

Analysing differences between algorithm configurations through ablation

Chris Fawcett · Holger H. Hoos

Received: date / Accepted: date

Abstract Developers of high-performance algorithms for hard computational problems increasingly take advantage of automated parameter tuning and algorithm configuration tools, and consequently often create solvers with many parameters and vast configuration spaces. However, there has been very little work to help these algorithm developers answer questions about the high-quality configurations produced by these tools, specifically about which parameter changes contribute most to improved performance. In this work, we present an automated technique for answering such questions by performing *ablation analysis* between two algorithm configurations. We perform an extensive empirical analysis of our technique on five scenarios from propositional satisfiability, mixed-integer programming and AI planning, and show that in all of these scenarios more than 95% of the performance gains between default configurations and optimised configurations obtained from automated configuration tools can be explained by modifying the values of a small number of parameters (1–4 in the scenarios we studied). We also investigate the use of our ablation analysis procedure for producing configurations that generalise well

C. Fawcett
University of British Columbia
Computer Science Department
201-2366 Main Mall
Vancouver, BC, V6T 1Z4, Canada
Tel.: +1 604 822-4599
Fax: +1 604 822-5485
E-mail: fawcettc@cs.ubc.ca

H.H. Hoos
University of British Columbia
Computer Science Department
201-2366 Main Mall
Vancouver, BC, V6T 1Z4, Canada
Tel.: +1 604 822-1964
Fax: +1 604 822-5485
E-mail: hoos@cs.ubc.ca

to previously-unseen problem domains, as well as for analysing the structure of the algorithm parameter response surface near and between high-performance configurations.

Keywords ablation analysis · parameter importance · automated algorithm configuration · empirical analysis

1 Introduction

High-performance solvers for hard computational problems, such as propositional satisfiability (SAT) or mixed-integer programming (MIP), are typically run by users on classes of problem instances from different application domains; for example, a SAT solver might be run on sets of random 3-SAT, hardware verification or software verification instances, and a MIP solver might be used to solve MIP-formulations of logistics, routing or production planning problems. The existence of such varied domains provides an incentive to the developers of such solvers to parameterise aspects of their implementation in order to be able to obtain good performance on each target problem domain. Finding good values manually for these algorithm parameters is difficult, as even human experts have trouble predicting which configurations will result in high performance due to interactions between parameters and the sheer size of the combinatorial configuration spaces involved.

While tools specifically designed for automatically tuning the parameters of such algorithms have been in use for at least a decade (see, *e.g.*, Birattari et al (2002)), the introduction of advanced procedures capable of dealing with dozens of parameters, such as ParamILS (Hutter et al, 2007b, 2009), GGA (Ansótegui et al, 2009), irace (López-Ibáñez et al, 2011) and SMAC (Hutter et al, 2011), has generated great interest in the area of automated algorithm configuration. The success of these automatic algorithm configurators in practice has inspired a software design paradigm called *Programming by Optimisation* (PbO) (Hoos, 2012), which encourages developers to expose design choices and actively seek alternatives for key parts of their algorithms, leading to highly parametric designs that are then automatically optimised for specific use contexts.

However, many configurations are sampled by these configuration tools, and developers are often left wondering *why* their algorithm parameters were set to specific values by the automated configuration process, or whether the modification of some parameters from their default settings was truly necessary to achieve substantially improved performance. Given a highly parameteric algorithm, after making many parameter changes as a result of automated configuration, how can an algorithm developer know which of the parameter changes were actually important? The ability to answer questions like these will allow developers to focus their efforts on the aspects of their solvers that are providing the most performance gains (or losses), in an iterative algorithm development process.

In this work, we introduce the concept of *ablation analysis*, a procedure investigating the path of configurations obtained by iteratively modifying parameter settings from a source configuration (*e.g.*, an expert-defined default) to those from a target configuration (*e.g.*, one obtained from an automatic configurator). Parameter values are modified one at a time, and at each stage the configuration with the best performance is retained.¹ We present a brute-force approach to this analysis, as well as an accelerated version that takes advantage of racing methods for algorithm selection. We demonstrate the effectiveness of this approach with an empirical study on five well-studied algorithm configuration scenarios that involve high-performance solvers for propositional satisfiability, mixed integer programming and AI planning problems, and we show that for these scenarios, more than 95% of the performance gains from automated configuration can be obtained by the modification of at most 4 (out of 26–76) algorithm parameters.

We also present the results of several additional experiments using our ablation analysis approach to produce algorithm configurations that generalise well to unseen problem instances from different domains. Finally, we carried out experiments to analyse the structure of the algorithm parameter response surface near and between configurations with high performance, by performing ablation analysis with these high-performance configurations as the source and target configurations.

The remainder of this paper is structured as follows: In Section 2, we place our contribution in context with related work in parameter importance, and we then provide an in-depth explanation of both variants of our ablation analysis procedure in Section 3. Section 4 presents the details of the experimental study that we performed, with the results of that study shown and discussed in Sections 5 and 6. We discuss opportunities for future work in this area in Section 7 and then conclude in Section 8.

2 Background and Related Work

Many individual applications of automated algorithm configuration to specific solvers include statements from the authors about the modified parameters, as a post-hoc subjective justification without formal analysis. Examples of this include the configuration of a state-of-the-art industrial SAT solver (Hutter et al, 2007a), as well as the automated design of general-purpose frameworks for AI planning (Vallati et al, 2011, 2013). However, there has been relatively little work on systematic techniques for assessing parameter importance. The most closely related area of related work is that of sensitivity analysis in statistics, especially analysis of variance (ANOVA) and func-

¹ Our use of the term *ablation* follows that of Aghaeepour and Hoos (2013) and loosely echoes its meaning in medicine, where it refers to the surgical removal of organs, organ parts or tissues. We ablate (*i.e.*, remove) changes in the settings of algorithm parameters to better understand the contribution of those changes to observed differences in algorithm performance.

tional ANOVA (Hooker, 2007) approaches to decomposing model or function response variance into low-order components. Furthermore, related work on interactive parameter exploration using contour plot visualization (Bartz-Beielstein, 2006), on evolutionary algorithms for parameter relevance estimation (Nannen and Eiben, 2007) and on experimental design for analysing optimization algorithms (Chiarandini and Goegebeur, 2010) can be found in the literature.

Chiarandini and Goegebeur (2010) present a thorough investigation of experimental design for analysing optimization algorithms, using linear mixed-effects models. Their analysis only includes parameter configuration spaces of low size and dimension, with fewer than 50 possible configurations. It is unclear whether this approach scales to larger configuration spaces (the scenarios considered in our study comprise up to 1.90×10^{47} possible configurations). Nannen and Eiben (2007) uses an entropy-based approach with evolutionary algorithms to perform parameter relevance estimation, but this approach also makes a smoothness assumption about the parameter response surface which eliminate categorical parameters from consideration, as well as an assumption that there will be few parameters (on the order of 10) in the problem. Finally, Bartz-Beielstein (2006) uses sequential parameter optimization (SPO) to estimate 1- and 2-parameter effects, as well as providing interactive contour plots in the configuration space. This approach also has difficulties with the discrete configuration spaces induced by the categorical parameters encountered in many algorithm configuration scenarios. As these techniques each have difficulties with the high dimensionality and/or the discrete nature of the configuration spaces of typical highly-parameterised algorithms, an approach overcoming both has high utility for algorithm developers.

Very recently, Hutter *et al.* have been using model-based techniques to investigate the problems of parameter importance and parameter interaction directly, using forward selection (Hutter et al, 2013) and functional ANOVA (Hutter et al, 2014). Both approaches require an initial data-gathering step to obtain algorithm performance data, which is then partitioned into training and test sets. In (Hutter et al, 2013), this data was obtained by sampling 1 000 – 10 000 pairs of configurations and instances uniformly at random, while in (Hutter et al, 2014) several experiments were performed using 10 000 randomly sampled runs, the algorithm runs performed during executions of the SMAC configurator (13 452 – 454 336 additional runs), as well as a combination of both.

In the forward selection approach by Hutter et al (2013), this performance data is used to iteratively build a regression model by greedily adding, at each iteration, the parameter or instance feature which results in a model with the lowest root mean squared error on the validation set. Hutter et al (2014), on the other hand, introduced an efficient technique for applying functional ANOVA to random forest models. This variance decomposition takes a random forest model constructed from the precomputed data, and expresses the performance variation in terms of components, with one component for every subset of

parameters of size up to k (for small k). These two contributions differ from our own in several fundamental ways.

The current version of the forward selection approach constructs models wholly or partially based on thousands of configurations sampled uniformly at random from the configuration space. The CPU time required to obtain this data, as well as the time required to build the models themselves, can be significant. The CPU time requirements for model construction are especially significant for forward selection, which typically requires the construction of thousands of models.

More importantly, this random sampling of configurations means that many of the configurations used to build the model are from parts of the configuration space that are unlikely to contain high quality configurations. Furthermore, both methods have so far been used only to measure parameter importance globally on expectation across the entire configuration space or, for functional ANOVA, to broad sets of samples restricted to: (1) configurations with better performance than the default configuration, or (2) configurations in the top 25% in terms of performance. The importance values derived from those experiments are still global measures, and can be averages across many regions of very different high-performance configurations. There is no guarantee that these importance measures apply to any individual algorithm configuration, specifically to any high-performance configuration and the local neighbourhoods around such configurations. Finally, the functional ANOVA work relies on the assumption that accurate models of algorithm performance can be obtained at a reasonable computational cost. This appears to be the case for the experiments reported by Hutter et al (2014), but there is no guarantee that on other scenarios, models with similar parameter importance accuracy can be practically obtained. Our approach does not require model construction, and is therefore not constrained by this assumption.

The most important distinguishing factor between our work presented here and these earlier studies lies in the fact that we aim to explain the importance of differences between two algorithm configurations that are of interest to an algorithm developer and user – for example, between the default configuration and one produced by applying an automated algorithm configuration tool, such as PARAMILS. Using ablation analysis, we can quantify the performance losses (or gains) along the “ablation path” (see Section 3.3) from one configuration to another. This allows algorithm developers or users to find a minimal set of parameter modifications from a given default configuration, while maintaining most or all of the performance gains achieved by automated algorithm configuration. We see this approach as complementary to the recent model-based techniques of Hutter et al (2013, 2014), as the local information provided by our approach can strengthen and validate (or invalidate) the results obtained with those techniques. We also believe that there are ways to combine the two lines of work (see Section 8).

