



HAL
open science

NGO-GM: Natural Gradient Optimization for Graphical Models

Eric Benhamou, Jamal Atifl, Rida Laraki, David Saltiel

► **To cite this version:**

Eric Benhamou, Jamal Atifl, Rida Laraki, David Saltiel. NGO-GM: Natural Gradient Optimization for Graphical Models. 2020. hal-02886514

HAL Id: hal-02886514

[**https://hal.science/hal-02886514v1**](https://hal.science/hal-02886514v1)

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.

NGO-GM: Natural Gradient Optimization for Graphical Models

Eric Benhamou
A.I Square Connect,*Lamsade†

Jamal Atif
Lamsade†

Rida Laraki
Lamsade†

David Saltiel
A.I Square Connect*

Abstract

This paper deals with estimating model parameters in graphical models. We reformulate it as an information geometric optimization problem and introduce a natural gradient descent strategy that incorporates additional meta parameters. We show that our approach is a strong alternative to the celebrated EM approach for learning in graphical models. Actually, our natural gradient based strategy leads to learning optimal parameters for the final objective function without artificially trying to fit a distribution that may not correspond to the real one. We support our theoretical findings with the question of trend detection in financial markets and show that the learned model performs better than traditional practitioner methods and is less prone to overfitting.

1 Introduction

One of the most challenging question in social sciences and in particular financial markets, is to be able from past observations to make some meaningful predictions. Part of the challenge comes from multiple difficulties. First of all, there is no universally established physical law and our model will be at best a simplified version and at worst a complete non sense. Second, sequential information does not imply stationarity of the process: data may incur regime changes. Third, we should find the right balance between not enough and too much modeling.

To be able to represent connection between our variables, an efficient and meaningful framework is graphical models. As stated in [Jordan, 2012], *graphical models are a*

marriage between probability theory and graph theory. They are powerful as they provide a condensed way to represent variables dependencies. The graphical representation allows not only compacting information, it also provides a powerful formalism for reasoning under uncertainty. It represents knowledge about the dynamics of the variables, their dependencies and their conditional distribution, hence sometimes called also Dynamic Bayesian Networks (DBN).

Graphical models exploit latent variables to make the model better in terms of explanation power. By defining a joint distribution over visible and latent variables, the corresponding distribution of the observed variables is obtained by marginalization. This has the nice property to express relatively complex distributions in terms of more tractable joint distributions over the expanded variable space. The canonical atomic example of hidden variable models is the mixture distribution in which the hidden variable is the discrete component label that provides the corresponding distribution for the observable variable. The static version leads to the Gaussian mixture model and in continuous space to the factor analysis model while the dynamic version leads respectively to HMM and the Kalman filter model, often referred to as the state space model and represented by figure 1. However, dynamic graphical models can be more complex as given by figure 2 with combination of the Kalman filter (KF) model and echo state networks (ESN) or by figure 3 with feedback from past observations to next step latent variables. We will exploit a combination of the latter in our numerical experiments.

To finalize our model, we need to solve the issue of learning the its parameters. The typical learning approach is the Expectation Maximization (EM) algorithm. It was initially developed for mixture models, in particular Gaussian mixtures and other natural laws such as Poisson, binomial, multinomial and exponential distributions, see [Hartley, 1958] and [Dempster et al., 1977]. It was only when the link between latent variables and Kalman fil-

* A.I Square Connect, Neuilly sur Seine, France

† Lamsade, Université Paris Dauphine, PSL, France




Figure 1: State Space model represented as a Bayesian Probabilistic Graphical model. Each vertical slice represents a time step. Nodes in gray (resp. white) represent observable (resp. non observable or latent) variables. This State Space model encompasses HMM and KF models




Figure 2: Example of another dynamic Bayesian network combining Kalman filter (KF) model and echo state networks (ESN). This is another example of a multi-input several multi-outputs (MISMO) forecasting model. It is used frequently in time series forecast (see for instance [Xiao et al., 2017])

ter models was made that it became clear that it could also be applied to Kalman and extended Kalman filters (see [Cappe et al., 2010] or [Einicke et al., 2010]). The EM method is so far the state of the art method for learning graphical model parameters as it provides an efficient way to find model parameters in a fraction of seconds (see for instance [Neal and Hinton, 1999], but also [Pfeifer and Protzel, 2018], [Li and McCormick, 2017], [Robin et al., 2017], [Levine, 2018]). Interestingly [Xu and Jordan, 1996] shows that the EM method can be viewed as a gradient descent where the decrement is computed as the projection of the gradient, making it a variable metric gradient ascent. Similarly, [Amari, 1995] proves that the E and M steps in EM can be interpreted as dual projected gradient flows under dual affine connections using information geometry. This advocates to find alternatives that are also gradient descent but in the natural space.

We argue that although EM enjoys several nice properties, it misses the point that graphical models are imprecise and simplified models for the reality especially when tackling complex problems such as time series forecasting. In particular, whenever we apply graphical models to economics and finance, we are forced to make some mod-




Figure 3: Example of a dynamic Bayesian network with connection between past observable variables and latent variables

eling assumptions about the state dynamics and the graph topology (the DBN structure). These assumptions may be incorrectly specified and add some bias compared to reality. Trying to use a best fit approach through maximum likelihood estimation and Kullback Leibler divergence optimization may miss this point and try to fit at all cost the model on data. It does not factor in the interdependence between our graphical model and the actions related to this graphical model. In the case of social sciences, if the graphical model is used to make a forecast which is then used to make a parametric action, the EM method does not take into account the interrelation between parameters of the graphical model denoted by θ and the ones of the action denoted by τ . To measure the efficiency of the full set of parameters (θ, τ) , we rely on a cost function that can contain a regularizer to lead to smooth parameters. In finance, this cost function can be for instance a measure of the performance of our prediction taking into account the risk in our actions.

We present here a new approach that takes a radical point of view and focuses on the final efficiency of our model. Graphical model parameters are now estimated in terms of their efficiency for the cost function together with the action parameters rather than just their distributional fit to data. We rely on information geometric optimization to find a local optimum for our final cost function. Our key features are the following:

- it is possible to directly optimize the cost function with a stochastic optimization approach and in particular the CMA-ES method;
- this approach computes a natural gradient in the implicit Fisher information Riemannian manifold;
- it is a good alternative to the EM approach as it does not fit at any cost the distribution of our graphical model to reality but rather looks at model efficiency measured by a loss function related to the problem under study;

- the estimation of our model parameters is performed jointly with the action strategy parameters;
- numerical results show that the overfitting issue of this approach due to local minima is less than the EM approach as it takes into account that the model dynamics is incorrectly specified.

The rest of the paper is organized as follows. Section 2 presents the overall framework and the resulting optimization problem. Section 3 provides some theoretical arguments that favor stochastic optimization approaches based on Information Geometric Optimization (IGO). Section 4 discusses an example in finance. Our method outperforms traditional trend following methods by far. We finally conclude about some possible extensions and further experiments.

2 Settings

Suppose we have determined an architecture for our Dynamic Bayesian Network. This may be inspired by combinations of simple network architectures such as those in Figures 1, 2, 3. This model is used for some specific goal. In our case, it is used to forecast some times series. But this is not our final objective! We are interested in using this forecast to perform a specific action. In the case of a financial market algorithmic trading strategy, we will use the forecast to make an informed decision and decide whether we should buy, sell or do nothing with a financial asset. To keep things simple, we will assume that when we take our decision, we target a pre-determined movement amplitude, materialized by a profit target and a stop loss level. The signal for the action is given by the difference between the forecast generated by our graphical model and the last price. To avoid false signals, we add an additional threshold parameter for our action and consider that there is an uptrend signal (respectively a downtrend signal) if the forecast is above the last price plus a threshold (respectively below the last price minus a threshold). Using a fixed price target, a stop loss and a threshold is quite realistic and is done by many practitioners as presented in various papers, e.g. [Labadie and Lehalle, 2010], [Giuseppe Di Graziano, 2014], [Fung, 2017], or [Vezeris et al., 2018]. The profit target ensures that the strategy locks the profit in real money and is technically corresponding to a limit order while the stop loss, technically corresponding to a stop order, safeguards the overall risk by limiting losses whenever the market backfires and contradicts the presumed direction.

The price target, stop loss and threshold are three additional parameters that govern our action. These parameters are closely interconnected to our network parameters.

If the network’s prediction is accurate enough, we should aim for a large price movement, hence a large price target and a small stop loss. If it is not accurate enough, on the contrary, we should reduce the price target and increase the stop loss. We hold the position until either the trade reaches the profit target and exits in a profit or it touches the stop loss level and exits with a loss. The objective function to measure the performance of the model is the eponymous Sharpe ratio (introduced in [Sharpe, 1966]) that corresponds to the ratio of the strategy net return over its standard deviation. We add an $L1$ regularization term to ensure sparsity of the dynamic graphical model parameters. Although this approach may recall reinforcement learning theory, the dynamic nature of our graphical model makes the usage of standard reinforcement learning tools such as Q-learning inappropriate.

On this simple example, we clearly see that optimizing graphical model parameters in a standalone fashion is not optimal and that these parameters should be evaluated together with the action parameters at the light of our loss function as represented by figure 4. We want to maximize the Sharpe ratio of our trading strategy over time. This optimization problem is non convex with potentially many local optima. Moreover, the binary nature of our action, i.e. we sell or buy as soon as our forecast hits the past level plus or minus the threshold, leads to discontinuities.

The fundamental difference between EM method and ours is to jointly optimize the model and the action parameters, learning (θ, τ) . Because model and action parameters are interconnected, this approach well fits our initial desiderata. It optimizes our final criteria, the loss function, and not the maximum likelihood of our graphical model. We will present in the next section some important results from information geometric optimization that allow us to tackle this complex optimization problem using a natural gradient descent.




Figure 4: Learning process for our graphical model. Decide an architecture for our graphical model with parameters θ , combine with a strategy with parameters τ to create a systematic algorithm. Select a loss function. Use natural gradient optimization to find best parameters (θ^*, τ^*)

3 Information Geometric Optimization

3.1 Natural gradient descent

Compared to the usual gradient descent method, the natural gradient takes into account the intrinsic geometric structure of the underlying Riemannian space. Since the seminal work of [Amari, 1998], it has been widely adopted by the machine learning community [Pascanu and Bengio, 2014], [Bernacchia et al., 2018]. In gradient descent, the usual update step is $\theta^{t+1} = \theta^t - \nu_t \frac{\partial U(\theta)}{\partial \theta^t}$ where θ^t is the optimization parameter at iteration t , ν_t the learning rate, and U the cost function. The method of natural gradients replacement proposes to change the gradient descent with $\theta^{t+1} = \theta^t - \nu_t G^{-1}(\theta^t) \frac{\partial U(\theta)}{\partial \theta^t}$ where $G(\theta^t)$ is a matrix that defines the Riemannian metric in the space of the parameters. This is very powerful as it computes the fastest descent when looking at the Riemannian metric induced by the parameters. In distributions space, the Riemannian metric is associated with the Fisher information matrix [Rao, 1945, Jeffreys, 1946]. In general, it is difficult to apply the natural gradient descent because we need to invert the Fisher information matrix, which is computationally heavy, except in some particular cases where there exists a closed form like for instance for exponential family distributions. We will follow the seminal works of [Ollivier et al., 2017] and [Akimoto, 2012] and present a stochastic optimization method that performs numerically natural gradients efficiently.

