
Bayesian Online Natural Gradient (BONG)

Matt Jones
U. Colorado, Boulder
mcj@colorado.edu

Peter Chang
MIT
gyuyoung@mit.edu

Kevin Murphy
Google DeepMind
kpmurphy@google.com

Abstract

We propose a novel approach to sequential Bayesian inference based on variational Bayes. The key insight is that, in the online setting, we do not need to add the KL term to regularize to the prior (which comes from the posterior at the previous timestep); instead we can optimize just the expected log-likelihood, performing a single step of natural gradient descent starting at the prior predictive. We prove this method recovers exact Bayesian inference if the model is conjugate, and empirically outperforms other online VB methods in the non-conjugate setting, such as online learning for neural networks, especially when controlling for computational costs.

1 Introduction

Bayesian methods for neural network (NN) training aim to minimize the Kullback-Leibler divergence between true and estimated posterior distributions. This is equivalent to minimizing the variational loss (or negative ELBO)

$$\mathcal{L}(\psi) = -\mathbb{E}_{\theta \sim q_\psi} [\log p(\mathcal{D}|\theta)] + D_{\text{KL}}(q_\psi \| p_0) \quad (1)$$

Here θ are the network parameters, ψ are the variational parameters of the approximate posterior $q_\psi(\theta)$, \mathcal{D} is the training dataset, and $p_0(\theta)$ is the prior. The two terms in the variational loss correspond to data fit and regularization to the prior, the latter being analogous to a regularizer $r(\theta) = -\log p_0(\theta)$ in traditional point estimation methods like SGD.

An important set of approaches learns the variational parameters by gradient descent on $\mathcal{L}(\psi)$ [Blundell et al., 2015]. More recently Khan and colleagues [Khan et al., 2018b, Khan and Rue, 2023, Shen et al., 2024] have proposed using the natural gradient $\mathbf{F}_\psi^{-1} \nabla_\psi \mathcal{L}(\psi)$ where \mathbf{F}_ψ is the Fisher information matrix of the variational family evaluated at q_ψ . Natural gradient descent (NGD) is often more efficient than vanilla GD because it accounts for the intrinsic geometry of the variational family [Amarri, 1998]. Khan and Rue [2023] call this approach the "Bayesian Learning Rule" or BLR. Using various choices for the variational distribution, generalized losses replacing negative log-likelihood, and other approximations, they reproduce many standard optimization methods such as Adam, and derive new ones.

We study Bayesian NN optimization in online learning, where the data are observed in sequence, $\mathcal{D}_t = \{(x_k, y_k)_{k=1}^t\}$, and the algorithm maintains an approximate posterior $q_{\psi_t}(\theta_t) \approx p(\theta_t | \mathcal{D}_t)$, which it updates at each step. Fast updates (in terms of both computational speed and statistical efficiency) are critical for many online learning applications [Zhang et al., 2024]. To allow for nonstationarity in the datastream, we include a time index on θ_t , to represent that the parameters may change over time, as is standard for approaches based on state-space models and the extended Kalman filter (see e.g., [Sarkka and Svensson, 2023]). The belief state is updated recursively using the prior $q_{\psi_{t|t-1}}$ derived from the previous time step so that the variational loss becomes

$$\mathcal{L}(\psi_t) = -\mathbb{E}_{\theta_t \sim q_{\psi_t}} [\log p(y_t | x_t, \theta_t)] + D_{\text{KL}}(q_{\psi_t} \| q_{\psi_{t|t-1}}) \quad (2)$$

One option for this online learning problem is to apply NGD on $\mathcal{L}(\psi_t)$ at each time step, iterating until ψ_t converges before consuming the next observation. Our first contribution is a proposal for

skipping this inner loop by (a) performing a single natural gradient step with unit learning rate and (b) omitting the D_{KL} term in Eq. (2) so that learning is based only on expected loglikelihood:

$$\psi_t = \psi_{t|t-1} + \mathbf{F}_{\psi_{t|t-1}}^{-1} \nabla_{\psi_{t|t-1}} \mathbb{E}_{\theta_t \sim q_{\psi_{t|t-1}}} [\log p(\mathbf{y}_t | \mathbf{x}_t, \theta_t)] \quad (3)$$

These two modifications work together: instead of regularizing toward the prior explicitly using $D_{\text{KL}}(q_{\psi_t} | q_{\psi_{t|t-1}})$, we do so implicitly by using $\psi_{t|t-1}$ as the starting point of our single natural gradient step. This may appear as a heuristic but we prove in Proposition 4.1 that it yields exact Bayesian inference when q_{ψ} and $p(\mathbf{y}|\mathbf{x}, \theta)$ are conjugate and q_{ψ} is an exponential family with natural parameter ψ . Thus our proposed update can be viewed as a relaxation of the Bayesian update to the non-conjugate variational case. We call Eq. (3) the Bayesian online natural gradient (BONG).

Our second contribution concerns the expectation in Eqs. (1) to (3). This is intractable for NNs, even for variational distributions that are easy to compute, since the likelihood takes the form $p(\mathbf{y}_t | \mathbf{x}_t, \theta_t) = p(\mathbf{y}_t | f(\mathbf{x}_t, \theta_t))$ with $f(\mathbf{x}_t, \theta_t)$ representing the function computed by the network. Many previous approaches have approximated the expected loglikelihood by sampling methods which add variance and computation time depending on the number of samples [Blundell et al., 2015] [Shen et al., 2024]. When the variational distribution is Gaussian (or a sub-family) we propose a deterministic, closed-form update derived in two equivalent ways. First, we use a local linear approximation of the network $f(\mathbf{x}_t, \theta_t) \approx \hat{f}_t(\theta_t)$ [Immer et al., 2021a] and a Gaussian approximation of the likelihood $\mathcal{N}(\mathbf{y}_t | f_t(\theta_t), \mathbf{R}_t)$ [Ollivier, 2018] [Tronarp et al., 2018]. Under these assumptions the expectation in Eq. (3) can be calculated analytically. Alternatively, we use a different linear approximation and a delta approximation $q_{\psi_{t|t-1}}(\theta_t) \approx \delta_{\mu_{t|t-1}}(\theta_t)$ where $\mu_{t|t-1} = \mathbb{E}_{q_{\psi_{t|t-1}}}[\theta_t]$ is the prior mean, so that the expectation in Eq. (3) is replaced by a plugin prediction $\log p(\mathbf{y}_t | f(\mathbf{x}_t, \mu_{t|t-1}))$. Assuming the likelihood is an exponential family with natural parameter $f(\mathbf{x}_t, \theta_t)$ (which is true for classification or regression), we prove in Proposition 4.2 that the linear-Gaussian and linear-plugin approximations yield the same update, which we call linearized BONG, or BONG-LIN.

Our BONG framework unifies several existing methods for Bayesian online learning, and it offers new algorithms based on alternative variational families or parameterizations. We define a large space of methods by crossing several factors in our theoretical framework and conduct experiments systematically testing how these factors affect online learning performance. We find support for all three principles of our approach—NGD, implicit regularization to the prior, and linearization—in terms of both statistical and computational efficiency.

2 Related work

Variational inference approximates the Bayesian posterior from within some suitable family in a way that bypasses the normalization term [Zellner, 1988] [Jordan et al., 1999]. A common choice for the variational family is a Gaussian. For online learning, the exact update equations for Gaussian variational filtering are given by the RVGA method of [Lambert et al., 2021]. This update is implicit but can be approximated by an explicit RVGA update which we show arises as a special case of BONG. Most applications of Gaussian VI use a mean-field approximation defined by diagonal covariance, which scales linearly with model size. More expressive but still linear in the model size are methods that express the covariance [Tomczak et al., 2020] or precision [Mishkin et al., 2018] [Lambert et al., 2023] [Chang et al., 2023] as a sum of diagonal and low rank matrices (DLR). In this paper, we consider variational families defined by full covariance, diagonal covariance, and DLR covariance.

For NNs and other complicated models, even the variational approximation can be intractable, so methods have been developed to approximately minimize the VI loss. Bayes by backprop (BBB) [Blundell et al., 2015] learns a variational distribution on NN weights by iterated GD on the VI loss of Eq. (1). They focus on mean-field Gaussian approximations but the approach also applies to other variational families. Here we adapt BBB to online learning to compare to our methods.

The Bayesian learning rule (BLR) replaces BBB’s GD with NGD [Khan and Rue, 2023]. Several variants of BLR have been developed including VON and VOGN for a mean-field Gaussian prior [Khan et al., 2018b] and SLANG for a DLR Gaussian [Mishkin et al., 2018]. BLR has also been used to derive versions of many classic optimizers including SGD, RMSprop and Adam [Khan et al., 2018a] [Khan and Rue, 2023] [Lin et al., 2024] [Shen et al., 2024]. Although BLR has been applied to online learning, we are particularly interested in Bayesian filtering including in nonstationary environments, where observations must be processed one at time and updates are based on the posterior from the

previous step, often in conjunction with SSM dynamics. We therefore develop filtering versions of BLR to compare to BONG, some of which reduce to VON, VOGN and SLANG in the batch setting, while others are novel. We also note BLR is a mature theory including several clever tricks we have not yet incorporated into our framework.

While BLR allows alternative losses in place of the NLL in Eq. (2), other work has further generalized VI by replacing the KL term with other divergences [Knoblauch et al., 2022]. Our approach fits within that framework in that it drops the divergence altogether. Our approach of implicitly regularizing to the prior using a single NGD step is also similar to the implicit MAP filter of [Bencomo et al., 2023] which performs truncated GD from the prior mode. The principal difference is they perform GD on model parameters (θ_t) while we do NGD on the variational parameters (ψ_t). Thus BONG maintains a full prior and posterior while IMAP is more concerned with how the choice of optimizer can substitute for explicit tracking of covariance.

EKF applications to NNs apply Bayesian filtering using a local linear approximation of the network, leading to simple closed form updates [Singhal and Wu, 1989; Puskorius and Feldkamp, 1991]. The classic EKF assumes a Gaussian observation distribution but it has been extended to other exponential families (e.g. for classification) by matching the mean and covariance in what we call the conditional moments EKF (CM-EKF) [Ollivier, 2018; Tronarp et al., 2018]. Applying a KL projection to diagonal covariance yields the variational diagonal EKF (VD-EKF) [Chang et al., 2022]. Alternatively, projecting to diagonal plus low rank covariance using SVD gives LO-FI [Chang et al., 2023]. We derive all these methods as special cases of BONG-LIN. Further developments in this direction include the method of [Titsias et al., 2023] which does Bayesian filtering on only the final weight layer, and WoLF [Duran-Martin et al., 2024] which achieves robustness to outliers through data-dependent weighting of the loglikelihood.

3 Background

We study online supervised learning where the agent receives input $\mathbf{x}_t \in \mathbb{R}^D$ and observation $\mathbf{y}_t \in \mathbb{R}^C$ on each time step, which it aims to model with a function $f(\mathbf{x}_t, \theta_t)$ such as a NN with weights $\theta_t \in \mathbb{R}^P$. The predictions for \mathbf{y}_t are given by some observation distribution $p(\mathbf{y}_t | f(\mathbf{x}_t, \theta_t))$. For example, f may compute the mean for regression or the class logits for classification.

We work in a Bayesian framework where the agent maintains an approximate posterior distribution over θ_t after observing data $\mathcal{D}_t = \{(\mathbf{x}_k, \mathbf{y}_k)\}_{k=1}^t$. The filtering posterior $q_{\psi_t}(\theta_t) \approx p(\theta_t | \mathcal{D}_t)$ is approximated within some parametric family indexed by the variational parameter ψ_t . We allow for nonstationarity by assuming θ changes over time according to some dynamic model $p(\theta_t | \theta_{t-1})$. By pushing the posterior from step $t-1$ through the dynamics we obtain a prior for step t given by $q_{\psi_{t|t-1}}(\theta_t) \approx p(\theta_t | \mathcal{D}_{t-1})$. For example if the variational family is Gaussian $q_{\psi_t}(\theta_t) = \mathcal{N}(\theta_t | \mu_t, \Sigma_t)$ and the dynamics are a Gaussian random walk $\theta_t \sim \mathcal{N}(\gamma \theta_{t-1}, \mathbf{Q})$ then the predict step would be $\mu_{t|t-1} = \gamma \mu_{t-1}$, $\Sigma_{t|t-1} = \gamma^2 \Sigma_{t-1} + \mathbf{Q}$. In general the predict step may require approximation to stay in the variational family (e.g., if the dynamics are nonlinear), but in this paper, our focus is the update step from $\psi_{t|t-1}$ to ψ_t upon observing $(\mathbf{x}_t, \mathbf{y}_t)$, so we assume constant (static) parameters, i.e., $p(\theta_t | \theta_{t-1}) = \delta(\theta_t - \theta_{t-1})$ for simplicity, i.e., we assume $\psi_{t|t-1} = \psi_{t-1}$, although our method is more general.

Variational inference seeks an approximate posterior that minimizes the KL divergence from the exact Bayesian update from the prior. In the online setting this becomes

$$\psi_t^* = \arg \min_{\psi} D_{\text{KL}}(q_{\psi}(\theta_t) | Z_t^{-1} q_{\psi_{t|t-1}}(\theta_t) p(\mathbf{y}_t | f(\mathbf{x}_t, \theta_t))) = \arg \min_{\psi} \mathcal{L}_t(\psi) \quad (4)$$

where \mathcal{L}_t is the online VI loss defined in Eq. (2), and the normalization term Z_t (which depends on \mathbf{x}_t) drops out as an additive constant. Our goal is an efficient approximate solution to this variational optimization problem.

We will sometimes assume q_{ψ} is an exponential family with natural parameter ψ so that $q_{\psi_t}(\theta_t) = \exp(\psi_t^\top T(\theta_t) - \Phi(\psi_t))$, with log-partition function Φ and sufficient statistics $T(\theta_t)$. Assuming Φ is strictly convex (which holds in the cases we study) there is a bijection between ψ and the dual (or expectation) parameter $\rho_t = \mathbb{E}_{\theta_t \sim q_{\psi_t}}[T(\theta_t)]$. Classical thermodynamic identities imply that the Fisher information matrix has the form $\mathbf{F}_{\psi_t} = \partial \rho_t / \partial \psi_t$. This has important implications for NGD on exponential families [Khan and Rue, 2023] because it implies that for any function ℓ defined on

in the variational parameter space the natural gradient wrt natural parameters ψ_t is the regular gradient wrt the dual parameters ρ_t , i.e., $\mathbf{F}_{\psi_t}^{-1} \nabla_{\psi_t} \ell = \nabla_{\rho_t} \ell$.

4 Methods

We propose to approximate the variational optimization problem in Eq. (4) using the BONG update in Eq. (3). When q_{ψ} is an exponential family, the fact that the natural gradient wrt the natural parameters ψ_t is the regular gradient wrt the dual parameters ρ_t implies an equivalent mirror descent form (see Appendix D for further analysis of BONG from the MD perspective):

$$\psi_t = \psi_{t|t-1} + \nabla_{\rho_{t|t-1}} \mathbb{E}_{\theta_t \sim q_{\psi_{t|t-1}}} [\log p(\mathbf{y}_t | \mathbf{x}_t, \theta_t)] \quad (5)$$

This is NGD with unit learning rate on the variational loss in Eq. (2) but ignoring the $D_{\text{KL}}(q_{\psi} \| q_{\psi_{t|t-1}})$ term. In this section we first prove this method is optimal when the model is conjugate and then describe extensions to more complex cases of practical interest.

4.1 Conjugate case

Our approach is motivated by the following result which states that BONG matches exact Bayesian inference when the variational distribution and the likelihood are conjugate exponential families:

Proposition 4.1. *Let the observation distribution be an exponential family with natural parameter θ_t (where for convenience \mathbf{y}_t is encoded as the sufficient statistics)*

$$p_t(\mathbf{y}_t | \theta_t) \propto \exp(\theta_t^\top \mathbf{y}_t - A(\theta_t) - b(\mathbf{y}_t)) \quad (6)$$

and let the prior be conjugate

$$q_{\psi_{t|t-1}}(\theta_t) = \exp\left(\psi_{t|t-1}^\top T(\theta_t) - \Phi(\psi_{t|t-1})\right) \quad (7)$$

with $T(\theta_t) = [\theta_t; -A(\theta_t)]$. Then the exact Bayesian update agrees with Eq. (5).

The proof is in Appendix C. Writing the natural parameters of the prior as $\psi_{t|t-1} = [\chi_{t|t-1}; \nu_{t|t-1}]$, we show the Bayesian update and BONG both yield $\chi_t = \chi_{t|t-1} + \mathbf{y}_t$ and $\nu_t = \nu_{t|t-1} + 1$. Intuitively, we are just accumulating a sum of the observed sufficient statistics, and a counter of the sample size (number of observations seen so far).

4.2 Variational case

In practical settings the conjugacy assumption of Proposition 4.1 will not be met, so Eqs. (3) and (5) will only approximate the Bayesian update. In this paper we restrict to Gaussian variational families. We refer to the unrestricted case as FC (full covariance), defined by the variational distribution

$$q_{\psi_{t|t-1}}(\theta_t) = \mathcal{N}(\theta_t | \mu_{t|t-1}, \Sigma_{t|t-1}) \quad (8)$$

where $\Sigma_{t|t-1}$ can be any PSD matrix. The natural and dual parameters are $\psi = (\Sigma^{-1}\mu, -\frac{1}{2}\text{vec}(\Sigma^{-1}))$ and $\rho = (\mu, \text{vec}(\mu\mu^\top + \Sigma))$. Appendix E.1.1 shows that Eq. (5) translated back to (μ, Σ) gives the following BONG update for the FC case:

$$\mu_t = \mu_{t|t-1} + \Sigma_t \mathbb{E}_{\theta_t \sim q_{\psi_{t|t-1}}} [\nabla_{\theta_t} \log p(\mathbf{y}_t | f_t(\theta_t))] \quad (9)$$

$$\Sigma_t^{-1} = \Sigma_{t|t-1}^{-1} - \mathbb{E}_{\theta_t \sim q_{\psi_{t|t-1}}} [\nabla_{\theta_t}^2 \log p(\mathbf{y}_t | f_t(\theta_t))] \quad (10)$$

which matches the explicit update in the RVGA method of [Lambert et al., 2021].

4.3 Monte Carlo approximation

The integrals over the prior $q_{\psi_{t|t-1}}$ in Eqs. (9) and (10) are generally intractable and must be approximated. One option is to use Monte Carlo, in what we call BONG-MC. Given M independent

samples $\hat{\theta}_t^{(m)} \sim q_{\psi_{t|t-1}}$, we estimate the expected gradient $\mathbf{g}_t = \mathbb{E}_{\theta_t \sim q_{\psi_{t|t-1}}} [\nabla_{\theta_t} \log p(\mathbf{y}_t | f_t(\theta_t))]$ and expected Hessian $\mathbf{G}_t = \mathbb{E}_{\theta_t \sim q_{\psi_{t|t-1}}} [\nabla_{\theta_t}^2 \log p(\mathbf{y}_t | f_t(\theta_t))]$ as the empirical means

$$\mathbf{g}_t^{\text{MC}} = \frac{1}{M} \sum_{m=1}^M \hat{\mathbf{g}}_t^{(m)}, \quad \hat{\mathbf{g}}_t^{(m)} = \nabla_{\theta_t=\hat{\theta}_t^{(m)}} \log p(\mathbf{y}_t | f_t(\theta_t)) \quad (11)$$

$$\mathbf{G}_t^{\text{MC-HESS}} = \frac{1}{M} \sum_{m=1}^M \hat{\mathbf{G}}_t^{(m)}, \quad \hat{\mathbf{G}}_t^{(m)} = \nabla_{\theta_t=\hat{\theta}_t^{(m)}}^2 \log p(\mathbf{y}_t | f_t(\theta_t)) \quad (12)$$

We use $\mathbf{G}^{\text{MC-HESS}}$ only for small models. Otherwise we use empirical Fisher (Section 4.5).

4.4 Linearized BONG

As an alternative to BONG-MC, we propose a linear approximation we call BONG-LIN that yields a deterministic and closed-form update. Assume the likelihood is an exponential family as in Proposition 4.1 but with natural parameter predicted by some function $f_t(\theta_t) = f(\mathbf{x}_t, \theta_t)$:

$$p(\mathbf{y}_t | \mathbf{x}_t, \theta_t) \propto \exp(f_t(\theta_t)^\top \mathbf{y}_t - A(f_t(\theta_t)) - b(\mathbf{y}_t)) \quad (13)$$

We also define the dual (moment) parameter of the likelihood as $h_t(\theta_t) = \mathbb{E}[\mathbf{y}_t | f_t(\theta_t)]$. In a NN, f_t and h_t are related by the final response layer. For example in classification f_t and h_t give the class logits and probabilities, with $h_t(\theta_t) = \text{softmax}(f_t(\theta_t))$, with \mathbf{y}_t being the one-hot encoding.

We now define two methods for approximating the expected gradient \mathbf{g}_t and expected Hessian \mathbf{G}_t , based on linearizing the predictive model at the prior mean $\mu_{t|t-1}$ in terms of either $f_t(\theta_t)$ or $h_t(\theta_t)$, and then prove their equivalence.

The **linear(h)-Gaussian** approximation [Ollivier, 2018, Tronarp et al., 2018] linearizes $h_t(\theta_t)$

$$\bar{h}_t(\theta_t) = \hat{\mathbf{y}}_t + \mathbf{H}_t(\theta_t - \mu_{t|t-1}) \quad (14)$$

$$\hat{\mathbf{y}}_t = h_t(\mu_{t|t-1}) \quad (15)$$

$$\mathbf{H}_t = \text{jac}(h_t(\cdot))(\mu_{t|t-1}) \quad (16)$$

and approximates the likelihood by a Gaussian with variance based at $\mu_{t|t-1}$

$$\bar{p}_t^{\text{LG}}(\mathbf{y}_t | \theta_t) = \mathcal{N}(\mathbf{y}_t | \bar{h}_t(\theta_t), \mathbf{R}_t), \quad \mathbf{R}_t = \mathbb{V}[\mathbf{y}_t | \theta_t = \mu_{t|t-1}] \quad (17)$$

The **linear(f)-delta approximation** linearizes $f_t(\theta_t)$ and maintains the original exponential family likelihood distribution in Eq. (13)

$$\bar{f}_t(\theta_t) = f_t(\mu_{t|t-1}) + \mathbf{F}_t(\theta_t - \mu_{t|t-1}) \quad (18)$$

$$\mathbf{F}_t = \text{jac}(f_t(\cdot))(\mu_{t|t-1}) \quad (19)$$

$$\bar{p}_t^{\text{LD}}(\mathbf{y}_t | \theta_t) \propto \exp(\bar{f}_t(\theta_t)^\top \mathbf{y}_t - A(f_t(\theta_t)) - b(\mathbf{y}_t)) \quad (20)$$

It also uses a plug-in approximation that replaces $q_{\psi_{t|t-1}}(\theta_t)$ with a point mass $\delta_{\mu_{t|t-1}}(\theta_t)$ so that the expected gradient and Hessian are approximated by their values at the prior mean: $\nabla_{\theta_t=\mu_{t|t-1}} \log \bar{p}_t^{\text{LD}}(\mathbf{y}_t | \theta_t)$ and $\nabla_{\theta_t=\mu_{t|t-1}}^2 \log \bar{p}_t^{\text{LD}}(\mathbf{y}_t | \theta_t)$.

Proposition 4.2. *Under a Gaussian variational distribution, the linear(h)-Gaussian and linear(f)-delta approximations yield the same values for the expected gradient and Hessian*

$$\mathbf{g}_t^{\text{LIN}} = \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (21)$$

$$\mathbf{G}_t^{\text{LIN-HESS}} = -\mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t \quad (22)$$

See Appendix C for the proof. The main idea for the $\mathbf{g}_t^{\text{LIN}}$ part is that the linear-Gaussian assumptions make the gradient linear in θ_t so the expected gradient equals the gradient at the mean. The main idea for the $\mathbf{G}_t^{\text{LIN-HESS}}$ part is that eliminating the Hessian of the NN requires different linearizing assumptions for the Gaussian and delta approximations, and the remaining nonlinear terms (from the log-likelihood in Eq. (13)) agree because of the property of exponential families that the Hessian of the log-partition A equals the conditional variance \mathbf{R}_t .

Applying Proposition 4.2 to Eqs. (9) and (10) gives the BONG-LIN update for a FC Gaussian prior

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \mathbf{K}_t(\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (23)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t|t-1} - \mathbf{K}_t \mathbf{H}_t \boldsymbol{\Sigma}_{t|t-1} \quad (24)$$

$$\mathbf{K}_t = \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^\top (\mathbf{R}_t + \mathbf{H}_t \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^\top)^{-1} \quad (25)$$

where \mathbf{K}_t is the Kalman gain matrix (see Appendix E.1.2). This matches the CM-EKF [Tronarp et al., 2018, Ollivier, 2018].

4.5 Empirical Fisher

The methods in Sections 4.3 and 4.4 require explicitly computing the Hessian of the loss (MC-HESS) or the Jacobian of the network (LIN-HESS). These are too expensive for large models or high-dimensional observations. Instead we can use an empirical Fisher approximation that replaces the Hessian with the outer product of the gradient (see e.g., [Martens, 2020]).

For the MC-EF variant, we make the following approximation:

$$\mathbf{G}_t^{\text{MC-EF}} = -\frac{1}{M} \hat{\mathbf{G}}_t^{(1:M)} \hat{\mathbf{G}}_t^{(1:M)\top} \quad (26)$$

where $\hat{\mathbf{G}}_t^{(1:M)} = [\hat{\mathbf{g}}_t^{(1)}, \dots, \hat{\mathbf{g}}_t^{(M)}]$ is the $P \times M$ matrix of gradients from the MC samples.

