



CEoptim: Cross-Entropy R Package for Optimization

Tim Benham	Qibin Duan	Dirk P. Kroese	Benoît Liquet
The University of Queensland, Brisbane	The University of Queensland, Brisbane	The University of Queensland, Brisbane	The University of Queensland, Brisbane

Abstract

The cross-entropy (CE) method is a simple and versatile technique for optimization, based on Kullback-Leibler (or cross-entropy) minimization. The method can be applied to a wide range of optimization tasks, including continuous, discrete, mixed and constrained optimization problems. The new package **CEoptim** provides the R implementation of the CE method for optimization. We describe the general CE methodology for optimization and well as some useful modifications. The usage and efficacy of **CEoptim** is demonstrated through a variety of optimization examples, including model fitting, combinatorial optimization, and maximum likelihood estimation.

Keywords: constrained optimization, continuous optimization, cross-entropy, discrete optimization, Kullback-Leibler divergence, lasso, maximum likelihood, R, regression.

1. Introduction

The cross-entropy (CE) method originates from an adaptive variance minimization algorithm in [Rubinstein \(1997\)](#) for the estimation of rare event probabilities in stochastic networks. It was realized in [Rubinstein \(1999\)](#) that many optimization problems could be converted into rare-event estimation problems, providing a rare-event based approach to optimization, where a sequence of probability densities is generated that converges to a degenerate density that concentrates its mass close to the optimizer.

Generally, the CE method involves two iterative phases:

1. Generation of a set of random samples (vectors, trajectories, etc.) according to a specified parameterized model.
2. Updating of the model parameters, based on the best samples generated in the previous step. This is done by Kullback-Leibler (also called cross-entropy) minimization.

Since the appearance of the CE monograph (Rubinstein and Kroese 2004) and tutorial (De Boer, Kroese, Mannor, and Rubinstein 2005), the CE method has continued to develop and has been successfully applied to a great variety of difficult optimization problems, including motion planning in robotic systems (Kobilarov 2012), electricity network generation, (Kothari and Kroese 2009), control of infectious diseases (Sani and Kroese 2008), buffer allocation (Alon, Kroese, Raviv, and Rubinstein 2005), Laguerre tessellation (Duan, Kroese, Brereton, Spetl, and Schmidt 2014), and network reliability (Kroese, Hui, and Nariai 2007). An extensive list of recent work can be found in Botev, Kroese, Rubinstein, and L'Ecuyer (2013). Websites that provide MATLAB (The MathWorks Inc. 2014) code include <http://www.cemethod.org/> and <http://www.montecarlohandbook.org/>. Since R (R Core Team 2016) has become an essential tool for statistical computation, it is useful to provide an accessible implementation of the CE method for R users, similar to R packages for simulated annealing (Xiang, Gubian, Suomela, and Hoeng 2013), evolutionary methods (Mullen, Ardia, Gil, Windover, and Cline 2011), and particle swarm optimization methods (Bendtsen 2012). Some advantages of the CE method are:

- The CE method is a global optimization method which is particularly useful when the objective function has many local optima.
- The CE method can be used to solve continuous, discrete, and mixed optimization problems, which may also include constraints.
- The CE code is extremely compact and is readily written in native R, making further development and modifications easy to implement.
- The CE method is based on rigorous mathematical and statistical principles.

Our aim is not to replace the standard optimization solvers such as `optim` and `nlm` but to provide a viable alternative in cases where standard gradient or simplex-based solvers are not applicable (e.g., when the optimization problem contains both discrete and continuous variables) or are expected to do poorly (e.g., when there are many local optima).

The rest of this paper is organized as follows. In Section 2, we sketch the general theory behind the CE method, which leads to the basic CE algorithm. In Section 3, we describe a variety of optimization scenarios, including continuous, discrete and constrained mixed problems, to which CE can be applied effectively. The description and usage of the **CEoptim** package (Benham, Duan, Kroese, and Lique 2017) are given in Section 4. Section 5 demonstrates the capability of the package through a range of numerical examples. In the final section we make concluding remarks for **CEoptim**.

2. CE method for optimization

Let \mathcal{X} be an arbitrary set of *states* and let S be a real-valued performance function on \mathcal{X} . Suppose the goal is to find the minimum of S over \mathcal{X} , and the corresponding minimizer \mathbf{x}^* (assuming, for simplicity, that there is only one). Denote the minimum by γ^* , so that

$$S(\mathbf{x}^*) = \gamma^* = \min_{\mathbf{x} \in \mathcal{X}} S(\mathbf{x}). \quad (1)$$

The CE methodology for optimization is adapted from the CE methodology for rare event estimation in the following way. Associate with the above problem (1) the estimation of the probability $\ell = \mathbb{P}(S(\mathbf{X}) \leq \gamma)$, where \mathbf{X} has some probability density $f(\mathbf{x}; \mathbf{u})$ on \mathcal{X} (for example corresponding to the uniform distribution on \mathcal{X}) depending on a parameter \mathbf{u} and a level γ . Thus, for optimization problems randomness is deliberately introduced in order to make the model stochastic. If γ is chosen close to the unknown γ^* , then ℓ is typically a rare-event probability. One of the most effective ways to estimate rare-event probabilities is to use *importance sampling*. In particular, to estimate $\ell = \mathbb{P}(S(\mathbf{X}) \leq \gamma)$ one can use the importance sampling estimator

$$\hat{\ell} = \frac{1}{N} \sum_{i=1}^N \frac{f(\mathbf{X}_i)}{g(\mathbf{X}_i)} \mathbf{I}\{S(\mathbf{X}_i) \leq \gamma\},$$

where $\mathbf{X}_1, \dots, \mathbf{X}_N$ are iid samples from a well-chosen importance sampling density g . The optimal importance sampling density is in this case $g^*(\mathbf{x}) = f(\mathbf{x})\mathbf{I}\{S(\mathbf{x}) \leq \gamma\}/\ell$, which gives a zero-variance estimator, but depends on the unknown quantity ℓ . The main idea behind the CE method for estimation is to adaptively determine an importance sampling pdf $f(\mathbf{x}; \mathbf{v}^*)$ – hence within the same family as the original distribution – that is close to g^* in Kullback-Leibler sense. Specifically, a parameter \mathbf{v}^* is sought that minimizes the cross-entropy distance

$$\mathcal{D}(g^*, f(\cdot; \mathbf{v})) = \mathbb{E}_{g^*} \left[\log \frac{g^*(\mathbf{X})}{f(\mathbf{X}; \mathbf{v})} \right] = \int g^*(\mathbf{x}) \log g^*(\mathbf{x}) \, d\mathbf{x} - \int g^*(\mathbf{x}) \log f(\mathbf{x}; \mathbf{v}) \, d\mathbf{x}.$$

This is equivalent to maximizing, with respect to \mathbf{v} ,

$$\int f(\mathbf{x}; \mathbf{u}) \mathbf{I}\{S(\mathbf{x}) \leq \gamma\} \log f(\mathbf{x}; \mathbf{v}) \, d\mathbf{x} = \mathbb{E}_{\mathbf{u}} [\mathbf{I}\{S(\mathbf{X}) \leq \gamma\} \log f(\mathbf{X}; \mathbf{v})],$$

which in turn can be estimated by maximizing the sample average

$$\frac{1}{N} \sum_{i=1}^N [\mathbf{I}\{S(\mathbf{X}_i) \leq \gamma\} \log f(\mathbf{X}_i; \mathbf{v})], \quad (2)$$

where $\mathbf{X}_1, \dots, \mathbf{X}_N$ is an iid sample from $f(\mathbf{x}; \mathbf{u})$. This is, in essence, maximum likelihood estimation. In particular, (2) gives the maximum likelihood estimator of \mathbf{v} based on only the samples $\mathbf{X}_1, \dots, \mathbf{X}_N$ that have a function value less than or equal to γ . These are the so-called *elite samples*.

The relevance to optimization is that when γ is close to the (usually unknown) minimum γ^* , then the importance sampling density g^* concentrates most of its mass in the vicinity of the minimizer \mathbf{x}^* . Sampling from such a distribution thus produces optimal or near-optimal states. The CE method for optimization produces a sequence of levels (γ_t) and reference parameters (\mathbf{v}_t) determined from (2) such that the former tends to the optimal γ^* and the latter to the optimal reference vector \mathbf{v}^* , where $f(\mathbf{x}; \mathbf{v}^*)$ corresponds to the point mass at \mathbf{x}^* ; see, e.g., [Rubinstein and Kroese \(2017, page 251\)](#).

