

True Asymptotic Natural Gradient Optimization

Yann Ollivier

Abstract

We introduce a simple algorithm, True Asymptotic Natural Gradient Optimization (TANGO), that converges to a true natural gradient descent in the limit of small learning rates, without explicit Fisher matrix estimation.

For quadratic models the algorithm is also an instance of averaged stochastic gradient, where the parameter is a moving average of a “fast”, constant-rate gradient descent. TANGO appears as a particular de-linearization of averaged SGD, and is sometimes quite different on non-quadratic models. This further connects averaged SGD and natural gradient, both of which are arguably optimal asymptotically.

In large dimension, small learning rates will be required to approximate the natural gradient well. Still, this shows it is possible to get arbitrarily close to exact natural gradient descent with a lightweight algorithm.

Let $p_\theta(y|x)$ be a probabilistic model for predicting output values y from inputs x ($x = \emptyset$ for unsupervised learning). Consider the associated log-loss

$$\ell(y|x) := -\ln p_\theta(y|x) \quad (1)$$

Given a dataset \mathcal{D} of pairs (x, y) , we optimize the average log-loss over θ via a momentum-like gradient descent.

DEFINITION 1 (TANGO). *Let $\delta t_k \leq 1$ be a sequence of learning rates and let $\gamma > 0$. Set $v_0 = 0$. Iterate the following:*

- Select a sample (x_k, y_k) at random in the dataset \mathcal{D} .
- Generate a pseudo-sample \tilde{y}_k for input x_k according to the predictions of the current model, $\tilde{y}_k \sim p_\theta(\tilde{y}_k|x_k)$ (or just $\tilde{y}_k = y_k$ for the “outer product” variant). Compute gradients

$$g_k \leftarrow \frac{\partial \ell(y_k|x_k)}{\partial \theta}, \quad \tilde{g}_k \leftarrow \frac{\partial \ell(\tilde{y}_k|x_k)}{\partial \theta} \quad (2)$$

- Update the velocity and parameter via

$$v_k = (1 - \delta t_{k-1})v_{k-1} + \gamma g_k - \gamma(1 - \delta t_{k-1})(v_{k-1}^\top \tilde{g}_k) \tilde{g}_k \quad (3)$$

$$\theta_k = \theta_{k-1} - \delta t_k v_k \quad (4)$$

TANGO is built to approximate Amari’s *natural gradient* descent, namely, a gradient descent preconditioned by the inverse of the Fisher information matrix of the probabilistic model p_θ (see definitions below). The natural gradient arguably provides asymptotically optimal estimates of the parameter θ [Ama98]. However, its use is unrealistic for large-dimensional models due to the computational cost of storing and inverting the Fisher matrix, hence the need for approximations. One of its key features is its invariance to any change of variable in the parameter θ (contrary to simple gradient descent). The natural gradient is also a special case of the *extended Kalman filter* from estimation theory [Oll17], under mild conditions.

In TANGO, $\delta t/\gamma$ should be small for a good natural gradient approximation.

For stability of the update (3) of v , γ should be taken small enough; but a small γ brings slower convergence to the natural gradient. A conservative, theoretically safe choice is setting $\gamma = 1/\max \|\tilde{g}\|^2$ using the largest norm of \tilde{g} seen so far. This may produce a too small γ if gradients are unbounded. If the gradients follow a Gaussian distribution (with any covariance matrix), then $\gamma = 1/\mathbb{E}[3\|\tilde{g}\|^2]$ is theoretically safe; the average can be estimated on past gradients. In general, $\gamma \leq \mathbb{E}[\|\tilde{g}\|^2]/\mathbb{E}[\|\tilde{g}\|^4]$ is a necessary but not sufficient condition; this may be used as a starting point. (See discussion after Theorem 5.)

TANGO enjoys the following properties:

1. TANGO converges to an *exact* natural gradient trajectory when the learning rate δt tends to 0 with γ fixed, namely, to the trajectory of the ordinary differential equation $d\theta/dt = -J(\theta)^{-1}\mathbb{E}[\partial\ell/\partial\theta]$ with J the Fisher matrix at θ (Theorem 3).
2. For $\delta t = 1$ TANGO is an ordinary gradient descent with constant learning rate γ .
3. For quadratic losses, TANGO is an instance of *averaged stochastic gradient descent* with additional noise (Proposition 2): a “fast” stochastic gradient descent with constant learning rate is performed, and the algorithm returns a moving average of this trajectory (updated by a factor δt_k at each step). However, for non-quadratic losses, TANGO can greatly differ from averaged SGD (Fig. 1).

Thus, TANGO smoothly interpolates between ordinary and natural gradient descent when the learning rate decreases.

To illustrate the convergence to the natural gradient in an informal way, take $\delta t = 0$. Then θ does not move, and the average of g is the gradient of the expected loss at θ . Then the average of v over time converges to $(\mathbb{E}\tilde{g}\tilde{g}^\top)^{-1}\mathbb{E}g$, the exact natural gradient direction at θ . Indeed, this is the only fixed point of (3) in expectation. Actually, (3) is a way of solving for

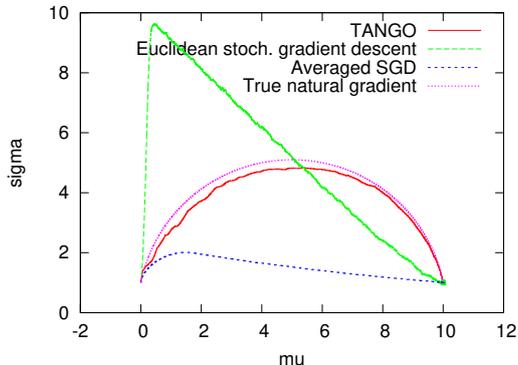


Figure 1: Learning a Gaussian model $\mathcal{N}(\mu, \sigma^2)$ with unknown μ and σ , via gradient descent on $(\mu, \ln \sigma)$. The initial point is $\mathcal{N}(0, 1)$ and the data are $\mathcal{N}(10, 1)$. The Fisher metric is isometric to the hyperbolic plane (μ, σ) , whose geodesics are circles, so that the true natural gradient starts by increasing variance so that μ moves faster. Plotted are trajectories of SGD with learning rate 10^{-3} , and TANGO and averaged SGD with $\gamma = 10^{-2}$ and $\delta t = 10^{-4}$.

$(\mathbb{E}\tilde{g}\tilde{g}^\top)v = \mathbb{E}g$ by stochastic gradient descent on v . The Fisher matrix J is $\mathbb{E}\tilde{g}\tilde{g}^\top$ by definition.

Acknowledgments. The author would like to thank Léon Bottou, Guillaume Charpiat, Fabrice Debbasch, Aaron Defazio, Gaétan Marceau-Caron and Corentin Tallec for helpful discussions and comments around these ideas.

Related work. Three different lines of work lead to TANGO-like algorithms. Averaged SGD [PJ92, Rup88] uses a “fast” gradient descent with large learning rate γ (here on the variable v), with an averaging operation on top (here by accumulation into θ). For linear problems γ can be kept constant.