We can also consider a similar approach for the LIN-EF variant that is Jacobian-free. Note that if $\hat{\mathbf{y}}_t$ were the true value of $\mathbb{E}[\mathbf{y}_t | \mathbf{x}_t]$ (i.e., if the model were correct) then we would have $\mathbb{E}[(\mathbf{y}_t - \hat{\mathbf{y}}_t)(\mathbf{y}_t - \hat{\mathbf{y}}_t)^\top] = \mathbf{R}_t$, implying $\mathbb{E}[\mathbf{g}_t^{\text{LIN}} (\mathbf{g}_t^{\text{LIN}})^\top] = -\mathbf{G}_t^{\text{LIN-HESS}}$. This suggests using

$$\mathbf{g}_t^{\text{LIN}} = \nabla_{\boldsymbol{\theta}_t=\boldsymbol{\mu}_{t|t-1}} \left[-\frac{1}{2} (\mathbf{y}_t - h_t(\boldsymbol{\theta}_t))^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - h_t(\boldsymbol{\theta}_t)) \right] \quad (27)$$

$$\mathbf{G}_t^{\text{LIN-EF}} = -\mathbf{g}_t^{\text{LIN}} (\mathbf{g}_t^{\text{LIN}})^\top \quad (28)$$

where Eq. (27) exactly implements Eq. (21) as the gradient of a scalar objective (by the chain rule) and Eq. (28) is the EF approximation to Eq. (22). We expect LIN-EF to be much faster than LIN-HESS with high-dimensional observations (since it avoids computing the Jacobian), but we do not report experimental results on this combination and leave its implementation to future work.

A more accurate EF approximation is possible by sampling virtual observations $\tilde{\mathbf{y}}_t$ from $p(\cdot | f_t(\hat{\boldsymbol{\theta}}_t^{(m)}))$ or $p(\cdot | f_t(\boldsymbol{\mu}_{t|t-1}))$ and using them for the gradients in Eq. (26) or Eq. (28) (respectively) [Martens, 2020, Kunstner et al., 2020]. However, in our experiments we use the actual observations \mathbf{y}_t which is faster and follows previous work (e.g., [Khan et al., 2018b]).

4.6 Algorithm variants

In addition to the four ways of approximating the expected Hessian (summarized in Table 1), we also consider four variants of BONG, based on what kind of loss we optimize and what kind of update we perform, as we describe below. See Table 2 for a summary. Note that BONG uses one-step update with a fixed learning rate of $\alpha_t = 1$ (which is optimal for conjugate models), whereas the other variants all require tuning the learning rate.

BOG (Bayesian online gradient) performs one step of GD (instead of NGD) on the expected log-likelihood. We include a learning rate α because GD does not have the scale-invariance of NGD:

$$\boldsymbol{\psi}_{t,i} = \boldsymbol{\psi}_{t,i-1} + \alpha_t \nabla_{\boldsymbol{\psi}_{t,i-1}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] \quad (29)$$

BLR (Bayesian learning rule, [Khan and Rue, 2023]) uses NGD (like BONG) but optimizes the VI loss using multiple iterations, instead of optimizing the expected NLL with a single step. When modified to the online setting, BLR starts an inner loop at each time step with $\boldsymbol{\psi}_{t,0} = \boldsymbol{\psi}_{t|t-1}$ and iterates

$$\boldsymbol{\psi}_{t,i} = \boldsymbol{\psi}_{t,i-1} + \alpha_t \mathbf{F}_{\boldsymbol{\psi}_{t,i-1}} \nabla_{\boldsymbol{\psi}_{t,i-1}} \left(\mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] - D_{\text{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) \right) \quad (30)$$

Name	Eqs.
MC-HESS	(11), (12)
LIN-HESS	(21), (22)
MC-EF	(11), (26)
LIN-EF	(27), (28)

Table 1: The 4 Hessian approximations.

Name	Loss	Update
BONG	$\mathbb{E}[\text{NLL}]$	NGD($I = 1$)
BOG	$\mathbb{E}[\text{NLL}]$	GD($I = 1$)
BLR	VI	NGD($I \geq 1$)
BBB	VI	GD($I \geq 1$)

Table 2: The 4 update algorithms.

Method	Family and parameterization				
	FC, natural	FC, moment	Diag, natural	Diag, moment	DLR
BONG-MC-EF	$O(MP^2)^*$ [RVGA]	$O(MP^2)^*$	$O(MP)^*$	$O(MP)^*$	$O((R+M)^2P)^*$
BLR-MC-EF	$O(IP^3)$	$O(IP^3)$	$O(IMP)^*$ [VON]	$O(IMP)^*$	$O(I(R+M)^2P)^*$ [SLANG]
BOG-MC-EF	$O(P^3)$	$O(MP^2)$	$O(MP)^*$	$O(MP)^*$	$O(RMP)^*$
BBB-MC-EF	$O(IP^3)$	$O(IP^3)$	$O(IMP)^*$	$O(IMP)^*$ [BBB]	$O(IR(R+M)P)^*$
BONG-LIN-HESS	$O(CP^2)$ [CM-EKF]	$O(CP^2)$	$O(C^2P)$ [VD-EKF]	$O(C^2P)$	$O((R+C)^2P)$ [LO-FI]
BLR-LIN-HESS	$O(IP^3)$	$O(IP^3)$	$O(IC^2P)$	$O(IC^2P)$	$O(I(2R+C)^2P)$
BOG-LIN-HESS	$O(P^3)$	$O(CP^2)$	$O(C^2P)$	$O(C^2P)$	$O(C(C+R)P)$
BBB-LIN-HESS	$O(IP^3)$	$O(IP^3)$	$O(IC^2P)$	$O(IC^2P)$	$O(I(C+R)RP)$

Table 3: Time complexity of the algorithms. We assume $P \gg \{I, C, R, M\}$ so display only the terms of leading order in P . Cases where MC-EF is asymptotically faster than MC-HESS are denoted with a * superscript; otherwise they are equal. Complexities of LIN-EF methods can be found by replacing C with 1. Named cells correspond to the following existing methods (or variants thereof) in the literature: RVGA: [Lambert et al., 2021] (explicit update version); VON: [Khan et al., 2018b] (modified for online); SLANG: [Mishkin et al., 2018] (modified for online); BBB: [Blundell et al., 2015] (modified for online); CM-EKF: [Ollivier, 2018, Tronarp et al., 2018]; VD-EKF: [Chang et al., 2022]; LO-FI: [Chang et al., 2023].

For an exponential variational family this can be written in mirror descent form

$$\psi_{t,i} = \psi_{t,i-1} + \alpha_t \nabla_{\psi_{t,i-1}} \left(\mathbb{E}_{\theta_t \sim q_{\psi_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] - D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) \right) \quad (31)$$

BBB (Bayes By Backprop, [Blundell et al., 2015]) is like BLR but uses GD instead of NGD. When adapted to online learning, it starts each time step at $\psi_{t,0} = \psi_{t|t-1}$ and iterates with GD:

$$\psi_{t,i} = \psi_{t,i-1} + \alpha_t \nabla_{\psi_{t,i-1}} \left(\mathbb{E}_{\theta_t \sim q_{\psi_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] - D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) \right) \quad (32)$$

4.7 Variational families and their parameterizations

We investigate five variational families for the latent distribution: (1) FC Gaussian using natural parameters $\psi = (\Sigma^{-1}\boldsymbol{\mu}, -\frac{1}{2}\Sigma^{-1})$, (2) FC Gaussian using central moment parameters $\psi = (\boldsymbol{\mu}, \Sigma)$, (3) diagonal Gaussian using natural parameters $\psi = (\boldsymbol{\sigma}^{-2}\boldsymbol{\mu}, -\frac{1}{2}\boldsymbol{\sigma}^{-2})$, (4) diagonal Gaussian using central moment parameters $\psi = (\boldsymbol{\mu}, \boldsymbol{\sigma}^2)$, and (5) DLR Gaussian with parameters $\psi = (\boldsymbol{\mu}, \boldsymbol{\Upsilon}, \mathbf{W})$ and precision $\Sigma^{-1} = \boldsymbol{\Upsilon} + \mathbf{W}\mathbf{W}^\top$ where $\boldsymbol{\Upsilon} \in \mathbb{R}^{P \times P}$ is diagonal and $\mathbf{W} \in \mathbb{R}^{P \times R}$ with $R \ll P$. The moment parameterizations are included to test the importance of using natural parameters per Proposition 4.1. The diagonal family allows learning of large models because it scales linearly in the model size P . DLR also scales linearly but is more expressive than diagonal, maintaining some of the correlation information between parameters that is lost in the mean field (diagonal) approximation [Lambert et al., 2023, Mishkin et al., 2018, Chang et al., 2023].

Optimizing the BONG objective wrt $(\boldsymbol{\mu}, \boldsymbol{\Upsilon}, \mathbf{W})$ using NGD methods is challenging because the Fisher information matrix in this parameterization cannot be efficiently inverted. Instead we first derive the update wrt the FC natural parameters (leveraging the fact that the prior $\Sigma_{t|t-1}^{-1}$ is DLR to make this efficient), and then use SVD to project the posterior precision back to low-rank form. However, if we omit the Fisher preconditioning matrix and use GD as in BOG and BBB, we can directly optimize the objective wrt $(\boldsymbol{\mu}, \boldsymbol{\Upsilon}, \mathbf{W})$ (see Appendix E.5).

4.8 Overall space of methods

Crossing the four algorithms in Table 2, the four methods of approximating the Hessian in Table 1, and the five variational families yields 80 algorithms. Table 3 lists 40 of the methods (focusing on the MC-EF and LIN-HESS Hessian approximations that we use in our experiments), and states their time complexity, in terms of P , the dimensionality of the parameter space; M , the number of MC samples (not relevant for the linearized variants); I , the number of iterations per time step (we fix $I = 1$ for BONG and BOG variants); R , the rank of the DLR approximation (not relevant for full or diagonal covariance); and C , the size of the output vector \mathbf{y}_t (relevant only for the LIN-HESS variants). Update equations for all 80 algorithms are derived in Appendix E. Pseudocode is given in Appendix A.

5 Experiments

This section presents our primary experimental results. These are based on MNIST ($D = 784$, $N_{\text{train}} = 60k$, $N_{\text{test}} = 10k$, $C = 10$ classes) [LeCun et al., 2010]. See Appendix B for more details on these experiments, and more results on MNIST and other datasets. We focus on training on a prefix of the first $T = 2000$ examples from each dataset, since our main interest is in online learning from potentially nonstationary distributions, where rapid adaptation of a model in response to a small number of new data points is critical.

Our primary evaluation objective is the negative log predictive density (NLPD) of the test set as a function of the number of training points observed so far.¹ It is defined as $\text{NLPD}_t = -\frac{1}{N_{\text{test}}} \sum_{i \in \mathcal{D}_{\text{test}}} \log \left[\int p(\mathbf{y}_i | f(\mathbf{x}_i, \boldsymbol{\theta}_t)) q_{\psi_t}(\boldsymbol{\theta}_t) d\boldsymbol{\theta}_t \right]$. We approximate this integral in two main ways: (1) using Monte Carlo sampling² or (2) using a plugin approximation, which we denote by NLL, where we replace the posterior $q_{\psi_t}(\boldsymbol{\theta}_t)$ with a delta function centered at the mean, $\delta(\boldsymbol{\theta}_t - \boldsymbol{\mu}_t)$.

For methods that require a learning rate (i.e., all methods except BONG), we optimize it wrt mid-way or final performance on a holdout validation set, using Bayesian optimization on NLL. All methods require specifying the prior belief state, $p(\boldsymbol{\theta}_0) = \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0 = \sigma_0^2 \mathbf{I})$. We optimize over σ_0^2 and sample $\boldsymbol{\mu}_0$ from a standard NN initializer. As Hessian approximations, we use MC-EF with $M = 100$ samples and LIN-HESS. (In the appendix, we also study MC-HESS but find that it works very poorly, even with $M = 1000$.)

In Fig. 1 we compare the 4 main algorithms using DLR representation with $R = 10$. We apply these to a CNN with two convolutional layers (each with 16 features and a (5,5) kernel), followed by two linear layers (one with 64 features and the final one with 10 features), for a total of 57,722 parameters. In Fig. 2 we use the same model to compare BONG using different variational families. Shaded regions in these and all other plots indicate ± 1 SE, based on independent trials randomly varying in the prior mean $\boldsymbol{\mu}_0$, data ordering, and MC sampling.

We draw the following conclusions, which are further supported by additional results in Appendix B.

- Linearization helps: BONG-LIN-HESS outperforms all other methods, and BLR-MC-EF (based on the standard BLR) is outperformed by the new linearized BLR-LIN-HESS.
- NGD helps: BOG does well but not as well as BONG.
- Implicit regularization helps: BONG outperforms BLR.
- DLR outperforms diagonal, and greater rank is better. (We find in results not reported here that rank 5–10 often gives results as good as FC.)
- The LIN-HESS approximation outperforms MC-EF.
- DIAG_MOM and DIAG perform comparably to each other.

Finally, in Fig. 3 we report the runtimes from the experiments in Figs. 1 and 2. We see significant speed advantages for the one-step methods (BONG and BOG vs. BLR and BBB) and for the linearized methods (LIN-HESS vs. MC-EF).

6 Conclusions, limitations and future work

Our experiment results show benefits of BONG’s three main principles: NGD, implicit regularization to the prior, and linearization. The clear winner across datasets and variational families is BONG-LIN-HESS, which embodies all three principles. BLR-LIN-HESS nearly matches its performance but is much slower. Several of the best-performing algorithms are previously known (notably CM-EKF and LO-FI) but we explain these results within a systematic theory that also offers new methods (including BLR-LIN-HESS).

BONG is motivated by Proposition 4, which applies only in the idealized setting of a conjugate prior. Nevertheless we find it performs well in non-conjugate settings. On the other hand our experiments

¹We assume the training and test sets are drawn from the same static distribution. Alternatively, if there is only one stream of data coming from a potential notstationary source, we can use the prequential or one-step-ahead log predictive density [Gama et al., 2013]. We leave studying the non-stationary case to future work.

²That is, we compute $S = 100$ posterior samples $\boldsymbol{\theta}_t^s \sim p(\boldsymbol{\theta}_t | \mathcal{D}_{1:t}^{\text{train}})$ and then use $p(\mathbf{y} | \mathbf{x}) \approx \frac{1}{S} \sum_{s=1}^S p(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta}_t^s)$. For efficiently sampling from a DLR Gaussian we follow [Mishkin et al., 2018].

are based on relatively small models and datasets. It will be important to test how our methods scale up, especially using the promising DLR representation.

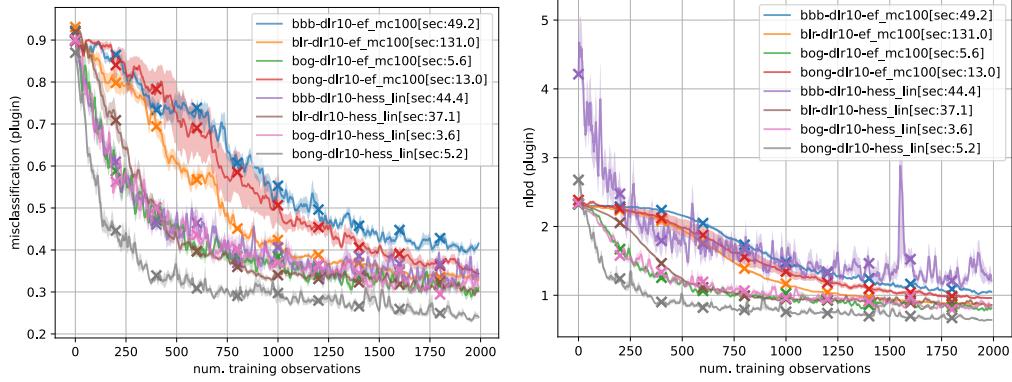


Figure 1: Performance on MNIST of plug-in predictive computed using BONG, BOG, BBB and BLR (each with LIN-HESS and MC-EF) with DLR-10 variational family.

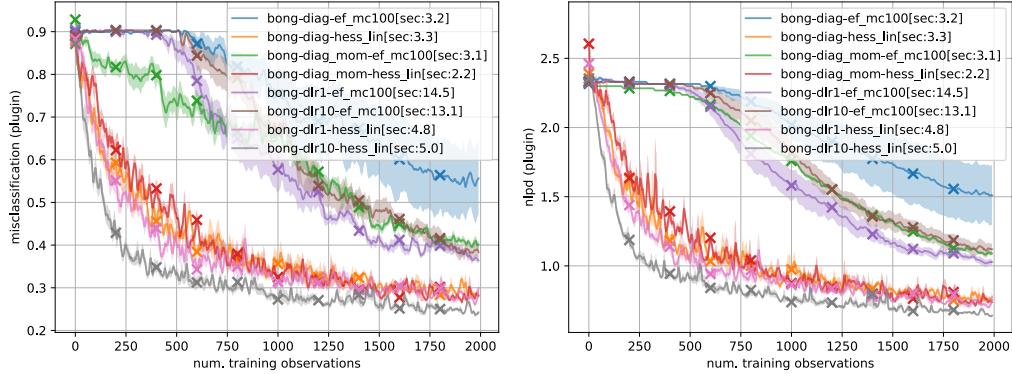


Figure 2: Performance on MNIST of plug-in predictive computed using BONG (with LIN-HESS and MC-EF) with different variational families, namely diagonal, diagonal-moment, DLR-1, DLR-10.

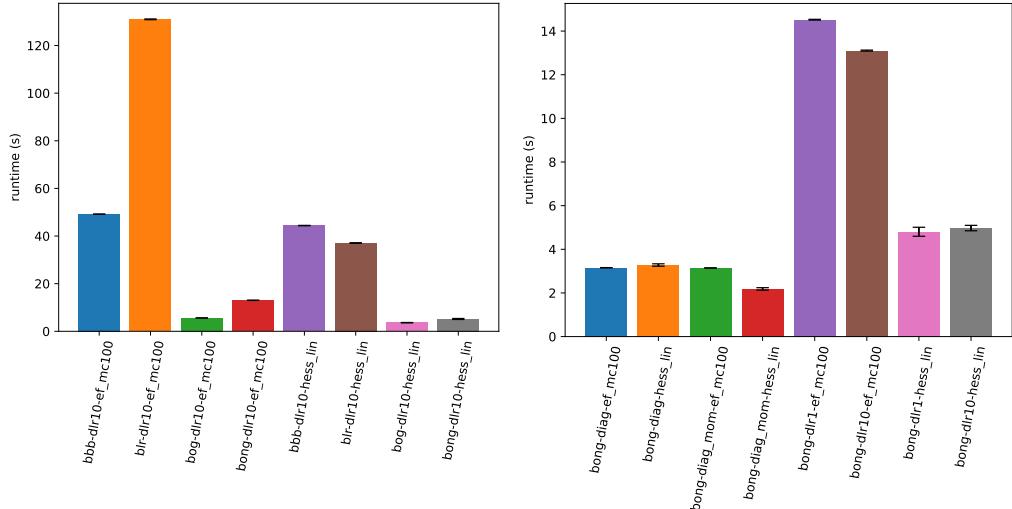


Figure 3: Runtimes for methods. Left: Corresponding to Fig. 1 using different algorithms on DLR-10 family. Right: Corresponding to Fig. 2 using BONG on different families.

Acknowledgments and Disclosure of Funding

Thanks to Gerardo Duràñ-Martín and Alex Shestopaloff for helpful input. MJ was supported by NSF grant 2020-906.

References

S Amari. Natural gradient works efficiently in learning. *Neural Comput.*, 10(2):251–276, 1998. URL <http://dx.doi.org/10.1162/089976698300017746>.

Gianluca M Bencomo, Jake C Snell, and Thomas L Griffiths. Implicit maximum a posteriori filtering via adaptive optimization. *arXiv preprint arXiv:2311.10580*, 2023.

Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. In *ICML*, 2015. URL <http://arxiv.org/abs/1505.05424>.

Georges Bonnet. Transformations des signaux aléatoires à travers les systèmes non linéaires sans mémoire. In *Annales des Télécommunications*, volume 19, pages 203–220. Springer, 1964.

Peter G Chang, Kevin Patrick Murphy, and Matt Jones. On diagonal approximations to the extended kalman filter for online training of bayesian neural networks. In *Continual Lifelong Learning Workshop at ACML 2022*, December 2022. URL <https://openreview.net/forum?id=asgeEt25kk>.

Peter G Chang, Gerardo Durán-Martín, Alexander Y Shestopaloff, Matt Jones, and Kevin Murphy. Low-rank extended Kalman filtering for online learning of neural networks from streaming data. In *COLLAS*, May 2023. URL <http://arxiv.org/abs/2305.19535>.

Gerardo Duran-Martin, Matias Altamirano, Alexander Y Shestopaloff, Leandro Sánchez-Betancourt, Jeremias Knoblauch, Matt Jones, François-Xavier Briol, and Kevin Murphy. Outlier-robust kalman filtering through generalised bayes. *arXiv preprint arXiv:2405.05646*, 2024.

João Gama, Raquel Sebastião, and Pedro Pereira Rodrigues. On evaluating stream learning algorithms. *MLJ*, 90(3):317–346, March 2013. URL <https://tinyurl.com/mrxfk4ww>.

Nathan Halko, Per-Gunnar Martinsson, and Joel A Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2):217–288, 2011.

Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.

Alexander Immer, Maciej Korzepa, and Matthias Bauer. Improving predictions of bayesian neural nets via local linearization. In Arindam Banerjee and Kenji Fukumizu, editors, *AISTATS*, volume 130 of *Proceedings of Machine Learning Research*, pages 703–711. PMLR, 2021a. URL <https://proceedings.mlr.press/v130/immer21a.html>.

Alexander Immer, Maciej Korzepa, and Matthias Bauer. Improving predictions of bayesian neural nets via local linearization. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, pages 703–711. PMLR, 2021b. URL <https://proceedings.mlr.press/v130/immer21a.html>.

Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37:183–233, 1999.

Mohammad Khan, Didrik Nielsen, Voot Tangkaratt, Wu Lin, Yarin Gal, and Akash Srivastava. Fast and scalable bayesian deep learning by weight-perturbation in adam. In *International conference on machine learning*, pages 2611–2620. PMLR, 2018a.

Mohammad Emtyiaz Khan and Håvard Rue. The bayesian learning rule. *J. Mach. Learn. Res.*, 2023. URL <http://arxiv.org/abs/2107.04562>.

Mohammad Emtiyaz Khan, Didrik Nielsen, Voot Tangkaratt, Wu Lin, Yarin Gal, and Akash Srivastava. Fast and scalable bayesian deep learning by Weight-Perturbation in adam. In *ICML*, 2018b. URL <http://arxiv.org/abs/1806.04854>.

Jeremias Knoblauch, Jack Jewson, and Theodoros Damoulas. An optimization-centric view on bayes' rule: Reviewing and generalizing variational inference. *Journal of Machine Learning Research*, 23(132):1–109, 2022.

Frederik Kunstner, Lukas Balles, and Philipp Hennig. Limitations of the empirical fisher approximation for natural gradient descent, 2020.

Marc Lambert, Silvère Bonnabel, and Francis Bach. The recursive variational gaussian approximation (R-VGA). *Stat. Comput.*, 32(1):10, December 2021. URL <https://hal.inria.fr/hal-03086627/document>.

Marc Lambert, Silvère Bonnabel, and Francis Bach. The limited-memory recursive variational gaussian approximation (l-rvga). *Statistics and Computing*, 33(3):70, 2023.

Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.

Wu Lin, Felix Dangel, Runa Eschenhagen, Juhan Bae, Richard E Turner, and Alireza Makhzani. Can we remove the square-root in adaptive gradient methods? a second-order perspective. *arXiv preprint arXiv:2402.03496*, 2024.

James Martens. New insights and perspectives on the natural gradient method. *Journal of Machine Learning Research*, 21(146):1–76, 2020.

Aaron Mishkin, Frederik Kunstner, Didrik Nielsen, Mark Schmidt, and Mohammad Emtiyaz Khan. SLANG: Fast structured covariance approximations for bayesian deep learning with natural gradient. In *NIPS*, pages 6245–6255. Curran Associates, Inc., 2018.

Yann Ollivier. Online natural gradient as a kalman filter. *Electron. J. Stat.*, 12(2):2930–2961, 2018. URL <https://projecteuclid.org/euclid.ejs/1537257630>.

Robert Price. A useful theorem for nonlinear devices having gaussian inputs. *IRE Transactions on Information Theory*, 4(2):69–72, 1958.

G V Puskorius and L A Feldkamp. Decoupled extended kalman filter training of feedforward layered networks. In *International Joint Conference on Neural Networks*, volume i, pages 771–777 vol.1, 1991. URL <http://dx.doi.org/10.1109/IJCNN.1991.155276>.

Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

Simo Sarkka and Lennart Svensson. *Bayesian Filtering and Smoothing (2nd edition)*. Cambridge University Press, 2023.

Yuesong Shen, Nico Daheim, Bai Cong, Peter Nickl, Gian Maria Marconi, Clement Bazan, Rio Yokota, Iryna Gurevych, Daniel Cremers, Mohammad Emtiyaz Khan, and Thomas Möllenhoff. Variational learning is effective for large deep networks. *arXiv preprint arXiv:2402.17641*, 2024.

Sharad Singhal and Lance Wu. Training multilayer perceptrons with the extended kalman algorithm. In *NIPS*, volume 1, 1989.

Michalis K Titsias, Alexandre Galashov, Amal Rannen-Triki, Razvan Pascanu, Yee Whye Teh, and Jorg Bornschein. Kalman filter for online classification of non-stationary data. *arXiv preprint arXiv:2306.08448*, 2023.

Marcin B Tomczak, Siddharth Swaroop, and Richard E Turner. Efficient low rank gaussian variational inference for neural networks. In *NIPS*, 2020. URL <https://proceedings.neurips.cc/paper/2020/file/310cc7ca5a76a446f85c1a0d641ba96d-Paper.pdf>.

Filip Tronarp, Ángel F García-Fernández, and Simo Särkkä. Iterative filtering and smoothing in nonlinear and Non-Gaussian systems using conditional moments. *IEEE Signal Process. Lett.*, 25(3):408–412, 2018. URL https://acris.aalto.fi/ws/portalfiles/portal/17669270/cm_parapub.pdf.