The generic steps for CE optimization are specified in [Algorithm 1](#).

To run the algorithm, one needs to provide the class of sampling densities $\{f(\cdot; \mathbf{v})\}$, the initial vector \mathbf{v}_0 , the sample size N , the rarity parameter ρ , and the stopping criterion. It is prudent to keep track of the overall best function value and corresponding state, and report these at

Algorithm 1 Generic CE algorithm

Require: Initial parameter vector \mathbf{v}_0 . Sample size N . Rarity parameter ρ .

Ensure: Sequence of levels $(\gamma_t)_{t=1}^T$ and parameters $(\mathbf{v}_t)_{t=1}^T$.

- 1: Let $N^e = \lceil \rho N \rceil$ (number of elite samples) and set $t = 1$ (level counter).
- 2: **while** the sampling distribution is not degenerate **do**
- 3: Generate $\mathbf{X}_1, \dots, \mathbf{X}_N \sim_{\text{iid}} f(\cdot; \mathbf{v}_{t-1})$. Calculate the performances $S(\mathbf{X}_i)$ for all i , and order them from smallest to largest: $S_{(1)} \leq \dots \leq S_{(N)}$. Let γ_t be the sample ρ -quantile of performances; that is, $\gamma_t = S_{(N^e)}$.
- 4: Use the *same* sample $\mathbf{X}_1, \dots, \mathbf{X}_N$ and solve the stochastic program

$$\max_{\mathbf{v}} \sum_{k=1}^N \mathbb{I}\{S(\mathbf{X}_k) \leq \gamma_t\} \log f(\mathbf{X}_k; \mathbf{v}). \quad (3)$$

Denote the solution by \mathbf{v}_t . Increase t by 1.

5: **end while**

the end of the algorithm as the optimal value and optimizer, respectively. The progression of level parameter γ_t gives an indication how well the algorithm converges.

As (3) is simply a maximum likelihood estimation step involving only the elite samples, it is possible to derive easy parameter updates for standard sampling distributions. The following two special cases are of particular importance.

1. *Multivariate normal distribution.*

Suppose each \mathbf{X} is sampled from an n -dimensional multivariate normal distribution with independent components. The parameter vector \mathbf{v} in the CE algorithm can be taken as the $2n$ -dimensional vector of means and standard deviations. In each iteration these means and standard deviations are updated according to the sample mean and sample standard deviation of the elite samples.

2. *Multivariate Bernoulli distribution.*

Suppose each \mathbf{X} is sampled from an n -dimensional Bernoulli distribution with independent components. The parameter vector \mathbf{v} in the CE algorithm can be taken as the n -dimensional vector of success probabilities. In each iteration the i th success probability is updated according to the mean number of successes (1s) at the i th position of the elite samples.

Remark 1 (Parameter Smoothing) Various modifications of the basic CE algorithm have been proposed in recent years. One such modification is *parameter smoothing*, where at the t th iteration the sampling parameter is updated via

$$\mathbf{v}_t = \alpha \tilde{\mathbf{v}}_t + (1 - \alpha) \mathbf{v}_{t-1}, \quad (4)$$

where $\tilde{\mathbf{v}}_t$ is the solution to (3) and $0 \leq \alpha \leq 1$ is a fixed smoothing parameter.

Smoothed updating can prevent the sampling distribution from converging too quickly to a sub-optimal degenerate distribution. This is especially relevant for the multivariate Bernoulli case where, once a success probability reaches 0 or 1, it can no longer change.

It is also possible to use different smoothing parameters for different components of the parameter vector (e.g., the means and the standard deviations). Note that the smoothing parameters for means and standard deviations should not be very small. When setting the smoothing parameters too small, the means would be improved very slowly and it would take much longer to run before the standard deviation converged.

Remark 2 (Choice of sampling densities) Although sampling distributions with independent components are the most convenient to use in a CE implementation, it is sometimes advantageous to consider more complex sampling models, such as mixture models. In this case the updating of parameters (maximum likelihood estimation) may no longer be trivial, but one can instead employ fast methods such as the EM algorithm to determine the parameter updates.

Remark 3 (Choice of the CE parameters) The CE method is fairly robust with respect to the choice of the parameters. The rarity parameter ρ is typically chosen between 0.01 and 0.1. The number of elite samples $N^e = \lceil \rho N \rceil$ should be large enough to obtain a reliable parameter update in (3). For example, if the dimension of \mathbf{v} is d , the number of elites should be in the order of $10d$ or higher.

3. Optimization scenarios

In this section we consider a number of optimization scenarios to which **CEoptim** could be applied.

3.1. Continuous optimization

Consider a continuous optimization problem with state space $\mathcal{X} = \mathbb{R}^n$. The sampling distribution on \mathbb{R}^n can be quite arbitrary and does not need to be related to the objective function S . Usually, the random vector $\mathbf{X} = (X_1, \dots, X_n)^\top \in \mathbb{R}^n$ is generated from a Gaussian distribution with independent components, characterized by a vector $\boldsymbol{\mu}$ of means and a vector $\boldsymbol{\sigma}$ of standard deviations. At each iteration of the CE method, these vectors of parameters are updated as the means and standard deviations of the elite samples. During the course of the algorithm a sequence of $(\boldsymbol{\mu}_t)$ and $(\boldsymbol{\sigma}_t)$ are generated, such that $\boldsymbol{\mu}_t$ tends to the optimizer \mathbf{x}^* , while the vector of standard deviations tends to the zero vector. At the end of the algorithm one should obtain a degenerate probability density with mean $\boldsymbol{\mu}_T$ approximately equal to the optimizer \mathbf{x}^* and all standard deviations close to 0. A possible stopping criterion is to stop when all components in $\boldsymbol{\sigma}_T$ are smaller than some ε . This scheme is referred to as *normal updating*.

CEoptim implements the normal updating scheme for continuous optimization.

3.2. Discrete optimization

If the state space \mathcal{X} is finite, the optimization problem is often referred to as a discrete or combinatorial optimization problem, where \mathcal{X} could be the space of combinatorial objects, such as binary vectors, trees, graphs, etc. To apply the CE method to a discrete optimization problem, one needs a conveniently parameterized random mechanism to generate samples.

For discrete optimization **CEoptim** implements sampling from state spaces \mathcal{X} of the form $\{0, 1, \dots, c_1 - 1\} \times \dots \times \{0, 1, \dots, c_n - 1\}$, where the $\{c_i\}$ are strictly positive integers. The components of the random vector $\mathbf{X} = (X_1, \dots, X_n)^\top \in \mathcal{X}$ are taken to be independent, so that its distribution is determined by a sequence of probability vectors $\mathbf{p}_1, \dots, \mathbf{p}_n$, with the j th component of \mathbf{p}_i corresponding to $p_{ij} = \mathbb{P}(X_i = j)$. For a given elite sample set \mathcal{E} of size N^e , the CE updating formulas for these probabilities are

$$p_{ij} = \frac{\sum_{\mathbf{x} \in \mathcal{E}} \mathbb{I}\{X_i = j\}}{N^e}, \quad i = 1, \dots, n, \quad j = 0, \dots, c_n - 1, \quad (5)$$

where \mathbb{I} denotes the indicator function. Hence at each iteration probability p_{ij} is updated simply as the average number of times that the i th component of the elite vectors is equal to j . A possible stopping rule for a discrete optimization problem is to stop when the overall best objective value does not change over a number of iterations. Alternatively, one could stop when the sampling distribution has degenerated sufficiently; for example, when all $\{p_{ij}\}$ are no further than ϵ away from either 0 or 1.

3.3. Constrained optimization

The general optimization problem (1) also covers constrained optimization, where the search space \mathcal{X} could, for example, be defined by a system of inequalities:

$$G_i(\mathbf{x}) \leq 0, \quad i = 1, \dots, k. \quad (6)$$

One way to deal with constraints is to use *acceptance-rejection*: Generate a random vector \mathbf{X} on a simple search space that contains \mathcal{X} , and accept or reject it based on whether the sample falls into \mathcal{X} or not. Alternatively, one could try to sample directly from a truncated distribution on \mathcal{X} , e.g., using Gibbs sampling.

