
Generalization Bounds for Stochastic Gradient Descent via Localized ε -Covers

Sejun Park
Korea University
sejun.park000@gmail.com

Umut Şimşekli
DI ENS, Ecole Normale Supérieure, Université PSL, CNRS, INRIA
umut.simsekli@inria.fr

Murat A. Erdogdu
University of Toronto & Vector Institute
erdogdu@cs.toronto.edu

Abstract

In this paper, we propose a new covering technique localized for the trajectories of SGD. This localization provides an algorithm-specific complexity measured by the covering number, which can have dimension-independent cardinality in contrast to standard uniform covering arguments that result in exponential dimension dependency. Based on this localized construction, we show that if the objective function is a finite perturbation of a piecewise strongly convex and smooth function with P pieces, i.e. non-convex and non-smooth in general, the generalization error can be upper bounded by $O(\sqrt{(\log n \log(nP))/n})$, where n is the number of data samples. In particular, this rate is independent of dimension and does not require early stopping and decaying step size. Finally, we employ these results in various contexts and derive generalization bounds for multi-index linear models, multi-class support vector machines, and K -means clustering for both hard and soft label setups, improving the known state-of-the-art rates.

1 Introduction

We consider the following stochastic optimization problem

$$\min_{\theta \in \Theta} \left\{ F(\theta) := \mathbb{E}_Z[f(\theta; Z)] \right\}, \quad (1.1)$$

where θ represents the optimization parameter, $\Theta \subset \mathbb{R}^d$ is a convex parameter domain, $f(\cdot; z)$ is a possibly non-convex loss incurred by a single data point $z \in \mathcal{Z}$, and Z is a random variable on \mathcal{Z} following the data distribution. Since the distribution of Z is unknown in general, the following proxy based on independent and identically distributed (i.i.d.) samples z_1, \dots, z_n of Z is optimized instead

$$\min_{\theta \in \Theta} \left\{ \hat{F}(\theta) := \frac{1}{n} \sum_{i=1}^n f(\theta; z_i) \right\}. \quad (1.2)$$

Given a learning algorithm $A(\cdot)$ mapping samples z_1, \dots, z_n to an approximate solution of (1.2), bounding the *generalization error*¹ $\hat{F}(A(z_1, \dots, z_n)) - F(A(z_1, \dots, z_n))$ is a fundamental problem

¹This quantity is also referred to as the generalization gap or the estimation error in the literature.

in learning theory. Classical *algorithm-independent* results rely on uniform convergence over the entire domain $\Theta \subset \mathbb{R}^d$; thus, they apply to any learning algorithm. However, these bounds often increase with the dimension d [SSSSS09, SSSBD14, Fel16], becoming vacuous in the modern over-parameterized regime [ZBH⁺21]. To derive dimension-independent bounds, researchers have been investigating *algorithm-dependent* generalization properties, especially for commonly used training methods such as stochastic gradient descent (SGD) [HRS15, SHN⁺18, YKM21].

Notably, *algorithmic stability* is a technique for deriving generalization bounds based on the properties of a specific learning algorithm, which leverages that if a parameter learned by an algorithm is robust under a perturbation of samples z_1, \dots, z_n , then the generalization error at that parameter must be small [BE02]. Based on this principle, several works proposed dimension-independent generalization bounds for SGD and its variants under various setups [HRS15, Lon17, FV19, LLQ20, BFGT20, LY20, FO21, LLY21, KWS22]. Bounds derived using algorithmic stability is optimal for strongly convex and smooth functions [SSSSS09] and convex and non-smooth functions [BFGT20, AKL21]. Nevertheless, without global (strong) convexity, early stopping, and/or decaying step size, generalization bounds based on algorithmic stability often diverge with the number of SGD iterations [HRS15, LLQ20], failing to explain the empirical observations.

To obtain (ambient) dimension-independent bounds that do not diverge with the number of iterations, recent works proposed to utilize the low-dimensional fractal structures generated by the SGD iterates whose complexity can be measured by a notion called the Hausdorff dimension [Fal14]. In this context, [SSDE20] showed that, under a continuous-time surrogate for SGD, the generalization error can be bounded by $\tilde{O}(\sqrt{d_H/n})$, where d_H denotes the Hausdorff dimension of the optimization trajectory. This result was later extended to discrete-time iterated function systems by [CDE⁺21]. Here, the Hausdorff dimension can be smaller than the ambient dimension [CDE⁺21], ultimately providing improved generalization bounds. However, both of these results are inherently asymptotic, and rely on opaque assumptions that are hard to verify in practice.

In this paper, we propose a new framework for deriving generalization bounds for the projected SGD with a constant step size and without requiring early stopping. Inspired by the works [SSDE20, CDE⁺21], our framework is based on a complexity measure of the trajectory of SGD, which can be quantified under standard verifiable conditions. Our contributions are as follows.

- **Localized ε -covers for SGD.** Our first principle contribution is a covering technique localized for the trajectories of SGD. This localization provides an algorithm-specific complexity measured by the covering number, which can have dimension-independent cardinality in contrast to standard covering arguments that result in exponential dimension dependency.
- **Generalization bounds for SGD.** Based on this localized covering, we establish dimension-independent generalization bounds for SGD, for a rich class of non-convex loss functions f whose gradients can be approximated by that of a piecewise strongly convex and smooth function h , i.e. $\|\nabla f(\theta; z) - \nabla h(\theta; z)\| \leq \xi$ for some ξ . In particular, with high probability, we prove the bound

$$\left| \hat{F}(\theta^{(t)}) - F(\theta^{(t)}) \right| = O \left(\sqrt{\frac{\log n \log(nP)}{n}} + \xi \right), \quad (1.3)$$

where $\theta^{(t)}$ denotes the parameter generated by t SGD iterations for a sufficiently large t , and P denotes the number of strongly convex pieces needed to approximate f . We further show that the gradient of any (piecewise) smooth function f can be approximated with that of a piecewise strongly convex and smooth function, demonstrating the wide applicability of the bound (1.3). Finally in the special case $P = 1$ and $\xi = 0$, our result reduces to a non-asymptotic bound where the complexity is captured by the Hausdorff dimension of the invariant measure of SGD.

- **Improved bounds in specific models.** We employ the above result to derive generalization bounds in several statistical models trained by SGD, including multi-index linear models, multi-class support vector machines, and K -means clustering with both hard and soft label setups, improving the previously known state-of-the-art generalization error bounds in this context.

Notation and problem setup. For $k \in \mathbb{N}$, we denote $[k] := \{1, \dots, k\}$. We use $\|\cdot\|$ to denote the ℓ_2 -norm. For $\varepsilon > 0$ and $\theta \in \mathbb{R}^d$, we use $\mathcal{B}_\varepsilon^d(\theta)$ to denote the d -dimensional closed ℓ_2 -ball of radius ε , centered at θ . Given a set $\mathcal{S} \subseteq \mathbb{R}^d$ and $\varepsilon > 0$, we say $\mathcal{C}_\varepsilon \subseteq \mathbb{R}^d$ is an “ ε -cover” of \mathcal{S} if $\mathcal{S} \subseteq \bigcup_{\theta \in \mathcal{C}_\varepsilon} \mathcal{B}_\varepsilon^d(\theta)$.

Given a step size $\eta > 0$, an initial parameter $\theta^{(0)} \in \Theta$, and a randomly sampled index $i_t \in [n]$, the t -th iteration of the projected SGD performs the following update on the parameters

$$\theta^{(t)} = g_{i_t}(\theta^{(t-1)}) := \Pi_{\Theta}(\theta^{(t-1)} - \eta \nabla f(\theta^{(t-1)}; z_{i_t})) \quad \text{for } t = 1, 2, \dots, \quad (1.4)$$

where $\Pi_{\Theta}(\theta) := \arg \min_{\theta' \in \Theta} \|\theta' - \theta\|$ denotes the Euclidean projection. The domain Θ is convex; thus, the projection operation is unique. We note that the projection is not needed under the presence of an ℓ_2 -regularizer and Lipschitz continuity; see Section 4 for more details. Throughout the paper, we use $\theta^{(t)} := g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)})$ for possibly random indices i_1, \dots, i_t .

Lastly, we recall a few standard notions. $f : \Theta \rightarrow \mathbb{R}$ is called “ α -strongly convex” and “ β -smooth” respectively if for all $\theta, \theta' \in \Theta$, the following conditions are satisfied

$$f(\theta) - f(\theta') - \nabla f(\theta')^\top (\theta - \theta') \geq \frac{\alpha}{2} \|\theta - \theta'\|^2, \quad \text{and} \quad \|\nabla f(\theta) - \nabla f(\theta')\| \leq \beta \|\theta - \theta'\|.$$

The function f is called “convex” if it is 0-strongly convex.

2 Main results

We demonstrate our covering construction in Section 2.1 in the classical strongly convex and smooth case in which the localization argument can be simplified by the contractivity of SGD. We present our main generalization result on non-convex losses in Section 2.2, and its implications in Section 3.

2.1 A localized covering construction: Strongly convex and smooth case

To motivate our approach, let us first briefly discuss the limitations of prior methods that are based on uniform convergence of empirical processes. Given a set of parameters $\Theta \subseteq \mathcal{B}_R^d(0)$, let \mathcal{C}_ε be an ε -cover of Θ , i.e. $\Theta \subseteq \bigcup_{\phi \in \mathcal{C}_\varepsilon} \mathcal{B}_\varepsilon^d(\phi)$. Uniform convergence over Θ can be established with high probability by simply applying the union bound over \mathcal{C}_ε , which yields a generalization error bound depending on the cardinality of the cover $\sqrt{\log |\mathcal{C}_\varepsilon|}$. However, \mathcal{C}_ε is typically independent of the algorithm being used, and standard ε -covers for Θ yield $|\mathcal{C}_\varepsilon| = (R/\varepsilon)^{\Omega(d)}$; thus, bounds based on covering numbers often grow with \sqrt{d} , which can be loose if d is large. To overcome this issue, we use the contractive properties of SGD in the strongly convex and smooth case and *localize* the ε -cover. Namely, instead of covering the entire feasible set Θ , we construct a cover that contains only the points that can be reached by SGD trajectories, resulting in a covering number that is independent of the ambient dimension d . We introduce the following sets produced by SGD trajectories.

Definition 1. We define the following two subsets of Θ .

- The set of points that can be reached by T SGD iterations initialized at $\theta^{(0)} \in \Theta$,

$$\Psi_T(\theta^{(0)}) := \{g_{i_T} \circ \dots \circ g_{i_1}(\theta^{(0)}) : i_1, \dots, i_T \in [n]\}.$$

- The set of points that can be reached by any $t \geq T$ SGD iterations initialized at $\theta^{(0)} \in \Theta$,

$$\Psi_{\geq T}(\theta^{(0)}) := \bigcup_{t \geq T} \Psi_t(\theta^{(0)}).$$

For $\gamma \in (0, 1)$, a function $g : \Theta \rightarrow \Theta$ is called “ γ -contractive” if for all $\theta, \theta' \in \Theta$, it satisfies $\|g(\theta) - g(\theta')\| \leq \gamma \|\theta - \theta'\|$. In the constant step-size case, SGD iterates converge to a distribution instead of a single point [DDB20], but their contractivity can still provide the following localization: all possible SGD iterates after sufficiently many iterations can be ε -covered by $n^{O(\log(1/\varepsilon))}$ points.

Lemma 2.1. Suppose that g_1, g_2, \dots are γ -contractive for some $\gamma \in (0, 1)$. Then, for any initialization $\theta^{(0)} \in \Theta \subseteq \mathcal{B}_R^d(0)$ and for any $\varepsilon > 0$, for $T := T_\varepsilon = \max \left\{ \left\lceil \frac{\log(R/\varepsilon)}{\log(1/\gamma)} \right\rceil, 0 \right\}$, we have

$$\Psi_{\geq T}(\theta^{(0)}) \subseteq \bigcup_{\phi \in \Psi_T(0)} \mathcal{B}_\varepsilon^d(\phi).$$

Proof. For $t \geq T$, let $\theta^{(t)} \in \Psi_{\geq T}(\theta^{(0)})$ such that $\theta^{(t)} := g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)})$ for some $i_1, \dots, i_t \in [n]$. Let $\phi := g_{i_t} \circ \dots \circ g_{i_{t-T+1}}(\theta^{(0)})$ and notice that $\phi \in \Psi_T(0)$ by construction. Then for $T := T_\varepsilon$ and $\theta^{(t-T)} := g_{i_{t-T}} \circ \dots \circ g_{i_1}$, we have

$$\|\theta^{(t)} - \phi\| \leq \gamma^T \|\theta^{(t-T)} - 0\| \leq \gamma^T R \leq \varepsilon,$$

since each g_i is γ -contractive and $\|\theta^{(t-T)}\| \leq R$ by (1.4). This implies $\theta^{(t)} \in \bigcup_{\phi \in \Psi_T(0)} \mathcal{B}_\varepsilon^d(\phi)$. \square

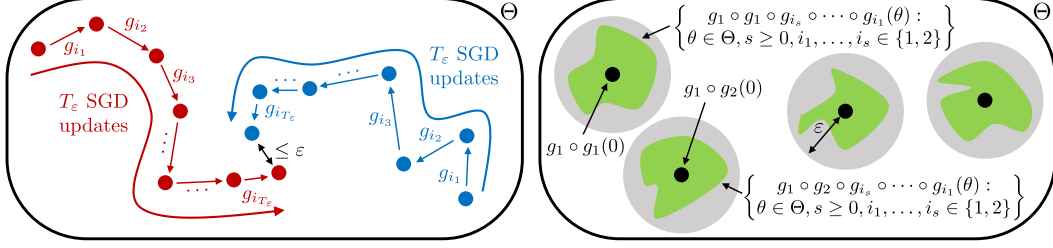


Figure 1: Left: illustration of $T = T_\epsilon$ coupled projected SGD updates from two distinct points. Right: illustration of our localized cover Ψ_T covering $\Psi_{\geq T}$ with $T = 2$ and $n = 2$ and $|\Psi_T| = n^T = 4$.

For strongly convex and smooth $f(\cdot; z)$, an SGD update g_i with a sufficiently small step size is contractive [DDB20]; that is, applying g_i to any two points decreases the distance between them (see Appendix A.1 for a formal derivation). Therefore, for any $\epsilon > 0$, there exists $T := T_\epsilon$ such that applying T *synchronously coupled* SGD updates that use the same sample at each iteration can make the distance between the initial points smaller than ϵ ; see Figure 1 (left). This observation implies that the set $\Psi_T(0)$ of all parameters that can be generated by T SGD updates when initialized at the origin ϵ -covers the set $\Psi_{\geq T}(\theta^{(0)})$ of all parameters that can be generated by any $t \geq T$ SGD updates for an arbitrary initialization $\theta^{(0)}$; see Figure 1 (right). In contrast to algorithm-independent covers of Θ that scale with $|\mathcal{C}_\epsilon| = (R/\epsilon)^{\Omega(d)} = e^{\tilde{\Omega}(d)}$, we obtain $|\Phi_T(0)| \leq n^T = e^{\tilde{O}(1)}$ which is independent of the dimension d and only polynomial in the number of samples n .

We make the following additional assumptions on the loss function f .

Assumption 1 (Weak Lipschitz continuity). *For $L > 0$, there exists $h : \Theta \rightarrow \mathbb{R}$ such that for all $\theta, \theta' \in \Theta$ and $z \in \mathcal{Z}$, $|f(\theta; z) - h(\theta) - (f(\theta'; z) - h(\theta'))| \leq L\|\theta - \theta'\|$.*

Notice that $h = 0$ reduces to the classical Lipschitz continuity, but the above condition is more general.

