Rethinking Exponential Averaging of the Fisher

Constantin Octavian Puiu[0000−0002−1724−4533] ✉

University of Oxford, Mathematical Institute,
constantin.puiu@maths.ox.ac.uk

Abstract. In optimization for Machine learning (ML), it is typical that curvature-matrix (CM) estimates rely on an exponential average (EA) of local estimates (giving EA-CM algorithms). This approach has little principled justification, but is very often used in practice. In this paper, we draw a connection between EA-CM algorithms and what we call a *"Wake of Quadratic models"*. The outlined connection allows us to understand what EA-CM algorithms are doing from an optimization perspective. Generalizing from the established connection, we propose a new family of algorithms, *KL-Divergence Wake-Regularized Models* (KLD-WRM). We give three different practical instantiations of KLD-WRM, and show numerically that these outperform K-FAC on MNIST.

Keywords: Optimization · Natural Gradient · KL Divergence · Overfit.

1 Introduction

Recent research in optimization for ML has focused on finding tractable approximations for curvature matrices. In addition to employing ingenious approximating structures ([1,2,3]), curvature matrices are typically estimated as an exponential average (EA) of the previously encountered local estimates ([1,3,4,5,6]). This is in particular true for Natural Gradient (NG) algorithms, which we focus on here. These exponential averages emerge rather heuristically, and have so far only been given incomplete motivations. The main such motivation is *"allowing curvature information to depend on much more data, with exponentially less dependence on older data"* [1]. However, it remains unclear what such an EA algorithm is actually doing from an optimization perspective. In this paper, we show that such EA algorithms can be seen as solving a sequence of local quadratic models whose construction obeys a recursive relationship - which we refer to as a *"Wake of Quadratic Models"*. Inspired by this recursion, we consider a similar, more principled and general recursion for our local models - which we refer to as a *"KL-Divergence Wake-Regularized Models"* (KLD-WRM). We show that under *suitable approximations*, KLD-WRM gives very similar optimization steps to EA-NG. This equivalence raises the hope that using better approximations in our proposed class of algorithms (KLD-WRM) might lead to algorithms which outperform EA-NG. We propose three practical instantiations of KLD-WRM (of increasing approx. accuracy) and compare them with K-FAC, the most widely used practical implementation of NG for DNNs. Numerically, KLD-WRM outperforms K-FAC on MNIST: higher test accuracy, lower test loss, lower variance.

2   Preliminaries

2.1   Neural Networks, Supervised Learning and Notation

We very briefly look at *fully-connected (FC) nets* (we omit CNN for simplicity). We have $\bar{a}_0 := [x, 1]^T$ (x is the input to the net). The *pre-activation* at layer l: $z_l = W_l \bar{a}_{l-1}$ for $l \in \{1, 2, ..., n_L\}$. The *post-activation* at layer l: $a_l = \phi_l(z_l)$ for $l \in \{1, 2, ..., n_L\}$. The *augmented post-activation* at layer l: $\bar{a}_l = [a_l, 1]^T$ for $l \in \{1, 2, ..., n_L - 1\}$ (we augment the post-activations to incorporate the bias into the weight matrix W_{l+1} w.l.o.g.; this is standard practice for K-FAC [1]).

In the above, we consider n_L layers, and $\phi_l(\cdot)$ are nonlinear activation functions. We collect all parameters, namely $\{W_l\}_{l=1}^{n_L}$, in a parameter vector $\theta \in \mathbb{R}^d$. Let us consider $N_o \in \mathbb{Z}^+$ neurons in the output layer. The output of the net is $h_\theta(x) := a_{n_L} \in \mathbb{R}^{N_o}$. The predictive (model) distribution of $y|x$ is $p_\theta(y|x) = p_\theta(y|h_\theta(x))$, that is $p(y|x)$ depends on x only through $h_\theta(x)$.

In supervised learning, where we have labeled pairs of datasets $\mathcal{D} = \{(x_i, y_i)\}_i$. Our objective to minimize is typically some regularized modification of

$$f(\theta) = -\log p(\mathcal{D}|\theta) = \sum_{(x_i, y_i) \in \mathcal{D}} \left(-\log p(y_i | h_\theta(x_i)) \right). \tag{1}$$

Thus, we have our loss function $L(y_i, x_i; \theta) := -\log p(y_i | h_\theta(x_i))$, and $f(\theta) = \sum_{(x_i, y_i) \in \mathcal{D}} L(y_i, x_i; \theta)$. This is what we will focus on here, for simplicity of exposition, but our ideas directly apply to different ML paradigms, such as Variational Inference (VI) in Bayesian Neural Networks (BNNs), and to RL as well.

Let us consider the iterates $\{\theta_k\}_{k=0,1,...}$, with θ_0 initialized in standard manner (perhaps on the edge of Chaos)[1]. Let us denote the optimization steps taken at θ_k as s_k, that is $\theta_{k+1} = \theta_k + s_k$. Let $g(\theta)$ and $H(\theta)$ be the gradient and hessian of our objective $f(\theta)$. We will use the notation $g_k := g(\theta_k)$ and $H_k := H(\theta_k)$.

2.2   KL-Divergence and Fisher Information Matrix

The *symmetric*[2] KL-Divergence is a distance measure between two distributions:

$$\mathbb{D}_{KL}(p, q) := \frac{1}{2}\left[\mathbb{D}_{KL}(q \,||\, p) + \mathbb{D}_{KL}(p \,||\, q) \right] = \frac{1}{2}\mathbb{E}_{x \sim p}\left[\log \frac{p(x)}{q(x)}\right] + \frac{1}{2}\mathbb{E}_{x \sim q}\left[\log \frac{q(x)}{p(x)}\right]. \tag{2}$$

In our case, we are interested in the symmetric KL-divergence (SKL) between the data joint distribution with one parameter value and the data joint distribution with a different parameter value. We have

$$\mathbb{D}_{KL}(\theta_1, \theta_2) := \mathbb{D}_{KL}\left(p_{\theta_1}(x, y), p_{\theta_2}(x, y) \right). \tag{3}$$

Since we only model the conditional distribution, we let $p_{\theta_i}(x, y) = \hat{q}(x)p(y|h_{\theta_i}(x))$, where $\hat{q}(x)$ is the marginal (empirical) data distribution of x. This gives

---

[1] Note that the initialization issue is orthogonal to our purpose.

[2] For convenience, we will refer to the *symmetric KL-Divergence* as *SKL-Divergence*.

$$\mathbb{D}_{KL}(\theta_1, \theta_2) = \frac{1}{N} \sum_{x_i \in \mathcal{D}} \mathbb{D}_{KL}\big(p_{\theta_1}(y|x_i), p_{\theta_2}(y|x_i)\big). \tag{4}$$

The Fisher has multiple definitions depending on the situation, but here we have

$$F_k := F(\theta_k) := \mathbb{E}_{\substack{x \sim \hat{q}(x) \\ y \sim p_\theta(y|x)}} \left[\nabla_\theta \log p_\theta(y|x) \nabla_\theta \log p_\theta(y|x)^T \right]. \tag{5}$$

This is the Fisher of the joint distribution $p_\theta(x, y) = \hat{q}(x)p(y|h_\theta(x))$. We let $F(\theta_k, x) := \mathbb{E}_{y \sim p_\theta(y|x)}[\nabla_\theta \log p_\theta(y|x)\nabla_\theta \log p_\theta(y|x)^T]$ be the Fisher of the conditional distribution $p_\theta(y|x)$. Then, we have $F_k = \mathbb{E}_{x \sim \hat{q}}[F(\theta_k, x)]$. Using the Fisher, we have the following approximation (see [1,7]) for the SKL-Divergence

$$\mathbb{D}_{KL}\big(p_{\theta_1}(y|x), p_{\theta_2}(y|x)\big) \approx \frac{1}{2}(\theta_2 - \theta_1)^T F(\theta_1, x)(\theta_2 - \theta_1), \tag{6}$$

which is exact as $\theta_2 \to \theta_1$. By plugging (6) into (4) and using linearity we get

$$\mathbb{D}_{KL}(\theta_1, \theta_2) \approx \frac{1}{2}(\theta_2 - \theta_1)^T F_1(\theta_2 - \theta_1). \tag{7}$$

2.3 Natural Gradient and K-FAC

The natural gradient (NG) is defined as [7]

$$\nabla_{NG} f(\theta_k) = F_k^{-1} g_k. \tag{8}$$

The NG descent (NGD) step is then taken to be $s_k^{(NGD)} = -\alpha_k \nabla_{NG} f(\theta_k)$, for some stepsize α_k. NGD has favorable properties, the most notable of which is reparametrization invariance (when the step-size is infinitesimally small) [8]. The NGD step can be expressed as the solution to the quadratic problem (see [1,8])

$$s_k^{(NGD)} := \arg \min_s s^T g_k + \frac{1}{2\alpha_k} s^T F_k s. \tag{9}$$

K-FAC Storing and inverting the Fisher is prohibitively expensive. K-FAC is a practical implementation of NG which bypasses this problem by approximating the Fisher as a block-diagonal[3] matrix, with each diagonal block further approximated as the Kronecker product of two smaller matrices [1]. That is, we have

$$F_k^{(KFAC)} := \text{blockdiag}\big(\{A_k^{(l)} \otimes \Gamma_k^{(l)}\}_{l=1,\ldots,n_L}\big), \tag{10}$$

where each block corresponds to a layer and $\otimes$ denotes the Kronecker product [1]. For example, for FC nets, the Kronecker factors are given by

$$A_k^{(l)} := \mathbb{E}_{x,y \sim p}[\bar{a}_{l-1}\bar{a}_{l-1}^T], \quad \Gamma_k^{(l)} := \mathbb{E}_{x,y \sim p}[\nabla_{z_l} L \nabla_{z_l} L^T]. \tag{11}$$

Note that the Kronecker factors depend on θ_k which influences both the forward and backward pass. Also, note they can be efficiently worked with since $(A_k^{(l)} \otimes \Gamma_k^{(l)})^{-1} = [A_k^{(l)}]^{-1} \otimes [\Gamma_k^{(l)}]^{-1}$, and $\big([A_k^{(l)}]^{-1} \otimes [\Gamma_k^{(l)}]^{-1}\big)v = \text{vec}\big([\Gamma_k^{(l)}]^{-1}V[A_k^{(l)}]^{-1}\big)$, where v maps to V in the same way $\text{vec}(W_l)$ maps to W_l [1].

---

[3] Block tri-diagonal approximation is also possible - but this lies outside our scope.

2.4   Curvature in Practice: Exponential Averaging

In practice, many algorithms (including ADAM and K-FAC) do not use the curvature matrix estimate as computed. Instead, they maintain an exponential average (EA) of it (eg. [3,4,5,6]). In the case of NG, this EA is

$$\bar{F}_k := \rho^k F_0 + (1 - \rho) \sum_{i=1}^{k} \rho^{k-i} F_i, \tag{12}$$

where $\rho \in [0, 1)$ is the exponential decay parameter. Let us refer to NG algorithms which replace the Fisher, F_k, with its exponential average, $\bar{F}_k$, as EA-NG.

In a similar spirit, K-FAC maintains an EA of the Kronecker factors

$$\bar{A}_k^{(l)} := \rho^k A_0^{(l)} + (1 - \rho) \sum_{i=1}^{k} \rho^{k-i} A_i^{(l)}, \quad \bar{\Gamma}_k^{(l)} := \rho^k \Gamma_0^{(l)} + (1 - \rho) \sum_{i=1}^{k} \rho^{k-i} \Gamma_i^{(l)}, \tag{13}$$

and in practice, $A_k^{(l)}$ and $\Gamma_k^{(l)}$ in (10) are replaced with $\bar{A}_k^{(l)}$ and $\bar{\Gamma}_k^{(l)}$ respectively. We will refer to the practical implementation of K-FAC which uses EA for the K-FAC matrices as EA-KFAC, to emphasize the presence of the EA aspect. However, this is the norm in practice rather than an exception, and virtually any algorithm referred to as K-FAC (or as using K-FAC) is in fact an EA-KFAC algorithm.

3   A Wake of Quadratic Models (WoQM)

The idea behind WoQM is simple. Instead of taking a step s_k, at θ_k which relies only on the local quadratic model, $s_k = \arg\min_s g^T s + (1/2)s^T B_k s$, we take a step which relies on an EA of all previous local models. Formally, let us define

$$M_k^{(Q)}(s) := g_k^T s + \frac{\lambda_k}{2} s^T B_k s, \tag{14}$$

for an arbitrary *symmetric-positive definite* curvature matrix B_k (typically an approximation of H_k or F_k). Note that $1/\lambda_k$ is a step-size (or learning rate) parameter. Our WoQM step, s_k, is then defined as the solution to

$$\min_s \sum_{i=0}^{k} \rho^{k-i} M_i^{(Q)}\left(s + \sum_{j=i}^{k-1} s_j\right), \quad \text{with} \tag{15}$$

$$\lambda_0 = \lambda, \text{ and } \lambda_k = (1 - \rho)\lambda, \quad \forall k \in \mathbb{Z}^+, \tag{16}$$

where we set $\sum_{j=k}^{k-1} s_j = 0$ by convention, $\rho \in [0, 1)$ is an exponential decay parameter, and $\lambda > 0$ is a hyperparameter. While (15) does not appear to be a proper exponential averaging, missing a $1 - \rho$ factor in terms where $i \geq 1$, it can be easily rearranged as such by slightly modifying the definition of $M_0^{(Q)}$ (to disobey (14)). To see this, multiply (15) by $(1 - \rho)$ and then absorb the $1 - \rho$

factor in the definition of $M_0^{(Q)}$. Our stated definition makes the exposition more compact while preserving intuition.

For our choice of model (14), the WoQM step s_k (at θ_k) is the solution of

$$\min_s s^T \left[\sum_{i=0}^{k} \rho^{k-i} \left(g_i + \lambda \kappa(i) B_i \sum_{j=i}^{k-1} s_j \right) \right] + \frac{\lambda}{2} s^T \left[\sum_{i=0}^{k} \kappa(i) \rho^{k-i} B_i \right] s, \qquad (17)$$

where we dropped all the constant terms, and have $\kappa(i) := \exp(\mathbb{I}_{\{i>0\}} \log(1-\rho))$, where $\mathbb{I}_{\mathcal{E}}$ is the indicatior function of event $\mathcal{E}$. We use the $\kappa(i)$ term for notational compactness. Note that definition (15) can be used for general models M_i (rather than quadratic $M_i^{(Q)}$), leading to a larger family of algorithms for which WoQM represents a particular instantiation: the *Wake of Models (WoM)* family.

3.1 Connection with Exponential-Averaging in Curvature Matrices

We now look at how the WoQM step relates to Exponential-Averaging Curvature matrices (EA-CM). EA-CM is standard practice in stochastic optimization, and in particular in training DNNs (see for example [1,3,4,5,6]). Formally, using EA-CM boils down to taking a step based on (14), but with B_k replaced by

$$\bar{B}_k := \rho^k B_0 + (1-\rho) \sum_{i=1}^{k} \rho^{k-i} B_i = \sum_{i=0}^{k} \kappa(i) \rho^{k-i} B_i. \qquad (18)$$

Note that we used the same exponential decay parameter for convenience, but this is not required.

It is obvious from (17) that we can get an analytic solution for WoQM step s_k, as a function of $\{s_j\}_{j=0}^{k-1}$, $\{(g_j, B_j)\}_{j=0}^{k}$, ρ and λ. Thus, by using the relationship recursively we can get an analytic solution for s_k as a function of $\{(g_j, B_j)\}_{j=0}^{k}$, ρ and λ. When doing this, the connection between WoQM and EA-CM is revealed. The result is presented in *Proposition 3.1*.

Proposition 3.1: Analytic Solution of WoQM step. *The* WoQM *step s_k at iterate θ_k can be expressed as*

$$s_k = -\lambda^{-1} \bar{B}_k^{-1} g_k, \quad \forall k \in \mathbb{Z}^+. \qquad (19)$$

Proof. Relies on simple inductive argument. See supplementary material. □

Proposition 3.1 tells us that the WoQM step is exactly the step obtained by using an EA curvature matrix $\bar{B}_k$ in a simple quadratic model of the form (14). That is, WoQM (15)-(16) is the principled optimization formulation of a EA-CM step[4], when the EA-CM stepsize is constant and equal to $1/\lambda$. Thus, we see what EA-CM is actually doing from an optimization perspective: instead of perfectly solving for the local quadratic model, it solves for a trade-off between all previously encountered models, where the weights of the trade-off are (almost[5]) given by an exponential average (older models receive exponentially less "attention").

---

[4] WoQM with $\rho = 0$ and B_k based on quantities at θ_k only is also the principled optimization formulation of no-EA CM algorithms, with steps of the form (14).

[5] Can modify the definition of $M_0^{(Q)}$ s.t. WoQM is a proper EA of quadratic models.

There are two observations to make at this point. First, we began by noting that the justification for EA-CM is largely heuristic, but we ended up explaining EA-CM through some optimization model which involved an EA of local models (the WoQM model). Since the EA was the difficult part to justify in the first place, it might seem that we are sweeping the problem from under one rug to another. However, this is not the case. The WoQM formulation aims to reveal a different perspective on EA-CM algorithms, rather than explain the presence of EA itself. It is indeed true that we did not justify why one should use EA and thus get the WoQM family, but this is not required to enhance our understanding and draw conclusions. We can draw conclusions purely based on the established equivalence. This leads us to the second observation, which is why EA-CM improves stability from a stochastic optimization perspective. Rather than conferring stability because it *"uses more data"* ([1,4]), EA-CM can alternatively be thought of as conferring stability because it uses the collection of all previous noisy local models to build a better model[6] (in terms of both noise and functional form).

Note that what we have discussed so far applies when using any curvature matrix B_k. In particular, *Proposition 3.1* can be directly applied to establish an equivalence between EA-NG algorithms and FISHER-WOQM algorithms (WoQM algorithms with $B_k = F_k$). In fact, we can replace the Fisher by any approximation and still have the equivalence holding, *if the EA is done as in (12)*.