Averaged SGD achieves the asymptotically optimal Cramer–Rao bound involving the inverse Fisher matrix, although “no explicit Hessian inversion has been performed” [MB11, PJ92]. TANGO may clarify how the implicit Hessian or Fisher matrix inversion occurs.

Later work on averaged SGD focussed on non-asymptotic behavior (especially, forgetting of the starting point), on somewhat dimension-independent bounds, and on larger γ for linear models [MB11, BM13, Mar14, DB15, DFB16]. A constant, large γ provides the most benefits; yet for nonlinear models, averaged SGD with constant γ leads to biases, hence the need for TANGO. Our analysis of the dynamics of v in TANGO and in Theorem 5 below follows this line of work.

Previous work on approximating the natural gradient for large-dimensional

models, such as TONGA and others [LMB07, Oll15, MG15, DSP⁺15, MCO16], did not provide an arbitrarily good approximation to the Fisher matrix, as it relied on structural matrix approximations (diagonal, block-diagonal, diagonal plus small-rank...) An exception is [DPCB13] for Boltzmann machines, directly transposed from the Hessian-free Newton method of [Mar10, MS11, MS12]: at each step, a large number of auxiliary conjugate gradient steps are performed to solve for Fisher matrix inversion, before the main update of the parameter occurs. From this viewpoint, TANGO performs the main gradient descent on θ and the auxiliary gradient descent at the same time.

For quasi-Newton methods in the convex case, auxiliary gradient descents to approximate the inverse Hessian have been suggested several times; see [ABH16, Mar10, MS11, MS12] and the references therein. Second-order methods for neural networks have a long history, see e.g. [LBOM98].¹

Third, “two-timescale” algorithms in reinforcement learning use updates reminiscent of TANGO, where the “fast” timescale is used to approximate a value function over a linear basis via a least squares method, and the “slow” timescale is used to adapt the parameters of a policy. For instance, the main results of [Tad04] or [KB17] deal with convergence of updates generalizing (3)–(4). However, these results crucially assume that both δt and γ tend to 0. This would be too slow in our setting. A constant γ can be used in TANGO (and in averaged SGD for linear least squares) thanks to the linearity of the update of v , but this requires a finer analysis of noise.

Discussion and shortcomings. Critical to TANGO is the choice of the parameter γ : the larger γ is, the faster the trajectory will resemble natural gradient (as v converges faster to $(\mathbb{E}\tilde{g}\tilde{g}^\top)^{-1}\mathbb{E}g$). However, if γ is too large the update for v is numerically unstable. For averaged SGD on quadratic losses, the choice of γ is theoretically well understood [DB15], but the situation is less clear for non-quadratic losses. We provide some general guidelines below.

The algorithmic interest of using TANGO with respect to direct Fisher matrix computation is not clear. Indeed, for $\delta t = 0$, the update equation (3) on v actually solves $v = (\mathbb{E}\tilde{g}\tilde{g}^\top)^{-1}\mathbb{E}g$ by stochastic gradient descent on v . The speed of convergence is heavily dimension-dependent, a priori. Similar Hessian-free Newton algorithms that rely on an auxiliary gradient descent to invert the Hessian, e.g., [Mar10], need a large number of auxiliary gradient iterations. In this case, the interest of TANGO may be its ease of implementation.

¹Technically the natural gradient is not a second-order method, as the Fisher matrix represents a Riemannian metric tensor rather than a Hessian of the loss. It can be computed from squared gradients, and the natural gradient is well-defined even if the loss is flat or concave. The Fisher matrix coincides with the Hessian of the loss function only asymptotically at a local minimum, provided the data follow the model.

Still, averaged SGD is proved to accelerate convergence for quadratic problems [PJ92]. So TANGO-like algorithms bring benefits in some regimes.

For linear models, [DFB16] study situations in which the convergence of (3) happens faster than suggested by the dimension of the problem, depending on the eigenvalues of the Hessian. For non-linear problems, this may be the case if the data clusters naturally in a few groups (e.g., classification with few labels): sampling a value of \tilde{y} in each of the clusters may already provide an interesting low-rank approximation of the Fisher matrix $\mathbb{E}\tilde{g}\tilde{g}^\top$. In such a situation, v may converge reasonably fast to an approximate natural gradient direction.

Implementation remarks: minibatches, preconditioned TANGO.

If \tilde{g} is computed as the average over a minibatch of size B , namely $\tilde{g} = \frac{1}{B} \sum_{i=1}^B \tilde{g}_i$ with \tilde{g}_i the gradient corresponding to output sample \tilde{y}_i in the minibatch, then the equation for v has to be modified to

$$v_k = (1 - \delta t_{k-1})v_{k-1} + \gamma g_k - \gamma B(1 - \delta t_{k-1})(v_{k-1}^\top \tilde{g}_k)\tilde{g}_k \quad (5)$$

because the expectation of $\tilde{g}\tilde{g}^\top$ is $\frac{1}{B}$ times the Fisher matrix.

Preconditioned TANGO (e.g., à la RMSProp) can be obtained by choosing a positive definite matrix C and iterating

$$v_k = (1 - \delta t_{k-1})v_{k-1} + \gamma C g_k - \gamma(1 - \delta t_{k-1})(v_{k-1}^\top \tilde{g}_k)C\tilde{g}_k \quad (6)$$

$$\theta_k = \theta_{k-1} - \delta t_k v_k \quad (7)$$

(This is TANGO on the variable $C^{-1/2}\theta$.) The matrix C may help to improve conditioning of gradients and of the matrix $C\mathbb{E}\tilde{g}\tilde{g}^\top$. Choices of C may include RMSProp (the entrywise reciprocal of the root-mean-square average of gradients) or the inverse of the diagonal Fisher matrix, $C^{-1} = \text{diag}(\mathbb{E}\tilde{g}^{\odot 2})$. These options will require different adjustments for γ .

Quadratic output losses can be seen as the log-loss of a probabilistic model, $\ell(y|x) = \frac{\|y - f_\theta(x)\|^2}{2\sigma^2}$ for any value of σ^2 . However, σ^2 should be set to the actual mean square error on the outputs, for the natural gradient descent to work best. The choice of σ^2 affects both the scaling of gradients g and \tilde{g} , and the sampling of pseudo-samples \tilde{y} , whose law is $\mathcal{N}(f_\theta(x), \sigma^2)$.

TANGO as an instance of averaged SGD for quadratic losses. Averaged SGD maintains a fast-moving parameter with constant learning rate, and returns a moving average of the fast trajectory. It is known to have excellent asymptotic properties for quadratic models.

For quadratic losses, TANGO can be rewritten as a form of averaged SGD, despite TANGO only using gradients evaluated at the “slow” parameter θ . This is specific to gradients being a linear function of θ .

Thus TANGO can be considered as a non-linearization of averaged SGD, written using gradients at θ only. Even for simple nonlinear models, the difference can be substantial (Fig. 1). For nonlinear models, averaged SGD with a fixed learning rate γ can have a bias of size comparable to γ , even with small δt .² TANGO does not exhibit such a bias.

