

Training recurrent networks online without backtracking

Yann Ollivier Corentin Tallec Guillaume Charpiat

Abstract

We introduce the “NoBackTrack” algorithm to train the parameters of dynamical systems such as recurrent neural networks. This algorithm works in an online, memoryless setting, thus requiring no backpropagation through time, and is scalable, avoiding the large computational and memory cost of maintaining the full gradient of the current state with respect to the parameters.

The algorithm essentially maintains, at each time, a single search direction in parameter space. The evolution of this search direction is partly stochastic and is constructed in such a way to provide, at every time, an *unbiased* random estimate of the gradient of the loss function with respect to the parameters. Because the gradient estimate is unbiased, on average over time the parameter is updated as it should.

The resulting gradient estimate can then be fed to a lightweight Kalman-like filter to yield an improved algorithm. For recurrent neural networks, the resulting algorithms scale linearly with the number of parameters.

Small-scale experiments confirm the suitability of the approach, showing that the stochastic approximation of the gradient introduced in the algorithm is not detrimental to learning. In particular, the Kalman-like version of NoBackTrack is superior to backpropagation through time (BPTT) when the time span of dependencies in the data is longer than the truncation span for BPTT.

Consider the problem of training the parameters θ of a dynamical system over a variable $h \in \mathbb{R}^n$ subjected to the evolution equation

$$h(t+1) = f(h(t), x(t), \theta) \tag{1}$$

where f is a fixed function of h and of an input signal $x(t)$, depending on parameters θ . The goal is online minimization of a loss function $\sum_t \ell_t(\hat{y}(t), y(t))$ between a desired output $y(t)$ at time t and a prediction¹

$$\hat{y}(t) = Y(h(t), \varphi) \tag{2}$$

¹The prediction \hat{y} may not live in the same set as y . Often, \hat{y} encodes a probability distribution over the possible values of y , and the loss is the logarithmic loss $\ell = -\log p_{\hat{y}}(y)$.

computed from $h(t)$ and additional parameters φ .

A typical example we have in mind is a recurrent neural network, with activities $a_i(t) := \text{sigm}(h_i(t))$ and evolution equation $h_i(t+1) = b_i + \sum_k r_{ki} x_k(t) + \sum_j W_{ji} a_j(t)$, with parameter $\theta = (b_i, r_{ki}, W_{ji})_{i,j,k}$.

If the full target sequence $y(t)_{t \in [0;T]}$ is known in advance, one strategy is to use the backpropagation through time algorithm (BPTT, see e.g. [Jae02]) to compute the gradient of the total loss $L_T := \sum_{t=0}^T \ell_t$ with respect to the parameters θ and φ , and use gradient descent on θ and φ .

However, if the data $y(t+1)$ arrive one at a time in a streaming fashion, backpropagation through time would require making a full backward computation from time $t+1$ to time 0 after each new data point becomes available. This results in an $\Omega(t^2)$ complexity and in the necessity to store past states, inputs, and outputs. A possible strategy is to only backtrack by a finite number of time steps [Jae02] rather than going back all the way to $t=0$. But this provides biased gradient estimates and may impair detection of time dependencies with a longer range than the backtracking time range.

By contrast, methods which are fully online are typically not scalable. One strategy, known as *real-time recurrent learning* (RTRL) in the recurrent network community,² maintains the full gradient of the current state with respect to the parameters:

$$G(t) := \frac{\partial h(t)}{\partial \theta} \quad (3)$$

which satisfies the evolution equation

$$G(t+1) = \frac{\partial f(h(t), x(t), \theta)}{\partial h} G(t) + \frac{\partial f(h(t), x(t), \theta)}{\partial \theta} \quad (4)$$

(by differentiating (1)). Knowing $G(t)$ allows to minimize the loss via a stochastic gradient descent on the parameters θ , namely,³

$$\theta \leftarrow \theta - \eta_t \frac{\partial \ell_t}{\partial \theta}^\top \quad (5)$$

with learning rate η_t . Indeed, the latter quantity can be computed from G_t and from the way the predictions depend on $h(t)$, via the chain rule

$$\frac{\partial \ell_t}{\partial \theta} = \frac{\partial \ell_t(Y(h(t), \varphi), y(t))}{\partial h} G(t) \quad (6)$$

However, the full gradient $G(t)$ is an object of dimension $\dim h \times \dim \theta$. This prevents computing or even storing $G(t)$ for moderately large-dimensional dynamical systems, such as recurrent neural networks.

²This amounts to applying forward automatic differentiation.

³We use the standard convention for Jacobian matrices, namely, $\partial x / \partial y$ is the matrix with entries $\partial x_i / \partial y_j$. Then the chain rule writes $\frac{\partial x}{\partial y} \frac{\partial y}{\partial z} = \frac{\partial x}{\partial z}$. This makes the derivatives $\partial \ell_t / \partial \theta$ into *row* vectors so that gradient descent is $\theta \leftarrow \theta - (\partial \ell_t / \partial \theta)^\top$.

Algorithms using a Kalman filter on θ also⁴ rely on this derivative $\frac{\partial \ell_t}{\partial \theta}$ (see [Hay04, Jae02] for the case of recurrent networks). So any efficient way of estimating this derivative can be fed, in turn, to a Kalman-type algorithm.

Algorithms suggested to train hidden Markov models online (e.g., [Cap11], based on expectation-maximization instead of gradient descent) share the same algebraic structure and suffer from the same problem.

1 The NoBackTrack algorithm

1.1 The rank-one trick: an expectation-preserving reduction

We propose to build an approximation $\tilde{G}(t)$ of $G(t)$ with a more sustainable algorithmic cost; $\tilde{G}(t)$ will be random with the property $\mathbb{E}\tilde{G}(t) = G(t)$ for all t . Then the stochastic gradient (5) based on $\tilde{G}(t)$ will introduce noise, but *no bias*, on the learning of θ : the average change in θ after a large number of time steps will reflect the true gradient direction. (This is true only if the noises on $\tilde{G}(t)$ at different times t are sufficiently decorrelated. This is the case if the dynamical system (1) is sufficiently ergodic.) Such unbiasedness does not hold, for instance, if the gradient estimate is simply projected onto the nearest small-rank or diagonal plus small-rank approximation.⁵

The construction of an unbiased \tilde{G} is based on the following “rank-one trick”.

PROPOSITION 1 (RANK-ONE TRICK). *Given a decomposition of a matrix A as a sum of rank-one outer products, $A = \sum_i v_i w_i^\top$, and independent uniform random signs $\varepsilon_i \in \{-1, 1\}$, then*

$$\tilde{A} := (\sum_i \varepsilon_i v_i) (\sum_j \varepsilon_j w_j)^\top \quad (7)$$

satisfies

$$\mathbb{E}\tilde{A} = \sum_i v_i w_i^\top = A \quad (8)$$

that is, \tilde{A} is an expectation-preserving rank-one approximation of A .

Moreover, one can minimize the variance of \tilde{A} by taking advantage of additional degrees of freedom in this decomposition, namely, one may first replace v_i and w_i with $\rho_i v_i$ and w_i/ρ_i for any $\rho_i \in \mathbb{R}^*$. The choice of ρ_i which yields minimal variance of \tilde{A} is when the norms of v_i and w_i become equal, namely, $\rho_i = \sqrt{\|w_i\|/\|v_i\|}$.

⁴One may use Kalman filtering either on θ alone or on the pair (θ, h) . In the first case, $\frac{\partial \ell_t}{\partial \theta}$ is explicitly needed. In the second case, all the information about how θ influences the current state $h(t)$ is contained in the covariance between θ and h , which the algorithm must maintain, and which is as costly as maintaining $G(t)$ above.

⁵We tried such methods first, with less satisfying results. In practice, consecutive projections tend to interact badly and reduce too much the older contributions to the gradient.

