

Use of function `tuneParameters` in `SamplingStrata` package

Package version 1.0.1

Giulio Barcaroli

January 9, 2013

Abstract

The package `SamplingStrata` allows to determine the best stratification of a target population, the one that ensures the minimum sample cost necessary to satisfy precision constraints in a multivariate and multidomain case. The underlying algorithm is based on a non deterministic evolutionary approach, making use of the genetic algorithm paradigm. The performance of the optimization heavily depends on the values given to the different parameters of the function `optimizeStrata`. In order to let the user choose a convenient combination of these values, the function `tuneParameters` available in the package permits to run a number of optimization tasks by changing in a planned way the parameters values, and to compare the results in order to choose the best combination.

1 Tuning of parameters and improvement of solutions

The performance of the optimization step depends highly on the values of the parameters that are given to the function `optimizeStrata`. It is worth while to run this step a number of times, by varying the values of the parameters in a planned way, and then to analyze the obtained results in order to choose the best values of the parameters. The function `tuneParameters` allows to run many optimization steps and compare the results in terms of the minimal sample cost required, and also in terms of the CV's that may be expected once having chosen a particular solution. `tuneParameters` can be applied only to a given domain per time. The parameters of this function are of the same nature than those of the function `optimizeStrata`, but they are given in a vectorial format, whose length is given by the number of optimizations to be run: it is therefore possible to define different combination of values of the parameters for each execution of `optimizeStrata`. After each optimization run, from the corresponding optimized frame a given number of samples are drawn. For each of them, the estimates of the target variables Y's are computed, and the absolute differences between the values of the estimates and the true values in the population, in order to permit a compared evaluation of the different solutions found in the different runs. As the optimal solution is stored for each run, after the evaluation it is possible to use it directly, or as a "suggestion" for a new optimization with more iterations (in order to improve it). First, we have to prepare the data, taking into account the fact that it is possible to apply the function `tuneParameters` to a single domain per time:

```
> library(SamplingStrata)
> data(swissstrata)
> data(swisserrors)
> data(swissframe)
> frame <- swissframe[swissframe$domainvalue == 1,]
> strata <- swissstrata[swissstrata$DOM1 == 1,]
> errors <- swisserrors[swisserrors$domainvalue == 1,]
```

In this case, we have selected the first domain from the `swissframe` dataframe and, consequently, the related strata and constraints on the CV's. Then, we have to set the different parameters necessary to the function. The first two parameters are peculiar to this function:

1. `noptim` defines the number of different optimization runs;
2. `nsampl` indicates how many samples have to be drawn from the population frame whose stratification has been optimized by the current solution.

The remaining parameters are the same utilized by the function `optimizeStrata`. The only difference is that they must be passed in a vectorial format, whose length is determined by the number of optimization runs. In this way, in

the i -th run `optimizeStrata` will receive the combination of parameters values corresponding to the i -th position in all vectors. It is possible to give the same values of some of the parameters for all the runs. In the example, only the parameter related to the `initialStrata` is varying, from a minimum of 10 to a maximum of all of them. We choose to do so because many experiences showed that the optimization step is very sensitive to the values of this parameter. The other parameter which is very important is the 'mutation chance'. So, we define the values of the parameters in this way:

```
> # Number of runs
> noptim <- 8
> # Number of samples to be drawn after each optimization
> nsampl <- 500
> # Number of initial strata
> initialStrata <- ceiling(c(1:noptim)*0.1*(nrow(strata)))
> # Rate for increasing the number of initial strata
> addStrataFactor <- rep(0.01,noptim)
> # Minimum number of units per stratum
> minnumstr <- rep(2,noptim)
> # Number of iterations for each optimization
> iter <- rep(200,noptim)
> # Number of solutions for each iteration
> pops <- rep(20,noptim)
> # Mutation chance
> mut_chance <- rep(0.004,noptim)
> # Elitism rate
> elitism_rate <- rep(0.2,noptim)
```