There is an interesting conceptual connection between our ablation approach and that of *path relinking*, a general-purpose method for combining the diversification and intensification stages in heuristic search (Glover, 1994,

Algorithm 1: Ablation ($\mathcal{A}, \theta_{\text{source}}, \theta_{\text{target}}, I, m$)

Input: Parameterised algorithm \mathcal{A} , two parameter configurations of \mathcal{A} , θ_{source} and θ_{target} , benchmark instance set I , performance metric m

Output: An ordered list $(\theta_0, \theta_1, \theta_2, \theta_3, \dots, \theta_l)$ of configurations of \mathcal{A} chosen during each round of ablation. $\theta_0 = \theta_{\text{source}}$ and $\theta_l = \theta_{\text{target}}$

```

 $\theta \leftarrow \theta_{\text{source}}$ 
activeParameters  $\leftarrow$  set of parameters of  $\mathcal{A}$  with different values in  $\theta_{\text{source}}$  and  $\theta_{\text{target}}$ 
ablationRoundsBest  $\leftarrow (\theta_{\text{source}})$ 
while activeParameters  $\neq \emptyset$  do
     $\mathcal{A}' \leftarrow$  set of algorithms with configurations obtained from  $\theta$  via flipping 1
    parameter in activeParameters to the value in  $\theta_{\text{target}}$ , ignoring configurations that
    are prohibited in the configuration space or that are equal to  $\theta$  due to parameter
    conditionality.
     $\theta' \leftarrow \text{determine\_best}(\mathcal{A}', I, m)$ 
     $\theta \leftarrow \theta'$ 
    activeParameters  $\leftarrow$  set of parameters of  $\mathcal{A}$  with different values in  $\theta$  and  $\theta_{\text{target}}$ 
    append (ablationRoundsBest,  $\theta'$ )
end
return ablationRoundsBest

```

1997; Glover and Laguna, 1997). Search strategies based on path relinking typically maintain a population of solutions and create new candidate solutions by constructing paths in the search space between existing solutions in the population. Each point in the search space lying on these paths is a potential solution, often with properties similar to the start or end points of the path. For further information on the history and application of path relinking, we refer the interested reader to two surveys by Glover et al (2000, 2003).

Our ablation analysis approach can be seen as an application of path relinking in the configuration space of a given parameterised algorithm, as we are constructing specific paths between a source and target configuration (as such, to the best of our knowledge, we are describing the first application of ideas from path relinking in the context of algorithm configuration). However, ablation analysis pursues a different goal from path relinking: While the latter is used to identify candidate configurations during search, the former aims to identify individual parameter settings responsible for the performance of a given target algorithm. As we demonstrate in Section 6, ablation analysis can also be used to find configurations that generalise well to different types of problem instances solvable by a given target algorithm, and this application of ablation analysis is even more closely related to path relinking.

3 Ablation Analysis

Given a parameterised algorithm \mathcal{A} with d parameters and configuration space Θ , along with a source and target configuration ($\theta_{\text{source}}, \theta_{\text{target}} \in \Theta$) of that algorithm, our ablation procedure works as follows. Given a set of benchmark instances I and a performance metric m (e.g., penalised average runtime or mean solution quality), we first compute the set of parameters whose val-

ues differ between θ_{source} and θ_{target} . Then, beginning from θ_{source} , we proceed through a series of rounds: in each round, we use a subprocedure `determine_best` to choose a configuration from the set of all configurations obtained by flipping one parameter in the current configuration to its value in θ_{target} . Algorithm 1 further outlines the details of this procedure.

In each round of ablation, *i.e.*, in each iteration of the while-loop in Algorithm 1, the procedure `determine_best` (\mathcal{A}', I, m) selects the configuration in \mathcal{A}' with the best performance on I w.r.t. m . In the case where the source configuration has better performance on I than the target configuration, each configuration selected by `determine_best` (\mathcal{A}', I, m) will be the one with *minimum loss* compared to the configuration θ from the previous round. Conversely, when θ_{target} has better performance than θ_{source} on I , the configuration selected by `determine_best` will be that with *maximum gain* over the previous θ . Some parameterised algorithms have conditional parameters, *i.e.*, parameters that only exist (or whose values only affect algorithm performance) if one or more other parameters (parents) are set to specific values. Configurations obtained by modifying the values of inactive conditional parameters are ignored in our procedure, as these configurations are by definition identical to the configuration from which they were produced. This means that modification of conditional parameters may be delayed until late in the ablation process if the corresponding parent parameters have little effect on algorithm performance. We discuss this issue further in Sections 5 and 7.

In the experiments we present in Section 4, we perform ablation analysis in both directions for every pair of configurations. By performing ablation in the direction of minimum loss, we can gauge the relative extent (by number of parameter modifications) of the local area around θ_{source} with roughly equal performance. In the direction of maximum gain, we find the minimal number of parameter modifications required to achieve roughly equal performance to θ_{target} . As a greedy approach (not unlike forward selection), ablation in either direction may produce suboptimal results at any distance except 1 from θ_{source} . In light of this, performing the analysis in two directions provides additional robustness. In the following, we describe two variants of `determine_best` (\mathcal{A}', I, m): a naïve brute-force method, which is easy to implement but slow, and a greatly accelerated version based on a racing method.

3.1 Brute-Force Ablation

Our brute-force implementation of `determine_best` (\mathcal{A}', I, m) involves performing a full empirical performance evaluation for every configuration in \mathcal{A}' , by running each configuration in \mathcal{A}' on every instance in I and recording the value of the performance metric m thus obtained. The configuration in \mathcal{A}' with the best performance according to m is selected and returned by `determine_best`.

Given that one parameter is eliminated from consideration in every round, ablation on instance set I with p differing parameters between θ_{source} and θ_{target} using this brute-force approach will require up to $|I| \cdot p \cdot (p + 1) / 2$ individual

runs of algorithm \mathcal{A} . (In the presence of conditional parameters, more than one parameter can be eliminated in one round of ablation if an inactive parameter is set to its default value in the target configuration.) Therefore, this procedure can be extremely time-consuming in the presence of high runtime cutoffs or large instance sets. Consider a typical case of ablation between a source and target configuration with 25 differing parameters, and an instance set I with 1 000 benchmark instances. Over the course of ablation using the brute-force method, 325 000 algorithm runs will be performed. Even with a mean CPU time of only 30 seconds per run of \mathcal{A} for any instance from I for all configurations considered in the analysis, this implies an overall runtime requirement of 9 750 000 CPU seconds or 112 CPU days. We note that while, by parallelizing runs across a cluster of machines (as we do in our experiments), this does not necessarily render ablation using this method completely impractical, it represents a formidable computational burden. Clearly, a more efficient ablation procedure would be highly desirable.

3.2 Acceleration via Racing

Based on early work for solving the model selection problem in memory-based supervised learning (Maron and Moore, 1994), F-Race is a prominent racing method for algorithm selection (Birattari et al, 2002). Given a benchmark instance set and performance metric, F-Race takes a set of candidate algorithms (or configurations of a parameterised algorithm) and iterates between gathering performance data by running the candidate algorithms on benchmark instances, and eliminating candidates once there is enough statistical evidence to justify removing them. The algorithms remaining at the end of the procedure are the winners of the race.

We apply F-Race to determine the best configurations in each round of ablation analysis, adhering very closely to the statistical framework described by Birattari et al (2002). In this context, F-Race starts with a set of candidate configurations containing all configurations in \mathcal{A}' and subsequently performs a sequence of stages. In stage k , the remaining candidate configurations $C = (c_1, c_2, \dots, c_n)$ are evaluated on a new instance $i_k \in I$, and the results are then combined with the results of the previous stages for each configuration. These results are then organised into k blocks, with the j^{th} block containing the n performance metric values resulting from running the configurations in C on instance i_j .

On these blocks, a Friedman two-way analysis of variance by ranks, also known as the Friedman test, is performed (Conover, 1999). If the null hypothesis of this test is rejected, we can conclude that at least one configuration in C has statistically significantly better performance than at least one other configuration. In this case, we proceed to pairwise testing to identify which configurations should be removed from the candidate list C . We use the same pairwise test here as described by Birattari *et al.* for F-Race, by comparing the configuration with the best sum of ranks across all blocks with the other $n - 1$

configurations in C using a modified t -test with $n - 1$ degrees of freedom (Birattari et al, 2002). After culling any configurations deemed to be statistically significantly worse, we proceed to the next round of the race. The race terminates when only one configuration remains, or when a specified maximum number of rounds have been performed. In the latter case, the configuration with the best mean performance (according to m) across all rounds is selected as the winner; in the case of further ties, tie-breaking is performed uniformly at random. (Further details on the statistical tests used in this procedure can be found in Birattari et al (2002); Conover (1999)).

Table 1 Results from performing ablation analysis on two scenarios with three ablation variants: the brute force approach, racing with the maximum number of rounds equal to the number of instances, and racing with the maximum number of rounds equal to 200. The SPEAR and CPLEX scenarios themselves are identical to those described in detail in Section 4. For both racing rows, we note that a full empirical analysis was performed on the chosen benchmark instances after each racing round in order to calculate the performance of every configuration on the ablation path. Run counts and CPU times for racing alone are approximately half of the reported values.

Scenario	SPEAR SWV		CPLEX CORLAT	
	runs	CPU time (s)	runs	CPU time (s)
Brute force	70 366	472 686.17	212 000	2 801 674.25
Racing (max #rounds = $ I $)	43 982	107 418.36	115 665	921 487.05
Racing (max #rounds = 200)	28 544	108 036.14	51 349	399 092.70

To measure the speed-ups of this racing approach to ablation analysis compared to the brute-force approach described earlier, we performed experiments using two of the scenarios described in Section 4. Using our SPEAR and CPLEX scenarios, we performed ablation analysis from the solver default to the configuration obtained from PARAMILS using three ablation analysis approaches. The brute-force approach required approximately 5 CPU days for the SPEAR scenario, and 32 CPU days for the CPLEX scenario. Using our racing approach with the maximum number of rounds set to the size of the benchmark set I reduced these requirements to just over 1 CPU day for SPEAR, and 10.7 CPU days for CPLEX. By limiting the maximum number of racing rounds to 200 (rather than $|I| = 302$ for SPEAR and 1000 for CPLEX), we achieved further reductions to 4.6 CPU days for CPLEX. We also note that every algorithm run inside a single stage of our racing approach can be performed in parallel, resulting in further reductions in the wall-clock time required. Table 1 gives a summary of these performance gains.

Determining the optimal value for the maximum number of racing rounds parameter is not straightforward, but in general it can be expected to depend on the number of instances in the given set of benchmark instances for which algorithm performance is highly correlated. We determined our conservative choice of 200 by subsampling runs from the full instance sets and choosing the lowest value that did not change the distribution of runtime over the resulting set in any substantial way for any of our scenarios.

