like MAXSAT or MAXCSP should not be difficult. Since SLS algorithms for decision problems like SAT implicitly solve the corresponding optimisation problem, we expect that this line of future work will provide interesting insights into SLS behaviour in general.

Building on previous work on stochastic local search, our goal in this work was to enhance the general understanding of SLS algorithms and their behaviour, to evaluate their performance, and, ultimately, to improve existing algorithms. While, hopefully, we reached this goal to some extent, many fundamental questions regarding SLS algorithms and their applications remain unanswered. From our point of view, the most important of these are the precise understanding of the factors which determine the efficiency of SLS algorithms and a general assessment of the role of SLS algorithms in the context of the generic problem solving approach discussed in Chapter 1 of this thesis. While we cannot provide the ultimate solutions to these important problems, we hope that the contributions of this thesis might be regarded as a further step towards finding satisfying solutions to these problems and will prove to be useful for future SLS research.