
On the interplay between noise and curvature and its effect on optimization and generalization

Valentin Thomas
Mila, Université de Montréal

Fabian Pedregosa
Google Research, Brain Team

Bart van Merriënboer
Google Research, Brain Team

Pierre-Antoine Mangazol
Google Research, Brain Team

Yoshua Bengio
Mila, Université de Montréal
CIFAR Senior Fellow

Nicolas Le Roux
Google Research, Brain Team
Mila, McGill University

Abstract

The speed at which one can minimize an expected loss using stochastic methods depends on two properties: the curvature of the loss and the variance of the gradients. While most previous works focus on one or the other of these properties, we explore how their interaction affects optimization speed. Further, as the ultimate goal is good generalization performance, we clarify how both curvature and noise are relevant to properly estimate the generalization gap. Realizing that the limitations of some existing works stems from a confusion between these matrices, we also clarify the distinction between the Fisher matrix, the Hessian, and the covariance matrix of the gradients.

1 Introduction

Training a machine learning model is often cast as the minimization of a smooth function f over parameters θ in \mathbb{R}^d . More precisely, we aim at finding a minimum of an expected loss, i.e.

$$\theta^* \in \arg \min_{\theta} \mathbb{E}_p[\ell(\theta, x)], \quad (1)$$

where the expectation is under the data distribution $x \sim p$. In practice, we only have access to an empirical

distribution \hat{p} over x and minimize the training loss

$$\hat{\theta}^* \in \arg \min_{\theta} \mathbb{E}_{\hat{p}}[\ell(\theta, x)] \quad (2)$$

$$= \arg \min_{\theta} f(\theta). \quad (3)$$

To minimize this function, we assume access to an oracle which, for every value of θ and x , returns both $\ell(\theta, x)$ and its derivative with respect to θ , i.e., $\nabla_{\theta} \ell(\theta, x)$. Given this oracle, stochastic gradient iteratively performs the following update: $\theta_{t+1} = \theta_t - \alpha_t \nabla \ell(\theta_t, x)$ ¹ where $\{\alpha_t\}_{t \geq 0}$ is a sequence of stepsizes.

Two questions arise: First, how quickly do we converge to $\hat{\theta}^*$ and how is this speed affected by properties of ℓ and \hat{p} ? Second, what is $\mathbb{E}_{\hat{p}}[\ell(\hat{\theta}^*, x)]$?

It is known that the former is influenced by two quantities: the curvature of the function, either measured through its smoothness constant or its condition number, and the noise on the gradients, usually measured through a bound on $\mathbb{E}_{\hat{p}}[\|\nabla \ell(\theta, x)\|^2]$. For instance, when f is μ -strongly convex, L -smooth, and the noise is bounded, i.e. $\mathbb{E}_{\hat{p}}[\|\nabla \ell(\theta, x)\|^2] \leq c$, then stochastic gradient with a constant stepsize α will converge linearly to a ball (Schmidt, 2014). Calling Δ the suboptimality, i.e. $\Delta_k = f(\theta_k) - f(\hat{\theta}^*)$, we have

$$\mathbb{E}[\Delta_k] \leq (1 - 2\alpha\mu)^k \Delta_0 + \frac{L\alpha c}{4\mu}. \quad (4)$$

This implies that as $k \rightarrow \infty$, the expected suboptimality depends on both the curvature through μ and L , and on the noise through c . However, bounding curvature and noise using constants rather than full matrices hides the dependencies between these two quantities. We also observed that, because existing works replace the full noise matrix with a constant when deriving convergence rates, that matrix is poorly understood

Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108. Copyright 2020 by the author(s).

¹We omit the subscripts when clear from context.

and is often confused with a curvature matrix. This confusion remains when discussing the generalization properties of a model. Indeed, the generalization gap stems from a discrepancy between the empirical and the true data distribution. An estimator of this gap must thus include an estimate of that discrepancy in addition to an estimate of the impact of an infinitesimal discrepancy on the loss. The former can be characterized as noise and the latter as curvature. Hence, attempts at estimating the generalization gap using only the curvature (Keskar et al., 2017; Novak et al., 2018) are bound to fail as do not characterize the size or geometry of the discrepancy.

In this work, we make the following contributions:

- We provide theoretical and empirical evidence of the similarities and differences surrounding the curvatures matrices; the Fisher \mathbf{F} and the Hessian \mathbf{H} , and the noise matrix, \mathbf{C} ;
- We briefly expand the convergence results of Schmidt (2014), theoretically and empirically highlighting the importance of the relationship between noise and curvature for strongly convex functions and quadratics;
- We make the connection with an old estimator of the generalization gap, the Takeuchi Information Criterion, and show how its use of both curvature and noise yields a superior estimator to other commonly used ones, such as flatness or sensitivity for neural networks.

2 Information matrices: definitions, similarities, and differences

Before delving into the impact of the information matrices for optimization and generalization, we start by recalling their definitions. We shall see that, despite having similar formulations, they encode different information. We then provide insights on their similarities and differences.

We discuss here two information matrices associated with curvature, the Fisher matrix \mathbf{F} and the Hessian \mathbf{H} , and one associated with noise, the gradients’ uncentered covariance \mathbf{C} . In particular, while \mathbf{F} and \mathbf{H} are well understood, \mathbf{C} is often misinterpreted. For instance, it is often called “empirical Fisher” (Martens, 2014) despite bearing no relationship to \mathbf{F} , the true Fisher. This confusion can have dire consequences and optimizers using \mathbf{C} as approximation to \mathbf{F} can have arbitrarily poor performance (Kunstner et al., 2019).