The function is invoked in this way:

```
> tuneParameters (
+           noptim,
+           nsampl,
+           frame,
+           errors,
+           strata,
+           cens = NULL,
+           strcens = FALSE,
+           alldomains = FALSE,
+           dom = 1,
+           initialStrata,
+           addStrataFactor,
+           minnumstr,
+           iter,
+           pops,
+           mut_chance,
```

```
+          elitism_rate  
+          )
```

Input data have been checked and are compliant with requirements

Input data have been checked and are compliant with requirements

GA Settings

```
Population size      = 20  
Number of Generations = 200  
Elitism              = 4  
Mutation Chance      = 0.004
```

[1] TRUE




Figure 1: Results of the different optimizations in terms of number of strata and minimal sample cost required varying the values of the parameter 'initial strata'

As already said, for each optimization a number of samples are drawn from the frame stratified accordingly to the corresponding optimal stratification: this allows to calculate the sampling variance for each target variable, expressed in terms of their CV's. The computed CV's are stored in the file 'results.csv', together with the number of strata and the required sample cost:

```
> results <- read.csv("results_1.csv")
> results
```

	nsimul	nstrati	cost	CV1	CV2	CV3
1	1	12	45	0.07634010	0.07156764	0.07550874
2	2	22	46	0.07139469	0.06982127	0.07032400
3	3	28	54	0.06440841	0.06337302	0.06625533
4	4	31	64	0.06505086	0.06543973	0.06443282
5	5	29	62	0.07336880	0.07221909	0.07100968
6	6	32	63	0.06586622	0.06217473	0.06488736
7	7	30	63	0.06378003	0.06077725	0.06093486
8	8	31	58	0.06144615	0.06142480	0.06279591

```
CV4
1 0.08056853
2 0.07910664
```

```
3 0.06637331
4 0.06769633
5 0.07134749
6 0.06659784
7 0.06075709
8 0.06332415
```

```
> ind <- which(results$cost==min(results$cost))
```

The value of 'ind' indicates the best found solution, and from this we can recall the associated values of the parameters:

```
-----
```

```
Best optimization: run # 1
```

```
Required sample cost: 45
```

```
-----
```

```
Values of parameters
```

```
-----
```

```
initialStrata: 12
```

```
addStrataFactor: 0.01
```

```
minnumstr: 2
```

```
iter: 200
```

```
pops: 20
```

```
mut_chanc: 0.004
```

```
elitism_rate: 0.2
```

```
-----
```

Moreover, for each sample, the difference between the sample estimates and the true values of the parameters in the population are also calculated. The distributions of these differences are reported in figure 2.

By comparing the performance of the different optimization runs, the best solutions in terms of sample cost required to comply the constraints on the maximum CV's, can be found. It is related to a precise value of the parameter *initial strata*. We can decide to assume this value as the best one, and then proceed to the tuning of the parameter *mutation chance*. So, we leave unchanged the previous values of the other parameters (with the exception of *initial strata*), and let the values of *mutation chance* vary:




Figure 2: Distributions of the differences between the sampling estimates and the true value of the parameter in the population under each optimization run and for each target variable Y

```

> # Number of initial strata
> bestInitialStrata <- initialStrata[ind]
> initialStrata <- rep(bestInitialStrata,noptim)
> # Mutation chance
> mut_chance <- c(1:noptim)*0.002

```

Then we run again the function tuneParameters:

```

> tuneParameters (
+     noptim,
+     nsampl,
+     frame,
+     errors,
+     strata,
+     cens = NULL,
+     strcens = FALSE,
+     alldomains = FALSE,
+     dom = 1,
+     initialStrata,
+     addStrataFactor,
+     minnumstr,
+     iter,
+     pops,
+     mut_chance,
+     elitism_rate
+ )

```

Input data have been checked and are compliant with requirements

Input data have been checked and are compliant with requirements

GA Settings

```

Population size      = 20
Number of Generations = 200
Elitism              = 4
Mutation Chance      = 0.002

