

Reusing Samples in Variance Reduction

Yujia Jin

YUJIAJIN@STANFORD.EDU

Ishani Karmarkar

ISHANIK@STANFORD.EDU

Aaron Sidford

SIDFORD@STANFORD.EDU

Jiayi Wang

JYW@STANFORD.EDU

Editors: Matus Telgarsky and Jonathan Ullman

Abstract

We provide a general framework to improve trade-offs between the number of *full batch* and *sample* queries used to solve structured optimization problems. Our results apply to a broad class of randomized optimization algorithms that iteratively solve sub-problems to high accuracy. We show that such algorithms can be modified to *reuse the randomness* used to query the input across sub-problems. Consequently, we improve the trade-off between the number of gradient (full batch) and individual function (sample) queries for finite sum minimization, the number of matrix-vector multiplies (full batch) and random row (sample) queries for top-eigenvector computation, and the number of matrix-vector multiplies with the transition matrix (full batch) and generative model (sample) queries for optimizing Markov Decision Processes. To facilitate our analysis we introduce the notion of *pseudo-independent algorithms*, a generalization of pseudo-deterministic algorithms (Gat and Goldwasser, 2011), that quantifies how independent the output of a randomized algorithm is from a randomness source.

Keywords: Variance reduction, sample reuse, finite-sum minimization, matrix-vector games, Markov decision processes, top-eigenvector computation.

1. Introduction

Variance reduction is a technique for designing efficient algorithms for solving structured optimization problems. This technique consists of reducing the variance of stochastic or randomized access to a component of the input (what we call *sample queries*) by performing occasional, more expensive queries, which involve the entire input (what we call *full batch queries*). Variance reduction (Johnson and Zhang, 2013) has led to improved query complexities for many problems: finite-sum minimization (Frostig et al., 2015; Johnson and Zhang, 2013; Lin et al., 2015), top eigenvector computation (Garber et al., 2016), Markov decision processes (MDPs) (Sidford et al., 2023, 2018a; Jin et al., 2024), and matrix games (Carmon et al., 2019).

Variance reduction schemes often apply an iterative method—which we refer to as an *outer-solver*—to reduce an original problem to solving a sequence of *sub-problems*. The outer-solver runs n_{outer} iterations and, in each iteration, a sub-problem is carefully constructed and solved using what we call a *sub-solver*. This sub-solver is a randomized algorithm that uses n_{batch} full batch and n_{sample} *fresh* sample queries to the input to solve a sub-problem. This solves the original problem

with a total of $n_{\text{outer}}n_{\text{batch}}$ full batch and $n_{\text{outer}}n_{\text{sample}}$ sample queries. Often, it is possible to tune the outer loop to trade off how many sub-problems are solved (n_{outer}), with how challenging each sub-problem is to solve (n_{batch} and n_{sample}).

The central question in this work is: *can we improve this trade-off between the number of full batch and sample queries in prominent variance reduction settings?* We ask whether randomness in sample queries can be *reused* across *all* n_{outer} iterations of the outer-solver to reduce the total number of sample queries from $n_{\text{outer}}n_{\text{sample}}$ to just n_{sample} , without sacrificing correctness guarantees. We answer affirmatively by designing a *sample reuse framework* that reuses randomness in variance-reduction settings where (1) the outer-solver is robust to ℓ_∞ -bounded random noise in sub-problem solutions, and (2) the sub-solver uses randomness *obliviously*, i.e., sampling a random variable that is independent of the sub-problem. We show that for many structured optimization problems, both properties hold, and consequently the same sample queries *can* be reused across all n_{outer} iterations of the outer-solver.

Applying our sample reuse framework enables us to decrease the total number of sample queries by a factor of n_{outer} for finite-sum minimization, top eigenvector computation, and composite ℓ_2 - ℓ_2 matrix games. We also propose a new outer-solver framework for solving *discounted Markov Decision Processes* (DMDPs) that reduces solving a DMDP to solving a sequence of DMDPs of *lower* discount factor. This result—which may be of independent interest—allows us to obtain improved sample query complexities for discounted MDPs and average-reward MDPs as well as faster running times in certain cases. Finally, we show how to apply our framework to obtain sample query complexity improvements for ℓ_2 - ℓ_1 matrix games, where prior variance-reduction schemes use a mix of oblivious and non-oblivious sampling.

Organization. This introduction discusses a motivating, illustrative example of finite-sum minimization (Section 1.1), our main results for other structured optimization problems (Section 1.2), and preliminaries (Section 1.3). Section 2, describes our sample-reuse framework and a new notion of *pseudoindependent* algorithms, which generalizes the concept of *pseudodeterminism* introduced in prior work [Gat and Goldwasser \(2011\)](#) and enables our analysis of this sample-reuse framework. Details on how our framework can be applied to many other structured optimization problems can be found in Sections 3 through 7. Section 8 concludes with a summary and some directions for future work. Omitted proofs are in the Appendix.

1.1. Motivating example: convex finite-sum minimization (FSM)

As an illustrative, motivating example, consider the prototypical *finite-sum minimization (FSM problem)* (introduced in detail in Section 3). In the FSM problem, we are given convex, L -smooth functions $f_1, \dots, f_n : \mathbb{R}^d \rightarrow \mathbb{R}$. In the standard query model, we have an oracle that, when queried at $\mathbf{x} \in \mathbb{R}^d$ and $i \in [n]$, outputs $\nabla f_i(\mathbf{x})$. The goal of FSM is to minimize $F : \mathbb{R}^d \rightarrow \mathbb{R}$ defined as $F(\mathbf{x}) := \frac{1}{n} \sum_{i \in [n]} f_i(\mathbf{x})$ provided that F is μ -strongly convex. (The individual f_i need not be strongly convex.)

Near-optimal query complexities can be obtained for FSM by applying an outer-solver such as accelerated proximal point/Catalyst (APP) ([Frostig et al., 2015](#); [Lin et al., 2015](#)) with stochastic variance-reduced gradient descent (SVRG) as the sub-solver ([Johnson and Zhang, 2013](#)). APP solves the FSM problem by solving $n_{\text{outer}} = \tilde{O}(\sqrt{\alpha/\mu})$ regularized problems of the form $\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) + \frac{\alpha}{2} \|\mathbf{x} - \mathbf{x}_t\|_2^2$ where for each iteration $t \in [n_{\text{outer}}]$, \mathbf{x}_t depends on the solution to the $(t - 1)$ -th sub-problem

(Frostig et al., 2015).¹ Each sub-problem is solved using $\tilde{O}(n + L/\alpha)$ queries, e.g., using SVRG as the sub-solver (Johnson and Zhang, 2013). Taking $\alpha = \max\{L/n, \mu\}$ yields a query complexity of $\tilde{O}(n + \sqrt{nL/\mu})$, which is known to be the optimal, up to logarithmic factors (Woodworth and Srebro, 2016; Agarwal and Bottou, 2015).

The batch-sample model. In light of this prior work, to obtain further improvements, we *must go beyond* this standard query model. Consequently, in this paper we consider a more *fine-grained batch-sample query model*, which better captures computational trade-offs that could be present in some settings.

To motivate this model, a closer inspection of the aforementioned algorithm (APP with SVRG) reveals that solving each sub-problem requires $\tilde{O}(1)$ computations of the gradient of F and $\tilde{O}(L/\alpha)$ computations of the gradient of a component f_i , where i is chosen uniformly at random. Using the fact that the a gradient of F can be computed by querying each of the n ∇f_i once (n queries in the standard query model), yields the aforementioned near-optimal query complexity of $\tilde{O}(n + \sqrt{nL/\mu})$ in the standard query model.

However, treating all queries to the gradient of f_i as *computationally equivalent* can obscure the structure of the problem. For example, in some practical computational models, computing the gradient of F could be *much cheaper* than computing ∇f_i for n arbitrary f_i —for example, due to caching behavior or memory layout (Fischetti et al., 2018).² Likewise, querying the gradient of the same f_i multiple times could be cheaper than making the same number of queries to different f_i . In such settings, it could be helpful to optimally trade-off between the number of queries to a gradient of F (batch-queries) and the number of queries to a gradient of a random f_i (sample-queries), depending on their relative costs.

Moreover, in some classic problems of FSM—namely, empirical risk minimization—just $O(1)$ queries to the gradient of f_i suffices to *exactly recover* f_i and know ∇f_i at all points *without any* further queries to f_i . This occurs, for example in linear regression when $f_i(\mathbf{x}) = \frac{1}{2}(\mathbf{a}_i^\top \mathbf{x} - \mathbf{b}(i))^2$ for feature vectors $\mathbf{a}_i \in \mathbb{R}^d$ and (explicitly known) labels $\mathbf{b} \in \mathbb{R}^n$ and for generalized linear models when $f_i(\mathbf{x}) = \phi_i(\mathbf{a}_i^\top \mathbf{x} - \mathbf{b}(i))$ for known ϕ_i (see Section 3.1). Hence, *repeatedly querying the gradient of the same f_i* , say T times, can be much cheaper than querying T arbitrary f_i . For instance, this can be the case in distributed memory layouts Woodworth et al. (2020); Sallinen et al. (2016).

To capture these nuances, we consider a more fine-grained analysis of the complexity of FSM where we allow *smoothly trading off* between two types of queries, depending on their relative costs:

- **Full batch query:** for input \mathbf{x} computes $\nabla F(\mathbf{x})$, and
- **Sample query:** for input $i \in [n]$ allows $\nabla f_i(\mathbf{x})$ to be computed for *any* future input \mathbf{x} .

For example, recall from above that in the special case of linear regression, each $f_i(\mathbf{x}) = \frac{1}{2}(\mathbf{a}_i^\top \mathbf{x} - \mathbf{b}(i))^2$ where $\mathbf{a}_i \in \mathbb{R}^d$ is row i of a data matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and \mathbf{b} is a label. In this case, a batch oracle

1. As in prior work, throughout, we may use $\tilde{O}(\cdot)$ to hide poly-logarithmic factors in problem parameters.
 2. Fischetti et al. (2018) demonstrated empirically that the technique of *minbatch persistency*, which reuses the same minibatch for consecutive SGD iterations, better exploits modern GPUs by reducing the overhead related to data loading and moreover, may lead to faster convergence in practice. Previously, Sallinen et al. (2016) also studied empirical variants of SGD designed to reduce communication in distributed memory stsems and improve cache utilization while preserving statistical efficiency. These empirical findings motivate our theoretical analysis of a more fine-grained batch-sample query model which might better capture such practical considerations.

query is implementable just with the appropriate *matrix-vector multiply*, $\mathbf{A}^\top \mathbf{A}x$, and a sample query is implementable by outputting the appropriate row, \mathbf{a}_i (see Section 3).

This batch-query model captures (1) the fact that full batch queries can be cheaper than n -sample queries due to caching layout as well as (2) the fact that re-querying previously cached components is often cheaper than querying fresh components of F due to caching and communication costs between machines. Depending on the computational environment, one can now seek to optimally *trade-off* these two types of oracle queries, depending on their relative costs. From this perspective, state-of-the-art FSM algorithms consist of $n_{\text{outer}} = \tilde{O}(\sqrt{\alpha/\mu})$ outer loop iterations of the APP outer-solver. The sub-solver in each iteration (SVRG) is implementable with $n_{\text{batch}} = \tilde{O}(1)$ full batch queries and $n_{\text{sample}} = \tilde{O}(L/\alpha)$ -sample queries. Tuning α allows us to smoothly trade off between these two types of oracle queries.

As mentioned, prior work established that $\tilde{O}(n + \sqrt{nL/\mu})$ -sample queries is near-optimal for FSM if one *only uses sample queries* (Woodworth and Srebro, 2016). This lower bound, when $n = 1$, also implies $\tilde{O}(\sqrt{L/\mu})$ full batch queries is optimal when using *only* full batch queries. Thus, it is perhaps natural to expect that solving FSM with $\tilde{O}(\sqrt{\alpha/\mu})$ full batch queries and $\tilde{O}(L/\sqrt{\mu\alpha})$ -sample queries is the optimal query complexity trade-off—after all, it recovers the optimal rate for using *only* full batch queries when $\alpha = L, n = 1$ and the optimal sample complexity for using *only* sample queries if $\alpha = \max(L/n, \mu)$.) Still, we ask: *can this trade-off be improved?*

Our FSM results. We show that this trade-off can be improved! By developing and applying techniques for reusing randomness (Section 2), we provide a method that uses only $\tilde{O}(\sqrt{\alpha/\mu})$ -full batch queries and $\tilde{O}(L/\alpha)$ -sample queries for any $\alpha \geq \mu$ to solve FSM. That is, we decrease the number of sample queries by $\tilde{O}(n_{\text{outer}}) = \tilde{O}(\sqrt{\alpha/\mu})$. This yields corresponding improvements for regression and generalized linear models (see Section 3.1) and shows these problems can be solved *with less information than was known previously*. Although this doesn't directly yield a worst-case asymptotic-runtime improvement that we are aware of, it sheds new light on the information needed to solve FSM and could yield faster algorithms depending on caching and memory layout.

To obtain this improved trade-off, we observe that in the previously discussed variance-reduction approach, the sub-solver solves each sub-problem to high accuracy by querying the sample oracle for a uniformly random component i . Importantly, this distribution for choosing the i *does not depend* on the particular sub-problem, i.e., the samples are *oblivious* of the point at which the gradients are queried. We refer to such queries as *oblivious sample queries*. As mentioned earlier, we show that by solving the regularized sub-problems to high accuracy and adding a small amount of uniform noise to the output, it is possible to *reuse the same* i in each sub-problem. To facilitate this proof, we introduce a notion of *pseudo-independence* (Section 2), a generalization of pseudo-determinism (Gat and Goldwasser, 2011), and provide general theorems about pseudo-independence in Appendix B.

1.2. Our results for structured optimization

Here we explain our results for prominent variance reduction settings (including FSM). These results follow a similar approach as in Section 1.1, but differ in oracles and sub-problem structure. For each problem we consider using an outer-solver framework (e.g., APP) to iteratively reduce the original optimization problem to a sequence of sub-problems, which are solved to high-accuracy using a sub-solver (e.g., SVRG). (This is depicted by *Outer-Solver* and *Standard Sub-Solver* in Figure 1.)

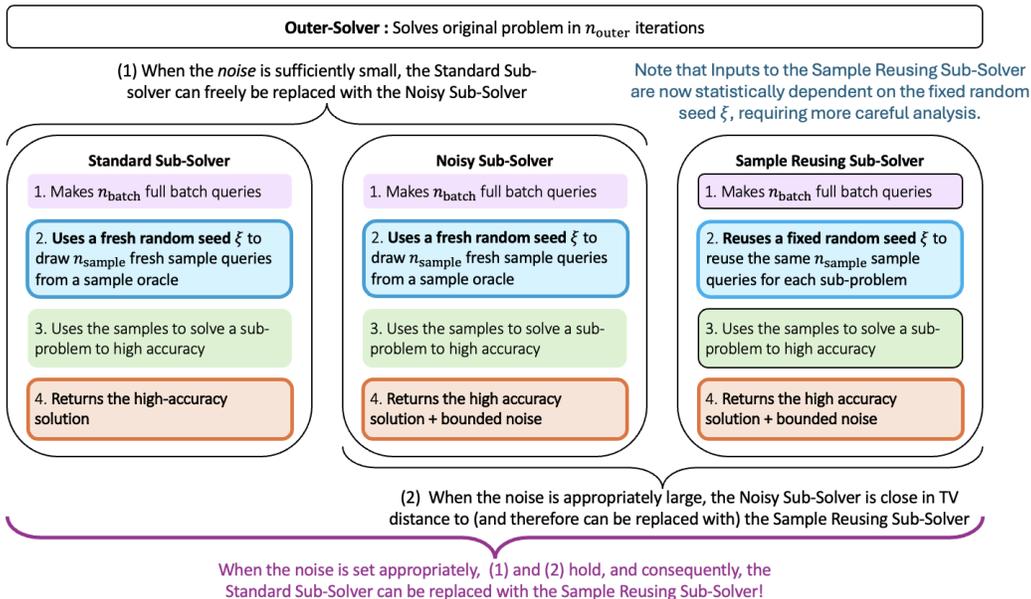


Figure 1: *Sample reuse framework*. The diagram summarizes our approach for replacing Standard Sub-Solvers with a Sample Reusing Sub-Solver, which reuses the same sample queries across all n_{outer} iterations of the Outer-Solver. In (2), TV abbreviates total variation distance. For a more detailed explanation of how the random seed is reused from the inner loop to the outer loop, please see Algorithm 1 and Table 3.

Two observations drive these results. First, in these problems, an outer-solver framework is guaranteed to (probably, approximately) solve the original problem, even if at each iteration, the solution to the t -th sub-problem is perturbed by some bounded random noise. (This is depicted by the *Noisy Sub-Solver* in Figure 1.) Second, when this random noise has a sufficiently large range, the random perturbations protect the randomness of the sample queries well enough so that even if we reuse the *same* sample queries across *all* iterations of the outer-solver (this is depicted by *Sample Reusing Sub-Solver* in Figure 1), the distribution over outputs does not change by much—as measured by total variation (TV) distance. To prove this fact, we describe and analyze a new notion of *pseudo-independence* of randomized algorithms in Section 2.

Combining these observations, we show that in many problems, when the noise is carefully scaled, the Outer-Solver can use the Sample-Reusing Sub-solver as the sub-problem solver (instead of the Standard Sub-Solver). This reduces the total number of sample queries by a multiplicative n_{outer} factor. This is perhaps surprising—it is well-known that randomized algorithms are *not* always amenable to reusing randomness across sequential invocations, see e.g., (Cherapanamjeri and Nelson, 2020; Cohen et al., 2024). However, our analysis shows randomness *can* be reused in many applications of variance reduction.