We denote by \mathbb{X} a metric space and associate to \mathbb{X} a Borel σ -field and a measure on \mathbb{X} denoted by \mathcal{F} and μ respectively. Typically \mathbb{X} is \mathbb{R}^d and μ is the Lebesgue measure on \mathbb{R}^d . We are interested in minimizing a μ -measurable real value function $f : \mathbb{X} \rightarrow \mathbb{R}$. In order to make our optimization invariant with respect to various standard transformations of f , instead of minimizing f , it is better to find the minimum of a *loss* (also referred to as an *invariant cost*) function that is invariant to any strictly increasing transformation of f . We define L_f as the μ -measured volume of the unit ball in the functional space f with radius $f(x)$, that is

$$L_f : x \mapsto \mu[y : f(y) \leq f(x)].$$

Said differently, L_f is the measure of all the elements whose value is less or equal to $f(x)$. We are interested in finding the optimum of the loss function with respect to a family of probability distributions P_θ on \mathbb{X} , $\theta \in \Theta$. Compared to standard optimization, we define a *quasi-objective* function, sometimes referred to as a utility function, U_f , on the parameter space Θ defined as the expected value of our loss function L_f over

the space of distributions to the power $2/d$: P_θ , namely $U_f(\theta) = \mathbb{E}_{X \sim P_\theta} [L_f^{2/d}(X)]$. The intuition behind the exponent $2/d$ is to ensure that the utility is in a sense dimensionless and homogeneous to the square of a Euclidean distance. This is because in dimension d , a unit ball is homogeneous to a radius of dimension d , where as the square of a Euclidean distance is of dimension 2. We will see in the practical example of a quadratic function (see subsection 3.5) that this choice of exponent makes the explicit computation of the natural gradient easy.

In [Ollivier et al., 2017], this utility function is defined as the opposite of the weighted quantile, $U_f^{IGO}(x) = -w(P_{\theta^t}[y : f(y) \leq f(x)])$, where $w : [0, 1] \mapsto \mathbb{R}$ is non-increasing weight function. This choice is not easy to analyze as the quantile $P_{\theta^t}[y : f(y) \leq f(x)]$ depends on the current parameter θ^t . At each iteration, the loss function $L_f^{IGO}(x)$ changes making it hard to analyze. In [Akimoto, 2012], the utility (referred to as the quasi objective) function is defined as $U_f^{NGA} = \mathbb{E}_{X \sim P_\theta} [\nu[y : f(y) \leq f(x)]]$ where ν represents any monotonically increasing set function on \mathcal{F} , i.e., $\nu(A) \leq \nu(B)$ for any $A, B \in \mathcal{F}$ s.t. $A \subseteq B$. Our function is somehow simpler as the loss function is just the expected volume of all the elements whose value is less or equal to $f(X)$ for $X \sim P_\theta$ while the Utility function is similar but homogeneous to a Euclidean distance squared.

We are interested in performing a natural gradient descent on a Riemannian manifold given by $(\Theta, G(\theta))$ equipped with the Fisher information metric G_θ . We choose the Fisher metric because it is the unique metric that does not depend on the choice of parameterization as explained in [Amari and Nagaoka, 2007]. The natural gradient is easy to obtain and is just the gradient with respect to the Fisher information matrix metric. One can “easily” compute it and get that it is provided by the product of the inverse of the Fisher information matrix G_θ with the standard or *vanilla* gradient $\nabla U_f(\theta)$ of the function to minimize. We can then apply a natural gradient descent as follows:

$$\theta^{t+1} = \theta^t - \nu_t G_{\theta^t}^{-1} \nabla U_f(\theta^t), \quad (1)$$

where ν_t is the learning rate. Compared to the vanilla gradient, we have an additional term given by $G_{\theta^t}^{-1}$. We present below an algorithm that allows performing the natural gradient descent in a general case. To make things clearer, we will present the algorithm in the case where we can explicitly compute the inverse of the Fisher information matrix and the standard gradient $\nabla U_f(\theta)$. We will call this the *Closed form Natural Gradient* or (NGD). We will then present the Monte-Carlo NGD that allows performing the natural gradient descent in a general framework efficiently.

Historically, this approach of computing a natural gradient in the space of distributions was done empirically and without strong theoretical arguments in the Evolution Strategies community. The most prominent approach belonging to this work which is indeed a natural gradient method is the CMA-ES algorithm [Hansen and Ostermeier, 2001]. Although CMA-ES has been state of the art in this line of research as shown by the various benchmarks of the Comparing Continuous Optimisers (COCO) INRIA platform for ill-posed functions, it was only later, after the works of [Akimoto et al., 2010], [Glamachers et al., 2010], [Akimoto et al., 2012], [Akimoto, 2012], and the deep theoretical study of [Ollivier et al., 2017]) that the community realizes that this algorithm performs a natural gradient descent. In the following we adapt these theoretical proofs to the setting of our graphical model learning problem. The CMA-ES algorithm has led to a large number of papers and articles and we refer the interested reader to [Hansen and Ostermeier, 2001], [Auger et al., 2004], [Igel et al., 2007], [Auger and Hansen, 2009], [Hansen and Auger, 2011], [Auger and Hansen, 2012], [Hansen and Auger, 2014], [Akimoto et al., 2015], [Akimoto et al., 2016], [Ollivier et al., 2017] and [Varelas et al., 2018] to cite a few.

3.2 Closed Form (CF) NGD Algorithm

In order to study the key property of the CMA-ES, we focus on the case where $\mathbb{X} = \mathbb{R}^d$ and suppose the measure μ to be the Lebesgue measure on \mathbb{R}^d , denoted μ_{Leb} , while \mathcal{F} denotes the Borel σ -field \mathcal{B}^d on \mathbb{R}^d .

Concerning the sampling distribution, we will work with the Gaussian distribution as it has the maximum entropy among distributions with known first two moments. Our distribution space, denoted by P_θ is parameterized by the parameter $\theta \in \Theta$. We adopt the traditional moment parameters, that is the mean vector $\mu(\theta)$, which is in \mathbb{R}^d and the covariance matrix $\Sigma(\theta)$, which is a symmetric and positive definite matrix of dimension $d \times d$, which leads that P_θ is indeed the normal written as $\mathcal{N}(\mu(\theta), \Sigma(\theta))$.

The loss function $L_f(x)$ is naturally associated to the Lebesgue measure μ_{Leb} and defined as $L_f(x) = \mu_{\text{Leb}}[y : f(y) \leq f(x)]$. In order to optimize our function, we look at the infimum of

$$U_f(\theta) = \mathbb{E}_{X \sim P_\theta} \left[L_f^{2/d}(X) \right].$$

If there is only one optimum, this translates into finding the domain Θ where the mean vector equals the global minimum of f while the covariance matrix is null.

It is worth noticing that the choice of the moment parameterization of Gaussian distributions does affect the

behavior of the natural gradient update (1) with finite learning rate ν^t . However, the steepest direction of U_f on the statistical manifold Θ is invariant under the choice of parameterization as explained in [Ollivier et al., 2017]. Using explicit computation for the normal, we get the natural gradient descent summarized by proposition 3.1 below.

Proposition 3.1. *If the loss function is squared integrable ($\mathbb{E}_{X \sim P_\theta} \left[L_f^2(X) \right] < \infty$), using two different learning rates (ν_μ^t, ν_Σ^t) for the mean vector and the covariance, the natural gradient descent (1) writes:*

$$\begin{aligned} \mu^{t+1} &= \mu^t - \nu_\mu^t \delta\mu^t, \\ \Sigma^{t+1} &= \Sigma^t - \nu_\Sigma^t \delta\Sigma^t, \end{aligned} \tag{2}$$

with $\delta\mu^t = \mathbb{E}_{X \sim P_{\theta^t}} [L_f^{2/d}(X)(X - \mu^t)]$

$$\delta\Sigma^t = \mathbb{E}_{X \sim P_{\theta^t}} [L_f^{2/d}(X)((X - \mu^t)(X - \mu^t)^T - \Sigma^t)]$$

Proof. Refer to A in the supplementary material for more details. \square

Equations (2) defines the closed-form or deterministic natural gradient descent (NGD) method, which is an ideal case where we know the optimum solution. This ideal case is useful for deriving fruitful properties of our optimization method but is useless in practice.

3.3 Monte-Carlo (MC) NGD Algorithm

We now turn to the real algorithm that is efficient in real-world situations. It tackles the issue of unknown and untractable gradient for our quasi objective function: $\nabla U_f(\theta)$. The last resort solution is to estimate the gradient with Monte Carlo simulations. We approximate the natural gradient and simulate the natural gradient descent with the algorithm 1.

This algorithm is the Monte Carlo version of our NGD algorithm. In [Ollivier et al., 2017], this is referred as the stochastic NGD algorithm. Compared to the closed-form NGD algorithm, we evaluate $\nabla U_f(\theta)$ thanks to a Monte Carlo simulation. This makes this algorithm stochastic in nature. In order to have the same behavior each time we run this algorithm, we freeze random seeds to ensure similar results for the normal random numbers. The core of the algorithm is to generate n samples x_i from a multivariate Gaussian distribution $\mathcal{N}(\mu^t, \Sigma^t)$, evaluate their value, computes the loss functions $L_f(x_i)$ afterwards. The estimates $\widehat{L}_f(x_i)$ are obtained as follows. In order to get the intuition of the Monte Carlo approximation, recall that

Algorithm 1 Monte-Carlo NGD Algorithm

while Not Converged **do**

Simulate n random normal vectors denoted by z_i according $\mathcal{N}(0, Id_d)$, infer $x_i = m_t + \sqrt{\Sigma^t} z_i$, and evaluate their values $f(x_1), \dots, f(x_n)$
Estimate the Loss function $L_f(x_i)$ as

$$\widehat{L}_f(x_i) = \frac{\sqrt{(2\pi)^d \det(\Sigma)}}{n} \sum_{j: f(x_j) \leq f(x_i)} \exp\left(-\frac{\|z_j\|^2}{2}\right)$$

Estimate the natural gradient $\delta\mu^t$ and $\delta\Sigma^t$ as

$$\begin{aligned} \widehat{\delta\mu}^t &= \frac{1}{n} \sum_{i=1}^n \left(\widehat{L}_f(x_i)\right)^{2/d} (x_i - \mu^t) \\ \widehat{\delta\Sigma}^t &= \frac{1}{n} \sum_{i=1}^n \left(\widehat{L}_f(x_i)\right)^{2/d} \left((x_i - \mu^t)(x_i - \mu^t)^\top - \Sigma^t\right) \end{aligned} \quad (3)$$