```

[1] TRUE




Figure 3: Results of the different optimizations in terms of number of strata and minimal sample size required varying the values of the parameter 'mutation chance'

These are the results:

	n simul	n strati	cost	CV1	CV2	CV3
1	1	12	50	0.07212738	0.07050995	0.07230257
2	2	12	45	0.07504061	0.07164662	0.07442796
3	3	13	41	0.07521874	0.07182623	0.07343629
4	4	12	44	0.06987886	0.06881587	0.06862635
5	5	14	38	0.07384147	0.07015846	0.07386063
6	6	12	34	0.07081392	0.06906242	0.06778959
7	7	13	39	0.07454734	0.06972845	0.07302507
8	8	13	36	0.07847949	0.07582692	0.07470683

CV4

1	0.07432539
2	0.08179287
3	0.07898535
4	0.07633604
5	0.07775678
6	0.07549415
7	0.07539589
8	0.07984569

```
-----  
Best optimization: run # 6
```

```
Required sample cost: 34  
-----
```

```
Values of parameters  
-----
```

```
initialStrata: 12
```

```
addStrataFactor: 0.01
```

```
minnumstr: 2
```

```
iter: 200
```

```
pops: 20
```

```
mut_chance: 0.012
```

```
elitism_rate: 0.2  
-----
```

So, we can read the best value of the parameter *mutation chance*. At this point we can decide to stop the tuning of the parameters (actually, the two that we considered are the most important, the others being much less influent on the results). As an option, we can improve the best solution so far obtained, by running again the function `optimizeStrata` while increasing the number of iterations. To speed up the convergence of the genetic algorithm, we give the solution previously obtained as a *suggestion* to the genetic algorithm. The solution is read from the `.txt` external file written at the end of the `ind`-th optimization run carried out by `tuneParameters`, and introduced in the first row of a matrix (this is the format required by the internal genetic algorithm).

```
> stmt <- paste("suggestions <- read.table(solution_dom1_iter",ind,".txt)",sep="")  
> eval(parse(text=stmt))  
> sugg <- matrix(nrow=1,ncol=nrow(suggestions),data=suggestions[,1])
```

Then, a new optimization step with the same values of the parameters, but with a greater number of iterations (the double), and with the use of the parameter `suggestion`, is run:

```

> bestMutationChance <- mut_chance[ind]
> solution <- optimizeStrata(
+   errors ,
+     strata ,
+     cens = NULL,
+     strcens = FALSE,
+     alldomains = FALSE,
+     dom = 1,
+     initialStrata = bestInitialStrata,
+     addStrataFactor = 0.01,
+     minnumstr = 2,
+     iter = 1000,
+     pops = 20,
+     mut_chance = bestMutationChance,
+     elitism_rate = 0.2,
+     highvalue = 1e+08,
+     suggestions = sugg,
+     realAllocation = TRUE,
+     writeFiles = TRUE
+   )

```

Input data have been checked and are compliant with requirements

GA Settings

```

Population size      = 20
Number of Generations = 1000
Elitism              = 4
Mutation Chance     = 0.012

```

Suggestions

```

1 = (11, 7, 3, 2, 11, 11, 4, 11, 7, 10, 6, 1, 3, 12, 1, 10, 10, 9, 9, 10, 9, 1, 6, 10, 9,

```




Figure 4: Convergence of the new optimization run

As we can see from the convergence graph, we have improved the previous solution. As usual, we can now proceed with the next operations (updating of the frame and selecting the sample):

```
> newstrata <- updateStrata(strata,solution)
> framew <- updateFrame(frame,newstrata)
> samp <- selectSample(framew,solution$aggr_strata)
```

```
*** Sample has been drawn successfully ***
 24 units have been selected from 12 strata
```

```
==> There have been 1 take-all strata
from which have been selected 2 units
```

Of course, we have to repeat these operations for all the different domains in the frame.