This sample reuse framework is formalized in Section 2. Although this formalism is somewhat abstract and technical—as it is intended to capture a wide range of variance-reduction settings—it enables us to obtain improved query complexity trade-offs for solving a broad range of prominent optimization and machine-learning problems. These improvements are summarized in Table 1. Note that in certain specialized problems, e.g., TopEV and special cases of FSM such as linear regression, there might be other approaches to obtain our improved trade-offs using preconditioning

(Cohen et al., 2020b, 2015). However, a strength of our sample reuse framework is its versatility—our approach applies even in problems where there is no clear preconditioning analog, e.g., MDPs or matrix games. In the next paragraphs, we describe the implications of our results for several problems beyond FSM as summarized in Table 2.

Markov Decision Processes. In Section 4 we consider the problem of computing an ϵ -optimal policy for a Markov Decision Process (MDP). We denote an MDP by $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbf{P})$ where \mathcal{S} denotes a finite state space, \mathcal{A} denotes the total set of all state-action pairs, and \mathbf{P} denotes the state-action-state transition matrix. We use \mathcal{A}_{tot} to denote the total number of state-action pairs (see Section 4 for more formal treatment.) We consider two well-studied variants of the problem: γ -discounted MDPs (DMDPs) and average-reward MDPs (AMDPs). In both settings, a *sample query* corresponds to drawing a *random state* from $\mathbf{p}(s, a)$ where $\mathbf{p}(s, a)$ is the transition probability distribution corresponding to (s, a) -th row of \mathbf{P} . Meanwhile, a *full batch* query corresponds to a matrix-vector product with the transition matrix \mathbf{P} .

Prior to our work, the state-of-the-art full batch versus sample query trade-off for DMDPs was obtained by Jin et al. (2024), who provided an algorithm to compute an ϵ -optimal policy for a γ -discounted DMDP in $\tilde{O}(1)$ -full batch queries and $\tilde{O}(\mathcal{A}_{\text{tot}}(1-\gamma)^{-2})$ -sample queries for any $\gamma \in (0, 1)$. On the other hand, classic value iteration obtains an ϵ -optimal policy for a γ -discounted DMDP in $\tilde{O}((1-\gamma)^{-1})$ -full batch queries and no sample queries. However, note that (in contrast to the *other* structured optimization problems in Table 2) prior work did *not* provide a smooth trade-off between the number of full batch versus sample queries.

To enable a smooth trade-off between the number of full batch and sample queries, we develop a new outer-solver framework which reduces solving a γ -discounted DMDP to solving a sequence of γ' -discounted DMDPs for $\gamma' < \gamma$ (Theorem 35); we believe this result may be of independent interest. Combining our new outer-solver framework for DMDPs with our sample reuse framework, we obtain a *tunable* trade-off between the number of full batch and sample queries needed to solve a γ -DMDP. We show that for any $\alpha \in [1-\gamma, 1)$, there is an algorithm that solves the problem using $\tilde{O}(\alpha(1-\gamma)^{-1})$ -full batch queries and $\tilde{O}(\mathcal{A}_{\text{tot}}\alpha^{-2})$ -sample queries. Interestingly, our new outer-solver framework for DMDPs also yields a new algorithm for solving γ -DMDPs to high-accuracy and improves the *runtime* of prior work (Jin et al., 2024) under appropriate conditions on the *sparsity* of the underlying transition matrix (Theorem 38.) Finally, to extend our results for DMDPs to AMDPs, we leverage a reduction of (Jin and Sidford, 2021).

Problem	Description
Finite-sum minimization (FSM)	$F(x) = \frac{1}{n} \sum_{i \in [n]} f_i(x)$ is a μ -strongly convex function where each f_i is L -smooth and convex. The goal is to reduce the error (with respect to the minimizer of F) of an initial point by a factor of $c > 1$ wp. $1 - \delta$. (Section 3 and, in greater generality, Section 6.)
Discounted MDP (DMDP)	The goal is to compute an ϵ -optimal policy for a γ -discounted infinite-horizon MDP wp. $1 - \delta$. We use \mathcal{A}_{tot} to denote the total number of state-action pairs in the MDP and assume that rewards and \mathcal{A}_{tot} are bounded. (Section 4)
Average-reward MDP (AMDP)	The goal is to compute an ϵ -optimal policy for an average-reward infinite-horizon MDP wp. $1 - \delta$. We use \mathcal{A}_{tot} to denote the total number of state-action pairs in the MDP and assume that rewards and \mathcal{A}_{tot} are bounded. (Section 4)
Composite ℓ_2 - ℓ_2 matrix games (ℓ_2 - ℓ_2)	The goal is to compute an ϵ -saddle point for an ℓ_2 - ℓ_2 matrix game in matrix $\mathbf{A} \in \mathbb{R}^{d \times n}$ wp. $1 - \delta$. $\ \mathbf{A}\ _F := (\sum_i \sum_j \mathbf{A}_{i,j}^2)^{1/2}$ is the Frobenius norm. (Section 5)
Top Eigenvector (TopEV)	The goal is to compute an ϵ -approximate top eigenvector of $\mathbf{A}^\top \mathbf{A} \in \mathbb{R}^{n \times d}$ with high probability in d . Here $\text{gap}(\mathbf{A})$ and $\text{sr}(\mathbf{A})$ are the relative eigen-gap and stable rank of \mathbf{A} respectively. (Section 7)

Table 1: Summary of structured optimization problems (and abbreviations) studied in this work. Throughout this paper, wp. abbreviates “with probability.”

Problem	Prior Work (\tilde{O})			Our Trade-off (\tilde{O})		Range
	FB	OS	Paper	FB	OS	
FSM	$\sqrt{\alpha/\mu}$	$L/\sqrt{\alpha\mu}$	Frostig et al. (2015)	$\sqrt{\alpha/\mu}$	L/α	$\alpha > \mu$
DMDP	1	$\mathcal{A}_{\text{tot}}(1 - \gamma)^{-2}$	Jin et al. (2024)	$\alpha(1 - \gamma)^{-1}$	$\mathcal{A}_{\text{tot}}\alpha^{-2}$	$1 > \alpha \geq 1 - \gamma$
AMDP	1	$\mathcal{A}_{\text{tot}} t_{\text{mix}}^2 \epsilon^{-2}$	Jin and Sidford (2021)	$\alpha t_{\text{mix}} \epsilon^{-1}$	$\mathcal{A}_{\text{tot}} \alpha^{-2}$	$1 > \alpha \geq \frac{\epsilon}{9t_{\text{mix}}}$
ℓ_2 - ℓ_2	$\alpha\epsilon^{-1}$	$\ \mathbf{A}\ _F^2 (\alpha\epsilon)^{-1}$	Carmon et al. (2019)	$\alpha\epsilon^{-1}$	$\ \mathbf{A}\ _F^2 \alpha^{-2}$	$\alpha > 0$
TopEV	$\sqrt{\frac{\alpha}{\text{gap}(\mathbf{A})}}$	$\text{sr}(\mathbf{A})(\alpha^3 \text{gap}(\mathbf{A}))^{-\frac{1}{2}}$	Garber et al. (2016)	$\sqrt{\frac{\alpha}{\text{gap}(\mathbf{A})}}$	$\text{sr}(\mathbf{A}) \alpha^{-2}$	$\alpha > \Theta(\text{gap}(\mathbf{A}))$

Table 2: *Main results.* FB and OS denote the required full batch and oblivious (non-adaptive) sample queries, respectively, with α tuning their trade-off. We compare our trade-offs with prior work. Importantly, the “Our trade-off” column always improves over the trade-off under “Prior Work” under the *same* assumptions and problem definitions made in the prior work.

Matrix games. In Section 5.2 we study the problem of computing an ϵ -Nash equilibrium for a composite ℓ_2 - ℓ_2 matrix game, which is the problem of finding an ϵ -approximate saddle point of the following minimax problem: $\min_{\mathbf{x} \in \mathbb{B}^n} \max_{\mathbf{y} \in \mathbb{B}^m} \mathbf{y}^\top \mathbf{A} \mathbf{x} + \phi(\mathbf{x}) - \psi(\mathbf{y})$ where \mathbb{B}^k is the k -dimensional Euclidean ball, and ϕ, ψ are (explicitly known) convex functions. In this setting, a *sample* query corresponds to sampling $\mathbf{A}_{i,j}$ while a *full batch* query corresponds to querying the pair of matrix-vector products $(\mathbf{y}^\top \mathbf{A}, \mathbf{A} \mathbf{x})$. Previously, Carmon et al. (2019) showed how to solve

the problem using $\tilde{O}(\alpha\epsilon^{-1})$ -full batch queries and $\tilde{O}(\|\mathbf{A}\|_F^2(\alpha\epsilon)^{-1})$ -sample queries, for any $\alpha > 0$. Using our sample reuse framework, we improve this trade-off and show how to solve the problem using $\tilde{O}(\alpha\epsilon^{-1})$ -full batch queries and $\tilde{O}(\|\mathbf{A}\|_F^2\alpha^{-2})$ -sample queries, for any $\alpha > 0$.

To highlight one interesting application of this, note that for composite ℓ_2 - ℓ_2 matrix games, *both* full batch queries and sample queries can be implemented by a (two-sided) matrix-vector oracle which outputs $(\mathbf{A}\mathbf{x}, \mathbf{A}^\top\mathbf{y})$ for a $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d \times \mathbb{R}^n$. With $\alpha = \|\mathbf{A}\|_F^{2/3} \epsilon^{1/3}$ we obtain an algorithm that solves ℓ_2 - ℓ_2 matrix games in only $\tilde{O}(\|\mathbf{A}\|_F^{2/3} \epsilon^{-2/3})$ -matrix-vector oracle queries. This improves the matrix-vector complexity of the problem over the $\tilde{O}(\|\mathbf{A}\|_F \epsilon^{-1})$ -matrix-vector oracle queries required in prior work (Carmon et al., 2019; Nemirovski, 2004) and is near-optimal, due to (Liu et al., 2023). *Special* cases of composite ℓ_2 - ℓ_2 matrix games reduce to ℓ_2 -regression, in which case this trade-off of $\tilde{O}(\|\mathbf{A}\|_F^{2/3} \epsilon^{-2/3})$ -matrix-vector oracle queries may also follow from preconditioning (Cohen et al., 2020b; Liu et al., 2023). *However*, our work is first to achieve near-optimal matrix-vector oracle query complexity for *general* composite ℓ_2 - ℓ_2 matrix games. Subsequent to our work, Karmarkar et al. (2025) recently showed how to match this matrix-vector oracle query complexity deterministically.

In Section 5.2 we also show how to apply our methods to another class of minimax problems, called composite ℓ_2 - ℓ_1 matrix games. A composite ℓ_2 - ℓ_1 matrix game corresponds to finding an ϵ -approximate saddle point of $\min_{\mathbf{x} \in \mathbb{B}^n} \max_{\mathbf{y} \in \Delta^k} \mathbf{y}^\top \mathbf{A}\mathbf{x} + \phi(\mathbf{x}) - \psi(\mathbf{y})$ where Δ^k is the k -dimensional probability simplex and ϕ, ψ are explicitly known convex functions. This setting captures, for example, the problem of computing a hard-margin support vector machine. Notably, in this setting, variance-reduced methods (Carmon et al., 2019) *combine* oblivious sample queries along with non-oblivious sample queries to solve the sub-problems induced by the outer-solver (conceptual proximal point (Nemirovski, 2004; Carmon et al., 2019)). Our sample reuse framework *still* applies and allows us to *reuse* the oblivious (non-adaptive) sample queries across all n_{outer} invocations of the sub-solver. We also discuss how this improves query complexity trade-offs for two computational geometry problems: the minimum enclosing and maximum inscribed ball problems.

Top eigenvector computation. In Section 7 we consider the problem of computing an ϵ -approximate top-eigenvector for a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$. We use $\text{gap}(\mathbf{A})$ and $\text{sr}(\mathbf{A})$ to denote the relative eigen-gap and stable rank of \mathbf{A} , respectively. In this setting, a *sample query* corresponds to querying an entry A_{ij} , while a *full batch* query corresponds to a matrix-vector product with \mathbf{A} . Prior work of Garber et al. (2016) showed how to build upon the approximate proximal point methods described for FSM Frostig et al. (2015); Lin et al. (2015) and the *shift-and-invert* power method Saad (2011) to compute an ϵ -approximate top-eigenvector using $\tilde{O}(\sqrt{\alpha/\text{gap}(\mathbf{A})})$ -full batch queries along with $\tilde{O}(\text{sr}(\mathbf{A})(\alpha^3 \text{gap}(\mathbf{A}))^{-1/2})$ -sample queries for any $\alpha > \Theta(\text{gap}(\mathbf{A}))$. Using our sample-reuse framework, we improve this trade-off to $\tilde{O}(\sqrt{\alpha/\text{gap}(\mathbf{A})})$ -full batch queries and $\tilde{O}(\text{sr}(\mathbf{A})\alpha^{-2})$ -sample queries.

1.3. Preliminaries

General notation. Bold lowercase letters are vectors in \mathbb{R}^d where $\mathbf{u}(i)$ is the i -th index of \mathbf{u} . Bold capital letters are matrices. The ℓ_p norm of \mathbf{u} is $\|\mathbf{u}\|_p$. For random variable A over the probability space (Ω, \mathcal{F}) , p_A is its probability measure. For event E , $p_{A|E}$ is the measure of A conditioned on E , $\neg E$ is the complement of E , and $\mathbb{1}[E]$ is the indicator of E . For random variables A, B , $A \stackrel{\mathcal{D}}{=} B$ denotes that A and B are equal in distribution. For measures p and q , $d_{TV}(p, q) := \max_{E \in \mathcal{F}} |p(E) - q(E)|$ is the total variation (TV) distance between p and q . The

support of distribution \mathcal{D} or measure q is denoted $\text{supp}(\mathcal{D})$ or $\text{supp}(q)$. Ber and Unif respectively denote Bernoulli and uniform distributions. $\text{Unif}^d(a, b)$ denotes a d -dimensional vector in which the i -th entry is an independently and identically distributed $\text{Unif}(a, b)$ random variable.

Randomized algorithm notation. We consider algorithms that use up to two independent sources of randomness. We use $\mathcal{A}_{\xi, \chi}$ to denote a randomized algorithm that takes an input $\mathbf{u} \in \mathbb{R}^d$ and independent random seeds $\xi \sim \mathcal{D}_\xi$ and $\chi \sim \mathcal{D}_\chi^{\mathbf{u}}$ where \mathcal{D}_ξ is *independent of, or oblivious with respect to, \mathbf{u}* and $\mathcal{D}_\chi^{\mathbf{u}}$ might be *dependent on, or be adaptive with respect to \mathbf{u}* . We also use the notation $\mathcal{D}_\chi = \{\mathcal{D}_\chi^{\mathbf{u}}\}_{\mathbf{u} \in \mathbb{R}^d}$ to denote the family of distributions parameterized by $\mathbf{u} \in \mathbb{R}^d$. On input $\mathbf{u} \in \mathbb{R}^d$, $\mathcal{A}_{\xi, \chi}$ outputs $\mathcal{A}_{\xi, \chi}(\mathbf{u})$ for randomly drawn $\xi \sim \mathcal{D}_\xi$ and $\chi \sim \mathcal{D}_\chi^{\mathbf{u}}$. At times, we analyze the output of $\mathcal{A}_{\xi, \chi}$ *conditioned* on the value of one or both random seeds. Specifically, $\mathcal{A}_{\xi=s, \chi}(\mathbf{u})$ and $\mathcal{A}_{\xi, \chi=c}(\mathbf{u})$ are used to denote the randomized algorithms obtained by fixing the value of $\xi = s \in \text{supp}(\mathcal{D}_\xi)$ or $\chi = c \in \text{supp}(\mathcal{D}_\chi^{\mathbf{u}})$, respectively. Analogously, $\mathcal{A}_{\xi=s, \chi=c}$ is a *deterministic* algorithm corresponding to conditioning on $\xi = s \in \text{supp}(\mathcal{D}_\xi)$ and $\chi = c \in \text{supp}(\mathcal{D}_\chi^{\mathbf{u}})$. We occasionally specify a decomposition of the seed χ into two sub-seeds $\chi = (\nu, \iota)$ where sub-seeds ν, ι are drawn independently from $\nu \sim \mathcal{D}_\nu^x$ and $\iota \sim \mathcal{D}_\iota$. We assume algorithms that output vectors in \mathbb{R}^d have runtime $\Omega(d)$. We use *high-accuracy* to refer to families of algorithms (parameterized by error and failure probability parameters) with at most *polylogarithmic* dependence on their accuracy and failure probability.

2. Sample reuse framework

Here we provide, contextualize, and analyze our sample-reuse framework. First, in Section 2.1 we introduce the framework, the main theorem regarding it (Theorem 6) and formally define *pseudoindependence* (Definition 5) which we use to prove the theorem. For additional context, we then compare pseudo-independence to related definitions in prior work in Section 2.2. We then conclude this section in Section 2.3 by analyzing the sample-reuse framework and proving Theorem 6.

2.1. Sample-reuse framework

Here, we introduce our sample-reuse framework. To describe the framework we provide a general *Meta-Algorithm* for solving an optimization problem. This Meta-Algorithm encompasses three different templates for variance reduction methods. Our framework relates these different templates and applying this relationship to different algorithms yields our improved query-complexity trade-offs.