The proof of the first statement is immediate.

The statement about minimizing variance is proven in Appendix A. Minimizing variance thanks to ρ_i is quite important in practice, see Section 1.3.

The rank-one trick also extends to tensors of arbitrary order; this may be useful in more complex situations.⁶

The rank-one reduction \tilde{A} depends, not only on the value of A , but also on the way A is decomposed as a sum of rank-one terms. In the applications to recurrent networks below, there is a natural such choice.⁷

We use this reduction operation at each step of the dynamical system, to build an approximation \tilde{G} of G . A key property is that the evolution equation (4) satisfied by G is *affine*, so that if $\tilde{G}(t)$ is an unbiased estimate of $G(t)$, then $\frac{\partial f(h(t), x(t), \theta)}{\partial h} \tilde{G}(t) + \frac{\partial f(h(t), x(t), \theta)}{\partial \theta}$ is an unbiased estimate of $G(t+1)$.

This leads to the NoBackTrack algorithm (Euclidean version) described in Algorithm 1. At each step, this algorithm maintains an approximation of G as

$$\tilde{G} = \bar{v}\bar{w}^\top + \sum_i e_i w_i^\top \quad (9)$$

where e_i is the i -th basis vector in space h , and $w_i := \frac{\partial f_i}{\partial \theta}^\top$ are sparse vectors.

To understand this structure, say that $\tilde{G}(t-1) = \bar{v}\bar{w}^\top$ is a rank-one unbiased approximation of $G(t-1)$. Then the evolution equation (4) for G yields $\left(\frac{\partial f}{\partial h}\right) (\bar{v}\bar{w}^\top) + \frac{\partial f}{\partial \theta} = \left(\frac{\partial f}{\partial h} \bar{v}\right) \bar{w}^\top + \sum_i e_i \frac{\partial f_i}{\partial \theta}$ as an approximation of $\tilde{G}(t)$. This new approximation is not rank-one any more, but it can be used to perform a gradient step on θ , and then reduced to a rank-one approximation before the next time step.

Note that handling $\frac{\partial f_i}{\partial \theta}$ is usually cheap: in many situations, only a small subset of the parameter θ directly influences each component $h_i(t+1)$ given $h(t)$, so that for each component i of the state space, $\frac{\partial f_i}{\partial \theta}$ has few non-zero

⁶The most symmetric way to do this is to use complex roots of unity, for instance, $\sum_i u_i \otimes v_i \otimes w_i = \mathbb{E} \operatorname{Re} \left((\sum_i \zeta_i u_i) \otimes (\sum_j \zeta_j v_j) \otimes (\sum_k \zeta_k w_k) \right)$ where each ζ_i is taken independently at random among $\{1, e^{\pm 2i\pi/3}\}$. This involves complex numbers but there is no need to complexify the original dynamical system (1). Another, complex-free possibility is to apply the rank-one trick recursively to tensors of smaller order, for instance, $\sum_i u_i \otimes v_i \otimes w_i \otimes x_i = \sum_i (u_i \otimes v_i) \otimes (w_i \otimes x_i) = \mathbb{E} \left[(\sum_i \varepsilon_i u_i \otimes v_i) \otimes (\sum_j \varepsilon_j w_j \otimes x_j) \right]$ and then apply *independent* rank-one decompositions in turn to $\sum_i \varepsilon_i u_i \otimes v_i$ and to $\sum_j \varepsilon_j w_j \otimes x_j$.

⁷The rank-one trick may also be performed using random Gaussian vectors, namely $A = \mathbb{E}[\xi(\xi^\top \Sigma^{-1} A)]$ with $\xi = \mathcal{N}(0, \Sigma)$. This version does not depend on a chosen decomposition of A , but depends on a choice of Σ . Variance can be much larger in this case: for instance, if $A = vv^\top$ is actually rank-one, then $(\varepsilon v)(\varepsilon w^\top) = \varepsilon v \varepsilon w^\top$ so that the rank-one trick with random signs is exact, whereas the Gaussian version yields $(\xi \xi^\top \Sigma^{-1}) v v^\top$ which is correct only in expectation. This case is particularly relevant because we are going to apply a reduction at each time step, thus working on objects that stay close to rank-one. The generalization to tensors is also more cumbersome in the Gaussian case.

components. For instance, for a recurrent neural network with activities $a_i(t) := \text{sigm}(h_i(t))$ and evolution equation $h_i(t+1) = b_i + \sum_k r_{ik}x_k(t) + \sum_j W_{ji}a_j(t)$, the derivative of $h_i(t+1)$ with respect to the parameter $\theta = (b, r, W)$ only involves the parameters b_i, r_{ik}, W_{ji} of unit i . In such situations, the total cost of computing and storing all the w_i 's is of the same order as the cost of computing $h(t+1)$ itself. See Section 1.3 for details on this example.

After the reduction step of Algorithm 1, \bar{w} may be interpreted as a “search direction” in parameter space θ , while \bar{v} is an estimate of the effect on the current state $h(t)$ of changing θ in the direction \bar{w} . The search direction \bar{w} evolves stochastically, but not fully at random, over time, so that on average $\bar{v}\bar{w}^\top$ is a fair estimate of the actual influence of the parameter θ .

Note that in Algorithm 1, the non-recurrent output parameters φ are trained according to their exact gradient. The rank-one trick is used only for the recurrent part of the system.

By construction, at each step of Algorithm 1, the quantity $\tilde{G}_t := \bar{v}\bar{w}^\top + \sum_i e_i w_i^\top$ satisfies $\mathbb{E}\tilde{G}_t = \frac{\partial h(t)}{\partial \theta}$. However, since the value of θ changes along the algorithm, we must be careful about the meaning of this statement. Intuitively, this derivative with respect to θ is taken along the actual trajectory of parameters θ_t realized by the algorithm.

More formally, let $\boldsymbol{\theta} = (\theta_0, \dots, \theta_t, \dots)$ be any sequence of parameters. Let f be any function depending on this sequence $\boldsymbol{\theta}$, such as the state of the system at time t (all functions considered below will depend only on a finite initial segment of $\boldsymbol{\theta}$). Define $\boldsymbol{\theta} + \varepsilon := (\theta_0 + \varepsilon, \dots, \theta_t + \varepsilon, \dots)$ and say that f has derivative $\frac{\partial f}{\partial \theta}$ with respect to $\boldsymbol{\theta}$ if $f(\boldsymbol{\theta} + \varepsilon) = f(\boldsymbol{\theta}) + \varepsilon \frac{\partial f}{\partial \theta} + O(\varepsilon^2)$ for small ε .

Thanks to this convention, the evolution equation (4) for the evolution of $G(t)$ holds for any sequence of parameters $\boldsymbol{\theta}$, with $G(t)$ defined as $\frac{\partial h(t)}{\partial \theta}$. The following statement is then easily proved by induction.

PROPOSITION 2 (UNBIASED RANK-ONE GRADIENT ESTIMATE FOR DYNAMICAL SYSTEMS). *At each time step t , the quantity $\tilde{G}_t := \bar{v}\bar{w}^\top + \sum_i e_i w_i^\top$ from Algorithm 1 is an unbiased estimate of the gradient of the state of the system with respect to the parameter:*

$$\mathbb{E}\tilde{G}_t = \frac{\partial h(t)}{\partial \theta} \tag{19}$$

where $\boldsymbol{\theta}$ is the sequence of parameters produced by the algorithm.

