

Implicit methods for iterative estimation with large data sets

A dissertation presented

by

Panagiotis Toulis

to

The Department of Statistics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Statistics

Harvard University

Cambridge, Massachusetts

February 2016

© 2016 Panagiotis Toulis

All rights reserved.

Dissertation Advisors:

Associate Professor Edoardo M. Airoidi
Profs. Donald B. Rubin & David C. Parkes

Author:

Panagiotis Toulis

Implicit methods for iterative estimation with large data sets

Abstract

The ideal estimation method needs to fulfill three requirements: (i) efficient computation, (ii) statistical efficiency, and (iii) numerical stability. The classical stochastic approximation of [Robbins and Monro \(1951\)](#) is an iterative estimation method, where the current iterate (parameter estimate) is updated according to some discrepancy between what is observed and what is expected assuming the current iterate has the true parameter value. Classical stochastic approximation undoubtedly meets the computation requirement, which explains its widespread popularity, for example, in modern applications of machine learning with large data sets, but cannot effectively combine it with efficiency and stability. Surprisingly, the stability issue can be improved substantially, if the aforementioned discrepancy is computed not using the current iterate, but using the conditional expectation of the next iterate given the current one. The computational overhead of the resulting *implicit* update is minimal for many statistical models, whereas statistical efficiency can be achieved through simple averaging of the iterates, as in classical stochastic approximation ([Ruppert, 1988](#)). Thus, implicit stochastic approximation is fast and principled, fulfills requirements (i-iii) for a number of popular statistical models including generalized linear models, M-estimation, and proportional hazards, and it is poised to become the workhorse of estimation with large data sets in statistical practice.

Contents

Abstract	iii
Acknowledgments	viii
Introduction	1
1 Implicit Stochastic Approximation	4
1.1 Overview	4
1.2 Introduction	5
1.3 Implicit stochastic approximation	6
1.4 Theory of implicit stochastic approximation	8
1.4.1 Convergence of implicit stochastic approximation	9
1.4.2 Non-asymptotic analysis	10
1.4.3 Asymptotic distribution	12
1.5 Algorithms for implicit stochastic approximation	12
1.6 Application in parameter estimation	14
1.6.1 Likelihood-based estimation	15
1.6.2 Estimation with likelihood known up to normalizing constant	16
1.6.3 Likelihood-free estimation	19
1.7 Conclusion	21
2 Implicit Stochastic Gradient Descent	23
2.1 Overview	23
2.2 Introduction	24
2.3 Overview of main results	27
2.3.1 Illustrative example	28
2.3.2 Experimental evaluation	29
2.3.3 Related work	30
2.4 Theory	32
2.4.1 Non-asymptotic bounds	34
2.4.2 Asymptotic variance and optimal learning rates	35
2.4.3 Optimality with averaging	40

2.4.4	Asymptotic normality	41
2.4.5	Stability	42
2.5	Applications	43
2.5.1	Efficient computation of implicit updates	44
2.5.2	Generalized linear models	44
2.5.3	Cox proportional hazards	46
2.5.4	M-Estimation	48
2.6	Simulation and data analysis	49
2.6.1	Validation of theory	50
2.6.2	Validation of performance	54
2.7	Discussion	62
2.8	Conclusion	66
References		67
Appendix A Appendix to Chapter 1		75
A.1	Theory of implicit stochastic approximation	75
A.1.1	Convergence	76
A.2	Non-asymptotic analysis	77
A.3	Asymptotic distribution	86
Appendix B Appendix to Chapter 2		88
B.1	R code	88
B.2	Useful lemmas	88
B.3	Non-asymptotic analysis	90
B.4	Asymptotic analysis	92
B.5	Stability	102
B.6	Applications	103
B.7	Additional experiments	105
B.7.1	Normality experiments with implicit SGD	105
B.7.2	Poisson regression	105
B.7.3	Experiments with <code>glmnet</code>	107
B.7.4	Support vector machines	108
B.7.5	Experiments with machine learning algorithms	109

List of Tables

2.1	Comparison of implicit SGD and R's <code>glm()</code> function	56
2.2	Comparison of implicit SGD and R's <code>biglm</code> package	57
2.3	Comparison of implicit SGD and R's <code>glmnet</code> package	58
B.1	Comparison between explicit and implicit SGD	108
B.2	Experiments comparing implicit SGD with <code>glmnet</code>	109
B.3	Implicit SGD implementation of support vector machine (SVM)	110
B.4	Comparison of implicit SGD with proximal variance reduction methods	112

List of Figures

1.1	Boxplots of implicit stochastic approximation in quantile estimation problem	22
2.1	Asymptotic variance of SGD methods and AdaGrad	51
2.2	Asymptotic normality of SGD estimates	53
2.3	Implicit SGD on Cox proportional hazards model	60
2.4	Implicit SGD on M-estimation task	61
2.5	Implicit SGD on generalized additive model	63
B.1	Search bounds for solving fixed-point equation of implicit SGD	104
B.2	Asymptotic normality of SGD estimates	106
B.3	Asymptotic normality of SGD estimates in logistic regression	107
B.4	Simulation of averaged implicit SGD	114
B.5	Comparison of SGD methods with proximal variance reduction methods . .	115

Acknowledgments

I was advised by three wonderful people: Professor David C. Parkes, Professor Edoardo M. Airoldi, and Professor Donald B. Rubin. I met David when I was still a visiting student in the Computer Science school at Harvard and, simply put, he has been my role model ever since. His academic prowess and sincere kindness have fueled my love for science and research. Edo was the person who drew me from Computer Science into Statistics as he has an inspiring natural talent for multidisciplinary research. As my primary Ph.D. advisor, he helped me immensely to develop my own independent research. To this end, I have also been greatly benefited from my collaboration with Don on problems of causal inference. Don has a magical analytic ability to dissect even the most complicated arguments down to simple pieces, which is an invaluable skill for research.

I would also like to thank the amazing people in the Department of Statistics at Harvard, especially Professor Xiao-Li Meng and Professor Natesh Pillai. Xiao-Li has been very supportive of my work, and has dedicated time and effort to discuss critical aspects of it. Natesh helped me greatly by giving invaluable feedback on early versions of this work. I am also thankful to the EconCS group in the Computer Science School at Harvard for fostering a collegial environment where I found support and inspiration.

My friends and colleagues at Harvard have been a crucial part of my graduate life and work. I have shared wonderful moments with Bobby, Alex and Rui, whom I met at the Harvard graduate dorms in my first year. I would also like to thank my colleague Edward Kao, with whom we have shared ideas, enthusiasm on research and, more importantly, a cherished friendship through all those years.

Outside of Harvard, I would be remiss if I didn't mention Jason Rennie (Google) since he was pivotal in exposing me to problems of learning with large-scale data sets, a critical theme of this thesis. I interned for two summers in Jason's team and I was immensely benefited from this experience. I would also like to give a special thanks to Leon Bottou (now at Facebook) who has been very helpful in giving feedback on my work. Leon is a world-renowned expert in stochastic gradient methods, and his input has been invaluable.

On a more personal level, I would like to thank my father Athanasios, my mother Sofia, and my brother Kostas. My parents are two wonderful giving people and I owe them everything. Kostas, my brother, is the person I always looked up to as I learned from him what it means to work hard, to dream big, and to persevere. This thesis would be impossible without his life-long mentoring and support.

Last but certainly not least, I wish to thank Chrysa, my beloved wife. My entire academic journey thus far would just be impossible without Chrysa's love and enduring support. She is my secret superpower, the one who made sure that every effort would be worth it. Above all, she gave me Athan, now my 16-month old beloved son, and for that I am eternally grateful. Athan, you make everything matter, you are and always will be in the center of our universe. With tremendous love I dedicate this work to you and your mom. I put passion and effort in this work; as you read these lines, I know we are around to support you in doing the same for anything you have chosen to do.

Dedicated to little Athan.

Introduction

Modern statistical applications typically involve high-dimensional models on large data sets. This has increased the need for estimation procedures that combine both computational and statistical efficiency. Standard popular procedures, such as the EM algorithm or iteratively-reweighted least squares, are computationally expensive and are thus inappropriate in this modern domain.

There is, however, one procedure that crosses the fine boundary between efficient computation and efficient statistical estimation, known as *stochastic gradient descent*, or SGD for short. SGD is an instance of the classical stochastic approximation of [Robbins and Monro \(1951\)](#). In this procedure the current iterate (parameter estimate) is updated according to some discrepancy between what is observed and what is expected assuming the current iterate has the true parameter value. This discrepancy is calculated on a single data point and not on the entire data set, which allows scaling up to very large data sets. However, there are two important issues with the classical SGD procedures:

1. Their statistical properties are largely unknown (e.g., asymptotic efficiency relative to maximum-likelihood)
2. They are numerically unstable and therefore unreliable for statistical analysis.

This can be illustrated through a (simplified) real-world problem that I encountered during my internship at Google, which also provided the original motivation for this research.

Example. Suppose $Y \sim \text{Poisson}(e^{\theta_*})$ represents # of bookings in a travel website. Suppose we have a stream of i.i.d. data points Y_1, Y_2, \dots , and we wish to estimate the parameter θ_* .

Estimation through stochastic gradient descent has the following intuitive form:

$$\theta_n = \theta_{n-1} + \gamma_n(Y_n - e^{\theta_{n-1}}),$$

where $\gamma_n = \gamma_1/n$, $\gamma_1 > 0$, and θ_0 is some initial estimate. Standard theory of stochastic approximations suggests that this procedure converges to a point θ_∞ such that

$$\mathbb{E}(Y_n) = e^{\theta_\infty} \Leftrightarrow \theta_\infty = \theta_*$$

However, in practice, the situation is quite different. Consider a data set where $Y_1 = Y_2 = 1001$, and suppose that $\theta_0 = 0$ and $\gamma_1 = 1$. Then, the first iteration yields

$$\theta_1 = 0 + 1 \cdot (1001 - 1) = 1000,$$

and so the second iteration yields

$$\theta_2 = \theta_1 + \frac{1}{2}(1001 - e^{1000}) = -\infty.$$

It follows that standard SGD procedures are numerically unstable. This instability, stemming from their sensitivity to parameters such as the *learning rate* γ_1 , is arguably the critical reason why they are not popular in standard statistical practice. But even if stability was not an issue, SGD procedures are not well-studied in terms of their statistical efficiency. In this example, we would like to know, say, the asymptotic variance of θ_n , i.e., the limit $n\text{Var}(\theta_n)$. A major result of this thesis is to derive this limit in closed form, more generally, in many popular families of statistical models.

Surprisingly, the stability issue can be improved substantially if the discrepancy, e.g., $Y_n - e^{\theta_{n-1}}$, is computed not using the current iterate, but using the conditional expectation of the next iterate given the current one, or an approximation of it. For instance, consider the procedure

$$\theta_n = \theta_{n-1} + \gamma_n(Y_n - e^{\theta_n})$$

This procedure is using an *implicit update* because the next iterate, θ_n , appears on both sides of the equation. Using the same data as before the new procedure behaves as follows.

The first iteration yields

$$\theta_1 = 0 + 1 \cdot (1001 - e^{\theta_1});$$

thus, $\theta_1 \approx \log(994) \approx 6.902$. The second iteration yields

$$\theta_2 = 6.902 + \frac{1}{2}(1001 - e^{\theta_2});$$

thus, $\theta_2 \approx \log(1001) \approx 6.909$. We see that the stability issue has been completely resolved, whereas there is no apparent cost in statistical efficiency—in fact, the estimate θ_2 in the implicit procedure is very close to the maximum-likelihood estimate ($\log 1001$). There is only a computational cost in computing the fixed-point equation of the implicit procedure, however, we will show later that this cost is negligible in a large family of statistical models.

The main result of this work is that the ensuing *implicit stochastic gradient descent* procedure (implicit SGD) maintains the desirable asymptotic properties of classical SGD (including convergence and statistical efficiency) but has a remarkable numerical stability. It is exactly this combination that makes implicit procedures appropriate for principled statistical analysis with large data sets.

The exposition is organized in two chapters. Chapter 1 provides the foundation of implicit SGD procedures by defining an implicit variant of the classical stochastic approximation procedure of [Robbins and Monro \(1951\)](#). There, we give a full theoretical analysis of the new approximation procedure and show its advantages over the classical one. In Chapter 2 we zoom into the implicit SGD procedure and analyze it vis-à-vis the standard SGD procedure that is currently widely used in practice.

Chapter 1

Implicit Stochastic Approximation

1.1 Overview

The need to carry out parameter estimation from massive data has reinvigorated interest in iterative estimation methods, in statistics and machine learning. Classical work includes deterministic gradient-based methods, such as quasi-Newton, and stochastic gradient descent and its variants, including adaptive learning rates, acceleration and averaging. Current work increasingly relies on methods that employ proximal operators, leading to updates defined through implicit equations, which need to be solved at each iteration. Such methods are especially attractive in modern problems with massive data because they are numerically stable and converge with minimal assumptions, among other reasons. However, while the majority of existing methods can be subsumed into the gradient-free stochastic approximation framework developed by [Robbins and Monro \(1951\)](#), there is no such framework for methods with implicit updates. Here, we conceptualize a gradient-free implicit stochastic approximation procedure, and develop asymptotic and non-asymptotic theory for it. This new framework provides a theoretical foundation for gradient-based procedures that rely on implicit updates, and opens the door to iterative estimation methods that do not require a gradient, nor a fully known likelihood.

1.2 Introduction

[Robbins and Monro \(1951\)](#) considered the problem of estimating the zero θ_* of a function $h : \mathbb{R}^p \rightarrow \mathbb{R}$. Specifically, for every fixed $\theta \in \mathbb{R}^p$, they assumed that the exact value $h(\theta)$ is unknown but can be unbiasedly estimated by a random variable W_θ , such that $\mathbb{E}(W_\theta) = h(\theta)$. Starting from an estimate θ_0^{rm} , [Robbins and Monro \(1951\)](#) recursively approximated θ_* as follows,

$$\theta_n^{\text{rm}} = \theta_{n-1}^{\text{rm}} - \gamma_n W_{\theta_{n-1}^{\text{rm}}}, \quad (1.1)$$

where $\{\gamma_n\}$ is a decreasing sequence of positive numbers, known as the learning rate sequence; typically, $\gamma_n \propto 1/n$ such that $\sum \gamma_i^2 < \infty$ to guarantee convergence, and $\sum \gamma_i = \infty$ to guarantee that convergence can be towards any point in \mathbb{R}^p . [Robbins and Monro \(1951\)](#) proved convergence in quadratic mean of procedure (1.1) —also known as the “Robbins-Monro procedure”—, i.e., $\mathbb{E}(\|\theta_n^{\text{rm}} - \theta_*\|^2) \rightarrow 0$. Since then, several other authors have explored its theoretical properties, for example, [Ljung *et al.* \(1992\)](#); [Kushner and Yin \(2003\)](#); [Borkar \(2008\)](#). Due to its remarkable simplicity and computational efficiency, the Robbins-Monro procedure has found widespread applications across scientific disciplines, including statistics ([Nevelson and Khasminski, 1973](#)), engineering ([Benveniste *et al.*, 1990](#)), and optimization ([Nesterov, 2004](#)).

Recently, the Robbins-Monro procedure has received renewed interest for its applicability in parameter estimation with large data sets, and its connections to stochastic optimization methods. In such settings, even though there is a finite data set D , the Robbins-Monro procedure can be applied with W_θ being, for example, the log-likelihood of θ at a single data point that is sampled with replacement from D . In this case, the theory of [Robbins and Monro \(1951\)](#) implies that $\mathbb{E}_D(\|\theta_n^{\text{rm}} - \hat{\theta}_n\|^2) \rightarrow 0$, where the expectation is now with respect to the empirical distribution of data points in D , and $\hat{\theta}_n$ is an estimator of θ_* , such as the maximum-likelihood estimator, or maximum a-posteriori if regularization is used. In this work, we will study the theoretical properties of a modified stochastic approximation method in the more general setting with a stream of data points, but will also discuss

applications to iterative estimation with a finite data set.

Despite its remarkable properties, a well-known issue with the Robbins-Monro procedure is that the sequence $\{\gamma_n\}$ crucially affects both its numerical stability and convergence. Regarding stability, consider a simple example where θ_n^{rm} in Eq. (1.1) is approximating a scale parameter, and thus needs to be positive. Clearly, even if the previous iterate θ_{n-1}^{rm} is positive, there is no guarantee that the update in Eq. (1.1) will provide a positive next iterate θ_n^{rm} . Regarding convergence, the Robbins-Monro procedure can be arbitrarily slow if γ_n is even slightly misspecified. For example, let $\gamma_n = \gamma_1/n$, and assume there exists the scalar potential $H : \mathbb{R}^p \rightarrow \mathbb{R}$, such that $\nabla H(\theta) = h(\theta)$, for all θ . If H is strongly convex with parameter c , then $\mathbb{E} (\|\theta_n^{\text{rm}} - \theta_*\|^2) = \mathcal{O}(n^{-\epsilon})$ if $\epsilon = 2c\gamma_1 < 1$ (Nemirovski *et al.*, 2009, Section 1); (Moulines and Bach, 2011, Section 3.1). In summary, large learning rates can make the iterates θ_n^{rm} diverge numerically, whereas small rates can make the iterates converge very slowly to θ_* . Importantly, these conflicting requirements for stability and fast convergence are very hard to reconcile in practice, especially in high-dimensional problems (Toulis and Airoidi, 2015b, Section 3.5), which renders the Robbins-Monro procedure, and all its derived methods, inapplicable without heuristic modifications.

1.3 Implicit stochastic approximation

Our idea to improve the Robbins-Monro procedure (1.1) is to transform its iteration into a stochastic fixed-point equation as follows,

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} - \gamma_n W_{\theta_n^{(*)}}, \text{ s.t.} \quad (1.2)$$

$$\theta_n^{(*)} = \theta_{n-1}^{\text{im}} - \gamma_n h(\theta_n^{(*)}) \stackrel{\text{def}}{=} \mathbb{E} (\theta_n^{\text{im}} | \mathcal{F}_{n-1}), \quad (1.3)$$

where \mathcal{F}_{n-1} is the σ -algebra adapted to $\{\theta_0^{\text{im}}, \theta_1^{\text{im}}, \dots, \theta_{n-1}^{\text{im}}\}$. The *implicit stochastic approximation* method of Eq. (1.2) and Eq. (1.3) can be motivated as the limit of a sequence of improved Robbins-Monro procedures, as follows. First, for convenience, drop the superscript *rm* from procedure (1.1), and fix the sample history \mathcal{F}_{n-1} . Then, $\mathbb{E} (\theta_n | \mathcal{F}_{n-1}) = \theta_{n-1} - \gamma_n h(\theta_{n-1}) \triangleq \theta_n^{(1)}$. Assuming that $\theta_n^{(1)}$ can be computed even though h is unknown, then it is reason-

able to expect that it is better to use $\theta_n^{(1)}$ instead of θ_{n-1} when drawing W_θ in procedure (1.1). This leads to update $\theta_n = \theta_{n-1} - \gamma_n W_{\theta_n^{(1)}}$. In turn, this update has expected value $\theta_{n-1} - \gamma_n h(\theta_n^{(1)}) \triangleq \theta_n^{(2)}$, and thus $\theta_n^{(2)}$ can be used instead of $\theta_n^{(1)}$, and so on. This argument can be repeated ad infinitum producing the following sequence of improved Robbins-Monro procedures,

$$\begin{aligned}\theta_n &= \theta_{n-1} - \gamma_n W_{\theta_{n-1}}, \\ \theta_n &= \theta_{n-1} - \gamma_n W_{\theta_n^{(1)}}, \\ \theta_n &= \theta_{n-1} - \gamma_n W_{\theta_n^{(2)}}, \\ &\dots \\ \theta_n &= \theta_{n-1} - \gamma_n W_{\theta_n^{(*)}},\end{aligned}$$

where $\theta_n^{(i)}$ is defined recursively as $\theta_n^{(i)} = \theta_{n-1} - \gamma_n h(\theta_n^{(i-1)})$. The initial iterate, $\theta_n^{(0)} = \theta_{n-1}$, corresponds to the classical Robbins-Monro procedure (1.1). The limit iterate, $\theta_n^{(\infty)} = \theta_n^{(*)}$, is the fixed point of Eq. (1.3), and corresponds to implicit stochastic approximation (1.2).

The improvement achieved by the fixed-point equation (1.3) can be explained through the following simple argument. First, take norms in Eq. (1.3) to obtain

$$\|\theta_{n-1}^{\text{im}} - \theta_\star\|^2 \geq \|\theta_n^{(*)} - \theta_\star\|^2 + 2\gamma_n h(\theta_n^{(*)})^\top (\theta_n^{(*)} - \theta_\star).$$

As before, suppose that the scalar potential H exists and is convex. Then, it follows $h(\theta)^\top (\theta - \theta_\star) \geq 0$, which implies that $\|\theta_n^{(*)} - \theta_\star\|^2 \leq \|\theta_{n-1}^{\text{im}} - \theta_\star\|^2$. Let $\zeta_n = W_{\theta_n^{(*)}} - h(\theta_n^{(*)})$, such that $\mathbb{E}(\zeta_n | \mathcal{F}_{n-1}) = 0$, then Eq. (1.2) can be written as $\theta_n^{\text{im}} = \theta_n^{(*)} - \gamma_n \zeta_n$. Thus, the fixed-point equation contracts $\theta_n^{(*)}$ towards θ_\star , and potentially contracts θ_n^{im} as well. This idea is also closely related to *proximal operators* in optimization because Eq. (1.3) can be written as

$$\theta_n^{(*)} = \arg \min_{\theta} \left\{ \frac{1}{2} \|\theta - \theta_{n-1}^{\text{im}}\|^2 + \gamma_n H(\theta) \right\}. \quad (1.4)$$

The right-hand side of Eq. (1.4) is a stochastic form of a *proximal operator* (Bauschke and Combettes, 2011), mapping θ_{n-1}^{im} to the intermediate point $\theta_n^{(*)}$, and has θ_\star as a fixed point.

Interest on proximal operators has surged in recent years because they are non-expansive and converge with minimal assumptions. Furthermore, they can be applied on non-smooth objectives, and can easily be combined in modular algorithms for optimization in large-scale and distributed settings (Parikh and Boyd, 2013).

On the practical side, the key advantage of the contractive property of implicit stochastic approximation (1.2) compared to classical stochastic approximation (1.1) is that it is no longer required to have small learning rates γ_n for stability. This resolves the conflicting requirements for stability and convergence in classical stochastic approximation, and provides valuable flexibility in choosing the learning rate γ_n . The theoretical results of the following section do confirm that implicit stochastic approximation converges under minimal conditions, maintains the asymptotic properties of classic stochastic approximation, but also has remarkable stability in short-term iterations.

On the downside, the implementation of implicit stochastic approximation is generally a challenging task because h is unknown. In specific, the point $\theta_n^{(*)}$ cannot be computed from Eq. (1.3) because function h is unknown, otherwise we could simply set $\theta_n^{\text{im}} = \theta_n^{(*)}$ in Eq. (1.2). However, as mentioned before, Eq. (1.2) can be written as $\theta_n^{\text{im}} = \theta_n^{(*)} - \gamma_n \xi_n$, where $\mathbb{E}(\xi_n | \mathcal{F}_{n-1}) = 0$, and therefore we only need a noisy estimate of the intermediate point $\theta_n^{(*)}$ in order to implement implicit stochastic approximation. We leverage this result to provide several concrete implementations in Section 1.5.

1.4 Theory of implicit stochastic approximation

The symbol $\|\cdot\|$ denotes the L_1 vector/matrix norm. We also define the error random variables $\xi_n \triangleq W_{\theta_{n-1}^{(*)}} - h(\theta_{n-1}^{(*)})$, such that $\mathbb{E}(\xi_n | \mathcal{F}_{n-1}) = 0$. The parameter space for θ will be \mathbb{R}^p without loss of generality. For a positive scalar sequence a_n , the sequence $b_n = \mathcal{O}(a_n)$ is such that $b_n \leq ca_n$, for some fixed $c > 0$, and every n ; the sequence $b_n = \mathcal{o}(a_n)$ is such that $b_n/a_n \rightarrow 0$ in the limit where $n \rightarrow \infty$. $b_n \downarrow 0$ denotes a positive sequence decreasing towards zero. We further assume that implicit stochastic approximation (1.2) operates under a combination of the following assumptions.

Assumption 1. It holds, $\gamma_n = \gamma_1 n^{-\gamma}$, $\gamma_1 > 0$ and $\gamma \in (1/2, 1]$.

Assumption 2. The regression function h is Lipschitz with parameter L , i.e., for all θ_1, θ_2 ,

$$\|h(\theta_1) - h(\theta_2)\| \leq L\|\theta_1 - \theta_2\|.$$

Assumption 3. For all n , function h satisfies either

(a) $(\theta - \theta_\star)^\top h(\theta) > 0$, for all θ , or

(b) $(\theta_n^{(*)} - \theta_\star)^\top h(\theta_n^{(*)}) \geq \delta_n \|\theta_n^{(*)} - \theta_\star\|^2$, where $\delta_n = \delta_1 n^{-\delta}$, $\delta_1 > 0$ and $0 < \gamma + \delta \leq 1$.

Assumption 4. There exists a scalar potential $H : \mathbb{R}^p \rightarrow \mathbb{R}$ such that $\nabla H(\theta) = h(\theta)$, for all θ .

Assumption 5. There exists fixed $\sigma^2 > 0$ such that, for every n ,

$$\mathbb{E} (\|\xi_n\|^2 | \mathcal{F}_{n-1}) \leq \sigma^2.$$

Assumption 6. Let $\Xi_n \stackrel{\text{def}}{=} \mathbb{E} (\xi_n \xi_n^\top | \mathcal{F}_{n-1})$, then $\|\Xi_n - \Xi\| = \mathcal{O}(1)$, and $\|\Xi_n - \Xi\| \rightarrow 0$ for fixed positive-definite matrix Ξ . Furthermore, if $\sigma_{n,s}^2 = \mathbb{E} \left(\mathbb{I}_{\|\xi_n\|^2 \geq s/\gamma_n} \|\xi_n\|^2 \right)$, then for all $s > 0$, $\sum_{i=1}^n \sigma_{i,s}^2 = \mathcal{O}(n)$ if $\gamma_n \propto n^{-1}$, or $\sigma_{n,s}^2 = \mathcal{O}(1)$ otherwise.

Assumption 3(a) is a typical convexity assumption. Assumption 3(b) is stronger than the convexity assumption, but weaker than the assumption of strong convexity that is typical in the literature. Assumption 5 was first introduced by [Robbins and Monro \(1951\)](#), and has since been standard in stochastic approximation analysis. Assumption 6 is a typical Lindeberg condition that is used to prove asymptotic normality of θ_n^{im} . Overall, our assumptions are weaker than the assumptions used in classical stochastic approximation; compare, for example, Assumptions 1-6 with assumptions (A1)-(A4) of [Borkar \(2008, Section 2.1\)](#), or assumptions of [Benveniste et al. \(1990, Theorem 15\)](#).

1.4.1 Convergence of implicit stochastic approximation

In [Theorem 1](#) we derive a proof of almost-sure convergence of implicit stochastic approximation, which relies on the supermartingale lemma of [Robbins and Siegmund \(1985\)](#). The

proof of all theorems in this document are given in the Appendix.

Theorem 1. *Suppose that Assumptions 1, 2, 3(a), and 5 hold. Then the iterates θ_n^{im} of the implicit stochastic approximation procedure (1.2) converge almost-surely to θ_* ; i.e., $\theta_n^{\text{im}} \rightarrow \theta_*$, such that $h(\theta_*) = 0$, almost-surely.*

The conditions for almost-sure convergence of implicit stochastic approximation are weaker than the conditions for classical stochastic approximation. For example, to show almost-sure convergence for standard stochastic approximation methods, it is typically assumed that the iterates θ_n^{rm} are almost-surely bounded (Borkar, 2008, Assumption (A4)).

1.4.2 Non-asymptotic analysis

In this section, we prove results on upper bounds for the deviance $\mathbb{E}(H(\theta_n^{\text{im}}) - H(\theta_*))$ and the errors $\mathbb{E}(\|\theta_n^{\text{im}} - \theta_*\|^2)$. This provides information on the rate of convergence, as well as the stability of implicit stochastic approximation methods. Theorem 2 on deviance uses Assumption 3(a), which only assumes non-strong convexity of H , whereas Theorem 3 on squared errors uses Assumption 3(b), which is weaker than strong convexity.

Theorem 2. *Suppose that Assumptions 1, 2, 3(a), 4, and 5 hold. Define $\Gamma^2 \triangleq \mathbb{E}(\|\theta_0^{\text{im}} - \theta_*\|^2) + \sigma^2 \sum_{i=1}^{\infty} \gamma_i^2 + \gamma_1^2 \sigma^2$. Then, if $\gamma \in (2/3, 1]$, there exists $n_{0,1} < \infty$ such that, for all $n > n_{0,1}$,*

$$\mathbb{E}(H(\theta_n^{\text{im}}) - H(\theta_*)) \leq \left[\frac{2\Gamma^2}{\gamma\gamma_1} + o(1) \right] n^{-1+\gamma}.$$

If $\gamma \in (1/2, 2/3)$, there exists $n_{0,2} < \infty$ such that, for all $n > n_{0,2}$,

$$\mathbb{E}(H(\theta_n^{\text{im}}) - H(\theta_*)) \leq \left[\Gamma\sigma\sqrt{L\gamma_1} + o(1) \right] n^{-\gamma/2}.$$

Otherwise, $\gamma = 2/3$ and there exists $n_{0,3} < \infty$ such that, for all $n > n_{0,3}$,

$$\mathbb{E}(H(\theta_n^{\text{im}}) - H(\theta_*)) \leq \left[\frac{3 + \sqrt{9 + 4\gamma_1^3 L\sigma^2 / \Gamma^2}}{2\gamma_1 / \Gamma^2} + o(1) \right] n^{-1/3}.$$

There are two main results in Theorem 2. First, the rates of convergence for the deviance are either $O(n^{-1+\gamma})$ or $O(n^{-\gamma/2})$, depending on the learning rate parameter γ . These

rates match standard stochastic approximation results under non-strong convexity; see, for example, Theorem 4 of [Moulines and Bach \(2011\)](#). Second, there is a uniform decay of the expected deviance towards zero, since the constants $n_{0,1}, n_{0,2}, n_{0,3}$ can be made small, depending on the desired accuracy in the constants of the upper-bounds in Theorem 2. In contrast, in standard stochastic approximation methods under non-strong convexity, there is a term $\exp(4L^2\gamma_1^2n^{1-2\gamma})$ ([Moulines and Bach, 2011](#), Theorem 4), which can amplify the initial conditions arbitrarily. Thus, implicit stochastic approximation has similar asymptotic properties to classical stochastic approximation, but is significantly more stable.

Theorem 3. *Suppose that Assumptions 1, 3(b), and 5 hold, and define $\zeta_n \triangleq \mathbb{E} (|\theta_n^{\text{im}} - \theta_\star|^2)$ and $\kappa \triangleq 1 + 2\gamma_1\delta_1$. Then, if $\gamma + \delta < 1$, for every $n > 0$ it holds,*

$$\zeta_n \leq e^{-\log \kappa \cdot n^{1-\gamma-\delta}} \zeta_0 + \sigma^2 \frac{\gamma_1 \kappa}{\delta_1} n^{-\gamma+\delta} + \mathcal{O}(n^{-\gamma+\delta-1}).$$

Otherwise, if $\gamma = 1, \delta = 0$, it holds,

$$\zeta_n \leq e^{-\log \kappa \cdot \log n} \zeta_0 + \sigma^2 \frac{\gamma_1 \kappa}{\delta_1} n^{-1} + \mathcal{O}(n^{-2}).$$

There are two main results in Theorem 3. First, if H is strongly convex ($\delta = 0$), then the rate of convergence of $\mathbb{E} (|\theta_n^{\text{im}} - \theta_\star|^2)$ is $\mathcal{O}(n^{-\gamma})$, which matches the rate of convergence for classical stochastic approximation under strong convexity ([Benveniste et al., 1990](#), Theorem 22, p.244). Second, there is an exponential discounting of initial conditions ζ_0 regardless of the specification of the learning rate parameter γ_1 and the Lipschitz parameter L . In stark contrast, in classical stochastic approximation there exists a term $\exp(L^2\gamma_1^2n^{1-2\gamma})$ in front of the initial conditions ζ_0 , which can make the approximation diverge numerically if γ_1 is misspecified with respect to the Lipschitz parameter L ([Moulines and Bach, 2011](#), Theorem 1). Thus, as in the non-strongly convex case of Theorem 2, implicit stochastic approximation has similar asymptotic rates to classical stochastic approximation, but is also significantly more stable.

1.4.3 Asymptotic distribution

Asymptotic distributions are well-studied in classical stochastic approximation. In this section, we leverage this theory to show that iterates from implicit stochastic approximation are asymptotically normal. The following theorem establishes this result using Theorem 1 of [Fabian \(1968a\)](#); see also ([Ljung et al., 1992](#), Chapter II.8).

Theorem 4. *Suppose that Assumptions 1, 2, 3(a), 5, and 6 hold. Suppose also that $(2\gamma_1 J_h(\theta_*) - I)$ is positive-definite, where $J_h(\theta)$ is the Jacobian of h at θ , and I is the $p \times p$ identity matrix. Then, the iterate θ_n^{im} of implicit stochastic approximation (1.2) is asymptotically normal, such that*

$$n^{\gamma/2}(\theta_n^{\text{im}} - \theta_*) \rightarrow \mathcal{N}_p(0, \Sigma).$$

The covariance matrix Σ is the unique solution of

$$(\gamma_1 J_h(\theta_*) - I/2)\Sigma + \Sigma(\gamma_1 J_h(\theta_*)^\top - I/2) = \Xi.$$

The asymptotic distribution of the implicit iterate θ_n^{im} is identical to the asymptotic distribution of θ_n^{rm} , as derived by [Fabian \(1968a\)](#). Intuitively, in the limit, $\theta_n^{(*)} \approx \theta_{n-1}^{\text{im}} + \mathcal{O}(\gamma_n)$ with high probability, and thus implicit stochastic approximation behaves like the classical one. We also note that if Ξ commutes with $J_h(\theta_*)$ then Σ can be derived in closed-form as $\Sigma = (2\gamma_1 J_h(\theta_*) - I)^{-1}\Xi$.

1.5 Algorithms for implicit stochastic approximation

In the previous section, we showed that implicit stochastic approximation has similar asymptotic properties to classical stochastic approximation, but it is also more stable. Therefore, implicit stochastic approximation is arguably a superior form of stochastic approximation.

However, the main drawback of implicit stochastic approximation is that the intermediate value $\theta_n^{(*)}$, which is necessary to compute the update in Eq. (1.3), is not readily available because it depends on the regression function $h(\cdot)$ that is unknown. Thus, every

implementation of implicit stochastic approximation needs to approximate or estimate $\theta_n^{(*)}$ at every iteration.

One strategy for such implementation is to approximate $\theta_n^{(*)}$ through a separate standard stochastic approximation procedure; i.e., at every n th iteration of procedure (1.2) run a Robbins-Monro procedure x_k for $k = 1, 2, \dots$, starting from $x_0 = \theta_{n-1}^{\text{im}}$ and updating as follows,

$$x_k = x_{k-1} - a_k(\gamma_n W_{x_{k-1}} + x_{k-1} - \theta_{n-1}^{\text{im}}), \quad (1.5)$$

where $a_k \propto 1/k$ satisfies the Robbins-Monro conditions, and $W_{x_{k-1}}$ are independent draws with fixed $\theta = x_{k-1}$. By the properties of the Robbins-Monro procedure, x_k converges to x_∞ that satisfies $\gamma_n h(x_\infty) + x_\infty - \theta_{n-1}^{\text{im}} = 0$, and thus $x_k \rightarrow \theta_n^{(*)}$, by Eq. (1.3). For practical reasons, it would be sufficient to apply algorithm (1.5) for a finite number of steps, say K , even before x_k has converged to $\theta_n^{(*)}$, and then use x_K in place of $\theta_n^{(*)}$ in procedure (1.2). We note that Eq. (1.5) is in fact a stochastic fixed-point iteration, where one critical necessary condition is that the mapping on the right-hand side of Eq. (1.5) is non-expansive (Borkar, 2008, Theorem 4, Section 10.3); this condition might be hard to validate in practice.