To present these matrices, we consider the case of maximum likelihood estimation (MLE). We have access to a set of samples $(x, y) \in \mathcal{X} \times \mathcal{Y}$ where x is the input and y the target. We define $p : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$ as the

data distribution and $q_\theta : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$ such that $q_\theta(x, y) = p(x)q_\theta(y|x)$ as the **model distribution**². For each sample $(x, y) \sim p$, our loss is the negative log-likelihood $\ell(\theta, y, x) = -\log q_\theta(y|x)$. Note that all the definitions and results in this section are valid whether we use the true data distribution p or the empirical \hat{p} .

Matrices \mathbf{H} , \mathbf{F} and \mathbf{C} are then defined as:

$$\mathbf{H}(\theta) = \mathbb{E}_p \left[\frac{\partial^2}{\partial \theta \partial \theta^\top} \ell(\theta, y, x) \right] \quad (5)$$

$$\mathbf{C}(\theta) = \mathbb{E}_p \left[\frac{\partial}{\partial \theta} \ell(\theta, y, x) \frac{\partial}{\partial \theta} \ell(\theta, y, x)^\top \right] \quad (6)$$

$$\mathbf{F}(\theta) = \mathbb{E}_{q_\theta} \left[\frac{\partial}{\partial \theta} \ell(\theta, y, x) \frac{\partial}{\partial \theta} \ell(\theta, y, x)^\top \right] \quad (7)$$

$$= \mathbb{E}_{q_\theta} \left[\frac{\partial^2}{\partial \theta \partial \theta^\top} \ell(\theta, y, x) \right]. \quad (8)$$

We observe the following: a) The definition of \mathbf{H} and \mathbf{C} involves the data distribution, in contrast with the definition of \mathbf{F} , which involves the model distribution; b) If $q_\theta = p$, all matrices are equal. Furthermore, as noted by Martens (2014), $\mathbf{H} = \mathbf{F}$ whenever the matrix of second derivatives does not depend on y , a property shared in particular by all generalized linear models.

As said above, \mathbf{H} , \mathbf{F} , and \mathbf{C} characterize different properties of the optimization problem. \mathbf{H} and \mathbf{F} are curvature matrices and describe the geometry of the space around the current point. \mathbf{C} , on the other hand, is a “noise matrix” and represents the sensitivity of the gradient to the particular sample.³

We now explore in more details their similarities and differences.

2.1 Bounds between \mathbf{H} , \mathbf{F} and \mathbf{C}

The following proposition bounds the distance between the information matrices:

Proposition 1 (Distance between \mathbf{H} , \mathbf{F} and \mathbf{C})

Assuming the second moments of the Fisher are bounded above, i.e. $\mathbb{E}_{q_\theta} [|\nabla_\theta^2 \ell(\theta, x, y)|^2] \leq \beta_1$ and $\mathbb{E}_{q_\theta} [|\nabla_\theta \ell(\theta, x, y) \nabla_\theta \ell(\theta, x, y)^\top|^2] \leq \beta_2$, we have

$$\begin{aligned} \|\mathbf{F} - \mathbf{H}\|^2 &\leq \beta_1 \mathcal{D}_{\chi^2}(p||q_\theta), \\ \|\mathbf{F} - \mathbf{C}\|^2 &\leq \beta_2 \mathcal{D}_{\chi^2}(p||q_\theta), \\ \|\mathbf{C} - \mathbf{H}\|^2 &\leq (\beta_1 + \beta_2) \mathcal{D}_{\chi^2}(p||q_\theta). \end{aligned}$$

² $q_\theta(y|x)$ are the softmax activations of a neural network in the classification setting.

³Technically, it is \mathbf{S} , the centered covariance matrix, rather than \mathbf{C} which plays that role but the two are similar close to a stationary point.

where $\mathcal{D}_{\chi^2}(p||q_\theta) = \iint \frac{(p(x,y)-q_\theta(x,y))^2}{q_\theta(x,y)} dydx$ is the χ^2 divergence and $\|\cdot\|$ is the Frobenius norm.

All the proofs are in the appendix.

In particular, when $p = q_\theta$ we recover that $\mathbf{F} = \mathbf{H} = \mathbf{C}$. At first glance, one could assume that, as the model is trained and q_θ approaches p , these matrices become more similar. The χ^2 divergence is however poorly correlated with the loss of the model. One sample where the prediction of the model is much smaller than the true distribution can dramatically impact the χ^2 divergence and thus the distance between the information matrices. We show in Section 5.3 how these distances evolve when training a deep network.

2.2 \mathbf{C} does not approximate \mathbf{F}

\mathbf{C} is often referred to as the ‘‘empirical Fisher’’ matrix, implying that it is an approximation to the true Fisher matrix \mathbf{F} (Martens, 2014). In some recent works (Liang et al., 2017; George et al., 2018) the empirical Fisher matrix is used instead of the Fisher matrix. However, in the general case, there is no guarantee that \mathbf{C} will approximate \mathbf{F} , even in the limit of infinite samples. We now give a simple example highlighting their different roles:

Example 1 (Mean regression) Let $X = (x_i)_{i=1,\dots,N}$ be an i.i.d sequence of random variables. The task is to estimate $\mu = \mathbb{E}[x]$ by minimizing the loss $\ell(\theta) = \frac{1}{2N} \sum_{n=1}^N \|x_n - \theta\|^2$. The minimum is attained at $\theta^{MLE} = \frac{1}{N} \sum_{n=1}^N x_n$. This estimator is consistent and converges to μ at rate $\mathcal{O}(\frac{1}{\sqrt{N}})$.