In particular, for learning rates η tending to 0, the parameter evolves slowly so that the derivative $\frac{\partial h(t)}{\partial \theta}$ is close to a derivative with respect to the current value θ_t of the parameter. Thus, in this regime, $\frac{\partial h(t)}{\partial \theta}$ tends to $\frac{\partial h(t)}{\partial \theta_t}$, and since \tilde{G}_t is an unbiased estimate of G_t , the situation gets closer and

Parameters: $h(0)$ (initial state), θ_0, φ_0 (initial value of the internal and output parameters), η_t (learning rate scheme);

Data: $x(t)$ (input signal), $y(t)$ (output signal);

Maintains: $h(t)$ (current state), θ, φ (internal and output parameters), \bar{v} (column vector of size $\dim h$), \bar{w} (column vector of size $\dim \theta$), w_i (sparse column vectors of size $\dim \theta$) for $i = 1, \dots, \dim h$.

Initialization: $\theta \leftarrow \theta_0, \varphi \leftarrow \varphi_0, \bar{v} \leftarrow 0, \bar{w} \leftarrow 0, w_i \leftarrow 0$;

for $t = 0$ **to** *end-of-time* **do**

Observation step: Compute prediction $\hat{y}(t) = Y(h(t), \varphi)$ from current state $h(t)$.

Observe $y(t)$ and incur loss $\ell_t(\hat{y}(t), y(t))$.

Update step: Compute derivative of loss with respect to output parameters, $\frac{\partial \ell_t}{\partial \varphi} = \frac{\partial \ell_t(Y(h(t), \varphi), y(t))}{\partial \varphi}$, and update output parameters:

$$\varphi \leftarrow \varphi - \eta_t \frac{\partial \ell_t^\top}{\partial \varphi} \quad (10)$$

Compute derivative of loss with respect to current state,

$$H \leftarrow \frac{\partial \ell_t(Y(h(t), \varphi), y(t))}{\partial h} \quad (11)$$

Update internal parameters θ :

$$\theta \leftarrow \theta - \eta_t (Hv)\bar{w} - \eta_t \sum_i H_i w_i \quad (12)$$

(this is a gradient step $\theta \leftarrow \theta - \eta_t (H\tilde{G})^\top$ using the current gradient estimate \tilde{G} from (9)).

Reduction step: Draw independent uniform random signs $\varepsilon_i = \pm 1$. Let e_i be the i -th basis vector in state space. Compute $\bar{\rho} := \sqrt{\|\bar{w}\| / \|\bar{v}\|}$ and $\rho_i := \sqrt{\|w_i\| / \|e_i\|}$ for each i . Update

$$\bar{v} \leftarrow \bar{\rho} \bar{v} + \sum_i \varepsilon_i \rho_i e_i \quad (13)$$

$$\bar{w} \leftarrow \bar{w} / \bar{\rho} + \sum_i \varepsilon_i w_i / \rho_i \quad (14)$$

$$w_i \leftarrow 0 \quad (15)$$

Transition step: Observe new value of input signal $x(t)$ and compute next state $h(t+1) = f(h(t), x(t), \theta)$. Update estimate \tilde{G} :

$$\bar{v} \leftarrow \frac{\partial f(h(t), x(t), \theta)}{\partial h} \bar{v} \quad (16)$$

$$w_i \leftarrow \frac{\partial f_i(h(t), x(t), \theta)}{\partial \theta}^\top \quad (17)$$

$$t \leftarrow t + 1 \quad (18)$$

end

Algorithm 1: NoBackTrack algorithm, Euclidean version.

closer to an ordinary stochastic gradient descent if η is small. Presumably this happens whenever the learning rate is small enough for θ not to change too much within a time range corresponding to a “forgetting time” of the dynamical system, although more work is needed here.

1.2 Feeding the gradient estimate to an extended Kalman filter

The Euclidean version of the NoBackTrack algorithm presented in Algorithm 1 is not enough to obtain good performance fast. Online estimation often yields best results when using filters from the Kalman family. We refer to [Hay04, Jae02] for a discussion of Kalman filtering applied to recurrent neural networks.

Kalman-based approaches rely on a covariance matrix estimate $P(t)$ on θ . After observing $y(t)$, the parameter θ gets adjusted via⁸

$$\theta \leftarrow \theta - P(t) \frac{\partial \ell_t^\top}{\partial \theta} \quad (20)$$

where the derivative of the loss with respect to θ is computed, as above, via the product of the derivative of the loss with respect to the current state $h(t)$, and the derivative $G(t) = \frac{\partial h(t)}{\partial \theta}$.

Maintaining a full covariance matrix on θ is usually too costly. However, having a good approximation of $P(t)$ is not as critical as having a good approximation of $\frac{\partial \ell_t}{\partial \theta}$. Indeed, given an unbiased approximation of $\frac{\partial \ell_t}{\partial \theta}$, *any* symmetric positive definite matrix $P(t)$ which changes slowly enough in time will yield an unbiased trajectory for θ .

Thus, we will use more aggressive matrix reduction techniques on $P(t)$, such as block-diagonal (as in [Hay04]) or quasi-diagonal [Oll15a] approximations. In our setting, the main point of using the covariance matrix is to get both a sensible scaling of the learning rate for each component of θ , and reparametrization-invariance properties [Oll15a].

In Kalman filtering, in the case when the “true” underlying parameter θ in the extended Kalman filter is constant, it is better to work with the inverse covariance matrix $J(t) := P(t)^{-1}$, and the extended Kalman filter on θ can be rewritten as

$$J(t) \leftarrow J(t-1) + \frac{\partial \hat{y}_t^\top}{\partial \theta} I_t \frac{\partial \hat{y}_t}{\partial \theta} \quad (21)$$

$$\theta \leftarrow \theta - J(t)^{-1} \frac{\partial \ell_t^\top}{\partial \theta} \quad (22)$$

⁸Indeed, in standard Kalman filter notation, one has $K_t R = P_t H_t^\top$, so that for the quadratic loss $\ell = \frac{1}{2}(\hat{y} - y)^\top R^{-1}(\hat{y} - y)$ (log-loss of a Gaussian model with covariance matrix R), the Kalman update for θ is equivalent to $\theta \leftarrow \theta - P(t) \frac{\partial \ell_t^\top}{\partial \theta}$.

where \hat{y}_t is the prediction at time t , where both $\frac{\partial \hat{y}_t}{\partial \theta}$ and $\frac{\partial \ell_t}{\partial \theta}$ can be computed from $h(t)$ via the chain rule if $G(t) = \frac{\partial h(t)}{\partial \theta}$ is known, and where I_t is the Fisher information matrix of \hat{y}_t as a probability distribution on y_t . (For exponential families this is just the Hessian $-\frac{\partial^2 \ell_t}{\partial \hat{y}_t^2}$ of the loss with respect to the prediction). This is the so-called *information filter*, because $J(t)$ approximates the Fisher information matrix on θ given the observations up to time t . This is basically a natural gradient descent on θ .

This approach is summarized in Algorithm 2, which we describe more loosely since matrix approximation schemes may depend on the application.

Algorithm 2 uses a decay factor $(1 - \gamma_t)$ on the inverse covariance matrices to limit the influence of old computations made with outdated values of θ . The factor γ_t also controls the effective learning rate of the algorithm, since, in line with Kalman filtering, we have not included a learning rate for the update of θ (namely, $\eta_t = 1$): the step size is adapted via the magnitude of J . For $\gamma_t = 0$, J grows linearly so that step size is $O(1/t)$.

Moreover, we have included a regularization term Λ for matrix inversion; in the Bayesian interpretation of Kalman filtering this corresponds to having a Gaussian prior on the parameters with inverse covariance matrix Λ . This is important to avoid fast divergence in the very first steps.

In practice we have used $\gamma_t = O(1/\sqrt{t})$ and $\Lambda = (\dim h) \cdot \text{Id}$.