Another implementation is possible if the analytic form of W_θ is known, even though the regression function $h(\theta)$ is not. The idea is, rather counter-intuitively, to use the next iterate θ_n^{im} as an *estimator* of $\theta_n^{(*)}$. This is reasonable because, by definition, $\mathbb{E}(\theta_n^{\text{im}} | \mathcal{F}_{n-1}) = \theta_n^{(*)}$, i.e., θ_n^{im} is an unbiased estimator of $\theta_n^{(*)}$. Thus, procedure (1.2) could be approximately implemented as follows,

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} - \gamma_n W_{\theta_n^{\text{im}}}. \quad (1.6)$$

The next iterate θ_n^{im} appears on both sides of Eq. (1.6) and thus Eq. (1.6) is a multi-dimensional fixed-point equation, also known as *implicit update*. Algorithm (1.6) appears to be the most practical implementation of implicit stochastic approximation, which justifies why the implicit update of algorithm (1.6) lends its name to implicit stochastic approximation. Stochastic approximation with implicit updates will be the focus of the next chapter of

this document.

Now, the main problem in applying algorithm (1.6) is to solve efficiently the fixed-point equation at every iteration and calculate θ_n^{im} . To solve this problem we distinguish two cases. First, let $W_\theta \equiv s(\theta)U$, where $s(\theta)$ is a scalar random variable that may depend on parameter θ , and U is unit-norm vector random variable that does not depend on θ . Then, it is easy to see that algorithm (1.6) is equivalent to

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} - \gamma_n \lambda_n s(\theta_{n-1}^{\text{im}}) U, \quad (1.7)$$

where the scalar λ_n satisfies

$$\lambda_n s(\theta_{n-1}^{\text{im}}) = s(\theta_{n-1}^{\text{im}} - \gamma_n \lambda_n s(\theta_{n-1}^{\text{im}}) U). \quad (1.8)$$

Therefore, at the n th iteration we can first draw U , then solve the one-dimensional fixed-point equation (1.8), given that the analytic form of $s(\theta)$ is known, and finally calculate θ_n^{im} through Eq. (1.7). The solution of the one-dimensional fixed-point equation is computationally efficient in a large family of statistical models when W_θ is the negated gradient of the log-likelihood; this family includes, for example, generalized linear models and convex M-estimation (Toulis and Airolidi, 2015b, Theorem 4.1). Second, consider the more general case $W_\theta = s(\theta)U_\theta$, where U_θ is a random vector that can depend on parameter θ . The problem is now is that we cannot sample $U_{\theta_n^{\text{im}}}$ because it depends on the next iterate. However, we can simply draw $U_{\theta_{n-1}^{\text{im}}}$ instead, which reduces to the previous case where $W_\theta = s(\theta)U$, and therefore algorithm (1.7) is again applicable.

1.6 Application in parameter estimation

A key application of stochastic approximation is in iterative estimation of the parameters of a statistical model (Nevelson and Khasminski, 1973, Chapter 8); (Ljung *et al.*, 1992, Chapter 10); (Borkar, 2008, Section 10.2). In particular, consider a stream of i.i.d. data points (X_n, Y_n) , $n = 1, 2, \dots$, where the outcome $Y_n \in \mathbb{R}^d$ is distributed conditional on covariates $X_n \in \mathbb{R}^p$ according to known density $f(Y_n; X_n, \theta_*)$, but unknown model parameters $\theta_* \in \mathbb{R}^p$. We will

consider two cases, one where the likelihood is fully known and can be calculated easily at each data point (X_n, Y_n) , and one where the likelihood is not fully known (e.g., it is known up to normalizing constant) and there is a finite data set instead of a stream. We will also discuss a third case of likelihood-free estimation.

1.6.1 Likelihood-based estimation

Consider the random variable $W_\theta \stackrel{\text{def}}{=} -\nabla \log f(Y_n; X_n, \theta)$, where the regression function $h(\theta) \stackrel{\text{def}}{=} \mathbb{E}(W_\theta)$ is still unknown. In many statistical models the log-likelihood $\log f$ is concave (Bickel and Doksum, 2015), and thus θ_* is the unique point for which $h(\theta_*) \stackrel{\text{def}}{=} -\mathbb{E}(\nabla \log f(Y_n; X_n, \theta_*)) = 0$. Therefore the standard stochastic approximation procedure (1.1) can be written as

$$\theta_n^{\text{rm}} = \theta_{n-1}^{\text{rm}} + \gamma_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{rm}}). \quad (1.9)$$

Stochastic approximation theory implies $\theta_n^{\text{rm}} \rightarrow \theta_*$ and therefore θ_n^{rm} is a consistent estimator of θ_* . Procedure (1.9) is known as *stochastic gradient descent* (SGD) in optimization and signal processing (Coraluppi and Young, 1969), and has been fundamental in modern machine learning with large data sets (Zhang, 2004; Bottou, 2010; Toulis and Airoidi, 2015a).

An implicit stochastic approximation version of procedure (1.9) can be implemented through algorithm (1.6), resulting in the iteration

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n \nabla \log f(Y_n; X_n, \theta_n^{\text{im}}). \quad (1.10)$$

Procedure (1.10) is known as the *incremental proximal method* in optimization (Bertsekas, 2011), or as *implicit stochastic gradient descent* in statistical estimation (Toulis et al., 2014), and has shown superior performance to standard stochastic gradient descent, both in theory and applications (Toulis et al., 2014; Toulis and Airoidi, 2015b). In particular, in accordance to the theoretical properties of their stochastic approximation counterparts, implicit SGD has identical asymptotic efficiency and convergence rate as standard SGD, but it is significantly more stable. The stochastic proximal gradient algorithm (Singer and Duchi, 2009; Rosasco

[et al., 2014](#)) is related to implicit SGD. However, in contrast to implicit SGD, the proximal gradient algorithm first makes a standard update (forward step), and then an implicit update (backward step), which may increase convergence speed, but may also introduce instability due to the forward step.

Example. Let $\theta_\star \in \mathbb{R}$ be the true parameter of a normal model with i.i.d. observations $Y_n | X_n \sim \mathcal{N}(X_n \theta_\star, 1)$, $X_n, Y_n \in \mathbb{R}$. Thus, $\log f(Y_n; X_n, \theta) = -\frac{1}{2}(Y_n - X_n \theta)^2$, and $\nabla \log f(Y_n; X_n, \theta) = (Y_n - X_n \theta) X_n$. Assume $\gamma_n = \gamma_1/n$ as the learning rate. Then, the SGD procedure (1.9) is given by

$$\theta_n^{\text{rm}} = (1 - \gamma_n X_n^2) \theta_{n-1}^{\text{rm}} + \gamma_n Y_n X_n. \quad (1.11)$$

Procedure (1.11) is known as the least mean squares filter (LMS) in signal processing, or as the Widrow-Hoff algorithm ([Widrow and Hoff, 1960](#)). The implicit SGD procedure for this problem, using update (1.10), is given by

$$\theta_n^{\text{im}} = \frac{1}{1 + \gamma_n X_n^2} \theta_{n-1}^{\text{im}} + \frac{\gamma_n}{1 + \gamma_n X_n^2} Y_n X_n. \quad (1.12)$$

Procedure (1.12) is also known as the normalized least mean squares filter (NLMS) in signal processing ([Nagumo and Noda, 1967](#)). From Eq. (1.11) we see that it is crucial for standard SGD to have a well-specified learning rate parameter γ_1 . For instance, assume fixed $X_n^2 = x^2$ for simplicity, then if $\gamma_1 x^2 \gg 1$ the iterate θ_n^{rm} will diverge to a value $O(2\gamma_1 x^2 / \sqrt{\gamma_1 x^2})$. In contrast, a very large γ_1 will not cause divergence in implicit SGD, but it will simply put more weight on the n th observation $Y_n X_n$ than the previous iterate θ_{n-1}^{im} . Moreover, from a statistical perspective, implicit SGD specifies the correct averaging by weighing the estimate and observation according to the inverse of information $(1 + \gamma_n X_n^2)$.

1.6.2 Estimation with likelihood known up to normalizing constant

We now consider the case where the likelihood is known up to a normalizing constant, which arises frequently in practice ([Gelman and Meng, 1998](#)). In such situations, a large family of estimation methods relies on being able to sample from the underlying model,

e.g., through Metropolis-Hastings, which does not require knowledge of the normalization constant. For example, the method of Markov chain Monte Carlo maximum-likelihood estimation (Geyer, 1991) uses simulations to estimate ratios of normalization constants, which appear in the maximization of the log-likelihood over a finite data set.

The same idea underlying simulation-based methods can be applied in iterative estimation with stochastic approximation. Consider a finite data set with N data points, and a complete sufficient statistic S_n that can be calculated on the n th data point. Let $\hat{S}(\theta; k)$ denote the averaged value of the statistic over k independent data points that are simulated conditional on θ . Then, the iteration

$$\begin{aligned} n &\sim \text{Uniform}\{1, 2, \dots, N\} \\ \theta_n &= \theta_{n-1} + \gamma_n (S_n - \hat{S}(\theta_{n-1}; k)), \end{aligned} \quad (1.13)$$

can converge to θ_* under typical conditions of stochastic approximation. The learning rates γ_n can be chosen as $\gamma_n = \gamma_1/n$, as it is common, and adaptive schemes are possible (Lai and Robbins, 1979). The constant k affects the variance of the stochastic part in update (1.13): higher k lowers the variance, but also increases the computational burden because more simulations are needed. Resolution of such statistical/computational trade-offs typically require problem-specific considerations.

The implicit counterpart of procedure (1.13) uses an update as follows,

$$\theta_n = \theta_{n-1} + \gamma_n (S_n - \lambda_n \hat{S}(\theta_{n-1}; k)), \quad (1.14)$$

where the scalar λ_n satisfies

$$\lambda_n \hat{S}(\theta_{n-1}; k) = \hat{S}(\theta_n; k) \stackrel{\text{def}}{=} \hat{S}(\theta_{n-1} + \gamma_n (S_n - \lambda_n \hat{S}(\theta_{n-1}; k)); k). \quad (1.15)$$

Solution of Eq. (1.15) is particularly challenging because one needs to find the scalar λ_n such that the update (1.14) leads to a new iterate θ_n with simulated statistics satisfying (1.15). One idea is to consider only $\lambda_n \in [0, 1]$, which is true under strong convexity; for example, in the normal linear model of Eq. (1.11), $\lambda_n = 1/(1 + \gamma_n X_n^2)$. More generally, this idea is

reasonable because λ_n usually acts as a shrinkage factor (Toulis and Airolidi, 2015b, Section 5). Then, we can repeat simulations on a grid of $m + 1$ values $[0, 1]_m \triangleq \{0, 1/m, 2/m, \dots, 1\}$, and set λ_n as follows,

$$\lambda_n = \arg \min_{\lambda \in [0, 1]_m} \|\lambda \hat{S}(\theta_{n-1}; k) - \hat{S}(\theta_n(\lambda); k)\|^2, \quad (1.16)$$

where $\theta_n(\lambda) \triangleq \theta_{n-1} + \gamma_n(S_n - \lambda \hat{S}(\theta_{n-1}; k))$. Procedure (1.14) uses $m \cdot k$ simulations per iteration from the underlying model, and is thus more computationally demanding than procedure (1.13), which uses only k simulations per iteration. Improvements in computation are possible if certain simulated moments are reused, similar to the idea of Bartz *et al.* (2009). This is reasonable because the grid of values $[0, 1]_m$ will increasingly yield similar simulated moments as the iteration counter n increases and $\gamma_n \rightarrow 0$.

Example. Suppose we observe N graphs $G_n, n = 1, 2, \dots, N$, and we want to fit an exponential random graph model (ERGM) (Frank and Strauss, 1986; Pattison and Wasserman, 1999; Robins *et al.*, 2007) with density of a graph G defined as

$$f(G; \theta) = \exp(\theta^\top X_G) / c(\theta), \quad (1.17)$$

where $X_G \in \mathbb{R}^p$ are properties of G , such as number of triangles, number of edges, etc., and $c(\theta)$ is an appropriate normalizing constant. Typically, $c(\theta)$ is hard to compute, and thus remains unknown.

In this case, it is still possible to use the stochastic approximation procedure (1.18), where we define $S_n \triangleq X_{G_n}$, and $\hat{S}(\theta_{n-1}; k) \triangleq (1/k) \sum_{i=1}^k X_{\tilde{G}_i}$, where \tilde{G}_i is an independent sample from the ERGM model (1.17) with parameter value fixed at θ_{n-1} . Obtaining such samples is computationally straightforward (Snijders, 2002; Bartz *et al.*, 2009). The implicit procedure (1.14) implemented through (1.16) has potential stability advantages, but a more efficient implementation remains an interesting open problem ahead.

1.6.3 Likelihood-free estimation

An important class of estimation methods in statistics does not rely on likelihood, such as the method of moments, or non-parametric methods. Typically, in such cases, there is a statistic S_n that can be calculated at every n th data point, and the regression function $T(\theta) \triangleq \mathbb{E}(S_n | \theta_*) = \theta$ is known. Then, the procedure

$$\theta_n = \theta_{n-1} + \gamma_n (S_n - T(\theta_{n-1})), \quad (1.18)$$

can converge to θ_* under typical stochastic approximation conditions (e.g., convexity of vector-valued function T). If the system of equations $\mathbb{E}(S_n) = T(\theta_*)$ is over-specified, convexification or regularization could be applied, similar, for example, to the way the generalized method of moments (Hall, 2005) resolves over-specification in moment conditions. Interestingly, the form of procedure (1.18), where the update depends on a discrepancy $S_n - T(\theta_{n-1})$ between an observed statistic and its expected value, also appears in likelihood-based estimation with SGD, e.g., in exponential family models —cf. Eq. (1.11). The implicit counterpart of procedure (1.18) has straightforward implementations through the algorithms of Section 1.5, which we illustrate in the following example.

Example. In their seminal paper, Robbins and Monro (1951) described an application of procedure (1.1) in iterative quantile estimation. In particular, consider a random variable Z with a cumulative distribution function F . An experimenter wants to find the point θ_* such that $F(\theta_*) = \alpha$, for a fixed $\alpha \in (0, 1)$. The experimenter cannot draw samples of Z , but has access to the random variable $W_\theta = \mathbb{I}\{Z \leq \theta\} - \alpha$, for any value of θ . Robbins and Monro (1951) argued that the procedure

$$\theta_n = \theta_{n-1} - \gamma_n W_{\theta_{n-1}}, \quad (1.19)$$

will converge to θ_* . Indeed, $\mathbb{E}(W_{\theta_*}) = \mathbb{E}(\mathbb{I}\{Z \leq \theta_*\}) - \alpha = F(\theta_*) - \alpha \stackrel{\text{def}}{=} 0$. Convexity or concavity of F is not required because F is nondecreasing and $F'(\theta_*) > 0$, which are the original conditions for convergence established by Robbins and Monro (1951).

Quantile estimation through implicit stochastic approximation can be accomplished

through algorithm (1.5), with the n th iteration defined as

$$\begin{aligned} x_k &= x_{k-1} - a_k(\gamma_n W_{x_{k-1}} + x_{k-1} - \theta_{n-1}^{\text{im}}), k = 1, 2, \dots, K, \\ \theta_n^{\text{im}} &= x_K, \end{aligned} \tag{1.20}$$

where $x_0 = \theta_{n-1}^{\text{im}}$, $a_k \propto 1/k$, and K is a small constant; we give concrete values to these constants in a numerical simulation at the end of this example.

The benefits of the implicit procedure over the classical one can be emphatically illustrated through the following numerical exercise. Suppose Z is a standard normal random variable, and $\alpha = 0.999$. In this case, $F(\theta_*) = 0.999$, and thus $\theta_* \approx 3.09$ is in a region of the parameter space where F is very flat; in fact, the region $[2.87, +\infty]$ is within ± 0.1 from the value 0.999 of $F(\theta_*)$. It is well-known that the convergence of the classical Robbins-Monro procedure depends on $F'(\theta_*)$: if the product $\gamma_1 F'(\theta_*)$ is small then convergence will be slow (Toulis and Airoidi, 2015a, Section 2.1). In this example, because $F'(\theta_*) \approx 0.0034$ is very small in the flat region, the learning rate parameter γ_1 needs to compensate by being large, otherwise convergence to θ_* will be very slow. In fact, standard theory suggests that it is optimal to select $\gamma_1 = 1/F'(\theta_*) \approx 294$. However, if γ_1 is large the classical procedure (1.19) will overshoot early on to a point $\theta_*^+ \gg \theta_*$ such that $F(\theta_*^+) \approx 1$. Return from that point will be very slow because the values of $W_{\theta_{n-1}}$ will be close to zero, since $W_{\theta_{n-1}} \approx 1 - \alpha = 0.001$. For example, suppose $\gamma_1 = 294$ and $\theta_*^+ = 10$, then to return back to $\theta_* \approx 3.09$ we will need m steps for which

$$\begin{aligned} \sum_{n=n_0+1}^{n_0+m} \gamma_n W_{\theta_{n-1}} &\approx (10 - 3.09) \\ \gamma_1 \sum_{n=n_0+1}^{n_0+m} 0.001/n &\approx 6.91 \\ \log m &\geq 6.91 \cdot 10^3 / \gamma_1 \approx 23. \end{aligned} \tag{1.21}$$

The classical procedure will therefore require at least e^{23} steps to return back after this simple overshoot from $\theta_* = 3.09$ to $\theta_*^+ = 10$. Smaller values of the learning rate, e.g., $\gamma_1 = 20$, will exacerbate this problem.

In stark contrast, the implicit procedure (1.20) can neither overshoot nor undershoot. The update (1.20) satisfies approximately $\theta_n^{\text{im}} + \gamma_n(F(\theta_n^{\text{im}}) - \alpha) \approx \theta_{n-1}^{\text{im}}$, which can be written as $\theta_n^{\text{im}} + \gamma_n F(\theta_n^{\text{im}}) \approx \theta_{n-1}^{\text{im}} + \gamma_n F(\theta_*)$. Because F is nondecreasing it follows that if $\theta_{n-1}^{\text{im}} < \theta_*$, then $\theta_n^{\text{im}} > \theta_{n-1}^{\text{im}}$. Similarly, if $\theta_{n-1}^{\text{im}} > \theta_*$ then $\theta_n^{\text{im}} < \theta_{n-1}^{\text{im}}$. Furthermore, a large value of γ_n will in fact push θ_n^{im} more towards θ_* . The implicit equation therefore bestows a remarkable stability property to the underlying stochastic approximation procedure.

To validate our theory, we ran 20 replications of both the Robbins-Monro procedure (1.19) and the implicit procedure (1.20) for quantile estimation. Both procedures had a learning rate $\gamma_n = \gamma_1/n$, where we picked multiple values of γ_1 to test the stability of the procedures. For the implicit procedure we also set $K = 50$ and $a_k = 10/k$. Each procedure was initialized at $\theta_0 = -10$ and was run for 5,000 iterations, at each replication. The final parameter estimate θ_n , $n = 5,000$, was stored for each procedure and replication. The results are shown in Figure 1.1.

We observe that, as the learning rate constant γ_1 increases, the classical RM procedure produces iterates that overshoot and get stuck. This is consistent with the theoretical analysis of this section, where we showed that the classical procedure cannot converge after overshooting because the true parameter value θ_* is in a region where the objective function F is flat. In stark contrast, the implicit procedure remains significantly robust, with iterates around the true parameter value across different learning rates. Importantly, the implicit procedure was only approximately implemented with algorithm (1.5), where the inner stochastic approximation for x_k was executed for only $K = 40$ iterations within each n th iteration of the main implicit procedure (1.20).

1.7 Conclusion

The need to estimate parameters of a statistical model from massive amounts of data has reinvigorated interest in approximate inference procedures. While most gradient-based procedures of current interest can be seen as special case of a gradient-free stochastic approximation method developed by [Robbins and Monro \(1951\)](#), there is no general gradient-




Figure 1.1: Boxplots of 20 replications of the Robbins-Monro (RM) (1.19) and the implicit procedure (1.20). Each replication yields one estimate θ_n , where $n = 5,000$. The true parameter value θ_* is depicted as a red horizontal line at $y = \Phi(0.999) \approx 3.09$, and both procedures start from $\theta_0 = -10$. When the learning rate constant is small (e.g., $\gamma_1 = 0.1$) both procedures are very slow and their iterates are far from the true value. As γ_1 increases, the classical RM procedure produces iterates that overshoot and get stuck (flat boxplots). In contrast, the implicit procedure remains robust, with final iterates around θ_* .

free method that can accommodate procedures with updates defined through implicit equations. Here, we conceptualize a gradient-free implicit stochastic approximation method, and develop asymptotic and non-asymptotic theory for it. This new approximation method provides the theoretical basis for gradient-based procedures that rely on proximal operators (implicit updates), and opens the door to new iterative estimation procedures that do not require access to a gradient or a fully-known likelihood function.

Chapter 2

Implicit Stochastic Gradient Descent

2.1 Overview

Stochastic optimization procedures, such as stochastic gradient descent, have gained popularity for parameter estimation from large data sets. However, standard stochastic optimization procedures cannot effectively combine numerical stability with statistical and computational efficiency. Here, we introduce an *implicit* stochastic gradient descent procedure, the iterates of which are implicitly defined. Intuitively, implicit iterates *shrink* the standard iterates. The amount of shrinkage depends on the observed Fisher information matrix, which does not need to be explicitly computed in practice, thus increasing stability without increasing the computational burden. When combined with averaging, the proposed procedure achieves statistical efficiency as well. We derive non-asymptotic bounds and characterize the asymptotic distribution of implicit procedures. Our analysis also reveals the asymptotic variance of a number of existing procedures. We demonstrate implicit stochastic gradient descent by further developing theory for generalized linear models, Cox proportional hazards, and M-estimation problems, and by carrying out extensive experiments. Our results suggest that the implicit stochastic gradient descent procedure is poised to become the workhorse of estimation with large data sets.

2.2 Introduction

Parameter estimation by optimization of an objective function, such as maximum likelihood and maximum a-posteriori, is a fundamental idea in statistics and machine learning (Fisher, 1922; Lehmann and Casella, 1998; Hastie *et al.*, 2011). However, widely used optimization-based estimation procedures, such as Fisher scoring, the EM algorithm or iteratively reweighted least squares (Fisher, 1925; Dempster *et al.*, 1977; Green, 1984), do not scale to modern data sets with millions of data points and hundreds or thousands of parameters (National Research Council, 2013). In this paper, we introduce and further develop iterative estimation procedures based on *stochastic gradient descent* (SGD) optimization, rooted in the early statistics literature on stochastic approximations (Robbins and Monro, 1951), which are computationally efficient and lead to estimates with good statistical properties.

We consider the problem of estimating the true vector of parameters $\theta_\star \in \mathbb{R}^p$ of a model that is assumed to produce i.i.d. data points (X_i, Y_i) , for $i = 1, 2, \dots, N$. Conditional on covariates $X_i \in \mathbb{R}^p$, the outcome $Y_i \in \mathbb{R}^d$ is distributed according to known density $f(Y_i; X_i, \theta_\star)$. The *expected Fisher information matrix* is $\mathcal{I}(\theta_\star) \stackrel{\text{def}}{=} \mathbb{E} (\nabla \log f(Y; X, \theta_\star) \nabla \log f(Y; X, \theta_\star)^\top)$, where (X, Y) denotes a data point. Properties of the Fisher information matrix are very important for the stability and efficiency of SGD procedures, as we show in Sections 2.4 and 2.4.5. Given a sample of N i.i.d. data points the estimation problem usually reduces to optimization, for example, finding the maximum likelihood estimate (MLE), defined as $\theta_N^{\text{mle}} = \arg \max_{\theta} \sum_{i=1}^N \log f(Y_i; X_i, \theta)$.¹

Traditional estimation procedures have a running time complexity that ranges between $O(Np^{1+\epsilon})$ and $O(Np^{2+\epsilon})$ in best cases and worst cases, respectively. Newton-Raphson, for instance, converges linearly to MLE (Kelley, 1999), however, matrix inversions and likelihood computations on the entire data set yield an algorithm with $O(Np^{2+\epsilon})$ complexity, per iteration, which makes it unsuitable for large data sets. Fisher scoring, a variant of Newton-

¹If a prior $\pi(\theta)$, also known as regularization, is assumed on the model parameters θ then the estimation problem reduces to finding the maximum a-posteriori estimate (MAP), defined as $\theta_N^{\text{map}} = \arg \max_{\theta} \sum_{i=1}^N \log f(Y_i; X_i, \theta) + \log \pi(\theta)$.

Raphson, is generally more stable but has similar computational properties (Lange, 2010). Quasi-Newton (QN) procedures are a powerful alternative and routinely used in practice² because they have $O(Np^2)$ complexity per-iteration, or $O(Np^{1+\epsilon})$ in certain favorable cases.(Hennig and Kiefel, 2013). Other general estimation algorithms, such as EM (Dempster *et al.*, 1977) or iteratively reweighted least squares (Green, 1984), involve computations (e.g. matrix inversions or maximizations between iterations) that are significantly more expensive than QN procedures.

In contrast, estimation with large data sets requires a running time complexity that is roughly $O(Np^{1-\epsilon})$, i.e., linear in data size N but sublinear in parameter dimension p . The first requirement on N is rather unavoidable because all data points carry information by the i.i.d. assumption, and thus all need to be considered by any iterative estimation procedure. Thus, sublinearity in p is crucial. Such computational requirements have recently sparked interest in stochastic optimization procedures, especially those only working with *first-order* information, i.e., gradients.

Stochastic optimization procedures of this kind are rooted in stochastic approximation methods (Robbins and Monro, 1951), where, interestingly, the estimation problem is formulated not as an optimization problem, but as a *characteristic equation*. In particular, if N is finite the characteristic equation is

$$\mathbb{E} \left(\nabla \log f(Y; X, \theta_N^{\text{mle}}) \right) = 0, \tag{2.1}$$

where the expectation is over the *empirical* distribution of (X, Y) on the finite data set. If N is infinite – a situation known as “stream of data” – the characteristic equation is

$$\mathbb{E} (\nabla \log f(Y; X, \theta_*) | X) = 0, \tag{2.2}$$

where the expectation is over the true conditional distribution of outcome Y given covariate X . Given a characteristic equation, stochastic approximations are procedures that iteratively approximate its solution (Ljung *et al.*, 1992; Benveniste *et al.*, 1990), i.e., they approximate

²For example, most implemented algorithms in R’s `optim()` function are Quasi-Newton.

θ_N^{mle} in Eq. (2.1) and θ_* in Eq. (2.2).³

A popular stochastic approximation procedure for estimation with large data sets is *stochastic gradient descent* (SGD), defined for $n = 1, 2, \dots$, as

$$\theta_n^{\text{sgd}} = \theta_{n-1}^{\text{sgd}} + \gamma_n C_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{sgd}}), \quad (2.3)$$

where $\gamma_n > 0$ is the *learning rate sequence*, typically defined as $\gamma_n = \gamma_1 n^{-\gamma}$, $\gamma_1 > 0$ is the *learning rate parameter*, $\gamma \in (1/2, 1]$, and C_n are $p \times p$ positive-definite matrices, also known as condition matrices.

When N is finite the data point (X_n, Y_n) in SGD (2.3) is a random sample with replacement from the finite data set. When N is infinite the data point (X_n, Y_n) is simply the n th data point in the data stream. For the rest of this paper we assume infinite N because it is the natural setting for stochastic approximations. This assumption is without loss of generality because all theoretical results for the infinite N case can be applied to finite N , where instead of estimating θ_* we estimate θ_N^{mle} (or MAP if there is regularization).

From a computational perspective, SGD (2.3) is appealing because it avoids expensive matrix inversions, as in Newton-Raphson, and the log-likelihood is evaluated at a single data point (X_n, Y_n) and not on the entire data set. From a theoretical perspective, SGD (2.3) is essentially a stochastic approximation procedure, and thus converges, under suitable conditions, to $\theta_\infty^{\text{sgd}}$ where $\mathbb{E} \left(\log f(Y; X, \theta_\infty^{\text{sgd}}) \mid X \right) = 0$. This condition satisfies both Eqs. (2.1) and (2.2), implying that SGD can be used on both finite and infinite data sets.

The main contribution of this paper is to provide a formal analysis and the statistical intuition behind *implicit* stochastic gradient descent, or *implicit SGD* for short, defined for $n = 1, 2, \dots$, as follows:

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n C_n \nabla \log f(Y_n; X_n, \theta_n^{\text{im}}), \quad (2.4)$$

where γ_n, C_n are defined as in standard SGD (2.3). To distinguish the two procedures, we

³The characteristic equations can have a set of multiple solutions. In this case, stochastic approximations converge to a point in that set, but the exact point to which they converge depends on initial conditions (Borkar, 2008).

will refer to SGD (2.3) as SGD with *explicit updates*, or *explicit SGD* for short, because the next iterate θ_n^{sgd} can be immediately computed given $\theta_{n-1}^{\text{sgd}}$ and the data point (X_n, Y_n) . In contrast, the update in Eq. (2.4) is *implicit* because the next iterate θ_n^{im} appears on both sides of the equation.

2.3 Overview of main results

Our first set of results relies on a non-asymptotic analysis of iterates θ_n^{im} of implicit SGD (2.4). In particular, we derive upper bounds for the mean-squared errors $\mathbb{E}(\|\theta_n^{\text{im}} - \theta_\star\|^2)$. Compared to non-asymptotic analyses of explicit SGD (Benveniste *et al.*, 1990; Ljung *et al.*, 1992; Moulines and Bach, 2011), implicit SGD is particularly robust to misspecification of the learning rate with respect to the problem characteristics (e.g., convexity). In contrast, in explicit SGD the initial conditions can be amplified arbitrarily if learning rates are misspecified.

Our second set of results relies on an asymptotic analysis of θ_n^{im} . In particular, we show that the asymptotic variance of θ_n^{im} is identical to the variance of θ_n^{sgd} . Both procedures lose statistical information compared, for example, to MLE θ_n^{mle} computed over n data points. However, this information loss can be quantified exactly (Theorem 6), which can be leveraged to design optimal learning rates (Eq. 2.9). Surprisingly, this information loss can be avoided by simply averaging the iterates θ_n^{im} (Theorem 7). This result matches similar results on averaging of explicit SGD procedures, first given by Ruppert (1988); Bather (1989); Polyak and Juditsky (1992a). Additionally, under typical Lindeberg conditions, we show that θ_n^{im} is asymptotically normal with known asymptotic variance, which can be leveraged for producing standard errors for implicit SGD estimates.

The combined results from the non-asymptotic and asymptotic analyses show that explicit SGD procedures cannot easily combine numerical stability with statistical efficiency (Section 2.4.5). In rough terms, for stability we need $\lambda_{\max}\gamma_1 < 1$, where γ_1 is the learning rate parameter, and λ_{\max} is the maximum eigenvalue of the Fisher information $\mathcal{I}(\theta_\star)$. For statistical efficiency we need $\lambda_{\min}\gamma_1 > 1/2$, where λ_{\min} is the minimum eigenvalue of

$\mathcal{I}(\theta_*)$. Thus, we need $\lambda_{\max} < 2\lambda_{\min}$ to achieve both stability and efficiency. This condition depends on the *condition number* of $\mathcal{I}(\theta_*)$ and is hard to satisfy in large data sets with high-dimensional parameters. In stark contrast, the stability condition is eliminated in implicit SGD because, effectively, any learning rate parameter γ_1 can yield a stable procedure (Theorem 5). With stability issues fixed, one has more freedom to select larger learning rates, with or without averaging, in order to speed up convergence and increase statistical efficiency, which sums up the benefits of implicit over explicit SGD.

Our third set of results relies on practical applications of implicit SGD on a wide family of statistical models. In particular, we devise a new algorithm for fast calculation of iterates θ_n^{im} (Algorithm 1) by solving efficiently the p -dimensional fixed-point equation (2.4) in the definition of implicit SGD. Algorithms 2, 4, and 5 are variants that specialize Algorithm 1 to generalized linear models, Cox proportional hazards, and M-estimation problems, respectively. Outside of such models, we show how implicit updates can be approximately implemented (Sections 2.6.2 & 2.7), and show empirical evidence that this approximation works well on Cox proportional hazards models (Section 2.6.2), and a big SVM model (Section B.7.4).

2.3.1 Illustrative example

We motivate our main results through the simple estimation problem first presented in Chapter 1. Let $\theta_* \in \mathbb{R}$ be the true parameter of a normal model with i.i.d. observations $Y_i|X_i \sim \mathcal{N}(X_i\theta_*, \sigma^2)$, where the variance σ^2 is assumed known for simplicity. The log-likelihood is $\log f(Y_i; X_i, \theta) = -\frac{1}{2\sigma^2}(Y_i - X_i\theta)^2$, and the score function (i.e., gradient of log-likelihood) is given by $\nabla \log f(Y_i; X_i, \theta) = \frac{1}{\sigma^2}(Y_i - X_i\theta)X_i$. Let X_i be distributed according to some unknown distribution with bounded second. Assume $\gamma_n = \gamma_1/n$, for some $\gamma_1 > 0$ as the learning rate, and an initial condition θ_0 . Then, the explicit SGD procedure (2.3) is

$$\begin{aligned}\theta_n^{\text{sgd}} &= \theta_{n-1}^{\text{sgd}} + \gamma_n(Y_n - \theta_{n-1}^{\text{sgd}}X_n)X_n \Rightarrow \\ \theta_n^{\text{sgd}} &= (1 - \gamma_n X_n^2)\theta_{n-1}^{\text{sgd}} + \gamma_n Y_n X_n.\end{aligned}\tag{2.5}$$

Procedure (2.5) is the least mean squares filter (LMS) in signal processing, also known as the Widrow-Hoff algorithm (Widrow and Hoff, 1960). The implicit SGD procedure can be derived in closed form in this problem using update (2.4) as

$$\begin{aligned}\theta_n^{\text{im}} &= \theta_{n-1}^{\text{im}} + \gamma_n(Y_n - X_n\theta_n^{\text{im}})X_n \Rightarrow \\ \theta_n^{\text{im}} &= \frac{1}{1 + \gamma_n X_n^2} \theta_{n-1}^{\text{im}} + \frac{\gamma_n}{1 + \gamma_n X_n^2} Y_n X_n.\end{aligned}\quad (2.6)$$

Procedure (2.6) is known as the normalized least mean squares filter (NLMS) in signal processing (Nagumo and Noda, 1967).

From Eq. (2.5) we see that it is crucial for explicit SGD to have a well-specified learning rate parameter γ_1 . For instance, if $\gamma_1 X_n^2 \gg 1$ then θ_n^{sgd} will diverge to a value at the order of $2\gamma_1 / \sqrt{\gamma_1}$, before converging to the true value (see Section 2.4.5, Lemma 1). In contrast, implicit SGD is more stable to misspecification of the learning rate parameter γ_1 . For example, a very large γ_1 will not cause divergence as in explicit SGD, but it will simply put more weight on the n th observation $Y_n X_n$ than the previous iterate θ_{n-1}^{im} . Assuming for simplicity $\theta_{n-1}^{\text{sgd}} = \theta_{n-1}^{\text{im}} = 0$, it also holds $\theta_n^{\text{im}} = \frac{1}{1 + \gamma_n X_n^2} \theta_n^{\text{sgd}}$, showing that implicit SGD iterates are shrunk versions of explicit ones (see also Section 2.7).

Let $v^2 \stackrel{\text{def}}{=} \mathbb{E}(X^2)$, then according to Theorem 6 the asymptotic variance of the explicit iterate θ_n^{sgd} (and the implicit θ_n^{im}) satisfies $n\text{Var}(\theta_n^{\text{sgd}}) \rightarrow \gamma_1^2 \sigma^2 v^2 / (2\gamma_1 v^2 - 1)$ if $2\gamma_1 v^2 - 1 > 0$. Since $\gamma_1^2 / (2\gamma_1 v^2 - 1) \geq 1/v^2$, it is best to set $\gamma_1 = 1/v^2$. In this case $n\text{Var}(\theta_n^{\text{sgd}}) \rightarrow \sigma^2/v^2$. Explicit SGD can thus be optimal by setting $\gamma_n = (\sum_{i=1}^n X_i^2)^{-1}$, which implies that $\theta_n^{\text{sgd}} = \sum_{i=1}^n Y_i X_i / \sum_{i=1}^n X_i^2$, i.e., it is the classical OLS estimator. The implicit SGD estimator θ_n^{im} (2.6) inherits the efficiency properties of θ_n^{sgd} , with the added benefit of being stable over a wide range of learning rates γ_1 . Overall, implicit SGD is a superior form of SGD.