3.2   Fisher-WoQM and Practical K-FAC Equivalence

We have seen (in *Section 2.4*) that K-FAC holds an EA for the *Kronecker factors* (see (13)), rather than for the K-FAC approximation to F_k (as in (12)). Thus, the EA scheme employed by K-FAC is *not* the same as (12) with $F_k \leftarrow F_k^{(KFAC)}$. Therefore, we cannot directly apply *Proposition 3.1* to obtain an equivalence between EA-KFAC and KFAC-WOQM (WoQM with $B_k \leftarrow F_k^{(KFAC)}$). We can loosely establish this equivalence by viewing the *EA over the Kronecker factors* as a convenient (but coarse) approximation to the *EA over $F_k^{(KFAC)}$* (see *Section 4* in the *supplementary material*). Indeed, carrying an EA for $F_k^{(KFAC)}$ is impractical.

Our equivalence reveals that EA-NG is actually solving an exponentially decaying wake of quadratic models (WoQM) where the curvature matrix (meant to be a Hessian approximation) is taken to be an approximate Fisher. This is in contrast with the typical EA-NG interpretation which says that we take NG steps by solving quadratic models of the form (14) with $B_k = F_k$, but then we further approximate F_k as an EA based on $\{\hat{F}_j\}_{j=1}^{k-1}$. While EA-KFAC does not do the exact same thing, it can be seen as a (very crude) approximation to it.

---

[6] Although it is still not clear why putting the previous local models together in an exponentially-averaged fashion is *"right"* for conferring further stability - and this aspect remains heuristic. While this could be informally and partly explained by *"older models should matter less"*, the complete explanation remains an open question.

Dealing with Infrequent Updates In practice, the curvature matrix is not computed at each location, and the EA update is typically performed every $N_u \approx 100$ steps[7] to save computation cost (at least in supervised learning[8]) [1]. In this circumstance, the final implementation of EA-NG would actually be equivalent to a WoQM algorithm where we set

$$B_k = \begin{cases} \hat{F}_k, & \text{if} \quad \mod(k, N_u) = 0 \\ \bar{B}_q, \text{ where } q := N_u \lfloor \frac{k}{N_u} \rfloor, \text{ otherwise} \end{cases} . \tag{20}$$

In (20), $\hat{F}_k$ represents an approximation for the Fisher F_k, whose computation has an associated cost. Recall that $\bar{B}_k = \sum_{j=0}^{k} \kappa(j)\rho^{k-j}B_j$. Note that (20) is well defined since $\bar{B}_0 = B_0 = \hat{F}_0$, then $B_1 = B_2 = ... = B_{N_u-1} = B_0$, and hence $\bar{B}_1 = \bar{B}_2 = ... = \bar{B}_{N_u-1} = B_0$, and so on. Further note that by (19), defining our WoQM curvature matirx as in (20) gives us the EA-NG matrix that we want, since we can easily see that: $\bar{B}_{N_u} = \rho\bar{B}_0 + (1-\rho)\hat{F}_{N_u} = \rho\hat{F}_0 + (1-\rho)\hat{F}_{N_u}$. We can then easily extend the argument to show $\bar{B}_{qN_u} = \sum_{j=0}^{q} \kappa(j)\rho^{q-j}\hat{F}_{jN_u}^{(KFAC)}$ which is exactly the form of EA that EA-NG employs when updating statistics every N_u steps. Note that we can apply *Proposition 3.1* irrespectively of our choice of B_k, so in particular, it must hold for B_k as defined in (20).

By extending our reasoning, we see that any heuristic which adapts N_u as a function of observations up until k can be transformed into an equivalent heuristic of picking between $B_k = \bar{B}_q$ (where q is now more generally the previous location where we computed $\hat{F}_k$) and $B_k = \hat{F}_k$. Thus, the equivalence between WoQM and EA-NG holds irrespectively of the heuristic which decides when to update the EA-NG matrix. The exact same reasoning holds for generic EA-CM algorithms.

3.3   Fisher-WoQM: A Different Perspective

Since we have the following approximation[9] for $k \geq i$ [8]

$$\mathbb{D}_{KL}(\theta_i, \theta_k + s) \approx \tilde{\mathbb{D}}_{KL}(\theta_i, \theta_k + s) := \frac{1}{2}\left(s + \sum_{j=i}^{k-1} s_j\right)^T F_i \left(s + \sum_{j=i}^{k-1} s_j\right), \quad (21)$$

one might think that a WoQM model with $B_k = F_k$ and $\rho \in (0,1)$ would in fact be some form of approximate *"KL-Divergence Wake-Regularized[10] model"*. This is indeed true, as we can write our FISHER-WoQM as

$$\min_s \left[\sum_{i=0}^{k} \rho^{k-i}g_i\right]^T s + \lambda \sum_{i=0}^{k} \kappa(i)\rho^{k-i}\tilde{\mathbb{D}}_{KL}(\theta_i, \theta_k + s). \tag{22}$$

Note that $\tilde{\mathbb{D}}_{KL}(\theta_i, \theta_k + s)$ in (21) is a second-order approximation of the

---

[7] More complicated heuristics can be designed, see [9].

[8] In RL we may prefer updating the K-FAC EA-matrix at each step [10].

[9] Which is exact in the limit as $\theta_k + s \to \theta_i$, but would nevertheless be very crude in practice, particularly for large $k - i$, since the steps taken might be relatively large.

[10] Regularizing w.r.t. SKL divergences relative to all previous distributions, as opposed to just the most recent one - as the quadratic model associated with NG step does.

SKL-Divergence between $p_{\theta_i}(x, y)$ and $p_{\theta_k+s}(x, y)$. Thus, we see that FISHER-WOQM (same as EA-NG by *Prop. 3.1*) is in fact solving a *regularized linear model*, where the model gradient is taken to be the *momentum gradient* (with parameter ρ), and the regularization term is an exponentially decaying wake of *(crudely) approximate* SKL-divergences relative to previously encountered distributions.

This equivalence raises scope for a new family of algorithms: perhaps using another model for the objective f, rather than a simple linear model based on momentum-gradient, and/or using a better approximation for the KL divergence could lead to better performance. This is the main topic of this paper, explored formally in *Section 4* and numerically in *Section 5*. We now note that WOQM (and thus also EA-NG by *Proposition 3.1*, and approximately so, EA-KFAC) is in fact a particular instantiation of the family proposed in the next section.

4   A KL-Divergence Wake-Regularized Models Approach

We now propose a new family of algorithms, which we call *KL-Divergence Wake-Regularized Models (*KLD-WRM*; reads: "Cold-Worm")*. At each location θ_k, the KLD-WRM step s_k is defined as the solution to the problem

$$\min_s M(s; \mathcal{F}_k) + \lambda \sum_{i=0}^{k} \zeta(i)\rho^{k-i}\mathbb{D}_{KL}(\theta_i, \theta_k + s), \tag{23}$$

where $\rho \in [0, 1)$, $M(s; \mathcal{F}_k)$ is a model of the objective which uses *at most* all the information ($\mathcal{F}_k$) encountered up until and including θ_k, and $\zeta(i)$ allows for different 'λ's' at different i's. Simple choices would be $\zeta(i) \equiv 1$, or $\zeta(i) = \kappa(i)$.

The motivation behind KLD-WRM is two-fold. First, a wake of SKL regularization allows us to stay close (in a KL sense) to all previously encountered distributions, rather than only to the most recent one. Thus, we might expect KLD-WRM to give more conservative steps in terms of distribution ($p_\theta(y|x)$) change. Second, KLD-WRM can be seen as a generalization[11] of EA-NG (which is also FISHER-WOQM), which also *"undoes"* the approximation[12] of SKL.

Note that (23) is the most general formulation of KLD-WRM, but in order to obtain practical algorithms we have to make further approximations and definitions (of M and ζ). For example, we could set

$$M(s; \mathcal{F}_k) = \sum_{i=0}^{k} \nu^{k-i} M_i\left(s + \sum_{j=i}^{k-1} s_j\right) \tag{24}$$

where the models $M_i(s)$ are general models (not necessary quadratic), constructed only based on local information at θ_i. This would be a *Wake of Models*

---

[11] The linear model $\left[\sum_{i=0}^{k} \rho^{k-i} g_i\right]^T s$ in (22) becomes arbitrary, and $\kappa(i)$ is replaced by a general $\zeta : \mathbb{Z}^+ \to \mathbb{R}$.

[12] For small $\|s\|$ we have $\mathbb{D}(\theta, \theta + s) \approx \mathbb{D}(\theta\|\theta + s) \approx \mathbb{D}(\theta + s\|\theta) \approx (1/2)s^T F(\theta)s$. Thus, the generalization towards our family could use $\xi\mathbb{D}(\theta\|\theta + s) + (1 - \xi)\mathbb{D}(\theta + s\|\theta)$ $\forall \xi \in \mathbb{R}$, instead of the SKL (i.e. $\xi = 1/2$). We choose SKL for simplicity.

KLD-WRM (WOM-KLD-WRM). We could make this even more particular, for example by considering instantiations where $\nu = 0$, but $\rho \in [0, 1)$ in (24). This would give a *Local-Model* KLD-WRM (LM-KLD-WRM), whose steps s_k are the solution to

$$\min_s M_k(s) + \lambda \sum_{i=0}^{k} \zeta(i) \rho^{k-i} \mathbb{D}_{KL}(\theta_i, \theta_k + s), \tag{25}$$

In this paper, we focus on LM-KLD-WRM instantiations (a particular sub-family of KLD-WRM), and leave the general case as future work. We give three instantiations of LM-KLD-WRM of increasing complexity, and discuss the links with already existing methods. We investigate their performance in *Section 5*.

4.1   Connection between KLD-WRM and Fisher-WoQM

It is easy to see that setting $\nu = \rho$, $\zeta(i) = \kappa(i)$, $M_i(s) = s^T g_i$ and approximating $\mathbb{D}_{KL}(\theta_i, \theta_k + s) \approx \tilde{\mathbb{D}}_{KL}(\theta_i, \theta_k + s)$ (defined in (21)) in a WOM-KLD-WRM gives the WOQM family. Note that WOQM is *not* an LM-KLD-WRM model as $M_k(s)$ can only include information local to θ_k (eg. cannot include g_{k-1}). However, the simplest instantiation of LM-KLD-WRM takes steps which are formally similar to FISHER-WOQM (and thus EA-NG) steps. We investigate this in *Section 4.2*.

4.2   Simplest KLD-WRM Instantiation: Smallest Order KLD-WRM

The simplest practical instantiation of KLD-WRM, *Smallest Order* KLD-WRM (SO-KLD-WRM), uses the most crude approximations to LM-KLD-WRM (25) and sets $\zeta(i) = \kappa(i)$. The SO-KLD-WRM step s_k (at θ_k) is given by

$$\min_s g_k^T s + \lambda \sum_{i=0}^{k} \kappa(i) \rho^{k-i} \tilde{\mathbb{D}}_{KL}(\theta_i, \theta_k + s), \tag{26}$$

It is trivial to see that SO-KLD-WRM differs from FISHER-WOQM (22) only through replacing $\left[\sum_{i=0}^{k} \rho^{k-i} g_i \right]$ with g_k. Since FISHER-WOQM is equivalent to EA-NG, one might expect that SO-KLD-WRM steps are *formally* similar to EA-NG steps. *Proposition 4.1* formalizes this result.

Proposition 4.1: Analytic Solution of SO-KLD-WRM step. *The* SO-KLD-WRM *step* s_k *at iterate* θ_k *can be expressed as*

$$s_k = -\lambda^{-1} \bar{F}_k^{-1} [g_k - \rho g_{k-1}], \tag{27}$$

$\forall k \in \mathbb{Z}^+$, *where* $\bar{F}_k := \sum_{i=0}^{k} \kappa(i) \rho^{k-i} F_i$ *and we set* $g_{-1} := 0$ *by convention.*

   Proof. By induction. See the *supplementary material*. $\square$

   Note that we set $g_{-1} = 0$ to avoid providing two separate cases (for $k = 0$ and for $k \geq 1$). *Proposition 4.1* tells us that the SO-KLD-WRM step is *formally similar* to the FISHER-WOQM step (which we have seen is the EA-NG step). The only (formal) difference between the SO-KLD-WRM step and the EA-NG step is that g_k gets replaced by $g_k - \rho g_{k-1}$ in (19). By *"formally"* here, we mean that the expressions look very similar. However, when considering the two implemented algorithms, the paths taken can be very different.

We have seen that EA-NG (being equivalent to FISHER-WOQM) algorithms are in fact a sub-family of the KLD-WRM family, and obviously SO-KLD-WRM is also a sub-family of KLD-WRM. Thus, the formal difference in the steps between SO-KLD-WRM and EA-NG tells us these two sub-families are distinct. The formal similarity between SO-KLD-WRM and EA-NG, combined with the fact that both are members of the KLD-WRM family raises hopes that more accurate instantiations KLD-WRM might lead to better performance than EA-NG.

Note that basing our SO-KLD-WRM step computation on *Proposition 4.1* gives a tractable algorithm, while solving (26) directly gives an intractable algorithm. To see this, review definition (21), and realize that solving equation (26) *directly* requires storing all previous $\bar{F}_j s_j$ matrix-vector products as a minimum (if the algorithm is written efficiently). Thus, we have an exploding number of vectors that need to be stored when solving (26) *directly*, which eventually will overflow the memory - giving an untractable algorithm. Conversely, using *Proposition 4.1* only requires storing at most 2 matrices ($\bar{F}_k$ and F_k) and 2 gradient-shaped vectors at any one moment in time. Thus, using *Proposition 4.1* with tractable approximations for F_k gives a tractable algorithm.

Reconsidering Gradient Momentum in EA-KFAC By comparing (22) and (26), we see that FISHER-WOQM (which is also EA-NG) can *almost* be seen as SO-KLD-WRM with added momentum for the gradient. The equivalence would be exact if we would have a $\kappa(i)$ inside the sum of the linear term of (22). In EA-KFAC, gradient momentum is not added in the standard fashion. A different way to add momentum is proposed[13] and presented as successful, perhaps because trying to add gradient momentum in the standard fashion gives worse performance [1]. From our discussion, this could be because EA-KFAC can be interpreted as an approximate EA-NG, and EA-NG is a SO-KLD-WRM algorithm which already has included gradient momentum. Thus, further adding momentum does not make sense. Note that by adding momentum to gradient in the *standard fashion* we mean replacing g_k by $\sum_{i=0}^{k} \kappa(i)\rho^{k-i}g_k$ (see [3]).

SO-KLD-WRM in practice When implementing SO-KLD-WRM in practice, we use the K-FAC approximation of the Fisher. Note that *Proposition 4.1* tells us that we need *not* store all previous K-FAC matrices. Instead, we can only save the EA-KFAC matrix. As we have discussed in *Section 3*, we can skip computing the K-FAC matrix at some locations to save on computation. To do that, we just pretend the new K-FAC matrix at θ_k is the EA-KFAC that we currently have stored. For example, if we want to compute K-FAC matrices only once in N_u steps, we use (20), but different heuristics can also be used (for eg. as in [9]). Of course, in practice we store an EA for the Kronecker factors (instead of an EA for the K-FAC matrix), as is standard with practical K-FAC implementation [1].

---

[13] More akin to a subspace method rather than a standard momentum method (see [1]).

4.3   Second KLD-WRM Instantiation: Q-KLD-WRM

The *Quadratic* KLD-WRM (Q-KLD-WRM) is an LM-KLD-WRM instantiation which uses a second-order approximation for both the Model and the SKLs, and sets $\zeta(i) = \kappa(i)$. The Q-KLD-WRM step s_k (at θ_k) solves

$$\min_s g_k^T s + \frac{1}{2} s^T B_k s + \lambda \sum_{i=0}^{k} \kappa(i) \rho^{k-i} \tilde{\mathbb{D}}_{KL}(\theta_i, \theta_k + s), \tag{28}$$

where B_k is a curvature matrix which aims to approximate the Hessian H_k. That is, Q-KLD-WRM sets M_k in (25) to be a quadratic model. Since (28) is overall a quadratic, we can obtian an *analytic solution* for the Q-KLD-WRM step.

Proposition 4.2: Analytic Solution of Q-KLD-WRM step. *The* Q-KLD-WRM *step s_k at location θ_k is given by the solution to the problem*

$$\min_s s^T \hat{g}_k + \frac{\lambda}{2} s^T \left[\bar{F}_k + \frac{1}{\lambda} B_k \right] s \tag{29}$$

where $\hat{g}_k$ is given by the one-step recursion

$$\hat{g}_{k+1} = g_{k+1} + \rho(I - \hat{M}_k)\hat{g}_k - \rho g_k, \quad \forall k \in \mathbb{Z}^+, \tag{30}$$

with $\hat{g}_0 := g_0$, $\bar{F}_k := \sum_{i=0}^{k} \kappa(i)\rho^{k-i} F_i$, and $\hat{M}_k := \left[I + \frac{1}{\lambda} B_k \bar{F}_k^{-1} \right]^{-1}$. That is, the Q-KLD-WRM *step is formally given by $s_k = -\frac{1}{\lambda}\left[\bar{F}_k + \frac{1}{\lambda} B_k \right]^{-1} \hat{g}_k$.*

Proof. By induction. See the *supplementary material.* □

By comparing *Propositions 4.1* and *4.2*, we see that, unlike SO-KLD-WRM, the Q-KLD-WRM step deviates significantly from the FISHER-WOQM step (also the EA-NG step). Note that setting $B_k = 0$ in *Proposition 4.2* gives *Proposition 4.1*. That is, SO-KLD-WRM is a particular case of Q-KLD-WRM (with $B_k \equiv 0$).

Note that $\{\hat{g}_k\}$ could explode with k, leading to divergence. A sufficient condition for $\{\hat{g}_k\}$ to stay bounded is $\rho \left\| I - \hat{M}_k \right\|_2 \leq \delta, \forall k \in \mathbb{Z}^+$, with $\delta \in (0, 1)$ and that $\|\nabla f(\theta)\|_2 \leq K_g, \forall \theta$. Under this condition, one can see that $\|\hat{g}_k\|_2 \leq \|g_k - \rho g_{k-1}\|_2 + \delta \|\hat{g}_{k-1}\|_2$. Applying this inequality recursively, and noting that $\|g_k - \rho g_{k-1}\|_2 \leq (1 + \rho)K_g$, we get that our sufficient condition yields $\|\hat{g}_k\|_2 \leq \frac{2}{1-\delta} K_g, \forall k \in \mathbb{Z}^+$. The condition $\rho \left\| I - \hat{M}_k \right\|_2 = \rho \left\| I - \left[I + \frac{1}{\lambda} B_k \bar{F}_k^{-1} \right]^{-1} \right\|_2 \leq \delta$ can always be achieved in practice, since taking $\lambda \to \infty$ or $\rho \to 0$ gives $\delta = 0$.