Zhewei Yao, Amir Gholami, Sheng Shen, Mustafa Mustafa, Kurt Keutzer, and Michael W Mahoney. ADAHESSIAN: An adaptive second order optimizer for machine learning. In *AAAI*, 2021. URL <http://arxiv.org/abs/2006.00719>.

Arnold Zellner. Optimal information processing and bayes’s theorem. *The American Statistician*, 42(4):278–280, 1988.

Wenxuan Zhang, Youssef Mohamed, Bernard Ghanem, Philip Torr, Adel Bibi, and Mohamed Elhoseiny. Continual learning on a diet: Learning from sparsely labeled streams under constrained computation. In *ICLR*, 2024. URL <https://openreview.net/pdf?id=Xvfz8NHmCj>.

A Abstract pseudocode

Algorithms 1 to 7 give pseudocode for applying the methods we study. The update step in Algorithm 3 calls one of four grad functions (Algorithms 4 to 7) that estimate the expected gradient and Hessian using either MC or linearization combined with either the direct Hessian (or Jacobian and observation covariance) or EF. The update step also calls an inner step function that implements BONG, BLR, BOG or BBB on some variational family corresponding to the encoding of ψ (not shown). In practice the grad-fn and inner-step-fn are not as cleanly separated because the full matrix $\bar{\mathbf{G}}_{t,i}$ is not passed between them except when the variational family is FC. When the family is diagonal, grad-fn only needs to return $\text{diag}(\bar{\mathbf{G}}_{t,i})$. When grad-fn uses EF, it only needs to return $\hat{\mathbf{G}}_{t,i}^{(1:M)}$ (grad-MC-EF) or $\mathbf{g}_{t,i}^{\text{LIN}}$ (grad-LIN-EF) and inner-step-fn will implicitly use the outer product of this output as $\bar{\mathbf{G}}_{t,1}$. Finally, note the expressions for $\mathbf{g}_{t,i}^{\text{LIN}}$ in Algorithms 6 and 7 are equivalent ways of computing the same quantity, as explained after Eq. (27).

```

for  $t = 1 : \infty$  do
   $\psi_{t|t-1} = \text{predict}(\psi_{t-1})$ 
   $\psi_t = \text{update}(\psi_{t|t-1}, \mathbf{x}_t, \mathbf{y}_t)$ 
end

```

Algorithm 1: Main loop.

```

def  $\text{predict}(\psi_{t-1})$ :
  Return  $\psi_{t|t-1} = \psi_{t-1}$ 

```

Algorithm 2: Predict step. We assume for simplicity that the parameters are constant over time.

```

def  $\text{update}(\psi_{t|t-1}, \mathbf{x}_t, \mathbf{y}_t, f(), A(), \text{grad-fn}, \text{inner-step-fn}, \alpha, I, M)$ :
   $\psi_{t,0} = \psi_{t|t-1}$ 
   $f_t(\theta_t) = f(\mathbf{x}_t, \theta_t)$ 
   $h_t(\theta_t) = E[\mathbf{y}_t | \mathbf{x}_t, \theta_t] = \nabla_{\eta=f_t(\theta_t)} A(\eta)$ 
   $\mathbf{V}_t(\theta_t) = \text{Cov}[\mathbf{y}_t | \mathbf{x}_t, \theta_t] = \nabla_{\eta=f_t(\theta_t)}^2 A(\eta)$ 
   $\ell_t(\theta_t) = \log p(\mathbf{y}_t | \mathbf{x}_t, \theta_t) = \log p(\mathbf{y}_t | f_t(\theta_t))$ 
  for  $i = 1 : I$  do
     $(\bar{\mathbf{g}}_{t,i}, \bar{\mathbf{G}}_{t,i}) = \text{grad-fn}(\psi_{t,i-1}, \ell_t, h_t, \mathbf{V}_t, M)$ 
     $\psi_{t,i} = \text{inner-step-fn}(\psi_{t|t-1}, \psi_{t,i-1}, \bar{\mathbf{g}}_{t,i}, \bar{\mathbf{G}}_{t,i}, \alpha)$ 
  end
  Return  $\psi_{t,I}$ 

```

Algorithm 3: Update step. The inner-step-fn is BONG, BLR, BOG or BBB (not shown).

def grad-MC-HESS($\psi_{t,i-1}, \ell_t, h_t = [], \mathbf{V}_t = [], M$):

for $m = 1 : M$ do

$$\left| \begin{array}{l} \hat{\theta}_{t,i}^{(m)} \sim q_{\psi_{t,i-1}}(\theta) \\ \hat{g}_{t,i}^{(m)} = \nabla_{\theta_t=\hat{\theta}_{t,i}^{(m)}} \ell_t(\theta_t) \\ \hat{\mathbf{G}}_{t,i}^{(m)} = \nabla_{\theta_t=\hat{\theta}_{t,i}^{(m)}}^2 \ell_t(\theta_t) \end{array} \right.$$

end

$$\begin{aligned} \mathbf{g}_{t,i}^{\text{MC}} &= \frac{1}{M} \sum_{m=1}^M \hat{g}_{t,i}^{(m)} \\ \mathbf{G}_{t,i}^{\text{MC-HESS}} &= \frac{1}{M} \sum_{m=1}^M \hat{\mathbf{G}}_{t,i}^{(m)} \\ \text{Return } &(\mathbf{g}_{t,i}^{\text{MC}}, \mathbf{G}_{t,i}^{\text{MC-HESS}}) \end{aligned}$$

Algorithm 4: MC gradient/Hessian estimator

def grad-MC-EF($\psi_{t,i-1}, \ell_t, h_t = [], \mathbf{V}_t = [], M$):

for $m = 1 : M$ do

$$\left| \begin{array}{l} \hat{\theta}_{t,i}^{(m)} \sim q_{\psi_{t,i-1}}(\theta) \\ \hat{g}_{t,i}^{(m)} = \nabla_{\theta_t=\hat{\theta}_{t,i}^{(m)}} \ell_t(\theta_t) \end{array} \right.$$

end

$$\begin{aligned} \mathbf{g}_{t,i}^{\text{MC}} &= \frac{1}{M} \sum_{m=1}^M \hat{g}_{t,i}^{(m)} \\ \hat{\mathbf{G}}_{t,i}^{(1:M)} &= [\hat{g}_{t,i}^{(1)}, \dots, \hat{g}_{t,i}^{(M)}] \\ \mathbf{G}_{t,i}^{\text{MC-EF}} &= -\frac{1}{M} \hat{\mathbf{G}}_{t,i}^{(1:M)} \hat{\mathbf{G}}_{t,i}^{(1:M)\top} \\ \text{Return } &(\mathbf{g}_{t,i}^{\text{MC}}, \mathbf{G}_{t,i}^{\text{MC-EF}}) \end{aligned}$$

Algorithm 5: MC gradient/Hessian estimator with Empirical Fisher

def grad-LIN-HESS($\psi_{t,i-1}, \ell_t = [], h_t, \mathbf{V}_t, M = []$):

$$\begin{aligned} \boldsymbol{\mu}_{t,i-1} &= E[\theta_t | \psi_{t,i-1}] \\ \hat{y}_{t,i} &= h_t(\boldsymbol{\mu}_{t,i-1}) \\ \mathbf{H}_{t,i} &= \text{jac}(h_t(\cdot))(\boldsymbol{\mu}_{t,i-1}) \\ \mathbf{R}_{t,i} &= \mathbf{V}_t(\boldsymbol{\mu}_{t,i-1}) \\ \mathbf{g}_{t,i}^{\text{LIN}} &= \mathbf{H}_{t,i}^{\top} \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{y}_{t,i}) \\ \mathbf{G}_{t,i}^{\text{LIN-HESS}} &= -\mathbf{H}_{t,i}^{\top} \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i} \\ \text{Return } &(\mathbf{g}_{t,i}^{\text{LIN}}, \mathbf{G}_{t,i}^{\text{LIN-HESS}}) \end{aligned}$$

Algorithm 6: Linearized gradient/Hessian estimator

def grad-LIN-EF($\psi_{t,i-1}, \ell_t = [], h_t, \mathbf{V}_t, M = []$):

$$\begin{aligned} \boldsymbol{\mu}_{t,i-1} &= E[\theta_t | \psi_{t,i-1}] \\ \mathbf{R}_{t,i} &= \mathbf{V}_t(\boldsymbol{\mu}_{t,i-1}) \\ \mathbf{g}_{t,i}^{\text{LIN}} &= \nabla_{\theta_t=\boldsymbol{\mu}_{t,i-1}} \left[-\frac{1}{2} (\mathbf{y}_t - h_t(\theta_t))^{\top} \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - h_t(\theta_t)) \right] \\ \mathbf{G}_{t,i}^{\text{LIN-EF}} &= -\mathbf{g}_{t,i}^{\text{LIN}} (\mathbf{g}_{t,i}^{\text{LIN}})^{\top} \\ \text{Return } &(\mathbf{g}_{t,i}^{\text{LIN}}, \mathbf{G}_{t,i}^{\text{LIN-EF}}) \end{aligned}$$

Algorithm 7: Linearized gradient/Hessian estimator with empirical Fisher

B Additional experiment results

In this section, we give a more thorough set of experimental results.

B.1 Further running time measures

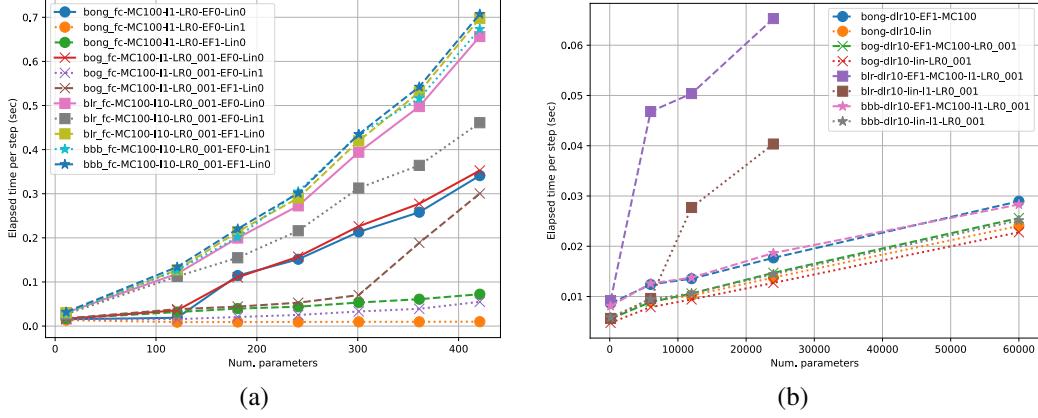


Figure 4: Running time (seconds) vs number of parameters P (size of state space). (a) Full Covariance representation. (b) DLR representation. The BLR plot is truncated due to out of memory problem.

The running times of the methods for the FC and DLR case are shown in Fig. 4. These were computed on a synthetic regression dataset. The slower speed of BLR (even with $I = 1$) relative to BONG is at least partly attributable to the fact that BLR must compute the SVD of a larger matrix (see Appendices E.5.3 and E.5.4).

B.2 Detailed results for CNN on MNIST

Here we report further metrics from the experiments in Figs. 1 and 2. In addition to plugin NLPD and MC NLPD, we evaluate using the linearized approximation proposed in Immer et al. (2021b), in which we replace $h(\mathbf{x}_i, \boldsymbol{\theta}_t)$ with $\hat{h}(\mathbf{x}_i, \boldsymbol{\theta}_t)$ defined in Eq. (14). The motivation for this approximation is the following. If we push posterior samples through a nonlinear predictive model, the results can be poor if the samples are far from the mean. However, if we linearize the predictive model, extrapolations away from the mean are more sensible. (This is true even if the posterior was not explicitly computed using a linear approximation.) For each of these three likelihood metrics, we also measure the corresponding misclassification rate based on picking the most probable predicted class from the corresponding predictive distribution. Results are shown in Figs. 5 and 6.

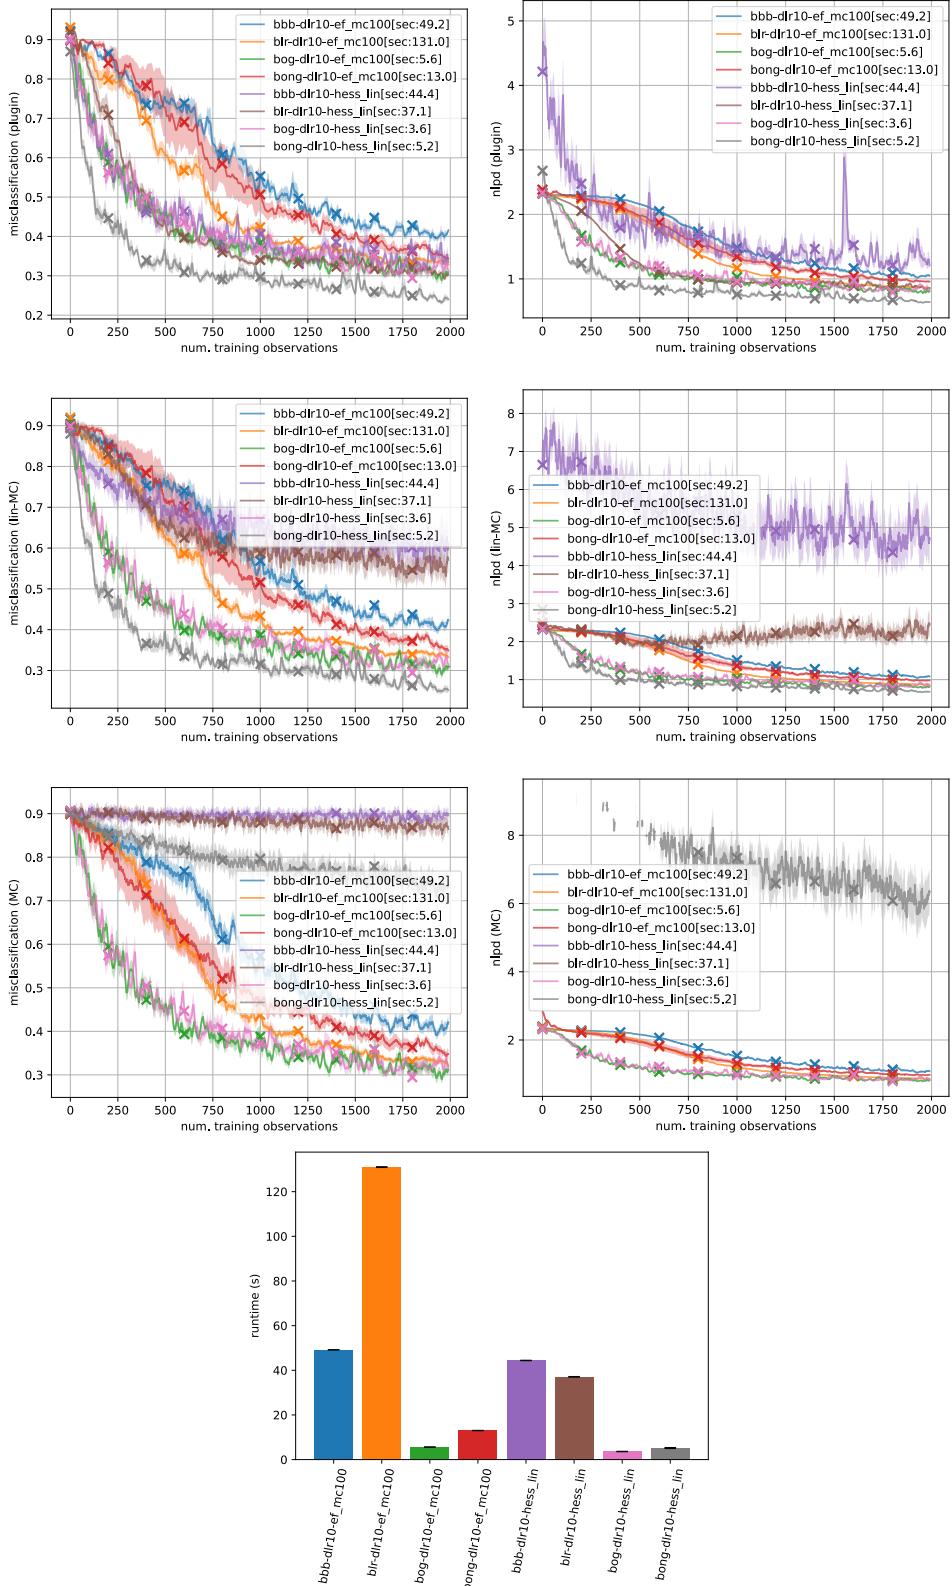


Figure 5: MNIST results for methods using DLR family.

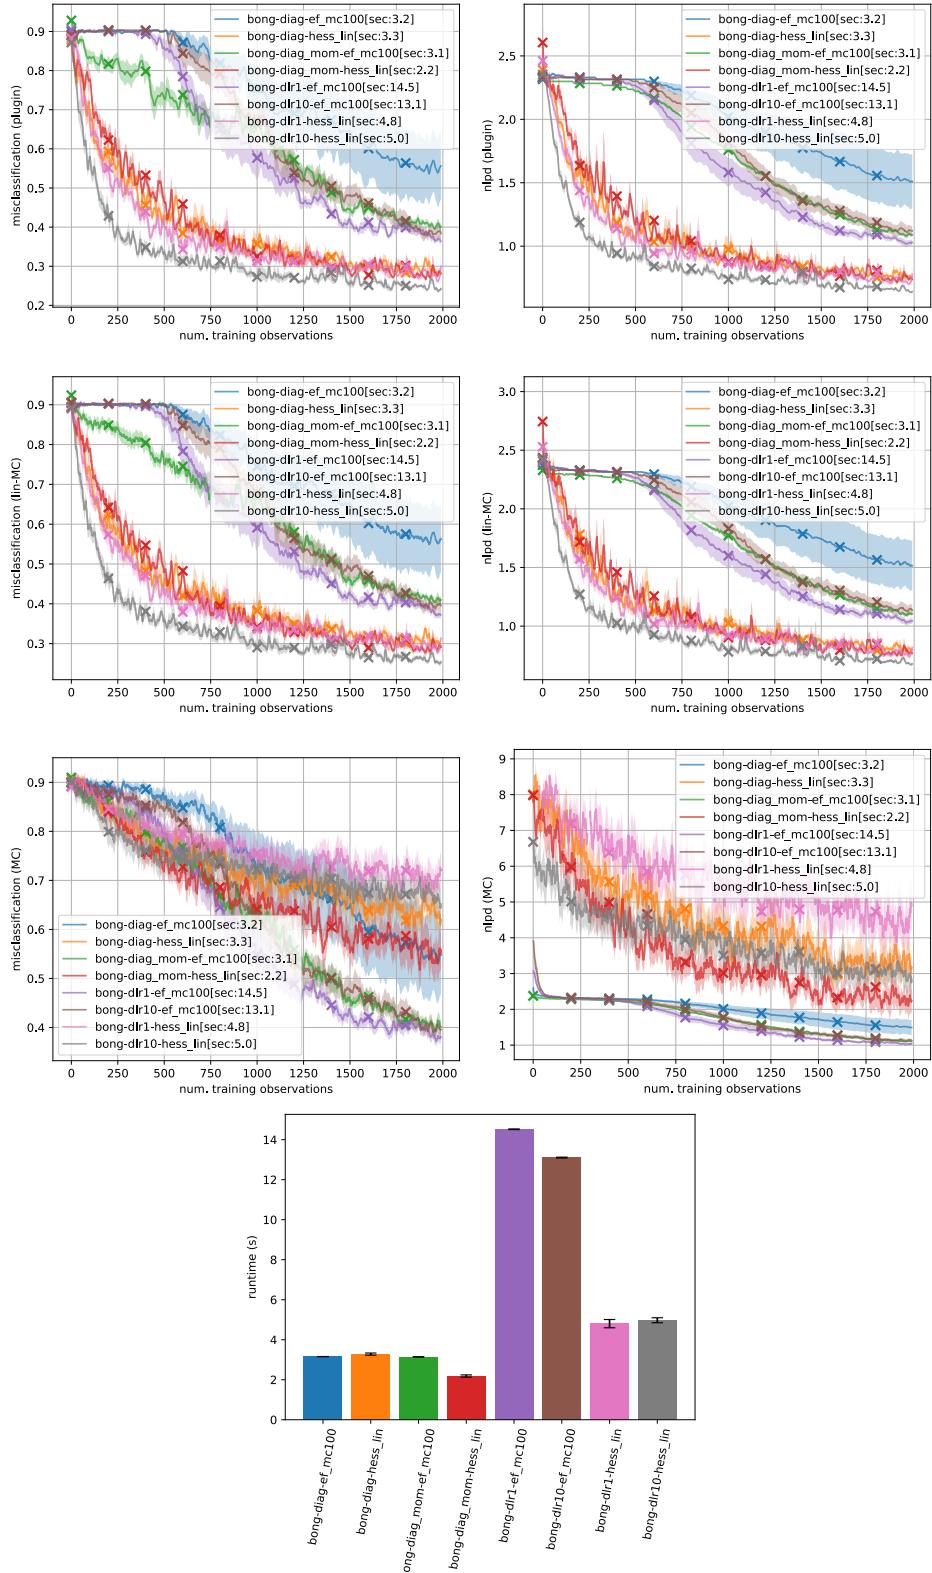


Figure 6: MNIST results for BONG variants.

B.3 SARCOS dataset

In addition to MNIST, we report experiments on the SARCOS regression dataset ($D = 22$, $N_{\text{train}} = 44,484$, $N_{\text{test}} = 4449$, $C = 1$). This dataset derives from an inverse dynamics problem for a seven degrees-of-freedom SARCOS anthropomorphic robot arm. The task is to map from a 21-dimensional input space (7 joint positions, 7 joint velocities, 7 joint accelerations) to the corresponding 7 joint torques. Following [Rasmussen and Williams \[2006\]](#), we pick a single target output dimension, so $C = 1$. The data is from <https://gaussianprocess.org/gpml/data/>.

We use a small MLP of size 21-20-20-1, so there are $P = 881$ parameters. For optimizing learning rates for SARCOS, we use grid search on NLPD-PI. We fix the variance of the prior belief state to $\sigma_0^2 = 1.0$, which represents a mild degree of regularization.³ We fix the observation variance to $R_t = 0.1\hat{R}$, where $\hat{R} = \text{Var}(y_{1:T})$ is the maximum likelihood estimate based on the training sequence; we can think of this as a simple empirical Bayes approximation, and the factor of 0.1 accounts for the fact that the variance of the residuals from the learned model will be smaller than from the unconditional baseline. We focus on DLR approximation of rank 10. This gives similar results to full covariance, but is much faster to compute. We also focus on the plugin approximation to NLPD, since the MC approximation gives much worse results (not shown).

B.3.1 Comparison of BONG, BLR, BBB and BOG

In Fig. 7 we show the results of using the LIN-HESS approximation. For 1 iteration per step, we see that BONG and BLR are indistinguishable in performance, and BBB and BOG are also indistinguishable, but much worse. For 10 iterations per step, we see that BBB improves significantly, and approaches BONG and BLR. However, BLR and BBB are now about 10 times slower. (In practice, the slowdown is less than 10, due to constant factors of the implementation.) (Note that BONG and BOG always use a single iteration, so their performance does not change.)

In Fig. 8 we show the results of using the MC-EF approximation with $MC = 100$ samples. The trends are similar to the LIN-HESS case. In particular, for $I = 1$, BONG and BLR are similar, with BONG having a slight edge; and for $I = 10$, BBB catches up with both BONG and BLR, with BOG always in last place. Finally, we see that the performance of MC-EF is slightly worse than LIN-HESS when $I = 1$, but catches up with $I = 10$. However, in larger scale experiments, we usually find that LIN-HESS is significantly better than MC-EF, even with $I = 10$.

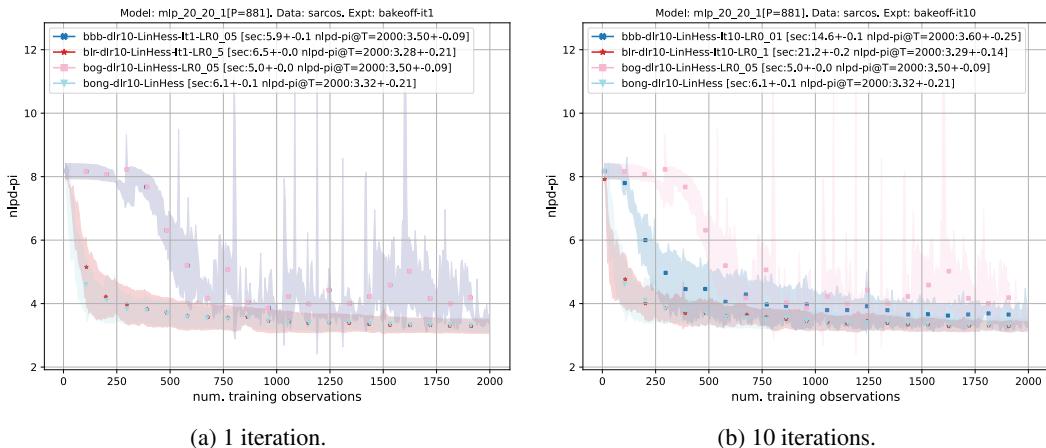


Figure 7: Predictive performance on SARCOS using MLP 21-20-20-1 with DLR rank 10. Error bars represent ± 1 standard deviation computed from 3 random trials, randomizing over data order and initial state μ_0 . (a) We show all 4 algorithms combined with LIN-HESS approximation and $I = 1$. (b) Same as (a) but with $I = 10$.

³This value was based on a small amount of trial and error. Using a smaller value of σ_0 results in underfitting relative to a linear least squares baseline, and using a much larger value results in unstable posterior covariances, causing the NLPD-MC samples to result in NaNs after a few hundred steps.

