

Scalable Kernel Methods via Doubly Stochastic Gradients

Bo Dai¹, Bo Xie¹, Niao He², Yingyu Liang¹, Anant Raj¹, Maria-Florina Balcan³, Le Song¹

¹ College of Computing, Georgia Institute of Technology
{bodai, bxie33, yliang39, araj34}@gatech.edu, lsong@cc.gatech.edu

² School of Industrial & Systems Engineering, Georgia Institute of Technology
nhe6@gatech.edu

³ School of Computer Science, Carnegie Mellon University
ninamf@cs.cmu.edu

September 24, 2014

Abstract

The general perception is that kernel methods are not scalable, and neural nets are the methods of choice for large-scale nonlinear learning problems. Or have we simply not tried hard enough for kernel methods? Here we propose an approach that scales up kernel methods using a novel concept called “*doubly stochastic functional gradients*”. Our approach relies on the fact that many kernel methods can be expressed as convex optimization problems, and we solve the problems by making *two unbiased* stochastic approximations to the functional gradient, one using random training points and another using random features associated with the kernel, and then descending using this noisy functional gradient. Our algorithm is simple, does *not* need to commit to a preset number of random features, and allows the flexibility of the function class to grow as we see more incoming data in the streaming setting. We show that a function learned by this procedure after t iterations converges to the optimal function in the reproducing kernel Hilbert space in rate $O(1/t)$, and achieves a generalization performance of $O(1/\sqrt{t})$. Our approach can readily scale kernel methods up to the regimes which are dominated by neural nets. We show that our method can achieve competitive performance to neural nets in datasets such as 2.3 million energy materials from MolecularSpace, 8 million handwritten digits from MNIST, and 1 million photos from ImageNet using convolution features.

1 Introduction

The general perception is that kernel methods are not scalable. When it comes to large-scale nonlinear learning problems, the methods of choice so far are neural nets where theoretical understanding remains incomplete. Are kernel methods really not scalable? Or is it simply because we have not tried hard enough, while neural nets have exploited sophisticated design of feature architectures, virtual example generation for dealing with invariance, stochastic gradient descent for efficient training, and GPUs for further speedup?

A bottleneck in scaling up kernel methods is the storage and computation of the kernel matrix, K , which is usually dense. Storing the matrix requires $O(n^2)$ space, and computing it takes $O(n^2d)$ operations, where n is the number of data points and d is the dimension. There have been many great attempts to scale up kernel methods, including efforts from numerical linear algebra, functional analysis, and numerical optimization perspectives.

A common numerical linear algebra approach is to approximate the kernel matrix using low-rank factors, $K \approx A^\top A$, with $A \in \mathbb{R}^{r \times n}$ and $\text{rank } r \leq n$. This low-rank approximation usually requires $O(nr^2 + nrd)$

operations, and then subsequent kernel algorithms can directly operate on A . Many works, such as Greedy basis selection techniques [1], Nyström approximation [2] and incomplete Cholesky decomposition [3], all followed this strategy. In practice, one observes that kernel methods with approximated kernel matrices often result in a few percentage of losses in performance. In fact, without further assumption on the regularity of the kernel matrix, the generalization ability after low-rank approximation is typically of the order $O(1/\sqrt{r} + 1/\sqrt{n})$ [4, 5], which implies that the rank needs to be nearly linear in the number of data points! Thus, in order for kernel methods to achieve the best generalization ability, the low-rank approximation based approaches quickly become impractical for big datasets due to their $O(n^3 + n^2d)$ preprocessing time and $O(n^2)$ memory requirement.

Random feature approximation is another popular approach for scaling up kernel methods [6, 7]. Instead of approximating the kernel matrix, the method directly approximates the kernel function using explicit feature maps. The advantage of this approach is that the random feature matrix for n data points can be computed in time $O(nrd)$ using $O(nr)$ memory, where r is the number of random features. Subsequent algorithms then only operate on an $O(nr)$ matrix. Similar to low-rank kernel matrix approximation approach, the generalization ability of random feature approach is of the order $O(1/\sqrt{r} + 1/\sqrt{n})$ [8, 9], which implies that the number of random features also needs to be $O(n)$. Another common drawback of these two approaches is that it is not easy to adapt the solution from a small r to a large r' . Often one is interested in increasing the kernel matrix approximation rank or the number of random features to obtain a better generalization ability. Then special procedures need to be designed to reuse the solution obtained from a small r , which is not straightforward.

Another approach that addresses the scalability issue rises from optimization perspective. One general strategy is to solve the dual forms of kernel methods using coordinate or block-coordinate descent (*e.g.*, [10, 11, 12]). By doing so, each iteration of the algorithm only incurs $O(nrd)$ computation and $O(nr)$ memory, where r is the size of the parameter block. A second strategy is to perform functional gradient descent by looking at a batch of data points at a time (*e.g.*, [13, 14]). Thus, the computation and memory requirements are also $O(nrd)$ and $O(nr)$ respectively in each iteration, where r is the batch size. These approaches can easily change to a different r without restarting the optimization and has no loss in generalization ability since they do not approximate the kernel matrix or function. However, a serious drawback of these approaches is that, without further approximation, all support vectors need to be kept for testing, which can be as big as the entire training set! (*e.g.*, kernel ridge regression and non-separable nonlinear classification problems.)

In summary, there exists a delicate trade-off between computation, memory and statistics if one wants to scale up kernel methods. Inspired by various previous efforts, we propose a simple yet general strategy to scale up many kernel methods using a novel concept called “*doubly stochastic functional gradients*”. Our method relies on the fact that most kernel methods can be expressed as convex optimization problems over functions in reproducing kernel Hilbert spaces (RKHS) and solved via functional gradient descent. Our algorithm proceeds by making *two unbiased* stochastic approximations to the functional gradient, one using random training points and the other one using random features associated with the kernel, and then descending using this noisy functional gradient. The key intuitions behind our algorithm originate from

- (i) the property of stochastic gradient descent algorithm that as long as the stochastic gradient is unbiased, the convergence of the algorithm is guaranteed [15]; and
- (ii) the property of pseudo-random number generators that the random samples can in fact be completely determined by an initial value (a seed).

We exploit these properties and enable kernel methods to achieve better balances between computation, memory and statistics. Our method interestingly combines kernel methods, functional analysis, stochastic optimization and algorithmic trick, and it possesses a number of desiderata:

Generality and simplicity. Our approach applies to many kernel methods, such as kernel ridge regression, support vector machines, logistic regression, two-sample test, and many different types of kernels, such as shift-invariant kernels, polynomial kernels, general inner product kernels, and so on. The algorithm can be summarized in just a few lines of code (Algorithm 1 and 2). For a different problem and kernel, we just need

to adapt the loss function and the random feature generator.

Flexibility. Different from previous uses of random features which typically prefix the number of features and then optimize over the feature weightings, our approach allows the number of random features, and hence the flexibility of the function class, to grow with the number of data points. This allows our method to be applicable to data streaming setting, which is not possible for previous random feature approach, and achieve the full potential of nonparametric methods.

Efficient computation. The key computation of our method is evaluating the doubly stochastic functional gradient, which involves the generation of the random features with specific random seeds and the evaluation of these random features on the small batch of data points. For iteration t , the computational complexity is $O(td)$.

Small memory. The doubly stochasticity also allows us to avoid keeping the support vectors which becomes prohibitive in large-scale streaming setting. Instead, we just need to keep a small program for regenerating the random features, and sample previously used random feature according to pre-specified random seeds. For iteration t , the memory needed is $O(t)$ independent of the dimension of the data.

Theoretical guarantees. We provide a novel and nontrivial analysis involving Hilbert space martingale and a newly proved recurrence relation, and show that the estimator produced by our algorithm, which might be outside of the RKHS, converges to the optimal RKHS function. More specifically, both in expectation and with high probability, our algorithm can estimate the optimal function in the RKHS in the rate of $O(1/t)$, and achieve a generalization bound of $O(1/\sqrt{t})$, which are indeed optimal [15]. The variance of the random features, introduced during our second approximation to the functional gradient, only contributes additively to the constant in the final convergence rate. These results are the first of the kind in kernel method literature, which could be of independent interest.

Strong empirical performance. Our algorithm can readily scale kernel methods up to the regimes which are previously dominated by neural nets. We show that our method compares favorably to other scalable kernel methods in medium scale datasets, and to neural nets in big datasets such as 8 million handwritten digits from MNIST, 2.3 million materials from MolecularSpace, and 1 million photos from ImageNet using convolution features. Our results suggest that kernel methods, theoretically well-grounded methods, can potentially replace neural nets in many large scale real-world problems where nonparametric estimation are needed.

In the remainder, we will first introduce preliminaries on kernel methods and functional gradients. We will then describe our algorithm and provide both theoretical and empirical supports.

2 Duality between Kernels and Random Processes

Kernel methods owe their name to the use of kernel functions, $k(x, x') : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$, which are symmetric positive definite (PD), meaning that for all $n > 1$, and $x_1, \dots, x_n \in \mathcal{X}$, and $c_1, \dots, c_n \in \mathbb{R}$, we have $\sum_{i,j=1}^n c_i c_j k(x_i, x_j) \geq 0$. There is an intriguing duality between kernels and stochastic processes which will play a crucial role in our later algorithm design. More specifically,

Theorem 1 (e.g., Devinatz [16]; Hein & Bousquet [17]) *If $k(x, x')$ is a PD kernel, then there exists a set Ω , a measure \mathbb{P} on Ω , and random feature $\phi_\omega(x) : \mathcal{X} \mapsto \mathbb{R}$ from $L_2(\Omega, \mathbb{P})$, such that $k(x, x') = \int_{\Omega} \phi_\omega(x) \phi_\omega(x') d\mathbb{P}(\omega)$.*

Essentially, the above integral representation relates the kernel function to a random process ω with measure $\mathbb{P}(\omega)$. Note that the integral representation may not be unique. For instance, the random process can be a Gaussian process on \mathcal{X} with the sample function $\phi_\omega(x)$, and $k(x, x')$ is simply the covariance function between two point x and x' . If the kernel is also continuous and shift invariant, i.e., $k(x, x') = k(x - x')$ for $x \in \mathbb{R}^d$, then the integral representation specializes into a form characterized by inverse Fourier transformation (e.g., [18, Theorem 6.6]),

Theorem 2 (Bochner) *A continuous, real-valued, symmetric and shift-invariant function $k(x - x')$ on \mathbb{R}^d is a PD kernel if and only if there is a finite non-negative measure $\mathbb{P}(\omega)$ on \mathbb{R}^d , such that $k(x - x') = \int_{\mathbb{R}^d} e^{i\omega^\top(x-x')} d\mathbb{P}(\omega) = \int_{\mathbb{R}^d \times [0, 2\pi]} 2 \cos(\omega^\top x + b) \cos(\omega^\top x' + b) d(\mathbb{P}(\omega) \times \mathbb{P}(b))$, where $\mathbb{P}(b)$ is a uniform distribution on $[0, 2\pi]$, and $\phi_\omega(x) = \sqrt{2} \cos(\omega^\top x + b)$.*

For Gaussian RBF kernel, $k(x - x') = \exp(-\|x - x'\|^2/2\sigma^2)$, this yields a Gaussian distribution $\mathbb{P}(\omega)$ with density proportional to $\exp(-\sigma^2\|\omega\|^2/2)$; for the Laplace kernel, this yields a Cauchy distribution; and for the Matern kernel, this yields the convolutions of the unit ball [19].

Similar representation where the explicit form of $\phi_\omega(x)$ and $\mathbb{P}(\omega)$ are known can also be derived for rotation invariant kernel, $k(x, x') = k(\langle x, x' \rangle)$, using Fourier transformation on sphere [19]. For polynomial kernels, $k(x, x') = (\langle x, x' \rangle + c)^p$, a random tensor sketching approach can also be used [20]. Explicit random features have been designed for many other kernels, such as dot product kernel [28], additive/multiplicative class of homogeneous kernels [29], *e.g.*, Hellinger's, χ^2 , exponentiated- χ^2 , Jensen-Shannon's and Intersection kernel, as well as kernels on Abelian semigroups [30].

Instead of finding the random process $\mathbb{P}(\omega)$ and function $\phi_\omega(x)$ given a kernel, one can go the reverse direction, and construct kernels from random processes and functions (*e.g.*, Wendland [18]).

Theorem 3 *If $k(x, x') = \int_{\Omega} \phi_\omega(x)\phi_\omega(x') d\mathbb{P}(\omega)$ for a nonnegative measure $\mathbb{P}(\omega)$ on Ω and $\phi_\omega(x) : \mathcal{X} \mapsto \mathbb{R}$ from $L_2(\Omega, \mathbb{P})$, then $k(x, x')$ is a PD kernel.*

For instance, $\phi_\omega(x) := \cos(\omega^\top \psi_\theta(x) + b)$, where $\psi_\theta(x)$ can be a random convolution of the input x parametrized by θ . Another important concept is the reproducing kernel Hilbert space (RKHS). An RKHS \mathcal{H} on \mathcal{X} is a Hilbert space of functions from \mathcal{X} to \mathbb{R} . \mathcal{H} is an RKHS if and only if there exists a $k(x, x') : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$ such that $\forall x \in \mathcal{X}, k(x, \cdot) \in \mathcal{H}$, and $\forall f \in \mathcal{H}, \langle f(\cdot), k(x, \cdot) \rangle_{\mathcal{H}} = f(x)$. If such a $k(x, x')$ exist, it is unique and it is a PD kernel. A function $f \in \mathcal{H}$ if and only if $\|f\|_{\mathcal{H}}^2 := \langle f, f \rangle_{\mathcal{H}} < \infty$, and its L_2 norm is dominated by RKHS norm $\|f\|_{L_2} \leq \|f\|_{\mathcal{H}}$.