PROPOSITION 2. *Assume that for each sample (x, y) , the log-loss $\ell(y|x)$ is a quadratic function of θ whose Hessian does not depend on y (e.g., linear regression $\ell(y|x) = \frac{1}{2} \|y - \theta^\top x\|^2$).*

Then TANGO is identical to the following trajectory averaging algorithm:

$$\theta_k^{\text{fast}} = \theta_{k-1}^{\text{fast}} - \gamma \frac{\partial \ell(y_k|x_k)}{\partial \theta_{k-1}^{\text{fast}}} + \gamma \xi_k \quad (8)$$

$$\theta_k = (1 - \delta t_k) \theta_{k-1} + \delta t_k \theta_k^{\text{fast}} \quad (9)$$

where ξ_k is some centered random variable whose law depends on $\theta_{k-1}^{\text{fast}}$ and θ_{k-1} . The identification with TANGO is via $v_k = \theta_{k-1} - \theta_k^{\text{fast}}$.

The proof (Appendix A) is mostly by direct algebraic manipulations. For quadratic losses, the gradients are a linear function of the parameter, so that the derivative at point θ^{fast} can be rewritten as the derivative at point θ plus a Hessian term; for quadratic losses, the Hessian is equal to the Fisher metric.

The additional noise ξ_k is multiplicative in v . This is standard for linear regression [DFB16]: indeed, in linear regression, the gradient from sample (x, y) is $-yx + xx^\top \theta$, and its expectation is $-\mathbb{E}(yx) + \mathbb{E}(xx^\top) \theta$ so that the gradient noise has a multiplicative component $(xx^\top - \mathbb{E}(xx^\top)) \theta$. (Treatments of gradient descent often assume additive noise instead, see discussion in [DFB16].)

Replacing the TANGO update of θ in (4) with $\theta_k = \theta_{k-1} - v_k$ would make TANGO equivalent to an *accelerated gradient* method with additional noise for quadratic functions.

Convergence of TANGO to the natural gradient. Let the Fisher matrix of the model be

$$J(\theta) := \mathbb{E} \tilde{g} \tilde{g}^\top = \mathbb{E}_{(x,y) \in \mathcal{D}} \mathbb{E}_{\tilde{y} \sim p_\theta(\tilde{y}|x)} \frac{\partial \ell(\tilde{y}|x)}{\partial \theta}^{\otimes 2} \quad (10)$$

where, for a column vector v , $v^{\otimes 2}$ is the outer product vv^\top .

²A bias of size γ is easy to see on the following example: Define a loss $\ell(x) = |x|$ for $|x| \geq \gamma/2$, and extend this loss in an arbitrary way on the interval $[-\gamma/2; \gamma/2]$. Since the gradients are ± 1 out of this interval, a gradient descent with fixed learning rate γ , initialized at a multiple of $\gamma/2$, will make jumps of size exactly γ and never visit the interior of the interval $[-\gamma/2; \gamma/2]$. Whatever the average parameter of this trajectory is, it is unrelated to the behavior of the loss on $[-\gamma/2; \gamma/2]$ and to the location of the minimum. Thus averaged SGD can have a bias of size $\approx \gamma$, whatever δt .

The stochastic natural gradient descent on θ , with learning rate δt , using the exact Fisher matrix $J(\theta)$, is

$$\theta^{t+\delta t} = \theta^t - \delta t J(\theta^t)^{-1} \frac{\partial \ell(y_k | x_k)}{\partial \theta^t} \quad (11)$$

where at each step (x_k, y_k) is a random sample from the dataset \mathcal{D} . In the limit of small learning rates $\delta t \rightarrow 0$, it converges to a “true” continuous-time natural gradient descent trajectory, driven by the differential equation

$$\frac{d\theta^t}{dt} = -J(\theta^t)^{-1} \mathbb{E}_{(x,y) \in \mathcal{D}} \frac{\partial \ell(y|x)}{\partial \theta^t} \quad (12)$$

THEOREM 3. *Make the following regularity assumptions: The second moment of gradients g is bounded over θ . The fourth moment of gradients \tilde{g} is bounded over θ . The lowest eigenvalue of the Fisher matrix $J(\theta)$, as a function of θ , is bounded away from 0. The Fisher matrix is a C^1 function of θ with bounded first derivatives.*

Let θ^T be the value of the exact natural gradient (12) at time T . Assume that the parameter γ in TANGO is smaller than some constant that depends on the moments of the gradients and the eigenvalues of the Fisher matrix.

Then the value of θ obtained after $T/\delta t$ iterations of TANGO converges in probability to θ^T , when $\delta t \rightarrow 0$.

The probability in this theorem refers to the random choice of samples x_k , y_k and \tilde{y}_k in TANGO.

Theorem 3 will be obtained as a corollary of the more general Theorem 5, which also provides quantitative versions of the choice of γ in TANGO.

To illustrate a key idea of the proof, we start with a simpler, noise-free situation.

PROPOSITION 4. *Consider the iteration of*

$$v_k = v_{k-1} + \gamma F(\theta_{k-1}) - \gamma A(\theta_{k-1}) v_{k-1} \quad (13)$$

$$\theta_k = \theta_{k-1} - \delta t v_k \quad (14)$$

initialized at $v_0 = 0$, where F is a vector field on θ and A is a field of symmetric positive definite matrices.

Assume that F and A are C^1 with bounded derivatives. Let $\lambda_{\min} := \inf_{\theta} \min \text{eigenvalues}(A(\theta))$ and $\lambda_{\max} := \sup_{\theta} \max \text{eigenvalues}(A(\theta))$, and assume $\lambda_{\min} > 0$ and $\lambda_{\max} < \infty$. Fix γ smaller than $1/\lambda_{\max}$.

Then when $\delta t \rightarrow 0$, the value θ of this system after $T/\delta t$ iterations converges to the value at time T of the ordinary differential equation with preconditioning A^{-1} ,

$$\frac{d\theta^t}{dt} = -A(\theta^t)^{-1} F(\theta^t) \quad (15)$$

initialized at $\theta^0 = \theta_0$. More precisely, $\theta_{T/\delta t} - \theta^T = O(\delta t)$.

PROOF.

We first deal with the case of constant $A(\theta) \equiv A$.

First, note that the sums of the contributions of v_1 to all future updates of θ is $\delta t \sum (\text{Id} - \gamma A)^k v_1 = \delta t \gamma^{-1} A^{-1} v_1$.

This suggests setting

$$z_{k+1} := \theta_k - \delta t \gamma^{-1} A^{-1} v_{k+1} \quad (16)$$

which contains “ θ_k plus all the known future updates from the terms $F(\theta_j)$, $j \leq k$, that are already present in v_k ”. Substituting for θ_k and v_{k+1} in z_{k+1} , one finds that the update for z is

$$z_{k+1} = \theta_{k-1} - \delta t v_k - \delta t \gamma^{-1} A^{-1} (v_k + \gamma F(\theta_k) - \gamma A v_k) \quad (17)$$

$$= z_k - \delta t A^{-1} F(\theta_k) \quad (18)$$

which only involves the new contribution from $F(\theta_n)$, and not v .

Moreover,

$$z_k = \theta_{k-1} - \delta t \gamma^{-1} A^{-1} v_k = \theta_k + \delta t v_k - \delta t \gamma^{-1} A^{-1} v_k = \theta_k + O(\delta t \|v_k\|) \quad (19)$$

since A^{-1} is bounded (its largest eigenvalue is $1/\lambda_{\min}$).