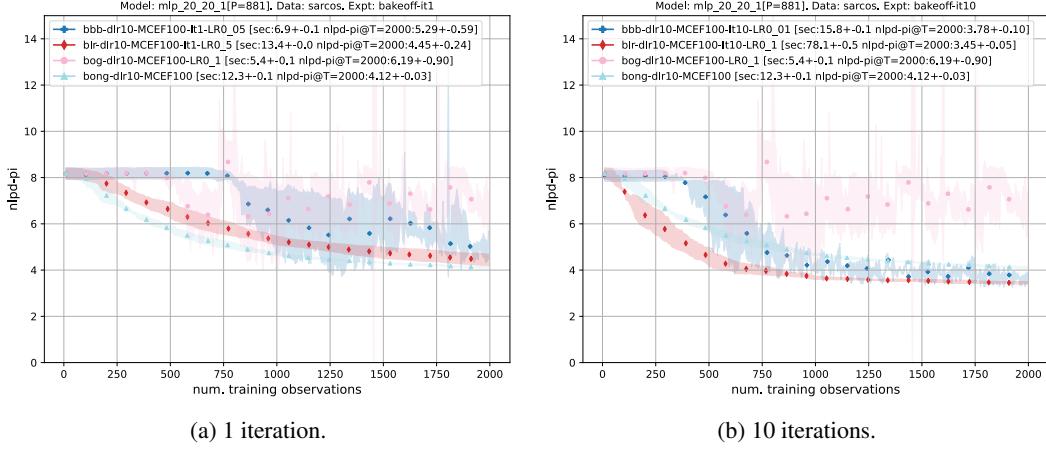


Figure 8: Same as Fig. 7 except we use MC-EF approximation with $MC = 100$.

B.3.2 Learning rate sensitivity

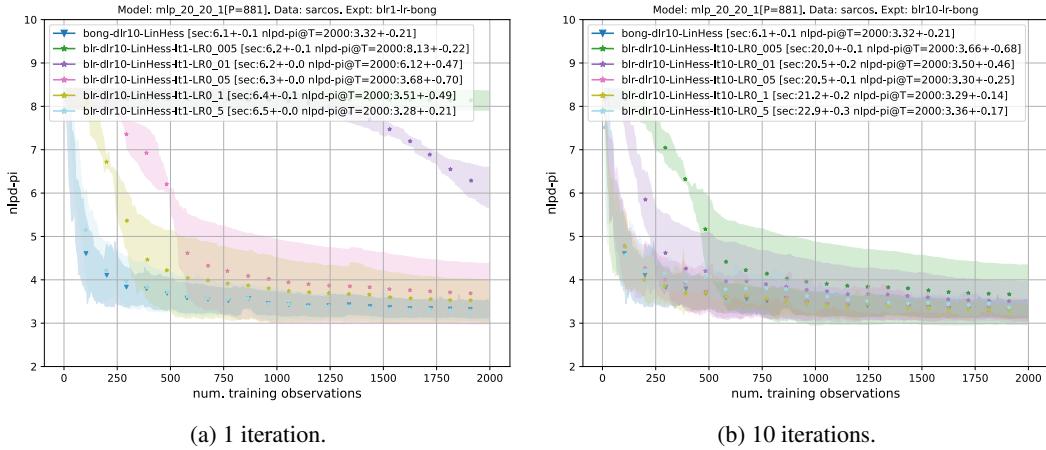


Figure 9: Same setup as Fig. 7, except now we plot performance for BLR for 5 different learning rates. We also show BONG as a baseline, which uses a fixed learning rate step size of 1.0.

In Fig. 9 we show the test set performance for BLR (with LIN-HESS approximation) for 5 different learning rates (namely 5×10^{-3} , 1×10^{-2} , 5×10^{-2} , 1×10^{-1} , and 5×10^{-1}).

When using 1 iteration per step, the best learning rate is $\alpha = 0.5$, which is also the value chosen based on validation set performance. With this value, BLR matches BONG. For other learning rates, BLR performance is much worse. When using 10 iterations per step, there are several learning rates all of which give performance as good as BONG.

In Fig. 10, we show the analogous plot for BBB. When using 1 iteration per step, all learning rates result in poor performance, with many resulting in NaNs. When using 10 iterations per step, there are some learning rates that enable BBB to get close to (but still not match) the performance of BONG.

Finally, in Fig. 11a, we show the analogous plot for BOG with LIN-HESS, and in Fig. 11b with MC-EF, where results are much worse.

Overall we conclude that all the methods (except BONG) are quite sensitive to the learning rate. In our experiments, we pick a value based on performance on a validation set, but in the truly online setting, where there is just a single data stream, picking an optimal learning rate is difficult, which is an additional advantage of BONG.

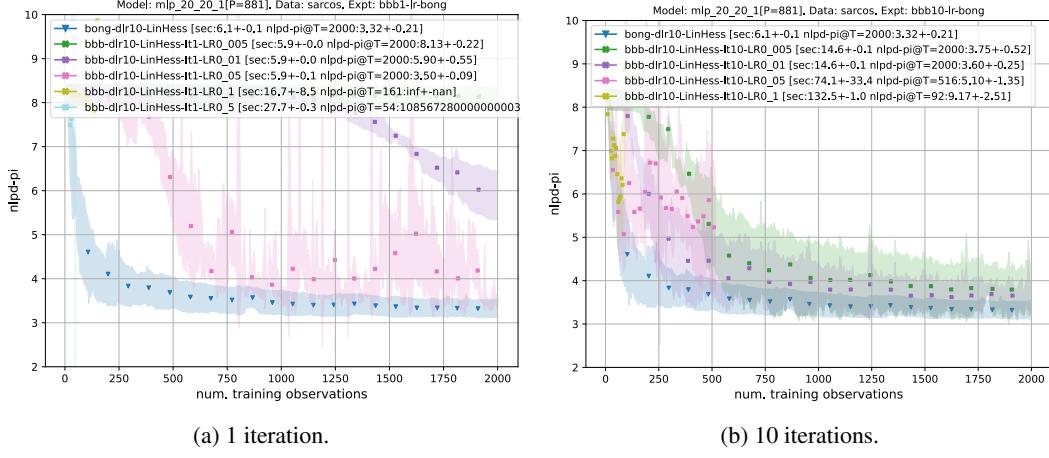


Figure 10: Same setup as Fig. 7 except now we plot performance for BBB for 5 different learning rates. We also show BONG as a baseline.

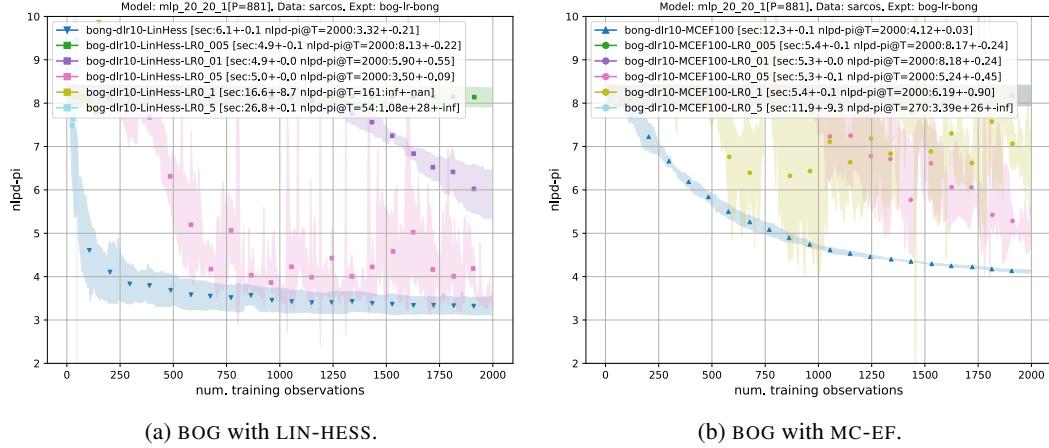


Figure 11: We plot performance for BOG for 5 different learning rates. We also show BONG as a baseline. (a) LIN-HESS approximation. (b) MC-EF approximation.

C Proof of Propositions 4.1 and 4.2

Proposition 4.1 To ease notation we write the natural parameters of the prior as $\psi_{t|t-1} = [\chi_{t|t-1}; \nu_{t|t-1}]$, which can be interpreted as the prior sufficient statistics and prior sample size. Note that \mathbf{x}_t can be omitted as a constant. Based on the prior $q_{\psi_{t|t-1}}(\theta_t)$ the exact posterior is

$$p(\theta_t | \mathcal{D}_t) \propto q_{\psi_{t|t-1}}(\theta_t) p_t(\mathbf{y}_t | \theta_t) \quad (33)$$

$$\propto \exp(\chi_{t|t-1}^\top \theta_t - \nu_{t|t-1} A(\theta_t)) \exp(\theta_t^\top \mathbf{y}_t - A(\theta_t)) \quad (34)$$

$$\propto q_{\psi_t}(\theta_t) \quad (35)$$

$$\psi_t = \begin{bmatrix} \chi_{t|t-1} + \mathbf{y}_t \\ \nu_{t|t-1} + 1 \end{bmatrix} \quad (36)$$

For BONG, we first note the dual parameter is given by

$$\rho_{t|t-1} = \mathbb{E}_{\theta_t \sim q_{\psi_{t|t-1}}} [T(\theta_t)] \quad (37)$$

$$= \mathbb{E}_{\theta_t \sim q_{\psi_{t|t-1}}} \left[\begin{bmatrix} \theta_t \\ -A(\theta_t) \end{bmatrix} \right] \quad (38)$$

Therefore the natural gradient in Eq. (5) is

$$\nabla_{\boldsymbol{\rho}_{t|t-1}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\psi_{t|t-1}}} [\log p(\mathbf{y}_t | \boldsymbol{\theta}_t)] = \nabla_{\boldsymbol{\rho}_{t|t-1}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\psi_{t|t-1}}} [\boldsymbol{\theta}_t^\top \mathbf{y}_t - A(\boldsymbol{\theta}_t)] \quad (39)$$

$$= \nabla_{\boldsymbol{\rho}_{t|t-1}} \boldsymbol{\rho}_{t|t-1}^\top \begin{bmatrix} \mathbf{y}_t \\ 1 \end{bmatrix} \quad (40)$$

$$= \begin{bmatrix} \mathbf{y}_t \\ 1 \end{bmatrix} \quad (41)$$

Therefore the BONG update yields $\psi_t = [\mathbf{x}_{t|t-1} + \mathbf{y}_t; \nu_{t|t-1} + 1]$ in agreement with Eq. (36). \square

Proposition 4.2 The intuition behind this proof is as follows. For the mean update in Eq. (9) $\nabla_{\boldsymbol{\theta}_t} \ell_t$ is linear in $\boldsymbol{\theta}_t$ so the expectation equals the value at the mean. For the covariance update in Eq. (10) $\nabla_{\boldsymbol{\theta}_t}^2 \ell_t$ is independent of $\boldsymbol{\theta}_t$ so we can drop the expectation operator. The tricky part is why we need different linearizations for the Hessians to agree. It has to do with making the Hessian of the NN disappear (as in GGN). In the Gaussian approximation this happens when the predicted mean ($\hat{\mathbf{y}}_t = f_t(\boldsymbol{\theta}_t)$) is linear in $\boldsymbol{\theta}_t$. In the plugin approximation it happens when the outcome-dependent part of the loglikelihood ($h_t(\boldsymbol{\theta}_t)^\top \mathbf{y}_t$) is linear in $\boldsymbol{\theta}_t$. In the latter case the only nonlinear term remaining in the log-likelihood is the log-partition A , and the two methods end up agreeing because of the property that the Hessian of the log-partition equals the conditional variance \mathbf{R}_t .

Formally, under the linear(h)-Gaussian approximation in Eqs. (14) and (17) the expected gradient and Hessian can be calculated directly:

$$\mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\psi_{t|t-1}}} [\nabla_{\boldsymbol{\theta}_t} \log \bar{p}_t(\mathbf{y}_t | \boldsymbol{\theta}_t)] = \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\psi_{t|t-1}}} [\mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t - \mathbf{H}_t(\boldsymbol{\theta}_t - \boldsymbol{\mu}_{t|t-1}))] \quad (42)$$

$$= \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (43)$$

$$\mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\psi_{t|t-1}}} [\nabla_{\boldsymbol{\theta}_t}^2 \log \bar{p}_t(\mathbf{y}_t | \boldsymbol{\theta}_t)] = \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\psi_{t|t-1}}} [-\mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t] \quad (44)$$

$$= -\mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t \quad (45)$$

For the linear(f)-delta approximation we use the properties of exponential families that (1) the gradient of the log-partition A with respect to the natural parameter $f_t(\boldsymbol{\theta}_t)$ equals the expectation parameter $h_t(\boldsymbol{\theta}_t)$, (2) the Hessian of the log-partition with respect to the natural parameter equals the conditional variance, and consequently (3) the Jacobian of the expectation parameter with respect to the natural parameter equals the conditional variance \mathbf{R}_t :

$$\nabla_{\boldsymbol{\eta} = f_t(\boldsymbol{\mu}_{t|t-1})} A(\boldsymbol{\eta}) = h_t(\boldsymbol{\mu}_{t|t-1}) = \hat{\mathbf{y}} \quad (46)$$

$$\nabla_{\boldsymbol{\eta} = f_t(\boldsymbol{\mu}_{t|t-1})}^2 A(\boldsymbol{\eta}) = \mathbb{V}[\mathbf{y}_t | \boldsymbol{\theta}_t = \boldsymbol{\mu}_{t|t-1}] = \mathbf{R}_t \quad (47)$$

$$\frac{\partial h_t(\boldsymbol{\theta}_t)}{\partial f_t(\boldsymbol{\theta}_t)} \Big|_{\boldsymbol{\theta}_t = \boldsymbol{\mu}_{t|t-1}} = \mathbf{R}_t \quad (48)$$

The last of these implies $\mathbf{F}_t = \mathbf{R}_t^{-1} \mathbf{H}_t$. Therefore the expected gradient and Hessian can be calculated as

$$\mathbb{E}_{\boldsymbol{\theta}_t \sim \delta_{\boldsymbol{\mu}_{t|t-1}}} [\nabla_{\boldsymbol{\theta}_t} \log \bar{p}_t(\mathbf{y}_t | \boldsymbol{\theta}_t)] = \nabla_{\boldsymbol{\theta}_t = \boldsymbol{\mu}_{t|t-1}} \log \bar{p}_t(\mathbf{y}_t | \boldsymbol{\theta}_t) \quad (49)$$

$$= \mathbf{F}_t^\top \mathbf{y}_t - \mathbf{F}_t^\top \nabla_{\boldsymbol{\eta} = f_t(\boldsymbol{\mu}_{t|t-1})} A(\boldsymbol{\eta}) \quad (50)$$

$$= \mathbf{F}_t^\top (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (51)$$

$$= \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (52)$$

$$\mathbb{E}_{\boldsymbol{\theta}_t \sim \delta_{\boldsymbol{\mu}_{t|t-1}}} [\nabla_{\boldsymbol{\theta}_t}^2 \log \bar{p}_t(\mathbf{y}_t | \boldsymbol{\theta}_t)] = \nabla_{\boldsymbol{\theta}_t = \boldsymbol{\mu}_{t|t-1}}^2 \log \bar{p}_t(\mathbf{y}_t | \boldsymbol{\theta}_t) \quad (53)$$

$$= -\mathbf{F}_t^\top \left(\nabla_{\boldsymbol{\eta} = f_t(\boldsymbol{\mu}_{t|t-1})}^2 A(\boldsymbol{\eta}) \right) \mathbf{F}_t \quad (54)$$

$$= -\mathbf{F}_t^\top \mathbf{R}_t \mathbf{F}_t \quad (55)$$

$$= -\mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t \quad (56)$$

\square

D Mirror descent formulation

In this section we give a more detailed derivation of BONG as mirror descent and use this to give two alternative interpretations of how BONG approximates exact VB: (1) by approximating the expected NLL as linear in the expectation parameter ρ , or (2) by replacing an implicit update with an explicit one.

Assume the variational family is an exponential one as introduced at the end of Section 3, with natural and dual parameters ψ and ρ , sufficient statistics $T(\theta)$, and log-partition $\Phi(\psi)$:

$$q_\psi(\theta) = \exp(\psi^\top T(\theta) - \Phi(\psi)) \quad (57)$$

$$\rho = \mathbb{E}_{\theta \sim q_\psi}[T(\theta)] \quad (58)$$

We first review how NGD on an exponential family is a special case of mirror descent [Khan and Rue, 2023, Martens, 2020]. The mirror map is the gradient of the log-partition, which satisfies the thermodynamic identity

$$\rho = \nabla \Phi(\psi) \quad (59)$$

This is a bijection when Φ is convex (which includes the cases we study), so we can implicitly treat ψ and ρ as functions of each other. Given a loss function $L(\psi)$, MD iteratively solves the local optimization problem

$$\psi_{i+1} = \arg \min_{\psi} \langle \nabla_{\rho_i} L(\psi_i), \rho \rangle + \frac{1}{\alpha} \mathbb{D}_\Phi(\psi_i, \psi) \quad (60)$$

The first term is a linear (in ρ) approximation of L about the previous iteration ψ_i and the second term is the Bregman divergence

$$\mathbb{D}_\Phi(\psi_i, \psi_{i+1}) = \Phi(\psi_i) - \Phi(\psi_{i+1}) - (\psi_i - \psi_{i+1})^\top \rho_{i+1} \quad (61)$$

The Bregman divergence acts as a regularizer toward ψ_i and captures the intrinsic geometry of the parameter space because of its equivalence with the (reverse) KL divergence

$$D_{\text{KL}}(q_{\psi_{i+1}} \| q_{\psi_i}) = \mathbb{E}_{\theta \sim q_{\psi_{i+1}}}[(\psi_{i+1} - \psi_i)^\top T(\theta) + \Phi(\psi_i) - \Phi(\psi_{i+1})] \quad (62)$$

$$= \mathbb{D}_\Phi(\psi_i, \psi_{i+1}) \quad (63)$$

Importantly, this recursive regularizer is not part of the loss and serves only to define an iterated algorithm that converges to a local minimum of L . Solving Eq. (60) by differentiating by ρ yields the MD update

$$\psi_{i+1} = \psi_i - \alpha \nabla_{\rho_i} L(\psi_i) \quad (64)$$

Because the Fisher matrix for an exponential family is $\mathbf{F}_\psi = \partial \rho / \partial \psi$, this is equivalent to NGD with respect to ψ . Khan and Rue [2023] offer this as a derivation of the BLR, when $L(\psi)$ is taken to be the variational loss from Eq. (1).

By applying this analysis to the online setting, our approach can be seen to follow from two insights. First, the online variational loss in Eq. (2) already includes KL divergence from the previous step, so we do not need the artificial regularizer in Eq. (60). That is, if we start from the online variational problem in Eq. (4) and define $L_t(\psi)$ as the expected NLL,

$$L_t(\psi) = -\mathbb{E}_{\theta_t \sim q_\psi}[\log p(\mathbf{y}_t | f_t(\theta_t))] \quad (65)$$

then replacing $L_t(\psi)$ with a linear approximation based at $\psi_{t|t-1}$ and applying Eq. (63) leads to

$$\psi_t = \arg \min_{\psi} \langle \nabla_{\rho_{t|t-1}} L_t(\psi_{t|t-1}), \rho \rangle + \mathbb{D}_\Phi(\psi_{t|t-1}, \psi) \quad (66)$$

By comparing to Eq. (60) we see this defines an MD algorithm with unit learning rate that works in a single step rather than by iterating. Paralleling the derivation of Eq. (64) from Eq. (60) we get

$$\psi_t = \psi_{t|t-1} - \nabla_{\rho_{t|t-1}} L_t(\psi_{t|t-1}) \quad (67)$$

which matches the BONG update in Eq. (5). Thus BONG can be seen as an approximate solution of the online variational problem in Eq. (4) based on linearizing the expected NLL wrt ρ . (Note this is different from the assumption underlying BONG-LIN that the NLL is linear in θ_t .)

Second, in the conjugate case, this linearity assumption is true: L_t is linear in ρ (see proof of Proposition 4.1). Therefore 66 is equivalent to solving Eq. 4 exactly:

$$\psi_t = \arg \min_{\psi} L_t(\psi) + \mathbb{D}_{\Phi}(\psi_{t|t-1}, \psi) \quad (68)$$

This recapitulates Proposition 4.1 that BONG is Bayes optimal in the conjugate case. In general the exact solution to Eq. (68) is

$$\psi_t = \psi_{t|t-1} - \nabla_{\rho_t} L_t(\psi_t) \quad (69)$$

This is an implicit update because the gradient is evaluated at the (unknown) posterior, whereas Eq. (67) is an explicit update because it evaluates the gradient at the prior. (In the Gaussian case these can be shown to match the implicit and explicit RVGA updates of [Lambert et al., 2021].) Therefore BONG can be also interpreted as an approximation of exact VB that replaces the implicit update, Eq. (69), with an explicit update, Eq. (67).

E Derivations

This section derives the update equations for all 80 algorithms we investigate (Table 3 plus the MC-HESS and LIN-EF variants). In Appendix E.6 we also translate the BLR algorithms from our online setting back to the batch setting used in [Khan and Rue, 2023].

For an exponential variational family with natural parameters ψ and dual parameters ρ , we can derive all 16 methods (BONG, BLR, BOG, BBB under all four Hessian approximations) from four quantities:

$$\nabla_{\rho_{t,i-1}} \mathbb{E}_{\theta_t \sim q_{\psi_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\theta_t))] \quad (70)$$

$$\nabla_{\rho_{t,i-1}} D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) \quad (71)$$

$$\nabla_{\psi_{t,i-1}} \mathbb{E}_{\theta_t \sim q_{\psi_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\theta_t))] \quad (72)$$

$$\nabla_{\psi_{t,i-1}} D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) \quad (73)$$

The NGD methods (BONG and BLR) use gradients with respect to $\rho_{t,i-1}$ while the GD methods (BOG and BBB) use gradients with respect to $\psi_{t,i-1}$. For BONG and BOG the D_{KL} term is not relevant, and there is no inner loop so $\psi_{t,i-1} = \psi_{t|t-1}$ and $\mathbf{g}_{t,i} = \mathbf{g}_t$, $\mathbf{G}_{t,i} = \mathbf{G}_t$.

When ψ is not the natural parameter of an exponential family we must explicitly compute the inverse-Fisher preconditioner for the NGD methods. Therefore the updates can be derived from these three quantities:

$$\mathbf{F}_{\psi_{t,i-1}} \quad (74)$$

$$\nabla_{\psi_{t,i-1}} \mathbb{E}_{\theta_t \sim q_{\psi_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\theta_t))] \quad (75)$$

$$\nabla_{\psi_{t,i-1}} D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) \quad (76)$$

We will frequently use Bonnet's and Price's theorems [Bonnet, 1964, Price, 1958]

$$\nabla_{\mu_{t,i-1}} \mathbb{E}_{\mathcal{N}(\mu_{t,i-1}, \Sigma_{t,i-1})} [\log p(\mathbf{y}_t | f_t(\theta_t))] = \mathbb{E}_{\mathcal{N}(\mu_{t,i-1}, \Sigma_{t,i-1})} [\nabla_{\theta_t} \log p(\mathbf{y}_t | f_t(\theta_t))] \quad (77)$$

$$= \mathbf{g}_{t,i} \quad (78)$$

$$\nabla_{\Sigma_{t,i-1}} \mathbb{E}_{\mathcal{N}(\mu_{t,i-1}, \Sigma_{t,i-1})} [\log p(\mathbf{y}_t | f_t(\theta_t))] = \frac{1}{2} \mathbb{E}_{\mathcal{N}(\mu_{t,i-1}, \Sigma_{t,i-1})} [\nabla_{\theta_t}^2 \log p(\mathbf{y}_t | f_t(\theta_t))] \quad (79)$$

$$= \frac{1}{2} \mathbf{G}_{t,i} \quad (80)$$

For diagonal Gaussians with covariance $\text{Diag}(\sigma^2)$, Price's theorem also implies⁴

$$\nabla_{\sigma_{t,i-1}^2} \mathbb{E}_{\theta_t \sim \mathcal{N}(\mu_{t,i-1}, \Sigma_{t,i-1})} [\log p(\mathbf{y}_t | f_t(\theta_t))] = \frac{1}{2} \text{diag}(\mathbf{G}_{t,i}) \quad (81)$$

Update equations for MC-HESS, MC-EF and LIN-EF methods are displayed together in the subsections that follow, because for the most part they differ only in the choice of $\mathbf{G}_t^{\text{MC-HESS}}$, $\mathbf{G}_t^{\text{MC-EF}}$ or $\mathbf{G}_t^{\text{LIN-EF}}$ to

⁴We use $\text{diag}(\mathbf{A})$ to denote the vector of diagonal elements of matrix \mathbf{A} and $\text{Diag}(\mathbf{v})$ to denote the matrix whose diagonal entries are \mathbf{v} and off-diagonal entries are 0.

approximate \mathbf{G}_t and \mathbf{g}_t^{MC} or $\mathbf{g}_t^{\text{LIN}}$ to approximate \mathbf{g}_t . We note cases where decomposing $\mathbf{G}_t^{\text{MC-EF}} = \hat{\mathbf{G}}_t^{(1:M)} \hat{\mathbf{G}}_t^{(1:M)\top}$ or $\mathbf{G}_t^{\text{LIN-EF}} = \mathbf{g}_t^{\text{LIN}} (\mathbf{g}_t^{\text{LIN}})^\top$ allows a more efficient update.