3 Doubly Stochastic Functional Gradients

Many kernel methods can be written as convex optimizations over functions in the RKHS and solved using the functional gradient methods [13, 14]. Inspired by these previous works, we will introduce a novel concept called “*doubly stochastic functional gradients*” to address the scalability issue. Let $l(u, y)$ be a scalar (potentially non-smooth) loss function convex of $u \in \mathbb{R}$. Let the subgradient of $l(u, y)$ with respect to u be $l'(u, y)$. Given a PD kernel $k(x, x')$ and the associated RKHS \mathcal{H} , many kernel methods try to find a function $f_* \in \mathcal{H}$ which solves the optimization problem

$$\operatorname{argmin}_{f \in \mathcal{H}} R(f) := \mathbb{E}_{(x,y)}[l(f(x), y)] + \frac{\nu}{2} \|f\|_{\mathcal{H}}^2 \iff \operatorname{argmin}_{\|f\|_{\mathcal{H}} \leq B(\nu)} \mathbb{E}_{(x,y)}[l(f(x), y)] \quad (1)$$

where $\nu > 0$ is a regularization parameter, $B(\nu)$ is a non-increasing function of ν , and the data (x, y) follow a distribution $\mathbb{P}(x, y)$. The functional gradient $\nabla R(f)$ is defined as the linear term in the change of the objective after we perturb f by ϵ in the direction of g , *i.e.*,

$$R(f + \epsilon g) = R(f) + \epsilon \langle \nabla R(f), g \rangle_{\mathcal{H}} + O(\epsilon^2). \quad (2)$$

For instance, applying the above definition, we have $\nabla f(x) = \nabla \langle f, k(x, \cdot) \rangle_{\mathcal{H}} = k(x, \cdot)$, and $\nabla \|f\|_{\mathcal{H}}^2 = \nabla \langle f, f \rangle_{\mathcal{H}} = 2f$.

Stochastic functional gradient. Given a data point $(x, y) \sim \mathbb{P}(x, y)$ and $f \in \mathcal{H}$, the stochastic functional gradient of $\mathbb{E}_{(x,y)}[l(f(x), y)]$ with respect to $f \in \mathcal{H}$ is

$$\xi(\cdot) := l'(f(x), y)k(x, \cdot), \quad (3)$$

which is essentially a single data point approximation to the true functional gradient. Furthermore, for any $g \in \mathcal{H}$, we have $\langle \xi(\cdot), g \rangle_{\mathcal{H}} = l'(f(x), y)g(x)$. Inspired by the duality between kernel functions and random processes, we can make an additional approximation to the stochastic functional gradient using a random feature $\phi_\omega(x)$ sampled according to $\mathbb{P}(\omega)$. More specifically,

Algorithm 1: $\{\alpha_i\}_{i=1}^t = \mathbf{Train}(\mathbb{P}(x, y))$

Require: $\mathbb{P}(\omega)$, $\phi_\omega(x)$, $l(f(x), y)$, ν .

- 1: **for** $i = 1, \dots, t$ **do**
 - 2: Sample $(x_i, y_i) \sim \mathbb{P}(x, y)$.
 - 3: Sample $\omega_i \sim \mathbb{P}(\omega)$ with **seed** i .
 - 4: $f(x_i) = \mathbf{Predict}(x_i, \{\alpha_j\}_{j=1}^{i-1})$.
 - 5: $\alpha_i = -\gamma_i l'(f(x_i), y_i) \phi_{\omega_i}(x_i)$.
 - 6: $\alpha_j = (1 - \gamma_j \nu) \alpha_j$ for $j = 1, \dots, i - 1$.
 - 7: **end for**
-

Algorithm 2: $f(x) = \mathbf{Predict}(x, \{\alpha_i\}_{i=1}^t)$

Require: $\mathbb{P}(\omega)$, $\phi_\omega(x)$.

- 1: Set $f(x) = 0$.
 - 2: **for** $i = 1, \dots, t$ **do**
 - 3: Sample $\omega_i \sim \mathbb{P}(\omega)$ with **seed** i .
 - 4: $f(x) = f(x) + \alpha_i \phi_{\omega_i}(x)$.
 - 5: **end for**
-

Doubly stochastic functional gradient. Let $\omega \sim \mathbb{P}(\omega)$, then the doubly stochastic gradient of $\mathbb{E}_{(x,y)}[l(f(x), y)]$ with respect to $f \in \mathcal{H}$ is

$$\zeta(\cdot) := l'(f(x), y) \phi_\omega(x) \phi_\omega(\cdot). \quad (4)$$

Note that the stochastic functional gradient $\xi(\cdot)$ is in RKHS \mathcal{H} but $\zeta(\cdot)$ may be outside \mathcal{H} , since $\phi_\omega(\cdot)$ may be outside the RKHS. For instance, for the Gaussian RBF kernel, the random feature $\phi_\omega(x) = \sqrt{2} \cos(\omega^\top x + b)$ is outside the RKHS associated with the kernel function.

However, these functional gradients are related by $\xi(\cdot) = \mathbb{E}_\omega[\zeta(\cdot)]$, which lead to unbiased estimators of the original functional gradient, *i.e.*,

$$\nabla R(f) = \mathbb{E}_{(x,y)}[\xi(\cdot)] + \nu f(\cdot), \quad \text{and} \quad \nabla R(f) = \mathbb{E}_{(x,y)} \mathbb{E}_\omega[\zeta(\cdot)] + \nu f(\cdot). \quad (5)$$

We emphasize that the source of randomness associated with the random feature is not present in the data, but artificially introduced by us. This is crucial for the development of our scalable algorithm in the next section. Meanwhile, it also creates additional challenges in the analysis of the algorithm which we will deal with carefully.

4 Doubly Stochastic Kernel Machines

The first key intuition behind our algorithm originates from the property of stochastic gradient descent algorithm that as long as the stochastic gradient is unbiased, the convergence of the algorithm is guaranteed [15]. In our algorithm, we will exploit this property and introduce *two* sources of randomness, one from data and another artificial, to scale up kernel methods.

The second key intuition behind our algorithm is that the random features used in the doubly stochastic functional gradients will be sampled according to *pseudo-random number generators*, where the sequences of apparently random samples can in fact be completely determined by an initial value (a seed). Although these random samples are not the “true” random sample in the purest sense of the word, however they suffice for our task in practice.

More specifically, our algorithm proceeds by making two unbiased stochastic approximation to the functional gradient in each iteration, and then descending using this noisy functional gradient. The overall algorithms for training and prediction is summarized in Algorithm 1 and 2. The training algorithm essentially just performs random feature sampling and doubly stochastic gradient evaluation, and maintains a collection of real number $\{\alpha_i\}$, which is computationally efficient and memory friendly. A crucial step in the algorithm is to sample the random features with “seed i ”. The seeds have to be aligned between training and prediction, and with the corresponding α_i obtained from each iteration. The learning rate γ_t in the algorithm needs to be chosen as $O(1/t)$, as shown by our later analysis to achieve the best rate of convergence. For now, we assume that we have access to the data generating distribution $\mathbb{P}(x, y)$. This can be modified to sample uniformly randomly from a fixed dataset, without affecting the algorithm and the later convergence analysis. Let the sampled data and random feature parameters be $\mathcal{D}^t := \{(x_i, y_i)\}_{i=1}^t$ and $\omega^t := \{\omega_i\}_{i=1}^t$ respectively after t iteration, the function obtained by Algorithm 1 is a simple additive form of the doubly

stochastic functional gradients

$$f_{t+1}(\cdot) = f_t(\cdot) - \gamma_t(\zeta_t(\cdot) + \nu f_t(\cdot)) = \sum_{i=1}^t a_t^i \zeta_i(\cdot), \quad \forall t > 1, \quad \text{and} \quad f_1(\cdot) = 0, \quad (6)$$

where $a_t^i = -\gamma_i \prod_{j=i+1}^t (1 - \gamma_j \nu)$ are deterministic values depending on the step sizes $\gamma_j (i \leq j \leq t)$ and regularization parameter ν . This simple form makes it easy for us to analyze its convergence.

We note that our algorithm can also take a mini-batch of points and random features at each step, and estimate an empirical covariance for preconditioning to achieve potentially better performance.

5 Theoretical Guarantees

In this section, we will show that, both in expectation and with high probability, our algorithm can estimate the optimal function in the RKHS with rate $O(1/t)$, and achieve a generalization bound of $O(1/\sqrt{t})$. The analysis for our algorithm has a new twist compared to previous analysis of stochastic gradient descent algorithms, since the random feature approximation results in an estimator which is outside the RKHS. Besides the analysis for stochastic functional gradient descent, we need to use martingales and the corresponding concentration inequalities to prove that the sequence of estimators, f_{t+1} , outside the RKHS converge to the optimal function, f_* , in the RKHS. We make the following standard assumptions ahead for later references:

- A. There exists an optimal solution, denoted as f_* , to the problem of our interest (1).
- B. Loss function $\ell(u, y) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is L -Lipschitz continuous in terms of the first argument,

$$\forall u, u', y \in \mathbb{R}, |\ell(u, y) - \ell(u', y)| \leq L|u - u'|.$$

- C. For any data $\{(x_i, y_i)\}_{i=1}^t$ and any trajectory $\{f_i(\cdot)\}_{i=1}^t$, there exists $M > 0$, such that $|\ell'(f_i(x_i), y_i)| \leq M$. Note in our situation M exists and $M < \infty$ since we assume bounded domain and the functions f_t we generate are always bounded as well.

- D. There exists $\kappa > 0$ and $\phi > 0$, such that $k(x, x') \leq \kappa$, $|\phi_\omega(x)\phi_\omega(x')| \leq \phi, \forall x, x' \in \mathcal{X}, \omega \in \Omega$. For example, when $k(\cdot, \cdot)$ is the Gaussian RBF kernel, we have $\kappa = 1$, $\phi = 2$.

We now present our main theorems as below. Due to the space restrictions, we will only provide a short sketch of proofs here. The full proofs for these theorems are given in the appendix.

Theorem 4 (Convergence in expectation) When $\gamma_t = \frac{\theta}{t}$ with $\theta > 0$ such that $\theta\nu \in (1, 2) \cup \mathbb{Z}_+$,

$$\mathbb{E}_{\mathcal{D}^t, \omega^t} [|f_{t+1}(x) - f_*(x)|^2] \leq \frac{2C^2 + 2\kappa Q_1^2}{t}, \quad \text{for any } x \in \mathcal{X}$$

where $Q_1 = \max \left\{ \|f_*\|_{\mathcal{H}}, (Q_0 + \sqrt{Q_0^2 + (2\theta\nu - 1)(1 + \theta\nu)^2 \theta^2 \kappa M^2}) / (2\nu\theta - 1) \right\}$, with $Q_0 = 2\sqrt{2}\kappa^{1/2}(\kappa + \phi)LM\theta^2$, and $C^2 = 4(\kappa + \phi)^2 M^2 \theta^2$.

Theorem 5 (Convergence with high probability) When $\gamma_t = \frac{\theta}{t}$ with $\theta > 0$ such that $\theta\nu \in \mathbb{Z}_+$, for any $x \in \mathcal{X}$, we have with probability at least $1 - 3\delta$ over $(\mathcal{D}^t, \omega^t)$,

$$|f_{t+1}(x) - f_*(x)|^2 \leq \frac{C^2 \ln(2/\delta)}{t} + \frac{2\kappa Q_2^2 \ln(2t/\delta) \ln^2(t)}{t},$$

where C is as above and $Q_2 = \max \left\{ \|f_*\|_{\mathcal{H}}, Q_0 + \sqrt{Q_0^2 + \kappa M^2 (1 + \theta\nu)^2 (\theta^2 + 16\theta/\nu)} \right\}$, with $Q_0 = 4\sqrt{2}\kappa^{1/2}M\theta(8 + (\kappa + \phi)\theta L)$.