Update parameters with ngd as

$$\begin{aligned} \mu^{t+1} &= \mu^t - \nu_\mu \widehat{\delta\mu}^t \\ \text{and } \Sigma^{t+1} &= \Sigma^t - \nu_\Sigma \widehat{\delta\Sigma}^t. \end{aligned}$$

end while

the loss function is given by:

$$\begin{aligned} L_f(x) &= \int \mathbf{1}_{\{f(y) \leq f(x)\}} dy \\ &= \int \frac{\mathbf{1}_{\{f(y) \leq f(x)\}}}{p_{\theta^t}(y)} p_{\theta^t}(y) dy \end{aligned}$$

In our Monte-Carlo algorithm, a first step consists in computing the integral as the corresponding Riemann sum:

$$\widehat{L}_f(x) = \frac{1}{n} \sum_{j=1}^n \frac{\mathbf{1}_{\{f(x_j) \leq f(x)\}}}{p_{\theta^t}(x_j)} \quad (4)$$

For the multivariate Gaussian, the probability weights are given by $p_{\theta^t}(x_j) = 1/\sqrt{(2\pi)^d \det(\Sigma)} \exp(-\|z_j\|^2/2)$ where the z_i are standardized normal $\mathcal{N}(0, Id_d)$ draws. It is worth noticing that these weights are not the traditional ones of the CMA-ES algorithm. CMA-ES relies rather on logarithmic weights [Hansen and Auger, 2014] and takes a fraction (denoted by μ , that means something different from our mean vector) out of λ candidates. In this setting, we use all the candidates to compute the new mean and covariance. From a theoretical point of view, the weights introduced in [Akimoto, 2012] are more meaningful and make the study of the property of this stochastic optimization algorithm simpler. The next step is to estimate the

natural gradient according to equations (2) as follows:

$$\begin{aligned} \mathbb{E}_{X \sim P_{\theta^t}} \left[L_f^{2/d}(X) (X - \mu^t) \right] \\ = \frac{1}{n} \sum_{i=1}^n \left(\widehat{L}_f(x_i)\right)^{2/d} (x_i - \mu^t) \end{aligned} \quad (5)$$

$$\begin{aligned} \mathbb{E}_{X \sim P_{\theta^t}} \left[L_f^{2/d}(X) \left((X - \mu^t)(X - \mu^t)^\top - \Sigma^t \right) \right] \\ = \frac{1}{n} \sum_{i=1}^n \left(\widehat{L}_f(x_i)\right)^{2/d} \left((x_i - \mu^t)(x_i - \mu^t)^\top - \Sigma^t \right) \end{aligned} \quad (6)$$

In the final natural gradient update, it is important to have both the mean and the covariance learning rates not too large. We will see in the next section the reason of the limitation of the learning rates when looking at the positivity of the covariance matrix at each step as well as when examining convergence properties.

3.4 Basic Properties

Invariance. Invariance properties are fundamental for the efficiency of optimization algorithm. Our stochastic optimization has two major invariance properties: under monotonic transformations of the objective function and under affine transformations of the search space. Invariance under monotonic transformation means the algorithm performs equally well on the function f and on $g \circ f$, the composition of f with any function g strictly increasing. This explains why it performs well on ill-conditioned functions whereas conventional gradient methods like Newton method relies on convexity properties of the objective function and handles this non-convex problems less well. This invariance property is a direct consequence of the loss function. Invariance under affine transformations of the search space is the key idea behind the Newton's method. The adaptation of the covariance matrix at each step explains the performance of this algorithm on ill-conditioned objective functions.

Positivity. The covariance matrix must be positive definite and symmetric at each iteration. Although [Krause et al., 2016] recently took the additional constraint to Cholesky decompose the covariance matrix, we can explicitly compute the condition on the learning rate ν_Σ^t to ensure that in our NGD algorithms the covariance matrix always remains positive definite symmetric. Proposition 3.2 shows that the learning rate has to be small enough. Should the learning rate be larger than the critical value, the covariance matrix would be non positive.

Proposition 3.2. *Provided that the matrix Σ^0 and the variances $I_d - \nu_\Sigma^t \sqrt{\Sigma^t}^{-1} \delta\Sigma^t \sqrt{\Sigma^t}^{-1}$ are positive definite, the covariance matrix Σ^t remains positive definite*

(and symmetric) for each t in the closed form NGD algorithm. Should the matrix $I_d - \nu_{\Sigma}^t \sqrt{\Sigma^t}^{-1} \delta \Sigma^t \sqrt{\Sigma^t}^{-1}$ not be positive (with eigen values negative or null), Σ^{t+1} would not be positive. In the Monte-Carlo NGD algorithm, the condition is modified into $I_d - \nu_{\Sigma}^t \sqrt{\Sigma^t}^{-1} \widehat{\delta \Sigma^t} \sqrt{\Sigma^t}^{-1}$ should be positive definite for any t .

Proof. Refer to B in the supplementary material for more details. \square

Remark 3.1. In case, the matrix $\sqrt{\Sigma^t}^{-1} \delta \Sigma^t \sqrt{\Sigma^t}^{-1}$ for the Closed Form case and $\sqrt{\Sigma^t}^{-1} \widehat{\delta \Sigma^t} \sqrt{\Sigma^t}^{-1}$ for the Monte-Carlo case has some positive eigen values, this result can be re-expressed trivially in terms of eigenvalues and state that the learning rate is bounded by $\nu_{\Sigma}^t < 1/\lambda_1(\sqrt{\Sigma^t}^{-1} \delta \Sigma^t \sqrt{\Sigma^t}^{-1})$ for the Closed Form case and $\nu_{\Sigma}^t < 1/\lambda_1(\sqrt{\Sigma^t}^{-1} \widehat{\delta \Sigma^t} \sqrt{\Sigma^t}^{-1})$ for the Monte Carlo case where $\lambda_1(M)$ denotes the largest eigenvalue of a matrix M . This is quite intuitive. The greater the critical largest eigenvalue is, the bigger the gradient is. Hence the learning rate should not be too large to go too far. Reciprocally the smaller the critical largest eigenvalue is, the smaller the gradient is. Should the matrices previously mentioned have only negative values, there would be no bound on the learning rate. This result is already presented in [Akimoto, 2012] but without mentioning that the two matrices shall not be necessarily positive. We will see in the trivial case of the quadratic function in subsection 3.5, that the matrix in the Closed Form case can be explicitly computed and given by $\kappa \Sigma^t H \Sigma^t$ with κ given by proposition 3.4. Should H not be positive, the matrix $\sqrt{\Sigma^t}^{-1} \delta \Sigma^t \sqrt{\Sigma^t}^{-1}$ would not be neither.

Convergence. Although the gradient estimator defined in (3) may not be necessarily unbiased, yet it converges to the true natural gradient as proved by proposition 3.3. This implies in particular that the Monte-Carlo NGD approximates well the closed-form NGD when we increase the sample size n . Let us denote by $\nabla : U \mapsto G_{\theta}^{-1} \nabla U$ the natural gradient operator, and by $\tilde{\nabla}_{\theta^n} : \widehat{U} \mapsto G_{\theta}^{-1} \nabla \widehat{U}$ the Monte Carlo estimate of the natural gradient.

Proposition 3.3. Let X_1, \dots, X_n be independent and identically distributed (i.i.d) random vectors following P_{θ} . Let us denote $\text{vec}(M)$ the vectorization of matrix M , that is the vector obtained by stacking matrix columns on top of one another. Let $\widehat{L}_f(x)$ and $\widehat{\nabla}_{\theta^n} = [(\widehat{\delta \mu}^t)^T, \text{vec}(\widehat{\delta \Sigma^t})^T]^T$ be the loss function estimator given by (4) and the natural gradient estimator where in (3) the x_1, \dots, x_n are replaced by X_1, \dots, X_n .

If the loss function is squared integrable: $\mathbb{E}[L_f^2(X)] < \infty$, then the gradient estimator of our utility function

converges almost surely: $\tilde{\nabla}_{\theta^n} \widehat{U} \xrightarrow{\text{a.s.}} \nabla U_f(\theta)$, where $\xrightarrow{\text{a.s.}}$ represents almost sure convergence.

Proof. Refer to C in the supplementary material. This result was proved as early as 2011 by [Ollivier et al., 2017] although their final manuscript was published a few years latter. Our approach though is quite original as we use the fundamental property of almost sure convergence that implies also Cesaro convergence. As Monte Carlo estimates are in nature Cesaro sums, this remarkable property ensures convergence of the MC NGD algorithm. Intrinsically, it means our Monte Carlo estimate is consistent. This is the theoretical foundation for our algorithm. The square integrability condition $\mathbb{E}[L_f^2(X)] < \infty$ is crucial to guarantee the convergence of the algorithm. \square

3.5 Convergence Properties

Since the MC NGD algorithm is quite complex, we are forced to simplify the problem and look at convergence properties only for the CF NGD algorithm. We additionally restrict ourselves to the canonical convex-quadratic function $f(x) = \frac{1}{2} x^T H x$, as introduced in [Auger et al., 2004], where H is a positive definite symmetric matrix. We also make assumptions on our learning rates: there exist $\nu_{\mu, \min} > 0$ and $\nu_{\Sigma, \min} > 0$ such that

$$\nu_{\mu, \min} \leq \nu_{\mu}^t \lambda_1((\Sigma^t)^{-1} \delta \Sigma^t) \leq 1, \quad (7)$$

$$\nu_{\Sigma, \min} \leq \nu_{\Sigma}^t \lambda_1((\Sigma^t)^{-1} \delta \Sigma^t) \leq 1/2. \quad (8)$$

A way to satisfy these inequalities is to set $\nu_{\mu}^t = \alpha_{\mu} / \lambda_1((\Sigma^t)^{-1} \delta \Sigma^t)$ (respectively $\nu_{\Sigma}^t = \alpha_{\Sigma} / \lambda_1((\Sigma^t)^{-1} \delta \Sigma^t)$) with α_{μ} (resp. α_{Σ}) being a learning rate constant for μ (resp. C) such that $0 < \alpha_{\mu} \leq 1$ (resp. $0 < \alpha_{\Sigma} \leq 1/2$).

With these assumptions, we can prove the global convergence of μ and C with linear convergence speed as stated by the proposition 3.4 below. The optimum of our quadratic test function is zero. In the following, we denote by $\|M\|$ the Frobenius norm of M , namely $\|M\| = \text{Tr}^{1/2}(M^T M)$.

