
Learning-to-Optimize with PAC-Bayesian Guarantees: Theoretical Considerations and Practical Implementation

Michael Sucker

*Department of Mathematics
University of Tübingen
Tübingen, Germany*

michael.sucker@math.uni-tuebingen.de

Jalal Fadili

*ENSICAEN
Normandie Université
CNRS, GREYC, France*

jalal.fadili@ensicaen.fr

Peter Ochs

*Department of Mathematics and Computer Science
Saarland University
Saarbrücken, Germany*

ochs@cs.uni-saarland.de

Abstract

We use the PAC-Bayesian theory for the setting of learning-to-optimize. To the best of our knowledge, we present the first framework to learn optimization algorithms with provable generalization guarantees (PAC-Bayesian bounds) and explicit trade-off between convergence guarantees and convergence speed, which contrasts with the typical worst-case analysis. Our learned optimization algorithms provably outperform related ones derived from a (deterministic) worst-case analysis. The results rely on PAC-Bayesian bounds for general, possibly unbounded loss-functions based on exponential families. Then, we reformulate the learning procedure into a one-dimensional minimization problem and study the possibility to find a global minimum. Furthermore, we provide a concrete algorithmic realization of the framework and new methodologies for learning-to-optimize, and we conduct four practically relevant experiments to support our theory. With this, we showcase that the provided learning framework yields optimization algorithms that provably outperform the state-of-the-art by orders of magnitude.

1 Introduction

Typically, optimization algorithms are derived by considering a *class of problems*, which consists of three ingredients: First, the *model*, which summarizes the imposed assumptions and common properties of the problems within that class. Second, the *oracle*, which describes the information the algorithm actually receives during run-time. And third, the *stopping criterion*, that specifies when the algorithm terminates. For example, consider all convex and continuously differentiable functions with Lipschitz continuous gradient (model), a first-order method that has access to the function value and its gradient in each iteration (oracle), with the overall goal to decrease the residual in function values up to a certain threshold (stopping-criterion).

Based on this class, one can perform a worst-case analysis to make sure that *in any case* a sufficient descent property, for example in function values, does hold along the iterations, which allows to prove convergence and derive corresponding convergence rates. Thus, such a worst-case analysis allows for theoretical guarantees by being applicable to any case in question. However, often this comes at the cost of slower convergence in the *typical case*, as considering the worst-case only is *inevitably* accompanied by a negligence of information and thus, an overly conservative estimation in most of the cases. Hence, there is an inherent trade-off between applicability in terms of imposed assumptions, theoretical guarantees and practical performance. A worst-case analysis aims for one of the extremes: Strong theoretical guarantees at the cost of potentially strong assumptions and practical performance. Doing so, we get an *uniform upper bound* on the performance. Yet, if the worst-case is actually equally likely as any other problem in that class, this indeed is the best one can hope for, as there are simple examples for which this upper bound then also matches the best possible lower bound, yielding the notion of an *optimal* algorithm. Therefore, to improve the performance of an optimization algorithm, one needs to have access to more information about the problems at hand. This can be done *analytically*, that is, by imposing further restrictive assumptions, or *statistically* under the assumption that the single problem instances are not equally likely to be observed. Here, the first approach boils down to reducing the class of functions in question, hence a change of the model. That means that one might get a faster algorithm, yet at the expense of less applicability. Therefore, it does not alter the underlying trade-off effectively. The second approach, however, rather aims for a change of the oracle and tries to use structure in the given class of functions, which is *explicitly* not analytically tractable, that is, not already encoded in the model. This means that, if there is actually more structure in the class of functions than the model describes analytically (the model is *larger* than the problem class at hand), fast algorithms can be obtained without making further restrictive assumptions. The premise for this is that one actually manages to access and use this structure in other ways than by pen and paper, for example, statistically. One particular way to do this is to *learn* the optimization algorithm on a given data set. For this, the oracle gets enriched by access to a data set that can be used for a training phase. By doing so, the algorithm becomes *data-dependent* and can in turn enrich its oracle by the statistical structure of the class of problems. Hence, learning actually alleviates the bounds of analytical tractability by introducing data-dependency. Furthermore, it also allows for new algorithms, which cannot be analyzed (or even derived) analytically and therefore were not tractable before. The whole process can be automated, such that, in total, one can access more information and train more complicated algorithms with actually less hand-crafting. Hence, there is a huge potential up-side to learning optimization algorithms. However, there is no free-lunch. Indeed, there is also a down-side: Since one insists on using structure and algorithms that are explicitly not analytically tractable, one loses some theoretical guarantees that were built upon exactly these analytic features. Hence, while learning alleviates the bounds of analytical tractability, it does it at the cost of losing the *typical* theoretical guarantees. Since the practical usefulness of an optimization algorithm without convergence guarantees is at least questionable, this is a major problem, and poses the first central question:

What kind of theoretical guarantees can be given for such a learned optimization algorithm? Are we able to *ensure* its usefulness?

Hence, in the first part of this paper we provide a theoretical framework for learning an optimization algorithm based on its performance on the training set, that is, based on the empirical risk together with a generalization bound for the (true) risk. A popular framework that provides such generalization bounds is the PAC-Bayesian approach to learning, which we apply to the setting of learning-to-optimize. This yields the following informal version of our main theoretical result (compare Theorem 4.1 together with Corollary 3.5):

Theorem 1.1 (Informal) *Under mild boundedness assumptions on the optimization algorithm, the \mathbb{Q} -average population loss \mathcal{R}_g of the last iterate can be bounded by the \mathbb{Q} -average empirical loss $\hat{\mathcal{R}}_g$ of the last iterate plus some remainder term r_N that vanishes with the size N of the data set, that is, for all $\varepsilon > 0$:*

$$\mathbb{P}_{\mathfrak{D}} \left\{ \forall \lambda \in \Lambda, \forall \mathbb{Q} \in \mathcal{P} : \mathbb{Q}[\mathcal{R}_g] \leq \mathbb{Q}[\hat{\mathcal{R}}_g] + r_N \right\} \geq 1 - \varepsilon.$$

This result states that it is unlikely to observe a data set for which the given bound on the risk of the last iterate would not hold uniformly in λ or \mathbb{Q} . Especially, the uniformity in \mathbb{Q} allows for *learning* such a distribution \mathbb{Q} . To prove this result, in Theorem 3.3 we provide a general PAC-Bayesian theorem for data-dependent exponential families, which holds under mild assumptions, thus being widely applicable. Then, in Corollary 3.5, by specifying the natural parameters and the sufficient statistics, that is, by choosing a specific exponential family, we directly deduce a PAC-Bayesian generalization bound. Here, since these two results are widely applicable and constructive in nature, they might be of independent interest. Finally, to be able to apply Theorem 3.3 and Corollary 3.5 to the setting of learning-to-optimize we identified suitable properties of optimization algorithms, such that the corresponding assumptions are satisfied (compare Theorem 4.1 and Theorem 4.9).

Taken together, this provides one possible answer to the central question posed above about the theoretical guarantees. However, while being a generalization bound, these guarantees are a statement about *relative values*, that is, how the true risk compares to the empirical risk. Though, they do not directly imply that these have to be small *absolute values*. Thus, one still has to train the optimization algorithm in such a way that the empirical risk is indeed small enough to be worth the effort. Hence, the second central question that arises is of a more practical nature and pertains to the actual training of such an algorithm:

How do we learn an optimization algorithm, so that its performance is *clearly* superior to the one achieved by a worst-case analysis?

Therefore, in the second part of this work, we develop a concrete algorithmic realization, which allows for learning an optimization algorithm and evaluating the corresponding theoretical guarantee. This involves several design choices and tricks that have not been used before and which are of interest in their own right. Furthermore, as empirical evaluation of our claims, we conduct four practically relevant experiments, all dealing with very different classes of functions, thereby demonstrating the wide applicability and strong practical performance of our approach. Examples are shown in Figure 1. Each subplot compares the performance of the learned algorithm to that of a standard algorithm on different problems: a smooth and strongly convex quadratic toy problem (upper left), a high-dimensional, smooth and convex image processing problem (upper right), a non-smooth and convex

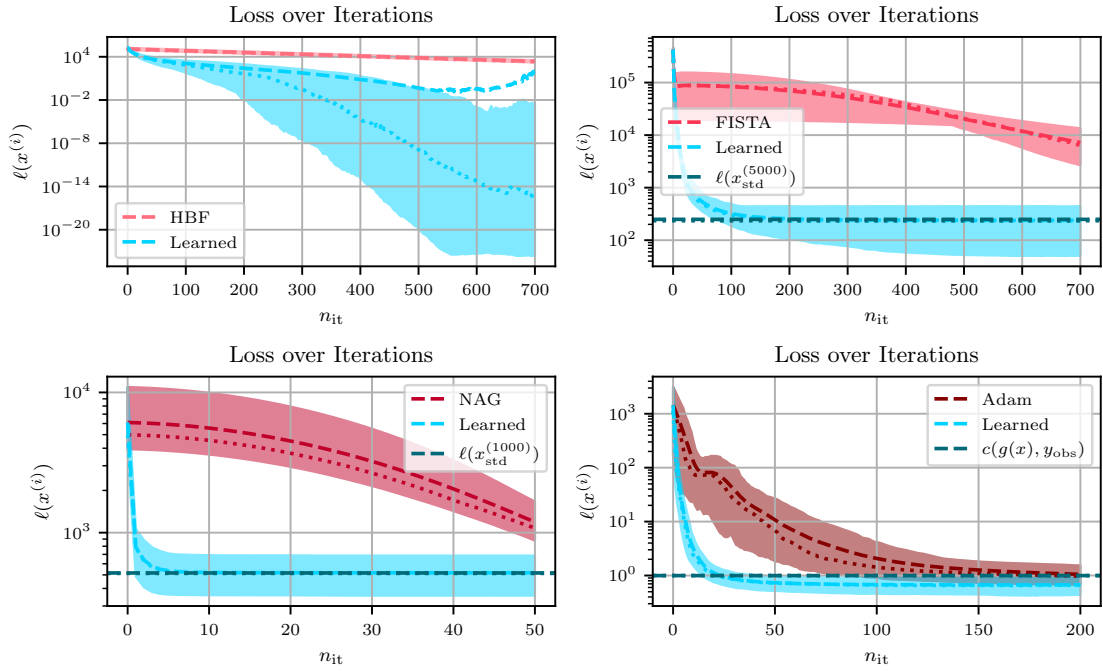


Figure 1: Some numerical results: Loss over iterations (mean as dashed and median as dotted line) of the learned algorithm compared to a standard choice.

LASSO problem (lower left), and a non-smooth and non-convex problem of training a neural network (lower right). Since the learned algorithm is clearly superior in each case, this provides a possible answer to the question about how to train optimization algorithms. Taken together, we provide a complete framework to train optimization algorithms with theoretical guarantees that are (in a certain sense) provably faster than their worst-case optimal counterparts. With this, we get closer to the ideal case of wide applicability, strong practical performance and sufficient theoretical guarantees. In particular, this work is a far reaching extension of our conference paper (Sucker and Ochs, 2023) by extending and clarifying the theoretical results in Sections 3, and by the algorithmic realization together with its evaluation in Sections 5 and 6, which additionally includes a probabilistic constraining procedure for sampling algorithms in Sections 5.2.

1.1 Related Work

The literature on both learning-to-optimize and the PAC-Bayes learning approach is vast. Hence, the discussion of learning-to-optimize will mainly focus on learning procedures for optimization algorithms and on approaches that provide some theoretical guarantees. Especially, this excludes many model-free approaches, which replace the whole update step with a learnable mapping such as a neural network. Chen et al. (2021) provide a good overview about the variety of approaches in learning-to-optimize, and good introductory references for the PAC-Bayesian approach are given by Guedj (2019) and Alquier (2021).

Broader Context of Learning-to-Optimize. Solving optimization problems is an integral part of machine learning. Thus, learning-to-optimize has significant overlap with the areas of meta-learning (or “learning-to-learn”) and AutoML. The first one is a subset of learning-to-optimize, as, while learning-to-optimize applies to general optimization problems, meta-learning is mostly concerned with determining parameters of machine learning models, that is, minimizing a training loss (Vilalta and Drissi, 2002; Hospedales et al., 2021). AutoML, however, more broadly refers to automating all steps necessary to create a machine learning application, which thus also involves the choice of an optimization algorithm and its hyperparameters (Yao et al., 2018; Hutter et al., 2019; He et al., 2021). Hence, most relevant for this work is model selection.

Learning-to-Optimize with Guarantees. Chen et al. (2021) point out that learned optimization methods may lack theoretical guarantees for the sake of convergence speed. That being said, there are applications where a convergence guarantee is of highest priority. To underline this, Moeller et al. (2019) provide an example where a purely learning-based approach fails to reconstruct the crucial details in a medical image. Also, they prove convergence of their method by restricting the output to descent directions, for which mathematical guarantees exist. The basic idea is to trace the learned object back to, or constrain it to, a mathematical object with convergence guarantees. Similarly, Sreehari et al. (2016) provide sufficient conditions under which the learned mapping is a proximal mapping. Related schemes, under different assumptions and guarantees, are given by Chan et al. (2016), Teodoro et al. (2017), Tirer and Giryes (2018), Buzzard et al. (2018), Ryu et al. (2019), Sun et al. (2019), Terris et al. (2021) and Cohen et al. (2021). A major advantage of these methods is the fact that the number of iterations is not restricted a priori. However, a major drawback is their restriction to specific algorithms and problems. Another approach, which limits the number of iterations, yet in principle can be applied to every iterative optimization algorithm, is unrolling, pioneered by Gregor and LeCun (2010). Xin et al. (2016) study the IHT algorithm and show that it is, under some assumptions, able to achieve a linear convergence rate. Likewise, Chen et al. (2018) establish a linear convergence rate for the unrolled ISTA. However, a difficulty in the theoretical analysis of unrolled algorithms is actually the notion of convergence itself, such that one rather has to consider the generalization performance. Only few works have addressed this: Either directly by means of Rademacher complexity (Chen et al., 2020b), or indirectly in form of a stability analysis (Kobler et al., 2020), as algorithmic stability is linked to generalization (Bousquet and Elisseeff, 2000, 2002; Shalev-Shwartz et al., 2010). *We consider a mixture of the iterative approach and the unrolling approach: The theoretical analysis corresponds to the approach of unrolling, that is, a fixed number of iterations, and the resulting PAC-Bayesian generalization bound applies to this particular number of iterations only. However, in our implementation and experiments, we stay more closely to the iterative approach of learning an update step that can be applied for an arbitrary number of iterations. Further, while the resulting algorithms do include neural networks in their update step, the construction is closely related to the model-based approach, as it is inspired by analytical algorithms.*

Typical Problems and Tricks in Learning-to-Optimize. The literature contains a vast amount of approaches, tricks and problems about learning procedures for optimization algorithms. We only touch briefly upon the ones that are most relevant for the present work:

Besides the actual optimizer, a typical design choice is that of the loss function. Here, either the final loss or a weighted sum of the losses incurred along the iterations is typically being used (Chen et al., 2021). This extends to choices about the computational graph, as one has to decide which gradients are computed during training. An assumption already made by Andrychowicz et al. (2016) and widely adopted thereafter, is to neglect higher-order derivatives. Importantly, if the theoretical analysis is coherent with the computational graph, one is allowed to do whatever works here.

A major problem of many learned optimization algorithms, especially the ones based on recurrent neural networks, is their restriction to a certain number of iterations, that is, they cannot be trained for an arbitrary number of iterations due to instabilities, such as vanishing or exploding gradients, and memory bottlenecks. Further, often they do not generalize well to more iterations (Andrychowicz et al., 2016; Chen et al., 2017; Lv et al., 2017; Chen et al., 2021). A typical way to mitigate this problem is to split the whole trajectory into smaller parts, for which the vanishing/exploding gradient phenomenon does not occur, which effectively also solves the memory problem. However, this often does not lead to fully satisfactory results either, such that other approaches have been proposed: Lv et al. (2017) introduce random scaling of the coordinates and the addition of a convex function to the objective. Metz et al. (2019) propose not to model the recurrent nature of optimization algorithms directly, and therefore to replace the recurrent neural network with a multilayer perceptron, to “smooth” the loss, and to dynamically use two unbiased gradient estimators instead of one. Doing so they manage to train algorithms that are faster in wall-clock time than standard ones like Adam. Wichrowska et al. (2017) introduce a hierarchical RNN architecture consisting of three RNNs, and additionally draw the number of unrollings and the unrolling length from a heavy-tailed exponential distribution. While achieving the needed generalization, this approach is computationally costly and does not achieve the same wall-clock time as simple optimization algorithms. Chen et al. (2020a) consider improved training techniques in general, and they introduce a progressive scheme that gradually increases the unrolling length, and an imitation learning approach by learning from analytic optimizers first. *We will introduce a new loss function for training optimization algorithms, motivated by an intuitive theoretical argument. To our best knowledge, this loss function is new. Here, we will typically also ignore higher-order derivatives in the computational graph, as the algorithms performance is analyzed independently of this step in the learning procedure. Further, we found that a combination of many of the above mentioned approaches works well: 1) We use a single learned update based on MLPs and analytic algorithms instead of an RNN. 2) We split the trajectory into subtrajectories and randomize its total length, however, in a different and new way, again motivated by theoretical considerations. 3) We start the learning procedure with imitation learning to find an initialization that performs similar to analytic algorithms.*

PAC-Bounds through Change-of-Measure. The PAC-Bayesian framework allows for giving high probability bounds on the risk, either as an oracle or as an empirical bound. The key ingredients are so-called change-of-measure inequalities. The choice of such an inequality strongly influences the corresponding bound. The one used most often is based on a variational representation of the Kullback–Leibler divergence due to Donsker and Varadhan (1975), employed, for example, by Catoni (2004, 2007). Yet, not all bounds are based

on a variational representation, that is, holding uniformly over all posterior distributions (Rivasplata et al., 2020). However, most bounds involve the Kullback–Leibler divergence as measure of proximity, for example, those by McAllester (2003a,b), Seeger (2002), Langford and Shawe-Taylor (2002), or the general PAC-Bayes bound of Germain et al. (2009). More recently, other divergences have been used: Honorio and Jaakkola (2014) prove an inequality for the χ^2 -divergence, which is also used by London (2017). Bégin et al. (2016) and Alquier and Guedj (2018) use the Renyi-divergence (α -divergence). Ohnishi and Honorio (2021) propose several PAC-bounds based on general f-divergences, which include the Kullback–Leibler-, α - and χ^2 -divergences. Even more recently, Amit et al. (2022) propose to replace the Kullback-Leibler divergence by so-called “integral probability metrics”, which encompass, for example, the Wasserstein distance that obeys many favorable properties and also captures the geometry of the underlying space (see Villani et al., 2009). Motivated by this, Haddouche and Guedj (2023) also investigate PAC-Bayesian generalization bounds for the Wasserstein distance and their interplay with the output of optimization algorithms. A major advantage of using the Wasserstein distances instead of the Kullback-Leibler divergence is the fact that it does not constrain the support of the distribution a-priori through the choice of the prior. On the other hand, it demands assumptions on the loss function, which are not necessarily satisfied in learning-to-optimize. *We give a general PAC-Bayesian theorem based on exponential families. Here, the role of prior, posterior, divergence and data dependence will be given naturally. Moreover, this approach allows us to implement a more abstract learning framework that can be applied to a wide variety of algorithms.*

Boundedness of the Loss Function. A major drawback of many of the existing PAC-Bayes bounds is the assumption of a bounded loss-function. However, this assumption is mainly there to apply some exponential moment inequality like the Hoeffding- or Bernstein-inequality (Rivasplata et al., 2020; Alquier, 2021). Several ways have been developed to solve this problem: Germain et al. (2009) explicitly include the exponential moment in the bound, Alquier et al. (2016) use so-called Hoeffding- and Bernstein-assumptions, Catoni (2004) restricts to the sub-Gaussian or sub-Gamma case. Another possibility to ensure the generalization or exponential moment bounds is to use properties of the algorithm in question. London (2017) uses algorithmic stability to provide PAC-Bayes bounds for SGD. *We consider suitable properties of optimization algorithms aside from algorithmic stability to ensure the exponential moment bounds. To the best of our knowledge, this is new.*

Minimization of the PAC-Bound. The PAC-bound relates the true risk to other terms such as the empirical risk. Yet, it does not directly say anything about the absolute numbers. Thus, in learning procedures based on the PAC-Bayesian theory one typically aims to minimize this bound: Langford and Caruana (2001) compute non-vacuous numerical generalization bounds through a combination of PAC-bounds with a sensitivity analysis. Dziugaite and Roy (2017) extend this by minimizing the PAC-bound directly. Pérez-Ortiz et al. (2021) also consider learning as minimization of the PAC-Bayes bound and provide tight generalization bounds. Thiemann et al. (2017) are able to solve the minimization problem resulting from their PAC-bound by alternating minimization. Further, they provide sufficient conditions under which the resulting minimization problem is quasi-convex. *We also follow this approach and consider learning as minimization of the PAC-Bayesian upper-bound, however, applied to the context of learning-to-optimize.*

Choice of the Prior. A common difficulty in learning with PAC-Bayesian bounds is the choice of the prior distribution, as it heavily influences the performance of the learned models and the generalization bound (Catoni, 2004; Dziugaite et al., 2021; Pérez-Ortiz et al., 2021). In part, and especially for the Kullback-Leibler divergence, this is due to the fact that the divergence term can dominate the bound, keeping the posterior close to the prior. This leads to the idea of choosing a data- or distribution-dependent prior (Seeger, 2002; Parrado-Hernández et al., 2012; Lever et al., 2013; Dziugaite and Roy, 2018; Pérez-Ortiz et al., 2021), which, by using an independent subset of the data set, gets optimized to yield a good performance. *The choice of the prior distribution is crucial for the performance of our learned algorithms. Thus, we use a data-dependent prior. Further, we show how the prior is essential in preserving crucial properties during learning. It is key to control the trade-off between convergence guarantee and convergence speed. Based on these insights, a major part of our algorithmic realization is concerned with constructing such a prior distribution.*

More Generalization Bounds. There are many areas of machine learning research that study generalization bounds and have not been discussed here. Importantly, the vast field of “stochastic optimization” (SO) provides generalization bounds for specific algorithms. The main differences to our setting are the learning approach and the assumptions made:

- In most of the cases, the concrete algorithms studied in SO generate a single point estimate by either minimizing the (regularized) empirical risk functional over a possibly large data set, or by repeatedly updating the point estimate based on a newly drawn (small) batch of samples. Then, one studies the properties of this point in terms of the stationarity measure of the true risk functional (Bottou et al., 2018; Davis and Drusvyatskiy, 2022; Bianchi et al., 2022). In our case, as we consider the PAC-Bayesian learning approach, the final object to be studied is a distribution over hyperparameters.
- As the setting in SO is more explicit, more assumptions have to be made. Typical assumptions are (weak) convexity (Shalev-Shwartz et al., 2009; Davis and Drusvyatskiy, 2019), bounded gradients (Défossez et al., 2022), bounded noise (Davis and Drusvyatskiy, 2022), or smoothness (Kavis et al., 2022). Since we consider an abstract optimization algorithm, and the problem of finding a distribution over its hyperparameters, all of these assumptions cannot be made without severely limiting the applicability of our results.

We provide a principled way to learn a distribution over general hyperparameters of an abstract algorithm under weak assumptions and go explicitly beyond analytically tractable quantities. Therefore, the methodology is independent of the chosen implementation.

2 Problem Setup & Assumptions

In this section we establish the notation, formalize the setting, and state the main assumptions that are used throughout the remainder of the text.