Table 2 Depots training and test set performance obtained when evaluating every parameter configuration in the 1-neighbourhood of the LPG default configuration.

parameter modified	training set performance	test set performance
default	30.499	22.140
cri_intermediate_levels	1.548	1.574
walkplan	14.923	15.633
triomemory	29.124	26.216
tabu_fct	29.225	23.607
tabu_length	29.293	24.958
donot_try_suspected_actions	30.497	24.798
dynoiseTabLen	30.587	24.937
dynoisecoefnum	30.622	23.588
tabu_act	30.624	24.959
no_cut	30.642	26.323
vicinato	31.964	22.238
maxnoise	33.387	24.963
inc_re	38.475	31.776
numrestart	39.073	29.187
static_noise	45.005	39.559
extended_effects_evaluation	49.412	51.042
numtry	70.004	51.778
weight_mutex_in_relaxed_plan	71.346	61.178
noise	132.844	121.984
noise_incr	175.963	148.953
ri_list	402.606	371.674
depots configured	0.642	0.637

While it is possible that racing will return different ablation results than the brute-force approach, we do not consider this to be a problem, since the brute-force approach is also not guaranteed to compute the optimal path between two configurations. Furthermore, in all of our experiments, the brute-force and racing results were closely aligned. We additionally tested how close our racing approach comes to the optimal parameter choice, by performing complete evaluations of the 1-neighbourhood of configurations around the default for each of our five scenarios. These neighbourhoods consist of those configurations obtained by modifying a single parameter from its value in the source configuration to its value in the target configuration. In all cases, the most important parameter selected by our racing approach for each scenario is the same as the best parameter in the 1-neighbourhood around the source configuration. The results for the LPG solver on our depots scenario are given in Table 2.

3.3 Ablation Paths

We call the path of configurations $(\theta_0, \dots, \theta_l)$ obtained between θ_{source} and θ_{target} computed by our ablation procedure along with the respective performance values on set I (or an independent test set of instances similar to those in I) an *ablation path*. These paths can take several qualitatively different forms, depending on the relative performance of θ_{source} and θ_{target} , and on characteristics of the response surface that captures the functional dependency of the performance of \mathcal{A} on its parameter settings. Figure 1 illustrates

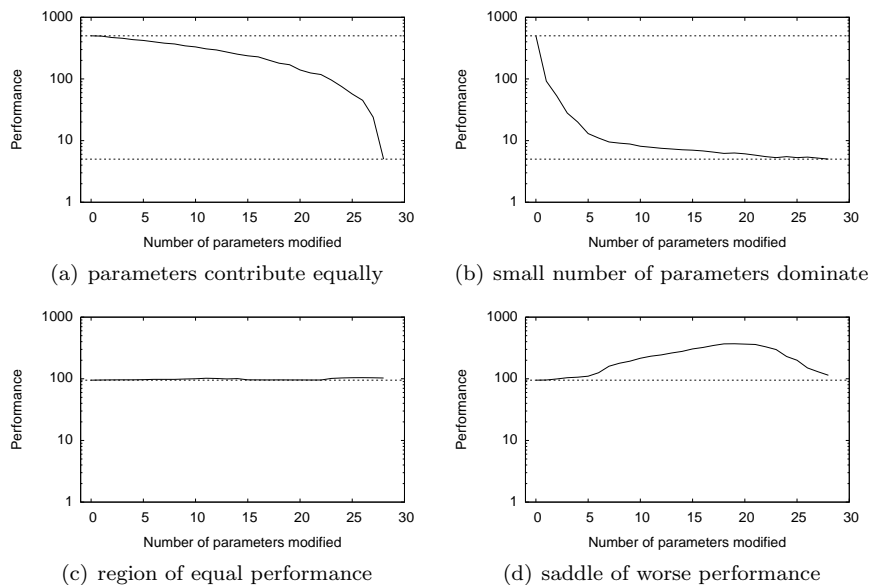


Fig. 1 Ablation paths can take several qualitatively different forms, illustrated by these (idealised) runtime examples. Lower values indicate better performance (for details, see text).

these cases; each point represents the performance of one of the configurations θ_i , from θ_{source} on the left-hand side to θ_{target} on the right. Figure 1(a) illustrates one extreme case, where θ_{source} and θ_{target} differ in performance on I and all parameters are of equal importance. A case at the opposite extreme (not shown) would be if the performance difference between θ_{source} and θ_{target} was fully explained by the modification of a single parameter.

A more realistic case lies between these extremes, with the modification of a small number of parameters explaining most of the difference in performance between θ_{source} and θ_{target} (Figure 1(b)). Figures 1(c) and 1(d) show two additional cases that may occur when the source and target configurations have roughly equal performance on I . In 1(c), θ_{source} and θ_{target} are connected by a path of configurations that all have the same performance; this could arise in a situation where both lie on a large plateau of the response surface. However, it is also possible that the two configurations lie in separate basins of the response surface, such that a “saddle” of worse performance must be surmounted along the ablation path from θ_{source} to θ_{target} , as illustrated in 1(d).² We note that the results from all of our experiments performing ablation from algorithm defaults to configurations obtained from automated configurators fall into case 1(b).

² The ablation path illustrated in Figures 1(d) resembles the folding pathways observed in biopolymers like RNA that have to overcome a thermodynamic barrier in order to change from one low-energy structure to another (Reidys, 2011)

4 Experiment design

In order to empirically evaluate our ablation methods, we performed experiments on five scenarios using state-of-the-art solvers for SAT, MIP and AI planning. We implemented the brute-force and racing-based ablation methods as plugins for HAL, a Java-based platform for distributed experiment execution and data management (Nell et al, 2011). All runs were performed using machines in the Compute-Calcul Canada / Westgrid Orcinus cluster, each equipped with two Intel Xeon X5650 6-core 2.66Ghz processors, 12MB of cache and 24GB of RAM, running CentOS 5 Linux. For each target algorithm run, we used a single core and enforced a maximum of 2GB (SPEAR and CPLEX) or 6GB (LPG) of RAM.

Using the existing PARAMILS plugin for HAL 1.1.6, we performed 10 independent runs of PARAMILS for each scenario, with each configurator run allocated 48 CPU hours of total runtime. In each case, we minimised penalised average runtime (PAR10), a standard performance metric for configuration and empirical analysis; under PAR10, each run that terminates successfully was assigned a score equal to the CPU time used, and runs that crash or do not produce a valid solution on a given instance were assigned a score of 10 times the runtime cutoff (this heavily penalised these cases to attempt to enforce good instance set coverage.) Of the 10 configurations produced in our PARAMILS runs, we selected the one with the best PAR10 performance on the full training set for that scenario. (This corresponds to one of the standard protocols for using PARAMILS.)

We then performed two ablation experiments using our racing-based ablation procedure on the training set for each scenario, with the maximum number of rounds set to 200. Ablation was performed in two directions, first from the default to the optimised configuration obtained through automated configuration using maximum gain, and then from the optimised configuration to the default using minimum loss. Each configuration on the resulting ablation paths for each scenario was subsequently evaluated using the independent test set for that scenario. (Performing ablation directly on test sets produced very similar results.)

SAT using SPEAR. The propositional satisfiability problem, or SAT, is the prototypical *NP*-hard problem with important real-world application, including circuit design as well as hardware and software verification. SAT has also been widely studied in the context of automated algorithm configuration Hutter et al (2007b, 2009, 2011). We chose to analyze the industrial SAT solver SPEAR 1.2.1, winner of one category of the 2007 Satisfiability Modulo Theories Competition (Hutter et al, 2007a); SPEAR has 26 configurable parameters, creating a space of 8.34×10^{17} configurations. SPEAR has also been used in two recent investigations of parameter importance using forward selection (Hutter et al, 2013) and functional ANOVA (Hutter et al, 2014). We analyzed the performance of SPEAR on the SWV software verification instance set used in several previous investigations. This set, consisting of 604 software verification condi-

tions produced by an automated static checker, is partitioned into a training set (used for configuration and ablation analysis) and test set (used for evaluation of the ablation paths) consisting of 302 instances each. Following previous work, we used a 300 CPU-second runtime cutoff for automated configuration and all analysis runs.

MIP using CPLEX. Mixed integer programming (MIP) is another widely-studied problem with many prominent real-world applications. IBM ILOG CPLEX is one of the most widely used MIP solvers, both in academia and industry, and has a highly-parameterised configuration space containing 76 configurable parameters that directly impact solver performance (a total of 1.90×10^{47} configurations). Automated configuration of CPLEX has proven successful in past work (Hutter et al, 2010, 2011), and CPLEX has also been used in the same parameter importance investigations as mentioned for SPEAR. We chose to use CPLEX 12.1 and the CORLAT instance set for this scenario (Gomes et al, 2008); CORLAT is a set of computational sustainability MIP instances based on real data used for wildlife corridor construction for grizzly bears in the Northern Rockies region. This set has been used in previous work on algorithm configuration and on parameter importance; it is partitioned into a training and test set containing 1 000 instances each. A 300 CPU-second runtime cutoff was used for all runs.

AI Planning using LPG. The design and automated configuration of highly-parameterised solvers has recently proven successful in the AI planning community, contributing to both the winner and runner-up in the Learning Track of the 7th International Planning Competition (IPC-2011) (Vallati et al, 2011, 2013). Highly-parameterised general-purpose planners represent ideal candidate scenarios for studying parameter importance, because intuitively, the benefits to be gained by exploiting the structure and differences between various planning domains suggest that high-performance configurations will vary widely between such domains. We chose to investigate the configuration space of LPG td-1.0, a state-of-the-art local search based planner, and a key component in the winner of the IPC-2011 Learning Track. LPG has 66 configurable parameters, with a total of 9.11×10^{36} possible configurations. We analyzed LPG’s performance on three planning domains: depots, satellite, and zenotrail. These three domains have been used in previous planning competitions, as well as in previous work on automated configuration for planning. Each instance set contains disjoint 2 000-instance training and test sets generated using the same parameter settings of a randomised instance generator. Consistent with previous work, a 60 CPU-second runtime cutoff was used for configuration, while a 300 CPU-second cutoff was used for all test-set evaluation and ablation analysis runs.

5 Results

Table 3 and Table 4 show the training and test set performance for the default configurations and automatically optimised configurations in all five scenarios

Table 3 Training set performance results for all 5 of our scenarios, for both the default configurations and those obtained from PARAMILS. Runtime cutoffs in all cases were 300 CPU seconds.

solver	instance set	Training set performance (PAR10, s)			
		q25	q50	q75	mean
SPEAR default	SWV	0.122	0.528	23.649	573.649
SPEAR configured	SWV	0.122	0.592	1.279	1.359
CPLEX default	CORLAT	0.101	3.563	90.596	556.531
CPLEX configured	CORLAT	0.110	1.220	5.812	5.511
LPG default	depots	0.551	1.086	8.182	43.245
LPG configured	depots	0.220	0.318	0.510	0.671
LPG default	satellite	15.232	17.580	20.595	17.962
LPG configured	satellite	4.827	5.645	6.404	5.662
LPG default	zenotravel	20.092	26.377	34.642	29.671
LPG configured	zenotravel	1.414	1.826	2.490	2.065

Table 4 Test set performance results for all 5 of our scenarios, for both the default configurations and those obtained from PARAMILS. Runtime cutoffs in all cases were 300 CPU seconds.

solver	instance set	Test set performance (PAR10, s)			
		q25	q50	q75	mean
SPEAR default	SWV	0.102	0.499	11.392	569.645
SPEAR configured	SWV	0.079	0.531	1.114	1.321
CPLEX default	CORLAT	0.097	3.551	70.602	471.722
CPLEX configured	CORLAT	0.112	1.238	5.650	5.411
LPG default	depots	0.535	1.055	7.194	38.097
LPG configured	depots	0.220	0.324	0.511	0.658
LPG default	satellite	15.173	17.575	20.514	17.940
LPG configured	satellite	4.943	5.760	6.529	5.783
LPG default	zenotravel	19.792	26.026	34.929	29.361
LPG configured	zenotravel	1.407	1.841	2.556	2.092

considered; as expected, and consistent with previously published results for these solvers, we observed 3- to 422-fold speedups after configuration.