Assumption 2 (Bounded deviation). *For $B > 0$, for all $\theta \in \Theta$, $|\sup_{z \in \mathcal{Z}} f(\theta; z) - \inf_{z \in \mathcal{Z}} f(\theta; z)| \leq B$.*

We note that Assumptions 1 & 2 (or their variants) both appear in several algorithmic stability-based results; see e.g. [HRS15, Thm 3.10]. We also highlight that these conditions are invariant to adding regularizers (e.g. consider $f(\theta; z) \leftarrow f(\theta; z) + \frac{\lambda}{2}\|\theta\|^2$) which will be useful in Section 3.

Theorem 2.1. *Suppose that Assumptions 1 & 2 hold and there exist $\alpha, \beta > 0$ such that $f(\cdot; z)$ is α -strongly convex and β -smooth on $\Theta \subseteq \mathcal{B}_R^d(0)$ for all $z \in \mathcal{Z}$. For any $\eta \in (0, 2/\beta)$, let $\gamma := \sqrt{1 - 2\alpha\eta + \alpha\beta\eta^2}$, $T := \max\left\{\left\lceil \frac{\log(2LRn)}{\log(1/\gamma)} \right\rceil, 0\right\}$. Then, with probability at least $1 - \delta$, for any $\theta^{(0)} \in \Theta$, $t \geq T$, and $i_1, \dots, i_t \in [n]$, SGD iterate $\theta^{(t)}$ satisfies*

$$\left| \hat{F}(\theta^{(t)}) - F(\theta^{(t)}) \right| \leq \frac{BT + 1}{n} + B\sqrt{\frac{T \log n + \log(2/\delta)}{2n}}. \quad (2.1)$$

Remark. A few remarks are in order. For stochastic convex optimization, the best achievable bound via uniform convergence over Θ is worse than the algorithmic stability-based bounds for SGD by a factor of \sqrt{d} [SSSS09, Fel16, FV19]; however, by localizing the uniform convergence argument, we are able to obtain a dimension-independent bound. We note that this high-probability result is still not directly comparable to most stability-based bounds which are given in expectation. An exception is [FV19, Thm 4.5] which obtains near-optimal bounds for SGD in the convex case by tuning the number of iterations to be taken, obtaining a rate of $O(\log(n/\delta)^2 \log(n)/\sqrt{n})$. In contrast, our bound reads $O(\sqrt{\log(n)^2 \log(1/\delta)/n})$ which is better by a logarithmic factor and holds for any sufficiently large number of SGD iterations. We emphasize that the bound (2.1) holds for *the union of all trajectories of SGD* generated by $t \geq T$ iterations, whereas the stability-based results often consider a single parameter generated by SGD. A similar setup is considered in Corollary 2.1 where we establish improved bounds on the generalization error removing the logarithmic factor. Lastly, we note that our localized covering can also be used for deriving bounds in expectation; see Appendix D.

The proof of Theorem 2.1 follows from three steps. *i)* Lemma 2.1 implies that any $\theta^{(t)}$ for $t \geq T$ can be approximated by some parameter $\theta^{(T)}$ generated by T SGD updates initialized at 0. *ii)* $\hat{F}(\theta^{(T)})$

concentrates around $F(\theta^{(T)})$ since $\theta^{(T)}$ is *almost* independent of the samples, i.e. it depends on at most $T = O(\log(n))$ of n samples. *iii*) The empirical process $|\hat{F}(\theta^{(T)}) - F(\theta^{(T)})|$ uniformly converges over the set $\Psi_T(0)$ which has a dimension-free cardinality. See Appendix B.2 for details.

While the bound in Theorem 2.1 holds uniformly over all possible initializations, it can be tightened by considering a single realization of $\theta^{(0)}$ in the following corollary.

Corollary 2.1. *Assume the setup in Theorem 2.1. Then, for any $\theta^{(0)} \in \Theta$, $t \geq T$, and $i_1, \dots, i_t \in [n]$, with probability at least $1 - \delta$, we have*

$$|\hat{F}(\theta^{(t)}) - F(\theta^{(t)})| \leq \frac{BT + 1}{n} + B\sqrt{\frac{\log(2/\delta)}{2n}}.$$

For sufficiently large n , the bound in Corollary 2.1 coincides with the tight concentration bound at a fixed $\theta \in \Theta$; however, it has the limitation that at least $t \geq T$ SGD updates are required. This is remedied in the next result which does not require strong convexity and smoothness.

Corollary 2.2. *Suppose that Assumption 2 holds. Then, for any $\theta^{(0)} \in \Theta$, $t \geq 0$, and $i_1, \dots, i_t \in [n]$, with probability at least $1 - \delta$, we have*

$$|\hat{F}(\theta^{(t)}) - F(\theta^{(t)})| \leq \frac{Bt}{n} + B\sqrt{\frac{\log(2/\delta)}{2n}}.$$

The bound in Corollary 2.2 requires early stopping since it diverges as $t \rightarrow \infty$. However, since Corollary 2.2 only requires Assumption 2, it can be easily combined with other results, e.g. Corollary 2.1, to provide a generalization bound that holds for any number of SGD iterations. The proofs of Corollaries 2.1 and 2.2 are presented in Appendices B.3 and B.4, respectively.

Relation to fractal dimension. There is an interesting connection to be made between Theorem 2.1 and the Hausdorff dimension of the support μ of the stationary distribution of (projected) SGD. For example, suppose that $\eta \in (0, 1)$, $\Theta = \mathcal{B}_R^d(0)$, $f(\theta; z) = \frac{1}{2}\|\theta - \theta_z^*\|^2$ for some $\theta_z^* \in \mathcal{B}_R^d(0)$, and $\|\theta_{z_i}^* - \theta_{z_j}^*\| \geq 2\gamma R$ for all $i \neq j$. Here, the last assumption can be satisfied with high probability if d, η are large enough and $\theta_z^* \sim N(0, \sigma^2 I)$, i.e. zero-mean Gaussian with covariance $\sigma^2 I$; see e.g. [PLYS21, Appendix C]. Under these assumptions, the bound in Theorem 2.1 can be reformulated as

$$|\hat{F}(\theta^{(t)}) - F(\theta^{(t)})| \leq \frac{BT + 1}{n} + B\sqrt{\frac{\lceil d_H + \frac{\log 2LR}{\log(1/\gamma)} \rceil \log n + \log(2/\delta)}{2n}}, \quad (2.2)$$

where d_H denotes the Hausdorff dimension of μ . The precise definition of d_H and the derivation of (2.2) are presented in Appendix B.5. In (2.2), d_H replaces the ambient dimension d that appear in algorithm-independent bounds [SSBD14]. We note that such a connection between the fractal dimension and generalization bounds has been also studied in [SSDE20, CDE⁺21]; however, their bounds are asymptotic and require non-trivial assumptions that are not easy to verify in practice.

Generalization due to contractivity. Contractivity of SGD has also been used to derive generalization bounds in the concurrent work [KWS22]. Although the localized covering construction in Lemma 2.1 relies on the same principle, our results differ from those in [KWS22] on two key aspects. First and foremost, as we shall see in Section 2.2, the localized covering argument used in Lemma 2.1 can also be applied to the non-convex case in which the uniform stability-based argument used in [KWS22] provably breaks down (Appendix F); thus, extending the results of [KWS22] to cover non-convex objectives is highly non-trivial. Further, the generalization bounds in [KWS22] are provided in expectation, i.e. $|\mathbb{E}[\hat{F}(\theta^{(t)}) - F(\theta^{(t)})]| \leq O(1/n)$, whereas we provide high-probability bounds, i.e. $|\hat{F}(\theta^{(t)}) - F(\theta^{(t)})| \leq \tilde{O}(1/\sqrt{n})$. When translated to bounds in expectation, our results read $|\mathbb{E}[\hat{F}(\theta^{(t)}) - F(\theta^{(t)})]| \leq \tilde{O}(1/n)$ and $\mathbb{E}[|\hat{F}(\theta^{(t)}) - F(\theta^{(t)})|] \leq \tilde{O}(1/\sqrt{n})$; see Appendices D & E.

2.2 Non-convex case: Perturbations of piecewise strongly convex and smooth functions

Algorithmic stability technique yields (near) optimal rates for strongly convex objectives; however, when applied to non-convex functions, the resulting bounds often diverge with the number of SGD iterations [HRS15, LLQ20]. The localized covering construction introduced in the previous section remedies this issue, providing more stable generalization bounds. Specifically, we establish a

dimension-independent generalization bound for functions that are finite perturbations of piecewise strongly convex and smooth functions. We further prove an approximation result and show that any smooth non-convex function can be approximated with a piecewise strongly convex function. Since piecewise strongly convex and smooth functions may not be differentiable on the entire Θ , we will use the auxiliary gradient in Definition 2 as a surrogate for the SGD update (1.4).

Definition 2. f is “piecewise α -strongly convex and β -smooth with P pieces on Θ ” if there exists a partition $\{\mathcal{P}_1, \dots, \mathcal{P}_P\}$ of Θ and α -strongly convex and β -smooth functions h_1, \dots, h_P on Θ such that $f = h_p$ on \mathcal{P}_p for all $p \in [P]$. We also define $\nabla f(\theta) := \nabla h_p(\theta)$ if $\theta \in \mathcal{P}_p$.

Piecewise strongly convex and smooth objectives are widely used in machine learning applications. For example, the objective of learning a single layer of a ReLU network with ℓ_2 -regularization is piecewise strongly convex and smooth. Furthermore, the objective of learning an entire ReLU network is also piecewise strongly convex and smooth on small loss regions [Mil19]. Such observations easily extend to more general settings, e.g. an objective function defined as the summation of a piecewise linear loss (e.g. hinge loss) and ℓ_2 -regularization. However, if we further allow for finite perturbations, any piecewise smooth non-convex function can be covered within this framework.

Proposition 2.1. For any piecewise β' -smooth f with Q pieces on $\Theta \subset \mathcal{B}_R^d(0)$, and for any $\xi > 0$ and $0 < \alpha \leq \beta$, there exists a piecewise α -strongly convex and β -smooth h with at most $Q(3(\beta + \beta')R/\xi)^d$ pieces such that $\|\nabla f(\theta) - \nabla h(\theta)\| \leq \xi$ for all $\theta \in \Theta$.

For this general class of non-convex functions, we derive the following generalization bound.

Theorem 2.2. Suppose that Assumptions 1 & 2 hold and $\Theta \subseteq \mathcal{B}_R^d(0)$. Suppose further that there exists $h : \Theta \times \mathcal{Z} \rightarrow \mathbb{R}$ such that for any $z \in \mathcal{Z}$, $h(\cdot; z)$ is piecewise α -strongly convex and β -smooth with P pieces on Θ satisfying

$$\|\nabla f(\theta; z) - \nabla h(\theta; z)\| \leq \xi, \quad (2.3)$$

for all $\theta \in \Theta$. For any $\eta \in (0, 2/\beta)$, let $\gamma := \sqrt{1 - 2\alpha\eta + \alpha\beta\eta^2}$. Then given $T \in \mathbb{N}$, with probability at least $1 - \delta$, for any $\theta^{(0)} \in \Theta$, $t \geq T$, and $i_1, \dots, i_t \in [n]$, we have

$$\left| \hat{F}(\theta^{(t)}) - F(\theta^{(t)}) \right| \leq \frac{BT}{n} + B\sqrt{\frac{T \log(nP) + \log(2/\delta)}{2n}} + 2L\left(\gamma^T R + \frac{1 - \gamma^T}{1 - \gamma} \eta \xi\right)$$

Remark. The bound above is stated in full generality and holds for any SGD iterate $t \geq T$ and any choice of $T \geq 1$. However, to obtain a meaningful generalization bound, one may choose, for example $T = O(\log(nR)/\log(\gamma^{-1}))$. In the case that $\gamma, B, L = \Theta(1)$, the bound simplifies to

$$\left| \hat{F}(\theta^{(t)}) - F(\theta^{(t)}) \right| \leq O\left(\sqrt{\frac{\log(nR) \log(nP) + \log(1/\delta)}{n}} + \xi\right). \quad (2.4)$$

Here, the first term in the bound is logarithmic in the number of pieces P and the last term scales linearly with the approximation error ξ . Thus, the generalization bound depends on the trade-off between the complexity of h through P , and how well ∇h approximates ∇f through ξ . In this regime, in contrast to algorithmic stability-based bounds, the above result does not grow with the number of SGD iterations, i.e. early stopping is not required for generalization in Theorem 2.2. We finally note that Theorem 2.2 holds for any f with an auxiliary gradient ∇f satisfying (2.3).

In light of Proposition 2.1, any piecewise smooth function can be approximated by a piecewise strongly convex and smooth function; in the worse case, the number of pieces P is at most $e^{\Omega(d)}$. Therefore in this pessimistic case, Theorem 2.2 recovers the classical algorithm-independent covering bound $\tilde{O}(B\sqrt{d/n})$ by choosing $\xi = O(1/(L\sqrt{n}))$ and $T = \Theta(\log(LRn))$. However, any value of P that is sub-exponential in dimension yields improved generalization bounds. In particular in the next section, we consider certain (non-convex) statistical models and carefully design h so that Theorem 2.2 improves the existing generalization error bounds. In contrast, uniform stability-based bounds for piecewise strongly convex and smooth functions are in general $\Omega(1)$ after sufficiently many SGD iterations; see Appendix F. That is, the contractivity-based bounds in [KWS22] cannot be directly extended to piecewise contractivity (2.3).

The proof of Theorem 2.2 relies on a modified version of the covering construction presented in Section 2.1. First, we define the auxiliary parameter update $g_{i,p}(\theta) := \theta - \nabla h_p(\theta; z_i)$ where $h_p(\cdot; z_i)$

denotes the strongly convex and smooth function satisfying $\nabla f(\cdot; z_i) \approx \nabla h_p(\cdot; z_i)$ on the p -th piece. We show that $\Psi'_T(0) := \{g_{i_T, p_T} \circ \dots \circ g_{i_1, p_1}(0) : i_1, \dots, i_T \in [n], p_1, \dots, p_T \in [P]\}$ ε -covers $\Psi_{\geq T}(\theta^{(0)})$ for any $\theta^{(0)} \in \Theta$ as in Lemma 2.1. Since $|\Psi'_T| \leq (nP)^T$, applying a concentration inequality and the union bound over the localized cover yields Theorem 2.2. Formal proofs of Proposition 2.1 and Theorem 2.2 are provided in Appendices B.6 and B.7, respectively.

3 Applications

In this section, we use our localized covering to prove generalization bounds for multi-index linear models, multi-class support vector machines, and K -means clustering for both hard and soft label setups, trained by SGD, improving the state-of-the-art results known for these models.

3.1 Multi-index linear models

Given a sample $z = (y, x) \in \mathcal{Y} \times \mathcal{B}_{R_x}^d(0)$, consider the ℓ_2 -regularized loss in a *multi-index model* with K indices parameterized by $\theta = (\theta_j)_{j=1}^K \in \mathbb{R}^{d \times K}$

$$f(\theta; z) := \ell(\theta_1^\top x, \dots, \theta_K^\top x; y) + \sum_{j=1}^K \frac{\lambda}{2} \|\theta_j\|^2. \quad (3.1)$$

Characterizing the generalization properties of multi-index models is an important problem with many applications including regression, classification, dimension reduction, and learning a single-layer of a neural network. In the following theorem, we derive a generalization bound for multi-index models by approximating each $f(\cdot; z)$ with a piecewise strongly convex and smooth function, and then applying Theorem 2.2. The proof is given in Appendix C.1.