Proof sketch: We focus on the convergence in expectation; the high probability bound can be established in a similar fashion. The main technical difficulty is that f_{t+1} may not be in the RKHS \mathcal{H} . The key of the proof is then to construct an intermediate function h_{t+1} , such that the difference between f_{t+1} and h_{t+1} and the difference between h_{t+1} and f_* can be bounded. More specifically,

$$h_{t+1}(\cdot) = h_t(\cdot) - \gamma_t(\xi_t(\cdot) + \nu h_t(\cdot)) = \sum_{i=1}^t a_t^i \xi_i(\cdot), \quad \forall t > 1, \quad \text{and} \quad h_1(\cdot) = 0, \quad (7)$$

where $\xi_t(\cdot) = \mathbb{E}_{\omega_t}[\zeta_t(\cdot)]$. Then for any x , the error can be decomposed as two terms

$$|f_{t+1}(x) - f_*(x)|^2 \leq 2 \underbrace{|f_{t+1}(x) - h_{t+1}(x)|^2}_{\text{error due to random features}} + 2\kappa \underbrace{\|h_{t+1} - f_*\|_{\mathcal{H}}^2}_{\text{error due to random data}}$$

For the error term due to random features, h_{t+1} is constructed such that $f_{t+1} - h_{t+1}$ is a martingale, and the stepsizes are chosen such that $|a_t^i| \leq \frac{\theta}{t}$, which allows us to bound the martingale. In other words, the choices of the stepsizes keep f_{t+1} close to the RKHS. For the error term due to random data, since $h_{t+1} \in \mathcal{H}$, we can now apply the standard arguments for stochastic approximation in the RKHS. Due to the additional randomness, the recursion is slightly more complicated, $e_{t+1} \leq (1 - \frac{2\nu\theta}{t})e_t + \frac{\beta_1}{t}\sqrt{\frac{e_t}{t}} + \frac{\beta_2}{t^2}$, where $e_{t+1} = \mathbb{E}_{\mathcal{D}^t, \omega^t}[\|h_{t+1} - f_*\|_{\mathcal{H}}^2]$, and β_1 and β_2 depends on the related parameters. Solving this recursion then leads to a bound for the second error term. ■

Theorem 6 (Generalization bound) *Let the true risk be $R_{true}(f) = \mathbb{E}_{(x,y)}[l(f(x), y)]$. Then with probability at least $1 - 3\delta$ over $(\mathcal{D}^t, \omega^t)$, and C and Q_2 defined as previously*

$$R_{true}(f_{t+1}) - R_{true}(f_*) \leq \frac{(C\sqrt{\ln(8\sqrt{e_t}/\delta)} + \sqrt{2\kappa}Q_2\sqrt{\ln(2t/\delta)\ln(t)})L}{\sqrt{t}}.$$

Proof By the Lipschitz continuity of $l(\cdot, y)$ and Jensen's Inequality, we have

$$R_{true}(f_{t+1}) - R_{true}(f_*) \leq L\mathbb{E}_x|f_{t+1}(x) - f_*(x)| \leq L\sqrt{\mathbb{E}_x|f_{t+1}(x) - f_*(x)|^2} = L\|f_{t+1} - f_*\|_2.$$

Again, $\|f_{t+1} - f_*\|_2$ can be decomposed as two terms $O(\|f_{t+1} - h_{t+1}\|_2^2)$ and $O(\|h_{t+1} - f_*\|_{\mathcal{H}}^2)$, which can be bounded similarly as in Theorem 5 (see Corollary 12 in the appendix). ■

Remarks. The overall rate of convergence in expectation, which is $O(1/t)$, is indeed optimal. Classical complexity theory (see, e.g. reference in [15]) shows that to obtain ϵ -accuracy solution, the number of iterations needed for the stochastic approximation is $\Omega(1/\epsilon)$ for strongly convex case and $\Omega(1/\epsilon^2)$ for general convex case. Different from the classical setting of stochastic approximation, our case imposes not one but two sources of randomness/stochasticity in the gradient, which intuitively speaking, might require higher order number of iterations for general convex case. However, the variance of the random features only contributes *additively* to the constant in the final convergence rate. Therefore, our method is still able to achieve the same rate as in the classical setting. The rate of the generalization bound is also nearly optimal up to log factors. Notice that these bounds are achieved by adopting the classical stochastic gradient algorithm, and they may be further refined with more sophisticated techniques and analysis. For example, techniques for reducing variance of SGD proposed in [31], mini-batch and preconditioning can be used to reduce the constant factors in the bound significantly. Theorem 4 also reveals bounds in L_∞ and L_2 sense as in Section A.2 in the appendix. The choices of stepsizes γ_t and the tuning parameters given in these bounds are only for sufficient conditions and simple analysis; other choices can also lead to bounds in the same order.

6 Computation, Memory and Statistics Trade-off

To investigate computation, memory and statistics trade-off, we will fix the desired L_2 error in the function estimation to ϵ , i.e., $\|f - f_*\|_2^2 \leq \epsilon$, and work out the dependency of other quantities on ϵ . These other quantities include the preprocessing time, the number of samples and random features (or rank), the number of iterations of each algorithm, and the computational cost and memory requirement for learning and prediction. We assume that the number of samples, n , needed to achieve the prescribed error ϵ is of the order $O(1/\epsilon)$, the same for all methods. Furthermore, we make no other regularity assumption about margin properties or the kernel matrix such as fast spectrum decay. Thus the required number of random feature (or ranks), r , will be of the order $O(n) = O(1/\epsilon)$ [4, 5, 8, 9].

We will pick a few representative algorithms for comparison, namely, (i) NORMA [13]: kernel methods trained with stochastic functional gradients; (ii) k-SDCA [12]: kernelized version of stochastic dual coordinate ascend; (iii) r-SDCA: first approximate the kernel function with random features, and then run stochastic dual coordinate ascend; (iv) n-SDCA: first approximate the kernel matrix using Nyström's method, and then run stochastic dual coordinate ascend; similarly we will combine Pegasos algorithm [21], stochastic block mirror descent (SBMD) [32], and random block coordinate descent (RBCD) [33] with random fea-

tures and Nyström’s method, and obtain (v) r-Pegasos, (vi) n-Pegasos, (vii) r-SBMD, (viii) n-SBMD, (ix) r-RBCD, and (x) n-RBCD, respectively. The comparisons are summarized below in Table. 1¹

Table 1: Comparison of Computation and Memory Requirements

Algorithms	Preprocessing Computation	Total Computation Cost		Total Memory Cost	
		Training	Prediction	Training	Prediction
Doubly SGD	$O(1)$	$O(d/\epsilon^2)$	$O(d/\epsilon)$	$O(1/\epsilon)$	$O(1/\epsilon)$
NORMA	$O(1)$	$O(d/\epsilon^2)$	$O(d/\epsilon)$	$O(d/\epsilon)$	$O(d/\epsilon)$
k-SDCA	$O(1)$	$O(d/\epsilon^2)$	$O(d/\epsilon)$	$O(d/\epsilon)$	$O(d/\epsilon)$
r-SDCA	$O(1)$	$O(d/\epsilon^2)$	$O(d/\epsilon)$	$O(1/\epsilon)$	$O(1/\epsilon)$
n-SDCA	$O(1/\epsilon^3)$	$O(d/\epsilon^2)$	$O(d/\epsilon)$	$O(1/\epsilon)$	$O(1/\epsilon)$
r-Pegasos	$O(1)$	$O(d/\epsilon^2)$	$O(d/\epsilon)$	$O(1/\epsilon)$	$O(1/\epsilon)$
n-Pegasos	$O(1/\epsilon^3)$	$O(d/\epsilon^2)$	$O(d/\epsilon)$	$O(1/\epsilon)$	$O(1/\epsilon)$
r-SBMD	$O(1)$	$O(d/\epsilon^2)$	$O(d/\epsilon)$	$O(1/\epsilon)$	$O(1/\epsilon)$
n-SBMD	$O(1/\epsilon^3)$	$O(d/\epsilon^2)$	$O(d/\epsilon)$	$O(1/\epsilon)$	$O(1/\epsilon)$
r-RBCD	$O(1)$	$O(d/\epsilon^2 \log(\frac{1}{\epsilon}))$	$O(d/\epsilon)$	$O(1/\epsilon)$	$O(1/\epsilon)$
n-RBCD	$O(1/\epsilon^3)$	$O(d/\epsilon^2 \log(\frac{1}{\epsilon}))$	$O(d/\epsilon)$	$O(1/\epsilon)$	$O(1/\epsilon)$

From Table 1, one can see that our method, r-SDCA, r-Pegasos, r-SBMD and r-RBCD achieve the best dependency on the dimension, d , of the data up to a log factor. However, often one is interested in increasing the number of random features as more data points are observed to obtain a better generalization ability, *e.g.*, **in streaming setting**. Then special procedures need to be designed for updating the r-SDCA, r-Pegasos, r-SBMD and r-RBCD solutions, which is not clear how to do easily and efficiently with theoretical guarantees. As a more refined comparison, our algorithm is also the cheapest in terms of per training iteration computation and memory requirement. We list the computational and memory requirements at a particular iteration $t < n$ for these five algorithms to achieve ϵ error in Table 2.

Table 2: Comparison of Computation and Memory Requirement Per Iteration. b denotes the block size in algorithms SBMD and RBCD.

Algorithms	Computation per Iteration	Memory per Iteration	Iteration #
Doubly SGD	$\Theta(dt + t + t)$	$\Theta(t)$	$O(1/\epsilon)$
r-SDCA	$\Theta(dn + n + n)$	$\Theta(n)$	$O(1/\epsilon)$
r-Pegasos	$\Theta(dn + n + n)$	$\Theta(n)$	$O(1/\epsilon)$
r-SBMD	$\Theta(dn + n + n/b)$	$\Theta(n)$	$O(b/\epsilon)$
r-RBCD	$\Theta(dn^2 + n^2 + n/b)$	$\Theta(n)$	$O(\log(1/\epsilon))$

7 Experiments

We show that our method compares favorably to other scalable kernel methods in medium scale datasets, and neural nets in large scale datasets. Below is a summary of the datasets used. A “yes” for the last column means that virtual examples are generated from datasets for training. K-ridge stands for kernel ridge regression; K-SVM stands for kernel SVM; K-logistic stands for kernel logistic regression.

¹We only considered general kernel algorithms in this section. For some specific loss functions, *e.g.*, hinge-loss, there are algorithms proposed to achieve better memory saving with extra training cost, such as support vector reduction technique [34].

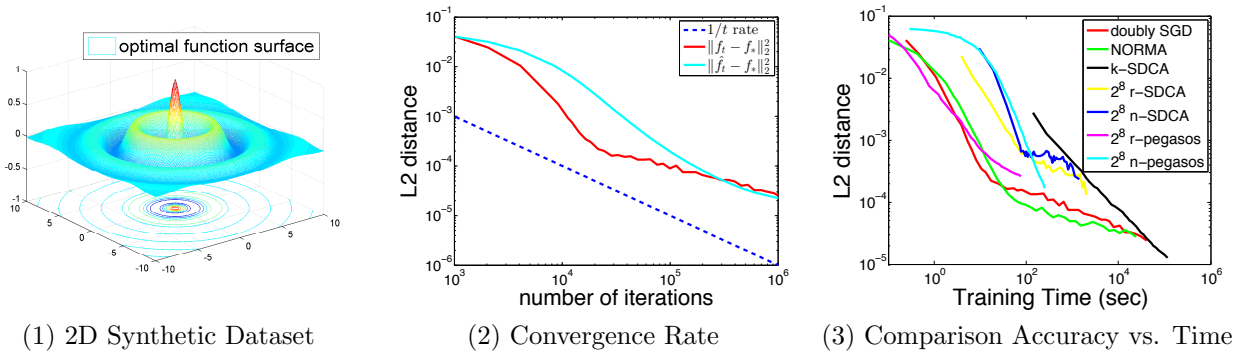


Figure 1: Experimental results for kernel ridge regression on synthetic dataset.

Table 3: Datasets

	Name	Model	# of samples	Input dim	Output range	Virtual
(1)	Synthetic	K-ridge	2^{20}	2	$[-1, 1.3]$	no
(2)	Adult	K-SVM	32K	123	$\{-1, 1\}$	no
(3)	MNIST 8M 8 vs. 6 [22]	K-SVM	1.6M	784	$\{-1, 1\}$	yes
(4)	Forest	K-SVM	0.5M	54	$\{-1, 1\}$	no
(5)	MNIST 8M [22]	K-logistic	8M	1568	$\{0, \dots, 9\}$	yes
(6)	CIFAR 10 [23]	K-logistic	60K	2304	$\{0, \dots, 9\}$	yes
(7)	ImageNet [24]	K-logistic	1.3M	9216	$\{0, \dots, 999\}$	yes
(8)	QuantumMachine [25]	K-ridge	6K	276	$[-800, -2000]$	yes
(9)	MolecularSpace [25]	K-ridge	2.3M	2850	$[0, 13]$	no

Experiment settings. For datasets (1) – (4), we compare with the first seven algorithms for solving kernel methods discussed in Table 1. For the algorithms based on low rank kernel matrix approximation and random features, *i.e.*, pegasos and SDCA, we set the rank r or number of random features r to be 2^8 . We use the same batch size for both our algorithms and the competitors. We adopted two stopping criteria for different purposes. We first stopped the algorithms when they pass through the entire dataset once (SC1). This stopping criterion is designed for justifying our motivation. By investigating the performances of these algorithms with different levels of random feature approximations but the same number of training samples, we could identify that the bottleneck of the performances of the vanilla methods with explicit feature will be their approximation ability. To further demonstrate the advantages of the proposed algorithm in computational cost, we also conduct experiments on datasets (2) – (4) running the competitors within the same time budget as the proposed algorithm (SC2). We do not count the preprocessing time of Nyström’s method for n-Pegasos and n-SDCA, though it takes substantial amount of time. The algorithms are executed on the machine with AMD 16 2.4GHz Opteron CPUs and 200G memory. It should be noticed that this gives advantage to NORMA and k-SDCA which could save all the data in the memory.