Proposition 3.4. For the canonical quadratic optimization, the natural gradient descent is explicit and given by

$$\begin{aligned} \mu^{t+1} &= \mu^t - \nu_{\mu}^t \kappa \Sigma^t H \mu^t \\ \text{and} \quad \Sigma^{t+1} &= \Sigma^t - \nu_{\Sigma}^t \kappa \Sigma^t H \Sigma^t \end{aligned} \quad (9)$$

where

$$\kappa = \frac{2\pi}{\left(\Gamma\left(\frac{d}{2} + 1\right) \sqrt{\det(H)}\right)^{2/d}} \quad (10)$$

with $\Gamma(z) = \int_0^{+\infty} t^{z-1} e^{-t} dt$ the regular Gamma function. Suppose (7) and (8) hold, then $\|\mu^t\|$ (resp. $\|\Sigma^t\|$)

converges to zero with a rate of convergence defined as the lim sup of the ratio of two consecutive terms upper bounded by $1 - \nu_{\mu, \min}$ (respectively $1 - \nu_{\Sigma, \min}$).

Proof. Refer to D in the supplementary material. The intuition of the result of the convergence appeared as early as in [Auger et al., 2004] and was progressively improved and more detailed in [Arnold and Hansen, 2010], [Akimoto, 2012], [Beyer, 2014] and lately [Hansen, 2016] among others. The novelty of our proof is the explicit computation of the natural gradient descent update equations with the true computation of the constant κ in the update equations (9). \square

Remark 3.2. Proposition 3.4 is remarkable because it indicates that the rate of convergence is linear, which is unexpectedly fast for this kind of algorithms and explains why the CMA-ES is currently the state of the art solution for several complex optimization settings.

4 Numerical Results

In order to test the efficiency of our optimization procedure that results in a CMA-ES optimization for Learning parameters of our graphical model and the price target and stop loss, we look at a trend following algorithm where we enter a long (respectively short) trade if the dynamic Bayesian network forecast is above (respectively below) the previous day closing price. For each comparison, we add an offset μ to avoid triggering false alarm signals. We set for each trade a pre-determined profit and stop loss target in ticks. The chosen architecture for our graphical model is a combination of model of Figure 2 and 3 and given in Figure 5.




Figure 5: Combination of a multi state Kalman filter and echo neural networks (ESN) with feedback effect of observable variable on next latent variables

The corresponding dynamic is composed of two latent variables representing short and long term effects, denote respectively: $x_{.1}$ and $x_{.2}$. We assume that long term effect impacts the next long term latent variable while there is no influence of short term latent variable on the next long

term effect latent variable. This results in the following equations:

$$\mathbf{x}_{t+1} = \Phi \mathbf{x}_t + \mathbf{c}_t + \mathbf{w}_t \quad (11)$$

$$\mathbf{z}_t = \mathbf{H} \mathbf{x}_t + \mathbf{v}_t \quad (12)$$

We assume that the observation noise \mathbf{v}_t follows a multivariate normal distribution with zero mean and covariance matrix given by \mathbf{R}_t : $\mathbf{v}_t \sim \mathcal{N}(0, \mathbf{R}_t)$. In addition, the initial state, and noise vectors at each step $\mathbf{x}_0, \mathbf{w}_1, \dots, \mathbf{w}_t, \mathbf{v}_1, \dots, \mathbf{v}_t$ are assumed to be all mutually independent. We also denote by $\mathbf{P}_t = \text{Cov}(\mathbf{x}_t)$ the covariance matrix of \mathbf{x}_t . We assume the following parameters:

$$\Phi(x) = \begin{bmatrix} p_1 & p_2 \\ 0 & p_3 \end{bmatrix}, \mathbf{H} = \begin{bmatrix} p_4 \\ p_5 \end{bmatrix}, \mathbf{Q}_{t=0} = \begin{bmatrix} p_6^2 & p_6 p_7 \\ p_7 p_6 & p_8^2 \end{bmatrix}, \\ \mathbf{R}_{t=0} = [p_9], \mathbf{P}_{t=0} = \begin{bmatrix} p_{10} & 0 \\ 0 & p_{11} \end{bmatrix}, \mathbf{c}_t = \begin{bmatrix} p_{12}(.5 + p_{13} - k_t) \\ p_{14}(.5 + p_{15} - k_t) \end{bmatrix} \quad (13)$$

The variable k_t is computed as the ratio of the difference between the last price observation and the running minimum over the difference between the running maximum and minimum observed over 20 periods. The pseudo code is given in Algorithm 2.

Algorithm 2 Graphical model Trend following algorithm

Initialize common trade details

SetProfitTarget(target)

SetStopLoss(stop_loss)

while Not In Position **do**

if DBN(p_1, \dots, p_n).Predict[0] \geq Close[0] + μ

then

 EnterLong()

else if DBN(p_1, \dots, p_n).Predict[0] \leq Close[0] + μ

then

 EnterShort()

end if

end while

Our resulting algorithm depends on the following parameters p_1, \dots, p_n , the parameters of our dynamic graphical model, the profit target, the stop loss and the signal threshold μ . We could estimate the graphical model parameters with the EM procedure, then optimize the profit target, stop loss and signal threshold separately. However, if by bad luck, the dynamics of the graphical model is incorrect, the noise induced by wrong specification will be factored in these three parameters.

We opt rather for a combined optimization of all the parameters. We use one minute data for the S&P 500 index futures from 01Jan2017 to 01Jan2018 but take daily decisions. We train our model on the first 6 months representing about 170,000 data points and test it on the next

six months. For a model given by equations (11) and (12) and parameters specified in (13), the optimization encompasses 18 parameters: p_1, \dots, p_{15} , the profit target, the stop loss and the signal offset μ , making it non trivial. We use the CMA-ES algorithm to find the optimal solution. In our optimization, we add a $L1$ penalty to enforce sparsity of our parameters. Obtained parameters are summarized in Table 1. Out of 18 parameters, only 10 are not null. Table 2 shows that statistics for train and test are very similar, which is a good sign for little overfitting. Figures 6 and 7 are very similar confirming our intuition of reduced overfitting.




Figure 6: Graphical model algorithm on train data set




Figure 7: Graphical model algorithm on test data set

Table 1: Optimal non null parameters

Parameters	p_{01}	p_{03}	p_{04}	p_{05}	p_{06}	p_{07}	p_{12}	threshold	stop	target
Value	24.8	11.8	46.2	77.5	67	99.9	99.9	5	80	150

We compare our algorithm with a classical EM and a traditional Moving Average (MA) crossover approach to test the efficiency of our graphical model for trend detection. The MA algorithm generates a buy (resp. sell) signal when the fast moving average crosses above (resp. below) the long moving average. Table 3 details the results of this comparison. We see that on the train dataset, the three algorithms share similar performances (similar first two columns) but they perform quite differently on the test set and only our graphical model manages to have a Sharpe ratio over 1 on the test period. Overall, the MA algorithm seems to suffer from overfitting as its performance in the test set is very different from the one on the train. This is confirmed by Figure 8 and 9, with sharp contrast between train and test.




Figure 8: Moving Average Crossover algorithm on train data set




Figure 9: Moving Average Crossover algorithm on test data set

Table 2: Train/Test statistics for NGD graphical model

Performance	Net Profit	Avg. Trade	Tot. Net Profit (%)	Ann. Net Profit (%)
Train	5,086 €	339.05 €	5.09%	10.59%
Test	4,266 €	284.38 €	4.27%	8.69%
Performance	Vol	Sharpe Ratio	Max. Drawdown	Recovery Factor
Train	6.54%	1.62	-2,941 €	3.510
Test	6.20%	1.40	-1,721 €	4.948
Performance	Profit Factor	# of Trades	# of Contracts Avg. Winning Trade	Max. consec. Winners
Train	1.75 €	15	1,692.09 €	3
Test	1.62 €	15	1,588.92 €	2

Table 3: Graphical Model versus EM and MA crossover

Algo	Train: Percent Profitable	Train: Total Net Profit	Test: Total Net Profit	Test: Recovery Factor
MA Cross over	54 %	5,260.00 €	935 €	0.32
EM optimized model	42 %	5,165.72 €	2,250 €	0.56
NGD Graphical model	48 %	5,085.79 €	4,266 €	2.48
Algo	Test: Sharpe Ratio	Total # of Trades	Overall Profit Factor	Test: Max. Drawdown
MA Cross over	0.41	26	1.13	-€2,889
EM optimized model	0.85	28	1.25	-€2,523
NGD Graphical model	1.40	30	1.62	-€1,721

5 Conclusion

In this paper, we presented a novel approach for learning parameters in dynamic graphical models that tackles the targeted usage of our dynamic graphical model and incorporates them in the natural gradient optimization of the loss function. Compared to EM method, this method performs better in real scenarios and is less prone to overfitting.

References

[Akimoto, 2012] Akimoto, Y. (2012). Analysis of a natural gradient algorithm on monotonic convex-quadratic-composite functions. In *Proceedings of the*