Notation. We will endow every topological space X with the corresponding Borel- σ -algebra $\mathcal{B}(X)$, and, given a product space $X \times Y$ of two measurable spaces (X, \mathcal{A}) and (Y, \mathcal{B}) , we endow it with the product- σ -algebra $\mathcal{A} \otimes \mathcal{B}$. To denote the product space of a larger number of spaces X_1, \dots, X_n , we will use $\prod_{i=1}^n X_i$ as the shorthand for $X_1 \times \dots \times X_n$.

and $\bigotimes_{i=1}^n \mathcal{B}(X_i)$ for the product- σ -algebra. Further, if $X = X_1 = \dots = X_n$, we also use $X^n = \prod_{i=1}^n X_i$. If we are given a multivariate function $f : X \times Y \rightarrow Z$, then $f(x, \cdot) : Y \rightarrow Z$ denotes the function $y \mapsto f(x, y)$ with fixed element $x \in X$. Similarly, for a set $C \subset X \times Y$ we denote the section of C for fixed $x \in X$ by $C_x := \{y \in Y : (x, y) \in C\}$. In general, generic sets are denoted in the typewriter font, for example A , and $\mathbb{1}_A$ denotes the function that is equal to one for $x \in A$ and zero else, while ι_A denotes the function that is equal to zero for $x \in A$ and $+\infty$ else.¹ For a measure μ on a measurable space X , and a measurable function f on X , $\mu[f]$ denotes the integral of f w.r.t. μ , while, if $f \geq 0$, $f \cdot \mu$ denotes the measure given by $(f \cdot \mu)[A] = \int_A f(x) \mu(dx)$, that is, $(f \cdot \mu)[X] = \mu[f]$ and $(f \cdot \mu)[A] = \mu[f \cdot \mathbb{1}_A]$. Hence, by definition, the measure $f \cdot \mu$ is absolutely continuous w.r.t. μ , written as $f \cdot \mu \ll \mu$, with f being the corresponding density. Here, the set of all measures on a measurable space (X, \mathcal{X}) will be denoted by $\mathcal{M}(X) := \{\mu : \mathcal{X} \rightarrow [0, \infty] : \mu \text{ is a measure}\}$, and the set of all probability measures that are absolutely continuous w.r.t. $\mu \in \mathcal{M}(X)$ are denoted by $\mathcal{P}(\mu) := \{\nu \in \mathcal{M}(X) : \|\nu\| = \nu[X] = 1 \text{ and } \nu \ll \mu\}$. Here, the Kullback-Leibler divergence between two measures ν and μ is defined as

$$D_{\text{KL}}(\mu \parallel \nu) = \begin{cases} \mu[\log(f)] = \int_X \log(f(x)) \mu(dx), & \mu \ll \nu \text{ with density } f, \\ +\infty, & \text{otherwise.} \end{cases}$$

If $\mu = \mathbb{P}$ is actually the specific probability measure of the underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, the corresponding expectation is denoted by $\mathbb{E}[f] := \mathbb{P}[f] = \int_\Omega f(\omega) \mathbb{P}(d\omega)$. Random variables, that is, measurable functions on $(\Omega, \mathcal{F}, \mathbb{P})$, are written in Fraktur-font, for example \mathfrak{X} . Given two random variables, say $\mathfrak{X} : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow X$ and $\mathfrak{Y} : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow Y$, integration of a function f on $X \times Y$ w.r.t. the induced probability measure $\mathbb{P}_{(\mathfrak{X}, \mathfrak{Y})}$ is specified by the subscript $(\mathfrak{X}, \mathfrak{Y})$, that is:

$$\mathbb{E}[f(\mathfrak{X}, \mathfrak{Y})] = \int_\Omega f(\mathfrak{X}, \mathfrak{Y})(\omega) \mathbb{P}(d\omega) = \int_{X \times Y} f(x, y) \mathbb{P}_{(\mathfrak{X}, \mathfrak{Y})}(dx, dy) = \mathbb{P}_{(\mathfrak{X}, \mathfrak{Y})}[f] = \mathbb{E}_{(\mathfrak{X}, \mathfrak{Y})}[f].$$

If we are given a regular version of the conditional distribution of \mathfrak{Y} , given \mathfrak{X} , denoted by $\mathbb{P}_{\mathfrak{Y}|\mathfrak{X}}$, their joint distribution $\mathbb{P}_{(\mathfrak{X}, \mathfrak{Y})}$ can be disintegrated into the product $\mathbb{P}_{\mathfrak{X}} \otimes \mathbb{P}_{\mathfrak{Y}|\mathfrak{X}}$ of the marginal $\mathbb{P}_{\mathfrak{X}}$ and the probability kernel $(x, B) \mapsto \mathbb{P}_{\mathfrak{Y}|\mathfrak{X}=x}[B]$, which allows us to use the notation:

$$\mathbb{E}[f(\mathfrak{X}, \mathfrak{Y})] = \int_X \int_Y f(x, y) \mathbb{P}_{\mathfrak{Y}|\mathfrak{X}=x}(dy) \mathbb{P}_{\mathfrak{X}}(dx) = \mathbb{E}_{\mathfrak{X}}[\mathbb{E}_{\mathfrak{Y}|\mathfrak{X}=x}[f(x, \cdot)]] .$$

Note that, in this case, naively changing the order of integration is not allowed. However, if \mathfrak{X} and \mathfrak{Y} are independent, their joint distribution $\mathbb{P}_{(\mathfrak{X}, \mathfrak{Y})}$ is given by the product measure $\mathbb{P}_{\mathfrak{X}} \otimes \mathbb{P}_{\mathfrak{Y}}$ for which Fubini's theorem is applicable, and the iterated integration is clarified by the corresponding subscripts $\mathfrak{X}, \mathfrak{Y}$:

$$\mathbb{E}[f(\mathfrak{X}, \mathfrak{Y})] = (\mathbb{P}_{\mathfrak{X}} \otimes \mathbb{P}_{\mathfrak{Y}})[f] = \int_X \int_Y f(x, y) \mathbb{P}_{\mathfrak{Y}}(dy) \mathbb{P}_{\mathfrak{X}}(dx) = \mathbb{E}_{\mathfrak{X}}[\mathbb{E}_{\mathfrak{Y}}[f(x, \cdot)]_{|x=\mathfrak{X}}] .$$

1. We omit the name here, as both $\mathbb{1}_A$ and ι_A are called ‘‘indicator function’’. The former in probability theory, the latter in optimization.

Furthermore, we will denote the extended real numbers by $\bar{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$. Since it holds that $\mathcal{B}(\bar{\mathbb{R}})|_{\mathbb{R}} = \mathcal{B}(\mathbb{R})$, any $\mathcal{B}(\mathbb{R})$ -measurable function can be identified with a $\mathcal{B}(\bar{\mathbb{R}})$ -measurable function (Klenke, 2013, p. 38, Corollary 1.87). Finally, our theoretical results will rely on the notions of probability kernels, Polish spaces, and exponential families, whose definitions are recalled in Appendix A.

2.1 Main Assumptions and Definitions

We assume that, inside the class of functions that is encoded analytically in the model, we are given a distribution over loss-functions with a specific structure, which is modelled by a random variable:

Assumption 2.1 *We are given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and, for some fixed $N \in \mathbb{N}$, we are given random variables $\mathfrak{S}, \mathfrak{S}_1, \dots, \mathfrak{S}_N : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow \Theta$, where Θ is a Polish space. Further, we are given a non-negative and measurable loss-function $\ell : \mathbb{R}^n \times \Theta \rightarrow [0, +\infty]$.*

Remark 2.2 *The assumption of a Polish space is not restrictive, yet suffices to ensure the existence of regular versions of the conditional probability distribution, that is, for the disintegration of a joint distribution into a marginal and a corresponding probability kernel (see Kallenberg, 2021, Thm. 8.5).*

Then, ideally, we would like to find a solution to each realization of the random objective:

$$\text{Find } x' : \Theta \rightarrow \mathbb{R}^n, \text{ s.t. } x'(\mathfrak{S}) \in \arg \min_{x \in \mathbb{R}^n} \ell(x, \mathfrak{S}) \quad a.s. \quad (1)$$

However, we will only solve a relaxed version of (1) and provide *generalization bounds* for the *average performance*. Implicitly, we assume that we are given several realizations $\ell(\cdot, \theta_1), \dots, \ell(\cdot, \theta_N)$ with similar structure, as $\theta_1, \dots, \theta_N \stackrel{iid}{\sim} \mathbb{P}_{\mathfrak{S}}$. This is close to the “standard” learning problem, where an algorithm gets trained on a data set and then applied to unseen data. For this, we need to have a data set:

Definition 2.3 *The measurable function $\mathfrak{D} : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow \Theta^N$, $\omega \mapsto (\mathfrak{S}_1, \dots, \mathfrak{S}_N)(\omega)$ is called a data set. Here, Θ^N is called the data-space. If the induced distribution $\mathbb{P}_{\mathfrak{D}} = \mathbb{P}_{(\mathfrak{S}_1, \dots, \mathfrak{S}_N)}$ factorizes into the product of the marginals, that is, if $\mathbb{P}_{\mathfrak{D}} = \bigotimes_{i=1}^N \mathbb{P}_{\mathfrak{S}_i}$, it is called independent and if, additionally, it satisfies $\mathbb{P}_{\mathfrak{D}} = \bigotimes_{i=1}^N \mathbb{P}_{\mathfrak{S}}$, it is called i.i.d.*

PAC-Bayesian generalization bounds involve a so-called posterior distribution, which usually is a “data-dependent distribution”. Often, this term is left unspecified. Yet, as also pointed out by Rivasplata et al. (2020), this is an instance of a probability kernel (also called a “stochastic”- or “Markov kernel”). Another commonly used name is “regular conditional probability”, following the definition of a regular conditional distribution through probability kernels (Catoni, 2004; Alquier, 2008). In our opinion, the notion of a probability kernel is the least ambiguous one, leading to the following definition:

Definition 2.4 *Let \mathfrak{D} be a data set with data-space Θ^N , and let X be a measurable space. A probability kernel from Θ^N to X is called a data-dependent distribution on X .*

For solving problem (1), for every realization θ of \mathfrak{S} , we apply an optimization algorithm \mathcal{A} to $\ell(\cdot, \theta)$. For this, we consider a similar setting as London (2017), that is, randomized algorithms are considered as deterministic algorithms with randomized hyperparameters:

Definition 2.5 *Let H be a Polish space and $n \in \mathbb{N}$. A measurable function*

$$\mathcal{A} : H \times \mathbb{R}^n \times \Theta \longrightarrow \mathbb{R}^n, \quad (\alpha, x^{(0)}, \theta) \mapsto \mathcal{A}(\alpha, x^{(0)}, \theta),$$

is called a parametric algorithm and H is called the hyperparameter space of \mathcal{A} . \mathbb{R}^n is the space of the optimization variable, Θ the space of the parameters of the problem instance, and H the space of the hyperparameters of the algorithm.

Remark 2.6 *\mathcal{A} corresponds to the whole algorithm, that is, its output is the final iterate of the optimization procedure, which has to be stopped after some iteration $k \in \mathbb{N}$, and could be contrasted with the notion of an estimator in statistics. However, one can also model a single update-step, and recover \mathcal{A} as k -fold concatenation of this update-step.*

Here, learning \mathcal{A} refers to finding a distribution \mathbb{Q} on H based on its performance on a data set \mathfrak{D} . In the PAC-Bayesian approach to learning and generalization, one needs a reference distribution, which can (and should) encode prior knowledge about suitable choices of hyperparameters, called the *prior*:

Assumption 2.7 *We are given a parametric optimization algorithm \mathcal{A} with Polish hyperparameter space H , and a (prior) distribution $\mathbb{P}_{\mathfrak{H}}$ on H that is induced by a random variable $\mathfrak{H} : (\Omega, \mathcal{F}, \mathbb{P}) \longrightarrow H$, which is independent of the data set \mathfrak{D} and \mathfrak{S} . Further, the initialization $x^{(0)} \in \mathbb{R}^n$ is given and fixed.*

Notation 2.8 *We will refer to the random variable \mathfrak{S} as the parameters of the loss-function ℓ and to \mathfrak{H} as the hyperparameters of the algorithm \mathcal{A} . Further, to simplify the notation, we also use the short-hand $\ell(\alpha, \theta) := \ell(\mathcal{A}(\alpha, \theta), \theta)$. Furthermore, if not needed explicitly, $x^{(0)}$ and $(\Omega, \mathcal{F}, \mathbb{P})$ will not be mentioned anymore in the following.*

Now, we define the risk of a randomized parametric optimization algorithm as usual:

Definition 2.9 *Suppose \mathfrak{S} and ℓ satisfy Assumption 2.1, and \mathcal{A} satisfies Assumption 2.7. Further, let $\mathfrak{D} = (\mathfrak{S}_1, \dots, \mathfrak{S}_N)$ be an i.i.d. data set with distribution $\bigotimes_{i=1}^N \mathbb{P}_{\mathfrak{S}}$ that is independent of \mathfrak{S} . The risk of \mathcal{A} is defined as the measurable function:*

$$\mathcal{R} : H \longrightarrow [0, +\infty], \quad \alpha \mapsto \mathbb{E}[\ell(\mathcal{A}(\alpha, \mathfrak{S}), \mathfrak{S})] = \mathbb{E}[\ell(\alpha, \mathfrak{S})] = \mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot)].$$

Similarly, the empirical risk of \mathcal{A} on \mathfrak{D} is defined as the measurable function:

$$\hat{\mathcal{R}} : H \times \Theta^N \longrightarrow [0, +\infty], \quad (\alpha, \mathfrak{D}) \mapsto \hat{\mathcal{R}}(\alpha, \mathfrak{D}) = \frac{1}{N} \sum_{i=1}^N \ell(\alpha, \mathfrak{S}_i).$$

The following theory is based on exponential families (see Definition A.4). This is a class of distributions on some space X , which is mainly build upon three parts: a function η on some index set Λ , a function T on X , and a distribution μ on X . Since we want to have a distribution over hyperparameters $\alpha \in H$, we will choose $X = H$. However, then, the exponential family $(\mathbb{Q}_{\lambda})_{\lambda \in \Lambda}$ is not data-dependent. Therefore, we will define *data-dependent exponential families* by extending the domain of T from X to $X \times \Theta^N$.

Definition 2.10 Let \mathfrak{D} be a data set with data-space Θ^N , and let Λ be a non-empty index set. A family of probability measures $(\mathbb{Q}_\lambda)_{\lambda \in \Lambda}$ on a measurable space X is called a data-dependent exponential family (in η and T), if there is a dominating probability measure μ , that is, $(\mathbb{Q}_\lambda)_{\lambda \in \Lambda} \subset \mathcal{P}(\mu)$, functions $\eta : \Lambda \rightarrow \mathbb{R}^k$, $A : \Lambda \rightarrow (0, +\infty)$, and measurable functions $T : X \times \Theta^N \rightarrow \mathbb{R}^k$, $h : X \rightarrow (0, +\infty)$, such that we have $\mathbb{Q}_\lambda(d) = hA(\lambda) \exp(\langle \eta(\lambda), T(\cdot, d) \rangle) \cdot \mu$ for every $\lambda \in \Lambda$, $d \in \Theta^N$, that is, $\mathbb{Q}_\lambda(d, \mathbf{B}) = \int_{\mathbf{B}} h(x)A(\lambda) \exp(\langle \eta(\lambda), T(x, d) \rangle) \mu(dx)$, $\mathbf{B} \in \mathcal{B}(X)$.

Remark 2.11 If η describes a lower-dimensional manifold in \mathbb{R}^k , $(\mathbb{Q}_\lambda)_{\lambda \in \Lambda}$ is called a curved exponential family (Efron, 1975), whose properties might differ from the ones for linear exponential families ($\eta(\lambda) = \lambda \in \mathbb{R}^k$), for example, convexity of the map $\lambda \mapsto A(\lambda)$.

We introduce data-dependency through T , since it strongly affects the shape of the distribution and, contrary to η , is defined on the underlying space X . Since we want to learn a distribution over hyperparameters $\alpha \in H$, we make the following assumption:

Assumption 2.12 On the hyperparameter space H , we are given a data-dependent exponential family $(\mathbb{Q}_\lambda)_{\lambda \in \Lambda}$ in η and T with dominating probability measure $\mu = \mathbb{P}_{\mathfrak{F}}$, such that the map $\alpha \mapsto h(\alpha) \exp(\langle \eta(\lambda), T(\alpha, d) \rangle)$ is non-trivial and integrable w.r.t. $\mathbb{P}_{\mathfrak{F}}$ for every $\lambda \in \Lambda$, $d \in \Theta^N$, that is, $\mathbb{E}_{\mathfrak{F}} [h \exp(\langle \eta(\lambda), T(\cdot, d) \rangle)] \in (0, \infty)$.

Remark 2.13 (i) In PAC-Bayes, the dominating measure $\mathbb{P}_{\mathfrak{F}}$ is usually referred to as prior and every distribution $\mathbb{Q} \in \mathcal{P}(\mathbb{P}_{\mathfrak{F}})$ is referred to as a posterior. This deviates from the standard definitions of prior and posterior in Bayesian statistics.

(ii) The integrability assumption is slightly restrictive, as it affects the choice of h, η and T . However, in Section 4 we will construct η and T such that this holds anyway.

The last integral in Assumption 2.12 will be of great interest in the following. Here, we will use a similar notation as in Barndorff–Nielsen (2014) and denote

$$\begin{aligned} c(\lambda, d) &:= \int_H h(\alpha) \exp(\langle \eta(\lambda), T(\alpha, d) \rangle) \mathbb{P}_{\mathfrak{F}}(d\alpha) = \mathbb{E}_{\mathfrak{F}} \{h \exp(\langle \eta(\lambda), T(\cdot, d) \rangle)\}, \\ \kappa(\lambda, d) &:= \log(c(\lambda, d)) = \log(\mathbb{E}_{\mathfrak{F}} \{h \exp(\langle \eta(\lambda), T(\cdot, d) \rangle)\}), \end{aligned} \quad (2)$$

or for short, $\kappa = \log(c)$. With this notation, it holds that $A(\lambda) = c(\lambda)^{-1}$. Through this, as shown in Lemma B.1, every member of the data-dependent exponential family is indeed a data-dependent distribution on H in the sense of Definition 2.4. Finally, we will restrict the space Λ to a compact set. This is needed to get a uniform bound in λ , and such ideas appeared in the literature (see Langford and Caruana, 2001; Catoni, 2007; Alquier, 2021).

Assumption 2.14 Λ is a compact set with finite covering $\mathcal{O} := \{\mathcal{O}_1, \dots, \mathcal{O}_K\}$, that is, $\Lambda \subset \bigcup_{i=1}^K \mathcal{O}_i$, such that there is a constant $C_{\mathcal{O}}$, which, for every $d \in \Theta^N$, allows for the bound $\max_{i=1, \dots, K} \sup_{\lambda, \lambda' \in \mathcal{O}_i} \kappa(\lambda, d) - \kappa(\lambda', d) \leq C_{\mathcal{O}}$.

Remark 2.15 The existence of a general finite covering for Λ is a consequence of the compactness. However, the existence of the constant $C_{\mathcal{O}}$ is non-trivial. It does hold, for example, if Λ is a finite set ($K = |\Lambda|$, $C_{\mathcal{O}} = 0$), or, if (Λ, ρ) is a compact metric space and κ is Lipschitz-continuous in λ (uniformly in d) with Lipschitz constant L , such that $C_{\mathcal{O}} = L \cdot \max_{i=1, \dots, K} \text{diam } \mathcal{O}_i$, where the diameter of a set \mathbf{A} is given by $\text{diam } \mathbf{A} = \sup_{x, y \in \mathbf{A}} \rho(x, y)$.

3 General PAC-Bayesian Theorem

In this section we prove the following general PAC-Bayesian bound for data-dependent exponential families, which then can be specialized into a generalization bound of the learned parametric optimization algorithm \mathcal{A} . It is based on the following two lemmas, whose proofs can be found in Appendix C and D, respectively. The first lemma is a form of the Donsker–Varadhan variational formulation and yields uniformity in the distributions \mathbb{Q} :

Lemma 3.1 *Suppose that Assumption 2.12 holds and define κ as in (2). Then for every $\lambda \in \Lambda$ and $d \in \Theta^N$:*

$$\kappa(\lambda, d) = \sup_{\mathbb{Q} \in \mathcal{P}(\mathbb{P}_{\mathfrak{H}})} \mathbb{Q}[\langle \eta(\lambda), T(\cdot, d) \rangle + \log(h)] - D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}_{\mathfrak{H}}).$$

Furthermore, for every $\lambda \in \Lambda$, the supremum is attained at $\mathbb{Q}_\lambda(d)$.

Similarly, the second lemma yields uniformity in $\lambda \in \Lambda$. For this, we need to be able to cover Λ , and to control $\lambda \mapsto \kappa(\lambda, d)$ for every $d \in \Theta^N$, as specified in Assumption 2.14.

Lemma 3.2 *Suppose that Assumption 2.14 holds and assume that $\mathbb{P}\{\kappa(\lambda, \mathfrak{D}) > s\} \leq \exp(-s)$ for all $s \in \mathbb{R}$ and $\lambda \in \Lambda$. Then $\mathbb{P}\{\sup_{\lambda \in \Lambda} \kappa(\lambda, \mathfrak{D}) \leq \log(K/\epsilon) + C_{\mathcal{O}}\} \geq 1 - \epsilon$.*

Taken together, these two lemmas allow us to state the PAC-Bayesian theorem in its general form:

Theorem 3.3 *Suppose that Assumptions 2.12 and 2.14 hold, and assume that $\mathbb{E}_{\mathfrak{D}}[c(\lambda, \cdot)] \leq 1$ for all $\lambda \in \Lambda$. Then, it holds that:*

$$\mathbb{P}\left\{\forall \lambda \in \Lambda, \forall \mathbb{Q} \in \mathcal{P}(\mathbb{P}_{\mathfrak{H}}) : \mathbb{Q}[\langle \eta(\lambda), T(\cdot, d) \rangle + \log(h)]|_{d=\mathfrak{D}} \leq D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}_{\mathfrak{H}}) + \log\left(\frac{K}{\epsilon}\right) + C_{\mathcal{O}}\right\} \geq 1 - \epsilon.$$

Proof For every $\lambda \in \Lambda$, $c(\lambda, \mathfrak{D})$ is a non-negative random variable, and $x \mapsto \exp(x)$ is monotonically increasing. Thus, Markov’s inequality yields for $s \in \mathbb{R}$:

$$\mathbb{P}\{c(\lambda, \mathfrak{D}) > \exp(s)\} \leq \frac{\mathbb{E}[c(\lambda, \mathfrak{D})]}{\exp(s)} \leq \exp(-s).$$

Thus, $\mathbb{P}\{\kappa(\lambda, \mathfrak{D}) > s\} \leq \exp(-s)$, because $c(\lambda, \mathfrak{D}) > \exp(s)$ is equivalent to $\kappa(\lambda, \mathfrak{D}) = \log(c(\lambda, \mathfrak{D})) > s$. Hence, Lemma 3.2 is applicable and gives:

$$\mathbb{P}\left\{\sup_{\lambda \in \Lambda} \kappa(\lambda, \mathfrak{D}) \leq \log\left(\frac{K}{\epsilon}\right) + C_{\mathcal{O}}\right\} \geq 1 - \epsilon.$$

Using Lemma 3.1 gives:

$$\mathbb{P}\left\{\sup_{\lambda \in \Lambda} \sup_{\mathbb{Q} \in \mathcal{P}(\mathbb{P}_{\mathfrak{H}})} \mathbb{Q}[\langle \eta(\lambda), T(\cdot, d) \rangle + \log(h)]|_{d=\mathfrak{D}} - D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}_{\mathfrak{H}}) \leq \log\left(\frac{K}{\epsilon}\right) + C_{\mathcal{O}}\right\} \geq 1 - \epsilon.$$

Simply rearranging and reformulating yields the result. ■

Remark 3.4 (i) Note that the statement is still true for a data-dependent prior $\mathbb{P}_{\mathfrak{S}}$: Denoting the additional data set by \mathfrak{D}' , one needs to assume that $\mathbb{E}[c(\lambda, (\mathfrak{D}, \mathfrak{D}'))] \leq 1$. Also Lemma 3.1 still applies “pointwise” (with the appropriate notational adjustments). Intuitively, this is due to the fact that the PAC-bound is a statement about relative, and not absolute values.

(ii) In Section 4 we provide sufficient conditions s.t. $\mathbb{E}[c(\lambda, \mathfrak{D})] \leq 1$ holds for all $\lambda > 0$.