The simplest and fastest way to approximate the Fisher matrix in Algorithm 2 is the outer product approximation (see discussion in [Oll15a]), which we have used in the experiments below. Namely, we simply use $I_t \leftarrow \frac{\partial \ell_t}{\partial \hat{y}_t} \frac{\partial \ell_t}{\partial \hat{y}_t}^\top$ so that the updates to J_φ and J_θ simplify and become rank-one outer product updates using the gradient of the loss, namely, $J_\theta \leftarrow (1 - \gamma_t)J_\theta + \frac{\partial \ell_t}{\partial \theta} \frac{\partial \ell_t}{\partial \theta}^\top$ and likewise for φ . Here the derivative $\frac{\partial \ell_t}{\partial \theta}$ is estimated from the current gradient estimate \tilde{G} .

For the matrix reductions, we have used a block-wise quasi-diagonal reduction as in [Oll15a]. This makes the cost of handling the various matrices linear in the number of parameters.

1.3 Examples

Let us show how Algorithm 1 works out on explicit examples.

The importance of norm rescaling. Let us first consider a simple dynamical system which illustrates the importance of rescaling the norms by $\bar{\rho}$ and ρ_i . Let $0 < \alpha < 1$ and consider the system

$$h(t+1) = (1 - \alpha)h(t) + \theta \tag{28}$$

with both h and θ in \mathbb{R}^n . This quickly converges towards θ/α . We have $\partial f/\partial h = (1 - \alpha)\text{Id}$ and $\partial f/\partial \theta = \text{Id}$ and so $\partial f_i/\partial \theta^\top = e_i$, the i -th basis

Parameters: $h(0)$ (initial state), θ_0, φ_0 (initial value of the parameters), $0 \leq \gamma_t < 1$ (covariance decay parameter scheme), Λ_φ and Λ_θ (inverse covariance matrix of the prior on the parameters);

Maintains: Same as Algorithm 1, plus a representation of matrices J_θ and J_φ allowing for efficient inversion;

Subroutines: A matrix reduction method `MatrixReduce(M)` which only evaluates a small number of entries of its argument M and returns an approximation of M that can be inverted efficiently;
A routine `FisherApprox(\hat{y}_t, y_t)` which returns either a positive definite approximation of the Fisher information matrix of \hat{y}_t as a probability distribution on y_t , or a positive definite approximation of the Hessian $-\frac{\partial^2 \ell_t}{\partial \hat{y}_t^2}$ of the loss with respect to the prediction.

Initialization: as in Algorithm 1, and $J_\theta \leftarrow 0, J_\varphi \leftarrow 0$;

for $t = 0$ **to** *end-of-time* **do**

Observation step: as in Algorithm 1.

Update step: Compute approximate Fisher information matrix w.r.t. \hat{y}_t :

$$I_t \leftarrow \text{FisherApprox}(\hat{y}_t, y_t) \quad (23)$$

 Compute derivative of prediction and of loss with respect to output parameters, $\frac{\partial \hat{y}_t}{\partial \varphi}$ and $\frac{\partial \ell_t}{\partial \varphi}$. Update inverse covariance matrix of output parameters φ :

$$J_\varphi \leftarrow (1 - \gamma_t)J_\varphi + \text{MatrixReduce} \left(\begin{array}{c} \frac{\partial \hat{y}_t^\top}{\partial \varphi} \quad I_t \quad \frac{\partial \hat{y}_t}{\partial \varphi} \end{array} \right) \quad (24)$$

 and update output parameters:

$$\varphi \leftarrow \varphi - (J_\varphi + \Lambda_\varphi)^{-1} \frac{\partial \ell_t^\top}{\partial \varphi} \quad (25)$$

 Compute derivative $\frac{\partial \hat{y}_t}{\partial h}$ of prediction with respect to current state $h(t)$. Update inverse covariance matrix of internal parameters θ :

$$J_\theta \leftarrow (1 - \gamma_t)J_\theta + \text{MatrixReduce} \left(\begin{array}{c} \tilde{G}^\top \quad \frac{\partial \hat{y}_t^\top}{\partial h} \quad I_t \quad \frac{\partial \hat{y}_t}{\partial h} \quad \tilde{G} \end{array} \right) \quad (26)$$

 and update internal parameters θ :

$$\theta \leftarrow \theta - (J_\theta + \Lambda_\theta)^{-1} \delta\theta \quad (27)$$

 where $\delta\theta := (Hv)\bar{w} - \sum_i H_i w_i$ is the update of θ from Algorithm 1.

Reduction step: Same as in Algorithm 1, but the norms used to compute $\bar{\rho}$ and ρ_i are derived from J_θ^{-1} (cf. Appendix B).

Transition step: Same as in Algorithm 1.

end

Algorithm 2: NoBackTrack algorithm, Kalman version.

vector. Then the reduction and transition steps in Algorithm 1, if the scalings ρ are not used, amount to

$$\bar{v}_{t+1} = (1 - \alpha) (\bar{v}_t + \sum_i \varepsilon_i(t) e_i) \quad (29)$$

$$\bar{w}_{t+1} = \bar{w}_t + \sum_i \varepsilon_i(t) e_i \quad (30)$$

with the $\varepsilon_i(t)$ independent at each step t . The resulting estimate of $\partial h(t)/\partial \theta$ is unbiased, but its variance grows linearly with time. Indeed, the dynamics of \bar{v}_t is stationary thanks to the factor $(1 - \alpha)$, but the dynamics of \bar{w}_t is purely additive so that w_t is just a d -dimensional random walk. On the other hand, if rescaling by ρ is used, then both \bar{v} and \bar{w} get rescaled by $\sqrt{1 - \alpha}$ at each step,⁹ so that their dynamics becomes stationary and variance does not grow.

Recurrent neural networks. The next example is a standard recurrent neural network (RNN). The state of the system is the set of pre-activation values $h_i(t)$, and the activities are $a_i(t) := \sigma(h_i(t))$ where σ is some activation function such as tanh or sigmoid. The recurrent dynamics of h is

$$h_i(t+1) = \sum_{j \rightarrow i} W_{ji} \sigma(h_j(t)) + \sum_l r_{li} x_l(t) \quad (31)$$

in which $h(t), h(t+1) \in \mathbb{R}^n$, $(W_{ji})_{j \rightarrow i}$ are a set of weights defining a graph on n nodes, and $(r_{li})_{(i,l)}$ are the input weights.¹⁰ The parameter is $\theta = (W, r)$. We hereby omit the output part of the network,¹¹ as it is of no use to analyze the estimation of $\partial h(t)/\partial \theta$.

(We have chosen the pre-activation values h , rather than the activities $a = \sigma(h)$, as the state of the system. This results in simpler expressions, especially for the input weights r .)

Thus, the function f defining the dynamical system for the variable h is (31). The derivatives of f are immediately computed as $\partial f_i / \partial W_{ji} = \sigma(h_j)$, $\partial f_i / \partial r_{li} = x_l$, $\partial f_i / \partial h_j = W_{ji} \sigma'(h_j)$, and all other derivatives are 0.

Algorithm 1 maintains, after the reduction step, an approximation $\frac{\partial h(t)}{\partial \theta} \approx \bar{v}(t) \bar{w}(t)^\top$. We can decompose $\bar{w}(t) = (\bar{W}(t), \bar{r}(t))$ into the components corresponding to the internal and input weights of the parameter $\theta = (W, r)$, so that

$$\frac{\partial h_i(t)}{\partial W_{kj}} \approx \bar{v}_i(t) \bar{W}_{kj}(t) \quad (32)$$

$$\frac{\partial h_i(t)}{\partial r_{lj}} \approx \bar{v}_i(t) \bar{r}_{lj}(t). \quad (33)$$

⁹Proof: By induction one has $\bar{v} = \bar{w}$ after the reduction step and $\bar{v} = (1 - \alpha) \bar{w}$ after the transition step, and $\bar{\rho} = 1/\sqrt{1 - \alpha}$.