This problem is an MLE problem if we define $q_\theta(x) = \mathcal{N}(x; \theta, \mathbf{I}_d)$. In this case, we have

$$\mathbf{H}(\theta^{MLE}) = \mathbf{F}(\theta^{MLE}) = \mathbf{I}_d, \quad \mathbf{C}(\theta^{MLE}) = \widehat{\Sigma}_x, \quad (9)$$

where $\widehat{\Sigma}_x$ is the empirical covariance of the x_i ’s. We see that, even in the limit of infinite data, the covariance \mathbf{C} does not converge to the actual Fisher matrix nor the Hessian. Hence we shall and will not refer to \mathbf{C} as the ‘‘empirical Fisher’’ matrix.

In some other settings, however, one can expect a stronger correlation between \mathbf{C} and \mathbf{H} :

Example 2 (Ordinary Least Squares) Let us assume we have a data distribution (x_n, y_n) so that

$$y_n = x_n^\top \theta^* + \epsilon_n, \quad \epsilon_n \sim \mathcal{N}(0, \Sigma) \quad (10)$$

with $y_n \in \mathbb{R}^p$ and $x_n, \epsilon_n \in \mathbb{R}^d$, $\theta \in \mathbb{R}^{p \times d}$. We train a model to minimize the sum of squares residual

$$\min_{\theta} \frac{1}{2N} \sum_{n=1}^N \|y_n - x_n^\top \theta\|^2. \quad (11)$$

In that case, for $\theta = \theta^*$

$$\mathbf{H} = \mathbf{F} = \mathbb{E}_p[xx^\top], \quad \mathbf{C} = \mathbb{E}_p[x\Sigma x^\top] \quad (12)$$

First, we observe that the Hessian and the Fisher are equal, for all θ . Second, when the input/output covariance matrix is isotropic $\Sigma = \sigma^2 \mathbf{I}$, then we have $\mathbf{C} \propto \mathbf{H} = \mathbf{F}$.

3 Information matrices in optimization

Now that the distinction between these matrices has been clarified, we can explain how their interplay affects optimization. In this section, we offer theoretical results, as well as a small empirical study, on the impact of the noise geometry on convergence rates.

3.1 Convergence rates

We start here by expanding the the result of Schmidt (2014) to full matrices, expliciting how the interplay of noise and curvature affects optimization. We then build a toy experiment to validate the results.

3.1.1 General setting

Proposition 2 (Function value) Let f be a twice-differentiable function. Assume f is μ -strongly convex and there are two matrices \mathbf{H} and \mathbf{S} such that for all θ, θ' :

$$f(\theta') \leq f(\theta) + \nabla f(\theta)^\top (\theta' - \theta) + \frac{1}{2} (\theta' - \theta)^\top \mathbf{H} (\theta' - \theta) \\ \mathbb{E}_p[\nabla \ell(\theta, x) \nabla \ell(\theta, x)^\top] \preceq \mathbf{S} + \nabla f(\theta) \nabla f(\theta)^\top.$$

Then stochastic gradient with stepsize α and positive definite preconditioning matrix \mathbf{M} satisfies

$$\mathbb{E}[\Delta_k] \leq (1 - 2\alpha\mu_M\mu)^k \Delta_0 + \frac{\alpha}{4\mu_M\mu} \text{Tr}(\mathbf{HMSM}),$$

where μ_M is the smallest eigenvalue of $\mathbf{M} - \frac{\alpha}{2} \mathbf{M}^\top \mathbf{H} \mathbf{M}$.

\mathbf{S} is the centered covariance matrix of the stochastic gradients and, if f is quadratic, then \mathbf{H} is the Hessian.

3.1.2 Centered and uncentered covariance

Proposition 2, as well as most results on optimization, uses a bound on the uncentered covariance of the gradients. The result is that the noise must be lower far away from the optimum, where the gradients are high. Thus, it seems more natural to define convergence rates as a function of the centered gradients’ covariance \mathbf{S} , although these results are usually weaker as a consequence of the relaxed assumption. For the remainder

of this section, focused on the quadratic case, we will use \mathbf{S} . Note that the two matrices are equal at any first-order stationary point.

Centered covariance matrices have been used in the past to derive convergence rates, for instance by Bach & Moulines (2013); Flammarion & Bach (2015); Dieuleveut et al. (2016). These works also include a dependence on the geometry of the noise since their constraint is of the form $\mathbf{S} \preceq \sigma^2 \mathbf{H}$. In particular, if \mathbf{S} and \mathbf{H} are not aligned, σ^2 must be larger for the inequality to hold.

3.1.3 Quadratic functions

Proposition 3 (Quadratic case) *Assuming we minimize a quadratic function*

$$f(\theta) = \frac{1}{2}(\theta - \hat{\theta}^*)^\top \mathbf{H}(\theta - \hat{\theta}^*)$$

only having access to a noisy estimate of the gradient $g(\theta) \sim \nabla f(\theta) + \epsilon$ with ϵ a zero-mean random variable with covariance $\mathbb{E}[\epsilon\epsilon^\top] = \mathbf{S}$, then the iterates obtained using stochastic gradient with stepsize α and preconditioning psd matrix \mathbf{M} satisfy

$$\mathbb{E}[\theta^k - \hat{\theta}^*] = (1 - \alpha \mathbf{M} \mathbf{H})^k (\theta^0 - \hat{\theta}^*).$$

Further, the covariance of the iterates $\Sigma_k = \mathbb{E}[(\theta^k - \hat{\theta}^)(\theta^k - \hat{\theta}^*)^\top]$ satisfies*

$$\Sigma_{k+1} = (\mathbf{I} - \alpha \mathbf{M} \mathbf{H}) \Sigma_k (\mathbf{I} - \alpha \mathbf{M} \mathbf{H})^\top + \alpha^2 \mathbf{M} \mathbf{S} \mathbf{M}^\top.$$