We derive updates for BONG-LIN-HESS and BOG-LIN-HESS from the corresponding BONG-MC-HESS and BOG-MC-HESS updates using Proposition 4.2 which entails substituting

$$\mathbf{g}_t \rightarrow \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (82)$$

$$\mathbf{G}_t \rightarrow -\mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t \quad (83)$$

For the algorithms with inner loops (BLR and BBB) we adapt the notation of Section 4.4 as follows:

$$\mathbf{y}_{t,i} = h_t(\boldsymbol{\mu}_{t,i-1}) \quad (84)$$

$$\mathbf{H}_{t,i} = \text{jac}(h_t(\cdot))(\boldsymbol{\mu}_{t,i-1}) \quad (85)$$

$$\mathbf{R}_{t,i} = \mathbb{V}[\mathbf{y}_t | \boldsymbol{\theta}_t = \boldsymbol{\mu}_{t,i-1}] \quad (86)$$

$$\mathbf{g}_{t,i} = \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) \quad (87)$$

$$\mathbf{G}_{t,i} = -\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i} \quad (88)$$

This corresponds to basing the linear(f)-Gaussian and linear(h)-delta approximations at $\boldsymbol{\mu}_{t,i-1}$ instead of $\boldsymbol{\mu}_{t|t-1}$. Thus the updates for BLR-LIN-HESS and BBB-LIN-HESS are obtained by substituting

$$\mathbf{g}_{t,i} \rightarrow \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) \quad (89)$$

$$\mathbf{G}_{t,i} \rightarrow -\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i} \quad (90)$$

E.1 Full covariance Gaussian, natural parameters

The natural and dual parameters for a general Gaussian are given by

$$\boldsymbol{\psi}_{t,i-1}^{(1)} = \boldsymbol{\Sigma}_{t,i-1}^{-1} \boldsymbol{\mu}_{t,i-1} \quad \boldsymbol{\rho}_{t,i-1}^{(1)} = \boldsymbol{\mu}_{t,i-1} \quad (91)$$

$$\boldsymbol{\psi}_{t,i-1}^{(2)} = -\frac{1}{2} \boldsymbol{\Sigma}_{t,i-1}^{-1} \quad \boldsymbol{\rho}_{t,i-1}^{(2)} = \boldsymbol{\mu}_{t,i-1} \boldsymbol{\mu}_{t,i-1}^\top + \boldsymbol{\Sigma}_{t,i-1} \quad (92)$$

Inverting these relationships gives

$$\boldsymbol{\mu}_{t,i-1} = -\frac{1}{2} \boldsymbol{\psi}_{t,i-1}^{(2)-1} \boldsymbol{\psi}_{t,i-1}^{(1)} = \boldsymbol{\rho}_{t,i-1}^{(1)} \quad (93)$$

$$\boldsymbol{\Sigma}_{t,i-1} = -\frac{1}{2} \boldsymbol{\psi}_{t,i-1}^{(2)-1} = \boldsymbol{\rho}_{t,i-1}^{(2)} - \boldsymbol{\rho}_{t,i-1}^{(1)} \boldsymbol{\rho}_{t,i-1}^{(1)\top} \quad (94)$$

The KL divergence in the VI loss is

$$\begin{aligned} D_{\mathbb{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) &= \frac{1}{2} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1})^\top \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \\ &\quad + \frac{1}{2} \text{Tr}(\boldsymbol{\Sigma}_{t|t-1}^{-1} \boldsymbol{\Sigma}_{t,i-1}) - \frac{1}{2} \log |\boldsymbol{\Sigma}_{t,i-1}| + \text{const} \end{aligned} \quad (95)$$

with gradients

$$\nabla_{\boldsymbol{\mu}_{t,i-1}} D_{\mathbb{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) = \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \quad (96)$$

$$\nabla_{\boldsymbol{\Sigma}_{t,i-1}} D_{\mathbb{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) = \frac{1}{2} (\boldsymbol{\Sigma}_{t|t-1}^{-1} - \boldsymbol{\Sigma}_{t,i-1}^{-1}) \quad (97)$$

Following Appendix C of [Khan et al., 2018a], for any scalar function ℓ the chain rule gives

$$\nabla_{\boldsymbol{\rho}_{t,i-1}^{(1)}} \ell = \frac{\partial \boldsymbol{\mu}_{t,i-1}}{\partial \boldsymbol{\rho}_{t,i-1}^{(1)}} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell + \frac{\partial \boldsymbol{\Sigma}_{t,i-1}}{\partial \boldsymbol{\rho}_{t,i-1}^{(1)}} \nabla_{\boldsymbol{\Sigma}_{t,i-1}} \ell \quad (98)$$

$$= \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell - 2 (\nabla_{\boldsymbol{\Sigma}_{t,i-1}} \ell) \boldsymbol{\mu}_{t,i-1} \quad (99)$$

$$\nabla_{\boldsymbol{\rho}_{t,i-1}^{(2)}} \ell = \frac{\partial \boldsymbol{\mu}_{t,i-1}}{\partial \boldsymbol{\rho}_{t,i-1}^{(2)}} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell + \frac{\partial \boldsymbol{\Sigma}_{t,i-1}}{\partial \boldsymbol{\rho}_{t,i-1}^{(2)}} \nabla_{\boldsymbol{\Sigma}_{t,i-1}} \ell \quad (100)$$

$$= \nabla_{\boldsymbol{\Sigma}_{t,i-1}} \ell \quad (101)$$

Therefore

$$\nabla_{\boldsymbol{\rho}_{t,i-1}^{(1)}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = \mathbf{g}_{t,i} - \mathbf{G}_{t,i} \boldsymbol{\mu}_{t,i-1} \quad (102)$$

$$\nabla_{\boldsymbol{\rho}_{t,i-1}^{(2)}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = \frac{1}{2} \mathbf{G}_{t,i} \quad (103)$$

and

$$\nabla_{\boldsymbol{\rho}_{t,i-1}^{(1)}} D_{\text{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) = \boldsymbol{\Sigma}_{t,i-1}^{-1} \boldsymbol{\mu}_{t,i-1} - \boldsymbol{\Sigma}_{t|t-1}^{-1} \boldsymbol{\mu}_{t|t-1} \quad (104)$$

$$\nabla_{\boldsymbol{\rho}_{t,i-1}^{(2)}} D_{\text{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) = \frac{1}{2} \left(\boldsymbol{\Sigma}_{t|t-1}^{-1} - \boldsymbol{\Sigma}_{t,i-1}^{-1} \right) \quad (105)$$

Following the same approach for ψ gives

$$\nabla_{\boldsymbol{\psi}_{t,i-1}^{(1)}} \ell = \frac{\partial \boldsymbol{\mu}_{t,i-1}}{\partial \boldsymbol{\psi}_{t,i-1}^{(1)}} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell + \frac{\partial \boldsymbol{\Sigma}_{t,i-1}}{\partial \boldsymbol{\psi}_{t,i-1}^{(1)}} \nabla_{\boldsymbol{\Sigma}_{t,i-1}} \ell \quad (106)$$

$$= -\frac{1}{2} \boldsymbol{\psi}_{t,i-1}^{(2)-1} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell \quad (107)$$

$$= \boldsymbol{\Sigma}_{t,i-1} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell \quad (108)$$

$$\nabla_{\boldsymbol{\psi}_{t,i-1}^{(2)}} \ell = \frac{\partial \boldsymbol{\mu}_{t,i-1}}{\partial \boldsymbol{\psi}_{t,i-1}^{(2)}} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell + \frac{\partial \boldsymbol{\Sigma}_{t,i-1}}{\partial \boldsymbol{\psi}_{t,i-1}^{(2)}} \nabla_{\boldsymbol{\Sigma}_{t,i-1}} \ell \quad (109)$$

$$= \frac{1}{2} \boldsymbol{\psi}_{t,i-1}^{(2)-1} (\nabla_{\boldsymbol{\mu}_{t,i-1}} \ell) \boldsymbol{\psi}_{t,i-1}^{(1)\top} \boldsymbol{\psi}_{t,i-1}^{(2)-1} + \frac{1}{2} \boldsymbol{\psi}_{t,i-1}^{(2)-1} (\nabla_{\boldsymbol{\Sigma}_{t,i-1}} \ell) \boldsymbol{\psi}_{t,i-1}^{(2)-1} \quad (110)$$

$$= 2 \boldsymbol{\Sigma}_{t,i-1} (\nabla_{\boldsymbol{\mu}_{t,i-1}} \ell) \boldsymbol{\mu}_{t,i-1}^\top + 2 \boldsymbol{\Sigma}_{t,i-1} (\nabla_{\boldsymbol{\Sigma}_{t,i-1}} \ell) \boldsymbol{\Sigma}_{t,i-1} \quad (111)$$

Therefore

$$\nabla_{\boldsymbol{\psi}_{t,i-1}^{(1)}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = \boldsymbol{\Sigma}_{t,i-1} \mathbf{g}_{t,i} \quad (112)$$

$$\nabla_{\boldsymbol{\psi}_{t,i-1}^{(2)}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = 2 \boldsymbol{\Sigma}_{t,i-1} \mathbf{g}_{t,i} \boldsymbol{\mu}_{t,i-1}^\top + \boldsymbol{\Sigma}_{t,i-1} \mathbf{G}_{t,i} \boldsymbol{\Sigma}_{t,i-1} \quad (113)$$

and

$$\nabla_{\boldsymbol{\psi}_{t,i-1}^{(1)}} D_{\text{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) = \boldsymbol{\Sigma}_{t,i-1} \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \quad (114)$$

$$\begin{aligned} \nabla_{\boldsymbol{\psi}_{t,i-1}^{(2)}} D_{\text{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) &= 2 \boldsymbol{\Sigma}_{t,i-1} \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \boldsymbol{\mu}_{t,i-1}^\top \\ &\quad + \boldsymbol{\Sigma}_{t,i-1} \left(\boldsymbol{\Sigma}_{t|t-1}^{-1} - \boldsymbol{\Sigma}_{t,i-1}^{-1} \right) \boldsymbol{\Sigma}_{t,i-1} \end{aligned} \quad (115)$$

E.1.1 BONG FC (explicit RVGA)

Substituting Eqs. (102) and (103) into Eq. (5) gives

$$\boldsymbol{\psi}_t^{(1)} = \boldsymbol{\psi}_{t|t-1}^{(1)} + \mathbf{g}_t - \mathbf{G}_t \boldsymbol{\mu}_{t|t-1} \quad (116)$$

$$\boldsymbol{\psi}_t^{(2)} = \boldsymbol{\psi}_{t|t-1}^{(2)} + \frac{1}{2} \mathbf{G}_t \quad (117)$$

Translating to $(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ gives the BONG-FC update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \boldsymbol{\Sigma}_t \mathbf{g}_t \quad (118)$$

$$\boldsymbol{\Sigma}_t^{-1} = \boldsymbol{\Sigma}_{t|t-1}^{-1} - \mathbf{G}_t \quad (119)$$

This is equivalent to the explicit update form of RVGA [Lambert et al., 2021]. Using $\mathbf{G}_t^{\text{MC-HESS}}$ this update takes $O(P^3)$ because of the matrix inversion. Using $\mathbf{G}_t^{\text{MC-EF}}$ and the Woodbury matrix identity we can write the update in a form that takes $O(MP^2 + M^3)$:

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \mathbf{K}_t \mathbf{1}_M \quad (120)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t|t-1} - \mathbf{K}_t \hat{\mathbf{G}}_t^{(1:M)\top} \boldsymbol{\Sigma}_{t|t-1} \quad (121)$$

$$\mathbf{K}_t = \boldsymbol{\Sigma}_{t|t-1} \hat{\mathbf{G}}_t^{(1:M)} \left(M \mathbf{I}_M + \hat{\mathbf{G}}_t^{(1:M)\top} \boldsymbol{\Sigma}_{t|t-1} \hat{\mathbf{G}}_t^{(1:M)} \right)^{-1} \quad (122)$$

Likewise using $\mathbf{G}_t^{\text{LIN-EF}}$ takes $O(P^2)$:

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \mathbf{K}_t \quad (123)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t|t-1} - \mathbf{K}_t (\mathbf{g}_t^{\text{LIN}})^\top \boldsymbol{\Sigma}_{t|t-1} \quad (124)$$

$$\mathbf{K}_t = \frac{\boldsymbol{\Sigma}_{t|t-1} \mathbf{g}_t^{\text{LIN}}}{1 + (\mathbf{g}_t^{\text{LIN}})^\top \boldsymbol{\Sigma}_{t|t-1} \mathbf{g}_t^{\text{LIN}}} \quad (125)$$

E.1.2 BONG-LIN-HESS FC (CM-EKF)

Applying Proposition 4.2 to Eqs. (118) and (119) gives the BONG-LIN-HESS-FC update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \boldsymbol{\Sigma}_t \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (126)$$

$$\boldsymbol{\Sigma}_t^{-1} = \boldsymbol{\Sigma}_{t|t-1}^{-1} + \mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t \quad (127)$$

This is equivalent to CM-EKF [Tronarp et al., 2018, Ollivier, 2018] and can be rewritten using the Woodbury identity in a form that takes $O(CP^2 + C^3)$:

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \mathbf{K}_t (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (128)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t|t-1} - \mathbf{K}_t \mathbf{H}_t \boldsymbol{\Sigma}_{t|t-1} \quad (129)$$

$$\mathbf{K}_t = \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^\top (\mathbf{R}_t + \mathbf{H}_t \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^\top)^{-1} \quad (130)$$

E.1.3 BLR FC

Substituting Eqs. (102) to (105) into Eq. (31) gives

$$\boldsymbol{\psi}_{t,i}^{(1)} = \boldsymbol{\psi}_{t,i-1}^{(1)} + \alpha \left(\mathbf{g}_{t,i} - \mathbf{G}_{t,i} \boldsymbol{\mu}_{t,i-1} - \boldsymbol{\Sigma}_{t,i-1}^{-1} \boldsymbol{\mu}_{t,i-1} + \boldsymbol{\Sigma}_{t|t-1}^{-1} \boldsymbol{\mu}_{t|t-1} \right) \quad (131)$$

$$\boldsymbol{\psi}_{t,i}^{(2)} = \boldsymbol{\psi}_{t,i-1}^{(2)} + \frac{\alpha}{2} \left(\mathbf{G}_{t,i} + \boldsymbol{\Sigma}_{t,i-1}^{-1} - \boldsymbol{\Sigma}_{t|t-1}^{-1} \right) \quad (132)$$

Translating to $(\boldsymbol{\mu}_{t,i}, \boldsymbol{\Sigma}_{t,i})$ gives the BLR-FC update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\Sigma}_{t,i} \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \boldsymbol{\Sigma}_{t,i} \mathbf{g}_{t,i} \quad (133)$$

$$\boldsymbol{\Sigma}_{t,i}^{-1} = (1 - \alpha) \boldsymbol{\Sigma}_{t,i-1}^{-1} + \alpha \boldsymbol{\Sigma}_{t|t-1}^{-1} - \alpha \mathbf{G}_{t,i} \quad (134)$$

This update takes $O(P^3)$ per iteration because of the matrix inversion. In Appendix E.1.1 we were able to use the Woodbury identity to exploit the low rank of $\mathbf{G}_t^{\text{MC-EF}}$ and $\mathbf{G}_t^{\text{LIN-EF}}$ and obtain BONG updates with complexity quadratic in P . This does not appear possible with Eq. (134) because of the extra precision term on the RHS (applying Woodbury would require inverting $(1 - \alpha) \boldsymbol{\Sigma}_{t,i-1}^{-1} + \alpha \boldsymbol{\Sigma}_{t|t-1}^{-1}$). Therefore unlike BONG-FC, BLR-FC requires time cubic in the model size, for reasons that can be traced back to the KL term in Eq. (31).

E.1.4 BLR-LIN-HESS FC

Applying Proposition 4.2 to Eqs. (133) and (134) gives the BLR-LIN-HESS-FC update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\Sigma}_{t,i} \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \boldsymbol{\Sigma}_{t,i} \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) \quad (135)$$

$$\boldsymbol{\Sigma}_{t,i}^{-1} = (1 - \alpha) \boldsymbol{\Sigma}_{t,i-1}^{-1} + \alpha \boldsymbol{\Sigma}_{t|t-1}^{-1} + \alpha \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i} \quad (136)$$

This update takes $O(P^3)$ per iteration because of the matrix inversion.

E.1.5 BOG FC

Substituting Eqs. (112) and (113) into Eq. (29) gives

$$\boldsymbol{\psi}_t^{(1)} = \boldsymbol{\psi}_{t|t-1}^{(1)} + \alpha \boldsymbol{\Sigma}_{t|t-1} \mathbf{g}_t \quad (137)$$

$$\boldsymbol{\psi}_t^{(2)} = \boldsymbol{\psi}_{t|t-1}^{(2)} + 2\alpha \boldsymbol{\Sigma}_{t|t-1} \mathbf{g}_t \boldsymbol{\mu}_{t|t-1}^\top + \alpha \boldsymbol{\Sigma}_{t|t-1} \mathbf{G}_t \boldsymbol{\Sigma}_{t|t-1} \quad (138)$$

Translating to $(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ gives the BOG-FC update

$$\boldsymbol{\mu}_t = \boldsymbol{\Sigma}_t \boldsymbol{\Sigma}_{t|t-1}^{-1} \boldsymbol{\mu}_{t|t-1} + \alpha \boldsymbol{\Sigma}_t \boldsymbol{\Sigma}_{t|t-1} \mathbf{g}_t \quad (139)$$

$$\boldsymbol{\Sigma}_t^{-1} = \boldsymbol{\Sigma}_{t|t-1}^{-1} - 4\alpha \boldsymbol{\Sigma}_{t|t-1} \mathbf{g}_t \boldsymbol{\mu}_{t|t-1}^\top - 2\alpha \boldsymbol{\Sigma}_{t|t-1} \mathbf{G}_t \boldsymbol{\Sigma}_{t|t-1} \quad (140)$$

This update takes $O(P^3)$ because of the matrix inversion. The greater cost of the BOG-FC update relative to BONG-FC can be traced to the difference between GD and NGD: the NLL gradients wrt $\boldsymbol{\psi}_{t|t-1}$ in Eqs. (112) and (113) are more complicated than the gradients wrt $\boldsymbol{\rho}_{t|t-1}$ in Eqs. (102) and (103).

E.1.6 BOG-LIN-HESS FC

Applying Proposition 4.2 to Eqs. (139) and (140) gives the BOG-LIN-HESS-FC update

$$\boldsymbol{\mu}_t = \boldsymbol{\Sigma}_t \boldsymbol{\Sigma}_{t|t-1}^{-1} \boldsymbol{\mu}_{t|t-1} + \alpha \boldsymbol{\Sigma}_t \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (141)$$

$$\boldsymbol{\Sigma}_t^{-1} = \boldsymbol{\Sigma}_{t|t-1}^{-1} - 4\alpha \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \boldsymbol{\mu}_{t|t-1}^\top + 2\alpha \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t \boldsymbol{\Sigma}_{t|t-1} \quad (142)$$

This update takes $O(P^3)$ because of the matrix inversion.

E.1.7 BBB FC

Substituting Eqs. (112) to (115) into Eq. (32) gives

$$\boldsymbol{\psi}_{t,i}^{(1)} = \boldsymbol{\psi}_{t,i-1}^{(1)} + \alpha \boldsymbol{\Sigma}_{t,i-1} \mathbf{g}_{t,i} - \alpha \boldsymbol{\Sigma}_{t,i-1} \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \quad (143)$$

$$\begin{aligned} \boldsymbol{\psi}_t^{(2)} = & \boldsymbol{\psi}_{t,i-1}^{(2)} + 2\alpha \boldsymbol{\Sigma}_{t,i-1} \mathbf{g}_{t,i} \boldsymbol{\mu}_{t,i-1}^\top + \alpha \boldsymbol{\Sigma}_{t|t-1} \mathbf{G}_t \boldsymbol{\Sigma}_{t|t-1} \\ & - 2\alpha \boldsymbol{\Sigma}_{t,i-1} \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \boldsymbol{\mu}_{t,i-1}^\top - \alpha \boldsymbol{\Sigma}_{t,i-1} \left(\boldsymbol{\Sigma}_{t|t-1}^{-1} - \boldsymbol{\Sigma}_{t,i-1}^{-1} \right) \boldsymbol{\Sigma}_{t,i-1} \end{aligned} \quad (144)$$

Translating to $(\boldsymbol{\mu}_{t,i}, \boldsymbol{\Sigma}_{t,i})$ gives the BBB-FC update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\Sigma}_{t,i} \boldsymbol{\Sigma}_{t,i-1}^{-1} \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\Sigma}_{t,i} \boldsymbol{\Sigma}_{t,i-1} \left(\mathbf{g}_{t,i} + \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) \right) \quad (145)$$

$$\boldsymbol{\Sigma}_{t,i}^{-1} = \boldsymbol{\Sigma}_{t,i-1}^{-1} - 2\alpha \boldsymbol{\Sigma}_{t,i-1} \left(\begin{array}{l} 2\boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) \boldsymbol{\mu}_{t,i-1}^\top + \mathbf{I}_P \\ + 2\mathbf{g}_{t,i} \boldsymbol{\mu}_{t,i-1}^\top + (\mathbf{G}_{t,i} - \boldsymbol{\Sigma}_{t|t-1}^{-1}) \boldsymbol{\Sigma}_{t,i-1} \end{array} \right) \quad (146)$$

This update takes $O(P^3)$ per iteration because of the matrix inversion.

E.1.8 BBB-LIN-HESS FC

Applying Proposition 4.2 to Eqs. (145) and (146) gives the BBB-LIN-HESS-FC update

$$\begin{aligned} \boldsymbol{\mu}_{t,i} = & \boldsymbol{\Sigma}_{t,i} \boldsymbol{\Sigma}_{t,i-1}^{-1} \boldsymbol{\mu}_{t,i-1} \\ & + \alpha \boldsymbol{\Sigma}_{t,i} \boldsymbol{\Sigma}_{t,i-1} \left(\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) + \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) \right) \end{aligned} \quad (147)$$

$$\boldsymbol{\Sigma}_{t,i}^{-1} = \boldsymbol{\Sigma}_{t,i-1}^{-1} - 2\alpha \boldsymbol{\Sigma}_{t,i-1} \left(\begin{array}{l} 2\boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) \boldsymbol{\mu}_{t,i-1}^\top + \mathbf{I}_P \\ + 2\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) \boldsymbol{\mu}_{t,i-1}^\top \\ - (\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i} + \boldsymbol{\Sigma}_{t|t-1}^{-1}) \boldsymbol{\Sigma}_{t,i-1} \end{array} \right) \quad (148)$$

This update takes $O(P^3)$ per iteration because of the matrix inversion.

E.2 Full covariance Gaussian, Moment parameters

The Bonnet and Price theorems give

$$\nabla_{\boldsymbol{\mu}_{t,i-1}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = \mathbf{g}_{t,i} \quad (149)$$

$$\nabla_{\boldsymbol{\Sigma}_{t,i-1}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = \frac{1}{2} \mathbf{G}_{t,i} \quad (150)$$

From Appendix E.1 we have

$$\nabla_{\boldsymbol{\mu}_{t,i-1}} D_{\text{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} \| q_{\boldsymbol{\psi}_{t|t-1}}) = \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \quad (151)$$

$$\nabla_{\boldsymbol{\Sigma}_{t,i-1}} D_{\text{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} \| q_{\boldsymbol{\psi}_{t|t-1}}) = \frac{1}{2} (\boldsymbol{\Sigma}_{t|t-1}^{-1} - \boldsymbol{\Sigma}_{t,i-1}^{-1}) \quad (152)$$

We write the Fisher with respect to the moment parameters $\boldsymbol{\psi} = (\boldsymbol{\mu}, \text{vec}(\boldsymbol{\Sigma}))$ as a block matrix:

$$\mathbf{F} = \begin{bmatrix} \mathbf{F}_{\boldsymbol{\mu}, \boldsymbol{\mu}} & \mathbf{F}_{\boldsymbol{\mu}, \boldsymbol{\Sigma}} \\ \mathbf{F}_{\boldsymbol{\Sigma}, \boldsymbol{\mu}} & \mathbf{F}_{\boldsymbol{\Sigma}, \boldsymbol{\Sigma}} \end{bmatrix} \quad (153)$$

The blocks can be calculated by the second-order Fisher formula

$$\mathbf{F}_{\boldsymbol{\mu}, \boldsymbol{\mu}} = -\mathbb{E}_{q_{\boldsymbol{\psi}}} [\nabla_{\boldsymbol{\mu}, \boldsymbol{\mu}} \log q_{\boldsymbol{\psi}}(\boldsymbol{\theta})] \quad (154)$$

$$= \boldsymbol{\Sigma}^{-1} \quad (155)$$

$$\mathbf{F}_{\boldsymbol{\mu}, \boldsymbol{\Sigma}} = -\mathbb{E}_{q_{\boldsymbol{\psi}}} [\nabla_{\boldsymbol{\mu}, \boldsymbol{\Sigma}} \log q_{\boldsymbol{\psi}}(\boldsymbol{\theta})] \quad (156)$$

$$= -\mathbb{E}_{q_{\boldsymbol{\psi}}} [(\nabla_{\boldsymbol{\Sigma}} \boldsymbol{\Sigma}^{-1}) (\boldsymbol{\theta} - \boldsymbol{\mu})] \quad (157)$$

$$= \mathbf{0} \quad (158)$$

$$\mathbf{F}_{\boldsymbol{\Sigma}, \boldsymbol{\Sigma}} = -\mathbb{E}_{q_{\boldsymbol{\psi}}} [\nabla_{\boldsymbol{\Sigma}, \boldsymbol{\Sigma}} \log q_{\boldsymbol{\psi}}(\boldsymbol{\theta})] \quad (159)$$

$$= -\mathbb{E}_{q_{\boldsymbol{\psi}}} [\nabla_{\boldsymbol{\Sigma}} (\frac{1}{2} \boldsymbol{\Sigma}^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu}) (\boldsymbol{\theta} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} - \frac{1}{2} \boldsymbol{\Sigma}^{-1})] \quad (160)$$

$$= -\frac{1}{2} \mathbb{E}_{q_{\boldsymbol{\psi}}} \left[+\boldsymbol{\Sigma}^{-1} (\boldsymbol{\theta} - \boldsymbol{\mu}) (\boldsymbol{\theta} - \boldsymbol{\mu})^T (\nabla_{\boldsymbol{\Sigma}} \boldsymbol{\Sigma}^{-1}) - (\nabla_{\boldsymbol{\Sigma}} \boldsymbol{\Sigma}^{-1}) \right] \quad (161)$$

$$= -\frac{1}{2} \nabla_{\boldsymbol{\Sigma}} \boldsymbol{\Sigma}^{-1} \quad (162)$$

$$= \frac{1}{2} \boldsymbol{\Sigma}^{-1} \otimes \boldsymbol{\Sigma}^{-1} \quad (163)$$

In the final line we used

$$\nabla_{\Sigma_{k\ell}} (\boldsymbol{\Sigma}^{-1})_{ij} = -(\boldsymbol{\Sigma}^{-1})_{ik} (\boldsymbol{\Sigma}^{-1})_{j\ell} \quad (164)$$

$$= -(\boldsymbol{\Sigma}^{-1} \otimes \boldsymbol{\Sigma}^{-1})_{ij, k\ell} \quad (165)$$

with ij and $k\ell$ treated as composite indices in $[P^2]$. Therefore the preconditioner in the NGD methods is

$$\mathbf{F}_{\boldsymbol{\psi}_{t,i-1}}^{-1} = \begin{bmatrix} \boldsymbol{\Sigma}_{t,i-1} & \mathbf{0} \\ \mathbf{0} & 2\boldsymbol{\Sigma}_{t,i-1} \otimes \boldsymbol{\Sigma}_{t,i-1} \end{bmatrix} \quad (166)$$

E.2.1 BONG FC, Moment

Substituting Eqs. (149), (150) and (166) into Eq. (3) gives the BONG-FC_MOM update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \boldsymbol{\Sigma}_{t|t-1} \mathbf{g}_t \quad (167)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t|t-1} + \boldsymbol{\Sigma}_{t|t-1} \mathbf{G}_t \boldsymbol{\Sigma}_{t|t-1} \quad (168)$$

This update takes $O(P^3)$ using $\mathbf{G}_t^{\text{MC-HESS}}$, $O(MP^2)$ using $\mathbf{G}_t^{\text{MC-EF}}$, and $O(P^2)$ using $\mathbf{G}_t^{\text{LIN-EF}}$. The update is similar to the BONG-FC update except that Eq. (168) ignores the $\hat{\mathbf{G}}_t^{(1:M)^T} \boldsymbol{\Sigma}_{t|t-1} \hat{\mathbf{G}}_t^{(1:M)}$ term in Eq. (122) or the $(\mathbf{g}_t^{\text{LIN}})^T \boldsymbol{\Sigma}_{t|t-1} \mathbf{g}_t^{\text{LIN}}$ term in Eq. (125) which estimate the epistemic part of predictive uncertainty.