(iii) $\log(K)$ bears the intrinsic dimension of Λ and thus, in full generality, might be large. For our generalization bound, however, it has only a minor influence, since there $\Lambda \subset \mathbb{R}$, and the the empirical risk is typically much larger.

For the rest of the paper, we will have $h \equiv 1$, such that $\log(h) \equiv 0$. The following corollary, whose proof is given in Appendix E, shows how to transform Theorem 3.3 into a high-probability bound on the risk.

Corollary 3.5 (PAC-Bayesian Generalization Bound) Denote the coordinates of T and η by $T = (T^{(1)}, T^{(2)}, \dots, T^{(k)}) =: (T^{(1)}, T^{(r)})$ and $\eta = (\eta^{(1)}, \eta^{(2)}, \dots, \eta^{(k)}) =: (\eta^{(1)}, \eta^{(r)})$. If $T^{(1)} = \mathcal{R} - \hat{\mathcal{R}}$ and $\eta^{(1)} > 0$, the following are equivalent for any $\lambda \in \Lambda$, $d \in \Theta^N$, $\mathbb{Q} \in \mathcal{P}(\mathbb{P}_{\mathfrak{S}})$:

$$(i) \quad \mathbb{Q}[\langle \eta(\lambda), T(\cdot, d) \rangle] \leq D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}_{\mathfrak{S}}) + \log\left(\frac{K}{\varepsilon}\right) + C_{\mathcal{O}},$$

$$(ii) \quad \mathbb{Q}[\mathcal{R}] \leq \mathbb{Q}[\hat{\mathcal{R}}(\cdot, d)] + \frac{1}{\eta^{(1)}(\lambda)} \left(D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}_{\mathfrak{S}}) + \log\left(\frac{K}{\varepsilon}\right) + C_{\mathcal{O}} - \mathbb{Q}[\langle \eta^{(r)}(\lambda), T^{(r)}(\cdot, d) \rangle] \right).$$

In particular, if Theorem 3.3 applies, we can replace (i) with (ii).

4 Learning-to-Optimize with Guarantees

Here, we consider properties of optimization algorithms, that assert the necessary condition $\mathbb{E}[c(\lambda, \mathfrak{D})] \leq 1$ for all $\lambda \in \Lambda$ to employ the PAC-Bayesian bound from Section 3.

4.1 Guaranteed Convergence

In the next theorem, the additional assumption on \mathcal{A} is sufficient to ensure the assumptions of Theorem 3.3. Essentially, it requires the loss of the algorithm’s output to be bounded. Yet, as shown in Section 4.2, it is too restrictive to learn hyperparameters that allow for a significant acceleration compared to the standard choices from a worst-case analysis.

Theorem 4.1 Suppose that \mathfrak{S} and ℓ satisfy Assumption 2.1, and suppose that \mathcal{A} satisfies Assumption 2.7. Further, assume that there is a constant $C \geq 0$ and a measurable function $\rho : H \rightarrow [0, \infty)$, such that for every $\alpha \in H$ it holds that $\ell(\alpha, \cdot) \leq C\rho(\alpha)\ell(x^{(0)}, \cdot)$ $\mathbb{P}_{\mathfrak{S}}$ -a.s. Furthermore, let \mathfrak{D} be a corresponding i.i.d. data set of size $N \in \mathbb{N}$. Finally, assume that $\mathbb{E}[\ell(x^{(0)}, \mathfrak{S})^2] < \infty$, and define $\eta : (0, \infty) \rightarrow \mathbb{R}^2$ and $T : H \times \Theta^N \rightarrow \mathbb{R}^2$ through:

$$\eta(\lambda) := \left(\lambda, -\frac{\lambda^2}{2} \right), \quad T(\alpha, d) := \left(\mathcal{R}(\alpha) - \hat{\mathcal{R}}(\alpha, d), \frac{\rho^2(\alpha)C^2}{N} \mathbb{E}[\ell(x^{(0)}, \mathfrak{S})^2] \right).$$

Then it holds that $\mathbb{E}[c(\lambda, \mathfrak{D})] \leq 1$ for all $\lambda > 0$.

Proof Since \mathfrak{H} and \mathfrak{D} are independent, their joint distribution is given by the product measure $\mathbb{P}_{\mathfrak{D}} \otimes \mathbb{P}_{\mathfrak{H}}$. Further, $\exp \geq 0$, such that Fubini's theorem is applicable and allows to do the integration iteratively, that is:

$$\mathbb{E} \left[\exp \left(\lambda (\mathcal{R}(\mathfrak{H}) - \hat{\mathcal{R}}(\mathfrak{H}, \mathfrak{D})) \right) \right] = \mathbb{E} \left\{ \mathbb{E} \left[\exp \left(\lambda (\mathcal{R}(\alpha) - \hat{\mathcal{R}}(\alpha, \mathfrak{D})) \right) \right] \Big|_{\alpha=\mathfrak{H}} \right\}.$$

Hence, first consider the inner integral for a fixed $\alpha \in H$. Then, by definition and the i.i.d. assumption one gets:

$$\begin{aligned} \mathbb{E} \left[\exp \left(\lambda (\mathcal{R}(\alpha) - \hat{\mathcal{R}}(\alpha, \mathfrak{D})) \right) \right] &= \mathbb{E} \left\{ \exp \left(-\frac{\lambda}{N} \sum_{i=1}^N (\ell(\alpha, \mathfrak{S}_i) - \mathbb{E}[\ell(\alpha, \mathfrak{S})]) \right) \right\} \\ &= \prod_{i=1}^N \mathbb{E}_{\mathfrak{S}} \left\{ \exp \left(-\frac{\lambda}{N} (\ell(\alpha, \cdot) - \mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot)]) \right) \right\}. \end{aligned}$$

The loss-function is non-negative and, by assumption, \mathcal{A} can be bounded $\mathbb{P}_{\mathfrak{S}}$ -a.s. Thus, for every $\alpha \in H$, $\ell(\alpha, \mathfrak{S})$ is a non-negative random variable with finite second-moment, as $\mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot)^2] \leq C^2 \rho(\alpha)^2 \mathbb{E}_{\mathfrak{S}}[\ell(x^{(0)}, \cdot)^2] < \infty$. Thus, by Lemma B.2, the boundedness assumption on \mathcal{A} , and the monotonicity of the exponential function, one gets the bound:

$$\begin{aligned} \mathbb{E}_{\mathfrak{S}} \left\{ \exp \left(-\frac{\lambda}{N} (\ell(\alpha, \cdot) - \mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot)]) \right) \right\} &\leq \exp \left(\frac{\lambda^2}{2N^2} \mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot)^2] \right) \\ &\leq \exp \left(\frac{\lambda^2}{2N^2} C^2 \rho(\alpha)^2 \mathbb{E}_{\mathfrak{S}}[\ell(x^{(0)}, \cdot)^2] \right). \end{aligned}$$

Therefore we have the following bound:

$$\begin{aligned} \mathbb{E} \left[\exp \left(\lambda (\mathcal{R}(\alpha) - \hat{\mathcal{R}}(\alpha, \mathfrak{D})) \right) \right] &\leq \prod_{i=1}^N \exp \left(\frac{\lambda^2}{2N^2} C^2 \rho(\alpha)^2 \mathbb{E}_{\mathfrak{S}}[\ell(x^{(0)}, \cdot)^2] \right) \\ &= \exp \left(\frac{\lambda^2}{2N} C^2 \rho(\alpha)^2 \mathbb{E}_{\mathfrak{S}}[\ell(x^{(0)}, \cdot)^2] \right). \end{aligned}$$

This can be rearranged into $\mathbb{E} \left\{ \exp \left(\lambda (\mathcal{R}(\alpha) - \hat{\mathcal{R}}(\alpha, \mathfrak{D})) - \frac{\lambda^2}{2} \frac{C^2}{N} \rho(\alpha)^2 \mathbb{E}_{\mathfrak{S}}[\ell(x^{(0)}, \cdot)^2] \right) \right\} \leq 1$, as the right-hand side does not depend on \mathfrak{D} . Since \mathfrak{H} and \mathfrak{D} are independent, and $\alpha \in H$ was arbitrary, this inequality does hold $\mathbb{P}_{\mathfrak{H}}$ -a.s. Therefore, one directly gets the bound $\mathbb{E} \left\{ \exp \left(\lambda (\mathcal{R}(\mathfrak{H}) - \hat{\mathcal{R}}(\mathfrak{H}, \mathfrak{D})) - \frac{\lambda^2}{2} \frac{C^2}{N} \rho(\mathfrak{H})^2 \mathbb{E}_{\mathfrak{S}}[\ell(x^{(0)}, \cdot)^2] \right) \right\} \leq 1$. Now, again by Fubini's theorem, one can also switch the order of integration to get:

$$\mathbb{E} \left\{ \mathbb{E} \left[\exp \left(\lambda (\mathcal{R}(\mathfrak{H}) - \hat{\mathcal{R}}(\mathfrak{H}, d)) - \frac{\lambda^2}{2} \frac{C^2}{N} \rho(\mathfrak{H})^2 \mathbb{E}_{\mathfrak{S}}[\ell(x^{(0)}, \cdot)^2] \right) \right] \Big|_{d=\mathfrak{D}} \right\} \leq 1.$$

Inserting the definition of η and T gives $\mathbb{E} \{ \mathbb{E} [\exp (\langle \eta(\lambda), T(\mathfrak{H}, d) \rangle)] |_{d=\mathfrak{D}} \} \leq 1$. Here, the inner term is the same as $\mathbb{E} [\exp (\langle \eta(\lambda), T(\mathfrak{H}, d) \rangle)] = \int_H \exp (\langle \eta(\lambda), T(\alpha, d) \rangle) \mathbb{P}_{\mathfrak{H}}(d\alpha) = c(\lambda, d)$. Hence, this is the same as $\mathbb{E}[c(\lambda, \mathfrak{D})] \leq 1$. \blacksquare

Remark 4.2 *The argument still works for a data-dependent prior, if the corresponding data sets \mathcal{D}' and \mathcal{D} are independent: While interchanging the integration w.r.t. \mathcal{D}' and \mathfrak{H} is not allowed, an interchange w.r.t. \mathfrak{H} and \mathcal{D} is still valid (under the integral), that is, for a function f it would hold $\mathbb{E}[f(\mathfrak{H}, \mathcal{D}, \mathcal{D}')] = \mathbb{E}_{\mathcal{D}'} \{ \mathbb{E}_{\mathfrak{H}|\mathcal{D}'=d'} [\mathbb{E}_{\mathcal{D}} [f(\alpha, \cdot, d') |_{\alpha=\mathfrak{H}}]] \} = \mathbb{E}_{\mathcal{D}'} \{ \mathbb{E}_{\mathcal{D}} [\mathbb{E}_{\mathfrak{H}|\mathcal{D}'=d'} [f(\cdot, d, d') |_{d=\mathcal{D}}]] \}$, and the inner term is ≤ 1 in any case. The premise here is that the boundedness assumption on \mathcal{A} holds.*

4.2 Conditioning on Convergence

Most of the time, the previous approach learns hyperparameters that ensure convergence, since the boundedness assumption on \mathcal{A} implicitly requires weak theoretical worst-case estimates *almost surely*. If, for example, the considered class of functions $(\ell(\cdot, \theta))_{\theta \in \Theta}$ is that of general quadratic functions, the convergence behaviour is accurately represented by analytic quantities from a worst-case analysis. Thereby, the boundedness prevents “aggressive” step-size parameters that lie outside the worst-case convergence regime, as they would lead to a diverging behaviour, which increases the incurred empirical risk dramatically. Thus, to motivate the upcoming discussion, consider the following thought-experiment:

Example 4.3 *Consider $\ell(x, \theta) := \frac{\theta}{2}x^2$ and assume that the chosen algorithm is gradient descent, that is $x^{(k+1)} = x^{(k)} - \tau \ell'(x^{(k)}, \theta)$. For a given θ , the optimal step-size is $\tau = \frac{1}{\theta}$, which gives convergence in one step. However, if θ is given by samples from the distribution $\mathbb{P}_{\Theta} = 0.99\delta_1 + 0.01\delta_{100}$, a worst-case analysis would suggest to take $\tau_w = \frac{1}{100}$. In this case, we would have an algorithm that converges in a single step for 1% of the problem instances, while having a linear convergence rate of $(\frac{99}{100})^k$ for the other 99%. Another choice is to take $\tau_d = 1$, which leads to an algorithm that does converge in a single step for 99% of the problem instances, but diverges in 1% of the cases. By restricting to the 99% of the cases where convergence does occur, the overall difference in speed is drastic.*

Hence, in this section, a different approach is taken: We actually allow for divergence, if it only occurs in rare cases with a controllable probability, that is, “almost surely” is relaxed to “with a sufficiently large probability”. Essentially, we only consider the loss for all those hyperparameters, where the loss is bounded by a certain constant, as well as the probability for that to occur. Then, in Section 4.2.1, we develop a technique that allows the user to actually control this probability. Clearly, a stronger guarantee trades for convergence speed.

Definition 4.4 *Given a measurable function $g : \Theta \rightarrow \mathbb{R}$, the (parametric) sublevel set $\mathbf{L}_g \subset H \times \Theta$ is defined as $\mathbf{L}_g = \{(\alpha, \theta) \in H \times \Theta : \ell(\alpha, \theta) \leq g(\theta)\}$. The sections of \mathbf{L}_g for fixed $\alpha \in H$ will be denoted by $\mathbf{L}_{g,\alpha}$.*

Remark 4.5 *The concrete choice of g will influence the PAC-Bayesian bound. Again, for proving the result, the essential property is that the loss after applying \mathcal{A} can be bounded.*

In Lemma B.3 we show that \mathbf{L}_g is indeed a measurable set. This is not obvious, as the loss function and the algorithm are composed in a non-standard way. This result further implies that the sections $\mathbf{L}_{g,\alpha}$ are measurable, too. Since H and Θ are Polish spaces, the product $H \times \Theta$ is again Polish. Hence, there exists a regular version of the conditional probability of \mathfrak{S} , given \mathfrak{H} , that is, a kernel $H \rightarrow \Theta$, $(\alpha, \mathbf{B}) \mapsto \mathbb{P}_{\mathfrak{S}|\mathfrak{H}=\alpha}[\mathbf{B}]$. By Witting (2013, Thm. 1.122,

p.124), this determines a regular version of the conditional probability of $(\mathfrak{H}, \mathfrak{S})$, given \mathfrak{H} , through $H \rightarrow H \times \Theta$, $(\alpha, \mathbf{B}) \mapsto \mathbb{P}_{(\mathfrak{H}, \mathfrak{S})|\mathfrak{H}=\alpha}[\mathbf{B}] := \mathbb{P}_{\mathfrak{S}|\mathfrak{H}=\alpha}[\mathbf{B}_\alpha]$. Here, for any measurable function $f : H \times \Theta \rightarrow \mathbb{R}$ for which $\mathbb{E}[f(\mathfrak{H}, \mathfrak{S})]$ exists, it holds $\mathbb{P}_{\mathfrak{H}}$ -a.s.:

$$\mathbb{E}\{f(\mathfrak{H}, \mathfrak{S}) \mid \mathfrak{H} = \alpha\} = \int_{\Theta} f(\alpha, \theta) \mathbb{P}_{\mathfrak{S}|\mathfrak{H}=\alpha}(d\theta). \quad (3)$$

Thus, we get $\mathbb{P}_{\mathfrak{H}}$ -a.s. the equality $\mathbb{P}\{(\mathfrak{H}, \mathfrak{S}) \in \mathbf{B} \mid \mathfrak{H} = \alpha\} = \mathbb{P}_{\mathfrak{S}|\mathfrak{H}=\alpha}[\mathbf{B}_\alpha]$. In particular, this applies to the sublevel set \mathbf{L}_g , and the map $\alpha \mapsto \mathbb{P}_{\mathfrak{S}|\mathfrak{H}=\alpha}[\mathbf{L}_{g,\alpha}]$ is measurable.

Definition 4.6 *Let \mathbf{L}_g be a parametric sublevel set. Define the sublevel probability as the measurable function $\alpha \mapsto p(\alpha) := \mathbb{P}_{\mathfrak{S}|\mathfrak{H}=\alpha}[\mathbf{L}_{g,\alpha}]$.*

Lemma 4.7 *Suppose Assumption 2.7 holds, and let $\mathbb{P}_{\mathfrak{S}|\mathfrak{H}=\alpha}[\mathbf{L}_{g,\alpha}] > 0$ for every $\alpha \in H$. Then we have $\mathbb{P}_{\mathfrak{H}}$ -a.s.:*

$$(i) \quad p(\alpha) = \mathbb{P}_{\mathfrak{S}}[\mathbf{L}_{g,\alpha}],$$

$$(ii) \quad \mathbb{E}\{\ell(\mathfrak{H}, \mathfrak{S}) \cdot \mathbf{1}_{\mathbf{L}_g}(\mathfrak{H}, \mathfrak{S}) \mid \mathfrak{H} = \alpha\} = \mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot) \mathbf{1}_{\mathbf{L}_{g,\alpha}}] = p(\alpha) \cdot \mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot) \mid \mathbf{L}_{g,\alpha}].$$

Proof By (3) and the independence of \mathfrak{S} and \mathfrak{H} , we have $\mathbb{E}\{\ell(\mathfrak{H}, \mathfrak{S}) \mathbf{1}_{\mathbf{L}_g}(\mathfrak{H}, \mathfrak{S}) \mid \mathfrak{H} = \alpha\} = \int_{\Theta} \ell(\alpha, \theta) \mathbf{1}_{\mathbf{L}_{g,\alpha}}(\theta) \mathbb{P}_{\mathfrak{S}}(d\theta) = \mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot) \mathbf{1}_{\mathbf{L}_{g,\alpha}}] \mathbb{P}_{\mathfrak{H}}$ -a.s., which shows the first equality of (ii). Since $\mathbb{P}_{\mathfrak{S}}[\mathbf{L}_{g,\alpha}] > 0$, the elementary conditional expectation is defined as $\mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot) \mid \mathbf{L}_{g,\alpha}] = \frac{\mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot) \mathbf{1}_{\mathbf{L}_{g,\alpha}}]}{\mathbb{P}_{\mathfrak{S}}[\mathbf{L}_{g,\alpha}]}$. Again by independence we have $\mathbb{P}_{\mathfrak{H}}$ -a.s. the equality $\mathbb{P}_{\mathfrak{S}}[\mathbf{L}_{g,\alpha}] = \mathbb{P}_{\mathfrak{S}|\mathfrak{H}=\alpha}[\mathbf{L}_{g,\alpha}] = p(\alpha)$, which shows (i) and the second equality of (ii). \blacksquare

This construction allows us to give a more fine-grained analysis of the algorithm, as it allows to trade the boundedness assumption for the sublevel probability. This basically extends a worst-case analysis, which would correspond to an uniform upper bound. Motivated by Lemma 4.7, to actually analyze the algorithm in this setting, we define the *sublevel risk* and its empirical counterpart as the expect loss *conditioned* on the sublevel set:

Definition 4.8 *Let \mathbf{L}_g be a parametric sublevel set. Then the sublevel risk $\mathcal{R}_g : H \rightarrow [0, +\infty]$ is defined as the conditional expectation of the loss given $\mathbf{L}_{g,\alpha}$:*

$$\alpha \mapsto \mathcal{R}_g(\alpha) := \mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot) \mid \mathbf{L}_{g,\alpha}] = \begin{cases} \frac{1}{p(\alpha)} \mathbb{E}_{\mathfrak{S}}[\ell(\alpha, \cdot) \mathbf{1}_{\mathbf{L}_{g,\alpha}}], & \text{if } p(\alpha) > 0; \\ 0, & \text{otherwise.} \end{cases}$$

Given a data set $\mathfrak{D} = (\mathfrak{S}_1, \dots, \mathfrak{S}_N)$, the empirical sublevel risk $\hat{\mathcal{R}}_g : H \times \Theta^N \rightarrow [0, +\infty]$ is defined as $(\alpha, \mathfrak{D}) \mapsto \hat{\mathcal{R}}_g(\alpha, \mathfrak{D}) := \frac{1}{p(\alpha)} \frac{1}{N} \sum_{i=1}^N \mathbf{1}_{\mathbf{L}_{g,\alpha}}(\mathfrak{S}_i) \ell(\alpha, \mathfrak{S}_i)$.

The following theorem is a direct generalization of Theorem 4.1. Especially, note that the additional assumption on \mathcal{A} is not needed anymore.

Theorem 4.9 *Suppose that \mathfrak{S} and ℓ satisfy Assumption 2.1, and suppose that \mathcal{A} satisfies Assumption 2.7. Further, let \mathfrak{D} be a corresponding i.i.d. data set of size $N \in \mathbb{N}$, and let*

\mathbf{L}_g be a parametric sublevel set with sublevel probability p . Assume that $\mathbb{P}_{\mathfrak{H}}\{p > 0\} = 1$ and $\mathbb{E}_{\mathfrak{G}}[g^2] < \infty$. Define $\eta : (0, \infty) \rightarrow \mathbb{R}^2$ and $T : H \times \Theta^N \rightarrow \mathbb{R}^2$ as

$$\eta(\lambda) := \left(\lambda, -\frac{\lambda^2}{2} \right), \quad T(\alpha, d) := \left(\mathcal{R}_g(\alpha) - \hat{\mathcal{R}}_g(\alpha, d), \frac{1}{p(\alpha)^2 N} \mathbb{E}_{\mathfrak{G}} [g^2 \mathbf{1}_{\mathbf{L}_{g,\alpha}}] \right).$$

Then, for all $\lambda > 0$, it holds that $\mathbb{E}[c(\lambda, \mathfrak{D})] \leq 1$.