2.3.2 Experimental evaluation

Our experiments are split into two sets. First, in Section 2.6.1 we perform experiments to validate our theoretical findings. In particular we test our results on asymptotic variance of first-order or second-order SGD procedures (Section 2.4.2, Theorem 6), and our results on

asymptotic normality (Theorem 8).

Second, in Section 2.6.2, we carry out an extensive set of experiments on real and simulated data to compare the performance of implicit SGD against optimal deterministic optimization procedures. In particular, we compare implicit SGD against deterministic procedures on statistical models presented in Sections 2.5.2-2.5.4. Subsequently, we focus on generalized linear models and compare implicit SGD against Fisher scoring (using the R `glm()` function) on various generalized linear models, and against a popular alternative to scale Fisher scoring to large data sets (using the R `biglm` package). We also compare implicit SGD against the elastic net by Friedman *et al.* (2010) for sparse regularized estimation (using the R `glmnet` package). Finally, we work on a variant of generalized linear models, where we re-analyze a large data set from the National Morbidity-Mortality Air Pollution study (Samet *et al.*, 2000; Dominici *et al.*, 2002), and compare against recently published methods that were specifically designed for that task.

Additional experiments are presented in the Appendix. In particular, we perform additional experiments to compare implicit SGD with known adaptive or proximal stochastic optimization methods, including AdaGrad (Duchi *et al.*, 2011a), Prox-SVRG (Xiao and Zhang, 2014), and Prox-SAG (Schmidt *et al.*, 2013), on popular machine learning tasks. Overall, our results provide strong evidence that the family of implicit SGD procedures provides a superior form of stochastic gradient descent, with comparable statistical efficiency to explicit SGD, but with much improved stability. Combined with results on asymptotic variance and normality, implicit SGD emerges as a principled estimation method that can become the workhorse of statistical estimation with large data sets in statistical practice.

2.3.3 Related work

Historically, the duo of explicit-implicit updates originate from the numerical methods introduced by Euler (ca. 1770) for approximating solutions of ordinary differential equations (Hoffman and Frankel, 2001). The explicit SGD procedure was first proposed by Sakrison (1965) as a recursive statistical estimation method and it is theoretically based on the

stochastic approximation method of [Robbins and Monro \(1951\)](#). Statistical estimation with explicit SGD is a straightforward generalization of Sakrison’s method and has recently attracted attention in the machine learning community as a fast learning method for large-scale problems ([Zhang, 2004](#); [Bottou, 2010](#); [Toulis and Airoldi, 2015a](#)). Applications of explicit SGD procedures in massive data problems can be found in many diverse areas such as large-scale machine learning ([Zhang, 2004](#); [Bottou, 2010](#)), online EM algorithm ([Cappé and Moulines, 2009](#); [Balakrishnan et al., 2014](#)), image analysis ([Lin et al., 2011](#)), deep learning ([Dean et al., 2012](#); [Erhan et al., 2010](#)) and MCMC sampling ([Welling and Teh, 2011](#)).

The implicit SGD procedure is less known and not well-understood. In optimization, implicit methods have recently attracted attention under the guise of proximal methods, or mirror-descent methods ([Nemirovski, 1983](#); [Beck and Teboulle, 2003](#)). In fact, the implicit SGD update (2.4) can be expressed as a proximal update as the solution of

$$\theta_n^{\text{im}} = \arg \max_{\theta} \left\{ -\frac{1}{2} \|\theta - \theta_{n-1}^{\text{im}}\|^2 + \gamma_n \log f(Y_n; X_n, \theta) \right\}. \quad (2.7)$$

From a Bayesian perspective, θ_n^{im} of the implicit procedure is the posterior mode of a model with the standard multivariate normal $\mathcal{N}(\theta_{n-1}^{\text{im}}, \gamma_n I)$ as the prior, and $\log f(Y_n; X_n, \theta)$ as the log-likelihood of θ for observation (X_n, Y_n) . Arguably, the normalized least mean squares (NLMS) filter ([Nagumo and Noda, 1967](#)), introduced in Eq. (2.6), was the first statistical model that used an implicit update as in Equation (2.4) and was shown to be consistent and robust to input noise ([Slock, 1993](#)). From an optimization perspective, update (2.7) corresponds to a stochastic version of the proximal point algorithm by [Rockafellar \(1976\)](#) which has been generalized through the idea of splitting algorithms ([Lions and Mercier, 1979](#); [Beck and Teboulle, 2009](#); [Singer and Duchi, 2009](#); [Duchi et al., 2011b](#)); see, also, the comprehensive review of proximal methods in optimization by [Parikh and Boyd \(2013\)](#). Additional intuition of implicit methods have been given by [Krakowski et al. \(2007\)](#) and [Nemirovski et al. \(2009\)](#) who have argued that proximal methods can fit better in the geometry of the parameter space. [Bertsekas \(2011\)](#) derived an asymptotic rate for an implicit procedure (2.4) on a fixed data set and compared convergence rates between randomly

sampling data (X_n, Y_n) and simply cycling through them. [Toulis et al. \(2014\)](#) derived the asymptotic variance of θ_n^{im} as estimator of θ_* in the family of generalized linear models, and provided an algorithm to efficiently compute the update (2.4). [Rosasco et al. \(2014\)](#) derived non-asymptotic bounds and convergence for a stochastic proximal gradient algorithm, which is a forward-backward procedure that first makes a stochastic explicit update (2.3), and then a deterministic implicit update.

In the online learning literature, “regret analyses” of implicit methods have been given by [Kivinen et al. \(2006\)](#) and [Kulis and Bartlett \(2010\)](#); [Schuurmans and Caelli \(2007\)](#) have further applied implicit methods on learning with kernels. Furthermore, the proximal update (2.7) is related to the importance weight updates proposed by [Karampatziakis and Langford \(2010\)](#), but the two updates have important differences ([Karampatziakis and Langford, 2010](#), Section 5).

Two recent stochastic proximal methods are `PROX-SVRG` ([Xiao and Zhang, 2014](#)) and `PROX-SAG` ([Schmidt et al., 2013](#), Section 6). Working in a finite data set, the main idea in both methods is to periodically compute an estimate of the full gradient averaged over all data points to reduce the variance of stochastic gradients. This periodic update is also controlled by additional hyperparameters, whereas `PROX-SAG` typically requires storage of the full gradient at every iteration. We compare those methods with implicit SGD in Appendix B.7.

2.4 Theory

The norm $\|\cdot\|$ denotes the L_2 norm. If a positive scalar sequence a_n is nonincreasing and $a_n \rightarrow 0$, we write $a_n \downarrow 0$. For two positive scalar sequences a_n, b_n , equation $b_n = O(a_n)$ denotes that b_n is bounded above by a_n , i.e., there exists a fixed $c > 0$ such that $b_n \leq ca_n$, for all n . Furthermore, $b_n = o(a_n)$ denotes that $b_n/a_n \rightarrow 0$. Similarly, for a sequence of vectors (or matrices) X_n , we write $X_n = O(a_n)$ if there is a fixed $c' > 0$ such that $\|X_n\| \leq c'a_n$, and $X_n = o(a_n)$ if $\|X_n\|/a_n \rightarrow 0$. For two matrices A, B $A \prec B$ denotes that $B - A$ is positive-definite. The set of eigenvalues of a matrix A is denoted by $\text{eig}(A)$; for example,

$A \succ 0$ if and only if $\lambda > 0$ for every $\lambda \in \text{eig}(A)$.

Every theoretical result of this section is stated under a combination of the following assumptions.

Assumption 7. *The explicit SGD procedure (2.3) and the implicit SGD procedure (2.4) operate under a combination of the following assumptions.*

- (a) *The sequence $\{\gamma_n\}$ is defined as $\gamma_n = \gamma_1 n^{-\gamma}$, where $\gamma_1 > 0$ is the learning parameter, and $\gamma \in (0.5, 1]$.*
- (b) *For the log-likelihood $\log f(Y; X, \theta)$ there exists function ℓ such that $\log f(Y; X, \theta) \equiv \ell(X^\top \theta; Y)$, which depends on θ only through the natural parameter $X^\top \theta$.*
- (c) *Function ℓ is concave, twice differentiable almost-surely with respect to natural parameter $X^\top \theta$ and Lipschitz continuous with constant L_0 .*
- (d) *The observed Fisher information matrix $\hat{\mathcal{I}}_n(\theta) \stackrel{\text{def}}{=} -\nabla^2 \ell(X_n^\top \theta; Y_n)$ has non-vanishing trace, i.e., there exists constant $b > 0$ such that $\text{trace}(\hat{\mathcal{I}}_n(\theta)) \geq b$ almost-surely, for all θ in the parameter space. The Fisher information matrix $\mathcal{I}(\theta_*) \stackrel{\text{def}}{=} \mathbb{E}(\hat{\mathcal{I}}_n(\theta_*))$ has minimum eigenvalue $\underline{\lambda}_f \geq b/2$ and maximum eigenvalue $\overline{\lambda}_f < \infty$. Typical regularity conditions hold (Lehmann and Casella, 1998, Theorem 5.1, p.463).*
- (e) *Every condition matrix C_n is a fixed positive-definite matrix, such that $C_n = C + \mathcal{O}(\gamma_n)$, where $C \succ 0$ and symmetric, and C commutes with $\mathcal{I}(\theta_*)$. For every C_n , $\min \text{eig}(C_n) \geq \underline{\lambda}_c > 0$, and $\max \text{eig}(C_n) \leq \overline{\lambda}_c < \infty$.*
- (f) *Let $\Xi_n \stackrel{\text{def}}{=} \mathbb{E}(\nabla \log f(Y_n; X_n, \theta_*) \nabla \log f(Y_n; X_n, \theta_*)^\top | \mathcal{F}_{n-1})$, then $\|\Xi_n - \Xi\| = \mathcal{O}(1)$ for all n , and $\|\Xi_n - \Xi\| \rightarrow 0$, for a symmetric positive-definite Ξ . Let $\sigma_{n,s}^2 \stackrel{\text{def}}{=} \mathbb{E}(\mathbb{I}_{\|\xi_n(\theta_*)\|^2 \geq s/\gamma_n} \|\xi_n(\theta_*)\|^2)$, then for all $s > 0$, $\sum_{i=1}^n \sigma_{i,s}^2 = o(n)$ if $\gamma = 1$, and $\sigma_{n,s}^2 = o(1)$ otherwise.*

Remarks. Assumption 7(a) is typical in stochastic approximation as it implies $\sum_i \gamma_i = \infty$ and $\sum_i \gamma_i^2 < \infty$, which were the original conditions given by Robbins and Monro (1951). Assumption 7(b) narrows our focus to models for which, conditional on covariate X , the likelihood depends on parameters θ through the linear combination $X^\top \theta$. This family of models is large and includes generalized linear models, Cox proportional hazards models, and M-estimation. Furthermore, in Section 2.7 we discuss an idea that allows application of implicit SGD on a wider family of models, and thus can significantly relax Assumption 7(b). Assumption 7(d) is equivalent to assuming strong convexity for the negative log-likelihood.

Such assumptions are typical for proving convergence in probability. The assumption on the observed Fisher information is less standard. Intuitively, this assumption posits that a minimum of statistical information is received from any data point, at least for certain model parameters. Making this assumption allows us to forgo boundedness assumptions on the errors of stochastic gradients that were originally used by [Robbins and Monro \(1951\)](#) and have since been standard in analysis of explicit SGD. Finally, Assumption 7(f) posits the typical Lindeberg conditions that are necessary to invoke the central limit theorem and prove asymptotic optimality; this assumption follows the conditions defined by [Fabian \(1968b\)](#) for the normality of explicit SGD procedures.

2.4.1 Non-asymptotic bounds

In this section we derive non-asymptotic upper-bounds for the errors $\|\theta_n^{\text{im}} - \theta_\star\|^2$ in expectation.

Theorem 5. *Let $\delta_n \triangleq \mathbb{E}(\|\theta_n^{\text{im}} - \theta_\star\|^2)$ and $\kappa \triangleq 1 + \gamma_1 \underline{\lambda}_c \underline{\lambda}_f \mu_0$, where $\mu_0 \in [1/(1 + \gamma_1 \underline{\lambda}_f \underline{\lambda}_c (p - 1)), 1]$. Suppose that Assumptions 7(a),(b),(c), (d) and (e) hold. Then, there exists constant n_0 such that,*

$$\delta_n \leq \frac{8L_0^2 \overline{\lambda}_c^{-2} \gamma_1 \kappa}{\underline{\lambda}_c \underline{\lambda}_f \mu_0} n^{-\gamma} + \exp(-\log \kappa \cdot \phi_\gamma(n)) [\delta_0 + \kappa^{n_0} \Gamma^2],$$

where $\Gamma^2 = 4L_0^2 \overline{\lambda}_c^{-2} \sum_i \gamma_i^2 < \infty$, and $\phi_\gamma(n) = n^{1-\gamma}$ if $\gamma < 1$, and $\phi_\gamma(n) = \log n$ if $\gamma = 1$, and n_0 is defined in [Corollary 1](#).

Not surprisingly, implicit SGD (2.4) matches the asymptotic rate of explicit SGD (2.3). In particular, the iterates θ_n^{im} have squared error $\mathcal{O}(n^{-\gamma})$, as seen in [Theorem 5](#), which is identical to the squared error of the explicit iterates θ_n^{sgd} ([Benveniste et al., 1990](#), [Theorem 22](#), p.244). Furthermore, we will show in the following section that both iterates have the same asymptotic efficiency when viewed as estimators of θ_\star .

However, the critical advantage of implicit SGD – more generally of implicit procedures – is their robustness to initial conditions and excess noise. This can be seen in [Theorem 5](#) where

the implicit procedure discounts the initial conditions $\mathbb{E}(\|\theta_0^{\text{im}} - \theta_\star\|^2)$ at an exponential rate through the term $\exp(-\log(1 + 2\gamma_1 \underline{\lambda}_c \underline{\lambda}_f) n^{1-\gamma})$, where γ_1 is the learning rate parameter, and $\underline{\lambda}_c, \underline{\lambda}_f$ are minimum eigenvalues of the condition and the Fisher information matrices, respectively. Importantly, the discounting of initial conditions happens regardless of the specification of the learning rate. In fact, large values of γ_1 can lead to faster discounting, and thus possibly to faster convergence, however at the expense of increased variance. The implicit iterates are therefore *unconditionally stable*, i.e., virtually any specification of the learning rate will lead to a stable discounting of the initial conditions.

In stark contrast, explicit SGD is known to be very sensitive to the learning rate, and can numerically diverge if the rate is misspecified. For example, [Moulines and Bach \(2011, Theorem 1\)](#) showed that there exists a term $\exp(L^2 \gamma_1^2 n^{1-2\gamma})$, where L is a Lipschitz constant for the gradient of the log-likelihood, amplifying the initial conditions $\mathbb{E}(\|\theta_0^{\text{sgd}} - \theta_\star\|^2)$ of explicit SGD, which can be catastrophic if the learning rate parameter γ_1 is misspecified.⁴ Thus, although implicit and explicit SGD have identical asymptotic performance, they are crucially different in their stability properties. This is confirmed in the experiments of [Section 2.6](#).

2.4.2 Asymptotic variance and optimal learning rates

In the previous section we showed that $\theta_n^{\text{im}} \rightarrow \theta_\star$ in quadratic mean, i.e., the implicit SGD iterates converge to the true model parameters θ_\star , similar to classical results for the explicit SGD iterates θ_n^{sgd} . Thus, θ_n^{im} and θ_n^{sgd} are consistent estimators of θ_\star . In the following theorem we show that both SGD estimators have the same asymptotic variance.

Theorem 6. *Consider SGD procedures (2.3) and (2.4), and suppose that Assumptions 7(a),(c),(d),(e) hold, where $\gamma = 1$. The asymptotic variance of the explicit SGD estimator (2.3) satisfies*

$$n\text{Var}\left(\theta_n^{\text{sgd}}\right) \rightarrow \gamma_1^2 (2\gamma_1 C\mathcal{I}(\theta_\star) - I)^{-1} C\mathcal{I}(\theta_\star)C.$$

⁴The Lipschitz conditions are different in the two works, however this does not affect our conclusion. Our result remains effectively unchanged if we assume Lipschitz continuity of the gradient $\nabla \ell$ instead of the log-likelihood ℓ , similar to [Moulines and Bach \(2011\)](#); see comment after proof of [Theorem 5](#).

The asymptotic variance of the implicit SGD estimator (2.4) satisfies

$$n\text{Var}(\theta_n^{\text{im}}) \rightarrow \gamma_1^2 (2\gamma_1 C\mathcal{I}(\theta_\star) - I)^{-1} C\mathcal{I}(\theta_\star)C.$$

Remarks. Although the implicit SGD estimator θ_n^{im} is significantly more stable than the explicit estimator θ_n^{sgd} (Theorem 5), both estimators have the same asymptotic efficiency in the limit according to Theorem 6. This implies that implicit SGD is a superior form of SGD, and should be preferred when the calculation of implicit updates (2.4) is computationally feasible. In Section 2.5 we show that this is possible in a large family of statistical models, and illustrate with several numerical experiments in Section 2.6.1.

Asymptotic variance results in stochastic approximation similar to Theorem 6 were first obtained by Chung (1954), Sacks (1958), and followed by Fabian (1968a), Polyak and Tsytkin (1979), and several other authors (see also Ljung *et al.*, 1992, Parts I, II). Our work is different in two important aspects. First, our asymptotic variance result includes implicit SGD, which is a stochastic approximation procedure with implicitly defined updates, whereas other works consider only explicit stochastic approximation procedures. Second, in our setting we estimate recursively the true parameters θ_\star of a statistical model, and thus we can exploit typical regularity conditions (Assumption 7(d)) to derive the asymptotic variance of θ_n^{im} (and θ_n^{sgd}) in a simplified closed-form; for example, under typical regularity conditions (see also Assumption 7(d)), $\text{Var}(\nabla \log f(Y; X, \theta_\star)) = \mathcal{I}(\theta_\star)$, which is used in the proof of Theorem 6. We illustrate the asymptotic variance results of Theorem 6 in experiments of Section 2.6.1.

Optimal learning rates

Crucially, the asymptotic variance formula of Theorem 6 depends on the limit of the sequence C_n used in the SGD procedures (2.3) and (2.4). We distinguish two classes of procedures, one where $C_n = I$ trivially, known as *first-order procedures*, and a second case where C_n is not trivial, known as *second-order procedures*.

In first-order procedures, $C_n = I$, i.e., only gradients are used in the SGD procedures.

Inevitably, no matter how we set the learning rate parameter γ_1 , first-order SGD procedures will lose statistical efficiency. We can immediately verify this by comparing the asymptotic variance in Theorem 6 with the asymptotic variance of the maximum likelihood estimator (MLE), denoted by θ_N^{mle} , on a data set with N data points $\{(X_n, Y_n)\}$, $n = 1, 2, \dots, N$. Under regularity conditions, the MLE is the asymptotically optimal unbiased estimator and $N\text{Var}(\theta_N^{\text{mle}} - \theta_*) \rightarrow \mathcal{I}(\theta_*)^{-1}$. By Theorem 6 and convergence of implicit SGD, it holds $N\text{Var}(\theta_n^{\text{im}} - \theta_*) \rightarrow \gamma_1^2(2\gamma_1\mathcal{I}(\theta_*) - I)^{-1}\mathcal{I}(\theta_*)$, which also holds for θ_n^{sgd} . For any $\gamma_1 > 0$ we have,

$$\gamma_1^2(2\gamma_1\mathcal{I}(\theta_*) - I)^{-1}\mathcal{I}(\theta_*) \succeq \mathcal{I}(\theta_*)^{-1}. \quad (2.8)$$

Therefore, both SGD estimators lose information and this loss can be quantified exactly by Ineq. (2.8). This inequality can also be leveraged to find the optimal choice for γ_1 given an appropriate objective. As demonstrated in the experiments in Section 2.6, this often suffices to achieve estimates that are comparable with MLE in statistical efficiency but with substantial computational gains. Assuming distinct eigenvalues λ_i for the matrix $\mathcal{I}(\theta_*)$, the eigenvalues of the variance matrix of both SGD estimators are equal to $\gamma_1^2\lambda_i/(2\gamma_1\lambda_i - 1)$, by Theorem 6. Therefore, one reasonable way to set the parameter γ_1 is to set it equal to γ_1^* , defined as

$$\gamma_1^* = \arg \min_{x > 1/2\underline{\lambda}_f} \sum_i x^2\lambda_i/(2x\lambda_i - 1). \quad (2.9)$$

Eq. (2.9) is under the constraint $x > 1/(2\underline{\lambda}_f)$, where $\underline{\lambda}_f = \min\{\lambda_i\}$, because Theorem 6 requires $2\gamma_1\mathcal{I}(\theta_*) - I$ to be positive-definite, and thus $(2x\lambda_i - 1)$ needs to be positive for every λ_i .

Of course, the eigenvalues λ_i 's are unknown in practice and need to be estimated from the data. This problem has received significant attention recently and several methods exist (see Karoui, 2008, and references within). We will use Eq. (2.9) extensively in our experiments (Section 2.6) in order to tune the SGD procedures. However, we note that in first-order SGD procedures, knowing the eigenvalues λ_i of $\mathcal{I}(\theta_*)$ does not necessarily

achieve statistical efficiency because of the spectral gap of $\mathcal{I}(\theta_*)$, i.e., the ratio between its maximum eigenvalue $\overline{\lambda}_f$ and minimum eigenvalue $\underline{\lambda}_f$; for instance, if $\underline{\lambda}_f = \overline{\lambda}_f$, then the choice of learning rate parameter (2.9) leads to statistically efficient first-order SGD procedures. However, this case is not typical in practice, especially in many dimensions.

In *second-order* procedures, we assume non-trivial condition matrices C_n . Such procedures are called second-order because they usually leverage information from the Fisher information matrix (or the Hessian of the log-likelihood), also known as curvature information. They are also known as *adaptive* procedures because they adapt their hyperparameters, i.e., learning rates γ_n or condition matrices C_n , according to observed data. For instance, let $C_n \equiv \mathcal{I}(\theta_*)^{-1}$ and $\gamma_1 = 1$. Plugging in $C_n = \mathcal{I}(\theta_*)^{-1}$ in Theorem 6, the asymptotic variance of the SGD estimators is

$$(1/n)\gamma_1^2(2\gamma_1\mathcal{I}(\theta_*)^{-1}\mathcal{I}(\theta_*) - I)^{-1}\mathcal{I}(\theta_*)^{-1}\mathcal{I}(\theta_*)\mathcal{I}(\theta_*)^{-1} = (1/n)\mathcal{I}(\theta_*)^{-1},$$

which is the theoretically optimal asymptotic variance of the MLE, i.e., the Cramér-Rao lower bound.

Therefore, to achieve asymptotic efficiency, second-order procedures need to estimate the Fisher information matrix. Because θ_* is unknown one can simply use $C_n = \mathcal{I}(\theta_n^{\text{im}})^{-1}$ (or $C_n = \mathcal{I}(\theta_{n-1}^{\text{sgd}})^{-1}$) as an iterative estimate of $\mathcal{I}(\theta_*)$, and the same optimality result holds. This approach in second-order explicit SGD was first studied by [Sakrison \(1965\)](#), and later by [Nevelson and Khasminski \(1973, Chapter 8, Theorem 5.4\)](#). It was later extended by [Fabian \(1978\)](#) and several other authors. Notably, [Amari \(1998\)](#) refers to the direction $\mathcal{I}(\theta_{n-1}^{\text{sgd}})^{-1}\nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{sgd}})$ as the “natural gradient” and uses information geometry arguments to prove statistical optimality.

An alternative way to implement second-order procedures is to use stochastic approximation to estimate $\mathcal{I}(\theta_*)$, in addition to the approximation procedure estimating θ_* . For

example, [Amari et al. \(2000\)](#) proposed the following second-order procedure,

$$\begin{aligned} C_n^{-1} &= (1 - a_n)C_{n-1}^{-1} + a_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{am}}) \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{am}})^\top \\ \theta_n^{\text{am}} &= \theta_{n-1}^{\text{am}} + \gamma_n C_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{am}}), \end{aligned} \quad (2.10)$$

where $a_n = a_1/n$ is a learning rate sequence, separate from γ_n . By standard stochastic approximation, C_n^{-1} converges to $\mathcal{I}(\theta_*)$, and thus procedure (2.10) is asymptotically optimal. However, there are two important problems with procedure (2.10). First, it is computationally costly because of matrix inversions. A faster way is to apply quasi-Newton ideas. SGD-QN developed by [Bordes et al. \(2009\)](#) is such a procedure where the first expensive matrix computations are substituted by typical secant conditions. Second, the stochastic approximation of $\mathcal{I}(\theta_*)$ is usually very noisy in high-dimensional problems and this affects the main approximation for θ_* . Recently, more robust variants of SGD-QN have been proposed ([Byrd et al., 2014](#)).

Another notable adaptive procedure is AdaGrad ([Duchi et al., 2011a](#)), which is defined as

$$\begin{aligned} C_n^{-1} &= C_{n-1}^{-1} + \text{diag} \left(\nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{ada}}) \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{ada}})^\top \right), \\ \theta_n^{\text{ada}} &= \theta_{n-1}^{\text{ada}} + \gamma_1 C_n^{1/2} \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{ada}}), \end{aligned} \quad (2.11)$$

where $\text{diag}(\cdot)$ takes the diagonal matrix of its matrix argument, and the learning rate is set constant to $\gamma_n \equiv \gamma_1$. AdaGrad can be considered a second-order procedure because it tries to approximate the Fisher information matrix, however it only uses gradient information so technically it is first-order. Under appropriate conditions, $C_n^{-1} \rightarrow \text{diag}(\mathcal{I}(\theta_*))$ and a simple modification in the proof of [Theorem 6](#) can show that the asymptotic variance of the AdaGrad estimate is given by

$$\sqrt{n} \text{Var}(\theta_n^{\text{ada}}) \rightarrow \frac{\gamma_1}{2} \text{diag}(\mathcal{I}(\theta_*))^{-1/2}. \quad (2.12)$$

This result reveals an interesting trade-off achieved by AdaGrad and a subtle contrast to first-order SGD procedures. The asymptotic variance of AdaGrad is $O(1/\sqrt{n})$, which indicates

significant loss of information. However, this rate is attained *regardless* of the specification of the learning rate parameter γ_1 .⁵ In contrast, as shown in Theorem 6, first-order SGD procedures require $2\gamma_1\mathcal{I}(\theta_*) - I \succ 0$ in order to achieve the $O(1/n)$ rate, and the rate is significantly worse if this condition is not met. For instance, Nemirovski *et al.* (2009) given an example of misspecification of γ_1 where the rate of first-order explicit SGD is $O(n^{-\epsilon})$, and ϵ can be arbitrarily small. The variance result (2.12) is illustrated in numerical experiments of Section 2.6.1.

2.4.3 Optimality with averaging

As shown in Section 2.4.2, Theorem 6 implies that first-order SGD procedures can be statistically inefficient, especially in many dimensions. One surprisingly simple idea to achieve statistical efficiency is to combine larger learning rates with averaging of the iterates. In particular, we consider the procedure

$$\begin{aligned}\theta_n^{\text{im}} &= \theta_{n-1}^{\text{im}} + \gamma_n \nabla \log f(Y_n; X_n, \theta_n^{\text{im}}), \\ \bar{\theta}_n^{\text{im}} &= \frac{1}{n} \sum_{i=1}^n \theta_i^{\text{im}},\end{aligned}\tag{2.13}$$

where θ_n^{im} are the typical implicit SGD iterates (2.4), and $\gamma_n = \gamma_1 n^{-\gamma}$, $\gamma \in [0.5, 1)$. Under suitable conditions, the iterates $\bar{\theta}_n^{\text{im}}$ are asymptotically efficient. This is formalized in the following theorem.

Theorem 7. *Consider SGD procedure (2.13) and suppose Assumptions 7(a),(c),(d), and (e) hold, where $\gamma \in [0.5, 1)$. Then, the iterate $\bar{\theta}_n$ converges to θ_* in probability and is asymptotically efficient, i.e.,*

$$n\text{Var}(\bar{\theta}_n) \rightarrow \mathcal{I}(\theta_*)^{-1}.$$

⁵ This follows from a technicality in Lemma 2. On a high-level, the term γ_{n-1}/γ_n is important for the variance rates of AdaGrad and SGD. When $\gamma_n \propto 1/n$, as in Theorem 6, it holds $\gamma_{n-1}/\gamma_n = 1 + \gamma_n/\gamma_1 + O(\gamma_n^2)$, which explains the quantity $2\mathcal{I}(\theta_*) - I/\gamma_1$ in first-order SGD. The rate $O(1/n)$ is attained only if $2\mathcal{I}(\theta_*) - I/\gamma_1 \succ 0$. When $\gamma_n \propto 1/\sqrt{n}$, as in AdaGrad, it holds $\gamma_{n-1}/\gamma_n = 1 + o(\gamma_n)$ and the rate $O(1/\sqrt{n})$ is attained without any additional requirement.

Remarks. In the context of explicit stochastic approximations, averaging was first proposed and analyzed by [Ruppert \(1988\)](#) and [Bather \(1989\)](#). [Ruppert \(1988\)](#) argued that larger learning rates in stochastic approximation uncorrelates the iterates allowing averaging to improve efficiency. [Polyak and Juditsky \(1992b\)](#) expanded the scope of averaging by proving asymptotic optimality in more general explicit stochastic approximations that operate under suitable conditions similar to [Theorem 7](#). [Polyak and Juditsky \(1992b\)](#) thus proved that slowly-converging stochastic approximations can be improved by using larger learning rates and averaging of the iterates. Recent work has analyzed explicit updates with averaging ([Zhang, 2004](#); [Xu, 2011](#); [Bach and Moulines, 2013](#); [Shamir and Zhang, 2012](#)), and has shown their superiority in numerous learning tasks.

2.4.4 Asymptotic normality

Asymptotic distributions, or more generally invariance principles, are well-studied in classical stochastic approximation ([Ljung *et al.*, 1992](#), Chapter II.8). In this section we leverage Fabian’s theorem ([Fabian, 1968a](#)) to show that iterates from implicit SGD are asymptotically normal.

Theorem 8. *Suppose that Assumptions [7\(a\)](#),[\(c\)](#),[\(d\)](#),[\(e\)](#),[\(f\)](#) hold. Then, the iterate θ_n^{im} of implicit SGD [\(2.4\)](#) is asymptotically normal, such that*

$$n^{\gamma/2}(\theta_n^{\text{im}} - \theta_*) \rightarrow \mathcal{N}_p(0, \Sigma),$$

where $\Sigma = \gamma_1^2 (2\gamma_1 C\mathcal{I}(\theta_*) - I)^{-1} C\mathcal{I}(\theta_*)C$.

Remarks. The combined results of [Theorems 5](#), [6](#), and [8](#) indicate that implicit SGD is numerically stable and has known asymptotic variance and distribution. Therefore, contrary to explicit SGD that has severe stability issues, implicit SGD emerges as a stable estimation procedure with known standard errors, which enables typical statistical tasks, such as confidence intervals, hypothesis testing, and model checking. We show empirical evidence supporting this claim in [Section 2.6.1](#).

2.4.5 Stability

To illustrate the stability, or lack thereof, of both SGD estimators in small-to-moderate samples, we simplify the SGD procedures and inspect the size of the biases $\mathbb{E}(\theta_n^{\text{sgd}} - \theta_*)$ and $\mathbb{E}(\theta_n^{\text{im}} - \theta_*)$. In particular, based on Theorem 5, we simply assume the Taylor expansion $\nabla \log f(Y_n; X_n, \theta_n) = -\mathcal{I}(\theta_*)(\theta_n - \theta_*) + O(\gamma_n)$; to simplify further we ignore the remainder term $O(\gamma_n)$.

Under this simplification, the SGD procedures (2.3) and (2.4) can be written as follows:

$$\mathbb{E}(\theta_n^{\text{sgd}} - \theta_*) = (I - \gamma_n \mathcal{I}(\theta_*)) \mathbb{E}(\theta_{n-1}^{\text{sgd}} - \theta_*) = P_1^n b_0, \quad (2.14)$$

$$\mathbb{E}(\theta_n^{\text{im}} - \theta_*) = (I + \gamma_n \mathcal{I}(\theta_*))^{-1} \mathbb{E}(\theta_{n-1}^{\text{im}} - \theta_*) = Q_1^n b_0, \quad (2.15)$$

where $P_1^n = \prod_{i=1}^n (I - \gamma_i \mathcal{I}(\theta_*))$, $Q_1^n = \prod_{i=1}^n (I + \gamma_i \mathcal{I}(\theta_*))^{-1}$, and b_0 denotes the initial bias of the two procedures from a common starting point θ_0 . Thus, the matrices P_1^n and Q_1^n describe how fast the initial bias decays for the explicit and implicit SGD respectively. In the limit, $P_1^n \rightarrow 0$ and $Q_1^n \rightarrow 0$ (see proof of Lemma 2), and thus both methods are *asymptotically stable*.

However, the explicit procedure has significant stability issues in small-to-moderate samples. By inspection of Eq. (2.14), the magnitude of P_1^n is dominated by $\overline{\lambda}_f$, the maximum eigenvalue of $\mathcal{I}(\theta_*)$. Furthermore, the rate of convergence is dominated by $\underline{\lambda}_f$, the minimum eigenvalue of $\mathcal{I}(\theta_*)$.⁶ For stability, it is desirable $|1 - \gamma_1 \lambda_i| < 1$, and for fast convergence $|1 - \gamma_1 \lambda_i| \approx 0$, for all eigenvalues $\lambda_i \in \text{eig}(\mathcal{I}(\theta_*))$. This roughly implies the requirements $\gamma_1 < 2/\overline{\lambda}_f$ for stability, and $\gamma_1 > 1/\underline{\lambda}_f$ for convergence. This is problematic in high-dimensional settings because $\overline{\lambda}_f$ is typically orders of magnitude larger than $\underline{\lambda}_f$. Thus, the requirements for stability and speed of convergence are in conflict: to ensure stability we need a small learning rate parameter γ_1 , thus paying a high price in convergence which will be at the order of $O(n^{-\gamma_1 \underline{\lambda}_f})$, and vice versa.

In contrast, the implicit procedure is *unconditionally stable*. The eigenvalues of Q_1^n are $\lambda'_i = \prod_{j=1}^n 1/(1 + \gamma_j \lambda_i/j) = O(n^{-\gamma_1 \lambda_i})$. Critically, it is no longer required to have a small γ_1

⁶To see this, note that the eigenvalues of P_1^n are $\lambda'_i = \prod_j (1 - \gamma_j \lambda_i/j) = O(n^{-\gamma_1 \lambda_i})$ if $0 < \gamma_1 \lambda_i < 1$. See also proof of Lemma 1.

for stability because the eigenvalues of Q_1^n are always less than one. We summarize these findings in the following lemma.

Lemma 1. *Let $\bar{\lambda}_f = \max \text{eig}(\mathcal{I}(\theta_*)$), and suppose $\gamma_n = \gamma_1/n$ and $\gamma_1 \bar{\lambda}_f > 1$. Then, then the maximum eigenvalue of P_1^n satisfies*

$$\max_{n>0} \max\{\text{eig}(P_1^n)\} = \Theta(2^{\gamma_1 \bar{\lambda}_f} / \sqrt{\gamma_1 \bar{\lambda}_f}).$$

For the implicit method,

$$\max_{n>0} \max\{\text{eig}(Q_1^n)\} = O(1).$$

Remark. Lemma 1 shows that in the explicit SGD procedure the effect from the initial bias can be amplified in an arbitrarily large way before fading out if the learning rate is misspecified (i.e., if $\gamma_1 \gg 1/\bar{\lambda}_f$). This sensitivity of explicit SGD is well-known and requires problem-specific considerations to be avoided in practice e.g., pre-processing, small-sample tests, projections, truncation (Chen *et al.*, 1987). In fact, there exists voluminous work, continuing up-to-date, in designing learning rates to stabilize explicit SGD; see, for example, a review by George and Powell (2006). Implicit procedures render such ad-hoc designs obsolete because they remain stable regardless of learning rate design, and still maintain the asymptotic convergence and efficiency properties of explicit SGD.

2.5 Applications

In this section we show how to apply implicit SGD (2.4) for estimation in generalized linear models, Cox proportional hazards, and more general M-estimation problems. We start by developing an algorithm that efficiently computes the implicit updates (2.4), which is generally applicable to all aforementioned applications.

2.5.1 Efficient computation of implicit updates

The main difficulty in applying implicit SGD is the solution of the multidimensional fixed point equation (2.4). In a large family of models where the likelihood depends on the parameter θ_* only through the natural parameter $X_n^\top \theta_*$, the solution of the fixed-point equation is feasible and computationally efficient. We prove the general result in Theorem 9, which depends on Assumption 7(b).