Now, the update for v_k is $(1 - \gamma \lambda_{\min})$ -contracting, because the condition $\gamma < 1/\lambda_{\max}$ implies that the eigenvalues of γA lie between $\gamma \lambda_{\min}$ and 1. Since $\lambda_{\min} > 0$ and F is bounded, it is easy to show by induction that $\|v_k\| \leq (\sup \|F\|)/\lambda_{\min}$ so that v is bounded.

Therefore, $z_k = \theta_k + O(\delta t)$. Then, given the regularity assumptions on F , one has $F(\theta_k) = F(z_k) + O(\delta t)$ and

$$z_{k+1} = z_k - \delta t A^{-1} F(z_k) + O(\delta t^2) \quad (20)$$

since A^{-1} is bounded. This does not involve v any more.

But this update for z_k is just a Euler numerical scheme for the differential equation $\dot{z} = -A^{-1} F(z)$. So by the standard theory of approximation of ordinary differential equations, when $\delta t \rightarrow 0$, $z_{T/\delta t}$ converges to the solution at time T of this equation, within an error $O(\delta t)$. Since $\theta_k - z_k$ is $O(\delta t)$ as well, we get the same conclusion for θ .

For the case of variable A , set

$$z_{k+1} := \theta_k - \delta t \gamma^{-1} A^{-1}(\theta_k) v_{k+1} \quad (21)$$

and substituting for θ_k and v_{k+1} in this definition, one finds

$$z_{k+1} = \theta_{k-1} - \delta t \gamma^{-1} A(\theta_k)^{-1} v_k - \delta t A(\theta_k)^{-1} F(\theta_k) \quad (22)$$

$$= z_k + \delta t \gamma^{-1} (A(\theta_{k-1})^{-1} - A(\theta_k)^{-1}) v_k - \delta t A(\theta_k)^{-1} F(\theta_k) \quad (23)$$

Now, under our eigenvalue assumptions, A^{-1} is bounded. Since A has bounded derivatives, so does A^{-1} thanks to $\partial_\theta A^{-1} = -A^{-1}(\partial_\theta A)A^{-1}$. Therefore we can apply a Taylor expansion of A^{-1} so that

$$A(\theta_{k-1})^{-1} - A(\theta_k)^{-1} = O(\theta_{k-1} - \theta_k) = O(\delta t \|v_k\|) \quad (24)$$

so that

$$z_{k+1} = z_k - \delta t A(\theta_k)^{-1} F(\theta_k) + O(\delta t^2 \|v_k\|^2) \quad (25)$$

after which the proof proceeds as for the case of constant A , namely: $z_k - \theta_k$ is $O(\delta t \|v_k\|)$ so that

$$z_{k+1} = z_k - \delta t A(z_k)^{-1} F(z_k) + O(\delta t^2 \|v_k\| + \delta t^2 \|v_k\|^2) \quad (26)$$

and v_k is bounded by induction. So the update for z_k is a Euler numerical scheme for the differential equation $\dot{z} = -A(z)^{-1}F(z)$, which ends the proof. \square

We now turn to the stochastic version of Proposition 4. This provides a generalization of Theorem 3: Theorem 3 is a corollary of Theorem 5 using $\hat{F}_k = g_k$ and $\hat{A}_k = (1 - \delta t)\tilde{g}_k\tilde{g}_k^\top + \frac{\delta t}{\gamma} \text{Id}$.

For numerical simulations of stochastic differential equations, the usual rate of convergence is $O(\sqrt{\delta t})$ rather than $O(\delta t)$ [KP92].

THEOREM 5. *Consider the iteration of*

$$v_k = v_{k-1} + \gamma \hat{F}_k - \gamma \hat{A}_k v_{k-1} \quad (27)$$

$$\theta_k = \theta_{k-1} - \delta t v_k \quad (28)$$

initialized at $v_0 = 0$, where \hat{F}_k is a vector-valued random variable and \hat{A}_k is a symmetric-matrix-valued random variable.

Let \mathcal{F}_k be the sigma-algebra generated by all variables up to time k , and abbreviate \mathbb{E}_k for $\mathbb{E}[\cdot | \mathcal{F}_k]$. Let

$$F_k := \mathbb{E}_{k-1} \hat{F}_k, \quad A_k := \mathbb{E}_{k-1} \hat{A}_k \quad (29)$$

and assume that these depend on θ_{k-1} only, namely, that exist functions $F(\theta)$ and $A(\theta)$ such that

$$F_k = F(\theta_{k-1}), \quad A_k = A(\theta_{k-1}) \quad (30)$$

Assume that the functions F and A are C^1 with bounded derivatives. Let $\lambda := \inf_\theta \min \text{eigenvalues}(A(\theta))$, and assume $\lambda > 0$.

Assume the following variance control: for some $\sigma^2 \geq 0$ and $R^2 \geq 0$,

$$\mathbb{E}_{k-1} \left\| \hat{F}_k \right\|^2 \leq \sigma^2, \quad \mathbb{E}_{k-1} \left[\hat{A}_k^\top \hat{A}_k \right] \preceq R^2 A_k \quad (31)$$

where $A \preceq B$ means $B - A$ is positive semidefinite.

Fix $0 < \gamma \leq 1/R^2$.

Then when $\delta t \rightarrow 0$, the value θ of this system after $T/\delta t$ iterations converges in probability to the value at time T of the ordinary differential equation with preconditioning A^{-1} ,

$$\frac{d\theta^t}{dt} = -A(\theta^t)^{-1}F(\theta^t) \quad (32)$$

initialized at $\theta^0 = \theta_0$.

More precisely, for any $\varepsilon > 0$, with probability $\geq 1 - \varepsilon$ one has $\theta_{T/\delta t} - \theta^T = O(\sqrt{\delta t})$ when the constant in $O(\cdot)$ depends on $\varepsilon, T, \lambda, \gamma, \sigma^2, R^2$, and the derivatives of $F(\theta)$ and $A(\theta)$. The bounds are uniform for T in compact intervals.

The variance assumption on \hat{A} directly controls the maximum possible value via $\gamma \leq 1/R^2$, and, consequently, the speed of convergence to A^{-1} . This assumption appears in [BM13, DB15, DFB16] for $\hat{A} = \tilde{g}\tilde{g}^\top$, where the value of R^2 for typical cases is discussed.