In a similar fashion to the role of *Proposition 4.1*, the role of *Proposition 4.2* (besides highlighting any similarity or dissimilarity to EA-NG) is to give a tractable algorithm. Again, as with SO-KLD-WRM, Q-KLD-WRM is not tractable if we implement it by solving (28) directly for the same reasons. On the other hand, implementing *Proposition 4.2* requires simultaneous storage of *at most* 3 matrices and 3 gradient-shaped vectors at any one point in time - that is, the storage cost does not explode with k. However, because we now have two different sets of matrices involved: $\{F_k\}$ and $\{B_k\}$, the situation is more subtle.

In particular, if B_k and F_k[14] are block-diagonal, then we can see that all involved matrices are block-diagonal, and thus tractably storable and invertable (eg. choose $B_k = F_k$, and approximate $F_k \approx \hat{F}_k^{(KFAC)}$[15]). On the contrary, B_k and F_k might be tractably storable and invertible in isolation, but if they have different structures, computing s_k from *Proposition 4.2* might be intractable.

Q-KLD-WRM in practice In practice, we can employ one of two approximations for B_k. The first option is to use a *BFGS* approximation ([11], [12]). The second option is to replace B_k by the K-FAC matrix. This latter approximation can be justified through the qualified equivalence between the *Fisher* and the *Generalized Gauss-Newton (GGN)* matrix, the latter of which is an approximation to the Hessian. However, in order for the qualified equivalence to hold, we need our predictive distribution $p(y|h_\theta(x))$ to be in the exponential family with natural parameters $h_\theta(x)$ (thinking of each conditional $y|x$ as a different distribution with its own parameters; see [8]). This qualified equivalence holds for most practically used models (see [8]), so is not of huge practical concern.

As we have discussed, it is not obvious how one could get a tractable algorithm from *Proposition 4.2* if the structures of B_k and F_k are dissimilar. Thus, in this paper we focus on instantiations where $B_k := F_k \approx F_k^{(KFAC)}$. In our experiments, we will choose $p(y|h_\theta(x))$ s.t. the qualified equivalence between Fisher and GGN matrix holds, and thus use $B_k \approx F_k^{(KFAC)}$, $F_k \approx F_k^{(KFAC)}$. As is typical with K-FAC, we also maintain an EA for the Kronecker factors, rather than for the block-diagonal matrix. With these choices, the Q-KLD-WRM step can be efficiently computed form *Proposition 4.2* (see the *supplementary material*).

4.4   Third KLD-WRM Instantiation: QE-KLD-WRM

The *Quadratic Objective Approximation Exact SKL*[16] KLD-WRM (QE-KLD-WRM) is the final instantiation of LM-KLD-WRM which we propose here. The QE-KLD-WRM step s_k (at θ_k) solves

$$\min_s g_k^T s + \frac{1}{2} s^T B_k s + \lambda \sum_{i=0}^{k} \zeta(i) \rho^{k-i} \mathbb{D}_{KL}(\theta_i, \theta_k + s), \qquad (31)$$

where B_k is a curvature matrix at θ_k, treated the same as we did in Q-KLD-WRM.

Practical QE-KLD-WRM for Regression To be able to work with the exact SKL, we restrict ourselves to a class of $p_\theta(y|x)$ models where the SKL can be expressed in terms of euclidean norms of differences in the network output space. Consider equation (4). For predictive distributions of the form (which are used in regression)

$$p(y|h_{\theta_k}(x_j)) = \mathcal{N}(y|h_{\theta_k}(x_j); I), \qquad (32)$$

---

[14] Of course, $\bar{F}_k$ will have the same structure as F_k.

[15] Using K-FAC further reduces the storage and computation cost through the K-factors.

[16] Note there is no *tilde* on $\mathbb{D}$ in (31

we have a special form for $D_{KL}\big(p(y|h_{\theta_1}(x_j)), p(y|h_{\theta_2}(x_j))\big)$, namely

$$\mathbb{D}_{KL}\big(p(y|h_{\theta_1}(x_j)), p(y|h_{\theta_2}(x_j))\big) = \frac{1}{2}\left\|h_{\theta_1}(x_j) - h_{\theta_2}(x_j)\right\|_2^2. \tag{33}$$

See the *supplementary material* for derivation details. Note that predictive distributions of the form (32) are the most frequently used in practice with regression. For this choice of predictive distribution, $\mathbb{D}_{KL}(\theta_1, \theta_2)$ becomes

$$\mathbb{D}_{KL}(\theta_1, \theta_2) = \frac{1}{2N}\sum_{j=0}^{N}\left\|h_{\theta_1}(x_j) - h_{\theta_2}(x_j)\right\|_2^2. \tag{34}$$

Thus, the QE-KLD-WRM step solves

$$\min_s g_k^T s + \frac{1}{2}s^T B_k s + \frac{\lambda}{2N}\sum_{i=0}^{k}\zeta(i)\rho^{k-i}\sum_{j=0}^{N}\left\|h_{\theta_i}(x_j) - h_{\theta_k+s}(x_j)\right\|_2^2. \tag{35}$$

Practical QE-KLD-WRM for Classification A similar practical computation of $\mathbb{D}_{KL}$ is also available for classification (see the *supplementary material*).

QE-KLD-WRM in Practice While the KL-regularization term in (35) is in principle computable, the amount of associated storage would explode with k (storing old parameters). To bypass this problem, we have two options. We could either choose to discard very old regularization terms, or model them through the approximation (21) (as we did for Q-KLD-WRM, but now only do so for old terms). The latter approach, while convenient, is not well principled - we should really use second-order approximation when we are close (so for recent distributions), not when we are far away. This can be improved, but is left as future work. In this paper, we focus on implementations that use the former approach. Note that we now need to iteratively solve (35) (for eg. with SGD), and get an approximate QE-KLD-WRM step. Further note that the Q-KLD-WRM step given by *Proposition 4.2* is an approximation of the solution to (35), where the SKL-divergence is approximated by its second-order Taylor expansion (21). Thus, the Q-KLD-WRM step is a good initial guess for (35), and we exploit this fact in practice.

Extension to Variable Stepsize For WOQM, SO-KLD-WRM and Q-KLD-WRM, we have so far considered only cases when the *"step-size"* $1/\lambda$ is fixed across different locations θ_k. Indeed, our established equivalence between FISHER-WOQM and EA-NG holds for fixed λ only. However, we may desire variable $\lambda \leftarrow \lambda^{(k)}$ in our KLD-WRM algorithms. To incorporate variable $\lambda^{(k)}$ in Q-KLD-WRM, all one has to do is to merely replace (30) with $\hat{g}_{k+1} = g_{k+1} + (\lambda^{(k+1)}/\lambda^{(k)})\rho[\hat{g}_k - g_k - \hat{M}_k\hat{g}_k]$ and all λ's in *Proposition 4.2* with $\lambda^{(k)}$. The version of *Proposition 4.2* which includes variable $\lambda^{(k)}$ can be found in the *supplementary material*. *Proposition 4.1* with variable $\lambda^{(k)}$ then follows trivially as a particular case with $B_k \equiv 0$.

4.5   Connection with Second Order Methods (SOM)

The particular KLD-WRM instantiation in equation (28), most simply illustrates the connection between our proposed class of algorithms and SOM. Setting $\lambda \leftarrow 0$ in (28) reverses Q-KLD-WRM back to a simple second order algorithm. More generally, from (23) we see that relative to SOM, KLD-WRM generalizes the local second order model to an arbitrary model that may also include previous information (in principle), and more importantly, it adds a wake of $\mathbb{D}_{KL}$ regularization. The connection between NG and Generalized Gauss-Newton can be found in [8].

5   Numerical Results

We compare our proposed KLD-WRM instantiations (SO, Q and QE) with K-FAC, on the MNIST classification problem. We investigated 4 different hyper-parameter settings for each solver, but only present the best ones here. The complete results, implementation details, as well as more in depth discussion can be found in *Section 8* of the *supplementary material*. *Figures 1 and 2* show *test loss* and *test accuracy* for our considered solvers. Ten runs were performed for each solver. *Table 1* shows important summary statistics of these result.

Hyper-parameters: The values of ρ and λ are specified in *Figures 1* and *2*. We used $\zeta(i) = \kappa(i)/330$ for QE, and $\zeta(i) = \kappa(i)$ for the other solvers, because the QE estimates of the SKL term were larger. The QE-specific hyper-parameters are: ω/λ – the learning rate of the inner SGD solving (35), N_{IS} – the number of inner SGD steps per iteration, and N_{CAP} – the total number of networks stored. For the results presented here, these were set to $7 \cdot 10^{-4}$, 10 and 4 respectively.

Test Accuracy and Loss: All KLD-WRM variants exceed 97.5% mean test accuracy, outperforming K-FAC by about 1.5%. The SD is 4-5 times lower for KLD-WRM variants, which is desirable. Analogous observations hold for the test loss. Since higher variance may more frequently yield *"favorable"* outliers, we show relevant metrics for this aspect in *Columns 2-6* of *Table 1*. We see that even from the *favorable outliers* point of view, KLD-WRM is mostly preferable.

Robustness and Overall Performance: *If we can only run the training a few times (perhaps once)*, it is preferable (in terms of epochs; both from a *test accuracy* and *test-loss* point of view) to use KLD-WRM rather than K-FAC.

Table 1. MNIST results summary. SO, Q and QE refer to KLD-WRM variants. *"Accuracy"* and *"loss"* are the test-set ones. μ_{acc} and σ_{acc} are the mean and SD of the empirical distribution of *accuracy* at the end of epoch 50. Notation is analogous for μ_{loss} and σ_{loss}, which refer to the *loss*. $\mathcal{N}_{\mathcal{C}}$ is the number of runs for which condition $\mathcal{C}$ is satisfied.

|	$\mathcal{N}_{acc \geq 98\%}$	$\mathcal{N}_{acc > 98\%}$	$\mathcal{N}_{acc \geq 98.5\%}$	$\mathcal{N}_{loss \leq 0.25}$	$\mathcal{N}_{loss \leq 0.2}$	μ_{acc}	σ_{acc}	μ_{loss}	σ_{loss}
K-FAC	3	3	**1**	5	**2**	96.19%	3.2%	0.26	0.10
SO	4	2	0	7	0	97.60%	0.85%	0.25	0.04
Q	5	**4**	0	4	0	97.69%	0.69%	0.26	0.04
QE	**8**	**7**	**2**	**9**	0	**98.01%**	**0.64%**	**0.23**	**0.02**

Fig. 1. MNIST test-loss results for K-FAC, and our three KLD-WRM variants.

That is because all our KLD-WRM variants more robustly achieve good test metrics. Conversely, *if we can run the training many times* (and choose the best run) K-FAC's large variance *may* eventually play to our advantage (not guaranteed).
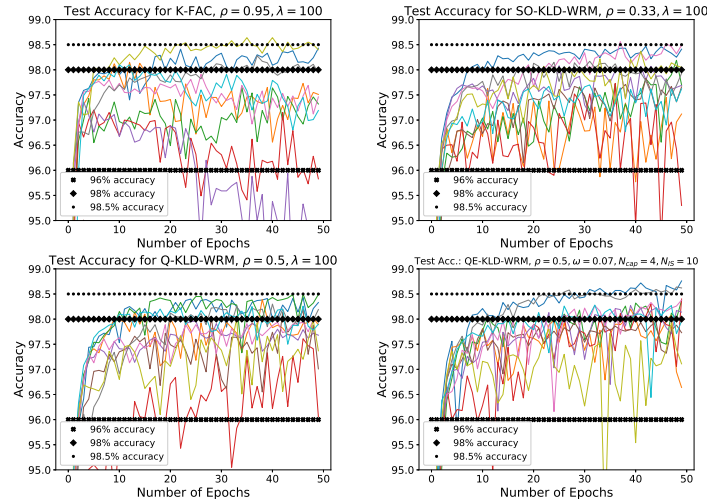
KLD-WRM variant selection: SO-KLD-WRM and Q-KLD-WRM have virtually the same computation cost per epoch as K-FAC. Conversely, QE-KLD-WRM has the same *data acquisition* and *linear algebra* costs as K-FAC, but 3-10 times higher oracle cost (fwd. and bwd. pass cost), owing to approximately solving (35). Thus, when *data cost is relatively low*, SO-KLD-WRM and Q-KLD-WRM will be preferable, as they will have the smallest wall-time per epoch (while having almost the same performance per epoch as QE-KLD-WRM). Conversely, when *data cost dominates*, all 4 solvers will have the same wall-time per epoch. In this case, QE-KLD-WRM is preferable as it gives the best performance per epoch[17].

6   Conclusions and Future Work

We established an equivalence between EA-CM algorithms (typically used in ML) and WOQM algorithms (which we defined in *Section 3*). The equivalence revealed what EA-CM algorithms are doing from a model-based optimization point of view. Generalizing from WOQM, we defined a broader class of algorithms in *Section 4*: KLD-WRM. We then focused our attention on a different subclass of KLD-WRM, LM-KLD-WRM, and provided three practical instantiations of it. Numerical results on MNIST showed that performance-metrics distributions have better mean and lower variance for our KLD-WRM algorithms, indicating they are preferable to K-FAC in practice due to higher robustness.

Future work: (a) KLD-WRM for VI BNNs and RL; (b) convergence theory; (c) investigate Q and QE variants when B_k and F_k have different structures; (d) consider KLD-WRM algorithms outside the LM-KLD-WRM subfamily (include info. at $\{\theta_j\}_{j<k}$ in $M(s; \mathcal{F}_k)$; see (23)); (e) consider arbitrary $\xi \in \mathbb{R}$ (see footnote 12).

---

[17] Codes repo: https://github.com/ConstantinPuiu/Rethinking-EA-of-the-Fisher

Fig. 2. MNIST test-accuracy results for K-FAC, and our three KLD-WRM variants.

References

1. Martens, J.; Grosse, R. Optimizing neural networks with Kronecker-factored approximate curvature, In: arXiv:1503.05671 (2015).
2. Yang, M.; Xu, D; Wen, Z.; Chen, M.; Xu, P. Sketchy empirical natural gradient methods for deep learning, In: arXiv:2006.05924 (2021).
3. Ba, J.; Kingma, D. Adam: A method for stochastic optimization, ICLR (2015).
4. LeCun, Y.; Bottou, L.; Orr, G.; Muller, K. Efficient backprop. Neural networks: Tricks of the trade, pages 546-546 (1998).
5. Schaul, T.; Zhang, S.; LeCun, Y. No more pesky learning rates. In ICML (2013).
6. Park, H.; Amari, S.-I.; Fukumizu, K. Adaptive natural gradient learning algorithms for various stochastic models. Neural Networks, 13(7):755-764 (2000).
7. Amari, S. I. Natural gradient works efficiently in learning, Neural Computation, 10(20), pp. 251-276 (1998).
8. Martens, J. New insights and perspectives on the natural gradient method, arXiv:1412.1193 (2020).
9. Osawa, K.; Yuichiro Ueno, T.; Naruse, A.; Foo, C.-S.; Yokota, R. Scalable and practical natural gradient for large-scale deep learning, arXiv:2002.06015 (2020).
10. Wu, Y.; Mansimov, E.; Grosse, R. B.; Liao, S.; Ba, J. Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation. In Advances in neural information processing systems, pages 5285-5294 (2017).
11. Bottou, L.; Curtis, F. E.; Nocedal, J. Optimization methods for large-scale machine learning (2018).
12. Goldfarb, D.; Ren, Y.; Bahamou, A. Practical Quasi-Newton methods for training deep neural networks, arXiv:2006.08877, (2021).

Supplementary Material: Rethinking Exponential Averaging of the Fisher

Constantin Octavian Puiu

1. Proof of Proposition 3.1

We reiterate the statement of *Proposition 3.1* and the *WoQM* definition for convenience. Recall that the *WoQM* step at θ_k solves

$$
\min_s s^T \left[\sum_{i=0}^{k} \rho^{k-i} \left(g_i + \lambda \kappa(i) B_i \sum_{j=i}^{k-1} s_j \right) \right]
$$
$$
+ \frac{\lambda}{2} s^T \left[\rho^k B_0 + (1-\rho) \sum_{i=1}^{k} \rho^{k-i} B_i \right] s, \tag{1}
$$

Proposition 3.1: Analytic Solution of *WoQM* step. *The WoQM step s_k at iterate θ_k can be expressed as*

$$
s_k = -\frac{1}{\lambda} \left(\rho^k B_0 + (1-\rho) \sum_{j=1}^{k} \rho^{k-j} B_j \right)^{-1} g_k, \quad \forall k \in \mathbb{N}. \tag{2}
$$

Proof. We use induction.

Induction check. First, check (2) holds for $k = 0$, $k = 1$ and $k = 2$. Since we have new sum terms starting to appear[1] at $k = 1$ and also at $k = 2$ (and keep on being present thereafter), we also need to check for $k = 1$.

For $k = 0$ we have the model in (1) reducing to

$$
\min_s s^T g_0 + \frac{\lambda}{2} s^T B_0 s, \tag{3}
$$

which has solution

$$
s_0 = -\frac{1}{\lambda} B_0^{-1} g_0, \tag{4}
$$

which satisfies (2).

For $k = 1$ we have the model in (1) reducing to

$$
\min_s s^T [\rho g_0 + g_1 + \rho \lambda B_0 s_0] + \frac{\lambda}{2} s^T [\rho B_0 + (1-\rho) B_1] s. \tag{5}
$$

But, $\rho g_0 + \rho \lambda B_0 s_0 = 0$ by (4). Thus, we have

$$
s_1 = -\frac{1}{\lambda} [\rho B_0 + (1-\rho) B_1]^{-1} g_1, \tag{6}
$$

---

[1] Because $\kappa(i) = 1$ for $i = 0$, and $\kappa(i) = 1 - \rho$ for $i \geq 1$ but the terms $(1-\rho) B_i \sum_{j=i}^{k-1} s_j$ only start appearing after $k \geq 2$.

which satisfies (2).