More formally, the Meta-Algorithm 1 iteratively reduces solving a problem instance X to solving a sequence of sub-problems. First, the algorithm initializes $\mathbf{u}_0 \in \mathbb{R}^d$; an oblivious distribution \mathcal{D}_ξ for sampling a random seed ξ ; and a *family of non-oblivious* distributions $\mathcal{D}_\chi = \{\mathcal{D}_\chi^x\}_{x \in \mathbb{R}^d}$. It iteratively runs one of *three* possible sub-routines (corresponding to the three different templates), which we describe below, and outputs a convex combination of the iterates. Table 3 summarizes the notation.

The standard sub-routine. The first template is $\text{Outer}(X; \text{Standard})$, i.e., Meta-Algorithm 1 with sub-routine StandardLoop (in the prose, we omit the τ since it has no effect when $\text{Type} = \text{Standard}$.) This template formalizes the ‘‘Standard Sub-Solver’’ panel in Figure 1 and describes the standard variance reduction template that applies to many algorithms for structured optimization problems—including those cited under ‘‘Prior Work’’ in Table 2. In $\text{Outer}(X; \text{Standard})$, the for

loop (Line 1) iteratively reduces the original optimization problem to solving n_{outer} sub-problems, where the t -th sub-problem is determined by the $(t - 1)$ -th iterate, \mathbf{u}_{t-1} .

The sub-problem solver, or *sub-solver*, denoted $\mathcal{A}_{\xi, \chi}^{\text{sub}}$, is a randomized algorithm that solves the sub-problem at each iteration by (1) making some deterministic full batch queries to the full batch oracle; (2) using the random seed ξ to make randomized *oblivious* sample queries to a sample oracle; and (3) using the random seed χ to perform some additional randomization. Depending on the application (Table 1), (3) might involve adding some noise to the sub-problem or making additional *adaptive* sample queries to a sample oracle. $\mathcal{A}_{\xi, \chi}^{\text{sub}}$ then outputs an intermediate iterate $\mathbf{u}_{t-1/2} \in \mathbb{R}^p$. At the end of each iteration, the routine then applies a “post-process” ζ which performs some *deterministic* post-processing of the pair $(\mathbf{u}_t, \mathbf{u}_{t-1/2})$ (e.g., implementing acceleration/momentum.)

As an illustrative, concrete, example, consider the motivating example of FSM (Section 1.1). Meta-Algorithm 1 captures the previously described combination of APP and SVRG to solve FSM (Frostig et al., 2015; Lin et al., 2015). That is, APP reduces the original problem to solving a sequence of regularized sub-problems, while $\mathcal{A}_{\xi, \chi}^{\text{sub}}$ represents the sub-solver SVRG (Johnson and Zhang, 2013), which is used to solve the regularized sub-problems. In this case, ζ implements acceleration as in (Frostig et al., 2015), and \mathbf{w} is set such that the algorithm will simply return the last iterate, $\mathbf{u}_{n_{\text{outer}}}$. (See also Section 3.)

The noisy sub-routine. To motivate the second template, recall our motivating observation that for many classes of structured optimization problems (Table 1), there is an instantiation of $\text{Outer}(X; \text{Standard})$ that provably solves the problem instance X provided that: at each iteration $t \in [n_{\text{outer}}]$, the sub-solver $\mathcal{A}_{\xi, \chi}^{\text{sub}}$ solves the sub-problem induced by \mathbf{u}_{t-1} to *high-accuracy* in the ℓ_∞ norm. We formalize this observation with the following two definitions.

Definition 1 (Function approximation) We say that a randomized algorithm $\mathcal{A}_{\xi, \chi}$ is an (η, δ) -approximation of $f : \mathbb{R}^d \rightarrow \mathbb{R}^p$ if $\mathcal{A}_{\xi, \chi}$ takes an input $\mathbf{u} \in \mathbb{R}^d$ along with two random seeds $\xi \sim \mathcal{D}_\xi$ and $\chi \sim \mathcal{D}_\chi^{\mathbf{u}}$ and outputs $\mathcal{A}_{\xi, \chi}(\mathbf{u}) \in \mathbb{R}^p$ such that $\mathbb{P}_{\xi \sim \mathcal{D}_\xi, \chi \sim \mathcal{D}_\chi^{\mathbf{u}}} (\|\mathcal{A}_{\xi, \chi}(\mathbf{u}) - f(\mathbf{u})\|_\infty \leq \eta) \geq 1 - \delta$.

Definition 2 (ℓ_∞ -robust) We say $\text{Outer}(X; \text{Standard})$ is (η, β) -robust with respect to f^{sub} if there exists a deterministic function $f^{\text{sub}} : \mathbb{R}^d \rightarrow \mathbb{R}^p$ such that $\text{Outer}(X; \text{Standard})$ is guaranteed to output a β -accurate solution for X whenever, for all $t \in [n_{\text{outer}}]$, $\|\mathbf{u}_{t-1/2} - f^{\text{sub}}(\mathbf{u}_{t-1})\|_\infty \leq \eta$.

To interpret these definitions, suppose that $\text{Outer}(X; \text{Standard})$ is (η, β) -robust, $\tau \leq \eta/2$, and the sub-solver $\mathcal{A}_{\xi, \chi}^{\text{sub}}$ is an $(\eta/2, \delta)$ -approximation of f^{sub} . Then *even if*, at every iteration $t \in [n_{\text{outer}}]$, $\mathbf{u}_{t-1/2}$ were to be randomly perturbed by some τ -bounded random noise, the output would *still* be a β -accurate solution to X wp. $1 - n_{\text{outer}}\delta$.

Correspondingly our second sub-routine NoisyLoop_τ in Meta-Algorithm 1 (and the corresponding template $\text{Outer}(X; \text{Noisy})$) is *identical* to StandardLoop (correspondingly, the algorithmic template $\text{Outer}(X; \text{Standard})$) *except* that at each iteration, $\mathbf{u}_{t-1/2}$ is *perturbed* by a small amount of $\text{Unif}^p(-\tau, \tau)$ random noise. This formalizes the “Noisy Sub-Solver” in Figure 1. Following the previous intuition, the following lemma shows that when τ is sufficiently small, we can freely interchange $\text{Outer}(X; \text{Standard})$ with $\text{Outer}(X; \text{Noisy}, \tau)$, without losing correctness guarantees.

Lemma 3 Suppose $\text{Outer}(X; \text{Standard})$ is (η, β) -robust with respect to f^{sub} , $\tau \in (0, \eta/2)$, and $\mathcal{A}_{\xi, \chi}^{\text{sub}}$ is an $(\eta/2, \delta)$ -approximation of f^{sub} . Then $\text{Outer}(X; \text{Standard})$ and $\text{Outer}(X; \text{Noisy}, \tau)$, wp. $1 - n_{\text{outer}}\delta$, output a β -accurate solution to X .

Table 3: Parameter table for Meta Algorithm (Algorithm 1).

Symbol	Description
\mathcal{D}_ξ	This is an oblivious distribution governing the random variable ξ .
\mathcal{D}_χ	This is a family of distributions parameterized by $\mathbf{x} \in \mathbb{R}^d$. It governs the <i>adaptive</i> random variable χ . We use $\mathcal{D}_\chi^{\mathbf{x}}$ to denote the distribution in \mathcal{D}_χ corresponding to an $\mathbf{x} \in \mathbb{R}^d$.
$\mathcal{A}_{\xi,\chi}^{\text{sub}}$	$\mathcal{A}_{\xi,\chi}^{\text{sub}}$ is a sub-problem solver (sub-solver) that takes in $\mathbf{u} \in \mathbb{R}^d$ and seeds ξ and χ and outputs $\mathcal{A}_{\xi,\chi}^{\text{sub}}(\mathbf{u}) \in \mathbb{R}^p$. $\mathcal{A}_{\xi,\chi}^{\text{sub}}$ may make full batch queries and sample queries to X . Batch queries depend <i>only</i> on \mathbf{u} . The seed ξ is used to make <i>oblivious</i> sample queries. Meanwhile, the seed χ may optionally be used to make additional <i>adaptive</i> sample queries.
ζ	$\zeta : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}^d$ is a deterministic post-processing function; we call it an <i>outer-process</i> since it captures the role of the <i>outer-solver</i> discussed in Section 1 (Figure 1).
\mathbf{w}	The algorithm outputs a convex combination of the \mathbf{u}_t given by coefficient vector \mathbf{w} .
τ	Noise parameter $\tau > 0$ controls the amount of noise to be added.

Proof The proof is immediate from Definition 2 and union bound over all n_{outer} iterations. ■

Algorithm 1: Variance-Reduction Meta-Algorithm $\text{Outer}(X; \text{Type}, \tau)$

Parameter: Loop specification: $\text{Type} \in \{\text{Standard}, \text{Noisy}, \text{Reuse}\}$ // Types defined below

Initialize $\mathbf{u}_0 \in \mathbb{R}^d$

// If $S = \text{Standard}$, τ may be omitted, in which case Line 1 is skipped and τ is omitted in Line 1

// The next line builds a noisy version of $\mathcal{A}_{\xi,\chi}^{\text{sub}}$, which is denoted $\mathcal{A}'_{\xi,\chi'}$

Let $\mathcal{D}_{\chi'}^{\mathbf{u}_{t-1}} := \mathcal{D}_\chi^{\mathbf{u}_{t-1}} \times \text{Unif}^p(-\tau, \tau)$ and $\mathcal{A}'_{\xi=s, \chi'=(c,e)}(\mathbf{u}_{t-1}) := \mathcal{A}_{\xi=s, \chi=c}^{\text{sub}}(\mathbf{u}_{t-1}) + \mathbf{e}$

Draw $s_1, \dots, s_{n_{\text{outer}}} \sim \mathcal{D}_\xi$ // Draw seeds from the oblivious distribution

for each $t \in [n_{\text{outer}}]$ **do call** $\langle \text{Type} \rangle \text{Loop}_\tau$

return: $\sum_{t \in [n_{\text{outer}}]} \mathbf{w}(t) \cdot \mathbf{u}_t$

// Different loops that can be called on Line 1

StandardLoop $_\tau$: Draw $c_t \sim \mathcal{D}_\chi^{\mathbf{u}_{t-1}}$, $\mathbf{u}_{t-\frac{1}{2}} \leftarrow \mathcal{A}_{\xi=s_t, \chi=c_t}^{\text{sub}}(\mathbf{u}_{t-1})$, $\mathbf{u}_t \leftarrow \zeta(\mathbf{u}_{t-1}, \mathbf{u}_{t-\frac{1}{2}})$

NoisyLoop $_\tau$: Draw $(c_t, \mathbf{e}_t) \sim \mathcal{D}_{\chi'}^{\mathbf{u}_{t-1}}$, $\mathbf{u}_{t-\frac{1}{2}} \leftarrow \mathcal{A}'_{\xi=s_t, \chi'=(c_t, \mathbf{e}_t)}(\mathbf{u}_{t-1})$,

$\mathbf{u}_t \leftarrow \zeta(\mathbf{u}_{t-1}, \mathbf{u}_{t-\frac{1}{2}})$

ReuseLoop $_\tau$: Draw $(c_t, \mathbf{e}_t) \sim \mathcal{D}_{\chi'}^{\mathbf{u}_{t-1}}$, $\mathbf{u}_{t-\frac{1}{2}} \leftarrow \mathcal{A}'_{\xi=s_1, \chi'=(c_t, \mathbf{e}_t)}(\mathbf{u}_{t-1})$,

$\mathbf{u}_t \leftarrow \zeta(\mathbf{u}_{t-1}, \mathbf{u}_{t-\frac{1}{2}})$

The sample-reuse sub-routine. At this point, it is perhaps natural to wonder: *why* is it helpful to work with the `NoisyLoop $_\tau$` subroutine as opposed to the standard, `StandardLoop` sub-routine in Meta-Algorithm 1? After all, for any τ , both `StandardLoop` and `NoisyLoop $_\tau$` require the same number of batch and sample queries, so there is no obvious computational advantage to using `NoisyLoop $_\tau$` .

As discussed in our second observation in Section 1.2, the key idea is the following: the random noise at each iteration ensures that the iterates \mathbf{u}_t in `NoisyLoop τ` are *almost independent* of the randomness in the sample queries induced by $\xi \sim \mathcal{D}_\xi$ in the following sense: *even if we were to reuse the same realization of ξ in all iterations $t \in [n_{\text{outer}}]$ of `NoisyLoop τ` , the returned output would be close—in total variation (TV) distance—to the output of `Outer(X ; NoisyLoop, τ)`. This brings us to our third and final sub-routine.*

The third and final sub-routine `ReuseLoop τ` in Meta-Algorithm 1 (and its corresponding algorithmic template `Outer(X ; Reuse)`) is *identical* to the `NoisyLoop τ` sub-routine (correspondingly, the algorithmic template `Outer(X ; Noisy)`) *except that* `ReuseLoop τ` *reuses* a single realization $s_1 \sim \mathcal{D}_\xi$ in all iterations $t \in [n_{\text{outer}}]$. This formalizes the “Sample Reusing Sub-Solver” in Figure 1. Reusing this realization across all iterations corresponds to reusing oblivious sample queries to X . Consequently, for any $\tau > 0$, `Outer(X ; Reuse, τ)` has lower sample query complexity (by an n_{outer} multiplicative factor) than the original `Outer(X ; Noisy, τ)` or `Outer(X ; Standard)`.

Pseudo-independence. To obtain improved query-complexity tradeoffs, our goal is to derive conditions when `Outer(X ; Reuse, τ)` can be used in place of `Outer(X ; Noisy, τ)`—without sacrificing correctness. We achieve this goal by introducing a new notion of *pseudo-independence*, which we use to bound the total variation (TV) distance between the outputs of `Outer(X ; Noisy, τ)` and `Outer(X ; Reuse, τ)`. Indeed, if this TV distance is small then as long as `Outer(X ; Noisy, τ)` is correct, with high probability, so is `Outer(X ; Reuse, τ)`. To define pseudo-independence we first define a *smoothing* of a randomized algorithm as follows.

Definition 4 ((ϵ, δ)-smoothing) Let $\mathcal{A}_{\xi, \chi}$ be a randomized algorithm which takes an input $\mathbf{u} \in \mathbb{R}^d$ and two random seeds $\xi \sim \mathcal{D}_\xi, \chi \sim \mathcal{D}_\chi^u$. We say that algorithm $\bar{\mathcal{A}}_\chi$ is an (ϵ, δ)-smoothing of $\mathcal{A}_{\xi, \chi}$ with respect to ξ if it takes as an input $\mathbf{u} \in \mathbb{R}^d$ along with one random seed $\chi \sim \mathcal{D}_\chi^u$ and for all $\mathbf{u} \in \mathbb{R}^d, \mathbb{P}_{s \sim \mathcal{D}_\xi} [d_{TV}(p_{\mathcal{A}_{\xi=s, \chi}}(\mathbf{u}), p_{\bar{\mathcal{A}}_\chi}(\mathbf{u})) \leq \epsilon] \geq 1 - \delta$.

We say a randomized algorithm is (ϵ, δ)-pseudo-independent if we can guarantee *existence* of an (ϵ, δ)-smoothing. Importantly, this smoothing need not be implementable or even explicitly known.

Definition 5 ((ϵ, δ)-pseudo-independence) Let $\mathcal{A}_{\xi, \chi}$ be as in Definition 4. $\mathcal{A}_{\xi, \chi}$ is (ϵ, δ)-pseudo-independent of ξ if it admits an (ϵ, δ)-smoothing with respect to ξ .

Intuitively, $\mathcal{A}_{\xi, \chi}$ is (ϵ, δ)-pseudo-independent of ξ if, wp. $1 - \delta$ over the draw of $s \sim \mathcal{D}_\xi$, $\mathcal{A}_{\xi, \chi}$ is *almost independent* of the first source of randomness—in the sense that wp. $1 - \epsilon$ over the draw of χ , $\mathcal{A}_{\xi=s, \chi}(\mathbf{u})$ is equal in distribution to a random variable that is *independent* of ξ (see also Fact 7 for further discussion.)

We use this notion of pseudo-independence to bound the TV distance between `Outer(X ; Noisy, τ)` and `Outer(X ; Standard)`. First, in Appendix A we show that when the noise parameter τ is set appropriately the sub-solver algorithm $\mathcal{A}_{\xi, \chi'}^{\text{sub}}$ in Meta-Algorithm 1 is (ϵ, δ)-pseudo-independent of ξ . Intuitively, this means that $\mathcal{A}_{\xi, \chi'}^{\text{sub}}$ is *almost independent* of ξ and consequently, it should be possible to reuse the same realization of the seed ξ in each invocation of $\mathcal{A}_{\xi, \chi'}^{\text{sub}}$.

To formalize this, we observe that `NoisyLoop τ` repeatedly composes ζ with $\mathcal{A}_{\xi, \chi'}^{\text{sub}}$ where $\xi_t \sim \mathcal{D}_\xi$ is drawn fresh in each iteration $t \in [n_{\text{outer}}]$. Meanwhile, `ReuseLoop τ` repeatedly composes ζ with $\mathcal{A}_{\xi, \chi'}^{\text{sub}}$ where the same realization $s_1 \sim \mathcal{D}_\xi$ is *reused* in each iteration $t \in [n_{\text{outer}}]$. In

Appendix B we provide a general theorems about pseudoindependence, which allow us to leverage the fact that $\mathcal{A}_{\xi, \chi}^{\text{sub}}$ is (ϵ, δ) -pseudo-independent of ξ to bound the TV distance between these repeated compositions as a function of $n_{\text{outer}}, \delta, \epsilon$. Thus, when τ, ϵ, δ are set appropriately, the TV distance between the outputs of $\text{Outer}(X; \text{Noisy})$ and $\text{Outer}(X; \text{Reuse})$ is small. This TV distance bound allows us to prove the following theorem, which in turn yields all the results in Table 2.