Theorem 3.1. *Let $\theta = (\theta_j)_{j=1}^K$ and $\Theta = \prod_{j=1}^K \Theta_j$ for some convex $\Theta_j \subseteq \mathcal{B}_R^d(0)$. Suppose that f satisfies (3.1), Assumptions 1 & 2 hold, and $\ell(\cdot, \dots, \cdot; y)$ is piecewise β -smooth with Q pieces on Θ for all $y \in \mathcal{Y}$. For any $\eta \in (0, 2/\lambda)$, let $\gamma := |1 - \eta\lambda|$, $T := \max\left\{\left\lceil \frac{\log(3LRn)}{\log(1/\gamma)} \right\rceil, 0\right\}$, and $P := \text{poly}(\beta, \eta, K, L, R, R_x, T, n)$. Then, with probability at least $1 - \delta$, for any $\theta^{(0)} \in \Theta$, $t \geq T$, and $i_1, \dots, i_t \in [n]$, we have*

$$\left| \hat{F}(\theta^{(t)}) - F(\theta^{(t)}) \right| \leq \frac{BT + 1}{n} + B \sqrt{\frac{T \log(nP^K Q) + \log(2/\delta)}{2n}}.$$

Generalization behavior of multi-index models has received considerable attention; a subset of notable results include [Gue02, Zha04, JKZ⁺12, CMR13, LDZK19]. Under the same conditions of Theorem 3.1, existing state-of-the-art bounds scale at least linearly in \sqrt{d} or K , while our result is dimension-free and scales with \sqrt{K} . Specifically, the result in [LDZK19] translated to our setting reads (ℓ is L -Lipschitz, and $f(\theta; z) \in [0, B]$ which are stronger conditions than Assumptions 1 & 2)

$$\left| \hat{F}(\theta) - F(\theta) \right| \leq \tilde{O}\left(\frac{KLRR_x + B}{\sqrt{n}}\right), \quad (3.2)$$

for all $\theta \in \Theta$, which is obtained via [LDZK19, Cor 3 & 9] with $\Lambda \leftarrow R$ and $p \leftarrow \infty$. While the bound in (3.2) is linear in $KLRR_x$, our bound is linear in \sqrt{K} and logarithmic in L, R, R_x . The significance of this improvement can be better seen, for example, when learning the first layer of neural networks, for which L can be very large. Moreover, if $x \sim N(0, I)$, then $\|x\| = \tilde{\Omega}(\sqrt{d})$ with high probability; that is, (3.2) is linear in $R_x = \Omega(\sqrt{d})$ but ours only scales with $\sqrt{\log d}$, which can make a significant difference especially in the overparameterized regime. We should note that compared to our bound, (3.2) does not require the smoothness of ℓ .

For a specific application of Theorem 3.1, consider the multi-class support vector machines, i.e. for $\mathcal{Y} = [K]$ and $\rho : [-RR_x, RR_x] \rightarrow [0, B]$ is L -Lipschitz and β -smooth, the objective is given as

$$f(\theta; z) = \max_{y' \neq y} \rho\left(\theta_y^\top x - \theta_{y'}^\top x\right) + \frac{\lambda}{2} \sum_{j=1}^K \|\theta_j\|^2.$$

Theorem 3.1 provides a generalization bound of $\tilde{O}(\sqrt{K/n})$ that improves the existing bounds which scale at least linearly with K and/or \sqrt{d} [Zha04, DSBDS15, LDBK15, LDZK19]; see e.g. [LDZK19, Sec 2.1] for similar related results.

3.2 K -means clustering

We can also provide generalization error bounds for the problem of K -means clustering in both hard and soft label setups. As before, we consider SGD as an optimizer which has been the focus of many works in this context [BB94, Scu10, TM17]. We assume throughout this section that the samples are supported on a bounded domain, and the SGD iterations are applied without projection.

Generalization properties of K -means clustering has been studied for decades [Ant05, Lev13, TTJC15, TM16]. While most existing bounds are at least linear in K , [LL21] recently provided an improved bound in the hard label setup, which is of order $(K \log^3 n/n)^{1/2}$ for bounded inputs. However, state-of-the-art bounds in the soft label setup were still linear in K . Below, we establish a bound of order $\sqrt{K \log n/n}$ in the soft label setup, which is the first bound that is sublinear in K . We further improve the bounds of [LL21] in the hard label setup, but by a logarithmic factor.

Soft K -means clustering. While the hard K -means clustering assigns exactly one cluster to each point using one-hot encoding, the soft K -means clustering allows *soft labels* within the probability simplex. Specifically, given $\zeta > 0$ and the samples z_1, \dots, z_n , the soft K -means clustering algorithm alternates between updating the soft labels $(w_j(z_i; \theta))_{j=1}^K$ of z_i for all $i \in [n]$, and estimating the cluster centers $\theta = (\theta_j)_{j=1}^K \in (\mathcal{B}_R^d(0))^K$ using the classical update rule [Mac03]

$$w_j(z_i; \theta) \leftarrow \frac{\exp(-\zeta \|\theta_j - z_i\|^2)}{\sum_{k=1}^K \exp(-\zeta \|\theta_k - z_i\|^2)}, \quad \theta_j \leftarrow \frac{\sum_{i=1}^n w_j(z_i; \theta) z_i}{\sum_{i=1}^n w_j(z_i; \theta)}. \quad (3.3)$$

Here, this procedure is equivalent to a special case of expectation-maximization algorithm, which converges to a local minimum of the following objective (see Appendix C.2 for details)

$$\hat{F}(\theta) = \frac{1}{n} \sum_{i=1}^n \left\{ f(\theta; z_i) := -\frac{1}{\zeta} \log \left(\sum_{j=1}^K \exp(-\zeta \|\theta_j - z_i\|^2) \right) \right\}. \quad (3.4)$$

Instead of running the standard alternating procedure (3.3), we directly minimize (3.4) using SGD and derive generalization bounds for the soft K -means clustering by approximating the objective (3.4) with piecewise strongly convex and smooth functions. The proof is presented in Appendix C.3.

Theorem 3.2. For $B := 4(R+1)^2$ and for any $\eta \in (0, Ke^{-\zeta B})$, let $\gamma := \sqrt{1 - \frac{4\eta e^{-\zeta B}}{K} + \frac{4\eta^2}{K^2}}$, $L := \frac{4R}{\sqrt{K}} e^{\zeta B}$, $T := \max \left\{ \left\lceil \frac{\log(3LRn)}{\log(1/\gamma)} \right\rceil, 0 \right\}$, and $P := \text{poly}(\eta, \zeta, B, K, R, T, n, e^{\zeta B})$. Then, with probability at least $1 - \delta$, for any $\theta^{(0)} \in \Theta$, $t \geq T$, and $i_1, \dots, i_t \in [n]$, we have

$$\left| \hat{F}(\theta^{(t)}) - F(\theta^{(t)}) \right| \leq \frac{BT+1}{n} + B \sqrt{\frac{T \log(nPK) + \log(2/\delta)}{2n}}.$$

The above generalization bound scales with \sqrt{K} , which improves the previously known bound $\tilde{O}(K/\sqrt{n})$ for clustering with soft labels [LL21]. Theorem 3.2 also implies $\tilde{O}(\sqrt{K/n})$ bound under $R, \zeta = \Theta(1)$ and $\eta = \Theta(K)$, which coincides with the best known rate in the hard label setup. Here, the choice of $\eta = \Theta(K)$ may look odd; however, it is indeed practical since the derivative of the objective function f scales with $1/K$, i.e. $\|\partial f(\theta; z)/\partial \theta_j\| \leq 4Re^{\zeta B}/K$.

Hard K -means clustering. Under the same setup as before, we minimize the objective function

$$\hat{F}(\theta) = \frac{1}{n} \sum_{i=1}^n \left\{ f(\theta; z_i) := \min_{j \in [K]} \|\theta_j - z_i\|^2 \right\}. \quad (3.5)$$

Note that (3.5) coincides with the soft K -means objective (3.4) as $\zeta \rightarrow \infty$ if $\|\theta_j - z_i\| \neq \|\theta_k - z_i\|$ for all i, j, k . Since $f(\theta; z) = \min_{j \in [K]} \|\theta_j - z\|^2$ may not be differentiable, we use an auxiliary gradient at a non-differentiable θ , i.e. for a randomly or deterministically chosen $\mathcal{S} \subseteq \arg \min_{j \in [K]} \|\theta_j - z\|^2$,

$$\frac{\partial}{\partial \theta_j} f(\theta; z) := \begin{cases} 0 & \text{if } j \notin \mathcal{S} \\ 2(\theta_j - z) & \text{if } j \in \mathcal{S} \end{cases}.$$

For example, one may choose a single cluster index ($|S|=1$) and compute the corresponding gradient. The following result characterizes the generalization of hard K -means clustering trained by SGD using this auxiliary gradient. Its proof is deferred to Appendix C.4.

Theorem 3.3. *For any $\eta \in (0, 1)$, let $B := 4R^2$, $\gamma := |1 - 2\eta|$, $T := \max \left\{ \left\lceil \frac{\log(16\sqrt{KR^2n})}{\log(1/\gamma)} \right\rceil, 0 \right\}$.*

Then, given $\theta^{(0)} \in \Theta$, the following bound holds with probability at least $1 - \delta$, for any $t \geq 0$ and $i_1, \dots, i_t \in [n]$, we have

$$\left| \hat{F}(\theta^{(t)}) - F(\theta^{(t)}) \right| \leq \frac{BKT + 1}{n} + B\sqrt{\frac{KT \log(2n) + \log(2/\delta)}{2n}}.$$

We note that under $\eta, R = \Theta(1)$ and $K = O(n)$, Theorem 3.3 provides $O((K \log^2 n/n)^{1/2})$ generalization bound which improves the bound given by [LL21], but this time by a logarithmic factor.

4 Discussions

Beyond (piecewise) strong convexity and smoothness. The localized covering construction we utilized is essentially based on the (piecewise) contractivity of SGD updates for (piecewise) strongly convex and smooth functions. However, local strong convexity is by no means necessary and this analysis can be extended to a broader class of objective functions.

Consider, for example, an objective function f that is uniformly convex [DN21] and has Hölder continuous gradient [Nes15]. Then, for any $\varepsilon > 0$, there exists a step size $\eta > 0$ and $T = T_\varepsilon$ such that applying T synchronously coupled SGD updates on any two points makes the distance between them smaller than ε . Consequently, an analog of Lemma 2.1 can be established under this setup as well, for a properly chosen η, ε, T . We highlight that this class of functions is already covered in our framework via Proposition 2.1; however, an analysis based on the actual curvature (as opposed to a piecewise approximation) of the objective may provide tighter generalization bounds.

Extension to SGD without projection. Projection operation is only needed to ensure that the SGD iterates stay bounded. Indeed, the projection is not needed in the presence of explicit regularization, or under a dissipativity-type condition on the objective [RRT17, EMS18, YBVE21, EH21]. For example, if $f(\cdot; z)$ is Lipschitz continuous, it is a straightforward exercise to show that SGD iterations are bounded in the presence of ℓ_2 -regularization and bounded initialization.

Extension to mini-batch setup and different sampling schemes. The results we presented can be easily extended to the mini-batch setting. To see this, note that the localization is based on the (piecewise) contractivity of a single iteration of the algorithm, and we have that if all g_1, \dots, g_n are contractive, then an average of their subset is also contractive. Moreover, our generalization bounds hold for any sampling scheme, e.g. sampling without replacement, random shuffling, data-dependent sampling, since the covering construction is based on an ε -cover of the union of all possible trajectories of SGD, which is independent of the underlying sampling scheme.

Extension to contractive stochastic optimization methods. In this paper, our main focus was the SGD algorithm and its generalization properties. However, it is straightforward to adapt our framework to a general stochastic optimization setup. For example, the next result follows from the identical steps leading to Theorem 2.2; hence, its proof is omitted.

Definition 3. *An “iterative stochastic algorithm” using $g : \Theta \times \mathcal{Z} \rightarrow \Theta$ performs the update $\theta^{(t)} = g_{i_t}(\theta^{(t-1)}) := g(\theta^{(t-1)}; z_{i_t})$ at iteration t , for a random sample $i_t \in [n]$.*

Definition 4. *$g : \Theta \rightarrow \Theta$ is “piecewise γ -contractive with P pieces on Θ ” if there exists a partition $\mathcal{P}_1, \dots, \mathcal{P}_P$ of Θ and γ -contractive h_1, \dots, h_P on Θ such that $g = h_p$ on \mathcal{P}_p for all $p \in [P]$.*

Theorem 4.1. *Suppose that Assumptions 1 & 2 hold, $\Theta \subset \mathcal{B}_R^d(0)$, and an iterative stochastic algorithm g is given. Suppose further that there exists $h : \Theta \times \mathcal{Z} \rightarrow \Theta$ such that for any $z \in \mathcal{Z}$, $h(\cdot; z)$ is piecewise γ -contractive with P pieces on Θ satisfying*

$$\|g(\theta; z) - h(\theta; z)\| \leq \xi$$

for all $\theta \in \Theta$. Choose $T \in \mathbb{N}$. Then with probability at least $1 - \delta$, for any $\theta^{(0)} \in \Theta$, $t \geq T$, and $i_1, \dots, i_t \in [n]$, we have

$$\left| \hat{F}(\theta^{(t)}) - F(\theta^{(t)}) \right| \leq \frac{BT}{n} + B\sqrt{\frac{T \log(nP) + \log(2/\delta)}{2n}} + 2L \left(\gamma^T R + \frac{1 - \gamma^T}{1 - \gamma} \xi \right).$$

We note that a concurrent work by [KWS22] requires *global* contractivity and their bounds are in expectation, whereas Theorem 4.1 is a high-probability statement which only requires *piecewise* contractivity as stated in Definition 4.

Limitation of our results. We outline a few limitations of our current analysis. In this paper, we only considered the generalization error (also referred to as the generalization gap), and our main result in the non-convex regime cannot be easily translated to a bound on the excess risk. This additional step would require a bound on the optimization error; nevertheless, establishing such bounds in the non-convex regime is highly non-trivial. It is worth highlighting that without assuming convexity, algorithmic stability-based bounds also suffer from this limitation [HRS15, FV19, BKZ20].

Because of the shared contractivity parameter γ across different iterations and pieces, our method cannot be easily extended to varying step size. This is indeed a limitation of our analysis as varying step size schedules are oftentimes used in practice. We leave this extension as a future work.

Acknowledgments

SP was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-00079, Artificial Intelligence Graduate School Program, Korea University) and Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2022R1F1A1076180). US was funded in part by the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), and the European Research Council Starting Grant DYNASTY – 101039676. MAE was supported by NSERC Grant [2019-06167], CIFAR AI Chairs program, and CIFAR AI Catalyst grant.

