

Probabilistic Line Searches for Stochastic Optimization

Maren Mahsereci

MMAHSERECI@TUE.MPG.DE

Philipp Hennig

PHENNIG@TUE.MPG.DE

*Max Planck Institute for Intelligent Systems
Max-Planck-Ring 4, 72076 Tübingen, Germany*

Editor: Mark Schmidt

Abstract

In deterministic optimization, line searches are a standard tool ensuring stability and efficiency. Where only stochastic gradients are available, no direct equivalent has so far been formulated, because uncertain gradients do not allow for a strict sequence of decisions collapsing the search space. We construct a probabilistic line search by combining the structure of existing deterministic methods with notions from Bayesian optimization. Our method retains a Gaussian process surrogate of the univariate optimization objective, and uses a probabilistic belief over the Wolfe conditions to monitor the descent. The algorithm has very low computational cost, and no user-controlled parameters. Experiments show that it effectively removes the need to define a learning rate for stochastic gradient descent.

Keywords: stochastic optimization, learning rates, line searches, Gaussian processes, Bayesian optimization

1. Introduction

This work substantially extends the work of Mahsereci and Hennig (2015) published at NIPS 2015. Stochastic gradient descent (SGD, Robbins and Monro, 1951) is currently the standard in machine learning for the optimization of highly multivariate functions if their gradient is corrupted by noise. This includes the online or mini-batch training of neural networks, logistic regression (Zhang, 2004; Bottou, 2010) and variational models (e.g. Hoffman et al., 2013; Hensman et al., 2012; Broderick et al., 2013). In all these cases, noisy gradients arise because an exchangeable loss-function $\mathcal{L}(x)$ of the optimization parameters $x \in \mathbb{R}^D$, across a large dataset $\{d_i\}_{i=1, \dots, M}$, is evaluated only on a subset $\{d_j\}_{j=1, \dots, m}$:

$$\mathcal{L}(x) := \frac{1}{M} \sum_{i=1}^M \ell(x, d_i) \approx \frac{1}{m} \sum_{j=1}^m \ell(x, d_j) =: \hat{\mathcal{L}}(x) \quad m \ll M. \quad (1)$$

If the indices j are i.i.d. draws from $[1, M]$, by the Central Limit Theorem, the error $\hat{\mathcal{L}}(x) - \mathcal{L}(x)$ is unbiased and approximately normal distributed. Despite its popularity and its low cost per step, SGD has well-known deficiencies that can make it inefficient, or at least tedious to use in practice. Two main issues are that, first, the gradient itself, even without noise, is not the optimal search *direction*; and second, SGD requires a *step size* (learning rate) that has drastic effect on the algorithm's efficiency, is often difficult to choose well, and virtually never optimal for each individual descent step. The former issue, adapting the search direction, has been addressed by many authors (see George and Powell, 2006, for

an overview). Existing approaches range from lightweight ‘diagonal preconditioning’ like ADAM (Kingma and Ba, 2014), ADAGRAD (Duchi et al., 2011), and ‘stochastic meta-descent’ (Schraudolph, 1999), to empirical estimates for the natural gradient (Amari et al., 2000) or the Newton direction (Roux and Fitzgibbon, 2010), to problem-specific algorithms (Rajesh et al., 2013), and more elaborate estimates of the Newton direction (Hennig, 2013). Most of these algorithms also include an auxiliary adaptive effect on the learning rate. Schaul et al. (2013) provided an estimation method to explicitly adapt the learning rate from one gradient descent step to another. Several very recent works have proposed the use of reinforcement learning and ‘learning-to-learn’ approaches for parameter adaption (Andrychowicz et al., 2016; Hansen, 2016; Li and Malik, 2016). Mostly these methods are designed to work well on a specified subset of optimization problems, which they are also trained on; they thus need to be re-learned for differing objectives. The corresponding algorithms are usually orders of magnitude more expensive than the low-level black box proposed here, and often require a classic optimizer (e.g. SGD) to tune their internal hyper-parameters.

None of the mentioned algorithms change the size of the *current* descent step. Accumulating statistics across steps in this fashion requires some conservatism: If the step size is initially too large, or grows too fast, SGD can become unstable and ‘explode’, because individual steps are not checked for robustness at the time they are taken.

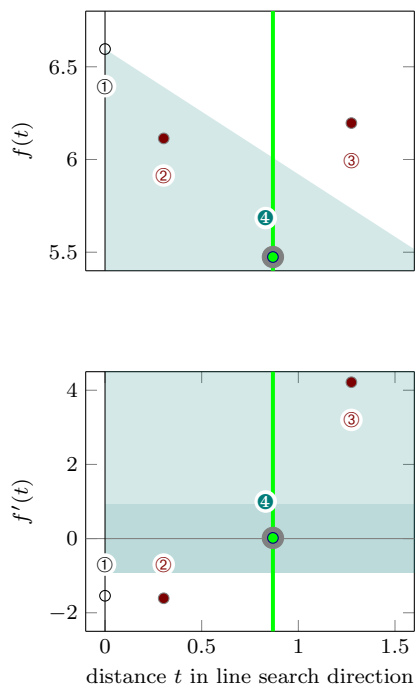


Figure 1: Sketch: The task of a *classic line search* is to tune the step taken by an optimization algorithm along a univariate search direction. The search starts at the endpoint ① of the previous line search, at $t = 0$. The *upper* plot shows function values, the *lower* plot corresponding gradients. A sequence of extrapolation steps ②,③ finds a point of positive gradient at ③. It is followed by interpolation steps until an acceptable point ④ is found. Points of insufficient decrease, above the line $f(0) + c_1 t f'(0)$ (white area in upper plot) are excluded by the Armijo condition W-I, while points of steep negative gradient (white area in lower plot) are excluded by the curvature condition W-II (the strong extension of the Wolfe conditions also excludes the light green area in the lower plot). Point ④ is the first to fulfill both conditions, and is thus accepted.

In essence, the same problem exists in deterministic (noise-free) optimization problems. There, providing stability is one of several tasks of the *line search* subroutine. It is a standard constituent of algorithms like the classic nonlinear conjugate gradient (Fletcher and Reeves, 1964) and BFGS (Broyden, 1969; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970)

methods (Nocedal and Wright, 1999, §3).¹ In the noise-free case, line searches are considered a solved problem (Nocedal and Wright, 1999, §3). But the methods used in deterministic optimization are not stable to noise. They are easily fooled by even small disturbances, either becoming overly conservative or failing altogether. The reason for this brittleness is that existing line searches take a sequence of hard decisions to shrink or shift the search space. This yields efficiency, but breaks hard in the presence of noise. Section 3 constructs a probabilistic line search for noisy objectives, stabilizing optimization methods like the works cited above. As line searches only change the length, not the direction of a step, they could be used in combination with the algorithms adapting SGD’s direction, cited above. In this paper we focus on parameter tuning of the SGD algorithm and leave other search directions to future work.

2. Connections

2.1 Deterministic Line Searches

There is a host of existing line search variants (Nocedal and Wright, 1999, §3). In essence, though, these methods explore a univariate domain ‘to the right’ of a starting point, until an ‘acceptable’ point is reached (Figure 1). More precisely, consider the problem of minimizing $\mathcal{L}(x) : \mathbb{R}^D \rightarrow \mathbb{R}$, with access to $\nabla \mathcal{L}(x) : \mathbb{R}^D \rightarrow \mathbb{R}^D$. At iteration i , some ‘outer loop’ chooses, at location x_i , a search direction $s_i \in \mathbb{R}^D$ (e.g. by the BFGS-rule, or simply $s_i = -\nabla \mathcal{L}(x_i)$ for gradient descent). It will *not* be assumed that s_i has unit norm. The line search operates along the univariate domain $x(t) = x_i + ts_i$ for $t \in \mathbb{R}_+$. Along this direction it collects scalar function values and projected gradients that will be denoted $f(t) = \mathcal{L}(x(t))$ and $f'(t) = s_i^\top \nabla \mathcal{L}(x(t)) \in \mathbb{R}$. Most line searches involve an initial extrapolation phase to find a point t_r with $f'(t_r) > 0$. This is followed by a search in $[0, t_r]$, by interval nesting or by interpolation of the collected function and gradient values, e.g. with cubic splines.²

2.1.1 THE WOLFE CONDITIONS FOR TERMINATION

As the line search is only an auxiliary step within a larger iteration, it need not find an exact root of f' ; it suffices to find a point ‘sufficiently’ close to a minimum. The *Wolfe conditions* (Wolfe, 1969) are a widely accepted formalization of this notion; they consider t acceptable if it fulfills

$$f(t) \leq f(0) + c_1 t f'(0) \quad (\text{W-I}) \quad \text{and} \quad f'(t) \geq c_2 f'(0) \quad (\text{W-II}), \quad (2)$$

using two constants $0 \leq c_1 < c_2 \leq 1$ chosen by the designer of the line search, not the user. W-I is the *Armijo* or *sufficient decrease* condition (Armijo, 1966). It encodes that acceptable functions values should lie below a linear extrapolation line of slope $c_1 f'(0)$. W-II is the *curvature condition*, demanding a decrease in slope. The choice $c_1 = 0$ accepts any value below $f(0)$, while $c_1 = 1$ rejects all points for convex functions. For the curvature condition,

1. In these algorithms, another task of the line search is to guarantee certain properties of the surrounding estimation rule. In BFGS, e.g., it ensures positive definiteness of the estimate. This aspect will not feature here.

2. This is the strategy in `minimize.m` by C. Rasmussen, which provided a model for our implementation. At the time of writing, it can be found at <http://learning.eng.cam.ac.uk/car1/code/minimize/minimize.m>

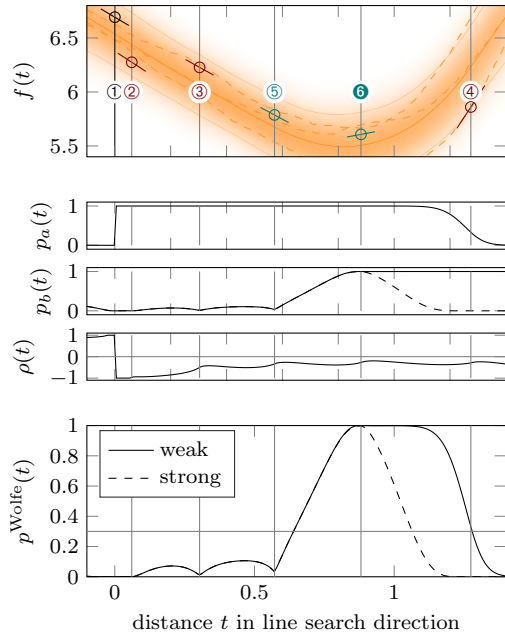


Figure 2: Sketch of a *probabilistic line search*. As in Fig. 1, the algorithm performs extrapolation (②,③,④) and interpolation (⑤,⑥), but receives unreliable, noisy function and gradient values. These are used to construct a GP posterior (top: solid posterior mean, thin lines at 2 standard deviations, local pdf marginal as shading, three dashed sample paths). This implies a bivariate Gaussian belief (§3.3) over the validity of the weak Wolfe conditions (middle three plots. $p_a(t)$ is the marginal for W-I, $p_b(t)$ for W-II, $\rho(t)$ their correlation). Points are considered acceptable if their joint probability $p^{\text{Wolfe}}(t)$ (bottom) is above a threshold (gray). An approximation (§3.3.1) to the strong Wolfe conditions is shown dashed.

$c_2 = 0$ only accepts points with $f'(t) \geq 0$; while $c_2 = 1$ accepts any point of greater slope than $f'(0)$. W-I and W-II are known as the *weak* form of the Wolfe conditions. The *strong* form replaces W-II with $|f'(t)| \leq c_2|f'(0)|$. This guards against accepting points of low function value but large positive gradient. Figure 1 shows a conceptual sketch illustrating the typical process of a line search, and the weak and strong Wolfe conditions. The exposition in §3.3 will initially focus on the weak conditions, which can be precisely modeled probabilistically. Section 3.3.1 then adds an approximate treatment of the strong form.

2.2 Bayesian Optimization

A recently blossoming sample-efficient approach to global optimization revolves around modeling the objective f with a probability measure $p(f)$; usually a Gaussian process (GP). Searching for extrema, evaluation points are then chosen by a utility functional $u[p(f)]$. Our line search borrows the idea of a Gaussian process surrogate, and a popular acquisition function, *expected improvement* (Jones et al., 1998). Bayesian optimization (BO) methods are often computationally expensive, thus ill-suited for a cost-sensitive task like a line search. But since line searches are governors more than information extractors, the kind of sample-efficiency expected of a Bayesian optimizer is not needed. The following sections develop a lightweight algorithm which adds only minor computational overhead to stochastic optimization.

3. A Probabilistic Line Search

We now consider minimizing $f(t) = \hat{\mathcal{L}}(x(t))$ from Eq. 1. That is, the algorithm can access only noisy function values and gradients y_t, y'_t at location t , with Gaussian likelihood

$$p(y_t, y'_t | f) = \mathcal{N} \left(\begin{bmatrix} y_t \\ y'_t \end{bmatrix}; \begin{bmatrix} f(t) \\ f'(t) \end{bmatrix}, \begin{bmatrix} \sigma_f^2 & 0 \\ 0 & \sigma_{f'}^2 \end{bmatrix} \right). \quad (3)$$

The Gaussian form is supported by the Central Limit argument at Eq. 1. The function value y_t and the gradient y'_t are assumed independent for simplicity; see §3.4 and Appendix A regarding estimation of the variances $\sigma_f^2, \sigma_{f'}^2$, and some further notes on the independence assumption of y and y' . Each evaluation of $f(t)$ uses a newly drawn mini-batch.

Our algorithm is modeled after the classic line search routine `minimize.m`² and translates each of its building blocks one-by-one to the language of probability. The following table illustrates these four ingredients of the probabilistic line search and their corresponding classic parts.

building block	classic	probabilistic
1) 1D surrogate for objective $f(t)$	piecewise cubic splines	GP where the mean are piecewise cubic splines
2) candidate selection	<i>one</i> local minimizer of cubic splines <i>xor</i> extrapolation	local minimizers of cubic splines <i>and</i> extrapolation
3) choice of best candidate	—————	BO acquisition function
4) acceptance criterion	classic Wolfe conditions	prob. Wolfe conditions

The table already motivates certain design choices, for example the particular choice of the GP-surrogate for $f(t)$, which strongly resembles the classic design. Probabilistic line searches operate in the same scheme as classic ones: 1) they construct a surrogate for the underlying 1D-function 2) they select candidates for evaluation which can interpolate between datapoints or extrapolate 3) a heuristic chooses among the candidate locations and the function is evaluated there 4) the evaluated points are checked for Wolfe-acceptance. The following sections introduce all of these building blocks with greater detail: A robust yet lightweight Gaussian process surrogate on $f(t)$ facilitating analytic optimization (§ 3.1); a simple Bayesian optimization objective for exploration (§ 3.2); and a probabilistic formulation of the Wolfe conditions as a termination criterion (§ 3.3). Appendix D contains a detailed pseudocode of the probabilistic line search; Algorithm 1 very roughly sketches the structure of the probabilistic line search and highlights its essential ingredients.

3.1 Lightweight Gaussian Process Surrogate

We model information about the objective in a probability measure $p(f)$. There are two requirements on such a measure: First, it must be robust to irregularity (low and high variability) of the objective. And second, it must allow analytic computation of discrete candidate points for evaluation, because a line search should not call yet another optimization subroutine itself. Both requirements are fulfilled by a once-integrated Wiener process, i.e. a zero-mean Gaussian process prior $p(f) = \mathcal{GP}(f; 0, k)$ with covariance function

$$k(t, t') = \theta^2 \left[1/3 \min^3(\tilde{t}, \tilde{t}') + 1/2 |t - t'| \min^2(\tilde{t}, \tilde{t}') \right]. \quad (4)$$

Algorithm 1 PROBLINESEARCHSKETCH($f, y_0, y'_0, \sigma_{f_0}, \sigma_{f'_0}$)

$GP \leftarrow \text{INITGP}(y_0, y'_0, \sigma_{f_0}, \sigma_{f'_0})$
 $T, Y, Y' \leftarrow \text{INITSTORAGE}(0, y_0, y'_0)$ ▷ for observed points
 $t \leftarrow 1$ ▷ scaled position of initial candidate

while budget not used and no Wolfe-point found do

$[y, y'] \leftarrow f(t)$ ▷ evaluate objective
 $T, Y, Y' \leftarrow \text{UPDATESTORAGE}(t, y, y')$
 $GP \leftarrow \text{UPDATEGP}(t, y, y')$
 $P^{\text{Wolfe}} \leftarrow \text{PROBWOLFE}(T, GP)$ ▷ compute Wolfe probability at points in T

if any P^{Wolfe} above Wolfe threshold c_W then

return Wolfe-point

else

$T_{\text{cand}} \leftarrow \text{COMPUTECANDIDATES}(GP)$ ▷ positions of new candidates
 $EI \leftarrow \text{EXPECTEDIMPROVEMENT}(T_{\text{cand}}, GP)$
 $PW \leftarrow \text{PROBWOLFE}(T_{\text{cand}}, GP)$
 $t \leftarrow$ where $(PW \odot EI)$ is maximal ▷ find best candidate among T_{cand}

end if

end while

return observed point in T with lowest GP mean since no Wolfe-point found

Here $\tilde{t} := t + \tau$ and $\tilde{t}' := t' + \tau$ denote a shift by a constant $\tau > 0$. This ensures this kernel is positive semi-definite, the precise value τ is irrelevant as the algorithm only considers positive values of t (our implementation uses $\tau = 10$). See §3.4 regarding the scale θ^2 . With the likelihood of Eq. 3, this prior gives rise to a GP posterior whose mean function is a cubic spline³ (Wahba, 1990). We note in passing that regression on f and f' from N observations of pairs (y_t, y'_t) can be formulated as a filter (Särkkä, 2013) and thus performed in $\mathcal{O}(N)$ time. However, since a line search typically collects < 10 data points, generic GP inference, using a Gram matrix, has virtually the same, low cost.

Because Gaussian measures are closed under linear maps (Papoulis, 1991, §10), Eq. 4 implies a Wiener process (linear spline) model on f' :

$$p(f; f') = \mathcal{GP} \left(\begin{bmatrix} f \\ f' \end{bmatrix}; \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} k & k^\partial \\ \partial_k & \partial_k^\partial \end{bmatrix} \right), \quad (5)$$

3. Eq. 4 can be generalized to the ‘natural spline’, removing the need for the constant τ (Rasmussen and Williams, 2006, §6.3.1). However, this notion is ill-defined in the case of a single observation, as in the line search.




