

FAQ for SMAC version v2.10.02-master

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

July 5, 2015

FAQ

1	Troubleshooting Questions	2
1.1	My target algorithm fails to run, what should I do?	2
1.2	SMAC gets a java.io.IOException: error=12, Cannot allocate memory when trying to execute the target algorithm	3
1.3	Why does SMAC do weird things to my console (change colors, set the title, etc...)?	3
1.4	SMAC seems to take so much longer than tunerTimeout to execute, what can I do?	3
1.5	How can I report a bug or documentation error?	3
1.6	My environment can be flaky, I don't want to ABORT but don't want to blindly continue with CRASHED	3
1.7	SMAC occasionally experiences a SEGV in the garbage collector	4
2	Usage Questions	4
2.1	System Compatibility	4
2.1.1	Can I run SMAC on Windows?	4
2.1.2	Can I run SMAC on Mac OS X or other Unix environments?	4
2.2	I don't have any instance features, is there any benefit I can get from running SMAC?	4
2.3	Validation	4
2.3.1	Does SMAC support validation?	4
2.3.2	How can I disable Validation?	4
2.3.3	Does SMAC support validating more than the last incumbent?	4
2.3.4	How can I change the amount of memory SMAC uses?	5
2.4	Can I change the number of cores used in post-SMAC Validation?	5
2.5	Can I change the Target Algorithm Evaluator used in post-SMAC Validation?	5
2.6	Parameter Configuration Space	5
2.6.1	Should I make a parameter categorical or numeric integral?	5
2.7	Should I make a parameter ordinal or numeric integer?	5
2.7.1	How should I encode a parameter interdependency in a PCS file (for example $x^2 + y^2 \leq 1$)?	5
2.8	What are the best options for running SMAC?	6

2.9	Saving and Restoring State	6
2.9.1	Can I continue a finished run of SMAC (<i>e.g.</i> , after a crash or running out of time)?	6
2.9.2	SMAC's state files eat up alot of disk space, can I turn them off?	6
2.9.3	Saving state takes a long time what are my options?	7
2.9.4	Can I save more than just the final set of runs?	7
2.9.5	How can I see what iterations I can restore SMAC to?	7
2.9.6	Can I change the settings of SMAC when restoring state (<i>e.g.</i> ,add more instances, change objectives)?	7
2.10	How do I use the preliminary bash auto-completion with SMAC?	7
3	Target Algorithm & Wrapper Questions	7
3.1	Should I use a wrapper or modify my program to execute directly?	7
3.2	How can I test that my wrapper will work with SMAC easily?	8
3.3	What is the difference between SAT and UNSAT in the algorithm responses?	8
3.4	What are the performance consequences of executing in parallel (<i>i.e.</i> setting <code>--maxConcurrentAlgoExecs > 1</code>)?	8
4	Customization and Developer Questions	9
4.1	There is a lack of logging options available, I need something more specific	9
4.2	I'd like to make some changes to SMAC where should I start?	9
4.3	I would like to optimize something that doesn't lend itself to simple command line execution, and/or would like to use something more advanced (<i>e.g.</i> , database caching, running on remote machines, etc...)	9
4.4	I would like to use different objective functions than are provided	9
4.5	Why doesn't SMAC support as many objective functions as ParamILS?	9
4.6	What commits generated this version of SMAC?	9
5	Miscellaneous Questions	10
5.1	What are the differences between SMAC and ParamILS?	10
5.2	What are the relevant articles / papers about SMAC?	10
5.3	I'm using SMAC in academic work, what article should I cite?	10
5.4	I'd like to use SMAC for commercial purposes, what should I do?	11
5.5	What does the name after the dash in version refer to?	11
5.6	Is the source code for SMAC available?	11
5.7	Can I use a different logger for SMAC?	11

1 Troubleshooting Questions

1.1 My target algorithm fails to run, what should I do?

Consult the logs, and see what error messages are recorded. In some cases you may have to turn on additional options such as `--logAllProcessOutput`, and see what information is available.

If it's not clear what the problem is consult the logs to determine the call string used and try executing the target algorithm directly on the environment that it failed on. If direct execution fails, then the problem is specific to the target algorithm and/or your scenario options *i.e.* The algorithm executable cannot be found .

Using the `smac-algotest` utility may make it easier to reproduce and diagnose the problem (see section 3.2).

If the call succeeds, the next step would be to replace the target algorithm that is being executed with a debug script/program that outputs all of the environment variables, and the call string it sees directly from SMAC. If nothing here seems to be relevant to your problem, then you will probably want to contact us for more advice.