E.2.2 BONG-LIN-HESS FC, Moment

Applying Proposition 4.2 to Eqs. (167) and (168) gives the BONG-LIN-HESS-FC_MOM update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^T \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (169)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t|t-1} - \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^T \mathbf{R}_t^{-1} \mathbf{H}_t \boldsymbol{\Sigma}_{t|t-1} \quad (170)$$

This update takes $O(CP^2)$.

E.2.3 BLR FC, Moment

Substituting Eqs. (149) to (152) and (166) into Eq. (30) gives the BLR-FC_MOM update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\Sigma}_{t,i-1} \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \boldsymbol{\Sigma}_{t,i-1} \mathbf{g}_{t,i} \quad (171)$$

$$\boldsymbol{\Sigma}_{t,i} = (1 + \alpha) \boldsymbol{\Sigma}_{t,i-1} + \alpha \boldsymbol{\Sigma}_{t,i-1} \left(\mathbf{G}_{t,i} - \boldsymbol{\Sigma}_{t|t-1}^{-1} \right) \boldsymbol{\Sigma}_{t,i-1} \quad (172)$$

This update takes $O(P^3)$ per iteration because of the matrix inversion. Using $\mathbf{G}_t^{\text{MC-EF}} = -\frac{1}{M} \hat{\mathbf{G}}_t^{(1:M)} \hat{\mathbf{G}}_t^{(1:M)\top}$ or $\mathbf{G}^{\text{LIN-EF}} = -\mathbf{g}_t^{\text{LIN}} (\mathbf{g}_t^{\text{LIN}})^\top$ allows the BONG-FC_MOM covariance update in Eq. (168) to scale quadratically, but this does not help here. Instead the BLR-FC_MOM update scales cubically because of the additional $\boldsymbol{\Sigma}_{t|t-1}^{-1}$ term that comes from the KL divergence in the VI objective.

E.2.4 BLR-LIN-HESS FC, Moment

Applying Proposition 4.2 to Eqs. (171) and (172) gives the BLR-LIN-HESS-FC_MOM update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\Sigma}_{t,i-1} \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \boldsymbol{\Sigma}_{t,i-1} \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) \quad (173)$$

$$\boldsymbol{\Sigma}_{t,i} = (1 + \alpha) \boldsymbol{\Sigma}_{t,i-1} - \alpha \boldsymbol{\Sigma}_{t,i-1} \left(\boldsymbol{\Sigma}_{t|t-1}^{-1} + \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i} \right) \boldsymbol{\Sigma}_{t,i-1} \quad (174)$$

This update takes $O(P^3)$ per iteration because of the matrix inversion in the $\boldsymbol{\Sigma}_{t|t-1}^{-1}$ term that comes from the KL divergence in the VI objective.

E.2.5 BOG FC, Moment

Substituting Eqs. (149) and (150) into Eq. (29) gives the BOG-FC_MOM update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \alpha \mathbf{g}_t \quad (175)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t|t-1} + \frac{\alpha}{2} \mathbf{G}_t \quad (176)$$

Note the mean update is vanilla online gradient descent (OGD) and does not depend on the covariance. This update takes $O(MP^2)$ using $\mathbf{G}_t^{\text{MC-HESS}}$ or $\mathbf{G}_t^{\text{MC-EF}}$ and $O(P^2)$ using $\mathbf{G}_t^{\text{LIN-EF}}$.

E.2.6 BOG-LIN-HESS FC, Moment

Applying Proposition 4.2 to Eqs. (175) and (176) gives the BOG-LIN-HESS-FC_MOM update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \alpha \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (177)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t|t-1} - \frac{\alpha}{2} \mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t \quad (178)$$

This update takes $O(CP^2)$.

E.2.7 BBB FC, Moment

Substituting Eqs. (149) to (152) into Eq. (32) gives the BBB-FC_MOM

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \mathbf{g}_{t,i} \quad (179)$$

$$\boldsymbol{\Sigma}_{t,i} = \boldsymbol{\Sigma}_{t,i-1} + \frac{\alpha}{2} \left(\boldsymbol{\Sigma}_{t,i-1}^{-1} - \boldsymbol{\Sigma}_{t|t-1}^{-1} + \mathbf{G}_{t,i} \right) \quad (180)$$

This update takes $O(P^3)$ per iteration because of the matrix inversion, which traces back to the VI objective. Comparing to the BOG-FC_MOM update in Eqs. (175) and (176) (which has quadratic complexity in P), the extra terms here come from the KL part of Eq. (32).

E.2.8 BBB-LIN-HESS FC, Moment

Applying Proposition 4.2 to Eqs. (179) and (180) gives the BBB-LIN-HESS-FC_MOM update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\Sigma}_{t|t-1}^{-1} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) \quad (181)$$

$$\boldsymbol{\Sigma}_t = \boldsymbol{\Sigma}_{t,i-1} + \frac{\alpha}{2} \left(\boldsymbol{\Sigma}_{t,i-1}^{-1} - \boldsymbol{\Sigma}_{t|t-1}^{-1} - \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i} \right) \quad (182)$$

This update takes $O(P^3)$ per iteration because of the matrix inversion. Comparing to the BOG-LIN-HESS-FC_MOM update in Eqs. (177) and (178) (which has quadratic complexity in P), the extra terms here come from the KL part of Eq. (32).

E.3 Diagonal Gaussian, Natural parameters

Throughout this subsection, vector multiplication and exponents are elementwise.

The natural and dual parameters for a diagonal Gaussian are given by

$$\psi_{t,i-1}^{(1)} = \sigma_{t,i-1}^{-2} \mu_{t,i-1} \quad \rho_{t,i-1}^{(1)} = \mu_{t,i-1} \quad (183)$$

$$\psi_{t,i-1}^{(2)} = -\frac{1}{2} \sigma_{t,i-1}^{-2} \quad \rho_{t,i-1}^{(2)} = \mu_{t,i-1} \mu_{t,i-1}^\top + \sigma_{t,i-1}^2 \quad (184)$$

Inverting these relationships gives

$$\mu_{t,i-1} = -\frac{1}{2} \left(\psi_{t,i-1}^{(2)} \right)^{-1} \psi_{t,i-1}^{(1)} = \rho_{t,i-1}^{(1)} \quad (185)$$

$$\sigma_{t,i-1}^2 = -\frac{1}{2} \left(\psi_{t,i-1}^{(2)} \right)^{-1} = \rho_{t,i-1}^{(2)} - \left(\rho_{t,i-1}^{(1)} \right)^2 \quad (186)$$

The KL divergence in the VI loss is

$$D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) = \frac{1}{2} (\mu_{t,i-1} - \mu_{t|t-1})^2 \sigma_{t|t-1}^{-2} + \frac{1}{2} \sum \left(\sigma_{t|t-1}^{-2} \sigma_{t,i-1}^2 - \log \sigma_{t,i-1}^2 \right) + \text{const} \quad (187)$$

with gradients

$$\nabla_{\mu_{t,i-1}} D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) = \sigma_{t|t-1}^{-2} (\mu_{t,i-1} - \mu_{t|t-1}) \quad (188)$$

$$\nabla_{\sigma_{t,i-1}^2} D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) = \frac{1}{2} \left(\sigma_{t|t-1}^{-2} - \sigma_{t,i-1}^{-2} \right) \quad (189)$$

For any scalar function ℓ the chain rule gives

$$\nabla_{\rho_{t,i-1}^{(1)}} \ell = \frac{\partial \mu_{t,i-1}}{\partial \rho_{t,i-1}^{(1)}} \nabla_{\mu_{t,i-1}} \ell + \frac{\partial \sigma_{t,i-1}^2}{\partial \rho_{t,i-1}^{(1)}} \nabla_{\sigma_{t,i-1}^2} \ell \quad (190)$$

$$= \nabla_{\mu_{t,i-1}} \ell - 2 \mu_{t,i-1} \nabla_{\sigma_{t,i-1}^2} \ell \quad (191)$$

$$\nabla_{\rho_{t,i-1}^{(2)}} \ell = \frac{\partial \mu_{t,i-1}}{\partial \rho_{t,i-1}^{(2)}} \nabla_{\mu_{t,i-1}} \ell + \frac{\partial \sigma_{t,i-1}^2}{\partial \rho_{t,i-1}^{(2)}} \nabla_{\sigma_{t,i-1}^2} \ell \quad (192)$$

$$= \nabla_{\sigma_{t,i-1}^2} \ell \quad (193)$$

Therefore

$$\nabla_{\rho_{t,i-1}^{(1)}} \mathbb{E}_{\theta_t \sim q_{\psi_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = \mathbf{g}_{t,i} - \text{diag}(\mathbf{G}_{t,i}) \mu_{t,i-1} \quad (194)$$

$$\nabla_{\rho_{t,i-1}^{(2)}} \mathbb{E}_{\theta_t \sim q_{\psi_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = \frac{1}{2} \text{diag}(\mathbf{G}_{t,i}) \quad (195)$$

and

$$\nabla_{\rho_{t,i-1}^{(1)}} D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) = \sigma_{t,i-1}^{-2} \mu_{t,i-1} - \sigma_{t|t-1}^{-2} \mu_{t|t-1} \quad (196)$$

$$\nabla_{\rho_{t,i-1}^{(2)}} D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) = \frac{1}{2} \left(\sigma_{t|t-1}^{-2} - \sigma_{t,i-1}^{-2} \right) \quad (197)$$

Following the same approach for ψ gives

$$\nabla_{\psi_{t,i-1}^{(1)}} \ell = \frac{\partial \boldsymbol{\mu}_{t,i-1}}{\partial \psi_{t,i-1}^{(1)}} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell + \frac{\partial \boldsymbol{\sigma}_{t,i-1}^2}{\partial \psi_{t,i-1}^{(1)}} \nabla_{\boldsymbol{\Sigma}_{t,i-1}} \ell \quad (198)$$

$$= -\frac{1}{2} \left(\psi_{t,i-1}^{(2)} \right)^{-1} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell \quad (199)$$

$$= \boldsymbol{\sigma}_{t,i-1}^2 \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell \quad (200)$$

$$\nabla_{\psi_{t,i-1}^{(2)}} \ell = \frac{\partial \boldsymbol{\mu}_{t,i-1}}{\partial \psi_{t,i-1}^{(2)}} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell + \frac{\partial \boldsymbol{\sigma}_{t,i-1}^2}{\partial \psi_{t,i-1}^{(2)}} \nabla_{\boldsymbol{\sigma}_{t,i-1}^2} \ell \quad (201)$$

$$= \frac{1}{2} \left(\psi_{t,i-1}^{(2)} \right)^{-2} \psi_{t,i-1}^{(1)} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell + \frac{1}{2} \left(\psi_{t,i-1}^{(2)} \right)^{-2} \nabla_{\boldsymbol{\sigma}_{t,i-1}^2} \ell \quad (202)$$

$$= 2\boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\mu}_{t,i-1} \nabla_{\boldsymbol{\mu}_{t,i-1}} \ell + 2\boldsymbol{\sigma}_{t,i-1}^4 \nabla_{\boldsymbol{\sigma}_{t,i-1}^2} \ell \quad (203)$$

Therefore

$$\nabla_{\psi_{t,i-1}^{(1)}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\psi_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = \boldsymbol{\sigma}_{t,i-1}^2 \mathbf{g}_{t,i} \quad (204)$$

$$\nabla_{\psi_{t,i-1}^{(2)}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\psi_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}_t))] = 2\boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\mu}_{t,i-1} \mathbf{g}_{t,i} + \boldsymbol{\sigma}_{t,i-1}^4 \text{diag}(\mathbf{G}_{t,i}) \quad (205)$$

and

$$\nabla_{\psi_{t,i-1}^{(1)}} D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) = \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \quad (206)$$

$$\begin{aligned} \nabla_{\psi_{t,i-1}^{(2)}} D_{\text{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) &= 2\boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} \boldsymbol{\mu}_{t,i-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \\ &\quad + \boldsymbol{\sigma}_{t,i-1}^4 (\boldsymbol{\sigma}_{t|t-1}^{-2} - \boldsymbol{\sigma}_{t,i-1}^{-2}) \end{aligned} \quad (207)$$

Our implementations often make use of the following trick: Suppose $\mathbf{A} \in \mathbb{R}^{n \times m}$ and $\mathbf{B} \in \mathbb{R}^{m \times n}$. Then we can efficiently compute $\text{diag}(\mathbf{AB})$ in $O(mn)$ time using $(\mathbf{AB})_{ii} = \sum_{j=1}^m A_{ij} B_{ji}$.

For MC-HESS methods, we approximate the diagonal of the Hessian for each MC sample $\hat{\theta}_t^{(m)}$ using Hutchinson's trace estimation method [Hutchinson, 1989] which has been used in other DNN optimization papers such as adahessian [Yao et al., 2021]. This involves an extra inner loop with size denoted N .

E.3.1 BONG Diag

Substituting Eqs. (194) and (195) into Eq. (5) gives

$$\psi_t^{(1)} = \psi_{t|t-1}^{(1)} + \mathbf{g}_t - \text{diag}(\mathbf{G}_t) \boldsymbol{\mu}_{t|t-1} \quad (208)$$

$$\psi_t^{(2)} = \psi_{t|t-1}^{(2)} + \frac{1}{2} \text{diag}(\mathbf{G}_t) \quad (209)$$

Translating to $(\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t^2)$ gives the BONG-DIAG update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \boldsymbol{\sigma}_t^2 \mathbf{g}_t \quad (210)$$

$$\boldsymbol{\sigma}_t^{-2} = \boldsymbol{\sigma}_{t|t-1}^{-2} - \text{diag}(\mathbf{G}_t) \quad (211)$$

This update takes $O(MP)$ to estimate \mathbf{G}_t using $\mathbf{G}_t^{\text{MC-EF}}$, $O(NMP)$ using $\mathbf{G}_t^{\text{MC-HESS}}$ and Hutchinson's method, and $O(P)$ using $\mathbf{G}_t^{\text{LIN-EF}}$.

E.3.2 BONG-LIN-HESS Diag (VD-EKF)

Applying Proposition 4.2 to Eqs. (208) and (209) gives the BONG-LIN-HESS-DIAG update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \boldsymbol{\sigma}_t^2 (\mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t)) \quad (212)$$

$$\boldsymbol{\sigma}_t^{-2} = \boldsymbol{\sigma}_{t|t-1}^{-2} + \text{diag}(\mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t) \quad (213)$$

This update is equivalent to VD-EKF [Chang et al., 2022] and takes $O(C^2 P)$.

E.3.3 BLR Diag (VON)

Substituting Eqs. (194) to (197) into Eq. (31) gives

$$\psi_{t,i}^{(1)} = \psi_{t,i-1}^{(1)} + \alpha \left(\mathbf{g}_{t,i} - \text{diag}(\mathbf{G}_{t,i}) \boldsymbol{\mu}_{t,i-1} - \boldsymbol{\sigma}_{t,i-1}^{-2} \boldsymbol{\mu}_{t,i-1} + \boldsymbol{\sigma}_{t|t-1}^{-2} \boldsymbol{\mu}_{t|t-1} \right) \quad (214)$$

$$\psi_{t,i}^{(2)} = \psi_{t,i-1}^{(2)} + \frac{\alpha}{2} \left(\text{diag}(\mathbf{G}_{t,i}) + \boldsymbol{\sigma}_{t,i-1}^{-2} - \boldsymbol{\sigma}_{t|t-1}^{-2} \right) \quad (215)$$

Translating to $(\boldsymbol{\mu}_{t,i}, \boldsymbol{\sigma}_{t,i}^2)$ gives the BLR-DIAG update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\sigma}_{t,i}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \boldsymbol{\sigma}_{t,i}^2 \mathbf{g}_{t,i} \quad (216)$$

$$\boldsymbol{\sigma}_{t,i}^{-2} = (1 - \alpha) \boldsymbol{\sigma}_{t,i-1}^{-2} + \alpha \boldsymbol{\sigma}_{t|t-1}^{-2} - \alpha \text{diag}(\mathbf{G}_{t,i}) \quad (217)$$

This update takes $O(MP)$ per iteration to estimate \mathbf{G}_t using $\mathbf{G}_t^{\text{MC-EF}}$, $O(NMP)$ per iteration using $\mathbf{G}_t^{\text{MC-HESS}}$ and Hutchinson's method, and $O(P)$ per iteration using $\mathbf{G}_t^{\text{LIN-EF}}$.

The MC-HESS and MC-EF versions of this update are respectively equivalent to VON and VOGN [Khan et al. [2018b]] in the batch setting where we replace $q_{\psi_{t|t-1}}$ with a spherical prior $\mathcal{N}(\mathbf{0}, \lambda^{-1} \mathbf{I}_P)$ (see Appendix E.6).

E.3.4 BLR-LIN-HESS Diag

Applying Proposition 4.2 to Eqs. (216) and (217) gives the BLR-LIN-HESS-DIAG update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\sigma}_{t,i}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \boldsymbol{\sigma}_{t,i}^2 (\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i})) \quad (218)$$

$$\boldsymbol{\sigma}_{t,i}^{-2} = (1 - \alpha) \boldsymbol{\sigma}_{t,i-1}^{-2} + \alpha \boldsymbol{\sigma}_{t|t-1}^{-2} + \alpha \text{diag}(\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i}) \quad (219)$$

This update takes $O(C^2 P)$ per iteration.

E.3.5 BOG Diag

Substituting Eqs. (204) and (205) into Eq. (29) gives

$$\psi_t^{(1)} = \psi_{t|t-1}^{(1)} + \alpha \boldsymbol{\sigma}_{t|t-1}^2 \mathbf{g}_t \quad (220)$$

$$\psi_t^{(2)} = \psi_{t|t-1}^{(2)} + 2\alpha \boldsymbol{\sigma}_{t|t-1}^2 \boldsymbol{\mu}_{t|t-1} \mathbf{g}_t + \alpha \boldsymbol{\sigma}_{t|t-1}^4 \text{diag}(\mathbf{G}_t) \quad (221)$$

Translating to $(\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t^2)$ gives the BOG-DIAG update

$$\boldsymbol{\mu}_t = \boldsymbol{\sigma}_t^2 \boldsymbol{\sigma}_{t|t-1}^{-2} \boldsymbol{\mu}_{t|t-1} + \alpha \boldsymbol{\sigma}_t^2 \boldsymbol{\sigma}_{t|t-1}^2 \mathbf{g}_t \quad (222)$$

$$\boldsymbol{\sigma}_t^{-2} = \boldsymbol{\sigma}_{t|t-1}^{-2} - 4\alpha \boldsymbol{\sigma}_{t|t-1}^2 \boldsymbol{\mu}_{t|t-1} \mathbf{g}_t - 2\alpha \boldsymbol{\sigma}_{t|t-1}^4 \text{diag}(\mathbf{G}_t) \quad (223)$$

This update takes $O(MP)$ to estimate \mathbf{G}_t using $\mathbf{G}_t^{\text{MC-EF}}$, $O(NMP)$ using $\mathbf{G}_t^{\text{MC-HESS}}$ and Hutchinson's method, and $O(P)$ using $\mathbf{G}_t^{\text{LIN-EF}}$.

E.3.6 BOG-LIN-HESS Diag

Applying Proposition 4.2 to Eqs. (222) and (223) gives the BOG-LIN-HESS-DIAG update

$$\boldsymbol{\mu}_t = \boldsymbol{\sigma}_t^2 \boldsymbol{\sigma}_{t|t-1}^{-2} \boldsymbol{\mu}_{t|t-1} + \alpha \boldsymbol{\sigma}_t^2 \boldsymbol{\sigma}_{t|t-1}^2 (\mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t)) \quad (224)$$

$$\boldsymbol{\sigma}_t^{-2} = \boldsymbol{\sigma}_{t|t-1}^{-2} - 4\alpha \boldsymbol{\sigma}_{t|t-1}^2 \boldsymbol{\mu}_{t|t-1} (\mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t)) + 2\alpha \boldsymbol{\sigma}_{t|t-1}^4 \text{diag}(\mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t) \quad (225)$$

This update takes $O(C^2 P)$.

E.3.7 BBB Diag

Substituting Eqs. (204) to (207) into Eq. (32) gives

$$\psi_{t,i}^{(1)} = \psi_{t,i-1}^{(1)} + \alpha \boldsymbol{\sigma}_{t,i-1}^2 \mathbf{g}_{t,i} - \alpha \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \quad (226)$$

$$\psi_{t,i}^{(2)} = \psi_{t,i-1}^{(2)} + 2\alpha \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\mu}_{t,i-1} \mathbf{g}_{t,i} + \alpha \boldsymbol{\sigma}_{t,i-1}^4 \text{diag}(\mathbf{G}_{t,i}) \quad (227)$$

$$- 2\alpha \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} \boldsymbol{\mu}_{t,i-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) - \alpha \boldsymbol{\sigma}_{t,i-1}^4 (\boldsymbol{\sigma}_{t|t-1}^{-2} - \boldsymbol{\sigma}_{t,i-1}^{-2})$$

Translating to $(\boldsymbol{\mu}_{t,i}, \boldsymbol{\Sigma}_{t,i})$ gives the BBB-DIAG update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\sigma}_{t,i}^2 \boldsymbol{\sigma}_{t,i-1}^{-2} \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\sigma}_{t,i}^2 \boldsymbol{\sigma}_{t,i-1}^2 \mathbf{g}_{t,i} + \alpha \boldsymbol{\sigma}_{t,i}^2 \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) \quad (228)$$

$$\begin{aligned} \boldsymbol{\sigma}_{t,i}^{-2} = & \boldsymbol{\sigma}_{t,i-1}^{-2} - 4\alpha \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\mu}_{t,i-1} \mathbf{g}_{t,i} - 2\alpha \boldsymbol{\sigma}_{t,i-1}^4 \text{diag}(\mathbf{G}_{t,i}) \\ & + 4\alpha \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} \boldsymbol{\mu}_{t,i-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) + 2\alpha \boldsymbol{\sigma}_{t,i-1}^4 (\boldsymbol{\sigma}_{t|t-1}^{-2} - \boldsymbol{\sigma}_{t,i-1}^{-2}) \end{aligned} \quad (229)$$

This update takes $O(MP)$ per iteration to estimate \mathbf{G}_t using $\mathbf{G}_t^{\text{MC-EF}}$, $O(NMP)$ per iteration using $\mathbf{G}_t^{\text{MC-HESS}}$ and Hutchinson's method, and $O(P)$ per iteration using $\mathbf{G}_t^{\text{LIN-EF}}$.

E.3.8 BBB-LIN-HESS Diag

Applying Proposition 4.2 to Eqs. (228) and (229) gives the BBB-LIN-HESS-DIAG update

$$\begin{aligned} \boldsymbol{\mu}_{t,i} = & \boldsymbol{\sigma}_{t,i}^2 \boldsymbol{\sigma}_{t,i-1}^{-2} \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\sigma}_{t,i}^2 \boldsymbol{\sigma}_{t,i-1}^2 (\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i})) \\ & + \alpha \boldsymbol{\sigma}_{t,i}^2 \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) \end{aligned} \quad (230)$$

$$\begin{aligned} \boldsymbol{\sigma}_{t,i}^{-2} = & \boldsymbol{\sigma}_{t,i-1}^{-2} - 4\alpha \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\mu}_{t,i-1} (\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i})) + 2\alpha \boldsymbol{\sigma}_{t,i-1}^4 \text{diag}(\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i}) \\ & + 4\alpha \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} \boldsymbol{\mu}_{t,i-1} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) + 2\alpha \boldsymbol{\sigma}_{t,i-1}^4 \boldsymbol{\sigma}_{t|t-1}^{-2} - 2\alpha \boldsymbol{\sigma}_{t,i-1}^2 \end{aligned} \quad (231)$$

This update takes $O(C^2 P)$ per iteration.