Figure 3: *Integrated Wiener process*: GP marginal posterior of function values; posterior mean in solid orange and, two standard deviations in thinner solid orange, local pdf marginal as shading; function value observations as gray circles (corresponding gradients not shown). Classic interpolation by piecewise cubic spline in dark blue. *Left*: observations are exact; the mean of the GP and the cubic spline interpolator of a classic line search coincide. *Right*: same observations with additive Gaussian noise (error-bars indicate ± 1 standard deviations); noise free interpolator in dashed gray for comparison. The classic interpolator in dark blue, which exactly matches the observations, becomes unreliable; the GP reacts robustly to noisy observations; the GP-mean still consists of piecewise cubic splines.

with (using the indicator function $\mathbb{I}(x) = 1$ if x , else 0)

$$\begin{aligned}
 k_{tt'}^{\partial} &:= \frac{\partial k(t, t')}{\partial t'} = \theta^2 \left[\mathbb{I}(t < t') \frac{\tilde{t}'^2}{2} + \mathbb{I}(t \geq t') \left(\tilde{t}\tilde{t}' - \frac{\tilde{t}'^2}{2} \right) \right] \\
 \partial_{k_{tt'}} &:= \frac{\partial k(t, t')}{\partial t} = \theta^2 \left[\mathbb{I}(t' < t) \frac{\tilde{t}^2}{2} + \mathbb{I}(t' \geq t) \left(\tilde{t}\tilde{t}' - \frac{\tilde{t}^2}{2} \right) \right] \\
 \partial_k^{\partial}{}_{tt'} &:= \frac{\partial^2 k(t, t')}{\partial t' \partial t} = \theta^2 \min(\tilde{t}, \tilde{t}').
 \end{aligned} \tag{6}$$

Given a set of evaluations $(\mathbf{t}, \mathbf{y}, \mathbf{y}')$ (vectors, with elements t_i, y_{t_i}, y'_{t_i}) with independent likelihood 3, the posterior $p(f | \mathbf{y}, \mathbf{y}')$ is a GP with posterior mean function μ and covariance function \tilde{k} as follows:

$$\begin{aligned}
 \begin{bmatrix} \mu(t) \\ \mu'(t) \end{bmatrix} &= \underbrace{\begin{bmatrix} k_{tt} & k_{tt}^{\partial} \\ \partial_{k_{tt}} & \partial_{k^{\partial}{}_{tt}} \end{bmatrix} \begin{bmatrix} k_{tt} + \sigma_f^2 \mathbf{I} & k_{tt}^{\partial} \\ \partial_{k_{tt}} & \partial_{k^{\partial}{}_{tt}} + \sigma_{f'}^2 \mathbf{I} \end{bmatrix}^{-1}}_{=: \mathbf{g}^{\top}(t)} \begin{bmatrix} \mathbf{y} \\ \mathbf{y}' \end{bmatrix} \\
 \begin{bmatrix} \tilde{k}(t, t') & \tilde{k}^{\partial}(t, t') \\ \tilde{\partial}_{\tilde{k}}(t', t) & \tilde{\partial}_{\tilde{k}^{\partial}}(t, t') \end{bmatrix} &= \begin{bmatrix} k_{tt'} & k_{tt'}^{\partial} \\ \partial_{k_{t't}} & \partial_{k^{\partial}{}_{t't}} \end{bmatrix} - \mathbf{g}^{\top}(t) \begin{bmatrix} k_{tt'} & k_{tt'}^{\partial} \\ \partial_{k_{t't}} & \partial_{k^{\partial}{}_{t't}} \end{bmatrix}
 \end{aligned} \tag{7}$$

The posterior marginal variance will be denoted by $\mathbb{V}(t) = \tilde{k}(t, t)$. To see that μ is indeed piecewise cubic (i.e. a cubic spline), we note that it has at most three non-vanishing

derivatives⁴, because

$$\begin{aligned} \partial^2 k_{tt'} &:= \frac{\partial^2 k(t, t')}{\partial t^2} = \theta^2 \mathbb{I}(t \leq t') & \partial^3 k_{tt'} &:= \frac{\partial^3 k(t, t')}{\partial t^3} = \theta^2 \mathbb{I}(t \leq t')(t' - t) \\ \partial^2 k_{t't} &:= \frac{\partial^2 k(t, t')}{\partial t'^2} = -\theta^2 \mathbb{I}(t \leq t') & \partial^3 k_{t't} &:= \frac{\partial^3 k(t, t')}{\partial t'^3} = 0. \end{aligned} \quad (8)$$

This piecewise cubic form of μ is crucial for our purposes: having collected N values of f and f' , respectively, all local minima of μ can be found analytically in $\mathcal{O}(N)$ time in a single sweep through the ‘cells’ $t_{i-1} < t < t_i$, $i = 1, \dots, N$ (here $t_0 = 0$ denotes the start location, where (y_0, y'_0) are ‘inherited’ from the preceding line search. For typical line searches $N < 10$, c.f. §4. In each cell, $\mu(t)$ is a cubic polynomial with at most one minimum in the cell, found by an inexpensive quadratic computation from the three scalars $\mu'(t_i), \mu''(t_i), \mu'''(t_i)$. This is in contrast to other GP regression models—for example the one arising from a squared exponential kernel—which give more involved posterior means whose local minima can be found only approximately. Another advantage of the cubic spline interpolant is that it does not assume the existence of higher derivatives (in contrast to the Gaussian kernel, for example), and thus reacts robustly to irregularities in the objective. In our algorithm, after each evaluation of (y_N, y'_N) , we use this property to compute a short list of *candidates* for the next evaluation, consisting of the $\leq N$ local minimizers of $\mu(t)$ and one additional *extrapolation* node at $t_{\max} + \alpha$, where t_{\max} is the currently largest evaluated t , and α is an extrapolation step size starting at $\alpha = 1$ and doubled after each extrapolation step.⁵

A conceptual (rather than algorithmic) motivation for using the integrated Wiener process as surrogate for the objective, as well as for the described candidate selection, are classic line searches. There, the 1D-objective is modeled by piecewise cubic interpolations between neighboring datapoints. In a sense, this is a non-parametric approach, since a new spline is defined, when a datapoint is added. Classic line searches always only deal with one spline at a time, since they are able to collapse all other parts of the search space. Indeed, for noise free observations, the mean of the posterior GP is identical to the classic cubic interpolations, and thus candidate locations are identical as well; this is illustrated in Figure 3. The non-parametric approach also prevents issues of over-constrained surrogates for more than two datapoints. For example, unless the objective is a perfect cubic function, it is impossible to fit a parametric third order polynomial to it, for more than two noise free observations. All other variability in the objective would need to be explained away by artificially introducing noise on the observations. An integrated Wiener process very naturally extends its complexity with each newly added datapoint without being overly assertive – the encoded assumption is that the objective has *at least* one derivative (which is also observed in this case).

4. There is no well-defined probabilistic belief over f'' and higher derivatives—sample paths of the Wiener process are almost surely non-differentiable almost everywhere (Adler, 1981, §2.2). But $\mu(t)$ is always a member of the reproducing kernel Hilbert space induced by k , thus piecewise cubic (Rasmussen and Williams, 2006, §6.1).

5. For the integrated Wiener process and heteroscedastic noise, the variance *always* attains its maximum exactly at the mid-point between two evaluations; including the variance into the candidate selection biases the existing candidates towards the center (additional candidates might occur between evaluations without local minimizer, even for noise free observations/classic line searches). We did not explore this further since the algorithm showed very good sample efficiency already with the adopted scheme.




Figure 4: *Candidate selection* by Bayesian optimization. *Top*: GP marginal posterior of function values. Posterior mean in solid orange and, two standard deviations in thinner solid orange, local pdf marginal as shading. The red and the blue are evaluations of the objective function, collected by the line search. *Middle*: GP marginal posterior of corresponding gradients. Colors same as in top plot. In all three plots the locations of the *two* candidate points (§3.1) are indicated as vertical dark red lines. The left one at about $t_1^{\text{cand}} \approx 1.54$ is a local minimum of the posterior mean in between the red and blue point (the mean of the gradient belief (solid orange, middle plot) crosses through zero here). The right one at $t_2^{\text{cand}} = 4$ is a candidate for extrapolation. *Bottom*: Decision criterion in arbitrary scale: The expected improvement u_{EI} (Eq. 9) is shown in dashed light blue, the Wolfe probability p^{Wolfe} (Eq. 14 and Eq. 16) in light red and their decisive product in solid dark blue. For illustrative purposes all criteria are plotted for the whole t -space. In practice solely the values at t_1^{cand} and t_2^{cand} are computed, compared, and the candidate with the higher value of $u_{\text{EI}} \cdot p^{\text{Wolfe}}$ is chosen for evaluation. In this example this would be the candidate at t_1^{cand} .

3.2 Choosing Among Candidates

The previous section described the construction of $< N + 1$ discrete candidate points for the next evaluation. To decide at which of the candidate points to actually call f and f' , we make use of a popular acquisition function from Bayesian optimization. *Expected improvement* (Jones et al., 1998) is the expected amount, under the GP surrogate, by which the function $f(t)$ might be smaller than a ‘current best’ value η (we set $\eta = \min_{i=0, \dots, N} \{\mu(t_i)\}$, where t_i

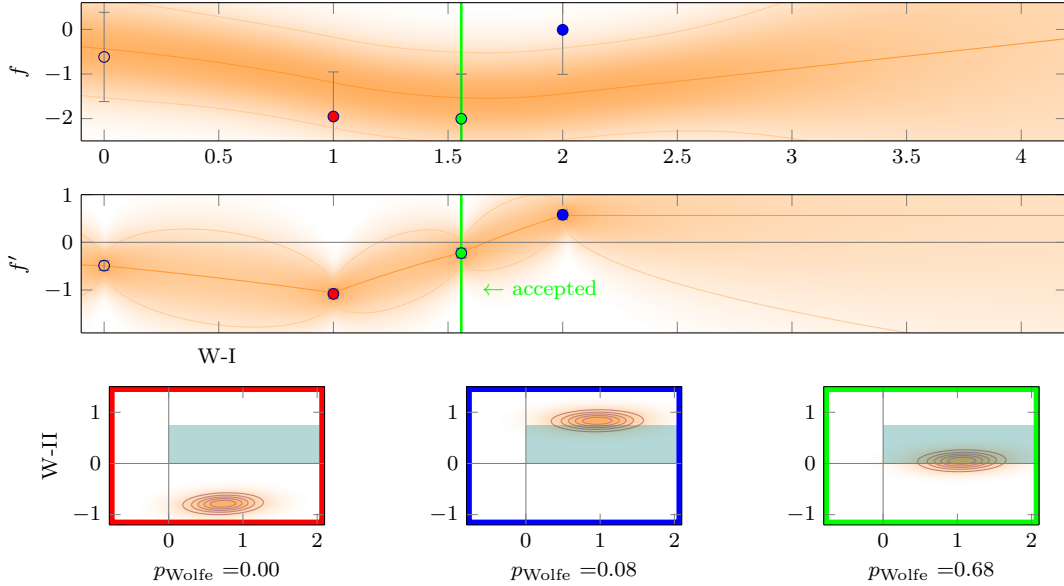


Figure 5: *Acceptance procedure.* *Top and middle:* plot and colors as in Figure 4 with an additional ‘green’ observation. *Bottom:* Implied bivariate Gaussian belief over the validity of the Wolfe conditions (Eq. 11) at the red, blue and green point respectively. Points are considered acceptable if their Wolfe probability p_t^{Wolfe} is above a threshold $c_W = 0.3$; this means that at least 30% of the orange 2D Gauss density must cover greenish shaded area. Only the green point fulfills this condition and is therefore accepted.

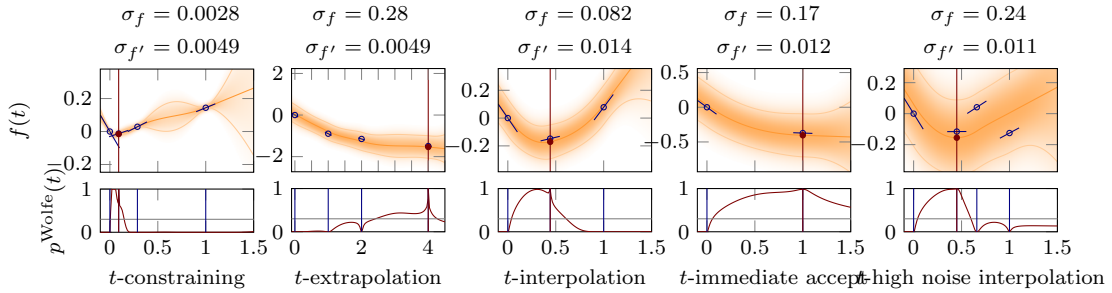


Figure 6: *Curated snapshots* of line searches (from N-I on MNIST), showing variability of the objective’s shape and the decision process. *Top row:* GP marginal posterior of function values and evaluations, *bottom row:* approximate p^{Wolfe} over strong Wolfe conditions. Accepted point marked red.

are observed locations),

$$\begin{aligned} u_{\text{EI}}(t) &= \mathbb{E}_{p(f_t | \mathbf{y}, \mathbf{y}')} [\min\{0, \eta - f(t)\}] \\ &= \frac{\eta - \mu(t)}{2} \left(1 + \operatorname{erf} \frac{\eta - \mu(t)}{\sqrt{2\mathbb{V}(t)}} \right) + \sqrt{\frac{\mathbb{V}(t)}{2\pi}} \exp \left(-\frac{(\eta - \mu(t))^2}{2\mathbb{V}(t)} \right). \end{aligned} \quad (9)$$

The next evaluation point is chosen as the candidate maximizing the product of Eq. 9 and Wolfe probability p^{Wolfe} , which is derived in the following section. The intuition is that p^{Wolfe} precisely encodes properties of desired points, but has poor exploration properties; u_{EI} has better exploration properties, but lacks the information that we are seeking a point with low curvature; u_{EI} thus puts weight on (by W-II) clearly ruled out points. An illustration of the candidate proposal and selection is shown in Figure 4.

In principle other acquisition functions (e.g. the upper-confidence bound, GP-UCB (Srinivas et al., 2010)) are possible, which might have a stronger explorative behavior; we opted for u_{EI} since exploration is less crucial for line searches than for general BO and some (e.g. GP-UCB) had one additional parameter to tune. We tracked the sample efficiency of u_{EI} instead and it was very good (low); the experimental Subsection 4.3 contains further comments and experiments on the alternative choices of u_{EI} and p^{Wolfe} as standalone acquisition functions; they performed equally well (in terms of loss and sample efficiency) to their product on the tested setups.

3.3 Probabilistic Wolfe Conditions for Termination

The key observation for a probabilistic extension of the Wolfe conditions W-I and W-II is that they are positivity constraints on two variables a_t, b_t that are both linear projections of the (jointly Gaussian) variables f and f' :

$$\begin{bmatrix} a_t \\ b_t \end{bmatrix} = \begin{bmatrix} 1 & c_1 t & -1 & 0 \\ 0 & -c_2 & 0 & 1 \end{bmatrix} \begin{bmatrix} f(0) \\ f'(0) \\ f(t) \\ f'(t) \end{bmatrix} \geq 0. \quad (10)$$

The GP of Eq. (5) on f thus implies, at each value of t , a bivariate Gaussian distribution

$$p(a_t, b_t) = \mathcal{N} \left(\begin{bmatrix} a_t \\ b_t \end{bmatrix}; \begin{bmatrix} m_t^a \\ m_t^b \end{bmatrix}, \begin{bmatrix} C_t^{aa} & C_t^{ab} \\ C_t^{ba} & C_t^{bb} \end{bmatrix} \right), \quad (11)$$

$$\text{with } m_t^a = \mu(0) - \mu(t) + c_1 t \mu'(0)$$

$$m_t^b = \mu'(t) - c_2 \mu'(0) \quad (12)$$

$$\text{and } C_t^{aa} = \tilde{k}_{00} + (c_1 t)^2 \partial \tilde{k}_{00}^\partial + \tilde{k}_{tt} + 2[c_1 t (\tilde{k}_{00}^\partial - \partial \tilde{k}_{0t}) - \tilde{k}_{0t}]$$

$$C_t^{bb} = c_2^2 \partial \tilde{k}_{00}^\partial - 2c_2 \partial \tilde{k}_{0t}^\partial + \partial \tilde{k}_{tt}^\partial \quad (13)$$

$$C_t^{ab} = C_t^{ba} = -c_2 (\tilde{k}_{00}^\partial + c_1 t \partial \tilde{k}_{00}^\partial) + c_2 \partial \tilde{k}_{0t}^\partial + \partial \tilde{k}_{t0}^\partial + c_1 t \partial \tilde{k}_{0t}^\partial - \tilde{k}_{tt}^\partial.$$

The quadrant probability $p_t^{\text{Wolfe}} = p(a_t > 0 \wedge b_t > 0)$ for the Wolfe conditions to hold, is an integral over a bivariate normal probability,

$$p_t^{\text{Wolfe}} = \int_{-\frac{m_t^a}{\sqrt{C_t^{aa}}}}^{\infty} \int_{-\frac{m_t^b}{\sqrt{C_t^{bb}}}}^{\infty} \mathcal{N} \left(\begin{bmatrix} a \\ b \end{bmatrix}; \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho_t \\ \rho_t & 1 \end{bmatrix} \right) da db, \quad (14)$$

with correlation coefficient $\rho_t = C_t^{ab} / \sqrt{C_t^{aa}C_t^{bb}}$. It can be computed efficiently (Drezner and Wesolowsky, 1990), using readily available code.⁶ The line search computes this probability for all evaluation nodes, after each evaluation. If any of the nodes fulfills the Wolfe conditions with $p_t^{\text{Wolfe}} > c_W$, greater than some threshold $0 < c_W \leq 1$, it is accepted and returned. If several nodes simultaneously fulfill this requirement, the most recently evaluated node is returned; there are additional safeguards for cases where e.g. no Wolfe-point can be found, which can be deduced from the pseudo-code in Appendix D; they are similar to standard safeguards of classic line search routines (e.g. returning the node of lowest mean). Section 3.4.1 below motivates fixing $c_W = 0.3$. The acceptance procedure is illustrated in Figure 5.