- 14th Annual Conference on Genetic and Evolutionary Computation, GECCO '12, pages 1293–1300, New York, NY, USA. ACM.
- [Akimoto et al., 2012] Akimoto, Y., Auger, A., and Hansen, N. (2012). Convergence of the continuous time trajectories of isotropic evolution strategies on monotonic c_2 -composite functions. In Coello, C. A. C., Cutello, V., Deb, K., Forrest, S., Nicosia, G., and Pavone, M., editors, *Volume 7491 of Lecture Notes in Computer Science*, pages 2–51. Springer.
- [Akimoto et al., 2015] Akimoto, Y., Auger, A., and Hansen, N. (2015). Continuous optimization and CMA-ES. In *Genetic and Evolutionary Computation Conference, GECCO 2015, Madrid, Spain, July 11-15, 2015, Companion Material Proceedings*, pages 313–344.
- [Akimoto et al., 2016] Akimoto, Y., Auger, A., and Hansen, N. (2016). CMA-ES and advanced adaptation mechanisms. In *Genetic and Evolutionary Computation Conference, GECCO 2016, Denver, CO, USA, July 20-24, 2016, Companion Material Proceedings*, pages 533–562.
- [Akimoto et al., 2010] Akimoto, Y., Nagata, Y., Ono, I., and Kobayashi, S. (2010). Bidirectional relation between cma evolution strategies and natural evolution strategies. *Lecture Notes in Computer Science - Proceedings of Parallel Problem Solving from Nature - PPSN XI*, 6238:154–163.
- [Amari and Nagaoka, 2007] Amari, S. and Nagaoka, H. (2007). *Methods of Information Geometry*. American Mathematical Society.
- [Amari, 1995] Amari, S.-I. (1995). Information Geometry of the EM and em Algorithms for Neural Networks. *Neural Netw.*, 8(9):1379–1408.
- [Amari, 1998] Amari, S.-I. (1998). Natural gradient works efficiently in learning. *Neural Comput.*, 10(2):251–276.
- [Arnold and Hansen, 2010] Arnold, D. V. and Hansen, N. (2010). Active covariance matrix adaptation for the (1+1)-cma-es. In *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation, GECCO '10*, pages 385–392, New York, NY, USA. ACM.
- [Auger and Hansen, 2009] Auger, A. and Hansen, N. (2009). Benchmarking the (1+1)-CMA-ES on the BBOB-2009 noisy testbed. In *Genetic and Evolutionary Computation Conference, GECCO 2009, Proceedings, Montreal, Québec, Canada, July 8-12, 2009, Companion Material*, pages 2467–2472.
- [Auger and Hansen, 2012] Auger, A. and Hansen, N. (2012). Tutorial CMA-ES: evolution strategies and covariance matrix adaptation. In *Genetic and Evolutionary Computation Conference, GECCO '12, Philadelphia, PA, USA, July 7-11, 2012, Companion Material Proceedings*, pages 827–848.
- [Auger et al., 2004] Auger, A., Schoenauer, M., and Vanhaecke, N. (2004). LS-CMA-ES: A second-order algorithm for covariance matrix adaptation. In *Parallel Problem Solving from Nature - PPSN VIII, 8th International Conference, Birmingham, UK, September 18-22, 2004, Proceedings*, pages 182–191.
- [Bernacchia et al., 2018] Bernacchia, A., Lengyel, M., and Hennequin, G. (2018). Exact natural gradient in deep linear networks and its application to the nonlinear case. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems 31*, pages 5945–5954. Advances in Neural Information Processing Systems 31.
- [Beyer, 2014] Beyer, H.-G. (2014). Convergence analysis of evolutionary algorithms that are based on the paradigm of information geometry. *Evol. Comput.*, 22(4):679–709.
- [Cappe et al., 2010] Cappe, O., Moulines, E., and Ryden, T. (2010). *Inference in Hidden Markov Models*. Springer Publishing Company, Incorporated.
- [Dempster et al., 1977] Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B*, 39(1):1–38.
- [Einicke et al., 2010] Einicke, G., Falco, G., and J.T. M. (2010). Em algorithm state matrix estimation for navigation. *IEEE SIGNAL PROCESSING LETTERS*, 17:437–440.
- [Fung, 2017] Fung, S. P. Y. (2017). Optimal online two-way trading with bounded number of transactions. *CoRR*.
- [Giuseppe Di Graziano, 2014] Giuseppe Di Graziano, D. B. A. (2014). Optimal trading stops and algorithmic trading. *SSRN*.
- [Glasmachers et al., 2010] Glasmachers, T., Schaul, T., Sun, Y., Wierstra, D., and Schmidhuber, J. (2010). Exponential natural evolution strategies. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation GECCO-2010*, pages 393–400. ACM.
- [Hansen, 2016] Hansen, N. (2016). The CMA evolution strategy: A tutorial. *CoRR*, abs/1604.00772.
- [Hansen and Auger, 2011] Hansen, N. and Auger, A. (2011). CMA-ES: evolution strategies and covariance matrix adaptation. In *13th Annual Genetic and Evolutionary Computation Conference, GECCO 2011, Companion Material Proceedings, Dublin, Ireland, July 12-16, 2011*, pages 991–1010.
- [Hansen and Auger, 2014] Hansen, N. and Auger, A. (2014). Evolution strategies and CMA-ES (covariance matrix adaptation). In *Genetic and Evolutionary Computation Conference, GECCO '14, Vancouver, BC, Canada, July 12-16, 2014, Companion Material Proceedings*, pages 513–534.
- [Hansen and Ostermeier, 2001] Hansen, N. and Ostermeier, A. (2001). Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195.
- [Hartley, 1958] Hartley, H. (1958). Maximum likelihood estimation from incomplete data. *Biometrics*, 14:174–194.
- [Igel et al., 2007] Igel, C., Hansen, N., and Roth, S. (2007). Covariance matrix adaptation for multi-objective optimization. *Evol. Comput.*, 15(1):1–28.
- [Jeffreys, 1946] Jeffreys, H. (1946). n invariant form for the prior probability in estimation problems. *A. Proc. Roy.Soc. London. Ser. A.*, 186:453–461.
- [Jordan, 2012] Jordan, M. (2012). *Learning in Graphical Models*. Nato Science Series D:. Springer Netherlands.
- [Krause et al., 2016] Krause, O., Arbonès, D. R., and Igel, C. (2016). Cmaes with optimal covariance update and storage complexity. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems 29*, pages 370–378. Curran Associates, Inc.
- [Labadie and Lehalle, 2010] Labadie, M. and Lehalle, C.-A. (2010). Optimal algorithmic trading and market microstructure. Working papers, HAL.
- [Levine, 2018] Levine, S. (2018). Reinforcement learning and control as probabilistic inference: Tutorial and review. *CoRR*, abs/1805.00909.
- [Li and McCormick, 2017] Li, Z. and McCormick, T. H. (2017). An Expectation Conditional Maximization approach for Gaussian graphical models. *arXiv e-prints*.
- [Mirsky, 1975] Mirsky, L. (1975). A trace inequality of john von neumann. *Monatshfte für Mathematik*, 4:303–306.
- [Neal and Hinton, 1999] Neal, R. M. and Hinton, G. E. (1999). Learning in graphical models. In Jordan, M. I., editor, *Learning in Graphical Models*, chapter A View of the EM Algorithm That Justifies Incremental, Sparse, and Other Variants, pages 355–368. MIT Press, Cambridge, MA, USA.
- [Ollivier et al., 2017] 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.
- [Pascanu and Bengio, 2014] Pascanu, R. and Bengio, Y. (2014). Revisiting natural gradient for deep networks. In *International Conference on Learning Representations*.
- [Pfeifer and Protzel, 2018] Pfeifer, T. and Protzel, P. (2018). Expectation-maximization for adaptive mixture models in graph optimization. *CoRR*, abs/1811.04748.
- [Rao, 1945] Rao, C. R. (1945). Information and the accuracy attainable in the estimation of statistical parameters. *Bull. Calcutta Math. Soc.*, 37:81–91.
- [Robin et al., 2017] Robin, G., Ambroise, C., and Robin, S. (2017). Incomplete graphical model inference via latent tree aggregation. *arXiv e-prints*.
- [Sharpe, 1966] Sharpe, W. F. (1966). Mutual fund performance. *Journal of Business*, pages 119–138.

- [Varelas et al., 2018] 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. In *Parallel Problem Solving from Nature - PPSN XV - 15th International Conference, Coimbra, Portugal, September 8-12, 2018, Proceedings, Part I*, pages 3–15.
- [Vezeris et al., 2018] Vezeris, D., Kyrgos, T., Schinas, C. T. P., and Loss, S. (2018). Trading strategies comparison in combination with an macd trading system. *J. Risk Financial Manag*, 11:56.
- [Xiao et al., 2017] Xiao, Q., Chaoqin, C., and Li, Z. (2017). Time series prediction using dynamic bayesian network. *Optik - International Journal for Light and Electron Optics*, 135.
- [Xu and Jordan, 1996] Xu, L. and Jordan, M. I. (1996). On convergence properties of the em algorithm for gaussian mixtures. *Neural Computation*, 8(1):129–151.

Supplementary Materials

A Proof of proposition 3.1

Proof. The regularity condition that the expectation of our loss function is square integrable ($\mathbb{E}_{X \sim P_\theta} [L_f^2(X)] < \infty$) implies that for any $d \geq 1$, we also have $\mathbb{E}_{X \sim P_\theta} [L_f^{2/d}(X)] < \infty$.

Hence we can interchange the order of integration and differentiation and compute the natural gradient update as

$$\nabla U_f(\theta) = \mathbb{E}_{X \sim P_\theta} [L_f^{2/d}(X) S(\theta, x)], \quad (14)$$

where $S(\theta, x) = \nabla \ln l(\theta, x)$ is the score function, defined as the gradient of the log-likelihood $L(\theta, x) = \ln l(\theta, x) = \ln p_\theta(x)$ of the multivariate Gaussian distribution with respect to its moment parameters.

For a multi variate normal distribution, the score function is given by

$$S(\theta, x) = \begin{bmatrix} \Sigma^{-1}(x - \mu) \\ -\frac{1}{2} \text{vec}(\Sigma^{-1} - \Sigma^{-1}(x - \mu)(x - \mu)^T \Sigma^{-1}) \end{bmatrix}$$

For a multi variate normal distribution, the Fisher matrix information has a closed form given by

$$G_\theta = \begin{bmatrix} \Sigma^{-1} & 0 \\ 0 & \frac{1}{2}(\Sigma^{-1} \otimes \Sigma^{-1}) \end{bmatrix}, \quad (15)$$

where \otimes denotes the Kronecker product operator. The natural gradient writes

$$\theta^{t+1} = \theta^t - \nu_t G_{\theta^t}^{-1} \nabla U_f(\theta^t) \quad (16)$$

Combining (14), (15) and (16), and splitting the learning rate ν_t among the mean vector and the covariance matrix, with two different learning rates, we obtained

$$\begin{aligned} \mu^{t+1} &= \mu^t - \nu_\mu^t \delta m^t, \\ \Sigma^{t+1} &= \Sigma^t - \nu_\Sigma^t \delta \Sigma^t. \end{aligned} \quad (17)$$

with

$$\begin{aligned} \delta \mu^t &= \mathbb{E}_{X \sim P_{\theta^t}} [L_f^{2/d}(X)(X - \mu^t)] \\ \delta \Sigma^t &= \mathbb{E}_{X \sim P_{\theta^t}} [L_f^{2/d}(X)((X - \mu^t)(X - \mu^t)^T - \Sigma^t)] \end{aligned}$$

which concludes the proof. \square

B Proof of proposition 3.2

Proof. The symmetry is obvious. Let us prove that the matrix Σ_t is positive definite for each t by mathematical induction. The result is trivially true for $t = 0$ as Σ_0 is positive definite. Suppose that at step t the result is true and that the matrix C^t is positive definite and symmetric. Notice that the update 2 writes

$$\Sigma^{t+1} = \sqrt{\Sigma^t} (I_d - \nu_\Sigma^t \sqrt{\Sigma^t}^{-1} \delta \Sigma^t \sqrt{\Sigma^t}^{-1}) \sqrt{\Sigma^t} \quad (18)$$

for the closed form (C) NGD and

$$\Sigma^{t+1} = \sqrt{\Sigma^t} ((I_d - \nu_\Sigma \sqrt{\Sigma^t}^{-1} \widehat{\delta \Sigma^t} \sqrt{\Sigma^t}^{-1}) \widehat{\Sigma^t}) \sqrt{\Sigma^t} \quad (19)$$

for the Monte Carlo case. Using the result that if A and B are symmetric positive definite, then $\sqrt{AB}\sqrt{A}$ is also positive definite, we get that Σ^{t+1} stays positive definite. This concludes the proof that the matrix Σ_t is positive definite for each t . Should the learning rate be too large so that the inner term in 18 $I_d - \nu_\Sigma^t \sqrt{\Sigma^t}^{-1} \delta \Sigma^t \sqrt{\Sigma^t}^{-1}$ for the closed form case, and $I_d - \nu_\Sigma \sqrt{\Sigma^t}^{-1} \widehat{\delta \Sigma^t} \sqrt{\Sigma^t}^{-1}$ for the Monte Carlo case are not definite positive, the next iteration Σ^{t+1} would also be not definitive positive. Should the previous matrices be non positive, the next iteration Σ^{t+1} would non positive, which concludes the proof. \square

C Proof of proposition 3.3

We will decompose this proof into various elementary lemmas. Usually, when proving convergence for Monte Carlo, one relies on traditional statistical tools like the strong law of large numbers as well as some measure tools like Lebesgue's dominated convergence, Cauchy Schwartz inequality, etc... However, we take an original approach and rather use connection between single term convergence and Cesaro convergence. A first lemma that is intuitive is that almost sure convergence implies almost sure (a.s.) Cesaro convergence. To make things more precise, let X be a random variable of \mathbb{R}^d and $f_n : \mathbb{R}^d \rightarrow \mathbb{R}$ a series of real valued functions.