References

- [AKL21] Idan Amir, Tomer Koren, and Roi Livni, *SGD Generalizes Better Than GD (And Regularization Doesn’t Help)*, Conference on Learning Theory, 2021.
- [Ant05] András Antos, *Improved Minimax Bounds on the Test and Training Distortion of Empirically Designed Vector Quantizers*, IEEE Transactions on Information Theory (2005).
- [BB94] Leon Bottou and Yoshua Bengio, *Convergence Properties of the K-Means Algorithms*, Advances in Neural Information Processing Systems, 1994.
- [BE02] Olivier Bousquet and André Elisseeff, *Stability and Generalization*, Journal of Machine Learning Research (2002).
- [BFGT20] Raef Bassily, Vitaly Feldman, Cristóbal Guzmán, and Kunal Talwar, *Stability of stochastic gradient descent on nonsmooth convex losses*, Advances in Neural Information Processing Systems, 2020.
- [BKZ20] Olivier Bousquet, Yegor Klochkov, and Nikita Zhivotovskiy, *Sharper Bounds for Uniformly Stable Algorithms*, Conference on Learning Theory, 2020.
- [CDE⁺21] Alexander Camuto, George Deligiannidis, Murat A. Erdogdu, Mert Gürbüzbalaban, Umut Şimşekli, and Lingjiong Zhu, *Fractal Structure and Generalization Properties of Stochastic Optimization Algorithms*, Advances in Neural Information Processing Systems, 2021.
- [CMR13] Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh, *Multi-Class Classification with Maximum Margin Multiple Kernel*, International Conference on Machine Learning, 2013.
- [DDB20] Aymeric Dieuleveut, Alain Durmus, and Francis Bach, *Bridging the Gap between Constant Step Size Stochastic Gradient Descent and Markov Chains*, The Annals of Statistics (2020).
- [DN21] Nikita Doikov and Yurii Nesterov, *Minimizing uniformly convex functions by cubic regularization of newton method*, Journal of Optimization Theory and Applications (2021).
- [DSBDSS15] Amit Daniely, Sivan Sabato, Shai Ben-David, and Shai Shalev-Shwartz, *Multiclass Learnability and the ERM Principle*, Journal of Machine Learning Research (2015).
- [EH21] Murat A Erdogdu and Rasa Hosseinzadeh, *On the convergence of langevin monte carlo: The interplay between tail growth and smoothness*, Conference on Learning Theory, PMLR, 2021, pp. 1776–1822.

- [EMS18] Murat A Erdogdu, Lester Mackey, and Ohad Shamir, *Global non-convex optimization with discretized diffusions*, Advances in Neural Information Processing Systems **31** (2018).
- [Fal14] Kenneth Falconer, *Fractal geometry: mathematical foundations and applications*, Wiley, 2014.
- [Fel16] Vitaly Feldman, *Generalization of ERM in Stochastic Convex Optimization: The Dimension Strikes Back*, Advances in Neural Information Processing Systems (2016).
- [FO21] Farzan Farnia and Asuman Ozdaglar, *Train simultaneously, generalize better: Stability of gradient-based minimax learners*, International Conference on Machine Learning, 2021.
- [FV19] Vitaly Feldman and Jan Vondrak, *High probability generalization bounds for uniformly stable algorithms with nearly optimal rate*, Conference on Learning Theory, 2019.
- [Gue02] Yann Guermeur, *Combining Discriminant Models with New Multi-Class SVMs*, Pattern Analysis & Applications (2002).
- [Hoe63] Wassily Hoeffding, *Probability Inequalities for Sums of Bounded Random Variables*, Journal of the American Statistical Association (1963).
- [HRS15] Moritz Hardt, Benjamin Recht, and Yoram Singer, *Train faster, generalize better: Stability of stochastic gradient descent*, arXiv preprint arXiv:1509.01240 (2015).
- [JKZ⁺12] Robert Jenssen, Marius Kloft, Alexander Zien, Sören Sonnenburg, and Klaus-Robert Müller, *A Scatter-Based Prototype Framework and Multi-Class Extension of Support Vector Machines*, PLoS one (2012).
- [KWS22] Leo Kozachkov, Patrick M. Wensing, and Jean-Jacques Slotine, *Generalization in Supervised Learning Through Riemannian Contraction*, arXiv preprint arXiv:2201.06656 (2022).
- [LDBK15] Yunwen Lei, Ürün Dogan, Alexander Binder, and Marius Kloft, *Multi-class SVMs: From Tighter Data-Dependent Generalization Bounds to Novel Algorithms*, Advances in Neural Information Processing Systems, 2015.
- [LDZK19] Yunwen Lei, Ürün Dogan, Ding-Xuan Zhou, and Marius Kloft, *Data-dependent generalization bounds for multi-class classification*, IEEE Transactions on Information Theory (2019).
- [Lev13] Clément Levrard, *Fast rates for empirical vector quantization*, Electronic Journal of Statistics (2013).
- [LL21] Shaojie Li and Yong Liu, *Sharper Generalization Bounds for Clustering*, International Conference on Machine Learning, 2021.
- [LLQ20] Jian Li, Xuanyuan Luo, and Mingda Qiao, *On Generalization Error Bounds of Noisy Gradient Methods for Non-Convex Learning*, International Conference on Learning Representations, 2020.
- [LLY21] Yunwen Lei, Mingrui Liu, and Yiming Ying, *Generalization Guarantee of SGD for Pairwise Learning*, International Conference on Artificial Intelligence and Statistics, 2021.
- [Lon17] Ben London, *A PAC-Bayesian analysis of randomized learning with application to stochastic gradient descent*, Advances in Neural Information Processing Systems (2017).
- [LY20] Yunwen Lei and Yiming Ying, *Fine-Grained Analysis of Stability and Generalization for Stochastic Gradient Descent*, International Conference on Machine Learning, 2020.
- [Mac03] David J.C. MacKay, *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, 2003.
- [Mil19] Tristan Milne, *Piecewise Strong Convexity of Neural Networks*, Advances in Neural Information Processing Systems, 2019.
- [Nes15] Yu Nesterov, *Universal gradient methods for convex optimization problems*, Mathematical Programming (2015).
- [PLYS21] Sejun Park, Jaeho Lee, Chulhee Yun, and Jinwoo Shin, *Provable Memorization via Deep Neural Networks using Sub-linear Parameters*, Conference on Learning Theory, 2021.
- [Pri12] Simon J.D. Prince, *Computer vision: models, learning, and inference*, Cambridge University Press, 2012.

- [RRT17] Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky, *Non-convex learning via Stochastic Gradient Langevin Dynamics: a nonasymptotic analysis*, Conference on Learning Theory, 2017.
- [Scu10] David Sculley, *Web-Scale K-Means Clustering*, International Conference on World Wide Web, 2010.
- [SHN⁺18] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro, *The Implicit Bias of Gradient Descent on Separable Data*, The Journal of Machine Learning Research (2018).
- [SSBD14] Shai Shalev-Shwartz and Shai Ben-David, *Understanding Machine Learning: From Theory to Algorithms*, Cambridge university press, 2014.
- [SSDE20] Umut Şimşekli, Ozan Sener, George Deligiannidis, and Murat A. Erdogdu, *Hausdorff Dimension, Heavy Tails, and Generalization in Neural Networks*, Advances in Neural Information Processing Systems, 2020.
- [SSSS09] Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan, *Stochastic Convex Optimization*, Conference on Learning Theory, 2009.
- [TM16] Cheng Tang and Claire Monteleoni, *On Lloyd’s Algorithm: New Theoretical Insights for Clustering in Practice*, International Conference on Artificial Intelligence and Statistics, 2016.
- [TM17] ———, *Convergence Rate of Stochastic k-means*, International Conference on Artificial Intelligence and Statistics, 2017.
- [TTJC15] Matthew Thorpe, Florian Theil, Adam M Johansen, and Neil Cade, *Convergence of the k-Means Minimization Problem using Γ -Convergence*, SIAM Journal on Applied Mathematics (2015).
- [Ver18] Roman Vershynin, *High-dimensional probability: An introduction with applications in data science*, Cambridge University Press, 2018.
- [YBVE21] Lu Yu, Krishnakumar Balasubramanian, Stanislav Volgushev, and Murat A. Erdogdu, *An Analysis of Constant Step Size SGD in the Non-convex Regime: Asymptotic Normality and Bias*, Advances in Neural Information Processing Systems, 2021.
- [YKM21] Chulhee Yun, Shankar Krishnan, and Hossein Mobahi, *A Unifying View on Implicit Bias in Training Linear Neural Networks*, International Conference on Learning Representations, 2021.
- [ZBH⁺21] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals, *Understanding deep learning (still) requires rethinking generalization*, Communications of the ACM (2021).
- [Zha04] Tong Zhang, *Statistical Analysis of Some Multi-Category Large Margin Classification Methods*, Journal of Machine Learning Research (2004).

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [No] This paper only contains theoretical results on the generalization properties of SGD.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Technical results

All results presented in this section are standard. However, we provide their proofs for convenience.

Lemma A.1. *Let \mathcal{S} be a Hilbert space associated with the norm $\|\cdot\|$ induced by the inner product $\langle \cdot, \cdot \rangle$. Let $\mathcal{C} \subset \mathcal{S}$ be a convex set and $\Pi_{\mathcal{C}}(x) := \arg \min_{z \in \mathcal{C}} \|x - z\|$ be a projection of x onto \mathcal{C} . Then,*

$$\|\Pi_{\mathcal{C}}(x) - \Pi_{\mathcal{C}}(y)\| \leq \|x - y\| \quad \forall x, y \in \mathcal{S}.$$

Proof. Since \mathcal{C} is convex, $\Pi_{\mathcal{C}}(x)$ is well-defined. From the definition of the projection and the convexity of \mathcal{C} , we have

$$\begin{aligned} \|\Pi_{\mathcal{C}}(x) - x\|^2 &\leq \|(1-t)\Pi_{\mathcal{C}}(x) + t\Pi_{\mathcal{C}}(y) - x\|^2 \\ &= \|\Pi_{\mathcal{C}}(x) - x + t(\Pi_{\mathcal{C}}(y) - \Pi_{\mathcal{C}}(x))\|^2 \end{aligned} \quad (\text{A.1})$$

for all $t \in [0, 1]$. Since we have

$$\frac{\partial}{\partial t} \|\Pi_{\mathcal{C}}(x) - x + t(\Pi_{\mathcal{C}}(y) - \Pi_{\mathcal{C}}(x))\|^2 = 2\langle \Pi_{\mathcal{C}}(x) - x, \Pi_{\mathcal{C}}(y) - \Pi_{\mathcal{C}}(x) \rangle + 2t\|\Pi_{\mathcal{C}}(y) - \Pi_{\mathcal{C}}(x)\|^2 \quad (\text{A.2})$$

and $\frac{\partial}{\partial t} \|\Pi_{\mathcal{C}}(x) - x + t(\Pi_{\mathcal{C}}(y) - \Pi_{\mathcal{C}}(x))\|^2 \geq 0$ at $t = 0$ by (A.1), we have

$$\langle \Pi_{\mathcal{C}}(x) - x, \Pi_{\mathcal{C}}(y) - \Pi_{\mathcal{C}}(x) \rangle \geq 0. \quad (\text{A.3})$$

Likewise, we have

$$\langle \Pi_{\mathcal{C}}(y) - y, \Pi_{\mathcal{C}}(x) - \Pi_{\mathcal{C}}(y) \rangle \geq 0. \quad (\text{A.4})$$

Now, consider the function

$$f(t) := \|\Pi_{\mathcal{C}}(x) - \Pi_{\mathcal{C}}(y) + t(x - \Pi_{\mathcal{C}}(x) - y + \Pi_{\mathcal{C}}(y))\|^2,$$

i.e. $f(0) = \|\Pi_{\mathcal{C}}(x) - \Pi_{\mathcal{C}}(y)\|^2$ and $f(1) = \|x - y\|^2$. Then $f(0) \leq f(1)$ since the following inequality holds:

$$\begin{aligned} \frac{d}{dt} f(t) &= 2\langle \Pi_{\mathcal{C}}(x) - \Pi_{\mathcal{C}}(y), x - \Pi_{\mathcal{C}}(x) - y + \Pi_{\mathcal{C}}(y) \rangle + 2t\|x - \Pi_{\mathcal{C}}(x) - y + \Pi_{\mathcal{C}}(y)\|^2 \\ &= 2\langle \Pi_{\mathcal{C}}(x) - x, \Pi_{\mathcal{C}}(y) - \Pi_{\mathcal{C}}(x) \rangle + 2\langle \Pi_{\mathcal{C}}(y) - y, \Pi_{\mathcal{C}}(x) - \Pi_{\mathcal{C}}(y) \rangle \\ &\quad + 2t\|x - \Pi_{\mathcal{C}}(x) - y + \Pi_{\mathcal{C}}(y)\|^2 \\ &\geq 0 \quad \forall t \in [0, 1] \end{aligned}$$

where the inequality follows from (A.3) and (A.4). This completes the proof of Lemma A.1. \square

Lemma A.2 (Hoeffding's inequality [Hoe63]). *Let X be a random variable on $[a, b]$ and X_1, \dots, X_n be independent copies of X . Then, it holds that*

$$\mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n X_i - \mathbb{E}[X] \right| \geq \varepsilon \right) \leq 2 \exp \left(-\frac{2n\varepsilon^2}{(b-a)^2} \right).$$

Lemma A.3. *Suppose that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex and differentiable. Then, the following conditions are equivalent.*

1. f is β -smooth, i.e. $\|\nabla f(\theta) - \nabla f(\theta')\| \leq \beta\|\theta - \theta'\|$ for all $\theta, \theta' \in \mathbb{R}^d$.
2. $\frac{\beta}{2}\|\theta\|^2 - f(\theta)$ is convex for all $\theta \in \mathbb{R}^d$.
3. $f(\theta) - f(\theta') - \nabla f(\theta')^\top(\theta - \theta') \leq \frac{\beta}{2}\|\theta - \theta'\|^2$ for all $\theta, \theta' \in \mathbb{R}^d$.
4. $(\nabla f(\theta) - \nabla f(\theta'))^\top(\theta - \theta') \geq \frac{1}{\beta}\|\nabla f(\theta) - \nabla f(\theta')\|^2$ for all $\theta, \theta' \in \mathbb{R}^d$.

Proof.

1 \Rightarrow 2. Let $g(\theta) := \frac{\beta}{2}\|\theta\|^2 - f(\theta)$. Then for any $\theta \neq \theta'$,

$$\begin{aligned} (\nabla g(\theta) - \nabla g(\theta'))^\top (\theta - \theta') &= (\beta(\theta - \theta') - (\nabla f(\theta) - \nabla f(\theta'))^\top)^\top (\theta - \theta') \\ &= \beta\|\theta - \theta'\|^2 - (\nabla f(\theta) - \nabla f(\theta'))^\top (\theta - \theta') \\ &\geq \beta\|\theta - \theta'\|^2 - \|\nabla f(\theta) - \nabla f(\theta')\| \cdot \|\theta - \theta'\| \\ &\geq 0. \end{aligned}$$

In addition, by the mean value theorem, there exists $s \in (0, 1)$ such that for $\theta_s = s\theta + (1-s)\theta'$,

$$g(\theta') - g(\theta) = \nabla g(\theta_s)^\top (\theta' - \theta).$$

Using this, we can derive the following inequality

$$\begin{aligned} 0 &\leq (\nabla g(\theta_s) - \nabla g(\theta'))^\top (\theta_s - \theta') \\ &= s(\nabla g(\theta_s) - \nabla g(\theta'))^\top (\theta - \theta') \\ &= s(g(\theta) - g(\theta') - \nabla g(\theta')^\top (\theta - \theta')), \end{aligned}$$

which implies that g is convex.

2 \Leftrightarrow 3. The following relation shows the equivalence of the second and the third statements:

$$g(\theta) \geq g(\theta') + \nabla g(\theta')^\top (\theta - \theta') \Leftrightarrow f(\theta) - f(\theta') - \nabla f(\theta')^\top (\theta - \theta') \leq \frac{\beta}{2}\|\theta - \theta'\|^2.$$

2,3 \Rightarrow 4. We first introduce the following claim.