3.3.1 APPROXIMATION FOR STRONG CONDITIONS:

As noted in Section 2.1.1, deterministic optimizers tend to use the strong Wolfe conditions, which use $|f'(0)|$ and $|f'(t)|$. A precise extension of these conditions to the probabilistic setting is numerically taxing, because the distribution over $|f'|$ is a non-central χ -distribution, requiring customized computations. However, a straightforward variation to 14 captures the spirit of the strong Wolfe conditions that large positive derivatives should not be accepted: Assuming $f'(0) < 0$ (i.e. that the search direction is a descent direction), the strong second Wolfe condition can be written exactly as

$$0 \leq b_t = f'(t) - c_2 f'(0) \leq -2c_2 f'(0). \quad (15)$$

The value $-2c_2 f'(0)$ is bounded to 95% confidence by

$$-2c_2 f'(0) \lesssim 2c_2 (|\mu'(0)| + 2\sqrt{\mathbb{V}'(0)}) =: \bar{b}. \quad (16)$$

Hence, an approximation to the strong Wolfe conditions can be reached by replacing the infinite upper integration limit on b in Eq. 14 with $(\bar{b} - m_t^b) / \sqrt{C_t^{bb}}$. The effect of this adaptation, which adds no overhead to the computation, is shown in Figure 2 as a dashed line.

3.4 Eliminating Hyper-parameters

As a black-box inner loop, the line search should not require any tuning by the user. The preceding section introduced six so-far undefined parameters: $c_1, c_2, c_W, \theta, \sigma_f, \sigma_{f'}$. We will now show that c_1, c_2, c_W , can be fixed by hard design decisions: θ can be eliminated by standardizing the optimization objective within the line search; and the noise levels can be estimated at runtime with low overhead for finite-sum objectives of the form in Eq. 1. The result is a parameter-free algorithm that effectively *removes* the one most problematic parameter from SGD—the learning rate.

3.4.1 DESIGN PARAMETERS c_1, c_2, c_W

Our algorithm inherits the Wolfe thresholds c_1 and c_2 from its deterministic sibling. We set $c_1 = 0.05$ and $c_2 = 0.5$. This is a standard setting that yields a ‘lenient’ line search, i.e. one

6. e.g. <http://www.math.wsu.edu/faculty/genz/software/matlab/bvn.m>

that accepts most descent points. The rationale is that the stochastic aspect of SGD is not always problematic, but can also be helpful through a kind of ‘annealing’ effect.

The acceptance threshold c_W is a new design parameter arising only in the probabilistic setting. We fix it to $c_W = 0.3$. To motivate this value, first note that in the noise-free limit, all values $0 < c_W < 1$ are equivalent, because p^{Wolfe} then switches discretely between 0 and 1 upon observation of the function. A back-of-the-envelope computation, assuming only two evaluations at $t = 0$ and $t = t_1$ and the same fixed noise level on f and f' (which then cancels out), shows that function values barely fulfilling the conditions, i.e. $a_{t_1} = b_{t_1} = 0$, can have $p^{\text{Wolfe}} \sim 0.2$ while function values at $a_{t_1} = b_{t_1} = -\epsilon$ for $\epsilon \rightarrow 0$ with ‘unlucky’ evaluations (both function and gradient values one standard-deviation from true value) can achieve $p^{\text{Wolfe}} \sim 0.4$. The choice $c_W = 0.3$ balances the two competing desiderata for precision and recall. Empirically (Fig. 6), we rarely observed values of p^{Wolfe} close to this threshold. Even at high evaluation noise, a function evaluation typically either clearly rules out the Wolfe conditions, or lifts p^{Wolfe} well above the threshold. A more in-depth analysis of c_1 , c_2 , and c_W is done in the experimental Section 4.2.1.

3.4.2 SCALE θ

The parameter θ of Eq. 4 simply scales the prior variance. It can be eliminated by scaling the optimization objective: We set $\theta = 1$ and scale $y_i \leftarrow (y_i - y_0)/|y'_0|$, $y'_i \leftarrow y'_i/|y'_0|$ within the code of the line search. This gives $y(0) = 0$ and $y'(0) = -1$, and typically ensures the objective ranges in the single digits across $0 < t < 10$, where most line searches take place. The division by $|y'_0|$ causes a non-Gaussian disturbance, but this does not seem to have notable empirical effect.

3.4.3 NOISE SCALES $\sigma_f, \sigma_{f'}$

The likelihood 3 requires standard deviations for the noise on both function values (σ_f) and gradients ($\sigma_{f'}$). One could attempt to learn these across several line searches; but the resulting estimator would be biased. In exchangeable models as captured by Eq. 1, however, the variance of the loss and its gradient can be estimated locally and unbiased, directly for the mini-batch, at low computational overhead—an approach already advocated by Schaul et al. (2013). We collect the empirical statistics

$$\hat{S}(x) := \frac{1}{m} \sum_j \ell^2(x, y_j), \quad \text{and} \quad \hat{\nabla}S(x) := \frac{1}{m} \sum_j \nabla \ell(x, y_j)^{\odot 2} \quad (17)$$

(where $\odot 2$ denotes the element-wise square) and estimate, at the beginning of a line search from x_k ,

$$\sigma_f^2 = \frac{1}{m-1} \left(\hat{S}(x_k) - \hat{\mathcal{L}}(x_k)^2 \right) \quad \text{and} \quad \sigma_{f'}^2 = s_i^{\odot 2\top} \left[\frac{1}{m-1} \left(\hat{\nabla}S(x_k) - (\nabla \hat{\mathcal{L}}(x_k))^{\odot 2} \right) \right]. \quad (18)$$

This amounts to the assumption that noise on the gradient is independent (see also Appendix A). We finally scale the two empirical estimates as described in Section §3.4.2: $\sigma_f \leftarrow \sigma_f/|y'(0)|$, and ditto for $\sigma_{f'}$. The overhead of this estimation is small if the computation of $\ell(x, y_j)$ itself is more expensive than the summation over j . In the neural network

examples N-I and N-II of the experimental Section 4, the additional steps added only $\sim 1\%$ cost overhead to the evaluation of the loss. This is rather at the lower end for these models; A more general statement about memory and time requirements for neural networks can be found in Sections 3.6 and 3.7. Estimating noise separately for each input dimension captures the often inhomogeneous structure among gradient elements, and its effect on the noise along the projected direction. For example, in multi-layer models, gradient noise is typically higher on weights between the input and first hidden layer, hence line searches along the corresponding directions are noisier than those along directions affecting higher-level weights.

3.4.4 PROPAGATING STEP SIZES BETWEEN LINE SEARCHES

As will be demonstrated in §4, the line search can find good step sizes even if the length of the direction s_i is mis-scaled. Since such scale issues typically persist over time, it would be wasteful to have the algorithm re-fit a good scale in each line search. Instead, we propagate step lengths from one iteration of the search to another: We set the initial search direction to $s_0 = -\alpha_0 \nabla \hat{\mathcal{L}}(x_0)$ with some initial learning rate α_0 . Then, after each line search ending at $x_i = x_{i-1} + t_* s_i$, the next search direction is set to $s_{i+1} = -\alpha_{\text{ext}} \cdot t_* \alpha_0 \nabla \hat{\mathcal{L}}(x_i)$ (with $\alpha_{\text{ext}} = 1.3$). Thus, the next line search starts its extrapolation at 1.3 times the step size of its predecessor (Section 4.2.2 for details).

3.5 Relation to Bayesian Optimization and Noise-Free Limit

The probabilistic line search algorithm is closely related to Bayesian optimization (BO) since it approximately minimizes a 1D-objective under potentially noisy function evaluations. It thus uses notions of BO (e.g. a GP-surrogate for the objective, and an acquisition function to discriminate locations for the next evaluation of the loss), but there are some differences concerning the aim, requirements on computational efficiency, and termination condition, which are shortly discussed here: (i) *Performance measure*: The final performance measure in BO is usually the lowest found value of the objective function. Line searches are subroutines inside of a greedy, iterative optimization machine, which usually performs several thousand steps (and line searches); many, very approximate steps often performs better than taking less, but preciser steps. (ii) *Termination*: The termination condition of a line search is imposed from the outside in the form of the Wolfe conditions. Stricter Wolfe conditions do not usually improve the performance of the outer optimizer, thus, no matter if a better (lower) minimum could be found, any Wolfe-point is acceptable at all times. (iii) *Sample efficiency*: Since the last evaluation from the previous line search can be re-used in the current line search, only one additional value and gradient evaluation is enough to terminate the procedure. This ‘immediate-accept’ is the desired behavior if the learning rate is currently well calibrated. (iv) *Locations for evaluation*: BO, usually calls an optimizer to maximize some acquisition function, and the preciseness of this optimization is crucial for performance. Line searches just need to find a Wolfe-acceptable point; classic line searches suggest that it is enough to look at plausible locations, like minimizer of a local interpolator, or some rough extrapolation point; this inexpensive heuristic usually works rather well. (v) *Exploration*: BO needs to solve an intricate trade-off problem in between exploring enough of the parameters space for possible locations of minima, and exploiting locations around them further. Since line searches are only concerned with finding

a Wolfe-point, they do not need to explore the parameter space of possible step sizes to that extend; crucial features are rather the possibility to explore somewhat larger steps than previous ones (which is done by extrapolation-candidates), and likewise to shorted steps (which is done by interpolation-candidates).

In the limit of noise free observed gradients and function values ($\sigma_f = \sigma_{f'} = 0$) the probabilistic line search behaves like its classic parent, except for very slight variations in the candidate choice (building block 3): The GP-mean reverts to the classic interpolator; all candidate locations are thus identical, but the probabilistic line search might propose a second option, since (even if there is a local minimizer) it *always* also proposes an extrapolation candidate. This is illustrated in the following table.

building block	classic	probabilistic (noise free)
1) 1D surrogate for objective $f(t)$	piecewise cubic splines	GP-mean identical to classic interpolator
2) candidate selection	local minimizer of cubic splines <i>xor</i> extrapolation	local minimizer of cubic splines <i>or</i> extrapolation
3) choice of best candidate	—————	BO acquisition function
4) acceptance criterion	classic Wolfe conditions	p^{Wolfe} identical to classic Wolfe conditions

3.6 Computational Time Overhead

The line search routine itself has little memory and time overhead; most importantly it is independent of the dimensionality of the optimization problem. After every call of the objective function, the GP (§3.1) needs to be updated which, at most, is at the cost of inverting a $2N \times 2N$ -matrix, where N usually is equal to 1, 2, or 3 but never > 10 . In addition, the bivariate normal integral p_t^{Wolfe} of Eq. 14 needs to be computed at most N times. On a laptop, one evaluation of p_t^{Wolfe} costs about 100 microseconds. For the choice among proposed candidates (§3.2), again at most N , for each, we need to evaluate p_t^{Wolfe} and $u_{\text{EI}}(t)$ (Eq. 9) where the latter comes at the expense of evaluating two error functions. Since all of these computations have a fixed cost (in total some milliseconds on a laptop), the relative overhead becomes less the more expensive the evaluation of $\nabla \hat{\mathcal{L}}(x)$.

The largest overhead actually lies outside of the actual line search routine. In case the noise levels σ_f and $\sigma_{f'}$ are not known, we need to estimate them. The approach we took is described in Section 3.4.3 where the variance of $\nabla \hat{\mathcal{L}}$ is estimated using the sample variance of the mini-batch, each time the objective function is called. Since in this formulation the variance estimation is about half as expensive as one backward pass of the net, the time overhead depends on the relative cost of the feed forward and backward passes (Balles et al., 2017). If forward and backward pass are the same cost, the most straightforward implementation of the variance estimation would make each function call < 1.3 times as expensive. This is an *upper bound* and the actual cost is usually lower.⁷ At the same time though, *all* exploratory experiments which very considerably increase the time spend when

7. It is desirable to decrease this value in the future reusing computation results or by approximation but this is beyond this discussion.

using SGD with a hand tuned learning rate schedule need not be performed anymore. In Section 4.1 we will also see that SGD using the probabilistic line search often needs less function evaluations to converge, which might lead to overall faster convergence in wall clock time than classic SGD in a single run.

3.7 Memory Requirement

Vanilla SGD, at all times, keeps around the current optimization parameters $x \in \mathbb{R}^D$ and the gradient vector $\nabla \hat{\mathcal{L}}(x) \in \mathbb{R}^D$. In addition to this, the probabilistic line search needs to store the estimated gradient variances $\Sigma'(x) = (1 - m)^{-1}(\hat{\nabla} S(x) - \nabla \hat{\mathcal{L}}(x))^{\odot 2}$ (Eq. 18) of same size. The memory requirement of SGD+PROBLS is thus comparable to ADAGRAD or ADAM. If combined with a search direction other than SGD always one additional vector of size D needs to be stored.

4. Experiments

This section reports on an extensive set of experiments to characterise and test the line search. The overall evidence from these tests is that the line search performs well and is relatively insensitive to the choice of its internal hyper-parameters as well the mini-batch size. We performed experiments on two multi-layer perceptrons N-I and N-II; both were trained on two well known datasets MNIST and CIFAR-10.

- N-I: fully connected net with 1 hidden layer and 800 hidden units + biases, and 10 output units, sigmoidal activation functions and a cross entropy loss. Structure without biases: 784-800-10. Many authors used similar nets and reported performances.⁸
- N-II: fully connected net with 3 hidden layers and 10 output units, tanh-activation functions and a squared loss. Structure without biases: 784-1000-500-250-10. Similar nets were also used for example in Martens (2010) and Sutskever et al. (2013).
- MNIST (LeCun et al., 1998): multi-class classification task with 10 classes: hand-written digits in gray-scale of size 28×28 (numbers ‘0’ to ‘9’); training set size 60 000, test set size 10 000.
- CIFAR-10 (Krizhevsky and Hinton, 2009): multi-class classification task with 10 classes: color images of natural objects (horse, dog, frog, ...) of size 32×32 ; training set size 50 000, test set size 10 000; like other authors, we only used the “batch 1” sub-set of CIFAR-10 containing 10 000 training examples.

In addition we train logistic regressors with sigmoidal output (N-III) on the following binary classification tasks:

- Wisconsin Breast Cancer Dataset (WDBC) (Wolberg et al., 2011): binary classification of tumors as either ‘malignant’ or ‘benign’. The set consist of 569 examples of which we used 169 to monitor generalization performing; thus 400 remain for the training set; 30 features describe for example radius, area, symmetry, et cetera. In comparison

8. <http://yann.lecun.com/exdb/mnist/>

to the other datasets and networks, this yields a very low dimensional optimization problem with only 30 (+1 bias) input parameters as well as just a small number of datapoints.

- GISETTE (Guyon et al., 2005): binary classification of the handwritten digits ‘4’ and ‘9’. The original 28×28 images are taken from the MNIST dataset; then the feature set was expanded and consists of the original normalized pixels, plus a randomly selected subset of products of pairs of features, which are slightly biased towards the upper part of the image; in total there are 5000 features, instead of 784 as in the original MNIST. The size of the training set and test set is 6000 and 1000 respectively.
- EPSILON: synthetic dataset from the PASCAL Challenge 2008 for binary classification. It consists of 400 000 training set datapoint and 100 000 test set datapoints, each having 2000 features.

In the text and figures, SGD using the probabilistic line search will occasionally be denoted as SGD+PROBLS. Section 4.1 contains experiments on the sensitivity to varying gradient noise levels (mini-batch sizes) performed on both multi-layer perceptrons N-I and N-II, as well as on the logistic regressor N-III. Section 4.2 discusses sensitivity to the hyper-parameters choices introduced in Section 3.4 and Section 4.3 contains additional diagnostics on step size statistics. Each single experiment was performed 10 times with different random seeds that determined the starting weights and the mini-batch selection and seeds were shared across all experiments. We report all results of the 10 instances as well as means and standard deviations.

4.1 Varying Mini-batch Sizes

The noise level of the gradient estimate $\nabla \hat{\mathcal{L}}(x)$ and the loss $\hat{\mathcal{L}}(x)$ is determined by the mini-batch size m and ultimately there should exist an optimal m that maximizes the optimizer’s performance in wall-clock-time. In practice of course the cost of computing $\nabla \hat{\mathcal{L}}(x)$ and $\hat{\mathcal{L}}(x)$ is not necessarily linear in m since it is upper bounded by the memory capacity of the hardware used. We assume here that the mini-batch size is chosen by the user; thus we test the line search with the default hyper-parameter setting (see Sections 3.4 and 4.2) on four different mini-batch sizes:

- $m = 10, 100, 200$ and 1000 (for MNIST, CIFAR-10, and EPSILON)
- $m = 10, 50, 100$, and 400 (for WDBC and GISETTE)

which correspond to increasing signal-to-noise ratios. Since the training set of WDBC only consists of 400 datapoints, the run with the largest mini-batch size of 400 in fact runs full-batch gradient descent on WDBC; this is not a problem, since—as discussed above—the probabilistic line search can also handle noise free observations.⁹ We compare to SGD-runs using a fixed step size (which is typical for these architectures) and an annealed step size with annealing schedule $\alpha_t = \alpha_0/t$. Because annealed step sizes performed much worse than

9. Since the dataset size M of WDBC is very small, we used the factor $\binom{M-m}{mM}$ instead of $1/m$ to scale the sample variances of Eq. 17. The former encodes sampling mini-batches \mathcal{B} with replacement, the latter without replacement; for $m \ll M$ both factors are nearly identical.