In particular, the stationary distribution of the iterates has a covariance Σ_∞ which verifies the equation

$$\Sigma_\infty \mathbf{H} \mathbf{M} + \mathbf{M} \mathbf{H} \Sigma_\infty = \alpha \mathbf{M} (\mathbf{S} + \mathbf{H} \Sigma_\infty \mathbf{H}) \mathbf{M}.$$

Recently, analyzing the stationary distribution of SGD by modelling its dynamics as a stochastic differential equation (SDE) has gained traction in the machine learning community (Chaudhari & Soatto, 2017; Jastrzebski et al., 2017). Worthy of note, Mandt et al. (2017); Zhu et al. (2018) do not make assumptions about the structure of the noise matrix \mathbf{S} . Our proposition above extends and corrects some of their results as it does not rely on the continuous-time approximation of the dynamics of SGD. Indeed, as pointed out in Yaida (2018), most of the works using the continuous-time approximation implicitly make the confusion between centered \mathbf{S} and uncentered \mathbf{C} covariance matrices of the gradients.

Proposition 4 (Limit cycle of SG) *If f is a quadratic function and \mathbf{H} , \mathbf{C} and \mathbf{M} are simultaneously diagonalizable, then stochastic gradient with symmetric positive definite preconditioner \mathbf{M} and stepsize α yields*

$$\mathbb{E}[\Delta_t] = \frac{\alpha}{2} \text{Tr}((2\mathbf{I} - \alpha \mathbf{M} \mathbf{H})^{-1} \mathbf{M} \mathbf{S}) + \mathcal{O}(e^{-t}). \quad (13)$$

Rather than using a preconditioner, another popular method to reduce the impact of the curvature is Polyak momentum, defined as

$$v_0 = 0 \quad , \quad v_t = \gamma v_{t-1} + \nabla_\theta \ell(\theta_t, x_t) \\ \theta_{t+1} = \theta_t - \alpha v_t.$$

Proposition 5 (Limit cycle of momentum) *If f is a quadratic function and \mathbf{H} and \mathbf{C} are simultaneously diagonalizable, then Polyak momentum with parameter γ and stepsize α yields*

$$\mathbb{E}[\Delta_t] = \frac{\alpha}{2} \frac{(1+\gamma)}{(1-\gamma)} \text{Tr}((2(1+\gamma)\mathbf{I} - \alpha \mathbf{H})^{-1} \mathbf{S}) + \mathcal{O}(e^{-t}). \quad (14)$$

4 Generalization

So far, we focused on the impact of the interplay between curvature and noise in the optimization setting. However, optimization, i.e. reaching low loss on the training set, is generally not the ultimate goal as one would rather reach a low test loss. The difference between training and test loss is called the generalization gap and estimating it has been the focus of many authors (Keskar et al., 2017; Neyshabur et al., 2017; Liang et al., 2017; Novak et al., 2018; Rangamani et al., 2019).

We believe there is a fundamental misunderstanding in several of these works, stemming from the confusion between curvature and noise. Rather than proposing a new metric, we empirically show how the Takeuchi information criterion (TIC: Takeuchi, 1976) addresses these misunderstandings. It makes use of both the Hessian of the loss with respect to the parameters, \mathbf{H} , and the uncentered covariance of the gradients, \mathbf{C} . While the former represents the curvature of the loss, i.e., the sensitivity of the gradient to a change in parameter space, the latter represents the sensitivity of the gradient to a change in inputs. As the generalization gap is a direct consequence of the discrepancy between training and test sets, the influence of \mathbf{C} is natural. Thus, our result further reinforces the idea that the Hessian cannot by itself be used to estimate the generalization gap, an observation already made by Dinh et al. (2017), among others.

4.1 Takeuchi information criterion

In the simplest case of a well specified least squares regression problem, an unbiased estimator of the generalization gap is the AIC (Akaike, 1974), which is simply the number of degrees of freedom divided by the number of samples: $\hat{\mathcal{G}}(\theta) = \frac{1}{N} d$ where d is the dimensionality of θ . This estimator is valid locally around the maximum likelihood parameters computed on the training data. However, these assumptions do

not hold in most cases, leading to the number of parameters being a poor predictor of the generalization gap (Novak et al., 2018). When dealing with maximum likelihood estimation (MLE) in misspecified models, a more general formula for estimating the gap is given by the Takeuchi information criterion (TIC: Takeuchi, 1976):

$$\hat{\mathcal{G}} = \frac{1}{N} \text{Tr}(\mathbf{H}(\hat{\theta}^*)^{-1} \mathbf{C}(\hat{\theta}^*)), \quad (15)$$

where $\hat{\theta}^*$ is a local optimum. Note that \mathbf{H} and \mathbf{C} here are the hessian and covariance of the gradients matrices computed on the *true data distribution*.

This criterion is not new in the domain of machine learning. It was rediscovered by Murata et al. (1994) and similar criteria have been proposed since then (Beirami et al., 2017; Wang et al., 2018). However, as far as we know, no experimental validation of this criterion has been carried out for deep networks. Indeed, for deep networks, \mathbf{H} is highly degenerate, most of its eigenvalues being close to 0 (Sagun et al., 2016). In this work, the Takeuchi information criterion is computed, in the degenerate case, by only taking into account the eigenvalues of the Hessian of significant magnitude. In practice, we cut all the eigenvalues smaller than a constant times the biggest eigenvalue and perform the inversion on that subspace. Details can be found in appendix B.1.