For the rest of this section we will treat $\ell(X^\top \theta; Y)$ as a function of the natural parameter $X^\top \theta$ (or $X_n^\top \theta$ referring to the n th data point) for a fixed outcome Y . Thus, $\ell'(X^\top \theta; Y)$ will refer to the first derivative of ℓ with respect to $X^\top \theta$ with fixed Y .

Theorem 9. *Suppose Assumption 7(b) holds. Then, the gradient for the implicit update (2.4) is a scaled version of the gradient at the previous iterate, i.e.,*

$$\nabla \log f(Y_n; X_n, \theta_n^{\text{im}}) = \lambda_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}}), \quad (2.16)$$

where the scalar λ_n satisfies,

$$\lambda_n \ell'(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) = \ell'(X_n^\top \theta_{n-1}^{\text{im}} + \gamma_n \lambda_n \ell'(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) X_n^\top C_n X_n; Y_n). \quad (2.17)$$

Remarks. Theorem 9 shows that the gradient $\nabla \log f(Y_n; X_n, \theta_n^{\text{im}})$ in the implicit update (2.4) is a scaled gradient $\nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}})$ calculated at the previous iterate θ_{n-1}^{im} . Therefore, computing the implicit update reduces to finding the scale factor λ_n in Theorem 9. Under Assumption 7(b) narrow search bounds for λ_n are available. This leads to Algorithm 1 that is a generic implementation of implicit SGD for models satisfying Assumption 7(b). This implementation is fast because the interval B_n of search bounds in Algorithm 1 has size $O(\gamma_n)$.

2.5.2 Generalized linear models

In this section, we apply implicit SGD to estimate generalized linear models (GLMs). In such models, Y_n follows an exponential distribution conditional on X_n , and $\mathbb{E}(Y_n | X_n) = h(X_n^\top \theta_*)$, where h is the *transfer function* of the GLM model (Nelder and Wedderburn, 1972; Dobson

Algorithm 1: Efficient implementation of implicit SGD (2.4)

```

1: for all  $n \in \{1, 2, \dots\}$  do
2:   # compute search bounds  $B_n$ 
3:    $r_n \leftarrow \gamma_n \ell' (X_n^\top \theta_{n-1}^{\text{im}}; Y_n)$ 
4:    $B_n \leftarrow [0, r_n]$ 
5:   if  $r_n \leq 0$  then
6:      $B_n \leftarrow [r_n, 0]$ 
7:   end if
8:   # solve fixed-point equation by a root-finding method
9:    $\xi = \gamma_n \ell' (X_n^\top \theta_{n-1}^{\text{im}} + \xi X_n^\top C_n X_n; Y_n)$ ,  $\xi \in B_n$ 
10:   $\lambda_n \leftarrow \xi / r_n$ 
11:  # following update is equivalent to update (2.4)
12:   $\theta_n^{\text{im}} \leftarrow \theta_{n-1}^{\text{im}} + \gamma_n \lambda_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}})$ 
13: end for

```

and Barnett, 2008). Furthermore, the gradient of the GLM log-likelihood for parameter value θ at data point (X_n, Y_n) is given by

$$\nabla \log f(Y_n; X_n, \theta) = [Y_n - h(X_n^\top \theta)] X_n. \quad (2.18)$$

The conditional variance of Y_n is $\text{Var}(Y_n | X_n) = h'(X_n^\top \theta_*) X_n X_n^\top$, and thus the Fisher information matrix is $\mathcal{I}(\theta) = \mathbb{E}(h'(X_n^\top \theta) X_n X_n^\top)$.

Thus, SGD procedures (2.3) and (2.4) can be written as

$$\theta_n^{\text{sgd}} = \theta_{n-1}^{\text{sgd}} + \gamma_n C_n [Y_n - h(X_n^\top \theta_{n-1}^{\text{sgd}})] X_n, \quad (2.19)$$

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n C_n [Y_n - h(X_n^\top \theta_n^{\text{im}})] X_n. \quad (2.20)$$

Implementation of explicit SGD is straightforward. Implicit SGD can be implemented through Algorithm 1 because the conditions of Assumption 7(b) are fulfilled. In particular, $\log f(Y; X, \theta) \equiv \ell(X^\top \theta; Y)$ where $\ell(\eta; Y) = Y - h(\eta)$. In typical models h is twice-differentiable and also $h'(\eta) \geq 0$ because it is proportional to the conditional variance of Y given X , thus fulfilling Assumption 7(b). In the simplified case where $C_n = I$, the identity matrix, for all n , Algorithm 1 simplifies to Algorithm 2, which was first derived by Toulis *et al.* (2014). We make extensive experiments using Algorithm 2 in Section 2.6.2.

Algorithm 2: Estimation of GLMs with implicit SGD

```
1: for all  $n \in \{1, 2, \dots\}$  do
2:    $r_n \leftarrow \gamma_n [Y_n - h(X_n^\top \theta_{n-1}^{\text{im}})]$ 
3:    $B_n \leftarrow [0, r_n]$ 
4:   if  $r_n \leq 0$  then
5:      $B_n \leftarrow [r_n, 0]$ 
6:   end if
7:    $\zeta = \gamma_n [Y_n - h(X_n^\top \theta_{n-1}^{\text{im}} + \zeta \|X_n\|^2)], \zeta \in B_n$ 
8:    $\theta_n^{\text{im}} \leftarrow \theta_{n-1}^{\text{im}} + \zeta X_n$ 
9: end for
```

2.5.3 Cox proportional hazards

In this section, we apply SGD to estimate a Cox proportional hazards model, which is a popular model in survival analysis of censored failure times (Cox, 1972; Klein and Moeschberger, 2003). Multiple variations of the model exist, but we will analyze one simple variation that is popular in practice (Davison, 2003). Consider N individuals, indexed by i , with observed survival times Y_i , failure indicators d_i , and covariates X_i . The survival times can be assumed ordered, $Y_1 < Y_2 \dots < Y_N$, whereas $d_i = 1$ denotes failure (e.g., death) and $d_i = 0$ indicates censoring (e.g., patient dropped out of study). Given a failure for unit i ($d_i = 1$) at time Y_i , the *risk set* \mathcal{R}_i is defined as the set of individuals that could possibly fail at Y_i , i.e., all individuals except those who failed or were censored before Y_i . In our simplified model, $\mathcal{R}_i = \{i, i + 1, \dots, N\}$. Define $\eta_i(\theta) \stackrel{\text{def}}{=} \exp(X_i^\top \theta)$, then the log-likelihood ℓ for θ is given by (Davison, 2003, Chapter 10)

$$\ell(\theta; X, Y) = \sum_{i=1}^N [d_i - H_i(\theta) \eta_i(\theta)] X_i, \quad (2.21)$$

where $H_i(\theta) = \sum_{j:i \in \mathcal{R}_j} d_j (\sum_{k \in \mathcal{R}_j} \eta_k(\theta))^{-1}$. In an online setting, where N is infinite and data points (X_i, Y_i) are observed one at a time, future observations affect the likelihood of previous ones, as can be seen by inspection of Eq. (2.21). Therefore, we apply SGD assuming fixed N to estimate the MLE θ_N^{mle} . As mentioned in Section 2.2, our theory in Section 2.4 can be applied unchanged if we only substitute θ_* , the true parameter, with the MLE θ_N^{mle} .

A straightforward implementation of explicit SGD (2.3) for the Cox model is shown in

Algorithm 3. For implicit SGD (2.4) we have the update

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n [d_i - H_i(\theta_n^{\text{im}}) \eta_i(\theta_n^{\text{im}})] X_i, \quad (2.22)$$

which is similar to the implicit procedure for GLMs (2.20). However, the log-likelihood term $d_i - H_i(\theta_n^{\text{im}}) \eta_i(\theta_n^{\text{im}})$ does not satisfy the conditions of Assumption 7(b). Although the term $\eta_i(\theta)$ is increasing with respect to $X_i^\top \theta$, $H_i(\theta)$ may be increasing or may be decreasing because it depends on terms $X_j^\top \theta, j \neq i$, as well. Thus, Theorem 9 cannot be applied. One way to circumvent this problem and apply Theorem 9 is to simply compute $H_i(\cdot)$ on the previous update θ_{n-1}^{im} instead of the current θ_n^{im} . Then, update (2.22) becomes,

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n [d_i - H_i(\theta_{n-1}^{\text{im}}) \eta_i(\theta_n^{\text{im}})] X_i, \quad (2.23)$$

which now satisfies Assumption 7(b) as $H_i(\theta_{n-1}^{\text{im}})$ is constant with respect to θ_n^{im} .⁷ The implicit SGD procedure for fitting Cox models is shown in Algorithm 4. We run experiments with implicit SGD on Cox models with simulated data in Section 2.6.2.

Algorithm 3: Explicit SGD for Cox model	Algorithm 4: Implicit SGD for Cox model
<pre> 1 for $n = 1, 2, \dots$ do 2 $i \leftarrow \text{sample}(1, N)$ 3 $\hat{H}_i \leftarrow \sum_{j:i \in \mathcal{R}_j} \frac{d_j}{\sum_{k \in \mathcal{R}_j} \eta_k(\theta_{n-1}^{\text{sgd}})}$ 4 $w_{n-1} \leftarrow [d_i - \hat{H}_i \eta_i(\theta_{n-1}^{\text{sgd}})]$ 5 $\theta_n^{\text{sgd}} = \theta_{n-1}^{\text{sgd}} + \gamma_n w_{n-1} C_n X_i$ </pre>	<pre> 1 for $n = 1, 2, \dots$ do 2 $i \leftarrow \text{sample}(1, N)$ 3 $\hat{H}_i \leftarrow \sum_{j:i \in \mathcal{R}_j} \frac{d_j}{\sum_{k \in \mathcal{R}_j} \eta_k(\theta_{n-1}^{\text{im}})}$ 4 $w(\theta) \stackrel{\text{def}}{=} d_i - \hat{H}_i \eta_i(\theta)$ 5 $W_n \leftarrow w(\theta_{n-1}^{\text{im}}) C_n X_i$ 6 $\lambda_n w(\theta_{n-1}^{\text{im}}) = w(\theta_{n-1}^{\text{im}} + \gamma_n \lambda_n W_n)$ 7 $\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n \lambda_n W_n$ </pre>

⁷This idea can be more generally used in order to apply implicit SGD on models that do not satisfy Assumption 7(b); see Section 2.7 for a discussion.

2.5.4 M-Estimation

Given N observed data points (X_i, Y_i) and a convex function $\rho : \mathbb{R} \rightarrow \mathbb{R}^+$, the *M-estimator* is defined as

$$\hat{\theta}^m = \arg \min_{\theta} \sum_{i=1}^N \rho(Y_i - X_i^\top \theta), \quad (2.24)$$

where it is assumed $Y_i = X_i^\top \theta_* + \epsilon_i$, and ϵ_i are i.i.d. zero mean-valued noise. M-estimators are especially useful in robust statistics (Huber *et al.*, 1964; Huber, 2011) because appropriate choice of ρ can reduce the influence of outliers in data. Typically, ρ is twice-differentiable around zero. In this case,

$$\mathbb{E} (\rho'(Y_i - X_i^\top \hat{\theta}^m) X_i) = 0, \quad (2.25)$$

where the expectation is over the empirical data distribution. Thus, according to Section 2.2, SGD procedures can be applied to approximate the M-estimator $\hat{\theta}^m$.⁸ There has been increased interest in the literature for fast approximation of M-estimators due to the robustness properties of such estimators (Donoho and Montanari, 2013; Jain *et al.*, 2014).

The implicit SGD procedure for approximating M-estimators is defined in Algorithm 5, and is a simple adaptation of Algorithm 1.

Algorithm 5: Implicit SGD for M-estimation

```

1 for  $n = 1, 2, \dots$  do
2    $i \leftarrow \text{sample}(1, N)$ 
3    $w(\theta) \stackrel{\text{def}}{=} \rho'(Y_i - X_i^\top \theta)$ 
4    $\lambda_n w(\theta_{n-1}^{\text{im}}) \leftarrow w(\theta_{n-1}^{\text{im}} + \gamma_n \lambda_n w(\theta_{n-1}^{\text{im}}) C_n X_i)$  # implicit update
5    $\theta_n^{\text{im}} \leftarrow \theta_{n-1}^{\text{im}} + \gamma_n \lambda_n w(\theta_{n-1}^{\text{im}}) C_n X_i$ 

```

Importantly, ρ is convex and thus $\rho'' \geq 0$ and therefore the conditions of Assumption 7(b) are met. Thus, Step 4 of Algorithm 5 is a straightforward application of Algorithm 1 by simply setting $\ell'(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) \equiv \rho'(Y_n - X_n^\top \theta_{n-1}^{\text{im}})$. The asymptotic variance of θ_n^{im} is also easy

⁸ It is also typical to assume that the density of ϵ_i is symmetric around zero. Therefore, it also holds $\mathbb{E} (\rho'(Y_i - X_i^\top \theta_*) X_i) = 0$, where the expectation is over the true data distribution. According to Section 2.2 SGD procedures can also be used to estimate θ_* in the case of infinite stream of observations ($N = \infty$). In this section we only consider the case of finite N , but it is trivial to adapt our procedures to infinite N .

to derive. If $S \stackrel{\text{def}}{=} \mathbb{E}(X_n X_n^\top)$, $C_n \rightarrow C > 0$ such that S and C commute, $\psi^2 \stackrel{\text{def}}{=} \mathbb{E}(\rho'(\epsilon_i)^2)$, and $v(z) \stackrel{\text{def}}{=} \mathbb{E}(\rho'(\epsilon_i + z))$, Theorem 6 can be leveraged to show that

$$n\text{Var}(\theta_n^{\text{im}}) \rightarrow \psi^2(2v'(0)CS - I)^{-1}CSC. \quad (2.26)$$

Historically, one of the first applications of explicit stochastic approximation procedures in robust estimation was due to [Martin and Masreliez \(1975\)](#). The asymptotic variance (2.26) was first derived, only for the explicit SGD case, by [Poljak and Tsykin \(1980\)](#) using stochastic approximation theory from [Nevelson and Khasminski \(1973\)](#).

2.6 Simulation and data analysis

In this section, we demonstrate the computational and statistical advantages of SGD estimation procedures (2.3) and (2.4). For our experiments we developed a new R package, namely `sgd`, which has been published on CRAN.⁹

The experiments are split into three sets. In the first set, presented in Section 2.6.1, we aim to validate the theoretical results of Section 2.4. In the second set, presented in Section 2.6.2, we aim to validate the performance of SGD procedures. We focus particularly on implicit SGD because it is a more stable estimation procedure, whereas for explicit SGD we could not devise a learning rate design that worked uniformly well in all experiments. The third set of experiments, presented in the Appendix, focuses on comparisons of implicit SGD, including implicit SGD with averaging, against popular machine learning methods, such as averaged explicit SGD on a large support vector machine (SVM) model, `prox-SAG` ([Schmidt *et al.*, 2013](#)), `prox-SVRG` ([Xiao and Zhang, 2014](#)), and averaged explicit SGD ([Xu, 2011](#); [Bach and Moulines, 2013](#)), on typical machine learning tasks.

⁹Our R package resides on CRAN at <http://cran.r-project.org/web/packages/sgd/index.html>, and has been co-authored with Dustin Tran and Kuang Ye. All experiments were conducted on a single laptop running Linux Ubuntu 13.x with 8 cores@2.4GHz, 16Gb of RAM memory and 256Gb of physical storage with SSD technology.

2.6.1 Validation of theory

In this section we aim to validate the theoretical results of Section 2.4, namely the result on asymptotic variance (Theorem 6 and asymptotic normality (Theorem 8) of SGD procedures.

Asymptotic variance

In this experiment we use a normal linear model following the experimental setup of Xu (2011) to check the asymptotic variance results of Theorem 6. The procedures we test are explicit SGD (2.3), implicit SGD (2.4), and AdaGrad (2.11). For simplicity we use first-order SGD where $C_n \equiv I$.

In the experiment we calculate the empirical variance of said procedures for 25 values of their common learning rate parameter γ_1 in the interval $[1.2, 10]$. For every value of γ_1 we calculate the empirical variances through the following process, repeated for 150 times. First, we set $\theta_\star = (1, 1, \dots, 1)^\top \in \mathbb{R}^{20}$ as the true parameter value. For iterations $n = 1, 2, \dots, 1500$, we sample covariates as $X_n \sim \mathcal{N}_p(0, S)$, where S is diagonal with elements uniformly in $[0.5, 5]$. The outcome Y_n is then sampled as $Y_n | X_n \sim \mathcal{N}(X_n^\top \theta_\star, 1)$. In every repetition we store the iterate θ_{1500} for every tested procedure and then calculate the empirical variance of stored iterates over all 150 repetitions.

For any fixed learning rate parameter γ_1 , we set the learning rates as follows. For implicit SGD we set $\gamma_n = \gamma_1/n$, and for AdaGrad we set $\gamma_n = \gamma_1$, as it is typical. However, for explicit SGD we set $\gamma_n = \min(0.3, \gamma_1/(n + \|X_n\|^2))$ in order to stabilize its updates. This trick is necessary by the analysis of Section 2.4.5. In particular, the Fisher information matrix is $\mathcal{I}(\theta_\star) = \mathbb{E}(X_n X_n^\top) = S$, and thus the minimum eigenvalue is $\underline{\lambda}_f = 0.5$ and the maximum is $\overline{\lambda}_f = 5$. Therefore, for stability we require $\gamma_1 < 1/\overline{\lambda}_f = 0.2$ and for fast convergence we require $\gamma_1 < 1/(2\underline{\lambda}_f) = 1$. The two requirements are incompatible, which indicates that explicit SGD can have serious stability issues.

For given $\gamma_1 > 1$, the asymptotic variance of SGD procedures after n iterations is $(1/n)\gamma_1^2(2\gamma_1 S - I)^{-1}S$, by Theorem 6. The asymptotic variance of AdaGrad after n iterations is equal to $(\gamma_1/2\sqrt{n})S^{-1/2}$, by the analysis of Section 2.4.2. The log traces of the empirical




Figure 2.1: Simulation with normal model. The x-axis corresponds to learning rate parameter γ_1 ; the y-axis curves corresponds to log trace of the empirical variance of tested procedures (explicit/implicit SGD, AdaGrad). Theoretical asymptotic variances of SGD and AdaGrad are plotted as well. Implicit SGD is stable and its empirical variance is very close to its asymptotic value. Explicit SGD becomes unstable at large γ_1 . AdaGrad is statistically inefficient but remains stable to large learning rates.

variance of the SGD procedures and AdaGrad in this experiment are shown in Figure B.4. The x-axis corresponds to different values of the learning rate parameter γ_1 , and the y-axis corresponds to the log trace of the empirical variance of the iterates for all three different procedures. We also include curves for the aforementioned theoretical values of the empirical variances.

We see that our theory predicts well the empirical variance of all methods. Explicit SGD performs on par with implicit SGD for moderate values of γ_1 , however, it required a modification in its learning rate to make it work. Furthermore, explicit SGD quickly becomes unstable at larger values of γ_1 (see, for example, its empirical variance for $\gamma_1 = 10$), and in several instances, not considered in Figure B.4, it numerically diverged. On the other hand, AdaGrad is stable to the specification of γ_1 and tracks its theoretical variance well. However, it gives inefficient estimators because their variance has order $O(1/\sqrt{n})$. Implicit SGD effectively combines stability and good statistical efficiency. First, it remains very

stable to the entire range of the learning rate parameter γ_1 . Second, its empirical variance is $O(1/n)$ and it tracks closely the theoretical value predicted by Theorem 6 for all γ_1 .

Asymptotic normality

In this experiment we use the normal linear model in the setup of Section 2.6.1 to check the asymptotic normality result of Theorem 8. We only test first-order implicit SGD (2.4) and first-order explicit SGD for simplicity. Asymptotic normality for explicit stochastic approximations has been shown Fabian (1968b) and several other authors (see also Ljung *et al.*, 1992, Parts I, II).

In the experiment we define a set learning rates $(0.5, 1, 3, 5, 6, 7)$. For every learning rate and for every method, implicit or explicit SGD, we take 400 samples of $N(\theta_N^{\text{im}} - \theta_\star)^\top \Sigma^{-1} (\theta_N^{\text{im}} - \theta_\star)$, where $N = 1200$; i.e., we run each SGD procedure for 1200 iterations. The matrix Σ is the asymptotic variance matrix in Theorem 8, and $\theta_\star \stackrel{\text{def}}{=} 10 \exp(-2 \cdot (1, 2, \dots, p))$, is the true parameter value. We use the ground-truth values both for Σ and θ_\star , as we are only interested to test normality of the iterates in this experiment. We also tried $p = 5, 10, 100$ as the parameter dimension. Because the explicit SGD was very unstable across experiments we only report results for $p = 5$. Results on the implicit procedure for larger p are given in Appendix B.7.1, where we also include results for a logistic regression model.

By Theorem 8 the quantity $N(\theta_N^{\text{im}} - \theta_\star)^\top \Sigma^{-1} (\theta_N^{\text{im}} - \theta_\star)$ is a chi-squared random variable with p degrees of freedom. Thus, for every procedure we plot this quantity against independent samples from a χ_p^2 distribution and visually check for deviations. As before, we tried to stabilize explicit SGD as much as possible by setting $\gamma_n = \min(0.3, \gamma_1 / (n + \|X_n\|^2))$ as the learning rate. This worked in many iterations, but not for all. Iterations for which explicit SGD diverged were not considered. For implicit SGD we simply set $\gamma_n = \gamma_1 / n$ without additional tuning.

The results of this experiment are shown in Figure 2.2. The vertical axis on the grid corresponds to different values of the learning rate parameter γ_1 , and the horizontal axis has histograms of $N(\theta_N - \theta_\star)^\top \Sigma^{-1} (\theta_N - \theta_\star)$ for both implicit and explicit procedures, and




Figure 2.2: Simulation with normal model. The x-axis corresponds to the SGD procedure (explicit or implicit) for various values of the learning rate parameter, $\gamma_1 \in \{0.5, 1, 3, 5, 7\}$. The histograms (x-axis) for the SGD procedures are 500 replications of SGD where at each replication we only store the quantity $N(\theta_N - \theta_*)^\top \Sigma^{-1}(\theta_N - \theta_*)$, for every method ($N = 1200$); the theoretical covariance matrix Σ is different for every learning rate and is given in Theorem 6. The data generation model is the same as in Section 2.6.1. We observe that implicit SGD is stable and follows the nominal chi-squared distribution. Explicit SGD becomes unstable at larger γ_1 and its distribution does not follow the nominal one well. In particular, the distribution of $N(\theta_N^{\text{sgd}} - \theta_*)^\top \Sigma^{-1}(\theta_N^{\text{sgd}} - \theta_*)$ becomes increasingly heavy-tailed as the learning rate parameter gets larger, and eventually diverges for $\gamma_1 \geq 7$.

also includes samples from a χ_5^2 distribution for visual comparison.

We see that the distribution $N(\theta_N^{\text{im}} - \theta_*)^\top \Sigma^{-1}(\theta_N^{\text{im}} - \theta_*)$ of the implicit iterates follows the nominal chi-squared distribution. This also seems to be unaffected by the size learning rate parameter. However, the distribution $N(\theta_N^{\text{sgd}} - \theta_*)^\top \Sigma^{-1}(\theta_N^{\text{sgd}} - \theta_*)$ of the explicit iterates does not follow a chi-squared distribution, except for small learning rate parameter values. For example, as the learning rate parameter increases, the distribution becomes more heavy-tailed (e.g., for $\gamma_1 = 6$), indicating that explicit SGD becomes unstable. Particularly for $\gamma_1 = 7$ explicit SGD diverged in almost all replications, and thus a histogram could not be constructed.

2.6.2 Validation of performance

In this section we aim to validate the performance of implicit SGD estimation against deterministic estimation procedures that are optimal. In Section 2.6.2, we compare implicit SGD with R's `glm()` function, which calculates the theoretically optimal MLE. In this experiment, we wish to test the computational efficiency of implicit SGD in terms of the problem size, i.e., parameter dimension p , and number of data points N , and its statistical efficiency in terms of MSE. We perform similar tests in Section 2.6.2 but on larger problem sizes, where we compare with R's `biglm` package, which is popular in estimating GLM models in large data sets (large N , small p). In Section 2.6.2 we compare implicit SGD with `glmnet` of Friedman *et al.* (2010), which is an efficient implementation of the elastic net for several GLMs for problems with $N > p$. In Sections 2.6.2, we re-analyze data from the NMMAPS study (Samet *et al.*, 2000) and show how implicit SGD can be naturally extended to fit a large generalized additive model (GAM) (Hastie and Tibshirani, 1990) to estimate the effects of air pollution on public health. We also compare with recently published statistical methods that are specifically designed to fit large-scale GAMs, and demonstrate SGD's superior performance.

In summary, our experimental results demonstrate that implicit SGD performs on par with optimal deterministic methods, even in problems with moderate-to-large data sets where stochastic methods are usually avoided. Combined its strong theoretical guarantees that were validated in Section 2.6.1, implicit SGD emerges as a principled estimation method that can be the workhorse of efficient estimation with large data sets in statistical practice.

Experiments with `glm()` function

The built-in function `glm()` in R¹⁰ performs deterministic maximum-likelihood estimation through iterative reweighted least squares. In this experiment, we wish to compare computing time and MSE between first-order implicit SGD and `glm()`. Our simulated data set is a

¹⁰Documentation is available at <http://stat.ethz.ch/R-manual/R-patched/library/stats/html/glm.html>.

simple normal linear model constructed as follows. First, we sample a binary $p \times p$ design matrix $X = (x_{ij})$ such that $x_{i1} = 1$ (intercept) and $P(x_{ij} = 1) = s$ i.i.d, where $s \in (0, 1)$ determines the sparsity of X . We set $s = 0.08$ indicating that roughly 8% of the X matrix will be nonzero. We generate θ_* by sampling p elements from $(-1, -0.35, 0, 0.35, 1)$ with replacement. The outcomes are $Y_i = X_i^\top \theta_* + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, 1)$ i.i.d., and $X_i = (x_{ij})$ is the $p \times 1$ vector of i 's covariates. By GLM properties,

$$\mathcal{I}(\theta_*) = \mathbb{E} (h'(X_i^\top \theta_*) X_i X_i^\top) = \begin{pmatrix} 1 & s & s & \cdots & s \\ s & s & s^2 & \cdots & s^2 \\ s & s^2 & s & s^2 & \cdots \\ \cdots & s^2 & \cdots & s & \cdots \\ s & s^2 & \cdots & \cdots & s \end{pmatrix}.$$

Slightly tedious algebra can show that the eigenvalues of $\mathcal{I}(\theta_*)$ are $s(1-s)$ with multiplicity $(p-2)$ and the two solutions of $x^2 - A(s)x + B(s) = 0$, where $A(s) = 1 + s + s^2(p-2)$ and $B(s) = s(1-s)$. It is thus possible to use the analysis of Section 2.4.2 and Eq. (2.9) to derive a theoretically optimal learning rate. We sample 200 pairs (p, N) for the problem size, uniformly in the ranges $p \sim [10, 500]$ and $N \sim [500, 50000]$, and obtain running times and MSE of the estimates from implicit SGD and $\text{glm}()$. Finally, we then run a regression of computing time and MSE against the problem size (N, p) .

The results are shown in Table 2.1. We observe that implicit SGD scales better in both sample size N , and especially in the model size p . We also observe, that this significant computational gain does not come with much efficiency loss. In fact, averaged over all samples, the MSE of the implicit SGD is on average 10% higher than the MSE of $\text{glm}()$ function with a standard error of ± 0.005 . Furthermore, the memory requirements (not reported in Table 2.1) are roughly $O(Np^2)$ for $\text{glm}()$ and only $O(p)$ for implicit SGD.

Table 2.1: Parameters from regressing computation time and MSE against (N, p) in log-scale for $\text{glm}()$ and implicit GLM. Computation time for $\text{glm}()$ is roughly $O(p^{1.47}N)$ and for implicit SGD, it is $O(p^{0.2}N^{0.9})$. Implicit SGD scales better in parameter dimension p , whereas MSE for both methods are comparable, at the order of $O(\sqrt{p/N})$.

METHOD	TIME(SEC)		MSE	
	log p (SE)	log N (SE)	log p (SE)	log N (SE)
GLM () FUNCTION	1.46 (0.019)	1.03 (0.02)	0.52 (0.007)	-0.52 (0.006)
IMPLICIT SGD	0.19 (0.012)	0.9 (0.01)	0.58 (0.007)	-0.53 (0.006)

Experiments with `biglm`

The package `biglm` is a popular choice to fitting GLMs in large data sets (large N , small p), and is part of the High-Performance Computing (HPC) task view of the CRAN project.¹¹ It works in an iterative way by splitting the data set in many parts, and by updating the model parameters using incremental QR decomposition (Miller, 1992), which results in only $O(p^2)$ memory requirement. In this experiment, we compare implicit SGD with `biglm` on larger data sets of Section 2.6.2. In particular, we focus on a few cases with small p and with large N such that $N \cdot p$ remains roughly constant.

The results are shown in Table 2.2. We observe that implicit SGD is significantly faster at a very small efficiency loss. The difference is more dramatic at large p ; for example, when $p = 10^3$ or $p = 10^4$, `biglm` quickly runs out of memory, whereas implicit SGD works without problems.

Experiments with `glmnet`

The `glmnet` package in R (Friedman *et al.*, 2010) is a deterministic optimization algorithm for generalized linear models that uses the elastic net. It performs a component-wise update of the parameter vector, utilizing thresholding from the regularization penalties for more computationally efficient updates. One update over all parameters costs roughly $O(Np)$

¹¹See <http://cran.r-project.org/web/packages/biglm/index.html> for the `biglm` package. The HPC view of the CRAN project is here <http://cran.r-project.org/web/views/HighPerformanceComputing.html>.

Table 2.2: Comparison of implicit SGD with *biglm*. MSE is defined as $\|\theta_N - \theta_*\|/\|\theta_0 - \theta_*\|$. Values “*” indicate out-of-memory errors. *biglm* was run in combination with the *ffdf* package to map big data files to memory. Implicit SGD used a similar but slower ad-hoc method. The table reports computation times excluding file access.

p	N	SIZE (GB)	PROCEDURE			
			BIGLM		IMPLICIT SGD	
			TIME(SECS)	MSE	TIME(SECS)	MSE
1E2	1E5	0.021	2.32	0.028	2.4	0.028
1E2	5E5	0.103	8.32	0.012	7.1	0.012
1E2	1E6	0.206	16	0.008	14.7	0.009
1E2	1E7	2.1	232	0.002	127.9	0.002
1E2	1E8	20.6	*	*	1397	0.00
1E3	1E6	2.0	*	*	31.38	0.153
1E4	1E5	2.0	*	*	25.05	0.160

operations. Additional computational gains are achieved when the design matrix is sparse because fewer components are updated per each iteration.

In this experiment, we compare implicit SGD with `glmnet` on a subset of experiments in the original package release (Friedman *et al.*, 2010). In particular, we implement the experiment of subsection 5.1 in that paper as follows. First, we sample the design matrix $X \sim \mathcal{N}_p(0, \Sigma)$, where $\Sigma = b^2U + I$ and U is the $p \times p$ matrix of ones. The parameter $b = \sqrt{\rho/(1-\rho)}$, where ρ is the target correlation of columns of X , is controlled in the experiments. The outcomes are $Y = X\theta_* + \sigma\epsilon$, where $\theta_j^* = (-1)^j \exp(-2(j-1)/20)$, and ϵ is a standard p -variate normal. The parameter σ is tuned to achieve a pre-defined signal-noise ratio. We report average computation times in Table B.2 over 10 replications, which expands Table 1 of Friedman *et al.* (2010).

First, we observe that implicit SGD is consistently faster than the `glmnet` method. In particular, the SGD method scales better at larger p following a sublinear growth as noted in Section 2.6.2. Interestingly, it is also not affected by covariate correlation, whereas `glmnet` gets slower as more components need to be updated at every iteration. For example, with correlation $\rho = 0.9$ and $N = 1e5$, $p = 200$, the SGD method is almost 10x faster.

Table 2.3: Comparing implicit SGD with `glmnet`. Table reports running times (in secs.) and MSE for both procedures. The MSE of `glmnet` is calculated as the median MSE over the 100 grid values of regularization parameter computed by default (Friedman et al., 2010).

METHOD	METRIC	CORRELATION (ρ)			
		0	0.2	0.6	0.9
<hr/> N = 1000, p = 10 <hr/>					
GLMNET	TIME(SEC)	0.005	0.005	0.008	0.022
	MSE	0.083	0.085	0.099	0.163
SGD	TIME(SEC)	0.011	0.011	0.011	0.011
	MSE	0.042	0.042	0.049	0.053
<hr/> N = 5000, p = 50 <hr/>					
GLMNET		0.058	0.067	0.119	0.273
		0.044	0.046	0.057	0.09
SGD		0.059	0.056	0.057	0.057
		0.019	0.02	0.023	0.031
<hr/> N = 100000, p = 200 <hr/>					
GLMNET		2.775	3.017	4.009	10.827
		0.017	0.017	0.021	0.033
SGD		1.475	1.464	1.474	1.446
		0.004	0.004	0.004	0.006

Second, to compare `glmnet` with implicit SGD in terms of MSE, we picked the median MSE produced by the grid of regularization parameters computed by `glmnet`. Because of regularization, it is reasonable to expect overall a slightly better performance for `glmnet` in situations where N is small compared to p . However, implicit SGD seems to perform better against the median performance of `glmnet`. This is reasonable because implicit SGD performs indirect regularization, as discussed in Section 2.7. Furthermore, Table B.2 indicates a clear trend where, for bigger dimensions p and higher correlation ρ , implicit SGD is performing better than `glmnet` in terms of efficiency as well. We obtain similar results in a comparison on a logistic regression model, presented in Appendix B.7.3.

Cox proportional hazards

In this experiment we test the performance of implicit SGD on estimating the parameters of a Cox proportional hazards model in a setup that is similar to the numerical example of [Simon *et al.* \(2011, Section 3\)](#).

We consider $N = 1000$ units with covariates $X \sim \mathcal{N}_N(0, \Sigma)$, where $\Sigma = 0.2U + I$, and U is the matrix of ones. We sample times as $Y_i \sim \text{Expo}(\eta_i(\theta_*))$, where $\eta_i(\theta) = \exp(X_i^\top \theta)$, and $\theta_* = (\theta_{*,k})$ is a vector with $p = 20$ elements defined as $\theta_{*,k} = 2(-1)^{-k} \exp(-0.1k)$. Time Y_i is censored, and thus $d_i = 0$, according to probability $(1 + \exp(-a(Y_i - b)))^{-1}$, where b is a quantile of choice (set here as $b = 0.8$), and a is set such that $\min\{Y_i\}$ is censored with a prespecified probability (set here as 0.1%). We replicate 50 times the following process. First, we run implicit SGD for $2N$ iterations, and then measure $\text{MSE} \|\theta_n^{\text{im}} - \theta_*\|^2$, for all $n = 1, 2, \dots, 2N$. To set the learning rates we use Eq. (2.9), where the Fisher matrix is diagonally approximated, through the AdaGrad procedure (2.11). We then take the 5%, 50% and 95% quantiles of MSE across all repetitions and plot them against iteration number n .

The results are shown in Figure 2.3. In the figure we also plot (horizontal dashed lines) the 5% and 95% quantiles of the MSE of the MLE, assumed to be the best MSE achievable for SGD. We observe that Implicit SGD performs well compared to MLE, even in this small-sized problem. In particular, implicit SGD, under the aforementioned generic tuning of learning rates, converges to the region of optimal MLE in a few thousands of iterations. In experiments with explicit SGD we were not able to replicate this performance because of numerical instability. To our best knowledge, there are no standard implementations of explicit SGD for estimating Cox proportional hazards models.

M-estimation

In this experiment we test the performance of implicit SGD, in particular Algorithm 5, on a M-estimation problem in a setup that is similar to the simulation example of [Donoho and Montanari \(2013, Example 2.4\)](#).




Figure 2.3: *Implicit SGD on Cox proportional hazards model. The different curves correspond to the 5%, 50%, and 95% quantiles of the MSE for 50 replications of implicit SGD iterates θ_{2000}^m . The top and bottom dashed horizontal lines correspond, respectively, to the 95% quantile and the 5% quantile over the same 50 replications of the best possible MSE.*

We set $N = 1000$ data points and $p = 200$ as the parameter dimension. We sample θ_* as a random vector with norm $\|\theta_*\| = 6\sqrt{p}$, and sample the design matrix as $X \sim \mathcal{N}_p(0, (1/N)I)$. The outcomes are sampled i.i.d. from a “contaminated” normal distribution, i.e., with probability 95%, $Y_n \sim \mathcal{N}(X_n^T \theta_*, 1)$, and $Y_n = 10$ with probability 5%.

