

0.0.1 SMAC Options

General Options for Running SMAC

--adaptiveCapping Use Adaptive Capping

Aliases: --adaptiveCapping

Default Value: Defaults to true when --intraInstanceObjective is RUNTIME, false otherwise

Domain: {*true, false*}

--capAddSlack amount to increase computed adaptive capping value of challengers by (post scaling)

Aliases: --capAddSlack

Default Value: 1.0

Domain: (0, ∞)

--capSlack amount to scale computed adaptive capping value of challengers by

Aliases: --capSlack

Default Value: 1.3

Domain: (0, ∞)

--consoleLogLevel default log level of console output (this cannot be more verbose than the logLevel)

Aliases: --consoleLogLevel

Default Value: INFO

Domain: {*TRACE, DEBUG, INFO, WARN, ERROR, OFF*}

--countSMACTimeAsTunerTime include the CPU Time of SMAC as part of the tunerTimeout

Aliases: --countSMACTimeAsTunerTime

Default Value: true

Domain: {*true, false*}

--doValidation perform validation when SMAC completes

Aliases: --doValidation, --validation

Default Value: true

Domain: {*true, false*}

--executionMode execution mode of the automatic configurator

Aliases: --executionMode

Default Value: SMAC

Domain: {*SMAC, ROAR*}

--expectedImprovementFunction expected improvement function to use during local search

Aliases: --expectedImprovementFunction
Default Value: EXPONENTIAL
Domain: {*EXPONENTIAL, SIMPLE*}

--experimentDir root directory for experiments Folder

Aliases: --experimentDir, -e
Default Value: <current working directory>

--help show help

Aliases: --help, -?, /?, -h

--imputationIterations amount of times to impute censored data when building model

Aliases: --imputationIterations
Default Value: 2
Domain: [0, 2147483647]

--intensificationPercentage percent of time to spend intensifying versus model learning

Aliases: --intensificationPercentage, --frac_rawruntime
Default Value: 0.5
Domain: (0, 1)

--logLevel Log Level for SMAC

Aliases: --logLevel
Default Value: DEBUG
Domain: {*TRACE, DEBUG, INFO, WARN, ERROR, OFF*}

--maskInactiveConditionalParametersAsDefaultValue build the model treating inactive conditional values as the default value

Aliases: --maskInactiveConditionalParametersAsDefaultValue
Default Value: true
Domain: {*true, false*}

--maxIncumbentRuns maximum number of incumbent runs allowed

Aliases: --maxIncumbentRuns, --maxRunsForIncumbent

Default Value: 2000

Domain: (0, 2147483647]

--numChallengers number of challengers needed for local search

Aliases: --numChallengers, --numberOfChallengers

Default Value: 10

Domain: (0, 2147483647]

--numEIRandomConfigs number of random configurations to evaluate during EI search

Aliases: --numEIRandomConfigs, --numberOfRandomConfigsInEI, --numRandomConfigsInEI, --numberOfEIRandomConfigs

Default Value: 10000

Domain: [0, 2147483647]

--numIterations limits the number of iterations allowed during automatic configuration phase

Aliases: --numIterations, --numberOfIterations

Default Value: 2147483647

Domain: (0, 2147483647]

--numPCA number of principal components features to use when building the model

Aliases: --numPCA

Default Value: 7

Domain: (0, 2147483647]

--numRun number of this run (and seed)

REQUIRED

Aliases: --numRun, --seed

Default Value: 0

Domain: [0, 2147483647]

--optionFile read options from file

Aliases: --optionFile

--optionFile2 read options from file

Aliases: --optionFile2

--restoreStateFrom location of state to restore

Aliases: --restoreStateFrom

Default Value: N/A (No state is being restored)

--restoreStateIteration iteration of the state to restore

Aliases: --restoreStateIteration, --restoreIteration

Default Value: N/A (No state is being restored)

--runGroupName name of subfolder of outputdir to save all the output files of this run to

Aliases: --runGroupName

Default Value: RunGroup-<current date and time>

--runtimeLimit limits the total wall-clock time allowed during the automatic configuration phase

Aliases: --runtimeLimit, --wallClockLimit

Default Value: 2147483647

Domain: (0, 2147483647]

--seedOffset offset of numRun to use from seed (this plus --numRun should be less than LONG_MAX)

Aliases: --seedOffset

Default Value: 0

--showHiddenParameters show hidden parameters that no one has use for, and probably just break SMAC
(no-arguments)

Aliases: --showHiddenParameters

--stateDeserializer determines the format of the files that store the saved state to restore