For datasets (5) – (7), we compare with neural nets for images (“jointly-trained”). In order to directly compare the performance of nonlinear classifiers rather than feature learning abilities, we also use the convolution layers of a trained neural net to extract features, then apply our algorithm and a nonlinear neural net on top to learn classifiers (“fixed”). For datasets (8) and (9), we compare with the neural net described in [25] and use exactly the same input. In all the experiments, we select the batch size so that for each update, the computation resources could be utilized efficiently.

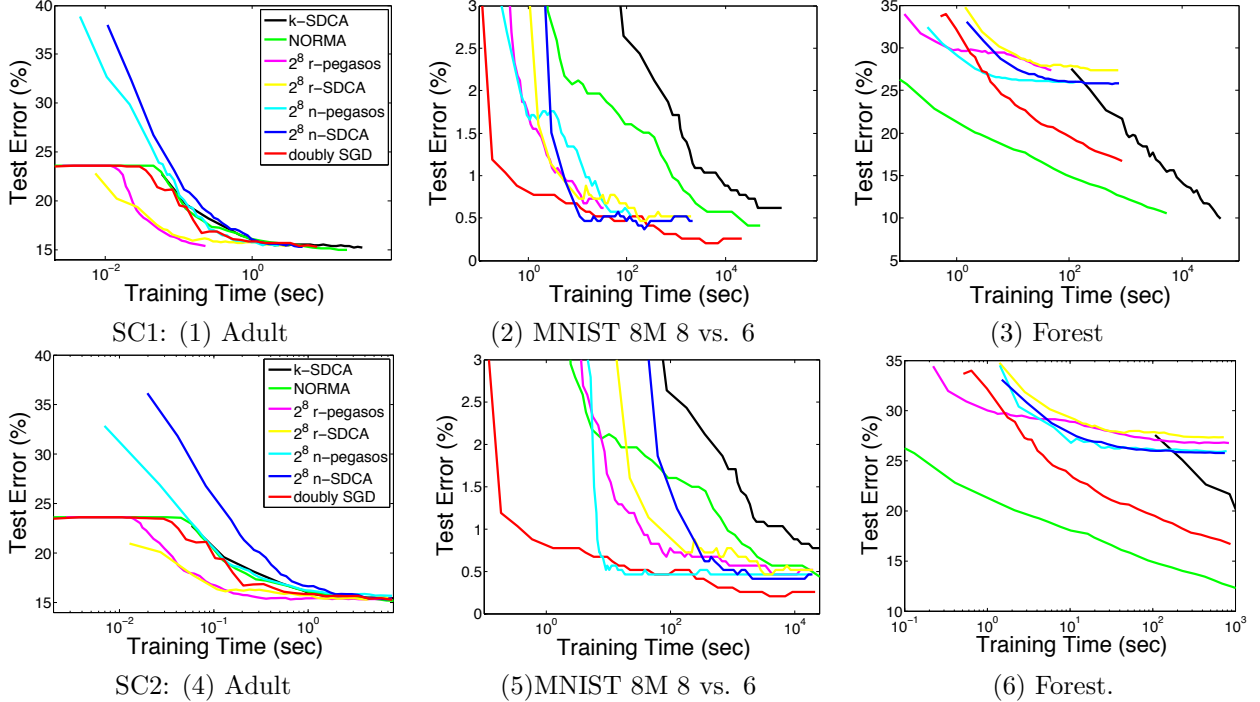


Figure 2: Comparison with other kernel SVM solvers on datasets (2) – (4) with two different stopping criteria.

7.1 Regression Comparisons on Synthetic Dataset

In this section, we compare our approach with alternative algorithms for kernel ridge regression on 2D synthetic dataset. The data are generated by

$$y = \cos(0.5\pi\|x\|_2) \exp(-0.1\pi\|x\|_2) + 0.1e$$

where $x \in [-10, 10]^2$ and $e \sim \mathcal{N}(0, 1)$. We use Gaussian RBF kernel with kernel bandwidth σ chosen to be 0.1 times the median of pairwise distances between data points (median trick). The regularization parameter ν is set to be 10^{-6} . The batch size and feature block are set to be 2^{10} .

The results are shown in Figure 1. In Figure 1(1), we plot the optimal functions generating the data. We justify our proof of the convergence rate in Figure 1(2). The blue dotted line is a convergence rate of $1/t$ as a guide. \hat{f}_t denotes the average solution after t -iteration, *i.e.*, $\hat{f}_t(x) = \frac{1}{t} \sum_{i=1}^t f_i(x)$. It could be seen that our algorithm indeed converges in the rate of $O(1/t)$. In Figure 1 (3), we compare the first seven algorithms listed in the Table 1 for solving the kernel ridge regression.

The comparison on synthetic dataset demonstrates the advantages of our algorithm clearly. Our algorithm achieves comparable performance with NORMA, which uses full kernel, in similar time but only costs $O(n)$ memory while NORMA costs $O(dn)$. The pegasos and SDCA using 2^8 random or Nyström features perform worse.

7.2 Classification Comparisons with Kernel SVM Algorithms

We evaluate our algorithm solving kernel SVM on three datasets (2)–(4) comparing with other several algorithms listed in Table 1 using stopping criteria SC1 and SC2.

Adult. We use Gaussian RBF kernel with kernel bandwidth obtained by median trick. The regularization parameter ν is set to be $1/(100n)$ where n is the number of training samples. We set the batch size to be 2^6 and feature block to be 2^5 . After going through the whole dataset one pass, the best error rate is

achieved by NORMA and k-SDCA which is 15% while our algorithm achieves comparable result 15.3%. The performances are illustrated in Figure 2(1). Under the same time budget, all the algorithms perform similarly in Figure 2(4). The reason of flat region of r-pegasos, NORMA and the proposed method on this dataset is that Adult dataset is unbalanced. There are about 24% positive samples while 76% negative samples.

MNIST 8M 8 vs. 6. We first reduce the dimension to 50 by PCA and use Gaussian RBF kernel with kernel bandwidth $\sigma = 9.03$ obtained by median trick. The regularization parameter ν is set to be $1/n$ where n is the number of training samples. We set the batch size to be 2^{10} and feature block to be 2^8 . The results are shown in Figure 2(2) and (5) under SC1 and SC2 respectively. Under both these two stopping criteria, our algorithm achieves the best test error 0.26% using similar training time.

Forest. We use Gaussian RBF kernel with kernel bandwidth obtained by median trick. The regularization parameter ν is set to be $1/n$ where n is the number of training samples. We set the batch size to be 2^{10} and feature block to be 2^8 . In Figure 2(3), we shows the performances of all algorithms using SC1. NORMA and k-SDCA achieve the best error rate, which is 10%, while our algorithm achieves around 15%, but still much better than the pegasos and SDCA with 2^8 features. In the same time budget, the proposed algorithm performs better than all the alternatives except NORMA in Figure 2(6).

As seen from the performance of pegasos and SDCA on Adult and MNIST, using fewer features does not deteriorate the classification error. This might be because there are cluster structures in these two binary classification datasets. Thus, they prefer low rank approximation rather than full kernel. Different from these two datasets, in the forest dataset, algorithms with full kernel, *i.e.*, NORMA and k-SDCA, achieve best performance. With more random features, our algorithm performs much better than pegasos and SDCA under both SC1 and SC2. Our algorithm is preferable for this scenario, *i.e.*, huge dataset with sophisticated decision boundary. Although utilizing full kernel could achieve better performance, the computation and memory requirement for the kernel on huge dataset are costly. To learn the sophisticated boundary while still considering the computational and memory cost, we need to efficiently approximate the kernel in $O(\frac{1}{\epsilon})$ with $O(n)$ random features. Our algorithm could handle so many random features efficiently in both computation and memory cost, while for pegasos and SDCA such operation is prohibitive.

7.3 Classification Comparisons to Convolution Neural Networks

We also compare our algorithm with the state-of-the-art neural network. In these experiments, the block size is set to be $O(10^4)$. Compared to the number of samples, $O(10^8)$, this block size is reasonable.

MNIST 8M. In this experiment, we compare to a variant of LeNet-5 [27], where all tanh units are replaced with rectified linear units. We also use more convolution filters and a larger fully connected layer. Specifically, the first two convolutions layers have 16 and 32 filters, respectively, and the fully connected layer contains 128 neurons. We use kernel logistic regression for the task. We extract features from the last max-pooling layer with dimension 1568, and use Gaussian RBF kernel with kernel bandwidth σ equaling to four times the median pairwise distance. The regularization parameter ν is set to be 0.0005.

The result is shown in Figure 3(1). As expected, the neural net with pre-learned features is faster to train than the jointly-trained one. However, our method is much faster compared to both methods. In addition, it achieves a lower error rate (0.5%) compared to the 0.6% error provided by the neural nets.

CIFAR 10. In this experiment, we compare to a neural net with two convolution layers (after contrast normalization and max-pooling layers) and two local layers that achieves 11% test error². The features are extracted from the top max-pooling layer from a trained neural net with 2304 dimension. We use kernel logistic regression for this problem. The kernel bandwidth σ for Gaussian RBF kernel is again four times the median pairwise distance. The regularization parameter ν is set to be 0.0005. We also perform a PCA (without centering) to reduce the dimension to 256 before feeding to our method.

²The specification is at <https://code.google.com/p/cuda-convnet/>

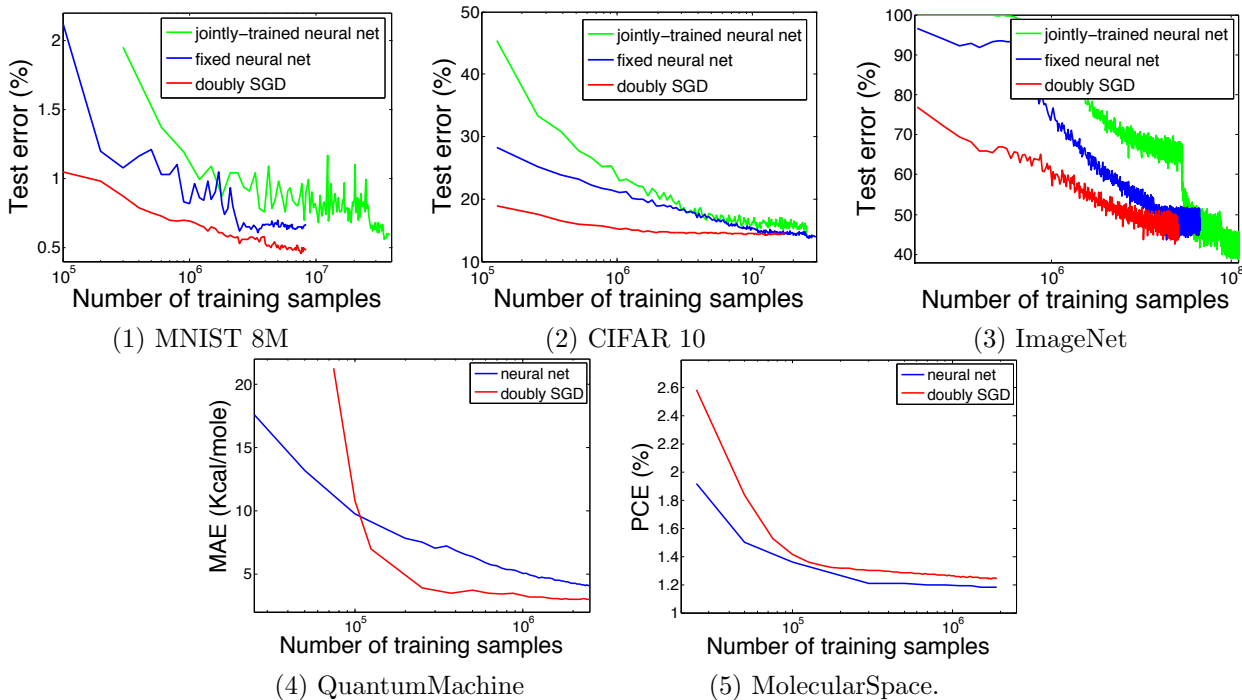


Figure 3: Comparison with Neural Networks on datasets (5) – (9).

The result is shown in Figure 3(2). The test error for our method drops significantly faster in the earlier phase, then gradually converges to that achieved by the neural nets. Our method is able to produce the same performance within a much restricted time budget.

ImageNet. In this experiment, we compare our algorithm with the neural nets on the ImageNet 2012 dataset, which contains 1.3 million color images from 1000 classes. Each image is of size 256×256 , and we randomly crop a 240×240 region with random horizontal flipping. The jointly-trained neural net is Alex-net [24]. The 9216 dimension features for our classifier and fixed neural net are from the last pooling layer of the jointly-trained neural net. The kernel bandwidth σ for Gaussian RBF kernel is again four times the median pairwise distance. The regularization parameter ν is set to be 0.0005.

Test error comparisons are shown in Figure 3(3). Our method achieves a test error of 44.5% by further max-voting of 10 transformations of the test set while the jointly-trained neural net arrives at 42% (without variations in color and illumination). At the same time, fixed neural net can only produce an error rate of 46% with max-voting. There may be some advantages to train the network jointly such that the layers work together to achieve a better objective. Although there is still a gap to the best performance by the jointly-trained neural net, our method comes very close with much faster convergence rate. Moreover, it achieves superior performance than the neural net trained with pre-learned features, both in accuracy and speed.