Claim 1. Suppose that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is convex and differentiable and have a global minimum θ^* . Then, $\frac{1}{2\beta}\|\nabla f(\theta)\|^2 \leq f(\theta) - f(\theta^*)$.

Proof. The statement of Claim 1 is a consequence of the following relation

$$f(\theta^*) = \inf_{\theta'} f(\theta') \leq \inf_{\theta'} f(\theta) + \nabla f(\theta)^\top (\theta' - \theta) + \frac{\beta}{2}\|\theta' - \theta\|^2 = f(\theta) - \frac{1}{2\beta}\|\nabla f(\theta)\|^2.$$

□

Let $f_\theta(\phi) := f(\phi) - \nabla f(\theta)^\top \phi$. Since $\frac{\beta}{2}\|\phi\|^2 - f_\theta(\phi)$ is convex and $\phi = \theta$ minimizes f_θ , from Claim 1, we have

$$\begin{aligned} f(\theta') - f(\theta) - \nabla f(\theta)^\top (\theta' - \theta) &= f_{\theta'}(\theta') - f_{\theta'}(\theta) \geq \frac{1}{2\beta}\|\nabla f(\theta) - \nabla f(\theta')\|^2, \\ f(\theta) - f(\theta') - \nabla f(\theta')^\top (\theta - \theta') &= f_{\theta'}(\theta) - f_{\theta'}(\theta') \geq \frac{1}{2\beta}\|\nabla f(\theta) - \nabla f(\theta')\|^2. \end{aligned}$$

Adding two above inequalities derive the fourth statement.

4 \Rightarrow 1. The following inequality is sufficient for deriving the first statement

$$\begin{aligned} \frac{1}{\beta}\|\nabla f(\theta) - \nabla f(\theta')\|^2 &\leq (\nabla f(\theta) - \nabla f(\theta'))^\top (\theta - \theta') \leq \|\nabla f(\theta) - \nabla f(\theta')\| \cdot \|\theta - \theta'\| \\ \Rightarrow \|\nabla f(\theta) - \nabla f(\theta')\| &\leq \beta\|\theta - \theta'\|. \end{aligned}$$

□

A.1 Contractivity of projected SGD for strongly convex and smooth objectives

Lemma A.4. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be α -strongly convex and β -smooth. Then for any $\eta \in (0, 2/\beta)$ and for any convex $\Theta \subset \mathbb{R}^d$, $\theta \mapsto \Pi_\Theta(\theta - \eta\nabla f(\theta))$ is $\sqrt{1 - 2\alpha\eta + \alpha\beta\eta^2}$ -contractive.

Proof. Let $g(\theta) := \theta - \eta \nabla f(\theta)$. Then for any $\theta, \theta' \in \mathbb{R}^d$,

$$\begin{aligned} \|g(\theta) - g(\theta')\|^2 &= \|\theta - \theta'\|^2 - 2\eta(\nabla f(\theta) - \nabla f(\theta'))^\top(\theta - \theta') + \eta^2\|\nabla f(\theta) - \nabla f(\theta')\|^2 \\ &\leq \|\theta - \theta'\|^2 - 2\eta \left(\left(1 - \frac{\beta\eta}{2}\right) \cdot \alpha\|\theta - \theta'\|^2 + \frac{\beta\eta}{2} \cdot \frac{1}{\beta}\|\nabla f(\theta) - \nabla f(\theta')\|^2 \right) + \eta^2\|\nabla f(\theta) - \nabla f(\theta')\|^2 \\ &= (1 - 2\alpha\eta + \alpha\beta\eta^2)\|\theta - \theta'\|^2. \end{aligned}$$

Here, the inequality is from the α -strong convexity and Lemma A.3, i.e.

$$\begin{aligned} (\nabla f(\theta) - \nabla f(\theta'))^\top(\theta - \theta') &\geq \alpha\|\theta - \theta'\|^2 \\ (\nabla f(\theta) - \nabla f(\theta'))^\top(\theta - \theta') &\geq \frac{1}{\beta}\|\nabla f(\theta) - \nabla f(\theta')\|^2. \end{aligned}$$

Using Lemma A.1 completes the proof of Lemma A.4. □

B Proofs of results in Section 2

B.1 Deriving generalization bounds using localized covers

In Lemma 2.1, we show that the union of all trajectories generated by $t \geq T = T_\varepsilon$ SGD iterations, i.e. $\bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq T}(\theta^{(0)})$, can be ε -covered by a set of points generated by exactly T SGD iterations, i.e. $\Psi_T(0)$. Here, each $g_{i_T} \circ \dots \circ g_{i_1}(0) \in \Psi_T(0)$ can be viewed as a *deterministic* algorithm mapping samples z_1, \dots, z_n to a parameter in Θ , where each $g_{i_T} \circ \dots \circ g_{i_1}(0)$ only depends on at most T samples z_{i_1}, \dots, z_{i_T} . Under this observation, in this section, we generalize our localized cover so that its elements are general deterministic algorithms depending on a small number of samples, i.e. not restricted to possible instances of SGD. Then, we derive generalization bounds using our (generalized) localized cover. To this end, we define algorithms depending on samples.

Definition 5. $\phi : \mathcal{Z}^n \rightarrow \Theta$ is a “(deterministic) algorithm depending on at most T samples” if there exists $\mathcal{I} \in \{\mathcal{S} \subset [n] : |\mathcal{S}| = T\}$ satisfying the following: for any $z_1, \dots, z_n, z'_1, \dots, z'_n \in \mathcal{Z}$ such that $z_i = z'_i$ for all $i \in \mathcal{I}$, $\phi(z_1, \dots, z_n) = \phi(z'_1, \dots, z'_n)$. Here, we refer to \mathcal{I} as the “set of sample indices determining ϕ .” The collection of all algorithms depending on at most T samples is denoted by \mathcal{A}_T .

In Definition 5, we define ϕ to be an algorithm depending on at most T samples if the value of ϕ can be fully determined by a subset of the samples $\{z_i\}_{i \in \mathcal{I}}$ for some $\mathcal{I} \in \{\mathcal{S} \subset [n] : |\mathcal{S}| = T\}$. We note that $\Psi_T(\theta^{(0)}) \subset \mathcal{A}_T$ for any $T \geq 0$ and $\theta^{(0)} \in \Theta$.

Now, we introduce the following theorem for deriving generalization bounds using a localized cover where each element in the cover is an algorithm depending on at most T samples.

Theorem B.1. *Suppose that Assumptions 1 & 2 hold and $\Theta \subset \mathbb{R}^d$. Let $\varepsilon > 0$, $\Phi_n \subseteq \mathcal{A}_n$, and $\Phi_{T,\varepsilon} \subseteq \mathcal{A}_T$ such that*

$$\psi(x_{1:n}) \in \bigcup_{\phi \in \Phi_{T,\varepsilon}} \mathcal{B}_\varepsilon^d(\phi(x_{1:n})) \quad \text{for all } x_{1:n} \in \mathcal{Z}^n \text{ and } \psi \in \Phi_n. \quad (\text{B.1})$$

Let μ be a distribution over \mathcal{Z} and $z_{1:n} = (z_1, \dots, z_n)$ is such that z_i 's are i.i.d. samples from μ . Then, with probability at least $1 - \delta$ over the sampling distribution of $z_{1:n}$, for any $\psi \in \Psi$,

$$\left| \hat{F}(\psi(z_{1:n})) - F(\psi(z_{1:n})) \right| \leq \frac{BT}{n} + B \sqrt{\frac{\log(2|\Phi_{T,\varepsilon}|/\delta)}{2n}} + 2L\varepsilon.$$

Theorem B.1 implies that if (i) Φ_n can be covered by a set $\Phi_{T,\varepsilon}$ of algorithms depending on at most T samples and (ii) B, L, T , and $|\Phi_{T,\varepsilon}|$ are independent of d , then a dimension-independent generalization bound can be derived. We note that the assumption (B.1) is a generalization of the observation in Lemma 2.1 since $\Psi_{\geq T} \subset \mathcal{A}_n$ and $\Psi_T \subset \mathcal{A}_T$.

Proof of Theorem B.1. From Assumption 1 and the assumption (B.1), for any $\psi \in \Phi_n$, we have

$$\left| \hat{F}(\psi(z_{1:n})) - F(\psi(z_{1:n})) \right| \leq \sup_{\phi \in \Phi_{T,\varepsilon}} \left| \hat{F}(\phi(z_{1:n})) - F(\phi(z_{1:n})) \right| + 2L\varepsilon. \quad (\text{B.2})$$

Namely, if we bound $|\hat{F}(\phi(z_{1:n})) - F(\phi(z_{1:n}))|$ for all $\phi \in \Phi_{T,\varepsilon}$, the bound in Theorem B.1 follows. Since the statement of Theorem B.1 is trivial if $T = n$ or $|\Phi_{T,\varepsilon}| = \infty$, we assume $T < n$ and $|\Phi_{T,\varepsilon}| < \infty$. Now, we derive the target bound: for $\phi \in \Phi_{T,\varepsilon}$, $\theta := \phi(z_{1:n}) \in \Theta$, and the set of indices \mathcal{I}_ϕ determining ϕ with $|\mathcal{I}_\phi| \leq T$,

$$\begin{aligned} \left| \hat{F}(\theta) - F(\theta) \right| &\leq \left| \frac{1}{n} \sum_{i \in \mathcal{I}_\phi} f(\theta; z_i) - F(\theta) \right| + \left| \frac{1}{n} \sum_{i \in [n] \setminus \mathcal{I}_\phi} f(\theta; z_i) - F(\theta) \right| \\ &\leq \frac{B|\mathcal{I}_\phi|}{n} + \left| \frac{1}{n} \sum_{i \in [n] \setminus \mathcal{I}_\phi} f(\theta; z_i) - F(\theta) \right| \\ &\leq \frac{B|\mathcal{I}_\phi|}{n} + \frac{n - |\mathcal{I}_\phi|}{n} \sqrt{\frac{B^2 \log(2|\Phi_{T,\varepsilon}|/\delta)}{2(n - |\mathcal{I}_\phi|)}} \quad \text{w.p. } 1 - \delta/|\Phi_{T,\varepsilon}| \\ &\leq \frac{BT}{n} + B \sqrt{\frac{\log(2|\Phi_{T,\varepsilon}|/\delta)}{2n}} \quad \text{w.p. } 1 - \delta/|\Phi_{T,\varepsilon}|. \end{aligned} \quad (\text{B.3})$$

For the first inequality in (B.3), we use the triangle inequality to upper bound $|\hat{F}(\theta) - F(\theta)|$ using two terms: we only utilize the samples determining θ in the first term while remaining samples independent of θ are considered in the second term. The second inequality directly follows from Assumption 2. To bound the second term in RHS of the second inequality, one can apply the Hoeffding's inequality (see Lemma A.2), which leads us to the third inequality in (B.3). Here, the last inequality naturally follows. By using (B.2), (B.3), and by applying the union bound for all $\phi \in \Phi_{T,\varepsilon}$, we obtain the bound in Theorem B.1. \square

B.2 Proof of Theorem 2.1

First, observe that g_1, \dots, g_n are γ -contractive by Lemma A.4. Let $\varepsilon = 1/(2Ln)$; then by Lemma 2.1, we have

$$\bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq T}(\theta^{(0)}) \subseteq \bigcup_{\phi \in \Psi_T(0)} \mathcal{B}_\varepsilon^d(\phi). \quad (\text{B.4})$$

Now, we apply Theorem B.1 with

$$\Phi_n \leftarrow \bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq T}(\theta^{(0)}), \Phi_{T,\varepsilon} \leftarrow \Psi_T(0), \varepsilon \leftarrow \varepsilon, T \leftarrow T, \delta \leftarrow \delta.$$

where the assumption (B.1) in Theorem B.1 is satisfied by (B.4). This completes the proof of Theorem 2.1.

B.3 Proof of Corollary 2.1

Let $\theta^{(t)} := g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)})$ and $\phi := g_{i_t} \circ \dots \circ g_{i_{t-T+1}}(0)$, i.e. ϕ is an algorithm depending on at most T samples. Since each g_i is γ -contractive by Lemma A.4, one can observe that

$$\|\theta^{(t)} - \phi\| \leq \gamma^T R \leq \frac{1}{2Ln} =: \varepsilon. \quad (\text{B.5})$$

Now, we apply Theorem B.1 with

$$\Phi_n \leftarrow \{\theta^{(t)}\}, \Phi_{T,\varepsilon} \leftarrow \{\phi\}, \varepsilon \leftarrow \varepsilon, T \leftarrow T, \delta \leftarrow \delta.$$

where the assumption (B.1) is satisfied by (B.5) and $|\Phi_\varepsilon| = 1$. This provides the bound in Corollary 2.1.

B.4 Proof of Corollary 2.2

The proof of Corollary 2.2 is simple. Since $\theta^{(0)}$, t , and i_1, \dots, i_t are fixed, $\theta^{(t)} = g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)})$ is an algorithm depending on at most t samples. Then by using Theorem B.1 with

$$\Phi_n, \Phi_{T,\varepsilon} \leftarrow \{\theta^{(t)}\}, \varepsilon \leftarrow 0, T \leftarrow t, \delta \leftarrow \delta,$$

we obtain the bound in Corollary 2.2.

B.5 Derivation of (2.2)

We first formally define the Hausdorff dimension.

Definition 6. Given $\mu \subset \mathbb{R}^d$, d_H defined below is the ‘‘Hausdorff dimension’’ of μ :

$$d_H := \inf\{s \geq 0 : h^s(\mu) = 0\},$$

$$h^s(\mu) := \liminf_{r \rightarrow 0} \left\{ \sum_{i=1}^k r_i^s : k \in \mathbb{N} \cup \{\infty\}, (r_i)_{i=1}^k \in (0, r)^k \right.$$

$$\left. \text{such that there exists } (\theta_i)_{i=1}^k \text{ satisfying } \mu \subset \bigcup_i \mathcal{B}_{r_i}(\theta_i) \right\}.$$

From the assumption on $f(\theta; z)$, one can observe that for all $i \in [n]$ and $\theta, \theta' \in \mathcal{B}_R^d(0)$,

$$\|g_i(\theta) - g_i(\theta')\| \leq \gamma \|\theta - \theta'\|$$

where $\gamma = 1 - \eta$ since $\alpha = \beta = 1$. Then, the assumption $\|\theta_{z_i}^* - \theta_{z_j}^*\| \geq 2\gamma R$ guarantees that $g_i(\text{int}(\mathcal{B}_R)) \cap g_j(\text{int}(\mathcal{B}_R)) = \emptyset$ where $\text{int}(\mathcal{S})$ denotes the interior of a set \mathcal{S} .

Finally, the following theorem shows that $d_H = \frac{\log n}{\log(1/\gamma)}$. Substituting d_H to the bound in Theorem 2.1 results in (2.2).