Table 5 Design space size and number of parameters changed from the default configuration by PARAMILS for all 5 scenarios. Note that the number of parameters changed from the default represents an upper bound on the number of ablation rounds, as conditional parameters can cause fewer rounds to be required if they are inactive in the source configuration and active with their default values in the target configuration.

solver	instance set	design space size	# parameters	# changed from default
SPEAR	SWV	8.34×10^{17}	26	21
CPLEX	CORLAT	1.90×10^{47}	76	20
LPG	depots	9.11×10^{36}	66	22
LPG	satellite	9.11×10^{36}	66	35
LPG	zenotravel	9.11×10^{36}	66	33

Table 5 details the design space size for each of the three solvers used in our empirical analysis, as well as the number of parameters that were changed from the default in each of our five scenarios. Interestingly, nearly every SPEAR parameter was changed from the default, while for the CPLEX and LPG scenarios, approximately one-third to one-half of the parameters were modified.

Table 6 Parameters chosen and resulting PAR10 performance on both the training and test sets for ablation analysis accelerated by racing (maximum of 200 rounds) for SPEAR on SWV training set. Ablation was performed using the SPEAR default configuration as the source, and the configuration produced by PARAMILS as the target.

round	parameter	performance	
		training	test
default	-	573.649	569.645
1	sp-var-dec-heur	2.041	1.800
2	sp-res-cutoff-cls	2.341	1.930
3	sp-rand-var-dec-scaling	1.554	1.287
4	sp-use-pure-literal-rule	1.350	1.276
5	sp-rand-var-dec-freq	1.365	1.288
6	sp-var-activity-inc	1.428	1.325
7	sp-clause-decay	1.413	1.310
8	sp-learned-size-factor	1.420	1.302
9	sp-res-cutoff-lits	1.432	1.319
10	sp-rand-phase-dec-freq	1.405	1.308
11	sp-clause-activity-inc	1.412	1.302
12	sp-orig-clause-sort-heur	1.418	1.341
13	sp-learned-clauses-inc	1.442	1.311
14	sp-max-res-runs	1.404	1.331
15	sp-learned-clause-sort-heur	1.352	2.018
16	sp-first-restart	1.273	1.214
17	sp-restart-inc	1.476	1.452
18	sp-res-order-heur	1.488	1.327
19	sp-variable-decay	1.417	1.377
20	sp-phase-dec-heur	12.038	1.602
21	sp-max-res-lit-inc	1.359	1.321
configured	-	1.359	1.321

Table 7 Parameters chosen and resulting PAR10 performance on both the training and test sets for ablation analysis accelerated by racing (maximum of 200 rounds) for CPLEX on the CORLAT training set. Ablation was performed using the CPLEX default configuration as the source, and the configuration produced by PARAMILS as the target.

round	parameter	performance	
		training	test
default	-	556.531	471.722
1	mip_cuts_covers	90.776	63.067
2	mip_strategy_heuristicfreq	16.599	21.385
3	simplex_dgradient	8.080	7.368
4	simplex_tolerances_markowitz	10.345	12.763
5	mip_limits_aggforcut	6.900	6.120
6	mip_strategy_variableselect	8.974	8.012
7	simplex_pgradient	5.799	8.037
8	mip_strategy_fpheur	8.699	5.155
9	lpmethod	8.700	5.151
10	barrier_crossover	8.662	5.146
11	preprocessing_symmetry	8.696	5.140
12	sifting_algorithm	8.695	5.145
13	barrier_limits_growth	8.678	5.137
14	mip_limits_gomorycand	8.704	5.143
15	mip_limits_cutsfactor	5.639	5.115
16	mip_cuts_gubcovers	5.672	5.132
17	mip_cuts_gomory	5.447	5.236
18	preprocessing_repeatpresolve	5.608	5.433
19	mip_strategy_presolvenode	5.766	5.288
20	mip_cuts_mircut	5.511	5.411
configured	-	5.511	5.411

Table 8 Parameters chosen and resulting PAR10 performance on both the training and test sets for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the depots training set. Ablation was performed using the LPG default configuration as the source, and the configuration produced by PARAMILS as the target.

round	parameter	performance	
		training	test
default	-	43.245	38.097
1	cri_intermediate_levels	1.610	1.520
2	triomemory	1.482	1.370
3	donot_try_suspected_actions	1.460	1.345
4	walkplan	1.310	1.162
5	weight_mutex_in_relaxed_plan	1.069	1.012
6	noise	1.143	1.033
7	static_noise	1.124	1.067
8	ri_list	1.132	1.069
9	numrestart	1.117	1.062
10	maxnoise	1.114	1.062
11	inc_re	1.086	1.008
12	extended_effects_evaluation	1.167	1.096
13	numtry	1.035	1.013
14	no_cut	1.035	1.015
15	vicinato	1.038	1.014
16	hpar_cut_neighb	0.671	0.658
configured	-	0.671	0.658

Tables 6 and 7 detail the full ablation paths using up to 200 rounds of racing, for the SPEAR and CPLEX scenarios. Tables 8, 9 and 10 detail the ablation paths for the three LPG scenarios. In all cases, we present the ablation paths using the default configuration as the source and the PARAMILS configuration as the target. The ablation paths in the reverse direction are qualitatively similar, with small deviations due to parameter conditionality in the case of our CPLEX and LPG scenarios.

Figure 2(a) illustrates the mean PAR10 score on the SWV test set for every configuration along the path found through racing-accelerated ablation analysis of SPEAR on the SWV training set. Expressing the performance gain from a single ablation round as a percentage of the total gain between θ_{source} and θ_{target} , 99.92% of the performance gain between the default configuration and the optimised configuration can be achieved by modifying the value of a single parameter, *sp-var-dec-heur*. This parameter controls the choice of the variable decision heuristic used by SPEAR, which is known to be an important parameter in most state-of-the-art SAT solvers. Furthermore, if we modify only four parameters (*sp-var-dec-heur*, *sp-rand-var-dec-scaling*, *sp-res-cutoff-cls*, and *sp-first-restart*) from their default values, we obtain a configuration with slightly *better* performance on the test set than the target configuration obtained with PARAMILS. In contrast, Hutter et al (2014) noted in their functional ANOVA work that *sp-var-dec-heur* was important, but only 76% of the improvement over the default could be attributed to single-parameter effects in their model. We hypothesize that *sp-var-dec-heur* is much more important in high-performance parts of the SPEAR configuration space, a bias that is not taken into account by the empirical performance models used by Hutter *et al.*

Table 9 Parameters chosen and resulting PAR10 performance on both the training and test sets for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the satellite training set. Ablation was performed using the LPG default configuration as the source, and the configuration produced by PARAMILS as the target.

round	parameter	performance	
		training	test
default	-	17.962	17.940
1	cri_intermediate_levels	7.316	7.577
2	criprecond	6.967	7.141
3	noise	7.029	6.997
4	no_cut	6.925	6.803
5	improve_reachability	6.943	6.837
6	donot_try_suspected_actions	6.834	6.793
7	numrestart	6.915	6.776
8	dynoisecoefnum	6.859	6.787
9	cri_update_iterations	6.856	6.795
10	vicinato	6.887	6.788
11	hpar_cut_neighb	5.357	5.317
12	triomemory	5.357	5.325
13	verifyinit	5.337	5.329
14	noise_incr	5.356	5.351
15	maxnoise	5.471	5.343
16	evaluate_mutex_for_action_remotion	5.360	5.325
17	walkplan	5.363	5.340
18	evaluation_function	5.217	5.339
19	dynoiseTabLen	5.239	5.341
20	static_noise	5.314	5.328
21	relaxed_examination	5.268	5.314
22	reset_extended_unsupported_facts	5.413	5.331
23	mutex_and_additive_effects	5.423	5.344
24	numtry	5.460	5.339
25	not_supported_preconds_evaluation	5.355	5.335
26	inc_re	5.379	5.327
27	zero_num_A	5.331	5.335
28	not_extended_unsupported_facts	5.372	5.344
29	fast_best_action_evaluation	5.371	5.341
30	no_mutex_with_additive_effects	5.375	5.331
31	lagrange	5.376	5.334
32	h	5.662	5.783
configured	-	5.662	5.783

Similarly, Figure 2(b) shows the performance of configurations on the path found through racing-accelerated ablation analysis of CPLEX on the CORLAT training set, evaluated on the test set. Here, 87.64% of the performance gain resulted from modifying the value of a single parameter, *mip_cuts_covers*, which controls whether or not to generate cover cuts. 99.58% of the gain can be achieved by modifying just three CPLEX parameters (*mip_cuts_covers*, *mip_strategy_heuristicfreq* and *simplex_dgradient*). We note that *simplex_dgradient* was not in the top 10 important parameters identified for this scenario by Hutter et al (2014), although 6 of the 10 most important CPLEX parameters as identified by their functional ANOVA approach were not changed from their default values in our experiments (this effect was also noted by Hutter *et al.*).

Finally, Figures 2(c), 2(d) and 2(e) illustrate the performance along the ablation paths for each of the three LPG scenarios: depots, satellite and zenotransport. For the depots and satellite scenarios, the top three parameters were the same. Modifying the value of *cri_intermediate_levels* resulted in 97.7% and

Table 10 Parameters chosen and resulting PAR10 performance on both the training and test sets for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the zenotravel training set. Ablation was performed using the LPG default configuration as the source, and the configuration produced by PARAMLS as the target.

round	parameter	performance	
		training	test
default	-	29.671	29.361
1	fast_best_action_evaluation	8.064	8.258
2	triomemory	5.789	5.912
3	noise	4.266	4.369
4	criprecond	3.701	3.713
5	stop_remove_act	3.458	3.446
6	verifyinit	3.418	3.435
7	weight_mutex_in_relaxed_plan	3.393	3.421
8	reset_extended_unsupported_facts	3.377	3.401
9	inc_re	3.348	3.387
10	twalkplan	3.347	3.392
11	choose_min_numA_fact	3.357	3.391
12	numrestart	3.375	3.393
13	lagrange	3.376	3.391
14	verifyAf	3.371	3.392
15	no_cut	3.368	3.397
16	zero_num_A	3.362	3.393
17	noise_incr	3.333	3.388
18	numtry	3.401	3.389
19	bestfirst	3.372	3.389
20	evaluation_function	3.343	3.388
21	no_pruning	3.367	3.391
22	h	3.385	3.392
23	donot_try_suspected_actions	3.367	3.388
24	cri_insertion_add_mutex	3.371	3.387
25	vicinato	3.370	3.381
26	hpar_cut_neighb	2.035	2.047
27	nonuniform_random	2.018	2.056
28	dynoiseTabLen	1.996	2.053
29	no_insert_threated_act_in_neighb	2.005	2.052
30	maxnoise	2.065	2.092
configured	-	2.065	2.092

84.55% of the target configuration performance over the default, respectively. Furthermore, modifying the values of three parameters (*cri_intermediate_levels*, *vicinato* (the neighbourhood choice), and *hpar_cut_neighb*) resulted in 99.22% of the performance gain for the depots scenario.