NOTE: If the call string that SMAC is using is not in the correct format, you probably need to use a wrapper to format it, and you should consult the *Algorithm executable / wrapper* in the manual.

NOTE: In certain Java implementations (such as Sun's Java 6) there is a bug¹ that causes the `LD_LIBRARY_PATH` variable to be modified internally by Java, and then the target algorithm sees the modification. In most cases this is benign, as it merely adds the location of the JVM libraries, and restates the location of the system libraries. However if library incompatibilities exist, this can cause problems. You will need to use a wrapper script to manually fix the `LD_LIBRARY_PATH` or upgrade Java.

1.2 SMAC gets a `java.io.IOException: error=12, Cannot allocate memory when trying to execute the target algorithm`

This is caused by a bug in earlier versions of Java 6 and Java 7. When SMAC executes the program in these versions, the system calls `fork()` and `exec()` which causes the entire address space footprint of SMAC to be duplicated, while this is done with traditional copy-on-write semantics, in systems which will not over commit resources, this causes the fork to fail. For more information see: http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=7034935

1.3 Why does SMAC do weird things to my console (change colors, set the title, etc...)?

SMAC outputs a dump of the environment variables to console when `DEBUG` logging is enabled. If any of these characters are terminal escape sequences they may unintentionally cause the terminal's behaviour to change.

1.4 SMAC seems to take so much longer than `tunerTimeout` to execute, what can I do?

Consider using the `--wallClockLimit` option to limit the amount of wall clock time that SMAC will execute for. Additionally consider looking at the `run_and_results` file for the run (see *State Files* in the SMAC Manual) and see how large the difference between the recorded wall clock time and the reported runtime is.

1.5 How can I report a bug or documentation error?

Please make a post on the SMAC forum: <https://groups.google.com/forum/#!forum/smac-forum>.

1.6 My environment can be flaky, I don't want to `ABORT` but don't want to blindly continue with `CRASHED`

Consider using the `--abortOnCrash` option in conjunction with `--retryTargetAlgorithmRunCount`. This will allow you to retry several times before ultimately giving up.

¹See http://bugs.sun.com/view_bug.do?bug_id=6670965

1.7 SMAC occasionally experiences a SEGFault in the garbage collector

We experienced this problem as well with SMAC and other applications, due to the rarity we don't know what exactly is causing it. This occurred on Java 1.7.0_45 and earlier. We can't comment directly other versions of Java but running with Java 1.8.0_11 seemed to have mitigated this. For Java 7 we recommend you try using version 1.7.0_65 or later.

2 Usage Questions

2.1 System Compatibility

2.1.1 Can I run SMAC on Windows?

We have done our best to ensure that SMAC works on windows. SMAC includes `.bat` files for execution that behave similarly to the unix shell scripts. Most wrappers require that you have ruby (which is available from: [ruby http://rubyinstaller.org/](http://rubyinstaller.org/)).

The only scenarios that will work on Windows is the `saps/saps-scenario.txt` scenario, which requires ruby, and the `branin/branin-scenario.txt` which requires python.

2.1.2 Can I run SMAC on Mac OS X or other Unix environments?

SMAC should work on any Unix environment that supports Java 7. The only included scenario that will support Mac OS X is the `saps/saps-scenario.txt`, which requires ruby, and the `branin/branin-scenario.txt` which requires python.

2.2 I don't have any instance features, is there any benefit I can get from running SMAC?

Yes, SMAC can also run without features, and sometimes continues to show state-of-the-art performance. Additionally, we have noticed that the performance of the default configuration on all instances is a good instance feature in many situations.

2.3 Validation

2.3.1 Does SMAC support validation?

Yes, see *Offline Validation* in the manual.

2.3.2 How can I disable Validation?

Yes, by setting `--validation` to false.

2.3.3 Does SMAC support validating more than the last incumbent?

Yes, by setting `--validateOnlyLastIncumbent` to false. Additionally see Validation Options in the SMAC Manual appendix for options that control which incumbents can be selected.

2.3.4 How can I change the amount of memory SMAC uses?

If you are using the supplied shell scripts, you can set the `SMAC_MEMORY` environment variable to a positive integer corresponding to amount of memory in MB SMAC will use. If you are using some other mechanism of execution, the `-Xmx` java option is the only other way.

2.4 Can I change the number of cores used in post-SMAC Validation?

Yes, using the `--validation-cores` options. This will only work if you are validating locally.

2.5 Can I change the Target Algorithm Evaluator used in post-SMAC Validation?

Not at this time.