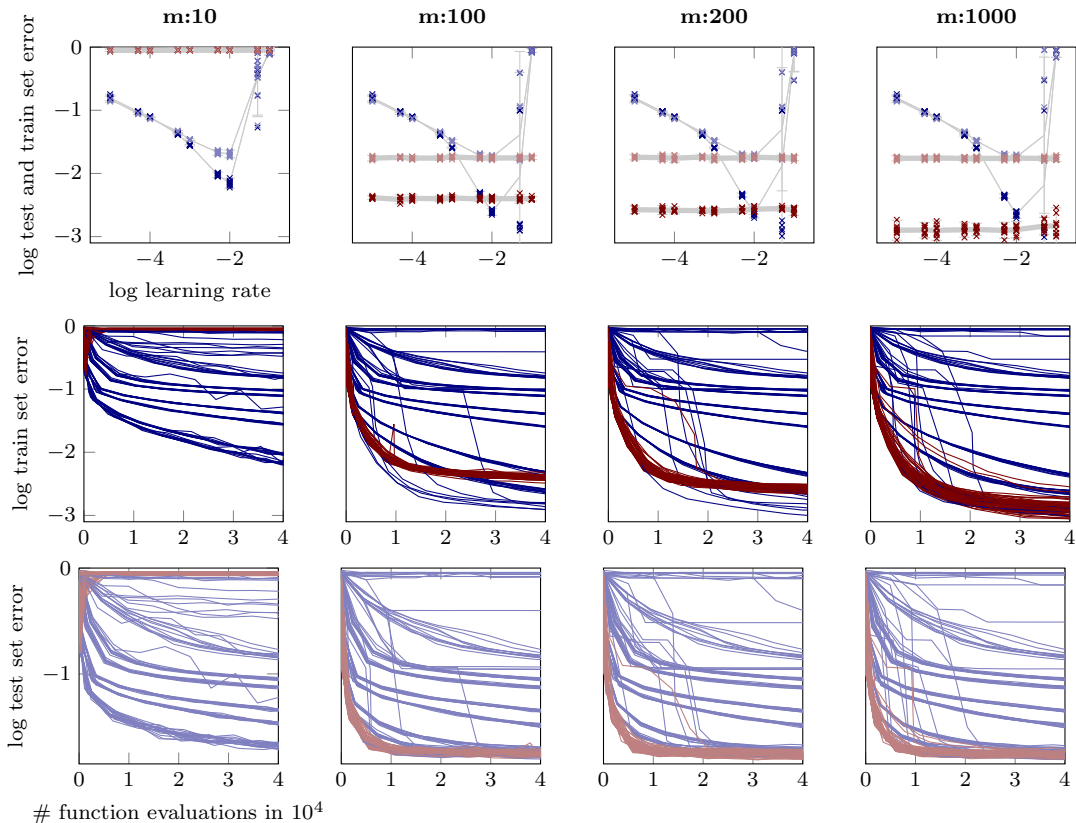


Figure 7: Performance of N-II on MNIST for *varying mini-batch sizes*. *Top*: final logarithmic test set and train set error after 40 000 function evaluations of training versus a large range of learning rates each for 10 different initializations. SGD-runs with fixed learning rates are shown in light blue (test set) and dark blue (train set); SGD+PROBLS-runs in light red (test set) and dark red (train set); means and two standard deviations for each of the 10 runs in gray. *Columns* from left to right refer to different mini-batch sizes m of 10, 100, 200 and 1000 which correspond to decreasing relative noise in the gradient observations. Not surprisingly the performance of SGD-runs with a fixed step size are very sensitive to the choice of this step size. SGD using the probabilistic line search adapts initially mis-scaled step sizes and performs well across the whole range of initial learning rates. *Middle and bottom*: Evolution of the logarithmic test and train set error respectively for all SGD-runs and SGD+PROBLS-runs versus $\#$ function evaluations (colors as in top plot). For mini-batch sizes of $m = 100, 200$ and 1000 *all* instances of SGD using the probabilistic line search reach the same best test set error. Similarly a good train set error is reached very fast by SGD+PROBLS. Only very few instances of SGD with a fixed learning rate reach a better train set error (and this advantage usually does not translate to test set error). For very small mini-batch sizes ($m = 10$, first column) the line search performs poorly on this architecture, most likely because of the variance estimation becoming too inaccurate.

SGD+fixed step size, we will only report on the latter results in the plots.¹⁰ Since classic SGD without the line search needs a hand crafted learning rate we search on exhaustive logarithmic grids of

$$\begin{aligned}\alpha_{\text{SGD}}^{\text{N-I}} &= [10^{-5}, 5 \cdot 10^{-5}, 10^{-4}, 5 \cdot 10^{-4}, 10^{-3}, 5 \cdot 10^{-3}, 10^{-2}, 5 \cdot 10^{-2}, 10^{-1}, 5 \cdot 10^{-1}] \\ \alpha_{\text{SGD}}^{\text{N-II}} &= [\alpha_{\text{SGD}}^{\text{N-I}}, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0] \\ \alpha_{\text{SGD}}^{\text{N-III}} &= [10^{-8}, 10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1, 10^2].\end{aligned}$$

We run 10 different initialization for each learning rate, each mini-batch size and each net and dataset combination ($10 \cdot 4 \cdot (2 \cdot 10 + 2 \cdot 17 + 3 \cdot 11) = 3480$ runs in total) for a large enough budget to reach convergence; and report all numbers. Then we perform the same experiments using the same seeds and setups with SGD using the probabilistic line search and compare the results. For SGD+PROBLS, α_{SGD} is the initial learning rate which is used in the very first step. After that, the line search automatically adapts the learning rate, and shows no significant sensitivity to its initialization.

Results of N-I and N-II on both, MNIST and CIFAR-10 are shown in Figures 7, 14, 15, and 16; results of N-III on WDBC, GISETTE and EPSILON are shown in Figures 18, 17, and 19 respectively. All instances (SGD and SGD+PROBLS) get the same computational budget (number of mini-batch evaluations) and not the same number of optimization steps. The latter would favour the probabilistic line search since, on average, a bit more than one mini-batch is evaluated per step. Likewise, all plots show performance measure versus the number of mini-batch evaluations, which is proportional to the computational cost.

All plots show similar results: While classic SGD is sensitive to the learning rate choice, the line search-controlled SGD performs as good, close to, or sometimes even better than the (in practice unknown) optimal classic SGD instance. In Figure 7, for example, SGD+PROBLS converges much faster to a good test set error than the best classic SGD instance. In all experiments, across a reasonable range of mini-batch sizes m and of initial α_{SGD} values, the line search quickly identified good step sizes α_t , stabilized the training, and progressed efficiently, reaching test set errors similar to those reported in the literature for tuned versions of these kind of architectures and datasets. The probabilistic line search thus effectively removes the need for exploratory experiments and learning-rate tuning.

Overfitting and training error curves: The training error of SGD+PROBLS often plateaus earlier than the one of vanilla SGD, especially for smaller mini-batch sizes. This does not seem to impair the performance of the optimizer on the test set. We did not investigate this further, since it seemed like a nice natural annealing effect; the exact causes are unclear for now. One explanation might be that the line search does indeed improve overfitting, since it tries to measure descent (by Wolfe conditions which rely on the *noise-informed* GP). This means that, if—close to a minimum—successive acceptance decisions can not identify a descent direction anymore, diffusion might set in.

4.2 Sensitivity to Design Parameters

Most, if not all, numerical methods make implicit or explicit choices about their hyper-parameters. Most of these are never seen by the user since they are either estimated at run

¹⁰. An example of annealed step size performance can be found in Mahsereci and Hennig (2015).




Figure 8: *Sensitivity to varying hyper-parameters θ_{reset} .* Plot and color coding as in Figure 9. Adopted parameter in dark red at $\theta_{\text{reset}} = 100$. Resetting the GP scale occurs very rarely. For example for $\theta_{\text{reset}} = 100$ the reset occurred in 0.02% of all line searches.

time, or set by design to a fixed, approximately insensitive value. Well known examples are the discount factor in ordinary differential equation solvers (Hairer et al., 1987, §2.4), or the Wolfe parameters c_1 and c_2 of classic line searches (§3.4.1). The probabilistic line search inherits the Wolfe parameters c_1 and c_2 from its classical counterpart as well as introducing two more: The Wolfe threshold c_W and the extrapolation factor α_{ext} . c_W does not appear in the classical formulation since the objective function can be evaluated exactly and the Wolfe probability is binary (either fulfilled or not). While c_W is thus a natural consequence of allowing the line search to model noise explicitly, the extrapolation factor α_{ext} is the result of the line search favoring shorter steps, which we will discuss below in more detail, but most prominently because of bias in the line search’s first gradient observation.

In the following sections we will give an intuition about the task of the most influential design parameters c_2 , c_W , and α_{ext} , discuss how they affect the probabilistic line search, and validate good design choices through exploring the parameter space and showing insensitivity to most of them. All experiments on hyper-parameter sensitivity were performed training N-II on MNIST with mini-batch size $m = 200$. For a full search of the parameter space c_W - c_2 - α_{ext} we performed 4950 runs in total with 495 different parameter combinations. All results are reported.

4.2.1 WOLFE II PARAMETER c_2 AND WOLFE THRESHOLD c_W

As described in Section 3.4, c_2 encodes the strictness of the curvature condition W-II. Pictorially speaking, a larger c_2 extends the range of acceptable gradients (green shaded are in the lower part of Figure 5) and leads to a lenient line search while a smaller value of c_2 shrinks this area, leading to a stricter line search. c_W controls how certain we want to be that the Wolfe conditions are actually fulfilled (pictorially, how much mass of the 2D-Gauss need to lie in the green shaded area). In the extreme case of complete uncertainty about the collected gradients and function values (roughly $\det[\text{cov}[a, b]] \rightarrow \infty$) p^{Wolfe} will always be ≤ 0.25 , if the strong Wolfe conditions are imposed. In the limit of certain observations ($\sigma_f, \sigma_{f'} \rightarrow 0$) p^{Wolfe} is binary and reverts to the classic Wolfe criteria. An overly strict line search, therefore (e.g. $c_W = 0.99$ and/ or $c_2 = 0.1$), will still be able to optimize the objective function well, but will waste evaluations at the expense of efficiency. Figure 10 explores the c_2 - c_W parameter space (while keeping α_{ext} fixed at 1.3). The left column shows final test and train set error, the right column the average number of function evaluations

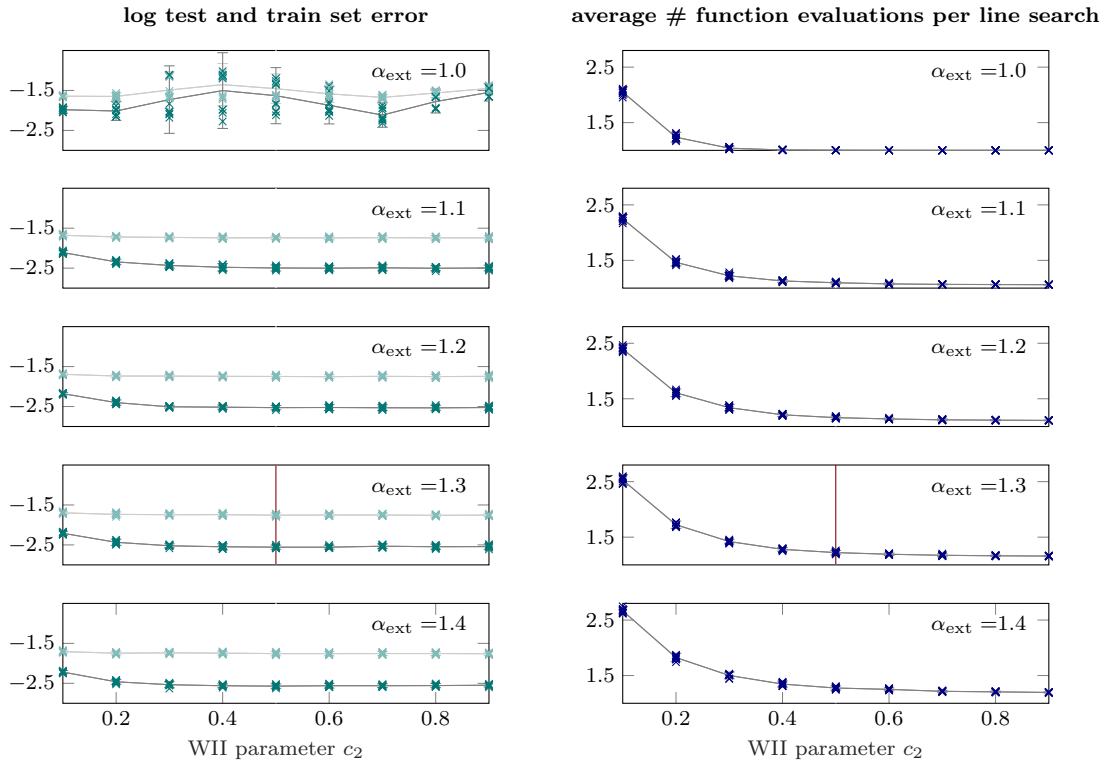


Figure 9: *Sensitivity to varying hyper-parameters c_2 , and α_{ext} .* Runs were performed training N-II on MNIST with mini-batch size $m = 200$. For each parameter setting 10 runs with different initializations were performed. *Left column:* logarithmic test set error (light green) and train set error (dark green) after 40 000 function evaluations; mean and \pm two standard deviations of the 10 runs in gray. *Right Column:* average number of function evaluations per line search. A low number indicates an efficient line search procedure (perfect efficiency at 1). For most parameter combinations this lies around $\approx 1.3 - 1.5$. Only at extreme parameter values, for example $\alpha_{\text{ext}} = 1.0$, which amounts to no extrapolation at all in between successive line searches, the line search performs poorer. The hyper-parameters adopted in the line search implementation are indicated as vertical dark red line at $\alpha_{\text{ext}} = 1.3$ and $c_2 = 0.5$.




Figure 10: *Sensitivity to varying hyper-parameters c_2 , and c_W .* Plot and color coding as in Figure 9 but this time for varying c_W instead of α_{ext} . *Right Column:* Again a low number indicates an efficient line search procedure (perfect efficiency at 1). For most parameter combinations this lies around $\approx 1.3 - 1.5$. Only at extreme parameter values for example $c_W = 0.99$, which amounts to imposing nearly absolute certainty about the Wolfe conditions, the line search becomes less efficient. Adopted parameters again in dark red at $c_W = 0.3$ and $c_2 = 0.5$

per line search, both versus different choices of Wolfe parameter c_2 . The left column thus shows the overall performance of the optimizer, while the right column is representative for the computational efficiency of the line search. Intuitively, a line search which is minimally invasive (only corrects the learning rate, when it is really necessary) is preferred. Rows in Figure 10 show the same plot for different choices of the Wolfe threshold c_W .

The effect of strict c_2 can be observed clearly in Figure 10 where for smaller values of $c_2 < \approx 0.2$ the average number of function evaluations spend in one line search goes up slightly in comparison to looser restrictions on c_2 , while still a very good performance is reached in terms of train and test set error. Likewise, the last row of Figure 10 for the extreme value of $c_W = 0.99$ (demanding 99% certainty about the validity if the Wolfe conditions), shows significant loss in computational efficiency having an average number of 7 function evaluations per line search. Besides losing efficiency, it is still optimizing the objective well. Lowering this threshold a bit to 90% increases the computational efficiency of the line search to be nearly optimal again.

Ideally, we want to trade off the desiderata of being strict enough to reject too small and too large steps that prevent the optimizer to converge, but being lenient enough to allow all other reasonable steps, thus increasing computational efficiency. The values $c_W = 0.3$ and $c_2 = 0.5$, which are adopted in our current implementation are marked as dark red vertical lines in Figure 10.

4.2.2 EXTRAPOLATION FACTOR α_{EXT}

The extrapolation parameter α_{ext} , introduced in Section 3.4.4, pushes the line search to try a larger learning rate first, than the one which was accepted in the previous step. Figure 9 is structured like Figure 10, but this time explores the line search sensitivity in the c_2 - α_{ext} parameter space (abscissa and rows respectively) while keeping c_W fixed at 0.3. Unless we choose $\alpha_{\text{ext}} = 1.0$ (no step size increase between steps) in combination with a lenient choice of c_2 the line search performs well. For now we adopt $\alpha_{\text{ext}} = 1.3$ as default value which again is shown as dark red vertical line in Figure 9.

The introduction of α_{ext} might seem arbitrary at first, but is a necessity and well-working fix because of a few shortcomings of the current design. First, the curvature condition W-II is the single condition that prevents too small steps and pushes optimization progress. On the other hand both W-I and W-II simultaneously penalize too large steps (see Figure 1 for a sketch). This is not a problem in case of deterministic observation ($\sigma_f, \sigma_{f'} \rightarrow 0$), where W-II undoubtedly decides if a gradient is still too negative. Unless W-II is chosen very tightly (small c_2) or c_W unnecessarily large (both choices, as discussed above, are undesirable), in the presence of noise, p^{Wolfe} will thus be more reliable in preventing overshooting than pushing progress. The first row of Figure 9 illustrates this behavior, where the performance drops somewhat if no extrapolation is done ($\alpha_{\text{ext}} = 1.0$) in combination with a looser version of W-II (larger c_2).

Another factor that contributes towards accepting small rather than larger learning rates is a bias introduced in the first observation of the line search at $t = 0$. Observations $y'(t)$ that the GP gets to see are projections of the gradient sample $\nabla \hat{\mathcal{L}}(t)$ onto the search direction $s = -\nabla \hat{\mathcal{L}}(0)$. Since the first observations $y'(0)$ is computed from the same mini-batch as the search direction (not doing this would double the optimizer’s computational cost) an

inevitable bias is introduced of approximate size of $\cos^{-1}(\gamma)$ (where γ is the expected angle between gradient evaluations from two independent mini-batches at $t = 0$). Since the scale parameter θ of the Wiener process is implicitly set by $y'(0)$ (§3.4.2), the GP becomes more uncertain at unobserved points than it needs to be; or alternatively expects the 1D-gradient to cross zero at smaller steps, and thus underestimates a potential learning rate. The posterior *at* observed positions is little affected. The over-estimation of θ rather pushes the posterior towards the likelihood (since there is less model to trust) and thus still gives a reliable measure for $f(t)$ and $f'(t)$. The effect on the Wolfe conditions is similar. With $y'(0)$ biased towards larger values, the Wolfe conditions, which measure the drop in projected gradient norm, are thus prone to accept larger gradients combined with smaller function values, which again is met by making small steps. Ultimately though, since candidate points at $t^{\text{cand}} > 0$ that are currently queried for acceptance, are always observed and unbiased, this can be controlled by an appropriate design of the Wolfe factor c_2 (§3.4.1 and §4.2.1) and of course α_{ext} .

4.2.3 FULL HYPER-PARAMETER SEARCH: c_W - c_2 - α_{EXT}

An exhaustive performance evaluation on the whole c_W - c_2 - α_{ext} -grid is shown in Appendix C in Figures 20-24 and Figures 25-35. As discussed above, it shows the necessity of introducing the extrapolation parameter α_{ext} and shows slightly less efficient performance for obviously undesirable parameter combinations. In a large volume of the parameter space, and most importantly in the vicinity of the chosen design parameters, the line search performance is stable and comparable to carefully hand tuned learning rates.