The results over 2000 iterations of implicit SGD are shown in Figure 2.4. In the figure we plot the 5% and 95% quantiles of MSE of implicit SGD over 100 replications of the experiment. We also plot (horizontal dashed line) the median of the MSE of the MLE estimator, computed using the `coxph` built-in command of R.¹² We observe that SGD converges steadily to the best possible MSE. Similar behavior was observed under various

¹²Documentation is available at <https://stat.ethz.ch/R-manual/R-devel/library/survival/html/coxph.html>.

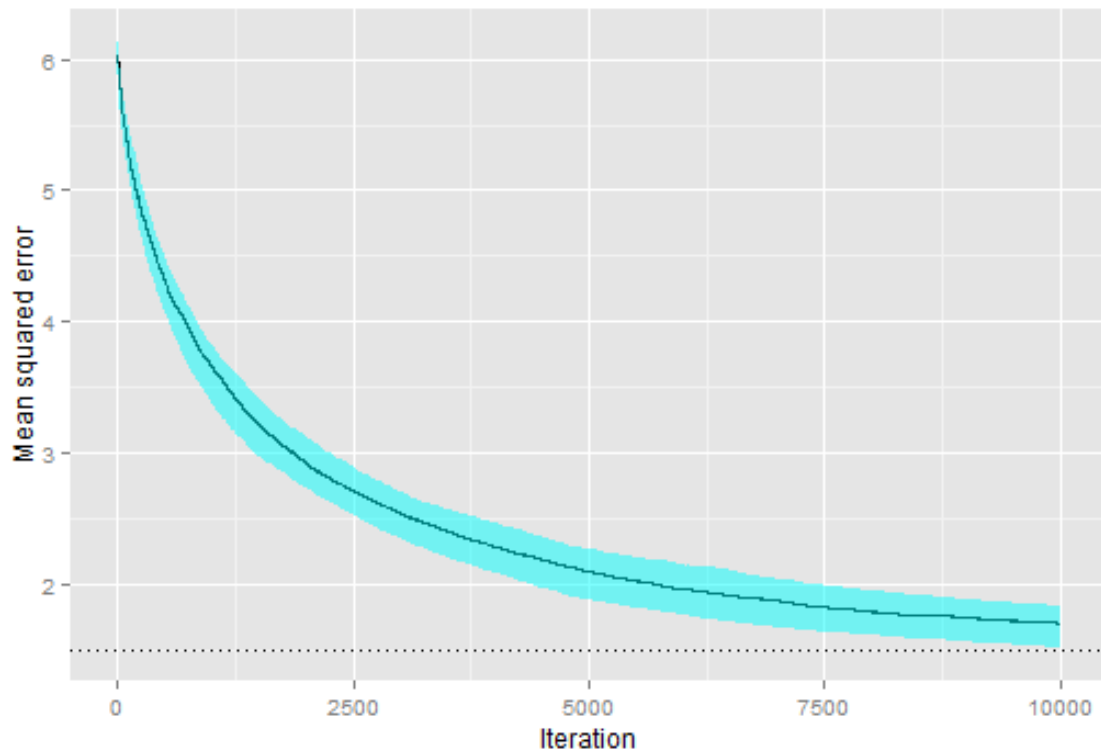


Figure 2.4: *Implicit SGD on M-estimation task. Top curve corresponds to 95% percentile of MSE, middle curve is median, and bottom curve is 5% of MSE over 100 replications of implicit SGD iterates $\theta_{2000}^{\text{im}}$. The dashed horizontal line corresponds to the median MSE of maximum-likelihood over the same 100 replications.*

modifications of the simulation parameters.

National Morbidity-Mortality Air Pollution study

The National Morbidity and Mortality Air Pollution (NMMAPS) study (Samet *et al.*, 2000; Dominici *et al.*, 2002) analyzed the risks of air pollution to public health. Several cities (108 in the US) are included in the study with daily measurements covering more than 13 years (roughly 5,000 days) including air pollution data (e.g. concentration of CO in the atmosphere) together with health outcome variables such as number of respiratory-related deaths.

The original study fitted a Poisson generalized additive model (GAM), separately for each city due to data set size. Recent recent research (Wood *et al.*, 2014) has developed procedures similar to `biglm`'s iterative QR decomposition to fit all cities simultaneously on the full data set with approximately $N = 1.2$ million observations and $p = 802$ covariates (7 Gb in size). In this experiment, we construct a GAM model using data from all cities in the NMMAPS study in a process that is very similar (but not identical) to the data set of Wood *et al.* (2014).

Our final data set has $N = 1,426,806$ observations and $p = 794$ covariates including all cities in the NMMAPS study (8.6GB in size), and is fit using the simplest first-order implicit SGD procedure, i.e., $C_n = I$ and $\gamma_1 = 1$. The time to fit the entire model with implicit SGD was roughly 123.4 seconds, which is almost 6x faster than the time reported by Wood *et al.* (2014) of about 12 minutes on a similar home computer. We cannot directly compare the estimates from the two procedures because different versions of the data sets were used. However, we can compare the estimates of our model with the estimates of `glm()` on a random small subset of the data. In particular, we subsampled $N = 50,000$ observations and $p = 50$ covariates (19.5MB in size) and fit the smaller data set using implicit SGD and `glm()`. A Q-Q plot of the estimates is shown in Figure 2.5. We observe that the estimates of the SGD procedure are very close to MLE. Further replications of the aforementioned testing process revealed the same pattern, indicating that implicit SGD converged on all replications.

2.7 Discussion

Our theory in Section 2.4 suggests that implicit SGD is numerically stable and has known asymptotic variance and asymptotic distribution. Our experiments in Section 2.6 showed that the empirical properties of SGD are well predicted by theory. In contrast, explicit SGD is unstable and cannot work well without problem-specific tuning. Thus, implicit SGD is a principled estimation procedure that is superior to classical explicit SGD.

Intuitively, implicit SGD leverages second-order information at every iteration, although




Figure 2.5: Estimates of implicit SGD (y-axis) and $glm(\cdot)$ (x-axis) on a subset of the NMMAPS data set with $N = 50,000$ observations and $p = 50$ covariates which is, roughly, 5% of the entire data set.

this quantity is not explicitly computed in Eq. (2.4). To see this, assume both explicit and implicit SGD are at the same estimate θ_0 . Then, using definitions (2.3) and (2.4), a Taylor approximation of the gradient $\nabla \log f(Y_n; X_n, \theta_n^{\text{im}})$ yields

$$\Delta \theta_n^{\text{im}} \approx [I + \gamma_n \hat{\mathcal{I}}(\theta_0; X_n, Y_n)]^{-1} \Delta \theta_n^{\text{sgd}}, \quad (2.27)$$

where $\Delta \theta_n^{\text{im}} = \theta_n^{\text{im}} - \theta_0$ and $\Delta \theta_n^{\text{sgd}} = \theta_n^{\text{sgd}} - \theta_0$, and the matrix $\hat{\mathcal{I}}(\theta_0; X_n, Y_n) = -\nabla^2 \log f(Y_n; X_n, \theta)|_{\theta=\theta_0}$ is the *observed* Fisher information at θ_0 . In other words, the implicit procedure is a *shrunked* version of the explicit one, where the shrinkage factor depends on the observed information.

Naturally, the implicit SGD iterate θ_n^{im} has also a Bayesian interpretation. In particular, θ_n^{im} is the posterior mode of a Bayesian model defined as

$$\begin{aligned}\theta|\theta_{n-1}^{\text{im}} &\sim \mathcal{N}(\theta_{n-1}^{\text{im}}, \gamma_n C_n) \\ Y_n|X_n, \theta &\sim f(\cdot; X_n, \theta).\end{aligned}\tag{2.28}$$

The explicit SGD update θ_n^{sgd} can be written as in Eq. (2.28), however f needs to be substituted with its linear approximation around $\theta_{n-1}^{\text{sgd}}$. Thus, Eq. (2.28) provides an alternative explanation why implicit SGD is more principled than explicit SGD. Furthermore, it indicates possible improvements for implicit SGD. For example, the prior in Eq. (2.28) should be chosen to fit better the parameter space (e.g., θ_* being on the simplex). [Krakowski *et al.* \(2007\)](#) and [Nemirovski *et al.* \(2009\)](#) have argued that appropriate implicit updates can fit better in the geometry of the parameter space, and thus converge faster. Setting up the parameters of the prior is also crucial. Whereas in explicit SGD there is no statistical intuition behind learning rates γ_n , Eq. (2.28) reveals that in implicit SGD the terms $(\gamma_n C_n)^{-1}$ encode the statistical information up to iteration n . It follows immediately that it is optimal, in general, to set $\gamma_n C_n = \mathcal{I}(\theta_*)^{-1}/n$, which is a special case of Theorem 6. Adaptive stochastic approximation methods in recursive estimation, typically being second-order as in Section 2.4.2, try to iteratively estimate the Fisher matrix $\mathcal{I}(\theta_*)$ along with estimation of θ_* .

The Bayesian formulation of Eq. (2.28) also explains the stability of implicit SGD. In Theorem 5 we showed that the initial conditions are discounted at an exponential rate, regardless of misspecification of the learning rates. This stability of implicit SGD allows several ideas for improvements. For example, constant learning rates could be used in implicit SGD to speed up convergence towards a region around θ_* . A sequential hypothesis test could decide on whether θ_n^{im} has reached that region or not, and switch to the theoretically optimal $1/n$ rate accordingly. Alternatively, we could run implicit SGD with AdaGrad learning rates (see Eq. (2.11)) and switch to $1/n$ rates when the theoretical $O(1/\sqrt{n})$ variance of AdaGrad becomes larger than the $O(1/n)$ variance of implicit SGD. Such schemes using constant rates with explicit SGD are very hard to do in practice because of instability.

Regarding statistical efficiency, a key technical result in this paper is that the asymptotic variance of implicit SGD can be quantified exactly using Theorem 6. Optimal learning rates were suggested in Eq. (2.9) that depend on the eigenvalues of the unknown Fisher matrix $\mathcal{I}(\theta_*)$. In this paper, we used second-order procedures of Section 2.4.2 to iteratively estimate the eigenvalues, however better methods are certainly possible and could improve the performance of implicit SGD. For example, it is known that typical iterative methods (as in Section 2.4.2) usually overestimate the largest eigenvalue and underestimate the smallest eigenvalue, in small-to-moderate samples. This crucially affects the behavior of stochastic approximations with learning rates that depend on sample eigenvalues. Empirical Bayes methods have been shown to be superior in iterative estimation of eigenvalues of large matrices and it would be interesting to apply such methods to design the learning rates of implicit SGD procedures (Haff, 1980; Dey, 1988; Ahmed, 1998; Mestre, 2008).

Regarding computational efficiency, we developed Algorithm 1 which implements implicit SGD on a large family of statistical models. However, the trick used in fitting the Cox proportional hazards model in Section 2.5.3 can be more generally applied to models outside this family. For example, assume a log-likelihood gradient of the form $s(X^\top\theta; Y)G(\theta; X, Y)$, where both its scale $s(\cdot)$ and direction $G(\cdot)$ depend on model parameters θ ; This violates conditions of Assumption 7(b). The implicit update (2.4) – where we set $C_n = I$, for simplicity – would be $\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n s(X_n^\top \theta_n^{\text{im}}; Y_n) G(\theta_n^{\text{im}}; X_n, Y_n)$, which cannot be computed by Algorithm 1. One way to circumvent this problem is to use an implicit update only on the scale and use an explicit update on the direction, i.e., $\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n s(X_n^\top \theta_n^{\text{im}}; Y_n) G(\theta_{n-1}^{\text{im}}; X_n, Y_n)$. This form of updates expands the applicability of implicit SGD.

Finally, hypothesis testing and construction of confidence intervals using SGD estimates is an important issue that has hitherto remained unexplored. In experiments of Section 2.6.1 we showed that implicit SGD is indeed asymptotically normal in several reasonable simulation scenarios. However, as SGD procedures are iterative, there needs to be a rigorous and general method to decide whether SGD iterates have converged to the asymptotic regime. Several methods, such as bootstrapping the data set, could be used for that.

Furthermore, conservative confidence intervals could be constructed through multivariate Chebyshev inequalities or other strategies (Marshall and Olkin, 1960).

2.8 Conclusion

In this paper, we introduced a new stochastic gradient descent procedure that uses implicit updates (i.e., solves fixed-point equations) at every iteration, which we termed implicit SGD. Equation (2.27) shows, intuitively, that the iterates of implicit SGD are a shrunked version of the standard iterates, where the shrinkage factor depends on the observed Fisher information matrix. Thus, implicit SGD combines the computational efficiency of first-order methods with the numerical stability of second-order methods.

In a theoretical analysis of implicit SGD, we derived non-asymptotic upper bounds for the mean-squared errors of its iterates, and their asymptotic variance and normality. Our analysis suggests principled strategies to calibrate a hyper-parameter that is common to both explicit and implicit SGD procedures, known as the learning rate. We illustrated the use of implicit SGD for statistical estimation in generalized linear models, Cox proportional hazards model, and general M-estimation problems. Implicit SGD can also be applied on several other models, such as support vector machines or generalized additive models, with minimal modifications.

Viewed as statistical estimation procedures, our results suggest that implicit SGD has the same asymptotic efficiency to explicit SGD. However, the implicit procedure is significantly more stable than the explicit one with respect to misspecification of the learning rate. In general, explicit SGD procedures are sensitive to outliers and to misspecification of the learning rates, making it impossible to apply without problem-specific tuning. In theory and in extensive experiments, implicit procedures emerge as principled iterative estimation methods because they are numerically stable, they are robust to tuning of hyper-parameters, and their standard errors are well-predicted by theory. Thus, implicit stochastic gradient descent is poised to become a workhorse of estimation with large data sets in statistical practice.

References

- AHMED, S. (1998). Large-sample estimation strategies for eigenvalues of a wishart matrix. *Metrika*, **47** (1), 35–45.
- AMARI, S.-I. (1998). Natural gradient works efficiently in learning. *Neural computation*, **10** (2), 251–276.
- , PARK, H. and FUKUMIZU, K. (2000). Adaptive method of realizing natural gradient learning for multilayer perceptrons. *Neural Computation*, **12** (6), 1399–1409.
- BACH, F. and MOULINES, E. (2013). Non-strongly-convex smooth stochastic approximation with convergence rate $o(1/n)$. In *Advances in Neural Information Processing Systems*, pp. 773–781.
- BALAKRISHNAN, S., WAINWRIGHT, M. J. and YU, B. (2014). Statistical guarantees for the em algorithm: From population to sample-based analysis. *arXiv preprint arXiv:1408.2156*.
- BARTZ, K., LIU, J. and BLITZSTEIN, J. (2009). Monte carlo maximum likelihood for exponential random graph models: From snowballs to umbrella densities.
- BATHER, J. (1989). *Stochastic approximation: A generalisation of the Robbins-Monro procedure*, vol. 89. Mathematical Sciences Institute, Cornell University.
- BAUSCHKE, H. H. and COMBETTES, P. L. (2011). *Convex analysis and monotone operator theory in Hilbert spaces*. Springer Science & Business Media.
- BECK, A. and TEBoulLE, M. (2003). Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, **31** (3), 167–175.
- and — (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, **2** (1), 183–202.
- BENVENISTE, A., PRIOURET, P. and MÉTIVIER, M. (1990). Adaptive algorithms and stochastic approximations.
- BERTSEKAS, D. P. (2011). Incremental proximal methods for large scale convex optimization. *Mathematical programming*, **129** (2), 163–195.
- BICKEL, P. J. and DOKSUM, K. A. (2015). *Mathematical Statistics: Basic Ideas and Selected Topics, volume I*, vol. 117. CRC Press.

- BLACKARD, J. (1998). *Comparison of neural networks and discriminant analysis in predicting forest cover types*. Ph.D. thesis, Department of Forest Sciences, Colorado State University.
- BORDES, A., BOTTOU, L. and GALLINARI, P. (2009). Sgd-qn: Careful quasi-newton stochastic gradient descent. *The Journal of Machine Learning Research*, **10**, 1737–1754.
- BORKAR, V. S. (2008). *Stochastic approximation*. Cambridge Books.
- BOTTOU, L. (2010). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, Springer, pp. 177–186.
- (2012). Stochastic Gradient Descent Tricks. In *Neural Networks: Tricks of the Trade*, vol. 1, pp. 421–436.
- BYRD, R. H., HANSEN, S., NOCEDAL, J. and SINGER, Y. (2014). A stochastic quasi-newton method for large-scale optimization. *arXiv preprint arXiv:1401.7020*.
- CAPPÉ, O. and MOULINES, E. (2009). On-line expectation–maximization algorithm for latent data models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **71** (3), 593–613.
- CHEN, H.-F., GUO, L. and GAO, A.-J. (1987). Convergence and robustness of the robbins-monro algorithm truncated at randomly varying bounds. *Stochastic Processes and their Applications*, **27**, 217–231.
- CHUNG, K. L. (1954). On a stochastic approximation method. *The Annals of Mathematical Statistics*, pp. 463–483.
- CORALUPPI, G. and YOUNG, T. Y. (1969). Stochastic signal representation. *Circuit Theory, IEEE Transactions on*, **16** (2), 155–161.
- COX, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 187–220.
- DAVISON, A. C. (2003). *Statistical models*, vol. 11. Cambridge University Press.
- DEAN, J., CORRADO, G., MONGA, R., CHEN, K., DEVIN, M., MAO, M., SENIOR, A., TUCKER, P., YANG, K., LE, Q. V. *et al.* (2012). Large scale distributed deep networks. In *Advances in Neural Information Processing Systems*, pp. 1223–1231.
- DEMPSTER, A., LAIRD, N. and RUBIN, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, **39**, 1–38.
- DEY, D. K. (1988). Simultaneous estimation of eigenvalues. *Annals of the Institute of Statistical Mathematics*, **40** (1), 137–147.
- DOBSON, A. J. and BARNETT, A. (2008). *An introduction to generalized linear models*. CRC press.
- DOMINICI, F., DANIELS, M., ZEGER, S. L. and SAMET, J. M. (2002). Air pollution and mortality: estimating regional and national dose-response relationships. *Journal of the American Statistical Association*, **97** (457), 100–111.

- DONOHO, D. and MONTANARI, A. (2013). High dimensional robust m-estimation: Asymptotic variance via approximate message passing. *arXiv preprint arXiv:1310.7320*.
- DUCHI, J., HAZAN, E. and SINGER, Y. (2011a). Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, **999999**, 2121–2159.
- , — and — (2011b). Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, **12**, 2121–2159.
- ERHAN, D., BENGIO, Y., COURVILLE, A., MANZAGOL, P.-A., VINCENT, P. and BENGIO, S. (2010). Why does unsupervised pre-training help deep learning? *The Journal of Machine Learning Research*, **11**, 625–660.
- FABIAN, V. (1968a). On asymptotic normality in stochastic approximation. *The Annals of Mathematical Statistics*, pp. 1327–1332.
- (1968b). On asymptotic normality in stochastic approximation. *The Annals of Mathematical Statistics*, pp. 1327–1332.
- (1978). On asymptotically efficient recursive estimation. *The Annals of Statistics*, pp. 854–866.
- FISHER, R. A. (1922). On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, **222**, 309–368.
- (1925). *Statistical Methods for Research Workers*. Edinburgh: Oliver and Boyd.
- FRANK, O. and STRAUSS, D. (1986). Markov graphs. *Journal of the American Statistical Association*, **81** (395), 832–842.
- FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, **33** (1), 1.
- GELMAN, A. and MENG, X.-L. (1998). Simulating normalizing constants: From importance sampling to bridge sampling to path sampling. *Statistical science*, pp. 163–185.
- GEORGE, A. P. and POWELL, W. B. (2006). Adaptive stepsizes for recursive estimation with applications in approximate dynamic programming. *Machine learning*, **65** (1), 167–198.
- GEYER, C. J. (1991). Markov chain monte carlo maximum likelihood.
- GLADYSHEV, E. (1965). On stochastic approximation. *Theory of Probability & Its Applications*, **10** (2), 275–278.
- GREEN, P. J. (1984). Iteratively reweighted least squares for maximum likelihood estimation, and some robust and resistant alternatives. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 149–192.
- HAFE, L. (1980). Empirical bayes estimation of the multivariate normal covariance matrix. *The Annals of Statistics*, pp. 586–597.

- HALL, A. R. (2005). *Generalized method of moments*. Oxford University Press Oxford.
- HASTIE, T., TIBSHIRANI, R. and FRIEDMAN, J. (2011). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edn.
- HASTIE, T. J. and TIBSHIRANI, R. J. (1990). *Generalized additive models*, vol. 43. CRC Press.
- HENNIG, P. and KIEFEL, M. (2013). Quasi-newton methods: A new direction. *The Journal of Machine Learning Research*, **14** (1), 843–865.
- HOFFMAN, J. D. and FRANKEL, S. (2001). *Numerical methods for engineers and scientists*. CRC press.
- HUBER, P. J. (2011). *Robust statistics*. Springer.
- *et al.* (1964). Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, **35** (1), 73–101.
- JAIN, P., TEWARI, A. and KAR, P. (2014). On iterative hard thresholding methods for high-dimensional m-estimation. In *Advances in Neural Information Processing Systems*, pp. 685–693.
- KARAMPATZIAKIS, N. and LANGFORD, J. (2010). Online importance weight aware updates. *arXiv preprint arXiv:1011.1576*.
- KAROUI, N. E. (2008). Spectrum estimation for large dimensional covariance matrices using random matrix theory. *The Annals of Statistics*, pp. 2757–2790.
- KELLEY, C. T. (1999). *Iterative methods for optimization*, vol. 18. Siam.
- KIVINEN, J., WARMUTH, M. K. and HASSIBI, B. (2006). The p-norm generalization of the lms algorithm for adaptive filtering. *Signal Processing, IEEE Transactions on*, **54** (5), 1782–1793.
- KLEIN, J. P. and MOESCHBERGER, M. L. (2003). *Survival analysis: techniques for censored and truncated data*. Springer Science & Business Media.
- KRAKOWSKI, K. A., MAHONY, R. E., WILLIAMSON, R. C. and WARMUTH, M. K. (2007). A geometric view of non-linear on-line stochastic gradient descent. *Author website*.
- KULIS, B. and BARTLETT, P. L. (2010). Implicit online learning. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 575–582.
- KUSHNER, H. J. and YIN, G. (2003). *Stochastic approximation and recursive algorithms and applications*, vol. 35. Springer Science & Business Media.
- LAI, T. L. and ROBBINS, H. (1979). Adaptive design and stochastic approximation. *The annals of Statistics*, pp. 1196–1221.
- LANGE, K. (2010). *Numerical analysis for statisticians*. Springer.
- LE CUN, Y., BOTTOU, L., BENGIO, Y. and HAFFNER, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of IEEE*, **86** (11), 2278–2324.

- LEHMANN, E. L. and CASELLA, G. (1998). *Theory of point estimation*, vol. 31. Springer Science & Business Media.
- LEWIS, D., YANG, Y., ROSE, T. and LI, F. (2004). Rcv1: A new benchmark collection for text categorization research. *The Journal of Machine Learning Research*, **5**, 361–397.
- LIN, Y., LV, F., ZHU, S., YANG, M., COUR, T., YU, K., CAO, L. and HUANG, T. (2011). Large-scale image classification: fast feature extraction and svm training. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, IEEE, pp. 1689–1696.
- LIONS, P.-L. and MERCIER, B. (1979). Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, **16** (6), 964–979.
- LJUNG, L., PFLUG, G. and WALK, H. (1992). *Stochastic approximation and optimization of random systems*, vol. 17. Springer.
- LYAPUNOV, A. M. (1992). The general problem of the stability of motion. *International Journal of Control*, **55** (3), 531–534.
- MARSHALL, A. W. and OLKIN, I. (1960). Multivariate chebyshev inequalities. *The Annals of Mathematical Statistics*, pp. 1001–1014.
- MARTIN, R. D. and MASRELIEZ, C. J. (1975). Robust estimation via stochastic approximation. *Information Theory, IEEE Transactions on*, **21** (3), 263–271.
- MESTRE, X. (2008). Improved estimation of eigenvalues and eigenvectors of covariance matrices using their sample estimates. *Information Theory, IEEE Transactions on*, **54** (11), 5113–5129.
- MILLER, A. J. (1992). Algorithm as 274: Least squares routines to supplement those of gentleman. *Applied Statistics*, pp. 458–478.
- MOULINES, E. and BACH, F. R. (2011). Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, pp. 451–459.
- NAGUMO, J.-I. and NODA, A. (1967). A learning method for system identification. *Automatic Control, IEEE Transactions on*, **12** (3), 282–287.
- NATIONAL RESEARCH COUNCIL (2013). *Frontiers in Massive Data Analysis*. Washington, DC: The National Academies Press.
- NELDER, J. and WEDDERBURN, R. (1972). Generalized linear models. *Journal of the Royal Statistical Society. Series A (General)*, pp. 370–384.
- NEMIROVSKI, A., JUDITSKY, A., LAN, G. and SHAPIRO, A. (2009). Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, **19** (4), 1574–1609.
- NEMIROVSKI, D., YUDIN (1983). *Problem complexity and method efficiency in optimization*. Wiley (Chichester and New York).

- NESTEROV, Y. (2004). *Introductory lectures on convex optimization*, vol. 87. Springer Science & Business Media.
- NEVELSON, M. B. and KHASHMINSKI, R. Z. (1973). *Stochastic approximation and recursive estimation*, vol. 47. Amer Mathematical Society.
- PARIKH, N. and BOYD, S. (2013). Proximal algorithms. *Foundations and Trends in optimization*, **1** (3), 123–231.
- PATTISON, P. and WASSERMAN, S. (1999). Logit models and logistic regressions for social networks: Ii. multivariate relations. *British Journal of Mathematical and Statistical Psychology*, **52** (2), 169–194.
- POLJAK, B. and TSYPKIN, J. Z. (1980). Robust identification. *Automatica*, **16** (1), 53–63.
- POLYAK, B. T. and JUDITSKY, A. B. (1992a). Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, **30** (4), 838–855.
- and — (1992b). Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, **30** (4), 838–855.
- and TSYPKIN, Y. Z. (1979). Adaptive estimation algorithms: convergence, optimality, stability. *Avtomatika i Telemekhanika*, (3), 71–84.
- ROBBINS, H. and MONRO, S. (1951). A stochastic approximation method. *The Annals of Mathematical Statistics*, pp. 400–407.
- and SIEGMUND, D. (1985). A convergence theorem for non negative almost supermartingales and some applications. In *Herbert Robbins Selected Papers*, Springer, pp. 111–135.
- ROBINS, G., PATTISON, P., KALISH, Y. and LUSHER, D. (2007). An introduction to exponential random graph (p^*) models for social networks. *Social networks*, **29** (2), 173–191.
- ROCKAFELLAR, R. T. (1976). Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, **14** (5), 877–898.
- ROSASCO, L., VILLA, S. and VŮ, B. C. (2014). Convergence of stochastic proximal gradient algorithm. *arXiv preprint arXiv:1403.5074*.
- RUPPERT, D. (1988). *Efficient estimations from a slowly convergent Robbins-Monro process*. Tech. rep., Cornell University Operations Research and Industrial Engineering.
- SACKS, J. (1958). Asymptotic distribution of stochastic approximation procedures. *The Annals of Mathematical Statistics*, **29** (2), 373–405.
- SAKRISON, D. J. (1965). Efficient recursive estimation; application to estimating the parameters of a covariance function. *International Journal of Engineering Science*, **3** (4), 461–483.
- SAMET, J. M., ZEGER, S. L., DOMINICI, F., CURRIERO, F., COURSAK, I., DOCKERY, D. W., SCHWARTZ, J. and ZANOBETTI, A. (2000). The national morbidity, mortality, and air pollution study. *Part II: morbidity and mortality from air pollution in the United States Res Rep Health Eff Inst*, **94** (pt 2), 5–79.

- SCHMIDT, M., LE ROUX, N. and BACH, F. (2013). *Minimizing finite sums with the stochastic average gradient*. Tech. rep., HAL 00860051.
- SCHUURMANS, L. C. S. and CAELLI, S. (2007). Implicit online learning with kernels. *Advances in neural information processing systems*, **19**, 249.
- SHAMIR, O. and ZHANG, T. (2012). Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. *arXiv preprint arXiv:1212.1824*.
- SIMON, N., FRIEDMAN, J., HASTIE, T., TIBSHIRANI, R. *et al.* (2011). Regularization paths for cox’s proportional hazards model via coordinate descent. *Journal of statistical software*, **39** (5), 1–13.
- SINGER, Y. and DUCHI, J. C. (2009). Efficient learning using forward-backward splitting. In *Advances in Neural Information Processing Systems*, pp. 495–503.
- SLOCK, D. T. (1993). On the convergence behavior of the lms and the normalized lms algorithms. *Signal Processing, IEEE Transactions on*, **41** (9), 2811–2825.
- SNIJEDERS, T. A. (2002). Markov chain monte carlo estimation of exponential random graph models. *Journal of Social Structure*, **3** (2), 1–40.
- SONNENBURG, S., FRANC, V., YOM-TOV, E. and SEBAG, M. (2008). Pascal large scale learning challenge.
- TOULIS, P. and AIROLDI, E. (2015a). Scalable estimation strategies based on stochastic approximations: classical results and new insights. *Statistics and Computing*, **25** (4), 781–795.
- , — and RENNIE, J. (2014). Statistical analysis of stochastic gradient methods for generalized linear models. *JMLR W&CP*, **32** (1), 667–675.
- and AIROLDI, E. M. (2015b). Implicit stochastic gradient descent for principled estimation with large datasets. *arXiv preprint arXiv:1408.2923*.
- WELLING, M. and TEH, Y. W. (2011). Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 681–688.
- WIDROW, B. and HOFF, M. E. (1960). Adaptive switching circuits. *Defense Technical Information Center*.
- WOOD, S. N., GOUDE, Y. and SHAW, S. (2014). Generalized additive models for large data sets. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*.
- XIAO, L. and ZHANG, T. (2014). A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, **24**, 2057–2075.
- XU, W. (2011). Towards optimal one pass large scale learning with averaged stochastic gradient descent. *arXiv preprint arXiv:1107.2490*.

ZHANG, T. (2004). Solving large scale linear prediction problems using gradient descent algorithms. In *Proceedings of the twenty-first international conference on Machine learning*, ACM, p. 116.

Appendix A

Appendix to Chapter 1

A.1 Theory of implicit stochastic approximation

The symbol $\|\cdot\|$ denotes the L_1 vector/matrix norm. We also define the error random variables $\xi_n \triangleq W_{\theta_{n-1}^{(*)}} - h(\theta_{n-1}^{(*)})$, such that $\mathbb{E}(\xi_n | \mathcal{F}_{n-1}) = 0$. The parameter space for θ will be \mathbb{R}^p without loss of generality. For a positive scalar sequence a_n , the sequence $b_n = O(a_n)$ is such that $b_n \leq ca_n$, for some fixed $c > 0$, and every n ; the sequence $b_n = o(a_n)$ is such that $b_n/a_n \rightarrow 0$ in the limit where $n \rightarrow \infty$. $b_n \downarrow 0$ denotes a positive sequence decreasing towards zero. We further assume that implicit stochastic approximation (1.2) operates under a combination of the following assumptions.

Assumption 1. *It holds, $\gamma_n = \gamma_1 n^{-\gamma}$, $\gamma_1 > 0$ and $\gamma \in (1/2, 1]$.*

Assumption 2. *The regression function h is Lipschitz with parameter L , i.e., for all θ_1, θ_2 ,*

$$\|h(\theta_1) - h(\theta_2)\| \leq L\|\theta_1 - \theta_2\|.$$

Assumption 3. *For all n , function h satisfies either*

- (a) $(\theta - \theta_*)^\top h(\theta) > 0$, for all θ , or
- (b) $(\theta_n^{(*)} - \theta_*)^\top h(\theta_n^{(*)}) \geq \delta_n \|\theta_n^{(*)} - \theta_*\|^2$, where $\delta_n = \delta_1 n^{-\delta}$, $\delta_1 > 0$ and $0 < \gamma + \delta \leq 1$.

Assumption 4. *There exists a scalar potential $H : \mathbb{R}^p \rightarrow \mathbb{R}$ such that $\nabla H(\theta) = h(\theta)$, for all θ .*

Assumption 5. There exists fixed $\sigma^2 > 0$ such that, for every n ,

$$\mathbb{E} (|\xi_n|^2 | \mathcal{F}_{n-1}) \leq \sigma^2.$$

Assumption 6. Let $\Xi_n \stackrel{\text{def}}{=} \mathbb{E} (\xi_n \xi_n^\top | \mathcal{F}_{n-1})$, then $\|\Xi_n - \Xi\| = \mathcal{O}(1)$, and $\|\Xi_n - \Xi\| \rightarrow 0$ for fixed positive-definite matrix Ξ . Furthermore, if $\sigma_{n,s}^2 = \mathbb{E} \left(\mathbb{I}_{|\xi_n|^2 \geq s/\gamma_n} \|\xi_n\|^2 \right)$, then for all $s > 0$, $\sum_{i=1}^n \sigma_{i,s}^2 = \mathcal{O}(n)$ if $\gamma_n \propto n^{-1}$, or $\sigma_{n,s}^2 = \mathcal{O}(1)$ otherwise.

A.1.1 Convergence

Theorem 1. Suppose that Assumptions 1, 2, 3(a), and 5 hold. Then the iterates θ_n^{im} of the implicit stochastic approximation procedure (1.2) converge almost-surely to θ_* ; i.e., $\theta_n^{\text{im}} \rightarrow \theta_*$, such that $h(\theta_*) = 0$, almost-surely.