Proof The proof is very similar to the proof of Theorem 4.1 and basically uses the same reasoning. Let $\ell_g(\alpha, \theta) := \mathbf{1}_{\mathbf{L}_{g,\alpha}}(\theta)\ell(\alpha, \theta)$. Since \mathfrak{H} and \mathfrak{D} are independent, and $\exp \geq 0$, one gets from Fubini's theorem:

$$\mathbb{E} \left[\exp(\lambda(\mathcal{R}_g(\mathfrak{H}) - \hat{\mathcal{R}}_g(\mathfrak{H}, \mathfrak{D}))) \right] = \mathbb{E} \left\{ \mathbb{E} \left[\exp(\lambda(\mathcal{R}_g(\alpha) - \hat{\mathcal{R}}_g(\alpha, \mathfrak{D}))) \right] \Big|_{\alpha=\mathfrak{H}} \right\}.$$

Thus, first consider a fixed $\alpha \in H$ with $p(\alpha) > 0$. Then, by definition and the i.i.d. assumption, it holds that:

$$\begin{aligned} \mathbb{E} \left[\exp(\lambda(\mathcal{R}_g(\alpha) - \hat{\mathcal{R}}_g(\alpha, \mathfrak{D}))) \right] &= \mathbb{E} \left\{ \exp \left(-\frac{\lambda}{Np(\alpha)} \sum_{i=1}^N (\ell_g(\alpha, \mathfrak{G}_i) - \mathbb{E}_{\mathfrak{G}}[\ell_g(\alpha, \cdot)]) \right) \right\} \\ &= \prod_{i=1}^N \mathbb{E}_{\mathfrak{G}} \left\{ \exp \left(-\frac{\lambda}{Np(\alpha)} (\ell_g(\alpha, \cdot) - \mathbb{E}_{\mathfrak{G}}[\ell_g(\alpha, \cdot)]) \right) \right\}. \end{aligned}$$

$\ell_g(\alpha, \cdot)$ is non-negative, and by definition of the parametric sublevel set has a finite second-moment, that is $\mathbb{E}_{\mathfrak{G}}[\ell_g(\alpha, \cdot)^2] \leq \mathbb{E}_{\mathfrak{G}}[g^2 \mathbf{1}_{\mathbf{L}_{g,\alpha}}] < \infty$. Hence, by Lemma B.2 we have the inequality $\mathbb{E}_{\mathfrak{G}} \left\{ \exp \left(-\frac{\lambda}{Np(\alpha)} (\ell_g(\alpha, \cdot) - \mathbb{E}_{\mathfrak{G}}[\ell_g(\alpha, \cdot)]) \right) \right\} \leq \exp \left(\frac{\lambda^2}{2N^2 p(\alpha)^2} \mathbb{E} [\ell_g(\alpha, \cdot)^2] \right)$. Thus:

$$\mathbb{E} \left[\exp(\lambda(\mathcal{R}_g(\alpha, \cdot) - \hat{\mathcal{R}}_g(\alpha, \mathfrak{D}))) \right] \leq \exp \left(\frac{\lambda^2}{2Np(\alpha)^2} \mathbb{E}_{\mathfrak{G}} [g^2 \mathbf{1}_{\mathbf{L}_{g,\alpha}}] \right).$$

This can be rearranged into $\mathbb{E} \left\{ \exp \left(\lambda(\mathcal{R}_g(\alpha) - \hat{\mathcal{R}}_g(\alpha, \mathfrak{D})) - \frac{\lambda^2}{2Np(\alpha)^2} \mathbb{E}_{\mathfrak{G}} [g^2 \mathbf{1}_{\mathbf{L}_{g,\alpha}}] \right) \right\} \leq 1$, since the right-hand side is independent of \mathfrak{D} . As this holds for any α with $p(\alpha) > 0$, which in turn does hold $\mathbb{P}_{\mathfrak{H}}$ -a.s., we get $\mathbb{E} \left\{ \exp \left(\lambda(\mathcal{R}_g(\mathfrak{H}) - \hat{\mathcal{R}}_g(\mathfrak{H}, \mathfrak{D})) - \frac{\lambda^2}{2Np(\mathfrak{H})^2} \mathbb{E}_{\mathfrak{G}} [g^2 \mathbf{1}_{\mathbf{L}_{g,\alpha}}] \Big|_{\alpha=\mathfrak{H}} \right) \right\} \leq 1$. Changing the order of integration with Fubini's theorem, we get:

$$\mathbb{E} \left\{ \mathbb{E} \left[\exp \left(\lambda(\mathcal{R}_g(\mathfrak{H}) - \hat{\mathcal{R}}_g(\mathfrak{H}, d)) - \frac{\lambda^2}{2Np(\mathfrak{H})^2} \mathbb{E}_{\mathfrak{G}} [g^2 \mathbf{1}_{\mathbf{L}_{g,\alpha}}] \Big|_{\alpha=\mathfrak{H}} \right) \right] \Big|_{d=\mathfrak{D}} \right\} \leq 1.$$

Using the definition of η and T , this is the same as $\mathbb{E} \left\{ \mathbb{E} [\exp(\langle \eta(\lambda), T(\mathfrak{H}, d) \rangle)] \Big|_{d=\mathfrak{D}} \right\} \leq 1$. The inner term can be rewritten as $\mathbb{E} [\exp(\langle \eta(\lambda), T(\mathfrak{H}, d) \rangle)] = \int_H \exp(\langle \eta(\lambda), T(\alpha, d) \rangle) \mathbb{P}_{\mathfrak{H}}(d\alpha) = c(\lambda, d)$. Hence, in total we get $\mathbb{E}[c(\lambda, \mathfrak{D})] \leq 1$. \blacksquare

Remark 4.10 Since $\mathbb{P}_{\mathfrak{H}}\{p > 0\} = 1$, under the prior the algorithm should be able to “reach” the sublevel set. This is a constraint on the support (see Definition A.5) of $\mathbb{P}_{\mathfrak{H}}$, which is not satisfied without further ado. Section 5 provides a construction for achieving this.

4.2.1 IMPLEMENTING THE NON-DIVERGENCE – SPEED TRADE-OFF

In Section 4.2, care has to be taken in the choice of the prior $\mathbb{P}_{\mathfrak{H}}$: Just minimizing the upper bound as much as possible can lead to a neglect of a high sublevel probability, that is, the algorithm is especially fast on a small subset of the parameters, while it diverges for the rest, because the term $\frac{1}{p(\alpha)}$ might not compensate for the smaller sublevel risk. Thus, if a certain sublevel probability $\varepsilon_{\text{conv}} \in [0, 1]$ has to be ensured, one has to enforce it. We propose to use absolute continuity:

Lemma 4.11 *Let $\varepsilon_{\text{conv}} \in [0, 1]$ and assume that $p(\mathfrak{H}) \geq \varepsilon_{\text{conv}}$ a.s. Then, for every $\mathbb{Q} \in \mathcal{P}(\mathbb{P}_{\mathfrak{H}})$ we have $\mathbb{Q}\{p < \varepsilon_{\text{conv}}\} = 0$.*

Proof Since $\{p \geq \varepsilon_{\text{conv}}\} = \{p < \varepsilon_{\text{conv}}\}^c$, we have $\mathbb{P}_{\mathfrak{H}}\{p < \varepsilon_{\text{conv}}\} = 0$. Thus, the result follows by definition of absolute continuity. ■

Though the proof is trivial, this lemma has a very important consequence, which we want to stress: If one can guarantee that a required property is satisfied for the prior, it will be *preserved* during the PAC-Bayesian learning process, that is, if the prior only puts mass on hyperparameters that ensure a certain sublevel probability, the posterior will do the same. We will enforce this constraint in the learning procedure, which is discussed next.

5 Learning Procedure

This section deals with the learning procedure, that is, how the abstract framework discussed in Sections 3 and 4 is actually implemented. Hence, this is also the beginning of the second part of the paper, which is of a more practical kind and less theoretical. The resulting learning procedure is visualized in Figure 2 and consists of four steps:

- (i) Find an initialization that is trainable in a stable way: This is due to the fact that one might include, for example, a neural network in the update step of \mathcal{A} , which, if initialized randomly, might predict points at the very beginning, that are so far off that one encounters numerical instabilities. Hence, we first train the algorithm to “follow” another algorithm \mathcal{A}' , just for the purpose of non-divergence.
- (ii) Locate the prior by finding a point $\alpha^{(0)} \in H$ that satisfies the constraint in Subsection 4.2 and yields a good performance. For this, we perform a constrained version of stochastic empirical risk minimization with a new, specifically designed loss function.
- (iii) Starting from $\alpha^{(0)}$, construct the prior distribution by running a constrained version of a sampling algorithm.
- (iv) Compute the posterior distribution by finding the optimal $\lambda^* \in \Lambda$ and performing the reweighting with the closed form for \mathbb{Q}_{λ^*} .

Hence, first, in Subsection 5.1 we will identify the optimal posterior \mathbb{Q}^* in the abstract setting. Second, in Section 5.2, we realize the constraining procedure, which allows to keep an optimization or sampling procedure inside an abstractly defined set \mathbf{A} . Lastly, we put things together in Section 5.3 and finalize the learning procedure by concretizing each single step: Section 5.3.1 deals with the pre-computation phase in (i), Section 5.3.2 deals with the

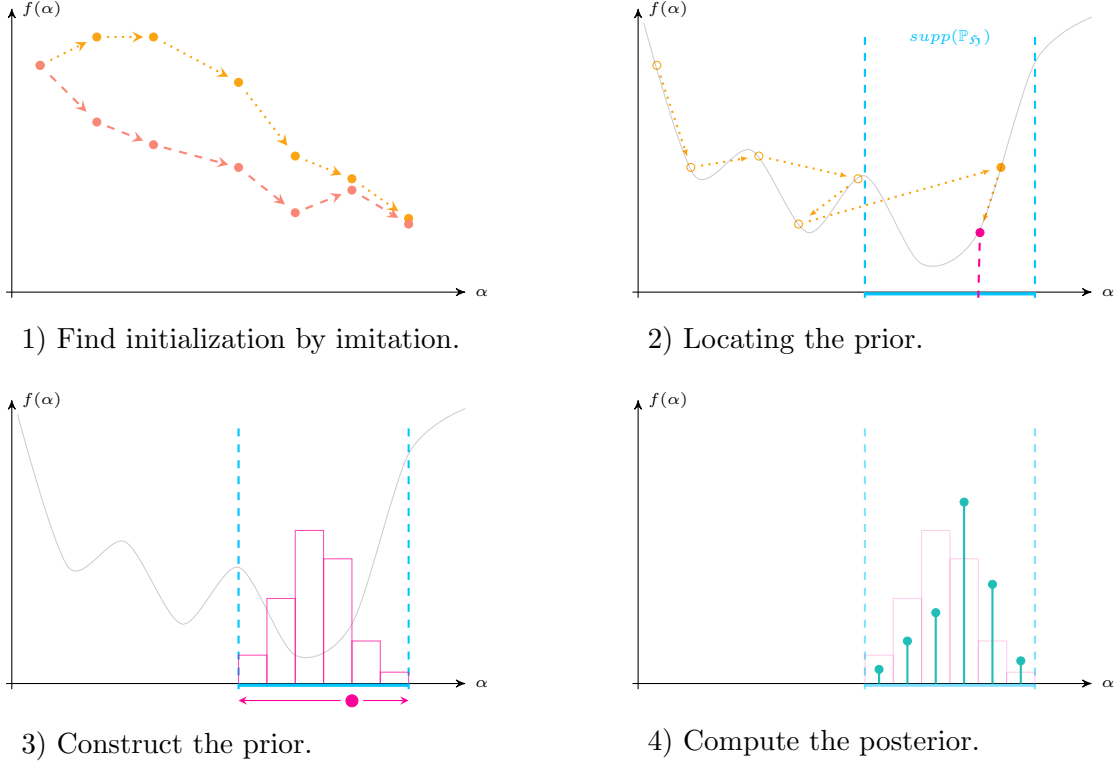


Figure 2: Learning procedure: **1)** Imitation learning. **2)** Probabilistically constrained stochastic empirical risk minimization. **3)** Construct prior through sampling. **4)** Compute posterior by performing the PAC-Bayesian learning step.

initialization in (ii), Section 5.3.3 provides the concrete choice of prior distribution in (iii), and Section 5.3.4 finally yields the posterior distribution in (iv).

5.1 Minimization of the PAC-Bound

Because learning is phrased as minimizing the PAC-Bayesian upper-bound, this is the first thing to do. Hence, in this subsection we consider η , T and item (ii) from Corollary 3.5. We seek for $\lambda \in \Lambda$ and $\mathbb{Q} \in \mathcal{P}(\mathbb{P}_S)$ that minimize the upper-bound in Corollary 3.5 (ii), that is, we want to solve:

$$\inf_{\lambda \in \Lambda} \inf_{\mathbb{Q} \in \mathcal{P}(\mathbb{P}_S)} \mathbb{Q}[\hat{\mathcal{R}}(\cdot, d)] + \frac{(D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}_S) + \log\left(\frac{K}{\varepsilon}\right) + C_{\mathcal{O}} - \mathbb{Q}[\langle \eta^{(r)}(\lambda), T^{(r)}(\cdot, d) \rangle])}{\eta^{(1)}(\lambda)}.$$

By factoring out $-\frac{1}{\eta^{(1)}(\lambda)}$ again, this is actually the same as:

$$\inf_{\lambda \in \Lambda} -\frac{1}{\eta^{(1)}(\lambda)} \left(\sup_{\mathbb{Q} \in \mathcal{P}(\mathbb{P}_S)} \mathbb{Q}[\langle \eta(\lambda), \tilde{T}(\cdot, d) \rangle] - D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}_S) - \log\left(\frac{K}{\varepsilon}\right) - C_{\mathcal{O}} \right),$$




Figure 3: Construction of $\tilde{\mathbb{P}}_{\mathfrak{X}}$: On the left, the set $C \subset X \times Y$ and two of its sections $C_x, C_{x'} \subset Y$ are visualized. On the right, the function $p(x) = \mathbb{P}_{\mathfrak{Y}}[C_x]$, the interval $[p_l, p_u]$, and the resulting support $\text{supp}(\tilde{\mathbb{P}}_{\mathfrak{X}})$ of $\tilde{\mathbb{P}}_{\mathfrak{X}}$ are visualized. Note that, contrary to the visualization here, p can actually be *highly discontinuous*.

where $\tilde{T}(\alpha, d) := (-\hat{\mathcal{R}}(\alpha, d), T^{(r)}(\alpha, d))$. Since $\log(K/\varepsilon) + C_{\mathcal{O}}$ is a constant, Lemma 3.1 shows that the term inside the brackets is given by $\tilde{\kappa}(\lambda, d) - \log(K/\varepsilon) - C_{\mathcal{O}}$, where $\tilde{\kappa}$ corresponds to the exponential family $(\tilde{\mathbb{Q}}_{\lambda})_{\lambda \in \Lambda}$ built upon \tilde{T} and η (with $h \equiv 1$). Furthermore, it shows that the (in this sense) optimal posterior distribution $\mathbb{Q} \in \mathcal{P}(\mathbb{P}_{\mathfrak{Y}})$ is given by the corresponding member of the data-dependent exponential family $\tilde{\mathbb{Q}}_{\lambda}(d) \propto \exp(\langle \eta(\lambda), \tilde{T}(\cdot, d) \rangle) \cdot \mathbb{P}_{\mathfrak{Y}}$, usually called the Gibbs posterior (Alquier, 2021). By denoting $F(\lambda, d) := -\frac{1}{\eta^{(1)}(\lambda)}(\tilde{\kappa}(\lambda, d) - \log(K/\varepsilon) - C_{\mathcal{O}})$, one is left with solving the following problem:

$$\inf_{\lambda \in \Lambda} F(\lambda, d), \quad (4)$$

which for $\Lambda \subset \mathbb{R}$ is one-dimensional. Based on Theorem 4.1 and Theorem 4.9, we have to restrict to $\Lambda \subset (0, +\infty)$, such that the solution to (4) can be seen as an approximation to the global minimum $\inf_{\lambda > 0} F(\lambda, d)$. For the latter one, one can show that the solution set lies in a compact interval $[\Lambda_{\min}, \Lambda_{\max}]$, with $\Lambda_{\min} > 0$, since $F(\lambda, d) \rightarrow \infty$ as $\lambda \rightarrow 0$ or $\lambda \rightarrow \infty$. Under our assumptions, $F(\cdot, d)$ is continuously differentiable. Hence, since Λ is compact, $F(\cdot, d)$ is Lipschitz-continuous on Λ and the minimum in (4) is attained. For a finite $\Lambda \subset [\Lambda_{\min}, \Lambda_{\max}]$, the optimization reduces to grid search. For $\Lambda = [\Lambda_{\min}, \Lambda_{\max}]$, we employ grid search as initialization for gradient-based optimization. However, we still might get stuck in a (close to optimal) local minimum. The computational bottleneck is given by evaluating $\lambda \mapsto \tilde{\kappa}(\lambda, d)$. In Sections 5.3.3 and 5.3.4 we will ensure that this is cheap.

5.2 Sampling under Probabilistic Constraints

In this section, we describe a methodology that allows for sampling from a distribution that is *probabilistically constrained* in the following sense: We are given two independent random variables $\mathfrak{X} : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow X$, $\mathfrak{Y} : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow Y$ taking values in the Polish spaces X and Y , with joint and marginal distributions $\mathbb{P}_{(\mathfrak{X}, \mathfrak{Y})}$, $\mathbb{P}_{\mathfrak{X}}$ and $\mathbb{P}_{\mathfrak{Y}}$, respectively. Further, we consider

a measurable set $C \subset X \times Y$, and we want to generate samples $\mathfrak{X} = x \in X$, such that the probability of $(\mathfrak{X}, \mathfrak{Y})$ lying in C , given $\mathfrak{X} = x$, takes values in a certain interval:

$$\mathbb{P}_{(\mathfrak{X}, \mathfrak{Y})|\mathfrak{X}=x}[C] = \mathbb{P}_{\mathfrak{Y}|\mathfrak{X}=x}[C_x] \in [p_l, p_u] \subset [0, 1].$$

By independence of \mathfrak{X} and \mathfrak{Y} , this is $\mathbb{P}_{\mathfrak{X}}$ -almost surely the same as $\mathbb{P}_{\mathfrak{Y}}[C_x] \in [p_l, p_u]$, and we will use the later formulation from now on. While having an abstract description of C , we do not have any *explicit* information about its geometrical or topological properties, let alone access to a distance or a projection onto it, and the marginal $\mathbb{P}_{\mathfrak{Y}}$ can only be accessed via samples $\mathfrak{Y}_1, \dots, \mathfrak{Y}_n, n \in \mathbb{N}$. However, since we are able to evaluate, for every $x \in X$ and $n \in \mathbb{N}$, whether $\mathfrak{Y}_n \in C_x$, this allows to define a function $p : X \rightarrow [0, 1]$, given by:

$$\begin{aligned} \mathbb{P}_{(\mathfrak{X}, \mathfrak{Y})}[C] &= (\mathbb{P}_{\mathfrak{X}} \otimes \mathbb{P}_{\mathfrak{Y}})[C] = \int_X \int_Y \mathbf{1}_C(x, y) \mathbb{P}_{\mathfrak{Y}}(dy) \mathbb{P}_{\mathfrak{X}}(dx) = \int_X \mathbb{P}_{\mathfrak{Y}}[C_x] \mathbb{P}_{\mathfrak{X}}(dx) \\ &=: \int_X p(x) \mathbb{P}_{\mathfrak{X}}(dx). \end{aligned}$$

Then p is indeed a measurable function, taking values in $[0, 1]$. Thus, for $p_l, p_u \in [0, 1]$ with $p_l < p_u$, we can define a measurable set $A := \{x \in X : \mathbb{P}_{\mathfrak{Y}}[C_x] \in [p_l, p_u]\}$, which yields a new measure $\tilde{\mathbb{P}}_{\mathfrak{X}}$ on X by restricting to A , that is, for a measurable set $B \subset X$ it holds:

$$\tilde{\mathbb{P}}_{\mathfrak{X}}[B] := ((\mathbf{1}_{[p_l, p_u]} \circ p) \cdot \mathbb{P}_{\mathfrak{X}})[B] = (\mathbf{1}_A \cdot \mathbb{P}_{\mathfrak{X}})[B] = \mathbb{P}_{\mathfrak{X}}[A \cap B].$$

Therefore, as stated before, we have the following goal:

Goal: Sample from $\tilde{\mathbb{P}}_{\mathfrak{X}}$, that is, get $\mathfrak{X}_1, \dots, \mathfrak{X}_K \sim \tilde{\mathbb{P}}_{\mathfrak{X}}$, such that $\mathbb{P}_{\mathfrak{Y}}[C_x]|_{x=\mathfrak{X}_i} \in [p_l, p_u]$.

This construction is depicted in Figure 3: The left figure visualizes the sections $\{C_x\}_{x \in X}$ of the set C , while the right figure shows the corresponding construction of the support of $\tilde{\mathbb{P}}_{\mathfrak{X}}$. Here, we assume that the imposed constraint is realizable.

Assumption 5.1 *The measure $\tilde{\mathbb{P}}_{\mathfrak{X}}$ has a non-empty support.*

Remark 5.2 *Assumption 5.1 includes two things: (i) There have to be points $x \in X$ (hyperparameters) for which the property in C does hold, that is, it has to be realizable. (ii) The support of $\mathbb{P}_{\mathfrak{X}}$ (prior) has to have a non-empty intersection with this region.*

Example 5.3 *Consider the random variables \mathfrak{S} and \mathfrak{H} from Section 4. By Lemma 4.11 we want to have $p(\mathfrak{H}) \in [\varepsilon_{\text{conv}}, 1]$, where the sublevel probability is given as $p(\alpha) = \mathbb{P}_{\mathfrak{S}}[\mathbf{L}_{\alpha}]$ (Lemma 4.7), and the sublevel set $\mathbf{L} \subset H \times \Theta$ is measurable by Lemma B.3. Thus, this corresponds to the identification $X = H, Y = \Theta$, and $p_l = \varepsilon_{\text{conv}}, p_u = 1$.*

5.2.1 INCORPORATION INTO A SAMPLING PROCEDURE

By definition it holds that $\mathbb{P}_{\mathfrak{X}} = \tilde{\mathbb{P}}_{\mathfrak{X}}$ on A . Hence, the only distinction between samples from $\tilde{\mathbb{P}}_{\mathfrak{X}}$ and samples from $\mathbb{P}_{\mathfrak{X}}$ is the restriction to the set A . Therefore, since many sampling algorithms sample from unnormalized measures anyways, it suffices to be able to sample from $\mathbb{P}_{\mathfrak{X}}$, if the restriction to A can be satisfied in other ways. Thus, we have to integrate

Algorithm 1 Iterative estimation of the probability p

Require: $p_l, p_u, \varepsilon \in [0, 1]$.

$a, b \leftarrow 1, 1$

▷ Initialize with uninformative prior.

while $Q_{a,b}(p_u) - Q_{a,b}(p_l) \geq \varepsilon$ **do**

▷ $Q_{a,b}$ is the quantile function for $\text{Beta}(a, b)$.

 Draw $\mathfrak{J} \sim \text{Ber}(p)$

$a \leftarrow a + \mathfrak{J}$ and $b \leftarrow b + (1 - \mathfrak{J})$

end while

this constraint into a sampling procedure for $\mathbb{P}_{\mathfrak{X}}$. Because we do not have any geometrical or topological information about the set \mathbf{C} , we have to resort to statistical information: We have access to i.i.d. samples $\mathfrak{Y}_1, \dots, \mathfrak{Y}_n \sim \mathbb{P}_{\mathfrak{Y}}$, and, for a given $x \in X$, we are able to evaluate the i.i.d. Bernoulli-distributed random variables $\mathfrak{J}_n := \mathbb{1}\{\mathfrak{Y}_n \in \mathbf{C}_x\}$, $n \in \mathbb{N}$. These have the parameter $\mathbb{P}\{\mathfrak{J}_n = 1\} = \mathbb{P}\{\mathfrak{Y}_n \in \mathbf{C}_x\} = \mathbb{P}_{\mathfrak{Y}}[\mathbf{C}_x] = p(x)$. Thus, by estimating $p(x)$ with an estimator $\hat{p}(x)$, we can approximate the constraint \mathbf{A} with $\hat{\mathbf{A}}$:

$$\mathbf{A} = \{x \in X : p(x) \in [p_l, p_u]\} \approx \{x \in X : \hat{p}(x) \in [p_l, p_u]\} =: \hat{\mathbf{A}}.$$

Given a sample $\mathfrak{X}_i \sim \mathbb{P}_{\mathfrak{X}}$, we resort to a simple accept-reject mechanism as in *Metropolis-Hastings*-type algorithms (Robert and Casella, 2004), to decide whether $\mathfrak{X}_i \in \mathbf{A}$, that is, whether \mathfrak{X}_i can actually be regarded as a sample from $\hat{\mathbb{P}}_{\mathfrak{X}}$.

Remark 5.4 *We believe that this is the only reasonable choice here, which keeps an iterative algorithm inside $\hat{\mathbf{A}}$. However, it does not provide a way into $\hat{\mathbf{A}}$, let alone \mathbf{A} .*

We estimate $p(x)$ in a Bayesian way, as it allows us to balance accuracy against computational complexity through uncertainty-quantification, which we use as a stopping criterion: We place a Beta-prior $\text{Beta}(a^{(0)}, b^{(0)})$ over the interval $[0, 1]$. As we do not have prior knowledge, and the map $x \mapsto p(x)$ can be discontinuous², we use a noninformative prior (Berger, 1985, Ch. 3.3), that is, $a_0 = b_0 = 1$. Since the Beta distribution is the conjugate prior for the Bernoulli distribution (Berger, 1985, p.130), that is, the posterior is again a Beta-distribution, after observing a sample \mathfrak{J}_{k+1} , the parameters $a^{(k)}, b^{(k)}$ get updated as:

$$a^{(k+1)} = a^{(k)} + \mathfrak{J}_{k+1}, \quad b^{(k+1)} = b^{(k)} + (1 - \mathfrak{J}_{k+1}).$$

Hence, we can perform the estimation iteratively: We only draw a new sample \mathfrak{J}_{n+1} as long as $Q^{(n)}(p_u) - Q^{(n)}(p_l) \geq \varepsilon$, where $Q^{(n)}$ denotes the quantile-function of the distribution $\text{Beta}(a^{(n)}, b^{(n)})$, and $p_u, p_l, \varepsilon \in [0, 1]$ are parameters that specify the accuracy of the estimation. Finally, one can use the posterior mean $\frac{a^{(n)}}{a^{(n)}+b^{(n)}}$ or posterior mode $\frac{a^{(n)}-1}{a^{(n)}+b^{(n)}-2}$ (provided $a^{(n)}, b^{(n)} > 1$) as point estimate \hat{p}_x . This procedure is summarized in Algorithm 1 and depicted in Figure 4.

