



HAL
open science

BCMA-ES II: revisiting Bayesian CMA-ES

Eric Benhamou, David Saltiel, Beatrice Guez, Nicolas Paris

► **To cite this version:**

Eric Benhamou, David Saltiel, Beatrice Guez, Nicolas Paris. BCMA-ES II: revisiting Bayesian CMA-ES. 2020. hal-02886508

HAL Id: hal-02886508

<https://hal.science/hal-02886508>

Preprint submitted on 1 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

BCMA-ES II: revisiting Bayesian CMA-ES

Eric Benhamou

A.I Square Connect and Lamsade, France
eric.benhamou@aisquareconnect.com

Beatrice Guez

A.I Square Connect, France
beatrice.guez@aisquareconnect.com

David Saltiel

A.I Square Connect and LISIC, France
david.saltiel@aisquareconnect.com

Nicolas Paris

A.I Square Connect, France
nicolas.paris@aisquareconnect.com

ABSTRACT

This paper revisits the Bayesian CMA-ES and provides updates for normal Wishart. It emphasizes the difference between a normal and normal inverse Wishart prior. After some computation, we prove that the only difference relies surprisingly in the expected covariance. We prove that the expected covariance should be lower in the normal Wishart prior model because of the convexity of the inverse. We present a mixture model that generalizes both normal Wishart and normal inverse Wishart model. We finally present various numerical experiments to compare both methods as well as the generalized method.

CCS CONCEPTS

• **Mathematics of computing** → *Probability and statistics*;

KEYWORDS

CMA ES, Bayesian, conjugate prior, normal Wishart, normal inverse Wishart, mixture models

1 INTRODUCTION

Bayesian statistics have revolutionized statistics like quantum mechanics have done for Newtonian mechanism. Like the latter, the usual frequentist statistics can be seen as a particular asymptotic case of the former. Indeed, the Cox Jaynes theorem ([8]) proves that under the four axiomatic assumptions given by:

- plausibility degrees are represented by real numbers (continuity of method),
- none of the possible data should be ignored (no retention)
- these values follow usual common sense rule as stated by the well known Laplace formula: *the probability theory is truly the common sense represented in calculus* (common sense),
- and states of equivalent knowledge should have equivalent degree of plausibility (consistency),

then, there exists a probability measure defined up to a monotonous function such that it follows the usual probability calculus and the fundamental rule of Bayes, that is:

$$\mathbb{P}(H, D) = \mathbb{P}(H|D)\mathbb{P}(D) = \mathbb{P}(D|H)\mathbb{P}(H) \quad (1)$$

where H and D are two members of the implied σ -algebra. The letters are not by chance. H stands for the hypothesis, which can be interpreted as an hypothesis on the parameters, while D stands for data.

The usual frequentist probabilities states that the probability of an observation $\mathbb{P}(D)$ is given certain hypothesis H on the state of the world. However, as the equation (1) is completely symmetric, nothing hinders us to change our point of view and state the inverse question. Given an observation of a data D , what is the plausibility of the hypothesis H . The Bayes rules trivially answers this question:

$$\mathbb{P}(H|D) = \mathbb{P}(D|H) \frac{\mathbb{P}(H)}{\mathbb{P}(D)} = \mathbb{P}(D|H)\mathbb{P}(H) \quad (2)$$

or equivalently,

$$\mathbb{P}(H|D) \propto \mathbb{P}(D|H)\mathbb{P}(H) \quad (3)$$

In the above equation, $\mathbb{P}(H)$ is called the prior probability or simply the prior while the conditional probability $\mathbb{P}(H|D)$ is called the posterior probability or simply the posterior. There are a few remarks to be made. First of all, the prior is not necessarily independent of the knowledge of the experience, on the contrary, a prior is often determined with some knowledge of previous experience in order to make a meaningful choice. Second, prior and posterior are not necessarily related to a chronological order but rather to a logical order.

After observing some data D , we revise the plausibility of H . it is interesting to see that the conditional probability $\mathbb{P}(D|H)$ considered as a function of H is indeed a likelihood for H . The Cox Jaynes theorem as presented in [18] gives the foundation for Bayesian calculus. Another important result is the De Finetti's theorem. Let us recall the definition of Infinite exchangeability.

DEFINITION 1.1. (*Infinite exchangeability*). We say that (x_1, x_2, \dots) is an infinitely exchangeable sequence of random variables if, for any n , the joint probability $p(x_1, x_2, \dots, x_n)$ is invariant to permutation of the indices. That is, for any permutation π ,

$$p(x_1, x_2, \dots, x_n) = p(x_{\pi 1}, x_{\pi 2}, \dots, x_{\pi n})$$

Equipped with this definition, the De Finetti's theorem as provided below states that exchangeable observations are conditionally independent relative to some latent variable.

THEOREM 1.1. (*De Finetti, 1930s*). A sequence of random variables (x_1, x_2, \dots) is infinitely exchangeable iff, for all n ,

$$p(x_1, x_2, \dots, x_n) = \int \prod_{i=1}^n p(x_i|\theta)P(d\theta),$$

for some measure P on θ .

This representation theorem 1.1 justifies the use of priors on parameters since for exchangeable data, there must exist a parameter θ , a likelihood $p(x|\theta)$ and a distribution π on θ . A proof of De Finetti theorem is for instance given in [23] (section 1.5). We will see that this Bayesian setting gives a powerful framework for revisiting black box optimization that is introduced below.

2 BLACK BOX OPTIMIZATION

We assume that we have a real value p -dimensional function $f : \mathbb{R}^p \rightarrow \mathbb{R}$. We examine the following optimization program:

$$\min_{x \in \mathbb{R}^p} f(x) \quad (4)$$

In contrast to traditional convex optimization theory, we do not assume that f is convex, neither continuous nor admits a global minimum. We are interested in the so called *Black box optimization* (BBO) settings where we only have access to the function f and nothing else. By nothing else, we mean we can not for instance

compute gradient. A practical way to do optimization in this very general and minimal setting is to do evolutionary optimization and in particular use the covariance matrix adaptation evolution strategy (CMA-ES) methodology. The CMA-ES [?] is arguably one of the most powerful real-valued derivative-free optimization algorithms, finding many applications in machine learning. It is a state-of-the-art optimizer for continuous black-box functions as shown by the various benchmarks of the COCO (COMparing Continuous Optimisers) INRIA platform for ill-posed functions. It has led to a large number of papers and articles and we refer the interested reader to [1, 2, 4–6, 11, 12, 16, 21?] and [24] to cite a few.

It has been successfully applied in many unbiased performance comparisons and numerous real-world applications. In particular, in machine learning, it has been used for direct policy search in reinforcement learning and hyper-parameter tuning in supervised learning ([10], [14, 15, 17]), and references therein, as well as hyperparameter optimization of deep neural networks [19].

In a nutshell, the (μ / λ) CMA-ES is an iterative black box optimization algorithm, that, in each of its iterations, samples λ candidate solutions from a multivariate normal distribution, evaluates these solutions (sequentially or in parallel) retains μ candidates and adjusts the sampling distribution used for the next iteration to give higher probability to good samples. Each iteration can be individually seen as taking an initial guess or *prior* for the multi variate parameters, namely the mean and the covariance, and after making an experiment by evaluating these sample points with the fit function updating the initial parameters accordingly. Although rethinking the CMA-ES in terms of a prior and posterior seems natural when coming over from Bayesian statistics, it is only recently that it has been explored [7].

Historically, the CMA-ES has been developed heuristically. It was done mainly by conducting experimental research and validating intuitions empirically.