Interestingly, the term $\text{Tr}(\mathbf{H}^{-1} \mathbf{C})$ appeared in several works before, whether as an upper bound on the suboptimality (Flammarion & Bach, 2015) or as a number of iterates required to reach a certain suboptimality (Bottou & Bousquet, 2008). Sadly, it is hard to estimate for large networks but we propose an efficient approximation in Section 5.5.1.

4.2 Limitations of flatness and sensitivity

We highlight here two commonly used estimators of the generalization gap as they provide good examples of failure modes that can occur when not taking the noise into account. This is not to mean that these estimators cannot be useful for the models that are common nowadays, rather that they are bound to fail in some cases.

Flatness (Hochreiter & Schmidhuber, 1997) links the spectrum of the Hessian at a local optimum with the generalization gap. This correlation, observed again by Keskar et al. (2017), was already shown to not hold in general (Dinh et al., 2017). As we showed in Section 2.2, the Hessian does not capture the covariance of the data, which is linked to the generalization gap through the central-limit theorem.

Sensitivity (Novak et al., 2018) links the generaliza-

tion gap to the derivative of the loss with respect to the input. The underlying idea is that we can expect some discrepancy between train and test data, which will induce changes in the output and a potentially higher test loss. However, penalizing the norm of the Jacobian assumes that changes between train and test data will be isotropic. In practice, we can expect data to vary more along some directions, which is not reflected in the sensitivity. In the extreme case where the test data is exactly the same as the training data, the generalization gap will be 0, which will again not be captured by the sensitivity. In practice, whitening the data makes the sensitivity appropriate, save for a scaling factor, as we will see in the experiments.

5 Experiments

We now provide experimental validation of all the results in this paper. We start by analyzing the distance between information matrices, first showing its poor correlation with the training loss, then showing that these matrices appear to be remarkably aligned, albeit with different scales, when training deep networks on standard datasets.

5.1 Discrepancies between \mathbf{C} , \mathbf{H} and \mathbf{F}

5.1.1 Experimental setup

For comparing the similarities and discrepancies between the information matrices, we tested

- 5 different architectures: logistic regression, a 1-hidden layer and 2-hidden layer fully connected network, and 2 small convolutional neural networks (CNNs, one with batch normalization (Ioffe & Szegedy, 2015) and one without);
- 3 datasets: MNIST, CIFAR-10, SVHN;
- 3 learning rates: 10^{-2} , $5 \cdot 10^{-3}$, and 10^{-3} , using SGD with momentum $\mu = 0.9$;
- 2 batch sizes: 64, 512;
- 5 dataset sizes: 5k, 10k, 20k, 25k, and 50k.

We train for 750k steps and compute the metrics every 75k steps. To be able to compute all the information matrices exactly, we reduced the input dimension by converting all images to greyscale and resizing them to 7×7 pixels. While this makes the classification task more challenging, our neural networks still exhibit the behaviour of larger ones by their ability to fit the training set, even with random labels. Details and additional figures can be found in appendix B.2.

5.2 Comparing Fisher and empirical Fisher

Figure 1 shows the squared Frobenius norm between \mathbf{F} and \mathbf{C} (on training data) for many architectures,

datasets, at various stages of the optimization. We see that, while the two matrices eventually coincide on the training set for some models, the convergence is very weak as even low training errors can lead to a large discrepancy between these two matrices. In practice, \mathbf{C} and \mathbf{F} might be significantly different, even when computed on the training set.




Figure 1: Squared Frobenius norm between $\bar{\mathbf{F}}$ and $\bar{\mathbf{C}}$ (computed on the training distribution). Even for some low training losses, there can be a significant difference between the two matrices.

5.3 Comparing \mathbf{H} , \mathbf{F} and \mathbf{C}

In this subsection, we analyze the similarities and differences between the information matrices. We will focus on the scale similarity r , defined as the ratio of traces, and the angle similarity s , defined as the cosine between matrices. Note that having both $r(\mathbf{A}, \mathbf{B}) = 1$ and $s(\mathbf{A}, \mathbf{B}) = 1$ implies $\mathbf{A} = \mathbf{B}$.




Figure 2: Scale and angle similarities between information matrices.

Figure 2 shows the scale (left) and angle (right) similarities between the three pairs of matrices during the optimization of all models used in figure 4. We can see that \mathbf{H} is not aligned with \mathbf{C} nor \mathbf{F} at the beginning of the optimization but this changes quickly. Then, all three matrices reach a very high cosine similarity, much higher than we would obtain for two random low-rank matrices. For the scaling, \mathbf{C} is “larger” than the other two while \mathbf{F} and \mathbf{H} are very close to each other. Thus, as in the least squares case, we have $\mathbf{C} \approx \mathbf{F} \approx \mathbf{H}$.

5.4 Impact of noise on second-order methods

Section 3 extended existing results to take the geometry of the noise and the curvature into account. Here, we show how the geometry of the noise, and in particular its relationship to the Hessian, can make or break second-order methods in the stochastic setting. To be clear, we assume here that we have access to the full Hessian and do not address the issue of estimating it from noisy samples.