For the zenotravel scenario (Figure 2e), we observed different choices for the two most important parameters, depending on the direction in which ablation was performed. Modifying *triomemory* and *fast_best_action_evaluation* from their default values resulted in 85.99% of the overall performance gain over the default, while modifying *vicinato* and *hpar_cut_neighb* (similar to the other two LPG scenarios) resulted in 88.09% of the total performance gain. Four parameter modifications (*vicinato*, *hpar_cut_neighb*, *triomemory*, and *noise*) accounted for 97.8% of the total performance gain.

It is interesting to note that there is a conditional parameter interaction between *hpar_cut_neighb* and *vicinato*, as *hpar_cut_neighb* is only active when *vicinato* takes certain values. In our experiments, modifying *vicinato* often did not produce large gains in performance by itself, but allowed for modification of *hpar_cut_neighb*, which in turn resulted in large performance improvements.

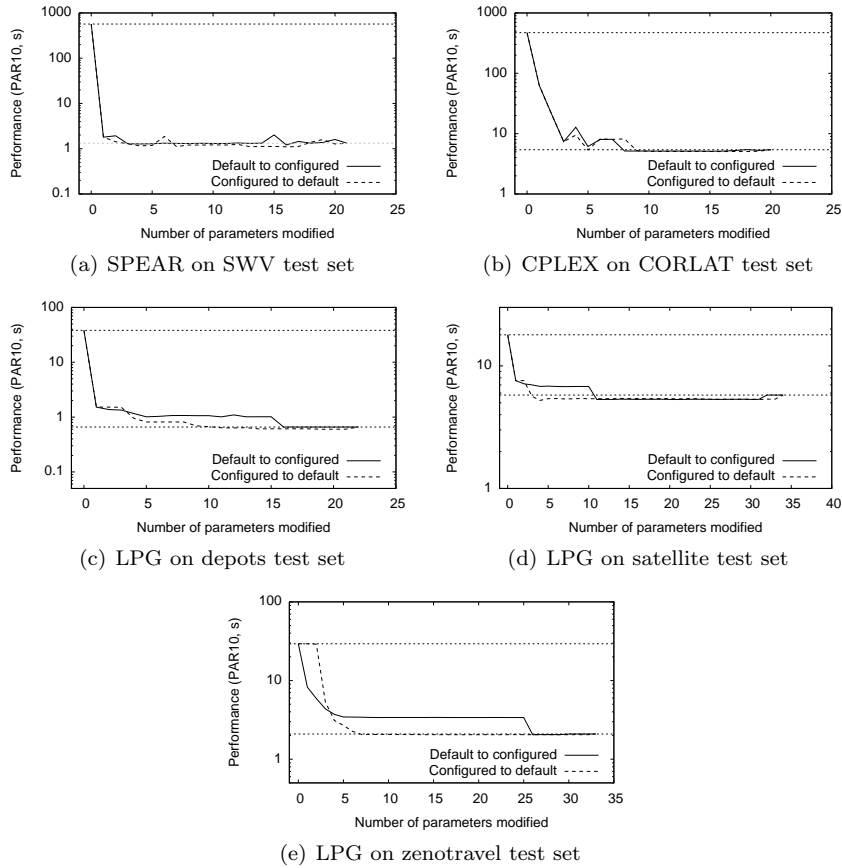


Fig. 2 Ablation paths determined using racing with up to 200 rounds on the training sets of the five configuration scenarios, with each configuration on the ablation path evaluated using the corresponding test set. The horizontal lines indicate the PAR10 scores of the default (source) and automatically-configured (target) configurations on the test set for each scenario.

The effect of conditional parameters can also be seen in the “late” performance improvements in the three LPG scenarios in the direction of maximum gain.

6 Additional experimental results

So far, we have focussed on results from straightforward applications of ablation analysis. However, ablation analysis can also be used to find configurations that generalise better than the default configuration to new domains (*i.e.*, problem instance classes) of interest, or to obtain information about the topology of the algorithm configuration space around and between high-

Table 11 Performance analysis of the LPG default, configurations obtained from PARAMILS, and the first 4 flips of ablation, for all 3 LPG scenarios, on the depots test set.

configuration	depots test set performance (PAR10)
default	38.097
depots ablation flip 1	1.520
depots ablation flip 2	1.390
depots ablation flip 3	1.261
depots ablation flip 4	1.037
depots configured	0.658
satellite ablation flip 1	10.413
satellite ablation flip 2	6.194
satellite ablation flip 3	224.976
satellite ablation flip 4	223.590
satellite configured	2643.840
zenotravel ablation flip 1	11.983
zenotravel ablation flip 2	11.790
zenotravel ablation flip 3	368.500
zenotravel ablation flip 4	1720.010
zenotravel configured	2807.531

Table 12 Performance analysis of the LPG default, configurations obtained from PARAMILS, and the first 4 flips of ablation, for all 3 LPG scenarios, on the satellite test set.

configuration	satellite test set performance (PAR10)
default	17.940
depots ablation flip 1	7.275
depots ablation flip 2	7.260
depots ablation flip 3	7.279
depots ablation flip 4	7.234
depots configured	7.617
satellite ablation flip 1	7.577
satellite ablation flip 2	7.141
satellite ablation flip 3	6.997
satellite ablation flip 4	6.997
satellite configured	5.783
zenotravel ablation flip 1	17.358
zenotravel ablation flip 2	17.381
zenotravel ablation flip 3	16.323
zenotravel ablation flip 4	10.102
zenotravel configured	7.952

performance configurations. As mentioned in Section 2, these extended uses of ablation analysis strongly echo ideas from path relinking.

In several of these experiments, we focussed on our three LPG scenarios, since these allowed us to investigate generalisation of performance between different classes of instances (here: planning domains).

Generalisation to other domains. After spending hours of CPU time to produce an algorithm configuration that is highly-optimised for a specific problem instance set, it would be ideal if a user could somehow use that configuration to find a different configuration with better performance than the default on a different, previously unseen problem instance set. To investigate the potential for ablation analysis to produce such configurations, we performed three sets of experiments, in order to investigate how configurations “near” the LPG default (along ablation paths toward a chosen PARAMILS configuration) perform on

Table 13 Performance analysis of the LPG default, configurations obtained from PARAMILS, and the first 4 flips of ablation, for all 3 LPG scenarios, on the zenotravel test set.

configuration	zenotravel test set performance (PAR10)
default	29.361
depots ablation flip 1	1114.835
depots ablation flip 2	908.578
depots ablation flip 3	896.241
depots ablation flip 4	894.766
depots configured	1028.118
satellite ablation flip 1	1467.274
satellite ablation flip 2	1468.614
satellite ablation flip 3	870.385
satellite ablation flip 4	867.646
satellite configured	2478.914
zenotravel ablation flip 1	8.258
zenotravel ablation flip 2	5.912
zenotravel ablation flip 3	4.369
zenotravel ablation flip 4	3.713
zenotravel configured	2.092

scenarios other than the one used to determine the original ablation paths. Tables 11, 12 and 13 illustrate the results of these experiments. For depots and satellite, configurations that are a small number of ablation rounds away from the default perform substantially better than the default configuration itself (2.5- to 6.1-fold speedups), before the ablation path hits more specialised configurations that do not generalise as well to new domains. We do not see the same effect in the zenotravel domain, where the ablation path does improve in quality before getting worse, but fails to exceed at any point the performance of the default configuration.

Ablation paths between high-quality configurations. Next, we investigated whether ablation analysis could be used to better understand the topology of the search space between high-quality configurations for each of our 5 scenarios. If the configurations found on ablation paths between two high-quality configurations do not differ in performance from the source and target, we might expect the existence of a large plateau of high-quality configurations in the search space. Conversely, if the configurations lying on such ablation paths have worse performance than the source and target, we have reason to believe that the source and target may lie in different basins of the parameter response surface.

We examined the ablation paths between two high-quality configurations for each of our five scenarios, defined as the best and second-best configurations found by PARAMILS, evaluated on the training set. We used our racing approach, with the maximum number of rounds again set to 200. Tables 14, 15, 16, 17 and 18 detail the qualities of each configuration along the ablation paths for each scenario, while Figure 3 gives a visual summary of the paths.

We see that for SPEAR, there appears to be at least one ridge of poor-quality configurations that must be surmounted in order to move from the source to the target configuration. In all of the other four scenarios, no such ridge is present, and every configuration along the ablation paths share the quality of the source/target configurations. We note that as ablation analysis is a greedy

Table 14 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for SPEAR on SWV test set, using the best PARAMILS configuration as the source and the second-best PARAMILS configuration as the target.

round	parameter	performance
best	-	1.321
1	sp-res-cutoff-lits	1.899
2	sp-learned-clauses-inc	11.104
3	sp-max-res-lit-inc	1.734
4	sp-phase-dec-heur	1.316
5	sp-max-res-runs	1.310
6	sp-restart-inc	1.501
7	sp-variable-decay	1.404
8	sp-clause-decay	1.397
9	sp-rand-var-dec-freq	11.892
10	sp-learned-size-factor	11.934
11	sp-rand-phase-dec-freq	1.1474
12	sp-use-pure-literal-rule	1.504
13	sp-clause-activity-inc	1.455
14	sp-clause-del-heur	1.481
15	sp-update-dec-queue	1.488
16	sp-learned-clause-sort-heur	11.327
17	sp-res-cutoff-cls	11.390
18	sp-rand-var-dec-scaling	93.552
19	sp-first-restart	2.256
second-best	-	2.256

process, the occurrence of such ridges on the ablation paths does not strictly imply the existence of barriers that must be overcome by all paths between the source and target configurations. However, in the cases where no such ridges are present in the ablation paths, at least one path of equal performance between the source and target configurations is guaranteed to exist.

Ablation between high-quality configurations for different domains. Next, we investigated the use of ablation paths between high-performance configurations for different domains, to see how different configurations “needed” to be before high performance was achieved on a domain different from that for which they were originally optimised. (We note that this question is closely related to the problem of transfer learning.) We performed ablation analysis experiments between the chosen PARAMILS configurations for two of the LPG scenarios, evaluated using the test set from one of those scenarios. Tables 19 and 20 show the full ablation paths between the depots and satellite configurations, evaluated on the depots and satellite test sets respectively.