Research was done without much focus on theoretical foundations because of the apparent complexity of this algorithm. It was only recently that [3], [9] and [21] made a breakthrough and provided a theoretical justification of CMA-ES updates thanks to information geometry. They proved that CMA-ES was performing a natural gradient descent in the Fisher information metric. The Bayesian formulation of the CMA-ES came effectively much later and has only been done so far with the normal inverse Wishart prior.

In this paper, we revisit the Bayesian CMA-ES formulation and show that there exists indeed an infinity of conjugate prior given by the convex combination of a normal Wishart and normal inverse Wishart Gaussian prior. We first prove that normal Wishart and normal inverse Wishart Gaussian priors have the same update equations except for the mean of the covariance matrix. We provide a theoretical argument to show that the inverse of a matrix should be lower than in the normal inverse Wishart Gaussian prior. We then introduce a new prior given by a mixture of normal Wishart and normal inverse Wishart Gaussian prior. Likewise, we derive the update equations. In section 5, we finally give numerical results to compare all these methods.

3 CONJUGATE PRIORS

A key concept in Bayesian statistics is conjugate priors that makes the computation really easy and is described below.

DEFINITION 3.1. *A prior distribution $\pi(\theta)$ is said to be a conjugate prior if the posterior distribution*

$$\pi(\theta|x) \propto p(x|\theta)\pi(\theta) \quad (5)$$

remains in the same distribution family as the prior.

At this stage, it is relevant to introduce exponential family distributions as this higher level of abstraction that encompasses the multi variate normal trivially solves the issue of founding conjugate priors. This will be very helpful for inferring conjugate priors for the multi variate Gaussian used in CMA-ES.

DEFINITION 3.2. *A distribution is said to belong to the exponential family if it can be written (in its canonical form) as:*

$$p(x|\eta) = h(x) \exp(\eta \cdot T(x) - A(\eta)), \quad (6)$$

where η is the natural parameter, $T(x)$ is the sufficient statistic, $A(\eta)$ is log-partition function and $h(x)$ is the base measure. η and $T(x)$ may be vector-valued. Here $a \cdot b$ denotes the inner product of a and b .

The log-partition function is defined by the integral:

$$A(\eta) \triangleq \log \int_{\mathcal{X}} h(x) \exp(\eta \cdot T(x)) dx. \quad (7)$$

Also, $\eta \in \Omega = \{\eta \in \mathbb{R}^m | A(\eta) < +\infty\}$ where Ω is the natural parameter space. Moreover, Ω is a convex set and $A(\cdot)$ is a convex function on Ω .

REMARK 3.1. *Not surprisingly, the normal distribution $\mathcal{N}(x; \mu, \Sigma)$ with mean $\mu \in \mathbb{R}^d$ and covariance matrix Σ belongs to the exponential family but with a different parameterization. Its exponential family form is given by:*

$$\eta(\mu, \Sigma) = \begin{bmatrix} \Sigma^{-1} \mu \\ \text{vec}(\Sigma^{-1}) \end{bmatrix}, \quad T(x) = \begin{bmatrix} x \\ \text{vec}(-\frac{1}{2}xx^T) \end{bmatrix}, \quad (8a)$$

$$h(x) = (2\pi)^{-\frac{d}{2}}, \quad A(\eta(\mu, \Sigma)) = \frac{1}{2} \mu^T \Sigma^{-1} \mu + \frac{1}{2} \log |\Sigma|. \quad (8b)$$

where in equations (8a), the notation $\text{vec}(\cdot)$ means we have vectorized the matrix, stacking each column on top of each other and hence can equivalently write for a and b , two matrices, the trace result $\text{Tr}(a^T b)$ as the scalar product of their vectorization $\text{vec}(a) \cdot \text{vec}(b)$ (see 7.2). We can remark the canonical parameters are very different from traditional (also called moment) parameters. We can notice that changing slightly the sufficient statistic $T(x)$ leads to change the corresponding canonical parameters η . In equation (8b), the notation $|\Sigma|$ means the determinant of the matrix: $\det(\Sigma)$.

For an exponential family distribution, it is particularly easy to characterize conjugate prior that belongs to the exponential family.

PROPOSITION 3.1. Characterization of exponential family conjugate prior. *If the observations have a density of the exponential family form $p(x|\theta, \kappa) = h(x) \exp(\eta(\theta, \kappa)^T T(x) - nA(\eta(\theta, \kappa)))$, with κ a set of hyper-parameters, then the conjugate priors that belong to the exponential family are necessarily of the form $\pi(\theta) \propto \exp(\lambda_1 \cdot \eta(\theta, \kappa) - \lambda_0 A(\eta(\theta, \kappa)))$ with $\lambda \triangleq (\lambda_0, \lambda_1)$.*

The proof is given in appendix subsection 7.1. As we can vary the parameterization of the likelihood, we can obtain multiple conjugate priors. Because of the conjugacy, if the initial parameters of the multi variate Gaussian follows the prior, the posterior is the true distribution given the information \mathcal{X} and stay in the same family making the update of the parameters really easy. Said differently, with conjugate prior, we make the optimal update.

A consequence of proposition 3.1 is that the various conjugate priors of the multi variate normal that belong to the exponential family can be determined. This is the subject of the corollary below.

COROLLARY 3.2. *The conjugate priors of the multi variate normal that belong to the exponential family are necessarily of the form :*

- *normal inverse Wishart distribution* $NIW(\mu_0, \lambda_0, \nu_0, \Psi_0)$ if the multivariate normal is described in terms of its mean vector μ and covariance matrix Σ .
- *normal Wishart distribution* $NW(\mu_0, \lambda_0, \nu_0, W_0)$ if the multivariate normal is described in terms of its mean vector μ and precision matrix Λ .

The proof is given in appendix subsection 7.3. As conjugate priors, the posterior of the two identified distributions of the corollary 3.2 are easy to derive and are given by the following proposition.

PROPOSITION 3.3. *For a likelihood of n points $(x)_{i=1..n}$ distributed according to a multi variate normal distribution whose parameters are given by the priors below:*

- (1) *the normal inverse Wishart distribution:*
 $NIW_0 = NIW(\mu_0, \lambda_0, \nu_0, \Psi_0)$
- (2) *the normal Wishart distribution:* $NW_0 = NW(\mu_0, \lambda_0, \nu_0, W_0)$
- (3) *the mixture of a normal inverse and normal Wishart with same parameters:* $wNIW_0 + (1-w)NW_0$ with $0 \leq w \leq 1$

The posterior is given by:

- (1) *the normal inverse Wishart distribution*

$$NIW_1 = NIW\left(\frac{\lambda_0 \mu_0 + n\bar{x}}{\lambda_0 + n}, \lambda_0 + n, \nu_0 + n, \Psi_0 + nC + nD\right) \quad (9)$$

- (2) *the normal Wishart distribution*

$$NW_1 = NW\left(\frac{\lambda_0 \mu_0 + n\bar{x}}{\lambda_0 + n}, \lambda_0 + n, \nu_0 + n, (W_0 + nC + nD)^{-1}\right) \quad (10)$$

- (3) *the mixture of a normal inverse and normal Wishart with same parameters:* $wNIW_1 + (1-w)NW_1$

where $\bar{x} = 1/n \sum_{i=1}^n x_i$ is the sample mean, $C = 1/n \sum_{i=1}^n (x_i - \bar{x})(x_i - \bar{x})^T$ the sample covariance and $D = \frac{\lambda_0 n}{n(\lambda_0 + n)}(\bar{x} - \mu_0)(\bar{x} - \mu_0)^T$.

The proof is given in appendix subsection 7.4.