Remark 5.5 *By adjusting p_l, p_u or ε , one can balance between accuracy and computational complexity. Yet, the number of iterations needed also depends crucially on the true probability: For $p \approx 0$ or $p \approx 1$, the uncertainty decreases significantly faster than for $p \approx 0.5$.*

2. Consider learning the step-size parameter $\alpha > 0$ for gradient descent on quadratic functions with largest eigenvalue L : The algorithm converges for $\alpha < \frac{2}{L}$ ($p(\alpha) = 1$) and diverges for $\alpha > \frac{2}{L}$ ($p(\alpha) = 0$).




Figure 4: Iterative estimation of $p(x)$: The black line shows the density $f_{a^{(n)}, b^{(n)}}$ of $\text{Beta}(a^{(n)}, b^{(n)})$ after having observed $\mathcal{I}_1, \dots, \mathcal{I}_n$. The red dotted line indicates the true probability p , which we are trying to estimate, and the blue dashed lines indicate the lower and upper quantiles q_l, q_u corresponding to p_l and p_u . The procedure stops as soon as $q_u - q_l < \varepsilon$, which is indicated by the double-headed arrow. Here, we use $q_l = 0.05, q_u = 0.95, \varepsilon = 0.15$ and $p = 0.8$.

5.2.2 BROADER CONTEXT

Different, yet conceptually similar ideas have appeared in the context of cutting the computational cost of Bayesian Markov Chain Monte Carlo algorithms through subsampling: Korattikara et al. (2014) introduce sequential hypothesis tests to reach the binary accept-reject decision in the Metropolis-Hastings algorithm on subsamples of the data set. Similarly, Bardenet et al. (2014) estimate the accept-reject step on a random subset of the data, in a way that guarantees to coincide with the true accept-reject step with a user-specified probability. Maclaurin and Adams (2014) introduce an auxiliary binary variable $z_n \in \{0, 1\}$, which allows to query only a subset of the data for the computation of the exact likelihood. And Quiroz et al. (2018) use a subsampling and bias-correction strategy to speed-up the sampling procedure. Here, Bardenet et al. (2017) provides a summary and discussion of different approximations and their biases (and errors). We leave the corresponding analysis for our proposed approximation to future work.

5.2.3 CHOICE OF THE SAMPLING PROCEDURE

Often, the hyperparameters $\alpha \in H$ of the learned algorithm \mathcal{A} are high-dimensional. Thus, we resort to using *stochastic gradient Langevin dynamics* (Welling and Teh, 2011) as the underlying sampling algorithm, which gets constrained to the set $\hat{\mathbf{A}}$ by use of the previously described procedure. This is summarized in Algorithm 2. However, if it fits the application, other sampling algorithms can be used, too. The computational overhead of the additional estimation depends on the cost of evaluating $\mathbf{1}\{\mathcal{Q}_n \in \mathbf{C}_x\}$. In our case it is rather expensive, as every sample \mathcal{I}_n requires to run the algorithm \mathcal{A} . However, this is to be expected, as the “prediction” of an optimization algorithm corresponds to approximating the solution of a minimization problem.

Remark 5.6 *Algorithm 2 requires to start in the set $\hat{\mathbf{A}}$. If such a point is not known, one can still run the algorithm and just “start” the accept-reject mechanism as soon as one has found a point $x \in \hat{\mathbf{A}}$. However, it is not guaranteed that such a point will actually be found.*

Algorithm 2 Probabilistically constrained sampling

Require: $p_l, p_u \in [0, 1]$, $n_{\max} \in \mathbb{N}$, $x_0 \in \hat{\mathcal{A}}$.

$n \leftarrow 0$ and $x \leftarrow x_0$

while $n \leq n_{\max}$ **do**

1) Draw a proposal x' with SGLD starting from x .

2) Estimate $p(x') = \mathbb{P}_{\mathfrak{H}}[C_{x'}]$ by $\hat{p}(x')$ with Algorithm 1.

if $\hat{p}(x') \in [p_l, p_u]$ **then**

$x \leftarrow x'$

else

Reject x' .

end if

end while

5.3 Putting Things Together

In the abstract learning procedure described in Sections 3 and 4 everything relies on the existence of a prior distribution $\mathbb{P}_{\mathfrak{H}}$ that a) gives a good performance of the algorithm in terms of a small risk, and b) ensures the user-specified sublevel probability. Both a) and b) have to be enforced through the *construction* of the prior $\mathbb{P}_{\mathfrak{H}}$, which is why often a *data-dependent* prior is used (Parrado-Hernández et al., 2012; Lever et al., 2013; Dziugaite and Roy, 2017, 2018; Dziugaite et al., 2021). Note that a) and b) are qualitatively different: a) asks for a certain location and concentration of $\mathbb{P}_{\mathfrak{H}}$, while b) puts hard constraints on its support, that is, it *excludes* many frequently used distributions with unconstrained support. Therefore, a significant part of the learning procedure deals with the *construction* of this prior only. Since the prior will be data-dependent, yet has to be independent of the data set that is used in the PAC-Bayesian learning procedure in Section 5.3.4, we have to split the data set \mathfrak{D} into independent parts $\mathfrak{D}_{\text{prior}}$, $\mathfrak{D}_{\text{val}}$, $\mathfrak{D}_{\text{train}}$ and $\mathfrak{D}_{\text{test}}$, where $\mathfrak{D}_{\text{prior}}$ and $\mathfrak{D}_{\text{val}}$ are used for the construction of the prior distribution, $\mathfrak{D}_{\text{train}}$ is used for the PAC-Bayesian learning step, and $\mathfrak{D}_{\text{test}}$ is the test set to evaluate the performance, that is, it is only needed for the experiments. Nevertheless, for notational simplicity, we will use the generic \mathfrak{D} , implicitly assuming the above partitioning. As the data set is *fixed* now, it will be described by the corresponding *realization* $D = \{\theta_1, \dots, \theta_N\}$ of the random variable $\mathfrak{D} = \{\mathfrak{S}_1, \dots, \mathfrak{S}_N\}$.

Remark 5.7 *Through the choice of the sampling algorithm, the concrete learning procedure described here mainly applies to the case $H = \mathbb{R}^d$, $d \in \mathbb{N}$. Nevertheless, the general methodology is still applicable to other Polish spaces, if this choice can be adjusted accordingly.*

5.3.1 FINDING A TRAINABLE INITIALIZATION

To increase numerical stability, we start with “imitation learning” (Chen et al., 2020a), that is, the algorithm \mathcal{A} should “follow” another algorithm \mathcal{A}' , for example, gradient descent. For this, we minimize the mean squared error between the iterates of the two algorithms: Given a starting point $x^{(0)} \in \mathbb{R}^n$, an iteration number $s \in \mathbb{N}$, and a parameter $\theta \in \Theta$, denote the first s iterates of $\mathcal{A}(\alpha, \theta, x^{(0)})$ by $x^{(1)}, \dots, x^{(s)} \in \mathbb{R}^n$ and the ones of $\mathcal{A}'(x^{(0)}, \theta)$ by

Algorithm 3 Procedure to find an initialization

Require: Data set $D_{\text{prior}}, x^{(0)} \in \mathbb{R}^n, s, n_{\text{init}} \in \mathbb{N}$ and $\varepsilon > 0$.

$m \leftarrow +\infty$ and sample $\theta \sim U_{D_{\text{prior}}}$.

while $\frac{1}{n_{\text{init}}}m > \varepsilon$ **do**

$m \leftarrow 0$

for $i = 1, \dots, n_{\text{init}}$ **do**

1) Compute $(x^{(1)}, y^{(1)}), \dots, (x^{(s)}, y^{(s)})$ with $\mathcal{A}(\alpha, \theta, x^{(0)})$ and $\mathcal{A}'(\theta, x^{(0)})$, resp.

2) Compute $\ell_{\text{init}}(\alpha, \theta, x^{(0)}, s) = \frac{1}{s} \sum_{k=1}^s \|x^{(k)} - y^{(k)}\|_2^2$.

3) Update $m \leftarrow m + \ell_{\text{init}}(\alpha, \theta, x^{(0)}, s)$

4) Update α by backpropagation and Adam. ▷ Other algorithms possible.

5) Update $\theta, x^{(0)}$ and s based on Section 5.3.2.

end for

end while

$y^{(1)}, \dots, y^{(s)} \in \mathbb{R}^n$. Then, define the loss as the mean squared error over these iterations:

$$\ell_{\text{init}}(\alpha, \theta, x^{(0)}, s) := \frac{1}{s} \sum_{k=1}^s \|x^{(k)} - y^{(k)}\|_2^2.$$

In each iteration, that is, each prediction of tuples $(x^{(1)}, y^{(1)}), \dots, (x^{(s)}, y^{(s)})$, the parameters $\theta, x^{(0)}$ and s are randomized as described in Section 5.3.2. We postpone the discussion of this step, as it is an integral part of the learning step described there, while being less important here. We terminate the process as soon as the average loss over $n_{\text{init}} \in \mathbb{N}$ iterations reaches a certain level of accuracy $\varepsilon > 0$, that is, as soon as $\frac{1}{n_{\text{init}}} \sum_{i=1}^{n_{\text{init}}} \ell_{\text{init}}(\alpha_i, \theta_i, x_i^{(0)}, s) < \varepsilon$. The choice of n_{init} and ε depends on the problem at hand. Usually, rough estimates suffice ($\varepsilon \leq 10$), as the purpose is to prevent divergence, and not actual imitation of \mathcal{A}' . The procedure is summarized in Algorithm 3.

5.3.2 LOCATING THE PRIOR

Empirically, the performance of the learned algorithm is significantly improved by the following two design choices, namely randomizing the trajectory length and using the ratio of successive losses. The motivation is to prevent overfitting and to learn a scale-independent contraction of the loss:

1) *Ratio of Losses:* Since the algorithm $\mathcal{A}(\alpha, \cdot)$ should achieve a small risk $\mathcal{R}(\alpha) = \mathbb{E}_{\mathfrak{G}}[\ell(\alpha, \cdot)]$, yet minimizing \mathcal{R} is intractable, the canonical loss function to be minimized is $\hat{\mathcal{R}}(\alpha, D_N) = \frac{1}{N} \sum_{i=1}^N \ell(\alpha, \theta_i)$. As H might be high-dimensional, we resort to stochastic empirical risk minimization, that is, in each iteration the observed loss would be of the form $\ell(\alpha, \theta_i)$. While this kind of loss was used extensively before, for learning-to-optimize it has a strong disadvantage: The overall outcome gets penalized only after the application of the *whole* algorithm, that is, after n_{max} iterations of the update-step. Thus, it does not take the trajectory into account. Further, often it is hard to minimize and does not lead to the desired performance. To circumvent this, Andrychowicz et al. (2016) proposed to use the sum over the incurred losses, that is, to consider $\tilde{\ell}_{\text{train}}(\alpha, \theta, x^{(0)}) := \sum_{i=1}^n \ell(x^{(i)}, \theta)$. Again, this formulation has a decisive flaw: Under most objectives, if the algorithm performs reasonably

Algorithm 4 Procedure to locate the prior

Require: Data sets $D_{\text{prior}}, D_{\text{val}}$, numbers $n_{\text{max}}, n, s \in \mathbb{N}$ with $s \leq n$, initialization $x^{(0)}$ and thresholds $p_l, p_u \in [0, 1]$ with $p_l < p_u$.

.....
 Set $x \leftarrow x^{(0)}$ and $b \leftarrow \mathbf{false}$, and sample $\theta \sim U_{D_{\text{prior}}}$ ▷ $b =$ Point inside constraint?

for $i = 1, \dots, n_{\text{max}}$ **do** ▷ Other stopping criteria possible.

1.a) Compute $x^{(1)}, \dots, x^{(s)}$ with $\mathcal{A}(\alpha, \theta, x^{(0)})$.

1.b) Compute $\ell_{\text{train}}(\alpha, \theta, x^{(0)}, s) = \sum_{i=1}^s \mathbb{1}_{\{\ell(x^{(i-1)}, \theta) > 0\}} \frac{\ell(x^{(i)}, \theta)}{\ell(x^{(i-1)}, \theta)}$.

1.c) Construct a proposal $\tilde{\alpha}$ by using backpropagation and Adam.

2) Estimate $p(\tilde{\alpha})$ by $\hat{p}(\tilde{\alpha})$ with Algorithm 1 on D_{val} .

if $\hat{p}(\tilde{\alpha}) \in [p_l, p_u]$ **then** ▷ If point inside constraint, just update.

$\alpha \leftarrow \tilde{\alpha}$ and $b \leftarrow \mathbf{true}$

else ▷ If not...

if $b = \mathbf{true}$ **then** ▷ ...reject moving outside constraint.

Reject $\tilde{\alpha}$, set $x^{(0)} \leftarrow x$, sample $\theta \sim U_{D_{\text{prior}}}$, and continue with 1).

else ▷ ...accept, if constraint has not been found yet.

$\alpha \leftarrow \tilde{\alpha}$

end if

end if

3) Draw $\tau \sim \text{Ber}(\frac{s}{n})$.

if $\tau = 0$ **then**

$x^{(0)} \leftarrow x^{(s)}$

else $\tau = 1$

$x^{(0)} \leftarrow x$ and sample $\theta \sim U_{D_{\text{prior}}}$

end if

end for

well, the loss at the beginning is *several orders* of magnitude larger than the loss at the end, that is, $\ell(x^{(1)}, \theta) \gg \ell(x^{(n)}, \theta)$. Hence, $\tilde{\ell}_{\text{train}}$ mainly penalizes the loss at the beginning, leading to an algorithm that minimizes the loss very fast in early iterations, but slows down *a lot* in later iterations. This is due to $\tilde{\ell}_{\text{train}}$ being *scale-sensitive*. Additionally, the incurred loss might vary strongly with the initialization $x^{(0)}$ alone, thereby introducing ambiguity into the incurred losses. We propose to use the *ratio of losses* of successive iterates:

$$\ell_{\text{train}}(\alpha, \theta, x^{(0)}, s) := \sum_{i=1}^s \mathbb{1}_{\{\ell(x^{(i-1)}, \theta) > 0\}} \frac{\ell(x^{(i)}, \theta)}{\ell(x^{(i-1)}, \theta)}, \quad s \in \mathbb{N}, \quad s \leq n.$$

This has several advantages: First, the loss is not scale-sensitive anymore, such that it favours hyperparameters that yield a good performance in each iteration. Second, there is

no ambiguity in the observed loss through the initialization, as the only criterion is a strong contraction of the loss (instead of a small loss). Intuitively, the algorithm \mathcal{A} gets trained to have a fast *convergence rate*. Third, the incurred losses do not vary too much, which empirically makes it easier to choose hyperparameters of the learning procedure. However, it also has a disadvantage: If the function values do indeed *converge* in a setting where the optimal loss is strictly greater than zero, this gets fully penalized, as then $\frac{\ell(x^{(i)}, \theta)}{\ell(x^{(i-1)}, \theta)} \equiv 1$. For now, we do not know how to get rid off this problem (apart from just stopping the iterations in case of convergence) while keeping the advantages, yet we (successfully) combat it with a decreasing step-size during training.

2) *Randomized Trajectory Length*: Training \mathcal{A} with fixed initialization $x^{(0)}$ and fixed trajectory length $n \in \mathbb{N}$, leads to overfitting, that is, applying \mathcal{A} at another starting point $\tilde{x}^{(0)}$, or applying it for more than n iterations typically leads to divergence. To get rid off this, we propose the following randomization: Given $\mathbb{N} \ni s \leq n$, set $p := \frac{s}{n}$ and $y := x^{(0)}$.

- 0) Sample a parameter θ uniformly at random from D_N .
- 1) Compute $x^{(1)}, \dots, x^{(s)}$ with $\mathcal{A}(\alpha, \theta, y)$ and the loss $\ell_{\text{train}}(\alpha, \theta, y, s)$, and update α .
- 2) Sample $\mathfrak{r}^{(k)} \sim \text{Ber}(p)$. If $\mathfrak{r}^{(k)} = 0$, set $y := x^{(s)}$ and go to step 1). If $\mathfrak{r}^{(k)} = 1$, set $y := x^{(0)}$ and go to step 0).

The random variable $\mathfrak{r}^{(k)}$ decides whether the algorithm gets restarted from $x^{(0)}$ with a new parameter $\tilde{\theta}$, or if one continues the current trajectory. To update α , we use backpropagation and Adam. The choice $p = \frac{s}{n}$ ensures that the expected trajectory-length equals n : Define $\tau := \inf\{k \in \mathbb{N} : \mathfrak{r}^{(k)} = 1\}$. Then, $\tau \sim \text{Geo}(p)$ is a geometrically distributed with expectation $\mathbb{E}[\tau] = \frac{1}{p}$. Therefore, the actual length $\mathfrak{L} = s \cdot \tau$ of the trajectory is (basically) a geometrically distributed random variable with $\mathbb{E}[\mathfrak{L}] = s\mathbb{E}[\tau] = s\frac{1}{p} = s\frac{n}{s} = n$.

Remark 5.8 (i) *Similarly to Andrychowicz et al. (2016), during training we omit the computation of second order derivatives. Additionally, and surprisingly, it usually suffices to consider single iterates, that is $s = 1$, where the iterates are treated independently in the computational graph. That amounts to learning an update step that is agnostic to the recurrent nature of the optimization algorithm and just learns to adapt to the local geometry of the loss function along the iterations.*

(ii) *Splitting the trajectory into several parts has been used before, especially in model-free approaches based on recurrent neural networks (Andrychowicz et al., 2016; Chen et al., 2017; Metz et al., 2019), where the main motivation are training instabilities and memory bottlenecks. Similarly, randomizing the trajectory length did appear in Wichrowska et al. (2017). Yet, to our knowledge, our specific randomization is new, theoretically motivated, and solves the problem of generalizing to more iterations quite effectively.*

Empirically, by using these design choices we are able to train \mathcal{A} in a way such that it performs equally well on several orders of magnitude of the loss function, and that it can be used for more iterations than it was actually trained for. The procedure is summarized in Algorithm 4 and consists of three steps:

Algorithm 5 Procedure to construct the prior

Require: Data sets D_{prior} (sampling) and D_{val} (constraint), $n_s \in \mathbb{N}$ and $\alpha \in \text{supp}(\tilde{\mathbb{P}}_{\mathfrak{H}})$.

- 1) Starting from α , run Algorithm 2 (with ℓ_{train} , and randomized trajectory length) to get the points $\alpha_1, \dots, \alpha_{n_s} \in H$.
 - 2) Evaluate φ_{prior} on $\{\alpha_1, \dots, \alpha_{n_s}\}$ by evaluating $\hat{\mathcal{R}}_g$ corresponding to D_{prior} .
 - 3) Compute $\mathbb{P}_{\mathfrak{H}}\{\alpha_j\}$, that is, $\mathbb{P}_{\mathfrak{H}}\{\alpha_j\} = \sigma(\varphi_{\text{prior}}(\alpha_1), \dots, \varphi_{\text{prior}}(\alpha_{n_s}))_j$.
-

- 1) Given $\theta, x^{(0)}$ and $s \in \mathbb{N}$, compute a new proposal $\tilde{\alpha}$ by performing one step of Adam on $\ell_{\text{train}}(\alpha, \theta, x^{(0)}, s)$, corresponding to s iterations of \mathcal{A} .
- 2) Check for the constraint $p(\tilde{\alpha}) \in [p_l, p_u]$ by estimating it with Algorithm 1. If the constraint is satisfied (or one did not find any point inside the constraint so far), accept $\tilde{\alpha}$. Otherwise, reject it and sample a new problem instance $\tilde{\theta} \sim U_{D_N}$.
- 3) If $\tilde{\alpha}$ got accepted, update θ and $x^{(0)}$ by randomizing the length of the trajectory.

5.3.3 CONSTRUCTING THE PRIOR

Besides the performance and the sublevel guarantees, the only assumption on the prior $\mathbb{P}_{\mathfrak{H}}$ is its independence of $\mathfrak{D}_{\text{train}}$. Further, by Lemma 3.1 the functional form of the posterior is fully specified, namely it is of the form:

$$\mathbb{Q}_{\lambda}(d) \propto \exp(\varphi_{\lambda}(\cdot, d)) \cdot \mathbb{P}_{\mathfrak{H}}, \quad \lambda \in \Lambda, \quad (5)$$

where the potential is given by $\varphi_{\lambda}(\alpha, d) = \langle \eta(\lambda), \tilde{T}(\alpha, d) \rangle$. Hence, for mathematical convenience, we will construct $\mathbb{P}_{\mathfrak{H}}$ by approximating the distribution \mathbb{P}' given by

$$\mathbb{P}' \propto \exp\left(-\hat{\mathcal{R}}_{g,\text{prior}} - \iota_{[p_l, p_u]} \circ p\right) \cdot \mu,$$

where μ denotes another measure on H , which allows to sample from \mathbb{P}' (possibly unnormalized). In our experiments it holds $H = \mathbb{R}^d$ and we choose $\mu = \lambda^d$, where λ^d is the d -dimensional Lebesgue measure. For sampling, we use (in principal) the *Stochastic Gradient Langevin Dynamics* algorithm. However, for ease of implementation and its practicability, we use the output of the backpropagation algorithm as proxy for the (sub)gradient, which is an element of a so-called *conservative set-valued field* (Bolte and Pauwels, 2021), which can be shown (under mild assumptions) to be λ^d -almost everywhere the same as the true gradient (Bianchi et al., 2022). Since we have to resort to a sampling algorithm to get points $\alpha_1, \dots, \alpha_{n_s} \in H$, $n_s \in \mathbb{N}$, we define the prior distribution directly as a discrete distribution, that is $\mathbb{P}_{\mathfrak{H}}\{\alpha\} := \frac{1}{Z} \sum_{i=1}^{n_s} a_i \delta_{\alpha_i}\{\alpha\}$. Thus, $\mathbb{P}_{\mathfrak{H}}$ is the suitable normalized discrete measure on H corresponding to $\alpha_1, \dots, \alpha_n$, where the coefficients and the normalization constant are given by $Z = \sum_{i=1}^{n_s} a_i$ with $a_i = \exp\left(-\hat{\mathcal{R}}_{g,\text{prior}}(\alpha_i) - \iota_{[p_l, p_u]}(\hat{p}(\alpha_i))\right)$. When $\alpha_1, \dots, \alpha_{n_s} \in H$ are given, the corresponding probabilities can equivalently be expressed with the so-called *softmax* function $\sigma(x_1, \dots, x_n)_j = \frac{\exp(x_j)}{\sum_{i=1}^{n_s} \exp(x_i)}$: Define the potential $\varphi_{\text{prior}} : H \rightarrow \mathbb{R}$ as $\varphi_{\text{prior}}(\alpha) = -\hat{\mathcal{R}}_{g,\text{prior}}(\alpha) - \iota_{[p_-, p_+]}(\hat{p}(\alpha))$ to get:

$$\mathbb{P}_{\mathfrak{H}}\{\alpha_j\} = \frac{\exp(\varphi_{\text{prior}}(\alpha_j))}{\sum_{i=1}^{n_s} \exp(\varphi_{\text{prior}}(\alpha_i))} = \sigma(\varphi_{\text{prior}}(\alpha_1), \dots, \varphi_{\text{prior}}(\alpha_{n_s}))_j.$$

Algorithm 6 Procedure to construct the posterior

Require: Points $\{\alpha_1, \dots, \alpha_{n_s}\}$, values $\{\varphi_{\text{prior}}(\alpha_1), \dots, \varphi_{\text{prior}}(\alpha_{n_s})\}$, data set D_{train} .