Theorem B.2 (Theorem 9.3 in [Fal14]). *Suppose that $g_i(\text{int}(\mathcal{B}_R)) \cap g_j(\text{int}(\mathcal{B}_R)) = \emptyset$ for all $i \neq j$ and $\|g_i(\theta) - g_i(\theta')\| = \gamma\|\theta - \theta'\|$ for all i . Then $d_H = \frac{\log n}{\log(1/\gamma)}$.*

B.6 Proof of Proposition 2.1

In this proof, we explicitly construct a piecewise β -strongly convex and β -smooth function, i.e. quadratic, which is always piecewise α -strongly convex and β -smooth for any $\alpha \in [0, \beta]$. Let $\{\mathcal{Q}_1, \dots, \mathcal{Q}_Q\}$ be the partition of $\mathcal{B}_R^d(0)$ such that $f = \ell_q$ on \mathcal{Q}_q for some smooth ℓ_q on $\mathcal{B}_R^d(0)$. Given $q \in [Q]$ and $\varepsilon > 0$, let $\mathcal{C}_{q,\varepsilon} = \{\phi_{q,1}, \dots, \phi_{q,P_q}\} \subset \Theta$ be an ε -cover of Θ with the minimum cardinality, i.e. $P_q \leq (3R/\varepsilon)^d$ for $\varepsilon \leq R$. Let $\{\mathcal{P}_{q,1}, \dots, \mathcal{P}_{q,P_q}\}$ be a partition of \mathcal{Q}_q where each $\mathcal{P}_{q,p}$ is defined as follows:

$$\mathcal{P}_{q,p} := \left\{ \theta \in \mathcal{Q}_q \setminus \bigcup_{r < p} \mathcal{P}_{q,r} : \|\phi_{q,p} - \theta\| \leq \min_{r > p} \|\phi_{q,r} - \theta\| \right\}.$$

For each $q \in [Q]$ and $p \in [P_q]$, we also define $h(\theta) := h_{q,p}(\theta)$ on $\mathcal{P}_{q,p}$ where

$$h_{q,p}(\theta) := \phi_{q,p} + \nabla \ell_q(\phi_{q,p})^\top (\theta - \phi_{q,p}) + \frac{\beta}{2} \|\theta - \phi_{q,p}\|^2$$

and ∇_j denotes the partial derivative with respect to the j -th entry. Then, one can observe that h is piecewise α -strongly convex and β -smooth with $Q(3R/\varepsilon)^d$ pieces for any $\alpha \leq \beta$. Furthermore, for any $\theta \in \mathcal{P}_{q,p}$, we have

$$\|\nabla f(\theta) - \nabla h(\theta)\| = \|\nabla \ell_q(\theta) - \nabla \ell_q(\phi_{q,p}) - \beta(\theta - \phi_{q,p})\| \leq (\beta + \beta')\varepsilon.$$

Choosing $\varepsilon := \xi/(\beta + \beta')$ completes the proof of Proposition 2.1.

B.7 Proof of Theorem 2.2

Given $z \in \mathcal{Z}$, let $\mathcal{P}_{z,1}, \dots, \mathcal{P}_{z,P}$ be a partition of Θ and $h_1(\cdot; z), \dots, h_P(\cdot; z)$ be α -strongly convex and β -smooth functions such that $h(\cdot; z) = h_p(\cdot; z)$ on $\mathcal{P}_{z,p}$ for all $p \in [P]$. Let $g_{i,p}(\theta) := \theta - \eta \nabla h_p(\theta; z_i)$, i.e. each $g_{i,p}$ is γ -contractive by Lemma A.4. Given $\psi^{(0)} \in \Theta$, $t \geq T$, and $i_1, \dots, i_t \in [n]$, let $\psi^{(s)} := g_{i_s} \circ \dots \circ g_{i_1}(\psi^{(0)})$ for all $s \in [t]$ and let $p_s \in [P]$ be an index satisfying $\psi^{(s-1)} \in \mathcal{P}_{z_{i_s}, p_s}$. Let $\phi := g_{i_t, p_t} \circ \dots \circ g_{i_{-T+1}, p_{-T+1}}(\psi^{(0)})$, i.e. ϕ is an algorithm depending on at most T samples. Then for any $\theta \in \Theta$, we have

$$\begin{aligned} \|\psi^{(s)} - g_{i_s, p_s}(\theta)\| &= \|\psi^{(s-1)} - \eta \nabla f(\psi^{(s-1)}; z_{i_s}) - g_{i_s, p_s}(\theta)\| \\ &= \|\psi^{(s-1)} - \eta \nabla h(\psi^{(s-1)}; z_{i_s}) + \eta \nabla h(\psi^{(s-1)}; z_{i_s}) - \eta \nabla f(\psi^{(s-1)}; z_{i_s}) - g_{i_s, p_s}(\theta)\| \\ &\leq \|g_{i_s, p_s}(\psi^{(s-1)}) - g_{i_s, p_s}(\theta)\| + \eta \|\nabla f(\psi^{(s-1)}; z_{i_s}) - \nabla h(\psi^{(s-1)}; z_{i_s})\| \\ &\leq \gamma \|\psi^{(s-1)} - \theta\| + \eta \xi \end{aligned}$$

This implies that

$$\|\psi^{(t)} - \phi\| \leq \gamma^T \|\psi^{(0)}\| + \eta \xi \sum_{t=0}^{T-1} \gamma^t \leq \gamma^T R + \frac{1 - \gamma^T}{1 - \gamma} \eta \xi =: \varepsilon. \quad (\text{B.6})$$

Now, we apply Theorem B.1 with

$$\begin{aligned} \Phi_n &\leftarrow \bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq T}(\theta^{(0)}), \quad \Phi_{T,\varepsilon} \leftarrow \{g_{i_T, p_T} \circ \dots \circ g_{i_1, p_1}(\theta) : i_1, \dots, i_T \in [n], p_1, \dots, p_T \in [P]\}, \\ \varepsilon &\leftarrow \varepsilon, \quad T \leftarrow T, \quad \delta \leftarrow \delta \end{aligned}$$

where the assumption (B.1) in Theorem B.1 is satisfied by (B.6) and $|\Phi_{T,\varepsilon}| \leq (nP)^T$. This leads us to the bound in Theorem 2.2.

C Proofs of results in Section 3

C.1 Proof of Theorem 3.1

Given $z = (y, x) \in \mathcal{Y} \times \mathcal{B}_{R_x}^d(0) = \mathcal{Z}$, let $\{\mathcal{Q}_{z,1}, \dots, \mathcal{Q}_{z,Q}\}$ be a partition of Θ such that $\ell(\cdot, \dots, \cdot; y)$ is β -smooth on $\mathcal{Q}_{z,q}$ for all $q \in [Q]$. For proving Theorem 3.1, we utilize Theorem 2.2 by approximating each $f(\cdot; z)$ using a piecewise strongly convex and smooth function $h(\cdot; z)$ with $P^K Q$ pieces. To this end, we first define P by

$$P := \left\lceil \frac{2RR_x}{\kappa} \right\rceil \quad \text{where} \quad \kappa := \frac{1}{12\beta\eta KLR_x Tn}.$$

Using P defined as above, for each $z = (y, x) \in \mathcal{Y} \times \mathcal{B}_{R_x}^d(0)$, we construct a partition $\{\mathcal{P}_{z,q,p_1,\dots,p_K} : p_1, \dots, p_K \in [P], q \in [Q]\}$ of Θ as follows

$$\mathcal{P}_{z,q,p_1,\dots,p_K} := \{(\theta_j)_{j=1}^K : (\theta_1^\top x, \dots, \theta_K^\top x) \in \mathcal{Q}_{z,q}, \theta_j^\top x \in \mathcal{T}_{p_j} \forall j \in [K]\}$$

where

$$\mathcal{T}_p := \begin{cases} [\mu_p, \mu_{p+1}] & \text{if } p \in [P-1] \\ [\mu_P, \mu_{P+1}] & \text{if } p = P \end{cases},$$

$$\mu_p := -RR_x + (p-1)\kappa \quad \forall p \in [P+1].$$

For the notational simplicity, let $u := (z, q, p_1, \dots, p_K) \in \mathcal{Z} \times [Q] \times [P]^K$. Now, we define our approximation as $h(\cdot; z) = h_u(\cdot)$ on \mathcal{P}_u if $\mathcal{P}_u \neq \emptyset$. Here, $h_u(\cdot)$ is defined as follows: for some fixed $\nu_u := (\nu_{u,1}, \dots, \nu_{u,K}) \in \mathcal{P}_u$,

$$h_u(\theta) := \ell(\nu_{u,1}^\top x, \dots, \nu_{u,K}^\top x; y) + \sum_{j=1}^K \nabla_j \ell(\nu_{u,1}^\top x, \dots, \nu_{u,K}^\top x; y) (\theta_j^\top x - \nu_{u,j}^\top x) + \sum_{j=1}^K \frac{\lambda}{2} \|\theta_j\|^2.$$

where ∇_j denotes the partial derivative with respect to the j -th argument. Namely, h_u is a first order approximation of ℓ at θ with ℓ_2 -regularization; hence, h_u is λ -strongly convex and λ -smooth. Then given $z \in \mathcal{Z}$ and $\theta \in \Theta$, for q and p_1, \dots, p_K satisfying $\theta \in \mathcal{P}_u$ for $u = (z, q, p_1, \dots, p_K)$, one can observe that

$$\begin{aligned} \|\nabla f(\theta; z) - \nabla h_u(\theta)\| &\leq K^{1/2} \max_{j \in [K]} \|\nabla_j \ell(\theta_1^\top x, \dots, \theta_K^\top x; y)x - \nabla_j \ell(\nu_{u,1}^\top x, \dots, \nu_{u,K}^\top x; y)x\| \\ &\leq \beta\kappa KR_x = \frac{1}{12\eta LTn}. \end{aligned}$$

Now, we apply Theorem 2.2 with

$$h \leftarrow h, P \leftarrow P^K Q, \xi \leftarrow \frac{1}{12\eta LTn}, T \leftarrow T, \alpha \leftarrow \lambda, \beta \leftarrow \lambda, L \leftarrow L, B \leftarrow B.$$

Then in the bound of Theorem 2.2, we have

$$2L \left(\gamma^T R + \frac{1 - \gamma^T}{1 - \gamma} \eta \xi \right) \leq \frac{1}{n}$$

since $\gamma^T R \leq 1/(3Ln)$ and $\frac{1 - \gamma^T}{1 - \gamma} = \sum_{j=0}^{T-1} \gamma^j \leq T$. This completes the proof of Theorem 3.1.

C.2 Equivalence between soft K -means algorithm and expectation-maximization for (3.4)

We first observe that applying the affine transformation $x \mapsto -nKx + \log(\frac{(\xi/\pi)^{d/2}}{K})$ to (3.4) results in the following objective:

$$\sum_{i=1}^n \log \left(\frac{1}{K} (\xi/\pi)^{d/2} \sum_{j=1}^K \exp(-\zeta \|\theta_j - z_i\|^2) \right) \quad (\text{C.1})$$

i.e. the expectation-maximization algorithm for (C.1) aims to find a local minimum of (3.4). Since (C.1) is the log-likelihood for the mixture of Gaussians $N(\theta_1, \frac{1}{2\zeta}I), \dots, N(\theta_K, \frac{1}{2\zeta}I)$ under the same cluster density $1/K$ with observations z_1, \dots, z_n , the expectation-maximization algorithm for (C.1) is identical to the alternative procedure (3.3) [Pri12].

C.3 Proof of Theorem 3.2

First, observe that for $\theta \in (\mathcal{B}_R^d(0))^K, z \in \mathcal{B}_R^d(0)$, we have

$$\begin{aligned}\frac{\partial}{\partial \theta_k} f(\theta; z) &= \frac{2(\theta_k - z_i) \exp(-\zeta \|\theta_k - z_i\|^2)}{\sum_{j=1}^K \exp(-\zeta \|\theta_j - z_i\|^2)}, \\ \frac{\partial^2}{\partial \theta_k \partial \theta_{k'}} f(\theta; z) &= \frac{4\zeta(\theta_k - z_i)(\theta_{k'} - z_i)^\top \exp(-\zeta \|\theta_k - z_i\|^2 - \zeta \|\theta_{k'} - z_i\|^2)}{(\sum_{j=1}^K \exp(-\zeta \|\theta_j - z_i\|^2))^2} \\ &\quad + \mathbf{1}_{k=k'} \times \frac{(2I - 4\zeta(\theta_k - z_i)(\theta_k - z_i)^\top) \exp(-\zeta \|\theta_k - z_i\|^2)}{\sum_{j=1}^K \exp(-\zeta \|\theta_j - z_i\|^2)}\end{aligned}$$

where $\mathbf{1}_{k=k'}$ is one if $k = k'$ and zero otherwise. Using this, we bound f and its first and second derivatives as follows:

$$f(\theta; z) \in [-B + \log K, \log K],$$

$$L := \frac{4R}{\sqrt{K}} \exp(\zeta B) = 4R \sqrt{\frac{K}{K^2 \exp(-\zeta B)^2}} \geq 4R \sqrt{\frac{\sum_{j=1}^K \exp(-\zeta \|\theta_j - z_i\|^2)^2}{(\sum_{j=1}^K \exp(-\zeta \|\theta_j - z_i\|^2))^2}}$$

$$\geq \sqrt{\sum_{k=1}^K \left\| \frac{\partial}{\partial \theta_k} f(\theta; z) \right\|^2} = \left\| \frac{\partial}{\partial \theta} f(\theta; z) \right\|,$$

$$\alpha := \frac{2}{K} \exp(-\zeta B) \leq \frac{1}{\|\theta_j - z\|} \left\| \frac{\partial}{\partial \theta_j} f(\theta; z) \right\| \quad \forall \theta \text{ s.t. } \theta_j \neq z,$$

$$\beta := \frac{2}{K} \exp(\zeta B) \geq \frac{1}{\|\theta_j - z\|} \left\| \frac{\partial}{\partial \theta_j} f(\theta; z) \right\| \quad \forall \theta \text{ s.t. } \theta_j \neq z,$$

$$\beta' := 4\zeta B \exp(\zeta B) + 4\zeta B + 2 \geq \|\nabla^2 f(\theta; z)\|_2.$$

To utilize Theorem 2.2, we approximate each $f(\cdot; z)$ using a piecewise strongly convex and smooth function $h(\cdot; z)$ with P^K pieces. To this end, we first construct a partition $\{\mathcal{P}_{z, p_1, \dots, p_K} : p_1, \dots, p_K \in [P]\}$ of Θ for each $z \in \mathcal{Z}$ as follows:

$$\mathcal{P}_{z, p_1, \dots, p_K} := \{(\theta_j)_{j=1}^K \in \Theta : \|\theta_j - z\| \in \mathcal{T}_{p_j} \quad \forall j \in [K]\}$$

where

$$\kappa := \frac{1}{12(\beta + \beta')\eta\sqrt{KLTn}},$$

$$P := \left\lceil \frac{2R}{\kappa} \right\rceil,$$

$$\mu_p := (p-1)\kappa \quad \forall p \in [P+1],$$

$$\mathcal{T}_p := \begin{cases} [\mu_p, \mu_{p+1}) & \text{if } p \in [P-1] \\ [\mu_P, \mu_{P+1}] & \text{if } p = P \end{cases}.$$

We define $h(\cdot; z) := h_{z, p_1, \dots, p_K}(\cdot)$ on $\mathcal{P}_{z, p_1, \dots, p_K}$ if $\mathcal{P}_{z, p_1, \dots, p_K} \neq \emptyset$. Here, note that $\mathcal{P}_{z, p_1, \dots, p_K} \setminus \{z\} \neq \emptyset$ if $\mathcal{P}_{z, p_1, \dots, p_K} \neq \emptyset$ from the definition of $\mathcal{P}_{z, p_1, \dots, p_K}$. Now, we define $h_{z, p_1, \dots, p_K}(\cdot)$ as follows: for some fixed $\nu_{z, p_1, \dots, p_K} = ((\nu_{z, p_1, \dots, p_K})_j)_{j=1}^K \in \mathcal{P}_{z, p_1, \dots, p_K} \setminus \{z\}$,

$$h_{z, p_1, \dots, p_K}(\theta) := \sum_{j=1}^K \frac{1}{2} \left\| \frac{1}{\|(\nu_{z, p_1, \dots, p_K})_j - z\|} \frac{\partial}{\partial \theta'_j} f(\theta'; z) \Big|_{\theta' = \nu_{z, p_1, \dots, p_K}} \right\| \cdot \|\theta_j - z\|^2.$$

Our construction of h_{z, p_1, \dots, p_K} has some nice properties. For example, for

$$\begin{aligned}a_{z, p_1, \dots, p_K} &:= \min_j \left\| \frac{1}{\|(\nu_{z, p_1, \dots, p_K})_j - z\|} \frac{\partial}{\partial \theta'_j} f(\theta'; z) \Big|_{\theta' = \nu_{z, p_1, \dots, p_K}} \right\|, \\ b_{z, p_1, \dots, p_K} &:= \max_j \left\| \frac{1}{\|(\nu_{z, p_1, \dots, p_K})_j - z\|} \frac{\partial}{\partial \theta'_j} f(\theta'; z) \Big|_{\theta' = \nu_{z, p_1, \dots, p_K}} \right\|,\end{aligned}$$

h_{z,p_1,\dots,p_K} is a_{z,p_1,\dots,p_K} -strongly convex and b_{z,p_1,\dots,p_K} -smooth and $\alpha \leq a_{z,p_1,\dots,p_K} \leq b_{z,p_1,\dots,p_K} \leq \beta$. Furthermore, for any $\nu' = (\nu'_j)_{j=1}^K$ satisfying

$$\|\nu'_j - z\| = \|(\nu_{z,p_1,\dots,p_K})_j - z\|, \quad (\text{C.2})$$

we have

$$\nabla f(\nu', z) = \nabla h_{z,p_1,\dots,p_K}(\nu') \quad (\text{C.3})$$

from the symmetry of f and h_{z,p_1,\dots,p_K} .