Aliases: --stateDeserializer

Default Value: LEGACY

Domain: {*NULL*, *LEGACY*}

--stateSerializer determines the format of the files to save the state in

Aliases: --stateSerializer

Default Value: LEGACY

Domain: {*NULL*, *LEGACY*}

--totalNumRunsLimit limits the total number of target algorithm runs allowed during the automatic configuration phase

Aliases: --totalNumRunsLimit, --numRunsLimit, --numberOfRunsLimit

Default Value: 9223372036854775807

Domain: (0, 9223372036854775807]

--treatCensoredDataAsUncensored builds the model as-if the response values observed for cap values, were the correct ones [NOT RECOMMENDED]

Aliases: --treatCensoredDataAsUncensored

Default Value: false

Domain: {*true, false*}

-v print version and exit

Aliases: -v, --version

0.0.2 Scenario Options

Standard Scenario Options for use with SMAC. In general consider using the `--scenarioFile` directive to specify these parameters and Algorithm Execution Options

--checkInstanceFilesExist check if instances files exist on disk

Aliases: --checkInstanceFilesExist

Default Value: false

Domain: {*true, false*}

--cutoffTime CPU time limit for an individual target algorithm run

REQUIRED

Aliases: --cutoffTime, --cutoff_time

Default Value: 0.0

Domain: (0, ∞)

--instanceFeatureFile file that contains the all the instances features

Aliases: --instanceFeatureFile, --feature_file

--instanceFile file containing a list of instances to use during the automatic configuration phase (see Instance File Format section of the manual)

REQUIRED

Aliases: --instanceFile, -i, --instance_file, --instance_seed_file

Default Value: null

--interInstanceObj objective function used to aggregate over multiple instances (that have already been aggregated under the Intra-Instance Objective)

Aliases: --interInstanceObj, --inter_instance_obj

Default Value: MEAN

Domain: {*MEAN*, *MEAN1000*, *MEAN10*}

--intraInstanceObj objective function used to aggregate multiple runs for a single instance

REQUIRED

Aliases: --intraInstanceObj, --overallObj, --overall_obj, --intra_instance_obj

Default Value: null

Domain: {*MEAN*, *MEAN1000*, *MEAN10*}

--outputDirectory Output Directory

Aliases: --outputDirectory, --outdir

Default Value: <current working directory>/smac-output

--runObj per target algorithm run objective type that we are optimizing for

REQUIRED

Aliases: --runObj, --run_obj

Default Value: null

Domain: {*RUNTIME*, *QUALITY*}

--scenarioFile scenario file

Aliases: --scenarioFile

--testInstanceFile file containing a list of instances to use during the validation phase (see Instance File Format section of the manual)

REQUIRED

Aliases: --testInstanceFile, --test_instance_file, --test_instance_seed_file

Default Value: null

--tunerTimeout limits the total cpu time allowed between SMAC and the target algorithm runs during the automatic configuration phase

Aliases: --tunerTimeout

Default Value: 2147483647

Domain: [0, 2147483647]

-p File containing algorithm parameter space information (see Algorithm Parameter File in the Manual)

REQUIRED

Aliases: -p, --paramFile, --paramfile

Default Value: null

0.0.3 Algorithm Execution Options

Options related to running the target algorithm

--abortOnCrash treat algorithm crashes as an ABORT (Useful if algorithm should never CRASH). NOTE: This only aborts if all retries fail.

Aliases: --abortOnCrash

Default Value: false

Domain: {*true, false*}

--abortOnFirstRunCrash if the first run of the algorithm CRASHED treat it as an ABORT, otherwise allow crashes.

Aliases: --abortOnFirstRunCrash

Default Value: true

Domain: {*true, false*}

--algoExec command string to execute algorithm with

REQUIRED

Aliases: --algoExec, --algo

Default Value: null

--deterministic treat the target algorithm as deterministic

Aliases: --deterministic

Default Value: false

Domain: {*true, false*}

--execDir working directory to execute algorithm in

REQUIRED

Aliases: --execDir, --execdir

Default Value: null

--logAllCallStrings log every call string

Aliases: --logAllCallStrings

Default Value: false

Domain: {*true, false*}

--logAllProcessOutput log all process output

Aliases: --logAllProcessOutput

Default Value: false

Domain: {*true, false*}

--numConcurrentAlgoExecs maximum number of concurrent target algorithm executions

Aliases: --numConcurrentAlgoExecs, --maxConcurrentAlgoExecs, --numberOfConcurrentAlgoExecs

Default Value: 1

--retryTargetAlgorithmRunCount number of times to retry an algorithm run before reporting crashed (NOTE: The original crashes DO NOT count towards any time limits, they are in effect lost). Additionally this only retries CRASHED runs, not ABORT runs, this is by design as ABORT is only for cases when we shouldn't bother further runs