4.2.4 SAFEGUARDING MIS-SCALED GPS: θ_{RESET}

For completeness, we performed an additional experiment on the threshold parameter which is denoted by θ_{reset} in the pseudo-code (Appendix D) and safeguards against GP mis-scaling. The introduction of noisy observations necessitates to model the variability of the 1D-function, which is described by the kernel scale parameter θ . Setting this hyper-parameter is implicitly done by scaling the observation input, assuming a similar scale than in the previous line search (§3.4.2). If, for some reason, the previous line search accepted an unexpectedly large or small step (what this means is encoded in θ_{reset}) the GP scale θ for the next line search is reset to an exponential running average of previous scales (α_{stats} in the pseudo-code). This occurs very rarely (for the default value $\theta_{\text{reset}} = 100$ the reset occurred in 0.02% of all line searches), but is necessary to safeguard against extremely mis-scaled GP's. θ_{reset} therefore is not part of the probabilistic line search model as such, but prevents mis-scaled GPs due to some unlucky observation or sudden extreme change in the learning rate. Figure 8 shows performance of the line search for $\theta_{\text{reset}} = 10, 100, 1000$ and 10 000 showing no significant performance change. We adopted $\theta_{\text{reset}} = 100$ in our implementation since this is the expected and desired multiplicative (inverse) factor to maximally vary the learning rate in one single step.

4.3 Candidate Selection and Learning Rate Traces

In the current implementation of the probabilistic line search, the choice among candidates for evaluation is done by evaluating an acquisition function $u_{\text{EI}}(t_i^{\text{cand}}) \cdot p^{\text{Wolfe}}(t_i^{\text{cand}})$ at every




Figure 11: *Different choices of acquisition function.* We compare between using expected improvement u_{EI} (blue), the Wolfe probability p^{Wolfe} (red) and their product $u_{EI} \cdot p^{Wolfe}$ (green) which is the default in our code. *Top:* evolution of logarithmic test and train set error. Different lines of the same color correspond to different seeds. *Rows 2-4:* learning rate traces of a single seed (colors same as in top plot). For plotting purposes the curves were smoothed and thinned out. The thick light green, light red and light blue horizontal lines show the mean of the raw (non-smoothed) values of accepted learning rates across the whole optimization process, the dotted lines show \pm two standard deviations and the gray solid lines mark a range of well performing constant learning rates.




Figure 12: Traces of *accepted logarithmic learning rates*. All runs are performed with default design parameters. Different rows show the same plot for different mini-batch sizes of $m = 100, 200$ and 1000 ; plots and smoothing as in rows 2-4 of Figure 11 (details in text).




Figure 13: Traces of *logarithmic noise levels* σ_f (top), $\sigma_{f'}$ (middle) and average number of function evaluations per line search (bottom). Setup and smoothing as in Figure 12. Different colors correspond to different mini-batch sizes (see legend). Curves of the same color correspond to different seeds (3 shown).

candidate point t_i^{cand} ; then choosing the one with the highest value for evaluation of the objective (§3.2). The Wolfe probability p^{Wolfe} actually encodes precisely what kind of point we want to find and incorporates both (W-I and W-II) conditions about the function value *and* to the gradient (§3.3). However p^{Wolfe} does not have very desirable exploration properties. Since the uncertainty of the GP grows to ‘the right’ of the last observation, the Wolfe probability quickly drops to a low, approximately constant value there (Figure 4). Also p^{Wolfe} is partially allowing for undesirably short steps (§4.2.2). The expected improvement u_{EI} , on the other hand, is a well studied acquisition function of Bayesian optimization trading off exploration and exploitation. It aims to globally find a point with a function value lower than a current best guess. Though this is a desirable property also for the probabilistic line search, it is lacking the information that we are seeking a point that also fulfills the W-II curvature condition. This is evident in Figure 4 where p^{Wolfe} significantly drops at points where the objective function is already evaluated but u_{EI} does not. In addition, we do not need to explore the positive t space to an extend, the expected improvement suggests, since the aim of a line search is just to find a good, acceptable point at positive t and not the globally best one. The product of both acquisition function $u_{\text{EI}} \cdot p^{\text{Wolfe}}$ is thus a trade-off between exploring enough, but still preventing too much exploitation in obviously undesirable regions. In practice, though, we found that all three choices ((i) $u_{\text{EI}} \cdot p^{\text{Wolfe}}$, (ii) u_{EI} only, (iii) p^{Wolfe} only) perform comparable. The following experiments were all performed training N-II on MNIST; only the mini-batch size might vary as indicated.

Figure 11 compares all three choices for mini-batch size $m = 200$ and default design parameters. The top plot shows the evolution of the logarithmic test and train set error (for plot and color description see Figure caption). All test and train set error curves respectively bundle up (only lastly plotted clearly visible). The choice of acquisition function thus does not change the performance here. Rows 2-4 of Figure 11 show learning rate traces of a single seed. All three curves show very similar global behavior. First the learning rate grows, then drops again, and finally settles around the best found constant learning rate. This is intriguing since *on average* a larger learning rate seems to be better at the beginning of the optimization process, then later dropping again to a smaller one. This might also explain why SGD+PROBLS in the first part of the optimization progress outperforms vanilla SGD (Figure 7). Runs that use just slightly larger constant learning rates than the best performing constant one (above the gray horizontal lines in Figure 11) were failing after a few steps. This shows that there is some non-trivial adaptation going on, not just globally, but locally at every step.

Figure 12 shows traces of accepted learning rates for different mini-batch sizes $m = 100, 200, 1000$. Again the global behavior is qualitatively similar for all three mini-batch sizes on the given architecture. For the largest mini-batch size $m = 1000$ (last row of Figure 12) the probabilistic line search accepts a larger learning rate (on average and in absolute value) than for the smaller mini-batch sizes $m = 100$ and 200 , which is in agreement with practical experience and theoretical findings (Hinton (2012, §4 and 7), Goodfellow et al. (2016, §9.1.3), Balles et al. (2016)).

Figure 13 shows traces of the (scaled) noise levels σ_f and $\sigma_{f'}$ and the average number of function evaluations per line search for different noise levels ($m = 100, 200, 1000$; same colors show the same setup but different seeds). The average number of function evaluations rises very slightly to $\approx 1.5 - 2$ for mini-batch size $m = 1000$ towards the end of the optimization

process, in comparison to ≈ 1.5 for $m = 100, 200$. This seems counter intuitive in a way, but since larger mini-batch sizes also observe smaller value and gradients (especially towards the end of the optimization process), the relative noise levels might actually be larger. (Although the curves for varying m are shown versus the same abscissa, the corresponding optimizers might be in different regions of the loss surface, especially $m = 1000$ probably reaches regions of smaller absolute gradients). At the start of the optimization the average number of function evaluations is high, because the initial default learning rate is small (10^{-4}) and the line search extends each step multiple times.

5. Conclusion

The line search paradigm widely accepted in deterministic optimization can be extended to noisy settings. Our design combines existing principles from the noise-free case with ideas from Bayesian optimization, adapted for efficiency. We arrived at a lightweight “black-box” algorithm that exposes no parameters to the user. Empirical evaluations so far show compatibility with the SGD search direction and viability for logistic regression and multi-layer perceptrons. The line search effectively frees users from worries about the choice of a learning rate: Any reasonable initial choice will be quickly adapted and lead to close to optimal performance. Our matlab implementation can be found at <http://tinyurl.com/probLineSearch>.

Acknowledgments

Thanks to Jonas Jazzkowic who prepared the base of the pseudo-code.

Appendix A. – Noise Estimation

Section 3.4.3 introduced the statistical variance estimators

$$\begin{aligned}\Sigma'(x) &= (1-m)^{-1}(\nabla\hat{S}(x) - \nabla\hat{\mathcal{L}}(x))^{\odot 2} \\ \Sigma(x) &= (1-m)^{-1}(\hat{S}(x) - \hat{\mathcal{L}}(x))^2\end{aligned}\tag{19}$$

of the function and gradient estimate $\hat{\mathcal{L}}(x)$ and $\nabla\hat{\mathcal{L}}(x)$ at position x . The underlying assumption is that $\hat{\mathcal{L}}(x)$ and $\nabla\hat{\mathcal{L}}(x)$ are distributed according to

$$\begin{bmatrix} \hat{\mathcal{L}}(x) \\ \nabla\hat{\mathcal{L}}(x) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \hat{\mathcal{L}}(x) \\ \nabla\hat{\mathcal{L}}(x) \end{bmatrix}; \begin{bmatrix} \mathcal{L}(x) \\ \nabla\mathcal{L}(x) \end{bmatrix}, \begin{bmatrix} \Sigma(x) & 0_{D \times 1} \\ 0_{1 \times D} & \text{diag } \Sigma'(x) \end{bmatrix}\right)\tag{20}$$

which implies Eq 3

$$\begin{bmatrix} \hat{\mathcal{L}}(x) \\ s(x)' \cdot \nabla\hat{\mathcal{L}}(x) \end{bmatrix} = \begin{bmatrix} y(x) \\ y'(x) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} f(x) \\ f'(x) \end{bmatrix}, \begin{bmatrix} \sigma_f(x) & 0 \\ 0 & \sigma_{f'}(x) \end{bmatrix}\right).\tag{21}$$

where $s(x)$ is the possibly new search direction at x . This is an approximation since the true covariance matrix is in general not diagonal. A better estimator for the projected gradient noise would be (dropping x from the notation)

$$\begin{aligned}\eta_{f'} &= s^\top \left[\frac{1}{m-1} \frac{1}{m} \sum_{k=1}^m (\nabla\ell^k - \nabla\hat{\mathcal{L}})(\nabla\ell^k - \nabla\hat{\mathcal{L}})^\top \right] s \\ &= \sum_{i,j=1}^D s_i s_j \frac{1}{m-1} \frac{1}{m} \sum_{k=1}^m (\nabla\ell_i^k - \nabla\hat{\mathcal{L}}_i) (\nabla\ell_j^k - \nabla\hat{\mathcal{L}}_j) \\ &= \frac{1}{m-1} \sum_{i,j=1}^D s_i s_j \left(\frac{1}{m} \sum_{k=1}^m \nabla\ell_i^k \nabla\ell_j^k - \nabla\hat{\mathcal{L}}_i \nabla\hat{\mathcal{L}}_j - \nabla\hat{\mathcal{L}}_j \nabla\hat{\mathcal{L}}_i + \nabla\hat{\mathcal{L}}_i \nabla\hat{\mathcal{L}}_j \right) \\ &= \frac{1}{m-1} \left(\frac{1}{m} \sum_{k=1}^m \sum_{i,j=1}^D s_i \nabla\ell_i^k s_j \nabla\ell_j^k - \sum_{i,j=1}^D s_j \nabla\hat{\mathcal{L}}_j s_i \nabla\hat{\mathcal{L}}_i \right) \\ &= \frac{1}{m-1} \left(\frac{1}{m} \sum_{k=1}^m (s' \cdot \nabla\ell^k)^2 - (s' \cdot \nabla\hat{\mathcal{L}})^2 \right).\end{aligned}\tag{22}$$

Comparing to $\sigma_{f'}$ yields

$$\begin{aligned}\eta_{f'} &= \frac{1}{m-1} \sum_{i,j=1}^D s_i s_j \left(\frac{1}{m} \sum_{k=1}^m \nabla\ell_i^k \nabla\ell_j^k - \nabla\hat{\mathcal{L}}_j \nabla\hat{\mathcal{L}}_i \right) \\ &= \frac{1}{m-1} \sum_{i=1}^D s_i^2 \left(\frac{1}{m} \sum_{k=1}^m (\nabla\ell_i^k)^2 - \nabla\hat{\mathcal{L}}_i^2 \right) \\ &\quad + \frac{1}{m-1} \sum_{i \neq j=1}^D s_i s_j \left(\frac{1}{m} \sum_{k=1}^m \nabla\ell_i^k \nabla\ell_j^k - \nabla\hat{\mathcal{L}}_j \nabla\hat{\mathcal{L}}_i \right) \\ \eta_{f'} &= \sigma_{f'} + \frac{1}{m-1} \sum_{i \neq j=1}^D s_i s_j \left(\frac{1}{m} \sum_{k=1}^m \nabla\ell_i^k \nabla\ell_j^k - \nabla\hat{\mathcal{L}}_j \nabla\hat{\mathcal{L}}_i \right).\end{aligned}\tag{23}$$

From Eq 22 we see that, in order to effectively compute $\eta_{f'}$, we need an efficient way of computing the inner product $(s' \cdot \nabla \ell^k)$ for all k . In addition, we need to know the search direction $s(x)$ of the potential next step (if x was accepted) at the time of computing $\eta_{f'}$. This is possible e.g. for the SGD search direction where $s(x) = -\frac{1}{m} \sum_{k=1}^m \nabla \ell^k(x)$ but potentially not possible or practical for arbitrary search directions. For all experiments in this paper we used the approximate variance estimator $\sigma_{f'}$.

The following paragraph is concerned with the independence assumption of gradient and function value y and y' (in contrast to independence among gradient elements). In general y and y' are not independent since the algorithm draws them from the same mini-batch; the likelihood including the correlation factor ρ reads

$$p(y_t, y'_t | f) = \mathcal{N} \left(\begin{bmatrix} y_t \\ y'_t \end{bmatrix}; \begin{bmatrix} f(t) \\ f'(t) \end{bmatrix}, \begin{bmatrix} \sigma_f^2 & \rho \\ \rho & \sigma_{f'}^2 \end{bmatrix} \right). \quad (24)$$

The noise covariance matrix enters the GP only in the inverse of the sum with the kernel matrix of the observations. We can compute it analytically for one datapoint at position t , since it is only a 2×2 matrix. For $\rho = 0$, define:

$$\begin{aligned} \det_{\rho=0} &:= [k_{tt} + \sigma_f^2][\partial k_{tt}^{\partial} + \sigma_{f'}^2] - k_{tt}^{\partial} \partial k_{tt} \\ G_{\rho=0}^{-1} &:= \begin{bmatrix} k_{tt} + \sigma_f^2 & k_{tt}^{\partial} \\ \partial k_{tt} & \partial k_{tt}^{\partial} + \sigma_{f'}^2 \end{bmatrix}^{-1} = \frac{1}{\det_{\rho=0}} \begin{bmatrix} \partial k_{tt}^{\partial} + \sigma_{f'}^2 & -k_{tt}^{\partial} \\ -\partial k_{tt} & k_{tt} + \sigma_f^2 \end{bmatrix}. \end{aligned} \quad (25)$$

For $\rho \neq 0$ we thus get:

$$\begin{aligned} \det_{\rho \neq 0} &:= [k_{tt} + \sigma_f^2][\partial k_{tt}^{\partial} + \sigma_{f'}^2] - [k_{tt}^{\partial} + \rho][\partial k_{tt} + \rho] \\ &= \det_{\rho=0} - \rho(\partial k_{tt} + k_{tt}^{\partial}) - \rho^2 \\ G_{\rho \neq 0}^{-1} &:= \begin{bmatrix} k_{tt} + \sigma_f^2 & k_{tt}^{\partial} + \rho \\ \partial k_{tt} + \rho & \partial k_{tt}^{\partial} + \sigma_{f'}^2 \end{bmatrix}^{-1} = \frac{1}{\det_{\rho \neq 0}} \begin{bmatrix} \partial k_{tt}^{\partial} + \sigma_{f'}^2 & -(k_{tt}^{\partial} + \rho) \\ -(\partial k_{tt} + \rho) & k_{tt} + \sigma_f^2 \end{bmatrix} \\ &= \frac{\det_{\rho=0}}{\det_{\rho \neq 0}} G_{\rho=0}^{-1} - \frac{\rho}{\det_{\rho \neq 0}} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \end{aligned} \quad (26)$$

The fraction $\det_{\rho=0}/\det_{\rho \neq 0}$ in the first term of the last row, is a positive scalar that scales all element of $G_{\rho=0}^{-1}$ equally (since $G_{\rho=0}$ and $G_{\rho \neq 0}$ are positive definite matrices, we know that $\det_{\rho=0} > 0$, $\det_{\rho \neq 0} > 0$). If $|\rho|$ is small in comparison to the determinant $\det_{\rho=0}$, then $\det_{\rho \neq 0} \approx \det_{\rho=0}$ and the scaling factor is approximately one. The second term corrects off-diagonal elements in $G_{\rho \neq 0}$ and is proportional to ρ ; if $|\rho| \ll \det_{\rho=0}$ this term is small as well.

In might be possible to estimate ρ as well from the mini-batch in a similar style to the estimation of σ_f and $\sigma_{f'}$; it is not clear, however, if the additional computational cost would justify the improvements in the GP-inference.

Appendix B. – Noise Sensitivity

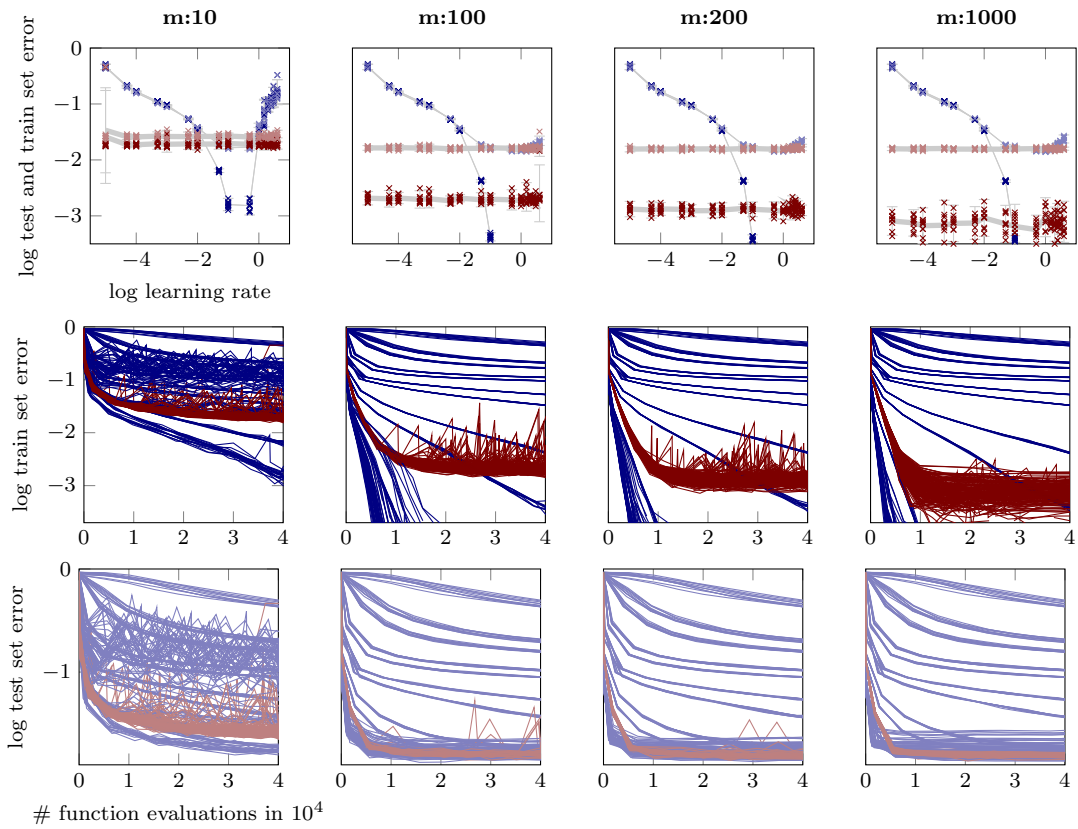


Figure 14: Performance of N-I on MNIST for *varying mini-batch sizes*; plots and colors same as in Figure 7 (middle plots cropped for readability).