With $\hat{A} = \tilde{g}\tilde{g}^\top$, the variance assumption on \hat{A} is always satisfied with $R^2 = \sup \|\tilde{g}\|^2$ if \tilde{g} is bounded.³ It is also satisfied with $R^2 = \mathbb{E}\|\tilde{g}\|^4/\lambda$, without bounded gradients. (Indeed, first, one has $\mathbb{E}\hat{A}^2 = \mathbb{E}(\|\tilde{g}\|^2 \tilde{g}\tilde{g}^\top) \leq (\sup \|\tilde{g}\|^2) \mathbb{E}\tilde{g}\tilde{g}^\top$; second, for any vector u , one has $u^\top \mathbb{E}[\tilde{g}\tilde{g}^\top \tilde{g}\tilde{g}^\top]u = \mathbb{E}[u^\top \tilde{g}\tilde{g}^\top \tilde{g}\tilde{g}^\top u] \leq \mathbb{E}[\|u\|^2 \|\tilde{g}\|^4] = \|u\|^2 \mathbb{E}\|\tilde{g}\|^4$ while $u^\top Au$ is at least $\lambda\|u\|^2$.) If the distribution of \tilde{g} has bounded curtosis κ in every direction, then the assumption is satisfied with $R^2 = \kappa \mathbb{E}\|\tilde{g}\|^2$ [DFB16]; in particular, for Gaussian \tilde{g} , with any covariance matrix, the assumption is satisfied with $R^2 = 3\mathbb{E}\|\tilde{g}\|^2$. All these quantities can be estimated based on past values of \tilde{g} .

Theorem 5 would still be valid with additional centered noise on θ and additional $o(\delta t)$ terms on θ ; for simplicity we did not include them, as they are not needed for TANGO.

LEMMA 6. *Under assumptions of Theorem 5, the largest eigenvalue of $A(\theta)$ is at most R^2 . The operator $(\text{Id} - \gamma A(\theta))$ is $(1 - \gamma\lambda)$ -contracting.*

Moreover, $\theta \mapsto A^{-1}(\theta)$ exists, is bounded, and is C^1 with bounded derivatives. The same holds for $\theta \mapsto A^{-1}(\theta)F(\theta)$.

PROOF.

First, for any vector u , one has $\|Au\|^2 = \|\mathbb{E}\hat{A}u\|^2 \leq \mathbb{E}\|\hat{A}u\|^2 = \mathbb{E}[u^\top \hat{A}^\top \hat{A}u] \leq R^2 u^\top Au$. Taking u an eigenvector associated with the largest eigenvalue λ_{\max} of A shows that $\lambda_{\max} \leq R^2$. Next, the eigenvalues of A lie between λ and

³ TANGO uses $\hat{A} = (1 - \delta t)\tilde{g}\tilde{g}^\top + \frac{\delta t}{\gamma} \text{Id}$ rather than $\hat{A} = \tilde{g}\tilde{g}^\top$. Actually it is enough to check the assumption with $\tilde{g}\tilde{g}^\top$. Indeed one checks that if $\tilde{g}\tilde{g}^\top$ satisfies the assumption with some R^2 , then $(1 - \delta t)\tilde{g}\tilde{g}^\top + \frac{\delta t}{\gamma} \text{Id}$ satisfies the assumption with $\max(R^2, 1/\gamma)$, and that $\gamma \leq 1/R^2$ implies $\gamma \leq 1/\max(R^2, 1/\gamma)$.

R^2 so that the eigenvalues of γA lie between $\gamma\lambda$ and 1. So the eigenvalues of $\text{Id} - \gamma A$ lie between 0 and $1 - \gamma\lambda$.

Since A is symmetric and its smallest eigenvalue is $\lambda > 0$, it is invertible with its inverse bounded by $1/\lambda$. Thanks to $\partial_\theta A^{-1} = -A^{-1}(\partial_\theta A)A^{-1}$, the derivatives of A^{-1} are bounded. \square

LEMMA 7. *Under the notation and assumptions of Theorem 5, for any k ,*

$$\mathbb{E} \|v_k\|^2 \leq \frac{4\sigma^2}{\lambda^2} \quad (33)$$

Up to the factor 4, this is optimal: indeed, when \hat{F} and \hat{A} have a distribution independent of k , the fixed point of v in expectation is $v = A^{-1}\mathbb{E}\hat{F}$, whose square norm is $(\mathbb{E}\hat{F})^\top A^{-2}\mathbb{E}\hat{F}$ which is $\|\mathbb{E}\hat{F}\|^2/\lambda^2$ if $\mathbb{E}\hat{F}$ lies in the direction of the eigenvalue λ .

PROOF.

The proof is a variant of arguments appearing in [BM13]; in our case A is not constant, \hat{F}_k is not centered, \hat{A}_k is not rank-one, and we do not use the norm associated with A on the left-hand-side. Let

$$w_k := (\text{Id} - \gamma\hat{A}_k)v_{k-1} \quad (34)$$

so that $v_k = w_k + \gamma\hat{F}_k$. Consequently

$$\|v_k\|^2 = \|w_k\|^2 + \|\gamma\hat{F}_k\|^2 + 2\gamma w_k \cdot \hat{F}_k \leq (1+\alpha)\|w_k\|^2 + (1+1/\alpha)\|\gamma\hat{F}_k\|^2 \quad (35)$$

for any $\alpha > 0$, thanks to $2ab = 2(\sqrt{\alpha}a)(b/\sqrt{\alpha}) \leq \alpha a^2 + b^2/\alpha$ for any $\alpha > 0$ and $a, b \in \mathbb{R}$.

Now

$$\|w_k\|^2 = \|v_{k-1}\|^2 - \gamma v_{k-1}^\top (\hat{A}_k + \hat{A}_k^\top) v_{k-1} + \gamma^2 v_{k-1}^\top \hat{A}_k^\top \hat{A}_k v_{k-1} \quad (36)$$

Take expectations conditionally to \mathcal{F}_{k-1} . Using $\mathbb{E}_{k-1} [\hat{A}_k^\top \hat{A}_k] \preceq R^2 A_k$ we find

$$\mathbb{E}_{k-1} \|w_k\|^2 \leq \|v_{k-1}\|^2 - \gamma(2 - \gamma R^2) v_{k-1}^\top A_k v_{k-1} \quad (37)$$

By the assumptions, $\gamma R^2 \leq 1$ and $v_{k-1}^\top A_k v_{k-1} \geq \lambda \|v_{k-1}\|^2$. Thus

$$\mathbb{E}_{k-1} \|w_k\|^2 \leq (1 - \gamma\lambda) \|v_{k-1}\|^2 \quad (38)$$

Taking $1 + \alpha = \frac{1 - \gamma\lambda/2}{1 - \gamma\lambda}$ we find

$$\mathbb{E}_{k-1} \|v_k\|^2 \leq (1 - \gamma\lambda/2) \|v_{k-1}\|^2 + (1 + 1/\alpha)\gamma^2\sigma^2 \quad (39)$$

$$\leq (1 - \gamma\lambda/2) \|v_{k-1}\|^2 + \frac{1 - \gamma\lambda/2}{\gamma\lambda/2} \gamma^2\sigma^2 \quad (40)$$

Taking unconditional expectations, we obtain

$$\mathbb{E} \|v_k\|^2 \leq (1 - \gamma\lambda/2)\mathbb{E} \|v_{k-1}\|^2 + \frac{1 - \gamma\lambda/2}{\gamma\lambda/2} \gamma^2 \sigma^2 \quad (41)$$

and by induction, starting at $v_0 = 0$, this implies

$$\mathbb{E} \|v_k\|^2 \leq \frac{1 - \gamma\lambda/2}{(\gamma\lambda/2)^2} \gamma^2 \sigma^2 \leq \frac{4\sigma^2}{\lambda^2} \quad (42)$$

□

COROLLARY 8. *Under the notation and assumptions of Theorem 5, for any n , for any $\varepsilon > 0$, with probability $\geq 1 - \varepsilon$ one has*

$$\sup_{0 \leq k \leq n} \|v_k\| \leq \frac{2\sigma}{\lambda} \sqrt{\frac{n}{\varepsilon}} \quad (43)$$

PROOF.