CEoptim implements linear constraints for continuous optimization of the form $A\mathbf{x} \leq \mathbf{b}$, where A is a matrix and \mathbf{b} a vector. The program will use either acceptance-rejection or Gibbs sampling to sample from the multivariate normal distribution truncated to the constraint set.

A second approach to handle constraints is to introduce a *penalty function*. For example, for the constraints (6), the objective function could be modified to

$$\tilde{S}(\mathbf{x}) = S(\mathbf{x}) + \sum_{i=1}^k H_i \max\{G_i(\mathbf{x}), 0\}, \quad (7)$$

where $H_i < 0$ measures the importance of the i th penalty. To use the penalty approach with **CEoptim** the user simply needs to modify the objective function according to (7). The choice of the penalty constants $\{H_i\}$ is problem specific and may need to be determined by trial and error.

4. CEoptim description

In this section we describe how to use **CEoptim**. The `CEoptim` function is the main function of the package **CEoptim**. It can be used to solve continuous and discrete optimization problems as well as mixtures thereof.

4.1. Usage

```
CEoptim(f, f.arg = NULL, maximize = FALSE, continuous = NULL,
  discrete = NULL, N = 100L, rho = 0.1, iterThr = 1e4L, noImproveThr = 5,
  verbose = FALSE)
```

4.2. Arguments

f : Function to be optimized. Can have continuous and discrete arguments.

f.arg : List of additional fixed arguments passed to function **f**.

maximize : Logical value determining whether to maximize or minimize the objective function.

continuous : List of arguments for the continuous optimization part, consisting of:

- **mean**: Vector of initial means.
- **sd**: Vector of initial standard deviations.
- **smoothMean**: Smoothing parameter for the vector of means. Default value 1 (no smoothing).
- **smoothSd**: Smoothing parameter for the standard deviations. Default value 1 (no smoothing).
- **sdThr**: Positive numeric convergence threshold. Check whether the maximum standard deviation is smaller than **sdThr**. Default value 0.001.
- **conMat**: Coefficient matrix of linear constraint $\text{conMat } \mathbf{x} \leq \text{conVec}$.
- **conVec**: Value vector of linear constraint $\text{conMat } \mathbf{x} \leq \text{conVec}$.

discrete : List of arguments for the discrete optimization part, consisting of:

- **categories**: Integer vector which defines the allowed values of the categorical variables. The *i*th categorical variable takes values in the set $\{0, 1, \dots, \text{categories}(i) - 1\}$.
- **probs**: List of initial probabilities for the categorical variables. Defaults to equal (uniform) probabilities.
- **smoothProb**: Smoothing parameter for the probabilities of the categorical sampling distribution. Default value 1 (no smoothing).
- **probThr**: Positive numeric convergence threshold. Check whether all probabilities in the categorical sampling distributions deviate less than **probThr** from either 0 or 1. Default value 0.001.

N : Integer representing the CE sample size.

rho : Value between 0 and 1 representing the elite proportion.

iterThr : Termination threshold on the largest number of iterations.

`noImproveThr` : Termination threshold on the largest number of iterations during which no improvement of the best function value is found.

`verbose` : Logical value set for CE progress output.

4.3. Value

`CEoptim` returns a list with the following components.

`optimum` : Optimal value of `f`.

`optimizer` : The location of the optimal value. A list consisting of:

- `continuous`: Continuous part of the optimizer.
- `discrete`: Discrete part of the optimizer.

`termination` : Termination information. A list consisting of:

- `niter`: Total number of iterations upon termination.
- `nfe`: Total number of function evaluations.
- `convergence`: One of the following termination statements:
 - `Not converged`, if the number of iterations reaches `iterThr`.
 - `The optimum did not change for noImproveThr iterations`, if the best value has not improved for `noImproveThr` iterations.
 - `Variances converged`, otherwise.

`states` : List of intermediate results computed at each iteration. It consists of the iteration number (`iter`), the best overall value (`optimum`) and the worst value of the elite samples, (`gammat`). The means (`mean`) and maximum standard deviation (`maxSd`) of the elite set are also included for continuous cases, and the maximum deviations (`maxProbs`) of the sampling probabilities to either 0 or 1 are included for discrete cases.

`states.probs` : List of categorical sampling probabilities computed at each iteration. Will only be returned for discrete and mixed cases.

4.4. Note

- Although partial parameter passing is allowed outside lists, it is recommended that parameter names be specified in full. Parameters inside lists have to be specified completely.
- Because `CEoptim` is function internally using pseudo-random number generation it is useful to (1) set the seed for the random number generator (for testing purposes), and (2) investigate the quality of the results by repeating the optimization a number of times.

5. Numerical examples

The following examples illustrate the use, flexibility, and efficacy of the `CEoptim` function from the package `CEoptim`.

5.1. Unconstrained minimization of the griewank function

To illustrate unconstrained optimization with `CEoptim`, we consider the minimization of the *griewank* function, which is widely used to test the convergence of optimization algorithms. The griewank function of order n is defined as

$$S(\mathbf{x}) = 1 + \frac{1}{4000} \sum_{i=1}^n x_i^2 - \prod_{i=1}^n \cos\left(\frac{x_i}{\sqrt{i}}\right), \quad (8)$$

where $\mathbf{x} = (x_1, \dots, x_n)^\top$ takes values in some subset of \mathbb{R}^n . The function has many local minima with (in the unconstrained case) a global minimum at $\mathbf{x}^* = (0, \dots, 0)$ of $S(\mathbf{x}^*) = 0$.

Suppose we are interested in minimizing the griewank function of order 5. To solve the problem with `CEoptim`, using normal updating, we must specify the vector of initial means $\boldsymbol{\mu}_0$ and standard deviations $\boldsymbol{\sigma}_0$ of the 5-dimensional Gaussian sampling distribution. The initial sampling distribution should cover, roughly, the region where the minimizer is thought to lie. As an example we take $\boldsymbol{\mu}_0 = (5, 5, 5, 5, 5)$ and $\boldsymbol{\sigma}_0 = (20, 20, 20, 20, 20)$. The important point is that the standard deviations are chosen large enough. Since this is a 5-dimensional problem, a larger sample size might be required, for example, we take $N = 1000$. For the other parameters we take their default values.

```
R> library("CEoptim")
R> griewank <- function(X) {
+   return(1 + sum(X^2) / 4000 - prod(cos(X / sqrt(1:length(X))))))
+ }
R> set.seed(1234)
R> mu0 <- c(5, 5, 5, 5, 5); sigma0 <- c(20, 20, 20, 20, 20)
R> res <- CEoptim(griewank, continuous = list(mean = mu0, sd = sigma0),
+   rho = 0.1, N = 1000L, verbose = TRUE, noImproveThr = Inf)
R> res
```

The output of this implementation is as below:

```
Optimizer for continuous part:
-5.397428e-05 -0.0002113244 2.584773e-05 -0.0002639864 7.067943e-05
Optimum:
2.197385e-08
Number of iterations:
56
Total number of function evaluations:
56000
Convergence:
Variance converged
```

Remark 4 (Extraction and print method) The user could easily extract any output using:

```
R> res$optimum
```

```
[1] 2.197385e-08
```

```
R> res$optimizer$continuous
```

```
[1] -5.397428e-05 -0.0002113244 2.584773e-05 -0.0002639864 7.067943e-05
```

The `print` method for ‘`CEoptim`’ objects provides by default the main description of the ‘`CEoptim`’ object including: `optimizer`; `optimum`; `termination`. To get the `states` and `states.probs` outputs, one should set the corresponding arguments to `TRUE`. The user could also specify the output to print:

```
R> print(res, optimum = TRUE)
```

```
Optimum:
2.197385e-08
```

5.2. Constrained minimization of the griewank function

To illustrate constrained optimization with `CEoptim`, we wish to minimize the griewank function of order 2 over the triangle with vertex points (1, 4), (4, 0), and (8, 4); see Figure 1.