4 ALGORITHM

The idea behind the algorithm is at each step to make use the previous iteration posterior as a prior, draw the likelihood and then update according to proposition (3.3) the posterior. In full generality, the prior is a distribution, so we would need to do a Monte Carlo of Monte Carlo. But in order to reduce the variance by this Monte Carlo of Monte Carlo, we make the simplification to use the mean value of the prior distribution. These values are given as follows:

- (1) for the normal inverse Wishart distribution, $\hat{\mu} = \mathbb{E}[\mu] = \mu_n$ and $\hat{\Sigma} = \mathbb{E}[\Sigma] = \Psi_n / (\nu_n - p - 1)$
- (2) for the normal Wishart distribution, $\hat{\mu} = \mathbb{E}[\mu] = \mu_n$ and $\hat{\Sigma} = \mathbb{E}[\Lambda^{-1}] = \Psi_n / \nu_n$ for $\Psi_n = W_n^{-1}$.
- (3) for the w mixture of the normal inverse and normal Wishart with same parameters, $\hat{\mu} = \mathbb{E}[\mu] = \mu_n$ and $\hat{\Sigma} = \mathbb{E}[\Sigma] = \frac{\nu_n - p - 1 + w\nu_n + w}{\nu_n(\nu_n - p - 1)} \Psi_n$

It is obvious that the expected value of the covariance matrix of the normal inverse Wishart $\hat{\Sigma} = \mathbb{E}[\Sigma]$ should be above the one of the normal Wishart distribution as the inverse of a matrix $Inv : S \rightarrow S^{-1}$ is a convex function in the domain \mathcal{S}_{++}^p of symmetric definite positive matrices. A proof is given in 7.5. To recover the true minimum, we design two strategies.

- we design a strategy where we rebuild our normal distribution but using sorted information of our X 's weighted by their normal density to ensure this is a true normal corrected from the Monte Carlo bias. We need to explicitly compute the

weights. For each simulated point X_i , we compute its assumed density denoted by $d_i = \mathcal{N}(\hat{\mu}, \hat{\Sigma})(X_i)$ where $\mathcal{N}(\hat{\mu}, \hat{\Sigma})(\cdot)$ denotes the p.d.f. of the multi-variate Gaussian. We divide these density by their sum to get weights $(w_i)_{i=1..k}$ that are positive and sum to one as follows. $w_j = d_j / \sum_{i=1}^k d_i$. Hence for k simulated points, we get $\{X_i, w_i\}_{i=1..k}$. We reorder jointly the uplets (points and density) in terms of their weights in decreasing order. To insist we take sorted value in decreasing order with respect to the weights $(w_i)_{i=1..k}$, we denote the order statistics $(i), w \downarrow$. This first sorting leads to k new uplets $\{X_{(i), w \downarrow}, w_{(i), w \downarrow}\}_{i=1..k}$. Using a *stable* sort (that keeps the order of the density), we sort jointly the uplets (points and weights) according to their objective function value (in increasing order this time) and get a k new uplets $\{X_{(i), f \uparrow}, w_{(i), w \downarrow}\}_{i=1..k}$. We can now compute a new mean as follows:

$$\hat{\mu} = \underbrace{\sum_{i=1}^k w_{(i), w \downarrow} \cdot X_{(i), f \uparrow}}_{\text{MC mean for } X_{f \uparrow}} - \underbrace{\left(\sum_{i=1}^k w_i X_i - \hat{\mu}\right)}_{\text{MC bias for } X} \quad (11)$$

The intuition of equation (11) is to compute in the left term the Monte Carlo mean using reordered points according to their objective value and correct our initial computation by the Monte Carlo bias computed as the right term, equal to the initial Monte Carlo mean minus the real mean. We call this **strategy one**.

- If we think for a minute about the strategy one, we get the intuition that when starting the minimization, it may not be optimal. This is because weights are proportional to $\exp(\frac{1}{2}(X - \hat{\mu})^T \hat{\Sigma}^{-1}(X - \hat{\mu}))$. When we start the algorithm, we use a large search space, hence a large covariance matrix $\hat{\Sigma}$ which leads to have weights which are quite similar. Hence even if we sort candidates by their fit, ranking them according to the value of f in increasing order, we will move our theoretical multi variate Gaussian little by little. A better solution is more to brutally move the center of our multi variate Gaussian to the best candidate seen so far, as follows:

$$\hat{\mu} = \arg \min_{X \in X} f(X) \quad (12)$$

We call this **strategy two**. Intuitively, strategy two should be best when starting the algorithm while strategy one would be better once we are close to the solution.

To recover the true variance, we can adapt what we did in strategy one as follows:

$$\hat{\Sigma} = \underbrace{\sum_{i=1}^k w_{(i), w \downarrow} \cdot \left(X_{(i), f \uparrow} - \bar{X}_{(\cdot), f \uparrow}\right) \left(X_{(i), f \uparrow} - \bar{X}_{(\cdot), f \uparrow}\right)^T}_{\text{MC covariance for } X_{f \uparrow}} - \underbrace{\left(\sum_{i=1}^k w_i \cdot \left(X_i - \bar{X}\right) \left(X_i - \bar{X}\right)^T - \hat{\Sigma}\right)}_{\text{MC covariance for simulated } X} \quad (13)$$

where $\bar{X}_{(\cdot), f \uparrow} = \sum_{i=1}^k w_{(i), w \downarrow} X_{(i), f \uparrow}$ and $\bar{X} = \sum_{i=1}^k w_i X_i$ are respectively the mean of the sorted and non sorted points.

Algorithm 1 Predict and Correct parameters at step k

- 1: **Simulate candidate**
- 2: Use mean values $\hat{\mu} = \mathbb{E}[\mu]$ and $\hat{\Sigma} = \mathbb{E}[\Sigma]$
- 3: Simulate k points $\mathcal{X} = \{X_i\}_{i=1..n} \sim \mathcal{N}(\hat{\mu}, \hat{\Sigma})$
- 4: Compute densities $(d_i)_{i=1..n} = (\mathcal{N}(\hat{\mu}, \hat{\Sigma})(X_i))_{i=1..n} =$
- 5: Sort in decreasing order with respect to d to get $\{X_{(i),d\downarrow}, d_{(i),d\downarrow}\}_{i=1..n}$
- 6: Stable Sort in increasing order with respect to $f(X_i)$ to get $\{X_{(i),f\uparrow}, d_{(i),d\downarrow}\}_{i=1..n}$
- 7:
- 8: **Correct $\hat{\mu}$ and $\hat{\Sigma}$**
- 9: Either Update $\hat{\mu}$ and $\hat{\Sigma}$ using (12) and (13) (**strategy two**)
- 10: Or Update $\hat{\mu}$ and $\hat{\Sigma}$ using (11) and (13) (**strategy one**)
- 11: Update $\mu_{k+1}, \lambda_{k+1}, v_{k+1}, \psi_{k+1}$ using proposition 3.3

5 NUMERICAL RESULTS

5.1 Functions examined

We have examined five functions to stress test our algorithm. They are listed in increasing order of complexity for our algorithm and correspond to different type of functions. They are all generalized function that can defined for any dimension n . For all, we present the corresponding equation for a variable $x = (x_1, x_2, \dots, x_n)$ of n dimension. Code is provided in supplementary materials. We have frozen seeds to have *reproducible of results*.