Proof. By definition (1.2) it follows

$$\|\theta_n^{\text{im}} - \theta_*\|^2 = \|\theta_{n-1}^{\text{im}} - \theta_*\|^2 - 2\gamma_n (\theta_{n-1}^{\text{im}} - \theta_*)^\top W_{\theta_n^{(*)}} + \gamma_n^2 \|W_{\theta_n^{(*)}}\|^2, \quad (\text{A.2})$$

where $(\theta_{n-1}^{\text{im}} - \theta_n^{(*)}) = \gamma_n h(\theta_n^{(*)})$ by Eq. (1.3). Setting $(\theta_{n-1}^{\text{im}} - \theta_*) = (\theta_n^{(*)} - \theta_*) + (\theta_{n-1}^{\text{im}} - \theta_n^{(*)})$, yields

$$\begin{aligned} R_n &\triangleq \mathbb{E} \left((\theta_{n-1}^{\text{im}} - \theta_*)^\top W_{\theta_n^{(*)}} \middle| \mathcal{F}_{n-1} \right) = (\theta_n^{(*)} - \theta_*)^\top h(\theta_n^{(*)}) + (\theta_{n-1}^{\text{im}} - \theta_n^{(*)})^\top h(\theta_n^{(*)}) \\ &= (\theta_n^{(*)} - \theta_*)^\top h(\theta_n^{(*)}) + \gamma_n \|h(\theta_n^{(*)})\|^2 > 0. \quad [\text{by Assumption 3(a)}] \end{aligned} \quad (\text{A.3})$$

Through Eq. (1.3) we obtain

$$\begin{aligned} \|\theta_{n-1}^{\text{im}} - \theta_*\|^2 &= \|\theta_n^{(*)} - \theta_*\|^2 + 2\gamma_n h(\theta_n^{(*)})^\top (\theta_n^{(*)} - \theta_*) + \gamma_n^2 \|h(\theta_n^{(*)})\|^2, \\ &> \|\theta_n^{(*)} - \theta_*\|^2. \quad [\text{by Assumption 3(a)}] \end{aligned} \quad (\text{A.4})$$

So, update (1.3) is non-expansive. Therefore,

$$\begin{aligned} \|h(\theta_n^{(*)})\| &= \|h(\theta_n^{(*)}) - h(\theta_*)\| \leq L \|\theta_n^{(*)} - \theta_*\| \quad [\text{by Assumption 2}] \\ &\leq L \|\theta_{n-1}^{\text{im}} - \theta_*\| \quad [\text{by non-expansiveness of update Eq. (1.3).}] \end{aligned} \quad (\text{A.5})$$

Furthermore,

$$\begin{aligned}
\mathbb{E} \left(\|W_{\theta_n^{(*)}}\|^2 \mid \mathcal{F}_{n-1} \right) &\stackrel{\text{def}}{=} \mathbb{E} \left(\|h(\theta_n^{(*)}) + \xi_n\|^2 \mid \mathcal{F}_{n-1} \right) \\
&= \|h(\theta_n^{(*)})\|^2 + \mathbb{E} \left(\|\xi_n\|^2 \mid \mathcal{F}_{n-1} \right) \\
&\leq L^2 \|\theta_{n-1}^{\text{im}} - \theta_\star\|^2 + \sigma^2. \quad [\text{by Eq. (A.5) and Assumption 5}] \quad (\text{A.6})
\end{aligned}$$

Taking expectations in Eq. (A.2) conditional on \mathcal{F}_{n-1} and using Eq. (A.3) and Ineq. (A.6), we obtain

$$\mathbb{E} \left(\|\theta_n^{\text{im}} - \theta_\star\|^2 \mid \mathcal{F}_{n-1} \right) \leq (1 + \gamma_n^2 L^2) \|\theta_{n-1}^{\text{im}} - \theta_\star\|^2 - 2\gamma_n R_n + \gamma_n^2 \sigma^2. \quad (\text{A.7})$$

We now use an argument—due to Gladyshev (1965)—that is also applicable to the classical Robbins-Monro procedure; see, for example, Benveniste *et al.* (1990, Section 5.2.2), or Ljung *et al.* (1992, Theorem 1.9). Random variable R_n is positive almost-surely by Ineq. (A.3), and $\sum \gamma_i = \infty$ and $\sum \gamma_i^2 < \infty$ by Assumption 1. Therefore, we can invoke the supermartingale lemma of Robbins and Siegmund (1985) to infer that $\|\theta_n^{\text{im}} - \theta_\star\|^2 \rightarrow B > 0$ and $\sum \gamma_n R_n < \infty$, almost-surely. If $B \neq 0$ then $\liminf \|\theta_n^{\text{im}} - \theta_\star\| > 0$, and thus the series $\sum_n \gamma_n R_n$ diverges by Ineq. (A.3) and $\sum \gamma_i = \infty$ (Assumption 1). Thus, $B = 0$. \square

A.2 Non-asymptotic analysis

Theorem 2. *Suppose that Assumptions 1, 2, 3(a), 4, and 5 hold. Define $\Gamma^2 \triangleq \mathbb{E} \left(\|\theta_0^{\text{im}} - \theta_\star\|^2 \right) + \sigma^2 \sum_{i=1}^{\infty} \gamma_i^2 + \gamma_1^2 \sigma^2$. Then, if $\gamma \in (2/3, 1]$, there exists $n_{0,1} < \infty$ such that, for all $n > n_{0,1}$,*

$$\mathbb{E} \left(H(\theta_n^{\text{im}}) - H(\theta_\star) \right) \leq \left[\frac{2\Gamma^2}{\gamma\gamma_1} + o(1) \right] n^{-1+\gamma}.$$

If $\gamma \in (1/2, 2/3)$, there exists $n_{0,2} < \infty$ such that, for all $n > n_{0,2}$,

$$\mathbb{E} \left(H(\theta_n^{\text{im}}) - H(\theta_\star) \right) \leq \left[\Gamma\sigma\sqrt{L\gamma_1} + o(1) \right] n^{-\gamma/2}.$$

Otherwise, $\gamma = 2/3$ and there exists $n_{0,3} < \infty$ such that, for all $n > n_{0,3}$,

$$\mathbb{E} (H(\theta_n^{\text{im}}) - H(\theta_\star)) \leq \left[\frac{3 + \sqrt{9 + 4\gamma_1^3 L\sigma^2/\Gamma^2}}{2\gamma_1/\Gamma^2} + \mathfrak{o}(1) \right] n^{-1/3}.$$

Proof. By Eq. (1.3) and Assumption 3(a), $\theta_n^{(*)} + \gamma_n h(\theta_n^{(*)}) = \theta_{n-1}^{\text{im}}$ is equivalent to minimization $\theta_n^{(*)} = \arg \min_{\theta} \{ \frac{1}{2\gamma_n} \|\theta - \theta_{n-1}^{\text{im}}\|^2 + H(\theta) \}$. Therefore, comparing the values of the expression for $\theta = \theta_n^{(*)}$ and $\theta = \theta_{n-1}^{\text{im}}$, we obtain

$$H(\theta_n^{(*)}) + \frac{1}{2\gamma_n} \|\theta_n^{(*)} - \theta_{n-1}^{\text{im}}\|^2 \leq H(\theta_{n-1}^{\text{im}}). \quad (\text{A.8})$$

Since $\theta_{n-1}^{\text{im}} - \theta_n^{(*)} = \gamma_n h(\theta_n^{(*)})$, Eq. (A.8) can be written as

$$H(\theta_{n-1}^{\text{im}}) - H(\theta_n^{(*)}) - \frac{1}{2} \gamma_n \|h(\theta_n^{(*)})\|^2 \geq 0. \quad (\text{A.9})$$

We also have

$$\begin{aligned} H(\theta_n^{(*)}) - H(\theta_\star) &\leq h(\theta_n^{(*)})^\top (\theta_n^{(*)} - \theta_\star) \quad [\text{by Assumption 3(a)}] \\ H(\theta_n^{(*)}) - H(\theta_\star) &\leq \|h(\theta_n^{(*)})\| \cdot \|\theta_n^{(*)} - \theta_\star\| \quad [\text{by Cauchy-Schwartz}] \\ [\mathbb{E} (H(\theta_n^{(*)}) - H(\theta_\star))]^2 &\leq \mathbb{E} (\|h(\theta_n^{(*)})\|^2) \mathbb{E} (\|\theta_n^{(*)} - \theta_\star\|^2) \quad [\text{by Jensen's inequality}]. \end{aligned} \quad (\text{A.10})$$

Furthermore,

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} - \gamma_n (h(\theta_n^{(*)}) + \xi_n) = \theta_n^{(*)} - \gamma_n \xi_n. \quad [\text{by Eq. (1.3)}] \quad (\text{A.11})$$

Therefore,

$$\begin{aligned} \mathbb{E} (\|\theta_n^{\text{im}} - \theta_\star\|^2) &= \mathbb{E} (\|\theta_n^{(*)} - \theta_\star\|^2) - 2\gamma_n \mathbb{E} ((\theta_n^{(*)} - \theta_\star)^\top \xi_n) + \gamma_n^2 \mathbb{E} (\|\xi_n\|^2) \\ &= \mathbb{E} (\|\theta_n^{(*)} - \theta_\star\|^2) + \gamma_n^2 \mathbb{E} (\|\xi_n\|^2) \\ &\leq \mathbb{E} (\|\theta_{n-1}^{\text{im}} - \theta_\star\|^2) + \gamma_n^2 \sigma^2. \quad [\text{by InEq. (A.4) and Assumption 5}] \\ &\leq \mathbb{E} (\|\theta_0^{\text{im}} - \theta_\star\|^2) + \sigma^2 \sum_{i=1}^n \gamma_i^2. \quad [\text{by induction.}] \end{aligned} \quad (\text{A.12})$$

For notational convenience, define $h_n \stackrel{\text{def}}{=} \mathbb{E} (H(\theta_n^{\text{im}}) - H(\theta_\star))$ and $h_n^{(*)} \stackrel{\text{def}}{=} \mathbb{E} (H(\theta_n^{(*)}) - H(\theta_\star))$. Our goal is to derive a bound for h_n . Now, by Eq. (A.11) $\theta_n^{(*)} = \theta_n^{\text{im}} + \gamma_n \xi_n$. Since $\mathbb{E}(\xi_n | \mathcal{F}_{n-1}) = 0$, it follows from Assumption 5, $\mathbb{E}(\|\theta_n^{(*)} - \theta_\star\|^2) \leq \mathbb{E}(\|\theta_n^{\text{im}} - \theta_\star\|^2) + \gamma_n^2 \sigma^2$. Hence, using Ineq. (A.12),

$$\mathbb{E}(\|\theta_n^{(*)} - \theta_\star\|^2) \leq \mathbb{E}(\|\theta_0^{\text{im}} - \theta_\star\|^2) + \sigma^2 \sum_{i=1}^{\infty} \gamma_i^2 + \gamma_n^2 \sigma^2 \leq \Gamma^2, \quad (\text{A.13})$$

by definition of Γ^2 . Furthermore, by convexity of H and Lipschitz continuity of h (Assumption 3(a)), and Assumption 5, we have

$$\begin{aligned} H(\theta_n^{\text{im}}) &= H(\theta_n^{(*)} - \gamma_n \xi_n) \\ H(\theta_n^{\text{im}}) &\leq H(\theta_n^{(*)}) - \gamma_n h(\theta_n^{(*)})^\top \xi_n + \gamma_n^2 \frac{L}{2} \|\xi_n\|^2 \quad [\text{by Lipschitz continuity}] \\ H(\theta_n^{\text{im}}) - H(\theta_\star) &\leq H(\theta_n^{(*)}) - H(\theta_\star) - \gamma_n h(\theta_n^{(*)})^\top \xi_n + \gamma_n^2 \frac{L}{2} \|\xi_n\|^2 \\ h_n &\leq h_n^{(*)} + \gamma_n^2 \frac{L\sigma^2}{2}. \quad [\text{by taking expectations.}] \end{aligned} \quad (\text{A.14})$$

From Ineq. (A.10) and Ineq. (A.13),

$$\mathbb{E}(\|h(\theta_n^{(*)})\|^2) \geq \frac{1}{\Gamma^2} [\mathbb{E}(H(\theta_n^{(*)}) - H(\theta_\star))]^2 \stackrel{\text{def}}{=} \frac{1}{\Gamma^2} h_n^{(*)2}. \quad (\text{A.15})$$

Now, in Ineq. (A.9), we subtract $H(\theta_\star)$ from the left-hand side, take expectations, and combine with (A.15) to obtain

$$h_{n-1} \geq h_n^{(*)} + \frac{1}{2\Gamma^2} \gamma_n h_n^{(*)2} \triangleq F_{\gamma_n}(h_n^{(*)}). \quad (\text{A.16})$$

The function $F_{\gamma_n}(x)$ is monotone increasing because $h_n \geq 0$ and $h_n^{(*)} \geq 0$, since $H(\theta_\star)$ is minimum. Let $F_{\gamma_n}^{-1}$ denote its inverse, which is also monotone increasing. Thus $h_n^{(*)} \leq F_{\gamma_n}^{-1}(h_{n-1})$. Using Eq. (A.16) we can rewrite (A.14) as

$$h_n \leq F_{\gamma_n}^{-1}(h_{n-1}) + \gamma_n^2 \frac{L\sigma^2}{2}. \quad (\text{A.17})$$

Ineq. (A.17) is our main recursion, since we want to upper-bound h_n . We will now try to find a base sequence $\{b_n\}$ such that $b_n \geq F_{\gamma_n}^{-1}(b_{n-1}) + \gamma_n^2 \frac{L\sigma^2}{2}$. Since one can take

b_n to be increasing arbitrarily, we will try to find the smallest possible sequence $\{b_n\}$. To make our analysis more tractable we will search in the family of sequences $b_n = b_1 n^{-\beta}$, for various values $b_1, \beta > 0$. Then, b_n will be an upper-bound for h_n . To see this inductively, assume that $h_{n-1} \leq b_{n-1}$ and that h_n satisfies (A.17). Then, $h_n \leq F_{\gamma_n}^{-1}(h_{n-1}) + \gamma_n^2 \frac{L\sigma^2}{2} \leq F_{\gamma_n}^{-1}(b_{n-1}) + \gamma_n^2 \frac{L\sigma^2}{2} \leq b_n$, where the first inequality follows from the monotonicity of F_{γ_n} , and the second inequality follows from definition of b_n .

Now, the condition for b_n can be rewritten as $b_{n-1} \leq F_{\gamma_n}(b_n - \gamma_n^2 \frac{L\sigma^2}{2})$, and by definition of F_{γ_n} we get

$$b_{n-1} \leq b_n - \gamma_n^2 \frac{L\sigma^2}{2} + \gamma_n \frac{1}{2\Gamma^2} (b_n - \gamma_n^2 \frac{L\sigma^2}{2})^2 \quad (\text{A.18})$$

Using $b_n = b_1 n^{-\beta}$ and $\gamma_n = \gamma_1 n^{-\gamma}$ (Assumption 1), we obtain

$$b_1 [(n-1)^{-\beta} - n^{-\beta}] + \frac{L\sigma^2 \gamma_1^2}{2} n^{-2\gamma} + \frac{L\sigma^2 \gamma_1^3 b_1}{2\Gamma^2} n^{-\beta-3\gamma} - \frac{\gamma_1 b_1^2}{2\Gamma^2} n^{-2\beta-\gamma} - \frac{L^2 \sigma^4 \gamma_1^5}{8\Gamma^2} n^{-5\gamma} \leq 0. \quad (\text{A.19})$$

We have $(n-1)^{-\beta} - n^{-\beta} < \frac{1}{1-\beta} n^{-1-\beta}$, for $n > 1$. Thus, it suffices to have

$$\frac{b_1}{1-\beta} n^{-1-\beta} + \frac{L\sigma^2 \gamma_1^2}{2} n^{-2\gamma} + \frac{L\sigma^2 \gamma_1^3 b_1}{2\Gamma^2} n^{-\beta-3\gamma} - \frac{\gamma_1 b_1^2}{2\Gamma^2} n^{-2\beta-\gamma} \leq 0, \quad (\text{A.20})$$

where we dropped the $n^{-5\gamma}$ term without loss of generality. The positive terms in Ineq. (A.20) are $n^{-1-\beta}$, $n^{-2\gamma}$, and $n^{-\beta-3\gamma}$, and the only negative term is of order $n^{-2\beta-\gamma}$. In order to find the largest possible β to satisfy (A.20), one needs to equate the term $n^{-2\beta-\gamma}$ with the slowest possible term with a positive coefficient, i.e., set $2\beta + \gamma = \min\{1 + \beta, \beta + 3\gamma, 2\gamma\}$. However, $\beta + 3\gamma > 1 + \beta$ and $\beta + 3\gamma > 2\gamma$, and thus $2\beta + \gamma = \min\{1 + \beta, 2\gamma\}$, which implies only three cases:

- (a) $1 + \beta < 2\gamma$, and thus $2\beta + \gamma = 1 + \beta$, which implies $\beta = 1 - \gamma$. Also, $1 + \beta < 2\gamma \Rightarrow 2 - \gamma < 2\gamma$, and thus $\gamma \in (2/3, 1]$. In this case, b_1 will satisfy (A.20) for all $n > n_{0,1}$, for some $n_{0,1}$, if

$$\frac{b_1}{1-\beta} < \frac{\gamma_1 b_1^2}{2\Gamma^2} \Leftrightarrow b_1 > \frac{2\Gamma^2}{\gamma\gamma_1}. \quad (\text{A.21})$$

(b) $2\gamma < 1 + \beta$, and thus $2\beta + \gamma = 2\gamma$, which implies $\beta = \gamma/2$. Also, $1 + \beta > 2\gamma \Rightarrow 1 + \gamma/2 > 2\gamma$, and thus $\gamma \in (1/2, 2/3)$. In this case, b_1 will satisfy (A.20) for all $n > n_{0,2}$, for some $n_{0,2}$, if

$$\frac{\gamma_1^2 L \sigma^2}{2} < \frac{\gamma_1 b_1^2}{2\Gamma^2} \Leftrightarrow b_1 > \Gamma \sigma \sqrt{L \gamma_1}. \quad (\text{A.22})$$

(c) $2\gamma = 1 + \beta$, and thus $2\gamma = 1 + \beta = 2\beta + \gamma$, which implies $\gamma = 2/3$ and $\beta = 1/3$. In this case,

$$\frac{b_1}{1 - \beta} + \frac{\gamma_1^2 L \sigma^2}{2} < \frac{\gamma_1 b_1^2}{2\Gamma^2}. \quad (\text{A.23})$$

Because all constants are positive in Ineq. (A.23), including b_1 , it follows that

$$b_1 > \frac{3 + \sqrt{9 + 4\gamma_1^3 L \sigma^2 / \Gamma^2}}{2\gamma_1 / \Gamma^2}. \quad (\text{A.24})$$

Remark. The constants $n_{0,1}, n_{0,2}, n_{0,3}$ depend on the problem parameters and the desired accuracy in the bounds of Theorem 2. Thus, it is straightforward to derive exact values for them. For example, consider case (a) and assume we picked b_1 such that $\frac{\gamma_1 b_1^2}{2\Gamma^2} - \frac{b_1}{1-\beta} = \epsilon > 0$. Ignoring the term $n^{-3\gamma-\beta}$ (for simplicity), Ineq. (A.20) becomes

$$\epsilon n^{-2+\gamma} \geq \frac{L \sigma^2 \gamma_1^2}{2} n^{-2\gamma} \Rightarrow n \geq \left(\frac{L \sigma^2 \gamma_1^2}{2\epsilon} \right)^c \equiv n_{0,1}, \quad (\text{A.25})$$

where $c = 1/(3\gamma - 2) > 0$ since $\gamma \in (2/3, 1]$. If the desired accuracy is small, then ϵ is large and so the value $n_{0,1}$ will become smaller. Similarly, we can derive expressions for $n_{0,2}$ and $n_{0,3}$. \square

Theorem 3. Suppose that Assumptions 1, 3(b), and 5 hold, and define $\zeta_n \triangleq \mathbb{E} (|\theta_n^{\text{im}} - \theta_*|^2)$ and $\kappa \triangleq 1 + 2\gamma_1 \delta_1$. Then, if $\gamma + \delta < 1$, for every $n > 0$ it holds,

$$\zeta_n \leq e^{-\log \kappa \cdot n^{1-\gamma-\delta}} \zeta_0 + \sigma^2 \frac{\gamma_1 \kappa}{\delta_1} n^{-\gamma+\delta} + \mathcal{O}(n^{-\gamma+\delta-1}).$$

Otherwise, if $\gamma = 1, \delta = 0$, it holds,

$$\zeta_n \leq e^{-\log \kappa \cdot \log n} \zeta_0 + \sigma^2 \frac{\gamma 1^K}{\delta_1} n^{-1} + \mathcal{O}(n^{-2}).$$

Proof. First we prove two lemmas that will be useful for Theorem 3.

Lemma 1. Consider a sequence b_n such that $b_n \downarrow 0$ and $\sum_{i=1}^{\infty} b_i = \infty$. Then, there exists a positive constant $K > 0$, such that

$$\prod_{i=1}^n \frac{1}{1+b_i} \leq \exp(-K \sum_{i=1}^n b_i). \quad (\text{A.26})$$

Proof. The function $x \log(1 + 1/x)$ is increasing-concave in $(0, \infty)$. From $b_n \downarrow 0$ it follows that $\log(1 + b_n)/b_n$ is non-increasing. Consider the value $K \stackrel{\text{def}}{=} \log(1 + b_1)/b_1$. Then, $(1 + b_n)^{-1} \leq \exp(-K b_n)$. Successive applications of this inequality yields Ineq. (A.26). The upper bound is more accurate when b_n takes lower values. \square

Lemma 2. Consider sequences $a_n \downarrow 0, b_n \downarrow 0$, and $c_n \downarrow 0$ such that, $a_n = o(b_n)$, $\sum_{i=1}^{\infty} a_i \stackrel{\text{def}}{=} A < \infty$, and there is n' such that $c_n/b_n < 1$ for all $n > n'$. Define,

$$\delta_n \stackrel{\text{def}}{=} \frac{1}{a_n} (a_{n-1}/b_{n-1} - a_n/b_n) \text{ and } \zeta_n \stackrel{\text{def}}{=} \frac{c_n}{b_{n-1}} \frac{a_{n-1}}{a_n}, \quad (\text{A.27})$$

and suppose that $\delta_n \downarrow 0$ and $\zeta_n \downarrow 0$. Pick a positive n_0 such that $\delta_n + \zeta_n < 1$ and $(1 + c_n)/(1 + b_n) < 1$, for all $n \geq n_0$.

Consider a positive sequence $y_n > 0$ that satisfies the recursive inequality,

$$y_n \leq \frac{1 + c_n}{1 + b_n} y_{n-1} + a_n. \quad (\text{A.28})$$

Then, for every $n > 0$,

$$y_n \leq K_0 \frac{a_n}{b_n} + Q_1^n y_0 + Q_{n_0+1}^n (1 + c_1)^{n_0} A, \quad (\text{A.29})$$

where $K_0 \stackrel{\text{def}}{=} (1 + b_1) (1 - \delta_{n_0} - \zeta_{n_0})^{-1}$, and $Q_i^n = \prod_{j=i}^n (1 + c_j)/(1 + b_j)$, such that $Q_i^n = 1$ if $n < i$, by definition.

Proof. We consider two separate cases, namely, $n < n_0$ and $n \geq n_0$, and then we will combine the respective bounds.

Analysis for $n < n_0$. We first find a crude bound for Q_{i+1}^n . It holds,

$$Q_{i+1}^n \leq (1 + c_{i+1})(1 + c_{i+2}) \cdots (1 + c_n) \leq (1 + c_1)^{n_0}, \quad (\text{A.30})$$

since $c_1 \geq c_n$ ($c_n \downarrow 0$ by definition) and there are no more than n_0 terms in the product. From Ineq. (A.28) we get

$$\begin{aligned} y_n &\leq Q_1^n y_0 + \sum_{i=1}^n Q_{i+1}^n a_i \quad [\text{by expanding recursive Ineq. (A.28)}] \\ &\leq Q_1^n y_0 + (1 + c_1)^{n_0} \sum_{i=1}^n a_i \quad [\text{using Ineq. (A.30)}] \\ &\leq Q_1^n y_0 + (1 + c_1)^{n_0} A. \end{aligned} \quad (\text{A.31})$$

This inequality holds also for $n = n_0$.

Analysis for $n \geq n_0$. In this case, we have for all $n \geq n_0$,

$$\begin{aligned} (1 + b_1)(1 - \delta_n - \zeta_n)^{-1} &\leq K_0 \quad [\text{by definition of } n_0, K_0] \\ K_0(\delta_n + \zeta_n) + 1 + b_1 &\leq K_0 \\ K_0(\delta_n + \zeta_n) + 1 + b_n &\leq K_0 \quad [\text{because } b_n \leq b_1, \text{ since } b_n \downarrow 0] \\ \frac{1}{a_n} K_0 \left(\frac{a_{n-1}}{b_{n-1}} - \frac{a_n}{b_n} \right) + \frac{1}{a_n} K_0 \frac{c_n a_{n-1}}{b_{n-1}} + 1 + b_n &\leq K_0 \quad [\text{by definition of } \delta_n, \zeta_n] \\ a_n(1 + b_n) &\leq K_0 a_n - K_0 \left(\frac{(1 + c_n) a_{n-1}}{b_{n-1}} - \frac{a_n}{b_n} \right) \\ a_n &\leq K_0 \left(\frac{a_n}{b_n} - \frac{1 + c_n}{1 + b_n} \frac{a_{n-1}}{b_{n-1}} \right). \end{aligned} \quad (\text{A.32})$$

Now we combine Ineqs. (A.32) and (A.28) to obtain

$$(y_n - K_0 \frac{a_n}{b_n}) \leq \frac{1 + c_n}{1 + b_n} (y_{n-1} - K_0 \frac{a_{n-1}}{b_{n-1}}). \quad (\text{A.33})$$

For brevity, define $s_n \stackrel{\text{def}}{=} y_n - K_0 a_n / b_n$. Then, from Ineq. (A.33), $s_n \leq \frac{1 + c_n}{1 + b_n} s_{n-1}$, where $\frac{1 + c_n}{1 + b_n} < 1$ since $n \geq n_0$. Assume n_1 is the smallest integer such that $n_1 \geq n_0$ and $s_{n_1} \leq 0$ (existence of n_1 is not crucial.) For all $n \geq n_1$, it follows $s_n \leq 0$, and thus $y_n \leq K_0 a_n / b_n$ for

all $n \geq n_1$. Alternatively, when $n_0 \leq n < n_1$, all s_n are positive. Using Ineq. (A.33) we have $s_n \leq (\prod_{i=n_0+1}^n \frac{1+c_i}{1+b_i})s_{n_0} \stackrel{\text{def}}{=} Q_{n_0+1}^n s_{n_0}$, and thus

$$\begin{aligned} y_n - K_0 \frac{a_n}{b_n} &\leq Q_{n_0+1}^n s_{n_0} \quad [\text{by definition of } s_n] \\ y_n &\leq K_0 \frac{a_n}{b_n} + Q_{n_0+1}^n y_{n_0} \quad [\text{because } s_n \leq y_n] \\ y_n &\leq K_0 \frac{a_n}{b_n} + Q_1^n y_0 + Q_{n_0+1}^n (1+c_1)^{n_0} A. \quad [\text{by Ineq. (A.31) on } y_{n_0}] \end{aligned} \quad (\text{A.34})$$

Combining this result with Ineqs. (A.31) and (A.34), we obtain

$$y_n \leq K_0 \frac{a_n}{b_n} + Q_1^n y_0 + Q_{n_0+1}^n (1+c_1)^{n_0} A, \quad (\text{A.35})$$

since $Q_i^n = 1$ for $n < i$, by definition. \square

Corollary 1. In Lemma 2 assume $a_n = a_1 n^{-\alpha}$ and $b_n = b_1 n^{-\beta}$, and $c_n = 0$, where $\alpha > \beta$, and $a_1, b_1, \beta > 0$ and $\alpha > 1$. Then,

$$y_n \leq 2 \frac{a_1(1+b_1)}{b_1} n^{-\alpha+\beta} + \exp(-\log(1+b_1)\phi_\beta(n))[y_0 + (1+b_1)^{n_0} A], \quad (\text{A.36})$$

where $n_0 > 0$, $A = \sum_i a_i < \infty$, and ϕ_β is defined as in Theorem 5; i.e., $\phi_\beta(n) = n^{1-\beta}$ if $\beta < 1$, and $\phi_\beta(n) = \log n$ if $\beta = 1$.

Proof. In this proof, we will assume that the inequalities $(n-1)^{-\gamma} - n^{-\gamma} \leq n^{-1-\gamma}$ and $\sum_{i=1}^n i^{-\gamma} \geq n^{1-\gamma}$ hold for $n \geq n'$, where $0 < \gamma < 1$, and n' is an appropriate threshold, that is generally small (e.g., $n' = 14$ if $\gamma = 0.1$, $n' = 5$ if $\gamma = 0.5$, and $n' = 9$ if $\gamma = 0.9$.)

By definition,

$$\delta_n \stackrel{\text{def}}{=} \frac{1}{a_n} \left(\frac{a_{n-1}}{b_{n-1}} - \frac{a_n}{b_n} \right) = \frac{1}{a_1 n^{-\alpha}} \frac{a_1}{b_1} ((n-1)^{-\alpha+\beta} - n^{-\alpha+\beta}) \leq \frac{1}{b_1} n^{-1+\beta}. \quad (\text{A.37})$$

Also, $\zeta_n = 0$ since $c_n = 0$. We can take $n_0 = \lceil (2/b_1)^{1/(1-\beta)} \rceil$, for which $\delta_{n_0} \leq 1/2$. Therefore, $K_0 \stackrel{\text{def}}{=} (1+b_1)(1-\delta_{n_0})^{-1} \leq 2(1+b_1)$; we can simply take $K_0 = 2(1+b_1)$. Since $c_n = 0$,

$Q_i^n = \prod_{j=i}^n (1 + b_j)^{-1}$. Thus,

$$\begin{aligned}
Q_1^n &\geq (1 + b_1)^{-n}, \text{ and} \\
Q_1^n &\leq \exp(-\log(1 + b_1) / b_1 \sum_{i=1}^n b_i), \quad [\text{by Lemma 1}] \\
Q_1^n &\leq \exp(-\log(1 + b_1) \phi_\beta(n)). \quad [\text{because } \sum_{i=1}^n i^{-\beta} \geq \phi_\beta(n)]
\end{aligned} \tag{A.38}$$

Lemma 2 and Ineqs. (A.38) imply

$$\begin{aligned}
y_n &\leq K_0 \frac{a_n}{b_n} + Q_1^n y_0 + Q_{n_0+1}^n (1 + c_1)^{n_0} A \quad [\text{by Lemma 2}] \\
&\leq 2 \frac{a_1(1 + b_1)}{b_1} n^{-\alpha+\beta} + Q_1^n [y_0 + (1 + b_1)^{n_0} A] \quad [\text{by Ineqs. (A.38), } c_1 = 0] \\
&\leq 2 \frac{a_1(1 + b_1)}{b_1} n^{-\alpha+\beta} + \exp(-\log(1 + b_1) n^{1-\beta}) [y_0 + (1 + b_1)^{n_0} A],
\end{aligned} \tag{A.39}$$

where the last inequality also follows from Ineqs. (A.38). \square

Proof of Theorem 3. Now we are ready to prove the main theorem. By definition (1.2),

$\theta_n^{\text{im}} = \theta_n^{(*)} - \gamma_n \tilde{\zeta}_n$, and thus, by Assumption 5,

$$\mathbb{E} (\|\theta_n^{\text{im}} - \theta_\star\|^2) \leq \mathbb{E} (\|\theta_n^{(*)} - \theta_\star\|^2) + \gamma_n^2 \sigma^2 \tag{A.40}$$

By definition (1.3), $\gamma_n h(\theta_n^{(*)}) + \theta_n^{(*)} = \theta_{n-1}^{\text{im}}$, and thus

$$\|\theta_{n-1}^{\text{im}} - \theta_\star\|^2 = \|\theta_n^{(*)} - \theta_\star\|^2 + 2\gamma_n (\theta_n^{(*)} - \theta_\star)^\top h(\theta_n^{(*)}) + \gamma_n^2 \|h(\theta_n^{(*)})\|^2. \tag{A.41}$$

Therefore,

$$\begin{aligned}
\|\theta_n^{(*)} - \theta_\star\|^2 + 2\gamma_n (\theta_n^{(*)} - \theta_\star)^\top h(\theta_n^{(*)}) &\leq \|\theta_{n-1}^{\text{im}} - \theta_\star\|^2 \\
\|\theta_n^{(*)} - \theta_\star\|^2 + 2\gamma_n \delta_n \|\theta_n^{(*)} - \theta_\star\|^2 &\leq \|\theta_{n-1}^{\text{im}} - \theta_\star\|^2 \quad [\text{by Assumption 3(b)}] \\
\|\theta_n^{(*)} - \theta_\star\|^2 &\leq \frac{1}{1 + 2\gamma_n \delta_n} \|\theta_{n-1}^{\text{im}} - \theta_\star\|^2.
\end{aligned} \tag{A.42}$$

Combining Ineq. (A.40) and Ineq. (A.42) yields

$$\begin{aligned}\mathbb{E} (\|\theta_n^{\text{im}} - \theta_\star\|^2) &= \mathbb{E} (\|\theta_n^{(*)} - \theta_\star\|^2) + \gamma_n^2 \sigma^2 \\ &\leq \frac{1}{1 + 2\gamma_n \delta_n} \mathbb{E} (\|\theta_{n-1}^{\text{im}} - \theta_\star\|^2) + \gamma_n^2 \sigma^2.\end{aligned}\quad (\text{A.43})$$

The final result of Theorem 3 is obtained through a direct application of Corollary 1 on recursion (A.43), by setting $y_n \equiv \mathbb{E} (\|\theta_n^{\text{im}} - \theta_\star\|^2)$, $b_n \equiv 2\gamma_n \delta_n$, and $a_n \equiv \gamma_n^2 \sigma^2$. \square

A.3 Asymptotic distribution

Theorem 4. *Suppose that Assumptions 1, 2, 3(a), 5, and 6 hold. Suppose also that $(2\gamma_1 J_h(\theta_\star) - I)$ is positive-definite, where $J_h(\theta)$ is the Jacobian of h at θ , and I is the $p \times p$ identity matrix. Then, the iterate θ_n^{im} of implicit stochastic approximation (1.2) is asymptotically normal, such that*

$$n^{\gamma/2}(\theta_n^{\text{im}} - \theta_\star) \rightarrow \mathcal{N}_p(0, \Sigma).$$

The covariance matrix Σ is the unique solution of

$$(\gamma_1 J_h(\theta_\star) - I/2)\Sigma + \Sigma(\gamma_1 J_h(\theta_\star)^\top - I/2) = \Xi.$$

Proof. Convergence of $\theta_n^{\text{im}} \rightarrow \theta_\star$ is established from Theorem 1. By definition of the implicit stochastic approximation procedure (1.2),

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} - \gamma_n (h(\theta_n^{(*)}) + \xi_n), \text{ and} \quad (\text{A.44})$$

$$\theta_n^{(*)} + \gamma_n h(\theta_n^{(*)}) = \theta_{n-1}^{\text{im}}. \quad (\text{A.45})$$

We use Eq. (A.45) and expand $h(\cdot)$ to obtain

$$\begin{aligned}h(\theta_n^{(*)}) &= h(\theta_{n-1}^{\text{im}}) - \gamma_n J_h(\theta_{n-1}^{\text{im}}) h(\theta_n^{(*)}) + \epsilon_n \\ h(\theta_n^{(*)}) &= (I + \gamma_n J_h(\theta_{n-1}^{\text{im}}))^{-1} h(\theta_{n-1}^{\text{im}}) + (I + \gamma_n J_h(\theta_{n-1}^{\text{im}}))^{-1} \epsilon_n,\end{aligned}\quad (\text{A.46})$$

where $\|\epsilon_n\| = O(\gamma_n^2)$ by Theorem 3. By Lipschitz continuity of $h(\cdot)$ (Assumption 3(a)) and the almost-sure convergence of θ_{n-1}^{im} to θ_\star , it follows $h(\theta_{n-1}^{\text{im}}) = J_h(\theta_\star)(\theta_{n-1}^{\text{im}} - \theta_\star) + o(1)$,

where $o(1)$ is a vector with vanishing norm. Therefore we can rewrite (A.46) as follows,

$$h(\theta_n^{(*)}) = A_n(\theta_{n-1}^{\text{im}} - \theta_*) + \mathcal{O}(\gamma_n^2), \quad (\text{A.47})$$

such that $\|A_n - J_h(\theta_*)\| \rightarrow 0$, and $\mathcal{O}(\gamma_n^2)$ denotes a vector with norm $\mathcal{O}(\gamma_n^2)$. Thus, we can rewrite (A.44) as

$$\theta_n^{\text{im}} - \theta_* = (I - \gamma_n A_n)(\theta_{n-1}^{\text{im}} - \theta_*) - \gamma_n \xi_n + \mathcal{O}(\gamma_n^2). \quad (\text{A.48})$$

The conditions for Fabian's theorem (Fabian, 1968a, Theorem 1) are now satisfied, and thus $\theta_n^{\text{im}} - \theta_*$ is asymptotically normal with mean zero, and variance that is given in the statement of Theorem 1 by Fabian (1968a). \square

Appendix B

Appendix to Chapter 2

B.1 R code

All experiments were run using our `sgd` R package that implements explicit SGD (2.3) and implicit SGD (2.4). The package is published at CRAN, <http://cran.r-project.org/web/packages/sgd/index.html>, and has been co-authored with Dustin Tran and Kuang Ye. The repository at <https://github.com/ptoulis/implicit-glms> contains legacy R code for explicit and implicit SGD. The folder `examples/aos2014` in that repository contains standalone code for the experiments in this paper, but it has been superceded by the `sgd` R package.

B.2 Useful lemmas

In this section, we will prove certain lemmas on recursions that will be useful for the non-asymptotic analysis of the implicit procedures. All following results are stated under a combination of Assumptions 7.

Lemma 1. Suppose Assumptions 7(b) (c), and (d) hold. Then, almost surely,

$$\lambda_n \leq \frac{1}{1 + \gamma_n \underline{\lambda}_c b'} \quad (\text{B.2})$$

$$\|\theta_n^{\text{im}} - \theta_{n-1}^{\text{im}}\|^2 \leq 4L_0^2 \gamma_n^2, \quad (\text{B.3})$$

where λ_n is defined in Theorem 9, and θ_n^{im} is the n th iterate of implicit SGD (2.4).