The constraint set can be written as the linearly constrained region $\{\mathbf{x} \in \mathbb{R}^2 : A\mathbf{x} \leq \mathbf{b}\}$ with

$$A = \begin{pmatrix} 0 & 1 \\ -1 & -1 \\ 1 & 1 \end{pmatrix} \text{ and } \mathbf{b} = \begin{pmatrix} 4 \\ -4 \\ 4 \end{pmatrix}.$$

To solve the problem with `CEoptim` we proceed as follows:

```
R> library("CEoptim")
R> set.seed(123)
R> griewank <- function(X) {
+   return(1 + sum(X^2) / 4000 - prod(cos(X / sqrt(1:length(X))))))
+ }
R> A <- rbind(c(0,1), c(-1,-1), c(1,-1))
R> b <- c(4, -4, 4)
R> res <- CEoptim(griewank, continuous = list(mean = c(0,0), sd = c(10,10),
+   conMat = A, conVec = b), rho = 0.1, N = 200L, verbose = TRUE,
+   noImproveThr = Inf)
R> cat("direct optimizer = ", res$optimizer$continuous, "\n")
R> cat("direct minimum = ", res$optimum, "\n")
```




Figure 1: Contour plot of the griewank function and the triangular constraint region. The optimal solution (indicated by a cross) lies on the boundary of the constraint region.

The corresponding output shows that the minimum is obtained at the boundary of the triangle.

```
direct minimizer = 3.139669 3.991955
direct minimum = 0.05685487
```

It is also possible to use a penalty approach for this problem. Here we take the penalty function

$$\tilde{S}(\mathbf{x}) = S(\mathbf{x}) + 100 \|\mathbf{Ax} - \mathbf{b}\|,$$

which can be implemented in the following way.

```
R> griewank.penalty <- function(X, A, b) {
+   fn <- griewank(X)
+   if (any(A %*% as.vector(X) > b)) {
+     penalty <- norm(A %*% as.vector(X) - b)
+     fn <- fn + 100 * penalty
+   }
+   return(fn)
+ }
```

The optimization now proceeds as follows (note that we have also changed `rho` and `N`):

```
R> set.seed(123)
R> res.pen <- CEoptim(griewank.penalty, f.arg = list(A, b),
+   continuous = list(mean = c(0, 0), sd = c(10, 10)), rho = 0.01,
+   N = 2000L, verbose = TRUE, noImproveThr = Inf)
```




Figure 2: Simulated data (points) and “unknown” true curve (red).

```
R> cat("penalty minimizer =", res.pen$optimizer$continuous, "\n")
R> cat("penalty minimum =", griewank(res.pen$optimizer$continuous), "\n")
```

This leads to practically the same result:

```
penalty minimizer = 3.139757 4
penalty minimum = 0.055103
```

5.3. Non-linear regression

We next consider a more complicated optimization task, involving data generated from the well-known *FitzHugh-Nagumo* differential equations:

$$\begin{aligned} \frac{dV_t}{dt} &= c \left(V_t - \frac{V_t^3}{3} + R_t \right), \\ \frac{dR_t}{dt} &= -\frac{1}{c} (V_t - a + bR_t), \end{aligned} \quad (9)$$

which model the behavior of certain types of neurons (Nagumo, Arimoto, and Yoshizawa 1962). Ramsay, Hooker, Campbell, and Cao (2007) consider estimating the parameters a , b , and c from noisy observations of (V_t) by using a generalized smoothing approach. The simulated data in Figure 2 (saved as `data("FitzHugh", package = "CEoptim")`) correspond to the values of V_t obtained from (9) at times $0, 0.05, \dots, 20.0$, adding Gaussian noise with standard deviation 0.5. That is, we use the non-linear regression model

$$Y_i = V_{0.05i}(\mathbf{x}) + \epsilon_i, \quad i = 1, \dots, 400, \quad (10)$$

where the $\{\epsilon_i\}$ are iid with a $N(0, \sigma^2)$ distribution, $V_{0.05i}(\mathbf{x})$ is the solution to (9) for time $t = 0.05i$, and $\mathbf{x} = (a, b, c, V_0, R_0)$ is the vector of parameters. The true parameter values are here $a = 0.2$, $b = 0.2$, and $c = 3$. The initial conditions are $V_0 = -1$ and $R_0 = 1$.

Estimation of the parameters via the CE method can be established by minimizing the least-squares performance

$$S(\mathbf{x}) = \sum_{i=0}^{400} (y_i - V_{0.05i}(\mathbf{x}))^2, \quad (11)$$

where the $\{y_i\}$ are the simulated data from the model (10). Note that we assume also that the initial conditions are unknown.

We use the **deSolve** package (Soetaert, Petzoldt, and Setzer 2010) to numerically solve the FitzHugh-Nagumo differential equations (9). Hereto, we first define the function FN.

```
R> FN <- function(t, state, parameters) {
+   with(as.list(c(state, parameters)), {
+     dV <- c * (V - V^3/3 + R)
+     dR <- -1 / c * (V - a + b * R)
+     list(c(dV, dR))
+   })
+ }
```

The following function `ssres` now implements the objective function in (11).

```
R> ssres <- function(x, fundf, times, y) {
+   parameters <- c(a = x[1], b = x[2], c = x[3])
+   state <- c(V = x[4], R = x[5])
+   out <- ode(y = state, times = times, func = fundf, parms = parameters)
+   return(sum((out[, 2] - y)^2))
+ }
```

`CEoptim` could be used with $\boldsymbol{\mu}_0 = (0, 0, 5, 0, 0)$ and $\boldsymbol{\sigma}_0 = (1, 1, 1, 1, 1)$. Constant smoothing parameters $\alpha = 0.9$ and $\beta = 0.5$ were used for the $\{\boldsymbol{\mu}_t\}$ and the $\{\boldsymbol{\sigma}_t\}$ respectively. To see the progress of the algorithm we set `verbose` to `TRUE`. The other arguments remain default.

```
R> library("deSolve")
R> library("CEoptim")
R> set.seed(123405)
R> times <- seq(0, 20, by = 0.05)
R> data("FitzHugh", package = "CEoptim")
R> res <- CEoptim(ssres, f.par = list(fundf = FN, times = times, y = ySim),
+   continuous = list(mean = c(0, 0, 5, 0, 0), sd = c(1, 1, 1, 1, 1),
+   smoothMean = 0.9, smoothSd = 0.5), verbose = TRUE)
```

The final output is as follows:

```
R> res
```

```
Optimizer for continuous part:
```

```
0.1959748 0.2395983 3.001453 -0.9938222 0.9791585
```

```
Optimum:
```

```
102.8005
```




Figure 3: The evolution of the sampling pdf for the first parameter a .

Number of iterations:

41

Total number of function evaluations:

4100

Convergence:

Variance converged

The output shows the estimates (notice that the initial condition was assumed to be unknown): $\hat{a} = 0.1959748$, $\hat{b} = 0.2395983$, $\hat{c} = 3.0014526$, $\hat{V}_0 = -0.9938222$, and $\hat{R}_0 = 0.9791585$, with the maximum likelihood estimate $\hat{\sigma} = \sqrt{102.8005/400} = 0.507$ for the residual standard deviation σ . The reader may confirm that the fitted curve is practically indistinguishable from the true one in Figure 2.

To illustrate how the sampling distributions change during the CE process, we have plotted in Figure 3 the evolution of the sampling pdf for the first parameter a , from the 15th to the final iteration. As can be seen from the figure, the sampling distribution converges to a point distribution around the optimal value for a .

5.4. Max-cut problem

The max-cut problem in graph theory can be formulated as follows. Given a weighted graph (V, E) with node set $V = \{1, \dots, n\}$ and edge set E , partition the nodes of the graph into two subsets V_1 and V_2 such that the sum of the (nonnegative) weights of the edges going from one subset to the other is maximized. Let $C = (C_{ij})$ be the matrix of weights. The objective is to maximize

$$\sum_{(i,j) \in V_1 \times V_2} (C_{ij} + C_{ji}) \quad (12)$$

over all cuts $\{V_1, V_2\}$. Such a cut can be conveniently represented by a binary *cut vector* $\mathbf{x} = (x_1, x_2, \dots, x_n)$, where $x_i = 1$ indicates that $i \in V_1$. Let \mathcal{X} be the set of cut vectors and let $S(\mathbf{x})$ be the value of the cut represented by \mathbf{x} , as given in (12).




Figure 4: Network of coappearances.