We assume a quadratic $\ell(\theta) = \frac{1}{2}\theta^\top \mathbf{H}\theta$ with $\theta \in \mathbb{R}^{20}$ and $\mathbf{H} \in \mathbb{R}^{20 \times 20}$ a diagonal matrix such that $\mathbf{H}_{ii} = i^2$ with a condition number $d^2 = 400$. At each timestep, we have access to an oracle that outputs a noisy gradient, $\mathbf{H}\theta_t + \epsilon$ with ϵ drawn from a zero-mean Gaussian with covariance \mathbf{S} . Note here that \mathbf{S} is the *centered* covariance of the gradients. We consider three settings: a) $\mathbf{S} = \alpha_1 \mathbf{H}$; b) $\mathbf{S} = \mathbf{I}$; c) $\mathbf{S} = \alpha_{-1} \mathbf{H}^{-1}$ where the constants α_1 and α_{-1} are chosen such that $\text{Tr}(\mathbf{S}) = d$. Hence, these three settings are indistinguishable from the point of view of the rate of Schmidt (2014).

In this simplified setting, we get an analytic formula for the variance at each timestep and we can compute the exact number of steps t such that $\mathbb{E}[\Delta_t]$ falls below a suboptimality threshold. To vary the impact of the noise, we compute the number of steps for three different thresholds: a) $\epsilon = 1$; b) $\epsilon = 0.1$; c) $\epsilon = 0.01$. For each algorithm and each noise, the stepsize is optimized to minimize the number of steps required to reach the threshold.

The results are in Table 1. We see that, while Stochastic gradient and momentum are insensitive to the geometry of the noise for small ϵ , Newton method is not and degrades when the noise is large in low curvature directions. For $\epsilon = 10^{-2}$ and $\mathbf{S} \propto \mathbf{H}^{-1}$, Newton is worse than SG, a phenomenon that is not captured by the bounds of Bottou & Bousquet (2008) since they do not take the structure of the curvature and the noise into account. We also see that the advantage of Polyak momentum over stochastic gradient disappears when the suboptimality is small, i.e. when the noise is large compared to the signal.

Also worthy of notice is the fixed stepsize required to achieve suboptimality ϵ , as shown in Table 2. While it hardly depends on the geometry of the noise for SG and Polyak, Newton method requires much smaller stepsizes when \mathbf{S} is anticorrelated with \mathbf{H} to avoid amplifying the noise.

5.5 The TIC and the generalization gap

We now empirically test the quality of the TIC as an estimator of the generalization gap in deep networks. Following Neyshabur et al. (2017) we assess the be-

ϵ	Method	$\beta = 1$	$\beta = 0$	$\beta = -1$
10^0	SG	44	43	42
	Newton	3	2	19
	Polyak	36	36	34
10^{-1}	SG	288	253	207
	Newton	3	28	225
	Polyak	119	111	97
10^{-2}	SG	2090	1941	1731
	Newton	29	315	2663
	Polyak	1743	1727	1705

Table 1: Number of updates required to reach suboptimality of ϵ for various methods and $\mathbf{S} \propto \mathbf{H}^\beta$.

ϵ	Method	$\beta = 1$	$\beta = 0$	$\beta = -1$
10^0	SG	$5 \cdot 10^{-3}$	$5 \cdot 10^{-3}$	$5 \cdot 10^{-3}$
	Newton	$1 \cdot 10^0$	$1 \cdot 10^0$	$2 \cdot 10^{-1}$
	Polyak	$5 \cdot 10^{-3}$	$4 \cdot 10^{-3}$	$5 \cdot 10^{-3}$
10^{-1}	SG	$4 \cdot 10^{-3}$	$4 \cdot 10^{-3}$	$5 \cdot 10^{-3}$
	Newton	$1 \cdot 10^0$	$2 \cdot 10^0$	$3 \cdot 10^{-2}$
	Polyak	$2 \cdot 10^{-3}$	$2 \cdot 10^{-3}$	$3 \cdot 10^{-3}$
10^{-2}	SG	$1 \cdot 10^{-3}$	$1 \cdot 10^{-3}$	$2 \cdot 10^{-3}$
	Newton	$2 \cdot 10^{-1}$	$2 \cdot 10^{-2}$	$3 \cdot 10^{-3}$
	Polyak	$3 \cdot 10^{-4}$	$3 \cdot 10^{-4}$	$3 \cdot 10^{-4}$

Table 2: Stepsizes achieving suboptimality ϵ in the fewest updates for various methods and $\mathbf{S} \propto \mathbf{H}^\beta$.

haviour of our generalization gap estimator by varying (1) the number of parameters in a model and (2) the label randomization ratio.

Experiments are performed using a fully connected feedforward network with a single hidden layer trained on a subset of 2k samples of SVHN (Netzer et al., 2011). In Figure 3a we vary the number of units in the hidden layer without label randomization while in Figure 3b we vary the label randomization ratio with a fixed architecture. Each point is computed using 3 different random number generator seeds. The neural networks are trained for 750k steps. The confidence intervals are provided using bootstrapping to estimate a 95% confidence interval. The Hessian, covariance matrices and sensitivity are computed on a subset of size 5k of the test data. Details can be found in Appendix B.2.

We now study the ability of the TIC across a wide variety of models, datasets, and hyperparameters. More specifically, we compare the TIC to the generalization gap for: The experiments of Figure 4 are performed with the experimental setup presented in subsection 5.1.1. Figure 4a shows that the TIC using \mathbf{H} and \mathbf{C} computed over the test set is an excellent estimator of the generalization gap. For comparison, we also show in Figure 4b the generalization gap as a function of \mathbf{H} computed over the test set. We see that,




Figure 3: Comparing the TIC to other estimators of the generalization gap on SVHN. The TIC matches the generalization gap more closely than both the AIC and the sensitivity.

even when using the test set, the correlation is much weaker than with the TIC.