Lemma C.1. *If $(f_n)_{n \in \mathbb{N}}$ converges almost surely to 0, then its Cesaro sum obtained for n independent realizations of X converges almost surely to 0. In other words,*

$$f_n(X) \xrightarrow{a.s.} 0 \implies \frac{1}{n} \sum_{i=1}^n f_n(X_i) \xrightarrow{a.s.} 0$$

where $X_i \sim X$ are i.i.d. variables.

Proof. For $m \leq n$, we have

$$\begin{aligned} \lim_{n \rightarrow \infty} \left| \frac{1}{n} \sum_{i=1}^n f_n(X_i) \right| &\leq \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \sup_{j \geq n} |f_j(X_i)| \\ &\leq \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \sup_{j \geq m} |f_j(X_i)|. \end{aligned} \quad (20)$$

By assumption, $f_n(x) \xrightarrow{\text{a.s.}} 0$ almost everywhere in x as n tends to infinity. This implies that $\sup_{j \geq m} |f_j(x)| \xrightarrow{\text{a.s.}} 0$ almost everywhere in x as μ tends to infinity. We can apply the Lebesgue's dominated convergence theorem to prove that $\mathbb{E}[\sup_{j \geq m} |f_j(X)|] \rightarrow 0$ as $m \rightarrow \infty$. This implies that $\mathbb{E}[\sup_{j \geq m} |f_j(X)|]$ is finite for μ sufficiently large. The strong law of large numbers (LLN) applies and proves that the right hand side of (20) converges to $\mathbb{E}[\sup_{j \geq m} |f_j(X)|]$ when n becomes large. The last expectation converges to 0 when μ grows to infinity. This concludes the proof. \square

Remark C.1. *That convergence implies Cesaro convergence relies strongly on the almost sure convergence nature. Should the initial convergence only be in probability, the convergence of our series of functions $(f_n)_{n \in \mathbb{N}}$ would not imply Cesaro convergence. Take for instance a standard Gaussian variable $X \sim \mathcal{N}(0, 1)$ and the series of function $(f_n)_{n \in \mathbb{N}}$ where f_n is equal to 0 everywhere except when $X \geq n$ and takes the value $1/(1 - \Phi(n))$ where Φ is the cumulative density function (c.d.f). Take X_1, X_2, \dots, X_n independent standard Gaussian variables, it is easy to see that the Cesaro sum $\frac{\sum_{k=1}^n f_n(X_k)}{n}$ does not converge to 0 as its mean is by construction equal to 1, while (f_n) converges in probability to 0. This is because the function f_n takes non zero value on a smaller and smaller interval measure: $X \geq n$ but with a value that explodes and is given by $1/(1 - \Phi(n))$ to compensate for the reduced interval on which the function f_n is not null.*

The second lemma concerns the almost sure consistency of our estimator of the loss function.

Lemma C.2. $\widehat{L}_f^{n, 2/d}(x)$ converges a.s. to $L_f^{2/d}(x)$

Proof. The estimator $\widehat{L}_f^{n, 2/d}(x)$ writes as

$$\widehat{L}_f^{n, 2/d}(x) = \left(\frac{1}{n} \sum_{j=1}^n \frac{\mathbf{1}_{\{f(x_j) \leq f(x)\}}}{p_{\theta^t}(x_j)} \right)^{2/d}.$$

hence, we have

$$\lim_{n \rightarrow \infty} \widehat{L}_f^{n, 2/d}(x) = \left(\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{j=1}^n \frac{\mathbf{1}_{f(X_j) \leq f(x)}}{p_{\theta}(X_j)} \right)^{2/d}.$$

We know that $\mathbb{E}[L_f^{2/d}(X)] < \infty$ which implies $\mathbb{E}[L_f^{2/d}(X)] < \infty$ for any dimension d . This implies in particular that $L_f^{2/d}(X) = (\mu_{\text{Leb}}[y : f(y) \leq f(x)])^{2/d}$ is finite almost everywhere. We can safely apply the strong law of large numbers to get

$$\lim_{n \rightarrow \infty} \widehat{L}_f^{n, 2/d}(x) = (\mu_{\text{Leb}}[y : f(y) \leq f(x)])^{2/d} = L_f^{2/d}(x)$$

almost surely and almost everywhere in x , which concludes the proof. \square

We are now ready to prove proposition 3.3.

Proof. The assumption in proposition 3.3 is that $L_f(X)$ is square integrable, that is $\mathbb{E}[L_f^2(X)] < \infty$, hence $L_f^{2/d}(X)$ is also integrable, which leads trivially that the expectation of $L_f^{2/d}(X)$ is finite

$$\mathbb{E}[L_f^{2/d}(X)] \leq \infty. \quad (21)$$

The variance of $L_f^{2/d}(X)$ being non negative, we have the trivial inequality between the first two moments:

$$\left(\mathbb{E}[L_f^{2/d}(X)] \right)^2 \leq \mathbb{E}[L_f^{4/d}(X)].$$

Moreover, Cauchy-Schwarz inequality gives us that

$$\mathbb{E}[\|L_f^{2/d}(X) \tilde{\nabla} L(\theta, x)\|^2] < \mathbb{E}[L_f^{4/d}(X)] \mathbb{E}[\|\tilde{\nabla} L(\theta, x)\|^2].$$

Using the relationship between the variance of the score and the information Fisher, we can notice that $\mathbb{E}[\|\tilde{\nabla} L(\theta, x)\|^2] = \text{Tr}(G_{\theta}^{-1}) < \infty$. This implies also that the expectation of $\|L_f^{2/d}(X) \tilde{\nabla} L(\theta, x)\|$ is finite

$$\mathbb{E}[\|L_f^{2/d}(X) \tilde{\nabla} L(\theta, x)\|] < \infty \quad (22)$$

Let us take the function series (f_n) as the difference between our estimated loss function and the loss function:

$$f_n(x) = \widehat{L}_f^{n, 2/d}(x) - L_f^{2/d}(X).$$

Let us decompose $\tilde{\nabla}_{\theta^n} \widehat{U}$ in three terms as follows:

$$\begin{aligned} \tilde{\nabla}_{\theta^n} \widehat{U} &= \frac{1}{n} \sum_{i=1}^n L_f^{2/d}(X_i) \tilde{\nabla} L(\theta; X_i) \\ &\quad - \underbrace{\left(\frac{1}{n} \sum_{i=1}^n \widehat{L}_f^{n, 2/d}(X_i) \right)}_{=b} \left(\frac{1}{n} \sum_{i=1}^n \tilde{\nabla} L(\theta; X_i) \right) \\ &\quad + \frac{1}{n} \sum_{i=1}^n f_n(X_i) \tilde{\nabla} L(\theta; X_i) \end{aligned} \quad (23)$$

We want to prove that $\tilde{\nabla}_{\theta^n} \widehat{U} \xrightarrow{\text{a.s.}} \nabla U_f(\theta)$ almost surely. We can examine successively the three terms.

Using (22) and the strong law of large numbers, the first term converges to $\mathbb{E}[L_f^{2/d}(X)\tilde{\nabla}L(\theta, x)] = \nabla U_f(\theta)$ almost surely as n tends to infinity.

We now examine the second and third terms in (23). To conclude the proof, we need to prove that these two terms converge a.s. to zero.

Let us tackle the second term. By lemma C.2, we know that our estimated loss function is almost surely consistent, which implies that $f_n(x)$ converges almost surely to 0. By lemma C.2, this implies that we have the almost surely Cesaro convergence of

$$\frac{1}{n} \sum_{i=1}^n f_n(X_i) \xrightarrow{\text{a.s.}} 0 \quad \text{as } n \rightarrow \infty. \quad (24)$$

The almost sure convergence (24) implies that the limit $\lim_{n \rightarrow \infty} \sum_{i=1}^n \widehat{L}_f^{2/d}(X_i)/n$ agrees with $\lim_{n \rightarrow \infty} \sum_{i=1}^n L_f^{2/d}(X_i)/n$. We also have from (21) and from the strong law of large numbers that the Cesaro sum $\sum_{i=1}^n \widehat{L}_f^{2/d}(X_i)/n$ converges almost surely to the expectation $\mathbb{E}[L_f^{2/d}(X)]$. The latter is finite by assumption. Also, by the strong law of large numbers we have that the Cesaro sum $\sum_{i=1}^n \tilde{\nabla}L(\theta; X_i)/n$ also converges to 0 when n tends to infinity. We can conclude that the second term of (23) converges to zero almost surely.

We are now left with the third term of (23). The proof of its almost sure convergence to zero is similar. A Cauchy-Schwarz inequality applied to the second term of (23) gives us a way to control the third term as follows:

$$\left| \sum_{i=1}^n \frac{f_n(X_i) \tilde{\nabla}L(\theta; X_i)}{n} \right|^2 \leq \sum_{i=1}^n \frac{f_n(X_i)^2}{n} \sum_{i=1}^n \frac{\|\tilde{\nabla}L(\theta; X_i)\|^2}{n}.$$

We can take of the two terms in this inequality successively. The strong law of large numbers proved that the second term of the right hand side converges to the expectation $\mathbb{E}[\|\tilde{\nabla}L(\theta, x)\|^2]$, which can be expressed in terms of the trace of the inverse of the Fisher information matrix: $\text{Tr}(G_\theta^{-1})$ which is finite. This means that the second term of the right hand side converges to a constant. We are left to prove that the first term on the right hand side converges to zero almost surely. By lemma C.2, we have that $f_n \xrightarrow{\text{a.s.}} 0$ almost everywhere in x as $n \rightarrow \infty$, which implies that $f_n^2 \xrightarrow{\text{a.s.}} 0$ almost everywhere in x as $n \rightarrow \infty$. By lemma C.1, this implies the Cesaro convergence of $\sum_{i=1}^n \frac{f_n(X_i)^2}{n}$ to 0 almost everywhere in x as $n \rightarrow \infty$ which concludes the proof. \square

D Proof of proposition 3.4

Like for proposition 3.3, we will first prove a few lemma and the result will become easy. Since our objective function is explicit and given by a parabola, which is a very simple function, we can explicitly compute the natural gradient ascent. This is the subject of the first lemma:

Lemma D.1. *The closed-form NGD algorithm on the benchmark quadratic function $\frac{1}{2}x^T H x$ writes as*

$$\mu^{t+1} = \mu^t - \nu_\mu^t \delta \mu^t, \quad (25)$$

$$\text{with } \delta \mu^t = \kappa \Sigma^t H \mu^t \quad (26)$$

$$\text{and } \Sigma^{t+1} = \Sigma^t - \nu_\Sigma^t \delta \Sigma^t, \quad (27)$$

$$\text{with } \delta \Sigma^t = \kappa \Sigma^t H \Sigma^t, \quad (28)$$

with the constant $\kappa > 0$ given by

$$\kappa = \frac{2\pi}{\left(\Gamma\left(\frac{d}{2} + 1\right) \sqrt{\det(H)}\right)^{2/d}}$$

where $\Gamma(z)$ the Gamma function.