For $k = 2$, we have the model in (1) reducing to

$$
\min_s \frac{\lambda}{2} s^T [\rho^2 B_0 + (1-\rho)\rho B_1 + (1-\rho) B_2] s +
$$
$$
s^T \left[\rho^2 g_0 + \rho g_1 + g_2 + \rho^2 \lambda B_0 (s_0 + s_1) + \rho \lambda (1-\rho) B_1 s_1 \right]. \tag{7}
$$

Now, $\rho^2 g_0 + \rho^2 \lambda B_0 s_0 = 0$ using (4). We also have $\rho g_1 + \rho \lambda (\rho B_0 + (1-\rho) B_1) s_1 = 0$ by (6). Thus, we get

$$
s_2 = -\frac{1}{\lambda} [\rho^2 B_0 + (1-\rho)\rho B_1 + (1-\rho) B_2]^{-1} g_2, \tag{8}
$$

which satisfies (2).

Induction main body. Assume for the **inductive hypothesis** that we have for some $k \geq 2$

$$
s_l = -\frac{1}{\lambda} \left(\rho^l B_0 + (1-\rho) \sum_{j=1}^{l} \rho^{l-j} B_j \right)^{-1} g_l, \quad \forall l \leq k \in \mathbb{N}. \tag{9}
$$

We **need to prove** that at $k + 1$, for all $k \geq 2$, we have

$$
s_{k+1} = -\frac{1}{\lambda} \left(\rho^{k+1} B_0 + (1-\rho) \sum_{j=1}^{k+1} \rho^{k+1-j} B_j \right)^{-1} g_{k+1}, \tag{10}
$$

and since the inductive check holds at $k \in \{0, 1, 2\}$, this would imply (2).

To begin, consider (1) evaluated at $k + 1$. We have that the *WoQM* step at $k + 1$, s_{k+1} is given by

$$
\min_s \frac{\lambda}{2} s^T \left[\rho^{k+1} B_0 + (1-\rho) \sum_{i=1}^{k+1} \rho^{k+1-i} B_i \right] s
$$
$$
+ s^T \left[\sum_{i=0}^{k+1} \rho^{k+1-i} \left(g_i + \lambda \kappa(i) B_i \sum_{j=i}^{k} s_j \right) \right]. \tag{11}
$$

Let us now consider the coefficient of the linear term in

(11). It reads

$$\sum_{i=0}^{k+1} \rho^{k+1-i}\left(g_i + \lambda\kappa(i)B_i\sum_{j=i}^{k} s_j\right) =$$

$$g_{k+1} + \rho\left[\sum_{i=0}^{k} \rho^{k-i}\left(g_i + \lambda\kappa(i)B_i\sum_{j=i}^{k} s_j\right)\right] =$$

$$g_{k+1} + \rho\left[\sum_{i=0}^{k} \rho^{k-i}g_i + \lambda\sum_{i=0}^{k}\kappa(i)\rho^{k-i}\left(B_i\sum_{j=i}^{k} s_j\right)\right].$$

$$(12)$$

Now we prove that the bracket on the third line of (12) is zero. Once we have that, we can very easily see from (11) that we have the desired result. To show the term is zero, consider

$$\lambda\sum_{i=0}^{k}\kappa(i)\rho^{k-i}\left(B_i\sum_{j=i}^{k} s_j\right) = \lambda\sum_{i=0}^{k}\sum_{j=i}^{k}\kappa(i)\rho^{k-i}B_i s_j$$

$$(13)$$

Now, inverting the i and j sums, we get

$$\lambda\sum_{i=0}^{k}\kappa(i)\rho^{k-i}\left(B_i\sum_{j=i}^{k} s_j\right) = \lambda\sum_{j=0}^{k}\sum_{i=0}^{j}\kappa(i)\rho^{k-i}B_i s_j =$$

$$\lambda\sum_{j=0}^{k}\rho^{k-j}\left(\sum_{i=0}^{j}\rho^{j-i}\kappa(i)B_i\right)s_j.$$

$$(14)$$

Now, explicitly writing that $\kappa(0) = 1$ and $\kappa(i) = 1 - \rho$ for all $i \geq 1$, we get that

$$\left(\sum_{i=0}^{j}\rho^{j-i}\kappa(i)B_i\right)s_j = \left(\rho^j B_0 + (1-\rho)\sum_{i=1}^{j}\rho^{j-i}B_i\right)s_j.$$

$$(15)$$

Substituting the inductive hypothesis (9) into (15), we get

$$\left(\sum_{i=0}^{j}\rho^{j-i}\kappa(i)B_i\right)s_j = -\frac{1}{\lambda}g_j \qquad (16)$$

Plugging (16) into (14), we see that

$$\lambda\sum_{i=0}^{k}\kappa(i)\rho^{k-i}\left(B_i\sum_{j=i}^{k} s_j\right) = -\sum_{j=0}^{k}\rho^{k-j}g_j. \qquad (17)$$

Finally, plugging (17) into the last line of (12), we see that

$$\sum_{i=0}^{k+1}\rho^{k+1-i}\left(g_i + \lambda B_i\sum_{j=i}^{k} s_j\right) = g_{k+1}, \qquad (18)$$

which means the linear coefficient in (11) is just g_{k+1}. Thus, differentiating (11) and setting to zero gives us the global solution[2] of the *WoQM* model as

_____

[2]Since B_k are positive definite by assumption.

$$s_{k+1} = -\frac{1}{\lambda}\left[\rho^{k+1}B_0 + (1-\rho)\sum_{i=1}^{k+1}\rho^{k+1-i}B_i\right]^{-1}g_{k+1}.$$

$$(19)$$

Equation (19) is the same as (10), which completes the proof $\square$.

2. Proof of Proposition 4.1

We reiterate the statement of *Proposition 4.1* and the *SO-KLD-WRM* definition for convenience. Recall that the *SO-KLD-WRM* step at θ_k solves

$$\min_s g_k^T s + \lambda\sum_{i=0}^{k}\kappa(i)\rho^{k-i}\tilde{\mathbb{D}}_{KL}(\theta_i||\theta_k + s), \qquad (20)$$

That is, the *SO-KLD-WRM* step at θ_k solves

$$\min_s s^T\left[g_k + \sum_{i=0}^{k-1}\rho^{k-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k-1} s_j\right)\right]$$

$$+ \frac{\lambda}{2}s^T\left[\sum_{i=0}^{k}\kappa(i)\rho^{k-i}F_i\right]s. \qquad (21)$$

Proposition 4.1: Analytic Solution of *SO-KLD-WRM* step. *The SO-KLD-WRM step s_k at iterate θ_k can be expressed as*

$$s_k = -\frac{1}{\lambda}\left(\rho^k F_0 + (1-\rho)\sum_{j=1}^{k}\rho^{k-j}F_j\right)^{-1}[g_k - \rho g_{k-1}],$$

$$(22)$$

$\forall k \in \mathbb{N}$, *where we set $g_{-1} := 0$ by convention.*

Proof. We use induction. The proof is very similar to the prof of *Proposition 3.1*.

Induction check. First, check (22) holds for $k = 0$, $k = 1$ and $k = 2$. Since we have new sum terms starting to appear[3] at $k = 1$ and also at $k = 2$ (and keep on being present thereafter), we also need to check for $k = 1$.

For $k = 0$ we have the model in (21) reducing to

$$\min_s s^T g_0 + \frac{\lambda}{2}s^T F_0 s, \qquad (23)$$

which has solution

$$s_0 = -\frac{1}{\lambda}F_0^{-1}g_0, \qquad (24)$$

which satisfies (22).

_____

[3]Because $\kappa(i) = 1$ for $i = 0$, and $\kappa(i) = 1 - \rho$ for $i \geq 1$ but the terms $(1-\rho)B_i\sum_{j=i}^{k-1} s_j$ only start appearing after $k \geq 2$.

For $k = 1$ we have the model in (1) reducing to

$$\min_s s^T[g_1 + \rho\lambda F_0 s_0] + \frac{\lambda}{2}s^T[\rho F_0 + (1 - \rho)F_1]s. \quad (25)$$

But, $g_1 + \rho\lambda B_0 s_0 = g_1 - \rho g_0$ by (24). Thus, we have

$$s_1 = -\frac{1}{\lambda}[\rho F_0 + (1 - \rho)F_1]^{-1}[g_1 - \rho g_0], \quad (26)$$

which satisfies (22).

For $k = 2$, we have the model in (1) reducing to

$$\min_s \frac{\lambda}{2}s^T[\rho^2 F_0 + (1 - \rho)\rho F_1 + (1 - \rho)F_2]s+ \\ s^T\left[g_2 + \rho^2\lambda F_0(s_0 + s_1) + \rho\lambda(1 - \rho)F_1 s_1\right]. \quad (27)$$

Now, $\rho^2\lambda F_0 s_0 = -\rho^2 g_0$ using (24). We also have $\rho\lambda[\rho F_0 + (1 - \rho)F_1]s_1 = -\rho[g_1 - \rho g_0]$ by (22). Thus, we get that the linear term of (27) is $s^T[g_2 - \rho g_1]$. That is,

$$s_2 = -\frac{1}{\lambda}[\rho^2 F_0 + (1 - \rho)\rho F_1 + (1 - \rho)F_2]^{-1}[g_2 - \rho g_1], \quad (28)$$

which satisfies (22).

Induction main body. Assume for the **inductive hypothesis** that we have for some $k \geq 2$

$$s_l = -\frac{1}{\lambda}\left(\sum_{j=0}^{l}\kappa(j)\rho^{l-j}F_j\right)^{-1}[g_l - \rho g_{l-1}], \quad \forall l \leq k. \quad (29)$$

We **need to prove** that at $k + 1$, for all $k \geq 2$, we have

$$s_{k+1} = -\frac{1}{\lambda}\left(\sum_{j=0}^{k+1}\kappa(j)\rho^{k+1-j}F_j\right)^{-1}[g_{k+1} - \rho g_k], \quad (30)$$

and since the inductive check holds at $k \in \{0, 1, 2\}$, this would imply (22).

To begin, consider (21) evaluated at $k+1$. We have that the *SO-KLD-WRM* step at $k + 1$, s_{k+1} is given by

$$\min_s \frac{\lambda}{2}s^T\left[\rho^{k+1}F_0 + (1 - \rho)\sum_{i=1}^{k+1}\rho^{k+1-i}F_i\right]s \\ + s^T\left[g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\right)\right]. \quad (31)$$

Let us now consider the coefficient of the linear term in (31). It reads

$$g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\right) = \\ g_{k+1} + \lambda\sum_{i=0}^{k}\sum_{j=i}^{k}\rho^{k+1-i}\kappa(i)F_i s_j = \quad (32) \\ g_{k+1} + \lambda\sum_{j=0}^{k}\left[\left(\sum_{i=0}^{j}\rho^{k+1-i}\kappa(i)F_i\right)s_j\right].$$

In the last line, we interchanged the summation over i and j. But the square bracket in the second term of (32) is given by

$$\left(\sum_{i=0}^{j}\rho^{k+1-i}\kappa(i)F_i\right)s_j = \rho^{k+1-j}\left(\sum_{i=0}^{j}\rho^{j-i}\kappa(i)F_i\right)s_j \quad (33)$$

Now, plugging the inductive hypothesis (29) in (33) - which we can do as the outer-sum over j in (32) is up until k - we get

$$\left(\sum_{i=0}^{j}\rho^{k+1-i}\kappa(i)F_i\right)s_j = -\frac{1}{\lambda}\rho^{k+1-j}[g_j - \rho g_{j-1}] \quad (34)$$

Finally, plugging (34) into (32), we get

$$g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\right) = \\ g_{k+1} + \lambda\sum_{j=0}^{k}\left[-\frac{1}{\lambda}\rho^{k+1-j}[g_j - \rho g_{j-1}]\right] = g_{k+1} - \rho g_k. \quad (35)$$

To get the last equality in (35) we solved the telescopic sum, and in doing that we recalled our convention $g_{-1} = 0$.

Finally, plugging (35) into (31), we see that

$$s_{k+1} = -\frac{1}{\lambda}\left[\sum_{i=0}^{k+1}\kappa(i)\rho^{k+1-i}F_i\right]^{-1}[g_{k+1} - \rho g_k]. \quad (36)$$

Equation (36) is the same as (30), which completes the proof $\square$.

3. Proof of Proposition 4.2

We reiterate the *Q-KLD-WRM* step definition as well as *Proposition 4.2* here for convenience. The *Q-KLD-WRM* step at θ_k is the solution to

$$\min_s g_k^T s + \frac{1}{2}s^T B_k s + \lambda\sum_{i=0}^{k}\kappa(i)\rho^{k-i}\tilde{\mathbb{D}}_{KL}(\theta_i||\theta_k + s), \quad (37)$$

That is, the *Q-KLD-WRM* step at θ_k solves

$$\min_s s^T\left[g_k + \sum_{i=0}^{k-1}\rho^{k-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k-1}s_j\right)\right] \\ + \frac{\lambda}{2}s^T\left[\sum_{i=0}^{k}\kappa(i)\rho^{k-i}F_i + \frac{1}{\lambda}B_k\right]s. \quad (38)$$

Of course, we have $\bar{F}_k = \sum_{i=0}^{k}\kappa(i)\rho^{k-i}F_i$, and we will use that notation most of the times.

Proposition 4.2 is given again below.

Proposition 4.2: Analytic Solution of *Q-KLD-WRM* step.
The Q-KLD-WRM step s_k at location θ_k is given by the solution to the problem

$$\min_s s^T \hat{g}_k + \frac{\lambda}{2} s^T \left[\bar{F}_k + \frac{1}{\lambda} B_k\right] s \qquad (39)$$

where $\hat{g}_k$ is given by the one-step recursion

$$\hat{g}_{k+1} = g_{k+1} + \rho(I - \hat{M}_k)\hat{g}_k - \rho g_k, \quad \forall k \in \mathbb{Z}^+, \quad (40)$$

with $\hat{g}_0 := g_0$, $\bar{F}_k := \sum_{i=0}^k \kappa(i)\rho^{k-i} F_i$, and

$$\hat{M}_k := \left[I + \frac{1}{\lambda} B_k \bar{F}_k^{-1}\right]^{-1}. \qquad (41)$$

That is, the Q-KLD-WRM step is formally given by

$$s_k = -\frac{1}{\lambda}\left[\bar{F}_k + \frac{1}{\lambda} B_k\right]^{-1} \hat{g}_k. \qquad (42)$$

Proof. By induction. The induction is mainly focused around finding a form for the linear term of the *Q-KLD-WRM* formulation at arbitrary k. The step analytic formula then follows easily.

Induction check. For $k = 0$, we have the model in (38) reducing to

$$\min_s s^T g_0 + \frac{\lambda}{2} s^T \left[\bar{F}_0 + \frac{1}{\lambda} B_0\right] s \qquad (43)$$

Now, noting that $\hat{g}_0 = g_0$ by definition, we see that the *Q-KLD-WRM* linear[4] term obeys *Proposition 4.2* at $k = 0$. Thus,

$$s_0 = -\frac{1}{\lambda}\left[\bar{F}_0 + \frac{1}{\lambda} B_0\right]^{-1} \hat{g}_0. \qquad (44)$$

For $k = 1$, we have the model in (38) reducing to

$$\min_s s^T [g_1 + \rho\lambda F_0 s_0] + \frac{\lambda}{2} s^T \left[\bar{F}_1 + \frac{1}{\lambda} B_1\right] s. \qquad (45)$$

Now, by using (44), we have that

$$\rho\lambda F_0 s_0 = -\rho[F_0^{-1}]^{-1}\left[\bar{F}_0 + \frac{1}{\lambda} B_0\right]^{-1} \hat{g}_0 \qquad (46)$$

$$\rho\lambda F_0 s_0 = -\rho\left[\bar{F}_0 F_0^{-1} + \frac{1}{\lambda} B_0 F_0^{-1}\right]^{-1} \hat{g}_0 = -\rho\hat{M}_0 \hat{g}_0 \qquad (47)$$

where the latter equality holds by the definition of $\hat{M}_k$ evaluated at $k = 0$ and the fact that $\bar{F}_0 = F_0$. Thus,

$$g_1 + \rho\lambda F_0 s_0 = g_1 - \rho\hat{M}_0 \hat{g}_0 = g_1 + \rho(I - \hat{M}_0)\hat{g}_0 - \rho g_0. \qquad (48)$$

---

[4]The Quadratic term obviously also obeys *Propositon 4.2* formulation - but this fact follows trivially from the original *Q-KLD-WRM* definition, as we can see.

The latter equality in (48) holds because $g_0 = \hat{g}_0$. Consequently, using the recursive definition of $\hat{g}_{k+1}$ for $k = 0$, we get

$$g_1 + \rho\lambda F_0 s_0 = \hat{g}_1, \qquad (49)$$

which shows us that the *Q-KLD-WRM* linear term obeys the formulation in *Proposition 4.2*. Thus the whole *Q-KLD-WRM* formulation obeys formulation (39). In particular, we have

$$s_1 = -\frac{1}{\lambda}\left[\bar{F}_1 + \frac{1}{\lambda} B_1\right]^{-1} \hat{g}_1. \qquad (50)$$

For $k = 2$, we have the model in (38) reducing to

$$\min_s \frac{\lambda}{2} s^T \left[\bar{F}_2 + \frac{1}{\lambda} B_2\right] s+$$
$$s^T \left[g_2 + \rho^2\lambda F_0(s_0 + s_1) + \rho\lambda(1 - \rho)F_1 s_1\right]. \qquad (51)$$

The linear term in (51) can be rearranged as

$$s^T\left[g_2 + \rho^2\lambda F_0 s_0 + \rho\lambda[\rho F_0 + (1 - \rho)F_1]s_1\right] =$$
$$s^T\left[g_2 + \rho^2\lambda\bar{F}_0 s_0 + \rho\lambda\bar{F}_1 s_1\right] = \qquad (52)$$
$$s^T\left[g_2 + \rho(g_1 + \rho\lambda\bar{F}_0 s_0) + \rho\lambda\bar{F}_1 s_1 - \rho g_1\right],$$

where to get the second equality we used the definition of $\bar{F}_k$ applied at $k = 0$ and $k = 1$. We have already seen when doing the check for $k = 1$ that

$$g_1 + \rho\lambda\bar{F}_0 s_0 = \hat{g}_1 \qquad (53)$$

Thus, we only need to work out the $\rho\lambda\bar{F}_1 s_1$ term. By using (50), this reads

$$\rho\lambda\bar{F}_1 s_1 = -\rho[\bar{F}_1^{-1}]^{-1}\left[\bar{F}_1 + \frac{1}{\lambda} B_1\right]^{-1} \hat{g}_1 \qquad (54)$$

$$\rho\lambda\bar{F}_1 s_1 = -\rho\left[\bar{F}_1\bar{F}_1^{-1} + \frac{1}{\lambda} B_1\bar{F}_1^{-1}\right]^{-1} \hat{g}_1 = -\rho\hat{M}_1 \hat{g}_1 \qquad (55)$$

where the latter equality holds by the definition of M_k evaluated at $k = 1$. Thus, plugging in (55) and (53)[5] in the linear term expression (52), we get that the linear term reads

$$s^T\left[g_2 + \rho^2\lambda F_0 s_0 + \rho\lambda[\rho F_0 + (1 - \rho)F_1]s_1\right] =$$
$$s^T\left[g_2 + \rho(I - \hat{M}_1)\hat{g}_1 - \rho g_1\right] = s^T\hat{g}_2, \qquad (56)$$

where in the last equality we used the recursive definition of $\hat{g}_{k+1}$ from *Proposition 4.2*, evaluated at $k = 1$. Thus,

---

[5]Or indeed can use (49), we wrote it again for clarity.