Figure 15: Performance of N-II on CIFAR-10 for *varying mini-batch sizes*; plots and colors same as in Figure 7, except the scaling of the y-axis which is not logarithmic here.




Figure 16: Performance of N-I on CIFAR-10 for *varying mini-batch sizes*; plots and colors same as in Figure 15.

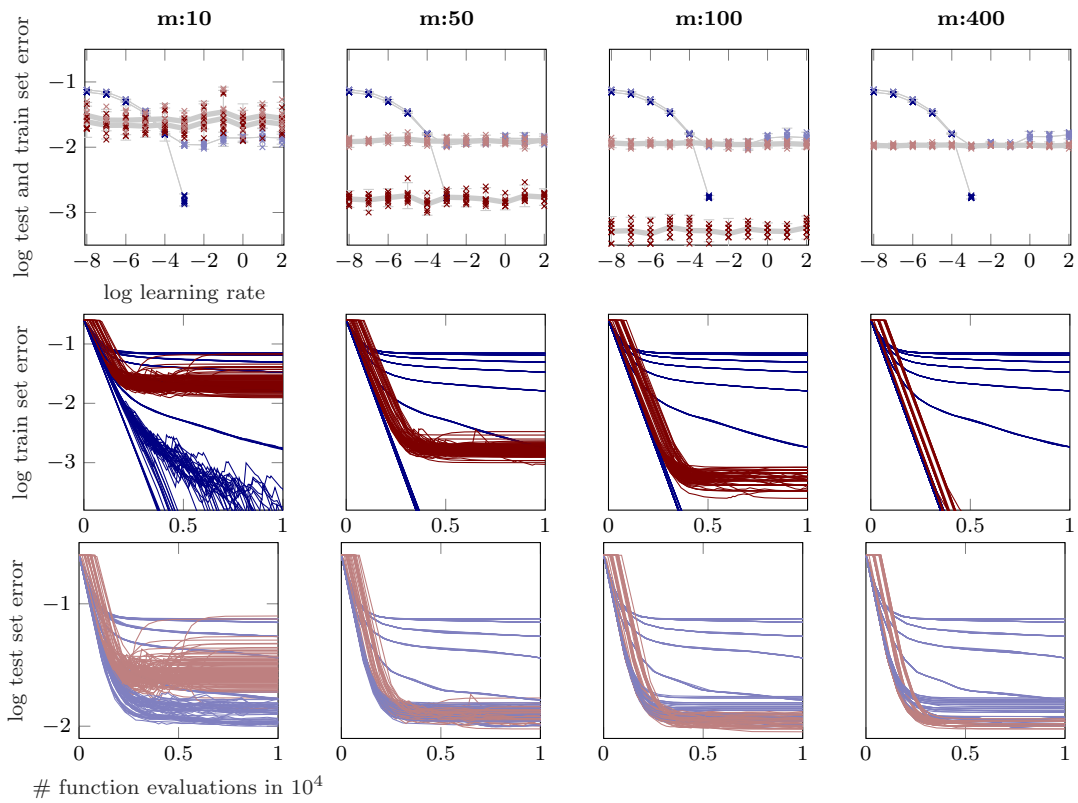


Figure 17: Performance of N-III on GISETTE for *varying mini-batch sizes*; plots and colors same as in Figure 15 (middle plots cropped for readability).




Figure 18: Performance of N-III on WDBC for *varying mini-batch sizes*; plots, colors, and description same as in Figure 17. *Remark:* since the training set is of size 400, the most right column ($m = 400$) in fact runs full-batch gradient descent; this is not a problem, since the probabilistic line search can handle noise free observations as well.

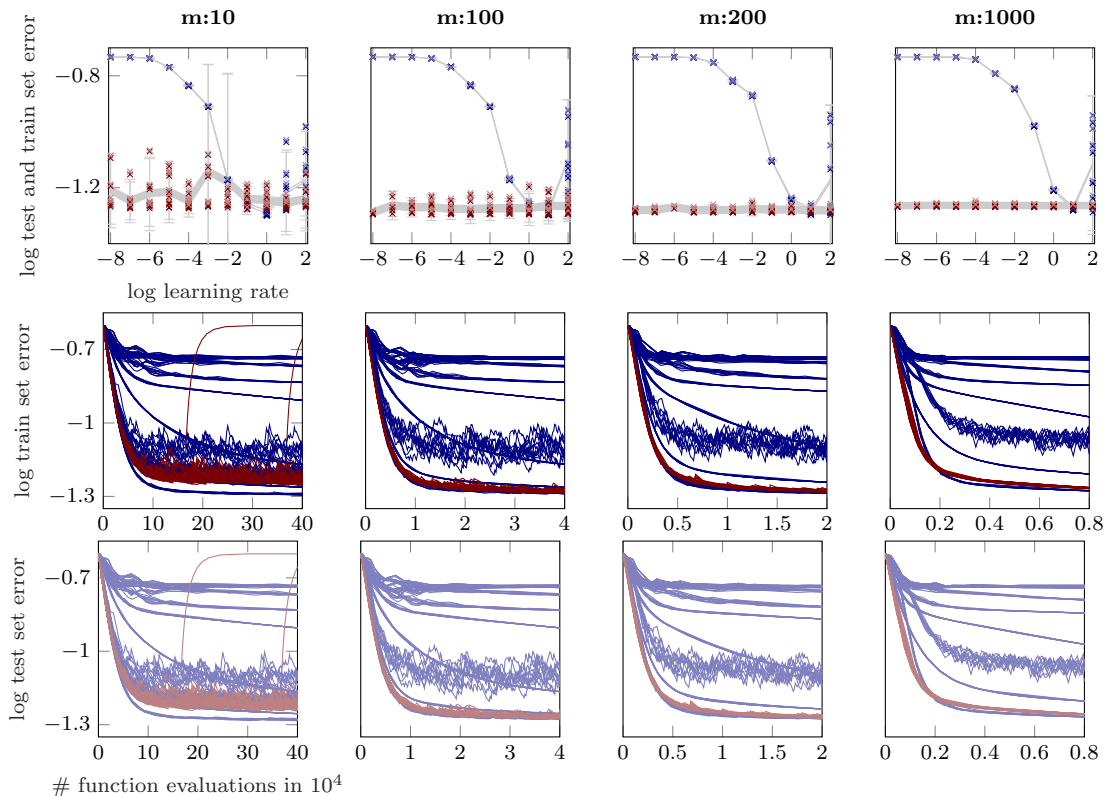


Figure 19: Performance of N-III on EPSILON for *varying mini-batch sizes*; plots, colors, and description same as in Figure 17. EPSILON is the largest dataset that was used in the experiments (400k samples); this did not seem to impair the performance of the line search or variance estimator.

Appendix C. – Parameter Sensitivity

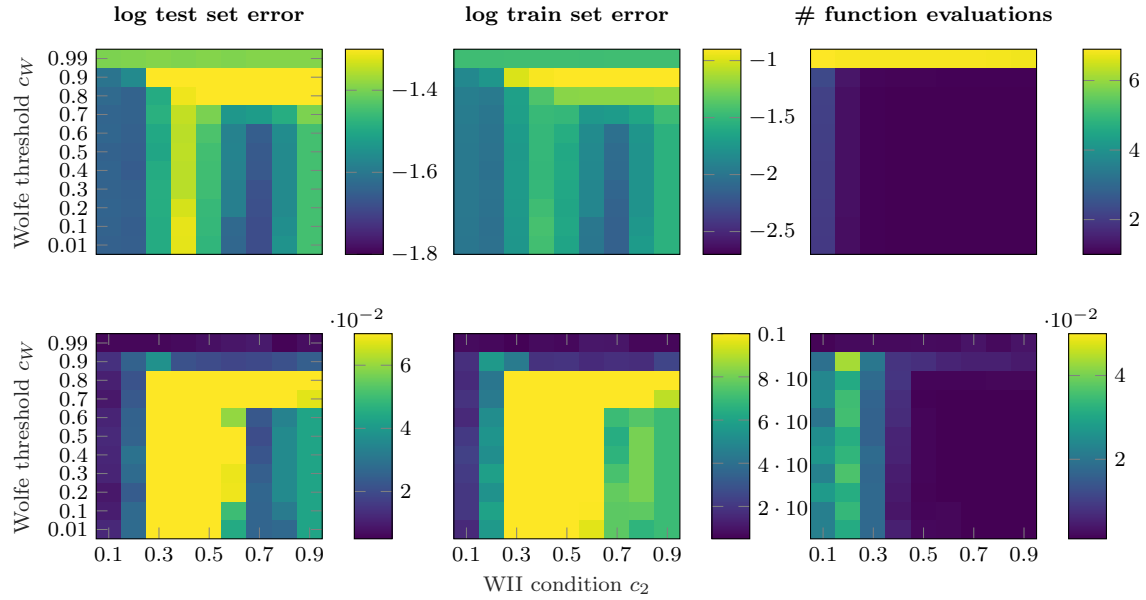


Figure 20: *Sensitivity to varying hyper-parameters c_2 , and c_W and fixed $\alpha_{\text{ext}} = 1.0$.* Experimental setup as in Figures 9 and 10. *Top row* from left to right: logarithmic test set error, train set error, and average number of function evaluations per line search averaged over 10 different initializations. *Bottom row*: corresponding relative standard deviations. In all plots darker colors are better. For extrapolation parameters $\alpha_{\text{ext}} > 1$ (see Figures 21, 22, 23, and 24) the different parameter combinations all result in similar good performance. Only at extreme choices, for example $\alpha_{\text{ext}} = 1.0$ (this figure), which amounts to no extrapolation at all in between successive line searches, the line search performs poorer. At the extreme value of $c_W = 0.99$, which amounts to imposing nearly absolute certainty about the Wolfe conditions, the line search becomes less efficient. In Figure 23 the default values adopted in the line search implementation ($c_W = 0.3$, $c_2 = 0.5$, and $\alpha_{\text{ext}} = 1.3$) are indicated as red dots.

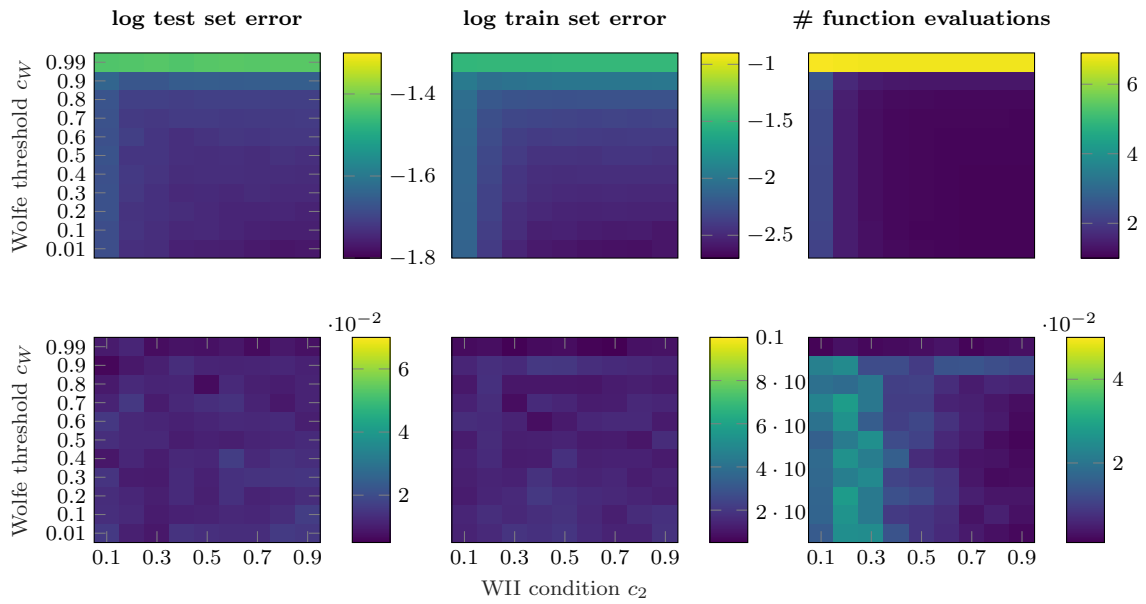


Figure 21: Same as Figure 20 but for *fixed* $\alpha_{\text{ext}} = 1.1$




Figure 22: Same as Figure 20 but for *fixed* $\alpha_{\text{ext}} = 1.2$




Figure 23: Same as Figure 20 but for *fixed* $\alpha_{\text{ext}} = 1.3$. The default values adopted in the line search implementation ($c_W = 0.3$, $c_2 = 0.5$, and $\alpha_{\text{ext}} = 1.3$) are indicated as red dots.




Figure 24: Same as Figure 20 but for *fixed* $\alpha_{\text{ext}} = 1.4$




Figure 25: *Sensitivity to varying hyper-parameters c_2 , and α_{ext} and fixed $c_W = 0.01$.* Experimental setup as in Figures 9 and 10 and plots like in Figure 20. In all plots darker colors are better. All choices of c_W result in good performance though very tight choices of $c_W = 0.99$ (Figure 35), which amounts to imposing nearly absolute certainty about the Wolfe conditions, are less efficient. As described in Figure 20, for a dropped extrapolation factor ($\alpha_{\text{ext}} \rightarrow 1$) in combination with a loose curvature condition (large c_2) the line searches performs poorer (top row, right half of columns in Figures 25–34). In Figure 28 the default values adopted in the line search implementation ($c_W = 0.3$, $c_2 = 0.5$, and $\alpha_{\text{ext}} = 1.3$) are indicated as red dots.




Figure 26: Same as Figure 25 but for *fixed* $c_W = 0.10$.

PROBABILISTIC LINE SEARCHES




Figure 27: Same as Figure 25 but for *fixed* $c_W = 0.20$.




Figure 28: Same as Figure 25 but for *fixed* $c_W = 0.30$. The default values as red dots.




Figure 29: Same as Figure 25 but for *fixed* $c_W = 0.40$.




Figure 30: Same as Figure 25 but for *fixed* $c_W = 0.50$.




Figure 31: Same as Figure 25 but for *fixed* $c_W = 0.60$.




Figure 32: Same as Figure 25 but for *fixed* $c_W = 0.70$.

PROBABILISTIC LINE SEARCHES




Figure 33: Same as Figure 25 but for *fixed* $c_W = 0.80$.




Figure 34: Same as Figure 25 but for *fixed* $c_W = 0.90$.




Figure 35: Same as Figure 25 but for *fixed* $c_W = 0.99$.

Appendix D. – Pseudocode

Algorithm 1 of Section 3 roughly sketches the structure of the probabilistic line search and its main ingredients. This section provides a detailed pseudocode which can be used for re-implementation. It is based on the code which was used for the experiments in this paper. A matlab implementation including a minimal example can be found at <http://tinyurl.com/probLineSearch>. The actual line search routine is called `PROBLINESEARCH` below and is quite short. Most of the pseudocode is occupied with comments, helper function that define the kernel of the GP, the GP-update or Gauss cdf and pdf which we printed here for completeness such that a detailed re-implementation is possible. For better readability of the pseudocode we use the following color coding:

- **blue**: comments
- **green**: variables of the integrated Wiener process.
- **red**: most recently evaluated observation (noisy loss and gradient). If the line search terminates, these will be returned as ‘accepted’.
- **orange**: inputs from the main solver procedure and unchanged during each line search.