¹⁰Biases are omitted; they can be treated by the inclusion of an always-activated unit i_0 with $a_{i_0}(t) \equiv 1$.

¹¹The experiments below use a softmax output with output parameters φ , see Section 2.

By plugging the values of the partial derivatives of f into Algorithm 1, we find the following update equations for the value of \bar{v} , \bar{W} and \bar{r} right after the reduction step:

$$\bar{v}_i(t+1) = \bar{\rho} \sum_{j \rightarrow i} W_{ji} \sigma'(h_j(t)) \bar{v}_j(t) + \varepsilon_i \rho_i \quad (34)$$

$$\bar{W}_{kj}(t+1) = \frac{\bar{W}_{kj}(t)}{\bar{\rho}} + \varepsilon_j \frac{\sigma(h_k(t))}{\rho_j} \quad (35)$$

$$\bar{r}_{lj}(t+1) = \frac{\bar{r}_{lj}(t)}{\bar{\rho}} + \varepsilon_j \frac{x_l(t)}{\rho_j} \quad (36)$$

where the ε_j are independent symmetric binary random variables, taking values ± 1 with probability $\frac{1}{2}$. Any non-zero choice of ρ_j leads to an unbiased estimation, though the values are to be optimized as mentioned above.

Applying this update has the same algorithmic cost as implementing one step (31) of the recurrent network itself.

Leaky recurrent neural networks. To capture long-term dependencies, in the experiments below we also use a *leaky* RNN, obtained via the addition of a direct feedback term:

$$h_i(t+1) = \alpha_i h_i(t) + \sum_l r_{li} x_l(t) + \sum_j W_{ji} a_j(t) \quad a_j(t) := \sigma(h_j(t)) \quad (37)$$

with $\alpha_i \in [0; 1]$ for all i . (See [Jae02] for similar models.) This feedback term reduces the impact of the vanishing gradient issue and keeps a longer memory of past inputs.

This only changes the derivative of f_i with respect to h_j , which becomes $\partial f_i / \partial h_j = W_{ji} \sigma'(h_j) + \alpha_i \delta_{ij}$. Consequently the update rules (35)–(36) for \bar{W} and \bar{r} are unchanged, while the update of \bar{v} becomes

$$\bar{v}_i(t+1) = \bar{\rho} \alpha_i \bar{v}_i(t) + \bar{\rho} \sum_{j \rightarrow i} W_{ji} \sigma'(h_j(t)) \bar{v}_j(t) + \varepsilon_i \rho_i \quad (38)$$

Multilayer recurrent neural networks. Let us now treat the case of a multilayer recurrent neural network with dynamics

$$h^{(1)}(t+1) = f^{(1)}(x(t), h^{(1)}(t), \theta_1) \quad (39)$$

$$h^{(2)}(t+1) = f^{(2)}(x(t), h^{(1)}(t+1), h^{(2)}(t), \theta_2) \quad (40)$$

$$\vdots \quad (41)$$

$$h^{(n)}(t+1) = f^{(n)}(x(t), h^{(n-1)}(t+1), h^{(n)}(t), \theta_n) \quad (42)$$

where each layer $h^{(i)}$ and $f^{(i)}$ define an RNN as in (31) above. Directly applying the rank-one approximation to the function $f = (f^{(1)}, f^{(2)}, \dots, f^{(n)})$

would be cumbersome: since the activity of a neuron of the i -th layer at time $t + 1$ depends on all parameters from the previous $i - 1$ layers, the derivative $\partial f / \partial \theta$ is not sparse.

To cope with this, a natural approach is to treat the dynamics in a “rolling” fashion and apply the rank-one approximation at each layer in turn. Formally, this amounts to defining the following model

$$\tilde{h}^{(1)}(nt + 1) = f^{(1)}(\tilde{x}(nt), \tilde{h}^{(1)}(nt), \theta_1) \quad (43)$$

$$\tilde{h}^{(2)}(nt + 2) = f^{(2)}(\tilde{x}(nt + 1), \tilde{h}^{(2)}(nt + 1), \theta_2) \quad (44)$$

$$\vdots \quad (45)$$

$$\tilde{h}^{(n)}(nt + n) = f^{(n)}(\tilde{x}(nt + n - 1), \tilde{h}^{(n)}(nt + n - 1), \theta_n) \quad (46)$$

with $\tilde{x}(t) := x(\lfloor t/n \rfloor)$, and where states not explicitly appearing in these equations stay unchanged ($h^{(i)}(nt + j) = h^{(i)}(nt + j - 1)$ for $i \neq j$). Thus, the transition function explicitly depends on time (more precisely, on time modulo the number of layers), and is sparse at each step. Indeed, at each step, applying the transition function amounts to applying one of the $f^{(i)}$ to the corresponding layer, and leaving the other layers unchanged. Thus the derivative of $f^{(i)}$ with respect to any θ_j , $j \neq i$, is zero; this leaves only the gradient of $f^{(i)}$ wrt θ_i to be dealt with, and Algorithm 1 or 2 can be applied at little cost.

1.4 Extensions

Rank- K reductions. A first obvious extension is to use higher-rank reductions. The simplest way to achieve this is to take several independent random rank-one $\tilde{v}_k \tilde{w}_k^\top$ reductions in (7) and average them. Note that w_i (Algorithm 1) has to be evaluated only once in this case. It might be slightly more efficient to first split the parameter components into K blocks (e.g., at random) so that the k -th term \tilde{w}_k^\top only involves parameters from the k -th block: indeed, applying the evolution equation for G preserves this structure so this requires less memory for storage of the \tilde{w}_k .

Algorithms similar to RTRL. Other algorithms have been proposed that have the same structure and shortcomings as real-time recurrent learning, for instance, the online EM algorithm for hidden Markov models from [Cap11]. In principle, the approach presented here can be extended to such settings.

Continuous-time systems. Another extension concerns continuous-time dynamical systems

$$\frac{dh(t)}{dt} = F(h(t), x(t), \theta) \quad (47)$$

which can be discretized as $h(t + \delta t) = h(t) + \delta t F(h(t), x(t), \theta)$. Thus this is analogous to the discrete-time case via $f = \text{Id} + \delta t F$, and Algorithm 1 may be applied to this discretization.

When performing the rank-one reduction (7), the scaling by $\rho_i = \sqrt{\|w_i\| / \|v_i\|}$ is important in this case: it ensures that both \bar{v} and \bar{w} change by $O(\sqrt{\delta t})$ times a random quantity at each step. This is the expected correct scaling for a continuous-time stochastic evolution equation, corresponding to the increment of a Wiener process during a time interval δt . (Without scaling by ρ_i , there will be no well-defined limit as $\delta t \rightarrow 0$, because \bar{v} would change by $O(1)$ at each step $t \leftarrow t + \delta t$, while \bar{w} would evolve by δt times a centered random quantity so that it would be constant in the limit.) Further work is needed to study this continuous-time limit.

2 Experiments

We report here a series of small-scale experiments on text prediction tasks. The experiments focus on two questions: First, does learning using the rank-one approximation \tilde{G} accurately reflect learning based on the actual gradient G computed exactly via RTRL, or is the noise introduced in this method detrimental to learning? Second, how does this approach compare to truncated backpropagation through time?

We used the RNN or leaky RNN models described above to predict a sequence of characters $y(t)$ in a finite alphabet \mathcal{A} , given the past observations $x(s) = y(s)$ for $1 \leq s \leq t - 1$. At each time, the network outputs a probability distribution on the next character z ; explicitly, the output at time t is $\hat{y}(t) \in \mathbb{R}^{\mathcal{A}}$ defined by

$$\hat{y}(t)_z := \varphi_z + \sum_i \varphi_{iz} a_i(t) \quad (48)$$

for each $z \in \mathcal{A}$, with parameters $\varphi = (\varphi_z, \varphi_{iz})_{i,z}$. The output $\hat{y} = (\hat{y}_y)_{y \in \mathcal{A}}$ defines a probability distribution on \mathcal{A} via a softmax $p_{\hat{y}}(y) := \frac{e^{\hat{y}_y}}{\sum_{z \in \mathcal{A}} e^{\hat{y}_z}}$, and the loss function is the log-loss on prediction of the next character, $\ell_t := -\log_2 p_{\hat{y}(t)}(y(t))$. The internal and output parameters θ and φ are trained according to Algorithms 1 and 2.