Proof. For the first part, from Theorem 9, the random variable λ_n satisfies

$$\ell(X_n^\top \theta_n^{\text{im}}; Y_n) = \lambda_n \ell(X_n^\top \theta_{n-1}^{\text{im}}; Y_n). \quad (\text{B.4})$$

Using definition (2.4),

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n \lambda_n \ell(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) C_n X_n. \quad (\text{B.5})$$

We use this definition of θ_n^{im} into Eq.(B.4) and perform a Taylor approximation on ℓ to obtain

$$\ell(X_n^\top \theta_n^{\text{im}}; Y_n) = \ell(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) + \tilde{\ell}'' \gamma_n \lambda_n X_n^\top C_n X_n, \quad (\text{B.6})$$

where $\tilde{\ell}'' = \ell''(\delta X_n^\top \theta_{n-1}^{\text{im}} + (1 - \delta) X_n^\top \theta_n^{\text{im}}; Y_n)$, and $\delta \in [0, 1]$. Thus, by definition,

$$\begin{aligned} (1 + \gamma_n \tilde{\ell}'' X_n^\top C_n X_n) \lambda_n &= 1 \\ (1 - \gamma_n \underline{\lambda}_c \tilde{\ell}'' \|X_n\|^2) \lambda_n &\leq 1 \quad [\text{by Assumption 7(e)}] \\ (1 + \gamma_n \underline{\lambda}_c b) \lambda_n &\leq 1 \quad [\text{by Assumption 7(d)}]. \end{aligned} \quad (\text{B.7})$$

For the second part, since the log-likelihood is differentiable (Assumption 7(b)) we can re-write the definition of implicit SGD (2.4) as

$$\theta_n^{\text{im}} = \arg \max \left\{ -\frac{1}{2\gamma_n} \|\theta - \theta_{n-1}^{\text{im}}\|^2 + \ell(X_n^\top \theta; Y_n) \right\}.$$

Therefore, setting $\theta \equiv \theta_{n-1}^{\text{im}}$ in the above equation,

$$\begin{aligned}
-\frac{1}{2\gamma_n} \|\theta_n^{\text{im}} - \theta_{n-1}^{\text{im}}\|^2 + \ell(X_n^\top \theta_n^{\text{im}}; Y_n) &\geq \ell(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) \\
\|\theta_n^{\text{im}} - \theta_{n-1}^{\text{im}}\|^2 &\leq 2\gamma_n (\ell(X_n^\top \theta_n^{\text{im}}; Y_n) - \ell(X_n^\top \theta_{n-1}^{\text{im}}; Y_n)) \\
\|\theta_n^{\text{im}} - \theta_{n-1}^{\text{im}}\|^2 &\leq 2\gamma_n L_0 \|\theta_n^{\text{im}} - \theta_{n-1}^{\text{im}}\|.
\end{aligned} \tag{B.8}$$

□

B.3 Non-asymptotic analysis

Theorem b. Let $\delta_n \triangleq \mathbb{E}(\|\theta_n^{\text{im}} - \theta_\star\|^2)$ and $\kappa \triangleq 1 + \gamma_1 \underline{\lambda}_c \underline{\lambda}_f \mu_0$, where $\mu_0 \in [1/(1 + \gamma_1 \underline{\lambda}_f \underline{\lambda}_c (p - 1)), 1]$. Suppose that Assumptions 7(a),(b),(c), (d) and (e) hold. Then, there exists constant n_0 such that,

$$\delta_n \leq \frac{8L_0^2 \bar{\lambda}_c^{-2} \gamma_1 \kappa}{\underline{\lambda}_c \underline{\lambda}_f \mu_0} n^{-\gamma} + \exp(-\log \kappa \cdot \phi_\gamma(n)) [\delta_0 + \kappa^{n_0} \Gamma^2],$$

where $\Gamma^2 = 4L_0^2 \bar{\lambda}_c^{-2} \sum_i \gamma_i^2 < \infty$, and $\phi_\gamma(n) = n^{1-\gamma}$ if $\gamma < 1$, and $\phi_\gamma(n) = \log n$ if $\gamma = 1$, and n_0 is defined in Corollary 1.

Proof. Starting from Procedure (2.4) we have

$$\begin{aligned}
\theta_n^{\text{im}} - \theta_\star &= \theta_{n-1}^{\text{im}} - \theta_\star + \gamma_n C_n \nabla \log f(Y_n; X_n, \theta_n^{\text{im}}) \\
\theta_n^{\text{im}} - \theta_\star &= \theta_{n-1}^{\text{im}} - \theta_\star + \gamma_n \lambda_n C_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}}) \quad [\text{By Theorem 9}] \\
\|\theta_n^{\text{im}} - \theta_\star\|^2 &= \|\theta_{n-1}^{\text{im}} - \theta_\star\|^2 \\
&\quad + 2\gamma_n \lambda_n (\theta_{n-1}^{\text{im}} - \theta_\star)^\top C_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}}) \\
&\quad + \gamma_n^2 \|C_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}})\|^2.
\end{aligned} \tag{B.9}$$

The last term can be simply bounded since $\nabla \log f(Y_n; X_n, \theta_n^{\text{im}}) = \theta_n^{\text{im}} - \theta_{n-1}^{\text{im}}$ by definition; thus,

$$\|C_n \nabla \log f(Y_n; X_n, \theta_n^{\text{im}})\|^2 \leq \bar{\lambda}_c^{-2} \|\theta_n^{\text{im}} - \theta_{n-1}^{\text{im}}\|^2 \leq 4L_0^2 \bar{\lambda}_c^{-2} \gamma_n^2, \tag{B.10}$$

which holds almost surely by Lemma 1-Eq.(B.3).

For the second-term we can bound its expectation as

$$\begin{aligned}
& \mathbb{E}(2\gamma_n\lambda_n(\theta_{n-1}^{\text{im}} - \theta_\star)^\top C_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}})) \\
& \leq \frac{2\gamma_n}{1 + \gamma_n \underline{\lambda}_c b} \mathbb{E}((\theta_{n-1}^{\text{im}} - \theta_\star)^\top C_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}})) \quad [\text{by Lemma 1}] \\
& \leq \frac{2\gamma_n}{1 + \gamma_n \underline{\lambda}_c b} \mathbb{E}((\theta_{n-1}^{\text{im}} - \theta_\star)^\top C_n \nabla h(\theta_{n-1}^{\text{im}})) \quad [\text{where } \nabla h(\theta_{n-1}^{\text{im}}) \stackrel{\text{def}}{=} \mathbb{E}(\nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}}) | \mathcal{F}_{n-1})] \\
& \leq \frac{2\gamma_n}{1 + \gamma_n \underline{\lambda}_c b} \mathbb{E}((z_{n-1} - z_\star)^\top \nabla h(z_{n-1})) \quad [\text{change of variables, } z_n \stackrel{\text{def}}{=} C_n^{1/2} \theta_n^{\text{im}}] \\
& \leq -\frac{2\gamma_n \underline{\lambda}_f}{1 + \gamma_n \underline{\lambda}_c b} \|z_{n-1} - z_\star\|^2 \quad [\text{by strong convexity, Assumption 7(d).}] \\
& \leq -\frac{2\gamma_n \underline{\lambda}_f \underline{\lambda}_c}{1 + \gamma_n \underline{\lambda}_c b} \|\theta_{n-1}^{\text{im}} - \theta_\star\|^2 \quad [\text{by change of variables, } z_n = C_n^{1/2} \theta_n^{\text{im}} \text{ and Assumption 7(e).}]
\end{aligned} \tag{B.11}$$

Taking expectations in Eq. (B.9) and substituting Eqs. (B.10) and (B.11) into Eq.(B.9) yields the recursion,

$$\mathbb{E}(\|\theta_n^{\text{im}} - \theta_\star\|^2) \leq \left(1 - \frac{2\gamma_n \underline{\lambda}_f \underline{\lambda}_c}{1 + \gamma_n \underline{\lambda}_c b}\right) \mathbb{E}(\|\theta_{n-1}^{\text{im}} - \theta_\star\|^2) + 4L_0^2 \overline{\lambda}_c^2 \gamma_n^2. \tag{B.12}$$

By Assumption 7(d) it follows $b \leq 2\underline{\lambda}_f$ and through simple algebra we obtain,

$$\left(1 - \frac{2\gamma_n \underline{\lambda}_f \underline{\lambda}_c}{1 + \gamma_n \underline{\lambda}_c b}\right) \leq \frac{1}{1 + 2\gamma_n \underline{\lambda}_f \underline{\lambda}_c}, \tag{B.13}$$

for all $n > 0$. Therefore we can write recursion (B.12) as

$$\mathbb{E}(\|\theta_n^{\text{im}} - \theta_\star\|^2) \leq \frac{1}{1 + 2\gamma_n \underline{\lambda}_f \underline{\lambda}_c} \mathbb{E}(\|\theta_{n-1}^{\text{im}} - \theta_\star\|^2) + 4L_0^2 \overline{\lambda}_c^2 \gamma_n^2. \tag{B.14}$$

We can now apply Corollary 1 with $a_n \equiv 4L_0^2 \overline{\lambda}_c^2 \gamma_n^2$ and $b_n \equiv 2\gamma_n \underline{\lambda}_f \underline{\lambda}_c$. \square

Note. Assuming Lipschitz continuity of the gradient $\nabla \ell$ instead of function ℓ would not critically alter the main result of Theorem 5. In fact, assuming Lipschitz continuity with constant L of $\nabla \ell$ and boundedness of $\mathbb{E}(\|\nabla \log f(Y_n; X_n, \theta_\star)\|^2) \leq \sigma^2$, as it is typical in the literature, would simply add a term $\gamma_n^2 L^2 \mathbb{E}(\|\theta_n^{\text{im}} - \theta_\star\|^2) + \gamma_n^2 \sigma^2$ in the right-hand side of Eq. (B.9). In this case the upper-bound is always satisfied for n such that $\gamma_n^2 L^2 > 1$,

which also highlights a difference of implicit SGD with explicit SGD, as in explicit SGD the term $\gamma_n^2 L^2 \|\theta_{n-1}^{\text{sgd}} - \theta_*\|^2$ increases the upper bound and can make $\|\theta_n^{\text{sgd}} - \theta_*\|^2$ diverge. For, $\gamma_n^2 L^2 < 1$, the discount factor for implicit SGD would be $(1 - \gamma_n^2 L^2)^{-1} (1 + \gamma_n \underline{\lambda}_f \underline{\lambda}_c)^{-1}$, which could then be bounded by a quantity $(1 + \gamma_n c)^{-1}$ for some constant c . This would lead to a solution that is similar to Theorem 5.

B.4 Asymptotic analysis

Here we prove the main result on the asymptotic variance of implicit SGD. First, we introduce linear maps $\mathbb{L}_B \{\cdot\}$ defined as $\mathbb{L}_B \{X\} = \frac{1}{2}(BX + XB)$, where B is symmetric positive-definite matrix and X is bounded. The identity map is denoted as \mathbb{I} and it holds $\mathbb{I}\{X\} = X$ for all X . Also, \mathbb{L}_0 is the null operator for which $\mathbb{L}_0 \{X\} = 0$ for all X . By the Lyapunov theorem (Lyapunov, 1992) the map \mathbb{L}_B is one-to-one and thus the inverse operator $\mathbb{L}_B^{-1} \{\cdot\}$ is well-defined. Furthermore, we define the norm of a linear map as $\|\mathbb{L}_B\| = \max_{\|X\|=1} \|\mathbb{L}_B \{X\}\|$. For bounded inputs X , it holds $\|\mathbb{L}_B\| = O(\|B\|)$.

Lemma 2. *Suppose that the sequence $\{\gamma_n\}$ satisfies Assumption 7(a). Consider the matrix recursions*

$$X_n = \mathbb{L}_{I - \gamma_n B_n} \{X_{n-1}\} + \gamma_n (C + D_n), \quad (\text{B.15})$$

$$Y_n = \mathbb{L}_{I + \gamma_n B_n}^{-1} \{X_{n-1} + \gamma_n (C + D_n)\}, \quad (\text{B.16})$$

such that

- (a) All matrices X_n, Y_n, B_n, D_n and C are bounded,
- (b) $B_n \rightarrow B$ is positive-definite and $\|B_n - B_{n-1}\| = O(\gamma_n^2)$,
- (c) C is a fixed matrix and $D_n \rightarrow 0$.

Then, both recursions approximate the matrix $\mathbb{L}_B^{-1} \{C\}$ i.e.,

$$\|X_n B + B X_n - 2C\| \rightarrow 0 \text{ and } \|Y_n B + B Y_n - 2C\| \rightarrow 0. \quad (\text{B.17})$$

If, in addition, B and C commute then $X_n, Y_n \rightarrow B^{-1}C$.

Proof. We make the following definitions.

$$\Gamma_n \stackrel{\text{def}}{=} I - \gamma_n B_n, \quad (\text{B.18})$$

$$P_i^n \stackrel{\text{def}}{=} \mathbb{L}_{\Gamma_n} \circ \mathbb{L}_{\Gamma_{n-1}} \circ \cdots \circ \mathbb{L}_{\Gamma_i}, \quad (\text{B.19})$$

where the symbol \circ denotes successive application of the linear maps, and $P_i^n = \mathbb{I}$ if $n < i$, by definition. It follows,

$$\|P_i^n\| = \mathcal{O}\left(\prod_{j=i}^n \|I - \gamma_j B_j\|\right) \leq K_0 e^{-K_1 \sum_{j=i}^n \gamma_j}, \quad (\text{B.20})$$

for suitable constants K_0, K_1 (see [Polyak and Juditsky, 1992b](#), Appendix, Part 3). Let $\Gamma(n) = K_1 \sum_{i=1}^n \gamma_i$. By Assumption 7(a), $\Gamma(n) \rightarrow \infty$ and thus $P_i^n \rightarrow \mathbb{L}_0$ as $n \rightarrow \infty$ and i is fixed. The matrix recursion in Lemma 2 can be rewritten as $X_n = \mathbb{L}_{\Gamma_n} \{X_{n-1}\} + \gamma_n C + \gamma_n D_n$. Solving the recursion yields

$$\begin{aligned} X_n &= \mathbb{L}_{\Gamma_n} \circ \mathbb{L}_{\Gamma_{n-1}} \circ \cdots \circ \mathbb{L}_{\Gamma_1} \{X_0\} + \gamma_n C + \gamma_n D_n \\ &\quad + a_{n-1} \mathbb{L}_{\Gamma_n} \{C\} + a_{n-1} \mathbb{L}_{\Gamma_n} \{D_{n-1}\} \\ &\quad + \cdots + \\ &\quad + a_1 \mathbb{L}_{\Gamma_n} \circ \mathbb{L}_{\Gamma_{n-1}} \circ \cdots \circ \mathbb{L}_{\Gamma_2} \{C\} + a_1 \mathbb{L}_{\Gamma_n} \circ \mathbb{L}_{\Gamma_{n-1}} \circ \cdots \circ \mathbb{L}_{\Gamma_2} \{D_1\} \\ &\stackrel{\text{def}}{=} P_1^n \{X_0\} + S_n \{C\} + \tilde{D}_n, \end{aligned} \quad (\text{B.21})$$

where we have defined the linear map $S_n = \sum_{i=1}^n \gamma_i P_{i+1}^n$ and the matrix $\tilde{D}_n = \sum_{i=1}^n \gamma_i P_{i+1}^n \{D_i\}$. Since $P_1^n \rightarrow \mathbb{L}_0$, our goal is to prove that $S_n \rightarrow \mathbb{L}_B^{-1}$ and $\tilde{D}_n \rightarrow 0$. By definition,

$$\sum_{i=1}^n \gamma_i P_{i+1}^n = \mathbb{L}_{B_n}^{-1} + \sum_{i=2}^n P_i^n (\mathbb{L}_{B_{i-1}}^{-1} - \mathbb{L}_{B_i}^{-1}) - P_1^n \mathbb{L}_{B_1}^{-1}. \quad (\text{B.22})$$

To see this, first note that $\gamma_n I = (I - \Gamma_n) B_n^{-1}$ for every n , and thus

$$\gamma_n \mathbb{I} = \mathbb{L}_{I - \Gamma_n} \circ \mathbb{L}_{B_n}^{-1}. \quad (\text{B.23})$$

Therefore, if we collect the coefficients of the terms $\mathbb{L}_{B_n}^{-1}$ in the right-hand side of (B.22), we

get

$$\begin{aligned}
& \mathbb{L}_{B_n}^{-1} + \sum_{i=2}^n P_i^n (\mathbb{L}_{B_{i-1}}^{-1} - \mathbb{L}_{B_i}^{-1}) - P_1^n \mathbb{L}_{B_1}^{-1} \\
&= (P_2^n - P_1^n) \mathbb{L}_{B_1}^{-1} + (P_3^n - P_2^n) \mathbb{L}_{B_2}^{-1} + \cdots + (P_{n+1}^n - P_n^n) \mathbb{L}_{B_n}^{-1} \\
&= P_2^n \circ \mathbb{L}_{I-\Gamma_1} \circ \mathbb{L}_{B_1}^{-1} + P_3^n \circ \mathbb{L}_{I-\Gamma_2} \circ \mathbb{L}_{B_2}^{-1} + \cdots + P_{n+1}^n \circ \mathbb{L}_{I-\Gamma_n} \circ \mathbb{L}_{B_n}^{-1} \\
&= P_2^n (\gamma_1 \mathbb{I}) + P_3^n (\gamma_2 \mathbb{I}) + \cdots + P_{n+1}^n (\gamma_n \mathbb{I}) \quad [\text{by Eq. (B.23)}] \\
&= \sum_{i=1}^n \gamma_i P_{i+1}^n,
\end{aligned}$$

where we used the identity $P_{i+1}^n - P_i^n = P_{i+1}^n \circ (\mathbb{I} - \mathbb{L}_{\Gamma_i}) = P_{i+1}^n \circ \mathbb{L}_{I-\Gamma_i}$. Furthermore, since B_i are bounded,

$$\begin{aligned}
\|\mathbb{L}_{B_{i-1}}^{-1} - \mathbb{L}_{B_i}^{-1}\| &= \|\mathbb{L}_{B_i}^{-1} \circ (\mathbb{L}_{B_i} - \mathbb{L}_{B_{i-1}}) \circ \mathbb{L}_{B_{i-1}}^{-1}\| = O(\|\mathbb{L}_{B_i} - \mathbb{L}_{B_{i-1}}\|) \\
&= O(\|B_i - B_{i-1}\|) = O(\gamma_i^2). \quad [\text{By assumption of Lemma 2}]
\end{aligned}$$

In addition, $\|\sum_{i=2}^n P_i^n \circ (\mathbb{L}_{B_{i-1}}^{-1} - \mathbb{L}_{B_i}^{-1})\| \leq K_0 e^{-\Gamma(n)} \sum_{i=2}^n e^{\Gamma(i)} O(\gamma_i^2)$. Since $\sum_i O(\gamma_i^2) < \infty$ and $e^{\Gamma(i)}$ is positive, increasing and diverging, we can invoke Kronecker's lemma and obtain $\sum_{i=2}^n e^{\Gamma(i)} O(\gamma_i^2) = o(e^{\Gamma(n)})$. Therefore

$$\sum_{i=2}^n P_i^n \circ (\mathbb{L}_{B_{i-1}}^{-1} - \mathbb{L}_{B_i}^{-1}) \rightarrow \mathbb{L}_0, \quad (\text{B.24})$$

and since $P_1^n \rightarrow \mathbb{L}_0$, we conclude from Equation (B.23) that

$$\lim_{n \rightarrow \infty} \sum_{i=1}^n \gamma_i P_{i+1}^n = \lim_{n \rightarrow \infty} \mathbb{L}_{B_n}^{-1} = \mathbb{L}_B^{-1}. \quad (\text{B.25})$$

Thus, $S_n \rightarrow \mathbb{L}_B^{-1}$, as desired. For \tilde{D}_n we have

$$\begin{aligned}
\tilde{D}_n &= \sum_{i=1}^n \gamma_i P_{i+1}^n \{D_i\} = \mathbb{L}_{B_n}^{-1} \{D_n\} + \sum_{i=2}^n P_i^n \circ (\mathbb{L}_{B_{i-1}}^{-1} \{D_{i-1}\} - \mathbb{L}_{B_i}^{-1} \{D_i\}) \\
&\quad + P_1^n \circ \mathbb{L}_{B_1}^{-1} \{D_1\}.
\end{aligned}$$

Since $\|D_n\| \rightarrow 0$ it follows that $\|\mathbb{L}_{B_n}^{-1} \{D_n\}\| \rightarrow 0$ and $\|(\mathbb{L}_{B_{i-1}}^{-1} \{D_{i-1}\} - \mathbb{L}_{B_i}^{-1} \{D_i\})\| = O(\gamma_i^2)$. Recall that $P_1^n \rightarrow \mathbb{L}_0$, and thus $\tilde{D}_n \rightarrow \mathbf{0}$. Finally, we substitute this result in Equation (B.23) to get $X_n \rightarrow \mathbb{L}_B^{-1} \{C\}$.

For the second recursion of the lemma,

$$Y_n = \mathbb{L}_{I+\gamma_n B_n}^{-1} \{Y_{n-1} + \gamma_n(C + D_n)\}, \quad (\text{B.26})$$

the proof is similar. First, we make the following definitions.

$$\begin{aligned} \Gamma_n &\stackrel{\text{def}}{=} I + \gamma_n B_n, \\ Q_i^n &\stackrel{\text{def}}{=} \mathbb{L}_{\Gamma_n}^{-1} \circ \mathbb{L}_{\Gamma_{n-1}}^{-1} \circ \cdots \circ \mathbb{L}_{\Gamma_i}^{-1}. \end{aligned}$$

As before, $Q_i^n \rightarrow \mathbb{L}_0$. Solving the recursion (B.26) yields

$$Y_n = Q_1^n \{Y_0\} + S_n \{C\} + \tilde{D}_n, \quad (\text{B.27})$$

where we defined $S_n \stackrel{\text{def}}{=} \sum_{i=1}^n \gamma_i Q_i^n$ and $\tilde{D}_n \stackrel{\text{def}}{=} \sum_{i=1}^n \gamma_i Q_i^n \{D_i\}$. The following identities can also be verified by the definition of the linear maps.

$$\mathbb{L}_{B_n}^{-1} \circ (\mathbb{I} - \mathbb{L}_{\Gamma_n}^{-1}) = \gamma_n \mathbb{L}_{\Gamma_n}^{-1}, \quad (\text{B.28})$$

$$\mathbb{L}_{B_n}^{-1} \mathbb{L}_{\Gamma_n}^{-1} = \mathbb{L}_{\Gamma_n}^{-1} \mathbb{L}_{B_n}^{-1}. \quad (\text{B.29})$$

It holds,

$$\begin{aligned} \mathbb{L}_{B_n}^{-1} + \sum_{i=1}^n Q_i^n \circ (\mathbb{L}_{B_{i-1}}^{-1} - \mathbb{L}_{B_i}^{-1}) &= \mathbb{L}_{B_n}^{-1} \circ (\mathbb{I} - \mathbb{L}_{\Gamma_n}^{-1}) + \mathbb{L}_{\Gamma_n}^{-1} \circ \mathbb{L}_{B_{n-1}}^{-1} \circ (\mathbb{I} - \mathbb{L}_{\Gamma_n}^{-1}) + \cdots \\ &= \gamma_n \mathbb{L}_{\Gamma_n}^{-1} + \gamma_{n-1} \mathbb{L}_{\Gamma_n}^{-1} \mathbb{L}_{\Gamma_{n-1}}^{-1} + \cdots = S_n, \end{aligned}$$

where the first line is obtained by Eq. (B.28) and the second line by Eq. (B.29). Thus, similar to the previously analyzed recursion, $S_n \rightarrow \mathbb{L}_B^{-1}$ and $\tilde{D}_n \rightarrow 0$. Therefore, $Y_n \rightarrow \mathbb{L}_B^{-1} \{C\}$.

For both cases, if B, C commute then $\mathbb{L}_B^{-1} \{C\} = X$ such that $BX + XB = 2C$. Setting $X = B^{-1}C$ is a solution since $BB^{-1}C + B^{-1}CB = C + B^{-1}BC = 2C$. By the Lyapunov theorem, this solution is unique. \square

Corollary 2. Consider the matrix recursions

$$X_n = \mathbb{L}_{I-\gamma_n B_n} \{X_{n-1}\} + \gamma_n^2(C + D_n), \quad (\text{B.30})$$

$$Y_n = \mathbb{L}_{I+\gamma_n B_n}^{-1} \{Y_{n-1} + \gamma_n^2(C + D_n)\}, \quad (\text{B.31})$$

where B_n, B, C, D_n satisfy the assumptions of Lemma 2. Moreover, suppose $\gamma_n = \gamma_1 n^{-1}$. If the matrix $B - I/\gamma_1$ is positive-definite, then

$$(1/\gamma_n)X_n, (1/\gamma_n)Y_n \rightarrow \mathbb{L}_{B-I/\gamma_1}^{-1} \{C\} \text{ i.e.,}$$

both matrices $(1/\gamma_n)X_n$ and $(1/\gamma_n)Y_n$ approximate the matrix $\mathbb{L}_{B-I/\gamma_1}^{-1} \{C\}$. If, in addition, B and C commute then $(1/\gamma_n)X_n, (1/\gamma_n)Y_n \rightarrow (B - I/\gamma_1)^{-1}C$.

Proof. Both $X_n, Y_n \rightarrow 0$ by direct application of Lemma (2). Let $\tilde{X}_n = (1/\gamma_n)X_n$. First, divide (B.30) by γ_n to obtain

$$\tilde{X}_n = \mathbb{L}_{I-\gamma_n B_n} \{\tilde{X}_{n-1}\} \frac{\gamma_{n-1}}{\gamma_n} + \gamma_n(C + D_n). \quad (\text{B.32})$$

By Assumption 7(a), $\gamma_{n-1}/\gamma_n = 1 + \gamma_n/\gamma_1 + O(\gamma_n^2)$. Then,

$$\mathbb{L}_{I-\gamma_n B_n} \{\tilde{X}_{n-1}\} \frac{\gamma_{n-1}}{\gamma_n} = \mathbb{L}_{I-\gamma_n B_n} \{\tilde{X}_{n-1}\} + \gamma_n \tilde{X}_{n-1} + O(\gamma_n^2). \quad (\text{B.33})$$

Therefore, we can rewrite Eq. (B.32) as

$$\tilde{X}_n = \mathbb{L}_{I-\gamma_n \Gamma_n} \{\tilde{X}_{n-1}\} + \gamma_n(C + D_n), \quad (\text{B.34})$$

where $\Gamma_n \stackrel{\text{def}}{=} B_n - I/\gamma_1 + O(\gamma_n)$. In the limit $\Gamma_n \rightarrow B - I/\gamma_1 > 0$. Furthermore, $\|\Gamma_{i-1} - \Gamma_i\| = O(\gamma_i^2)$ by assumptions of Corollary 2. Thus, we can apply Lemma (2) to conclude that $\tilde{X}_n \stackrel{\text{def}}{=} (1/\gamma_n)X_n \rightarrow \mathbb{L}_{B-I/\gamma_1}^{-1} \{C\}$. The proof for Y_n follows the same reasoning since $(I + \gamma_n B_n)^{-1}(\gamma_{n-1}/\gamma_n) = (I + \gamma_n \Gamma_n)^{-1}$, where $\Gamma_n \stackrel{\text{def}}{=} B_n - I/\gamma_1 + O(\gamma_n)$. \square

Theorem b. Consider SGD procedures (2.3) and (2.4), and suppose that Assumptions 7(a),(c),(d),(e) hold, where $\gamma = 1$. The asymptotic variance of the explicit SGD estimator (2.3) satisfies

$$n\text{Var} \left(\theta_n^{\text{sgd}} \right) \rightarrow \gamma_1^2 (2\gamma_1 C\mathcal{I}(\theta_*) - I)^{-1} C\mathcal{I}(\theta_*)C.$$

The asymptotic variance of the implicit SGD estimator (2.4) satisfies

$$n\text{Var}(\theta_n^{\text{im}}) \rightarrow \gamma_1^2 (2\gamma_1 C\mathcal{I}(\theta_\star) - I)^{-1} C\mathcal{I}(\theta_\star)C.$$

Proof. We begin with the implicit SGD procedure. For notational convenience we make the following definitions: $V_n \stackrel{\text{def}}{=} \text{Var}(\theta_n^{\text{im}})$, $S_n(\theta) \stackrel{\text{def}}{=} \nabla \log f(Y_n; X_n, \theta)$, Note, $\mathbb{E}(S_n(\theta)) \stackrel{\text{def}}{=} h(\theta)$. Let J_h denote the Jacobian of function h , then, under typical regularity conditions and by Theorem 5:

$$\begin{aligned} \mathbb{E}(S_n(\theta_\star) | X_n) &= 0 \\ \text{Var}(S_n(\theta_\star)) &= \mathbb{E}(\text{Var}(S_n(\theta_\star) | X_n)) \stackrel{\text{def}}{=} \mathcal{I}(\theta_\star) \\ J_h(\theta) &= -\mathcal{I}(\theta), \quad [\text{under regularity conditions}] \\ h(\theta_n^{\text{im}}) &= -\mathcal{I}(\theta_\star)(\theta_n^{\text{im}} - \theta_\star) + \mathcal{O}(\gamma_n) \quad [\text{by Theorem 5}], \\ \|\text{Var}(S_n(\theta) - S_n(\theta_\star))\| &\leq \mathbb{E}(\|S_n(\theta) - S_n(\theta_\star)\|^2) \leq L_0^2 \mathbb{E}(\|\theta - \theta_\star\|^2). \end{aligned} \quad (\text{B.35})$$

We can now rewrite Eq. (2.4) as follows,

$$\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n C_n S_n(\theta_n^{\text{im}}) = \theta_{n-1}^{\text{im}} + \gamma_n \lambda_n C_n S_n(\theta_{n-1}^{\text{im}}), \quad (\text{B.36})$$

where λ_n is defined in Theorem 9 and $\lambda_n = 1 - \mathcal{O}(\gamma_n)$ by Eq. (B.2). Then, taking variances on both sides of Eq. (B.36) yields

$$V_n = V_{n-1} + \gamma_n^2 C_n \text{Var}(S_n(\theta_n^{\text{im}})) C_n^\top + \gamma_n \text{Cov}(\theta_{n-1}^{\text{im}}, S_n(\theta_n^{\text{im}})) C_n^\top + \gamma_n C_n \text{Cov}(S_n(\theta_n^{\text{im}}), \theta_{n-1}^{\text{im}}). \quad (\text{B.37})$$

We can simplify all variance/covariance terms in Eq. (B.37) as follows.

$$\begin{aligned} C_n \text{Var}(S_n(\theta_n^{\text{im}})) C_n^\top &= C_n \text{Var}(S_n(\theta_\star) + [S_n(\theta_n^{\text{im}}) - S_n(\theta_\star)]) C_n^\top \\ &= C\mathcal{I}(\theta_\star)C^\top + \mathcal{O}(1), \quad [\text{by Eqs. (B.35), Theorem 5, and Assumption 7(e)}] \\ \text{Cov}(\theta_{n-1}^{\text{im}}, S_n(\theta_n^{\text{im}})) &= \text{Cov}(\theta_{n-1}^{\text{im}}, S_n(\theta_{n-1}^{\text{im}})) + \text{Cov}(\theta_{n-1}^{\text{im}}, (\lambda_n - 1)S_n(\theta_{n-1}^{\text{im}})) \\ &= \text{Cov}(\theta_{n-1}^{\text{im}}, h(\theta_{n-1}^{\text{im}})) + \mathcal{O}(\gamma_n) \\ &= V_{n-1} \mathcal{I}(\theta_\star) + \mathcal{O}(\gamma_n). \quad [\text{by Eq. (B.35), Theorem 5, Eq. (B.2)}]. \end{aligned}$$

Similarly, $\text{Cov}(h(\theta_n^{\text{im}}), \theta_{n-1}^{\text{im}}) = V_{n-1}\mathcal{I}(\theta_*) + \mathcal{O}(\gamma_n)$. We can now rewrite Eq. (B.37) as

$$V_n = \mathbb{L}_{I-\gamma_n B_n} \{V_{n-1}\} + \gamma_n^2 [C\mathcal{I}(\theta_*)C^\top + \mathfrak{o}(1)], \quad (\text{B.38})$$

where $B_n \stackrel{\text{def}}{=} 2C_n\mathcal{I}(\theta_*)$ and $B_n \rightarrow 2C\mathcal{I}(\theta_*)$. Corollary (2) on recursion (B.38) yields

$$(1/n)V_n \rightarrow \gamma_1^2 (2\gamma_1 C\mathcal{I}(\theta_*) - I)^{-1} C\mathcal{I}(\theta_*)C.$$

We obtain the asymptotic variance in closed-form since B and C commute, and C is symmetric. The regularity conditions (B.35) and the convergence rates of Theorem 5 that are crucial for this proof, also hold for the explicit procedure. Therefore for the proof for the asymptotic variance of explicit SGD (2.3) is identical. \square

Theorem c. Consider SGD procedure (2.13) and suppose Assumptions 7(a),(c),(d), and (e) hold, where $\gamma \in [0.5, 1)$. Then, the iterate $\theta_n^{(*)}$ converges to θ_* in probability and is asymptotically efficient, i.e.,

$$n\text{Var}\left(\theta_n^{(*)}\right) \rightarrow \mathcal{I}(\theta_*)^{-1}.$$

Proof. By Theorem 5 and Assumptions 7 (d), (c), we have

$$\nabla \log f(Y_n; X_n, \theta_n^{\text{im}}) = \nabla \log f(Y_n; X_n, \theta_*) - \mathcal{I}(\theta_*)(\theta_n^{\text{im}} - \theta_*) + \mathcal{O}(\gamma_n). \quad (\text{B.39})$$

Define, for convenience $\varepsilon_n \stackrel{\text{def}}{=} \nabla \log f(Y_n; X_n, \theta_*)$, $F \stackrel{\text{def}}{=} \mathcal{I}(\theta_*)$. Then, the first-order implicit SGD iteration becomes

$$\theta_n^{\text{im}} - \theta_* = (I + \gamma_n F)^{-1} (\theta_{n-1}^{\text{im}} - \theta_* + \gamma_n \varepsilon_n + \mathcal{O}(\gamma_n^2)). \quad (\text{B.40})$$

We make the following definitions.

$$\begin{aligned} e_i &\stackrel{\text{def}}{=} \gamma_i (I + \gamma_i F)^{-1} (\varepsilon_i + \mathcal{O}(\gamma_i^2)), \\ B_i^j &\stackrel{\text{def}}{=} \prod_{k=j}^i (I + \gamma_k F)^{-1}, \\ D_j^n &\stackrel{\text{def}}{=} \prod_{k=n-1}^i B_{j+1}^k = I + B_{j+1}^{j+1} + B_{j+1}^{j+2} + \dots + B_{j+1}^{n-1}. \end{aligned} \quad (\text{B.41})$$

Then, we can solve the recursion for $\theta_n^{(*)} - \theta_*$ to obtain

$$\theta_n^{(*)} - \theta_* = (1/n)D_0^n(\theta_0^{(*)} - \theta_*) + (1/n) \sum_i^{n-1} D_i^n e_i. \quad (\text{B.42})$$

Our proof is now split into proving the following two lemmas.

Lemma 3. *Under Assumption 7(a) $D_0^n = o(n)$.*

Proof. Matrix F is positive-definite by Assumption 7(d). Thus, if λ is some eigenvalue of F then the corresponding eigenvalue of D_0^n is $1 + \frac{1}{1+\gamma_1\lambda} + \frac{1}{1+\gamma_1\lambda} \frac{1}{1+\gamma_2\lambda} + \dots \leq \sum_{i=0}^n \exp(-K\lambda \sum_{k=1}^i \gamma_k)$, where the last inequality is obtained by Lemma 1. Because $\sum \gamma_i \rightarrow \infty$, the summands are $o(1)$, and thus D_0^n is $o(n)$. \square

Lemma 4. *Suppose Assumption 7(a) and Eq. (B.39) hold. Then,*

$$\gamma_i D_i^n (I + \gamma_i F)^{-1} = \Omega_i^n + F^{-1}, \quad (\text{B.43})$$

such that $\sum_{i=0}^{n-1} \Omega_i^n = o(n)$.