Proof. The natural gradient is directly related to the inverse Fisher information matrix times the gradient of the loss function:

$$\mu^{t+1} = \mu^t - \nu_\mu^t (G_\theta^{-1} \nabla U_f(\theta))_\mu \quad (29)$$

$$\Sigma^{t+1} = \Sigma^t - \nu_\Sigma^t (G_\theta^{-1} \nabla U_f(\theta))_\Sigma \quad (30)$$

To fully make explicit the above formula, we need to compute explicitly $G_\theta^{-1} \nabla U_f(\theta)$ and prove that it is given by

$$G_\theta^{-1} \nabla U_f(\theta) = \kappa \begin{bmatrix} \Sigma H \mu \\ \text{vec}(\Sigma H \Sigma) \end{bmatrix}. \quad (31)$$

We can remark that $\mu_{\text{Leb}}[y : f(y) \leq f(x)]$ is equivalent to the volume of the ellipsoid given $\{y : y^T H y \leq x^T H x\}$, or equivalently, $\{y : \frac{\|\sqrt{H}y\|}{\|\sqrt{H}x\|} \leq 1\}$. The last set is the image of unit ball $\{v : \|v\| \leq 1\}$ under the transformation given by $\frac{\sqrt{H}y}{\|\sqrt{H}x\|} = v$ or equivalently $y = \|\sqrt{H}x\| \sqrt{H}^{-1} v$. The change of volume under this linear map is given by the determinant of the transform which is $\|\sqrt{H}x\|^d \det(H)^{-1/2}$, hence,

$$\mu_{\text{Leb}}[y : f(y) \leq f(x)] = \|\sqrt{H}x\|^d \det(H)^{-1/2} V_d(1),$$

where $V_d(1)$ is the volume of the unit ball in dimension d . The utility or quasi objective function is defined as the

expectation of the loss function to the power $2/d$. Hence, we get

$$\begin{aligned} U_f(\theta) &:= \mathbb{E}_{X \sim P_\theta} \left[L_f^{2/d}(X) \right] \\ &= \mathbb{E}_{X \sim P_\theta} \left[(\mu_{\text{Leb}}[y : f(y) \leq f(x)])^{2/d} \right] \\ &= \left(\frac{V_d(1)}{\sqrt{\det(H)}} \right)^{2/d} \mathbb{E}_{X \sim P_\theta} [X^T H X] \\ &= \left(\frac{V_d(1)}{\sqrt{\det(H)}} \right)^{2/d} (\mu^T H \mu + \text{Tr}(H \Sigma)). \end{aligned}$$

Differentiating the both side of the above relation, we have

$$\nabla U_f(\theta) = 2 \left(\frac{V_d(1)}{\sqrt{\det(H)}} \right)^{2/d} \left[\frac{1}{2} \text{vec}(H) \right].$$

We can use the fact that the volume of the unit ball is given by

$$V_d(1) = \frac{\pi^{d/2}}{\Gamma(\frac{d}{2} + 1)},$$

to get

$$\nabla U_f(\theta) = \kappa \left[\frac{1}{2} \text{vec}(H) \right]. \quad (32)$$

with

$$\kappa = \frac{2\pi}{\left(\Gamma(\frac{d}{2} + 1) \sqrt{\det(H)} \right)^{2/d}}$$

For a multi variate Gaussian distribution, the inverse of the Fisher information matrix is given by

$$G_\theta = \begin{bmatrix} \Sigma & 0 \\ 0 & 2(\Sigma \otimes \Sigma) \end{bmatrix}, \quad (33)$$

where \otimes denotes the Kronecker product operator. Multiplying the gradient of our utility function (32) by the inverse of the Fisher information matrix (33), we get equation (31), which concludes the proof. \square

The second lemma relates to the condition number of a diagonal matrix series $(\Lambda_t)_{t \in \mathbb{N}}$ that satisfies a particular recursive relationship given by

$$\Lambda_{t+1} = \Lambda_t - \nu_t (\Lambda_t)^2. \quad (34)$$

We also take the convention that λ_i^t is the i th largest value of the matrix Λ_t . We recall that for a symmetric matrix M , its condition number denoted by $\text{Cond}(M)$ is defined by the ratio of its maximal to its minimal eigen value. For the specific case of a diagonal matrix, this is simply the ratio of its largest to its smallest value.

Lemma D.2. *Let a diagonal matrix series $(\Lambda_t)_{t \in \mathbb{N}}$ that satisfies (34), with the constraint on the learning rate at each step given by:*

$$0 < \nu_{\min} \leq \nu_t \lambda_1(\Lambda_t) \leq 1/2. \quad (35)$$

Then, the series of condition number of the matrix Λ_t satisfies for any $t > 0$ the following inequalities:

$$\frac{\text{Cond}(\Lambda_{t+1}) - 1}{\text{Cond}(\Lambda_t) - 1} \leq 1 - \nu_{\min}. \quad (36)$$

Proof. We can remark that $\nu_t(\lambda_i^t + \lambda_j^t) \leq 1$ for any i, j . This is trivially obtained by the fact that as by definition, λ_1 is the largest eigen value, we have $\lambda_i^t \leq \lambda_1$ for any i , which combined with the constraint (34) implies $\nu_t(\lambda_i^t + \lambda_j^t) \leq 2\nu_t \lambda_1 \leq 1$.

Let suppose without loss of generality that we have sorted the diagonal matrix elements by their value so that the first element of the diagonal matrix value is the largest eigen values, the second element the second largest value, etc.. Let us now prove that the mapping between the matrix position and the eigen values ranking order stays the same between step t and $t + 1$.

Let us prove the result by mathematical induction. The result is true for step $t = 0$. Suppose the result is true at step t and $\lambda_i^t \geq \lambda_j^t$. From (34) and the inequality $\nu_t(\lambda_i^t + \lambda_j^t) \leq 1$, we have

$$\begin{aligned} \lambda_i^{t+1} - \lambda_j^{t+1} &= \lambda_i^t(1 - \nu_t \lambda_i^t) - \lambda_j^t(1 - \nu_t \lambda_j^t) \\ &= (1 - \underbrace{\nu_t(\lambda_i^t + \lambda_j^t)}_{\leq 1}) \underbrace{(\lambda_i^t - \lambda_j^t)}_{\geq 0} \geq 0 \end{aligned}$$

with equality holding if and only if $\lambda_i^t = \lambda_j^t$. Therefore, if $\lambda_i^t > \lambda_j^t$, then $\lambda_i^{t+1} > \lambda_j^{t+1}$, which implies the result holds also at time $t + 1$.

As by convention there is mapping between diagonal element position and their eigen value rank, the condition number of matrix $\text{Cond}(\Lambda_t)$ is computed at each step as the ratio of the first and last element of the diagonal matrix Λ_t given by λ_1^t/λ_d^t . According to our recursive equation (34) we have

$$\begin{aligned} \frac{\lambda_1^{t+1} - \lambda_d^{t+1}}{\lambda_d^{t+1}} &= \frac{\lambda_1^t(1 - \nu_t \lambda_1^t) - \lambda_d^t(1 - \nu_t \lambda_d^t)}{\lambda_d^t(1 - \nu_t \lambda_d^t)} \\ &= \underbrace{\frac{\lambda_1^t - \lambda_d^t}{\lambda_d^t}}_{\text{Cond}(\Lambda_t) - 1} \frac{1 - \nu_t(\lambda_1^t + \lambda_d^t)}{(1 - \nu_t \lambda_d^t)}. \quad (37) \end{aligned}$$

We can notice that the second term in the right hand side is bounded as follows:

$$\frac{1 - \nu_t(\lambda_1^t + \lambda_d^t)}{1 - \nu_t \lambda_d^t} = 1 - \frac{\nu_t \lambda_1^t}{1 - \nu_t \lambda_1^t \text{Cond}^{-1}(\Lambda_t)} \leq 1 - \nu_t \lambda_1^t.$$

Because of the lower bound on the learning rate in (35), the right term in the above inequality is bounded by $1 - \nu_{\Sigma, \min}$. This leads to the result:

$$\frac{\text{Cond}(\Lambda^{t+1}) - 1}{\text{Cond}(\Lambda^t) - 1} \leq 1 - \nu_{\Sigma, \min}. \quad (38)$$

This concludes the proof. \square

Remark D.1. As the condition number defined as the ratio of the maximum with the minimum eigen value converges to 1, this implies that the diagonal matrix Λ^t converges to the identity matrix in the sense of the Frobenius norm times its maximal eigen value (or equivalent its minimal eigen value as they get similar as t tends to infinity). This is not by magic. The proportional convergence to the identity matrix is forced by the inequalities (35), satisfied by the learning rate at each step which forces all the matrix eigen values to progressively converge to a common number. We can also note that the largest eigen value converges to zero as t tends to infinity and is decreasing. This means in particular that in order to satisfy the condition (35), the learning rate ν_t needs to be not too small and specifically not below $\nu_{\min}/\lambda_1(\Lambda_t)$. This implies that the learning rate makes a trade-off between a small value to ensure that it is below $\frac{1}{2\lambda_1(\Lambda_t)}$ but above $\nu_{\min}/\lambda_1(\Lambda_t)$, with the latter becoming larger and larger as $\lambda_1(\Lambda_t)$ becomes smaller and smaller.

A third lemma concerns the relationship between largest singular values and matrix norm. We denote by $\sigma_i(M)$ the i th largest singular value of a matrix M , and $\sigma_1(M)$ its largest singular value.