5.1.1 Cone. The most simple function to optimize is the quadratic cone whose equation is given by (14) and represented in figure 1. It is also the standard Euclidean norm. It is obviously convex and is a good test of the performance of an optimization method.

$$f(x) = \left(\sum_{i=1}^n x_i^2 \right)^{1/2} = \|x\|_2 \tag{14}$$

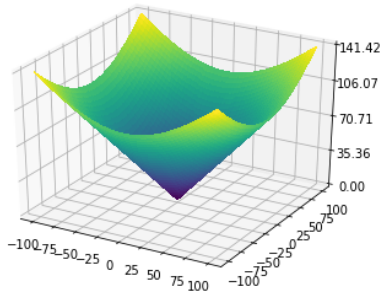


Figure 1: A simple convex function: the quadratic norm. Minimum in 0

5.1.2 Schwefel 2 function. A slightly more complicated function is the Schwefel 2 function whose equation is given by (15) and represented in figure 2. It is a piece wise linear function and validates the algorithm can cope with non convex function.

$$f(x) = \sum_{i=1}^n |x_i| + \prod_{i=1}^n |x_i| \tag{15}$$

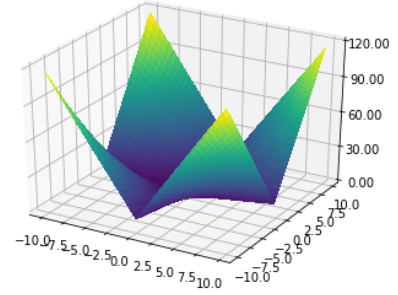


Figure 2: Schwefel 2 function: a simple piecewise linear function

5.1.3 Rastrigin. The Rastrigin function, first proposed by [22] and generalized by [20], is more difficult compared to the Cone and the Schwefel 2 function. Its equation is given by (16) and represented in figure 3. It is a non-convex function often used as a performance test problem for optimization algorithms. It is a typical example of non-linear multi modal function. Finding its minimum is considered a good stress test for an optimization algorithm, due to its large search space and its large number of local minima.

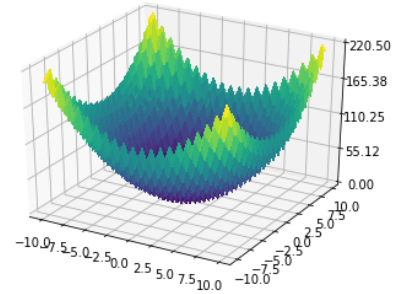


Figure 3: Rastrigin function: a non convex function multi-modal and with a large number of local minima

$$f(x) = 10 \times n + \sum_{i=1}^n [x_i^2 - 10 \cos(2\pi x_i)] \tag{16}$$

5.1.4 Schwefel 1 function. The Schwefel 1 function whose equation is given by (17) is a tricky function to optimize. It is represented in figure 4. It is sometimes only defined on $[-500, 500]^n$. The Schwefel 1 function shares similarities with the Rastrigin function. It is continuous, not convex, multi-modal and with a large number of local minima. The extra difficulty compared to the Rastrigin function, the local minima are more pronounced local bowl making the optimization even harder.

$$f(x) = 418.9829 \times n - \sum_{i=1}^n \left[x_i \sin(\sqrt{|x_i|}) \mathbb{1}_{|x_i| < 500} + 500 \sin(\sqrt{500}) \mathbb{1}_{|x_i| \geq 500} \right] \quad (17)$$




Figure 4: Schwefel 1 function: a non convex function multi-modal and with a large number of local pronounced bowls

5.1.5 *Eggholder function.* The Eggholder function whose equation is given by (18) is a difficult function to optimize, because of the large number of local minima. It is sometimes only defined on $[-512, 512]^n$. It shares similarities with the Schwefel1 function. It is continuous, not convex, multi-modal and with a large number of local minima.

$$f(x, y) = -(y + 47) \sin \sqrt{\left| \frac{x}{2} + (y + 47) \right|} - x \sin \sqrt{|x - (y + 47)|} \quad (18)$$




Figure 5: Eggholder function: a non convex function multi-modal and with a large number of local pronounced bowls

5.2 Convergence

For each of the functions, we compared our method using strategy one entitled *B-CMA-ES S1*: update $\hat{\mu}$ and $\hat{\Sigma}$ using (11) and (13) in orange with strategy two *B-CMA-ES S2*: same update but using (12) and (13), in blue and standard CMA-ES as provided by the

opensource python package *pycma* in green. We clearly see that strategy two outperforms standard CMA-ES and Bayesian CMA-ES S1. The convergence graphics that shows the error compared to the minimum are represented

- for the cone function by figure 6 (case of a convex function), with initial point (10, 10)
- for the Schwefel 2 function in figure 7 (case of piecewise linear function), with initial point (10, 10)
- for the Rastrigin function in figure 8 (case of a non convex function with multiple local minima), with initial point (10, 10)
- and for the Schwefel 1 function in figure 9 (case of a non convex function with multiple large bowl local minima), with initial point (10, 10)

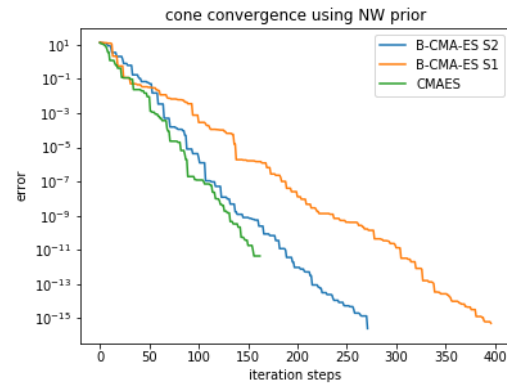
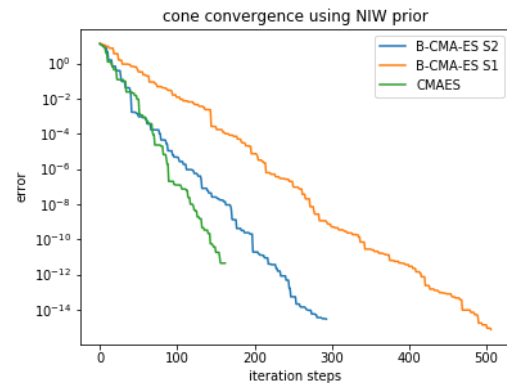


Figure 6: Convergence for the Cone function. B-CMA-ES S2 outperforms standard CMA-ES and B-CMA-ES S1.

For functions that are convex, our method performs similarly as standard CMA-ES. For function with harder local minima, the Bayesian CMA-ES is able to perform better. We conjecture that this is due to contraction dilatation mechanism that enables to avoid being trapped in a local minimum.

6 CONCLUSION

In this paper, we have revisited the CMA-ES algorithm and provided a Bayesian version of it. Taking conjugate priors, we can find optimal update for the mean and covariance of the multi variate Normal. We have provided the corresponding algorithm that

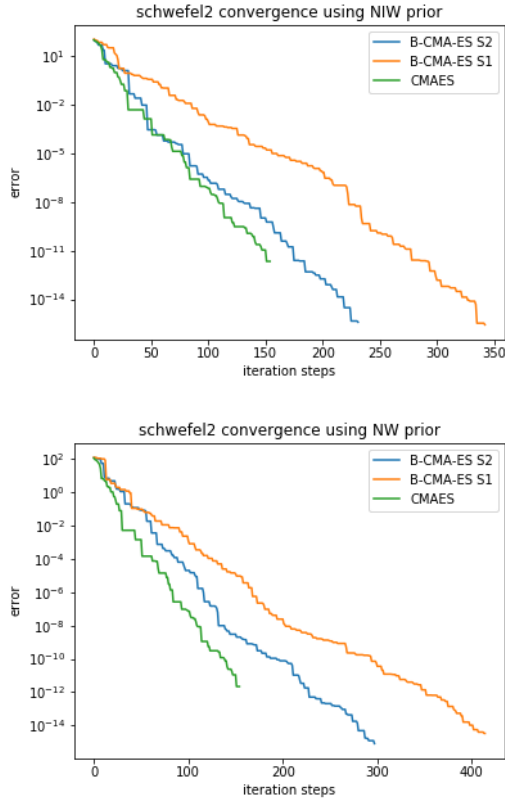


Figure 7: Convergence for the Schwefel 2 function. B-CMA-ES S2 outperforms standard CMA-ES and B-CMA-ES S1.