(56) shows us that the *Q-KLD-WRM* linear term obeys the formulation in *Proposition 4.2*. Thus the whole *Q-KLD-WRM* formulation obeys formulation (39). In particular, if we now plug the linear term in (56) back into the full model (51), we get that the *Q-KLD-WRM* step satisfies

$$s_2 = -\frac{1}{\lambda}\left[\bar{F}_2 + \frac{1}{\lambda}B_2\right]^{-1}\hat{g}_2. \tag{57}$$

We are done with the induction check.

Induction main body. Assume for the **inductive hypothesis** that we have for some $k \geq 2$

$$s_l = -\arg\min_s s^T\hat{g}_l + \frac{\lambda}{2}s^T\left[\bar{F}_l + \frac{1}{\lambda}B_l\right]s, \forall l \in \{0,..,k\} \tag{58}$$

$$\hat{g}_l = g_l + \rho(I - \hat{M}_{l-1})\hat{g}_{l-1} - \rho g_{l-1}, \quad \forall l \in \{1,...,k\}, \tag{59}$$

with $\hat{g}_0 := g_0$, $\bar{F}_k := \sum_{i=0}^{k}\kappa(i)\rho^{k-i}F_i$, and

$$\hat{M}_l := \left[I + \frac{1}{\lambda}B_l\bar{F}_l^{-1}\right]^{-1}. \tag{60}$$

We **need to prove** that at $k+1$, for all $k \geq 2$, we have

$$\hat{g}_{k+1} = g_{k+1} + \rho(I - \hat{M}_k)\hat{g}_k - \rho g_k, \tag{61}$$

with $\hat{M}_k$ defined as in (60), and

$$s_{k+1} = \arg\min_s s^T\hat{g}_{k+1} + \frac{\lambda}{2}s^T\left[\bar{F}_{k+1} + \frac{1}{\lambda}B_{k+1}\right]s. \tag{62}$$

Since the inductive check holds at $k \in \{0, 1, 2\}$, this would imply *Proposition 4.2* holds.

To begin, consider (38) evaluated at $k+1$. We have that the *Q-KLD-WRM* step at $k+1$, s_{k+1} is the solution to

$$\min_s \frac{\lambda}{2}s^T\left[\bar{F}_{k+1} + \frac{1}{\lambda}B_{k+1}\right]s$$
$$+ s^T\left[g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\right)\right]. \tag{63}$$

Let us now consider the coefficient of the linear term in (63). It reads

$$g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\right) =$$

$$g_{k+1} + \lambda\sum_{i=0}^{k}\sum_{j=i}^{k}\rho^{k+1-i}\kappa(i)F_is_j =$$

$$g_{k+1} + \lambda\sum_{j=0}^{k}\left[\left(\sum_{i=0}^{j}\rho^{k+1-i}\kappa(i)F_i\right)s_j\right] =$$

$$g_{k+1} + \lambda\sum_{j=0}^{k}\rho^{k+1-j}\bar{F}_js_j = g_{k+1} + \rho\sum_{j=0}^{k}\rho^{k-j}\lambda\bar{F}_js_j. \tag{64}$$

To get from the third to the fourth line in (64) we used the definition of $\bar{F}_k$ evaluated at $k = j$. We can further write the linear coeffcient of (63) as

$$g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\right) =$$

$$g_{k+1} + \rho\sum_{j=0}^{k}\rho^{k-j}\lambda\bar{F}_js_j = \tag{65}$$

$$g_{k+1} + \rho\sum_{j=0}^{k-1}\rho^{k-j}\lambda\bar{F}_js_j + \rho\lambda\bar{F}_ks_k.$$

Now, considering the same linear term of our *Q-KLD-WRM* formulation, *but at arbitrary* $l \in [1, k] \cap \mathbb{Z}$ (rather than at $k + 1$) we get can apply[6] equation (65) to get

$$g_l + \sum_{i=0}^{l-1}\rho^{l-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\right) = g_l + \sum_{j=0}^{l-1}\rho^{l-j}\lambda\bar{F}_js_j. \tag{66}$$

But the inductive hypothesis, (58) and (59), tells us that this linear term must also be $\hat{g}_l$. That is, we have

$$\hat{g}_l = g_l + \sum_{j=0}^{l-1}\rho^{l-j}\lambda\bar{F}_js_j, \tag{67}$$

and therefore

$$\sum_{j=0}^{l-1}\rho^{l-j}\lambda\bar{F}_js_j = \hat{g}_l - g_l. \tag{68}$$

In particular (68) must also hold for $l = k$, because the induction hypothesis holds up until $\hat{g}_k$. We have

$$\sum_{j=0}^{k-1}\rho^{k-j}\lambda\bar{F}_js_j = \hat{g}_k - g_k, \tag{69}$$

Now, plugging (69) into (65), we have

$$g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\right) = \tag{70}$$
$$g_{k+1} + \rho(\hat{g}_k - g_k) + \rho\lambda\bar{F}_ks_k,$$

so all is left to is to compute the term $\rho\lambda\bar{F}_ks_k$. To do that, we use the induction hypothesis again. In particular, we use (58) for $l = k$ and solve exactly for s_k, to get

$$s_k = -\frac{1}{\lambda}\left[\bar{F}_k + \frac{1}{\lambda}B_k\right]^{-1}\hat{g}_k. \tag{71}$$

---

[6]We can do that because our calculations (63) - (65) hold for arbitrary $k \in \mathbb{Z}^+$, and so does the *Q-KLD-WRM* formulation (65).

Using (71), we have

$$\rho\lambda\bar{F}_k s_k = -\rho\big[\bar{F}_k^{-1}\big]^{-1}\Big[\bar{F}_k + \frac{1}{\lambda}B_k\Big]^{-1}\hat{g}_k = -\rho\hat{M}_k\hat{g}_k. \tag{72}$$

The last equality of (72) is obtained by combining the two inverses in the inverse of a single matrix, and then using the definition of $\hat{M}_k$.

Plugging (72) into (70) gives us that our linear term coefficient of the *Q-KLD-WRM* formulation at $k+1$ is given by

$$g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\Big(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\Big) = \\ g_{k+1} + \rho(\hat{g}_k - g_k) - \rho\hat{M}_k\hat{g}_k = \\ g_{k+1} + \rho(I - \hat{M}_k)\hat{g}_k - \rho g_k. \tag{73}$$

Equation (73), combined with (63) shows us that the induction target (61)-(62) is satisfied. Thus, the proof is complete $\square$.

4. The error between EA over K-Factors and EA over K-FAC approximated Fisher

Our discussion is based on the case when we hold an EA of the Fisher (or any approximation thereof). However, when using *K-FAC* (and also our *KLD-WRM* instantiations which use *K-FAC* as a platform), we do *not* hold an EA over the *K-FAC* approximated Fisher like

$$\bar{F}_k^{(KFAC)} = \sum_{i=0}^{k}\kappa(i)\rho^{k-i}F_k^{(KFAC)} \tag{74}$$

but rather hold an EA over each of the K-Factors as:

$$\bar{A}_k^{(l)} := \sum_{i=0}^{k}\kappa(i)\rho^{k-i}A_i^{(l)}, \\ \bar{\Gamma}_k^{(l)} := \sum_{i=1}^{k}\kappa(i)\rho^{k-i}\Gamma_i^{(l)}, \tag{75}$$

where $\kappa(i) = 1 - \rho$ if $i \geq 1$, and $\kappa(0) = 1$. This is standard implementation procedure as using (74) would undo any benefit from the Kronecker factorization approach.

Thus, strictly speaking, *K-FAC*, as well as the proposed *KLD-WRM* instantiation algorithms lay outside the KLD-WRM family. However, they do approximate algorithms in the *KLD-WRM* family. We now briefly analyze the error. Formal bounds are future work.

4.1. The Error when using EA over Kronecker Factors

Let us focus on a single layer l. We drop the layer index superscript for convenience. The error between using an

EA over the K-FAC approximated Fisher and using an EA over the Kronecker factors is

$$\text{err}_k = \sum_{i=0}^{k}\rho^{k-i}\kappa(i)A_i \otimes G_i \\ - \Big(\sum_{i=0}^{k}\rho^{k-i}\kappa(i)A_i\Big) \otimes \Big(\sum_{j=0}^{k}\rho^{k-j}\kappa(j)G_j\Big) \tag{76}$$

Using the distributive property of the Kronecker Factor, we have

$$\text{err}_k = \sum_{i=0}^{k}\rho^{k-i}\kappa(i)A_i \otimes \Big(G_i - \sum_{j=0}^{k}\rho^{k-j}\kappa(j)G_j\Big) \\ = \sum_{j=0}^{k}\rho^{k-j}\kappa(j)\Big(A_j - \sum_{i=0}^{k}\rho^{k-i}\kappa(i)A_i\Big) \otimes G_j. \tag{77}$$

Thus, we see that if either one of the following holds: (1) A_k does not change much, (2) G_k does not change much, or (3) $\rho \approx 0$, then the error is small (going to zero if either A_i is constant, G_i is constant or $\rho \to 0$). While one may expect (1) and/or (2) to hold in practice for very small step-size (at least if the whole[7] data set is used to approximate the expectation), (3) is unrealistic. Further formal investigation of this aspect is required and will be future work.

5. Efficient Computation of Q-KLD-WRM Step when $B_k = F_k^{(KFAC)}$

5.1. Computation of $\hat{g}_{k+1}$

To compute the *Q-KLD-WRM* step (s_{k+1}) using *Proposition 4.2*, we must first compute

$$\hat{g}_{k+1} = g_{k+1} + \rho(I - \hat{M}_k)\hat{g}_k - \rho g_k, \tag{78}$$

where

$$\hat{M}_k := \Big[I + \frac{1}{\lambda}B_k\bar{F}_k^{-1}\Big]^{-1}. \tag{79}$$

Clearly, the bottleneck lies in computing

$$\hat{M}_k\hat{g}_k = \Big[I + \frac{1}{\lambda}B_k\bar{F}_k^{-1}\Big]^{-1}\hat{g}_k, \tag{80}$$

as the remaining operations in (78) are just vector additions. First, note that $\hat{M}_k$ can be written as

$$\hat{M}_k = \bar{F}_k\Big[\bar{F}_k + \frac{1}{\lambda}B_k\Big]^{-1} \tag{81}$$

---

[7]To avoid changes in A_i and G_i due to picking different batches, in which case the change will only be determined by the change in θ. Of course, when using a batch estimate to A_i and G_i, different batches give different mean estimations. However this change would merely be induced by sub-sampling, and would not be present in the exact quantities (or the distribution of the estimates). It is only the change in the distribution of estimates that is of interest, rather than the change in their realization.

We hold an EA of the Kronecker-factors, rathar than an EA of the K-FAC matrix, which means $\bar{F}_k$ is of the form

$$\bar{F}_k = \text{block-diag}\left(\{\bar{A}_k^{(l)} \otimes \bar{\Gamma}_k^{(l)}\}_{l \in \{1,2,\dots N_l\}}\right), \quad (82)$$

where N_l is the number of layers $\bar{A}_k^{(l)}$, $\bar{\Gamma}_k^{(l)}$ are the exponential averages of the standard Kronecker factors (see [1]). Similarly, we chose $B_k = F_k^{(KFAC)}$, so

$$\bar{B}_k = \text{block-diag}\left(\{A_k^{(l)} \otimes \Gamma_k^{(l)}\}_{l \in \{1,2,\dots N_l\}}\right). \quad (83)$$

Thus (80) reduces to performing an inversion of the form (for each layer)

$$\left[\hat{M}_k \hat{g}_k\right]_l = (\bar{A}_k^{(l)} \otimes \bar{\Gamma}_k^{(l)})\left[\bar{A}_k^{(l)} \otimes \bar{\Gamma}_k^{(l)} + \frac{1}{\lambda}(A_k^{(l)} \otimes \Gamma_k^{(l)})\right]^{-1} v_l, \quad (84)$$

There are two ways to approach solving (84).

The first one is to recognize that computing

$$u_l = \left[\bar{A}_k^{(l)} \otimes \bar{\Gamma}_k^{(l)} + \frac{1}{\lambda}(A_k^{(l)} \otimes \Gamma_k^{(l)})\right]^{-1} v_l \quad (85)$$

amounts to solving a generalized Stein equation, and then use the methods proposed in Appendix B of [1] to get u_l. We then obtain $\left[\hat{M}_k \hat{g}_k\right]_l = (\bar{A}_k^{(l)} \otimes \bar{\Gamma}_k^{(l)})u_l = \Gamma_k^{(l)} U_l \bar{A}_k^{(l)}$, with $u_l = \text{vec}(U_l)$.

The second approach uses another approximation - in the same spirit as with replacing an EA for the Fisher with an EA over the Kronecker factors. For this approach, we note that when $B_k = F_k$, $\bar{F}_k + \frac{1}{\lambda}B_k = \sum_i^{k-1} \kappa(i)\rho^{k-i}F_i + \frac{(1-\rho)\lambda+1}{\lambda}F_k$ is just a re-weighting of the terms in the $\bar{F}_k$ sum. Thus, instead of carrying EA-averaged Kronecker factors for $\bar{F}_k$, and then try to work with $\bar{F}_k + \frac{1}{\lambda}B_k$, we can merely edit the weight of the last Kronecker factor and get a Kronecker-decomposed estimate for $\bar{F}_k + \frac{1}{\lambda}B_k$ directly. This would be much easier to work with. That is, we save EA-Kronecker factors for two previous levels, and use

$$\hat{A}_k^{(l)} = \rho\bar{A}_{k-1}^{(l)} + \frac{(1-\rho)\lambda+1}{\lambda}A_k^{(l)}, \quad (86)$$

$$\hat{\Gamma}_k^{(l)} = \rho\bar{\Gamma}_{k-1}^{(l)} + \frac{(1-\rho)\lambda+1}{\lambda}\Gamma_k^{(l)} \quad (87)$$

These should be contrasted with the running EA for Kronecker factors

$$\bar{A}_k^{(l)} = \rho\bar{A}_{k-1}^{(l)} + (1-\rho)A_k^{(l)}, \quad (88)$$

$$\bar{\Gamma}_k^{(l)} = \rho\bar{\Gamma}_{k-1}^{(l)} + (1-\rho)\Gamma_k^{(l)}. \quad (89)$$

We thus have the approximation[8]

$$\left[\hat{M}_k \hat{g}_k\right]_l \approx (\bar{A}_k^{(l)} \otimes \bar{\Gamma}_k^{(l)})\left[\hat{A}_k^{(l)} \otimes \hat{\Gamma}_k^{(l)}\right]^{-1} v_l. \quad (90)$$

We have

$$\begin{aligned}\left[\hat{M}_k \hat{g}_k\right]_l &\approx (\bar{A}_k^{(l)} \otimes \bar{\Gamma}_k^{(l)})\left[\hat{A}_k^{(l)} \otimes \hat{\Gamma}_k^{(l)}\right]^{-1} v_l = \\ &(\bar{A}_k^{(l)} \otimes \bar{\Gamma}_k^{(l)})\left[(\hat{A}_k^{(l)})^{-1} \otimes (\hat{\Gamma}_k^{(l)})^{-1}\right] v_l = \\ &\text{vec}\left[\Gamma_k^{(l)}(\hat{\Gamma}_k^{(l)})^{-1} V_l (\hat{A}_k^{(l)})^{-1}(A_k^{(l)})\right],\end{aligned} \quad (91)$$

where $v_l = \text{vec}(V_l)$ - i.e. $V_l \in \mathbb{R}^{n_l \times n_{l-1}}$ is the matrix format of v_l. That is, v_l maps to V_l in the same way $\text{vec}(W_l)$ maps to W_l, where W_l is the weight matrix[9] at layer l.

Equation (91) gives us an efficient way of estimating $\hat{M}_k g_k$. Thus, we can efficiently assemble $\hat{g}_{k+1}$ using (78).

5.2. Computing s_{k+1}

All that is now left, is to compute

$$s_{k+1} = -\frac{1}{\lambda}\left[\bar{F}_{k+1} + \frac{1}{\lambda}B_{k+1}\right]^{-1} \hat{g}_{k+1}. \quad (92)$$

In the previous subsection, we have already discussed about the two possible approaches one can take to invert $F_{k+1} + \frac{1}{\lambda}B_{k+1}$ when $B_{k+1} = F_k^{(KFAC)}$ and we approximate $\bar{F}_{k+1}$ by holding an EA for the Kronecker factors. Any of the two approaches apply directly. For example, taking the second approach, we get

$$[s_{k+1}]_l \approx = -\frac{1}{\lambda}\text{vec}\left[(\hat{\Gamma}_{k+1}^{(l)})^{-1}\hat{G}_{k+1}^{(l)}(\hat{A}_{k+1}^{(l)})^{-1}\right], \quad (93)$$

where $\text{vec}(\hat{G}_{k+1}^{(l)}) = [\hat{g}_{k+1}]_l$. That is, $[\hat{g}_{k+1}]_l$ maps to $\hat{G}_{k+1}^{(l)}$ in the same way $\text{vec}(W_l)$ maps to W_l, where W_l is the weight matrix at layer l.