Theorem 6 *Suppose $\text{Outer}(X; \text{Standard})$ is (η, β) -robust with respect to f^{sub} and $\delta \in (0, 1)$. Let $\eta' := \min(\eta/2, \eta\delta)$. Suppose $\mathcal{A}_{\xi, \chi}^{\text{sub}}$ is an (η', δ) -approximation of f^{sub} . Then we have that wp. $1 - 5n_{\text{outer}}^2\delta$, $\text{Outer}(X; \text{Reuse}, \eta'/(2\delta))$ outputs a β -accurate solution to X .*

In the problems described in Table 2, the standard variance-reduced algorithms are instantiations of $\text{Outer}(X; \text{Standard})$ and (η, β) -robust, for η scaling polynomially in β and the problem parameters. This ensures that the blowups in error and failure probability in Theorem 6 are at most *polylogarithmic* in all problem parameters. Moreover, for all of our applications, the sub-solvers (e.g., SVRG) are *high-accuracy* algorithms, meaning that they have query- and time-complexity *polylogarithmic* in the accuracy and failure probability parameters. Thus, the polynomial blow-ups in accuracy and failure probability in Theorem 6 imply at most *polylogarithmic* growth in complexity.

2.2. Comparisons of pseudoindependence to prior work

In this section, we discuss *pseudo-independence*, which generalizes the well-studied notion of pseudo-determinism (Definition 8) (Goldwasser et al., 2017; Goldwasser and Grossman, 2017; Braverman et al., 2023; Grossman et al., 2023; Dixon et al., 2022) and is also related to *reproducibility* (Impagliazzo et al., 2022). First, we state the following Fact 7, which will be helpful in our analysis.

Fact 7 (Lemma 4.1.13 of (Roch, 2024), restated) *Let $\epsilon > 0$ and A and B be random variables. Then $d_{TV}(p_A, p_B) \leq \epsilon$ if and only if there exist random variables C, D, F and an independent event E such that $\mathbb{P}\{E\} = 1 - \epsilon$; $A \stackrel{D}{=} C\mathbb{1}_E + D\mathbb{1}_{\neg E}$; and $B \stackrel{D}{=} C\mathbb{1}_E + F\mathbb{1}_{\neg E}$.*

This fact implies the following intuitive interpretation of pseudo-independence. $\mathcal{A}_{\xi, \chi}$ is (ϵ, δ) -pseudo-independent of ξ if wp. $1 - \delta$ over the draw of $s \sim \mathcal{D}_\xi$, $\mathcal{A}_{\xi, \chi}$ is *almost* independent of the first source of randomness—in the sense that wp. $1 - \epsilon$ over the draw of χ , $\mathcal{A}_{\xi=s, \chi}(\mathbf{u})$ is equal in distribution to a random variable that is *independent* of ξ .

Now, we briefly compare pseudo-independence to pseudo-determinism and reproducibility. The term *pseudo-deterministic algorithm*—introduced by Gat and Goldwasser (2011)—describes an algorithm that outputs a deterministic value with high probability. This is formalized with the following definition.

Definition 8 (δ -pseudo-deterministic algorithm) *Let \mathcal{A}_ξ be a randomized algorithm which takes an input $\mathbf{u} \in \mathbb{R}^d$ and a random seed $\xi \sim \mathcal{D}_\xi$. \mathcal{A}_ξ is δ -pseudo-deterministic if there exists a function h over \mathbb{R}^d such that $\mathbb{P}_{s \sim \mathcal{D}_\xi} \{\mathcal{A}_{\xi=s}(\mathbf{u}) = h(\mathbf{u})\} \geq 1 - \delta$.*

Intuitively, in a δ -pseudo-deterministic algorithm \mathcal{A}_ξ , the role of h is similar to that of a smoothing (Definition 4) in the definition of pseudo-independence (Definition 5); however, h must be deterministic whereas as a smoothing can be a randomized algorithm. To formalize this intuition we

need to introduce some additional notation because pseudo-determinism is defined in terms of a single source of randomness, whereas pseudo-independence (Definition 5) is defined in terms of two sources of randomness. Thus, to compare pseudo-determinism to pseudo-independence, we define a simple way to *lift* a single-seed randomized algorithm \mathcal{A}_ξ to two sources of randomness.

Definition 9 Let \mathcal{A}_ξ be a randomized algorithm which takes an input $\mathbf{u} \in \mathbb{R}^d$ and a random seed $\xi \sim \mathcal{D}_\xi$. Let $\mathcal{A}_{\xi,\chi}^{\text{pd}}(\mathbf{u})$ be the randomized algorithm which takes input $\mathbf{u} \in \mathbb{R}^d$ and seeds $\xi \sim \mathcal{D}_\xi$ and $\chi \sim \mathcal{D}_\chi^x$ where $\text{supp}(\mathcal{D}_\chi^x) = \{0\}$; and maps $\mathbf{u} \mapsto \mathcal{A}_\xi(\mathbf{u})$.

That is, $\mathcal{A}_{\xi,\chi}^{\text{pd}}$ is identical to \mathcal{A}_ξ ; however, it accepts an additional 0-bit “dummy” random seed χ . The next lemma explains how pseudo-independence captures pseudo-determinism as a special case.

Lemma 10 \mathcal{A}_ξ is δ -pseudo-deterministic if and only if $\mathcal{A}_{\xi,\chi}^{\text{pd}}$ is $(0, \delta)$ -pseudo-independent of ξ .

Proof First, suppose that \mathcal{A}_ξ is δ -pseudo-deterministic. Then, there exists a function h such that

$$1 - \delta \geq \mathbb{P}_{s \sim \mathcal{D}_\xi} \{ \mathcal{A}_{\xi=s}(\mathbf{u}) = h(\mathbf{u}) \} = \mathbb{P}_{s \sim \mathcal{D}_\xi} \left\{ d_{TV} \left(p_{\mathcal{A}_{\xi=s,\chi}^{\text{pd}}}(\mathbf{u}), p_{h(\mathbf{u})} \right) = 0 \right\}.$$

Hence, if \mathcal{A}'_χ is the algorithm that deterministically maps $\mathbf{u} \mapsto h(\mathbf{u})$ then \mathcal{A}'_χ is a $(0, \delta)$ -smoothing of $\mathcal{A}_{\xi,\chi}^{\text{pd}}$ with respect to ξ . Thus, $\mathcal{A}_{\xi,\chi}^{\text{pd}}$ is $(0, \delta)$ -pseudo-independent with respect to ξ .

Conversely, suppose that $\mathcal{A}_{\xi,\chi}^{\text{pd}}$ is $(0, \delta)$ -pseudo-independent with respect to ξ and let \mathcal{A}'_χ be a $(0, \delta)$ -smoothing of $\mathcal{A}_{\xi,\chi}^{\text{pd}}$ with respect to ξ . Then,

$$1 - \delta \geq \mathbb{P}_{s \sim \mathcal{D}_\xi} \left\{ d_{TV} \left(p_{\mathcal{A}_{\xi=s,\chi}^{\text{pd}}}(\mathbf{u}), p_{\mathcal{A}'_\chi(\mathbf{u})} \right) = 0 \right\} = \mathbb{P}_{s \sim \mathcal{D}_\xi} \{ \mathcal{A}_{\xi=s}(\mathbf{u}) = \mathcal{A}'_{\chi=0}(\mathbf{u}) \}.$$

Letting h to be the function mapping $\mathbf{u} \mapsto \mathcal{A}'_{\chi=0}(\mathbf{u})$, we see that \mathcal{A}_ξ is δ -pseudo-deterministic. ■

On the other hand, *reproducibility* is a related notion introduced in [Impagliazzo et al. \(2022\)](#).

Definition 11 ((δ, \mathcal{D})-reproducible algorithm) Let \mathcal{A}_ξ be a randomized algorithm which takes an input $\mathbf{u} \in \mathbb{R}^d$ and a random seed $\xi \sim \mathcal{D}_\xi$. Let \mathcal{D} be a distribution over \mathbb{R}^d . \mathcal{A}_ξ is a (δ, \mathcal{D}) -reproducible algorithm if there is a function h over $\text{supp}(\mathcal{D})$ such that $\mathbb{P}_{\mathbf{u} \sim \mathcal{D}, s \sim \mathcal{D}_\xi} \{ \mathcal{A}_{\xi=s}(\mathbf{u}) = h(s) \} \geq 1 - \delta$.

Reproducibility asks that with high probability over the draw of an input sample $\mathbf{u} \sim \mathcal{D}$ and of the random seed $s \sim \mathcal{D}_\xi$, the algorithm outputs a deterministic function of the realized seed: $h(s)$. On the other hand, pseudo-determinism asks that for *every* input \mathbf{u} , with high probability over the random draw $s \sim \mathcal{D}_\xi$, the randomized algorithm outputs a deterministic function of the input $h(\mathbf{u})$. Finally, pseudo-independence asks that for *every* input \mathbf{u} , a randomized algorithm should be *almost independent* of one of its random seeds in the sense that with high probability over the draw of $s \sim \xi$, its output is close in total-variation distance to a randomized algorithm which is completely oblivious of s .

2.3. Analysis of the sample-reuse framework

In this section, we give an outline of our proof of Theorem 6. Our discussion here expands on the intuitive explanation in Section 2.1. To reduce notational clutter, throughout this section, we omit the superscript sub from $\mathcal{A}'_{\xi, \chi'}$ and refer to it just as $\mathcal{A}'_{\xi, \chi'}$.

Inducing pseudo-independence in the noisy solver. Recall that in Section 2.1, we said that the first step in our proof of Theorem 6 would be to show that $\mathcal{A}'_{\xi, \chi'}$ in Meta-Algorithm 1 is pseudo-independent of ξ when the noise parameter τ is set appropriately. Concretely, we prove the following theorem in Appendix A.

Theorem 12 *Let $\epsilon, \delta \in (0, 1)$, $\eta > 0$, $\eta' = \min(\eta/2, \eta\epsilon/p)$ and let $\mathcal{A}_{\xi, \chi}$ be a randomized algorithm that is an $(\eta/2, \delta)$ -approximation of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}^p$. Let $\mathcal{A}'_{\xi, \chi'}$ be the randomized algorithm defined as follows. $\mathcal{A}'_{\xi, \chi'}$ takes input $\mathbf{u} \in \mathbb{R}^d$, random seed $\xi \sim \mathcal{D}_\xi$, and χ' where $\chi' = (\chi, \nu) \sim \mathcal{D}_{\chi'}^{\mathbf{u}}$ is the concatenation of an independently drawn seed $\chi \sim \mathcal{D}_\chi^{\mathbf{u}}$ and seed $\nu \sim \mathcal{D}_\nu := \text{Unif}^p(-\tau, \tau)$ for $\tau := \max\left(\frac{\eta'}{2}, \frac{\eta'p}{2\epsilon}\right)$. For any realization s, c, \mathbf{e} of ξ, χ, ν and any $\mathbf{u} \in \mathbb{R}^d$, we define $\mathcal{A}'_{\xi=s, \chi'=(c, \mathbf{e})}(\mathbf{u}) = \mathcal{A}_{\xi=s, \chi=c}(\mathbf{u}) + \mathbf{e}$. Then, $\mathcal{A}'_{\xi, \chi'}$ is an (ϵ, δ) -pseudo-independent of ξ and an (η, δ) -approximation of f with the same runtime and query complexities as $\mathcal{A}_{\xi, \chi}$ up to an additive $O(p)$ in runtime.*

Proof Let $\mathcal{A}_{\xi, \chi}$ be the randomized algorithm which takes input $\mathbf{x} \in \mathbb{R}^d$, random seed $\xi \sim \mathcal{D}_\xi$, and χ where $\chi = (\chi', \nu) \sim \mathcal{D}_\chi^{\mathbf{x}}$ is the concatenation of an independently drawn seed $\chi' \sim \mathcal{D}_{\chi'}^{\mathbf{x}}$ and seed $\nu \sim \mathcal{D}_\nu := \text{Unif}^p(-t, t)$, where we use Unif^p to denote the distribution of a p -dimensional independent uniform random vector, and t is some parameter to be specified later in this proof.

For any realization s, c, \mathbf{e} of ξ, χ', ν , let $\mathcal{A}'_{\xi=s, \chi'=(c, \mathbf{e})}(\mathbf{x}) = \mathcal{A}_{\xi=s, \chi=c}(\mathbf{x}) + \mathbf{e}$. That is, on a given input \mathbf{x} , $\mathcal{A}'_{\xi, \chi'}(\mathbf{x})$ has the same distribution of a random variable which is equal to $\mathcal{A}_{\xi, \chi}(\mathbf{x})$ plus some independent $\text{Unif}(-t, t)$ random noise in each coordinate.

First, we construct a smoothing for $\mathcal{A}'_{\xi, \chi'}$. Let $\bar{\mathcal{A}}_\chi$ be the randomized algorithm which takes input $\mathbf{x} \in \mathbb{R}^d$ and a random seed $\chi \sim \mathcal{D}_\chi^{\mathbf{x}}$. For any realization (c, \mathbf{e}) of $\chi' = (\chi, \nu)$, let $\bar{\mathcal{A}}_{\chi'}(\mathbf{x}) = f(\mathbf{x}) + \nu$. Now, using that $\mathcal{A}_{\xi, \chi}$ is an (η', δ) -approximation of f , we can conclude that

$$\mathbb{P}_{\xi \sim \mathcal{D}_\xi, \chi \sim \mathcal{D}_\chi^{\mathbf{x}}} (\|\mathcal{A}_{\xi, \chi}(\mathbf{x}) - f(\mathbf{x})\|_\infty \leq \eta') \geq 1 - \delta, \quad (1)$$

and hence

$$\mathbb{P}_{\xi \sim \mathcal{D}_\xi, \chi' \sim \mathcal{D}_{\chi'}^{\mathbf{x}}} (\|\mathcal{A}'_{\xi, \chi'}(\mathbf{x}) - f(\mathbf{x})\|_\infty \leq (\eta' + t)) \geq 1 - \delta. \quad (2)$$

Thus, $\mathcal{A}'_{\xi, \chi'}$ is an $(\eta' + t, \delta)$ -approximation of f . In particular, setting $t = \tau = \max\left(\frac{\eta'}{2}, \frac{\eta'p}{2\epsilon}\right)$ and recalling $\eta' = \min\left(\frac{\eta}{2}, \frac{\eta\epsilon}{p}\right)$, we have

$$\eta' + t \leq \eta' \max\left(\frac{3}{2}, \frac{p}{\epsilon}\right) \leq \eta$$

and hence $\mathcal{A}'_{\xi, \chi'}$ is an (η, δ) -approximation of f . To bound the total variation distance, note that by (24), we have that with probability $1 - \delta$, $d_{TV}\left(p_{\mathcal{A}'_{\xi=s, \chi'}(\mathbf{x})}, p_{\bar{\mathcal{A}}_{\chi'}(\mathbf{x})}\right)$ is simply the total variation

distance between q, q' where q is the distribution of a $\text{Unif}^p[-t, t]$ random variable and q' is the distribution of a random variable whose i -th entry is distributed as $\text{Unif}[-\alpha_i, \alpha_i]$ for some $|\alpha_i| \leq \eta'$. Now, for $t \geq \eta'/2$, we have

$$d_{TV}(q, q') = 1 - \prod_{i \in [p]} \left(\frac{\max(0, 2t - |\alpha_i|)}{2t} \right) \leq 1 - \left(1 - \frac{\eta'}{2t} \right)^p \leq \frac{\eta' p}{2t},$$

where the final step used Bernoulli's inequality. Thus, setting $t = \tau = \max\left(\frac{\eta'}{2}, \frac{\eta' p}{2\epsilon}\right)$, we have that with probability $1 - \delta$,

$$\mathbb{P}_{s \sim \mathcal{D}_\chi} \left\{ d_{TV} \left(p_{A'_{\xi=s, \chi'}}(x), p_{\bar{A}_\chi(x)} \right) \leq \epsilon \right\} \geq 1 - \delta.$$

For the second guarantee, note that $\mathcal{A}'_{\xi, \chi'}$ has the same runtime and query complexities as $\mathcal{A}_{\xi, \chi}$ up to an additive $O(p)$ increase in the runtime due to the cost of adding the p -dimensional uniform random noise induced by ν . \blacksquare

We discuss implications to numerical stability further in Appendix A.

Theorem 12 show how to convert standard high-accuracy structured optimization algorithms into *pseudo-independent* high-accuracy structured optimization algorithms (Section 2) by adding noise to the output. Next, we will prove these pseudo-independent versions are amenable to *reusing* samples across multiple iterations of an outer-solver (Figure 1.) We briefly remark that this technique of adding noise to an algorithm's output also appears in differential privacy (Asi and Duchi, 2020; Ghazi et al., 2022) and in the design and analysis of algorithms that are robust to adaptive adversaries (Chen et al., 2022; Lee and Sidford, 2015; Beimel et al., 2022; van den Brand et al., 2022).

Analyzing repeated compositions of pseudo-independent functions Once we know that $\mathcal{A}'_{\xi, \chi'}$ in Meta-Algorithm 1 is *pseudo-independent* of ξ for appropriately chosen τ , we move on to proving Theorem 6.

Recall that, as outlined in Section 2, our goal will be to show that if $\mathcal{A}_{\xi, \chi'}$ in Meta-Algorithm 1 is (ϵ, δ) -pseudo-independent of ξ , then the distribution over outputs of $\text{Outer}(X; \text{Noisy}, \tau)$ and $\text{Outer}(X; \text{Reuse}, \tau)$ are close in TV. To prove this, first observe that NoisyLoop_τ repeatedly composes ζ with $\mathcal{A}'_{\xi, \chi'}$. Definition 13 introduces notation for these compositions.