Similarly, Tables 21 and 22 show the results for ablation between the depots and zenotravel configurations, and Tables 23 and 24 show the results for ablation between the satellite and zenotravel configurations. Figure 4 presents these ablation paths visually.

We also performed additional ablation experiments between the configurations obtained from PARAMILS for each pair of LPG scenarios, evaluated using the test set of the third scenario. This set of experiments further tests the use of ablation paths to find configurations that generalise well to previously un-

Table 15 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for CPLEX on CORLAT test set, using the best PARAMILS configuration as the source and the second-best PARAMILS configuration as the target.

round	parameter	performance
best	-	5.411
1	preprocessing_repeatpresolve	5.117
2	mip_cuts_gomory	5.027
3	simplex_limits_singularity	5.014
4	simplex_pgradient	4.983
5	preprocessing_boundstrength	4.989
6	mip_limits_cutsfactor	5.026
7	mip_strategy_dive	4.749
8	barrier_crossover	4.753
9	lpmethod	4.765
10	mip_cuts_pathcut	4.751
11	preprocessing_symmetry	4.760
12	simplex_pricing	4.760
13	barrier_limits_growth	4.767
14	mip_cuts_zerohalfcut	4.747
15	mip_cuts_cliques	4.751
16	sifting_algorithm	4.745
17	mip_strategy_presolvenode	4.780
18	mip_strategy_variableselect	5.114
19	mip_limits_cutpasses	5.658
20	mip_strategy_probe	5.592
21	simplex_tolerances_markowitz	5.485
22	mip_cuts_flowcovers	6.286
23	emphasis_mip	6.017
24	mip_cuts_mircut	5.096
25	mip_limits_aggforcut	4.823
26	mip_cuts_covers	4.812
27	mip_strategy_branch	5.036
28	simplex_dgradient	5.024
29	mip_strategy_heuristicfreq	5.932
second-best	-	5.932

Table 16 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the depots test set, using the best PARAMILS configuration as the source and the second-best PARAMILS configuration as the target.

round	parameter	performance
best	-	0.658
1	extended_effects_evaluation	0.603
2	static_noise	0.614
3	numrestart	0.617
4	evaluation_function	0.612
5	triomemory	0.611
6	dynoiseTabLen	0.616
7	ri_list	0.617
8	cri_update_iterations	0.621
9	noise_incr	0.608
10	dynoisecoefnum	0.627
11	numtry	0.643
12	weight_mutex_in_relaxed_plan	0.583
13	walkplan	0.689
14	tabu_length	0.585
second-best	-	0.585

Table 17 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the satellite test set, using the best PARAMILS configuration as the source and the second-best PARAMILS configuration as the target.

round	parameter	performance
best	-	5.783
1	static_noise	5.293
2	cri_update.iterations	5.293
3	inc_re	5.239
4	fast_best_action_evaluation	5.256
5	evaluation_function	5.233
6	evaluate_threated_supported_preconds_of_neighb_action	5.234
7	weight_mutex_in_relaxed_plan	5.213
8	walkplan	5.257
9	choose_min_numA_fact	5.211
10	dynoiseTabLen	5.281
11	mutex_and_additive_effects	5.269
12	maxnoise	5.161
13	not_supported_preconds_evaluation	5.220
14	no_mutex_with_additive_effects	5.187
15	dynoisecoefnum	5.214
16	no_pruning	5.210
17	noise_incr	5.189
18	numtry	5.189
19	no_insert_threated_act_in_neighb	5.145
20	zero_num_A	5.144
21	avoid_best_action_cycles	5.182
22	noise	5.261
23	twalkplan	5.170
24	numrestart	5.198
25	ri_list	5.117
26	extended_effects_evaluation	5.157
27	not_extended_unsupported_facts	5.356
second-best	-	5.356

Table 18 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the zenotravel test set, using the best PARAMILS configuration as the source and the second-best PARAMILS configuration as the target.

round	parameter	performance
best	-	2.092
1	no_pruning	2.040
2	criprecond	2.019
3	choose_min_numA_fact	2.029
4	inc_re	2.022
5	remove_act_next_step	2.023
6	nonuniform_random	2.016
7	verifyAf	2.013
8	numtry	2.009
9	maxnoise	2.016
10	numrestart	2.006
11	no_insert_threated_act_in_neighb	2.019
12	bestfirst	2.009
13	dynoisecoefnum	2.013
14	dynoiseTabLen	2.019
15	verifyinit	2.025
16	noise_incr	2.007
17	cri_insertion_add_mutex	2.014
18	evaluation_function	2.014
19	extended_effects_evaluation	2.105
20	h	2.119
second-best	-	2.119

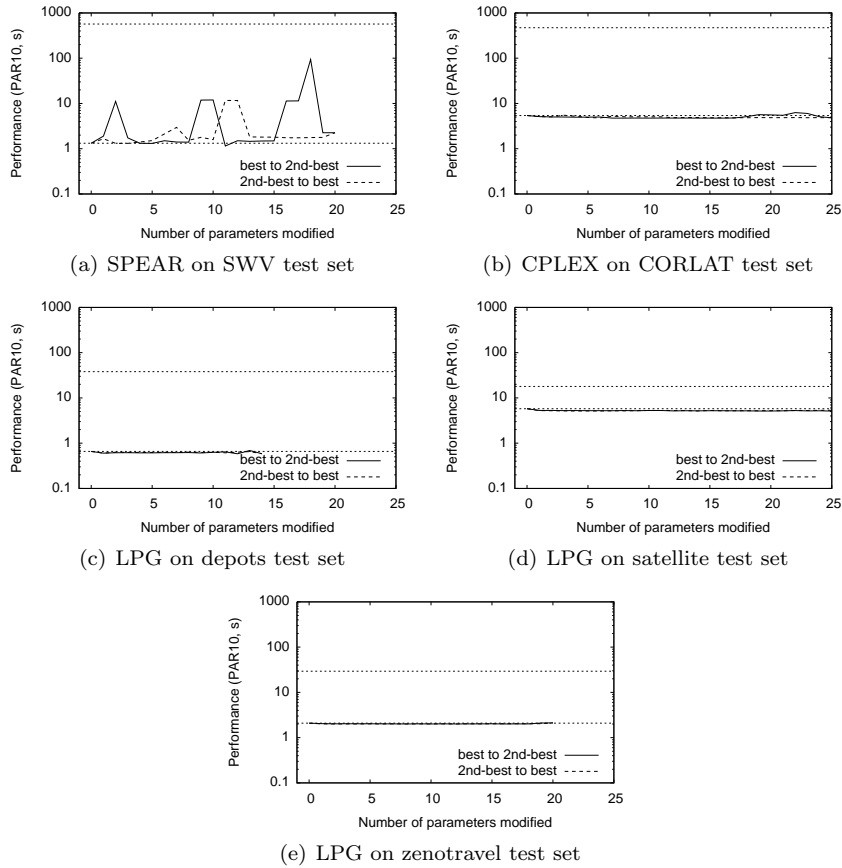


Fig. 3 Ablation paths determined using racing with up to 200 rounds on the test sets of the five configuration scenarios. Ablation was performed from the best PARAMILS incumbent to the second-best incumbent on each scenario. The horizontal lines indicate the PAR10 scores of the default (source) and automatically-configured (target) configurations.

seen problem instance sets. The results of these experiments are detailed in Tables 25, 26 and 27, and are shown visually in Figure 5.

For the ablation from the depots configuration to the satellite configuration, both the source and target configurations are significantly worse than the default on the zenotravel test set. However, there are configurations along the resulting ablation path with better performance than the default.

For the ablation from the depots configuration to the zenotravel configuration, both the source and target configurations are better than the default, but we again see improvement along the ablation path. In fact, many of the configurations along the ablation path are better than the satellite PARAMILS configuration, evaluated on the satellite test set.

Table 19 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the depots test set, using the depots configuration as the source and the satellite configuration as the target.

round	parameter	performance
depots	-	0.658
1	extended_effects_evaluation	0.603
2	criprecond	0.640
3	ri_list	0.637
4	verifyinit	0.642
5	relaxed_examination	0.637
6	fast_best_action_evaluation	0.593
7	not_supported_preconds_evaluation	0.673
8	inc_re	0.677
9	zero_num_A	0.679
10	cri_update_iterations	0.679
11	maxnoise	0.678
12	evaluation_function	0.674
13	h	0.672
14	numrestart	0.670
15	improve_reachability	0.686
16	mutex_and_additive_effects	0.626
17	weight_mutex_in_relaxed_plan	0.617
18	numtry	0.631
19	evaluate_mutex_for_action_remotion	0.697
20	not_extended_unsupported_facts	0.773
21	no_mutex_with_additive_effects	0.866
22	reset_extended_unsupported_facts	0.959
23	noise	19.654
24	cri_intermediate_levels	1005.004
25	hpar_cut_neighb	2643.840
satellite	-	2643.840

Finally, for the ablation from the satellite configuration to the zenotravel configuration evaluated on depots, again we see that the source and target configurations are orders of magnitude worse than the default configuration. As with the other two experiments, the configurations along the ablation path have better performance on the depots test set than either the source or target configuration, although none reach the performance of the default configuration.

7 Future work

This work can be extended in various directions, with the primary extension being improved handling of conditional effects and interdependencies between parameters. As we outlined in Section 3, for an inactive conditional parameter to be flipped in a given ablation round, all of its dependent parent parameters that were inactive in the source configuration must have been flipped in previous ablation rounds. This dependency can cause an important conditional parameter to be flipped much later in the ablation process than it should be. We observed this with the *vicinato* and *hpar_cut_neighb* parameters in all three of our LPG scenarios, for example.

We propose to solve this issue by allowing the modification of more than one parameter in each ablation round. In this case, rather than ignoring a

Table 20 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the satellite test set, using the satellite configuration as the source and the depots configuration as the target.

round	parameter	performance
depots	-	7.617
1	hpar_cut_neighb	5.997
2	not_extended_unsupported_facts	5.805
3	reset_extended_unsupported_facts	5.710
4	criprecond	5.340
5	extended_effects_evaluation	5.290
6	evaluate_mutex_for_action_remotion	5.559
7	relaxed_examination	5.345
8	inc_re	5.313
9	ri_list	5.532
10	not_supported_preconds_evaluation	5.307
11	cri_intermediate_levels	5.368
12	numrestart	5.405
13	zero_num_A	5.376
14	lagrange	5.376
15	evaluation_function	5.375
16	h	5.675
17	noise	5.602
18	cri_update_iterations	5.708
19	weight_mutex_in_relaxed_plan	5.515
20	improve_reachability	5.462
21	no_mutex_with_additive_effects	5.427
22	maxnoise	5.473
23	numtry	5.668
24	fast_best_action_evaluation	5.640
25	mutex_and_additive_effects	5.297
26	verifyinit	5.783
satellite	-	5.783

parameter that is currently active, we would include as a candidate the configuration obtained by flipping that parameter and all of its inactive parent parameters to their value in the target configuration. We have implemented a prototype for this extension, and in preliminary experiments with SPEAR and LPG we observed ablation paths qualitatively similar to those in Section 5. However, as expected, in the LPG scenarios *vicinato* and *hpar_cut_neighb* are now flipped together much earlier in the ablation path, giving a better estimate of their importance.