Equation (93) gives us an efficient way to estimate the *Q-KLD-WRM* step provided we have $\hat{g}_{k+1}$ (which we can efficiently get as in (91)).

---

[8]Note that the approximation would be exact if holding EA for the Kronecker factors was the same as holding EA for the big matrix - which is never true in practice.

[9]As ever with K-FAC, bias is included in the weight matrix through appending a 1 at each layer output - see [1].

6. KL-Divergence in terms of Network Output

6.1. Regression with Gaussian Predictive Distributions of Fixed Variance

In this section, we derive the fact that the symmetric KL-Divergence is

$$\mathbb{D}_{KL}\big(p(y|h_{\theta_1}(x_j)), p(y|h_{\theta_2}(x_j))\big) = \frac{1}{2}\left\|h_{\theta_1}(x_j) - h_{\theta_2}(x_j)\right\|_2^2. \tag{94}$$

when the predictive distribution of our net is of the form

$$p(y|h_{\theta_k}(x_j)) = \mathcal{N}(y|h_{\theta_k}(x_j); I). \tag{95}$$

That is, $p(y|h_{\theta_k}(x_j))$ is a normal distribution with mean $h_{\theta_1}(x_j)$ and covariance matrix I. This boils down to computing the KL-divergence between two multivariate Gaussians having the same covariance matrix. Let us denote for convenience $\mu_{1,j} := h_{\theta_1}(x_j)$, $\mu_{2,j} := h_{\theta_2}(x_j)$. Then, we have the forward KL divergence

$$\mathbb{D}_{KL}\big(p(y|\mu_{1,j})\,\big|\big|\,p(y|\mu_{2,j})\big) =$$
$$\mathbb{E}_{y \sim \mathcal{N}(\mu_{1,j}|I)}\left[-\frac{1}{2}\left\|y - \mu_{1,j}\right\|_2^2 + \frac{1}{2}\left\|y - \mu_{2,j}\right\|_2^2\right] =$$
$$-\frac{1}{2}\mathbb{E}_{y \sim \mathcal{N}(\mu_{1,j}|I)}\left[\mu_{1,j}^T \mu_{1,j} + 2y^T(\mu_{2,j} - \mu_{1,j}) - \mu_{2,j}^T \mu_{2,j}\right] =$$
$$-\frac{1}{2}\left[\mu_{1,j}^T \mu_{1,j} + 2\mu_{1,j}^T(\mu_{2,j} - \mu_{1,j}) - \mu_{2,j}^T \mu_{2,j}\right]. \tag{96}$$

Thus, we have

$$\mathbb{D}_{KL}\big(p(y|\mu_{1,j})\,\big|\big|\,p(y|\mu_{2,j})\big) = \frac{1}{2}\left\|\mu_{2,j} - \mu_{1,j}\right\|_2^2. \tag{97}$$

By exchanging the indices in (97) we see that the backwards KL-divergence is equal to the forwards KL divergence

$$\mathbb{D}_{KL}\big(p(y|\mu_{1,j})\,\big|\big|\,p(y|\mu_{2,j})\big) = \mathbb{D}_{KL}\big(p(y|\mu_{2,j})\,\big|\big|\,p(y|\mu_{1,j})\big) = \frac{1}{2}\left\|\mu_{2,j} - \mu_{1,j}\right\|_2^2. \tag{98}$$

Thus, we have

$$\mathbb{D}_{KL}\big(p(y|\mu_{1,j}), p(y|\mu_{2,j})\big) = \mathbb{D}_{KL}\big(p(y|\mu_{1,j})\,\big|\big|\,p(y|\mu_{2,j})\big)$$
$$= \frac{1}{2}\left\|\mu_{2,j} - \mu_{1,j}\right\|_2^2 \tag{99}$$,

which is our claim $\square$.

6.2. Classification: Categorical Distribution

In this case, our predictive distribution is given by

$$p(y|h_{\theta_k}(x_j)) = \mathcal{S}\big(h_{\theta_k}(x_j)\big)[y], \tag{100}$$

where $h_{\theta_k}(x_j) \in \mathbb{R}^{n_c}$ is the output of the net, $\mathcal{S}(\cdot)$ is the Softmax function, n_c is the number of possible classes, and $\mathcal{S}[y]$ represents the y^{th} entry of (probability) vector $\mathcal{S}$. We have the forward KL-Divergence

$$\mathbb{D}_{KL}\big(p(y|h_{\theta_1}(x_j)\,\big|\big|\,p(y|h_{\theta_2}(x_j)) =$$
$$\sum_{i=1}^{n_c} p(i|h_{\theta_1}(x_j)) \log \frac{p(i|h_{\theta_1}(x_j))}{p(i|h_{\theta_2}(x_j))}, \tag{101}$$

and the forward KL-Divergence

$$\mathbb{D}_{KL}\big(p(y|h_{\theta_2}(x_j)\,\big|\big|\,p(y|h_{\theta_1}(x_j)) =$$
$$\sum_{i=1}^{n_c} p(i|h_{\theta_2}(x_j)) \log \frac{p(i|h_{\theta_2}(x_j))}{p(i|h_{\theta_1}(x_j))}. \tag{102}$$

Since we can get the vectors $h_{\theta_1}(x_j)$ and $h_{\theta_2}(x_j)$ by merely saving the old parameter θ_1 (and we have θ_2 as the current parameter) and performing 2 forward passes, we can compute both the forward KL divergence in (101) and the backward KL-divergence in (102) easily (as n_c is typically very small). We can then obtain the symmetric KL divergence as

$$\mathbb{D}_{KL}\big(p(y|h_{\theta_1}(x_j), p(y|h_{\theta_1}(x_j)) =$$
$$\frac{1}{2}\sum_{i=1}^{n_c} \mathcal{S}\big(h_{\theta_1}(x_j)\big)[i] \cdot \log \frac{\mathcal{S}\big(h_{\theta_1}(x_j)\big)[i]}{\mathcal{S}\big(h_{\theta_2}(x_j)\big)[i]} \tag{103}$$
$$+\frac{1}{2}\sum_{i=1}^{n_c} \mathcal{S}\big(h_{\theta_2}(x_j)\big)[i] \cdot \log \frac{\mathcal{S}\big(h_{\theta_2}(x_j)\big)[i]}{\mathcal{S}\big(h_{\theta_1}(x_j)\big)[i]}.$$

7. Extension of KLD-WRM to variable step-size

Note that for WOQM, SO-KLD-WRM and Q-KLD-WRM, the parameter $1/\lambda$ plays the role of a step-size. In the main body, we considered only cases when λ is fixed across different locations θ_k, that is, a fixed step-size. For KLD-WRM, a fixed λ at all steps can also be interpreted as a fixed 'strength' of KL-wake-regularization.

Our established equivalence between FISHER-WOQM and EA-NG holds for fixed step-size only, and in fact, one can check that it does not hold if we consider varying λ with θ_k. However, one might want to consider variable λ (thus variable step-size/KL-wake-regularization strength) in our KLD-WRM algorithms. That is, one might want to consider replacing λ in *equation (23) of the main body* with $\lambda^{(k)}$, whose schedule can be set as the inverse of the desired

learning-rate schedule. It turns out that if we do this, *Propositions 4.1* and *4.2* only change mildly. To incorporate variable $\lambda^{(k)}$ in Q-KLD-WRM, all one has to do is to replace (40) with $\hat{g}_{k+1} = g_{k+1} + (\lambda^{(k+1)}/\lambda^{(k)})\rho[\hat{g}_k - g_k - \hat{M}_k\hat{g}_k]$ and all λ's in *Proposition 4.2* with $\lambda^{(k)}$. This fact allows us to easily use variable $\lambda^{(k)}$ in practice. The change to *Proposition 4.1* can be directly obtained by noting that it is a particular case of *Proposition 4.2* with $B_k = 0$. *Proposition 4.2* modified to allow for variable $\lambda^{(k)}$, as well as its proof are presented in the next subsection.

7.1. Extension of Proposition 4.2 to Variable λ

Before we begin, we restate the *Q-KLD-WRM* formulation with variable λ for clarity. The *Q-KLD-WRM* step at θ_k, with variable $\lambda^{(k)}$, is the solution to

$$\min_s g_k^T s + \frac{1}{2} s^T B_k s + \lambda^{(k)} \sum_{i=0}^{k} \kappa(i)\rho^{k-i}\widetilde{\mathbb{D}}_{KL}(\theta_i||\theta_k + s),$$
(104)

where B_k is a curvature matrix which aims to approximate the Hessian H_k, and $\lambda^{(k)}$ is the level of KL-wake-regularization (or the inverse step-size) which is allowed to vary with k. Equation (104) tells us that the *Q-KLD-WRM* step at θ_k (with variable $\lambda^{(k)}$) solves

$$\min_s s^T \left[g_k + \sum_{i=0}^{k-1} \rho^{k-i}\left(\lambda^{(k)}\kappa(i)F_i \sum_{j=i}^{k-1} s_j\right)\right]$$
$$+ \frac{\lambda^{(k)}}{2} s^T \left[\sum_{i=0}^{k} \kappa(i)\rho^{k-i}F_i + \frac{1}{\lambda^{(k)}}B_k\right]s.$$
(105)

Proposition A.1 (Extension of Proposition 4.2): Analytic Solution of *Q-KLD-WRM* step with variable $\lambda^{(k)}$. *The Q-KLD-WRM step s_k at location θ_k, with variable $\lambda^{(k)}$, is given by the solution to the problem*

$$\min_s s^T \hat{g}_k + \frac{\lambda^{(k)}}{2} s^T \left[\bar{F}_k + \frac{1}{\lambda^{(k)}}B_k\right]s$$
(106)

where $\hat{g}_k$ is given by the one-step recursion

$$\hat{g}_{k+1} = g_{k+1} + \frac{\lambda^{(k+1)}}{\lambda^{(k)}}\rho\left[\hat{g}_k - g_k - \hat{M}_k\hat{g}_k\right], \ \forall k \in \mathbb{Z}^+,$$
(107)

with $\hat{g}_0 := g_0$, $\bar{F}_k := \sum_{i=0}^{k} \kappa(i)\rho^{k-i}F_i$, and

$$\hat{M}_k := \left[I + \frac{1}{\lambda^{(k)}}B_k\bar{F}_k^{-1}\right]^{-1}.$$
(108)

That is, the Q-KLD-WRM step is formally given by

$$s_k = -\frac{1}{\lambda^{(k)}}\left[\bar{F}_k + \frac{1}{\lambda^{(k)}}B_k\right]^{-1}\hat{g}_k.$$
(109)

Proof. By induction. The induction is very similar to the proof of *Proposition 4.2*.

Induction check. For $k = 0$, we have the model in (105) reducing to

$$\min_s s^T g_0 + \frac{\lambda^{(0)}}{2}s^T\left[\bar{F}_0 + \frac{1}{\lambda^{(0)}}B_0\right]s$$
(110)

Now, noting that $\hat{g}_0 = g_0$ by definition, we see that the *Q-KLD-WRM* linear[10] term obeys *Proposition A.1* at $k = 0$. Thus,

$$s_0 = -\frac{1}{\lambda^{(0)}}\left[\bar{F}_0 + \frac{1}{\lambda^{(0)}}B_0\right]^{-1}\hat{g}_0.$$
(111)

For $k = 1$, we have the model in (105) reducing to

$$\min_s s^T[g_1 + \rho\lambda^{(1)}F_0 s_0] + \frac{\lambda^{(1)}}{2}s^T\left[\bar{F}_1 + \frac{1}{\lambda^{(1)}}B_1\right]s.$$
(112)

Now, by using (111), we have that

$$\rho\lambda^{(1)}F_0 s_0 = -\rho\frac{\lambda^{(1)}}{\lambda^{(0)}}\left[F_0^{-1}\right]^{-1}\left[\bar{F}_0 + \frac{1}{\lambda^{(0)}}B_0\right]^{-1}\hat{g}_0$$
(113)

$$\rho\lambda^{(1)}F_0 s_0 = -\rho\frac{\lambda^{(1)}}{\lambda^{(0)}}\left[\bar{F}_0 F_0^{-1} + \frac{1}{\lambda^{(0)}}B_0 F_0^{-1}\right]^{-1}\hat{g}_0 =$$
$$-\rho\frac{\lambda^{(1)}}{\lambda^{(0)}}\hat{M}_0\hat{g}_0$$
(114)

where the latter equality holds by the definition of $\hat{M}_k$ evaluated at $k = 0$ and the fact that $\bar{F}_0 = F_0$. Thus,

$$g_1 + \rho\lambda^{(1)}F_0 s_0 = g_1 - \rho\frac{\lambda^{(1)}}{\lambda^{(0)}}\hat{M}_0\hat{g}_0,$$
(115)

$$g_1 + \rho\lambda^{(1)}F_0 s_0 = g_1 + \rho\frac{\lambda^{(1)}}{\lambda^{(0)}}\left[\hat{g}_0 - g_0 - \hat{M}_0\hat{g}_0\right]$$
(116)

The latter equality in (116) holds because $g_0 = \hat{g}_0$. Consequently, using the recursive definition of $\hat{g}_{k+1}$ for $k = 0$, we get

$$g_1 + \rho\lambda^{(1)}F_0 s_0 = \hat{g}_1,$$
(117)

which shows us that the *Q-KLD-WRM* linear term obeys the formulation in *Proposition A.1*. Thus the whole *Q-KLD-WRM* formulation obeys formulation (106). In particular, we have

$$s_1 = -\frac{1}{\lambda^{(1)}}\left[\bar{F}_1 + \frac{1}{\lambda^{(1)}}B_1\right]^{-1}\hat{g}_1.$$
(118)

---

[10]The Quadratic term obviously also obeys *Propositon A.1* formulation - but this fact follows trivially from the original *Q-KLD-WRM* definition, as we can see.

For $k = 2$, we have the model in (105) reducing to

$$\min_s \frac{\lambda^{(2)}}{2} s^T \left[\bar{F}_2 + \frac{1}{\lambda^{(2)}} B_2 \right] s +$$
$$s^T \left[g_2 + \rho^2 \lambda^{(2)} F_0 (s_0 + s_1) + \rho \lambda^{(2)} (1 - \rho) F_1 s_1 \right]. \quad (119)$$

The linear term in (119) can be rearranged as

$$s^T \left[g_2 + \rho^2 \lambda^{(2)} F_0 s_0 + \rho \lambda^{(2)} [\rho F_0 + (1 - \rho) F_1] s_1 \right] =$$
$$s^T \left[g_2 + \rho^2 \lambda^{(2)} \bar{F}_0 s_0 + \rho \lambda^{(2)} \bar{F}_1 s_1 \right], \quad (120)$$

where to get the second equality we used the definition of $\bar{F}_k$ applied at $k = 0$ and $k = 1$. We have already seen when doing the check for $k = 1$ (from (117)) that

$$\rho^2 \bar{F}_0 s_0 = \rho \frac{1}{\lambda^{(1)}} (\hat{g}_1 - g_1) \quad (121)$$

Thus, we have

$$\rho^2 \lambda^{(2)} \bar{F}_0 s_0 = \rho \frac{\lambda^{(2)}}{\lambda^{(1)}} (\hat{g}_1 - g_1). \quad (122)$$

So we only need to work out the $\rho \lambda^{(2)} \bar{F}_1 s_1$ term in (120). By using (118), this reads

$$\rho \lambda^{(2)} \bar{F}_1 s_1 = -\rho \frac{\lambda^{(2)}}{\lambda^{(1)}} [\bar{F}_1^{-1}]^{-1} \left[\bar{F}_1 + \frac{1}{\lambda^{(1)}} B_1 \right]^{-1} \hat{g}_1 \quad (123)$$

$$\rho \lambda^{(2)} \bar{F}_1 s_1 = -\rho \frac{\lambda^{(2)}}{\lambda^{(1)}} \left[\bar{F}_1 \bar{F}_1^{-1} + \frac{1}{\lambda^{(1)}} B_1 \bar{F}_1^{-1} \right]^{-1} \hat{g}_1 =$$
$$- \rho \frac{\lambda^{(2)}}{\lambda^{(1)}} \hat{M}_1 \hat{g}_1 \quad (124)$$

where the latter equality holds by the definition of M_k evaluated at $k = 1$. Thus, plugging in (124) and (122) in the linear term expression (120), we get that the linear term reads

$$s^T \left[g_2 + \rho^2 \lambda F_0 s_0 + \rho \lambda [\rho F_0 + (1 - \rho) F_1] s_1 \right] =$$
$$s^T \left[g_2 + \rho \frac{\lambda^{(2)}}{\lambda^{(1)}} (\hat{g}_1 - g_1 - \hat{M}_1 \hat{g}_1) \right] = s^T \hat{g}_2, \quad (125)$$

where in the last equality we used the recursive definition of $\hat{g}_{k+1}$ from *Proposition A.1*, evaluated at $k = 1$. Thus, (56) shows us that the *Q-KLD-WRM* linear term obeys the formulation in *Proposition 4.2*. Thus the whole *Q-KLD-WRM* formulation obeys formulation (106). In particular,

if we now plug the linear term in (125) back into the full model (119), we get that the *Q-KLD-WRM* step satisfies

$$s_2 = -\frac{1}{\lambda} \left[\bar{F}_2 + \frac{1}{\lambda} B_2 \right]^{-1} \hat{g}_2. \quad (126)$$

We are done with the induction check.

Induction main body. Assume for the **inductive hypothesis** that we have for some $k \geq 2$

$$s_l = -\arg \min_s s^T \hat{g}_l + \frac{\lambda^{(l)}}{2} s^T \left[\bar{F}_l + \frac{1}{\lambda^{(l)}} B_l \right] s, \ \forall l \in \{0, .., k\} \quad (127)$$

$$\hat{g}_l = g_l + \frac{\lambda^{(l)}}{\lambda^{(l-1)}} \rho \left[\hat{g}_{l-1} - g_{l-1} - \hat{M}_{l-1} \hat{g}_{l-1} \right] \ \forall l \in \{1, ..., k\}, \quad (128)$$

with $\hat{g}_0 := g_0$, $\bar{F}_k := \sum_{i=0}^{k} \kappa(i) \rho^{k-i} F_i$, and

$$\hat{M}_l := \left[I + \frac{1}{\lambda} B_l \bar{F}_l^{-1} \right]^{-1}. \quad (129)$$

We **need to prove** that at $k + 1$, for all $k \geq 2$, we have

$$\hat{g}_{k+1} = g_{k+1} + \frac{\lambda^{(k+1)}}{\lambda^{(k)}} \rho \left[\hat{g}_k - g_k - \hat{M}_k \hat{g}_k \right] \quad (130)$$

with $\hat{M}_k$ defined as in (129), and

$$s_{k+1} = \arg \min_s s^T \hat{g}_{k+1}$$
$$+ \frac{\lambda^{(k+1)}}{2} s^T \left[\bar{F}_{k+1} + \frac{1}{\lambda^{(k+1)}} B_{k+1} \right] s. \quad (131)$$

Since the inductive check holds at $k \in \{0, 1, 2\}$, this would imply *Proposition A.1* holds.