This follows from Lemma 7 by the Markov inequality and a union bound. □

The next two lemmas result from standard martingale arguments; the detailed proofs are given in the Appendix.

LEMMA 9. *Under the notation and assumptions of Theorem 5, let ξ be the noise on F ,*

$$\xi_k := \hat{F}_k - F_k \quad (44)$$

Let (M_k) be any sequence of operators such that M_k is \mathcal{F}_{k-1} -measurable and $\|M_k\|_{\text{op}} \leq \Lambda$ almost surely.

Then

$$\mathbb{E} \sum_{j=1}^n \|M_j \xi_j\|^2 \leq n\Lambda^2 \sigma^2 \quad (45)$$

and moreover for any n , for any $\varepsilon > 0$, with probability $\geq 1 - \varepsilon$, for any $k \leq n$ one has

$$\left\| \sum_{j=k}^n M_j \xi_j \right\| \leq 2\sqrt{\frac{n\Lambda^2 \sigma^2}{\varepsilon}} \quad (46)$$

LEMMA 10. *Under the notation and assumptions of Theorem 5, set*

$$\zeta_k := (\hat{A}_k - A_k)v_{k-1} \quad (47)$$

Let (M_k) be any sequence of operators such that M_k is \mathcal{F}_{k-1} -measurable and $\|M_k\|_{\text{op}} \leq \Lambda$ almost surely. Let $\lambda_{\max} = \sup_{\theta} \max \text{eigenvalues}(A_k)$, which is finite by Lemma 6.

Then

$$\mathbb{E} \sum_{j=1}^n \|M_j \zeta_j\|^2 \leq 4nR^2 \lambda_{\max} \Lambda^2 \sigma^2 / \lambda^2 \quad (48)$$

and moreover, for any n , for any $\varepsilon > 0$, with probability $\geq 1 - \varepsilon$, for any $k \leq n$,

$$\left\| \sum_{j=k}^n M_j \zeta_j \right\| \leq 4 \sqrt{\frac{n R^2 \lambda_{\max} \Lambda^2 \sigma^2}{\varepsilon \lambda^2}} \quad (49)$$

PROOF OF THEOREM 5.

Let $n := T / \delta t$ be the number of discrete steps corresponding to continuous time T . All the constants implied in $O()$ notation below depend on T and on the assumptions of the theorem (R^2 , γ , λ , etc.), and we study the dependency on δt .

Similarly to Proposition 4, set

$$z_k := \theta_{k-1} - \delta t \gamma^{-1} B_k v_k \quad (50)$$

where B_k is a matrix to be defined later (equal to A^{-1} for the case of constant A). Informally, z contains θ plus the future updates to be made to θ based on the current value of v .

Substituting $\theta_{k-1} = \theta_{k-2} - \delta t v_{k-1}$ and $v_k = v_{k-1} + \gamma \hat{F}_k - \gamma A_k v_{k-1} - \gamma \zeta_k$ into the definition of z_k , one finds

$$z_k = \theta_{k-2} - \delta t v_{k-1} - \delta t \gamma^{-1} B_k (v_{k-1} + \gamma \hat{F}_k - \gamma A_k v_{k-1} - \gamma \zeta_k) \quad (51)$$

$$= \theta_{k-2} - \delta t B_k (\hat{F}_k - \zeta_k) - \delta t (\text{Id} + \gamma^{-1} B_k - B_k A_k) v_{k-1} \quad (52)$$

$$= z_{k-1} - \delta t B_k (\hat{F}_k - \zeta_k) - \delta t (\text{Id} - B_k A_k + \gamma^{-1} (B_k - B_{k-1})) v_{k-1} \quad (53)$$

Now define B_k in order to cancel the v_{k-1} term, namely

$$B_{k-1} := B_k + \gamma (\text{Id} - B_k A_k) \quad (54)$$

initialized with $B_n := A_n^{-1}$. (If A is constant, then $B = A^{-1}$.) Then $\delta t \gamma^{-1} B_k v_k$ represents all the future updates to θ stemming from the current value v_k .

With this choice, the update for z is

$$z_k = z_{k-1} - \delta t B_k (\hat{F}_k - \zeta_k) = z_{k-1} - \delta t B_k (F_k + \xi_k - \zeta_k) \quad (55)$$

Remove the noise by defining

$$y_k := z_k - \delta t \sum_{j=k+1}^n B_j (\xi_j - \zeta_j) \quad (56)$$

so that

$$y_k = y_{k-1} - \delta t B_k F_k \quad (57)$$

Assume for now that $B_k = A^{-1}(\theta_{k-1}) + O(\sqrt{\delta t})$. Then

$$y_k = y_{k-1} - \delta t A^{-1}(\theta_{k-1}) F(\theta_{k-1}) + O(\delta t^{3/2}) \quad (58)$$

Since $A^{-1}F$ is Lipschitz (Lemma 6), we have

$$y_k = y_{k-1} - \delta t A^{-1}(y_{k-1})F(y_{k-1}) + O(\delta t \|y_{k-1} - \theta_{k-1}\|) + O(\delta t^{3/2}) \quad (59)$$

If we prove that $y_{k-1} - \theta_{k-1} = O(\sqrt{\delta t})$ then we find

$$y_k = y_{k-1} - \delta t A^{-1}(y_{k-1})F(y_{k-1}) + O(\delta t^{3/2}) \quad (60)$$

so that y_k is a Euler numerical scheme for the differential equation $\dot{y} = -A^{-1}(y)F(y)$, and thus converges to the natural gradient trajectory up to $O(\sqrt{\delta t})$, uniformly on the time interval $[0; T]$.

Since we assumed that $\theta_k - y_k = O(\sqrt{\delta t})$, this holds for θ_k as well.

We still have to prove the two assumptions that $y_{k-1} - \theta_{k-1} = O(\sqrt{\delta t})$ and that $B_k = A^{-1}(\theta_{k-1}) + O(\sqrt{\delta t})$.

LEMMA 11. *Define $B_{k-1} := B_k + \gamma(\text{Id} - B_k A_k)$ initialized with $B_n := A_n^{-1}$. Then for any $\varepsilon > 0$, with probability $\geq 1 - \varepsilon$, one has $\sup_k \|B_k - A_k^{-1}\|_{\text{op}} = O(\sqrt{\delta t})$.*

PROOF OF LEMMA 11.

With this definition one has

$$B_{k-1} - A_{k-1}^{-1} = (B_k - A_k^{-1})(\text{Id} - \gamma A_k) + A_k^{-1} - A_{k-1}^{-1} \quad (61)$$

by a direct computation.