- 1) Evaluate $\tilde{T}(\alpha_i, d)$, $i = 1, \dots, n_s$, on D_{train} .
 - 2) Setup $\{\varphi_\lambda(\alpha_1, d), \dots, \varphi_\lambda(\alpha_{n_s}, d)\}$ as functions in λ .
 - 3) Solve $\lambda^* \in \arg \min_{\lambda \in \Lambda} F(\lambda, d)$. $\triangleright F(\lambda^*, d)$ is the predicted PAC-bound.
 - 4) Compute $\mathbb{Q}_{\lambda^*}(d, \{\alpha_j\}) = \sigma(\varphi_{\lambda^*}(\alpha_1, d), \dots, \varphi_{\lambda^*}(\alpha_{n_s}, d))_j$, $j = 1, \dots, n_s$.
 - 5) Optional: Choose $\alpha^* = \arg \max_{i=1, \dots, n_s} \mathbb{Q}_{\lambda^*}(d, \{\alpha_i\})$ as final point-estimate.
-

Hence, the prior distribution is given by evaluating the softmax-function, where the potentials φ_{prior} have to be computed *only once* for every α_i , $i = 1, \dots, n_s$. This is summarized in Algorithm 5. Note that for sampling $\alpha_1, \dots, \alpha_{n_s}$, the loss function ℓ_{train} is used again.

Remark 5.9 *As one would approximate the intractable integrals with Monte-Carlo estimates anyway, choosing a discrete measure is less restrictive than it seems, and it has the additional advantage of allowing for exact instead of approximate inference.*

5.3.4 COMPUTING THE POSTERIOR

Given a discrete prior $\mathbb{P}_{\mathfrak{F}}$, every posterior $\mathbb{Q} \in \mathcal{P}(\mathbb{P}_{\mathfrak{F}})$ is also discrete with the same support $\{\alpha_1, \dots, \alpha_{n_s}\}$. Then, by the closed-form solution (5) for the supremum, for every $\lambda \in \Lambda$ the optimal posterior $\mathbb{Q}_\lambda(d)$ is given by:

$$\mathbb{Q}_\lambda(d, \{\alpha_j\}) = \frac{\exp(\langle \eta(\lambda), \tilde{T}(\alpha_j, d) \rangle + \varphi_{\text{prior}}(\alpha_j))}{\sum_{i=1}^{n_s} \exp(\langle \eta(\lambda), \tilde{T}(\alpha_i, d) \rangle + \varphi_{\text{prior}}(\alpha_i))} = \sigma(\varphi_\lambda(\alpha_1, d), \dots, \varphi_\lambda(\alpha_{n_s}, d))_j,$$

with the potential $\varphi_\lambda : H \times \Theta^N \rightarrow \mathbb{R}$, $(\alpha, d) \mapsto \langle \eta(\lambda), \tilde{T}(\alpha, d) \rangle + \varphi_{\text{prior}}(\alpha)$. Thus, to get the distribution $\mathbb{Q}_\lambda(d)$ as a function of λ , one has to compute $\tilde{T}(\alpha_i, d)$ *only once* for every $i = 1, \dots, n_s$, such that it can be evaluated with the softmax function. Hence, the only missing ingredient is the optimal $\lambda^* \in \Lambda$, which is found as described in Section 5.1. After evaluating the potentials $\varphi_\lambda(\cdot, d)$, which has to be done anyways, evaluating $\tilde{\kappa}(\cdot, d)$ in λ is cheap. The process is summarized in Algorithm 6.

6 Experiments

We consider the *smooth and strongly convex* problem of minimizing quadratic functions with varying strong convexity and smoothness constants, a *high-dimensional* image processing problem, the *non-smooth* Lasso problem, and the *non-smooth and non-convex* problem of training a neural network. We use the following training procedure in all experiments: $N = N_{\text{prior}} + N_{\text{train}} + N_{\text{val}} + N_{\text{test}}$ denotes the total number of datapoints, and we use $N_{\text{prior}} = \dots = N_{\text{test}} = 250$. (Sub)Gradients are defined by the output of backpropagation as it is implemented in PyTorch (Paszke et al., 2019), and we use $g(\theta) := a\ell(x^{(0)}, \theta)^b$, $a, b > 0$, to define the sublevel set \mathbf{L}_g . In Algorithm 1, we use $p_l = 0.95$, $p_u = 1.0$, $q_l = 0.01$, $q_u = 0.99$, and $\varepsilon = 0.075$. Thus, the algorithm should reach \mathbf{L}_g in at least 95% of the cases, and for the estimation of the sublevel probability it should concentrate 98% of the mass within a distance of 0.075. In Algorithm 2, we use *stochastic gradient Langevin dynamics* to

draw 10^2 samples, where we decay the step-size starting from $\alpha = 10^{-6}$. In Algorithm 3, we use Adam with an initial step-size of 10^{-3} , which gets reduced by a factor of 0.5 every 200 iterations, until an accuracy of $\varepsilon = 10^{-2}$ is reached, or for at most $n_{\text{init}} = 10^3$ iterations. In Algorithm 4, we use Adam with an initial step-size of 10^{-4} , which gets reduced by a factor of 0.5 every $2 \cdot 10^4$ iterations, for a total of $n_{\text{max}} = 2 \cdot 10^5$ iterations. We use a trajectory length of $s = 1$, that is, only single points, and update the constraint only every $2 \cdot 10^4$ iterations (with a reset to previous iterates, if we have left the set \hat{A}). In Algorithm 6, we use a finite Λ with $|\Lambda| = 75 \cdot 10^3$, and an accuracy (of the PAC-bound) of $\varepsilon = 0.05$. As we contrast the learned algorithm to first-order methods, in each iteration \mathcal{A} has access to iterates, (sub)gradients, and function values, and the update is solely based on these. Here, we perform preprocessing: The (sub)gradient is split into its norm $\|\nabla\ell(x^{(k)})\|$ and the corresponding unit vector $d_1^{(k)}$. Further, the norm is transformed to $n_1^{(k)} := \log(1 + \|\nabla\ell(x^{(k)})\|)$ to be less scale-sensitive. The iterates $x^{(k)}, x^{(k-1)}$ are combined into the momentum term $m^{(k)} := x^{(k)} - x^{(k-1)}$, which also is split into the unit vector $d_2^{(k)}$ and the transformed norm $n_2^{(k)}$. In the evaluation, we will always show the *loss over the iterations* in the upper left plot, the *performance in terms of computation time* in the upper right plot, the *loss histogram with predicted PAC-bound* in the lower left plot, and the *final estimate for the sublevel probability*, that is, whether the probabilistically constrained optimization/sampling procedure did work correctly, in the lower right plot. Finally, note that we always show the performance of a single sample α^* , and not the mean performance.

Remark 6.1 (i) *We always use the output of the backpropagation algorithm instead of exact (sub-)gradients, that is, the learned algorithms do not rely on smoothness.*

(ii) *We use 100 samples only, as they are very costly: To evaluate the potentials φ_{prior} and $\varphi_{\lambda}(\cdot, d)$ on a single sample $\alpha \in H$, one has to compute all losses $\ell(\alpha, \theta_i)$, $i = 1, \dots, N_{\text{prior}} + N_{\text{train}}$, that is, “solving” $N_{\text{prior}} + N_{\text{train}}$ optimization problems.*

6.1 Quadratics

As first problem we consider strongly convex quadratic functions with varying strong convexity, varying smoothness and varying right-hand side, that is, each optimization problem is of the form:

$$\min_{x \in \mathbb{R}^n} \frac{1}{2} \|Ax - b\|^2, \quad A \in \mathbb{R}^{n \times n}, b \in \mathbb{R}^n.$$

Thus, the parameters are given by $\theta = (A, b) \in \mathbb{R}^{n^2+n} =: \Theta$, while the optimization variable is $x \in \mathbb{R}^n$, where we use $n = 200$.

Construction of the Parameters. To control the strong-convexity and smoothness of ℓ , we specify intervals $[m_-, m_+], [L_-, L_+] \subset (0, +\infty)$, and sample $m_1, \dots, m_N \stackrel{iid}{\sim} \mathcal{U}_{[m_-, m_+]}$, $L_1, \dots, L_N \stackrel{iid}{\sim} \mathcal{U}_{[L_-, L_+]}$. Then, the matrices A_j , $j = 1, \dots, N$, are created as *diagonal matrices* with entries $a_{ii}^j = \sqrt{m_j} + i \cdot \frac{\sqrt{L_j} - \sqrt{m_j}}{n}$, $i = 1, \dots, n$, that is, we linearly interpolate from $\sqrt{m_j}$ to $\sqrt{L_j}$. Hence, the matrix $A_j^T A_j$ has smallest and largest eigenvalue m_j and L_j , respectively. To change the relative position between the ellipsoid of the quadratic and the




Figure 5: Update step of \mathcal{A} for quadratic problems: The directions $d_1^{(k)}$, $d_2^{(k)}$ and $d_1^{(k)} \odot d_2^{(k)}$ are inserted as different channels into the Conv2d -block, which performs 1×1 “convolutions”, that is, the algorithm acts coordinate-wise on the input. The scales $n^{(1)}$, $n^{(2)}$ get transformed separately by the fully-connected block.

initialization, we randomize the right-hand side by sampling $b_1, \dots, b_N \stackrel{iid}{\sim} \mathcal{N}(\mu, \Sigma)$, where we create μ and $\Sigma = C^T C$ by sampling $\mu_i, C_{i,k} \stackrel{iid}{\sim} \text{U}_{[-5,5]}$, $i, k = 1, \dots, n$.

Baseline. Assuming that it is not feasible to compute L_θ and m_θ for each problem instance ℓ_θ separately, the given class of functions is L_+ -smooth and m_- -strongly convex. Hence, as baseline we use *heavy-ball with friction* (HBF) (Polyak, 1964), whose update is given by $x^{(k+1)} = x^{(k)} - \tau \nabla f(x^{(k)}) + \beta (x^{(k)} - x^{(k-1)})$, where the optimal worst-case convergence rate is attained for $\tau = \left(\frac{2}{\sqrt{L_+ + \mu_-}} \right)^2$, $\beta = \left(\frac{\sqrt{L_+} - \sqrt{\mu_-}}{\sqrt{L_+} + \sqrt{\mu_-}} \right)^2$ (Nesterov, 2018).

Algorithm. The algorithmic update of the learned algorithm \mathcal{A} is visualized in Figure 5 and consists of two blocks: The update-block combines the gradient direction $d_1^{(k)}$, the momentum direction $d_2^{(k)}$, and their “interaction” $d_1^{(k)} \odot d_2^{(k)}$ into the new update-direction $d^{(k)}$, while the other block computes a step-size based on the corresponding logarithmically transformed norms $n^{(1)}$ and $n^{(2)}$. Interestingly, the theoretically redundant concatenation of two linear layers (1x1 convolutional layers, resp.) does seem to increase training stability.

Results. The upper left plot of Figure 6 shows that the learned algorithm outperforms HBF by orders of magnitude and, despite being trained for $n_{\text{train}} = 350$ iterations, one can use it for at least 700 iterations in most cases ($\geq 90\%$). However, the mean indicates that there are single instances for which instabilities occur, and, by comparing it to the median, one observes that the mean is far from being representative of the “typical” performance. Further, the algorithm performs well on very different orders of magnitude, ranging from $1e5$ to $1e-20$. The upper right plot confirms that also in terms of computation time the learned algorithm is way faster than heavy-ball, and the lower left plot shows that the predicted PAC-bound is quite tight. Lastly, the lower right plot shows that the algorithm will not converge in all of the cases, yet satisfies the specified constraints p_l and p_u .




Figure 6: **Upper left:** The dashed lines represent the mean losses, while the dotted lines represent the median losses over the iterations $x^{(i)}$, $i = 0, \dots, 700$. The shaded regions represent the 10th to 90th percentile (80 % of the test data). The learned algorithm \mathcal{A} is shown in blue, while HBF is shown in red. **Upper right:** The different lines indicate the cumulative computation time the two algorithms need to solve all the test problems up to a certain accuracy (in function values) measured by $\ell(x^{(i)}, \theta) < \varepsilon$. However, note that, since the learned algorithm is not guaranteed to stop by itself, both algorithms are run for maximally $n_{\max} = 1e4$ iterations. **Lower left:** Loss histogram after $n_{\text{train}} = 350$ iterations with predicted PAC-bound. **Lower right:** Estimated sublevel probability. The orange dashed line shows the point estimate, while the orange solid line shows the Beta-posterior arising from a non-informative prior. The gray dotted lines indicate the constraints p_l, p_u and show the feasible region as green line.

6.2 Image Processing

We consider (gray-scale) *image denoising/deblurring* with a regularizer given as smooth approximation to the L_1 -norm of the image derivative, that is, problems of the form:

$$\min_{x \in \mathbb{R}^n} \frac{1}{2} \|Ax - b\|^2 + \lambda \sum_{i,j=1}^n \sqrt{(D_h x)_{i,j}^2 + (D_w x)_{i,j}^2 + \varepsilon^2} \quad \lambda \in \mathbb{R}, A, D_h, D_w \in \mathbb{R}^{n \times n}, b \in \mathbb{R}^n.$$

The matrix A describes the “blurring” of the image, while D_h and D_w are the discrete image derivatives in h- and w-direction, respectively, which are used to penalize local changes in

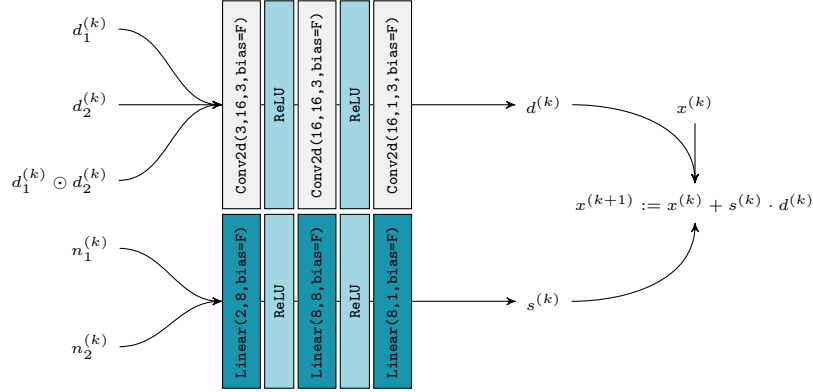


Figure 7: Update step of \mathcal{A} for the image-processing problems: The directions $d_1^{(k)}$, $d_2^{(k)}$ and $d_1^{(k)} \odot d_2^{(k)}$ are inserted as different channels (in the shape of the image) into the Conv2d-block, which performs a 3×3 -convolution. The scales $n^{(1)}$, $n^{(2)}$ get transformed separately by the fully-connected block.

the image. We use images of height $N_h = 250$ and width $N_w = \text{int}(0.75 \cdot N_h) = 187$. Thus, the dimension n of the optimization space is given by $n = 46750$. Further, as parameters θ we use the observed image and the regularization parameter, that is, $\theta = (b, \lambda) \in \mathbb{R}^{n+1} =: \Theta$.

Construction of the Parameters. Throughout, we use $\varepsilon = 0.01$. For computational efficiency, the matrices A, D_h, D_w are implemented through the convolution of the image x with a corresponding kernel (with reflective boundary conditions). For A , we use a standard (5×5) -Gaussian kernel, while D_h and D_w are given through the kernels:

$$k_h = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 1 & 0 \end{pmatrix} \in \mathbb{R}^{3 \times 3} \quad \text{and} \quad k_w = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -1 & 1 \\ 0 & 0 & 0 \end{pmatrix} \in \mathbb{R}^{3 \times 3}.$$

Additionally, after blurring an image with A , we add centered Gaussian noise $\varepsilon_{i,j}$ with standard deviation $\sigma = \frac{25}{256}$ to each pixel, that is, $b_{i,j} = (Ax^*)_{i,j} + \varepsilon_{i,j}$ with $\varepsilon_{i,j} \stackrel{iid}{\sim} \mathcal{N}(0, \sigma)$, $i = 1, \dots, N_h, j = 1, \dots, N_w$. The regularization parameters $\lambda_i \in \mathbb{R}$, $i = 1, \dots, N$, are given by sampling uniformly, that is, $\lambda_i \stackrel{iid}{\sim} U_{[\lambda_-, \lambda_+]}$, where we use $\lambda_- = 0.05$ and $\lambda_+ = 0.5$.

Baseline. Since the problem is smooth and convex, yet not strongly convex, the baseline algorithm is given by the *accelerated gradient descent* algorithm due to Nesterov (1983). Its update is given by first computing $y^{(k+1)} = x^{(k)} + \frac{t_k - 1}{t_{k+1}}(x^{(k)} - x^{(k-1)})$ followed by setting $x^{(k+1)} = y^{(k)} - \tau \nabla f(y^{(k+1)})$. We use the optimal choices $t_{k+1} = \frac{1}{2} \left(1 + \sqrt{1 + 4t_k^2} \right)$ and $\tau = \frac{1}{L}$. Here, the smoothness constant L is given by the largest eigenvalue of $A^T A + \frac{\lambda}{\varepsilon} D^T D$, where $D \in \mathbb{R}^{2n \times n}$ is given by “stacking” D_h and D_w , that is, $D = (D_h \ D_w)^T$.

Algorithm. The algorithmic update of \mathcal{A} is visualized in Figure 7 and consists of an update-block, which combines $d_1^{(k)}$, $d_2^{(k)}$ and their “interaction” $d_1^{(k)} \odot d_2^{(k)}$ into the new

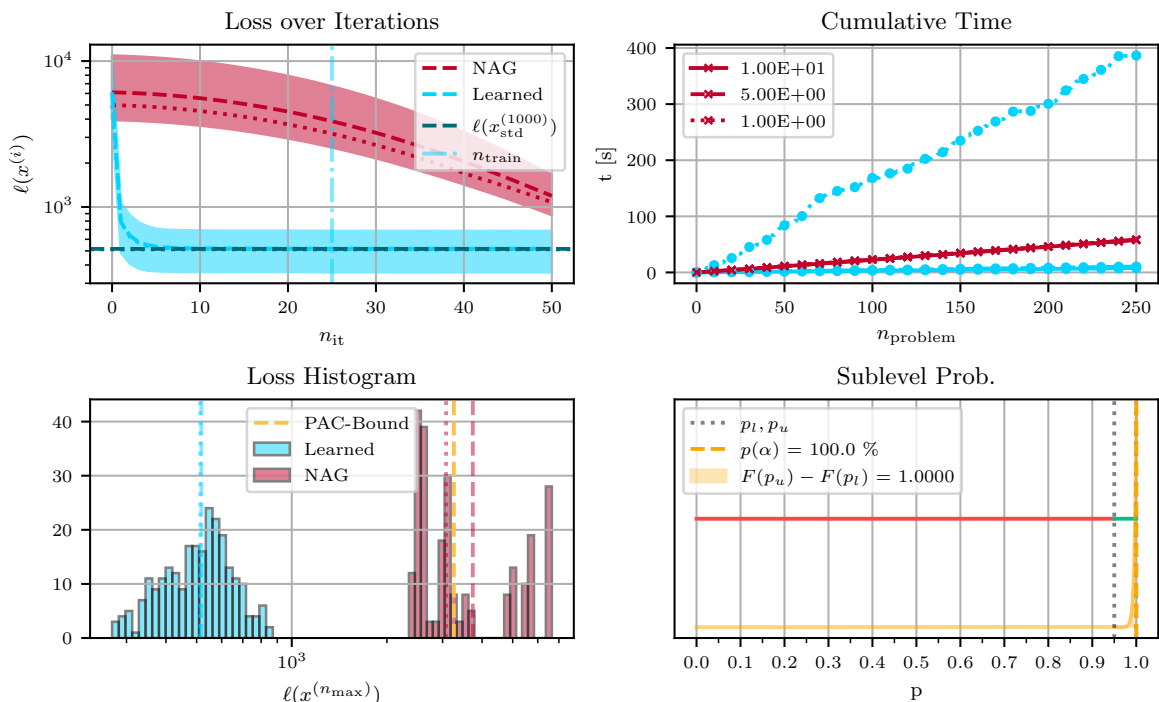


Figure 8: **Upper left:** Dashed lines represent mean losses, while dotted lines show median losses over the iterations $x^{(i)}$, $i = 0, \dots, 50$. The shaded regions represent the 10th to 90th percentile (80 % of the test data). The learned algorithm \mathcal{A} is shown in blue, while NAG is shown in red. **Upper right:** The different lines show the cumulative computation time the two algorithms need to solve all the test problems up to a certain accuracy (in function values) measured by $\ell(x^{(i)}, \theta) - \ell(x_{\text{std}}^{(1000)}, \theta) < \varepsilon$. However, note that both algorithms are run for maximally $n_{\text{max}} = 1000$ iterations. **Lower left:** Loss histogram after $n_{\text{train}} = 50$ iterations with predicted PAC-bound. **Lower right:** Estimated sublevel probability: The orange dashed line shows the point estimate, while the orange solid line shows the Beta-posterior arising from a non-informative prior. The gray dotted lines indicate to the constraints p_l, p_u and show the feasible region as green line.

update direction $d^{(k)}$, and a block to compute a step-size from the norms of the gradient and momentum term. Note that we use 3×3 -convolutions this time, that is, we incorporate the knowledge about an image-processing problem into the design of the optimization algorithm.

Results. The results of this experiment are summarized in Figure 8. The upper left plot shows that the learned algorithm is much faster than NAG in terms of the loss over the iterations, reaching a loss close to the ground-truth after only 5 iterations. The upper right plot confirms this finding also in terms of computation time. Yet, one can observe a strong increase in computation time for the dotted line (loss per pixel of about $\frac{1}{46750}$), indicating that the learned algorithm often is not able to reach this accuracy. In the lower left plot, one

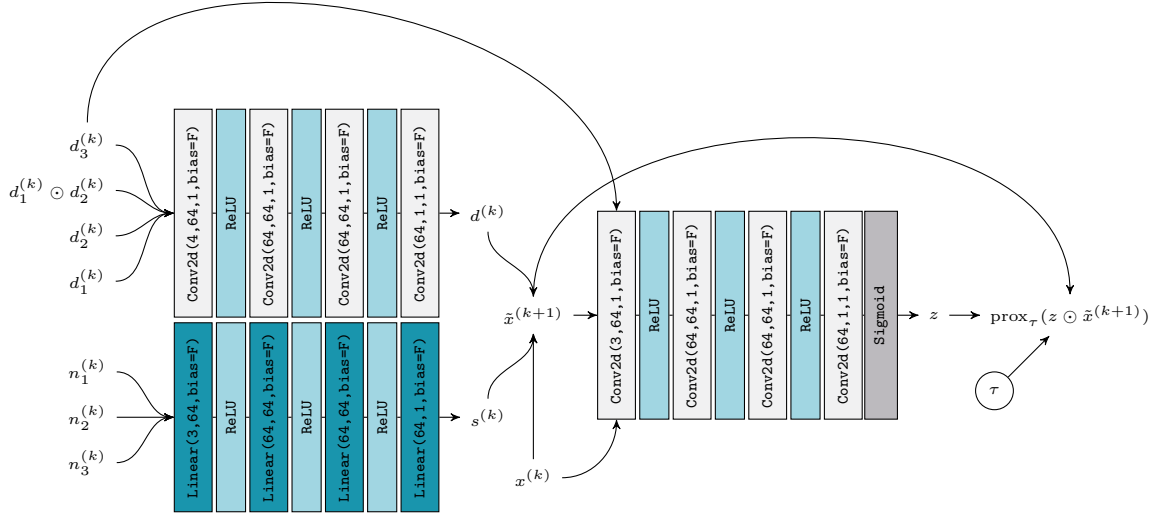


Figure 9: Algorithmic update for the Lasso-problem: Additionally to the directions $d_1^{(k)}$, $d_2^{(k)}$ and $d_1^{(k)} \odot d_2^{(k)}$, we add the subgradient $\partial\|x^{(k)}\|_1$ as input $d_3^{(k)}$. Hence, also the lower block, which computes the step-size $s^{(k)}$, gets an additional input $n_3^{(k)}$, namely the (transformed) norm of the subgradient. Then, $d^{(k)}$ and $s^{(k)}$ get used to compute a new point $\tilde{x}^{(k+1)} := x^{(k)} + s^{(k)} \cdot d^{(k)}$, which, together with the current point $x^{(k)}$ and the subgradient, gets fed into the “sparsity-block”, which computes a vector in $[0, 1]^d$ that is used to decrease the coordinates of $\tilde{x}^{(k+1)}$. Finally, we apply the proximal-map with learned parameter τ to get the next iterate.

can observe that the predicted PAC-bound is not perfectly tight, yet provides the guarantee to outperform NAG. Finally, the lower right plot shows that the probabilistically constraint optimization/sampling procedure did work correctly, as the point estimate is equal to one, that is, the algorithm did reach the sublevel set in 100% of the test cases.