To begin, consider (105) evaluated at $k + 1$. We have that the *Q-KLD-WRM* step at $k + 1$, s_{k+1} is the solution to

$$\min_s \frac{\lambda^{(k+1)}}{2} s^T \left[\bar{F}_{k+1} + \frac{1}{\lambda^{(k+1)}} B_{k+1} \right] s$$
$$+ s^T \left[g_{k+1} + \sum_{i=0}^{k} \rho^{k+1-i} \left(\lambda^{(k+1)} \kappa(i) F_i \sum_{j=i}^{k} s_j \right) \right]. \quad (132)$$

Let us now consider the coefficient of the linear term in

(63). It reads

$$g_{k+1} + \sum_{i=0}^{k} \rho^{k+1-i}\left(\lambda^{(k+1)}\kappa(i)F_i\sum_{j=i}^{k}s_j\right) =$$

$$g_{k+1} + \lambda^{(k+1)}\sum_{i=0}^{k}\sum_{j=i}^{k}\rho^{k+1-i}\kappa(i)F_is_j =$$

$$g_{k+1} + \lambda^{(k+1)}\sum_{j=0}^{k}\left[\left(\sum_{i=0}^{j}\rho^{k+1-i}\kappa(i)F_i\right)s_j\right] = \quad (133)$$

$$g_{k+1} + \lambda^{(k+1)}\sum_{j=0}^{k}\rho^{k+1-j}\bar{F}_js_j =$$

$$g_{k+1} + \rho\sum_{j=0}^{k}\rho^{k-j}\lambda^{(k+1)}\bar{F}_js_j.$$

To get from the third to the fourth line in (133) we used the definition of $\bar{F}_k$ evaluated at $k = j$. We can further write the linear coeffcient of (132) as

$$g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda^{(k+1)}\kappa(i)F_i\sum_{j=i}^{k}s_j\right) =$$

$$g_{k+1} + \rho\sum_{j=0}^{k}\rho^{k-j}\lambda^{(k+1)}\bar{F}_js_j =$$

$$g_{k+1} + \rho\frac{\lambda^{(k+1)}}{\lambda^{(k)}}\sum_{j=0}^{k-1}\rho^{k-j}\lambda^{(k)}\bar{F}_js_j + \rho\lambda^{(k+1)}\bar{F}_ks_k.$$

$$(134)$$

Now, considering the same linear term of our *Q-KLD-WRM* formulation, *but at arbitrary* $l \in [1, k] \cap \mathbb{Z}$ (rather than at $k + 1$) we get can apply[11] equation (134) to get

$$g_l + \sum_{i=0}^{l-1}\rho^{l-i}\left(\lambda^{(l)}\kappa(i)F_i\sum_{j=i}^{k}s_j\right) = g_l + \sum_{j=0}^{l-1}\rho^{l-j}\lambda^{(l)}\bar{F}_js_j.$$

$$(135)$$

But the inductive hypothesis, (127) and (128), tells us that this linear term must also be $\hat{g}_l$. That is, we have

$$\hat{g}_l = g_l + \sum_{j=0}^{l-1}\rho^{l-j}\lambda^{(l)}\bar{F}_js_j, \quad (136)$$

and therefore

$$\sum_{j=0}^{l-1}\rho^{l-j}\lambda^{(l)}\bar{F}_js_j = \hat{g}_l - g_l. \quad (137)$$

---

[11]We can do that because our calculations (132) - (134) hold for arbitrary $k \in \mathbb{Z}^+$, and so does the *Q-KLD-WRM* formulation (134).

In particular (137) must also hold for $l = k$, because the induction hypothesis holds up until $\hat{g}_k$. We have

$$\sum_{j=0}^{k-1}\rho^{k-j}\lambda^{(k)}\bar{F}_js_j = \hat{g}_k - g_k, \quad (138)$$

Now, plugging (138) into (134), we have

$$g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda^{(k+1)}\kappa(i)F_i\sum_{j=i}^{k}s_j\right) =$$
$$(139)$$
$$g_{k+1} + \rho\frac{\lambda^{(k+1)}}{\lambda^{(k)}}(\hat{g}_k - g_k) + \rho\lambda^{(k+1)}\bar{F}_ks_k,$$

so all is left is to compute the term $\rho\lambda^{(k+1)}\bar{F}_ks_k$. To do that, we use the induction hypothesis again. In particular, we use (127) for $l = k$ and solve exactly for s_k, to get

$$s_k = -\frac{1}{\lambda^{(k)}}\left[\bar{F}_k + \frac{1}{\lambda^{(k)}}B_k\right]^{-1}\hat{g}_k. \quad (140)$$

Using (71), we have

$$\rho\lambda^{(k+1)}\bar{F}_ks_k = -\rho\frac{\lambda^{(k+1)}}{\lambda^{(k)}}\left[\bar{F}_k^{-1}\right]^{-1}\left[\bar{F}_k + \frac{1}{\lambda^{(k)}}B_k\right]^{-1}\hat{g}_k$$

$$= -\rho\frac{\lambda^{(k+1)}}{\lambda^{(k)}}\hat{M}_k\hat{g}_k.$$

$$(141)$$

The last equality of (141) is obtained by combining the two inverses in the inverse of a single matrix, and then using the definition of $\hat{M}_k$, (108).

Plugging (141) into (139) gives us that our linear term coefficient of the *Q-KLD-WRM* formulation at $k + 1$ is given by

$$g_{k+1} + \sum_{i=0}^{k}\rho^{k+1-i}\left(\lambda\kappa(i)F_i\sum_{j=i}^{k}s_j\right) =$$

$$g_{k+1} + \rho\frac{\lambda^{(k+1)}}{\lambda^{(k)}}(\hat{g}_k - g_k) - \rho\frac{\lambda^{(k+1)}}{\lambda^{(k)}}\hat{M}_k\hat{g}_k = \quad (142)$$

$$g_{k+1} + \rho\frac{\lambda^{(k+1)}}{\lambda^{(k)}}\left[\hat{g}_k - g_k - \hat{M}_k\hat{g}_k\right].$$

Equation (142), combined with (132) shows us that the induction target (130)-(131) is satisfied. Thus, the proof is complete □.

8. Further Numerical Experiments

In this section, we first give the implementation details separately for each of the four considered solvers (K-FAC, SO-KLD-WRM, Q-KLD-WRM, QE-KLD-WRM). We then give the network architecture and finally present our numerical results in more detail than in the main body (looking at more metrics and more in-depth comments). The considered problem was MNIST classification.

8.1. Optimizer Implementation Details

Kronecker Factors Regularization: Typically, a Levenberg-Marquardt style regularization is used with the Kronecker factors in K-FAC to implicitly decide the step size (see *Section 6* in [1]) . For our network architecture, the Kronecker factors were full rank as long as the batch size was reasonably large. We found the optimal batch size to be 512, and this was well above the limit at which the Kronecker factors became low rank. Thus, we avoided using the complicated Kronecker-factor regularization method described in the original K-FAC paper ([1]).

Instead, we simply used a constant learning rate, and added a constant regularization factor of 0.01 to the Kronecker factor's eigenvalues (when doing the eigen-decomposition) in order to avoid infinite stepsize. This performed well enough in practice for all solvers including K-FAC (note that all our examined KLD-WRM implementations are based on the Kronecker factors in K-FAC).

Learning Rates and λ: The learning rate for *K-FAC* was set to 0.01. In all KLD-WRM cases, we set $\lambda = 100$ (which is roughly equivalent[12] to a learning rate of 0.01).

Step/Gradient Clipping: For K-FAC and SO-KLD-WRM the clipping strategy of our *"K-FAC library codes[13]"* was used, with a clip parameter of 0.1. For Q-KLD-WRM and QE-KLD-WRM we used a slightly different clipping strategy: clip each parameter group individually if $\|\theta_{G_i}\|_2 / \sqrt{|\theta_{G_i}|} > \tau$, where θ_{G_i} is the parameters in group G_i and τ is a clipping threshold. In our experiments we set $\tau = 2$ throughout.

The reason for using different strategies with different solvers was to try altering the step by the least amount when clipping in each case, while avoiding blow-ups. From our numerical trials, it seemed that the outlined choices fit our purpose. Note that clipping matters very little for our results. One reason for that is because it only very rarely activates, and whenever it does so, it will be in the initial phase of the training. We hereby focus on performance metrics that have to do much more with mid and end-training results - in other words, we do not claim (nor investigate whether) KLD-WRM implementations perform better than K-FAC in extremely low-epoch budget regime.

Updating the Kronecker Factors and Computing their Inverse: It is typical that the eigen-decomposition (i.e. computing inverses) of the Kronecker factors is performed with lower frequency than updating the Kronecker factors

(the statistics) for maximal computation efficiency [1, 2]. However, we performed both of them with the same frequency (every 30 steps) across all our four investigated solvers. We chose to do so in order to avoid the complications which arise with KLD-WRM instantiations when having these two frequencies different (which we did not discuss here). While K-FAC would perform better if we increase the statistics update frequency and keep the eigen-decomposition frequency constant, this would also happen for KLD-WRM. Thus, we believe that our comparison is fair, even if we set the two update frequencies to be equal (which is empirically observed to be sub-optimal).

Other Details: We used *weight decay* with parameter of 0.001, and a batch size of 512 throughout.

8.1.1. QE-KLD-WRM Specific Implementation Details

For K-FAC, SO-KLD-WRM and Q-KLD-WRM there are no specific implementation details, and all these fall under the generic details we have just described. However, for QE-KLD-WRM we need to (approximately) solve the QE-KLD-WRM sub-problem in equation (36) of the main body, which defines the QE-KLD-WRM step. This brings in further hyper-parameters: initial guess, *"inner"* optimizer choice, and the number of inner optimization steps.

We use the Q-KLD-WRM step as the initial guess. This way, even if our inner optimizer makes little progress, we can still get a reasonable step. We choose the inner optimizer to be SGD with no momentum. We let ω be the ratio between the inner optimizer learning rate and outer optimizer "learning rate[14]". Further, we let N_{IS} and N_{CAP} be the number of inner optimization steps, and the maximum number of *"old"* networks saved[15] respectively. Note that when we have more than N_{CAP} networks stored (including the current one), we delete the oldest one.

We have to choose N_{CAP} large enough s.t. we have a good approximation of the $\mathbb{D}_{KL}$ term in equation (36) of the main body, but small enough s.t. the memory does not over-flow. A simple rule (which is the one we followed) is to pick N_{IS} to be the smallest positive integer s.t. $\rho^{N_{IS}} < 0.1$. The over-flow problem might limit us from making this choice[16] for larger nets if TPUs are not available. However, for our net, it was not a concern even on a standard CPU/GPU. We let ρ, ω, N_{IS} be variable, and try four different parameters configurations (see *Figures 1,2,3* and *4*). Note that ρ is not specific to QE, but is present in all solvers.

---

[12]We can see that by thinking of K-FAC as a practical implementation of EA-NG, which is of the form shown in equation (22) in the main body: in this case $1/\lambda$ is the learning rate.

[13]These are the codes we built our codes on. In particular, we used the *kfac.py* file at https://github.com/ikostrikov/pytorch-a2c-ppo-acktr-gail/blob/master/a2c_ppo_acktr/algo/kfac.py.

[14]Which is $1/\lambda = 0.01$ here for all solvers.

[15]To approximate the $\mathbb{D}_{KL}$ term in equation (36) of the main body.

[16]In this case, we have to choose between reducing rho or using a coarser approximation of $\mathbb{D}_K L$.

Since the exact KL divergence estimate tended to be 2.5 orders o magnitude larger than the loss, we used $\zeta(i) = \kappa(i)/330$ for QE-KLD-WRM, rather than $\zeta(i) = \kappa(i)$ as for the other solvers.

8.2. Implementation Detail: Network Architecture

Conv1: seven 5×5 *filters, ReLu activation,* 2×2 *MaxPool; Conv2: five* 5×5 *filters, ReLu activation,* 2×2 *MaxPool; Fully-Connected 1: 112 nodes and ReLu activation; Fully-Connected 2: 30 nodes.* The output dimension is equal to the number of classes, which is $n_c = 10$. This gives a total of 4712 parameters. We also add dropout at the very last layer to reduce over-fit. The output of the net represents the logits, which are then passed through softmax activation function to obtain class probabilities. The loss is taken to be cross-entropy (using the softmax output), as is typical in classification.

Both the network architecture and the optimizer details can be directly observed from the codes available at: GIVE CODE LINK.

8.3. Further Numerical Results and Comments

In this section, we present numerical results for all investigated parameters (as opposed to only the best ones as in the main body). Even though we do not focus on the training loss, we also show it here for completeness.

8.3.1. CONSIDERED METRICS EXPLANATION

We present results for K-FAC, SO-KLD-WRM, Q-KLD-WRM and QE-KLD-WRM in *Figures 1, 2, 3* and *4* respectively. Four different hyper-parameter settings are considered for each solver, and these vary across columns for each figure. *Test accuracy*, *test loss* and *train loss* are considered going down the rows. Ten runs are considered for each *[optimizer, parameter setting]* pair. To simplify the interpretation of these results, we consider extracting some *summary statistics*[17] out of the raw results (shown in *Figures 1 - 4*). These summary statistics are shown in *Table 1*.

Firstly, we consider the distribution of *test accuracy* and *test loss* after 50 training epochs. We show the *empirical mean* and *standard deviation* in the final columns of *Table 1*. In principle, low standard deviation is desirable, because it indicates more robustness (w.r.t. randomness) in obtaining good results. However, having *"favourable" outliers* (which are more likely to happen under large variance) is beneficial.

Secondly, we count the *number of runs that exceed a certain goodness threshold* for both *test accuracy* (higher is better) and *test loss* (lower is better). These metrics are

---

[17]Which are also performance metrics.

meant to give an idea about *"favourable" outliers*, as well as about the robustness of achieving good results with a solver. There are 2 important points about the counting procedure:

1. *Counting is done even if the metric degrades afterwards.* This is because we assume one would periodically save checkpoints - and choose the best at the end of training (this is easily implementable in the code at the expense of some communication cost and disk storage cost);

2. *Counting is not done unless the metric* either exceeds the threshold significantly once, or does so multiple times. We do this to try to reduce the noise.

8.3.2. INTERPRETATION OF RESULTS

Test Accuracy: For K-FAC, mean accuracy is around 96%, whereas for our proposed solvers it gets as high as 97.5% for all of them, with some getting close to, and even exceeding, 98%. The standard deviation is also much smaller for our proposed algorithms (desirable). This can also be observed by looking at *Figures 1-4*, and counting the number of runs that exceed a certain threshold (avoid counting a line if it appears to be crossed only due to noise). These counting results are also summarized in *Table 1*.

We see that, while K-FAC typically has one outlier with more than 98.5% test accuracy, it only has at most 4 runs where the test accuracy is greater than or equal to 98%, and 3 runs where this is strictly greater than 98%. In contrast, our proposed KLD-WRM instantiations have at least 6 runs where the test accuracy is greater than or equal to 98%, and at least 4 runs where this is strictly greater than 98% for most hyper-parameter values. In fact, by observing columns 2 and 3 in *Table 1*, we see that for some hyper-parameter values, our proposed solvers can get as many as 7-8 runs where the test accuracy is greater than or equal to 98%, and 5 runs where this is strictly greater than 98%. Furthermore, even in terms of favourable outliers, we have a KLD-WRM variant that performs on par with K-FAC: QE-KLD-WRM (see column 4 in *Table 1*).

Test Loss: Test loss results are roughly analogous to the test accuracy results. There are three main points to note here. First, there are 1-2 K-FAC *"favourable" outliers* (test loss better than 0.2) rather than just 1, as was the case for test accuracy. Second, there are no such outliers for KLD-WRM variants (not even for QE, as was the case with test accuracy results). Third, very good test loss is not always equivalent to very good test accuracy. This observation applies to all four solvers, but it seems that the equivalence is slightly more fragile for KLD-WRM variants. The meaning and reasons behind this are outside our scope and touch ideas like the fact that *the test set is just a sample from the distribution we care about*, and *comparing accuracy versus*

robustness to adversarial attacks.

Robustness and Overall Performance: Overall, our results show that if we can only afford to run the training a few times (perhaps once), it is preferable (in terms of epochs; both from a *test accuracy* and *test-loss* point of view) to use one of our KLD-WRM variants rather than K-FAC, because these more robustly[18] achieve good results. Conversely, if we can afford to run the algorithms many times (and take the best solution), the high variance of K-FAC results metrics *may eventually* play to our advantage (but this is not guaranteed). Thus, we believe that based on the presented results, it is preferable (in terms of epochs) to use a KLD-WRM variant rather than using K-FAC.

Choosing an Appropriate KLD-WRM Variant and Winners: Which variant of KLD-WRM should we choose depends on the problem. Recall that SO-KLD-WRM and Q-KLD-WRM have virtually the same computational[19] cost per step as K-FAC. Conversely, QE-KLD-WRM has the same data acquisition and linear-algebra[20] costs as K-FAC, but higher oracle[21] cost by a factor of $N_{IS} + 1$, due to the SGD steps taken to (approximately) solve the QE-KLD-WRM optimization subproblem (which defines the QE step).

Thus, for *problems where the data acquisition cost is relatively low*[22], one should choose between Q-KLD-WRM and SO-KLD-WRM, as these would (most of the time) give the best results with virtually identical wall-time as K-FAC. The best picks are SO-KLD-WRM with $\rho = 0.33$ and Q-KLD-WRM with $\rho \in \{0.33, 0.5\}$. The latter performs slightly better, but the former gives almost the same performance for a negligible implementation effort (assuming K-FAC codes are available) and might thus be preferred for prototyping.