E.4 Diagonal Gaussian, Moment parameters

Throughout this subsection, vector multiplication and exponents are elementwise.

The Bonnet and Price theorems give

$$\nabla_{\boldsymbol{\mu}_{t,i-1}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}))] = \mathbf{g}_{t,i} \quad (232)$$

$$\nabla_{\boldsymbol{\sigma}_{t,i-1}^2} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | f_t(\boldsymbol{\theta}))] = \frac{1}{2} \text{diag}(\mathbf{G}_{t,i}) \quad (233)$$

From Appendix E.3 we have

$$\nabla_{\boldsymbol{\mu}_{t,i-1}} D_{\mathbb{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} \| q_{\boldsymbol{\psi}_{t|t-1}}) = \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \quad (234)$$

$$\nabla_{\boldsymbol{\sigma}_{t,i-1}^2} D_{\mathbb{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} \| q_{\boldsymbol{\psi}_{t|t-1}}) = \frac{1}{2} (\boldsymbol{\sigma}_{t|t-1}^{-2} - \boldsymbol{\sigma}_{t,i-1}^{-2}) \quad (235)$$

We write the Fisher with respect to the moment parameters $\boldsymbol{\psi} = (\boldsymbol{\mu}, \boldsymbol{\sigma}^2)$ as a block matrix:

$$\mathbf{F}_{\boldsymbol{\psi}} = \begin{bmatrix} \mathbf{F}_{\boldsymbol{\mu}, \boldsymbol{\mu}} & \mathbf{F}_{\boldsymbol{\mu}, \boldsymbol{\sigma}^2} \\ \mathbf{F}_{\boldsymbol{\sigma}^2, \boldsymbol{\mu}} & \mathbf{F}_{\boldsymbol{\sigma}^2, \boldsymbol{\sigma}^2} \end{bmatrix} \quad (236)$$

The blocks can be calculated by the second-order Fisher formula

$$\mathbf{F}_{\boldsymbol{\mu}, \boldsymbol{\mu}} = -\mathbb{E}_{q_{\boldsymbol{\psi}}} [\nabla_{\boldsymbol{\mu}, \boldsymbol{\mu}} \log q_{\boldsymbol{\psi}}(\boldsymbol{\theta})] \quad (237)$$

$$= \text{Diag}(\boldsymbol{\sigma}^{-2}) \quad (238)$$

$$\mathbf{F}_{\boldsymbol{\mu}, \boldsymbol{\sigma}^2} = -\mathbb{E}_{q_{\boldsymbol{\psi}}} [\nabla_{\boldsymbol{\mu}, \boldsymbol{\sigma}^2} \log q_{\boldsymbol{\psi}}(\boldsymbol{\theta})] \quad (239)$$

$$= \mathbb{E}_{q_{\boldsymbol{\psi}}} [\text{Diag}((\boldsymbol{\theta} - \boldsymbol{\mu}) \boldsymbol{\sigma}^{-4})] \quad (240)$$

$$= \mathbf{0} \quad (241)$$

$$\mathbf{F}_{\boldsymbol{\sigma}^2, \boldsymbol{\sigma}^2} = -\mathbb{E}_{q_{\boldsymbol{\psi}}} [\nabla_{\boldsymbol{\sigma}^2, \boldsymbol{\sigma}^2} \log q_{\boldsymbol{\psi}}(\boldsymbol{\theta})] \quad (242)$$

$$= -\mathbb{E}_{q_{\boldsymbol{\psi}}} [\text{Diag}(-(\boldsymbol{\mu} - \boldsymbol{\theta})^2 \boldsymbol{\sigma}^{-6} + \frac{1}{2} \boldsymbol{\sigma}^{-4})] \quad (243)$$

$$= \frac{1}{2} \text{Diag}(\boldsymbol{\sigma}^{-4}) \quad (244)$$

Therefore the preconditioner for the NGD methods is

$$\mathbf{F}_{\boldsymbol{\psi}_{t,i-1}}^{-1} = \begin{bmatrix} \text{Diag}(\boldsymbol{\sigma}_{t,i-1}^2) & \mathbf{0} \\ \mathbf{0} & 2\text{Diag}(\boldsymbol{\sigma}_{t,i-1}^4) \end{bmatrix} \quad (245)$$

E.4.1 BONG Diag, Moment

Substituting Eqs. (232), (233) and (245) into Eq. (3) gives the BONG-DIAG_MOM update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \boldsymbol{\sigma}_{t|t-1}^2 \mathbf{g}_t \quad (246)$$

$$\boldsymbol{\sigma}_t^2 = \boldsymbol{\sigma}_{t|t-1}^2 + \boldsymbol{\sigma}_{t|t-1}^4 \text{diag}(\mathbf{G}_t) \quad (247)$$

This update takes $O(MP)$ to estimate \mathbf{G}_t using $\mathbf{G}_t^{\text{MC-EF}}$, $O(NMP)$ using $\mathbf{G}_t^{\text{MC-HESS}}$ and Hutchinson's method, and $O(P)$ using $\mathbf{G}_t^{\text{LIN-EF}}$.

E.4.2 BONG-LIN-HESS Diag, Moment

Applying Proposition 4.2 to Eqs. (246) and (247) gives the BONG-LIN-HESS-DIAG_MOM update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \boldsymbol{\sigma}_{t|t-1}^2 (\mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t)) \quad (248)$$

$$\boldsymbol{\sigma}_t^2 = \boldsymbol{\sigma}_{t|t-1}^2 - \boldsymbol{\sigma}_{t|t-1}^4 \text{diag}(\mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t) \quad (249)$$

This update takes $O(C^2 P)$.

E.4.3 BLR Diag, Moment

Substituting Eqs. (232) to (235) and (245) into Eq. (30) gives the BLR-DIAG_MOM update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \boldsymbol{\sigma}_{t,i-1}^2 \mathbf{g}_{t,i} \quad (250)$$

$$\boldsymbol{\sigma}_{t,i}^2 = \boldsymbol{\sigma}_{t,i-1}^2 + \alpha \boldsymbol{\sigma}_{t,i-1}^4 (\boldsymbol{\sigma}_{t,i-1}^{-2} - \boldsymbol{\sigma}_{t|t-1}^{-2}) + \alpha \boldsymbol{\sigma}_{t,i-1}^4 \text{diag}(\mathbf{G}_{t,i}) \quad (251)$$

This update takes $O(MP)$ per iteration to estimate \mathbf{G}_t using $\mathbf{G}_t^{\text{MC-EF}}$, $O(NMP)$ per iteration using $\mathbf{G}_t^{\text{MC-HESS}}$ and Hutchinson's method, and $O(P)$ per iteration using $\mathbf{G}_t^{\text{LIN-EF}}$.

E.4.4 BLR-LIN-HESS Diag, Moment

Applying Proposition 4.2 to Eqs. (250) and (251) gives the BLR-LIN-HESS-DIAG_MOM update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\sigma}_{t,i-1}^2 \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \boldsymbol{\sigma}_{t,i-1}^2 (\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i})) \quad (252)$$

$$\boldsymbol{\sigma}_{t,i}^2 = \boldsymbol{\sigma}_{t,i-1}^2 + \alpha \boldsymbol{\sigma}_{t,i-1}^4 (\boldsymbol{\sigma}_{t,i-1}^{-2} - \boldsymbol{\sigma}_{t|t-1}^{-2}) - \alpha \boldsymbol{\sigma}_{t,i-1}^4 \text{diag}(\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i}) \quad (253)$$

This update takes $O(C^2 P)$ per iteration.

E.4.5 BOG Diag, Moment

Substituting Eqs. (232) and (233) into Eq. (29) gives the BOG-DIAG_MOM update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \alpha \mathbf{g}_t \quad (254)$$

$$\boldsymbol{\sigma}_t^2 = \boldsymbol{\sigma}_{t|t-1}^2 + \frac{\alpha}{2} \text{diag}(\mathbf{G}_t) \quad (255)$$

This update takes $O(MP)$ to estimate \mathbf{G}_t using $\mathbf{G}_t^{\text{MC-EF}}$, $O(NMP)$ using $\mathbf{G}_t^{\text{MC-HESS}}$ and Hutchinson's method, and $O(P)$ using $\mathbf{G}_t^{\text{LIN-EF}}$.

E.4.6 BOG-LIN-HESS Diag, Moment

Applying Proposition 4.2 to Eqs. (254) and (255) gives the BOG-LIN-HESS-DIAG_MOM update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \alpha \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (256)$$

$$\boldsymbol{\sigma}_t^2 = \boldsymbol{\sigma}_{t|t-1}^2 - \frac{\alpha}{2} \text{diag}(\mathbf{H}_t^\top \mathbf{R}_t^{-1} \mathbf{H}_t) \quad (257)$$

This update takes $O(C^2 P)$.

E.4.7 BBB Diag, Moment

Substituting Eqs. (232) to (235) into Eq. (32) gives the BBB-DIAG_MOM

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \mathbf{g}_{t,i} \quad (258)$$

$$\boldsymbol{\sigma}_t^2 = \boldsymbol{\sigma}_{t,i-1}^2 + \frac{\alpha}{2} (\boldsymbol{\sigma}_{t,i-1}^{-2} - \boldsymbol{\sigma}_{t|t-1}^{-2}) + \frac{\alpha}{2} \text{diag}(\mathbf{G}_{t,i}) \quad (259)$$

This update takes $O(MP)$ per iteration to estimate \mathbf{G}_t using $\mathbf{G}_t^{\text{MC-EF}}$, $O(NMP)$ per iteration using $\mathbf{G}_t^{\text{MC-HESS}}$ and Hutchinson's method, and $O(P)$ per iteration using $\mathbf{G}_t^{\text{LIN-EF}}$.

This is similar to the original diagonal Gaussian method in [Blundell et al., 2015] except (1) they reparameterize $\boldsymbol{\sigma} = \log(1 + \exp(\boldsymbol{\rho}))$ (elementwise) and do GD on $(\boldsymbol{\mu}, \boldsymbol{\rho})$ instead of $(\boldsymbol{\mu}, \boldsymbol{\sigma}^2)$, and (2) they use the reparameterization trick instead of Price's theorem for calculating the gradient with respect to $\boldsymbol{\rho}$ (via $\boldsymbol{\sigma}$).

E.4.8 BBB-LIN-HESS Diag, Moment

Applying Proposition 4.2 to Eqs. (258) and (259) gives the BBB-LIN-HESS-DIAG_MOM update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \boldsymbol{\sigma}_{t|t-1}^{-2} (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) \quad (260)$$

$$\boldsymbol{\sigma}_t^2 = \boldsymbol{\sigma}_{t,i-1}^2 + \frac{\alpha}{2} (\boldsymbol{\sigma}_{t,i-1}^{-2} - \boldsymbol{\sigma}_{t|t-1}^{-2}) - \frac{\alpha}{2} \text{diag}(\mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i}) \quad (261)$$

This update takes $O(C^2P)$ per iteration.

E.5 Diagonal plus low rank

Assume the prior is given by

$$q_{\boldsymbol{\psi}_{t|t-1}}(\boldsymbol{\theta}_t) = \mathcal{N}\left(\boldsymbol{\theta}_t | \boldsymbol{\mu}_{t|t-1}, \left(\boldsymbol{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top\right)^{-1}\right) \quad (262)$$

with $\mathbf{W} \in \mathbb{R}^{P \times R}$ and diagonal $\boldsymbol{\Upsilon}_{t|t-1} \in \mathbb{R}^{P \times P}$. We sometimes abuse notation by writing $\boldsymbol{\Upsilon}_{t|t-1}$ for the vector $\text{diag}(\boldsymbol{\Upsilon}_{t|t-1})$ when the meaning is clear from context.

Substituting the DLR form in the gradients for the KL divergence derived in Appendix E.2 gives

$$\nabla_{\boldsymbol{\mu}_{t,i-1}} D_{\text{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) = \left(\boldsymbol{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top\right) (\boldsymbol{\mu}_{t,i-1} - \boldsymbol{\mu}_{t|t-1}) \quad (263)$$

$$\nabla_{\boldsymbol{\Sigma}_{t,i-1}} D_{\text{KL}}(q_{\boldsymbol{\psi}_{t,i-1}} | q_{\boldsymbol{\psi}_{t|t-1}}) = \frac{1}{2} \left(\boldsymbol{\Upsilon}_{t|t-1} - \boldsymbol{\Upsilon}_{t,i-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top - \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top\right) \quad (264)$$

For any function ℓ the chain rule gives

$$\nabla_{\boldsymbol{\Upsilon}_{t|t-1}} \ell = -\text{diag}\left(\left(\boldsymbol{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top\right)^{-1} (\nabla_{\boldsymbol{\Sigma}_{t|t-1}} \ell) \left(\boldsymbol{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top\right)^{-1}\right) \quad (265)$$

$$\nabla_{\mathbf{W}_{t|t-1}} \ell = -2 \left(\boldsymbol{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top\right)^{-1} (\nabla_{\boldsymbol{\Sigma}_{t|t-1}} \ell) \left(\boldsymbol{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top\right)^{-1} \mathbf{W}_{t|t-1} \quad (266)$$

Therefore the gradients we need are

$$\nabla_{\boldsymbol{\mu}_{t,i-1}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | h_t(\boldsymbol{\theta}_t))] = \mathbf{g}_{\boldsymbol{\psi}_{t,i}} \quad (267)$$

$$\nabla_{\boldsymbol{\Upsilon}_{t,i-1}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | h_t(\boldsymbol{\theta}_t))] = -\frac{1}{2} \text{diag} \left(\begin{array}{l} \left(\boldsymbol{\Upsilon}_{t,i-1} + \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top\right)^{-1} \mathbf{G}_{t,i} \\ \times \left(\boldsymbol{\Upsilon}_{t,i-1} + \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top\right)^{-1} \end{array} \right) \quad (268)$$

$$\begin{aligned} \nabla_{\mathbf{W}_{t,i-1}} \mathbb{E}_{\boldsymbol{\theta}_t \sim q_{\boldsymbol{\psi}_{t,i-1}}} [\log p(\mathbf{y}_t | h_t(\boldsymbol{\theta}_t))] &= -\left(\boldsymbol{\Upsilon}_{t,i-1} + \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top\right)^{-1} \mathbf{G}_{t,i} \\ &\times \left(\boldsymbol{\Upsilon}_{t,i-1} + \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top\right)^{-1} \mathbf{W}_{t,i-1} \end{aligned} \quad (269)$$

and

$$\begin{aligned} \nabla_{\Upsilon_{t,i-1}} D_{\mathbb{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) = \\ -\frac{1}{2} \text{diag} \left(\begin{array}{l} (\Upsilon_{t,i-1} + \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top)^{-1} \\ \times (\Upsilon_{t|t-1} - \Upsilon_{t,i-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top - \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top) \\ \times (\Upsilon_{t,i-1} + \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top)^{-1} \end{array} \right) \end{aligned} \quad (270)$$

$$\begin{aligned} \nabla_{\mathbf{W}_{t,i-1}} D_{\mathbb{KL}}(q_{\psi_{t,i-1}} | q_{\psi_{t|t-1}}) = -(\Upsilon_{t,i-1} + \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top)^{-1} \\ \times (\Upsilon_{t|t-1} - \Upsilon_{t,i-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top - \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top) \\ \times (\Upsilon_{t,i-1} + \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top)^{-1} \mathbf{W}_{t,i-1} \end{aligned} \quad (271)$$

The Fisher matrix can be decomposed as a block-diagonal with blocks for $\mu_{t,i}$ and for $(\Upsilon_{t,i}, \mathbf{W}_{t,i})$, but (in contrast to the FC Gaussian case in Eq. (166)) we have not found an efficient way to analytically invert the latter block, which has size $P + RP$. To avoid the $O(R^3 P^3)$ cost of brute force inversion, we use a different strategy for BONG and BLR of performing the update on the natural parameters of the FC Gaussian as in Appendix E.1 and projecting the result back to rank R using SVD. Specifically, assume the updated precision from applying Eq. (119) for BONG or Eq. (134) for BLR can be written as

$$\tilde{\Sigma}_{t,i}^{-1} = \tilde{\Upsilon}_{t,i} + \tilde{\mathbf{W}}_{t,i} \tilde{\mathbf{W}}_{t,i}^\top \quad (272)$$

and let the SVD of $\tilde{\mathbf{W}}_{t,i}$ be

$$\tilde{\mathbf{W}}_{t,i} = \mathbf{U}_{t,i} \mathbf{\Lambda}_{t,i} \mathbf{V}_{t,i}^\top \quad (273)$$

with $\mathbf{U}_{t,i}, \mathbf{V}_{t,i}$ orthogonal and $\mathbf{\Lambda}_{t,i}$ rectangular-diagonal. Following [Mishkin et al., 2018, Chang et al., 2023] we define the update

$$\mathbf{W}_{t,i} = \mathbf{U}_{t,i}[:, :R] \mathbf{\Lambda}_{t,i}[:, R, :] \quad (274)$$

$$\Upsilon_{t,i} = \tilde{\Upsilon}_{t,i} + \text{diag}(\tilde{\mathbf{W}}_{t,i} \tilde{\mathbf{W}}_{t,i}^\top - \mathbf{W}_{t,i} \mathbf{W}_{t,i}^\top) \quad (275)$$

so that $\mathbf{W}_{t,i}$ contains the top R singular vectors and values of the FC posterior and the diagonal is preserved: $\text{diag}(\Upsilon_{t,i} + \mathbf{W}_{t,i} \mathbf{W}_{t,i}^\top) = \text{diag}(\tilde{\Sigma}_{t,i}^{-1})$. This approach works for MC-EF and LIN-EF methods but not MC-HESS, which we omit.

Finally, in the MC-EF methods we sample from the DLR prior using the routine in [Mishkin et al., 2018, Lambert et al., 2023] which takes $O(R(R + M)P)$.

E.5.1 BONG DLR

Substituting the DLR prior from Eq. (262) into the FC precision update from Eq. (119) and using the MC-EF approximation yields the posterior precision

$$\tilde{\Sigma}_t^{-1} = \Upsilon_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top - \mathbf{G}_t^{\text{MC-EF}} \quad (276)$$

$$= \tilde{\Upsilon}_t + \tilde{\mathbf{W}}_t \tilde{\mathbf{W}}_t^\top \quad (277)$$

$$\tilde{\Upsilon}_t = \Upsilon_{t|t-1} \quad (278)$$

$$\tilde{\mathbf{W}}_t = \left[\mathbf{W}_{t|t-1}, \frac{1}{\sqrt{M}} \hat{\mathbf{G}}_t^{(1:M)} \right] \quad (279)$$

Note $\tilde{\mathbf{W}}_t \in \mathbb{R}^{P \times (R+M)}$. Using this posterior precision in the mean update from Eq. (118) yields

$$\mu_t = \mu_{t|t-1} + (\Upsilon_{t|t-1} + \tilde{\mathbf{W}}_t \tilde{\mathbf{W}}_t^\top)^{-1} \mathbf{g}_t \quad (280)$$

$$= \mu_{t|t-1} + \left(\Upsilon_{t|t-1}^{-1} - \Upsilon_{t|t-1}^{-1} \tilde{\mathbf{W}}_t \left(\mathbf{I}_{R+M} + \tilde{\mathbf{W}}_t^\top \Upsilon_{t|t-1}^{-1} \tilde{\mathbf{W}}_t \right)^{-1} \tilde{\mathbf{W}}_t^\top \Upsilon_{t|t-1}^{-1} \right) \mathbf{g}_t \quad (281)$$

where the second line comes from the Woodbury identity and can be computed in $O((R + M)^2 P + (R + M)^3)$. Applying the SVD projection gives

$$\mathbf{W}_t = \mathbf{U}_t[:, :R] \mathbf{\Lambda}_t[:, :R] \quad (282)$$

$$\mathbf{\Upsilon}_t = \mathbf{\Upsilon}_{t|t-1} + \text{diag}(\tilde{\mathbf{W}}_t \tilde{\mathbf{W}}_t^\top - \mathbf{W}_t \mathbf{W}_t^\top) \quad (283)$$

$$(\mathbf{U}_t, \mathbf{\Lambda}_t, _) = \text{SVD}(\tilde{\mathbf{W}}_t) \quad (284)$$

which takes $O((R + M)^2 P)$ for the SVD. Therefore the BONG-MC-EF-DLR update is defined by Eqs. (279) and (281) to (284) and takes $O((R + M)^2 P + (R + M)^3)$.

The BONG-LIN-EF-DLR update comes from replacing $\frac{1}{\sqrt{M}} \hat{\mathbf{G}}_t^{(1:M)}$ with $\mathbf{g}_t^{\text{LIN}}$ in Eq. (279) and replacing \mathbf{I}_{R+M} with \mathbf{I}_{R+1} in Eq. (281). This update takes $O((R + 1)^2 P + (R + 1)^3)$.

E.5.2 BONG-LIN-HESS DLR (LO-FI)

Applying Proposition 4.2 to Eqs. (279) and (281) gives the BONG-LIN-HESS-DLR update:

$$\begin{aligned} \boldsymbol{\mu}_t &= \boldsymbol{\mu}_{t|t-1} + \left(\mathbf{\Upsilon}_{t|t-1}^{-1} - \mathbf{\Upsilon}_{t|t-1}^{-1} \tilde{\mathbf{W}}_t \left(\mathbf{I}_{R+C} + \tilde{\mathbf{W}}_t^\top \mathbf{\Upsilon}_{t|t-1}^{-1} \tilde{\mathbf{W}}_t \right)^{-1} \tilde{\mathbf{W}}_t^\top \mathbf{\Upsilon}_{t|t-1}^{-1} \right) \\ &\quad \times \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \end{aligned} \quad (285)$$

$$\mathbf{W}_t = \mathbf{U}_t[:, :R] \mathbf{\Lambda}_t[:, :R] \quad (286)$$

$$\mathbf{\Upsilon}_t = \mathbf{\Upsilon}_{t|t-1} + \text{diag}(\tilde{\mathbf{W}}_t \tilde{\mathbf{W}}_t^\top - \mathbf{W}_t \mathbf{W}_t^\top) \quad (287)$$

$$(\mathbf{U}_t, \mathbf{\Lambda}_t, _) = \text{SVD}(\tilde{\mathbf{W}}_t) \quad (288)$$

$$\tilde{\mathbf{W}}_t = [\mathbf{W}_{t|t-1}, \mathbf{H}_t^\top \mathbf{A}_t^\top] \quad (289)$$

$$\mathbf{A}_t = \text{chol}(\mathbf{R}_t^{-1}) \quad (290)$$

This is equivalent to LO-FI [Chang et al., 2023] and takes $O((R + C)^2 P + (R + C)^3)$.

E.5.3 BLR DLR (SLANG)

Substituting DLR forms for $q_{\psi_{t|t-1}}$ and $q_{\psi_{t,i-1}}$ into the FC precision update from Eq. (134) and using the MC-EF approximation yields the posterior precision

$$\tilde{\Sigma}_{t,i}^{-1} = (1 - \alpha) (\mathbf{\Upsilon}_{t,i-1} + \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top) + \alpha (\mathbf{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top) - \alpha \mathbf{G}_{t,i}^{\text{MC-EF}} \quad (291)$$

$$= \tilde{\mathbf{\Upsilon}}_{t,i} + \tilde{\mathbf{W}}_{t,i} \tilde{\mathbf{W}}_{t,i}^\top \quad (292)$$

$$\tilde{\mathbf{\Upsilon}}_{t,i} = (1 - \alpha) \mathbf{\Upsilon}_{t,i-1} + \alpha \mathbf{\Upsilon}_{t|t-1} \quad (293)$$

$$\tilde{\mathbf{W}}_{t,i} = \left[\sqrt{1 - \alpha} \mathbf{W}_{t,i-1}, \sqrt{\alpha} \mathbf{W}_{t|t-1}, \sqrt{\frac{\alpha}{M}} \hat{\mathbf{G}}_{t,i}^{(1:M)} \right] \quad (294)$$

Note $\tilde{\mathbf{W}}_t \in \mathbb{R}^{P \times (2R+M)}$. Using this posterior precision in the mean update from Eq. (133) yields

$$\begin{aligned} \boldsymbol{\mu}_{t,i} &= \boldsymbol{\mu}_{t,i-1} + \alpha \left(\tilde{\mathbf{\Upsilon}}_{t,i} + \tilde{\mathbf{W}}_{t,i} \tilde{\mathbf{W}}_{t,i}^\top \right)^{-1} \\ &\quad \times \left((\mathbf{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top) (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \mathbf{g}_{t,i} \right) \end{aligned} \quad (295)$$

$$\begin{aligned} &= \boldsymbol{\mu}_{t,i-1} + \alpha \left(\tilde{\mathbf{\Upsilon}}_{t,i}^{-1} - \tilde{\mathbf{\Upsilon}}_{t,i}^{-1} \tilde{\mathbf{W}}_{t,i} \left(\mathbf{I}_{2R+M} + \tilde{\mathbf{W}}_{t,i}^\top \tilde{\mathbf{\Upsilon}}_{t,i}^{-1} \tilde{\mathbf{W}}_{t,i} \right)^{-1} \tilde{\mathbf{W}}_{t,i}^\top \tilde{\mathbf{\Upsilon}}_{t,i}^{-1} \right) \\ &\quad \times \left((\mathbf{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top) (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \mathbf{g}_{t,i} \right) \end{aligned} \quad (296)$$

where the second line comes from the Woodbury identity and can be computed in $O((2R + M)^2 P + (2R + M)^3)$. Applying the SVD projection gives

$$\mathbf{W}_{t,i} = \mathbf{U}_{t,i}[:, :R] \mathbf{\Lambda}_{t,i}[:, R:] \quad (297)$$

$$\mathbf{\Upsilon}_{t,i} = \tilde{\mathbf{\Upsilon}}_{t,i} + \text{diag}(\tilde{\mathbf{W}}_{t,i} \tilde{\mathbf{W}}_{t,i}^\top - \mathbf{W}_{t,i} \mathbf{W}_{t,i}^\top) \quad (298)$$

$$(\mathbf{U}_{t,i}, \mathbf{\Lambda}_{t,i}, _) = \text{SVD}(\tilde{\mathbf{W}}_{t,i}) \quad (299)$$

which takes $O((2R + M)^2 P)$ for the SVD. Therefore the BLR-MC-EF-DLR update is defined by Eqs. (293), (294) and (296) to (299) and takes $O((2R + M)^2 P + (2R + M)^3)$ per iteration. Notice $\tilde{\mathbf{W}}_t$ has larger rank for BLR than for BONG ($2R + M$ vs. $R + M$) because of the extra $\sqrt{\alpha} \mathbf{W}_{t|t-1}$ term that originates in the KL part of the VI loss. This difference will slow BLR-MC-EF-DLR relative to BONG-MC-EF-DLR especially when R is not small relative to M .