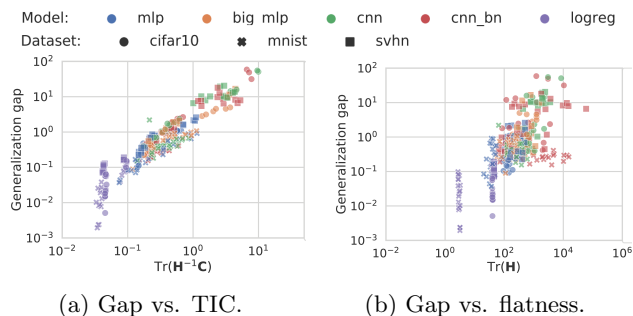


Figure 4: Generalization gap as a function of the Takeuchi information criterion (*left*) and the trace of the Hessian on the test set (*right*) for many architectures, datasets, and hyperparameters. Correlation is perfect if all points lie on a line. We see that the Hessian cannot by itself capture the generalization gap.

5.5.1 Efficient approximations to the TIC

Although the TIC is a good estimate of the generalization gap, it can be expensive to compute on large models. Following our theoretical and empirical analysis of the proximity of \mathbf{H} and \mathbf{F} , we propose two approximations to the TIC: $\text{Tr}(\mathbf{F}^{-1}\mathbf{C})$ and $\text{Tr}(\mathbf{C})/\text{Tr}(\mathbf{F})$. They are easier to compute as the \mathbf{F} is in general easier to compute than \mathbf{H} and the second does not require any matrix inversion.

Using the same experimental setting as in 5.5, we observe in Figure 5 that the replacing \mathbf{H} with \mathbf{F} leads to almost no loss in predictive performance. On the other hand, the ratio of the traces works best when the generalization gap is high and tends to overestimate it when it is small.

Intuition on $\text{Tr}(\mathbf{C})/\text{Tr}(\mathbf{F})$: it is not clear right away




Figure 5: Generalization gap as a function of two approximations to the Takeuchi Information Criterion: $\text{Tr}(\mathbf{F}^{-1}\mathbf{C})$ (left) and $\text{Tr}(\mathbf{C})/\text{Tr}(\mathbf{F})$ (right).

why the ratio of traces might be an interesting quantity. However, as observed in figure 2, \mathbf{C} and \mathbf{F} are remarkably aligned, but there remains a scaling factor. If we had $\mathbf{C} = \alpha\mathbf{F}$, then $\text{Tr}(\mathbf{F}^{-1}\mathbf{C}) = k\alpha$ where k is the dimension of the invertible subspace of \mathbf{F} and $\text{Tr}(\mathbf{C})/\text{Tr}(\mathbf{F}) = d\alpha$ where d is the dimensionality of θ . So, up to a multiplicative constant (or an offset in log scale), we can expect these two quantities to exhibit similarities. Notice that on figure 5, this offset does appear and is different for every dataset (MNIST has the smallest one, then SVHN and CIFAR10, just slightly bigger).

5.6 The importance of the noise in estimating the generalization gap

For a given model, the generalization gap captures the discrepancy that exists between the training set and the data distribution. Hence, estimating that gap involves the evaluation of the uncertainty around the data distribution. The TIC uses \mathbf{C} to capture that uncertainty but other measures probably exist. However, estimators which do not estimate it are bound to have failure modes. For instance, by using the square norm of the derivative of the loss with respect to the input, the sensitivity implicitly assumes that the uncertainty around the inputs is isotropic and will fail should the data be heavily concentrated in a low-dimensional subspace. It would be interesting to adapt the sensitivity to take the covariance of the inputs into account.

Another aspect worth mentioning is that estimators such as the margin assume that the classifier is fixed but the data is a random variable. Then, the margin quantifies the probability that a new datapoint would fall on the other side of the decision boundary. By contrast, the TIC assumes that the data are fixed but that the classifier is a random variable. It estimates the probability that a classifier trained on slightly different data would classify a training point incorrectly. In that, it echoes the uniform stability theory (Bousquet

& Elisseeff, 2002), where a full training with a slightly different training set has been replaced with a local search.

6 Conclusion and open questions

We clarified the relationship between information matrices used in optimization. While their differences seem obvious in retrospect, the widespread confusion makes these messages necessary. Indeed, several well-known algorithms, such as Adam (Kingma & Ba, 2014), claiming to use second-order information about the loss to accelerate training seem instead to be using the covariance matrix of the gradients. Equipped with this new understanding of the difference between the curvature and noise information matrices, one might wonder if the success of these methods is not due to variance reduction instead. If so, one should be able to combine variance reduction and geometry adaptation, an idea attempted by Le Roux et al. (2011).

We also showed how, in certain settings, the geometry of the noise could affect the performance of second-order methods. While Polyak momentum is affected by the scale of the noise, its performance is independent of the geometry, similar to stochastic gradient but unlike Newton method. However, empirical results indicate that common loss functions are in the regime favorable to second-order methods.

Finally, we investigated whether the Takeuchi information criterion is relevant for estimating the generalization gap in neural networks. We provided evidence that this complexity measures involving the information matrices is predictive of the generalization performance.

We hope this study will clarify the interplay of the noise and curvature in common machine learning settings, potentially giving rise to new optimization algorithms as well as new methods to estimate the generalization gap.

Acknowledgments

We would like to thank Gauthier Gidel, Reyhane Askari and Giancarlo Kerg for reviewing an earlier version of this paper. We also thank Aristide Baratin for insightful discussions. Valentin Thomas acknowledges funding from the Open Philanthropy project.