Notation and operators:

operator or function	definition
$A \odot B$	elementwise multiplication
$A \oslash B$	elementwise division
$A^{\odot b}$	elementwise power of b
A'	transpose of A
$A \cdot B$	scalar-scalar, scalar-matrix or matrix-matrix multiplication
A/B	right matrix division, the same as $A \cdot B^{-1}$
$A \setminus B$	left matrix division, the same as $A^{-1} \cdot B$
$\text{sign}(a)$	sign of scalar a
$\text{erf}(x)$	error function $\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$
$\max(A)$	maximum element in A
$\min(A)$	minimum element in A
$ a $	absolute value of scalar a
$A < B$	elementwise ‘less’ comparison
$A \leq B$	elementwise ‘less-or-equal’ comparison
$A > B$	elementwise ‘greater’ comparison
$A \geq B$	elementwise ‘greater-or-equal’ comparison
$[a, b, c] \leftarrow f(x)$	function f called at x returns the values a, b and c

For better readability and to avoid confusion with transposes, we denote derivatives for example as dy and df (instead of y' and f' as in the main text).

```

1: function SGDSOLVER( $f$ )
2:   ▶  $f$  – function handle to objective. Usage:  $[y, dy, \Sigma_f, \Sigma_{df}] \leftarrow f(x)$ .
3:
4:   ▶ initial weights
5:    $x \leftarrow$  initial weights
6:
7:   ▶ initial step size (rather small to avoid over-shooting in very first step)
8:    $\alpha \leftarrow e.g. \approx 10^{-4}$ 
9:    $\alpha_{stats} \leftarrow \alpha$ 
10:
11:  ▶ initial function evaluation at  $x$ 
12:   $[y, dy, \Sigma_f, \Sigma_{df}] \leftarrow f(x)$ 
13:
14:  ▶ initial search direction
15:   $d \leftarrow -dy$ 
16:
17:  ▶ loop over line searches
18:  while budget not used do
19:
20:    ▶ line search finds step size
21:     $[\alpha, \alpha_{stats}, x, y, dy, \Sigma_f, \Sigma_{df}] \leftarrow \text{PROBLINESEARCH}(x, d, y, dy, \Sigma_f, \Sigma_{df}, \alpha, \alpha_{stats}, f)$ 
22:
23:    ▶ set new search direction
24:     $d \leftarrow -dy$ 
25:  end while
26:
27:  return  $x$ 
28: end function

```

```

1: function PROBLINESEARCH( $x_0, d, f_0, df_0, \Sigma_{f_0}, \Sigma_{df_0}, \alpha_0, \alpha_{stats}, f$ )
2:   ▶  $x_0$  – current weights  $[D \times 1]$ 
3:   ▶  $f$  – function handle to objective.
4:   ▶
5:   ▶  $d$  – search direction  $[D \times 1]$  (does not need to be normalized)
6:   ▶  $f_0$  – function value at start,  $f_0 = f(x_0)$ 
7:   ▶  $df_0$  – gradient at start,  $df_0 = \nabla f(x_0)$   $[D \times 1]$ 
8:   ▶  $\Sigma_{f_0}$  – sample variance of  $f_0$ 
9:   ▶  $\Sigma_{df_0}$  – sample variances of  $df_0$ ,  $[D \times 1]$ 
10:  ▶  $\alpha_0$  – initial step size
11:
12:  ▶ set maximum # of  $f$  evaluations per line search
13:   $L \leftarrow 6$ 

```

```

14:   ▶ scaling and noise level of GP
15:    $\beta \leftarrow |d' \cdot \Sigma_{df_0}|$  ▷ scale factor
16:    $\sigma_f \leftarrow \sqrt{\Sigma_{f_0} / (\alpha_0 \cdot \beta)}$  ▷ scaled sample variance of  $f_0$ 
17:    $\sigma_{df} \leftarrow \sqrt{((d^{\odot 2})' \cdot \Sigma_{df_0}) / \beta}$  ▷ scaled and projected sample variances of  $df_0$ 
18:
19:   ▶ initialize counter and non-fixed parameters
20:    $N \leftarrow 1$  ▷ size of GP =  $2 \cdot N$ 
21:    $t_{\text{ext}} \leftarrow 1$  ▷ scaled step size for extrapolation
22:    $tt \leftarrow 1$  ▷ scaled position of first function evaluation
23:
24:   ▶ initialize storage for GP. Dynamic arrays of maximum size  $[L + 1 \times 1]$ 
25:    $T \leftarrow [0]$  ▷ scaled positions along search direction
26:    $Y \leftarrow [0]$  ▷ scaled function values at  $T$ 
27:    $dY \leftarrow [(df_0' \cdot d) / \beta]$  ▷ scaled projected gradients at  $T$ 
28:
29:   ▶ initialize GP with observation at start
30:    $[G, A] \leftarrow \text{UPDATEGP}(T, Y, dY, N, \sigma_f, \sigma_{df})$ 
31:
32:   ▶ loop until budget is used or acceptable point is found
33:   for  $N$  from 2 to  $L + 1$  do
34:
35:       ▶ evaluate objective function at  $tt$ .
36:        $[y, dy, \Sigma_f, \Sigma_{df}, T, Y, dY, N] \leftarrow \text{EVALUATEOBJECTIVE}(tt, x_0, \alpha_0, d, T, Y, dY, N, \beta, f)$ 
37:
38:       ▶ update the GP which is now of size  $2 \cdot N$ .
39:        $[G, A] \leftarrow \text{UPDATEGP}(T, Y, dY, N, \sigma_f, \sigma_{df})$ 
40:
41:       ▶ initialize storage for candidates. Dynamic arrays of maximum size  $[N \times 1]$ .
42:        $T_{\text{cand}} \leftarrow []$  ▷ scaled position of candidates
43:        $M_{\text{cand}} \leftarrow []$  ▷ GP mean of candidates
44:        $S_{\text{cand}} \leftarrow []$  ▷ GP standard deviation of candidates
45:
46:       ▶ current point is above the Wolfe threshold? If yes, accept point and return.
47:       if  $\text{PROBWOLFE}(tt, T, A, G)$  then
48:            $output \leftarrow \text{RESCALEOUTPUT}(x_0, f_0, \alpha_0, d, tt, y, dy, \Sigma_f, \Sigma_{df}, \beta)$ 
49:           return  $output$ 
50:       end if
51:
52:       ▶ Wolfe conditions not satisfied at this point.
53:       ▶ find suitable candidates for next evaluation.
54:
55:       ▶ GP mean of function values and corresponding gradients at points in  $T$ .
56:        $M \leftarrow \text{map function } M(-, T, A) \text{ over } T$ 
57:        $dM \leftarrow \text{map function } D1M(-, T, A) \text{ over } T$ 

```

```

58:     ▶ candidates 1: local minima of GP mean.
59:      $T_{\text{sorted}} \leftarrow \text{sort } T \text{ in ascending order}$ 
60:      $T_{\text{Wolfes}} \leftarrow []$                                      ▷ prepare list of acceptable points
61:
62:     ▶ iterate through all  $N - 1$  cells, compute locations of local minima.
63:     for  $n$  from 1 to  $N - 1$  do
64:          $T_n \leftarrow \text{value of } T_{\text{sorted}} \text{ at } n$ 
65:          $T_{n+1} \leftarrow \text{value of } T_{\text{sorted}} \text{ at } n + 1$ 
66:
67:         ▶ add a little offset for numerical stability
68:          $t_{\text{rep}} \leftarrow T_n + 10^{-6} \cdot (T_{n+1} - T_n)$ 
69:
70:         ▶ compute location of cubic minimum in  $n^{\text{th}}$  cell
71:          $t_{\text{cubMin}} \leftarrow \text{CUBICMINIMUM}(t_{\text{rep}}, T, A, N)$ 
72:
73:         ▶ add point to candidate list if minimum lies in between  $T_n$  and  $T_{n+1}$ 
74:         if  $t_{\text{cubMin}} > T_n$  and  $t_{\text{cubMin}} < T_{n+1}$  then
75:             if (not isnanOrIsinf( $t_{\text{cubMin}}$ )) and ( $t_{\text{cubMin}} > 0$ ) then
76:                  $T_{\text{cand}} \leftarrow \text{append } t_{\text{cubMin}}$ 
77:                  $M_{\text{cand}} \leftarrow \text{append } M(t_{\text{cubMin}}, T, A)$ 
78:                  $S_{\text{cand}} \leftarrow \text{append } V(t_{\text{cubMin}}, T, G)$ 
79:             end if
80:         else
81:
82:             ▶ most likely uphill? If yes, break.
83:             if  $n = 1$  and  $D1M(0, T, A) > 0$  then
84:                  $r \leftarrow 0.01$ 
85:                  $tt \leftarrow r \cdot (T_n + T_{n+1})$ 
86:
87:                 ▶ evaluate objective function at  $tt$  and return.
88:                  $[y, dy, \Sigma_f, \Sigma_{df}, T, Y, dY, N] \leftarrow \text{EVALUATEOBJECTIVE}(tt, x_0, \alpha_0, d, T, Y, dY, N, \beta, f)$ 
89:
90:                  $output \leftarrow \text{RESCALEOUTPUT}(x_0, f_0, \alpha_0, d, tt, y, dy, \Sigma_f, \Sigma_{df}, \beta)$ 
91:                 return  $output$ 
92:             end if
93:         end if
94:
95:         ▶ check whether there is an acceptable point among the old evaluations
96:         if  $n > 1$  and  $\text{PROBWOLFE}(T_n, T, A, G)$  then
97:              $T_{\text{Wolfes}} \leftarrow \text{append } T_n$ 
98:         end if
99:     end for
100:
101:     ▶ check if acceptable points exists and return
102:     if  $T_{\text{Wolfes}}$  is not empty then

```

```

103:     ► if last evaluated point is among acceptable ones, return it.
104:     if  $tt$  in  $T_{\text{Wolfes}}$  then
105:          $output \leftarrow \text{RESCALEOUTPUT}(x_0, f_0, \alpha_0, d, tt, y, dy, \Sigma_f, \Sigma_{df}, \beta)$ 
106:         return  $output$ 
107:     end if
108:
109:     ► else, choose the one with the lowest GP mean and re-evaluate its gradient.
110:      $M_{\text{Wolfes}} \leftarrow \text{map } M(-, T, A) \text{ over } T_{\text{Wolfes}}$ 
111:      $tt \leftarrow \text{value of } T_{\text{Wolfes}} \text{ at index of } \min(M_{\text{Wolfes}})$ 
112:
113:     ► evaluate objective function at  $tt$ .
114:      $[y, dy, \Sigma_f, \Sigma_{df}, T, Y, dY, N] \leftarrow \text{EVALUATEOBJECTIVE}(tt, x_0, \alpha_0, d, T, Y, dY, N, \beta, f)$ 
115:
116:      $output \leftarrow \text{RESCALEOUTPUT}(x_0, f_0, \alpha_0, d, tt, y, dy, \Sigma_f, \Sigma_{df}, \beta)$ 
117:     return  $output$ 
118: end if
119:
120: ► candidates 2: one extrapolation step
121:  $T_{\text{cand}} \leftarrow \text{append } \max(T) + t_{\text{ext}}$ 
122:  $M_{\text{cand}} \leftarrow \text{append } M(\max(T) + t_{\text{ext}}, T, A)$ 
123:  $S_{\text{cand}} \leftarrow \text{append } V(\max(T) + t_{\text{ext}}, T, G)^{\frac{1}{2}}$ 
124:
125: ► find minimal mean among  $M$ .
126:  $\mu_{\text{EI}} \leftarrow \text{minimal value of } M$ 
127:
128: ► compute expected improvement and Wolfe probabilities at  $T_{\text{cand}}$ 
129:  $EI_{\text{cand}} \leftarrow \text{EXPECTEDIMPROVEMENT}(M_{\text{cand}}, S_{\text{cand}}, \mu_{\text{EI}})$ 
130:  $PW_{\text{cand}} \leftarrow \text{map } \text{PROBWOLFE}(-, T, A, G) \text{ over } T_{\text{cand}}$ 
131:
132: ► choose point among candidates that maximizes  $EI_{\text{cand}} \wedge PW_{\text{cand}}$ 
133:  $i_{\text{bestCand}} \leftarrow \text{index of } \max(EI_{\text{cand}} \odot PW_{\text{cand}})$ 
134:  $tt_{\text{bestCand}} \leftarrow \text{value of } T_{\text{cand}} \text{ at } i_{\text{bestCand}}$ 
135:
136: ► extend extrapolation step if necessary
137: if  $tt_{\text{bestCand}}$  is equal to  $tt + t_{\text{ext}}$  then
138:      $t_{\text{ext}} \leftarrow 2 \cdot t_{\text{ext}}$ 
139: end if
140:
141: ► set location for next evaluation
142:  $tt \leftarrow tt_{\text{bestCand}}$ 
143: end for
144:
145: ► limit reached: evaluate a final time and return the point with lowest GP mean
146:  $[y, dy, \Sigma_f, \Sigma_{df}, T, Y, dY, N] \leftarrow \text{EVALUATEOBJECTIVE}(tt, x_0, \alpha_0, d, T, Y, dY, N, \beta, f)$ 

```

```

147:   ▶ update the GP which is now of size  $2 \cdot N$ .
148:    $[G, A] \leftarrow \text{UPDATEGP}(T, Y, dY, N, \sigma_f, \sigma_{df})$ 
149:
150:   ▶ check last point for acceptance
151:   if  $\text{PROBWOLFE}(tt, T, A, G)$  then
152:      $output \leftarrow \text{RESCALEOUTPUT}(x_0, f_0, \alpha_0, d, tt, y, dy, \Sigma_f, \Sigma_{df}, \beta)$ 
153:     return  $output$ 
154:   end if
155:
156:   ▶ at the end of budget return point with the lowest GP mean
157:   ▶ compute GP means at  $T$ 
158:    $M \leftarrow \text{map } M(-, T, A)$  over  $T$ 
159:    $i_{\text{lowest}} \leftarrow$  index of minimal value in  $M$ 
160:    $t_{\text{lowest}} \leftarrow$  value of  $T$  at  $i_{\text{lowest}}$ 
161:
162:   ▶ if  $t_{\text{lowest}}$  is the last evaluated point, return
163:   if  $t_{\text{lowest}}$  is equal to  $tt$  then
164:      $output \leftarrow \text{RESCALEOUTPUT}(x_0, f_0, \alpha_0, d, tt, y, dy, \Sigma_f, \Sigma_{df}, \beta)$ 
165:     return  $output$ 
166:   end if
167:
168:   ▶ else, re-evaluate its gradient and return
169:    $tt \leftarrow$  value of  $t_{\text{lowest}}$ 
170:
171:   ▶ evaluate objective function at  $tt$ .
172:    $[y, dy, \Sigma_f, \Sigma_{df}, T, Y, dY, N] \leftarrow \text{EVALUATEOBJECTIVE}(tt, x_0, \alpha_0, d, T, Y, dY, N, \beta, f)$ 
173:
174:    $output \leftarrow \text{RESCALEOUTPUT}(x_0, f_0, \alpha_0, d, tt, y, dy, \Sigma_f, \Sigma_{df}, \beta)$ 
175:   return  $output$ 
176: end function

```

```

1: function  $\text{RESCALEOUTPUT}(x_0, f_0, \alpha_0, d, tt, y, dy, \Sigma_f, \Sigma_{df}, \beta, \alpha_{\text{stats}})$ 
2:   ▶ design parameters
3:    $\alpha_{\text{ext}} \leftarrow 1.3$  ▷ extrapolation parameter
4:    $\theta_{\text{reset}} \leftarrow 100$  ▷ reset threshold for GP scale
5:
6:   ▶ rescale accepted step size
7:    $\alpha_{\text{acc}} \leftarrow tt \cdot \alpha_0$ 
8:
9:   ▶ update weights
10:   $x_{\text{acc}} \leftarrow x_0 + \alpha_{\text{acc}} \cdot d$ 
11:
12:  ▶ rescale accepted function value
13:   $f_{\text{acc}} \leftarrow y \cdot (\alpha_0 \cdot \beta) + f_0$ 

```

```

14:   ▶ accepted gradient
15:    $df_{\text{acc}} \leftarrow dy$ 
16:
17:   ▶ sample variance of  $f_{\text{acc}}$ 
18:    $\Sigma_{f_{\text{acc}}} \leftarrow \Sigma_f$ 
19:
20:   ▶ sample variances of  $df_{\text{acc}}$ 
21:    $\Sigma_{df_{\text{acc}}} \leftarrow \Sigma_{df}$ 
22:
23:   ▶ update exponential running average of scalings
24:    $\gamma \leftarrow 0.95$ 
25:    $\alpha_{\text{stats}} \leftarrow \gamma \cdot \alpha_{\text{stats}} + (1 - \gamma) \cdot \alpha_{\text{acc}}$ 
26:
27:   ▶ next initial step size
28:    $\alpha_{\text{next}} \leftarrow \alpha_{\text{acc}} \cdot \alpha_{\text{ext}}$ 
29:
30:   ▶ if new GP scaling is drastically different than previous ones reset it.
31:   if ( $\alpha_{\text{next}} < \alpha_{\text{stats}}/\theta_{\text{reset}}$ ) or ( $\alpha_{\text{next}} > \alpha_{\text{stats}} \cdot \theta_{\text{reset}}$ ) then
32:        $\alpha_{\text{next}} \leftarrow \alpha_{\text{stats}}$ 
33:   end if
34:
35:   ▶ compressed output for readability of pseudocode
36:    $output \leftarrow [\alpha_{\text{next}}, \alpha_{\text{stats}}, x_{\text{acc}}, f_{\text{acc}}, df_{\text{acc}}, \Sigma_{f_{\text{acc}}}, \Sigma_{df_{\text{acc}}}]$ 
37:
38:   return  $output$ 
39: end function

```

```

1: function EVALUATEOBJECTIVE( $tt, x_0, \alpha_0, d, T, Y, dY, N, \beta, f$ )
2:   ▶ evaluate objective function at  $tt$ 
3:    $[y, dy, \Sigma_f, \Sigma_{df}] \leftarrow f(x_0 + tt \cdot \alpha_0 \cdot d)$ 
4:
5:   ▶ scale output
6:    $y \leftarrow (y - f_0)/(\alpha_0 \cdot \beta)$ 
7:    $dy \leftarrow (dy' \cdot d)/\beta$ 
8:
9:   ▶ storage
10:   $T \leftarrow \text{append } tt$ 
11:   $Y \leftarrow \text{append } y$ 
12:   $dY \leftarrow \text{append } dy$ 
13:   $N \leftarrow N + 1$ 
14:
15:  return  $[y, dy, \Sigma_f, \Sigma_{df}, T, Y, dY, N]$ 
16: end function

```

```

1: function CUBICMINIMUM( $t, T, A, N$ )
2:   ▶ compute necessary derivatives of GP mean at  $t$ 
3:    $d1m_t \leftarrow D1M(t, T, A)$ 
4:    $d2m_t \leftarrow D2M(t, T, A)$ 
5:    $d3m_t \leftarrow D3M(t, T, A, N)$ 
6:    $a \leftarrow 0.5 \cdot d3m_t$ 
7:    $b \leftarrow d2m_t - t \cdot d3m_t$ 
8:    $c \leftarrow d1m_t - d2m_t \cdot t + 0.5 \cdot d3m_t \cdot t^2$ 
9:
10:  ▶ third derivative is almost zero  $\rightarrow$  essentially a quadratic, single extremum
11:  if  $|d3m_t| < 1^{-9}$  then
12:     $t_{\text{cubMin}} \leftarrow -(d1m_t - t \cdot d2m_t)/d2m_t$ 
13:    return  $t_{\text{cubMin}}$ 
14:  end if
15:
16:  ▶ roots are complex, no extremum
17:   $\lambda \leftarrow b^2 - 4 \cdot a \cdot c$ 
18:  if  $\lambda < 0$  then
19:     $t_{\text{cubMin}} \leftarrow +\infty$ 
20:    return  $t_{\text{cubMin}}$ 
21:  end if
22:
23:  ▶ compute the two possible roots
24:   $LR \leftarrow (-b - \text{sign}(a) \cdot \sqrt{\lambda})/(2 \cdot a)$  ▷ left root
25:   $RR \leftarrow (-b + \text{sign}(a) \cdot \sqrt{\lambda})/(2 \cdot a)$  ▷ right root
26:
27:  ▶ calculate the two values of the cubic at those points (up to a constant)
28:   $dt_L \leftarrow LR - t$  ▷ distance to left root
29:   $dt_R \leftarrow RR - t$  ▷ distance to right root
30:   $CV_L \leftarrow d1m_t \cdot dt_L + 0.5 \cdot d2m_t \cdot dt_L^2 + (d3m_t \cdot dt_L^3)/6$  ▷ left cubic value
31:   $CV_R \leftarrow d1m_t \cdot dt_R + 0.5 \cdot d2m_t \cdot dt_R^2 + (d3m_t \cdot dt_R^3)/6$  ▷ right cubic value
32:
33:  ▶ find the minimum and return it.
34:  if  $CV_L < CV_R$  then
35:     $t_{\text{cubMin}} \leftarrow LR$ 
36:  else
37:     $t_{\text{cubMin}} \leftarrow RR$ 
38:  end if
39:
40:  return  $t_{\text{cubMin}}$ 
41:
42: end function