In what follows, Eq. (3) represents an algorithm that repeatedly applies T iterations, where in each iteration ξ, χ are fresh random variables drawn from their respective distributions (as in NoisyLoop_τ). Eq. (4) is the analogous algorithm where in the i -th iteration a fresh χ' is drawn from $\mathcal{D}_{\chi'}^u$, but $\xi = s$ is *reused* across iterations (as in ReuseLoop_τ .) Eq. (5) is the analogous expression to (3) with the randomized algorithm $\bar{\mathcal{A}}_{\chi'}$, which takes only *one* source of randomness, χ' .

Definition 13 (Composition of randomized algorithms) Let $\mathcal{A}'_{\xi, \chi'}$ be a randomized algorithm which takes an input $\mathbf{u} \in \mathbb{R}^d$ and random seeds $\xi \sim \mathcal{D}_\xi, \chi' \sim \mathcal{D}_{\chi'}^x$ and outputs a point in \mathbb{R}^p . Let $\bar{\mathcal{A}}_{\chi'}$ be a randomized algorithm that takes an input $\mathbf{u} \in \mathbb{R}^d$ and one random seed $\chi' \sim \mathcal{D}_{\chi'}^x$ and outputs a point in \mathbb{R}^p . Let $\zeta : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}^d$ be a deterministic function. For $T \geq 1$, let:

$$\begin{aligned} \Phi_{\mathcal{A}'}^1(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}) &:= \zeta(\mathbf{u}, \mathcal{A}'_{\xi, \chi'}(\mathbf{u})), \\ \Phi_{\mathcal{A}'}^{T+1}(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}) &:= \zeta\left(\Phi_{\mathcal{A}'}^T(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}), \mathcal{A}'_{\xi, \chi'}(\Phi_{\mathcal{A}'}^T(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}))\right) \end{aligned} \quad (3)$$

$$\begin{aligned}\Phi_{\mathcal{A}'}^1(\mathbf{u}; s, D_{\mathcal{X}'}) &:= \zeta(\mathbf{u}, \mathcal{A}'_{\xi=s, \mathcal{X}'}(\mathbf{u})), \\ \Phi_{\mathcal{A}'}^{T+1}(\mathbf{u}; s, D_{\mathcal{X}'}) &:= \zeta\left(\Phi_{\mathcal{A}'}^T(\mathbf{u}; s, D_{\mathcal{X}'}), \mathcal{A}'_{\xi=s, \mathcal{X}'}(\Phi_{\mathcal{A}'}^T(\mathbf{u}; s, D_{\mathcal{X}'}))\right)\end{aligned}\quad (4)$$

$$\begin{aligned}H_{\bar{\mathcal{A}}}^1(\mathbf{u}; D_{\mathcal{X}'}) &:= \zeta(\bar{\mathcal{A}}_{\mathcal{X}'}(\mathbf{u})), \\ H_{\bar{\mathcal{A}}}^{T+1}(\mathbf{u}; D_{\mathcal{X}'}) &:= \zeta\left(\Phi_{\bar{\mathcal{A}}}^T(\mathbf{u}; D_{\mathcal{X}'}), \bar{\mathcal{A}}_{\mathcal{X}'}(\Phi_{\bar{\mathcal{A}}}^T(\mathbf{u}; D_{\mathcal{X}'}))\right)\end{aligned}\quad (5)$$

In Appendix B, we show the following general theorem about pseudo-independence and repeated composition.

Lemma 14 *Let $\mathcal{A}'_{\xi, \mathcal{X}'}$ be randomized algorithm which takes an input $\mathbf{u} \in \mathbb{R}^d$ and two random seeds $\xi \sim \mathcal{D}_\xi, \mathcal{X}' \sim \mathcal{D}_{\mathcal{X}'}$ and is (ϵ, δ) -pseudo-independent of ξ . Then,*

$$d_{TV}(p_{\Phi_{\mathcal{A}'}^T(\mathbf{u}; s, D_{\mathcal{X}'})}, p_{\Phi_{\mathcal{A}'}^T(\mathbf{u}; \mathcal{D}_\xi, D_{\mathcal{X}'})}) \leq 2T(\delta + \epsilon).$$

Finally, to prove Theorem 6, we can combine the previous results.

Theorem 6 *Suppose $\text{Outer}(X; \text{Standard})$ is (η, β) -robust with respect to f^{sub} and $\delta \in (0, 1)$. Let $\eta' := \min(\eta/2, \eta\delta)$. Suppose $\mathcal{A}'_{\xi, \mathcal{X}'}$ is an (η', δ) -approximation of f^{sub} . Then we have that wp. $1 - 5n_{\text{outer}}^2\delta$, $\text{Outer}(X; \text{Reuse}, \eta'/(2\delta))$ outputs a β -accurate solution to X .*

Proof For notational convenience, set $\epsilon = \delta$ and $\tau = \eta'/(2\epsilon)$. Notice that

$$\text{Outer}(X; \text{Noisy}, \tau) \stackrel{\mathcal{D}}{=} \sum_{t \in [n_{\text{outer}}]} w(t) \cdot \Phi_{\mathcal{A}'}^t(\mathbf{u}_0; \mathcal{D}_\xi, D_{\mathcal{X}'}), \quad (6)$$

$$\text{Outer}(X; \text{Reuse}, \tau) \stackrel{\mathcal{D}}{=} \sum_{t \in [n_{\text{outer}}]} w(t) \cdot \Phi_{\mathcal{A}'}^t(\mathbf{u}_0; s_1, D_{\mathcal{X}'}), \quad (7)$$

where $s_1 \sim \mathcal{D}_\xi$. Since $\tau \leq \eta/2$ and $\eta' \leq \eta/2$, by Lemma 3, (6) is a β -accurate solution to X wp. $1 - n_{\text{outer}}\delta$.

Also, by Theorem 12, $\mathcal{A}'_{\xi, \mathcal{X}'}$ is (ϵ, δ) -pseudo-independent of ξ . Thus, by Lemma 14, the TV distance between the t -th summand of (6) and the t -th summand of (7) is bounded by $2t(\delta + \epsilon)$.

As (6) is β -accurate wp. $1 - n_{\text{outer}}\delta$, we can apply Fact 7 and take union bound over all $t \in [n_{\text{outer}}]$ to conclude that wp. $1 - 2n_{\text{outer}}^2(\delta + \epsilon) - n_{\text{outer}}\delta \geq 1 - 5n_{\text{outer}}^2\delta$, (7) is also β -accurate. ■

3. Application: Finite-sum minimization

In this section, we apply our sample reuse framework to obtain improved full batch versus sample query trade-offs for finite sum minimization (FSM). As a special case, we discuss implications for generalized linear models.

We focus on this setting for FSM due to its simplicity as it is perhaps the most illustrative and foundational setting. We consider a more general variant of FSM in Appendix 6, where we obtain rates dependent on the (potentially non-uniform) smoothness L_i of each f_i .

FSM arises in many machine learning settings, such as *empirical risk minimization*, where each f_i might be a loss function for a sampled data point $i \in [n]$ and the goal is to minimize the average

loss across all the data. In this context, a query \mathbf{x} to a gradient oracle for F corresponds to making a pass over the *full batch* of data $i \in [n]$ to compute the gradient at \mathbf{x} . Meanwhile, querying a component oracle with i only requires accessing information pertaining to the i -th data point.

Below, we formally define the problem and the batch-sample oracle model and provide a brief sketch of how we obtain our result in Table 2 for FSM (Section 3).

Definition 15 (FSM problem) *In the FSM problem we are given $c > 1$ and $\mathbf{x}_0 \in \mathbb{R}^d$ and must output $\hat{\mathbf{x}} \in \mathbb{R}^d$ such that $F(\hat{\mathbf{x}}) - \min_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x}) \leq 1/c \cdot (F(\mathbf{x}_0) - \min_{\mathbf{z} \in \mathbb{R}^d} F(\mathbf{z}))$ where $F(\mathbf{x}) := \frac{1}{n} \sum_{i \in [n]} f_i(\mathbf{x})$, F is μ strongly-convex, and each $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is L -smooth and convex.*

Definition 16 (Gradient oracle - FSM batch oracle) *When queried with $\mathbf{x} \in \mathbb{R}^d$, a gradient oracle for differentiable $F : \mathbb{R}^d \rightarrow \mathbb{R}$ returns $\nabla F(\mathbf{x}) \in \mathbb{R}^d$.*

Definition 17 (Component oracle - FSM sample oracle) *When queried with $i \in [n]$, a component oracle for $F(\mathbf{x}) = \frac{1}{n} \sum_{i \in [n]} f_i(\mathbf{x})$ for differentiable $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ returns a gradient oracle for f_i .*

The standard outer-solver and sub-solver. As discussed in Section 1.1, state-of-the-art query complexities for FSM can be achieved by with accelerated proximal point/Catalyst (APP) as an outer-solver (Frostig et al., 2015; Lin et al., 2015), ℓ_2 -regularized FSM for the sub-problems, and stochastic variance-reduced gradient descent (SVRG) as a sub-solver (Johnson and Zhang, 2013). Formally, each iteration of APP approximately solves a λ -regularized *sub-problem*, as follows.

Definition 18 (FSM sub-problem) *Let F be as in Definition 15 and $\lambda \geq \mu$. For any $\mathbf{u} = (\mathbf{x}, \mathbf{v}) \in \mathbb{R}^d \times \mathbb{R}^d$, let $\rho = (\mu + 2\lambda)/\mu$ and $\mathbf{y}_{\mathbf{u}} := 1/(1 + \rho^{-1/2})\mathbf{x} + \rho^{-1/2}/(1 + \rho^{-1/2})\mathbf{v}$. We define*

$$f^{\text{sub}}(\mathbf{u}) = \operatorname{argmin}_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2.$$

Moreover, we say that $\mathbf{a}' = (\mathbf{x}', \mathbf{v}')$ is a solution to the (\mathbf{u}, λ, c) -sub-problem if

$$F(\mathbf{x}') - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2 \leq \frac{1}{c} \left(F(\mathbf{y}_{\mathbf{u}}) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2 \right).$$

APP uses the following post-process (recall Meta-Algorithm 1). This is essentially intended to implement *acceleration* (sometimes called *momentum*) on the iterates.

Definition 19 (FSM post-process) *Let F be as in Definition 15 and $\lambda > 0$. For any $\mathbf{u} = (\mathbf{x}, \mathbf{v}) \in \mathbb{R}^d \times \mathbb{R}^d$ and $\mathbf{u}' = (\mathbf{x}', \mathbf{v}') \in \mathbb{R}^d \times \mathbb{R}^d$, let $\rho = (\mu + 2\lambda)/\mu$ and $\iota = 2/\mu + 1/\lambda$, and define*

$$\zeta(\mathbf{u}, \mathbf{u}') := (1 - \rho^{-1/2})\mathbf{v} + \rho^{-1/2} (\mathbf{y}_{\mathbf{x}, \mathbf{v}} - \iota\lambda(\mathbf{y}_{\mathbf{x}, \mathbf{v}} - \mathbf{x}')).$$

Now, APP is known to solve the original problem, given approximate solutions to these sub-problems, in the following sense.

Theorem 20 (APP, Theorem 1.1 of Frostig et al. (2015), restated - FSM outer-solver) *Let F be a μ -strongly-convex function and $\lambda \geq \mu$. There is a $c' = \text{poly}(\lambda, \mu, c, d)$ such that the following holds. Suppose that in each iteration $t \in [n_{\text{outer}}]$ of Algorithm 2, $\mathbf{x}_{t-1/2}$ is a solution to the $(\mathbf{u}_{t-1}, \lambda, c')$ -sub-problem; then $\mathbf{u}_{n_{\text{outer}}}$ is a solution to the FSM problem.*

Algorithm 2: APP-SVRG $_{\lambda,\delta}(\mathbf{x}_0, F, c)$ Pseudocode

Input: Initial point $\mathbf{x}_0 \in \mathbb{R}^d$, gradient oracle and component oracle for F , error factor c
Parameters: Failure probability δ , and $\lambda \geq \mu$.
 Initialize $\mathbf{v}_0 \leftarrow \mathbf{0} \in \mathbb{R}^d$
 Initialize $\mathbf{u}_0 \leftarrow (\mathbf{x}_0, \mathbf{v}_0) \in \mathbb{R}^d \times \mathbb{R}^d$
 Set sufficiently large $n_{\text{outer}} = \tilde{O}(\sqrt{\lambda/\mu})$
 Set sufficiently large $c' = \text{poly}(\lambda, \mu, c, d)$
for each $t \in [n_{\text{outer}}]$ **do**
 // We draw realizations of the random seed ξ according to the following distribution.
 Draw $s_t := \{i_1, \dots, i_T\} \sim \mathcal{D}_\xi$ where each $i_j \sim \text{Unif}[n]$ for sufficiently large $T = \tilde{O}(L/\lambda)$
 // We include χ , a zero-bit random bit for technical consistency with Section 2
 Draw $c_t \sim \mathcal{D}_\chi := \text{Unif}[0]$
 $\mathbf{u}_{t-1/2} = (\mathbf{x}_{t-1/2}, \mathbf{v}_{t-1/2}) \leftarrow \mathcal{A}_{\xi=s_t, \chi=c_t}^{\text{SVRG}|\lambda, c', \delta/n_{\text{outer}}}(\mathbf{u}_{t-1})$.
 // The post-process implements momentum, as described in Definition 19
 $\mathbf{u}_t = (\mathbf{x}_t, \mathbf{v}_t) \leftarrow \zeta(\mathbf{u}_t, \mathbf{u}_{t-1/2})$
end
return: $\mathbf{x}_{n_{\text{outer}}}$

In particular, SVRG is a commonly used *sub-solver* to solve the sub-problems required in APP (Theorem 20) efficiently. We restate this result below.

Theorem 21 (SVRG, Theorem 2.2 of Frostig et al. (2015), restated - FSM sub-problem solver)

Let F be as in Definition 15 and $\lambda \geq \mu$. There is a randomized algorithm $\mathcal{A}_{\xi, \chi}^{\text{SVRG}|\lambda, c, \delta}$ such that the following holds.

- The algorithm takes in the random seeds ξ, χ distributed as follows. The first random seed $\xi \sim \{i_1, \dots, i_T\}$ where each $i_j \sim \text{Unif}[n]$ for some $T = \tilde{O}(L/\lambda)$. The second random seed $\chi \sim \text{Unif}[0]$ is a 0-bit random seed.
- If $\mathbf{x}' = \mathcal{A}_{\xi, \chi}^{\text{SVRG}|\lambda, c, \delta}(\mathbf{u})$, then w.p. $1 - \delta$ over the draw of the random seeds ξ and χ , \mathbf{x}' is a solution to the (\mathbf{u}, λ, c) -sub-problem.
- $\mathcal{A}_{\xi, \chi}^{\text{SVRG}|\lambda, c, \delta}$ makes $\tilde{O}(1)$ batch queries and makes sample queries only on the indices contained within ξ .

By combining Theorem 20 (APP) with Theorem 21 (SVRG) and taking a union bound over all n_{outer} iterations in Algorithm 2, we obtain an algorithm for solving FSM problems. This algorithm obtains a trade-off of $\tilde{O}(\sqrt{\lambda/\mu})$ full batch (gradient oracle) queries and $\tilde{O}(L/\sqrt{\lambda\mu})$ sample (component oracle) queries for any $\lambda \geq \mu$, as reported in the ‘‘Prior Work’’ column of Table 2.

However, using our sample reuse framework from Section 2, we can improve this trade-off!

The sample reusing sub-solver Observe that Algorithm 2 *exactly* fits into the Meta-Algorithm 1 framework from Section 2. In particular, Algorithm 2 implements $\text{Outer}(X, \text{Standard})$ where $X = (\mathbf{x}_0, F, c)$ is an FSM problem instance as per Definition 15. To show that we can replace the standard sub-solver with the sample-reusing sub-solver (i.e., $\text{Outer}(X, \text{Reuse})$), we need to invoke Theorem 6. To do so, we need to show (1) that APP satisfies robustness with respect to f^{sub}

in the sense of Definition 2 and (2) that SVRG can compute an approximation in the ℓ_∞ norm, as per Definition 1. We prove this below

Lemma 22 (FSM outer-solver robustness) *Let F be as in Definition 15 and $\lambda > 0$. There is a $c' = \text{poly}(\lambda, \mu, c, d)$ such that the following holds. Suppose that in each iteration $t \in [n_{\text{outer}}]$ of the algorithm,*

$$\left\| \mathbf{u}_{t-1/2} - f^{\text{sub}}(\mathbf{y}_{\mathbf{u}_{t-1}}) \right\|_\infty^2 \leq \frac{1}{c'} \left(F(\mathbf{y}_{\mathbf{u}_{t-1}}) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}_{t-1}}\|_2^2 \right),$$

Then $\mathbf{u}_{n_{\text{outer}}}$ is a solution to the FSM problem.

