

Quickstart Guide for SMAC version v2.04.02-master

Frank Hutter & Steve Ramage
Department of Computer Science
University of British Columbia
Vancouver, BC V6T 1Z4, Canada
{hutter, seramage}@cs.ubc.ca

July 6, 2013

1 Introduction

This document is the manual for SMAC [?] (an acronym for *Sequential Model-based Algorithm Configuration*). SMAC aims to solve the following *algorithm configuration* problem: Given a binary of a parameterized algorithm \mathcal{A} , a set of instances \mathcal{S} of the problem \mathcal{A} solves, and a performance metric m , find parameter settings of \mathcal{A} optimizing m across \mathcal{S} . The goal of this quickstart guide is to get you off the ground quickly; for more detailed information, please see the comprehensive manual.

1.1 License

SMAC is to be released under a dual usage license. Academic & non-commercial usage is permitted free of charge. Please contact us to discuss commercial usage.

2 Usage

2.1 System Requirements

SMAC itself requires only Java 6 or newer to run. The included scripts are currently only available for Unix-compatible operating systems. The included example scenarios require Ruby.

2.2 Running SMAC

Download and extract the SMAC archive from:

<http://www.cs.ubc.ca/labs/beta/Projects/SMAC/smac-VERSION.tar.gz>

This will create a new folder named `smac-VERSION` containing SMAC and several example scenarios. From this folder, you can run SMAC using the simple bash script `./smac`. If bash is not available (e.g., on Windows), you can run SMAC directly using `java`.¹

¹The main entry point class is `ca.ubc.cs.beta.smac.executorsAutomaticConfigurator` and you need to include every jar in the folder in the classpath. You would call it using: `java -cp <jar files> ca.ubc.cs.beta.smac.executorsAutomaticConfigurator`.

2.3 Basic Usage Example

SMAC comes with some small configuration scenarios you can run to ensure everything is set up properly. We will use one of them as an example:

```
example_spear/scenario-Spear-QCP-sat-small-train-small-test-mixed.txt
```

This configures the systematic tree search algorithm Spear for a small training set of 5 SAT instances. (Normally, one would use a much larger set of training instances to avoid over-tuning to the specific instances in the training set – we usually use 1 000 instances if available; here, we did not include as many instances to keep the release small in size.) From `smac-v2.04.02-master-483`, you can start SMAC on this scenario with the following command line:

```
./smac --scenarioFile  
example_spear/scenario-Spear-QCP-sat-small-train-small-test-mixed.txt  
--runGroupName example_out --numRun 0 --numberOfValidationRuns 10
```

This command executes SMAC on this simple configuration scenario (which should take just above 30 seconds) and writes several output files to folder

```
smac-VERSION/example_out
```

The `numRun` parameter controls SMAC's random seed and output filenames in that folder; since SMAC is highly randomized and its performance differs across runs, we recommend to perform several parallel runs (with different values for `numRun`) and use the one with the best training performance (see [?] for details). The parameter `numberOfValidationRuns` controls the number of runs to perform for validating the final incumbent SMAC finds at the end of its runtime budget (again, here we set this to a small number of 10 to facilitate a quick example run, but in practice one would use a much larger number, *e.g.* the default of 1 000 validation runs). Note that validation runs are carried out on a benchmark test set disjoint from the training set; this is done to guard against over-tuning since performance on the training set is a biased (optimistic) estimator of generalization performance.

2.4 Interpreting output files

SMAC writes several output files; not all of these are important in everyday use, but if you encounter a problem please zip up and send all of them. The most important output file for users is the log file, here:

```
smac-VERSION/example_out/log-run0.txt
```

This file lists a lot of information about the run and at its end includes an example command line call for the final incumbent, as well as its training and test performance. For example, in the run we performed, this output was:

```
[INFO ] Total Objective of Final Incumbent 13 (0x2923B)  
on training set: 0.09066666666666667; on test set: 0.053  
[INFO ] Sample Call for Final Incumbent 13 (0x2923B)  
cd /ubc/cs/project/arrow/hutter/smac/smac-VERSION/example_spear;  
ruby spear_wrapper.rb example_data/SWGCP-satisfiable-instances/SWlin2006.8287.cnf  
0 5.0 2147483647 3724427 -sp-clause-activity-inc '0.5463890824415055'  
-sp-clause-decay '1.052825346704275' -sp-clause-del-heur '2' -sp-first-restart
```

```
'162' -sp-learned-clause-sort-heur '12' -sp-learned-clauses-inc '1.2595926598325478'
-sp-learned-size-factor '0.24332560173797435' -sp-max-res-lit-inc '2.5150171126927'
-sp-max-res-runs '4' -sp-orig-clause-sort-heur '0' -sp-phase-dec-heur
'4' -sp-rand-phase-dec-freq '0.001' -sp-rand-phase-scaling '1.0144220679252771'
-sp-rand-var-dec-freq '0.005' -sp-rand-var-dec-scaling '0.6040334190120095'
-sp-res-cutoff-cls '2' -sp-res-cutoff-lits '101' -sp-res-order-heur '13'
-sp-resolution '2' -sp-restart-inc '1.36773341049033' -sp-update-dec-queue
'1' -sp-use-pure-literal-rule '0' -sp-var-activity-inc '0.5541449022276423'
-sp-var-dec-heur '17' -sp-variable-decay '1.0871444238524568'
```

Note that due to SMAC's randomization, the result will likely differ when you execute this command. Also note that training and test performance are rather different from each other here because of the small number of training/test instances used. With larger training and test sets sampled from the same underlying instance distribution, these numbers should be much closer. A better training than test performance can be a sign of over-tuning.