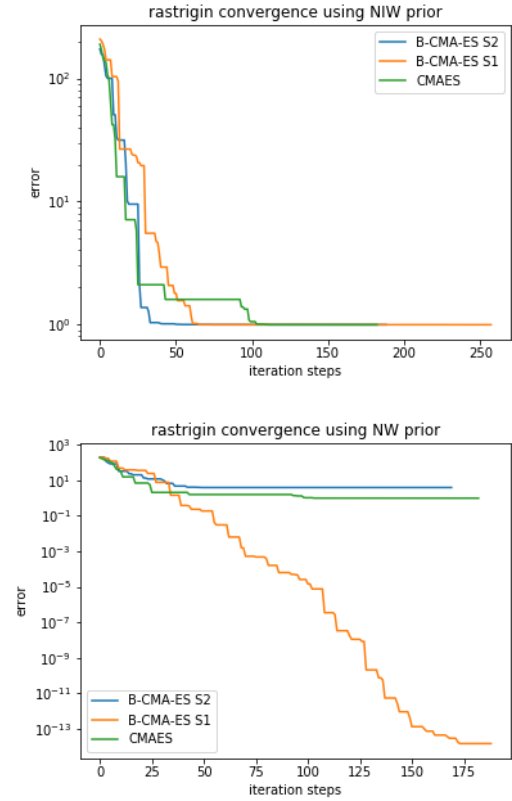


Figure 8: Convergence for the Rastrigin function. B-CMA-ES S2 outperforms standard CMA-ES and B-CMA-ES S1.

is a new version of CMA-ES. First numerical experiments show this new version is comparable to standard CMA-ES on traditional functions such as cone, Schwefel 1, Rastrigin and Schwefel 2. The similar convergence can be explained on a theoretical side from the optimal update of the prior (thanks to Bayesian update) and the use of the best candidate seen at each simulation to shift the mean of the multi-variate Gaussian likelihood. We envisage further works to benchmark our algorithm to traditional CMA-ES and other evolutionary algorithms, in particular to use the COCO platform to provide more meaningful tests and confirm the theoretical intuition of good performance of this new version of CMA-ES, and to test the importance of the prior choice.

7 APPENDIX

7.1 Conjugate priors

PROOF. Consider n independent and identically distributed (IID) measurements $\mathcal{X} \triangleq \{\mathbf{x}^j \in \mathbb{R}^d | 1 \leq j \leq n\}$ and assume that these variables have an exponential family density. The likelihood $p(\mathcal{X}|\theta, \kappa)$, writes simply as the product of each individual likelihood:

$$p(\mathcal{X}|\theta, \kappa) = \left(\prod_{j=1}^n h(\mathbf{x}^j) \right) \exp \left(\eta(\theta, \kappa)^T \sum_{j=1}^n T(\mathbf{x}^j) - nA(\eta(\theta, \kappa)) \right). \quad (19)$$

If we start with a prior $\pi(\theta)$ of the form $\pi(\theta) \propto \exp(\mathcal{F}(\theta))$ for some function $\mathcal{F}(\cdot)$, its posterior writes:

$$\begin{aligned} \pi(\theta|\mathcal{X}) &\propto p(\mathcal{X}|\theta) \exp(\mathcal{F}(\theta)) \\ &\propto \exp \left(\eta(\theta, \kappa) \cdot \sum_{j=1}^n T(\mathbf{x}^j) - nA(\eta(\theta, \kappa)) + \mathcal{F}(\theta) \right). \quad (20) \end{aligned}$$

It is easy to check that the posterior (20) is in the same exponential family as the prior iff $\mathcal{F}(\cdot)$ is in the form

$$\mathcal{F}(\theta) = \lambda_1 \cdot \eta(\theta, \kappa) - \lambda_0 A(\eta(\theta, \kappa)) \quad (21)$$

for some $\lambda \triangleq (\lambda_0, \lambda_1)$, such that

$$p(\mathcal{X}|\theta, \kappa) \propto \exp \left(\left(\lambda_1 + \sum_{j=1}^n T(\mathbf{x}^j) \right)^T \eta(\theta, \kappa) - (n + \lambda_0) A(\eta(\theta, \kappa)) \right). \quad (22)$$

Hence, the conjugate prior for the likelihood (19) is parametrized by λ and given by

$$p(\mathcal{X}|\theta, \kappa) = \frac{1}{Z} \exp(\lambda_1 \cdot \eta(\theta, \kappa) - \lambda_0 A(\eta(\theta, \kappa))), \quad (23)$$

where $Z = \int \exp(\lambda_1 \cdot \eta(\theta, \kappa) - \lambda_0 A(\eta(\theta, \kappa))) \, d\mathbf{x}$. \square

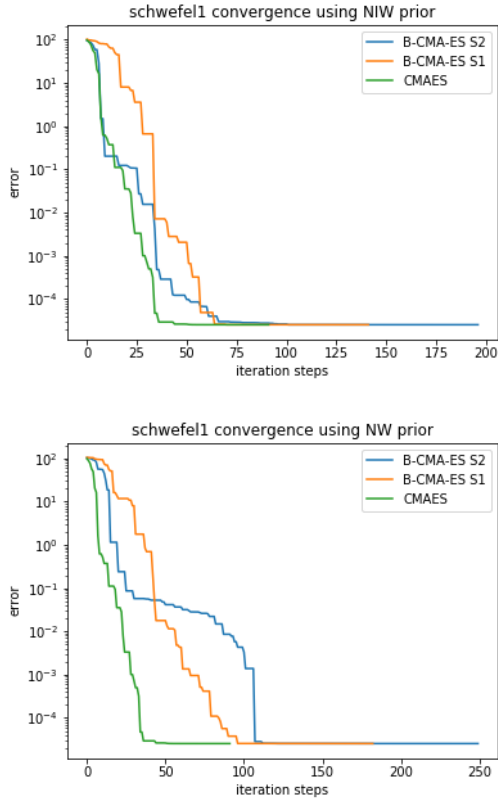


Figure 9: Convergence for the Schwefel 1 function. B-CMA-ES S2 outperforms standard CMA-ES and B-CMA-ES S1.

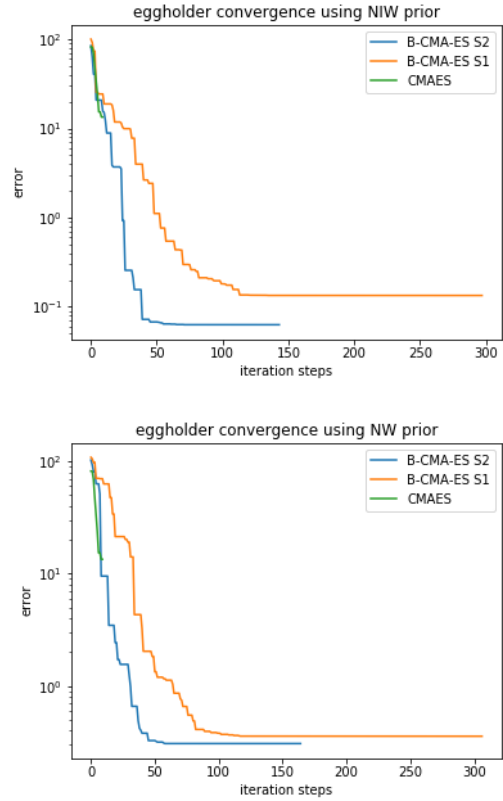


Figure 10: Convergence for the Eggholder function

7.2 Multivariate Canonical form

In the case of the multi variate normal, the canonical form for this distribution writes as