7.4 Regression Comparisons to Neural Networks

We test our algorithm for kernel ridge regression with neural network proposed in [25] on two large-scale real-world regression datasets, (8) and (9) in Table 3. To our best knowledge, this is the first comparison between kernel ridge regression and neural network on the dataset MolecularSpace.

QuantumMachine. In this experiment, we use the same binary representations converted based on random Coulomb matrices as in [25]. We first generate a set of randomly sorted coulomb matrices for each

molecule. And then, we break each dimension of the Coulomb matrix apart into steps and convert them to the binary predicates. Predictions are made by taking average of all prediction made on various Coulomb matrices of the same molecule. For this experiment, 40 sets of randomly permuted matrices are generated for each training example and 20 for each test example. We use Gaussian kernel with kernel bandwidth $\sigma = 60$ obtained by median trick. The batch size is set to be 50000 and the feature block is 2^{11} . The total dimension of random features is 2^{20} .

The results are shown in Figure 3(4). In QuantumMachine dataset, our method achieves Mean Absolute Error (MAE) of 2.97 kcal/mole, outperforming neural nets results, 3.51 kcal/mole. Note that this result is already close to the 1 kcal/mole required for chemical accuracy.

MolecularSpace. In this experiment, the task is to predict the power conversion efficiency (PCE) of the molecule. This dataset of 2.3 million molecular motifs is obtained from the Clean Energy Project Database. We use the same feature representation as for “QuantumMachine” dataset [25]. We set the kernel bandwidth of Gaussian RBF kernel to be 290 by median trick. The batch size is set to be 25000 and the feature block is 2^{11} . The total dimension of random features is 2^{20} .

The results are shown in Figure 3(5). It could be seen that our method is comparable with neural network on this 2.3 million dataset.

8 Conclusion

Our work contributes towards making kernel methods scalable for large-scale datasets. Specifically, by introducing artificial randomness associated with kernels besides the random data samples, we propose doubly stochastic functional gradient for kernel machines which makes the kernel machines efficient in both computation and memory requirement. Our algorithm successfully reduces the memory requirement of kernel machines from $O(dn)$ to $O(n)$. Meanwhile, we also show that our algorithm achieves the optimal rate of convergence, $O(1/t)$, for stochastic optimization. We compare our algorithm on both classification and regression problems with the state-of-the-art neural networks on several large-scale datasets. With our efficient algorithm, kernel methods could perform comparable to sophisticated-designed neural network empirically.

Acknowledgement

M.B. is supported in part by NSF grant CCF-1101283, AFOSR grant FA9550-09-1-0538, a Microsoft Faculty Fellowship, and a Raytheon Faculty Fellowship. L.S. is supported in part by NSF IIS-1116886, NSF/NIH BIGDATA 1R01GM108341, NSF CAREER IIS-1350983, and a Raytheon Faculty Fellowship.

References

- [1] A. J. Smola and B. Schölkopf. Sparse greedy matrix approximation for machine learning. In *ICML*, pages 911–918, San Francisco, 2000. Morgan Kaufmann Publishers.
- [2] C. K. I. Williams and M. Seeger. Using the Nystrom method to speed up kernel machines. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *NIPS*, 2000.
- [3] S. Fine and K. Scheinberg. Efficient SVM training using low-rank kernel representations. *JMLR*, 2:243–264, 2001.
- [4] P. Drineas and M. Mahoney. On the nystrom method for approximating a gram matrix for improved kernel-based learning. *JMLR*, 6:2153–2175, 2005.

- [5] Corinna Cortes, Mehryar Mohri, and Ameet Talwalkar. On the impact of kernel approximation on learning accuracy. In *AISTATS*, pages 113–120, 2010.
- [6] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In J.C. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *NIPS*. MIT Press, Cambridge, MA, 2008.
- [7] Q.V. Le, T. Sarlos, and A. J. Smola. Fastfood — computing hilbert space expansions in loglinear time. In *ICML*, 2013.
- [8] Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. In *NIPS*, 2009.
- [9] David Lopez-Paz, Suvrit Sra, A. J. Smola, Zoubin Ghahramani, and Bernhard Schölkopf. Randomized nonlinear component analysis. In *ICML*, 2014.
- [10] John C. Platt. Sequential minimal optimization: A fast algorithm for training support vector machines. Technical Report MSR-TR-98-14, Microsoft Research, 1998.
- [11] T. Joachims. Making large-scale SVM learning practical. In B. Schölkopf, C. J. C. Burges, and A. J. Smola, editors, *Advances in Kernel Methods — Support Vector Learning*, pages 169–184, Cambridge, MA, 1999. MIT Press.
- [12] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss. *JMLR*, 14(1):567–599, 2013.
- [13] J. Kivinen, A. J. Smola, and R. C. Williamson. Online learning with kernels. *IEEE Transactions on Signal Processing*, 52(8), Aug 2004.
- [14] N. Ratliff and J. Bagnell. Kernel conjugate gradient for fast kernel machines. In *IJCAI*, volume 20, January 2007.
- [15] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM J. on Optimization*, 19(4):1574–1609, January 2009.
- [16] A. Devinatz. Integral representation of pd functions. *Trans. AMS*, 74(1):56–77, 1953.
- [17] M. Hein and O. Bousquet. Kernels, associated structures, and generalizations. Technical Report 127, Max Planck Institute for Biological Cybernetics, 2004.
- [18] H. Wendland. *Scattered Data Approximation*. Cambridge University Press, Cambridge, UK, 2005.
- [19] Bernhard Schölkopf and A. J. Smola. *Learning with Kernels*. MIT Press, Cambridge, MA, 2002.
- [20] N. Pham and R. Pagh. Fast and scalable polynomial kernels via explicit feature maps. In *KDD*. ACM, 2013.
- [21] Shai Shalev-Shwartz, Yoram Singer, and Nathan Srebro. Pegasos: Primal estimated sub-gradient solver for SVM. In *ICML*, 2007.
- [22] G. Loosli, S. Canu, and L. Bottou. Training invariant support vector machines with selective sampling. In L. Bottou, O. Chapelle, D. DeCoste, and J. Weston, editors, *Large Scale Kernel Machines*, pages 301–320. MIT Press, 2007.
- [23] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [24] A. Krizhevsky, I. Sutskever, and G. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

- [25] Grégoire Montavon, Katja Hansen, Siamac Fazli, Matthias Rupp, Franziska Biegler, Andreas Ziehe, Alexandre Tkatchenko, Anatole von Lilienfeld, and Klaus-Robert Müller. Learning invariant representations of molecules for atomization energy prediction. In *NIPS*, pages 449–457, 2012.
- [26] Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. In *ICML*, pages 449–456, 2012.
- [27] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, November 1998.
- [28] Purushottam Kar and Harish Karnick. Random feature maps for dot product kernels. In Neil D. Lawrence and Mark A. Girolami, editors, *AISTATS-12*, volume 22, pages 583–591, 2012.
- [29] Andrea Vedaldi and Andrew Zisserman. Efficient additive kernels via explicit feature maps. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(3):480–492, 2012.
- [30] Jiyang Yang, Vikas Sindhwani, Quanfu Fan, Haim Avron, and Michael W. Mahoney. Random laplace feature maps for semigroup kernels on histograms. In *CVPR*, 2014.
- [31] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *NIPS*, pages 315–323, 2013.
- [32] Cong D. Dang and Guanghui Lan. Stochastic block mirror descent methods for nonsmooth and stochastic optimization. Technical report, University of Florida, 2013.
- [33] Yurii Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.
- [34] Andrew Cotter, Shai Shalev-Shwartz, and Nati Srebro. Learning optimally sparse support vector machines. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 266–274, 2013.

Appendix

A Proof Details

A.1 Convergence Rate

We first provide specific bounds and detailed proofs for the two error terms appeared in Theorem 4 and Theorem 5.

A.1.1 Error due to random features

Lemma 7 *We have*

(i) For any $x \in \mathcal{X}$, $\mathbb{E}_{\mathcal{D}^t, \omega^t} [|f_{t+1}(x) - h_{t+1}(x)|^2] \leq B_{1,t+1}^2 := 4M^2(\kappa + \phi)^2 \sum_{i=1}^t |a_i^i|^2$.

(ii) For any $x \in \mathcal{X}$, with probability at least $1 - \delta$ over $(\mathcal{D}^t, \omega^t)$,

$$|f_{t+1}(x) - h_{t+1}(x)|^2 \leq B_{2,t+1}^2 := 2M^2(\kappa + \phi)^2 \ln\left(\frac{2}{\delta}\right) \sum_{i=1}^t |a_i^i|^2$$

Proof Let $V_i(x) = V_i(x; \mathcal{D}^i, \omega^i) := a_t^i (\zeta_i(x) - \xi_i(x))$. Since $V_i(x)$ is a function of $(\mathcal{D}^i, \omega^i)$ and

$$\mathbb{E}_{\mathcal{D}^i, \omega^i} [V_i(x) | \omega^{i-1}] = a_t^i \mathbb{E}_{\mathcal{D}^i, \omega^i} [\zeta_i(x) - \xi_i(x) | \omega^{i-1}] = a_t^i \mathbb{E}_{\mathcal{D}^i, \omega^i} [\mathbb{E}_{\omega^i} [\zeta_i(x) - \xi_i(x) | \omega^{i-1}]] = 0,$$

we have that $\{V_i(x)\}$ is a martingal difference sequence. Further note that

$$|V_i(x)| \leq c_i = 2M(\phi + \kappa) |a_t^i|.$$

Then by Azuma's Inequality, for any $\epsilon > 0$,

$$\Pr_{\mathcal{D}^t, \omega^t} \left\{ \left| \sum_{i=1}^t V_i(x) \right| \geq \epsilon \right\} \leq 2 \exp \left\{ -\frac{2\epsilon^2}{\sum_{i=1}^t c_i^2} \right\}$$

which is equivalent as

$$\Pr_{\mathcal{D}^t, \omega^t} \left\{ \left(\sum_{i=1}^t V_i(x) \right)^2 \geq \ln(2/\delta) \sum_{i=1}^t c_i^2 / 2 \right\} \leq \delta.$$

Moreover,

$$\mathbb{E}_{\mathcal{D}^t, \omega^t} \left[\left(\sum_{i=1}^t V_i(x) \right)^2 \right] = \int_0^\infty \Pr_{\mathcal{D}^t, \omega^t} \left\{ \left(\sum_{i=1}^t V_i(x) \right)^2 \geq \epsilon \right\} d\epsilon = \int_0^\infty 2 \exp \left\{ -\frac{2\epsilon}{\sum_{i=1}^t c_i^2} \right\} d\epsilon = \sum_{i=1}^t c_i^2$$

Since $f_{t+1}(x) - h_{t+1}(x) = \sum_{i=1}^t V_i(x)$, we immediately obtain the two parts of the lemma. ■

Lemma 8 *Suppose $\gamma_i = \frac{\theta}{i}$ ($1 \leq i \leq t$) and $\theta\nu \in (1, 2) \cup \mathbb{Z}_+$. Then we have*

(1) $|a_t^i| \leq \frac{\theta}{i}$. Consequently, $\sum_{i=1}^t (a_t^i)^2 \leq \frac{\theta^2}{t}$.

(2) $\sum_{i=1}^t \gamma_i |a_t^i| \leq \begin{cases} \frac{\theta^2(\ln(t)+1)}{t}, & \text{if } \theta\nu \in [1, 2), \\ \frac{\theta^2}{t}, & \text{if } \theta\nu \in [2, +\infty) \cap \mathbb{Z}_+ \end{cases}$.