3 Configuring your own algorithms

We will refer to the algorithm to be configured as the *target algorithm*. The quickest way to set up everything needed to configure a new target algorithm is to copy-paste and edit one of the provided examples. The following subsections briefly describe the parts of a configuration scenario and give an example for each. Full details are given in Section 4 of the SMAC manual. All filenames below are relative to the SMAC folder `smac-VERSION/`.

3.1 Scenario File

This is the central file describing all ingredients of a configuration scenario. This file is not actually required since all its contents can also be specified on the command line, but it is still often useful to have as a central specification saved in this file. It references an executable for calling the target algorithm (and the path to execute it from), a file describing the target algorithm's parameters, a file containing a list of problem instances, a performance metric, an optional file specifying characteristics for each training instance, and some additional information for the configuration procedure. In the Spear example above, this file is

```
example_spear/scenario-Spear-QCP-sat-small-train-small-test-mixed.txt
```

3.2 Parameter File

This file, referenced in the scenario file, specifies the configurable parameters your target algorithm accepts as an input, along with their domains. In the SPEAR example above, this file is

```
example_spear/spear-params-mixed.txt
```

In this file, notice the different syntax for specifying the domain of categorical parameters (a set of possible values in curly braces) and of numerical parameters (an interval of possible values). For both of these parameter types, the default is given after the domain specification. Also notice the two optional modifiers of numerical parameters given right after the default; 'i' stands for integer parameters and 'l' stands for parameters that naturally vary on a log-scale. For example, Spear parameter *sp-first-restart* is an integer parameter that would naturally be discretized to 25, 50, ..., 3200, with default 100. As a categorical parameter, this would be written as `{25, 50, 100, 200, 400, 800, 1600, 3200}[100]`; as an integer parameter, it is now written as `[25, 3200][100]i1`.

3.3 Algorithm wrapper

The scenario file references an executable of the target algorithm. Since this executable has to follow a rigid input/output format that we don't want to impose on algorithm developers, we typically specify a *wrapper* around the target algorithm that translates between SMAC's and the target algorithm's input/output format. The input specifies (in this order) instance name, instance specifics (which is often empty but could, *e.g.*, hold an optimal solution quality to reach), a cutoff time for the run, a cutoff length for the run (often empty), a seed, and a list of parameter value pairs in the format `-param value`. In the Spear example above, the wrapper is the Ruby script

```
example_spear/spear_wrapper.rb.
```

Out of its inputs, this wrapper script only uses the instance name, cutoff time, seed, and parameter values; it ignores the two inputs instance specifics and cutoff length. It then runs Spear with the given parameters on the named instance, instructing it to terminate unsuccessfully after the given cutoff time. It then parses Spear's output and prints the following output line, from which SMAC parses the run's performance (in this case `measured_runtime`):

```
Result for ParamILS: <solved>, <measured_runtime>, <measured_runlength>,
<best_solution>, <seed>
```

Note that while this output line is compatible with our previous ParamILS software, SMAC can also parse other result lines. See Section 4.5 in the manual for this and further details on the wrapper.

3.3.1 Wrapper Changes

Compared to ParamILS, we allow a new possible result code *ABORT* that is used to signify that not only did the program fail, but that if this occurs the configuration procedure should abort.

3.4 Instance File

This file, referenced in the scenario file, specifies the list of training instances to configure the target algorithm on. In the Spear example above, this file is

```
example_spear/example_data/SPEAR-QCP-train.txt
```

There is also a file with test instances that will not be used at configuration time, but are used for offline validation of the final incumbent SMAC found; in the Spear example above, this file is

```
example_spear/example_data/SPEAR-QCP-test.txt
```

3.5 Feature file

In order to optimize performance across a set or distribution of instances, SMAC accepts an optional input file that contains additional information for each instance in the training set. If such a feature file is given it has to contain an entry for every instance in the training set. Further, the file has to hold the same features (in the same order) for each instance. Any domain-specific instance characteristics (or "features") that might correlate with performance can be useful. If the file contains entries for additional instances these will be ignored by SMAC. (This feature enables users to, *e.g.*, have one file containing the features for all their SAT instances.) In the Spear example above, the feature file is

```
example_spear/SPEAR-QCP-train-features.csv
```

4 Differences Between SMAC and ParamILS

There are various differences between SMAC and ParamILS. The following differences are described in some more detail in Section 2 of the SMAC manual.

- SMAC natively supports continuous and integer parameters.
- SMAC so far only implements ParamILS's most popular run objectives.
- The order of instances in SMAC's instance file is not important.
- SMAC counts its own overhead time as part of the configuration budget and allows for wall clock timeouts. This can fix ParamILS's large overheads in cases where target algorithm runs are extremely fast.
- SMAC can resume previous runs.
- SMAC accepts optional feature files as input.
- SMAC supports new options for the wrappers.
- SMAC auto-detects instance files vs. instance/seed files.