Proof We prove the lemma by invoking Theorem 20. That is, it suffices to show that for any c, λ , there exists a $c' = \text{poly}(\lambda, \mu, c, d)$ such that whenever \mathbf{u}' satisfies

$$\left\| \mathbf{u}' - f^{\text{sub}}(\mathbf{y}_{\mathbf{u}}) \right\|_\infty^2 \leq \frac{1}{c'} \left(F(\mathbf{y}_{\mathbf{u}}) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2 \right), \quad (8)$$

we have that

$$F(\mathbf{u}') - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2 \leq \frac{1}{c} \cdot \left(F(\mathbf{y}_{\mathbf{u}}) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2 \right). \quad (9)$$

To prove this we will use the smoothness of the function $F_{\lambda, \mathbf{u}}$, which is defined as follows:

$$F_{\lambda, \mathbf{u}}(\tilde{\mathbf{x}}) := F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2.$$

First, we have that (8) implies

$$\left\| \mathbf{u}' - f^{\text{sub}}(\mathbf{y}_{\mathbf{u}}) \right\|_2^2 \leq d \left\| \mathbf{u}' - f^{\text{sub}}(\mathbf{y}_{\mathbf{u}}) \right\|_\infty^2 \leq \frac{2d}{c'} \left(F(\mathbf{y}_{\mathbf{u}}) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2 \right).$$

Thus, by L -smoothness, whenever $c' > cd(L + \lambda)$, (9) holds. The result follows by Theorem 20. ■

Lemma 23 (FSM sub-problem solver high-precision) *Let F be as in Definition 15. There is a randomized algorithm $\mathcal{A}_{\xi, \chi}^{\text{SVRG-HP}|\lambda, c, \delta}$ such that the following holds.*

- The algorithm takes in the random seeds ξ, χ distributed as follows. The first random seed $\xi \sim \{i_1, \dots, i_T\}$ where each $i_j \sim \text{Unif}[n]$ and $T = \tilde{O}(L/\lambda)$. The second random seed $\chi \sim \text{Unif}[0]$ is a 0-bit random seed.
- If $\mathbf{x}' = \mathcal{A}_{\xi, \chi}^{\text{SVRG-HP}|\lambda, c, \delta}(\mathbf{u})$, then wp. $1 - \delta$ over the draw of the random seeds ξ and χ ,

$$\left\| \mathbf{x}' - f^{\text{sub}}(\mathbf{y}_{\mathbf{u}}) \right\|_\infty^2 \leq \frac{1}{c} \cdot \left(F(\mathbf{y}_{\mathbf{u}}) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2 \right).$$

- $\mathcal{A}_{\xi, \chi}^{\text{SVRG-HP}|\lambda, c, \delta}$ makes $\tilde{O}(1)$ batch queries and makes sample queries only on indices in ξ .

Proof We prove the lemma by invoking Theorem 21. It is enough to show that there is a $c' = \text{poly}(\lambda, \mu, c, d)$ such that whenever \mathbf{x}' satisfies

$$F(\mathbf{x}') - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_u\|_2^2 \leq \frac{1}{c'} \cdot \left(F(\mathbf{y}_u) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_u\|_2^2 \right),$$

we also have that

$$\left\| \mathbf{x}' - f^{\text{sub}}(\mathbf{y}_u) \right\|_\infty^2 \leq \frac{1}{c} \cdot \left(F(\mathbf{y}_u) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_u\|_2^2 \right),$$

We will use the strong convexity of the function $F_{\lambda, \mathbf{u}}$, which is defined as follows:

$$F_{\lambda, \mathbf{x}}(\tilde{\mathbf{x}}) := F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{x}\|_2^2.$$

Now, by μ -strong convexity,

$$\left\| \mathbf{x}' - f^{\text{sub}}(\mathbf{y}_u) \right\|_\infty^2 \leq \left\| \mathbf{x}' - f^{\text{sub}}(\mathbf{y}_u) \right\|_2^2 \leq \frac{\mu}{2c} \cdot \left(F(\mathbf{y}_u) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_u\|_2^2 \right),$$

thus it suffices to set $c' = c\mu/2$. ■

Algorithm 3: APP-SVRG-Reuse $_{\lambda, \delta}(\mathbf{z}_0, F, c)$ Pseudocode

Input: Initial point $\mathbf{z}_0 \in \mathbb{R}^d$, gradient oracle and component oracle for F , error factor c

Parameters: Failure probability δ , and $\lambda \geq \mu$.

Initialize $\mathbf{u}_0 \leftarrow (\mathbf{x}_0, \mathbf{0}) \in \mathbb{R}^d \times \mathbb{R}^d$

Set sufficiently large $n_{\text{outer}} = \tilde{O}(\sqrt{\lambda/\mu})$

// We draw a realization of the random seed ξ according to the following distribution.

Draw $s_1 := \{i_1, \dots, i_T\} \sim \mathcal{D}_\xi$ where each $i_j \sim \text{Unif}[n]$ for sufficiently large $T = \tilde{O}(L/\lambda)$

for each $t \in [n_{\text{outer}}]$ **do**

// Draw c'_t from the noisy distribution $\mathcal{D}_{\mathcal{X}'}$ as in Meta-Algorithm 1.

Draw $c'_t \sim \mathcal{D}_{\mathcal{X}'}$

// Implement the noisy analog of $\mathcal{A}^{\text{SVRG-HP}|\lambda, \text{poly}(\lambda, \mu, c), \delta/(5n_{\text{outer}}^2)}$ as in Meta-Algorithm 1.

$\mathbf{u}_{t-1/2} = (\mathbf{x}_{t-1/2}, \mathbf{v}_{t-1/2}) \leftarrow \mathcal{A}'_{\xi=s_1, \mathcal{X}'=c'_t}{}^{\text{SVRG-HP}|\lambda, \text{poly}(\lambda, \mu, c), \delta/(5n_{\text{outer}}^2)}(\mathbf{u}_{t-1})$.

// The post-process implements momentum, as described in Definition 19

$\mathbf{u}_t = (\mathbf{x}_t, \mathbf{v}_t) \leftarrow \zeta(\mathbf{u}_t, \mathbf{u}_{t-1/2})$

end

return: $\mathbf{x}_{n_{\text{outer}}}$

By combining Lemma 23 with Lemma 22 and applying Theorem 6, we immediately obtain our improved trade-off.

Theorem 24 (FSM improvement) *There is an algorithm (Algorithm 3) that for F as in Definition 15, $\lambda \geq \mu$, and $\delta \in (0, 1)$, makes $\tilde{O}(\sqrt{\lambda/\mu})$ -batch queries and $\tilde{O}(L/\lambda)$ -sample queries and solves the FSM problem wp. $1 - \delta$.*

Proof We apply Theorem 6 using $\mathcal{A}'_{\xi, \mathcal{X}'}{}^{\text{SVRG-HP}|\lambda, \delta/(5n_{\text{outer}}^2)}$ as the sub-solver to solve the $(\mathbf{u}_{t-1}, \lambda, c)$ sub-problem in each iteration, where c' is as required by Lemma 22. The lower failure probability $\delta/(5n_{\text{outer}}^2)$ is used in order to counter the blowup in failure probability in Theorem 6. ■

3.1. Applications to regression with generalized linear models (GLM)

Regression with generalized linear models is a special case of FSM where we have n data vectors $\{\mathbf{a}_i\}_{i=1}^n \in \mathbb{R}^d$ with n corresponding, *explicitly known*, labels $\{b_i\}_{i=1}^n \in \mathbb{R}$, and $f_i = \phi_i(\mathbf{a}_i^\top \mathbf{x})$ for some explicit function $\phi_i(z)$ (e.g., for least-squares regression, $\phi_i(z) = 1/2(z - b_i)^2$ and for logistic regression, $\phi_i(z) = \log(1 + \exp(-zb_i))$).

Because the gradient mapping $g : z \mapsto \nabla \phi_i(z)$ is known explicitly, component-oracles for f_i are easily implementable: we can simply query $i \sim \mathbf{p}$, lookup \mathbf{a}_i , and return the function $\nabla \phi_i(\mathbf{a}_i^\top \mathbf{x}) = \mathbf{a}_i \cdot g(\mathbf{a}_i^\top \mathbf{x})$ explicitly for *any* future point \mathbf{x} .

So, in these settings, our improved sample query complexity corresponds to an improved trade-off between the number of full passes over $\{\mathbf{a}_i\}_{i=1}^n$ and the number of random samples $\mathbf{a}_{i \sim \text{Unif}[n]}$ required to train a GLM. The implications of this result are two-fold:

- First, this improvement means that our results prove that problems such as fitting a generalized linear model can be solved with *less information* than was known previously.
- Second, our results show that one can *reuse* the same samples $\mathbf{a}_{i \sim \text{Unif}[n]}$ across all iterations of the outer solver. In distributed memory settings or in settings where caching significantly affects computational costs, this ability to reuse the same cached samples might be beneficial in reducing memory retrieval or communication costs.

3.2. Extension to non-uniform smoothness

Our method can also be extended to non-uniformly smooth f_i , where we relax the assumption that each f_i is L -smooth in Definition 15 to assume that each f_i is L_i -smooth. In this setting, can develop improved trade-offs that depend on the distribution of the L_i 's rather than the worst-case smoothness $\max_i L_i$ by instantiating APP with the primal-dual FSM sub-problem solver of Jin et al. (2022). We defer a discussion of this setting to Section 6.

A note on pseudocode in the remainder of this paper. We include the pseudocode Algorithm 2 and Algorithm 3 in this section in order to provide a clear example of a concrete instantiation of Meta-Algorithm 1 for FSM.

However, for the sake of brevity, in later appendix sections, we will not rewrite the instantiation of Meta-Algorithm 1 for each application. Instead, we will simply define the initializations; setting for n_{outer} ; post-processing functions; sub-solvers, distributions, and target functions f^{sub} —as this fully characterizes the instantiations of Meta-Algorithm 1 for each application. Additionally, we include thorough references to related work which contains application-specific pseudocodes for the outer-solvers and sub-solvers related to each application.

An exception is in the case of discounted MDPs (Section 4.2), where we *will* include pseudocode. This is because our outer-solver for DMDPs is a new contribution of our work, so we feel it is helpful to include the full pseudocode for completeness.

4. Application: Infinite-horizon Markov Decision Processes (MDPs)

In this section, we discuss how our framework can be applied to infinite-horizon Markov Decision Processes (MDPs). We primarily focus on the discounted case (DMDPs) in Section 4.2 and then in Section 4.3 extend our results to the average-reward case (AMDPs) using a standard reduction due to Jin and Sidford (2021). We begin with preliminaries in Section 4.1.

4.1. Preliminaries

We denote an MDP by $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbf{P})$ where \mathcal{S} denotes a finite state space, \mathcal{A} denotes the total set of all state-action pairs. Concretely, we use \mathcal{A}_s to denote the set of available actions in state $s \in \mathcal{S}$ and define $\mathcal{A} = \{(s, a) : s \in \mathcal{S}, a \in \mathcal{A}_s\}$ as well as $\mathcal{A}_{\text{tot}} := |\mathcal{A}|$. The state-action-state transition matrix is given by $\mathbf{P} \in \mathbb{R}^{\mathcal{A} \times \mathcal{S}}$. We use $\mathbf{p}(s, a) \in \Delta^{\mathcal{S}}$ to denote the (s, a) -th row of \mathbf{P} . A policy π is a mapping from states to actions, i.e., $\pi : s \mapsto \pi(s) \in \mathcal{A}_s$. We use Π to denote the set of all possible policies.

Throughout this section, for vectors \mathbf{a}, \mathbf{b} , we use $\mathbf{a} \geq \mathbf{b}$ to mean entrywise inequality (and use $\leq, >, <$ analogously.) For \mathcal{A}_{tot} -dimensional vectors, say $\mathbf{r} \in \mathbb{R}_{\text{tot}}^{\mathcal{A}}$, we use the notation $\mathbf{r}(s, a)$ to denote the (s, a) -th entry of \mathbf{r} . We assume rewards are known a-priori, as in prior works in this setting (see, e.g., Jin et al. (2024) and references therein.)

For any policy π , we use $\mathbf{r}^\pi \in \mathbb{R}_{\geq 0}^{\mathcal{S}}$ to denote the $|\mathcal{S}|$ -dimensional vector with s -th entry given by $\mathbf{r}^\pi(s, \pi(s))$. Likewise, we use $\mathbf{P}^\pi \in \mathbb{R}^{\mathcal{S} \times \mathcal{S}}$ to denote the sub-matrix of \mathbf{P} where the s -th row of \mathbf{P}^π is $\mathbf{p}(s, \pi(s))$.

An agent interacts with the MDP in *timestep* as follows. At each timestep $t \geq 0$, an agent begins in state s_t , chooses an available action $a_t \in \mathcal{A}_{s_t}$, and collects a (finite) reward $\mathbf{r}(s_t, a_t)$. The agent then (stochastically) transitions to the next state s_{t+1} where $s_{t+1} \sim \mathbf{p}(s_t, a_t)$. The agent’s goal is to compute a *policy* for selecting actions in each state that maximizes the agent’s infinite-horizon expected utility over the set of all policies, denoted Π .

This objective is often formalized in two settings: the discounted setting (DMDP) and the average-reward setting (AMDP), which we will discuss in Section 4.2 and Section 4.3 respectively. However, we will first define the full batch and sample oracle access for \mathcal{M} .

Definition 25 (Matrix-vector oracle - DMDP/AMDP batch oracle) When queried with $\mathbf{x} \in \mathbb{R}^{\mathcal{S}}$, a matrix-vector oracle for \mathbf{P} returns $\mathbf{P}\mathbf{x}$.

Definition 26 (Simulator oracle - DMDP/AMDP sample oracle) When queried with $(s, a) \in \mathcal{A}$, a simulator oracle returns a sample $s' \sim \mathbf{p}(s, a)$.

The oracle described in Definition 26 is often called a *generative model* in the reinforcement learning theory literature (see, for example, Kearns and Singh (1998)).

4.2. Discounted MDP

We begin by describing MDPs in the *discounted setting*.

Discounted MDP (DMDP). In a DMDP, the objective is to compute an ϵ -optimal policy for maximizing the infinite-horizon γ -discounted reward for some known constant $\gamma \in (0, 1)$.

Definition 27 (Discounted MDP (DMDP)) Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbf{P})$ be an MDP, $\gamma \in (0, 1)$, and $\mathbf{r} \in [0, 1]^{\mathcal{A}}$. We denote the associated γ -discounted-MDP by $\mathcal{M}_{\gamma, \mathbf{r}} = (\mathcal{M}; \gamma, \mathbf{r})$.

In a DMDP, the value of a policy π is defined as follows.

Definition 28 (DMDP policy value) Let $\mathcal{M}_{\gamma,r}$ be a DMDP and π be a policy. The value of policy π , denoted $\mathbf{v}_{\gamma,r}^\pi$, is defined as

$$\mathbf{v}_{\gamma,r}^\pi(s) := \mathbb{E} \left[\sum_{t \geq 0} \gamma^t \mathbf{r}(s_t, \pi(s_t)) \mid s_0 = s \right].$$

We can alternatively define the value of a policy in terms of the Bellman operator.

Definition 29 (Bellman operator) Let $\mathcal{M}_{\gamma,r}$ be a DMDP. Given a policy π , and vector $\mathbf{v} \in \mathbb{R}^S$ we define the Bellman operator associated with π as $\mathcal{T}_{\gamma,r}^\pi[\mathbf{v}] \in \mathbb{R}^S$ as

$$\mathcal{T}_{\gamma,r}^\pi[\mathbf{v}](s) := \mathbf{r}(s, \pi(s)) + \gamma \mathbf{p}(s, \pi(s))^\top (s, \pi(s)).$$

The value of policy π is the unique vector such that $\mathbf{v}_{\gamma,r}^\pi = \mathcal{T}_{\gamma,r}^\pi[\mathbf{v}_{\gamma,r}^\pi]$. We also define the Bellman operator $\mathcal{T}_{\gamma,r}$ to be the operator that maps $\mathbf{v} \mapsto \mathcal{T}_{\gamma,r}[\mathbf{v}] \in \mathbb{R}^S$ where

$$\mathcal{T}_{\gamma,r}[\mathbf{v}](s) := \max_{a \in \mathcal{A}_s} \mathbf{r}(s, a) + \gamma \mathbf{p}(s, a)^\top \mathbf{v}.$$

Fact 30 (Bellman optimality conditions) The value of the optimal policy $\pi_{\gamma,r}^*$ for $\mathcal{M}_{\gamma,r}$ satisfies

$$\mathbf{v}_{\gamma,r}^{\pi_{\gamma,r}^*}(s) = \max_{\pi \in \Pi} \mathbb{E} \left[\sum_{t \geq 0} \gamma^t \mathbf{r}(s_t, \pi(s_t)) \mid s_0 = s \right],$$

and $\mathbf{v}_{\gamma,r}^{\pi_{\gamma,r}^*}$ is the unique vector such that $\mathbf{v}_{\gamma,r}^{\pi_{\gamma,r}^*} = \mathcal{T}_{\gamma,r}[\mathbf{v}_{\gamma,r}^{\pi_{\gamma,r}^*}]$. We use the shorthand $\mathbf{v}_{\gamma,r}^* = \mathbf{v}_{\gamma,r}^{\pi_{\gamma,r}^*}$ and refer to it as the optimal value vector for $\mathcal{M}_{\gamma,r}$.

Equipped with these definitions of values in DMDPs, we define the DMDP problem as follows.

Definition 31 (DMDP problem) Let $\mathcal{M}_{\gamma,r}$ be a γ -discounted MDP and $\epsilon \in (0, (1 - \gamma)^{-1}]$ be given. We must compute a value \mathbf{v} and a policy π such that

$$\mathbf{0} \leq \mathbf{v}_{\gamma,r}^* - \mathbf{v} \leq \epsilon \mathbf{1}, \quad \text{and} \quad \mathbf{v}_{\gamma,r}^* - \epsilon \mathbf{1} \leq \mathbf{v}_{\gamma,r}^\pi \leq \mathbf{v}_{\gamma,r}^*.$$

Such a policy π and value \mathbf{v} is called an ϵ -optimal policy and ϵ -optimal value for $\mathcal{M}_{\gamma,r}$.