```

```

1: function UPDATEGP( $T, Y, dY, N, \sigma_f, \sigma_{df}$ )
2:   ► initialize kernel matrices
3:    $k_{TT} \leftarrow [N \times N]$  matrix with zeros           ▷ covariance of function values
4:    $kd_{TT} \leftarrow [N \times N]$  matrix with zeros       ▷ covariance of function values and gradients
5:    $dkd_{TT} \leftarrow [N \times N]$  matrix with zeros      ▷ covariance of gradients
6:
7:   ► fill kernel matrices
8:   for  $i = 1$  to  $N$  do
9:     for  $j = 1$  to  $N$  do
10:       $k_{TT}(i, j) \leftarrow K(T(i), T(j))$ 
11:       $kd_{TT}(i, j) \leftarrow KD(T(i), T(j))$ 
12:       $dkd_{TT}(i, j) \leftarrow DKD(T(i), T(j))$ 
13:     end for
14:   end for
15:
16:   ► build diagonal covariance matrix of Gaussian likelihood  $[2N \times 2N]$ .
17:    $\Lambda \leftarrow \begin{bmatrix} \text{diag}(\sigma_f^2)_{N \times N} & 0_{N \times N} \\ 0_{N \times N} & \text{diag}(\sigma_{df}^2)_{N \times N} \end{bmatrix}$ 
18:
19:    $G \leftarrow \begin{pmatrix} k_{TT} & kd_{TT} \\ kd'_{TT} & dkd_{TT} \end{pmatrix} + \Lambda$            ▷  $[2N \times 2N]$  matrix
20:
21:   ► residual between observed and predicted data
22:    $\Delta \leftarrow \begin{pmatrix} Y \\ dY \end{pmatrix}$            ▷  $[2N \times 1]$  vector
23:
24:   ► compute weighted observations  $A$ .
25:    $A \leftarrow G \backslash \Delta$            ▷  $[2N \times 1]$  vector
26:
27:   return  $[G, A]$ 
28:
29: end function

```

```

1: function M( $t, T, A$ )
2:   ► posterior mean at  $t$ 
3:   return [K( $t, T'$ ), KD( $t, T'$ )] ·  $A$ 
4: end function
5:
6: function D1M( $t, T, A$ )
7:   ► first derivative of mean at  $t$ 
8:   return [DK( $t, T'$ ), DKD( $t, T'$ )] ·  $A$ 
9: end function
10:
11: function D2M( $t, T, A$ )
12:   ► second derivative of mean at  $t$ 
13:   return [DDK( $t, T'$ ), DDKD( $t, T'$ )] ·  $A$ 
14: end function
15:
16: function D3M( $t, T, A, N$ )
17:   ► third derivative of mean at  $t$ 
18:   return [DDDK( $t, T'$ ), ZEROS(1, N)] ·  $A$ 
19: end function
20:
21: function V( $t, T, G$ )
22:   ► posterior variance of function values at  $t$ 
23:   return K( $t, t$ ) - [K( $t, T'$ ), KD( $t, T'$ )] · ( $G \setminus$  [K( $t, T'$ ), KD( $t, T'$ )]')
24: end function
25:
26: function VD( $t, T, G$ )
27:   ► posterior variance of function values and derivatives at  $t$ 
28:   return KD( $t, t$ ) - [K( $t, T'$ ), KD( $t, T'$ )] · ( $G \setminus$  [DK( $t, T'$ ), DKD( $t, T'$ )]')
29: end function
30:
31: function DVD( $t, T, G$ )
32:   ► posterior variance of derivatives at  $t$ 
33:   return DKD( $t, t$ ) - [DK( $t, T'$ ), DKD( $t, T'$ )] · ( $G \setminus$  [DK( $t, T'$ ), DKD( $t, T'$ )]')
34: end function
35:
36: function V0F( $t, T, G$ )
37:   ► posterior covariances of function values at  $t = 0$  and  $t$ 
38:   return K(0,  $t$ ) - [K(0,  $T'$ ), KD(0,  $T'$ )] · ( $G \setminus$  [K( $t, T'$ ), KD( $t, T'$ )]')
39: end function
40:
41: function VD0F( $t, T, G$ )
42:   ► posterior covariance of gradient and function value at  $t = 0$  and  $t$  respectively
43:   return DK(0,  $t$ ) - [DK(0,  $T'$ ), DKD(0,  $T'$ )] · ( $G \setminus$  [K( $t, T'$ ), KD( $t, T'$ )]')
44: end function

```

```

45: function V0DF( $t, T, G$ )
46:     ► posterior covariance of function value and gradient at  $t = 0$  and  $t$  respectively
47:     return  $\text{KD}(0, t) - [\text{K}(0, T'), \text{KD}(0, T')] \cdot (G \setminus [\text{DK}(t, T'), \text{DKD}(t, T')])'$ 
48: end function
49:
50: function VD0DF( $t, T, G$ )
51:     ► same as V0F(-) but for gradients
52:     return  $\text{DKD}(0, t) - [\text{DK}(0, T'), \text{DKD}(0, T')] \cdot (G \setminus [\text{DK}(t, T'), \text{DKD}(t, T')])'$ 
53: end function

```

```

1: -----
2: ► all following procedures use the same design parameter:
3:  $\tau \leftarrow 10$ 
4: -----
5: function K( $a, b$ )
6:     ► Wiener kernel integrated once in each argument
7:     return  $1/3 \odot \min(a + \tau, b + \tau)^{\odot 3} + 0.5 \odot |a - b| \odot \min(a + \tau, b + \tau)^{\odot 2}$ 
8: end function
9:
10: function KD( $a, b$ )
11:     ► Wiener kernel integrated in first argument
12:     return  $0.5 \odot (a < b) \odot (a + \tau)^{\odot 2} + (a \geq b) \odot ((a + \tau) \cdot (b + \tau) - 0.5 \odot (b + \tau)^{\odot 2})$ 
13: end function
14:
15: function DK( $a, b$ )
16:     ► Wiener kernel integrated in second argument
17:     return  $0.5 \odot (a > b) \odot (b + \tau)^{\odot 2} + (a \leq b) \odot ((a + \tau) \cdot (b + \tau) - 0.5 \odot (a + \tau)^{\odot 2})$ 
18: end function
19:
20: function DKD( $a, b$ )
21:     ► Wiener kernel
22:     return  $\min(a + \tau, b + \tau)$ 
23: end function
24:
25: function DDK( $a, b$ )
26:     ► Wiener kernel integrated in second argument and 1x derived in first argument
27:     return  $(a \leq b) \odot (b - a)$ 
28: end function
29:
30: function DDKD( $a, b$ )
31:     ► Wiener kernel 1x derived in first argument
32:     return  $(a \leq b)$ 
33: end function

```

```

34: function DDDK( $a, b$ )
35:   ► Wiener kernel 2x derived in first argument and integrated in second argument
36:   return  $-(a \leq b)$ 
37: end function

```

```

1: function PROBWOLFE( $t, T, A, G$ )
2:   ► design parameters
3:    $c_1 \leftarrow 0.05$  ▷ constant for Armijo condition
4:    $c_2 \leftarrow 0.5$  ▷ constant for curvature condition
5:    $c_W \leftarrow 0.3$  ▷ threshold for Wolfe probability
6:
7:   ► mean and covariance values at start position ( $t = 0$ )
8:    $m_0 \leftarrow M(0, T, A)$ 
9:    $dm_0 \leftarrow D1M(0, T, A)$ 
10:   $V_0 \leftarrow V(0, T, G)$ 
11:   $Vd_0 \leftarrow VD(0, T, G)$ 
12:   $dVd_0 \leftarrow DVD(0, T, G)$ 
13:
14:  ► marginal mean and variance for Armijo condition
15:   $m_a \leftarrow m_0 - M(t, T, A) + c_1 \cdot t \cdot dm_0$ 
16:   $V_{aa} \leftarrow V_0 + (c_1 \cdot t)^2 \cdot dVd_0 + V(t) + 2 \cdot (c_1 \cdot t \cdot (Vd_0 - VD0F(t)) - V0F(t))$ 
17:
18:  ► marginal mean and variance for curvature condition
19:   $m_b \leftarrow D1M(t) - c_2 \cdot dm_0$ 
20:   $V_{bb} \leftarrow c_2^2 \cdot dVd_0 - 2 \cdot c_2 \cdot VD0DF(t) + DVD(t)$ 
21:
22:  ► covariance between conditions
23:   $V_{ab} \leftarrow -c_2 \cdot (Vd_0 + c_1 \cdot t \cdot dVd_0) + c_2 \cdot VD0F(t) + V0DF(t) + c_1 \cdot t \cdot VD0DF(t) - VD(t)$ 
24:
25:  ► extremely small variances → very certain (deterministic evaluation)
26:  if  $V_{aa} \leq 10^{-9}$  and  $V_{bb} \leq 10^{-9}$  then
27:     $p_{Wolfe} \leftarrow (m_a \geq 0) \cdot (m_b \geq 0)$ 
28:
29:    ► accept?
30:     $p_{acc} \leftarrow p_{Wolfe} > c_W$ 
31:    return  $p_{acc}$ 
32:  end if
33:
34:  ► zero or negative variances (maybe something went wrong?)
35:  if  $V_{aa} \leq 0$  or  $V_{bb} \leq 0$  then
36:    return 0
37:  end if

```

```

38:   ▶ noisy case (everything is alright)
39:   ▶ correlation
40:    $\rho \leftarrow V_{ab} / \sqrt{V_{aa} \cdot V_{bb}}$ 
41:
42:   ▶ lower and upper integral limits for Armijo condition
43:    $low_a \leftarrow -m_a / \sqrt{V_{aa}}$ 
44:    $up_a \leftarrow +\infty$ 
45:
46:   ▶ lower and upper integral limits for curvature condition
47:    $low_b \leftarrow -m_b / \sqrt{V_{bb}}$ 
48:    $up_b \leftarrow (2 \cdot c_2 \cdot (|dm_0| + 2 \cdot \sqrt{dVd_0}) - m_b) / \sqrt{V_{bb}}$ 
49:
50:   ▶ compute Wolfe probability
51:    $p_{Wolfe} \leftarrow \text{BVN}(low_a, up_a, low_b, up_b, \rho)$ 
52:
53:   ▶ accept?
54:    $p_{acc} \leftarrow p_{Wolfe} > c_W$ 
55:   return  $p_{acc}$ 
56:
57:   

58:       The function  $\text{BVN}(low_a, up_a, low_b, up_b, \rho)$  evaluates the 2D-integral
59:       
$$\int_{low_a}^{up_a} \int_{low_b}^{up_b} \mathcal{N} \left( \begin{bmatrix} a \\ b \end{bmatrix}; \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \right) da db.$$

60:


61:
62: end function

```

```

1: function GAUSSCDF( $z$ )
2:   ▶ Gauss cumulative density function
3:   return  $0.5 \odot (1 + \text{erf}(z / \sqrt{2}))$ 
4: end function
5:
6: function GAUSSPDF( $z$ )
7:   ▶ Gauss probability density function
8:   return  $\exp(-0.5 \odot z^{\odot 2}) \oslash \sqrt{2\pi}$ 
9: end function
10:
11: function EXPECTEDIMPROVEMENT( $m, s, \eta$ )
12:   ▶ Jones et al. (1998)
13:   return  $(\eta - m) \odot \text{GAUSSCDF}((\eta - m) \oslash s) + s \odot \text{GAUSSPDF}((\eta - m) \oslash s)$ 
14: end function

```

References

- R.J. Adler. *The Geometry of Random Fields*. Wiley, 1981.
- S.-I. Amari, H. Park, and K. Fukumizu. Adaptive method of realizing natural gradient learning for multilayer perceptrons. *Neural Computation*, 12(6):1399–1409, 2000.
- M. Andrychowicz, M. Denil, S. Gomez-Colmenarejo, M.W. Hoffman, D. Pfau, T. Schaul, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. *CoRR*, abs/1606.04474, 2016. URL <http://arxiv.org/abs/1606.04474>.
- L. Armijo. Minimization of functions having Lipschitz continuous first partial derivatives. *Pacific Journal of Mathematics*, 16(1):1–3, 1966.
- L. Balles, J. Romero, and P. Hennig. Coupling Adaptive Batch Sizes with Learning Rates. *ArXiv e-prints*, December 2016.
- L. Balles, M. Mahsereci, and P. Hennig. Automating stochastic optimization with gradient variance estimates. *ICML AutoML Workshop*, 2017.
- L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of the 19th Int. Conf. on Computational Statistics (COMPSTAT)*, pages 177–186. Springer, 2010.
- T. Broderick, N. Boyd, A. Wibisono, A.C. Wilson, and M.I. Jordan. Streaming variational Bayes. In *Advances in Neural Information Processing Systems (NIPS 26)*, pages 1727–1735, 2013.
- C.G. Broyden. A new double-rank minimization algorithm. *Notices of the AMS*, 16(4):670, 1969.
- Z. Drezner and G.O. Wesolowsky. On the computation of the bivariate normal integral. *Journal of Statistical Computation and Simulation*, 35(1-2):101–107, 1990.
- J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.
- R. Fletcher. A new approach to variable metric algorithms. *The Computer Journal*, 13(3):317, 1970.
- R. Fletcher and C.M. Reeves. Function minimization by conjugate gradients. *The Computer Journal*, 7(2):149–154, 1964.
- A.P. George and W.B. Powell. Adaptive stepsizes for recursive estimation with applications in approximate dynamic programming. *Machine Learning*, 65(1):167–198, 2006.
- D. Goldfarb. A family of variable metric updates derived by variational means. *Math. Comp.*, 24(109):23–26, 1970.
- I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- I. Guyon, S. Gunn, A. Ben-Hur, and G.n Dror. Result analysis of the nips 2003 feature selection challenge. In L. K. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 545–552. MIT Press, 2005.
- E. Hairer, S.P. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations I – Nonstiff Problems*. Springer, 1987.
- S. Hansen. Using deep q-learning to control optimization hyperparameters. *CoRR*, abs/1602.04062, 2016. URL <http://arxiv.org/abs/1602.04062>.

- P. Hennig. Fast Probabilistic Optimization from Noisy Gradients. In *30th International Conference on Machine Learning (ICML)*, 2013.
- J. Hensman, M. Rattray, and N.D. Lawrence. Fast variational inference in the conjugate exponential family. In *Advances in Neural Information Processing Systems (NIPS 25)*, pages 2888–2896, 2012.
- G.E. Hinton. *A Practical Guide to Training Restricted Boltzmann Machines*, pages 599–619. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012. ISBN 978-3-642-35289-8. doi: 10.1007/978-3-642-35289-8_32. URL http://dx.doi.org/10.1007/978-3-642-35289-8_32.
- M.D. Hoffman, D.M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(1):1303–1347, 2013.
- D.R. Jones, M. Schonlau, and W.J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998.
- D.P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL <http://arxiv.org/abs/1412.6980>.
- A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. 2009. URL <https://www.cs.toronto.edu/~kriz/cifar.html>.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- K. Li and J. Malik. Learning to optimize. *CoRR*, abs/1606.01885, 2016. URL <http://arxiv.org/abs/1606.01885>.
- M. Mahsereci and P. Hennig. Probabilistic line searches for stochastic optimization. In *Advances in Neural Information Processing Systems 28*, pages 181–189, 2015.
- J. Martens. Deep learning via Hessian-free optimization. In *International Conference on Machine Learning (ICML)*, 2010.
- J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer Verlag, 1999.
- A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, New York, 3rd ed. edition, 1991.
- R. Rajesh, W. Chong, D. Blei, and E. Xing. An adaptive learning rate for stochastic variational inference. In *30th International Conference on Machine Learning (ICML)*, pages 298–306, 2013.
- C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. MIT, 2006.
- H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22(3):400–407, Sep. 1951.
- N.L. Roux and A.W. Fitzgibbon. A fast natural Newton method. In *27th International Conference on Machine Learning (ICML)*, pages 623–630, 2010.
- S. Särkkä. *Bayesian filtering and smoothing*. Cambridge University Press, 2013.
- T. Schaul, S. Zhang, and Y. LeCun. No more pesky learning rates. In *30th International Conference on Machine Learning (ICML-13)*, pages 343–351, 2013.
- N.N. Schraudolph. Local gain adaptation in stochastic gradient descent. In *Ninth International Conference on Artificial Neural Networks (ICANN) 99*, volume 2, pages 569–574, 1999.

- D.F. Shanno. Conditioning of quasi-Newton methods for function minimization. *Math. Comp.*, 24 (111):647–656, 1970.
- N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design. In *International Conference on Machine Learning (ICML)*, 2010.
- I. Sutskever, J. Martens, G. Dahl, and G. Hinton. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, volume 28. JMLR Workshop and Conference Proceedings, 2013.
- G. Wahba. *Spline models for observational data*. Number 59 in CBMS-NSF Regional Conferences series in applied mathematics. SIAM, 1990.
- W.H. Wolberg, W.N Street, and O.L. Mangasarian. *UCI Machine Learning Repository: Breast Cancer Wisconsin (Diagnostic) Data Set*, January 2011. URL [http://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Diagnostic\)](http://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Diagnostic)).
- P. Wolfe. Convergence conditions for ascent methods. *SIAM Review*, pages 226–235, 1969.
- T. Zhang. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Twenty-first International Conference on Machine Learning (ICML 2004)*, 2004.