6.3 Lasso-Problem

Here we consider the Lasso problem (Tibshirani, 1996), that is, a non-smooth problem of the form:

$$\min_{x \in \mathbb{R}^n} \frac{1}{2} \|Ax - b\|_2^2 + \lambda \|x\|_1 \quad A \in \mathbb{R}^{p \times n}, b \in \mathbb{R}^p,$$

with $p \leq n$. Thus, we are solving an overdetermined system of linear equations with an additional ℓ_1 -regularization term, which promotes sparsity in the solution (see Hastie et al., 2009). Hence, the optimization variable is given by $x \in \mathbb{R}^n$.

Construction of the Parameters. The same matrix $A \in \mathbb{R}^{p \times n}$ with dimensions $n = 350$ and $p = 70$ is used for all problem instances. Here, we sample each entry uniformly, that

is, $a_{i,j} \stackrel{iid}{\sim} \mathcal{U}_{[-10,10]}$, $i = 1, \dots, p$, $j = 1, \dots, n$. Thus, the parameters θ are given by the right-hand side and the regularization parameter, that is, $\theta = (b, \lambda) \in \mathbb{R}^{p+1} =: \Theta$. For this, the regularization parameter λ is also sampled uniformly, that is, $\lambda_i \stackrel{iid}{\sim} \mathcal{U}_{[\lambda_-, \lambda_+]}$, $i = 1, \dots, N$, with $\lambda_- = 0.01$ and $\lambda_+ = 0.5$, while the right-hand side is sampled from a multivariate normal distribution, that is, $b_i \stackrel{iid}{\sim} \mathcal{N}(\mu, \Sigma)$, $i = 1, \dots, N$, where we first construct μ and $\Sigma = C^T C$ by sampling each entry of μ and C uniformly at random in $[-5, 5]$.

Baseline. We use the FISTA algorithm (Beck and Teboulle, 2009) as baseline, which performs an extrapolation step followed by a proximal gradient step, that is, abbreviating $h(x) := \frac{1}{2} \|Ax - b\|^2$ and $g(x) := \lambda \|x\|_1$, the update is given by first computing $y^{(k)} = x^{(k)} + \beta^{(k)} (x^{(k)} - x^{(k-1)})$ followed by setting $x^{(k+1)} = \text{prox}_{\tau g}(y^{(k)} - \tau \nabla h(y^{(k)}))$. Here, the proximal mapping is defined as $\hat{x} = \text{prox}_{\tau g}(\bar{x}) = \arg \min_{x \in \mathbb{R}^n} \lambda \|x\|_1 + \frac{1}{2\tau} \|x - \bar{x}\|_2^2$, and can be computed efficiently in closed-form yielding the so-called *soft-thresholding operator*:

$$\hat{x}_i = \begin{cases} \bar{x}_i - \tau \lambda \frac{\bar{x}_i}{|\bar{x}_i|}, & |\bar{x}_i| > \tau \lambda; \\ 0, & \text{otherwise,} \end{cases}, \quad i = 1, \dots, n.$$

We choose $\tau = 1/L$, where L is the largest eigenvalue of $A^T A$, that is, the smoothness parameter of h , while $\beta^{(k)}$ is set to $\beta^{(k)} := (t_k - 1)/t_{k+1}$ with $t_{k+1} = (1 + \sqrt{1 + 4t_k^2})/2$.

Algorithm. The algorithmic update of the learned algorithm \mathcal{A} for this experiment is visualized in Figure 9 and consists of three blocks: First, an update-block, which combines the directions $d_1^{(k)}$, $d_2^{(k)}$, their interaction $d_1^{(k)} \odot d_2^{(k)}$, and the direction $d_3^{(k)}$ (of the L_1 -norm) into an update direction $d^{(k)}$. Second, a block to compute a step-size $s^{(k)}$ based on the norms of the gradient, subgradient and momentum term. And a final block, which is able to scale the single coordinates of a given point down, and, together with the proximal map, is used to promote sparsity in the solution. However, by appropriate scaling afterwards we ensure that the proximal map does not change the norm of the given point, to avoid that the network reduces the loss by just putting coordinates to zero.

Results. The results of this experiment are summarized in Figure 10. The upper left plot shows that the learned algorithm outperforms FISTA by several orders of magnitude, achieving a loss that is slightly smaller than the one of $x_{\text{std}}^{(5000)}$ after only 200 iterations, and one can observe that the learned algorithm can be used for more iterations than it was trained for. The upper right plot confirms that it is also way faster in terms of computation time. Yet, the increase in computation time for the dotted line indicates an increased difficulty to reach the level of accuracy, that is, \mathcal{A} might not reach arbitrary levels of accuracy. The lower left plot shows that the predicted PAC-bound is not perfectly tight here, yet guarantees that \mathcal{A} will outperform FISTA on average. And the lower right plot indicates that the probabilistically constrained optimization/sampling procedure did work as intended, since the algorithm did reach the sublevel set in all of the test cases.

6.4 Training Neural Networks

This experiment considers the problem of training a neural network on a regression problem, that is, \mathcal{A} is trained to predict the parameters $\beta \in \mathbb{R}^p$ of a neural network \mathbb{N}_β , which then is used to predict a function $g : \mathbb{R} \rightarrow \mathbb{R}$. Hence, the optimization variable is given by $\beta \in \mathbb{R}^p$.

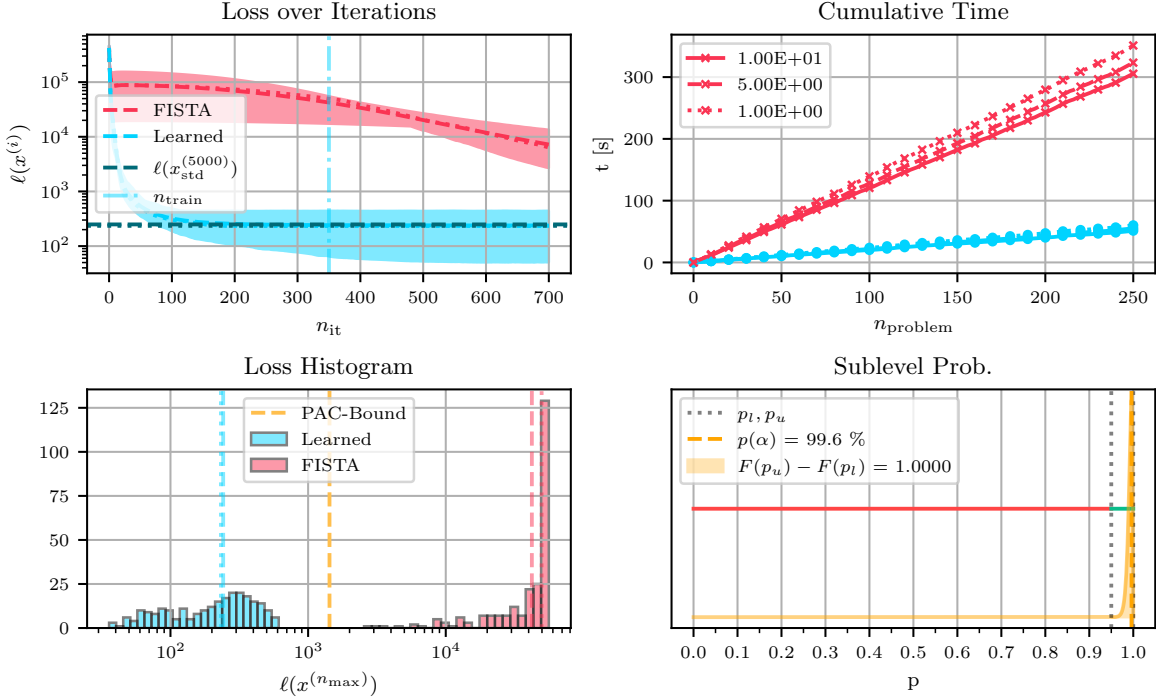


Figure 10: **Upper left:** Dashed lines represent the mean losses, while dotted lines represent the median losses over the iterations $x^{(i)}$, $i = 0, \dots, 700$. The shaded region represents the 10th to 90th percentile (80 % of the test data). The green horizontal lines represent the loss achieved by FISTA after 5000 iterations, which serves as approximation for the solution, while the vertical dash-dotted line indicates the number of iterations the algorithm was trained for. FISTA is shown in red and the learned algorithm in blue. **Upper right:** The different lines show the cumulative computation time the two algorithms need to solve all the test problems up to a certain accuracy (in function-values) measured by $\ell(x^{(i)}, \theta) - \ell(x_{std}^{(5000)}, \theta) < \varepsilon$. However, note that both algorithms are run for maximally $n_{max} = 5000$ iterations. **Lower left:** Loss histogram (after $n_{train} = 350$ iterations) with the predicted PAC-bound. **Lower right:** Estimated sublevel probability: The orange dashed line shows the point estimate, while the orange solid line shows the Beta-posterior arising from a non-informative prior. The gray dotted lines indicate the constraints p_l, p_u and show the feasible region as green line.

Construction of the Parameters. We assume that the neural network should learn a function $g : \mathbb{R} \rightarrow \mathbb{R}$ from noisy observations y_j , that is $y_j = g(x_j) + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, 1)$. For this, we construct polynomials g_i , $i = 1, \dots, N$, of degree $d = 5$ as follows: For every function g_i , we sample points $\{x_{i,j}\}_{j=1}^K$ (here: $K = 50$) uniformly in $[-2, 2]$, that is, $x_{i,j} \stackrel{iid}{\sim} U[-2, 2]$, $i = 1, \dots, N$, $j = 1, \dots, K$. Then, we sample the coefficients $(c_{i,0}, \dots, c_{i,5})$ of g_i uniformly in




Figure 11: Algorithmic update for training the neural network: Based on the two norms $n_1^{(k)}$ and $n_2^{(k)}$, the scalar product $s^{(k)} := \langle d_1^{(k)}, d_2^{(k)} \rangle$, and the current and previous (logarithmically scaled) loss $\ell^{(k)}$, $\ell^{(k-1)}$, we compute four scalars s_1, \dots, s_4 , which are used for weighting $d_1^{(k)}$, $d_2^{(k)}$, and their corresponding pre-conditioned versions $g \odot d_1^{(k)}$ and $m \odot d_2^{(k)}$. Together with the current and previous point, they get fed (as separate channels) into the first layer of a 1x1-convolutional block, which computes an update direction $d^{(k)}$. Then, we update $x^{(k+1)} := x^{(k)} + d^{(k)}$.

$[-5, 5]$, that is, $c_{i,l} \stackrel{iid}{\sim} \mathcal{U}[-5, 5]$, $i = 1, \dots, N$, $l = 0, \dots, 5$. Lastly, we get the values $y_{i,j}$ as:

$$y_{i,j} = g_i(x_{i,j}) + \varepsilon_{i,j} \quad \text{with} \quad \varepsilon_{i,j} \stackrel{iid}{\sim} \mathcal{N}(0, 1), \quad i = 1, \dots, N, \quad j = 1, \dots, 50.$$

For every function $g_i : \mathbb{R} \rightarrow \mathbb{R}$ the neural network is trained on the data set $\theta_i := \{X_i, Y_i\}$ with $X_i = (x_{i,1}, \dots, x_{i,K}) \in \mathbb{R}^K$ and $Y_i = (y_{i,1}, \dots, y_{i,K}) \in \mathbb{R}^K$. Hence, the data set will serve as the parameter θ of the loss function $\ell : \mathbb{R}^p \times \Theta \rightarrow \mathbb{R}_{\geq 0}$, such that the parameter space Θ can be identified as the space of these data sets, that is, $\Theta = \mathbb{R}^{K \times 2}$.

Loss Function. Since the mean square error is the standard choice for training models on regression tasks, the loss is given by $\ell(\beta, \theta_i) := c(\mathbb{N}(\beta, X_i), Y_i) := \frac{1}{K} \sum_{j=1}^K (\mathbb{N}_\beta(x_{i,j}) - y_{i,j})^2$.

Architecture. As model we use a fully-connected two layer neural network with ReLU-activation functions. To have more features in the input layer, the input x is transformed into the vector (x, x^2, \dots, x^5) . Hence, the parameters $\beta \in \mathbb{R}^p$ are given by the weights $A_1 \in \mathbb{R}^{50 \times 5}$, $A_2 \in \mathbb{R}^{1 \times 50}$ and biases $b_1 \in \mathbb{R}^{50}$, $b_2 \in \mathbb{R}$ of the two fully-connected layers. Therefore, the optimization space is of dimension $p = (5 \cdot 50) + (1 \cdot 50) + 50 + 1 = 351$.

Baseline. As baseline we use Adam (Kingma and Ba, 2015) (as it is implemented in PyTorch), which is a widely used optimization algorithms for training neural networks. For tuning, we perform a grid search over 100 step-size parameters, such that its performance is best for the given n iterations. Note that originally Adam was introduced for stochastic optimization, while we use it in the “full-batch setting” here, that is, without stochasticity.

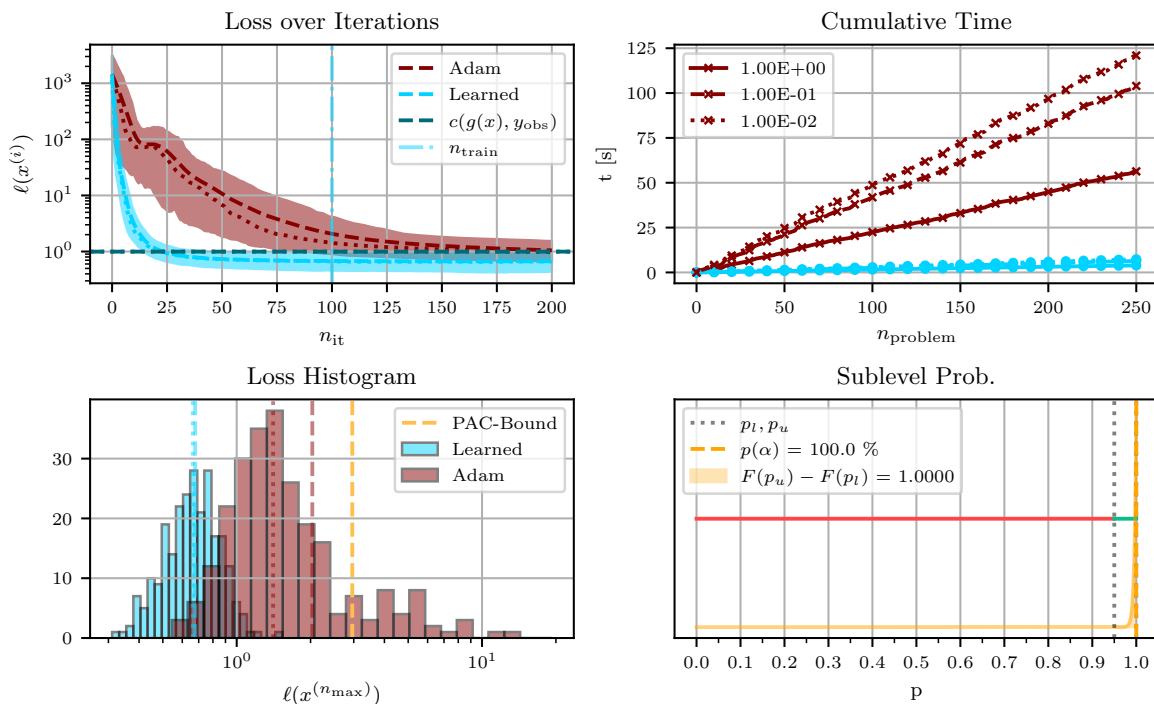


Figure 12: **Upper left:** Dashed lines represent the mean losses, while dotted lines represent the median losses over the iterations $\beta^{(i)}$, $i = 0, \dots, 200$. The shaded regions indicate the 10th to 90th percentile of the losses (80 % of the test data). The vertical, dashdotted line shows the number of iterations \mathcal{A} was trained for, and the horizontal, green, dashed line represents the average loss of the ground-truth function g (equal to one, as we added standard Gaussian noise). **Upper right:** The different lines show the cumulative computation time of the two algorithms to solve all the test problems up to a certain accuracy (in function values) measured by $\ell(\beta^{(i)}, \theta) - c(X_i, Y_i) < \varepsilon$. However, note that both algorithms are run for maximally $n_{\max} = 5000$ iterations. **Lower left:** Loss histogram (after n_{train} iterations) and PAC-bound. **Lower right:** Estimated sublevel probability: The orange dashed line shows the point estimate, while the orange solid line shows the Beta-posterior arising from a non-informative prior. Here, the gray dotted lines indicate the constraints p_l, p_u and show the feasible region as green line.

Algorithm. The algorithmic update in Figure 11 consists of two blocks: A weighting block, which computes four weights s_1, \dots, s_4 based on the norms $n_1^{(k)}, n_2^{(k)}$, the losses $\ell(x^{(k)}, \theta), \ell(x^{(k-1)}, \theta)$, and the scalar product $\langle d_1^{(k)}, d_2^{(k)} \rangle$. Each of these gets multiplied with $d_1^{(k)}, d_2^{(k)}$, or the “pre-conditioned” versions, which we compute by pointwise multiplication with the learned vectors g and m . Then, additionally to the $x^{(k)}$ and $x^{(k-1)}$, these weighted directions get fed into an update-block, which computes the final update direction $d^{(k)}$.

Results. The upper left plot of Figure 12 shows that the learned algorithm clearly outperforms Adam, reaching the ground-truth loss after about 25 iterations, while Adam is not able to reach it within 200 iterations. Further, while the algorithm was trained for 100 iterations, it can be applied for more. The upper right plot confirms that also in terms of computation time \mathcal{A} is way faster in training the neural network than Adam, and, as the learned algorithm goes below the ground-truth loss at around 25 iterations, all three levels of accuracy are reached similarly fast. The lower left plot shows that the predicted PAC-bound is not perfectly tight, yet yields a reasonable bound on the average performance, and guarantees to perform at least as good as Adam (on average). The lower right plot indicates that the probabilistically constrained optimization/sampling procedure did work correctly, as the algorithm did reach the sublevel set in all test cases.

7 Discussion and Limitations

The motivation for this paper was to use more structure in a given problem than is analytically tractable. For this, we considered a distribution over parametric loss functions and formulated the (ultimate) goal in (1), that is, to find a solution to each realization from this distribution. Under reasonable assumptions, this problem is too general to be solved. This led to the formulation of the performance of an algorithm in terms of its risk. However, since this is intractable, we derived PAC-Bayesian generalization bounds relating the risk to the empirically observable performance on a data set. This resulted in the formulation of a training objective, which relies heavily on the existence of a prior distribution satisfying our assumptions and yielding a good performance. As such a distribution is typically not known, we derived a procedure to construct it. This involved several key choices, including the loss-function, specific randomization steps, and, especially, the probabilistic constraints. Finally, we validated the resulting learning procedure on four practically relevant problems and showed that it yields a superior performance. While these experimental results are promising, we nevertheless see four main limitations of our work. First, the only guarantee that is provided by the PAC-Bayesian bound is an upper bound on the function value after a specified number of iterations. In particular, it does not guarantee that the function values, the iterates, or the gradient norm actually do converge. Second, our learning procedure is not *guaranteed* to work and still involves many design choices. Third, one still has to find a good architecture for each given problem, which can be time-consuming. Fourth, the presented algorithmic procedure is quite elaborate, and its implementation is tedious. Additionally, the training procedure (offline) has a high computational cost, which, at least in part, however, is due to the nature of learning-to-optimize. All these are promising directions of research that we leave to a future work.

Acknowledgments and Disclosure of Funding

M. Sucker and P. Ochs acknowledge funding by the German Research Foundation under Germany’s Excellence Strategy – EXC number 2064/1 – 390727645. Furthermore, J. Fadili and P. Ochs are supported by the ANR-DFG joint project TRINOM-DS under the numbers ANR-20-CE92-0037-01 and OC150/5-1.

Appendix A. Supplementary Definitions

Definition A.1 (Probability Kernel) Let $(X, \mathcal{X}), (Y, \mathcal{Y})$ be measurable spaces. A function $\mu : X \times \mathcal{Y} \rightarrow [0, \infty]$, $(x, A) \mapsto \mu(x, A)$ is called a kernel from X to Y , written as $\mu : X \rightarrow Y$, if for every set $A \in \mathcal{Y}$, the map $x \mapsto \mu(x, A)$ is measurable, and for every point $x \in X$, the map $A \mapsto \mu(x, A)$ is a measure. Furthermore, μ is called a probability kernel from X to Y , if $\mu(x, Y) = 1$ for every $x \in X$.

Definition A.2 (Polish space) A topological space is said to be Polish if it is separable and admits a complete metrization.

Remark A.3 Examples of Polish spaces are countable discrete spaces, \mathbb{R}^n with the standard topology, separable Banach spaces with the norm topology, or countable products of Polish spaces.

Definition A.4 (Exponential Family) Let Λ be a non-empty index set. A family of probability measures $(\mathbb{Q}_\lambda)_{\lambda \in \Lambda}$ on a measurable space X is called an exponential family (in η and T), if there is a dominating probability measure μ , that is, $(\mathbb{Q}_\lambda)_{\lambda \in \Lambda} \subset \mathcal{P}(\mu)$, functions $\eta : \Lambda \rightarrow \mathbb{R}^k$, $A : \Lambda \rightarrow (0, +\infty)$, and measurable functions $T : X \rightarrow \mathbb{R}^k$, $h : X \rightarrow (0, +\infty)$, such that for every $\lambda \in \Lambda$ we have $\mathbb{Q}_\lambda = hA(\lambda) \exp(\langle \eta(\lambda), T \rangle) \cdot \mu$, that is, $\mathbb{Q}_\lambda[B] = \int_B h(x)A(\lambda) \exp(\langle \eta(\lambda), T(x) \rangle) \mu(dx)$, $B \in \mathcal{B}(X)$.

Definition A.5 (Support of a Measure) Let S be a topological space, and let μ be a measure on S . The support of μ is defined as:

$$\text{supp}(\mu) := \{s \in S : \mu[B] > 0 \text{ for every neighborhood } B \text{ of } s\}.$$

Remark A.6 The support of a measure is closed, that is, Borel measurable, and satisfies $\mu(\text{supp}(\mu)^c) = 0$.

Appendix B. Supplementary Lemmas

Lemma B.1 Under Assumption 2.12, \mathbb{Q}_λ is a data-dependent distribution for every $\lambda \in \Lambda$.

Proof Denote the density of \mathbb{Q}_λ w.r.t. $\mathbb{P}_\mathfrak{F}$ by $f_\lambda(\alpha, d) := \frac{h(\alpha)}{c(\lambda, d)} \exp(\langle \eta(\lambda), T(\alpha, d) \rangle)$. The map $c(\lambda, \cdot) : \Theta^N \rightarrow [0, \infty]$ is $\mathcal{B}(\Theta^N)$ measurable, as T is measurable w.r.t. the product- σ -algebra and $\mathbb{P}_\mathfrak{F}$ is a finite measure (Kallenberg, 2021, Lemma 1.28, p.25). Hence, f_λ is measurable w.r.t. $\mathcal{B}(H) \otimes \mathcal{B}(\Theta^N)$, since $c(\lambda, d) \in (0, \infty)$. Thus, it holds that $\mathbb{Q}_\lambda = f_\lambda \cdot \mathbb{P}_\mathfrak{F}$ is a kernel from Θ^N to H (Kallenberg, 2021, Lem. 3.2, p.56). Finally, $\mathbb{Q}_\lambda : \Theta^N \rightarrow H$ is actually a probability kernel, since $c(\lambda, d)$ is the corresponding normalization constant. ■

Since the loss-function is neither assumed to be bounded nor to satisfy any self-bounding or bounded-difference property, the following result is needed. It states that non-negative random variables with finite second moment satisfy a one-sided sub-Gaussian inequality. It can be found as Exercise 2.9 on page 47 in Boucheron et al. (2013).