To maximize S via the CE method one can generate the random cut vectors by drawing each component (except the first one, which is set to 1) independently from a Bernoulli distribution, that is, $\mathbf{X} = (X_1, X_2, \dots, X_n) \sim \text{Ber}(\mathbf{p})$, where $\mathbf{p} = (1, p_2, \dots, p_n)$. In this case the updated success probability for the i th component is the mean of the i th components of the vectors in the elite set.

As an example, consider the network from Knuth (1993) describing the coappearances of 77 characters from Victor Hugo's novel *Les Misérables*. Each node of the network represents a selected character and edges connect any pair of characters who coappear. The weights of the edges are the number of such coappearances. Using **CEoptim**, the data can be loaded via the command `data("lesmis", package = "CEoptim")`. The network is displayed in Figure 4, using the graph analysis package **sna** (Butts 2008).

```
R> library("sna")
R> library("CEoptim")
R> data("lesmis")
R> gplot(lesmis, gmode = "graph")
```

For any fixed cost matrix `costs` and cut vector `x`, the objective function of the max-cut problem can be written as:

```
R> fmaxcut <- function(x, costs) {
+   v1 <- which(x == 1)
+   v2 <- which(x == 0)
```

```
+   return(sum(costs[v1, v2]))
+ }
```

To optimize this function with the **CEoptim** package, we specify the following arguments: `discrete$probs = {(0,1); (0.5,0.5);...;(0.5,0.5)}`, sample size `N = 3000` and optimization type: `maximize = TRUE`. To see the output we set `verbose = TRUE`. The other arguments are taken as default. Note that users only need to specify either `categories` or `probs`; if both of them are specified, then `categories` will be overridden.

```
R> set.seed(5)
R> p0 <- list()
R> for (i in 1:77) {
+   p0 <- c(p0, list(rep(0.5, 2)))
+ }
R> p0[[1]] <- c(0, 1)
R> res <- CEoptim(fmaxcut, f.arg = list(costs = lesmis), maximize = TRUE,
+   verbose = TRUE, discrete = list(probs = p0), N = 3000L)
R> ind <- res$optimizer$discrete
R> group1 <- colnames(lesmis)[which(ind == TRUE)]
R> group2 <- colnames(lesmis)[which(ind == FALSE)]
```

The output of **CEoptim** is as follows:

```
R> res

Optimizer for discrete part:
 1 0 1 0 0 0 0 0 0 0 1 0 0 1 1 1 0 0 1 1 0 1 0
 1 0 1 1 1 1 0 0 1 1 1 1 1 1 0 0 0 0 1 0 1 0 0
 1 0 1 1 1 0 1 1 1 0 0 0 0 1 1 0 1 0 0 1 1 1 1
 0 0 1 0 1 0 0 0

Optimum:
535
Number of iterations:
20
Total number of function evaluations:
60000
Convergence:
Optimum did not change for 5 iterations
```

Note that character 1 (Myriel) is always in `group1`. The initial probabilities for the other characters are 0.5. With `states.probs`, we can plot the evolution of the probabilities that each character belongs to `group1`; see Figure 5.

```
R> probs <- res$states.probs
R> X <- matrix(NA, nrow = length(probs), ncol = 77)
R> prob0 <- cbind(1, t(rep(0.5, 76)))
R> for (i in 1:length(probs)) {
```

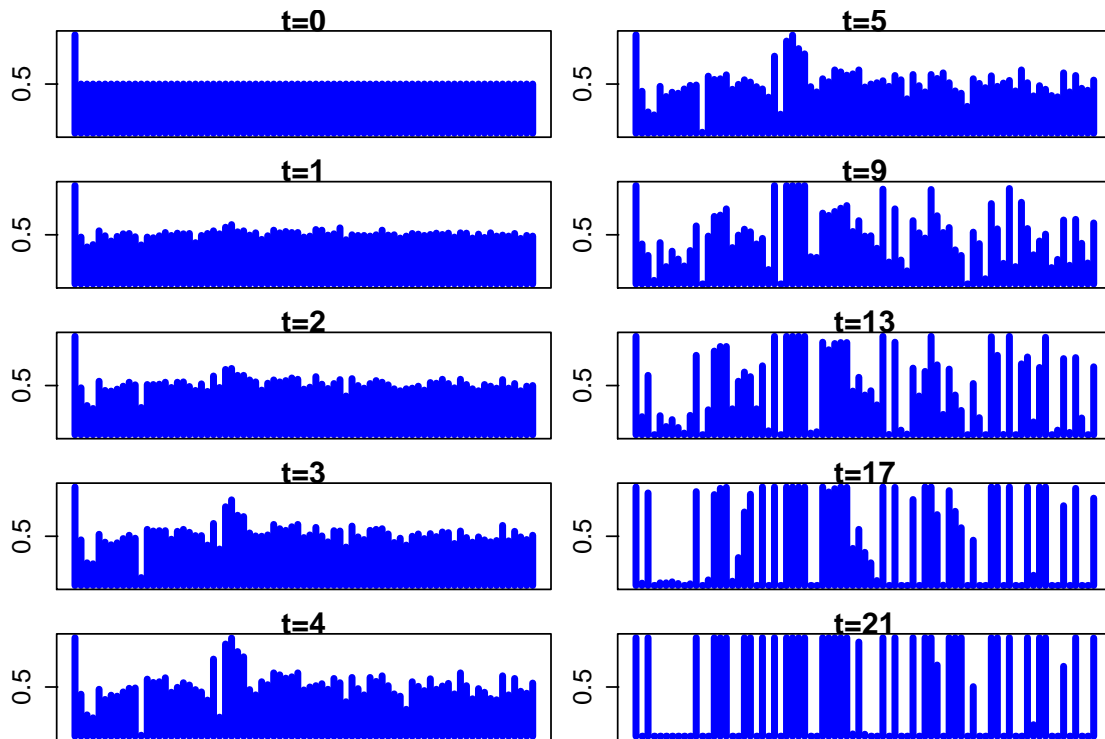



Figure 5: Evolution of categorical sampling probabilities that characters in group 1.

```

+   for (j in 1:77) {
+     X[i, j] <- res$states.probs[[i]][[j]][2]
+   }
+ }
R> X <- rbind(prob0, X)
R> par(mfcol = c(5, 2), mar = c(1, 1.5, 1, 1.5), oma = c(1, 1, 1, 1))
R> for (i in 1:5) {
+   plot(X[i, ], type = "h", lwd = 4, col = "blue", ylim = c(0, 1),
+     xaxt = "n", yaxt = "n", ylab = "",
+     main = paste("t=", i - 1, sep = ""))
+   axis(2, at = 0.5, labels = 0.5)
+ }
R> for (i in 1:5) {
+   plot(X[1 + 4 * i, ], type = "h", lwd = 4, col = "blue",
+     ylim = c(0, 1), xaxt = "n", yaxt = "n", ylab = "",
+     main = paste("t=", 1 + 4 * i, sep = ""))
+   axis(2, at = 0.5, labels = 0.5)
+ }

```

Based on the output above, the two groups of characters are indicated in Table 1: We ran the program 1000 times with different initial random number generator seeds. In 312 cases the optimal solution (535) was found. The frequency of the results of `CEoptim` is given in Figure 6.

group1	group2
Myriel, MlleBaptistine, Labarre, MmeDeR, Isabeau, Gervais, Fameuil, Blacheville , Dahlia, Fantine, Thenardier, Cosette, Javert, Fauchelevent, Simplicite, Scaufflaire, Oldwoman1, Judge, Champmathieu, Brevet, Eponine, Oldwoman2, Jondrette, Gavroche, Gillenormand, Magnon, MmePontmercy, MlleVaubois, LtGillenormand, Combeferre, Prouvaire, Courfeyrac, Joly, Grantaire, MotherPlutarch, Gueulemer, Montparnasse, Child1	Napoleon, MmeMagloire, CountessDeLo, Geborand, Champtercier, Cravatte, Count, OldMan, Valjean, Marguerite, Tholomyes, Listolier, Favourite, Zepherine, MmeThenardier, Bamatabois, Perpetue, Chenildieu, Cochepaille, Pontmercy, Boulatruelle, Anzelma, MotherInnocent, Gribier, MmeBurgon, MlleGillenormand, Marius, BaronessT, Mabeuf, Enjolras, Feuilly, Bahorel, Bossuet, Babet, Claquesous, Toussaint, Child2, Brujon, MmeHucheloup

Table 1: Two groups of characters providing a maximal cut.