Now, we bound $\|\nabla f(\theta; z) - \nabla h(\theta; z)\|$ to utilize Theorem 2.2. For $\theta = (\theta_j)_{j=1}^K \in (\mathcal{B}_R^d(0))^K$, let p_1, \dots, p_K be indices in $[P]$ satisfying $\theta \in \mathcal{P}_{z,p_1,\dots,p_K}$. Let ν' be a point on the line connecting z and θ such that ν' satisfies (C.2), i.e. $\|\nu'_j - \theta_j\| \leq \kappa$ for all $j \in [K]$. Then, we have

$$\begin{aligned} \|\nabla f(\theta; z) - \nabla h(\theta; z)\| &= \|\nabla f(\theta; z) - \nabla f(\nu'; z) + \nabla f(\nu'; z) - \nabla h(\theta; z)\| \\ &\leq \|\nabla f(\theta; z) - \nabla f(\nu'; z)\| + \|\nabla h(\theta; z) - \nabla h_{z,p_1,\dots,p_K}(\nu')\| \\ &\leq \beta' \kappa \sqrt{K} + \sqrt{K} \max_{j \in [K]} \left\| \left\| \frac{\partial}{\partial \theta_j} h_{z,p_1,\dots,p_K}(\theta) - \frac{\partial}{\partial \theta_j} h_{z,p_1,\dots,p_K}(\nu') \right\| \right\| \\ &\leq (\beta + \beta') \kappa \sqrt{K} \\ &\leq \frac{1}{12\eta L T n} \end{aligned}$$

The first inequality holds since $\nabla f(\nu'; z) = \nabla h_{z,p_1,\dots,p_K}(\nu')$ by (C.3) and other inequalities holds from the definitions of β and β' . Now, we apply Theorem 2.2 with

$$h \leftarrow h, P \leftarrow P^K, \xi \leftarrow \frac{1}{12\eta L T n}, T \leftarrow T, \alpha \leftarrow \alpha, \beta \leftarrow \beta, L \leftarrow L, B \leftarrow B.$$

Then in the bound of Theorem 2.2, we have

$$2L \left(\gamma^T R + \frac{1 - \gamma^T}{1 - \gamma} \eta \xi \right) \leq \frac{1}{n}$$

since $\gamma^T R \leq 1/(3Ln)$ and $\frac{1 - \gamma^T}{1 - \gamma} = \sum_{j=0}^{T-1} \gamma^j \leq T$. This completes the proof of Theorem 3.2.

C.4 Proof of Theorem 3.3

For $i \in [n]$ and $j \in [K]$, we first define $g_{i,j}$ as

$$g_{i,j}(\theta) := (\theta'_k)_{k=1}^K \quad \text{where} \quad \theta'_k = \begin{cases} \theta_k - 2\eta(\theta_k - z_i) & \text{if } k = j \\ \theta_k & \text{if } k \neq j \end{cases}.$$

For $\varepsilon := \frac{1}{8Rn}$, We define

$$\Phi_\varepsilon := \{(\theta'_j)_{j=1}^K : \theta'_j = (g_{i_j,t_j,j} \circ \dots \circ g_{i_1,t_1,j}(\theta^{(0)}))_j, t_j \in [T] \cup \{0\}, i_{j,1}, \dots, i_{j,t_j} \in [n], \forall j \in [K]\}.$$

Given $\theta^{(0)}$, since $g_{i,j}$ is γ -contractive by Lemma A.4, we have for any $t \geq 0, i_1, \dots, i_t \in [n]$, and $j \in [K]$,

$$(g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)}))_j \in \bigcup_{\phi \in \Phi_\varepsilon} \mathcal{B}_{K-1/2\varepsilon}^d(\phi_j),$$

i.e.

$$g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)}) \in \bigcup_{\phi \in \Phi_\varepsilon} \mathcal{B}_\varepsilon^d(\phi), \quad (\text{C.4})$$

Namely, Φ_ε contains all possible SGD parameters starting from $\theta^{(0)}$ where each element in Φ_ε is an algorithm depending on at most KT samples. Since each $f(\cdot; z)$ is $(4R)$ -Lipschitz on Θ , we apply Theorem B.1 with

$$\begin{aligned} \Phi_n &\leftarrow \{g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)}) : t \geq 0, i_1, \dots, i_t \in [n]\} \\ \Phi_{T,\varepsilon} &\leftarrow \Phi_\varepsilon, \varepsilon \leftarrow \varepsilon, T \leftarrow KT, \delta \leftarrow \delta, L \leftarrow 4R, B \leftarrow B. \end{aligned}$$

where

$$|\Phi_\varepsilon| \leq \sum_{t_1, \dots, t_K \in [T] \cup \{0\}} n^{\sum_{k=1}^K t_k} \leq \sum_{t_1, \dots, t_K \in [T] \cup \{0\}} n^{KT} \leq (T+1)^K \cdot n^{KT} \leq (2n)^{KT}$$

and the assumption (B.1) in Theorem B.1 is satisfied by (C.4). This leads us to the bound in Theorem 3.3.

D Expected generalization gap: Strongly convex and smooth case

In this section, we bound the expected generalization gap using our localized cover for a contractive iterative stochastic optimizer g (see Definition 3), i.e. each $g(\cdot; z)$ is contractive for all $z \in \mathcal{Z}$. Here, we use \mathcal{S} for denoting the set of samples $\{z_1, \dots, z_n\}$ and g_i for denoting $g(\cdot; z_i)$.

Theorem D.1. *Suppose that Assumptions 1 & 2 hold and there exists $\gamma \in (0, 1)$ such that $g(\cdot; z)$ is γ -contractive on $\Theta \subseteq \mathcal{B}_R^d(0)$ for all $z \in \mathcal{Z}$. Let $T := \max \left\{ \left\lceil \frac{\log(2LRn)}{\log(1/\gamma)} \right\rceil, 0 \right\}$. Then for any $\theta^{(0)} \in \Theta$, $t \geq 0$, and $i_1, \dots, i_t \in [n]$, the t -th iterate $\theta^{(t)} := g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)})$ satisfies*

$$|\mathbb{E}_{\mathcal{S}}[\hat{F}(\theta^{(t)}) - F(\theta^{(t)})]| \leq \frac{BT + 1}{n}.$$

Compared to the uniform stability-based bound for γ -contractive $\frac{1}{n} \sum_{i=1}^n g_i$ under L' -Lipschitz $f(\cdot; z)$ and $\|g(\theta; z)\| \leq K$ [KWS22]

$$|\mathbb{E}_{\theta^{(0)}, i_1, \dots, i_t, \mathcal{S}}[\hat{F}(\theta^{(t)}) - F(\theta^{(t)})]| \leq \frac{2KL'}{(1-\gamma)n}, \quad (\text{D.1})$$

our bound in Theorem D.1 is linear in B , has an additional $\log n$ factor, and requires contractive $g(\cdot; z)$ but does not depend on L' and K . Here, note that our bound uniformly holds for all parameters generated by SGD while the bound (D.1) requires the expectation over $\theta^{(0)}$ and the indices i_1, \dots, i_t used for the t updates using the iterative stochastic optimizer.

Although the bounds in Theorem D.1 and (D.1) provide some information about the expected generalization gap, these bounds cannot be used for understanding the absolute deviation $|\hat{F}(\theta^{(t)}) - F(\theta^{(t)})|$, which can be of practical interest. In the following theorem, we provide a bound on the expectation of the absolute generalization gap where $\Psi_{\geq T}(\theta)$ and $\Psi_T(\theta)$ are defined for general g_i , analogous to Definition 1:

$$\begin{aligned} \Psi_T(\theta^{(0)}) &:= \{g_{i_T} \circ \dots \circ g_{i_1}(\theta^{(0)}) : i_1, \dots, i_T \in [n]\}, \\ \Psi_{\geq T}(\theta^{(0)}) &:= \bigcup_{t \geq T} \Psi_t(\theta^{(0)}). \end{aligned} \quad (\text{D.2})$$

Theorem D.2. *Assume the setup in Theorem D.1. Then there exists some absolute constant $C > 0$ such that*

$$\mathbb{E}_{\mathcal{S}} \left[\sup_{\theta^{(t)} \in \bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq T}(\theta^{(0)})} |\hat{F}(\theta^{(t)}) - F(\theta^{(t)})| \right] \leq \frac{BT + 1}{n} + CB \sqrt{\frac{T \log n}{n}}.$$

To our knowledge, the bound in Theorem D.2 is the first bound on $\mathbb{E}_{\mathcal{S}}[|\hat{F}(\theta^{(t)}) - F(\theta^{(t)})|]$ for the contractive case, including strongly convex and smooth functions. Compared to Theorem D.1, the bound in Theorem D.2 loses $O(\sqrt{T}/(n \log n))$ factor for taking the absolute value inside the expectation. Nevertheless, this bound can be improved by removing the supremum inside the expectation.

Corollary D.1. *Assume the setup in Theorem 2.1. Then there exists some absolute constant $C > 0$ such that for any $\theta^{(0)} \in \Theta$, $t \geq 0$, and $i_1, \dots, i_t \in [n]$, we have*

$$\mathbb{E}_{\mathcal{S}} \left[|\hat{F}(\theta^{(t)}) - F(\theta^{(t)})| \right] \leq \frac{BT + 1}{n} + CB \sqrt{\frac{1}{n}}.$$

Compared to Theorem D.2, the bound in Corollary D.1 does not have $\sqrt{T \log n}$ factor but the LHS in the bound is weaker.

D.1 Proof of Theorem D.1

For any $\theta^{(t)} \in \bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq 0}(\theta^{(0)})$, let $\theta^{(t)} = g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)})$. Let $\phi = g_{i_t} \circ \dots \circ g_{i_{t-T+1}}(0)$ and $\mathcal{I} = \{i_{t-T+1}, \dots, i_t\} \subset [n]$ if $t > T$ and $\phi = g_{i_t} \circ \dots \circ g_{i_1}(0)$ and $\mathcal{I} = \{i_1, \dots, i_t\} \subset [n]$ otherwise. Then as in the proof of Lemma 2.1, we have

$$\theta^{(t)} \in \mathcal{B}_{\varepsilon}^d(\phi) \quad (\text{D.3})$$

for $\varepsilon = 1/(2Ln)$, regardless of the choice of \mathcal{S} . Using this, we have

$$\begin{aligned}
|\mathbb{E}_{\mathcal{S}}[\hat{F}(\theta^{(t)}) - F(\theta^{(t)})]| &\leq |\mathbb{E}_{\mathcal{S}}[\hat{F}(\theta^{(t)}) - \hat{F}(\phi)]| + |\mathbb{E}_{\mathcal{S}}[F(\theta^{(t)}) - F(\phi)]| + |\mathbb{E}_{\mathcal{S}}[\hat{F}(\phi) - F(\phi)]| \\
&\leq 2L\varepsilon + |\mathbb{E}_{\mathcal{S}}[\hat{F}(\phi) - F(\phi)]| \\
&\leq 2L\varepsilon + \left| \mathbb{E}_{\mathcal{S}} \left[\frac{1}{n} \sum_{i \in [n] \setminus \mathcal{I}} (\hat{f}(\phi; z_i) - F(\phi)) \right] \right| + \left| \mathbb{E}_{z_i: i \in \mathcal{I}} \left[\frac{1}{n} \sum_{i \in \mathcal{I}} (\hat{f}(\phi; z_i) - F(\phi)) \right] \right| \\
&= 2L\varepsilon + \left| \mathbb{E}_{z_i: i \in \mathcal{I}} \left[\frac{1}{n} \sum_{i \in \mathcal{I}} (\hat{f}(\phi; z_i) - F(\phi)) \right] \right| \\
&\leq 2L\varepsilon + \frac{BT}{n} = \frac{BT+1}{n}.
\end{aligned}$$

The first inequality is from the triangle inequality and the second inequality is from (D.3). The third inequality is again from the triangle inequality and the first equality is from the fact that ϕ is independent of $\{z_i : i \in \mathcal{I}\}$. The last inequality is from Assumption 2. Since the above bound holds for any $\theta^{(0)} \in \Theta$ and \mathcal{I}_t , this completes the proof of Theorem D.1.