$$\begin{aligned}
 & \frac{1}{\sqrt{(2\pi)^d \det(\Sigma)}} \exp\left(-\frac{(X - \mu)^T \Sigma^{-1} (X - \mu)}{2}\right) \\
 = & \exp\left((\Sigma^{-1} \mu, \Sigma^{-1})^T \cdot (X, -\frac{1}{2} \text{vec}(XX^T))\right) \\
 & \frac{1}{(2\pi)^{d/2}} \exp\left(-\frac{1}{2} \mu^T \Sigma^{-1} \mu - \frac{1}{2} \log(\det(\Sigma))\right) \quad (24)
 \end{aligned}$$

which gives the following moment and canonical parameters:

$$\begin{aligned}
 \theta &= (\mu, \Sigma) \\
 T(X) &= \left(X, -\frac{1}{2} \text{vec}(XX^T)\right) \\
 \eta(\theta) &= \left(\Sigma^{-1} \mu, \Sigma^{-1}\right) \\
 A(\eta(\theta)) &= \frac{1}{2} \mu^T \Sigma^{-1} \mu + \frac{1}{2} \log(\det(\Sigma)) \\
 h(x) &= \frac{1}{(2\pi)^{d/2}} \quad (25)
 \end{aligned}$$

7.3 Conjugate priors determination

Using proposition 3.1 and the exponential family formulation of the multi variate normal (equations (25)), we have that any conjugate prior for the multi variate normal that belongs to the exponential family is given by

$$\begin{aligned}
 \pi(\theta) &\propto \exp(\lambda_1 \cdot \eta(\theta) - \lambda_0 A(\eta(\theta))) \\
 &\propto \exp\left(\lambda_1 \cdot \left(\Sigma^{-1} \mu, \Sigma^{-1}\right)\right) \\
 &\quad \exp\left(-\frac{1}{2} \mu^T \left(\frac{\Sigma}{\lambda_0}\right)^{-1} \mu - \frac{1}{2} \log(\det\left(\frac{\Sigma}{\lambda_0}\right))\right) \quad (26)
 \end{aligned}$$

If we write $\lambda_1 = (\lambda_0 \mu_0, \lambda_2)$ and $\Psi_0 = -2(\lambda_2 + \frac{\lambda_0}{2} \mu_0 \mu_0^T)$, we get

$$\begin{aligned}
 \pi(\theta) &\propto \exp\left(-\frac{1}{2} (\mu - \mu_0)^T \left(\frac{\Sigma}{\lambda_0}\right)^{-1} (\mu - \mu_0) - \frac{1}{2} \log(\det\left(\frac{\Sigma}{\lambda_0}\right))\right) \\
 &\quad \exp\left(-\frac{1}{2} \Psi_0 \cdot \Sigma^{-1}\right) \quad (27)
 \end{aligned}$$

The first term is a normal multi variate distribution. Its parameters are μ_0 and $\frac{\Sigma}{\lambda_0}$.

In the second term, we can recognize the proportional term of an inverse Wishart $\exp\left(-\frac{1}{2} \text{Tr}(\Psi_0 \Sigma^{-1})\right)$, with parameters ν_0, Ψ .

This shows the conjugate prior of the multi variate normal given by its mean vector μ and covariance matrix Σ is a normal inverse Wishart. Its parameters are $NIW(\mu_0, \lambda_0, \nu_0, \Psi_0)$ \square

If the multi variate normal is parametrized by its mean vector μ and its precision matrix Λ , the same reasoning gives

$$\pi(\theta) \propto \exp(\lambda_1 \cdot (\Lambda\mu, \Lambda)) \exp\left(-\frac{1}{2}\mu^T(\lambda_0\Lambda)\mu + \frac{1}{2}\log(\det(\frac{\Lambda}{\lambda_0}))\right) \quad (28)$$

The second term is a multi variate normal distribution given by $N(\mu_0, (\lambda_0\Lambda)^{-1})$ while the first one is the term of a Wishart distribution that is proportional to $\exp(\frac{1}{2}\text{Tr}(W^{-1}\Lambda))$ whose parameters are $\mathcal{W}(W_0, \nu_0)$. This shows that the conjugate prior of the multi variate normal described by its mean vector μ and precision matrix Λ is a normal Wishart distribution $NW(\mu_0, \lambda_0, \nu_0, W_0)$ \square

7.4 Posterior update

The posterior update is quite straightforward and very similar for the two cases: NIW and NW. We will detail only the calculation for the NIW case as it is very similar for the NW. Recall that the probability density function of a Normal inverse Wishart random variable is expressed as the product of a Normal and an Inverse Wishart probability density functions. Denoting by $p \times p$ the dimension of the covariance matrix Σ and using the Bayes rules, the posterior is proportional to the product of the prior and likelihood:

$$\begin{aligned} & \text{posterior} \\ & \propto \text{prior} \times \text{likelihood} \\ & \propto \sqrt{\frac{\lambda_0}{|\Sigma|}} \exp\left(-\frac{1}{2}(\mu - \mu_0)^T \left(\frac{\Sigma}{\lambda_0}\right)^{-1} (\mu - \mu_0)\right) \\ & \quad \times |\Psi|^{v_0/2} |\Sigma|^{-\frac{v_0+p+1}{2}} \exp\left(-\frac{1}{2}\text{Tr}(\Sigma^{-1}\Psi_0)\right) \\ & \quad \times |\Sigma|^{-n/2} \prod_{i=1}^n \exp\left(-\frac{1}{2}(x_i - \mu)^T \Sigma^{-1} (x_i - \mu)\right) \end{aligned} \quad (29)$$

First of all, we can regroup all terms in x_i as follows

$$\begin{aligned} & \prod_{i=1}^n \exp\left(-\frac{1}{2}(x_i - \mu)^T \Sigma^{-1} (x_i - \mu)\right) \\ & = \exp\left(-\frac{1}{2} \sum_{i=1}^n (x_i - \mu)^T \Sigma^{-1} (x_i - \mu)\right) \end{aligned} \quad (30)$$

and use the following remarkable identity:

$$\begin{aligned} & \sum_{i=1}^n (x_i - \mu)^T \Sigma^{-1} (x_i - \mu) \\ & = n \left[\frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^T \Sigma^{-1} (x_i - \bar{x}) + (\bar{x} - \mu)^T \Sigma^{-1} (\bar{x} - \mu) \right] \\ & = \text{Tr}(\Sigma^{-1} nC) + n(\bar{x} - \mu)^T \Sigma^{-1} (\bar{x} - \mu) \end{aligned} \quad (31)$$

where we have used the commutativity property of the trace operator $\text{Tr}(AB) = \text{Tr}(BA)$ and that for a real number, the number is equal to its trace and written $C = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^T \Sigma^{-1} (x_i - \bar{x})$ the sample covariance. Going further, we have

$$\begin{aligned} & (\mu - \mu_0)^T \left(\frac{\Sigma}{\lambda_0}\right)^{-1} (\mu - \mu_0) + n(\bar{x} - \mu)^T \Sigma^{-1} (\bar{x} - \mu) \\ & = (\lambda_0 + n)\mu^T \Sigma^{-1} \mu - 2(\lambda_0 \mu_0 + n\bar{x})^T \Sigma^{-1} \mu + \lambda_0 \mu_0^T \Sigma^{-1} \mu_0 \\ & \quad + n\bar{x}^T \Sigma^{-1} \bar{x} \end{aligned} \quad (32)$$