Figure 6: Frequency of best max-cut values found by CEoptim.

5.5. Dirichlet data

Suppose that we are given a random sample of data from a Dirichlet(α) distribution, where $\alpha = (\alpha_1, \dots, \alpha_{K+1})^\top$ is an unknown parameter vector satisfying $\alpha_i > 0$, $i = 1, \dots, K + 1$. Recall that the pdf of a random vector $\mathbf{Y} = (Y_1, \dots, Y_K) \sim \text{Dirichlet}(\alpha)$ is given by

$$f(\mathbf{y}; \alpha) = \frac{\Gamma(\sum_{i=1}^{K+1} \alpha_i)}{\prod_{i=1}^{K+1} \Gamma(\alpha_i)} \prod_{i=1}^K y_i^{\alpha_i-1} \left(1 - \sum_{i=1}^K y_i\right)^{\alpha_{K+1}-1},$$

for $y_i \geq 0$, $i = 1, \dots, K$ and $\sum_{i=1}^K y_i \leq 1$, where Γ is the gamma function. The conditions on α provide natural inequality constraints: $G_i(\alpha) \equiv -\alpha_i \leq 0$, $i = 1, \dots, K + 1$.

We will use the CE method to obtain the maximum likelihood estimate by direct maximization of the log-likelihood for the Dirichlet distribution given the data.




Figure 7: Evolution of the best value (upper line) and the worst value of the best (elite) samples (lower line).

As a concrete example, a sample of size $n = 100$ points was drawn from the Dirichlet(1, 2, 3, 4, 5) distribution with the assistance of the function `dirichletrnd` in the **CEoptim** package.

```
R> library("CEoptim")
R> set.seed(12345)
R> a <- 1:5
R> K <- length(a) - 1
R> n <- 100
R> y <- dirichletrnd(a, n)
```

To use **CEoptim** to solve the MLE problem, the objective function is written as follows:

```
R> dirichletLoglike <- function(alpha, Y, n, K) {
+   t <- apply(Y, MARGIN = 1, function(y)
+     sum((alpha[1:K] - 1) * log(y[1:K])) +
+     (alpha[K + 1] - 1) * log(1 - sum(y[1:K])))
+   out <- n * (lgamma(sum(alpha)) - sum(lgamma(alpha))) + sum(t)
+   return(out)
+ }
```

The CE parameters are the initial mean vector $\boldsymbol{\mu} = (0, 0, 0, 0, 0)$ and the standard deviation vector $\boldsymbol{\sigma} = (10, 10, 10, 10, 10)$. The sample size is $N = 10^4$ and the elite ratio is the default. To pass the linear constraints $\alpha_i > 0, i = 1, \dots, K + 1$, the constraint coefficient matrix is

$$A = \begin{pmatrix} -1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 \end{pmatrix},$$

and the constraint vector is $\mathbf{b} = (0, 0, 0, 0, 0)$. No smoothing parameter is applied to the mean vector, but a constant smoothing parameter of `smoothSd = 0.5` is applied to each of the standard deviations. This is a maximization problem, so we set `maximize = TRUE`.

```
R> mu0 <- rep(0, times = K + 1)
R> sigma0 <- rep(10, times = K + 1)
R> A <- matrix(rep(0, times = 25), nrow = 5)
R> diag(A) <- rep(-1, times = 5)
R> b <- rep(0, times = 5)
R> res <- CEoptim(dirichletLoglike, f.arg = list(Y = y, n = 100, K = 4),
+   maximize = TRUE, continuous = list(mean = mu0, sd = sigma0,
+   conMat = A, conVec = b, smoothSd = 0.5), N = 10000L, verbose = TRUE)
```

With the returned `states` variable, we can plot the evolution of optimal values per iteration, as shown in Figure 7, where the upper line indicates the best value found so far, while the lower line gives the worst value of the current elite sample.

```
R> par(mai = c(0.6, 1, 0.5, 0.2), oma = c(0, 0, 0, 1))
R> plot(res$states[, "iter"], res$states[, "gammat"], type = "s",
+   col = "blue", xlab = "", ylab = "")
R> lines(res$states[, "optimum"], type = "s", col = "red")
R> res
```

Optimizer for continuous part:

```
1.111656 2.000186 3.534268 3.983616 5.142336
```

Optimum:

```
486.2124
```

Number of iterations:

```
35
```

Total number of function evaluations:

```
350000
```

Convergence:

```
Variance converged
```

Maximum likelihood estimates for Dirichlet data can be computed to a high accuracy via the fixed-point techniques of Minka (2012). This requires sophisticated numerical techniques for inverting digamma functions. When applying this method to the same Dirichlet(1, 2, 3, 4, 5) data, we obtained the estimate $\hat{\alpha} = (1.111715, 2.000243, 3.534321, 3.983752, 5.142596)$, with a likelihood value of 486.2124, giving excellent agreement between the two approaches.

5.6. Lasso regression

Suppose that we observed some data from the following model:

$$Y_i = \mathbf{x}_i^\top \boldsymbol{\beta} + \epsilon_i, \quad i = 1, \dots, n,$$

where $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^\top$ is the p -vector of explanatory variables, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^\top$ is the p -vector of regression coefficients, and the $\{\epsilon_i\}$ are the noise terms with $\mathbb{E}[\epsilon_i] = 0$, $\text{Var}[\epsilon_i] = \sigma^2$,

for all i and $\text{Cov}(\epsilon_i, \epsilon_j) = 0$ ($\forall i \neq j$). Consider a lasso regression approach to estimate the regression vector β :

$$\begin{aligned}\hat{\beta}^{\text{lasso}} &= \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \frac{1}{2n} \sum_{i=1}^n (Y_i - \mathbf{x}_i^\top \beta)^2 + \lambda \sum_{j=1}^p |\beta_j| \\ &= \underset{\beta \in \mathbb{R}^p}{\text{argmin}} \underbrace{\frac{1}{2n} \|\mathbf{Y} - \mathcal{X}\beta\|_2^2}_{\text{Loss}} + \lambda \underbrace{\|\beta\|_1}_{\text{Penalty}},\end{aligned}$$

where $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$ and $\mathcal{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top$ is the $(n \times p)$ design matrix. The tuning parameter λ controls the amount of regularization.

For a given value of λ , we will use the CE method to obtain the lasso regression coefficient and compare our results with those obtained by the function `glmnet` from the package `glmnet` presented by [Friedman, Hastie, and Tibshirani \(2008\)](#).

We generate a set of test data of size $n = 150$, with $p = 60$ explanatory variables independently generated from a standard normal distribution. The true coefficients β are chosen such that 10 are large (between 0.5 and 1) and 50 are exactly 0. The variance of the noise is equal to 1.

```
R> set.seed(10)
R> n <- 150
R> p <- 60
R> beta <- c(runif(10, 0.5, 1), rep(0, 50))
R> X <- matrix(rnorm(n * p), ncol = 60)
R> Y <- X %*% matrix(beta, ncol = 1) + rnorm(n)
```

We first use the `glmnet` function to find the lasso regression coefficient that gives a *sparsity* of 10; that is, exactly 10 coefficients are non-zero.

```
R> library("glmnet")
R> res.glmnet <- glmnet(X, Y)
R> sparsity.10 <- which(res.glmnet$df == 10)
R> (lambda.10 <- res.glmnet$lambda[sparsity.10[1]])
```

0.2731371

```
R> beta.glmnet <- res.glmnet$beta[, sparsity.10[1]]
```

The corresponding indices are correctly identified by `glmnet`:

```
R> (ind.beta <- which(res.glmnet$beta[, sparsity.10[1]] != 0))
```

```
V1 V2 V3 V4 V5 V6 V7 V8 V9 V10
 1  2  3  4  5  6  7  8  9  10
```

The values of the non-zero coefficients (NZ) are given by:

```
R> (beta.glmnet.NZ <- res.glmnet$beta[ind.beta, sparsity.10[1]])
```

	V1	V2	V3	V4	V5	V6
0.39006188	0.39345242	0.40795534	0.57510345	0.18776598	0.19553092	
	V7	V8	V9	V10		
0.02929225	0.55435619	0.57656731	0.56279719			