Aliases: --retryTargetAlgorithmRunCount

Default Value: 0

Domain: [0, 2147483647]

--runHashCodeFile file containing a list of run hashes one per line: Each line should be: "Run Hash Codes: (Hash Code) After (n) runs". The number of runs in this file need not match the number of runs that we execute, this file only ensures that the sequences never diverge. Note the n is completely ignored so the order they are specified in is the order we expect the hash codes in this version. Finally note you can simply point this at a previous log and other lines will be disregarded

Aliases: --runHashCodeFile

--targetAlgorithmEvaluator Target Algorithm Evaluator to use when making target algorithm calls

Aliases: --targetAlgorithmEvaluator, --tae

Default Value: CLI

--targetAlgorithmEvaluatorSearchPath location to look for other target algorithm evaluators [See manual but generally you can ignore this]

Aliases: --targetAlgorithmEvaluatorSearchPath, --taeSP

Default Value: <current working directory>/plugins/ among others

0.0.4 Random Forest Options

Options used when building the Random Forests

--fullTreeBootstrap bootstrap all data points into trees

Aliases: --fullTreeBootstrap

Default Value: false

Domain: {*true, false*}

--ignoreConditionality ignore conditionality for building the model

Aliases: --ignoreConditionality

Default Value: false

Domain: {*true, false*}

--logModel store response values in log-normal form

Aliases: --logModel

Default Value: true

Domain: {*true, false*}

--minVariance minimum allowed variance

Aliases: --minVariance

Default Value: 1.0E-14

Domain: (0, ∞)

--numTrees number of trees to create in random forest

Aliases: --numTrees, --nTrees, --numberOfTrees

Default Value: 10

Domain: (0, 2147483647]

--penalizeImputedValues treat imputed values that fall above the cutoff time, and below the penalized max time, as the penalized max time

Aliases: --penalizeImputedValues

Default Value: false

Domain: {*true, false*}

--preprocessMarginal build random forest with preprocessed marginal

Aliases: --preprocessMarginal

Default Value: true

Domain: {*true, false*}

--ratioFeatures ratio of the number of features to consider when splitting a node

Aliases: --ratioFeatures

Default Value: 0.8333333333333334

Domain: (0, 1]

--shuffleImputedValues shuffle imputed value predictions between trees

Aliases: --shuffleImputedValues

Default Value: false

Domain: $\{true, false\}$

--splitMin minimum number of elements needed to split a node

Aliases: --splitMin

Default Value: 10

Domain: [0, 2147483647]

--storeDataInLeaves store full data in leaves of trees

Aliases: --storeDataInLeaves

Default Value: false

Domain: $\{true, false\}$

0.0.5 Validation Options

Options that control validation

--maxTimestamp maximum relative timestamp in the trajectory file to configure against. -1 means auto-detect

Aliases: --maxTimestamp

Default Value: Auto Detect

Domain: $[0, \infty) \cup \{-1\}$

--minTimestamp minimum relative timestamp in the trajectory file to configure against.

Aliases: --minTimestamp

Default Value: 0.0

Domain: $[0, \infty)$

--multFactor base of the geometric progression of timestamps to validate (timestamps selected are: $\maxTime \times \text{multFactor}^{-n}$ where n is $\{1, 2, 3, 4, \dots\}$ while $\text{timestamp} \geq \text{minTimestamp}$)

Aliases: --multFactor

Default Value: 2.0

Domain: $(0, \infty)$

--numSeedsPerTestInstance number of test seeds to use per instance during validation

Aliases: --numSeedsPerTestInstance, --numberOfSeedsPerTestInstance

Default Value: 1000

Domain: (0, 2147483647]

--numTestInstances number of instances to test against (will execute min of this, and number of instances in test Instance File)

Aliases: --numTestInstances, --numberOfTestInstances

Default Value: 2147483647

Domain: (0, 2147483647]

--numValidationRuns approximate number of validation runs to do

Aliases: --numValidationRuns, --numberOfValidationRuns

Default Value: 1000

Domain: (0, 2147483647]

--validateOnlyLastIncumbent validate only the last incumbent found

Aliases: --validateOnlyLastIncumbent

Default Value: true

Domain: {*true, false*}

--validationHeaders put headers on output CSV files for validation

Aliases: --validationHeaders

Default Value: true

Domain: {*true, false*}

--validationRoundingMode selects whether to round the number of validation (to next multiple of numTestInstances)

Aliases: --validationRoundingMode

Default Value: UP

Domain: {*UP, NONE*}