Proof. Our goal will be to compare the eigenvalues of $\gamma_i D_i^n$ and F . Any matrix D_i^n shares the same eigenvectors with F because F is positive-definite, and thus a relationship on eigenvalues will automatically establish a relationship on the matrices. For convenience, define $q_i^j \stackrel{\text{def}}{=} \prod_{k=i}^j (1 + \gamma_k \lambda)^{-1}$ for $\lambda > 0$; by convention, $q_{i-1}^i = 1$. Also let $s_i^j \stackrel{\text{def}}{=} \sum_{k=i}^j \gamma_k$ be the function of partial sums. By Lemma 1 $q_i^j = O(\exp(-K\lambda s_i^j))$, for some $K > 0$. For an eigenvalue $\lambda > 0$ of F the corresponding eigenvalue, say λ' , of matrix $\gamma_i D_i^n (I + \gamma_i F)^{-1}$ is equal to

$$\lambda' = \frac{\gamma_i}{1 + \gamma_i \lambda} (q_{i+1}^i + q_{i+1}^{i+1} + \dots + q_{i+1}^{n-1}). \quad (\text{B.44})$$

Thus,

$$\lambda' (1 + \gamma_i \lambda) = \sum_{k=i}^{n-1} \gamma_i q_{i+1}^k. \quad (\text{B.45})$$

Our goal will be to derive the relationship between λ and λ' . By definition

$$\begin{aligned}
\gamma_{i+1}\lambda q_{i+1}^{i+1} + q_{i+1}^{i+1} &= 1 \\
\gamma_{i+2}\lambda q_{i+1}^{i+2} + q_{i+1}^{i+2} &= q_{i+1}^{i+1} \\
&\dots\dots \\
\gamma_{n-2}\lambda q_{i+1}^{n-2} + q_{i+1}^{n-2} &= q_{i+1}^{n-3} \\
\gamma_{n-1}\lambda q_{i+1}^{n-1} + q_{i+1}^{n-1} &= q_{i+1}^{n-2}.
\end{aligned} \tag{B.46}$$

By summing over the terms we obtain:

$$\lambda \sum_{k=i+1}^{n-1} \gamma_k q_{i+1}^k + q_{i+1}^{n-1} = 1. \tag{B.47}$$

If we combine with (B.44) we obtain

$$\lambda \sum_{k=i}^{n-1} \gamma_i q_{i+1}^k + \lambda \sum_{k=i}^{n-1} (\gamma_k - \gamma_i) q_{i+1}^k + q_{i+1}^{n-1} = 1 + \gamma_i \lambda \Rightarrow \tag{B.48}$$

$$(1 + \gamma_i \lambda) \lambda \lambda' + \lambda \sum_{k=i}^{n-1} (\gamma_k - \gamma_i) q_{i+1}^k + q_{i+1}^{n-1} = 1 + \gamma_i \lambda. \tag{B.49}$$

We now focus on the second term. By telescoping the series we obtain

$$\begin{aligned}
\lambda \sum_{k=i}^{n-1} (\gamma_k - \gamma_i) q_{i+1}^k &= \lambda \sum_{k=i}^{n-1} \left[\sum_{j=i}^k (\gamma_{j+1} - \gamma_j) \right] q_{i+1}^k = \lambda \sum_{k=i}^{n-1} \left[\sum_{j=i}^k \gamma_j o(\gamma_j) \right] q_{i+1}^k \\
&\leq \lambda o(\gamma_i) \sum_{k=i}^{n-1} s_i^k q_{i+1}^k \triangleq q_i^n.
\end{aligned} \tag{B.50}$$

In Eq. (B.50) we used $(\gamma_{j+1} - \gamma_j)/\gamma_j = O(n^{-1-\gamma})/n^{-\gamma} = O(n^{-1}) = o(\gamma_j)$, by Assumption 7(a). Our goal is now to show $\sum_{i=0}^{n-1} q_i^n = o(n)$. Since $q_{i+1}^k = O(\exp(-K\lambda s_{i+1}^k))$ by Polyak (1992, p845, see A6 and A7) we obtain that $q_i^n \rightarrow 0$ for fixed i as $n \rightarrow \infty$. Therefore we can rewrite (B.48) as

$$\lambda' \lambda + q_i^n + O(q_{i+1}^n) = 1, \tag{B.51}$$

where $\sum_{i=0}^n q_{i+1}^n = o(n)$ and $\sum_{i=0}^{n-1} q_i^n = o(n)$. □

Our proof is now complete. By Eq. (B.42) and Lemmas 3 and 4 we have

$$\theta_n^{(*)} - \theta_\star = F^{-1} \sum_{i=1}^n \varepsilon_i + (1/n)\mathfrak{o}(n).$$

Because $\text{Var}(\varepsilon_i) = \mathcal{I}(\theta_\star)$, we finally obtain

$$n\text{Var}(\theta_n^{(*)} - \theta_\star) = \mathcal{I}(\theta_\star)^{-1}.$$

□

Theorem d. *Suppose that Assumptions 7(a),(c),(d),(e),(f) hold. Then, the iterate θ_n^{im} of implicit SGD (2.4) is asymptotically normal, such that*

$$n^{\gamma/2}(\theta_n^{\text{im}} - \theta_\star) \rightarrow \mathcal{N}_p(0, \Sigma),$$

where $\Sigma = \gamma_1^2(2\gamma_1 C\mathcal{I}(\theta_\star) - I)^{-1} C\mathcal{I}(\theta_\star)C$.

Proof. Let $S_n(\theta) \stackrel{\text{def}}{=} \nabla \log f(Y_n; X_n, \theta)$ as in the proof of Theorem 6. The conditions for Fabian's theorem – see Fabian (1968a, Theorem 1) – hold also for the implicit procedure. The goal is to show that

$$\theta_n^{\text{im}} - \theta_\star = (I - \gamma_n A_n)(\theta_{n-1}^{\text{im}} - \theta_\star) + \gamma_n \xi_n(\theta_\star) + \mathcal{O}(\gamma_n^2), \quad (\text{B.52})$$

where $A_n \rightarrow A \succeq 0$, and $\xi_n(\theta) = S_n(\theta) - h(\theta)$, and $h(\theta) = \mathbb{E}(S_n(\theta))$; note, $\xi_n(\theta_\star) = S_n(\theta_\star)$. Indeed, by a Taylor expansion on $S_n(\theta_n^{\text{im}})$ and considering that $\theta_n^{\text{im}} = \theta_{n-1}^{\text{im}} + \gamma_n S_n(\theta_n^{\text{im}})$, by definition, we have

$$(I + \gamma_n \hat{\mathcal{L}}_n(\theta_\star))(\theta_n^{\text{im}} - \theta_\star) = \theta_{n-1}^{\text{im}} - \theta_\star + \gamma_n S_n(\theta_\star), \quad (\text{B.53})$$

where $\hat{\mathcal{L}}_n(\theta_\star) = -\nabla^2 S_n(\theta_\star)$; note, $\mathbb{E}(\hat{\mathcal{L}}_n(\theta_\star)) = \mathcal{I}(\theta_\star)$. Because $(I + \gamma_n \hat{\mathcal{L}}_n(\theta_\star))^{-1} = I - \gamma_n \hat{\mathcal{L}}_n(\theta_\star) + \mathcal{O}(\gamma_n^2)$, we can rewrite Eq. (B.53) as

$$\theta_n^{\text{im}} - \theta_\star = (I - \gamma_n \hat{\mathcal{L}}_n(\theta_\star))(\theta_{n-1}^{\text{im}} - \theta_\star) + \gamma_n S_n(\theta_\star) + \mathcal{O}(\gamma_n^2). \quad (\text{B.54})$$

We can now apply Fabian's Theorem to derive asymptotic normality of θ_n^{im} . The variance matrix of the asymptotic normal distribution is derived in Theorem 4 under weaker

conditions. □

B.5 Stability

Theorem f. Let $\bar{\lambda}_f = \max \text{eig}(\mathcal{I}(\theta_\star))$, and suppose $\gamma_n = \gamma_1/n$ and $\gamma_1 \bar{\lambda}_f > 1$. Then, then the maximum eigenvalue of P_1^n satisfies

$$\max_{n>0} \max \{ \text{eig}(P_1^n) \} = \Theta(2^{\gamma_1 \bar{\lambda}_f} / \sqrt{\gamma_1 \bar{\lambda}_f}).$$

For the implicit method,

$$\max_{n>0} \max \{ \text{eig}(Q_1^n) \} = \mathcal{O}(1).$$

Proof. We will use the following intermediate result:

$$\max_{n>0} \left| \prod_{i=1}^n (1 - b/i) \right| \approx \begin{cases} 1 - b & \text{if } 0 < b < 1 \\ \frac{2^b}{\sqrt{2\pi b}} & \text{if } b > 1 \end{cases}$$

The first case is obvious. For the second case, $b > 1$, assume without loss of generality that b is an even integer. Then the maximum is given by

$$(b-1)(b/2-1)(b/3-1) \cdots (2-1) = \frac{1}{2} \binom{b}{b/2} = \Theta(2^b / \sqrt{2\pi b}), \quad (\text{B.55})$$

where the last approximation follows from Stirling's formula. The stability result on the explicit SGD updates of Lemma 1 follows immediately by using the largest eigenvalue $\bar{\lambda}_f$ of $\mathcal{I}(\theta_\star)$. For the implicit SGD updates, we note that the eigenvalues of $(I + \gamma_n \mathcal{I}(\theta_\star))^{-1}$ are less than one, for any $\gamma_n > 0$ and any Fisher matrix. □

B.6 Applications

Theorem g. Suppose Assumption 7(b) holds. Then, the gradient for the implicit update (2.4) is a scaled version of the gradient at the previous iterate, i.e.,

$$\nabla \log f(Y_n; X_n, \theta_n^{\text{im}}) = \lambda_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}}), \quad (\text{B.56})$$

where the scalar λ_n satisfies,

$$\lambda_n \ell'(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) = \ell'(X_n^\top \theta_{n-1}^{\text{im}} + \gamma_n \lambda_n \ell'(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) X_n^\top C_n X_n; Y_n). \quad (\text{B.57})$$

Proof. From the chain rule $\nabla \log f(Y_n; X_n, \theta) = \ell'(X_n^\top \theta; Y_n) X_n$, and thus $\nabla \log f(Y_n; X_n, \theta_n^{\text{im}}) = \ell'(X_n^\top \theta_n^{\text{im}}; Y_n) X_n$ and $\nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}}) = \ell'(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) X_n$, and thus the two gradients are colinear. Therefore there exists a scalar λ_n such that

$$\begin{aligned} \nabla \log f(Y_n; X_n, \theta_n^{\text{im}}) &= \lambda_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{im}}) \Rightarrow \\ \ell'(X_n^\top \theta_n^{\text{im}}; Y_n) X_n &= \lambda_n \ell'(X_n^\top \theta_{n-1}^{\text{im}}; Y_n) X_n. \end{aligned} \quad (\text{B.58})$$

We also have,

$$\begin{aligned} \theta_n^{\text{im}} &= \theta_{n-1}^{\text{im}} + \gamma_n C_n \log f(Y_n; X_n, \theta_n^{\text{im}}) \quad [\text{by definition of implicit SGD (2.4)}] \\ &= \theta_{n-1}^{\text{im}} + \gamma_n \lambda_n C_n \log f(Y_n; X_n, \theta_{n-1}^{\text{im}}). \quad [\text{by Eq. (B.58)}] \end{aligned} \quad (\text{B.59})$$

Substituting the expression for θ_n^{im} in Eq.(B.59) into Eq. (B.58) we obtain the desired result of the Theorem in Eq. (B.56).

We now prove the last claim of the theorem regarding the search bounds for λ_n . For notational convenience, define $a \stackrel{\text{def}}{=} X_n^\top \theta_{n-1}^{\text{im}}$, $g(x) \stackrel{\text{def}}{=} \ell'(x; Y_n)$, and $c = X_n^\top C_n X_n$, where $c > 0$ because C_n are positive-definite. Also let $x_\star \stackrel{\text{def}}{=} \gamma_n \lambda_n g(a)$, then the fixed-point equation (B.57) can be written as

$$x_\star = \gamma_n g(a + x_\star c). \quad (\text{B.60})$$

where g is decreasing by Assumption (b). If $g(a) = 0$ then $x_\star = 0$. If $g(a) > 0$ then $x_\star > 0$

and $\gamma_n g(a + xc) < \gamma_n g(a)$ for all $x > 0$, since $g(a + xc)$ is decreasing; taking $x = x_*$ yields $\gamma_n g(a) > \gamma_n g(a + x_*c) = x_*$, by the fixed-point equation (B.60). Thus, $0 < x_* < \gamma_n g(a)$. Similarly, if $g(a) < 0$ then $x_* < 0$ and $\gamma_n g(a + xc) > \gamma_n g(a)$ for all $x < 0$, since $g(a + xc)$ is decreasing; taking $x = x_*$ yields $\gamma_n g(a) < \gamma_n g(a + x_*c) = x_*$, by the fixed-point equation. Thus, $\gamma_n g(a) < x_* < 0$. In both cases $0 < \lambda_n < 1$. A visual proof is given Figure B.1.




Figure B.1: (Search bounds for solution of Eq. (B.60)) **Case $g(a) > 0$:** Corresponds to Curve (a) defined as $\gamma_n g(a + xc)$, $c > 0$. The solution x_* of fixed point equation (B.60) (corresponding to right triangle) is between 0 and $\gamma_n g(a)$ since Curve (a) is decreasing. **Case $g(a) < 0$:** Corresponds to Curve (b) also defined as $\gamma_n g(a + xc)$. The solution x_* of fixed point equation (B.60) (left triangle) is between $\gamma_n g(a)$ and 0 since Curve (b) is also decreasing.

□

B.7 Additional experiments

B.7.1 Normality experiments with implicit SGD

In Figure B.2 we plot the experimental results of Section 2.6.1 for $p = 50$ (parameter dimension). We see that explicit SGD becomes even more unstable in more dimensions as expected. In contrast, implicit SGD remains stable and validates the theoretical normal distribution for small learning rates. In larger learning rates we observe a divergence from the asymptotic chi-squared distribution (e.g., $\gamma_1 = 6$) because when the learning rate parameter is large there is more noise in the stochastic approximations, and thus more iterations are required for convergence. In this experiment we fixed the number of iterations for each value of the learning rate, but subsequent experiments verified that implicit SGD reaches the theoretical chi-squared distribution if the number of iterations is increased. Finally, in Figure B.3 we make a similar plot for a logistic regression model. In this case the learning rates need to be larger because with the same distribution of covariates for X_n , the Fisher information is smaller than in the linear normal model. Thus, in almost all experiments explicit SGD was unstable and could not converge whereas implicit SGD was stable and followed the theoretical chi-squared distribution.

B.7.2 Poisson regression

In this experiment, we illustrate our method on a bivariate Poisson model which is simple enough to derive the variance formula analytically. We assume binary features such that, for any iteration n , X_n is either $(0,0)^\top$, $(1,0)^\top$ or $(0,1)^\top$ with probabilities 0.6, 0.2 and 0.2 respectively. We set $\theta_\star = (\theta_1, \theta_2)^\top$ for some θ_1, θ_2 , and assume $Y_n \sim \text{Poisson}(\exp(X_n^\top \theta_\star))$, where the transfer function h is the exponential, i.e., $h(x) = \exp(x)$. It follows,

$$\mathcal{I}(\theta_\star) = \mathbb{E} \left(h'(X_n^\top \theta_\star) X_n X_n^\top \right) = 0.2 \begin{pmatrix} e^{\theta_1} & 0 \\ 0 & e^{\theta_2} \end{pmatrix}.$$

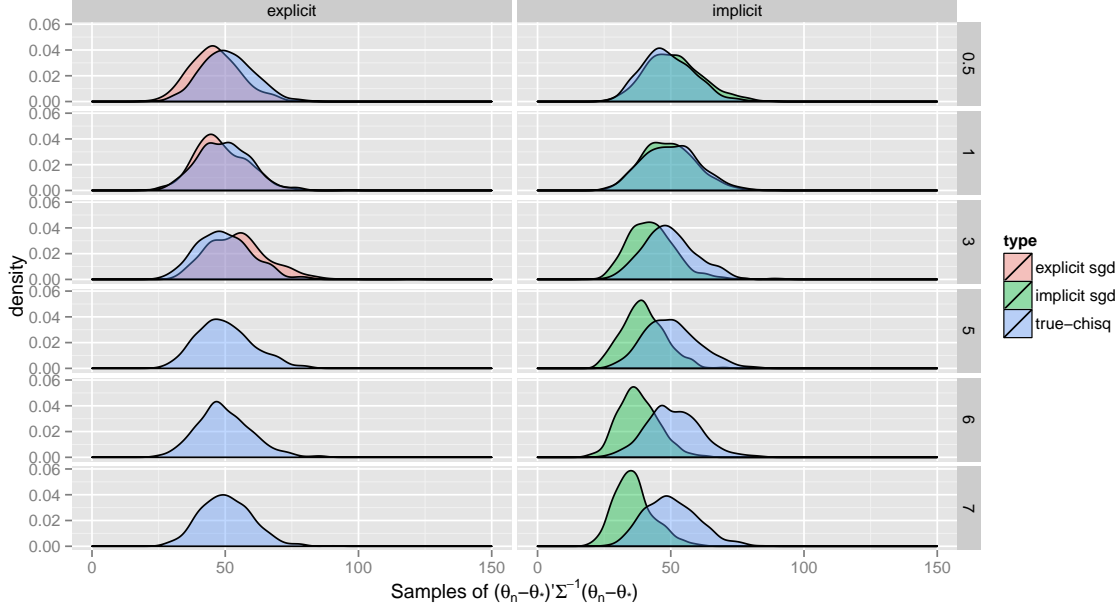


Figure B.2: Simulation with normal model for $p = 50$ parameters. Implicit SGD is stable and follows the nominal chi-squared distribution well, regardless of the particular learning rate. Explicit SGD becomes unstable at larger γ_1 and its distribution does not follow the theoretical distribution chi-squared distribution well. In particular, the distribution of $N(\theta_N^{\text{sgd}} - \theta_*)^\top \Sigma^{-1} (\theta_N^{\text{sgd}} - \theta_*)$ quickly becomes unstable for larger values of the learning rate parameter, and eventually diverges when $\gamma_1 > 3$.

We set $\gamma_n = 10/3n$ and $C_n = I$. Setting $\theta_1 = \log 2$ and $\theta_2 = \log 4$, the asymptotic variance Σ in Theorem 6 is equal to

$$\Sigma = \frac{2}{3} \begin{pmatrix} \frac{e^{\theta_1}}{(4/3)e^{\theta_1}-1} & 0 \\ 0 & \frac{e^{\theta_2}}{(4/3)e^{\theta_2}-1} \end{pmatrix} = \begin{pmatrix} 0.8 & 0 \\ 0 & 0.62 \end{pmatrix}. \quad (\text{B.61})$$

Next, we obtain 100 independent samples of θ_n^{sgd} and θ_n^{im} for $n = 20000$ iterations of procedures (2.3) and (2.4), and compute their empirical variances. We observe that the implicit estimates are particularly stable and have an empirical variance satisfying

$$(1/\gamma_n) \widehat{\text{Var}}(\theta_n^{\text{im}}) = \begin{pmatrix} 0.86 & -0.06 \\ -0.06 & 0.64 \end{pmatrix},$$

and that is close to the theoretical value (B.61). In contrast, the standard SGD estimates are quite unstable and their L_2 distance to the true values θ_* are orders of magnitude larger than the implicit ones (see Table B.1 for sample quantiles). By Lemma 1, such deviations are

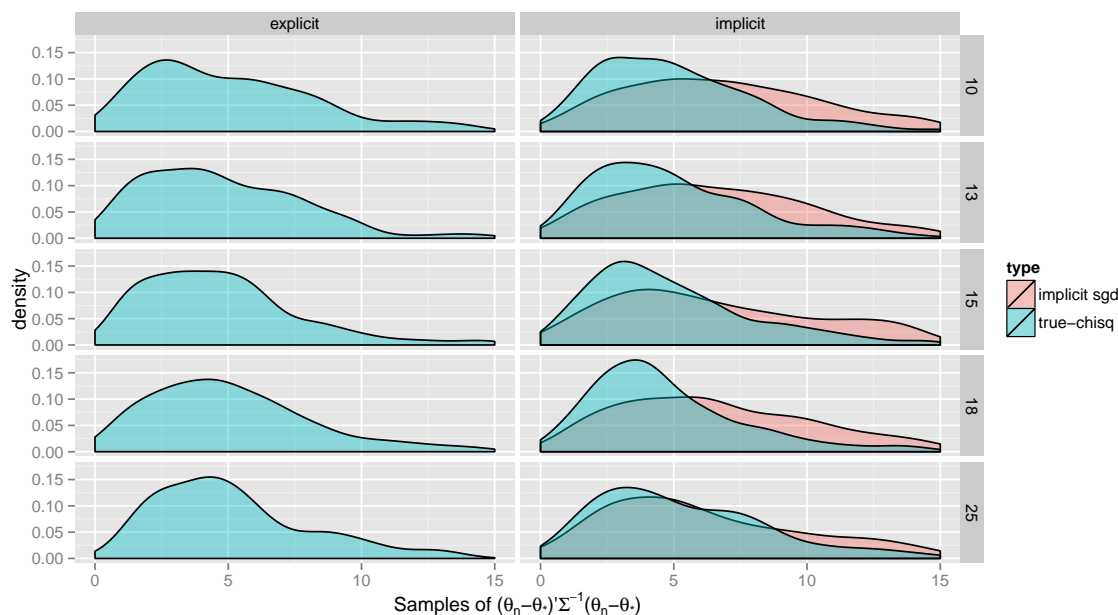


Figure B.3: Simulation with logistic regression model for $p = 5$. Learning rates are larger than in the linear normal model to ensure the asymptotic covariance matrix of Theorem 6 is positive definite. Implicit SGD is stable and follows the nominal chi-squared distribution regardless of the learning rate. Explicit SGD is unstable at virtually all replications of this experiment.

expected for standard SGD because the largest eigenvalue of $\mathcal{I}(\theta_*)$ is $\lambda_{(2)} = 0.8$ satisfying $\gamma_1 \lambda_{(2)} = 8/3 > 1$. Note that it is fairly straightforward to stabilize the standard SGD procedure in this problem, for example by modifying the learning rate sequence to $\gamma_n = \min\{0.15, 10/3n\}$. In general, when the optimization problem is well-understood, it is easy to determine the learning rate schedule that avoids out-of-bound explicit updates. In practice, however, we are working with problems that are not so well-understood and determining the correct learning rate parameters may take substantial effort, especially in multi-dimensional settings. The implicit method eliminates this overhead; a wide range of learning rates can lead to convergence on all problems.

B.7.3 Experiments with glmnet

In this section, we transform the outcomes in the original experiment Y through the logistic transformation and then fit a logistic regression model. The results are shown in Table B.2

Table B.1: Quantiles of $\|\theta_n^{\text{sgd}} - \theta_\star\|$ and $\|\theta_n^{\text{im}} - \theta_\star\|$. Values larger than $1e3$ are marked “*”.

METHOD	QUANTILES					
	25%	50%	75%	85%	95%	100%
SGD	0.01	1.3	435.8	*	*	*
IMPLICIT	0.00	0.01	0.02	0.02	0.03	0.04

which replicates and expands on Table 2 of (Friedman *et al.*, 2010). The implicit SGD method maintains a stable running time over different correlations and scales sub-linearly in the model size p . In contrast, `glmnet` is affected by the model size p and covariate correlation such that it remains 2x-10x slower across experiments. The SGD method is significantly slower in the logistic regression example compared to the normal case (Table B.2). This is because the implicit equation of Algorithm 1 needs to be solved numerically, whereas a closed-form solution is available in the normal case.

B.7.4 Support vector machines

In this experiment, we are interested to test the performance of the implicit procedure outside the family of GLMs. For that purpose, we implement an implicit online learning procedure for a SVM model and compare it to a standard SGD method on the RCV1 dataset, which is a typical large-scale machine learning benchmark.¹ Some results using variations on the loss functions and the regularization parameter are shown in Table B.3. A complete understanding of these results is still missing, however we do observe that the implicit method fares well compared to optimized explicit SGD and, at the same time, remains remarkably robust to misspecification. For example, in all experiments the standard SGD method degrades in performance for small or large regularization (in these experiments, the regularization parameter λ also affects the learning rate such that larger λ means larger learning rates). However, the implicit method maintains a more stable performance across

¹ We used Bottou’s SVM SGD implementation available at <http://leon.bottou.org/projects/sgd>. Our implicit SVM is available at the first author’s website.

Table B.2: Experiments comparing implicit SGD with `glmnet`. Covariates X are sampled as normal, with cross-correlation ρ , and the outcomes are sampled as $\mathbf{y} \sim \text{Binom}(\mathbf{p})$, $\text{logit}(\mathbf{p}) = \mathcal{N}(X\theta_*, \sigma^2 I)$. Running times (in secs) are reported for different values of ρ averaged over 10 repetitions.

METHOD	METRIC	CORRELATION (ρ)			
		0	0.2	0.6	0.9
<hr/> $N = 1000, p = 10$ <hr/>					
GLMNET	TIME(SECS)	0.02	0.02	0.026	0.051
	MSE	0.256	0.257	0.292	0.358
SGD	TIME(SECS)	0.058	0.058	0.059	0.062
	MSE	0.214	0.215	0.237	0.27
<hr/> $N = 5000, p = 50$ <hr/>					
GLMNET		0.182	0.193	0.279	0.579
		0.131	0.139	0.152	0.196
SGD		0.289	0.289	0.296	0.31
		0.109	0.108	0.116	0.14
<hr/> $N = 100000, p = 200$ <hr/>					
GLMNET		8.129	8.524	9.921	22.042
		0.06	0.061	0.07	0.099
SGD		5.455	5.458	5.437	5.481
		0.045	0.046	0.048	0.058

experiments and, interestingly, it achieves best performance under minimal regularization using the hinge loss.

B.7.5 Experiments with machine learning algorithms

In this section we perform additional experiments with related methods from the machine learning literature. We focus on averaged implicit SGD (2.13), which was shown to achieve optimality under suitable conditions, because most machine learning methods are also designed to achieve optimality (also known linear convergence rate), in the context of maximum-likelihood (or maximum a-posteriori) computation with a finite data set. In summary, our experiments include the following procedures:

Table B.3: Test errors of standard and implicit SGD methods on the RCV1 dataset benchmark. Training times are roughly comparable. Best scores, for a particular loss and regularization, are in bold.

LOSS		REGULARIZATION (λ)		
		1E-5	1E-7	1E-12
HINGE	SGD	4.65%	3.57%	4.85%
	IMPLICIT	4.68%	3.6%	3.46%
LOG	SGD	5.23%	3.87%	5.42%
	IMPLICIT	4.28%	3.69%	4.01%

- Explicit SGD procedure (2.3).
- Implicit SGD procedure (2.4).
- Averaged explicit SGD: Averaged stochastic gradient descent with explicit updates of the iterates (Xu, 2011; Shamir and Zhang, 2012; Bach and Moulines, 2013). This is equivalent to procedure (2.13) where the implicit update is replaced by an explicit one, $\theta_n^{\text{sgd}} = \theta_{n-1}^{\text{sgd}} + \gamma_n \nabla \log f(Y_n; X_n, \theta_{n-1}^{\text{sgd}})$.
- Prox-SVRG: A proximal version of the stochastic gradient descent with progressive variance reduction (SVRG) method (Xiao and Zhang, 2014).
- Prox-SAG: A proximal version of the stochastic average gradient (SAG) method (Schmidt *et al.*, 2013). While its theory has not been formally established, Prox-SAG has shown similar convergence properties to Prox-SVRG.²
- AdaGrad (Duchi *et al.*, 2011b) as defined in Eq. (2.11). We note that AdaGrad and similar adaptive methods effectively approximate the natural gradient by using a larger-dimensional learning rate. It has the added advantage of being less sensitive than first-order methods to tuning of hyperparameters.

²We note that the linear convergence rates for Prox-SVRG and Prox-SAG refer to convergence to the empirical minimizer (i.e., MLE or MAP), and not to the ground truth θ_* .

Averaged explicit SGD

In this experiment we validate the theory of statistical efficiency and stability of averaged implicit SGD. To do so, we follow a simple normal linear regression example from [Bach and Moulines \(2013\)](#), which is similar to the experiment in [Section 2.6.2](#). We set $N = 1e6$ as the number of observations, and $p = 20$ be the number of covariates. We also set $\theta_* = (0, 0, \dots, 0)^\top \in \mathbb{R}^{20}$ as the true parameter value. The random variables X_n are distributed i.i.d. as $X_n \sim \mathcal{N}_p(0, H)$, where H is a randomly generated symmetric matrix with eigenvalues $1/k$, for $k = 1, \dots, p$. The outcome Y_n is sampled from a normal distribution as $Y_n | X_n \sim \mathcal{N}(X_n^\top \theta_*, 1)$, for $n = 1, \dots, N$. We choose a constant learning rate $\gamma_n \equiv \gamma$ according to the average radius of the data $R^2 = \text{trace}(H)$, and for both averaged explicit and implicit SGD we collect iterates θ_n for $n = 1, \dots, N$, and keep the average $\bar{\theta}_n$. In [Figure B.4](#), we plot $(\theta_n - \theta_*)^\top H(\theta_n - \theta_*)$ for each iteration for a maximum of N iterations, i.e., a full pass over the data, in log-log space.

[Figure B.4](#) shows that averaged implicit SGD performs on par with averaged explicit SGD for the rates at which averaged explicit SGD is known to be optimal. Thus, averaged implicit SGD is also optimal. However, the benefit of the implicit procedure in averaged implicit SGD becomes clear as the learning rate deviates; notably, averaged implicit SGD remains stable for learning rates that are above the theoretical threshold, i.e., $\gamma > 1/R^2$, whereas averaged explicit SGD diverges in the case of $\gamma = 2/R^2$. This stable behavior is also exhibited in implicit SGD, but it converges at a slower rate than averaged implicit SGD, and thus cannot effectively combine stability with statistical efficiency. We note that stability of averaged implicit SGD is also observed in the same experiments using decaying learning rates.

Prox-SVRG and Prox-SAG

We now conduct a study of averaged implicit SGD's empirical performance on the standard benchmark of large scale linear classification on real data sets. For brevity, we display results on four data sets, although we have seen similar results on eight additional ones.

Table B.4: Summary of data sets and the L_2 regularization parameter λ used

	description	type	features	training set	test set	λ
covtype	forest cover type	sparse	54	464,809	116,203	10^{-6}
delta	synthetic data	dense	500	450,000	50,000	10^{-2}
rcv1	text data	sparse	47,152	781,265	23,149	10^{-5}
mnist	digit image features	dense	784	60,000	10,000	10^{-3}

Our datasets are summarized in Table B.4. The COVTYPE data set (Blackard, 1998) consists of forest cover types in which the task is to classify class 2 among 7 forest cover types. DELTA is synthetic data offered in the PASCAL Large Scale Challenge (Sonnenburg *et al.*, 2008) and we apply the default processing offered by the challenge organizers. The task in RCV1 is to classify documents belonging to class CCAT in the text dataset (Lewis *et al.*, 2004), where we apply the standard preprocessing provided by Bottou (2012). In the MNIST data set (Le Cun *et al.*, 1998) of images of handwritten digits, the task is to classify digit 9 against all others.

For averaged implicit SGD and averaged explicit SGD, we use the learning rate $\gamma_n = \eta_0(1 + \lambda\eta_0n)^{-3/4}$ prescribed in Xu (2011), where the constant η_0 is determined using a small subset of the data. Hyperparameters for other methods are set based on a computationally intensive grid search over the entire hyperparameter space: for PROX-SVRG, this includes the step size η in the proximal update and the inner iteration count m , and for PROX-SAG, the same step size η .

The results are shown in Figure B.5. We see that averaged implicit SGD achieves comparable performance with the tuned proximal methods PROX-SVRG and PROX-SAG, as well as AdaGrad. All methods have a comparable convergence rate and take roughly a single pass in order to converge. AdaGrad exhibits a larger variance in its estimate than the proximal methods, which can be explained from our theoretical results in Section 2.4.2. We also note that as averaged implicit SGD achieves comparable results to the

other proximal methods, it also requires no tuning while `Prox-SVRG` and `Prox-SAG` do require careful tuning of their hyperparameters. This was confirmed from separate sensitivity analyses (not reported in this paper), which indicated that `aisgd` is robust to fine-tuning of hyperparameters in the learning rate, whereas small perturbations of hyperparameters in averaged explicit SGD (the learning rate), `Prox-SVRG` (proximal step size η and inner iteration m), and `Prox-SAG` (proximal step size η), can lead to arbitrarily bad error rates.




Figure B.4: Loss of averaged implicit SGD, averaged explicit SGD, and plain implicit SGD (2.4) ($C_n = I$), on simulated multivariate normal data with $N = 1e6$ observations $p = 20$ features. The plot shows that averaged implicit SGD is stable regardless of the specification of the learning rate γ and without sacrificing performance. In contrast, explicit averaged SGD is very sensitive to misspecification of the learning rate.

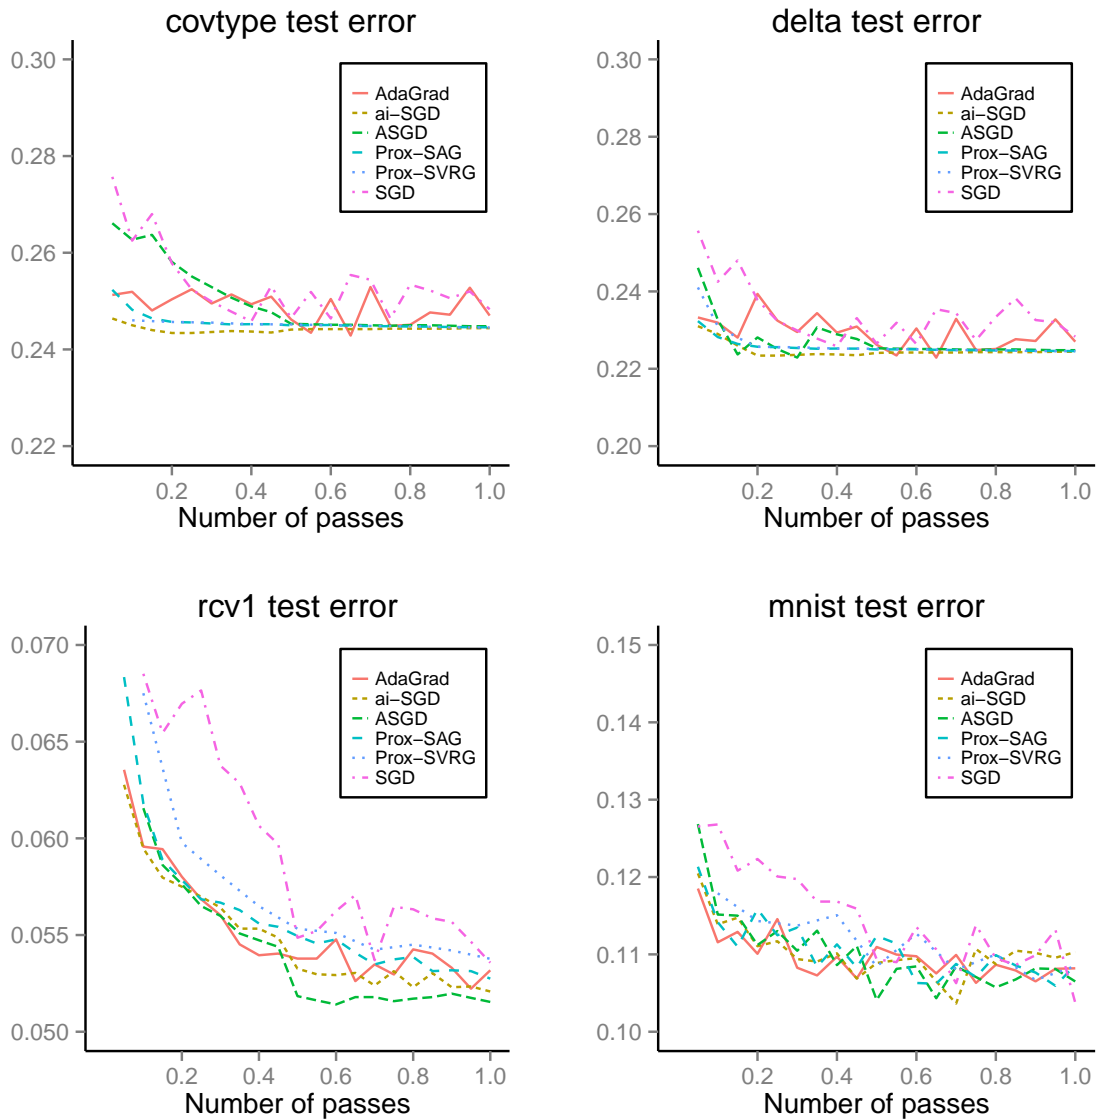


Figure B.5: Large scale linear classification with log loss on four data sets. Each plot indicates the test error of various stochastic gradient methods over a single pass of the data.