Lemma B.2 *Let \mathfrak{X} be a non-negative random variable with finite second moment. Then, for every $\lambda > 0$ it holds $\mathbb{E}[\exp(-\lambda(\mathfrak{X} - \mathbb{E}[\mathfrak{X}]))] \leq \exp\left(\frac{\lambda^2}{2}\mathbb{E}[\mathfrak{X}^2]\right)$.*

Lemma B.3 *The sublevel set L_g is measurable.*

Proof As g is assumed to be measurable, it suffices to show that the specific composition of ℓ and \mathcal{A} is measurable, that is, $\ell \circ \mathcal{A} : H \times \Theta \rightarrow [0, +\infty]$, $(\alpha, \theta) \mapsto \ell(\mathcal{A}(\alpha, \theta), \theta)$ is measurable w.r.t. $\mathcal{B}(H) \otimes \mathcal{B}(\Theta)$ and $\mathcal{B}([0, +\infty])$. Since $\ell \geq 0$ is measurable, there exists a sequence of simple³ functions ℓ_n with $\ell = \lim_{n \rightarrow \infty} \ell_n$. Thus, since limits of measurable functions are measurable, it suffices to consider the case of a simple function $\ell : \mathbb{R}^n \times \Theta \rightarrow \mathbb{R}$. Then, however, since addition and multiplication are measurable, it suffices to consider characteristic functions of the form $\mathbb{1}_A$ for a measurable set $A \in \mathcal{B}(\mathbb{R}^n) \otimes \mathcal{B}(\Theta)$. Since the product- σ -algebra is generated by cylinder sets of the form $B \times D$ with $B \in \mathcal{B}(\mathbb{R}^n)$, $D \in \mathcal{B}(\Theta)$, it actually suffices to consider the case $\ell = \mathbb{1}_{B \times D}$, that is, $(\ell \circ \mathcal{A})(\alpha, \theta) = \mathbb{1}_{B \times D}(\mathcal{A}(\alpha, \theta), \theta) = \mathbb{1}_B(\mathcal{A}(\alpha, \theta))\mathbb{1}_D(\theta)$. The second term is obviously measurable, and the first term is measurable as a (standard) composition of two measurable functions. Therefore, the conclusion follows. \blacksquare

Appendix C. Proof of Lemma 3.1

Proof Take any $\lambda \in \Lambda$ and $d \in \Theta^N$. First, let $\mathbb{Q} \in \mathcal{P}(\mathbb{P}_{\mathfrak{F}})$ be arbitrary. By the Radon-Nikodym theorem, there exists a measurable function $f \geq 0$, s.t. $\mathbb{Q} = f \cdot \mathbb{P}_{\mathfrak{F}}$. Since the convention $0 \cdot \infty = 0$ applies throughout measure theory, one has:

$$D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}_{\mathfrak{F}}) = \mathbb{Q}[\log(f)] = \mathbb{P}_{\mathfrak{F}}[f \log(f)] = \mathbb{P}_{\mathfrak{F}}[\mathbb{1}_{\{f>0\}} f \log(\mathbb{1}_{\{f>0\}} f)] .$$

Hence, w.l.o.g. we can assume that $f > 0$ $\mathbb{P}_{\mathfrak{F}}$ -a.s. Then, by Jensen's inequality, one gets:

$$\begin{aligned} & \mathbb{Q}[\langle \eta(\lambda), T(\cdot, d) \rangle + \log(h)] - \mathbb{Q}[\log(f)] = \mathbb{Q} \left[\log \left(\frac{h}{f} \exp(\langle \eta(\lambda), T(\cdot, d) \rangle) \right) \right] \\ & \leq \log \left(\mathbb{Q} \left[\frac{h}{f} \exp(\langle \eta(\lambda), T(\cdot, d) \rangle) \right] \right) = \log \left((f \cdot \mathbb{P}_{\mathfrak{F}}) \left[\frac{h}{f} \exp(\langle \eta(\lambda), T(\cdot, d) \rangle) \right] \right) \\ & = \log(\mathbb{P}_{\mathfrak{F}}[h \exp(\langle \eta(\lambda), T(\cdot, d) \rangle)]) = \kappa(\lambda, d) . \end{aligned}$$

Hence, we have $\kappa(\lambda, d) \geq \mathbb{Q}[\langle \eta(\lambda), T(\cdot, d) \rangle + \log(h)] - D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}_{\mathfrak{F}})$ for any probability measure $\mathbb{Q} \ll \mathbb{P}_{\mathfrak{F}}$. Now consider the member of the exponential family:

$$\begin{aligned} D_{\text{KL}}(\mathbb{Q}_\lambda(d) \parallel \mathbb{P}_{\mathfrak{F}}) &= \int_H \log(h(\alpha)A(\lambda, d) \exp(\langle \eta(\lambda), T(\alpha, d) \rangle)) \mathbb{Q}_\lambda(d, d\alpha) \\ &= \int_H \log(h(\alpha)) + \langle \eta(\lambda), T(\alpha, d) \rangle \mathbb{Q}_\lambda(d, d\alpha) + \log(A(\lambda, d)) \\ &= \int_H \log(h(\alpha)) + \langle \eta(\lambda), T(\alpha, d) \rangle \mathbb{Q}_\lambda(d, d\alpha) - \log(c(\lambda, d)) \\ &= \int_H \log(h(\alpha)) + \langle \eta(\lambda), T(\alpha, d) \rangle \mathbb{Q}_\lambda(d, d\alpha) - \kappa(\lambda, d) . \end{aligned}$$

3. A function is called simple, if it is of the form $\ell_n = \sum_{i=1}^K a_n^i \mathbb{1}_{A_n^i}$ with disjoint sets A_n^i .

Rearranging yields $\kappa(\lambda, d) = \int_H \log(h(\alpha)) + \langle \eta(\lambda), T(\alpha, d) \rangle \mathbb{Q}_\lambda(d, d\alpha) - D_{\text{KL}}(\mathbb{Q}_{\lambda, d} \parallel \mathbb{P}_\mathfrak{H})$. ■

Appendix D. Proof of Lemma 3.2

Proof W.l.o.g. assume that $\mathcal{O}_i \neq \emptyset$ and choose $\lambda_i \in \mathcal{O}_i$, $i = 1, \dots, K$. Then, for every $d \in \Theta^N$, it holds that:

$$\begin{aligned} \sup_{\lambda \in \Lambda} \kappa(\lambda, d) &\leq \max_{i=1, \dots, K} \sup_{\lambda \in \mathcal{O}_i} \kappa(\lambda, d) = \max_{i=1, \dots, K} \left\{ \kappa(\lambda_i, d) + \sup_{\lambda \in \mathcal{O}_i} (\kappa(\lambda, d) - \kappa(\lambda_i, d)) \right\} \\ &\leq \max_{i=1, \dots, K} \kappa(\lambda_i, d) + C_{\mathcal{O}}. \end{aligned}$$

Thus, in total one gets for $s \in \mathbb{R}$:

$$\begin{aligned} \mathbb{P} \left\{ \sup_{\lambda \in \Lambda} \kappa(\lambda, \mathfrak{D}) > s \right\} &\leq \mathbb{P} \left\{ \max_{i=1, \dots, K} \kappa(\lambda_i, \mathfrak{D}) + C_{\mathcal{O}} > s \right\} \leq \sum_{i=1}^K \mathbb{P} \{ \kappa(\lambda_i, \mathfrak{D}) + C_{\mathcal{O}} > s \} \\ &\leq \sum_{i=1}^K \exp(C_{\mathcal{O}} - s) = K \exp(C_{\mathcal{O}} - s). \end{aligned}$$

Taking $s = \log\left(\frac{K}{\epsilon}\right) + C_{\mathcal{O}}$ gives $\mathbb{P} \left\{ \sup_{\lambda \in \Lambda} \kappa(\lambda, \mathfrak{D}) > \log\left(\frac{K}{\epsilon}\right) + C_{\mathcal{O}} \right\} \leq \epsilon$. ■

Appendix E. Proof of Corollary 3.5

Proof The two formulas are simply rewritings of each other: By assumption, bilinearity and definition of the euclidean scalar product, and linearity of the integral, the term $\mathbb{Q}[\langle \eta(\lambda), T(\cdot, d) \rangle]$ can be split up as:

$$\begin{aligned} \mathbb{Q}[\langle \eta(\lambda), T(\cdot, d) \rangle] &= \mathbb{Q}[\eta^{(1)}(\lambda)(\mathcal{R} - \hat{\mathcal{R}}(\cdot, d))] + \mathbb{Q}[\langle \eta^{(r)}(\lambda), T^{(r)}(\cdot, d) \rangle] \\ &= \eta^{(1)}(\lambda) \mathbb{Q}[\mathcal{R}] - \eta^{(1)}(\lambda) \mathbb{Q}[\hat{\mathcal{R}}(\cdot, d)] + \mathbb{Q}[\langle \eta^{(r)}(\lambda), T^{(r)}(\cdot, d) \rangle]. \end{aligned}$$

Simply rearranging the terms then yields the result, as $\eta^{(1)} > 0$. ■

References

- Pierre Alquier. PAC-Bayesian bounds for randomized empirical risk minimizers. *Mathematical Methods of Statistics*, 17(4):279–304, 2008.
- Pierre Alquier. User-friendly introduction to PAC-Bayes bounds. *arXiv preprint arXiv:2110.11216*, 2021.
- Pierre Alquier and Benjamin Guedj. Simpler PAC-Bayesian bounds for hostile data. *Machine Learning*, 107(5):887–902, 2018.
- Pierre Alquier, James Ridgway, and Nicolas Chopin. On the properties of variational approximations of Gibbs posteriors. *Journal of Machine Learning Research*, 17(1):8374–8414, 2016.
- Ron Amit, Baruch Epstein, Shay Moran, and Ron Meir. Integral probability metrics PAC-Bayes bounds. *Advances in Neural Information Processing Systems*, 35:3123–3136, 2022.
- Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. Learning to learn by gradient descent by gradient descent. *Advances in neural information processing systems*, 29, 2016.
- Rémi Bardenet, Arnaud Doucet, and Chris Holmes. Towards scaling up markov chain monte carlo: an adaptive subsampling approach. In *International conference on machine learning*, pages 405–413. PMLR, 2014.
- Rémi Bardenet, Arnaud Doucet, and Chris Holmes. On Markov chain Monte Carlo methods for tall data. *Journal of Machine Learning Research*, 18(47), 2017.
- Ole Barndorff-Nielsen. *Information and exponential families: in statistical theory*. John Wiley & Sons, 2014.
- Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- Luc Bégin, Pascal Germain, François Laviolette, and Jean-François Roy. PAC-Bayesian bounds based on the Rényi divergence. In *Artificial Intelligence and Statistics*, pages 435–444. PMLR, 2016.
- James O. Berger. *Statistical decision theory and Bayesian analysis*. Springer-Verlag, 1985.
- Pascal Bianchi, Walid Hachem, and Sholom Schechtman. Convergence of constant step stochastic gradient descent for non-smooth non-convex functions. *Set-Valued and Variational Analysis*, pages 1–31, 2022.
- Jérôme Bolte and Edouard Pauwels. Conservative set valued fields, automatic differentiation, stochastic gradient methods and deep learning. *Mathematical Programming*, 188(1):19–51, 2021.
- Leon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.

- Stephane Boucheron, Gabor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- Olivier Bousquet and André Elisseeff. Algorithmic stability and generalization performance. *Advances in Neural Information Processing Systems*, 13, 2000.
- Olivier Bousquet and André Elisseeff. Stability and generalization. *Journal of Machine Learning Research*, 2:499–526, 2002.
- Gregery T Buzzard, Stanley H Chan, Suhas Sreehari, and Charles A Bouman. Plug-and-play unplugged: Optimization-free reconstruction using consensus equilibrium. *SIAM Journal on Imaging Sciences*, 11(3):2001–2020, 2018.
- Olivier Catoni. *Statistical learning theory and stochastic optimization: Ecole d’Eté de Probabilités de Saint-Flour, XXXI-2001*, volume 1851. Springer Science & Business Media, 2004.
- Olivier Catoni. PAC-Bayesian supervised classification: The thermodynamics of statistical learning. *Lecture Notes-Monograph Series*, 56:i–163, 2007.
- Stanley H Chan, Xiran Wang, and Omar A Elgandy. Plug-and-play ADMM for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging*, 3(1):84–98, 2016.
- Tianlong Chen, Weiyi Zhang, Zhou Jingyang, Shiyu Chang, Sijia Liu, Lisa Amini, and Zhangyang Wang. Training stronger baselines for learning to optimize. *Advances in Neural Information Processing Systems*, 33:7332–7343, 2020a.
- Tianlong Chen, Xiaohan Chen, Wuyang Chen, Howard Heaton, Jialin Liu, Zhangyang Wang, and Wotao Yin. Learning to optimize: A primer and a benchmark. *arXiv preprint arXiv:2103.12828*, 2021.
- Xiaohan Chen, Jialin Liu, Zhangyang Wang, and Wotao Yin. Theoretical linear convergence of unfolded ISTA and its practical weights and thresholds. *Advances in Neural Information Processing Systems*, 31, 2018.
- Xinshi Chen, Yufei Zhang, Christoph Reisinger, and Le Song. Understanding deep architecture with reasoning layer. *Advances in Neural Information Processing Systems*, 33:1240–1252, 2020b.
- Yutian Chen, Matthew W Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Timothy P Lillicrap, Matt Botvinick, and Nando Freitas. Learning to learn without gradient descent by gradient descent. In *International Conference on Machine Learning*, pages 748–756. PMLR, 2017.
- Regev Cohen, Michael Elad, and Peyman Milanfar. Regularization by denoising via fixed-point projection. *SIAM Journal on Imaging Sciences*, 14(3):1374–1406, 2021.
- Damek Davis and Dmitriy Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization*, 29(1):207–239, 2019.

- Damek Davis and Dmitriy Drusvyatskiy. Graphical convergence of subgradients in non-convex optimization and learning. *Mathematics of Operations Research*, 47(1):209–231, 2022.
- Alexandre Défossez, Leon Bottou, Francis Bach, and Nicolas Usunier. A simple convergence proof of adam and adagrad. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856.
- Monroe D Donsker and SR Srinivasa Varadhan. Asymptotic evaluation of certain Markov process expectations for large time, i. *Communications on Pure and Applied Mathematics*, 28(1):1–47, 1975.
- Gintare Karolina Dziugaite and Daniel M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence, UAI 2017, Sydney, Australia, August 11-15, 2017*. AUAI Press, 2017.
- Gintare Karolina Dziugaite and Daniel M Roy. Data-dependent PAC-Bayes priors via differential privacy. *Advances in neural information processing systems*, 31, 2018.
- Gintare Karolina Dziugaite, Kyle Hsu, Waseem Gharbieh, Gabriel Arpino, and Daniel Roy. On the role of data in PAC-Bayes bounds. In *International Conference on Artificial Intelligence and Statistics*, pages 604–612. PMLR, 2021.
- Bradley Efron. Defining the curvature of a statistical problem (with applications to second order efficiency). *The Annals of Statistics*, pages 1189–1242, 1975.
- Pascal Germain, Alexandre Lacasse, François Laviolette, and Mario Marchand. PAC-Bayesian learning of linear classifiers. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 353–360, 2009.
- Karol Gregor and Yann LeCun. Learning fast approximations of sparse coding. In *Proceedings of the 27th international conference on international conference on machine learning*, pages 399–406, 2010.
- Benjamin Guedj. A primer on PAC-Bayesian learning. In *Proceedings of the second congress of the French Mathematical Society*, volume 33, 2019.
- Maxime Haddouche and Benjamin Guedj. Wasserstein PAC-Bayes learning: Exploiting optimisation guarantees to explain generalisation, 2023.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning: data mining, inference and prediction*. Springer, 2 edition, 2009.
- Xin He, Kaiyong Zhao, and Xiaowen Chu. Automl: A survey of the state-of-the-art. *Knowledge-Based Systems*, 212:106622, 2021.
- Jean Honorio and Tommi Jaakkola. Tight bounds for the expected risk of linear classifiers and PAC-Bayes finite-sample guarantees. In *Artificial Intelligence and Statistics*, pages 384–392. PMLR, 2014.

- Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5149–5169, 2021.
- Frank Hutter, Lars Kotthoff, and Joaquin Vanschoren. *Automated machine learning: methods, systems, challenges*. Springer Nature, 2019.
- O. Kallenberg. *Foundations of Modern Probability*. Probability theory and stochastic modelling. Springer, 2021. ISBN 9783030618728.
- Ali Kavis, Kfir Yehuda Levy, and Volkan Cevher. High probability bounds for a class of nonconvex algorithms with adagrad stepsize. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=dSw0QtRMJk0>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015*, 2015.
- Achim Klenke. *Wahrscheinlichkeitstheorie*. Springer Spektrum, 2013.
- Erich Kobler, Alexander Effland, Karl Kunisch, and Thomas Pock. Total deep variation: A stable regularizer for inverse problems. *arXiv preprint arXiv:2006.08789*, 2020.
- Anoop Korattikara, Yutian Chen, and Max Welling. Austerity in MCMC land: Cutting the Metropolis-Hastings budget. In *International conference on machine learning*, pages 181–189. PMLR, 2014.
- John Langford and Rich Caruana. (Not) bounding the true error. In T. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001.
- John Langford and John Shawe-Taylor. PAC-Bayes and margins. *Advances in neural information processing systems*, 15, 2002.
- Guy Lever, François Laviolette, and John Shawe-Taylor. Tighter PAC-Bayes bounds through distribution-dependent priors. *Theoretical Computer Science*, 473:4–28, 2013.
- Ben London. A PAC-Bayesian analysis of randomized learning with application to stochastic gradient descent. *Advances in Neural Information Processing Systems*, 30, 2017.
- Kaifeng Lv, Shunhua Jiang, and Jian Li. Learning gradient descent: Better generalization and longer horizons. In *International Conference on Machine Learning*, pages 2247–2255. PMLR, 2017.
- Dougal Maclaurin and Ryan P Adams. Firefly Monte Carlo: Exact MCMC with subsets of data. In *30th Conference on Uncertainty in Artificial Intelligence, UAI 2014*, pages 543–552. AUAI Press, 2014.
- David McAllester. Simplified PAC-Bayesian margin bounds. In *Learning theory and Kernel machines*, pages 203–215. Springer, 2003a.

- David McAllester. PAC-Bayesian stochastic model selection. *Machine Learning*, 51(1):5–21, 2003b.
- Luke Metz, Niru Maheswaranathan, Jeremy Nixon, Daniel Freeman, and Jascha Sohl-Dickstein. Understanding and correcting pathologies in the training of learned optimizers. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 4556–4565. PMLR, 2019.
- Michael Moeller, Thomas Mollenhoff, and Daniel Cremers. Controlling neural networks via energy dissipation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3256–3265, 2019.
- Yurii Nesterov. A method for solving the convex programming problem with convergence rate $O(1/k^2)$. *Proceedings of the USSR Academy of Sciences*, 269:543–547, 1983.
- Yurii Nesterov. *Lectures on convex optimization*, volume 137. Springer, 2018.
- Yuki Ohnishi and Jean Honorio. Novel change of measure inequalities with applications to PAC-Bayesian bounds and Monte Carlo estimation. In *International Conference on Artificial Intelligence and Statistics*, pages 1711–1719. PMLR, 2021.
- Emilio Parrado-Hernández, Amiran Ambroladze, John Shawe-Taylor, and Shiliang Sun. PAC-Bayes bounds with data dependent priors. *Journal of Machine Learning Research*, 13(1):3507–3531, 2012.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- María Pérez-Ortiz, Omar Rivasplata, John Shawe-Taylor, and Csaba Szepesvári. Tighter risk certificates for neural networks. *Journal of Machine Learning Research*, 22(227):1–40, 2021.
- Boris T Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- Matias Quiroz, Robert Kohn, Mattias Villani, and Minh-Ngoc Tran. Speeding up MCMC by efficient data subsampling. *Journal of the American Statistical Association*, 2018.
- Omar Rivasplata, Ilja Kuzborskij, Csaba Szepesvári, and John Shawe-Taylor. PAC-Bayes analysis beyond the usual bounds. *Advances in Neural Information Processing Systems*, 33:16833–16845, 2020.
- C.P. Robert and G. Casella. *Monte Carlo statistical methods*. Springer New York, 2004.
- Ernest Ryu, Jialin Liu, Sicheng Wang, Xiaohan Chen, Zhangyang Wang, and Wotao Yin. Plug-and-play methods provably converge with properly trained denoisers. In *International Conference on Machine Learning*, pages 5546–5557. PMLR, 2019.

- Matthias Seeger. PAC-Bayesian generalisation error bounds for Gaussian process classification. *Journal of Machine Learning Research*, 3:233–269, 2002.
- Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Stochastic convex optimization. In *COLT*, volume 2, page 5, 2009.
- Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Learnability, stability and uniform convergence. *Journal of Machine Learning Research*, 11:2635–2670, 2010.
- Suhas Sreehari, S Venkat Venkatakrishnan, Brendt Wohlberg, Gregory T Buzzard, Lawrence F Drummy, Jeffrey P Simmons, and Charles A Bouman. Plug-and-play priors for bright field electron tomography and sparse interpolation. *IEEE Transactions on Computational Imaging*, 2(4):408–423, 2016.
- Michael Sucker and Peter Ochs. PAC-Bayesian learning of optimization algorithms. In *International Conference on Artificial Intelligence and Statistics*, pages 8145–8164. PMLR, 2023.
- Yu Sun, Brendt Wohlberg, and Ulugbek S Kamilov. An online plug-and-play algorithm for regularized image reconstruction. *IEEE Transactions on Computational Imaging*, 5(3):395–408, 2019.
- Afonso M Teodoro, José M Bioucas-Dias, and Mário AT Figueiredo. Scene-adapted plug-and-play algorithm with convergence guarantees. In *2017 IEEE 27th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6. IEEE, 2017.
- Matthieu Terris, Audrey Repetti, Jean-Christophe Pesquet, and Yves Wiaux. Enhanced convergent PNP algorithms for image restoration. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 1684–1688. IEEE, 2021.
- Niklas Thiemann, Christian Igel, Olivier Wintenberger, and Yevgeny Seldin. A strongly quasiconvex PAC-Bayesian bound. In *International Conference on Algorithmic Learning Theory*, pages 466–492. PMLR, 2017.
- Robert Tibshirani. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288, 1996.
- Tom Tirer and Raja Giryes. Image restoration by iterative denoising and backward projections. *IEEE Transactions on Image Processing*, 28(3):1220–1234, 2018.
- Ricardo Vilalta and Youssef Drissi. A perspective view and survey of meta-learning. *Artificial intelligence review*, 18:77–95, 2002.
- Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Max Welling and Yee W Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML)*, pages 681–688, 2011.

Olga Wichrowska, Niru Maheswaranathan, Matthew W Hoffman, Sergio Gomez Colmenarejo, Misha Denil, Nando Freitas, and Jascha Sohl-Dickstein. Learned optimizers that scale and generalize. In *International conference on machine learning*, pages 3751–3760. PMLR, 2017.

Hermann Witting. *Mathematische Statistik I: Parametrische Verfahren bei festem Stichprobenumfang*. Springer-Verlag, 2013.

Bo Xin, Yizhou Wang, Wen Gao, David Wipf, and Baoyuan Wang. Maximal sparsity with deep networks? *Advances in Neural Information Processing Systems*, 29, 2016.

Quanming Yao, Mengshuo Wang, Yuqiang Chen, Wenyuan Dai, Yu-Feng Li, Wei-Wei Tu, Qiang Yang, and Yang Yu. Taking human out of learning applications: A survey on automated machine learning. *arXiv preprint arXiv:1810.13306*, 2018.