We used three datasets. The first is a “text” representing synthetic music notation with several syntactic, rhythmic and harmonic constraints (Example 3 from [Oll15b]). The data was a file of length $\approx 10^5$ characters, after which the signal cycled over the same file. The second dataset is the classical $a^n b^n$ example, synthesized by repeatedly picking an integer n at random in some interval, then outputting a series of n a ’s followed by a line break, then n b ’s and another line break. This model tests the ability of a learning algorithm to learn precise timing and time dependencies. The

third example is the full set of Shakespeare’s works, obtained from Project Gutenberg.¹² The file is roughly 5.10^6 characters long.

The benchmarks included are `gzip`, a standard non-online compression algorithm, and context tree weighting (CTW) [BEYY04], a more advanced online text compression algorithm, as well as the actual entropy rate of the generative model for synthetic music and $a^n b^n$.

The code used in the experiments is available at http://www.yann-ollivier.org/rech/code/nobacktrack/code_nobacktrack_exp.tar.gz

Euclidean NoBackTrack. We first study whether the low rank approximation in the *Euclidean* version of NoBackTrack impacts the gradient descent. For this first set of experiments, we use a fully connected RNN with 20 units, as described above, on the synthetic music example. We compared RTRL, Euclidean rank-one NoBackTrack, and Euclidean NoBackTrack using rank-two and rank-ten reductions (obtained by averaging two or ten independent rank-one reductions, as discussed in Section 1.4).

The results are summed up in Figure 1 and Figure 2. All the models were trained using the same learning rate $\eta_t = 1/\sqrt{t}$ for Figure 1 and $\eta_t = 0.03/\sqrt{t}$ for Figure 2.

The various algorithms were run for the same amount of time. This is reflected in the different curve lengths for the different algorithms; in particular, the curve for RTRL is much shorter, reflecting its higher computational cost. (Note the log scale on the t axis: RTRL is roughly 20 times slower with 20 units.)

The impact of stochasticity of the low-rank approximation when using large learning rates is highlighted on Figure 1: Euclidean NoBackTrack with a large learning rate displays instabilities, even when increasing the rank of the approximation.

Smaller learning rates allow the algorithm to cope with this, as the noise in the gradients is averaged out over longer time spans. This is illustrated in Figure 2, in which the trajectories of Euclidean NoBackTrack track those of RTRL closely even with a rank-two approximation.

Kalman NoBackTrack. Next, we report the results of the Kalman version of NoBackTrack on the same experimental setup. A quasi-diagonal outer product (QDOP) approximation [Oll15a] of the full Kalman inverse covariance matrix is used, to keep complexity low.

We compare the low-rank approximations to RTRL. To make the comparison clear, for RTRL we also use a quasi-diagonal (QDOP) approximation of the Kalman filtering algorithm on top of the exact gradient computed by RTRL.

¹²www.gutenberg.org

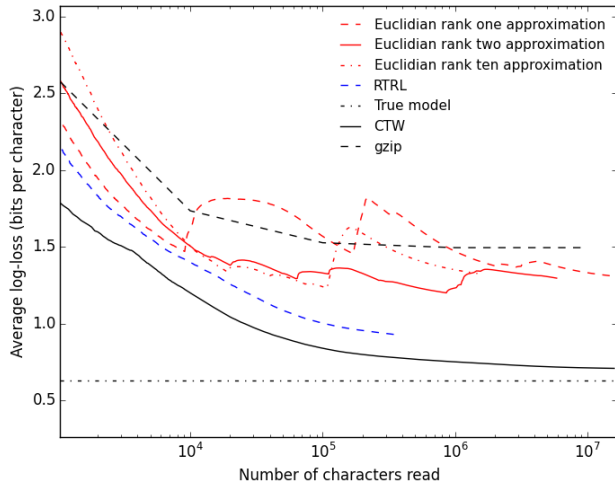


Figure 1: Average log-loss (bits per character) on synthetic music as a function of the number of characters read, for an RNN with 20 units trained with the Euclidean version of the NoBackTrack algorithm for different rank values and RTRL, with learning rate $\eta_t = 1/\sqrt{t}$, benchmarked against the true model entropy rate, `gzip`, and CTW.

Learning rates were set to 1 and all algorithms were run for the same amount of time.

The use of the QDOP-approximated Kalman inverse covariance appears to fully fix the unstable behaviour. Overall, low-rank approximations appear to be roughly on par with QDOP RTRL. There is no obvious gain, on this particular example, in using higher-rank approximations.

Still, on this particular task and with this particular network size, none of the RNN algorithms (including BPTT reported below) match the performance of Context Tree Weighting. RNNs beat CTW on this task if trained using a non-online, Riemannian gradient descent [Oll15b] (analogous to using the Kalman inverse covariance). So this is arguably an effect of imperfect online RNN training.

Kalman NoBackTrack and truncated BPTT. Our next set of experiments aims at comparing Kalman NoBackTrack to truncated BPTT, with truncation¹³ parameter $T = 15$. As BPTT truncates the full gradient by

¹³In the version of BPTT used here, the algorithm does not backtrack by T steps at every time step t ; rather, it waits for T steps between t and $t + T$, then backtracks by T steps and collects all gradients in this interval. Otherwise, truncated BPTT would be T times slower, which was unacceptable for our experiments.




Figure 2: Average log-loss (bits per character) on synthetic music as a function of the number of characters read, for an RNN with 20 units trained with the Euclidean version of the NoBackTrack algorithm for different rank values, and with RTRL, with learning rate $\eta_t = 0.03/\sqrt{t}$, benchmarked against the true model entropy rate, `gzip`, and `CTW`.

removing dependencies at distances longer than the truncation parameter, we expect Kalman NoBackTrack to learn better models on datasets presenting long term correlations.

The two algorithms are first compared on the synthetic music dataset, with the same experimental setup as above, for the same amount of time, with a learning rate $\eta_t = 1/\sqrt{t}$ for truncated BPTT and $\gamma_t = 1/\sqrt{t}$ for Kalman NoBackTrack.¹⁴ The results are shown in Figure 4.

On this example, truncated BPTT performs better than Kalman NoBackTrack, even though the two algorithms display broadly comparable performance. Noticeably, RTRL and truncated BPTT are roughly on par here, with truncated BPTT slightly outperforming RTRL in the end: apparently, maintaining long term dependencies in gradient calculations does not improve learning in this synthetic music example.

Next, to compare NoBackTrack and truncated BPTT on their specific ability to learn precise middle and long term dependencies, we present experiments on the $a^n b^n$ example. This will clearly illustrate the biased nature of the gradients computed by truncated BPTT.

¹⁴These learning rates have different meanings for Kalman NoBackTrack and truncated BPTT, and are not directly comparable.




Figure 3: Average log-loss (bits per character) on synthetic music as a function of the number of characters read, for an RNN with 20 units trained with the Kalman/QDOP version of the NoBackTrack algorithm for different rank values and Kalman/QDOP RTRL, benchmarked against the true model entropy rate, `gzip`, and CTW.

The $a^n b_{[k;l]}^n$ dataset is synthesized by sequentially picking a number n between k and l uniformly at random, then outputting a series of n a 's followed by a line break, then n b 's and another line break. The true entropy rate is $\frac{\log_2(l-k+1)}{l+k+2}$ in this example.¹⁵ A roughly 10^6 character long input sequence was synthesized, using $[k;l] = [1;32]$.