2.6 Parameter Configuration Space

2.6.1 Should I make a parameter categorical or numeric integral?

Numeric integral parameters are internally treated as numeric continuous parameters, and only when a call to the target algorithm is made, is the value rounded to an integer.

There is no hard and fast rule governing which method is better. If the value really is numeric and smooth, then using a numeric integral is likely to result in better predictions. If on the other hand the value is used as an index to a switch, then categorical should be better.

To illustrate using extreme examples of this:

If your value represents a bit-mask, used to control four separate options, where the most meaningful split would be based on the least significant bit (*i.e.*, the best split would be into even and odd numbers), a numeric integral value with domain $[0, 15]$ could never capture this, nor would the model ever detect it. On the other extreme, if the value really is numeric integral (*e.g.* $[0, 1024]$) then values that have never been seen will be randomly assigned to leafs, not capturing any smoothness.

To help inform the decision, from a SMAC internal point of view there are two places in the code where this distinction will matter. The first is during neighbourhood searches. All numeric parameters will only have 4 neighbours selected, and those neighbours will be sampled from a distribution $N(\mu, 0.2)$ where μ is the existing (*standardized* to $[0, 1]$) value. Any samples that fall outside of this range will be rejected.

The next distinction is within the random forest itself. Categorical values can be split arbitrarily but numeric values are always split at some constant c , with values $>$ than c going to the right, and values \leq going to the left.

2.7 Should I make a parameter ordinal or numeric integer?

If you are encoding the same set of values then it doesn't really matter as internally they will both be treated the same. The numeric integer format is probably more concise and less error-prone.

2.7.1 How should I encode a parameter interdependency in a PCS file (for example $x^2 + y^2 \leq 1$)?

There are a number of ways to do this. The first and most direct, if you are using SMAC v2.10 or later, is to simply use the advanced forbidden syntax, which allows you to specify any expression. Keep in mind that the expression indicates which parameters are forbidden, so to encode the example it would be:

```
x real [-1,1] [0]
y real [-1,1] [0]

{ x^2 + y^2 > 1 } #Forbidden if outside the circle.
```

See the manual section on the PCS format, for more information. Keep in mind that forbidden parameter settings can slow down the configuration scenario. In particular the slowdown is proportional to the region of the space that is forbidden over the total area of the space ignoring forbiddens. In the previous specification, changing x to have a range of $[-10, 10]$ would slow down configuration by $10\times$.

One can also in some cases encode the parameter space another way of encoding the parameter space is to use:

```
angle [0,360] [0]
radius [0,1] [0]
```

A third less obvious way is:

```
x real [-1,1] [0]
yScalePercent real [-100,100] [0]
```

In this example the x is chosen as normal, but y is a percentage of the allowed region, so the original y is: $y = \frac{yScalePercent}{100} \cdot (1 - \sqrt{|x|})$. This would be converted internally to the target algorithm or wrapper.

Which option is best unfortunately depends on the algorithm being optimized. In general you should chose an encoding of your parameter space that minimizes the interaction effects between terms.

2.8 What are the best options for running SMAC?

The defaults have been set intelligently enough and some settings will have their defaults changed based on the setting of others. The only setting disabled by default that might actually improved performance is **--share-run-data**. It is an advanced option that fundamentally changes how SMAC works. You should look in the manual for a discussion of how this mode works under the *Shared Model Model* section. In short if you are trying to tune your algorithm for best performance, it is appropriate to use, but if you are conducting a scientific inquiry into your algorithms performance there are some caveats that you should be aware of.

2.9 Saving and Restoring State

2.9.1 Can I continue a finished run of SMAC (e.g., after a crash or running out of time)?

Yes, see *State Restoration* in the manual.

2.9.2 SMAC's state files eat up alot of disk space, can I turn them off?

Yes, set the **--stateSerializer** to *NULL*. Beginning in version 2.02, the default was to only save the files needed to restore the final set of runs, and not necessarily every iteration before hand.

2.9.3 Saving state takes a long time what are my options?

SMAC's state saving mechanism should be fairly cheap to save states, but if you have many many instances (say 30,000) it can become very expensive to save state, even the quick states. You can of course turn state serialization off, or you could use the **--quick-saves**, **--intermediary-saves** to fine tune when SMAC actually makes saves.

2.9.4 Can I save more than just the final set of runs?

Yes, set the **--cleanOldStateOnSuccess** option to false, and a complete state will be saved on every 2^n $n \in \mathbb{N}$ iteration.