$$\begin{aligned} & = (\lambda_0 + n) \left(\mu - \frac{\lambda_0 \mu_0 + n\bar{x}}{\lambda_0 + n} \right)^T \Sigma^{-1} \left(\mu - \frac{\lambda_0 \mu_0 + n\bar{x}}{\lambda_0 + n} \right) \\ & \quad - \frac{1}{\lambda_0 + n} (\lambda_0 \mu_0 + n\bar{x})^T \Sigma^{-1} (\lambda_0 \mu_0 + n\bar{x}) \\ & \quad + \lambda_0 \mu_0^T \Sigma^{-1} \mu_0 + n\bar{x}^T \Sigma^{-1} \bar{x} \end{aligned} \quad (33)$$

$$\begin{aligned} & = (\lambda_0 + n) \left(\mu - \frac{\lambda_0 \mu_0 + n\bar{x}}{\lambda_0 + n} \right)^T \Sigma^{-1} \left(\mu - \frac{\lambda_0 \mu_0 + n\bar{x}}{\lambda_0 + n} \right) \\ & \quad + \text{Tr}(\Sigma^{-1} nD) \end{aligned} \quad (34)$$

$$\text{where } D = \frac{\lambda_0 n}{n(\lambda_0 + n)} (\bar{x} - \mu_0)(\bar{x} - \mu_0)^T \quad (35)$$

Hence, we can compute explicitly the posterior as follows:

$$\begin{aligned} & \text{posterior} \\ & \propto \sqrt{\frac{\lambda_1}{|\Sigma|}} \exp\left\{-\frac{1}{2}(\mu - \mu_1)^T \left(\frac{\Sigma}{\lambda_1}\right)^{-1} (\mu - \mu_1)\right\} \\ & \quad \times |\Psi_1|^{v_1/2} |\Sigma|^{-\frac{v_1+p+1}{2}} \exp\left\{-\frac{1}{2}\text{Tr}(\Sigma^{-1}\Psi_1)\right\} \end{aligned} \quad (36)$$

$$\begin{aligned} \text{with } \mu_1 & = \frac{\lambda_0 \mu_0 + n\bar{x}}{\lambda_0 + n} \\ \lambda_1 & = \lambda_0 + n \\ \nu_1 & = \nu_0 + n \\ \Psi_1 & = \Psi_0 + nC + nD \end{aligned} \quad (37)$$

which are exactly the equations provided in (9) \square

7.5 Convexity of the inverse of a matrix

We give here **six different proofs** of the convexity of the inverse of a matrix in the domain of symmetric definite positive matrices \mathcal{S}_{++}^p . The first and second proofs relies on the fact that the result is a consequence of proving that the matrix fractional function $f(X, y) = y^T X^{-1} y$ is convex on the domain $\text{dom} f = \mathcal{S}_{++}^p \times \mathbb{R}^p$. The implication comes from the fact that

f is convex

$$\begin{aligned} & \implies (1 - \lambda)f(M, y) + \lambda f(N, y) \geq f((1 - \lambda)M + \lambda N, y) \\ & \implies y^T \left[(1 - \lambda)M^{-1} + \lambda N^{-1} - ((1 - \lambda)M + \lambda N)^{-1} \right] y \geq 0 \end{aligned} \quad (38)$$

Since y is arbitrary, this implies the matrix within the square bracket in equation (38) is positive semi-definite. It is interesting to notice that matrix fractional function is in a sense an extension of the fact that the quadratic over linear function defined as $f(x, y) = x^2/y$ is convex on \mathbb{R}_+^2 .

PROOF. The **first proof** uses the property that the minimum of a convex function over a convex set is convex. For $\Sigma \in \mathcal{S}_{++}^n$, and for $u, y \in \mathbb{R}^n$ we can consider the quadratic function $f(u)$ defined by

$$f(u) = \frac{1}{2} u^T \Sigma u - u^T y$$

As $\Sigma \in \mathcal{S}_{++}^n$, this function is a obviously convex (quadratic function with its quadratic coefficient given by a definite positive matrix). Hence its minimum $\inf_{u \in \mathbb{R}^n} f(u)$ over a convex set is convex. Its easy to minimize a quadratic function and find its minimum given by the stationary point of its gradient $\frac{1}{2} y^T \Sigma^{-1} y$, which concludes the proof. \square

PROOF. A **second proof** is to show that the epigraph of f , denoted by $\text{epi}(f)$ is convex thanks to the link between positive semi definite cones and Schur complements. We have that

$$x \in \text{epi}(f) \Leftrightarrow y^T \Sigma^{-1} y \leq t \Leftrightarrow \begin{bmatrix} \Sigma & y \\ y^T & t \end{bmatrix} \geq 0 \quad (39)$$

This concludes the proof as the epigraph of f is convex as the inverse image of the positive semi definite cone S_{++}^{n+1} by the Schur complement that is an affine mapping. \square

PROOF. A **third proof** relies on the fundamental identity of the inverse of a matrix X : $XX^{-1} = I_p$, where I_p is the identity matrix with p rows (or columns). Take M, N two positive definite symmetric matrices and $\lambda \in [0, 1]$. Take $P_\lambda = (1 - \lambda)M + \lambda N$. P and P^{-1} are obviously symmetric positive definite. Denote by $(\cdot)'$ the derivative with respect to λ . We have:

$$PP^{-1} = I_p \implies P'P^{-1} + P(P^{-1})' = 0_p \implies (P^{-1})' = -P^{-1}P'P^{-1}$$

Notice that $P'' = 0_p$, since P is linear in λ . Differentiate one more time to get:

$$(P^{-1})'' = -(P^{-1})'P'P^{-1} - P^{-1}P'(P^{-1})' = 2P^{-1}P'P^{-1}P'P^{-1} \quad (40)$$

For any non-zero random vector y , define $v_\lambda = P'_\lambda P_\lambda^{-1} y$ and $\varphi_\lambda = y^T P_\lambda^{-1} y$. Equations (40) says that

$$\varphi_\lambda'' = y^T (P_\lambda^{-1})'' y = 2v_\lambda^T P_\lambda^{-1} v_\lambda \geq 0 \quad (41)$$

since P_λ^{-1} is positive definite. As the second order derivative is positive, we conclude that φ_λ is a convex function for λ over $[0, 1]$. As a result, for any $\lambda \in (0, 1)$, we have:

$$(1 - \lambda)\varphi(0) + \lambda\varphi(1) - \varphi_\lambda \geq 0 \\ \iff y^T [(1 - \lambda)M^{-1} + \lambda N^{-1} - ((1 - \lambda)M + \lambda N)^{-1}] y \geq 0 \quad (42)$$

Since y is arbitrary, this implies the matrix within the square bracket in (42) is positive semi-definite and hence:

$$(1 - \lambda)M^{-1} + \lambda N^{-1} \geq ((1 - \lambda)M + \lambda N)^{-1}$$

Please note that when $P' = N - M$ is invertible, v_λ is non-zero for non-zero y . The inequalities in (41) and (42) become strict and the matrix within the square bracket in (42) is positive definite instead of positive semi-definite. \square