Proof (1) follows by induction on i . $|a_t^t| \leq \frac{\theta}{t}$ is trivially true. We have

$$|a_t^i| = |a_t^{i+1} \frac{\gamma_i}{\gamma_{i+1}} (1 - \nu\gamma_{i+1})| = \frac{i+1}{i} \left| 1 - \frac{\nu\theta}{i+1} \right| \cdot |a_t^{i+1}| = \left| \frac{i+1-\nu\theta}{i} \right| \cdot |a_t^{i+1}|.$$

When $\nu\theta \in (1, 2)$, $i - 1 < i + 1 - \nu\theta < i$ for any $i \geq 1$, so $|a_t^i| < |a_t^{i+1}| \leq \frac{\theta}{t}$. When $\nu\theta \in \mathbb{Z}_+$, if $i > \nu\theta - 1$, then $|a_t^i| < |a_t^{i+1}| \leq \frac{\theta}{t}$; if $i \leq \nu\theta - 1$, then $|a_t^i| = 0$. For (2), when $\theta\nu \in [1, 2)$,

$$\sum_{i=1}^t \gamma_t |a_t^i| = \sum_{i=1}^t \frac{\theta^2}{i^2} \cdot \frac{i+1-\theta\nu}{i+1} \cdots \frac{t-\theta\nu}{t} \leq \sum_{i=1}^t \frac{\theta^2}{i^2} \cdot \frac{i}{i+1} \cdots \frac{t-1}{t} \leq \sum_{i=1}^t \frac{\theta^2}{it} \leq \frac{\theta^2(\ln(t)+1)}{t}.$$

When $\theta\nu \in \mathbb{Z}_+$ and $2 \leq \theta\nu \leq t$,

$$\sum_{i=1}^t \gamma_t |a_t^i| = \sum_{i=2}^t \frac{\theta^2}{i^2} \cdot \frac{i+1-\theta\nu}{i+1} \cdots \frac{t-\theta\nu}{t} \leq \sum_{i=1}^t \frac{\theta^2}{i^2} \cdot \frac{i-1}{i+1} \cdots \frac{t-2}{t} \leq \sum_{i=2}^t \frac{\theta^2(i-1)}{it(t-1)} \leq \frac{\theta^2}{t}.$$

■

A.1.2 Error due to random data

Lemma 9 Assume $l'(u, y)$ is L -Lipschitz continuous in terms of $u \in \mathbb{R}$. Let f_* be the optimal solution to our target problem. Then

(i) If we set $\gamma_t = \frac{\theta}{t}$ with θ such that $\theta\nu \in (1, 2) \cup \mathbb{Z}_+$, then

$$\mathbb{E}_{\mathcal{D}^t, \omega^t} \left[\|h_{t+1} - f_*\|_{\mathcal{H}}^2 \right] \leq \frac{Q_1^2}{t},$$

where

$$Q_1 = \max \left\{ \|f_*\|_{\mathcal{H}}, \frac{Q_0 + \sqrt{Q_0^2 + (2\theta\nu - 1)(1 + \theta\nu)^2 \theta^2 \kappa M^2}}{2\nu\theta - 1} \right\}, Q_0 = 2\sqrt{2}\kappa^{1/2}(\kappa + \phi)LM\theta^2.$$

Particularly, if $\theta\nu = 1$, we have $Q_1 \leq \max \left\{ \|f_*\|_{\mathcal{H}}, 4\sqrt{2}((\kappa + \phi)L + \nu) \cdot \frac{\kappa^{1/2}M}{\nu^2} \right\}$.

(ii) If we set $\gamma_t = \frac{\theta}{t}$ with θ such that $\theta\nu \in \mathbb{Z}_+$, then with probability at least $1 - 2\delta$ over $(\mathcal{D}^t, \omega^t)$,

$$\|h_{t+1} - f_*\|_{\mathcal{H}}^2 \leq Q_2^2 \frac{\ln(2t/\delta) \ln(t)}{t}.$$

where

$$Q_2 = \max \left\{ \|f_*\|_{\mathcal{H}}, Q_0 + \sqrt{Q_0^2 + \kappa M^2(1 + \theta\nu)^2(\theta^2 + 16\theta/\nu)} \right\}, Q_0 = 4\sqrt{2}\kappa^{1/2}M\theta(8 + (\kappa + \phi)\theta L).$$

Particularly, if $\theta\nu = 1$, we have $Q_2 \leq \max \left\{ \|f_*\|_{\mathcal{H}}, 8\sqrt{2}((\kappa + \phi)L + 9\nu) \cdot \frac{\kappa^{1/2}M}{\nu^2} \right\}$.

Proof For the sake of simple notations, let us first denote the following three different gradient terms, which are

$$\begin{aligned} g_t &= \xi_t + \nu h_t = l'(f_t(x_t), y_t)k(x_t, \cdot) + \nu h_t, \\ \hat{g}_t &= \hat{\xi}_t + \nu h_t = l'(h_t(x_t), y_t)k(x_t, \cdot) + \nu h_t, \\ \bar{g}_t &= \mathbb{E}_{\mathcal{D}_t} [\hat{g}_t] = \mathbb{E}_{\mathcal{D}_t} [l'(h_t(x_t), y_t)k(x_t, \cdot)] + \nu h_t. \end{aligned}$$

Note that by our previous definition, we have $h_{t+1} = h_t - \gamma_t g_t, \forall t \geq 1$.

Denote $A_t = \|h_t - f_*\|_{\mathcal{H}}^2$. Then we have

$$\begin{aligned} A_{t+1} &= \|h_t - f_* - \gamma_t g_t\|_{\mathcal{H}}^2 \\ &= A_t + \gamma_t^2 \|g_t\|_{\mathcal{H}}^2 - 2\gamma_t \langle h_t - f_*, g_t \rangle_{\mathcal{H}} \\ &= A_t + \gamma_t^2 \|g_t\|_{\mathcal{H}}^2 - 2\gamma_t \langle h_t - f_*, \bar{g}_t \rangle_{\mathcal{H}} + 2\gamma_t \langle h_t - f_*, \bar{g}_t - \hat{g}_t \rangle_{\mathcal{H}} + 2\gamma_t \langle h_t - f_*, \hat{g}_t - g_t \rangle_{\mathcal{H}} \end{aligned}$$

Because of the strongly convexity of (1) and optimality condition, we have

$$\langle h_t - f_*, \bar{g}_t \rangle_{\mathcal{H}} \geq \nu \|h_t - f_*\|_{\mathcal{H}}^2$$

Hence, we have

$$A_{t+1} \leq (1 - 2\gamma_t\nu)A_t + \gamma_t^2 \|g_t\|_{\mathcal{H}}^2 + 2\gamma_t \langle h_t - f_*, \bar{g}_t - \hat{g}_t \rangle_{\mathcal{H}} + 2\gamma_t \langle h_t - f_*, \hat{g}_t - g_t \rangle_{\mathcal{H}}, \forall t \geq 1 \quad (8)$$

Proof for (i): Let us denote $\mathcal{M}_t = \|g_t\|_{\mathcal{H}}^2$, $\mathcal{N}_t = \langle h_t - f_*, \bar{g}_t - \hat{g}_t \rangle_{\mathcal{H}}$, $\mathcal{R}_t = \langle h_t - f_*, \hat{g}_t - g_t \rangle_{\mathcal{H}}$. We first show that $\mathcal{M}_t, \mathcal{N}_t, \mathcal{R}_t$ are bounded. Specifically, we have for $t \geq 1$,

- (1) $\mathcal{M}_t \leq \kappa M^2(1 + \nu c_t)^2$, where $c_t = \sqrt{\sum_{i,j=1}^{t-1} |a_{t-1}^i| \cdot |a_{t-1}^j|}$ for $t \geq 2$ and $c_1 = 0$;
- (2) $\mathbb{E}_{\mathcal{D}^t, \omega^t}[\mathcal{N}_t] = 0$;
- (3) $\mathbb{E}_{\mathcal{D}^t, \omega^t}[\mathcal{R}_t] \leq \kappa^{1/2} L B_{1,t} \sqrt{\mathbb{E}_{\mathcal{D}^{t-1}, \omega^{t-1}}[A_t]}$, where $B_{1,t}^2 := 4M^2(\kappa + \phi)^2 \sum_{i=1}^{t-1} |a_{t-1}^i|^2$ for $t \geq 2$ and $B_{1,1} = 0$;

We prove these results separately in Lemma 10 below. Let us denote $e_t = \mathbb{E}_{\mathcal{D}^{t-1}, \omega^{t-1}}[A_t]$, given the above bounds, we arrive at the following recursion,

$$e_{t+1} \leq (1 - 2\gamma_t\nu)e_t + \kappa M^2 \gamma_t^2 (1 + \nu c_t)^2 + 2\kappa^{1/2} L \gamma_t B_{1,t} \sqrt{e_t}. \quad (9)$$

When $\gamma_t = \theta/t$ with θ such that $\theta\nu \in (1, 2) \cup \mathbb{Z}_+$, from Lemma 8, we have $|a_t^i| \leq \frac{\theta}{t}, \forall 1 \leq i \leq t$. Consequently, $c_t \leq \theta$ and $B_{1,t}^2 \leq 4M^2(\kappa + \phi) \frac{\theta^2}{t-1}$. Applying these bounds leads to the refined recursion as follows

$$e_{t+1} \leq \left(1 - \frac{2\nu\theta}{t}\right) e_t + \kappa M^2 \frac{\theta^2}{t^2} (1 + \nu\theta)^2 + 2\kappa^{1/2} L \frac{\theta}{t} \sqrt{4M^2(\kappa + \phi)^2 \frac{\theta^2}{t-1}} \sqrt{e_t}$$

that can be further written as

$$e_{t+1} \leq \left(1 - \frac{2\nu\theta}{t}\right) e_t + \frac{\beta_1}{t} \sqrt{\frac{e_t}{t}} + \frac{\beta_2}{t^2},$$

where $\beta_1 = 4\sqrt{2}\kappa^{1/2} L M(\kappa + \phi)\theta^2$ and $\beta_2 = \kappa M^2(1 + \nu\theta)^2\theta^2$. Invoking Lemma 14 with $\eta = 2\nu\theta > 1$, we obtain

$$e_t \leq \frac{Q_1^2}{t},$$

where $Q_1 = \max\left\{\|f_*\|_{\mathcal{H}}, \frac{Q_0 + \sqrt{Q_0^2 + (2\theta\nu - 1)(1 + \theta\nu)^2\theta^2\kappa M^2}}{2\nu\theta - 1}\right\}$, and $Q_0 = 2\sqrt{2}\kappa^{1/2}(\kappa + \phi) L M \theta^2$.

Proof for (ii): Cumulating equations (8) with $i = 1, \dots, t$, we end up with the following inequality

$$A_{t+1} \leq \prod_{i=1}^t (1 - 2\gamma_i\nu) A_1 + 2 \sum_{i=1}^t \gamma_i \prod_{j=i+1}^t (1 - 2\nu\gamma_j) \langle h_i - f_*, \bar{g}_i - \hat{g}_i \rangle_{\mathcal{H}} + 2 \sum_{i=1}^t \gamma_i \prod_{j=i+1}^t (1 - 2\nu\gamma_j) \langle h_i - f_*, \hat{g}_i - g_i \rangle_{\mathcal{H}} + \sum_{i=1}^t \gamma_i^2 \prod_{j=i+1}^t (1 - 2\nu\gamma_j) \|g_i\|_{\mathcal{H}}^2 \quad (10)$$

Let us denote $b_t^i = \gamma_i \prod_{j=i+1}^t (1 - 2\nu\gamma_j)$, $1 \leq i \leq t$, the above inequality is equivalent as

$$A_{t+1} \leq \prod_{i=1}^t (1 - 2\gamma_i\nu) A_1 + \sum_{i=1}^t \gamma_i b_t^i \mathcal{M}_i + 2 \sum_{i=1}^t b_t^i \mathcal{N}_i + 2 \sum_{i=1}^t b_t^i \mathcal{R}_i$$

We first show that

- (4) for any $0 < \delta < 1/e$ and $t \geq 4$, with probability $1 - \delta$ over $(\mathcal{D}^t, \omega^t)$,

$$\sum_{i=1}^t b_t^i \mathcal{N}_i \leq 2 \max\left\{4\kappa^{1/2} M \sqrt{\sum_{i=1}^t (b_t^i)^2 A_i}, \max_i |b_t^i| \cdot C_0 \sqrt{\ln(\ln(t)/\delta)}\right\} \sqrt{\ln(\ln(t)/\delta)},$$

where $C_0 = \frac{4 \max_{1 \leq i \leq t} \mathcal{M}_i}{\nu}$.

- (5) for any $\delta > 0$, with probability $1 - \delta$ over $(\mathcal{D}^t, \omega^t)$,

$$\sum_{i=1}^t b_t^i \mathcal{R}_i \leq \sum_{i=1}^t b_t^i \kappa^{1/2} L \hat{B}_{2,i} \sqrt{A_i},$$

where $\hat{B}_{2,i}^2 = 2M^2(\kappa + \phi)^2 \ln\left(\frac{2t}{\delta}\right) \sum_{j=1}^{i-1} |a_{i-1}^j|^2$.