2.9.5 How can I see what iterations I can restore SMAC to?

Using the `smac-possible-restores` utility. Point it to the directory that contains the state you'd like to restore.

2.9.6 Can I change the settings of SMAC when restoring state (e.g., add more instances, change objectives)?

Yes, the only thing that is restored with the state is the runs, and the state of the random objects. Everything else comes from the configuration supplied. You may want to look at the source code of `LegacyStateDeserializer` to see exactly what can be changed and what can't. NOTE: some invariants of SMAC may no longer be in force, and care must be taken. The incumbent may no longer be the best found under the new objectives, and/or the incumbent may no longer have the most number of runs ².

2.10 How do I use the preliminary bash auto-completion with SMAC?

You can load the auto completion information by running `./util/bash_autocomplete.sh` at a bash prompt. For instance:

```
$. ./util/bash_autocomplete.sh
```

3 Target Algorithm & Wrapper Questions

3.1 Should I use a wrapper or modify my program to execute directly?

Either will work. It may be easier to use a wrapper but this can introduce more overhead, which (depending on the runtime of the target algorithm) may be significant (we typically see overheads up to a second). When using a wrapper it is important not to poll the output stream of the process: if the target algorithm outputs lots of data this can result in a high degree of lock contention and significantly affect the runtime performance.

²If you don't include the `java_obj_dump` file the incumbent should be recalculated for you. The calculated incumbent will be the one that has the most runs, and performs the best on the objective. Additionally you also lose the random object state.

3.2 How can I test that my wrapper will work with SMAC easily?

You can use the `smac-algotest` utility. It provides the same options that SMAC does with respect to executing your algorithm, but lets you specify the instance, seed & configuration. For example the following tests a random configuration of `spear` with a cutoff time of 50 seconds and the `sp-var-activity-inc` variable set to `1.0924`:

```
./smac-algotest --execDir ./example_spear --paramFile ./example_spear/spear-params-mixed.txt
--algo "ruby spear_wrapper.rb" --cutoff_time 50 --instance
example_data/QCP-instances/qcplin2006.10408.cnf -Psp-var-activity-inc=1.0924
--config RANDOM
```

3.3 What is the difference between SAT and UNSAT in the algorithm responses?

The short answer is you can use either, as SMAC makes no distinction. The reason two responses are used is a historical hold over for debugging. Essentially it was used as a cross check to ensure that some unknown bug didn't crop up, as a result of the running configuration that would break the solver. SMAC itself doesn't do anything with this information, but it's possible to use the `run_and_results` file to check the responses against the instances, and ensure that no errors were made. If your instance specific information file follows the conventions outlined in the `--verifySAT` option SMAC will report an error if the algorithm ever reports differently. Additionally SMAC internally will keep track of this output and if it is ever inconsistent for a given instance it will report it.

3.4 What are the performance consequences of executing in parallel (i.e. setting `--maxConcurrentAlgoExec > 1`)?

We don't feel that this option is likely to give you much bang for the buck, and a better option of utilizing multiple cores see *Shared Model Mode* in the manual for more information.

It largely depends on how SMAC ends up exploring the configuration space. The largest change one could expect for a series of N algorithm executions, given infinite concurrent processors would be from $O(N)$ to $O(\lg(N))$. NOTE: This case would generally occur when we have lots of runs for the incumbent, and challengers perform very close to the incumbent's performance. If the configurations SMAC finds can be shown to perform poorly fairly quickly, then the runtime would still be about $O(N)$, and there would be no significant improvement.

Additionally when using adaptive capping with M processors, all parallel runs will be scheduled as if the other runs took 0 seconds. In other words, if we want to exploit the parallelism of running on M processors, we have to run each algorithm for M times as long, thus nullifying any advantage.

For validation however, given infinite processors the runtime would be $O(1)$, and depending on the scenario validation can be more than 50% of the runtime.

Note: If you execute SMAC on a multi-core system using 1 core, and then validate using many cores you are likely to skew the results. The reasons are that multi-core systems typically share some level of caches, and when SMAC is executing by itself it gets the full cache, but when validating it gets cut into a much smaller fraction. This is also a problem when running SMAC on a cluster, but it is presently an open problem of how to best utilize cores.

4 Customization and Developer Questions

4.1 There is a lack of logging options available, I need something more specific

SMAC uses slf4j (<http://www.slf4j.org/>), a library that allows for abstracting and replacing the logging system with ease, and uses logback (<http://logback.qos.ch/>) as the default logging system. While there is limited ability to change logging options via the command line (*e.g.*, `--log-level`, `--console-log-level`, `--log-all-call-strings`, `--log-all-process-output`). It is possible that by setting a system property³ you can override the configuration by using `-Dlogback.configurationFile=/path/to/config.xml`

NOTE: If you replace the logger in SMAC or modify the configuration file, the logging command line options may no longer work.