We now use CEoptim to estimate the lasso regression function for the given $\lambda = 0.2731371$.

```
R> library("CEoptim")
R> RSS.penalized <- function(x, X, Y, lambda) {
+   out <- (1/2) * mean((Y-X %*% matrix(x, ncol = 1, nrow = dim(X)[2],
+     byrow = TRUE))**2) + lambda * sum(abs(x))
+   return(out)
+ }
R> mu0 <- rep(0, times = p)
R> sigma0 <- rep(5, times = p)
R> N <- 1000
R> set.seed(1212)
R> res <- CEoptim(RSS.penalized, f.arg = list(X = X, Y = Y,
+   lambda = lambda.10), continuous = list(mean = mu0, sd = sigma0,
+   sdThr = 0.00001), N = N)
R> beta.CEoptim <- res$optimizer$continuous
R> (ind.beta.CEoptim.NZ <- which(abs(beta.CEoptim) > 0.000001))
```

```
[1] 1 2 3 4 5 6 7 8 9 10
```

```
R> beta.CEoptim.NZ <- beta.CEoptim[ind.beta.CEoptim.NZ]
R> (compare.beta.NZ <- rbind(beta.glmnet.NZ, beta.CEoptim.NZ))
```

	V1	V2	V3	V4	V5	V6
beta.glmnet.NE	0.3900619	0.3934524	0.4079553	0.5751035	0.1877660	0.1955309
beta.CEoptim.NE	0.3631798	0.3826273	0.4419025	0.6014707	0.1639559	0.1721753
	V7	V8	V9	V10		
beta.glmnet.NE	0.029292247	0.5543562	0.5765673	0.5627972		
beta.CEoptim.NE	0.005821685	0.5537388	0.5854710	0.6034473		

The two methods give similar values for the non-zero coefficients, although they are not exactly the same. Note, however, that of the two solutions the one found by CEoptim gives the smaller value for the objective function $\frac{\text{RSS}}{2n} + \lambda \|\beta\|_1$ (where residual sum of squares: $\text{RSS} = \|\mathbf{Y} - \mathcal{X}\beta\|_2^2$).

```
R> (RSS.penalized(beta.CEoptim, X = X, Y = Y, lambda = lambda.10))
```

```
[1] 1.990268
```

```
R> (RSS.penalized(beta.glmnet, X = X, Y = Y, lambda = lambda.10))
```

```
[1] 1.993622
```




Figure 8: Difference of the objective function values between `glmnet` and `CEoptim` for a sequence of 73 values of λ .

Further, we compare the results obtained by `CEoptim` with the ones given by `glmnet` for the sequence of values of the tuning parameter λ used by default in the `glmnet` function. Results given by `CEoptim` are slightly better than the `glmnet` optimizer (see Figure 8). In more than 90% of the cases (over the 73 values of λ investigated) `CEoptim` gives a lower value for the objective function than does `glmnet`. However, the coordinate descent algorithm (Friedman *et al.* 2008) used in `glmnet` is computationally less demanding than the CE approach.

5.7. AR(1) model with regime switching

As a final illustration of the use of `CEoptim`, we consider a model fitting problem involving both continuous and discrete variables.

Let Y_t be the added value of a stock at time t , at day $t = 1, 2, \dots, 300$; that is, the increase (which may be negative) in stock price relative to the price at time $t = 0$. Let X_t be the increment at day t . Hence,

$$Y_t = \sum_{i=1}^t X_i, \quad t = 1, \dots, 300.$$

We assume that the $\{X_i\}$ satisfy a zero-mean AR(1) model with three possibly different regimes. Specifically, we assume

$$X_i = \theta_i X_{i-1} + \epsilon_i, \quad i = 1, \dots, 300,$$

where

$$\theta_i = \begin{cases} \theta^{(1)}, & i = 1, \dots, r_1 \\ \theta^{(2)}, & i = r_1 + 1, \dots, r_2 \\ \theta^{(3)}, & i = r_2 + 1, \dots, 300, \end{cases} \quad (13)$$

$1 \leq r_1 < r_2 < 300$, $|\theta_i| \leq 1$, $i = 1, 2, 3$, and the error terms $\{\epsilon_i\}$ are iid and normally distributed with standard deviation σ . The model thus has two discrete and three continuous parameters, as well as a nuisance parameter σ . Define $\boldsymbol{\theta} = (\theta^{(1)}, \theta^{(2)}, \theta^{(3)})^\top$, $\mathbf{r} = (r_1, r_2)^\top$, and let x_1, \dots, x_{300} be the observed increments. We put $x_0 = 0$. We fit the parameters by

minimizing the least squares function

$$L(\boldsymbol{\theta}, \mathbf{r}) = \sum_{i=1}^{300} (x_i - \hat{x}_i)^2,$$

where \hat{x}_i is the fitted value $\theta_i x_{i-1}$, and θ_i is determined by $\boldsymbol{\theta}$ and \mathbf{r} via (13). The vector of fitted values, say $\hat{\mathbf{x}}$, can be written in matrix notation as $\hat{\mathbf{x}} = \mathcal{X}\boldsymbol{\theta}$, where \mathcal{X} is a 300×3 matrix where the elements in rows $1, \dots, r_1$ in the first column are equal to x_0, \dots, x_{r_1-1} ; the elements in rows $r_1 + 1, \dots, r_2$ in the second column are equal to $x_{r_1}, \dots, x_{r_2-1}$; the elements in rows $r_2 + 1, \dots, 300$ in the third column are equal to x_{r_2}, \dots, x_{299} ; and all other elements are 0. The implementation of the least squares function is given below. Note that the function requires input $\mathbf{r} - 1$ rather than \mathbf{r} , because each categorical variable used in **CEoptim** takes value in a set $\{0, \dots, c\}$ for some c .

```
R> sumsqrs <- function(theta, rm1, x) {
+   N <- length(x)
+   r <- 1 + sort(rm1)
+   if (r[1] == r[2]) {
+     return(Inf)
+   }
+   thetas <- rep(theta, times = c(r, N) - c(1, r + 1) + 1)
+   xhat <- c(0, head(x, -1)) * thetas
+   sum((x - xhat)^2)
+ }
```

The data were generated using the parameters $\boldsymbol{\theta} = (0.3, 0.9, -0.9)$, $\mathbf{r} = (100, 200)$ and $\sigma = 0.1$. The data are included in the package and are available by using:

```
R> data("yt", package = "CEoptim")
R> xt <- yt - c(0, yt[-300])
```

The following code implements the use of **CEoptim** for this constrained mixed problem.

```
R> A <- rbind(diag(3), -diag(3))
R> b <- rep(1, 6)
R> set.seed(123)
R> library("CEoptim")
R> res <- CEoptim(f = sumsqrs, f.arg = list(xt), continuous =
+   list(mean = c(0, 0, 0), sd = rep(1.0, 3), conMat = A, conVec = b),
+   discrete = list(categories = c(298L, 298L), smoothProb = 0.5),
+   N = 10000, rho = 0.001, verbose = TRUE)
```

The output is as follows:

```
R> res
```

```
Optimizer for continuous part:
0.2702714 0.8801672 -0.8975874
```





Figure 9: Assessment of the fit of the model: x_t (top) and y_t (bottom).

```

Optimizer for discrete part:
 99 199
Optimum:
 2.675727
Number of iterations:
 12
Total number of function evaluations:
120000
Convergence:
Variance converged

```

As the input to `CEoptim` is $\mathbf{r} - 1$, the optimal vector \mathbf{r} is given by

```
R> (est.r <- sort(res$optimizer$discrete) + 1)
```

```
[1] 100 200
```

which gives exactly the “true” boundaries for the regimes. From the estimates of the model, one can assess the fit of the model by comparing y_t with $\hat{y}_t = \sum_{i=1}^{300} \hat{x}_t$ and x_t against the fit \hat{x}_t . Figure 9 shows an excellent fit.

```

R> t <- 1:300
R> est.theta <- res$optimizer$continuous
R> est.thetas <- rep(est.theta, times = c(est.r, 300) - c(1, est.r + 1) + 1)
R> xfit <- c(0, head(xt, -1)) * est.thetas
R> par(mfrow = c(2,1))
R> plot(xt ~ t, type = "l", col = "blue")
R> lines(xfit, col = "red")