As standard RNN models do not seem to be able to deal with this example, whatever the training algorithm, we used a leaky RNN¹⁶ as presented in Section 1.3, again with 20 fully connected units. All the algorithms used a learning rate of $1/\sqrt{t}$. The results are reported on Figure 5, which also includes the entropy rate of the exact $a^n b^n$ model and the (twice larger) entropy rate of an $a^n b^p$ model with independent n and p .

Kalman NoBackTrack clearly outperforms truncated BPTT on this dataset. This was to be expected, as the typical time range of the temporal dependencies exceeds the truncation range for BPTT, so that the

¹⁵Indeed, $\log_2(l-k+1)$ bits are needed to encode the value of n in each $a^n b^n$ block (this is the entropy of a uniform distribution on $\{k, \dots, l\}$), and the average value of n is $(k+l)/2$ so that the average length of an $a^n b^n$ block, including the two newline symbols, is $2 \times (l+k)/2 + 2$.

¹⁶The parameter α of the LRNN can be learned, but this sometimes produces numerical instabilities unless cumbersome changes of variables are introduced. We just initialized α to a random value separately for each unit and kept it fixed.




Figure 4: Average log-loss (bits per character) on synthetic music as a function of the number of characters read, for an RNN with 20 units trained with the Kalman/QDOP version of the NoBackTrack algorithm for different rank values, Euclidean RTRL, and truncated BPTT, benchmarked against the true model entropy rate, `gzip`, and `CTW`.




Figure 5: Average log-loss (bits per character) on the $a^n b^n_{[1,32]}$ dataset, as a function of the number of characters read, for a leaky RNN with 20 units, trained with a Kalman/QDOP version of the NoBackTrack algorithm, RTRL, and BPTT.




Figure 6: Average log-loss (bits per character) on Shakespeare’s works as a function of the number of characters read, for an RNN with 20 units trained with the QDOP version of the NoBackTrack algorithm for different rank values, Euclidean RTRL and truncated BPTT, benchmarked against `gzip` and CTW.

approximated gradients computed by truncated BPTT are significantly biased.

Keeping track of the long term dependencies is key here, and RTRL outperforms all the algorithms epochwise, though it is still penalized by its high complexity. Truncated BPTT is unable to learn the full dependencies between a ’s and b ’s, and ends up closer to the entropy of an $a^n b^p$ model with independent values of n and p (presumably, it still manages to learn the $a^n b^n$ blocks where n is short). At some point the learning curve of truncated BPTT appears not to decrease anymore and even goes slightly up, which is consistent with a biased gradient estimate.

On the other hand, Kalman NoBackTrack seems to be mostly successful in learning the dependencies. This is confirmed by visual inspection of the output of the learned model. The small remaining gap between the true model and the learned model could be related to incomplete training, or to an imperfect modelling of the exact uniform law for $n \in [k; l]$.

Finally, we report performance of truncated BPTT and Kalman NoBackTrack on Shakespeare’s works. The same 20-unit RNN model is used, again with all algorithms run for the same amount of time using the same learning rate $1/\sqrt{t}$. The curves obtained are displayed in Figure 6.

On this example, RTRL, truncated BPTT, and Kalman NoBackTrack

with various ranks all have a similar performance; it is not clear whether the differences on Figure 6 are statistically significant. This proves, once more, that the stochasticity and rank reduction inherent to NoBackTrack are not detrimental to learning, and allow it to keep up with exact gradient algorithms.

All RNN algorithms have a significantly worse performance than CTW on this example, thus proving that a 20-unit RNN does not accurately model Shakespeare’s works.

Conclusion. We have introduced an algorithm that computes a stochastic, provably unbiased estimate of the derivative of the current state of a dynamical system with respect to its parameters, in a fully online fashion. For recurrent neural networks, the computational cost of this algorithm is comparable to that of running the network itself. Previously known algorithms were either not fully online or had a significantly higher computational cost.

In our experiments, this algorithm appears as a practical alternative to truncated backpropagation through time, especially in its Kalman version, while the Euclidean version requires smaller learning rates. The (unbiased) noise and rank reduction introduced in the gradient approximation do not appear to prevent learning. The interest of NoBackTrack with respect to truncated BPTT depends on the situation at hand, especially on the scale of time dependencies in the data (which results in biased gradient estimates for BPTT), and on whether the storage of past states and past data required by truncated BPTT is acceptable or not.

Acknowledgments. The authors would like to thank Hugo Larochelle for his helpful questions that resulted in several clarifications of the text.

A Variance of the rank-one trick

Keep the notation of Proposition 1 and let $\|\cdot\|$ be a Euclidean norm on the vector space in which the v_i and w_i live.

To measure the variance of \tilde{A} we use the Hilbert–Schmidt norm $\|\tilde{A}\|_{\text{HS}}^2 := \text{Tr}(\tilde{A}^\top \tilde{A})$. This norm satisfies $\|vw^\top\|_{\text{HS}} = \|v\| \|w\|$, and $\langle v_1 w_1^\top | v_2 w_2^\top \rangle_{\text{HS}} = \langle v_1 | v_2 \rangle \langle w_1 | w_2 \rangle$ for the associated scalar product.

Let us evaluate the variance of \tilde{A} in this norm. Since $\text{Var } \tilde{A} = \mathbb{E} \|\tilde{A}\|_{\text{HS}}^2 - \|\mathbb{E} \tilde{A}\|_{\text{HS}}^2$ and $\mathbb{E} \tilde{A} = A$ is fixed, it is enough to evaluate the second moment $\mathbb{E} \|\tilde{A}\|_{\text{HS}}^2$.

We claim that

$$\mathbb{E} \|\tilde{A}\|_{\text{HS}}^2 = \left(\sum_i \|v_i\|^2 \right) \left(\sum_j \|w_j\|^2 \right) + 2 \sum_i \sum_{j \neq i} \langle v_i | v_j \rangle \langle w_i | w_j \rangle \quad (49)$$

Indeed, $\tilde{A} = \sum_{ij} \varepsilon_i \varepsilon_j v_i w_j^\top$ so, by bilinearity of the Hilbert–Schmidt scalar product,

$$\mathbb{E} \|\tilde{A}\|_{\text{HS}}^2 = \mathbb{E} \langle \tilde{A} | \tilde{A} \rangle_{\text{HS}} = \mathbb{E} \sum_{ijkl} \varepsilon_i \varepsilon_j \varepsilon_k \varepsilon_l \langle v_i | v_j \rangle \langle w_k | w_l \rangle \quad (50)$$

Since $\mathbb{E} \varepsilon_i = 0$ and $\mathbb{E}(\varepsilon_i \varepsilon_j) = 0$ for $i \neq j$, the only cases to consider are:

1. $i = j$ and $k = l$ and $i \neq k$: contribution $\sum_i \sum_{k \neq i} \|v_i\|^2 \|w_k\|^2$
2. $i = k$ and $j = l$ and $i \neq j$: contribution $\sum_i \sum_{j \neq i} \langle v_i | v_j \rangle \langle w_i | w_j \rangle$
3. $i = l$ and $j = k$ and $i \neq j$: same contribution as the previous one
4. $i = j = k = l$: contribution $\sum_i \|v_i\|^2 \|w_i\|^2$
5. all other cases contribute 0.

The first and fourth contributions add up to $(\sum_i \|v_i\|^2)(\sum_k \|w_k\|^2)$. This proves (49).