4.2 I'd like to make some changes to SMAC where should I start?

The manual contains a *Developer Reference* section which should hopefully give you enough overview to start looking at the code. At the class and public method level most of SMAC's components are well documented. Internally much of SMAC's functionality is implemented by a library called AEATK, and a manual is under development for it as well. You can e-mail the developers above to get a copy of the current edition.

4.3 I would like to optimize something that doesn't lend itself to simple command line execution, and/or would like to use something more advanced (*e.g.*, database caching, running on remote machines, etc...)

You probably need to implement a `TargetAlgorithmEvaluator`. See the *Target Algorithm Evaluator* in the manual.

4.4 I would like to use different objective functions than are provided

You will have to modify the source code, namely the `RunObjective` and `OverallObjective` classes to support the new objective. Also see Question 4.5.

4.5 Why doesn't SMAC support as many objective functions as ParamILS?

The Random Forests currently only optimize for **MEAN** and so other objective functions may have to deal with a poorer model quality. As we do not currently have other scenarios, we were unable to measure how bad the performance was and have disabled these options in the code. When `--executionMode` is set to **ROAR** this shouldn't be an issue.

4.6 What commits generated this version of SMAC?

Project	Version	Commit	Dirty Flag
aeatk	v2.10.01-master-803	d97311a573a01ad81b40b818502fb29daa6e9f05	0
aeatk	v2.10.02-development-804	c4b45729e04087639b294c9783cf1261391befb2	0
aeatk	v2.10.02-master-805	9c0ad1b3eea679d95755e1444e3f4e38626aae4d	0
smac	v2.10.02-master-773	841b0af6f350bfc8ca6f478d509f240046834441	0

³You'd have to edit the start up script `smac` or `smac.bat`

NOTE: For non-master builds these commits may not contain everything in the build. (*i.e.*, non-master builds can be built with uncommitted changes). If the dirty flag is 0 that means the commit contains this exact copy, 1 means there were some uncommitted changes, and something else means some other error occurred when we tried to generate this.

5 Miscellaneous Questions

5.1 What are the differences between SMAC and ParamILS?

See *Differences Between SMAC and ParamILS* in the manual.

5.2 What are the relevant articles / papers about SMAC?

Several papers led up to SMAC, but the most important/up-to-date ones are (in bibtex format):

```
@InProceedings{HutHooLey11-SMAC,
  author = {F. Hutter and H.~H. Hoos and K. Leyton-Brown},
  title = {Sequential Model-Based Optimization for General
    Algorithm Configuration},
  booktitle = {LION-5},
  series = {LNCS},
  year = {2011},
  pages = {507--523}
}

@InProceedings{HutHooLey11-censoring,
  author = {F. Hutter and H.~H.~Hoos and K. Leyton-Brown},
  title = {Bayesian Optimization With Censored Response Data},
  booktitle = {2011 NIPS workshop on Bayesian Optimization,
    Sequential Experimental Design, and Bandits},
  year = {2011},
  Note = {Published online}
}

@INPROCEEDINGS{HutHooLey12-ParallelAC,
  author = {Frank Hutter and Holger~H. Hoos and Kevin Leyton-Brown},
  title = {Parallel Algorithm Configuration},
  booktitle = {LION-6},
  year = {2012},
  series = {LNCS},
  note = {To appear}
}
```

5.3 I'm using SMAC in academic work, what article should I cite?

Please cite the above LION-5 paper when you use SMAC. The above NIPS workshop paper is a reference for adaptive capping in the model-based framework, and the above LION-6 paper for a parallel version of

SMAC.

5.4 I'd like to use SMAC for commercial purposes, what should I do?

This version of SMAC is licensed under the AGPLv3. Please contact us to discuss other licensing options. Note that previous versions of SMAC were not available under this license.

5.5 What does the name after the dash in version refer to?

The name refers to the git branch, almost every public release should be *master*. Development releases are *development*.

5.6 Is the source code for SMAC available?

The source code for SMAC is included in the files `smac-src.jar`, `aclib-src.jar`, and `fastrf-src.jar`.

5.7 Can I use a different logger for SMAC?

Yes, you should be able to simply replace the `logback.jar`s with something else that `slf4j` supports and it should work out of the box.