For *problems where the data acquisition cost is sufficiently high*[23] *(dominating the linear algebra and oracle cost)*, the wall time per epoch will be virtually the same across all four solvers. In this case, one should choose QE-KLD-WRM. The best parameters we found (on MNIST classification) were $\rho = 0.5$, $N_{IS} = 10$, $\omega = 0.07$, $N_{\text{cap}} = 4$.

Best Hyper-parameter Settings for each Solver: By investigating *Figures 1, 2, 3, 4,* and *Table 1*, one would

(arguably) choose the following hyper-paratemers settings to be the best: *(1) K-FAC:* $\rho = 0.95$ (which is what also the original author recommends), *(2) SO-KLD-WRM:* $\rho = 0.33$, *(3) Q-KLD-WRM:* $\rho = 0.33$ or $\rho = 0.5$, *(4) QE-KLD-WRM:* $\rho = 0.5$, $N_{IS} = 10$, $\omega = 0.07$, $N_{\text{cap}} = 4$. These are the values we used for comparison in the main body.

References

[1] Martens, J.; Grosse, R. Optimizing neural networks with Kronecker-factored approximate curvature, *In: arXiv:1503.05671* (2015).

[2] Osawa, K.; Yuichiro Ueno, T.; Naruse, A.; Foo, C.-S.; Yokota, R. Scalable and practical natural gradient for large-scale deep learning, arXiv:2002.06015 (2020).

---

[18]Better mean and lower variance.

[19]Includes both data acquisition cost and optimizer cost (optimizer cost includes linear algebra and oracle cost - Forward pass and automatic differentiation cost).

[20]The linear algebra cost is marginally larger for QE-KLD-WRM than for K-FAC, but this has no real influence on the observed wall time (as with the SO and Q variants).

[21]Here we use oracle cost to refer to the forward pass and backward pass computation cost.

[22]Eg. classification problems with data on our disk.

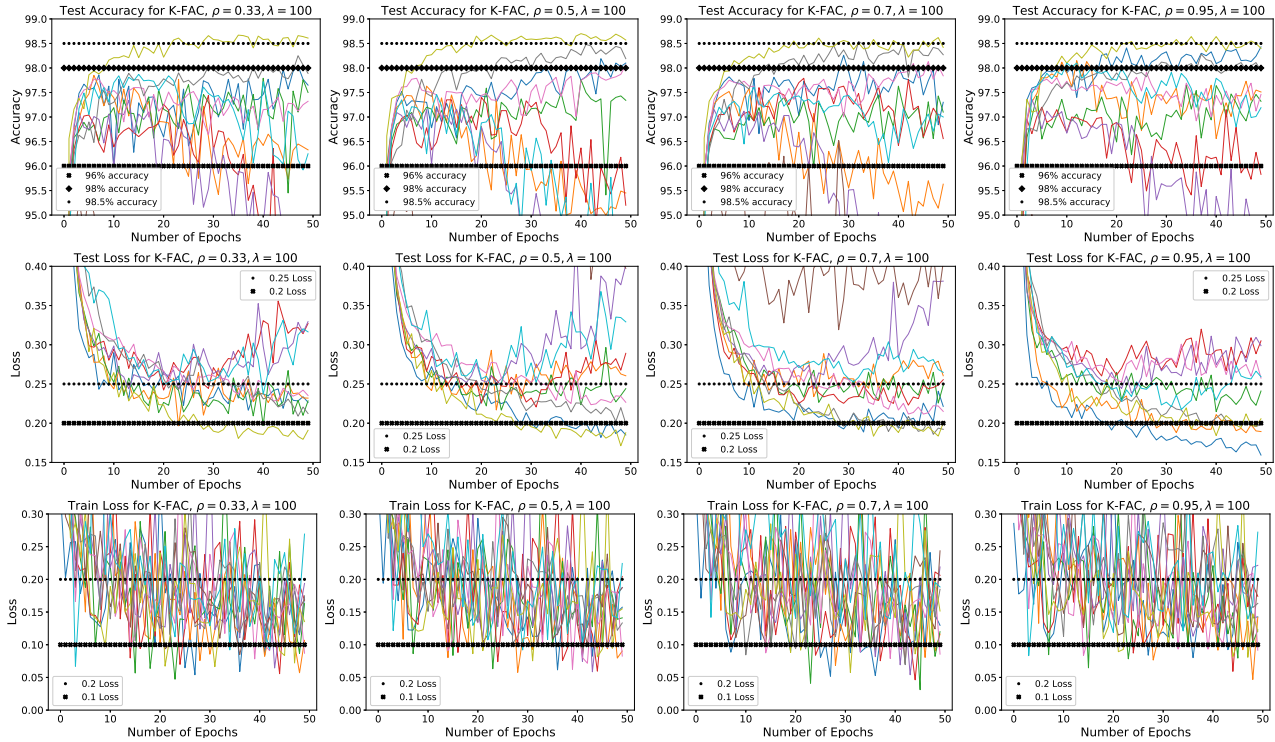[23]Eg. reinforcement learning with high data simulation cost.

Figure 1. **K-FAC Results (MNIST Classification)**: from top to bottom: Test Accuracy, Test Loss, Training Loss.From left to right: different ρ (0.33, 0.5, 0.7, 0.95). The unspecified hyper-parameters were the same across all runs, and stated in *Section 7.1*.

Figure 2. **SO-KLD-WRM Results (MNIST Classification)**: from top to bottom: Test Accuracy, Test Loss, Training Loss. From left to right: different ρ (0.33, 0.5, 0.7, 0.8). The unspecified hyper-parameters were the same across all runs, and stated in *Section 7.1*.
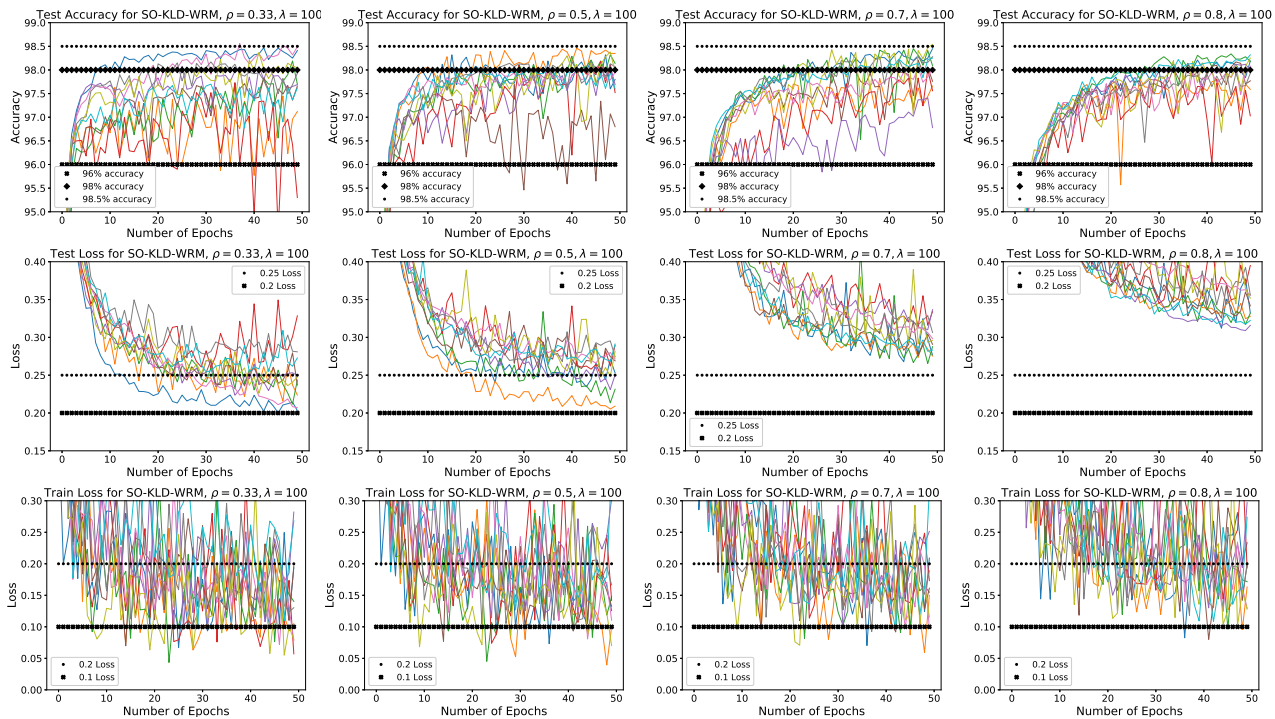
Figure 3. **Q-KLD-WRM Results (MNIST Classification)**: from top to bottom: Test Accuracy, Test Loss, Training Loss. From left to right: different ρ (0.33, 0.5, 0.7, 0.8). The unspecified hyper-parameters were the same across all runs, and stated in *Section 7.1*.
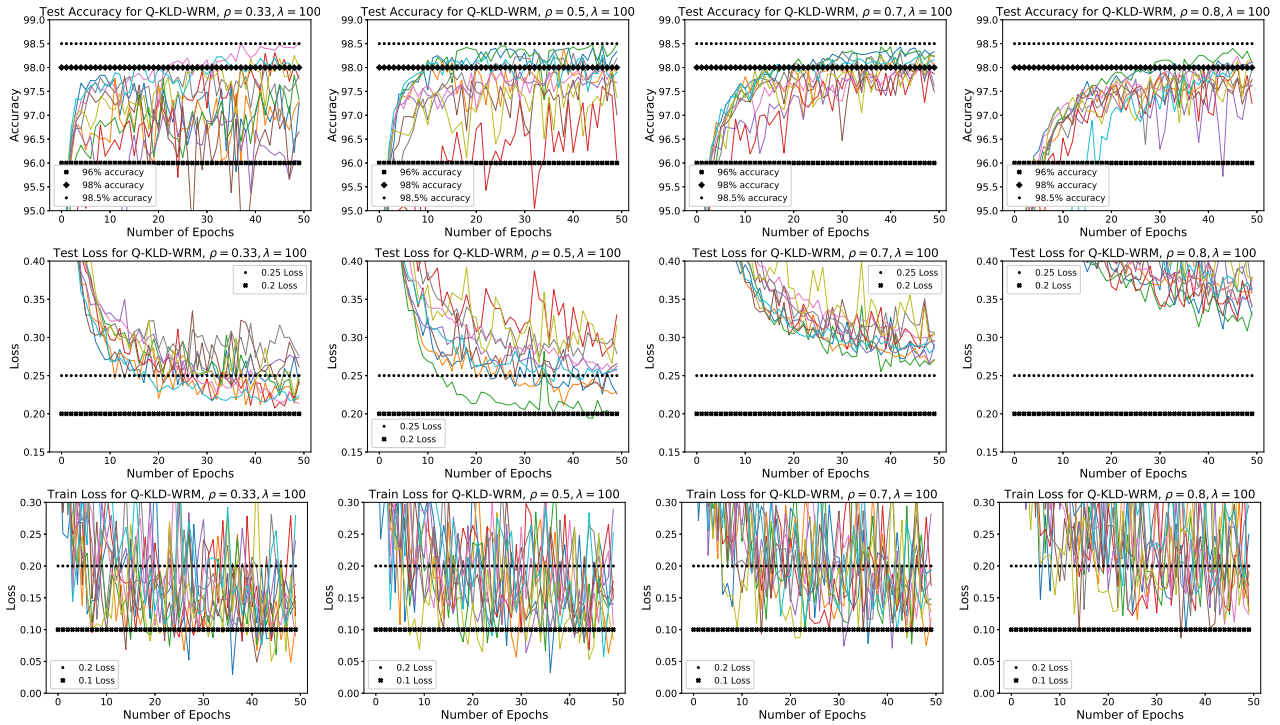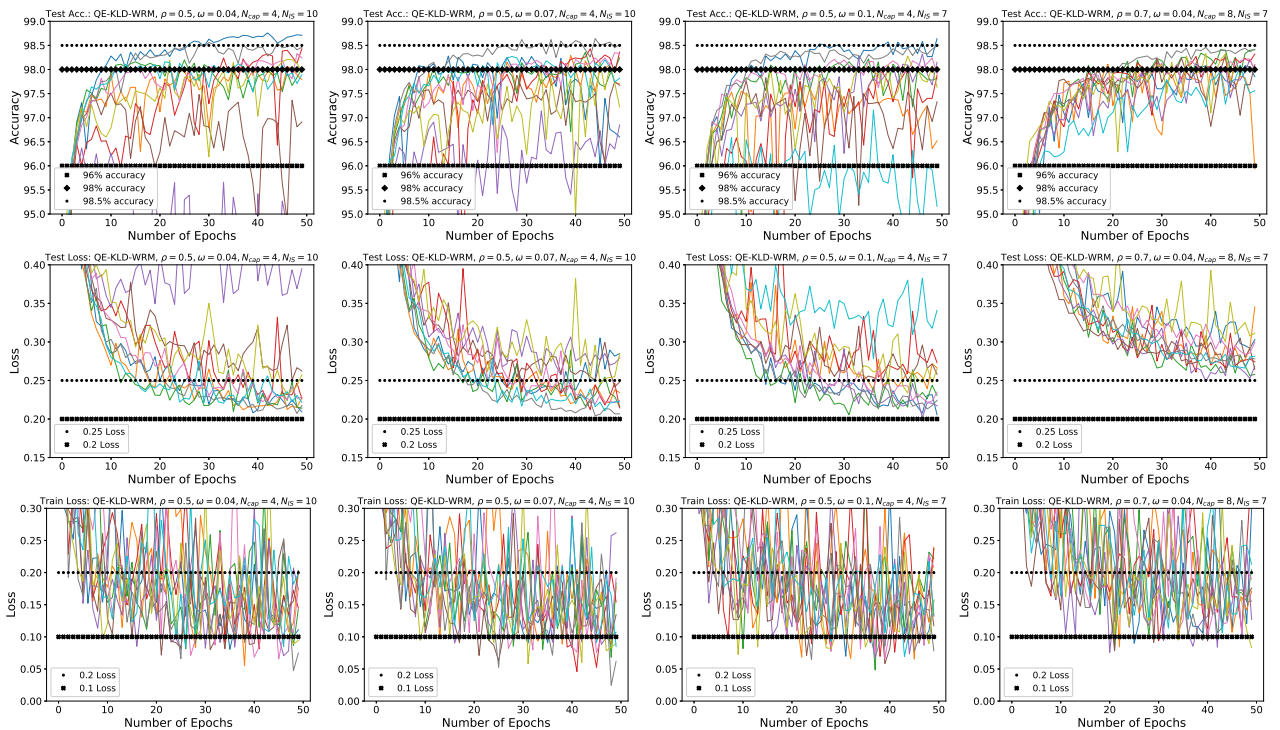

Figure 4. **QE-KLD-WRM Results (MNIST Classification)**: from top to bottom: Test Accuracy, Test Loss, Training Loss. Left to right: different QE hyper-parameter values. The unspecified hyper-parameters were the same across all runs, and stated in *Section 7.1*.

Table 1. MNIST results summary. SO-KLD-WRM, Q-KLD-WRM and QE-KLD-WRM are our proposed KLD-WRM variants and K-FAC is the benchmark. Results are displayed for multiple hyper-parameter values of the four optimizers: K-FAC, SO-KLD-WRM, Q-KLD-WRM, QE-KLD-WRM. A slash indicates counting is somewhat debateable. Bolded entries are the best ones in their column. Top to bottom: Different optimizers and hyper-parameters. Left to right: Performance metrics (out of 10 runs).

Optimizer / Metric	No. runs test acc. $\geq 98\%$	No. runs test acc. $> 98\%$	No. runs test acc. $\geq 98.5\%$	No. runs below 0.25 test loss	No. runs below 0.2 test loss	Mean test accuracy (epoch 50)	STD. test accuracy (epoch 50)	Mean test loss (epoch 50)	STD. test loss (epoch 50)
K-FAC: $\rho = 0.33$	3	1	1	6	1	95.42%	3.73%	0.28	0.09
K-FAC: $\rho = 0.5$	3	2	1	5	**2**	95.69%	2.85%	0.27	0.08
K-FAC: $\rho = 0.7$	4	2	1	5	**2**	96.13%	2.58%	0.26	0.07
K-FAC: $\rho = 0.95$	3	3	1	5	**2**	96.19%	3.2%	0.26	0.10
SO-KLD-WRM: $\rho = 0.33$	4	2	0	7	0	97.60%	0.85%	0.25	0.04
SO-KLD-WRM: $\rho = 0.5$	7 / 8	4 /5	0	3	0	97.79%	0.41%	0.26	0.02
SO-KLD-WRM: $\rho = 0.7$	6	5	0	0	0	97.94%	0.5%	0.30	0.02
SO-KLD-WRM: $\rho = 0.8$	5	3	0	0	0	97.93%	0.37%	0.34	0.02
Q-KLD-WRM: $\rho = 0.33$	3	1	0	7	0	97.47%	0.69%	**0.24**	0.02
Q-KLD-WRM: $\rho = 0.5$	5	4	0	4	0	97.69%	0.69%	0.26	0.04
Q-KLD-WRM: $\rho = 0.7$	6 / 7	4	0	0	0	**98.03%**	**0.27**%	0.29	**0.01**
Q-KLD-WRM: $\rho = 0.8$	4	1	0	0	0	97.88%	**0.27**%	0.35	0.02
QE-KLD-WRM: $\rho = 0.5$, $N_{IS} = 10$ $\omega = 0.04$, $N_{\text{cap}} = 4$	7 / 8	7	**2**	8	0	97.85%	0.67%	**0.24**	0.03
QE-KLD-WRM: $\rho = 0.5$, $N_{IS} = 10$, $\omega = 0.07$, $N_{\text{cap}} = 4$	8	7	**2**	**9**	0	**98.01%**	0.64%	**0.23**	0.02
QE-KLD-WRM: $\rho = 0.5$, $N_{IS} = 7$ $\omega = 0.1$, $N_{\text{cap}} = 4$	7	6	**2**	**9**	0	97.93%	0.46%	**0.23**	0.02
QE-KLD-WRM: $\rho = 0.7$, $N_{IS} = 10$ $\omega = 0.04$, $N_{\text{cap}} = 8$	**9 / 10**	**8**	0	0	0	**98.07%**	0.44%	0.28	0.02