Now $A_k^{-1} - A_{k-1}^{-1} = A^{-1}(\theta_{k-1}) - A^{-1}(\theta_{k-2}) = O(\theta_{k-1} - \theta_{k-2})$ because A^{-1} is Lipschitz. Moreover $\theta_{k-1} = \theta_{k-2} - \delta t v_{k-1}$. So $A_k^{-1} - A_{k-1}^{-1} = O(\delta t \|v_{k-1}\|)$. Thanks to Corollary 8, with probability $\geq 1 - \varepsilon$, $\sup_k \|v_{k-1}\| = O(\sqrt{n}) = O(1/\sqrt{\delta t})$ so that $A_k^{-1} - A_{k-1}^{-1}$ is $O(\sqrt{\delta t})$, uniformly in k .

Now, the operator $(\text{Id} - \gamma A_k)$ is $(1 - \gamma\lambda)$ -contracting. Therefore,

$$\|B_{k-1} - A_{k-1}^{-1}\|_{\text{op}} \leq (1 - \gamma\lambda) \|B_k - A_k^{-1}\|_{\text{op}} + O(\sqrt{\delta t}) \quad (62)$$

and $B_n - A_n^{-1}$ is 0, so by induction, $\|B_{k-1} - A_{k-1}^{-1}\|_{\text{op}} = O(\sqrt{\delta t})$, uniformly in k . \square

Back to the proof of Theorem 5. To prove that $y_k - \theta_k = O(\sqrt{\delta t})$, let us first prove that $y_k - z_k = O(\sqrt{\delta t})$. We have

$$z_k - y_k = \delta t \sum_{j=k+1}^n B_j (\xi_j - \zeta_j) \quad (63)$$

Thanks to Lemma 11, this rewrites as

$$z_k - y_k = \delta t \sum_{j=k+1}^n A_j^{-1} (\xi_j - \zeta_j) + O\left(\delta t^{3/2} \sum_{j=k+1}^n (\|\xi_j\| + \|\zeta_j\|)\right) \quad (64)$$

For the first term, note that $A_j^{-1} = A^{-1}(\theta_{j-1})$ is \mathcal{F}_{j-1} -measurable (while B_j is not, because it depends on θ_k for $k \geq j$). By Lemmas 9 and 10, $\sum A_j \xi_j$ and $\sum A_j \zeta_j$ are both $O(\sqrt{n}) = O(\sqrt{1/\delta t})$ with high probability. So the first term of $z_k - y_k$ is $O(\sqrt{\delta t})$.

For the second term,

$$\sum_{j=k+1}^n \|\xi_j\| \leq \sum_{j=1}^n \|\xi_j\| \leq \sqrt{n} \sqrt{\sum_{j=1}^n \|\xi_j\|^2} \quad (65)$$

by Cauchy–Schwarz. By Lemma 9, $\mathbb{E} \sum \|\xi_j\|^2$ is $O(n)$. So with probability $\geq 1 - \varepsilon$, thanks to the Markov inequality, $\sqrt{\sum \|\xi_j\|^2}$ is $O(\sqrt{n})$ where the constant in $O()$ depends on ε . Therefore, $\sum_{j=k}^n \|\xi_j\|$ is $O(n) = O(1/\delta t)$. The same argument applies to ζ thanks to Lemma 10.

Therefore, $z_k - y_k$ is $O(\sqrt{\delta t})$.

Finally, $z_k - \theta_k$ is $O(\delta t \|v_k\|)$ which is $O(\delta t \sqrt{n}) = O(\sqrt{\delta t})$ by Corollary 8. Therefore $y_k - \theta_k$ is $O(\sqrt{\delta t})$ as well. \square

A Additional proofs

PROOF OF PROPOSITION 2.

Start with the algorithm in Proposition 2, with any noise ξ_k . Under the update for θ_k one has

$$\theta_k - \theta_k^{\text{fast}} = (1 - \delta t_k)(\theta_{k-1} - \theta_k^{\text{fast}}) \quad (66)$$

Now set

$$v_k := \theta_{k-1} - \theta_k^{\text{fast}} \quad (67)$$

so that the update for θ_k is $\theta_k = \theta_{k-1} - \delta t_k \theta_{k-1} + \delta t_k \theta_k^{\text{fast}} = \theta_{k-1} - \delta t_k v_k$ by construction. To determine the update for v , remove θ_{k-1} from the update of θ_k^{fast} :

$$\theta_k^{\text{fast}} - \theta_{k-1} = \theta_{k-1}^{\text{fast}} - \theta_{k-1} - \gamma g_k^{\text{fast}} + \gamma \xi_k \quad (68)$$

where we abbreviate $g_k^{\text{fast}} := \frac{\partial \ell(y_k | x_k)}{\partial \theta_{k-1}^{\text{fast}}}$, the gradient of the loss at $\theta_{k-1}^{\text{fast}}$.

Let H_k be the Hessian of the loss on the k -th example with respect to the parameter. Since losses are quadratic, the gradient of the loss is a linear function of the parameter:

$$g_k^{\text{fast}} = g_k + H_k(\theta_{k-1}^{\text{fast}} - \theta_{k-1}) \quad (69)$$

where $g_k := \frac{\partial \ell(y_k | x_k)}{\partial \theta_{k-1}}$ is the gradient of the loss at θ_{k-1} .

Thus (68) rewrites as

$$v_k = -\theta_{k-1}^{\text{fast}} + \theta_{k-1} + \gamma g_k + \gamma H_k(\theta_{k-1}^{\text{fast}} - \theta_{k-1}) - \gamma \xi_k \quad (70)$$

and thanks to (66),

$$\theta_{k-1} - \theta_{k-1}^{\text{fast}} = (1 - \delta t_{k-1})v_{k-1} \quad (71)$$

so the above rewrites as

$$v_k = (1 - \delta t_{k-1})v_{k-1} + \gamma g_k - \gamma(1 - \delta t_{k-1})H_k v_{k-1} - \gamma \xi_k \quad (72)$$

If we set

$$\xi_k := (1 - \delta t_{k-1})(\tilde{g}_k \tilde{g}_k^\top - H_k)v_{k-1} \quad (73)$$

then this is identical to TANGO. However, we still have to prove that such a ξ_k is a centered noise, namely, $\mathbb{E}\xi_k = 0$. This will be the case if

$$H_k = \mathbb{E}\tilde{g}_k \tilde{g}_k^\top \quad (74)$$

where the expectation is with respect to the choice of the random output \tilde{y}_k given x_k . From the double definition of the Fisher matrix of a probabilistic model, we know that

$$\mathbb{E}_{\tilde{y} \sim p_\theta(\tilde{y}|x)} \frac{\partial \ell(\tilde{y}|x)}{\partial \theta} \frac{\partial \ell(\tilde{y}|x)}{\partial \theta}^\top = \mathbb{E}_{\tilde{y} \sim p_\theta(\tilde{y}|x)} \frac{\partial^2 \ell(\tilde{y}|x)}{\partial \theta^2} \quad (75)$$

Since we have assumed that this Hessian does not depend on \tilde{y} , it is equal to H_k .

Thus TANGO rewrites as averaged SGD with a particular model of noise on the fast parameter. \square

PROOF OF LEMMA 9.