Let us minimize variance over the degrees of freedom given by $v_i w_i^\top = (\rho_i v_i)(w_i / \rho_i)^\top$. ρ_i does not change the last contribution to $\mathbb{E} \|\tilde{A}\|_{\text{HS}}^2$ in (49), neither does it change the expectation $\mathbb{E} \tilde{A} = A$, so to minimize the variance we only have to minimize the first term $(\sum_i \|v_i\|^2)(\sum_k \|w_k\|^2)$. Applying the scaling, this term becomes

$$\left(\sum_i \|v_i\|^2 \rho_i^2 \right) \left(\sum_k \|w_k\|^2 / \rho_k^2 \right) \quad (51)$$

and, by differentiation with respect to a single ρ_i , one checks that this is minimal for

$$\rho_i \propto \sqrt{\|w_i\| / \|v_i\|} \quad (52)$$

(multiplying all ρ_i 's by a common factor does not change the result). So, after optimal scaling,

$$\tilde{A} = \left(\sum_i \varepsilon_i v_i \sqrt{\|w_i\| / \|v_i\|} \right) \otimes \left(\sum_i \varepsilon_i w_i \sqrt{\|v_i\| / \|w_i\|} \right) \quad (53)$$

Consequently, after scaling, the first term in the variance of \tilde{A} in (49) becomes $(\sum_i \|v_i\| \|w_i\|)^2$. The second term in (49) does not change.

Thus, after optimal scaling we find

$$\mathbb{E} \|\tilde{A}\|_{\text{HS}}^2 = \left(\sum_i \|v_i\| \|w_i\| \right)^2 + 2 \sum_i \sum_{j \neq i} \langle v_i | v_j \rangle \langle w_i | w_j \rangle \quad (54)$$

To obtain the variance of \tilde{A} , we just subtract the square norm of $\mathbb{E} \tilde{A} = A$, which is

$$\|A\|_{\text{HS}}^2 = \left\| \sum_i v_i w_i^\top \right\|_{\text{HS}}^2 = \sum_i \|v_i w_i^\top\|_{\text{HS}}^2 + \sum_i \sum_{j \neq i} \langle v_i w_i^\top | v_j w_j^\top \rangle_{\text{HS}} \quad (55)$$

(by bilinearity of the Hilbert–Schmidt scalar product)

$$= \sum_i \|v_i\|^2 \|w_i\|^2 + \sum_i \sum_{j \neq i} \langle v_i | v_j \rangle \langle w_i | w_j \rangle \quad (56)$$

This yields, after optimal scaling,

$$\text{Var } \tilde{A} = \left(\sum_i \|v_i\| \|w_i\| \right)^2 - \sum_i \|v_i\|^2 \|w_i\|^2 + \sum_i \sum_{j \neq i} \langle v_i | v_j \rangle \langle w_i | w_j \rangle \quad (57)$$

$$= \sum_i \sum_{j \neq i} \|v_i\| \|v_j\| \|w_i\| \|w_j\| + \langle v_i | v_j \rangle \langle w_i | w_j \rangle \quad (58)$$

B Invariant norms derived from the Kalman covariance

Algorithm 2 is built to offer invariance properties (a Kalman filter over a variable θ is invariant by affine reparameterization of θ , for instance). However, this only holds if the norms $\|\bar{v}\|$, $\|\bar{w}\|$, $\|v_i\|$, $\|w_i\|$, used to compute the scaling factors $\bar{\rho} = \sqrt{\|\bar{w}\| / \|\bar{v}\|}$ and $\rho_i = \sqrt{\|w_i\| / \|e_i\|}$, are themselves reparameterization-invariant.

This can be achieved if we decide to choose the scalings ρ as to minimize the variance of \tilde{G} computed in the (Mahalanobis) norm defined by the covariance matrix of θ and of h appearing in the Kalman filter.

Let C_θ be the covariance matrix of θ obtained in the Kalman filter; in Algorithm 2, C_θ is approximated by $C_\theta \approx J_\theta^{-1}$.

Any linear form on θ , such as \bar{w} and w_i , can be given a norm by

$$\|\bar{w}\|^2 := \bar{w}^\top C_\theta \bar{w} \approx \bar{w}^\top J_\theta^{-1} \bar{w} \quad (59)$$

and likewise for w_i . This norm is invariant under θ -reparameterization.

Given the covariance C_θ of θ and the dependency $G = \frac{\partial h}{\partial \theta}$ of h with respect to θ , the covariance of h is

$$C_h := G C_\theta G^\top \quad (60)$$

and its inverse $J_h := C_h^{-1}$ can be used to define a norm for a tangent vector v at state h via

$$\|v\|^2 := v^\top J_h v \quad (61)$$

which is also reparametrization-invariant. (We use J_θ^{-1} for the norm of w and J_h for the norm of v because v is a tangent vector (covariant) at point h , while w is a linear form (contravariant) at point θ .)

However, handling of full covariance matrices would be too costly. In Algorithm 2, the inverse covariance J_θ of θ is already an approximation (diagonal, quasi-diagonal...) via MatrixReduce. Moreover, here we only have access to an approximation \tilde{G} of G . Thus, we simply replace G with \tilde{G} in the definition of C_h , and use a diagonal reduction. This leads to $C_h \approx \text{Diag}(\tilde{G} J_\theta^{-1} \tilde{G}^\top)$ and

$$J_h \approx (\text{Diag}(\tilde{G} J_\theta^{-1} \tilde{G}^\top))^{-1} \quad (62)$$

where as usual \tilde{G} is the gradient approximation given by (9).

The diagonal reduction is necessary if \tilde{G} is low-rank, since $\tilde{G} J_\theta^{-1} \tilde{G}^\top$ will be low-rank as well, and thus non-invertible.

Then the scaling factors $\bar{\rho}$ and ρ_i can finally be computed as

$$\bar{\rho} = \sqrt{\frac{\|\bar{w}\|}{\|\bar{v}\|}} = \frac{(\bar{w}^\top J_\theta^{-1} \bar{w})^{1/4}}{(\sum_i (\tilde{G} J_\theta^{-1} \tilde{G}^\top)_{ii}^{-1} \bar{v}_i^2)^{1/4}} \quad (63)$$

and

$$\rho_i = \sqrt{\frac{\|w_i\|}{\|e_i\|}} = \frac{(w_i^\top J_\theta^{-1} w_i)^{1/4}}{((\tilde{G} J_\theta^{-1} \tilde{G}^\top)_{ii}^{-1})^{1/4}} \quad (64)$$

The particular structure of J_θ (if approximated by, e.g., a block-diagonal matrix) and of $\tilde{G} = \bar{v}\bar{w}^\top + \sum_i e_i w_i^\top$ make these computations efficient.

Note that even with the approximations above, \tilde{G} is still an unbiased estimate of G . Indeed, any choice of ρ has this property; we are simply approximating the optimal ρ which minimizes the variance of \tilde{G} .

In practice, small regularization terms are included in the denominator of every division and inversion to avoid numerical overflow.

References

- [BEYY04] Ron Begleiter, Ran El-Yaniv, and Golan Yona. On prediction using variable order markov models. *Journal of Artificial Intelligence Research*, pages 385–421, 2004.
- [Cap11] Olivier Cappé. Online EM algorithm for hidden Markov models. *J. Comput. Graph. Statist.*, 20(3):728–749, 2011.
- [Hay04] Simon Haykin. *Kalman filtering and neural networks*. John Wiley & Sons, 2004.
- [Jae02] Herbert Jaeger. Tutorial on training recurrent neural networks, covering BPTT, RTRL, EKF and the “echo state network” approach. Technical Report 159, German National Research Center for Information Technology, 2002.
- [Oll15a] Yann Ollivier. Riemannian metrics for neural networks I: feedforward networks. *Information and Inference*, 4(2):108–153, 2015.
- [Oll15b] Yann Ollivier. Riemannian metrics for neural networks II: recurrent networks and learning symbolic data sequences. *Information and Inference*, 4(2):154–193, 2015.