Lemma D.3. Let $M \in \mathbb{R}^{d \times d}$ be a positive definite matrix and $S \in \mathbb{R}^{d \times d}$ a positive definite symmetric matrix, then we have the following matrix norm inequality:

$$\|MS\|^2 \leq \sigma_1^2(M)\|S\|^2. \quad (39)$$

Proof. Using trace commutation property and relationship between the norm and the trace, we have

$$\|MS\|^2 = \text{Tr}(SM^TMS) = \text{Tr}(M^TMS^2) \quad (40)$$

Additionally, the J. von Neumann's trace inequality [Mirsky, 1975] gives us

$$|\text{Tr}(M_1M_2)| \leq \sum_{i=1}^d \sigma_i(M_1)\sigma_i(M_2) \leq \sigma_1(M_1) \sum_{i=1}^d \sigma_i(M_2), \quad (41)$$

where M_1 and M_2 are any matrices in $\mathbb{R}^{d \times d}$. Combining (40) and (41), we get

$$\|MS\|^2 \leq \sigma_1(M^T M) \sum_{i=1}^d \sigma_i(S^2) = \sigma_1^2(M)\|S\|^2. \quad \square$$

Remark D.2. Indeed, this lemma is quite obvious when looking at the variational definition of singular matrix, that is for a matrix M , the first or largest singular value is written as the solution of the following maximization program

$$\sigma_{\max}(M) = \sup_{x, y \neq 0} \frac{x^T M y}{\|x\|_2 \|y\|_2} = \sup_{y \neq 0} \frac{\|My\|_2}{\|y\|_2}$$

Said differently, it also says that the maximum singular value is the L_2 operator norm of the matrix M .

A powerful result that was initially proved in [Akimoto, 2012] is that the covariance matrix of the NGD algorithm converges proportionally to the inverse of the Hessian matrix of our quadratic problem which is given by H . To prove this convergence, we shows that the condition number of $\Sigma^t H$ converges to 1 with a linear speed. This is the subject of the following lemma:

Lemma D.4. If the learning rate ν_{Σ}^t for the covariance matrix satisfies inequalities (8), then the condition number of $\Sigma^t H$ converges to one with a linear rate of convergence given by

$$\limsup_{t \rightarrow \infty} \frac{\text{Cond}(\Sigma^{t+1} H) - 1}{\text{Cond}(\Sigma^t H) - 1} \leq \frac{1 - 2\nu_{\Sigma, \min}}{1 - \nu_{\Sigma, \min}}. \quad (42)$$

Remarkably, we also have that the condition number is upper bounded by

$$\text{Cond}(\Sigma^t H) \leq 1 + (1 - \nu_{\Sigma, \min})^t (\text{Cond}(\Sigma^0 H) - 1). \quad (43)$$

Finally, if the limit $\nu_{\Sigma, \lim} = \lim_{t \rightarrow \infty} \nu_{\Sigma}^t \lambda_1((\Sigma^t)^{-1} \delta \Sigma^t)$ exists, in equation (42), $\nu_{\Sigma, \min}$ can be replaced by $\nu_{\Sigma, \lim}$

Proof. We are interested in the condition number of the matrix $\Sigma^t H$. To make the form more symmetric, we can notice that it is better to look at the symmetrized term $\sqrt{H} \Sigma^t \sqrt{H}$. This term does exist because H is a positive definite and symmetric matrix, hence it admits a square root \sqrt{H} . As $\sqrt{H} \Sigma^t \sqrt{H}$ is a positive definite and symmetric matrix, we can decompose it in an orthogonal matrix O^t and a diagonal matrix D^t as follows: $O^t D^t (O^t)^T$, where the diagonal elements of $D^t = \text{diag}(D_1^t, \dots, D_d^t)$ are the eigenvalues of $\sqrt{H} \Sigma^t \sqrt{H}$ and each column of O^t is the eigenvector corresponding to each diagonal element of D^t . The orthogonal and diagonal matrix decomposition may not be unique but we will first prove that if we find an initial orthogonal matrix O^0 that diagonalizes $\sqrt{H} \Sigma^0 \sqrt{H}$, it will also diagonalize the matrix $\sqrt{H} \Sigma^t \sqrt{H}$ for any $t \geq 0$.

Let us prove this result by mathematical induction The result is true for step $t = 0$ as $\sqrt{H} \Sigma^0 \sqrt{H}$ is a positive definite and symmetric matrix, hence it admits an orthogonal

and diagonal matrix decomposition $O^0 D^0 (O^0)^T$, such that the diagonal elements of $D^0 = \text{diag}(D_1^0, \dots, D_d^0)$ are the eigenvalues of $\sqrt{H}\Sigma^0\sqrt{H}$ and each column of O^0 is the eigenvector corresponding to each diagonal element of D^t . Let us assume that at the result holds at step t , that is $\sqrt{H}\Sigma^t\sqrt{H}$ admits an orthogonal and diagonal matrix decomposition such that

$$\sqrt{H}\Sigma^t\sqrt{H} = O^t D^t (O^t)^T$$

Using first lemma D.1, we can use the covariance matrix update (27). If we multiply both the right and the left side of the covariance matrix update (27) by \sqrt{H} , we get:

$$\sqrt{H}\Sigma^{t+1}\sqrt{H} = \sqrt{H}\Sigma^t\sqrt{H} - \nu_\Sigma^t \kappa(\sqrt{H}\Sigma^t\sqrt{H})^2.$$

Reformulating the above equation in terms of the orthogonal and diagonal matrix decomposition, we get:

$$\sqrt{H}\Sigma^{t+1}\sqrt{H} = O^t (D^t - \nu_\Sigma^t \kappa(D^t)^2) (O^t)^T.$$

As $D^t - \nu_\Sigma^t \kappa(D^t)^2$ is trivially a diagonal matrix, the above equation proves that the orthogonal O^t also diagonalizes $\sqrt{H}\Sigma^{t+1}\sqrt{H}$. This proves that the result holds for step $t + 1$. In addition, we get as a by product that the next step diagonal matrix D^{t+1} is related to the previous step diagonal matrix D^t as follows:

$$D^{t+1} = D^t - \nu_\Sigma^t \kappa(D^t)^2. \quad (44)$$

which is very interesting and is almost the update equation of our lemma D.2. To get exactly the update equation, the trick is to introduce the diagonal matrix $\Lambda_t = cD^t$ as the constant $c > 0$ and remark that equation (44) is transformed into:

$$\Lambda^{t+1} = \Lambda^t - \nu_\Sigma^t (\Lambda^t)^2. \quad (45)$$

which is exactly the update of the lemma D.2. At this stage, we can notice various interesting remarks:

- the condition number of matrix $\Sigma^t H$ is also equal to the one of the matrix $\sqrt{H}\Sigma^t\sqrt{H}$ itself equal to the condition number of D^t itself equal to the condition number of Λ^t
- the diagonal matrix series $(\Lambda_t)_{t \in \mathbb{N}}$ satisfies the assumption of lemma D.2, hence we get

$$\frac{\text{Cond}(\Lambda_{t+1}) - 1}{\text{Cond}(\Lambda_t) - 1} \leq 1 - \nu_{\min} \quad (46)$$

or equivalently

$$\frac{\text{Cond}(\Sigma^{t+1}H) - 1}{\text{Cond}(\Sigma^t H) - 1} \leq 1 - \nu_{\min} \quad (47)$$

These remarks imply that the following upper bound for the condition number of $\Sigma^t H$

$$\text{Cond}(\Sigma^t H) \leq 1 + (1 - \nu_{\Sigma, \min})^t (\text{Cond}(\Sigma^0 H) - 1)$$

which exactly the equation (43). Moreover a trivial usage of squeeze (also known as the pinching or sandwich) theorem proves that

$$\lim_{t \rightarrow \infty} \text{Cond}(\Sigma^t H) = 1$$

as the condition number of a matrix is always lower bounded by 1. Using our first remark about the equality of the condition number of the matrix $\Sigma^t H$ with the one of matrix Λ_t and equation (37), we also get

$$\begin{aligned} \limsup_{t \rightarrow \infty} \frac{\text{Cond}(\Sigma^{t+1}H) - 1}{\text{Cond}(\Sigma^t H) - 1} &= \limsup_{t \rightarrow \infty} \frac{1 - 2\nu_\Sigma^t \lambda_1^t}{1 - \nu_\Sigma^t \lambda_1^t} \\ &\leq \frac{1 - 2\nu_{\Sigma, \min}}{1 - \nu_{\Sigma, \min}}. \end{aligned}$$

This proves the equation (42). Moreover, if the limit $\nu_{\Sigma, \lim}$ exists, from (36), we see that one can replace $\nu_{\Sigma, \min}$ by $\nu_{\Sigma, \lim}$ in (42). This concludes the proof. \square

Remark D.3. Lemma D.4 is remarkable as it proves that the covariance gradually adapts to the inverse of the Hessian of our optimization problem, performing something quite similar to a Newton decrement.

We can now prove the speed convergence part of proposition 3.4 as follows:

Proof. Using lemma D.1 and lemma D.3, we can apply the matrix and vector norm inequality $\|Mx\|^2 \leq \sigma_1(M)^2 \|x\|^2$ to (25) and (26), we get

$$\frac{\|\mu^{t+1}\|}{\|\mu^t\|} \leq \sigma_1(I - \kappa \nu_\mu^t \Sigma^t H).$$

Likewise, the same applies for the update of the covariance matrix Σ given by equations (27) and (28), leading to a similar inequality for the matrix norm of the covariance

$$\frac{\|\Sigma^{t+1}\|}{\|\Sigma^t\|} \leq \sigma_1(I - \kappa \nu_\Sigma^t \Sigma^t H).$$

Let us look at the maximum singular value as follows:

$$\begin{aligned} &\limsup_t \sigma_1(I - \kappa \nu_\mu^t \Sigma^t H) \\ &= \limsup_t \sigma_1 \left(I - \nu_\mu^t \lambda_1(\kappa \Sigma^t H) \frac{\Sigma^t H}{\lambda_1(\Sigma^t H)} \right) \end{aligned}$$

We can tackle these terms one by one. Using the lemma D.4, we have that

$$\lim_{t \rightarrow \infty} \frac{\Sigma^t H}{\lambda_1(\Sigma^t H)} = I_d$$

while assumptions (7) states that

$$\nu_\mu^t \lambda_1(\kappa \Sigma^t H) \geq \nu_{\mu, \min}$$

leading to

$$\begin{aligned} & \limsup_t \sigma_1(I - c\nu_\mu^t \Sigma^t H) \\ &= \limsup_t \sigma_1\left(I - \nu_{\mu, \min} I_d\right) \leq 1 - \nu_{\mu, \min}. \end{aligned}$$

This implies linear convergence of μ with rate of convergence at most $1 - \nu_{\mu, \min}$, that is

$$\limsup_t \frac{\|\mu^{t+1}\|}{\|\mu^t\|} \leq 1 - \nu_{\mu, \min}, \quad (48)$$

Likewise, the reasoning is almost the same for the covariance matrix, and we also get

$$\limsup_t \sigma_1(I - \kappa \nu_\Sigma^t \Sigma^t H) \leq 1 - \nu_{\Sigma, \min}.$$

which proves the linear convergence of Σ with rate of convergence at most $1 - \nu_{\Sigma, \min}$, that is

$$\limsup_t \frac{\|\Sigma^{t+1}\|}{\|\Sigma^t\|} \leq 1 - \nu_{\Sigma, \min}, \quad (49)$$

which ends the proof. \square