This method closely resembles SLANG [Mishkin et al., 2018] in the batch setting where we replace $q_{\psi_{t|t-1}}$ with a spherical prior $\mathcal{N}(\mathbf{0}, \lambda^{-1} \mathbf{I}_P)$ (see Appendix E.6).

We can also define a BLR-LIN-EF-DLR method by replacing $\sqrt{\frac{\alpha}{M}} \hat{\mathbf{G}}_t^{(1:M)}$ with $\sqrt{\alpha} \mathbf{g}_t^{\text{LIN}}$ in Eq. (294) and \mathbf{I}_{2R+M} with \mathbf{I}_{2R+1} in Eq. (296). This update takes $O((2R + 1)^2 P + (2R + 1)^3)$ per iteration.

E.5.4 BLR-LIN-HESS DLR

Applying Proposition 4.2 to Eqs. (294) and (296) gives the BLR-LIN-HESS-DLR update:

$$\begin{aligned} \boldsymbol{\mu}_{t,i} &= \boldsymbol{\mu}_{t,i-1} + \alpha \left(\tilde{\mathbf{\Upsilon}}_{t,i}^{-1} - \tilde{\mathbf{\Upsilon}}_{t,i}^{-1} \tilde{\mathbf{W}}_{t,i} \left(\mathbf{I}_{2R+C} + \tilde{\mathbf{W}}_{t,i}^\top \tilde{\mathbf{\Upsilon}}_{t,i}^{-1} \tilde{\mathbf{W}}_{t,i} \right)^{-1} \tilde{\mathbf{W}}_{t,i}^\top \tilde{\mathbf{\Upsilon}}_{t,i}^{-1} \right) \\ &\quad \times \left(\left(\mathbf{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top \right) (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \mathbf{H}_{t,i}^\top \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) \right) \end{aligned} \quad (300)$$

$$\mathbf{W}_{t,i} = \mathbf{U}_{t,i}[:, :R] \mathbf{\Lambda}_{t,i}[:, R:] \quad (301)$$

$$\mathbf{\Upsilon}_{t,i} = \tilde{\mathbf{\Upsilon}}_{t,i} + \text{diag}(\tilde{\mathbf{W}}_{t,i} \tilde{\mathbf{W}}_{t,i}^\top - \mathbf{W}_{t,i} \mathbf{W}_{t,i}^\top) \quad (302)$$

$$(\mathbf{U}_{t,i}, \mathbf{\Lambda}_{t,i}, _) = \text{SVD}(\tilde{\mathbf{W}}_{t,i}) \quad (303)$$

$$\tilde{\mathbf{W}}_{t,i} = [\sqrt{1-\alpha} \mathbf{W}_{t,i-1}, \sqrt{\alpha} \mathbf{W}_{t|t-1}, \sqrt{\alpha} \mathbf{H}_{t,i}^\top \mathbf{A}_{t,i}^\top] \quad (304)$$

$$\mathbf{A}_{t,i} = \text{chol}(\mathbf{R}_{t,i}^{-1}) \quad (305)$$

$$\tilde{\mathbf{\Upsilon}}_{t,i} = (1 - \alpha) \mathbf{\Upsilon}_{t,i-1} + \alpha \mathbf{\Upsilon}_{t|t-1} \quad (306)$$

This update takes $O((2R + C)^2 P + (2R + C)^3)$ per iteration. As with the EF versions of BLR-DLR, $\tilde{\mathbf{W}}_t$ has larger rank for BLR-LIN-HESS-DLR than for BONG-LIN-HESS-DLR ($2R + C$ vs. $R + C$). This difference will slow BLR-LIN-HESS-DLR relative to BONG-LIN-HESS-DLR especially when R is not small relative to C .

E.5.5 BOG DLR

Substituting Eqs. (267) to (269) into Eq. (29) gives the BOG-DLR update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \alpha \mathbf{g}_t \quad (307)$$

$$\mathbf{\Upsilon}_t = \mathbf{\Upsilon}_{t|t-1} - \frac{\alpha}{2} \text{diag} \left(\left(\mathbf{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top \right)^{-1} \mathbf{G}_t \left(\mathbf{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top \right)^{-1} \right) \quad (308)$$

$$\mathbf{W}_t = \mathbf{W}_{t|t-1} - \alpha \left(\mathbf{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top \right)^{-1} \mathbf{G}_t \left(\mathbf{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top \right)^{-1} \mathbf{W}_{t|t-1} \quad (309)$$

Using the EF approximation and Woodbury, the BOG-MC-EF-DLR update can be rewritten as

$$\mathbf{\Upsilon}_t = \mathbf{\Upsilon}_{t|t-1} + \frac{\alpha}{2M} \text{diag}(\mathbf{B}_t \mathbf{B}_t^\top) \quad (310)$$

$$\mathbf{W}_t = \mathbf{W}_{t|t-1} + \frac{\alpha}{M} \mathbf{B}_t \mathbf{B}_t^\top \mathbf{W}_{t|t-1} \quad (311)$$

$$\mathbf{B}_t = \left(\mathbf{\Upsilon}_{t|t-1}^{-1} - \mathbf{\Upsilon}_{t|t-1}^{-1} \mathbf{W}_{t|t-1} \left(\mathbf{I}_R + \mathbf{W}_{t|t-1}^\top \mathbf{\Upsilon}_{t|t-1}^{-1} \mathbf{W}_{t|t-1} \right)^{-1} \mathbf{W}_{t|t-1}^\top \mathbf{\Upsilon}_{t|t-1}^{-1} \right) \hat{\mathbf{G}}_t^{(1:M)} \quad (312)$$

which takes $O(RMP + MR^2 + R^3)$.

The BOG-LIN-EF-DLR update comes from replacing $\hat{\mathbf{G}}_t^{(1:M)}$ with $\mathbf{g}_t^{\text{LIN}}$ in Eq. (312) and dropping the M^{-1} factors in Eqs. (310) and (311). This update takes $O(RP + R^3)$.

E.5.6 BOG-LIN-HESS DLR

Applying Proposition 4.2 to Eqs. (307) and (310) to (312) gives the BOG-LIN-HESS-DLR update

$$\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t|t-1} + \alpha \mathbf{H}_t^\top \mathbf{R}_t^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_t) \quad (313)$$

$$\mathbf{\Upsilon}_t = \mathbf{\Upsilon}_{t|t-1} + \frac{\alpha}{2} \text{diag}(\mathbf{B}_t \mathbf{B}_t^\top) \quad (314)$$

$$\mathbf{W}_t = \mathbf{W}_{t|t-1} + \alpha \mathbf{B}_t \mathbf{B}_t^\top \mathbf{W}_{t|t-1} \quad (315)$$

$$\mathbf{B}_t = \left(\mathbf{\Upsilon}_{t|t-1}^{-1} - \mathbf{\Upsilon}_{t|t-1}^{-1} \mathbf{W}_{t|t-1} \left(\mathbf{I}_R + \mathbf{W}_{t|t-1}^\top \mathbf{\Upsilon}_{t|t-1}^{-1} \mathbf{W}_{t|t-1} \right)^{-1} \mathbf{W}_{t|t-1}^\top \mathbf{\Upsilon}_{t|t-1}^{-1} \right) \mathbf{H}_t^\top \mathbf{A}_t^\top \quad (316)$$

This update takes $O(C(C + R)P + CR^2 + R^3)$.

E.5.7 BBB DLR

Substituting Eqs. (263) and (267) to (271) into Eq. (32) gives the BBB-DLR update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \left(\mathbf{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top \right) (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \mathbf{g}_t \quad (317)$$

$$\begin{aligned} \mathbf{\Upsilon}_{t,i} &= \mathbf{\Upsilon}_{t,i-1} \\ &+ \frac{\alpha}{2} \text{diag} \left(\boldsymbol{\Sigma}_{t,i-1} \left(\mathbf{\Upsilon}_{t|t-1} - \mathbf{\Upsilon}_{t,i-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top - \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top - \mathbf{G}_{t,i} \right) \boldsymbol{\Sigma}_{t,i-1} \right) \end{aligned} \quad (318)$$

$$\begin{aligned} \mathbf{W}_{t,i} &= \mathbf{W}_{t,i-1} \\ &+ \alpha \boldsymbol{\Sigma}_{t,i-1} \left(\mathbf{\Upsilon}_{t|t-1} - \mathbf{\Upsilon}_{t,i-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^\top - \mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^\top - \mathbf{G}_{t,i} \right) \boldsymbol{\Sigma}_{t,i-1} \mathbf{W}_{t,i-1} \end{aligned} \quad (319)$$

The previous covariance can be written using Woodbury as

$$\boldsymbol{\Sigma}_{t,i-1} = \mathbf{\Upsilon}_{t,i-1}^{-1} - \mathbf{\Upsilon}_{t,i-1}^{-1} \mathbf{W}_{t,i-1} \left(\mathbf{I}_R + \mathbf{W}_{t,i-1}^\top \mathbf{\Upsilon}_{t,i-1}^{-1} \mathbf{W}_{t,i-1} \right)^{-1} \mathbf{W}_{t,i-1}^\top \mathbf{\Upsilon}_{t,i-1}^{-1} \quad (320)$$

The BBB-MC-EF-DLR update can be computed efficiently by expanding terms in Eqs. (318) to (320). For example the terms involving $\mathbf{G}_{t,i}$ can be calculated as

$$\text{diag}(-\Sigma_{t,i-1} \mathbf{G}_{t,i}^{\text{MC-EF}} \Sigma_{t,i-1}) = \frac{1}{M} \text{diag} \begin{pmatrix} \boldsymbol{\Upsilon}_{t,i-1}^{-1} \hat{\mathbf{G}}_{t,i}^{(1:M)} \hat{\mathbf{G}}_{t,i}^{(1:M)\top} \boldsymbol{\Upsilon}_{t,i-1}^{-1} \\ -\boldsymbol{\Upsilon}_{t,i-1}^{-1} \hat{\mathbf{G}}_{t,i}^{(1:M)} \mathbf{B}_{t,i}^{\top} \\ -\mathbf{B}_{t,i} \hat{\mathbf{G}}_{t,i}^{(1:M)\top} \boldsymbol{\Upsilon}_{t,i-1}^{-1} + \mathbf{B}_{t,i} \mathbf{B}_{t,i}^{\top} \end{pmatrix} \quad (321)$$

$$\begin{aligned} -\Sigma_{t,i-1} \mathbf{G}_{t,i}^{\text{MC-EF}} \Sigma_{t,i-1} \mathbf{W}_{t,i-1} &= \frac{1}{M} \boldsymbol{\Upsilon}_{t,i-1}^{-1} \hat{\mathbf{G}}_{t,i}^{(1:M)} \hat{\mathbf{G}}_{t,i}^{(1:M)\top} \boldsymbol{\Upsilon}_{t,i-1}^{-1} \mathbf{W}_{t,i-1} \\ &\quad - \frac{1}{M} \boldsymbol{\Upsilon}_{t,i-1}^{-1} \hat{\mathbf{G}}_{t,i}^{(1:M)} \mathbf{B}_{t,i}^{\top} \mathbf{W}_{t,i-1} \\ &\quad - \frac{1}{M} \mathbf{B}_{t,i} \hat{\mathbf{G}}_{t,i}^{(1:M)\top} \boldsymbol{\Upsilon}_{t,i-1}^{-1} \mathbf{W}_{t,i-1} + \frac{1}{M} \mathbf{B}_{t,i} \mathbf{B}_{t,i}^{\top} \mathbf{W}_{t,i-1} \end{aligned} \quad (322)$$

$$\begin{aligned} \mathbf{B}_{t,i} &= \boldsymbol{\Upsilon}_{t,i-1}^{-1} \mathbf{W}_{t,i-1} \left(\mathbf{I}_R + \mathbf{W}_{t,i-1}^{\top} \boldsymbol{\Upsilon}_{t,i-1}^{-1} \mathbf{W}_{t,i-1} \right)^{-1} \\ &\quad \times \mathbf{W}_{t,i-1}^{\top} \boldsymbol{\Upsilon}_{t,i-1}^{-1} \hat{\mathbf{G}}_{t,i}^{(1:M)} \end{aligned} \quad (323)$$

Using this strategy the update takes $O((R + M)RP + MR^2 + R^3)$.

The BBB-LIN-EF-DLR update comes from replacing $\hat{\mathbf{G}}_{t,i}^{(1:M)}$ with $\mathbf{g}_t^{\text{LIN}}$ and dropping the M^{-1} factors in Eqs. (321) to (323). This update takes $O(R^2P + R^3)$.

E.5.8 BBB-LIN-HESS DLR

Applying Proposition 4.2 to Eqs. (317) to (319) gives the BBB-LIN-HESS-DLR update

$$\boldsymbol{\mu}_{t,i} = \boldsymbol{\mu}_{t,i-1} + \alpha \left(\boldsymbol{\Upsilon}_{t|t-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^{\top} \right) (\boldsymbol{\mu}_{t|t-1} - \boldsymbol{\mu}_{t,i-1}) + \alpha \mathbf{H}_{t,i}^{\top} \mathbf{R}_{t,i}^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_{t,i}) \quad (324)$$

$$\boldsymbol{\Upsilon}_{t,i} = \boldsymbol{\Upsilon}_{t,i-1} + \frac{\alpha}{2} \text{diag} \left(\Sigma_{t,i-1} \begin{pmatrix} \boldsymbol{\Upsilon}_{t|t-1} - \boldsymbol{\Upsilon}_{t,i-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^{\top} \\ -\mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^{\top} + \mathbf{H}_{t,i}^{\top} \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i} \end{pmatrix} \Sigma_{t,i-1} \right) \quad (325)$$

$$\mathbf{W}_{t,i} = \mathbf{W}_{t,i-1} + \alpha \Sigma_{t,i-1} \begin{pmatrix} \boldsymbol{\Upsilon}_{t|t-1} - \boldsymbol{\Upsilon}_{t,i-1} + \mathbf{W}_{t|t-1} \mathbf{W}_{t|t-1}^{\top} \\ -\mathbf{W}_{t,i-1} \mathbf{W}_{t,i-1}^{\top} + \mathbf{H}_{t,i}^{\top} \mathbf{R}_{t,i}^{-1} \mathbf{H}_{t,i} \end{pmatrix} \Sigma_{t,i-1} \mathbf{W}_{t,i-1} \quad (326)$$

This can be computed in $O((C + R)^2P + CR^2 + R^3)$ using Eq. (320) and following a computational approach similar to Eqs. (321) to (323).

E.6 Batch BLR

It is interesting to translate the BLR updates derived here back to the batch setting where BLR was developed [Khan and Rue, 2023], by replacing $\mathcal{N}(\boldsymbol{\mu}_{t|t-1}, \Sigma_{t|t-1})$ with a centered spherical prior $\mathcal{N}(\mathbf{0}, \lambda^{-1} \mathbf{I}_P)$.

The batch BLR-FC update becomes

$$\boldsymbol{\mu}_i = \boldsymbol{\mu}_{i-1} + \alpha \Sigma_i (\mathbf{g}_i - \lambda \boldsymbol{\mu}_{i-1}) \quad (327)$$

$$\Sigma_i^{-1} = (1 - \alpha) \Sigma_{i-1}^{-1} + \alpha (\lambda \mathbf{I}_P - \mathbf{G}_i) \quad (328)$$

The batch BLR-LIN-HESS-FC update becomes

$$\boldsymbol{\mu}_i = \boldsymbol{\mu}_{i-1} + \alpha \Sigma_i (\mathbf{H}_i^{\top} \mathbf{R}_i^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_i) - \lambda \boldsymbol{\mu}_{i-1}) \quad (329)$$

$$\Sigma_i^{-1} = (1 - \alpha) \Sigma_{i-1}^{-1} + \alpha (\lambda \mathbf{I}_P + \mathbf{H}_i^{\top} \mathbf{R}_i^{-1} \mathbf{H}_i) \quad (330)$$

The batch BLR-FC_MOM update becomes

$$\boldsymbol{\mu}_i = \boldsymbol{\mu}_{i-1} + \alpha \Sigma_{i-1} (\mathbf{g}_i - \lambda \boldsymbol{\mu}_{i-1}) \quad (331)$$

$$\Sigma_i = (1 + \alpha) \Sigma_{i-1} + \alpha \Sigma_{i-1} (\mathbf{G}_i - \lambda \mathbf{I}_P) \Sigma_{i-1} \quad (332)$$

The batch BLR-LIN-HESS-FC_MOM update becomes

$$\boldsymbol{\mu}_i = \boldsymbol{\mu}_{i-1} + \alpha \boldsymbol{\Sigma}_{i-1} (\mathbf{H}_i^\top \mathbf{R}_i^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_i) - \lambda \boldsymbol{\mu}_{i-1}) \quad (333)$$

$$\boldsymbol{\Sigma}_i = (1 + \alpha) \boldsymbol{\Sigma}_{i-1} - \alpha \boldsymbol{\Sigma}_{i-1} (\lambda \mathbf{I}_P + \mathbf{H}_i^\top \mathbf{R}_i^{-1} \mathbf{H}_i) \boldsymbol{\Sigma}_{i-1} \quad (334)$$

The batch BLR-DIAG update becomes

$$\boldsymbol{\mu}_i = \boldsymbol{\mu}_{i-1} + \alpha \boldsymbol{\sigma}_i^2 (\mathbf{g}_i - \lambda \boldsymbol{\mu}_{i-1}) \quad (335)$$

$$\boldsymbol{\sigma}_i^{-2} = (1 - \alpha) \boldsymbol{\sigma}_{i-1}^{-2} + \alpha \text{diag}(\lambda \mathbf{I}_P - \mathbf{G}_i) \quad (336)$$

The MC-HESS version of this update is equivalent to VON [Khan et al., 2018b] if we use MC approximation with $M = 1$. The MC-EF version with $M = 1$ is equivalent to VOGN [Khan et al., 2018b].

The batch BLR-LIN-HESS-DIAG update becomes

$$\boldsymbol{\mu}_i = \boldsymbol{\mu}_{i-1} + \alpha \boldsymbol{\sigma}_i^2 (\mathbf{H}_i^\top \mathbf{R}_i^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_i) - \lambda \boldsymbol{\mu}_{i-1}) \quad (337)$$

$$\boldsymbol{\sigma}_i^{-2} = (1 - \alpha) \boldsymbol{\sigma}_{i-1}^{-2} + \alpha \text{diag}(\lambda \mathbf{I}_P + \mathbf{H}_i^\top \mathbf{R}_i^{-1} \mathbf{H}_i) \quad (338)$$

The batch BLR-DIAG_MOM update becomes

$$\boldsymbol{\mu}_i = \boldsymbol{\mu}_{i-1} + \alpha \boldsymbol{\sigma}_{i-1}^2 (\mathbf{g}_i - \lambda \boldsymbol{\mu}_{i-1}) \quad (339)$$

$$\boldsymbol{\sigma}_i^2 = (1 + \alpha) \boldsymbol{\sigma}_{i-1}^2 + \alpha \boldsymbol{\sigma}_{i-1}^4 \text{diag}(\mathbf{G}_i - \lambda \mathbf{I}_P) \quad (340)$$

The batch BLR-LIN-HESS-DIAG_MOM update becomes

$$\boldsymbol{\mu}_i = \boldsymbol{\mu}_{i-1} + \alpha \boldsymbol{\sigma}_{i-1}^2 (\mathbf{H}_i^\top \mathbf{R}_i^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_i) - \lambda \boldsymbol{\mu}_{i-1}) \quad (341)$$

$$\boldsymbol{\sigma}_i^2 = (1 + \alpha) \boldsymbol{\sigma}_{i-1}^2 - \alpha \boldsymbol{\sigma}_{i-1}^4 \text{diag}(\lambda \mathbf{I}_P + \mathbf{H}_i^\top \mathbf{R}_i^{-1} \mathbf{H}_i) \quad (342)$$

The batch BLR-DLR update becomes

$$\boldsymbol{\mu}_i = \boldsymbol{\mu}_{i-1} + \alpha \left(\tilde{\boldsymbol{\Upsilon}}_i^{-1} - \tilde{\boldsymbol{\Upsilon}}_i^{-1} \tilde{\mathbf{W}}_i \left(\mathbf{I}_{R+M} + \tilde{\mathbf{W}}_i^\top \tilde{\boldsymbol{\Upsilon}}_i^{-1} \tilde{\mathbf{W}}_i \right)^{-1} \tilde{\mathbf{W}}_i^\top \tilde{\boldsymbol{\Upsilon}}_i^{-1} \right) \times (\mathbf{g}_t - \lambda \boldsymbol{\mu}_{i-1}) \quad (343)$$

$$\mathbf{W}_i = \mathbf{U}_i[:, :R] \boldsymbol{\Lambda}_i[:, R, :] \quad (344)$$

$$\boldsymbol{\Upsilon}_i = \tilde{\boldsymbol{\Upsilon}}_i + \text{diag}(\tilde{\mathbf{W}}_i \tilde{\mathbf{W}}_i^\top - \mathbf{W}_i \mathbf{W}_i^\top) \quad (345)$$

$$(\mathbf{U}_i, \boldsymbol{\Lambda}_i, _) = \text{SVD}(\tilde{\mathbf{W}}_i) \quad (346)$$

$$\tilde{\mathbf{W}}_i = \left[\sqrt{1 - \alpha} \mathbf{W}_{i-1}, \sqrt{\frac{\alpha}{M}} \hat{\mathbf{G}}_i^{(1:M)} \right] \quad (347)$$

$$\mathbf{A}_i = \text{chol}(\mathbf{R}_i^{-1}) \quad (348)$$

$$\tilde{\boldsymbol{\Upsilon}}_i = (1 - \alpha) \boldsymbol{\Upsilon}_{i-1} + \alpha \lambda \mathbf{I}_P \quad (349)$$

This is equivalent to SLANG except for the following differences. SLANG processes a minibatch of M examples at each iteration, using a single sample $\boldsymbol{\theta} \sim q_{\psi_{i-1}}$ for each minibatch. It uses a different SVD routine which is slightly faster but stochastic, taken from [Halko et al., 2011]. Most significantly, SLANG applies the SVD before the mean update, meaning $\boldsymbol{\mu}_i$ is calculated using the rank- R \mathbf{W}_i and $\boldsymbol{\Upsilon}_i$ instead of the rank- $(R + M)$ $\tilde{\mathbf{W}}_i$ and $\tilde{\boldsymbol{\Upsilon}}_i$, thus ignoring the non-diagonal information in the M discarded singular vectors from $\tilde{\mathbf{W}}_i$.

The batch BLR-LIN-HESS-DLR update becomes

$$\begin{aligned}\boldsymbol{\mu}_i &= \boldsymbol{\mu}_{i-1} + \alpha \left(\tilde{\boldsymbol{\Upsilon}}_i^{-1} - \tilde{\boldsymbol{\Upsilon}}_i^{-1} \tilde{\mathbf{W}}_i \left(\mathbf{I}_{R+C} + \tilde{\mathbf{W}}_i^\top \tilde{\boldsymbol{\Upsilon}}_i^{-1} \tilde{\mathbf{W}}_i \right)^{-1} \tilde{\mathbf{W}}_i^\top \tilde{\boldsymbol{\Upsilon}}_i^{-1} \right) \\ &\quad \times \left(\mathbf{H}_i^\top \mathbf{R}_i^{-1} (\mathbf{y}_t - \hat{\mathbf{y}}_i) - \lambda \boldsymbol{\mu}_{i-1} \right)\end{aligned}\quad (350)$$

$$\mathbf{W}_i = \mathbf{U}_i[:, :R] \boldsymbol{\Lambda}_i[:, R, :R] \quad (351)$$

$$\boldsymbol{\Upsilon}_i = \tilde{\boldsymbol{\Upsilon}}_i + \text{diag} \left(\tilde{\mathbf{W}}_i \tilde{\mathbf{W}}_i^\top - \mathbf{W}_i \mathbf{W}_i^\top \right) \quad (352)$$

$$(\mathbf{U}_i, \boldsymbol{\Lambda}_i, _) = \text{SVD} \left(\tilde{\mathbf{W}}_i \right) \quad (353)$$

$$\tilde{\mathbf{W}}_i = [\sqrt{1-\alpha} \mathbf{W}_{i-1}, \sqrt{\alpha} \mathbf{H}_i^\top \mathbf{A}_i^\top] \quad (354)$$

$$\mathbf{A}_i = \text{chol} (\mathbf{R}_i^{-1}) \quad (355)$$

$$\tilde{\boldsymbol{\Upsilon}}_i = (1-\alpha) \boldsymbol{\Upsilon}_{i-1} + \alpha \lambda \mathbf{I}_P \quad (356)$$

This algorithm could be called SLANG-LIN-HESS and would be deterministic and faster than SLANG since it does not need MC sampling. We can also define SLANG-LIN-EF which would be even faster, by replacing $\sqrt{\frac{\alpha}{M}} \hat{\mathbf{G}}_i^{(1:M)}$ with $\sqrt{\alpha} \mathbf{g}_i^{\text{LIN}}$ in Eq. (347) and \mathbf{I}_{R+M} with \mathbf{I}_R in Eq. (343).