Again, the proofs of these results are given separately in Lemma 10. Applying the above bounds leads to the refined recursion as follows,

$$A_{t+1} \leq \prod_{i=1}^t (1 - 2\gamma_i \nu) A_1 + \sum_{i=1}^t \gamma_i b_t^i \mathcal{M}_i + 2 \sum_{i=1}^t b_t^i \kappa^{1/2} L B_{2,i} \sqrt{A_i} \\ + 4 \max \left\{ 4\kappa^{1/2} M \sqrt{\sum_{i=1}^t (b_t^i)^2 A_i}, \max_i |b_t^i| \cdot C_0 \sqrt{\ln(\ln(t)/\delta)} \right\} \sqrt{\ln(\ln(t)/\delta)}$$

with probability $1 - 2\delta$. When $\gamma_t = \theta/t$ with θ such that $\theta\nu \in \mathbb{Z}_+$, with similar reasons in Lemma 8, we have $|b_t^i| \leq \frac{\theta}{t}$, $1 \leq i \leq t$ and also we have $\prod_{i=1}^t (1 - 2\gamma_i \nu) = 0$, and $\sum_{i=1}^t \gamma_i b_t^i \leq \frac{\theta^2}{t}$. Therefore, we can rewrite the above recursion as

$$A_{t+1} \leq \frac{\beta_1}{t} + \beta_2 \sqrt{\ln(2t/\delta)} \cdot \sum_{i=1}^t \frac{\sqrt{A_i}}{t\sqrt{i}} + \beta_3 \sqrt{\ln(\ln(t)/\delta)} \frac{\sqrt{\sum_{i=1}^t A_i}}{t} + \beta_4 \ln(\ln(t/\delta)) \frac{1}{t} \quad (11)$$

where $\beta_1 = \kappa M^2(1 + \nu\theta)^2 \theta^2$, $\beta_2 = 2\sqrt{2}\kappa^{1/2} LM(\kappa + \phi)\theta^2$, $\beta_3 = 16\kappa^{1/2} M\theta$, $\beta_4 = 16\kappa M^2(1 + \theta\nu)^2 \theta/\nu$. Invoking Lemma 15, we obtain

$$A_{t+1} \leq \frac{Q_2^2 \ln(2t/\delta) \ln^2(t)}{t},$$

with the specified Q_2 . ■

Lemma 10 *In this lemma, we prove the inequalities (1)–(5) in Lemma 9.*

Proof Given the definitions of $\mathcal{M}_t, \mathcal{N}_t, \mathcal{R}_t$ in Lemma 9, we have

$$(1) \mathcal{M}_t \leq \kappa M^2 (1 + \nu \sqrt{\sum_{i,j=1}^{t-1} |a_{t-1}^i| \cdot |a_{t-1}^j|})^2;$$

This is because

$$\mathcal{M}_t = \|g_t\|_{\mathcal{H}}^2 = \|\xi_t + \nu h_t\|_{\mathcal{H}}^2 \leq (\|\xi_t\|_{\mathcal{H}} + \nu \|h_t\|_{\mathcal{H}})^2.$$

We have

$$\|\xi_t\|_{\mathcal{H}} = \|l'(f_t(x_t), y_t)k(x_t, \cdot)\|_{\mathcal{H}} \leq \kappa^{1/2} M,$$

and

$$\|h_t\|_{\mathcal{H}}^2 = \sum_{i=1}^{t-1} \sum_{j=1}^{t-1} a_{t-1}^i a_{t-1}^j l'(f_i(x_i), y_i) l'(f_j(x_j), y_j) k(x_i, x_j) \\ \leq \kappa M^2 \sum_{i=1}^{t-1} \sum_{j=1}^{t-1} |a_{t-1}^i| \cdot |a_{t-1}^j|.$$

$$(2) \mathbb{E}_{\mathcal{D}^t, \omega^t} [\mathcal{N}_t] = 0;$$

This is because $\mathcal{N}_t = \langle h_t - f_*, \bar{g}_t - \hat{g}_t \rangle_{\mathcal{H}}$,

$$\begin{aligned} \mathbb{E}_{\mathcal{D}^t, \omega^t} [\mathcal{N}_t] &= \mathbb{E}_{\mathcal{D}^{t-1}, \omega^t} [\mathbb{E}_{\mathcal{D}_t} [\langle h_t - f_*, \bar{g}_t - \hat{g}_t \rangle_{\mathcal{H}} | \mathcal{D}^{t-1}, \omega^t]] \\ &= \mathbb{E}_{\mathcal{D}^{t-1}, \omega^t} [\langle h_t - f_*, \mathbb{E}_{\mathcal{D}_t} [\bar{g}_t - \hat{g}_t] \rangle_{\mathcal{H}}] \\ &= 0. \end{aligned}$$

$$(3) \mathbb{E}_{\mathcal{D}^t, \omega^t} [\mathcal{R}_t] \leq \kappa^{1/2} L B_{1,t} \sqrt{\mathbb{E}_{\mathcal{D}^{t-1}, \omega^{t-1}} [A_t]}, \text{ where } B_{1,t}^2 := 4M^2(\kappa + \phi)^2 \sum_{i=1}^{t-1} |a_{t-1}^i|^2;$$

This is because $\mathcal{R}_t = \langle h_t - f_*, \hat{g}_t - g_t \rangle_{\mathcal{H}}$,

$$\begin{aligned}
\mathbb{E}_{\mathcal{D}^t, \omega^t}[\mathcal{R}_t] &= \mathbb{E}_{\mathcal{D}^t, \omega^t}[\langle h_t - f_*, \hat{g}_t - g_t \rangle_{\mathcal{H}}] \\
&= \mathbb{E}_{\mathcal{D}^t, \omega^t}[\langle h_t - f_*, [l'(f_t(x_t), y_t) - l'(h_t(x_t), y_t)]k(x_t, \cdot) \rangle_{\mathcal{H}}] \\
&\leq \mathbb{E}_{\mathcal{D}^t, \omega^t}[\|l'(f_t(x_t), y_t) - l'(h_t(x_t), y_t)\| \cdot \|k(x_t, \cdot)\|_{\mathcal{H}} \cdot \|h_t - f_*\|_{\mathcal{H}}] \\
&\leq \kappa^{1/2} L \cdot \mathbb{E}_{\mathcal{D}^t, \omega^t}[\|f_t(x_t) - h_t(x_t)\| \|h_t - f_*\|_{\mathcal{H}}] \\
&\leq \kappa^{1/2} L \sqrt{\mathbb{E}_{\mathcal{D}^t, \omega^t}[|f_t(x_t) - h_t(x_t)|^2]} \sqrt{\mathbb{E}_{\mathcal{D}^t, \omega^t}[\|h_t - f_*\|_{\mathcal{H}}^2]} \\
&\leq \kappa^{1/2} L B_{1,t} \sqrt{\mathbb{E}_{\mathcal{D}^{t-1}, \omega^{t-1}}[A_t]}
\end{aligned}$$

where the first and third inequalities are due to Cauchy–Schwarz Inequality and the second inequality is due to L -Lipschitz continuity of $l'(\cdot, \cdot)$ in the first parameter, and the last step is due to Lemma 7 and the definition of A_t .

(4) for any $0 < \delta < 1/e$ and $t \geq 4$, with probability at least $1 - \delta$ over $(\mathcal{D}^t, \omega^t)$,

$$\sum_{i=1}^t b_i^i \mathcal{N}_i \leq 2 \max \left\{ 4\kappa^{1/2} M \sqrt{\sum_{i=1}^t (b_i^i)^2 A_i}, \max_i |b_i^i| \cdot C_0 \sqrt{\ln(\ln(t)/\delta)} \right\} \sqrt{\ln(\ln(t)/\delta)},$$

where $C_0 = \frac{4 \max_{1 \leq i \leq t} \mathcal{M}_i}{\nu}$.

This result follows directly from Lemma 3 in [Alexander et.al., 2012]. Let us define $d_i = d_i(\mathcal{D}^i, \omega^i) := b_i^i \mathcal{N}_i = b_i^i \langle h_i - f_*, \bar{g}_i - \hat{g}_i \rangle_{\mathcal{H}}$, $1 \leq i \leq t$, we have

- $\{d_i\}_{i=1}^t$ is martingale difference sequence since $\mathbb{E}_{\mathcal{D}^i, \omega^i}[\mathcal{N}_i | \mathcal{D}^{i-1}, \omega^{i-1}] = 0$.
- $|d_i| \leq \max_i |b_i^i| \cdot C_0$, with $C_0 = \frac{4 \max_{1 \leq i \leq t} \mathcal{M}_i}{\nu}$, $\forall 1 \leq i \leq t$.
- $\text{Var}(d_i | \mathcal{D}^{i-1}, \omega^{i-1}) \leq 4\kappa M^2 |b_i^i|^2 A_i$, $\forall 1 \leq i \leq t$.

Plugging in these specific bounds in Lemma 3 in [Alexander et.al., 2012], which is,

$$\Pr \left(\sum_{i=1}^t d_t \geq 2 \max\{2\sigma_t, d_{max} \sqrt{\ln(1/\delta)}\} \sqrt{\ln(1/\delta)} \right) \leq \ln(t)\delta.$$

where $\sigma_t^2 = \sum_{i=1}^t \text{Var}_{i-1}(d_i)$ and $d_{max} = \max_{1 \leq i \leq t} |d_i|$, we immediately obtain the above inequality as desired.

(5) for any $\delta > 0$, with probability at least $1 - \delta$ over $(\mathcal{D}^t, \omega^t)$,

$$\sum_{i=1}^t b_i^i \mathcal{R}_i \leq \sum_{i=1}^t |b_i^i| \kappa^{1/2} L \hat{B}_{2,i} \sqrt{A_i},$$

where $\hat{B}_{2,i}^2 = 2M^2(\kappa + \phi)^2 \ln\left(\frac{2t}{\delta}\right) \sum_{j=1}^{i-1} |a_{i-1}^j|^2$.

This is because, for any $1 \leq i \leq t$, recall that from analysis in (3), we have $\mathcal{R}_i \leq \kappa^{1/2} L |f_t(x_t) - h_t(x_t)| \cdot \|h_t - f_*\|_{\mathcal{H}}$, therefore from Lemma 9,

$$\Pr(b_i^i \mathcal{R}_i \leq \kappa^{1/2} L |b_i^i| \hat{B}_{2,i} \sqrt{A_i}) \geq \Pr(|f_i(x_i) - h_i(x_i)|^2 \leq \hat{B}_{2,i}^2) \geq 1 - \delta/t.$$

Taking the sum over i , we therefore get

$$\Pr(\sum_{i=1}^t b_i^i \mathcal{R}_i \leq \sum_{i=1}^t |b_i^i| \kappa^{1/2} L B_{2,i} \sqrt{A_i}) \geq 1 - \delta.$$

■

Applying these lemmas immediately gives us Theorem 4 and Theorem 5, which implies pointwise distance between the solution $f_{t+1}(\cdot)$ and $f_*(\cdot)$. Now we prove similar bounds in the sense of L_∞ and L_2 distance.

A.2 L_∞ distance, L_2 distance, and generalization bound

Corollary 11 (L_∞ distance) *Theorem 4 also implies a bound in L_∞ sense, namely,*

$$\mathbb{E}_{\mathcal{D}^t, \omega^t} \|f_{t+1} - f_*\|_\infty^2 \leq \frac{2C^2 + 2\kappa Q_1^2}{t}.$$

Consequently, for the average solution $\hat{f}_{t+1}(\cdot) := \frac{1}{t} \sum_{i=1}^t f_i(\cdot)$, we also have

$$\mathbb{E}_{\mathcal{D}^t, \omega^t} \|\hat{f}_{t+1} - f_*\|_\infty^2 \leq \frac{(2C^2 + 2\kappa Q_1^2)(\ln(t) + 1)}{t}.$$

This is because $\|f_{t+1} - f_*\|_\infty = \max_{x \in \mathcal{X}} |f_{t+1}(x) - f_*(x)| = |f_{t+1}(x_*) - f_*(x_*)|$, where $x_* \in \mathcal{X}$ always exists since \mathcal{X} is closed and bounded. Note that the result for average solution can be improved without log factor using more sophisticated analysis (see also reference in [26]).

Corollary 12 (L_2 distance) *With the choices of γ_t in Lemma 9, we have*

$$(i) \quad \mathbb{E}_{\mathcal{D}^t, \omega^t} \|f_{t+1} - f_*\|_2^2 \leq \frac{2C^2 + 2\kappa Q_1^2}{t},$$

$$(ii) \quad \|f_{t+1} - f_*\|_2^2 \leq \frac{C^2 \ln(8\sqrt{et}/\delta) + 2\kappa Q_2^2 \ln(2t/\delta) \ln^2(t)}{t}, \text{ with probability at least } 1 - 3\delta \text{ over } (\mathcal{D}^t, \omega^t).$$

Proof (i) follows directly from Theorem 4. (ii) can be proved as follows. First, we have

$$\|f_{t+1} - f_*\|_2^2 = \mathbb{E}_x |f_{t+1}(x) - f_*(x)|^2 \leq 2\mathbb{E}_x |f_{t+1}(x) - h_{t+1}(x)|^2 + 2\kappa \|h_{t+1} - f_*\|_{\mathcal{H}}.$$

From Lemma 9, with probability at least $1 - 2\delta$, we have

$$\|h_{t+1} - f_*\|_{\mathcal{H}}^2 \leq \frac{Q_2^2 \ln(2t/\delta) \ln^2(t)}{t}. \quad (12)$$

From Lemma 7, for any $x \in \mathcal{X}$, we have

$$\Pr_{\mathcal{D}^t, \omega^t} \left\{ |f_{t+1}(x) - h_{t+1}(x)|^2 \geq \frac{2(\kappa + \phi)^2 M^2 \ln(\frac{2}{\epsilon}) \theta^2}{t} \right\} \leq \epsilon.$$

Since $C^2 = 4(\kappa + \phi)^2 M^2 \theta^2$, the above inequality can be written as

$$\Pr_{\mathcal{D}^t, \omega^t} \left\{ |f_{t+1}(x) - h_{t+1}(x)|^2 \geq \frac{C^2 \ln(\frac{2}{\epsilon})}{2t} \right\} \leq \epsilon.$$

which leads to

$$\Pr_{x \sim \mathbb{P}(x)} \Pr_{\mathcal{D}^t, \omega^t} \left\{ |f_{t+1}(x) - h_{t+1}(x)|^2 \geq \frac{C^2 \ln(\frac{2}{\epsilon})}{2t} \right\} \leq \epsilon.$$

By Fubini's theorem and Markov's inequality, we have

$$\Pr_{\mathcal{D}^t, \omega^t} \left\{ \Pr_{x \sim \mathbb{P}(x)} \left\{ |f_{t+1}(x) - h_{t+1}(x)|^2 \geq \frac{C^2 \ln(\frac{2}{\epsilon})}{2t} \right\} \geq \frac{\epsilon}{\delta} \right\} \leq \delta.$$

From the analysis in Lemma 7, we also have that $|f_{t+1}(x) - h_{t+1}(x)| \leq C^2$. Therefore, with probability at least $1 - \delta$ over $(\mathcal{D}^t, \omega^t)$, we have

$$\mathbb{E}_{x \sim \mathbb{P}(x)} [|f_{t+1}(x) - h_{t+1}(x)|^2] \leq \frac{C^2 \ln(\frac{2}{\epsilon})}{2t} (1 - \frac{\epsilon}{\delta}) + C^2 \frac{\epsilon}{\delta}$$

Let $\epsilon = \frac{\delta}{4t}$, we have

$$\mathbb{E}_{x \sim \mathbb{P}(x)} [|f_{t+1}(x) - h_{t+1}(x)|^2] \leq \frac{C^2}{2t} (\ln(8t/\delta) + \frac{1}{2}) = \frac{C^2 \ln(8\sqrt{et}/\delta)}{2t}. \quad (13)$$

Summing up equation (13) and (12), we have

$$\|f_{t+1} - f_*\|_2^2 \leq \frac{C^2 \ln(8\sqrt{et}/\delta) + 2\kappa Q_2^2 \ln(2t/\delta) \ln^2(t)}{t}$$

as desired. ■

From the bound on L_2 distance, we can immediately get the generalization bound.