The standard outer-solver and sub-solver. Now, we will show how to reduce the DMDP problem for a discount factor $\gamma > 0$ to solving a sequence of sub-problems. In this case, the sub-problem will essentially require solving a DMDP with a *smaller* discount factor γ' . Note that in the sub-problems, we relax the requirement from $\mathbf{r} \in [0, 1]^A$ to simply $\mathbf{r} \in \mathbb{R}_{\geq 0}^A$.

Definition 32 (DMDP sub-problem) Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbf{P})$ be an MDP and $\mathbf{r} \in \mathbb{R}_{\geq 0}^A$ be a reward vector. In the $(\mathcal{M}, \mathbf{r}, \gamma', \mathbf{v}, \epsilon)$ -sub-problem, we are given $\gamma' > 0$, $\mathbf{v} \in \mathbb{R}^S$ and $\epsilon \in (0, 1/(1 - \gamma'))$. We define $\mathbf{r}' := \mathbf{r} - (\gamma' - \gamma)\mathbf{P}\mathbf{v}$, and must compute a value \mathbf{v} such that $\|\mathbf{v}_{\gamma',r'}^* - \mathbf{v}\|_\infty \leq \epsilon/2$.

To enable computing approximately optimal policies in addition to approximately optimal values, we also introduce the following *policy*-sub-problem as follows.

Definition 33 (DMDP policy-sub-problem) Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbf{P})$ be an MDP and $\mathbf{r} \in \mathbb{R}_{\geq 0}^{\mathcal{S}}$ be a reward vector. In the $(\mathcal{M}, \mathbf{r}, \gamma', \mathbf{v}, \epsilon)$ -policy-sub-problem, we are given $\gamma' > 0$, $\mathbf{v} \in \mathbb{R}^{\mathcal{S}}$, $\epsilon \in (0, 1/(1 - \gamma'))]$

$$\mathbf{r}' := \mathbf{r} - (\gamma' - \gamma)\mathbf{P}\mathbf{v},$$

and we must compute a value v and a policy π such that $\mathbf{0} \leq \mathbf{v}_{\gamma', \mathbf{r}'}^* - \mathbf{v} \leq \epsilon \mathbf{1}$ and $\mathbf{v} \leq \mathbf{v}_{\gamma', \mathbf{r}'}^{\pi}$. That is, we must find an ϵ -optimal value and policy for $\mathcal{M}_{\gamma', \mathbf{r}'}$.

In comparison to the DMDP sub-problem (Definition 32), in the DMDP policy-sub-problem, we must not only output an ϵ -optimal value for $\mathcal{M}_{\gamma', \mathbf{r}'}$ but must *also* output a policy π that attains at least that value. For the DMDP problem, the outer process is defined (simply) as follows.

Definition 34 (DMDP post-process) Fix $\epsilon \in (0, 1/(1 - \gamma'))]$. For any $\mathbf{v}, \mathbf{v}' \in \mathbb{R}^{\mathcal{S}}$, we define $\zeta_{\epsilon}(\mathbf{v}, \mathbf{v}') := \mathbf{v}' - \epsilon \mathbf{1}$.

The outer process is intended to adjust (shift down the entries of) \mathbf{v}' to ensure that it is an *underestimate* of the optimal value.

We prove the following outer-solver for DMDPs, which we call the Proximal Reward Method (PRM), since it is in spirit similar to proximal point methods in convex optimization.

Theorem 35 (PRM - DMDP outer-solver) Let $\mathcal{M}_{\gamma, \mathbf{r}}$ be a DMDP, $\gamma' < \gamma$, and $\epsilon \in (0, 1/(1 - \gamma))]$. Let $\epsilon' = \epsilon/4 \cdot (1 - \gamma)/(1 - \gamma')$. Suppose that in each iteration $t \in [n_{\text{outer}}]$ of Algorithm 4, $\mathbf{v}_{t-1/2}$ is a solution to the $(\mathcal{M}, \mathbf{r}, \gamma', \mathbf{v}_{t-1}, \epsilon')$ -sub-problem. Suppose that $\mathbf{v}_{n_{\text{outer}}+1}, \pi_{n_{\text{outer}}+1}$ is a solution to the $(\mathcal{M}, \mathbf{r}, \gamma', \mathbf{v}_{n_{\text{outer}}}, \epsilon')$ -policy sub-problem. Then, $\mathbf{v}_{n_{\text{outer}}+1}, \pi_{n_{\text{outer}}+1}$ solves the DMDP problem.

We prove Theorem 35 in Section 4.4, and for now, focus on its application. To do so, we need to discuss high-precision algorithms for solving the DMDP sub-problem and policy-sub-problem.

Subproblem solvers. There are many high-accuracy algorithms for solving the DMDP subproblem as well as the DMDP-policy-sub-problem. These methods can broadly be divided into interior-point methods (e.g., (Sidford et al., 2018b; Lee and Sidford, 2014; van den Brand et al., 2021; Cohen et al., 2020a; Jiang et al., 2021)), classical value iteration (VI) (Littman et al., 1995; Tseng, 1990), variance-reduced variants of (VI), as well as policy-based methods such as policy iteration or policy gradient and variants (see e.g., Grondman et al. (2012); Bertsekas (2011) for a survey.) These variants (Sidford et al., 2018b; Jin et al., 2024) implement an approximate version of VI by trading-off between full batch queries (matrix-vector products in \mathbf{P}) and sample queries (simulator oracle queries) to obtain faster runtimes. Since we are interested in trade-offs between full batch and sample queries, we consider the Truncated Variance-Reduced Value Iteration (TVRVI) outer-solver from Jin et al. (2024) since it achieves the best trade-off between full batch and sample queries in this setting.

Theorem 36 (TVRVI, Theorem 1.2 of Jin et al. (2024), restated - DMDP sub-problem solver) Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbf{P})$ be an MDP, $\mathbf{r} \in \mathbb{R}_{\geq 0}^{\mathcal{S}}$ and $0 < \gamma' < \gamma < 1$. There are randomized algorithms $\mathcal{A}_{\xi, \mathcal{X}}^{\text{SVRG}|\gamma', \epsilon}(\mathbf{x})$ and $\mathcal{A}_{\xi, \mathcal{X}}^{\text{SVRG-Policy}|\gamma', \epsilon}(\mathbf{x})$ such that the following hold true.

- The algorithms take in the random seeds ξ, χ distributed as follows. The first random seed $\xi \sim \{i_1, \dots, i_T\}$ where each $i_j \in \mathcal{S}^A$ and each $i_j(s, a) \sim \mathbf{p}(s, a)$ and $T = \tilde{O}((1 - \gamma')^{-2})$. The second random seed $\chi \sim \text{Unif}[0]$ is a 0-bit random seed.
- If $\mathbf{v}' = \mathcal{A}_{\xi, \chi}^{\text{TVRVI}|\gamma', \epsilon, \delta}(\mathbf{v})$, then wp. $1 - \delta$ over the draw of the random seeds ξ and χ , \mathbf{v}' is a solution to the $(\mathcal{M}, \mathbf{r}, \gamma', \mathbf{v}, \epsilon)$ -sub-problem.
- If $\mathbf{v}', \pi = \mathcal{A}_{\xi, \chi}^{\text{TVRVI-Policy}|\gamma', \epsilon, \delta}(\mathbf{v})$, then wp. $1 - \delta$ over the draw of the random seeds ξ and χ , \mathbf{v}', π are a solution to the $(\mathcal{M}, \mathbf{r}, \gamma', \mathbf{v}, \epsilon)$ -policy sub-problem.
- The algorithms make only $\tilde{O}(1)$ batch queries and make only the sample queries that are encoded in ξ .

By instantiating the DMDP outer-solver PRM (Theorem 35) with TVRVI (Theorem 36) as the sub-solver, for $\gamma' \leq \gamma$, we obtain a full batch versus sample query trade-off of $\tilde{O}((1 - \gamma')/(1 - \gamma))$ full batch and $\tilde{O}((1 - \gamma')^{-1}(1 - \gamma)^{-1})$ sample queries per state-action pair. The pseudo-code is shown in Algorithm 4.

Algorithm 4: PRM-TV_{RVI} $_{\gamma', \delta}(\mathcal{M}, \mathbf{r}, \epsilon, \gamma)$ Pseudocode

Input: Matrix-vector oracle for \mathbf{P} , simulator oracle for \mathbf{P} , reward vector $\mathbf{r} \in [0, 1]^S$, error tolerance $\epsilon \in (0, 1/(1 - \gamma)]$, discount factor $\gamma \in (0, 1)$

Parameters: Failure probability δ , and $0 < \gamma' < \gamma < 1$.

Initialize $\mathbf{v}_0 \leftarrow \mathbf{0} \in \mathbb{R}^S$

Set sufficiently large $n_{\text{outer}} = \tilde{O}((1 - \gamma')/(1 - \gamma))$

Set $\epsilon' = \epsilon/4 \cdot (1 - \gamma)/(1 - \gamma')$

for each $t \in [n_{\text{outer}}]$ **do**

// We draw realizations of the random seed ξ according to the following distribution.

Draw $s_t := \{i_1, \dots, i_T\} \sim \mathcal{D}_\xi$ where each $i_j \in \mathcal{S}^A$ and each $i_j(s, a) \sim \mathbf{p}(s, a)$ for sufficiently large $T = \tilde{O}((1 - \gamma')^{-2})$

// We include a χ , a zero-bit random bit for technical consistency with Section 2

Draw $c_t \sim \mathcal{D}_\chi := \text{Unif}[0]$

$\mathbf{v}_{t-1/2} \leftarrow \mathcal{A}_{\xi=s_t, \chi=c_t}^{\text{TVRVI}|\gamma', \epsilon', \delta/n_{\text{outer}}}(\mathbf{v}_{t-1})$.

// The post-process is as described in Definition 34

$\mathbf{v}_t \leftarrow \zeta_{\epsilon'}(\mathbf{v}_{t-1}, \mathbf{v}_{t-1/2})$

end

// For the final iteration, we again draw a realization of the random seed ξ according to the following distribution.

Draw $s := \{i_1, \dots, i_T\} \sim \mathcal{D}_\xi$ where each $i_j \in \mathcal{S}^A$ and each $i_j(s, a) \sim \mathbf{p}(s, a)$ for sufficiently large $T = \tilde{O}((1 - \gamma')^{-2})$

// Again, we include a χ , a zero-bit random bit for technical consistency with Section 2

Draw $c_t \sim \mathcal{D}_\chi := \text{Unif}[0]$

$\mathbf{v}_{n_{\text{outer}}+1}, \pi_{n_{\text{outer}}+1} \leftarrow \mathcal{A}_{\xi=s, \chi=c}^{\text{TVRVI-Policy}|\gamma', \epsilon', \delta/n_{\text{outer}}}(\mathbf{v}_{n_{\text{outer}}})$.

return: $\mathbf{v}_{n_{\text{outer}}+1}, \pi_{n_{\text{outer}}+1}$

The sample reusing sub-solver Next, we observe that Algorithm 4 *exactly* fits into the Meta-Algorithm 1 framework from Section 2. In particular, we observe that Algorithm 2 implements $\text{Outer}(X, \text{Standard})$ where $X = (\mathcal{M}, \mathbf{r}, \gamma)$ is an DMDP problem instance as per Definition 31. To show that we can replace the standard sub-solver with the sample-reusing sub-solver (i.e., $\text{Outer}(X, \text{Reuse})$), we need to invoke Theorem 6.

However, we can also observe that the DMDP outer-solver PRM in Theorem 36 is *by definition* robust with respect to the target function f^{sub} (defined below), in the sense of Definition 2. Further, the sub-solver TVRVI *already* guarantees ℓ_∞ accuracy guarantees in the sense of Definition 1 for the target function

$$f^{\text{sub}}(\mathbf{v}) := \mathbf{r} - \mathbf{v}_{\gamma', (\gamma' - \gamma)}^* \mathbf{P} \mathbf{v}.$$

Consequently, we can *directly* apply our sample reuse framework to reuse samples across all n_{outer} iterations of the for loop in Algorithm 4 to obtain the following improved trade-off. In particular, by invoking Theorem 6, we obtain the following result.

Theorem 37 (DMDP trade-off improvement) *Let $\mathcal{M}_{\gamma, r}$ be a DMDP and $\gamma' < \gamma$. There is an algorithm (Algorithm 5) which makes $\tilde{O}((1 - \gamma')/(1 - \gamma))$ -batch queries and $\tilde{O}(\mathcal{A}_{\text{tot}}(1 - \gamma')^{-2})$ -sample queries and solves the DMDP problem w.p. $1 - \delta$.*

Faster algorithm for certain DMDPs Theorem 35 also yields another interesting implication regarding the runtime for solving a DMDP. In particular, Truncated Variance-Reduced Value Iteration Jin et al. (2024) is known to solve the $(\mathcal{M}, \gamma', \mathbf{r}', \epsilon, \delta)$ -sub-problem in $\tilde{O}(\text{nnz}(\mathbf{P}) + \mathcal{A}_{\text{tot}}(1 - \gamma')^{-2})$ -time (Theorem 1.1 of Jin et al. (2024)) for $\mathbf{r} \in [0, 1]^S$.³ Consequently, Theorem 36 solves the DMDP problem in

$$\tilde{O}\left(\frac{(1 - \gamma')}{(1 - \gamma)} \cdot (\text{nnz}(\mathbf{P}) + \mathcal{A}_{\text{tot}}(1 - \gamma')^{-2})\right) = \tilde{O}\left(\frac{\text{nnz}(\mathbf{P})(1 - \gamma')}{(1 - \gamma)} + \frac{\mathcal{A}_{\text{tot}}}{(1 - \gamma)(1 - \gamma')}\right)$$

time for $\gamma' \leq \gamma$. So, by selecting $1 - \gamma' = \max(\sqrt{\mathcal{A}_{\text{tot}}/\text{nnz}(\mathbf{P})}, 1 - \gamma)$ to minimize *runtime*, we obtain the following immediate corollary of Theorem 35.

Theorem 38 (Faster runtime for solving certain MDPs) *Suppose $\text{nnz}(\mathbf{P}) \leq \mathcal{A}_{\text{tot}}(1 - \gamma)^{-2}$. Then, there is an algorithm that solves the DMDP problem (Definition 31) in $\tilde{O}(\text{nnz}(\mathbf{P}) + \sqrt{\text{nnz}(\mathbf{P})\mathcal{A}_{\text{tot}}}(1 - \gamma)^{-1})$ -time.*

Proof If $\text{nnz}(\mathbf{P}) \leq \mathcal{A}_{\text{tot}}(1 - \gamma)^{-2}$, then taking

$$1 - \gamma' = \max(\sqrt{\mathcal{A}_{\text{tot}}/\text{nnz}(\mathbf{P})}, (1 - \gamma)) = \sqrt{\mathcal{A}_{\text{tot}}/\text{nnz}(\mathbf{P})},$$

in which case the runtime becomes

$$\tilde{O}\left(\text{nnz}(\mathbf{P}) + \sqrt{\text{nnz}(\mathbf{P})\mathcal{A}_{\text{tot}}}(1 - \gamma)^{-1}\right).$$

■

Theorem 38 improves on Jin et al. (2024) in the regime where $\text{nnz}(\mathbf{P}) = o((1 - \gamma)^{-2}\mathcal{A}_{\text{tot}})$ and improves upon vanilla value iteration (which runs in $\tilde{O}(\text{nnz}(\mathbf{P})(1 - \gamma)^{-1})$ in the regime where $\text{nnz}(\mathbf{P}) \geq \mathcal{A}_{\text{tot}}$.

3. We use $\text{nnz}(\mathbf{P})$ to denote the number of nonzero entries in \mathbf{P} .

Algorithm 5: PRM-TV_{RVI}-Reuse $_{\gamma',\delta}(\mathcal{M}, \mathbf{r}, \epsilon, \gamma)$ Pseudocode

Input: Matrix-vector oracle for \mathbf{P} , simulator oracle for \mathbf{P} , reward vector $\mathbf{r} \in [0, 1]^{\mathcal{S}}$, error tolerance $\epsilon \in (0, 1/(1 - \gamma)]$, discount factor $\gamma \in (0, 1)$

Parameters: Failure probability δ , and $0 < \gamma' < \gamma < 1$.

Initialize $\mathbf{v}_0 \leftarrow \mathbf{0} \in \mathbb{R}^{\mathcal{S}}$

Set sufficiently large $n_{\text{outer}} = \tilde{O}((1 - \gamma')/(1 - \gamma))$

Set $\epsilon' = \epsilon/4 \cdot (1 - \gamma)/(1 - \gamma')$

// We draw realizations of the random seed ξ according to the following distribution.

Draw $s_1 := \{i_1, \dots, i_T\} \sim \mathcal{D}_\xi$ where each $i_j \in \mathcal{S}^A$ and each $i_j(s, a) \sim \mathbf{p}(s, a)$ for sufficiently large $T = \tilde{O}((1 - \gamma')^{-2})$

for each $t \in [n_{\text{outer}} - 1]$ **do**

// Draw c'_t from the noisy distribution $\mathcal{D}_{\chi'}$ as in Meta-Algorithm 1.

Draw $c'_t \sim \mathcal{D}_{\chi'}$

// Implement the noisy analog of $\mathcal{A}^{\text{TVRVI}|\gamma',\epsilon',\delta/(5n_{\text{outer}}^2)}$ as in Meta-Algorithm 1.

$\mathbf{v}_{t-1/2} \leftarrow \mathcal{A}'_{\xi=s_t, \chi'=c'_t}{}^{\text{TVRVI}|\gamma',\epsilon',\delta/(5n_{\text{outer}}^2)}(\mathbf{v}_{t-1})$.