References

- Akaike, Hirotugu. A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6):716–723, 1974.
- Bach, Francis and Moulines, Eric. Non-strongly-convex smooth stochastic approximation with convergence

- rate $o(1/n)$. In *Advances in Neural Information Processing Systems*, pp. 773–781, 2013.
- Beirami, Ahmad, Razaviyayn, Meisam, Shahrampour, Shahin, and Tarokh, Vahid. On optimal generalizability in parametric learning. In *Advances in Neural Information Processing Systems*, pp. 3455–3465, 2017.
- Bottou, Léon and Bousquet, Olivier. The tradeoffs of large scale learning. In *Advances in neural information processing systems*, pp. 161–168, 2008.
- Bousquet, Olivier and Elisseeff, André. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002.
- Chaudhari, Pratik and Soatto, Stefano. Stochastic gradient descent performs variational inference, converges to limit cycles for deep networks. *arXiv preprint arXiv:1710.11029*, 2017.
- Dieuleveut, Aymeric, Bach, Francis, et al. Nonparametric stochastic approximation with large step-sizes. *The Annals of Statistics*, 44(4):1363–1399, 2016.
- Dinh, Laurent, Pascanu, Razvan, Bengio, Samy, and Bengio, Yoshua. Sharp minima can generalize for deep nets. *International Conference on Machine Learning (ICML)*, 2017.
- Flammarion, Nicolas and Bach, Francis. From averaging to acceleration, there is only a step-size. In *Conference on Learning Theory*, pp. 658–695, 2015.
- George, Thomas, Laurent, César, Bouthillier, Xavier, Ballas, Nicolas, and Vincent, Pascal. Fast approximate natural gradient descent in a kronecker factored eigenbasis. In *Advances in Neural Information Processing Systems*, pp. 9550–9560, 2018.
- Hochreiter, Sepp and Schmidhuber, Jürgen. Flat minima. *Neural Computation*, 9(1):1–42, 1997.
- Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Jastrzebski, Stanislaw, Kenton, Zachary, Arpit, Devansh, Ballas, Nicolas, Fischer, Asja, Bengio, Yoshua, and Storkey, Amos. Three factors influencing minima in sgd. *arXiv preprint arXiv:1711.04623*, 2017.
- Keskar, Nitish Shirish, Mudigere, Dheevatsa, Nocedal, Jorge, Smelyanskiy, Mikhail, and Tang, Ping Tak Peter. On large-batch training for deep learning: Generalization gap and sharp minima. *International Conference on Learning Representations (ICLR)*, 2017.
- Kingma, Diederik P and Ba, Jimmy. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kunstner, Frederik, Balles, Lukas, and Hennig, Philipp. Limitations of the empirical fisher approximation. *arXiv preprint arXiv:1905.12558*, 2019.
- Le Roux, Nicolas, Bengio, Yoshua, and Fitzgibbon, Andrew. Improving first and second-order methods by modeling uncertainty. *Optimization for Machine Learning*, pp. 403, 2011.
- Liang, Tengyuan, Poggio, Tomaso, Rakhlin, Alexander, and Stokes, James. Fisher-rao metric, geometry, and complexity of neural networks. *arXiv preprint arXiv:1711.01530*, 2017.
- Mandt, Stephan, Hoffman, Matthew D, and Blei, David M. Stochastic gradient descent as approximate bayesian inference. *The Journal of Machine Learning Research*, 18(1):4873–4907, 2017.
- Martens, James. New insights and perspectives on the natural gradient method. *arXiv preprint arXiv:1412.1193*, 2014.
- Murata, Noboru, Yoshizawa, Shuji, and Amari, Shun-ichi. Network information criterion-determining the number of hidden units for an artificial neural network model. *IEEE Transactions on Neural Networks*, 5(6):865–872, 1994.
- Netzer, Yuval, Wang, Tao, Coates, Adam, Bissacco, Alessandro, Wu, Bo, and Ng, Andrew Y. Reading digits in natural images with unsupervised feature learning. *Neural Information Processing Systems (NeurIPS)*, 2011.
- Neysshabur, Behnam, Bhojanapalli, Srinadh, McAllester, David, and Srebro, Nati. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pp. 5947–5956, 2017.
- Novak, Roman, Bahri, Yasaman, Abolafia, Daniel A, Pennington, Jeffrey, and Sohl-Dickstein, Jascha. Sensitivity and generalization in neural networks: an empirical study. *International Conference on Learning Representations (ICLR)*, 2018.
- Rangamani, Akshay, Nguyen, Nam H, Kumar, Abhishek, Phan, Dzung, Chin, Sang H, and Tran, Trac D. A scale invariant flatness measure for deep network minima. *arXiv preprint arXiv:1902.02434*, 2019.
- Sagun, Levent, Bottou, Léon, and LeCun, Yann. Eigenvalues of the hessian in deep learning: Singularity and beyond. *arXiv preprint arXiv:1611.07476*, 2016.
- Schmidt, Mark. Convergence rate of stochastic gradient with constant step size. Technical report, UBC, 2014.
- Takeuchi, Kei. The distribution of information statistics and the criterion of goodness of fit of models. *Mathematical Science*, 153:12–18, 1976.

Wang, Shuaiwen, Zhou, Wenda, Lu, Haihao, Maleki, Arian, and Mirrokni, Vahab. Approximate leave-one-out for fast parameter tuning in high dimensions. *arXiv preprint arXiv:1807.02694*, 2018.

Yaida, Sho. Fluctuation-dissipation relations for stochastic gradient descent. *arXiv preprint arXiv:1810.00004*, 2018.

Zhu, Zhanxing, Wu, Jingfeng, Yu, Bing, Wu, Lei, and Ma, Jinwen. The anisotropic noise in stochastic gradient descent: Its behavior of escaping from minima and regularization effects. *arXiv preprint arXiv:1803.00195*, 2018.