Theorem 6 (Generalization bound) *Let the true risk be $R_{true}(f) = \mathbb{E}_{(x,y)} [l(f(x), y)]$. Then with probability at least $1 - 3\delta$ over $(\mathcal{D}^t, \omega^t)$, and C and Q_2 defined as previously*

$$R_{true}(f_{t+1}) - R_{true}(f_*) \leq \frac{(C\sqrt{\ln(8\sqrt{e}t/\delta)} + \sqrt{2\kappa}Q_2\sqrt{\ln(2t/\delta)}\ln(t))L}{\sqrt{t}}.$$

Proof By the Lipschitz continuity of $l(\cdot, y)$ and Jensen's Inequality, we have

$$R_{true}(f_{t+1}) - R_{true}(f_*) \leq L\mathbb{E}_x |f_{t+1}(x) - f_*(x)| \leq L\sqrt{\mathbb{E}_x |f_{t+1}(x) - f_*(x)|^2} = L\|f_{t+1} - f_*\|_2.$$

Then the theorem follows from Corollary 12. ■

A.3 Suboptimality

For comprehensive purposes, we also provide the $O(1/t)$ bound for suboptimality.

Corollary 13 *If we set $\gamma_t = \frac{\theta}{t}$ with $\theta\nu = 1$, then the average solution $\hat{f}_{t+1} := \frac{1}{t} \sum_{i=1}^t f_i$ satisfies*

$$R(\mathbb{E}_{\mathcal{D}^t, \omega^t}[\hat{f}_{t+1}]) - R(f_*) \leq \frac{Q(\ln(t) + 1)}{t}.$$

where $Q = (4\kappa M^2 + 2\sqrt{2}\kappa^{1/2}LM(\kappa + \phi)Q_1)/\nu$, with Q_1 defined as in Lemma 9.

Proof From the analysis in Lemma 9, we have

$$\langle h_t - f_*, \bar{g}_t \rangle_{\mathcal{H}} = \frac{1}{2\gamma_t} A_t - \frac{1}{2\gamma_t} A_{t+1} + \gamma_t \mathcal{M}_t + \mathcal{N}_t + \mathcal{R}_t$$

Invoking strongly convexity of $R(f)$, we have $\langle h_t - f_*, \bar{g}_t \rangle \geq R(h_t) - R(f_*) + \frac{\nu}{2} \|h_t - f_*\|_{\mathcal{H}}^2$. Taking expectation on both side and use the bounds in last lemma, we have

$$\mathbb{E}_{\mathcal{D}^t, \omega^t} [R(h_t) - R(f_*)] \leq \left(\frac{1}{2\gamma_t} - \frac{\nu}{2}\right) e_t - \frac{1}{2\gamma_t} e_{t+1} + \gamma_t \kappa M^2 (1 + \nu c_t)^2 + \kappa^{1/2} L B_{1,t} \sqrt{e_t}$$

Assume $\gamma_t = \frac{\theta}{t}$ with $\theta = \frac{1}{\nu}$, then cumulating the above inequalities leads to

$$\sum_{i=1}^t \mathbb{E}_{\mathcal{D}^t, \omega^t} [R(h_i) - R(f_*)] \leq \sum_{i=1}^t \gamma_i \kappa M^2 (1 + \nu c_i)^2 + \sum_{i=1}^t \kappa^{1/2} L B_{1,i} \sqrt{e_i}$$

which can be further bounded by

$$\begin{aligned} \sum_{i=1}^t \mathbb{E}_{\mathcal{D}^t, \omega^t} [R(h_i) - R(f_*)] &\leq \sum_{i=1}^t \gamma_i \kappa M^2 (1 + \nu c_i)^2 + \sum_{i=1}^t \kappa^{1/2} L B_{1,i} \sqrt{e_i} \\ &\leq \frac{4\kappa M^2}{\nu} \sum_{i=1}^t \frac{1}{i} + \frac{2\sqrt{2}\kappa^{1/2}LM(\kappa + \phi)}{\nu} \sum_{i=1}^t \sqrt{\frac{e_i}{i}} \\ &\leq \frac{4\kappa M^2}{\nu} (\ln(t) + 1) + \frac{2\sqrt{2}\kappa^{1/2}LM(\kappa + \phi)}{\nu} Q_1 (\ln(t) + 1) \\ &= \frac{Q(\ln(t) + 1)}{t} \end{aligned}$$

By convexity, we have $\mathbb{E}_{\mathcal{D}^t, \omega^t} [R(\hat{h}_{t+1}) - R(f_*)] \leq \frac{Q(\ln(t)+1)}{t}$. The corollary then follows from the fact that $\mathbb{E}_{\mathcal{D}^t, \omega^t} [\hat{f}_{t+1}] = \mathbb{E}_{\mathcal{D}^t, \omega^t} [\hat{h}_{t+1}]$ and $R(\mathbb{E}_{\mathcal{D}^t, \omega^t} [\hat{h}_{t+1}]) \leq \mathbb{E}_{\mathcal{D}^t, \omega^t} [R(\hat{h}_{t+1})]$. ■

A.4 Technical lemma for recursion bounds

Lemma 14 Suppose the sequence $\{\Gamma_t\}_{t=1}^\infty$ satisfies $\Gamma_1 \geq 0$, and $\forall t \geq 1$

$$\Gamma_{t+1} \leq \left(1 - \frac{\eta}{t}\right) \Gamma_t + \frac{\beta_1}{t\sqrt{t}} \sqrt{\Gamma_t} + \frac{\beta_2}{t^2},$$

where $\eta > 1, \beta_1, \beta_2 > 0$. Then $\forall t \geq 1$,

$$\Gamma_t \leq \frac{R}{t}, \text{ where } R = \max\{\Gamma_1, R_0^2\}, R_0 = \frac{\beta_1 + \sqrt{\beta_1^2 + 4(\eta-1)\beta_2}}{2(\eta-1)}.$$

Proof The proof follows by induction. When $t = 1$, it always holds true by the definition of R . Assume the conclusion holds true for t with $t \geq 1$, i.e., $\Gamma_t \leq \frac{R}{t}$, then we have

$$\begin{aligned} \Gamma_{t+1} &\leq \left(1 - \frac{\eta}{t}\right) \Gamma_t + \frac{\beta_1}{t\sqrt{t}} \sqrt{\Gamma_t} + \frac{\beta_2}{t^2} \\ &= \frac{R}{t} - \frac{\eta R - \beta_1 \sqrt{R} - \beta_2}{t^2} \\ &\leq \frac{R}{t+1} + \frac{R}{t(t+1)} - \frac{\eta R - \beta_1 \sqrt{R} - \beta_2}{t^2} \\ &\leq \frac{R}{t+1} - \frac{1}{t^2} \left[-R + \eta R - \beta_1 \sqrt{R} - \beta_2\right] \\ &\leq \frac{R}{t+1} \end{aligned}$$

where the last step can be verified as follows.

$$\begin{aligned} (\eta-1)R - \beta_1 \sqrt{R} - \beta_2 &= (\eta-1) \left[\sqrt{R} - \frac{\beta_1}{2(\eta-1)} \right]^2 - \frac{\beta_1^2}{4(\eta-1)} - \beta_2 \\ &\geq (\eta-1) \left[R_0 - \frac{\beta_1}{2(\eta-1)} \right]^2 - \frac{\beta_1^2}{4(\eta-1)} - \beta_2 \geq 0 \end{aligned}$$

where the last step follows from the definition of R_0 . ■

Lemma 15 Suppose the sequence $\{\Gamma_t\}_{t=1}^\infty$ satisfies

$$\Gamma_{t+1} \leq \frac{\beta_1}{t} + \beta_2 \sqrt{\ln(2t/\delta)} \cdot \sum_{i=1}^t \frac{\sqrt{\Gamma_i}}{t\sqrt{i}} + \beta_3 \sqrt{\ln(\ln(t)/\delta)} \frac{\sqrt{\sum_{i=1}^t \Gamma_i}}{t} + \beta_4 \ln(\ln(t/\delta)) \frac{1}{t}$$

where $\beta_1, \beta_2, \beta_3, \beta_4 > 0$ and $\delta \in (0, 1/e)$. Then $\forall 1 \leq j \leq t (t \geq 4)$,

$$\Gamma_j \leq \frac{R \ln(2t/\delta) \ln^2(t)}{j}, \text{ where } R = \max\{\Gamma_1, R_0^2\}, R_0 = 2\beta_2 + 2\sqrt{2}\beta_3 + \sqrt{(2\beta_2 + 2\sqrt{2}\beta_3)^2 + \beta_1 + \beta_4}.$$

Proof The proof follows by induction. When $j = 1$ it is trivial. Let us assume it holds true for $1 \leq j \leq t-1$, therefore,

$$\begin{aligned}
\Gamma_{j+1} &\leq \frac{\beta_1}{j} + \beta_2 \sqrt{\ln(2j/\delta)} \cdot \sum_{i=1}^j \frac{\sqrt{\Gamma_i}}{j\sqrt{i}} + \beta_3 \sqrt{\ln(\ln(j)/\delta)} \frac{\sqrt{\sum_{i=1}^j \Gamma_i}}{j} + \beta_4 \ln(\ln(j/\delta)) \frac{1}{j} \\
&\leq \frac{\beta_1}{j} + \beta_2 \sqrt{\ln(2j/\delta)}/j \cdot \sum_{i=1}^j \frac{\sqrt{R \ln(2t/\delta) \ln^2(t)}}{i} \\
&\quad + \beta_3 \sqrt{\ln(\ln(j)/\delta)} \frac{\sqrt{\sum_{i=1}^j R \ln(2t/\delta) \ln^2(t)/i}}{j} + \beta_4 \ln(\ln(j/\delta)) \frac{1}{j} \\
&\leq \frac{\beta_1}{j} + \beta_2 \sqrt{\ln(2j/\delta)}/j \sqrt{R \ln(2t/\delta) \ln^2(t) (1 + \ln(j))} \\
&\quad + \beta_3 \sqrt{\ln(\ln(j)/\delta)}/j \sqrt{R \ln(2t/\delta) \ln^2(t) \sqrt{\ln(j) + 1}} + \beta_4 \ln(\ln(j/\delta)) \frac{1}{j} \\
&\leq \frac{\beta_1}{j} + 2\beta_2 \sqrt{R} \ln(2t/\delta) \ln^2(t)/j + \sqrt{2}\beta_3 \sqrt{R} \ln(2t/\delta) \ln^2(t)/j + \beta_4 \ln(2t/\delta) \frac{1}{j} \\
&\leq (2\beta_2 + \sqrt{2}\beta_3) \sqrt{R} \frac{\ln(2t/\delta) \ln^2(t)}{j} + (\beta_1 + \beta_4 \ln(2t/\delta)) \frac{1}{j} \\
&\leq \frac{\ln(2t/\delta) \ln^2(t)}{j} [(2\beta_2 + \sqrt{2}\beta_3) \sqrt{R} + \frac{\beta_1}{2} + \frac{\beta_4}{2}]
\end{aligned}$$

Since $\sqrt{R} \geq 2\beta_2 + 2\sqrt{2}\beta_3 + \sqrt{(2\beta_2 + 2\sqrt{2}\beta_3)^2 + \beta_1 + \beta_4}$, we have $(2\beta_2 + 2\sqrt{2}\beta_3) \sqrt{R} + \frac{\beta_1}{2} + \frac{\beta_4}{2} \leq R/2$. Hence, $\Gamma_{j+1} \leq \frac{R \ln(2t/\delta) \ln^2(t)}{j+1}$. ■