D.2 Proof of Theorem D.2

In this proof, we assume $T < n$ since the statement trivially follows otherwise. As in the statement of Lemma 2.1, one can observe that

$$\bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq T}(\theta^{(0)}) \subseteq \bigcup_{\phi \in \Psi_T(0)} \mathcal{B}_{\varepsilon}^d(\phi) \tag{D.4}$$

for $\varepsilon = 1/(2Ln)$. For each $\phi \in \Psi_T(0)$, let $\mathcal{I}_{\phi} := \{i_1, \dots, i_T\}$ such that $\phi = g_{i_T} \circ \dots \circ g_{i_1}(0)$. Using this we can derive the following bound: for $\Pi_{\Psi_T(0)}(\theta^{(t)}) := \arg \min_{\phi \in \Psi_T(0)} \|\theta^{(t)} - \phi\|$,

$$\begin{aligned}
\mathbb{E}_{\mathcal{S}} \left[\sup_{\theta^{(t)} \in \bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq T}(\theta^{(0)})} |\hat{F}(\theta^{(t)}) - F(\theta^{(t)})| \right] &\leq \mathbb{E}_{\mathcal{S}} \left[\sup_{\phi \in \Psi_T(0)} |\hat{F}(\phi) - F(\phi)| \right] \\
&+ \mathbb{E}_{\mathcal{S}} \left[\sup_{\theta^{(t)} \in \bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq T}(\theta^{(0)})} |\hat{F}(\theta^{(t)}) - \hat{F}(\Pi_{\Psi_T(0)}(\theta^{(t)}))| + |F(\theta^{(t)}) - F(\Pi_{\Psi_T(0)}(\theta^{(t)}))| \right] \\
&\leq 2L\varepsilon + \mathbb{E}_{\mathcal{S}} \left[\sup_{\phi \in \Psi_T(0)} |\hat{F}(\phi) - F(\phi)| \right] \\
&\leq \frac{1}{n} + \mathbb{E}_{\mathcal{S}} \left[\sup_{\phi \in \Psi_T(0)} \left| \frac{1}{n} \sum_{i \in \mathcal{I}_{\phi}} (f(\phi; z_i) - F(\phi)) \right| \right] + \mathbb{E}_{\mathcal{S}} \left[\sup_{\phi \in \Psi_T(0)} \left| \frac{1}{n} \sum_{i \in [n] \setminus \mathcal{I}_{\phi}} (f(\phi; z_i) - F(\phi)) \right| \right] \\
&\leq \frac{BT+1}{n} + \mathbb{E}_{\mathcal{S}} \left[\sup_{\phi \in \Psi_T(0)} \left| \frac{|[n] \setminus \mathcal{I}_{\phi}|}{n} \cdot \frac{1}{|[n] \setminus \mathcal{I}_{\phi}|} \sum_{i \in [n] \setminus \mathcal{I}_{\phi}} (f(\phi; z_i) - F(\phi)) \right| \right] \\
&\leq \frac{BT+1}{n} + CB \sqrt{\frac{T \log n}{n}}.
\end{aligned}$$

The first inequality is from the triangle inequality and the second inequality is from (D.4). The third inequality is again from the triangle inequality and the definition of ε , and the fourth inequality is from Assumption 2. Since each $\frac{1}{|[n] \setminus \mathcal{I}_{\phi}|} \sum_{i \in [n] \setminus \mathcal{I}_{\phi}} (f(\phi; z_i) - F(\phi))$ is sub-Gaussian with the sub-Gaussian norm bounded by $cB/\sqrt{|[n] \setminus \mathcal{I}_{\phi}|}$ for some absolute constant c (see Proposition 2.5.2 in [Ver18]), the last inequality follows from a standard upper bound for empirical processes (see Exercise 2.5.10 in [Ver18]). This completes the proof of Theorem D.2.

D.3 Proof of Corollary D.1

As in the proof of Theorem D.2, we assume that $T < n$ without loss of generality. The proof here is almost identical to that of Theorems D.1 & D.2.

For any $\theta^{(t)} \in \bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq 0}(\theta^{(0)})$, let $\theta^{(t)} = g_{i_t} \circ \dots \circ g_{i_1}(\theta^{(0)})$. Let $\phi = g_{i_t} \circ \dots \circ g_{i_{t-T+1}}(0)$ and $\mathcal{I} = \{i_{t-T+1}, \dots, i_t\} \subset [n]$ if $t > T$ and $\phi = g_{i_t} \circ \dots \circ g_{i_1}(0)$ and $\mathcal{I} = \{i_1, \dots, i_t\} \subset [n]$ otherwise. Then for $\varepsilon = 1/(2Ln)$, we have

$$\theta^{(t)} \in \mathcal{B}_\varepsilon(\phi). \quad (\text{D.5})$$

Using this, we derive the following inequality:

$$\begin{aligned} \mathbb{E}_{\mathcal{S}}[|\hat{F}(\theta^{(t)}) - F(\theta^{(t)})|] &\leq \mathbb{E}_{\mathcal{S}}[|(\hat{F}(\theta^{(t)}) - \hat{F}(\phi))|] + \mathbb{E}_{\mathcal{S}}[|F(\theta^{(t)}) - F(\phi)|] + \mathbb{E}_{\mathcal{S}}[|\hat{F}(\phi) - F(\phi)|] \\ &\leq 2L\varepsilon + \mathbb{E}_{\mathcal{S}}[|\hat{F}(\phi) - F(\phi)|] \\ &\leq 2L\varepsilon + \mathbb{E}_{\mathcal{S}}\left[\left|\frac{1}{n} \sum_{i \in [n] \setminus \mathcal{I}} (\hat{f}(\phi; z_i) - F(\phi))\right|\right] + \mathbb{E}_{z_i: i \in \mathcal{I}}\left[\left|\frac{1}{n} \sum_{i \in \mathcal{I}} (\hat{f}(\phi; z_i) - F(\phi))\right|\right] \\ &\leq 2L\varepsilon + \mathbb{E}_{\mathcal{S}}\left[\left|\frac{1}{n} \sum_{i \in [n] \setminus \mathcal{I}} (\hat{f}(\phi; z_i) - F(\phi))\right|\right] + \frac{BT}{n} \\ &\leq \frac{BT+1}{n} + \frac{|[n] \setminus \mathcal{I}|}{n} \mathbb{E}_{z_i: i \in [n] \setminus \mathcal{I}}\left[\left|\frac{1}{|[n] \setminus \mathcal{I}|} \sum_{i \in [n] \setminus \mathcal{I}} (\hat{f}(\phi; z_i) - F(\phi))\right|\right] \\ &\leq \frac{BT+1}{n} + \frac{|[n] \setminus \mathcal{I}|}{n} \left(CB \sqrt{\frac{1}{|[n] \setminus \mathcal{I}|}} \right) \\ &\leq \frac{BT+1}{n} + CB \sqrt{\frac{1}{n}}. \end{aligned}$$

The first and second inequality follows from the triangle inequality and (D.5), respectively. The third inequality is again from the triangle inequality while Assumption 2 gives us the fourth inequality. The sixth inequality is from Proposition 2.5.2 and Exercise 2.5.10 in [Ver18] where the last inequality naturally follows. Since the above bound holds for any $\theta^{(0)} \in \Theta$ and \mathcal{I}_t , this completes the proof of Corollary D.1

E Comparison with existing bounds

Table 1: Summary of generalization bounds for constant step-size SGD. In the column ‘‘Assumptions’’, Lip. assumes L' -Lipschitz $f(\cdot; z)$, Weak Lip. assumes Assumption 1, Bdd. assumes $f(\Theta; \mathcal{Z}) \subset [0, B']$, and Bdd. Dev. assumes Assumption 2. If LHS of a bound does not contain expectation, that bound is a high-probability bound. For simplicity, we hide values other than $d, t, T, L, L', B, B', P, \xi, n$ in \lesssim where $T = \log(LRn)$ and $\Theta \subset \mathcal{B}_R^d(0)$.

Reference	Objective	Assumptions	Stable as $t \rightarrow \infty$	Bound: $\Delta^{(t)} = \hat{F}(\theta^{(t)}) - F(\theta^{(t)})$
[HRS15]	Strongly convex & smooth	Lip.	✓	$ \mathbb{E}[\Delta^{(t)}] \lesssim (L')^2/n^*$
Thm D.1		Weak Lip. & Bdd. Dev.		$ \mathbb{E}[\Delta^{(t)}] \lesssim BT/n^{\parallel}$
Thm D.2				$\mathbb{E}[\sup \Delta^{(t)}] \lesssim B\sqrt{T \log n/n}^{\ddagger\parallel}$
Cor D.1				$\mathbb{E}[\Delta^{(t)}] \lesssim B\sqrt{1/n}^{\parallel}$
Thm 2.1				$\sup \Delta^{(t)} \lesssim B\sqrt{T \log n/n}^{\ddagger}$
Cor 2.1–2.2				$ \Delta^{(t)} \lesssim B\sqrt{1/n}$
[HRS15]	Convex & smooth		Lip.	✗
[FV19]		Lip. & Bdd.	$ \Delta^{(t)} \lesssim t \log^2 n/n + \sqrt{1/n}^{\P}$	
[BFGT20]	Convex & non-smooth	Lip.	$ \mathbb{E}[\Delta^{(t)}] \lesssim (L')^2(\sqrt{t} + t/n)^*$	
		Lip. & Bdd.	$ \Delta^{(t)} \lesssim (L' \log n)^2(\sqrt{t} + t/n) + B'\sqrt{1/n}$	
[HRS15]	Non-convex & smooth	Lip.	$ \mathbb{E}[\Delta^{(t)}] \lesssim (L')^{\frac{2}{\beta c+1}} t^{\frac{\beta c}{\beta c+1}}/n^{*\S}$	
Thm 2.2	Approx. piecewise strongly convex & smooth	Weak Lip. & Bdd. Dev.	✓	
				$\sup \Delta^{(t)} \lesssim B\sqrt{T \log(nP)/n} + L\xi^{\ddagger}$

Table 2: Summary of generalization bounds for a (piecewise) contractive optimizer with update functions g_1, \dots, g_n . The column ‘‘Non-cvx SGD’’ evaluates if a bound can be used for SGD on non-convex objectives without diverging as $t \rightarrow \infty$.

Reference	Optimizer	Assumptions	Non-cvx SGD	Bound: $\Delta^{(t)} = \hat{F}(\theta^{(t)}) - F(\theta^{(t)})$
[KWS22]	Contractive $\frac{1}{n} \sum_{i=1}^n g_i$	Lip. & $\ g_i(\Theta)\ \leq K$	✗	$ \mathbb{E}[\Delta^{(t)}] \lesssim KL'/n^*$
Thm D.1	Contractive g_i	Weak Lip. & Bdd. Dev.	✗	$ \mathbb{E}[\Delta^{(t)}] \lesssim BT/n^{\parallel}$
Thm D.2				$\mathbb{E}[\sup \Delta^{(t)}] \lesssim B\sqrt{T \log n/n}^{\ddagger\parallel}$
Cor D.1				$\mathbb{E}[\Delta^{(t)}] \lesssim B\sqrt{1/n}^{\parallel}$
Thm 4.1				$\sup \Delta^{(t)} \lesssim B\sqrt{T \log n/n}^{\ddagger}$
Cor 2.1–2.2 ^b				$ \Delta^{(t)} \lesssim B\sqrt{1/n}$
Thm 4.1	Approx. piecewise contractive g_i	Weak Lip. & Bdd. Dev.	✓	$\sup \Delta^{(t)} \lesssim B\sqrt{T \log(nP)/n} + L\xi^{\ddagger}$

*The expectation is taken over $\{z_1, \dots, z_n\}, \{i_1, \dots, i_t\}, \theta^{(0)}$.

[∥]The expectation is taken over $\{z_1, \dots, z_n\}$.

[‡]The supremum is taken over $\theta^{(t)} \in \bigcup_{\theta^{(0)} \in \Theta} \Psi_{\geq T}(\theta^{(0)})$ (see Definition 1 and (D.2)).

[§] β denotes the smoothness parameter and the adaptive learning rate must satisfy $\eta_t \leq c/t$.

[¶]This bound is under 1-Lipschitz continuity of $f(\cdot; z)$ and $f(\Theta, \mathcal{Z}) \subset [0, 1]$.

^bAlthough Corollaries 2.1–2.2 are for SGD, the same result holds for contractive g_i with the same proof.

F Lower bound on uniform stability-based bound for piecewise strongly convex and smooth objectives

In this section, we show that the uniform stability-based bound [HRS15] is $\Omega(1)$ for piecewise strongly convex and smooth functions after sufficiently many SGD iterations in general. To this end, we first introduce the formal definition of the uniform stability and a standard tool for showing that the stability implies generalization, which are from Definition 2.1 and Theorem 2.2 in [HRS15].

Definition 7 (Uniform stability). *A randomized algorithm ρ is “ ε -uniformly stable” if for all data sets $\mathcal{S}, \mathcal{S}' \in \mathcal{Z}^n$ such that \mathcal{S} and \mathcal{S}' differ in at most one example, we have*

$$\sup_{z \in \mathcal{Z}} \mathbb{E}_\rho [f(\rho(\mathcal{S}); z) - f(\rho(\mathcal{S}'); z)] \leq \varepsilon.$$

Theorem F.1. *If ρ is ε -uniformly stable, then*

$$|\mathbb{E}_{\mathcal{S}, \rho} [\hat{F}(\rho(\mathcal{S})) - F(\rho(\mathcal{S}))]| \leq \varepsilon.$$

For the remaining section, we provide an example that

$$\sup_{z \in \mathcal{Z}} \mathbb{E}_\rho [f(\rho(\mathcal{S}); z) - f(\rho(\mathcal{S}'); z)] = \Omega(1),$$

regardless of n where ρ denotes sufficiently many SGD updates. Namely, the uniform stability-based bound based on Theorem F.1 is $\Omega(1)$.

Let $\Theta = [0, 4]$, $\mathcal{Z} = \{0, 1\}$, $\theta^{(0)} \sim \text{Unif}([0, 4])$, $\mathbb{P}(z = 0) = \mathbb{P}(z = 1) = \frac{1}{2}$, $f(\cdot; 0) = \min\{(x - 1)^2, \frac{1}{2} + \frac{1}{2}(x - 3)^2\}$, $f(\cdot; 1) = (x - 1)^2$, and the auxiliary gradient $\nabla f(2; 0) = 2$ at the non-differentiable point $\theta = 2$, i.e., $f(\cdot; z)$ is piecewise 1-strongly convex and 2-smooth with the partition $\mathcal{P} = \{[0, 2], (2, 4]\}$. For simplicity, choose $\eta = 1/3$ which can be generalized to arbitrary $\eta \in (0, 1)$. Under this setup, using Theorem 2.2, one can easily derive a generalization bound that does not increase with the number of SGD iterations and converges to zero as n grows.

However, under the same setup and sufficiently many SGD iterations, the uniform stability-based bound is lower bounded by a constant regardless of n . To see this, let $\mathcal{S} = (0, \dots, 0)$ and $\mathcal{S}' = (1, 0, \dots, 0)$. Then, one can observe that $f([0, 2]; 0) \subset [0, 2]$ and $f((2, 4]; 0) \subset (2, 4]$, i.e. SGD iterates for \mathcal{S} converge to either 1 or 3 depending on whether $\theta^{(0)} \in [0, 2]$ or $\theta^{(0)} \in (2, 4]$. Since we assumed $\theta^{(0)} \sim \text{Unif}([0, 4])$, we have

$$\lim_{t \rightarrow \infty} \mathbb{E}_{\theta^{(t)}} [f(\theta^{(t)}; 1)] = \frac{1}{2}f(1; 1) + \frac{1}{2}f(3; 1) = 2 \quad (\text{F.1})$$

where $\theta^{(t)}$ denotes a random parameter generated by t SGD updates for \mathcal{S} , from $\theta^{(0)} \sim \text{Unif}([0, 4])$. Furthermore, we have $f(\Theta; 1) \subset [0, 2]$ and $f([0, 2]; 0) \subset [0, 2]$. This implies that if a single SGD update for \mathcal{S}' use the first sample in \mathcal{S}' (i.e. $z = 1$), which occurs with high probability under sufficiently many SGD iterations, then the SGD iterates will converge to 1 almost surely. In other words, we have

$$\lim_{t \rightarrow \infty} \mathbb{E}_{\phi^{(t)}} [f(\phi^{(t)}; 1)] = f(1; 1) = 0 \quad (\text{F.2})$$

where $\phi^{(t)}$ denotes a random parameter generated by t SGD updates for \mathcal{S}' , from $\phi^{(0)} \sim \text{Unif}([0, 4])$. Combining (F.1) and (F.2) implies a constant lower bound on the uniform stability-based bound, regardless of n , under sufficiently many SGD iterations. We note that the same conclusion can also be derived for any $\eta \in (0, 1)$ as long as the number of SGD iterations is large enough.