Another opportunity for future work is the investigation of alternatives and improvements to our racing-based ablation approach. For example, we have observed that in later ablation rounds, F-Race quickly eliminates all candidate configurations except for a small number of configurations with very similar performance. In these cases, the maximum number of racing rounds are used. We also observed a related case where the F-Race was quickly reduced to two candidate configurations, where the p-value of the Friedman test was sufficient to proceed to pairwise elimination, but the p-value of the subsequent *t*-test was insufficient to eliminate the configuration with worse performance. Although this is consistent with what is known in terms of the impact of violated normality assumptions on the power of tests involving the *t*-statistic, it would be beneficial to investigate further as a solution to this problem could substantially reduce the number of algorithm runs required. Possible candidates

Table 21 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the depots test set, using the depots configuration as the source and the zenotravel configuration as the target.

round	parameter	performance
depots	-	0.658
1	extended_effects_evaluation	0.603
2	no_insert_threated_act_in_neighb	0.586
3	lagrange	0.595
4	no_pruning	0.595
5	choose_min_numA_fact	0.576
6	evaluation_function	0.584
7	verifyinit	0.595
8	zero_num_A	0.590
9	nonuniform_random	0.601
10	h	0.596
11	bestfirst	0.578
12	ri_list	0.588
13	maxnoise	0.590
14	cri_insertion_add_mutex	0.582
15	triomemory	0.591
16	numrestart	0.587
17	static_noise	0.635
18	inc_re	0.538
19	criprecond	0.543
20	dynoiseTabLen	0.542
21	noise_incr	0.542
22	dynoisecoefnum	0.542
23	twalkplan	2.022
24	walkplan	2.009
25	numtry	0.535
26	verifyAf	0.524
27	reset_extended_unsupported_facts	0.526
28	fast_best_action_evaluation	8.115
29	hpar_cut_neighb	190.062
30	noise	1275.739
31	cri_intermediate_levels	2767.166
32	stop_remove_act	2807.531
zenotravel	-	2807.531

would be to use a non-parametric test instead, possibly in combination with multiple-testing correction (see Styles and Hoos (2013)).

Finally, we believe that our approach and the model-based techniques discussed in Section 2 are complementary and can be combined, *e.g.*, by building functional ANOVA models using configurations sampled along ablation paths or from the localised region between the two input configurations.

8 Conclusions

In this work, we have introduced a new procedure, ablation analysis, which allows developers of highly-parameterised algorithms to ascertain which of their parameters contribute most to performance differences between two algorithm configurations. Using ablation analysis, it is possible to determine which modifications of a given default configuration were truly necessary to achieve improved performance, and which modifications can essentially be considered spurious side effects of an automated (or manual) configuration process.

Table 22 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the zenotravel test set, using the zenotravel configuration as the source and the depots configuration as the target.

round	parameter	performance
depots	-	1028.118
1	ri_list	1021.959
2	cri_intermediate_levels	10.058
3	hpar_cut_neighb	2.755
4	criprecond	3.584
5	twalkplan	2.163
6	extended_effects_evaluation	2.130
7	noise	2.061
8	stop_remove_act	2.020
9	verifyAf	2.019
10	numrestart	2.026
11	triomemory	2.044
12	walkplan	2.046
13	verifyinit	2.046
14	bestfirst	2.031
15	cri_insertion_add_mutex	2.060
16	static_noise	2.054
17	zero_num_A	2.056
18	maxnoise	2.077
19	dynoiseTabLen	2.055
20	lagrange	2.055
21	no_insert_threated_act_in_neighb	2.077
22	inc_re	2.082
23	h	2.086
24	dynoisecoefnum	2.108
25	numtry	2.116
26	choose_min_numA_fact	2.067
27	evaluation_function	2.077
28	reset_extended_unsupported_facts	2.063
29	fast_best_action_evaluation	2.115
30	noise_incr	2.088
31	nonuniform_random	2.040
32	no_pruning	2.092
zenotravel	-	2.092

We validated our approach in an experimental study using five well-studied configuration scenarios from propositional satisfiability, mixed-integer programming and AI planning, with 26 to 76 configurable parameters. We showed that a variant of our approach accelerated by a racing method required 25% of the CPU time needed by the brute-force variant, while achieving qualitatively similar results. In all of these scenarios, we found that 95–99% of the performance improvements achieved by automated configuration of the given, highly-parametric solver could be obtained with the modification of only 1–4 parameters, a small fraction of total number of parameters for each algorithm. In two cases, we found that modification of a single parameter could achieve 99.92% and 87.64% of the performance gain between the default configuration and one found by PARAMILS. Similar results have been reported for the global impact of parameters previously (*e.g.*, by Hutter et al (2014)), but we show that this is true locally for high-performance configurations, and in some cases the locally-important parameters are different from those that are important globally. Overall, we believe that our ablation analysis approach can be of great use to help algorithm developers and end users understand more

Table 23 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the satellite test set, using the satellite configuration as the source and the zenotravel configuration as the target.

round	parameter	performance
satellite	-	5.783
1	static_noise	5.293
2	dynoisecoefnum	5.293
3	no_pruning	5.767
4	improve_reachability	5.485
5	twalkplan	5.402
6	not_extended_unsupported_facts	5.708
7	nonuniform_random	5.778
8	verifyAf	5.526
9	noise_incr	5.631
10	not_supported_preconds_evaluation	5.613
11	dynoiseTabLen	5.681
12	bestfirst	5.616
13	cri_update_iterations	5.223
14	cri_insertion_add_mutex	5.502
15	evaluate_mutex_for_action_remotion	5.306
16	inc_re	5.293
17	h	5.385
18	numtry	5.471
19	no_mutex_with_additive_effects	5.322
20	mutex_and_additive_effects	5.275
21	triomemory	5.232
22	stop_remove_act	5.222
23	relaxed_examination	5.238
24	weight_mutex_in_relaxed_plan	5.297
25	numrestart	5.293
26	walkplan	5.312
27	no_insert_threated_act_in_neighb	5.355
28	choose_min_numA_fact	5.291
29	maxnoise	5.357
30	cri_intermediate_levels	7.952
zenotravel	-	7.952

about which algorithm parameters (and therefore algorithm subsystems and behaviours) are most responsible for high performance on problem instances of interest. The implementation of our approach has additionally been made available for other researchers to use, please see the project page ³ for more details.

Acknowledgements The authors would like to thank the anonymous reviewers of the earlier version of this paper presented at MIC 2013, as well as members of the UBC BETA lab for fruitful discussion regarding our approach.

References

Aghaeepour N, Hoos HH (2013) Ensemble-based prediction of RNA secondary structures. *BMC Bioinformatics* 14:139

³ <http://www.cs.ubc.ca/labs/beta/Projects/Ablation/>

Table 24 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the zenotravel test set, using the zenotravel configuration as the source and the satellite configuration as the target.

round	parameter	performance
satellite	-	2478.914
1	cri_intermediate_levels	2.954
2	stop_remove_act	2.142
3	not_supported_preconds_evaluation	2.280
4	walkplan	2.084
5	not_extended_unsupported_facts	2.084
6	numtry	2.096
7	evaluate_mutex_for_action_remotion	2.121
8	cri_update_iterations	2.062
9	mutex_and_additive_effects	2.022
10	relaxed_examination	2.021
11	maxnoise	2.019
12	cri_insertion_add_mutex	2.020
13	weight_mutex_in_relaxed_plan	2.017
14	numrestart	2.020
15	no_mutex_with_additive_effects	2.031
16	triomemory	2.038
17	choose_min_numA_fact	2.035
18	verifyAf	2.051
19	improve_reachability	2.040
20	inc_re	2.056
21	h	2.097
22	no_insert_threated_act_in_neighb	2.050
23	twalkplan	2.086
24	static_noise	2.066
25	noise_incr	2.104
26	bestfirst	2.088
27	nonuniform_random	2.040
28	no_pruning	2.092
zenotravel	-	2.092

- Ansótegui C, Sellmann M, Tierney K (2009) A gender-based genetic algorithm for the automatic configuration of algorithms. In: Proceedings of CP 2009, pp 142–157
- Bartz-Beielstein T (2006) Experimental Research in Evolutionary Computation—The New Experimentalism. Natural Computing Series, Springer, Berlin, Heidelberg, New York
- Birattari M, Stützle T, Paquete L, Varrentrapp K (2002) A racing algorithm for configuring metaheuristics. In: Proceedings of GECCO’02, pp 11–18
- Chiarandini M, Goegebeur Y (2010) Mixed models for the analysis of optimization algorithms. Experimental Methods for the Analysis of Optimization Algorithms pp 225–264
- Conover W (1999) Practical Nonparametric Statistics. John Wiley and Sons, New York, NY, USA, third edition
- Glover F (1994) Genetic algorithms and scatter search: Unsuspected potentials. Statistics and Computing 4:131–140
- Glover F (1997) Tabu search and adaptive memory programming - advances, applications and challenges. In: Barr RS, Helgason RV, Kennington JL (eds) Interfaces in Computer Science and Operations Research, vol 7, Springer US, pp 1–75

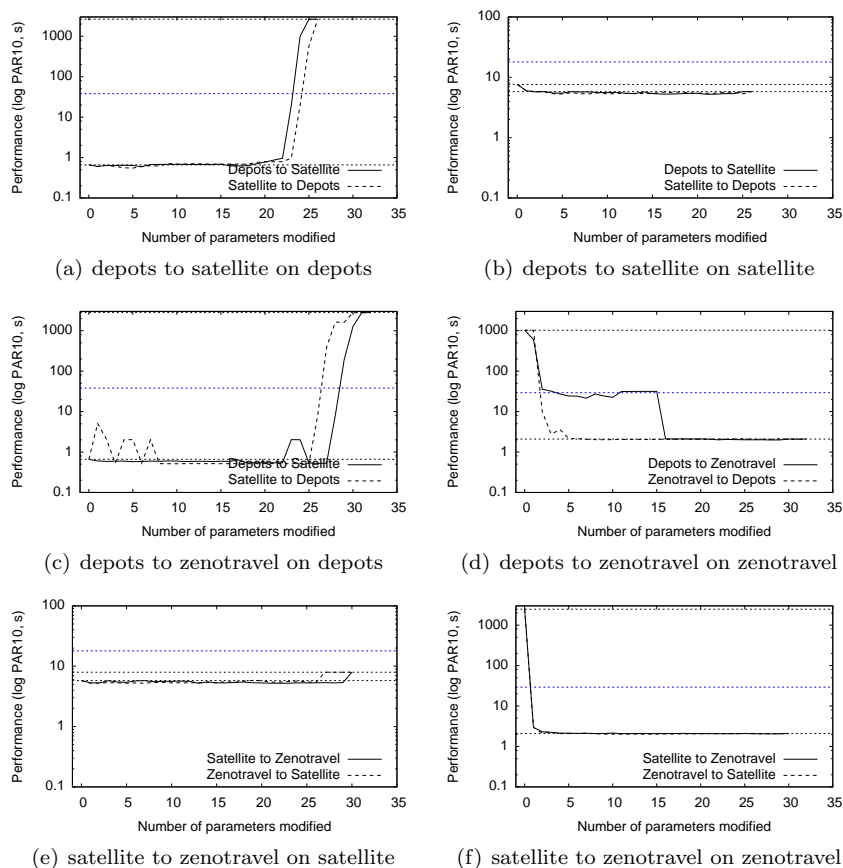


Fig. 4 Ablation paths determined using racing with up to 200 rounds on the test sets of the three LPG scenarios. Ablation was performed from the best PARAMILS incumbent on domain A to the best incumbent on domain B on each scenario, evaluated on domain A and B. The horizontal lines indicate the PAR10 scores of the default and the domain A/B incumbents.