PROOF. A **fourth proof** is to derive the convexity of the inverse of a matrix from the convexity of the function $f(t) = \frac{1}{t}$ for $t \geq 0$. Let $P = X^{-1/2} Y X^{-1/2}$. We want to prove that

$$(1 - \lambda)X^{-1} + \lambda Y^{-1} - ((1 - \lambda)X + \lambda Y)^{-1} \geq 0 \quad (43)$$

$$\iff X^{1/2} [(1 - \lambda)X^{-1} + \lambda Y^{-1} - ((1 - \lambda)X + \lambda Y)^{-1}] X^{1/2} \geq 0 \quad (44)$$

$$\iff (1 - \lambda)I + \lambda P^{-1} - ((1 - \lambda)I + \lambda P)^{-1} \geq 0 \quad (45)$$

where in inequality (43), we have left- and right- multiplied both sides by $X^{1/2}$. As P is positive definite, it can be unitary diagonalised and hence without loss of generality, we can assume that it is a diagonal matrix. So, the inequality reduces down to the scalar case $(1 - \lambda) + \lambda p_{ii}^{-1} \geq ((1 - \lambda) + \lambda p_{ii})^{-1}$, which is true using the fact that the function $f(t) = \frac{1}{t}$ is convex for $t \geq 0$ \square

The last two proofs relies on the fact that the result is also implied by the fact that the function $f(X) = \text{Tr}(X^{-1} y y^T) = \text{Tr}(y^T X^{-1} y)$ is convex for $X \in \mathcal{S}_{++}^p$ for any $y \in \mathbb{R}^p$. This comes from the nice

property that the Trace operator can commute and that the trace of a real number is itself.

PROOF. The **fifth proof** uses the fact that a positive second order derivative along any line is enough to prove convexity. Consider $S(t) = U + tV$ where U and V are symmetric positive definite. It is

enough to show that $\left. \frac{d^2}{dt^2} \text{Tr}(y^T S(t)^{-1} y) \right|_{t=0} \geq 0$ We have

$$S(t)^{-1} = (U(I + tU^{-1}V))^{-1} \\ = U^{-1} - tU^{-1}VU^{-1} + t^2U^{-1}VU^{-1}VU^{-1} + \dots \quad (46)$$

So

$$\left. \frac{d^2}{dt^2} \text{Tr}(S(t)^{-1}) \right|_{t=0} = 2\text{Tr}(U^{-1}VU^{-1}VU^{-1})$$

But $U^{-1}VU^{-1}VU^{-1} = WU^{-1}W^T$ where $W = U^{-1}V$ and U^{-1} is positive definite, so $WU^{-1}W^T$ is positive semi definite, which implies $\text{Tr}(WU^{-1}W^T) \geq 0$, which concludes the proof \square

PROOF. A final **sixth proof** is to relate this to eigen values. We can notice that the function $f(X) = \text{Tr}(X^{-1} y y^T)$ is indeed the sum of the inverse of eigen values denoted by λ_i .

$$\text{Tr}((y^T X^{-1} y)) = \sum_{i=1..n} \frac{1}{\lambda_i}$$

We know that the function that associates to a diagonal matrix with strictly positive terms its kth element (which turns out to be one of its eigen values but not necessarily its kth one) is linear, hence convex and concave. By the composition rules for convex function, with $g(x) = 1/x$, we can conclude that the inverse of the kth elements is convex for diagonal matrices with strictly positive term. Thus, the sum of the inverse of eigen values (defined as a sum of convex functions) is convex on the set of diagonal matrix with strictly positive term. We can conclude using the diagonalisation result of definite positive matrix (with $S = UDU^T$, U an orthonormal matrix, D a diagonal matrix with strictly positive term and $S \in \mathcal{S}_{++}^n$) to extend the convexity property to the set of \mathcal{S}_{++}^n and use also that $\text{Tr}(AB) = \text{Tr}(BA)$ \square

REFERENCES

- [1] Akimoto, Y., Auger, A., and Hansen, N. (2015). Continuous optimization and CMA-ES. *GECCO 2015, Madrid, Spain*, 1:313–344.
- [2] Akimoto, Y., Auger, A., and Hansen, N. (2016). CMA-ES and advanced adaptation mechanisms. *GECCO, Denver*, 2016:533–562.
- [3] Akimoto, Y., Nagata, Y., Ono, I., and Kobayashi, S. (2010). Bidirectional relation between cma evolution strategies and natural evolution strategies. *PPSN, XI(1)*:154–163.
- [4] Auger, A. and Hansen, N. (2009). Benchmarking the (1+1)-CMA-ES on the BBOB-2009 noisy testbed. *Companion Material, GECCO 2009*:2467–2472.
- [5] Auger, A. and Hansen, N. (2012). Tutorial CMA-ES: evolution strategies and covariance matrix adaptation. *Companion Material Proceedings*, 2012(12):827–848.
- [6] Auger, A., Schoenauer, M., and Vanhaccke, N. (2004). LS-CMA-ES: A second-order algorithm for covariance matrix adaptation. *PPSN VIII, 8th International Conference, Birmingham, UK, September 18-22, 2004, Proceedings*, 2004(2004):182–191.
- [7] Benhamou, E. and Saltiel, D. (2019). A Bayesian approach to CMA-ES. *arXiv e-prints*.
- [8] Cox, R. T. (1946). Probability, frequency, and reasonable expectation. *American Journal of Physics*, 14(2):1–13.
- [9] Glasmachers, T., Schaul, T., Yi, S., Wierstra, D., and Schmidhuber, J. (2010). Exponential natural evolution strategies. In: *Proceedings of Genetic and Evolutionary Computation Conference*, pp. 2010(2010):393–400.
- [10] Gomez, F., Schmidhuber, J., and Miikkulainen, R. (2008). Accelerated neural evolution through cooperatively coevolved synapses. *J. Mach. Learn. Res.*, 9:937–965.
- [11] Hansen, N. and Auger, A. (2011). CMA-ES: evolution strategies and covariance matrix adaptation. *GECCO 2011*, 2011(1):991–1010.
- [12] Hansen, N. and Auger, A. (2014). Evolution strategies and CMA-ES (covariance matrix adaptation). *GECCO Vancouver*, 2014(14):513–534.
- [13] Hansen, N. and Ostermeier, A. (2001). Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195.
- [14] Heidrich-Meisner, V. and Igel, C. (2009). Neuroevolution strategies for episodic reinforcement learning. *J. Algorithms*, 64(4):152–168.
- [15] Igel, C. (2010). Evolutionary kernel learning. In *Encyclopedia of Machine Learning and Data Mining*, pages 465–469. Springer, New-York.
- [16] Igel, C., Hansen, N., and Roth, S. (2007). Covariance matrix adaptation for multi-objective optimization. *Evol. Comput.*, 15(1):1–28.
- [17] Igel, C., Heidrich-Meisner, V., and Glasmachers, T. (2009). Shark. *Journal of Machine Learning Research*, 9:993–996.
- [18] Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge University Press, Cambridge.
- [19] Loshchilov, I. and Hutter, F. (2016). CMA-ES for Hyperparameter Optimization of Deep Neural Networks. *arXiv e-prints*, 1604(April):arXiv:1604.07269.
- [20] Mühlenbein, H., Schomisch, M., and Born, J. (1991). The parallel genetic algorithm as function optimizer. *Parallel Computing*, 17(6):619 – 632.
- [21] Ollivier, Y., Arnold, L., Auger, A., and Hansen, N. (2017). Information-geometric optimization algorithms: A unifying picture via invariance principles. *J. Mach. Learn. Res.*, 18(1):564–628.
- [22] Rastrigin, L. A. (1974). *Systems of extremal control*. Mir, Moscow.
- [23] Schervish, M. (1996). *Theory of Statistics*. Springer Series in Statistics. Springer, New York.
- [24] Varelas, K., Auger, A., Brockhoff, D., Hansen, N., ElHara, O. A., Semet, Y., Kassab, R., and Barbaresco, F. (2018). A comparative study of large-scale variants of CMA-ES. *PPSN XV - 15th International Conference, Coimbra, Portugal*, 15(2018):3–15.