This is a standard martingale argument. By the variance assumption on \hat{F}_k , one has $\mathbb{E}_{k-1} \|\xi_k\|^2 \leq \sigma^2$. Likewise, $\mathbb{E}_{k-1} \|M_k \xi_k\|^2 \leq \Lambda^2 \sigma^2$. This proves the first claim.

Moreover, since $\mathbb{E}_{k-1} \xi_k = 0$ and M_k is \mathcal{F}_{k-1} -measurable, $\mathbb{E}_{k-1} M_k \xi_k = 0$, namely, the $M_k \xi_k$ are martingale increments.

Setting $X_k := \left\| \sum_{j=1}^k M_j \xi_j \right\|^2$, we find $\mathbb{E}_k X_{k+1} = X_k + 2\mathbb{E}_k[(M_{k+1} \xi_{k+1}) \cdot \sum_{j=1}^k M_j \xi_j] + \mathbb{E}_k \|M_{k+1} \xi_{k+1}\|^2 = X_k + \mathbb{E}_k \|M_{k+1} \xi_{k+1}\|^2$.

Consequently, $\mathbb{E} X_n \leq n \Lambda^2 \sigma^2$. Moreover, $\mathbb{E}_k X_{k+1} \geq X_k$, so that X_k is a submartingale. Therefore, by Doob's martingale inequality, with probability $\geq 1 - \varepsilon$,

$$\sup_{0 \leq k \leq n} X_k \leq \frac{\mathbb{E} X_n}{\varepsilon} \leq \frac{n \Lambda^2 \sigma^2}{\varepsilon} \quad (76)$$

Finally, $\sum_{j=k}^n M_j \xi_j = \sum_{j=1}^n M_j \xi_j - \sum_{j=1}^{k-1} M_j \xi_j$, hence the conclusion by the triangle inequality. \square

PROOF OF LEMMA 10.

The argument is similar to the preceding lemma, together with the bound on $\mathbb{E} \|v_k\|^2$ from Lemma 7. Conditionally to \mathcal{F}_{k-1} one has $\mathbb{E}_{k-1} \|\zeta_k\|^2 = \mathbb{E}_{k-1} v_{k-1}^\top (\hat{A}_k - A_k) (\hat{A}_k - A_k) v_{k-1} = \mathbb{E}_{k-1} v_{k-1}^\top \hat{A}_k^2 v_{k-1} - v_{k-1}^\top A_k^2 v_{k-1} \leq R^2 v_{k-1}^\top A_k v_{k-1} \leq R^2 \lambda_{\max} \|v_{k-1}\|^2$. Therefore, $\mathbb{E} \|\zeta_k\|^2 \leq R^2 \lambda_{\max} \mathbb{E} \|v_{k-1}\|^2 \leq 4R^2 \lambda_{\max} \sigma^2 / \lambda^2$ by Lemma 7.

The operators M_k introduce an additional factor Λ^2 . Consequently, $\mathbb{E} \sum_{k=1}^n \|M_k \zeta_k\|^2 \leq 4nR^2 \Lambda^2 \lambda_{\max} \sigma^2 / \lambda^2$.

The rest of the proof is identical to Lemma 9. \square

References

- [ABH16] Naman Agarwal, Brian Bullins, and Elad Hazan. Second order stochastic optimization in linear time. *arXiv preprint arXiv:1602.03943*, 2016.
- [Ama98] Shun-ichi Amari. Natural gradient works efficiently in learning. *Neural Comput.*, 10:251–276, February 1998.
- [BM13] Francis Bach and Eric Moulines. Non-strongly-convex smooth stochastic approximation with convergence rate $o(1/n)$. In *Advances in neural information processing systems*, pages 773–781, 2013.
- [DB15] Alexandre Défossez and Francis Bach. Averaged least-mean-squares: Bias-variance trade-offs and optimal sampling distributions. In *Artificial Intelligence and Statistics*, pages 205–213, 2015.
- [DFB16] Aymeric Dieuleveut, Nicolas Flammarion, and Francis Bach. Harder, better, faster, stronger convergence rates for least-squares regression. *arXiv preprint arXiv:1602.05419*, 2016.
- [DPCB13] Guillaume Desjardins, Razvan Pascanu, Aaron Courville, and Yoshua Bengio. Metric-free natural gradient for joint-training of boltzmann machines. *arXiv preprint arXiv:1301.3545*, 2013.
- [DSP⁺15] Guillaume Desjardins, Karen Simonyan, Razvan Pascanu, et al. Natural neural networks. In *Advances in Neural Information Processing Systems*, pages 2071–2079, 2015.
- [KB17] Prasenjit Karmakar and Shalabh Bhatnagar. Two time-scale stochastic approximation with controlled markov noise and off-policy temporal-difference learning. *Mathematics of Operations Research*, 2017.

- [KP92] Peter E. Kloeden and Eckhard Platen. *Numerical solution of stochastic differential equations*, volume 23 of *Applications of Mathematics (New York)*. Springer-Verlag, Berlin, 1992.
- [LBOM98] Yann Le Cun, Léon Bottou, Genevieve B. Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural Networks, Tricks of the Trade*, Lecture Notes in Computer Science LNCS 1524. Springer Verlag, 1998.
- [LMB07] Nicolas Le Roux, Pierre-Antoine Manzagol, and Yoshua Bengio. Topmoumoute online natural gradient algorithm. In *Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3-6, 2007*, pages 849–856, 2007.
- [Mar10] James Martens. Deep learning via Hessian-free optimization. In Johannes Fürnkranz and Thorsten Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pages 735–742. Omnipress, 2010.
- [Mar14] James Martens. New insights and perspectives on the natural gradient method. *arXiv preprint arXiv:1412.1193*, 2014.
- [MB11] Éric Moulines and Francis R Bach. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, pages 451–459, 2011.
- [MCO16] Gaétan Marceau-Caron and Yann Ollivier. Practical riemannian neural networks. *arXiv preprint arXiv:1602.08007*, 2016.
- [MG15] James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International Conference on Machine Learning*, pages 2408–2417, 2015.
- [MS11] James Martens and Ilya Sutskever. Learning recurrent neural networks with Hessian-free optimization. In *ICML*, pages 1033–1040, 2011.
- [MS12] James Martens and Ilya Sutskever. Training deep and recurrent neural networks with Hessian-free optimization. In Grégoire Montavon, Geneviève B. Orr, and Klaus-Robert Müller, editors, *Neural Networks: Tricks of the Trade*, volume 7700 of *Lecture Notes in Computer Science*, pages 479–535. Springer, 2012.

- [Oll15] Yann Ollivier. Riemannian metrics for neural networks I: feedforward networks. *Information and Inference*, 4(2):108–153, 2015.
- [Oll17] Yann Ollivier. Online natural gradient as a kalman filter. *arXiv preprint arXiv:1703.00209*, 2017.
- [PJ92] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.
- [Rup88] David Ruppert. Efficient estimations from a slowly convergent robbins-monro process. Technical report, Cornell University Operations Research and Industrial Engineering, 1988.
- [Tad04] Vladislav B Tadic. Almost sure convergence of two time-scale stochastic approximation algorithms. In *American Control Conference, 2004. Proceedings of the 2004*, volume 4, pages 3802–3807. IEEE, 2004.