- Glover F, Laguna M (1997) Tabu Search. Kluwer Academic Publishers
- Glover F, Laguna M, Martí R (2000) Fundamentals of scatter search and path relinking. *Control and Cybernetics* 39:653–684
- Glover F, Laguna M, Martí R (2003) Scatter search and path relinking: Advances and applications. In: Glover F, Kochenberger GA (eds) *Handbook of Metaheuristics*, International Series in Operations Research & Management Science, vol 57, Springer US, pp 1–35
- Gomes CP, van Hoeve WJ, Sabharwal A (2008) Connections in networks: a hybrid approach. In: *Proceedings of CPAIOR 2008*, pp 303 – 307
- Hooker G (2007) Generalized functional ANOVA diagnostics for high dimensional functions of dependent variables. *Journal of Computational and*

Table 25 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the zenotravel test set, using the depots configuration as the source and the satellite configuration as the target.

round	parameter	performance
depots	-	1028.131
1	hpar_cut_neighb	605.424
2	extended_effects_evaluation	35.646
3	not_supported_preconds_evaluation	26.560
4	fast_best_action_evaluation	22.093
5	reset_extended_unsupported_facts	18.780
6	no_mutex_with_additive_effects	17.173
7	not_extended_unsupported_facts	18.690
8	improve_reachability	18.594
9	verifyinit	18.628
10	numrestart	18.599
11	maxnoise	18.579
12	cri_update_iterations	18.550
13	evaluation_function	18.511
14	h	18.508
15	ri_list	18.544
16	relaxed_examination	18.502
17	zero_num_A	18.492
18	inc_re	17.706
19	numtry	21.087
20	criprecond	19.495
21	evaluate_mutex_for_action_remotion	26.080
22	mutex_and_additive_effects	28.061
23	weight_mutex_in_relaxed_plan	31.302
24	noise	38.765
25	cri_intermediate_levels	2478.914
satellite	-	2478.914

Graphical Statistics 16:709–732

- Hoos HH (2012) Programming by optimization. *Communications of the ACM* 55(2):70–80
- Hutter F, Babić D, Hoos HH, Hu AJ (2007a) Boosting verification by automatic tuning of decision procedures. In: *Formal Methods in Computer-Aided Design*, pp 27–34
- Hutter F, Hoos HH, Stützle T (2007b) Automatic algorithm configuration based on local search. In: *AAAI 07*, pp 1152–1157
- Hutter F, Hoos HH, Leyton-Brown K, Stützle T (2009) ParamILS: An automatic algorithm configuration framework. *Journal of Artificial Intelligence Research* 36:267–306
- Hutter F, Hoos HH, Leyton-Brown K (2010) Automated configuration of mixed integer programming solvers. In: *Proceedings of CPAIOR 2010*, pp 186–202
- Hutter F, Hoos HH, Leyton-Brown K (2011) Sequential model-based optimization for general algorithm configuration. In: *Proceedings of LION-5*, pp 507–523
- Hutter F, Hoos HH, Leyton-Brown K (2013) Identifying key algorithm parameters and instance features using forward selection. In: *Proceedings of LION-7*, to appear

Table 26 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the satellite test set, using the depots configuration as the source and the zenotravel configuration as the target.

round	parameter	performance
depots	-	7.620
1	hpar_cut_neighb	5.997
2	criprecond	5.719
3	walkplan	5.552
4	lagrange	5.733
5	verifyAf	5.706
6	static_noise	5.567
7	dynoisecoefnum	5.567
8	noise	5.470
9	maxnoise	5.574
10	inc_re	5.584
11	triomemory	5.486
12	no_pruning	5.363
13	numtry	5.304
14	numrestart	5.245
15	reset_extended_unsupported_facts	5.311
16	noise_incr	5.417
17	verifyinit	5.382
18	twalkplan	5.335
19	dynoiseTabLen	5.310
20	no_insert_threated_act_in_neighb	5.460
21	evaluation_function	5.638
22	fast_best_action_evaluation	5.441
23	h	5.460
24	extended_effects_evaluation	5.496
25	bestfirst	5.421
26	zero_num_A	5.325
27	nonuniform_random	5.454
28	ri_list	5.399
29	cri_insertion_add_mutex	5.310
30	stop_remove_act	5.409
31	choose_min_numA_fact	5.568
32	cri_intermediate_levels	7.951
zenotravel	-	7.951

- Hutter F, Hoos HH, Leyton-Brown K (2014) An efficient approach for assessing hyperparameter importance. In: Proceedings of the 31st International Conference on Machine Learning (ICML 2014), to appear
- López-Ibáñez M, Dubois-Lacoste J, Stützle T, Birattari M (2011) The irace package, iterated race for automatic algorithm configuration. Tech. Rep. TR/IRIDIA/2011-004, IRIDIA, Université Libre de Bruxelles, Belgium
- Maron O, Moore A (1994) Hoeffding races: Accelerating model selection search for classification and function approximation. In: Advances in Neural Information Processing Systems, vol 6, pp 59–66
- Nannen V, Eiben A (2007) Relevance estimation and value calibration of evolutionary algorithm parameters. In: Proceedings of IJCAI 2007, pp 975–980
- Nell C, Fawcett C, Hoos HH, Leyton-Brown K (2011) HAL: A framework for the automated analysis and design of high-performance algorithms. In: Proceedings of LION-5, pp 600 – 615
- Reidys CM (2011) Combinatorial Computational Biology of RNA: Pseudoknots and Neutral Networks. Springer

Table 27 Parameters chosen and resulting test set performance for ablation analysis accelerated by racing (maximum of 200 rounds) for LPG on the depots test set, using the satellite configuration as the source and the zenotravel configuration as the target.

round	parameter	performance
satellite	-	2643.840
1	walkplan	1386.764
2	verifyAf	1306.732
3	numtry	1309.948
4	static_noise	1246.085
5	dynoiseTabLen	1174.005
6	cri_insertion_add_mutex	1173.921
7	noise_incr	1173.913
8	not_supported_preconds_evaluation	1181.307
9	nonuniform_random	1181.321
10	relaxed_examination	1181.324
11	no_pruning	1181.252
12	triomemory	1181.313
13	improve_reachability	1181.309
14	dynoisecoefnum	1181.309
15	maxnoise	1181.284
16	h	1181.300
17	bestfirst	1181.288
18	cri_update_iterations	1181.284
19	choose_min_numA_fact	1181.288
20	evaluate_mutex_for_action_remotion	1122.883
21	not_extended_unsupported_facts	1070.618
22	no_mutex_with_additive_effects	1074.977
23	weight_mutex_in_relaxed_plan	1116.059
24	mutex_and_additive_effects	1098.086
25	no_insert_threated_act_in_neighb	1181.602
26	numrestart	1701.370
27	inc_re	1843.723
28	stop_remove_act	2148.242
29	cri_intermediate_levels	2575.323
30	twalkplan	2807.531
zenotravel	-	2807.531

- Styles J, Hoos HH (2013) Ordered racing protocols for automatically configuring algorithms for scaling performance. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2013, pp 551–558
- Vallati M, Fawcett C, Gerevini AE, Hoos HH, Saetti A (2011) Automatic generation of efficient domain-optimized planners from generic parametrized planners. In: Proceedings of IJCAI RCRA Workshop 2011, pp 111–123
- Vallati M, Fawcett C, Gerevini AE, Hoos HH, Saetti A (2013) Automatic generation of efficient domain-specific planners from generic parametrized planners. In: Proceedings of the 6th Annual Symposium on Combinatorial Search (SoCS-13), pp 184–192

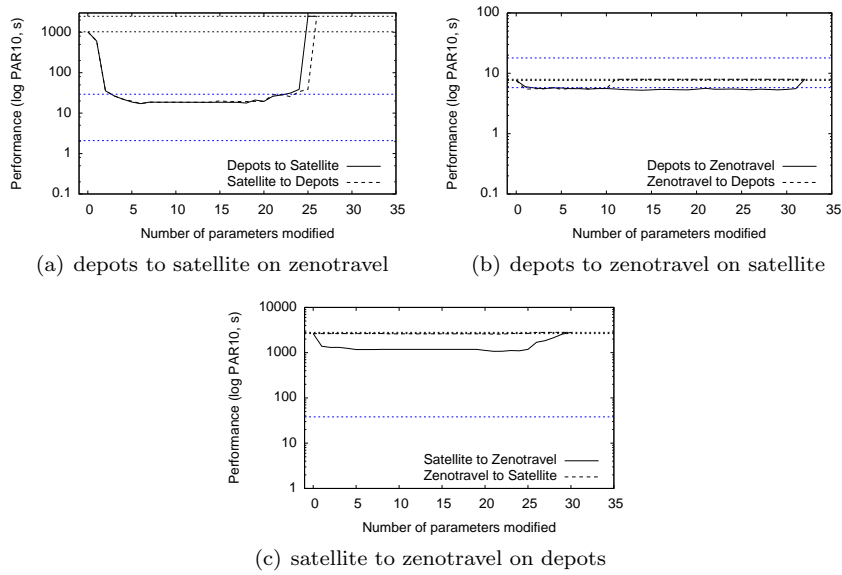


Fig. 5 Ablation paths determined using racing with up to 200 rounds on the test sets of the three LPG scenarios. Ablation was performed from the best PARAMILS incumbent on domain A to the best incumbent on domain B on each scenario, evaluated on domain C. The horizontal lines indicate the PAR10 scores of the default, the domain A/B incumbents, and the performance of the PARAMILS configuration for domain C.