```

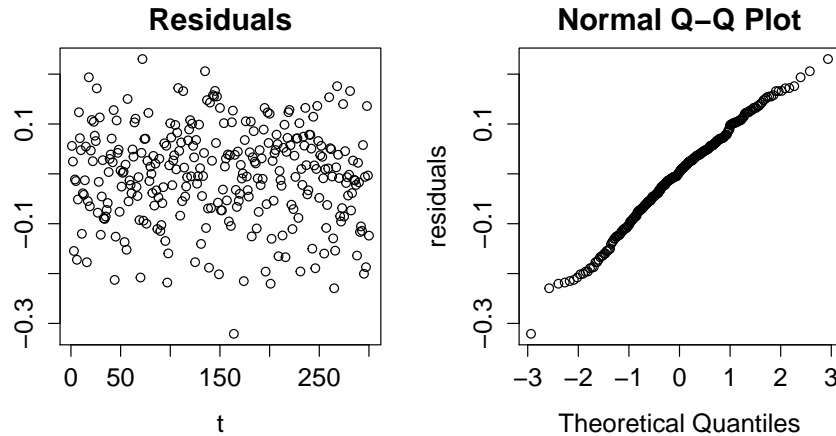


Figure 10: Diagnostic residuals of the model: scatterplot of the residuals (left) and quantile quantile normal plot(right).

```
R> abline(v = c(100, 200))
R> plot(yt, type = "l", col = "blue")
R> lines(cumsum(xfit), col = "red")
R> abline(v = c(100, 200))
```

A diagnostic of the residuals is presented in Figure 10, showing a normal quantile plot (left panel) and a scatterplot of the residuals (right panel).

```
R> par(mfrow = c(1, 2))
R> resid <- xfit - xt
R> plot(resid, ylab = "residuals", xlab = "t")
R> qqnorm(resid, ylab = "residuals")
```

6. Concluding remarks

CEoptim provides an R implementation of the cross-entropy method for optimization. The versatility and effectiveness of this new package have been illustrated through a variety of optimization examples involving continuous, discrete, mixed and constrained optimization problems. We have demonstrated how this simple algorithm can be of benefit in statistical inference, including model fitting, regression, maximum likelihood, and lasso methods. **CEoptim** is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=CEoptim>.

Acknowledgments

This work was supported by the Australian Research Council *Centre of Excellence for Mathematical and Statistical Frontiers* (ACEMS) under grant number CE140100049. Qibin Duan would also like to acknowledge the support from the University of Queensland through the UQ International Scholarships scheme.

References

- Alon G, Kroese DP, Raviv T, Rubinstein RY (2005). “Application of the Cross-Entropy Method to the Buffer Allocation Problem in a Simulation-Based Environment.” *The Annals of Operations Research*, **134**(1), 137–151. doi:10.1007/s10479-005-5728-8.
- Bendtsen C (2012). **pso**: Particle Swarm Optimization. R package version 1.0.3, URL <https://CRAN.R-project.org/package=pso>.
- Benham T, Duan Q, Kroese DP, Lique B (2017). **CEoptim**: Cross-Entropy R Package for Optimization. R package version 1.2, URL <https://CRAN.R-project.org/package=CEoptim>.
- Botev ZI, Kroese DP, Rubinstein RY, L’Ecuyer P (2013). “The Cross-Entropy Method for Optimization.” In *Handbook of Statistics – Machine Learning: Theory and Applications*, volume 31, pp. 35–59. Elsevier. doi:10.1016/b978-0-444-53859-8.00003-5.
- Butts CT (2008). “Social Network Analysis with **sna**.” *Journal of Statistical Software*, **24**(6), 1–51. doi:10.18637/jss.v024.i06.
- De Boer PT, Kroese DP, Mannor S, Rubinstein RY (2005). “A Tutorial on the Cross-Entropy Method.” *The Annals of Operations Research*, **134**(1), 19–67. doi:10.1007/s10479-005-5724-z.
- Duan Q, Kroese DP, Brereton T, Spetl A, Schmidt V (2014). “Inverting Laguerre Tessellations.” *The Computer Journal*, **57**(9), 1431–1440.
- Friedman J, Hastie T, Tibshirani R (2008). “Regularization Paths for Generalized Linear Models via Coordinate Descent.” *Journal of Statistical Software*, **33**(1), 1–22. doi:10.18637/jss.v033.i01.
- Knuth DE (1993). *The Stanford GraphBase: A Platform for Combinatorial Computing*. ACM Press, Reading.
- Kobilarov M (2012). “Cross-Entropy Motion Planning.” *The International Journal of Robotics Research*, **31**(7), 855–871. doi:10.1177/0278364912444543.
- Kothari RP, Kroese DP (2009). “Optimal Generation Expansion Planning via the Cross-Entropy Method.” In *Winter Simulation Conference*, pp. 1482–1491.
- Kroese DP, Hui KP, Nariai S (2007). “Network Reliability Optimization via the Cross-Entropy Method.” *IEEE Transactions on Reliability*, **56**(2), 275–287. doi:10.1109/tr.2007.895303.
- Minka TP (2012). “Estimating a Dirichlet Distribution.” *Technical report*, Machine Intelligence and Perception Group. URL <https://tminka.github.io/papers/dirichlet/minka-dirichlet.pdf>.
- Mullen KM, Ardia D, Gil DL, Windover D, Cline J (2011). “**DEoptim**: An R Package for Global Optimization by Differential Evolution.” *Journal of Statistical Software*, **40**(6), 1–26. doi:10.18637/jss.v040.i06.

- Nagumo J, Arimoto S, Yoshizawa S (1962). “An Active Pulse Transmission Line Simulating Nerve Axon.” *Proceedings of the IRE*, **50**(10), 2061–2070. doi:10.1109/jrproc.1962.288235.
- Ramsay JO, Hooker G, Campbell D, Cao J (2007). “Parameter Estimation for Differential Equations: A Generalized Smoothing Approach.” *Journal of the Royal Statistical Society B*, **69**(5), 741–796. doi:10.1111/j.1467-9868.2007.00610.x.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Rubinstein RY (1997). “Optimization of Computer Simulation Models with Rare Events.” *European Journal of Operational Research*, **99**(1), 89–112. doi:10.1016/s0377-2217(96)00385-2.
- Rubinstein RY (1999). “The Cross-Entropy Method for Combinatorial and Continuous Optimization.” *Methodology and Computing in Applied Probability*, **1**(2), 127–190. doi:10.1023/a:1010091220143.
- Rubinstein RY, Kroese DP (2004). *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning*. Springer-Verlag, New York.
- Rubinstein RY, Kroese DP (2017). *Simulation and the Monte Carlo Method*. 3rd edition. John Wiley & Sons, New York.
- Sani A, Kroese DP (2008). “Controlling the Number of HIV Infectives in a Mobile Population.” *Mathematical Biosciences*, **213**(2), 103–112. doi:10.1016/j.mbs.2008.03.003.
- Soetaert K, Petzoldt T, Setzer RW (2010). “Solving Differential Equations in R: Package **deSolve**.” *Journal of Statistical Software*, **33**(9), 1–25. doi:10.18637/jss.v033.i09.
- The MathWorks Inc (2014). *MATLAB – The Language of Technical Computing, Version R2014b*. Natick, Massachusetts. URL <http://www.mathworks.com/products/matlab/>.
- Xiang Y, Gubian S, Suomela B, Hoeng J (2013). “Generalized Simulated Annealing for Global Optimization: The **GenSA** Package.” *The R Journal*, **5**(1), 13–28. doi:10.1080/10556780008805777.

Affiliation:

Tim Benham, Qibin Duan, Dirk P. Kroese
School of Mathematics and Physics
The University of Queensland
Brisbane, Australia

E-mail: tim.j.benham@gmail.com, q.duan@uq.edu.au, kroese@maths.uq.edu.au

Benoît Liquez
School of Mathematics and Physics
The University of Queensland
ARC Centre of Excellence for Mathematical and Statistical Frontiers
Queensland University of Technology (QUT)
Brisbane, Australia
Laboratoire de Mathématiques et de leurs Applications,
Université de Pau et des Pays de l'Adour,
UMR CNRS 5142, Pau, France
E-mail: b.liquet@uq.edu.au (*corresponding author*)