// The post-process is as described in Definition 34

$\mathbf{v}_t \leftarrow \zeta_{\epsilon'}(\mathbf{v}_{t-1}, \mathbf{v}_{t-1/2})$

end

// For the final iteration, we again draw a realization of the random seed ξ according to the following distribution.

Draw $s := \{i_1, \dots, i_T\} \sim \mathcal{D}_\xi$ where each $i_j \in \mathcal{S}^A$ and each $i_j(s, a) \sim \mathbf{p}(s, a)$ for sufficiently large $T = \tilde{O}((1 - \gamma')^{-2})$

// Again, we include a χ , a zero-bit random bit for technical consistency with Section 2

Draw $c_t \sim \mathcal{D}_\chi := \text{Unif}[0]$

$\mathbf{v}_{n_{\text{outer}}+1}, \pi_{n_{\text{outer}}+1} \leftarrow \mathcal{A}'_{\xi=s, \chi=c}{}^{\text{TVRVI-Policy}|\gamma',\epsilon',\delta/(5n_{\text{outer}}^2)}(\mathbf{v}_{n_{\text{outer}}})$.

return: $\mathbf{v}_{n_{\text{outer}}+1}, \pi_{n_{\text{outer}}+1}$

4.3. Infinite-horizon Average-reward MDPs

In this section, we consider the average-reward setting, in which there is no discount factor γ .

Definition 39 (Average-reward MDP (AMDP)) Let $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathbf{P})$ be an MDP and $\mathbf{r} \in [0, 1]^{\mathcal{A}}$. We denote the associated AMDP by $\mathcal{M}_{\mathbf{r}} = (\mathcal{M}; \mathbf{r})$.

In the average-reward case, we define the value of a policy as follows.

Definition 40 (AMDP policy value) Let $\mathcal{M}_{\mathbf{r}}$ be an AMDP and π be a policy. The value of policy π , denoted $\mathbf{v}_{\mathbf{r}}^\pi$, is defined as

$$\mathbf{v}_{\mathbf{r}}^\pi(s) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t \geq 0} \mathbf{r}(s_t, \pi(s_t)) \mid s_0 = s \right].$$

As in the DMDP case, the goal is to find an approximately optimal policy.

Definition 41 (AMDP problem) Let $\mathcal{M}_{\mathbf{r}}$ be an AMDP and $\epsilon \in (0, 1)$ be given. We must compute, w.p. $1 - \delta$, a policy π such that $\|\mathbf{v}_{\mathbf{r}}^\pi - \mathbf{v}_{\mathbf{r}}^*\|_\infty \leq \epsilon$. Such a policy is called ϵ -optimal for $\mathcal{M}_{\mathbf{r}}$.

In order to extend our improved trade-offs for DMDPs to AMDPs, we leverage a result of [Jin and Sidford \(2021\)](#), which showed that the AMDP problem can be reduced to solving a DMDP with a sufficiently high discount factor. In the following, Δ denotes the probability simplex.

Definition 42 (Mixing time) *Let \mathcal{M}_r be an AMDP. \mathcal{M}_r is said to be mixing if there exists a stationary distribution $\nu \in \Delta^S$ such that*

$$t_{\text{mix}} := \max_{\pi \in \Pi} \operatorname{argmin}_{t \geq 1} \sum_{s \in \mathcal{S}} \max_{\mathbf{q} \in \Delta^S} \mathbf{p}(s, \pi(s))^\top \mathbf{q} - \nu < \infty.$$

The quantity t_{mix} is called the mixing time of \mathcal{M}_r .

Lemma 43 (Lemma 3 of [Jin and Sidford \(2021\)](#)) *Let \mathcal{M}_r be an AMDP. Suppose the mixing time of the \mathcal{M}_r is $t_{\text{mix}} < \infty$. Then, for any $\epsilon > 0$ and $\gamma \in (0, 1 - \epsilon/(9t_{\text{mix}}))$, an $\epsilon/(3(1 - \gamma))$ -optimal policy for the DMDP $\mathcal{M}_{\gamma, r}$ is also an ϵ -optimal policy for the AMDP \mathcal{M}_r .*

By combining Lemma 43 with Theorem 37, we immediately obtain the following full batch and sample query trade-off for AMDPs.

Theorem 44 (AMDP trade-off improvement) *Let \mathcal{M}_r be an AMDP. Suppose the mixing time of \mathcal{M}_r is $t_{\text{mix}} < \infty$. Then, for $\gamma' \leq 1 - \epsilon/(9t_{\text{mix}})$, there is an algorithm that solves the AMDP problem using $\tilde{O}((1 - \gamma')t_{\text{mix}}/\epsilon)$ full batch queries and only $\tilde{O}(\mathcal{A}_{\text{tot}}(1 - \gamma')^{-2})$ sample queries.*

4.4. Proximal Reward Method for DMDPs: Proof of Theorem 35

In this section we present the proof of Theorem 35. First, we prove a stability result regarding the optimal value of a DMDP under a reward perturbation.

Lemma 45 *Let $\mathbf{r}, \mathbf{r}' \in \mathbb{R}^A$ such that $\mathbf{r}' \leq \mathbf{r}$, and let $\gamma > 0$. Then, for any $\mathbf{v} \in \mathbb{R}^S$, we have that for all $s \in \mathcal{S}$,*

$$\mathbf{0} \leq \mathbf{v}_{\gamma, \mathbf{r}}^* - \mathbf{v}_{\gamma, \mathbf{r}'}^* \leq \frac{1}{1 - \gamma} \cdot \left(\max_{(s, a) \in \mathcal{A}} \mathbf{r}(s, a) - \mathbf{r}'(s, a) \right) \cdot \mathbf{1}.$$

Proof By the Bellman optimality conditions (Definition 29), for each $s \in \mathcal{S}$ we have

$$\begin{aligned} \mathbf{v}_{\gamma, \mathbf{r}}^*(s) &= \max_{\pi \in \Pi} \left((\mathbf{I} - \gamma \mathbf{P}^\pi)^{-1} \mathbf{r}^\pi \right)(s), \\ \mathbf{v}_{\gamma, \mathbf{r}'}^*(s) &= \max_{\pi \in \Pi} \left((\mathbf{I} - \gamma \mathbf{P}^\pi)^{-1} \mathbf{r}'^\pi \right)(s). \end{aligned}$$

Since $(\mathbf{I} - \gamma \mathbf{P}^\pi)^{-1}$ is a positive matrix, we have that for each $s \in \mathcal{S}$,

$$\mathbf{0} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*(s) - \mathbf{v}_{\gamma, \mathbf{r}'}^*(s) \leq \max_{\pi \in \Pi} \left((\mathbf{I} - \gamma \mathbf{P}^\pi)^{-1} (\mathbf{r}^\pi - \mathbf{r}'^\pi) \right)(s) \leq \frac{1}{1 - \gamma} \cdot \max_{(s, a) \in \mathcal{A}} \mathbf{r}(s, a) - \mathbf{r}'(s, a). \quad \blacksquare$$

Next, we show how to iteratively solve a sequence of γ' -discounted MDPs to solve a γ -discounted MDP. First, we prove the following lemma, which bounds the convergence rate of the scheme which solves $\mathcal{M}_{\gamma, r}$ by iteratively solving problems in $\mathcal{M}_{\gamma', r'}$ for a sequence of rewards $\mathbf{r}^{(1)}, \dots, \mathbf{r}^{(T)}$.

Lemma 46 Let \mathcal{M} be an MDP and $\mathbf{r} \in \mathbb{R}_{\geq 0}^A$ be a reward vector. Let $0 < \gamma' \leq \gamma < 1$, $\eta, \epsilon > 0$, and $T \geq 1$. Let $\mathbf{r}^{(0)} = \mathbf{r}$ and $\mathbf{v}^{(0)} = \mathbf{0}$. For each $t \geq 1$, define

$$\mathbf{r}^{(t)} := \mathbf{r} - (\gamma' - \gamma)\mathbf{P}\mathbf{v}^{(t-1)}. \quad (10)$$

Suppose that for each $t \geq 1$, $\mathbf{v}^{(t)}$ satisfies

$$\mathbf{0} \leq \mathbf{v}_{\gamma', \mathbf{r}^{(t)}}^* - \eta \mathbf{1} \leq \mathbf{v}^{(t)} \leq \mathbf{v}_{\gamma', \mathbf{r}^{(t)}}^*.$$

Then, for any $T \geq 1$,

$$\mathbf{0} \leq \mathbf{v}_{\gamma, \mathbf{r}}^* - \left(\left(\frac{(\gamma - \gamma')}{(1 - \gamma')} \right)^T \max_{s \in \mathcal{S}} \left(\mathbf{v}_{\gamma, \mathbf{r}}^* - \mathbf{v}^{(0)} \right)(s) + \frac{(1 - \gamma')}{(1 - \gamma)} \eta \right) \cdot \mathbf{1} \leq \mathbf{v}^{(T)} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*.$$

Consequently, if $\eta \leq \frac{(1 - \gamma)}{2(1 - \gamma')} \epsilon$ and $T = \tilde{\Omega} \left(\frac{(1 - \gamma')}{(1 - \gamma)} \right)$ then $\mathbf{0} \leq \mathbf{v}_{\gamma, \mathbf{r}}^* - \epsilon \mathbf{1} \leq \mathbf{v}^{(T)} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*$.

Proof First, we will induct on t to show that for $t \geq 0$,

$$\mathbf{0} \leq \mathbf{v}_{\gamma, \mathbf{r}}^* - \left(\left(\frac{(\gamma - \gamma')}{(1 - \gamma')} \right)^t \max_{s \in \mathcal{S}} \left(\mathbf{v}_{\gamma, \mathbf{r}}^* - \mathbf{v}^{(0)} \right)(s) + \sum_{j=1}^t \left(\frac{(\gamma - \gamma')}{(1 - \gamma')} \right)^{(t-j)} \eta \right) \cdot \mathbf{1} \leq \mathbf{v}^{(t)} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*. \quad (11)$$

Inductive proof of (11). In the base case, when $t = 0$ the claim reduces to $\mathbf{0} \leq \mathbf{v}^{(0)} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*$, which is trivially satisfied because $\mathbf{v}^{(0)} = \mathbf{0}$ and $\mathbf{r} \in \mathbb{R}_{\geq 0}^A$. For the inductive step, we have

$$\mathbf{0} \leq \mathbf{v}_{\gamma, \mathbf{r}}^* - \left(\left(\frac{(\gamma - \gamma')}{(1 - \gamma')} \right)^{t-1} \max_{s \in \mathcal{S}} \left(\mathbf{v}_{\gamma, \mathbf{r}}^* - \mathbf{v}^{(0)} \right) + \sum_{j=1}^{t-1} \left(\frac{(\gamma - \gamma')}{(1 - \gamma')} \right)^{(t-1-j)} \eta \right) \cdot \mathbf{1} \leq \mathbf{v}^{(t-1)} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*.$$

Next, observe that

$$\mathbf{v}_{\gamma, \mathbf{r}}^* - \mathbf{v}^{(t)} = \mathbf{v}_{\gamma, \mathbf{r}}^* - \mathbf{v}_{\gamma', \mathbf{r}^{(t)}}^* + \mathbf{v}_{\gamma', \mathbf{r}^{(t)}}^* - \mathbf{v}^{(t)}. \quad (12)$$

By the assumption on $\mathbf{v}^{(t)}$,

$$\mathbf{0} \leq \mathbf{v}_{\gamma', \mathbf{r}^{(t)}}^* - \mathbf{v}^{(t)} \leq \eta \mathbf{1}. \quad (13)$$

By the Bellman optimality conditions, for each $s \in \mathcal{S}$,

$$\mathbf{v}_{\gamma, \mathbf{r}}^*(s) = \max_{a \in \mathcal{A}_s} \mathbf{r}(s, a) + \gamma \mathbf{p}(s, a)^\top \mathbf{v}_{\gamma, \mathbf{r}}^* = \max_{a \in \mathcal{A}_s} \mathbf{r}(s, a) - (\gamma' - \gamma) \mathbf{p}(s, a)^\top \mathbf{v}_{\gamma, \mathbf{r}}^* + \gamma' \mathbf{p}(s, a)^\top \mathbf{v}_{\gamma, \mathbf{r}}^*.$$

Consequently, $\mathbf{v}_{\gamma, \mathbf{r}}^* = \mathbf{v}_{\gamma', \mathbf{r} - (\gamma' - \gamma)\mathbf{P}\mathbf{v}_{\gamma, \mathbf{r}}^*}^*$. By substituting into (12) and applying the definition of $\mathbf{r}^{(t)}$, Lemma 45 implies that

$$\mathbf{0} \leq \mathbf{v}_{\gamma, \mathbf{r}}^* - \mathbf{v}^{(t)} \leq \frac{(\gamma - \gamma')}{(1 - \gamma')} \cdot \max_{(s, a) \in \mathcal{A}} (\mathbf{P}(\mathbf{v}^{(t-1)} - \mathbf{v}_{\gamma, \mathbf{r}}^*))(s, a) \cdot \mathbf{1} + \eta \mathbf{1}$$

$$\leq \frac{(\gamma - \gamma')}{(1 - \gamma')} \cdot \max_{s \in \mathcal{S}} (\mathbf{v}^{(t-1)}(s) - \mathbf{v}_{\mathbf{r}, \gamma}^*(s)) \cdot \mathbf{1} + \eta \mathbf{1},$$

where in the last step we used that $\|\mathbf{P}\|_\infty = 1$ and \mathbf{P} is a positive matrix. Consequently,

$$\mathbf{0} \leq \mathbf{v}_{\gamma, \mathbf{r}}^* - \left(\left(\frac{(\gamma - \gamma')}{(1 - \gamma')} \right)^t \max_{s \in \mathcal{S}} (\mathbf{v}_{\gamma, \mathbf{r}}^* - \mathbf{v}^{(0)})(s) + \sum_{j=1}^t \left(\frac{(\gamma - \gamma')}{(1 - \gamma')} \right)^{t-j} \eta \right) \cdot \mathbf{1} \leq \mathbf{v}^{(t)} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*.$$

This completes the inductive argument.

Now, we bound the geometric series as follows

$$\sum_{j=1}^t \left(\frac{(\gamma - \gamma')}{(1 - \gamma')} \right)^{t-j} \eta = \sum_{j=1}^t \left(1 - \frac{(1 - \gamma)}{(1 - \gamma')} \right)^{t-j} \eta \leq \frac{(1 - \gamma')}{(1 - \gamma)} \eta.$$

Finally, note that when $T = \tilde{\Omega} \left(\frac{(1 - \gamma')}{(1 - \gamma)} \right)$ and $\eta \leq \frac{(1 - \gamma)}{2(1 - \gamma')} \epsilon$,

$$\mathbf{v}_{\gamma, \mathbf{r}}^* - \epsilon \leq \mathbf{v}^{(T)} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*.$$

■

Finally, we can leverage Lemma 46 to complete the proof of Theorem 35.

Theorem 35 (PRM - DMDP outer-solver) *Let $\mathcal{M}_{\gamma, \mathbf{r}}$ be a DMDP, $\gamma' < \gamma$, and $\epsilon \in (0, 1/(1 - \gamma))$. Let $\epsilon' = \epsilon/4 \cdot (1 - \gamma)/(1 - \gamma')$. Suppose that in each iteration $t \in [n_{\text{outer}}]$ of Algorithm 4, $\mathbf{v}_{t-1/2}$ is a solution to the $(\mathcal{M}, \mathbf{r}, \gamma', \mathbf{v}_{t-1}, \epsilon')$ -sub-problem. Suppose that $\mathbf{v}_{n_{\text{outer}}+1}, \pi_{n_{\text{outer}}+1}$ is a solution to the $(\mathcal{M}, \mathbf{r}, \gamma', \mathbf{v}_{n_{\text{outer}}}, \epsilon')$ -policy sub-problem. Then, $\mathbf{v}_{n_{\text{outer}}+1}, \pi_{n_{\text{outer}}+1}$ solves the DMDP problem.*

Proof Note that the outer process in Algorithm 4 (and Algorithm 5) ensures the following. Suppose each $\mathbf{v}_{t-1/2}$ is a solution to the $(\mathcal{M}, \mathbf{r}, \gamma', \mathbf{v}_{t-1}, \epsilon/4 \cdot (1 - \gamma)/(1 - \gamma'))$ -sub-problem. Then due to the post-process, each \mathbf{v}_t meets the conditions on $\mathbf{v}^{(t)}$ ins Lemma 46 for $\eta = \epsilon/2 \cdot (1 - \gamma)/(1 - \gamma')$.

Consequently, by Lemma 46, we have that for sufficiently large $n_{\text{outer}} = \tilde{O}((1 - \gamma')/(1 - \gamma))$,

$$\begin{aligned} \mathbf{0} &\leq \mathbf{v}_{\gamma, \mathbf{r}}^* - \frac{\epsilon}{2} \mathbf{1} \leq \mathbf{v}_{n_{\text{outer}}} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*, \\ \mathbf{0} &\leq \mathbf{v}_{\gamma, \mathbf{r}}^* - \frac{\epsilon}{2} \mathbf{1} \leq \mathbf{v}_{n_{\text{outer}}+1} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*, \end{aligned}$$

and by the definition of the policy-sub-problem, we can further conclude

$$\mathbf{0} \leq \mathbf{v}_{\gamma', \mathbf{r}^{n_{\text{outer}}+1}}^* - \frac{\epsilon}{2} \mathbf{1} \leq \mathbf{v}_{n_{\text{outer}}+1} \leq \mathbf{v}_{\gamma', \mathbf{r}^{n_{\text{outer}}+1}}^{\pi_{n_{\text{outer}}+1}} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*, \quad (14)$$

where

$$\mathbf{r}^{n_{\text{outer}}+1} := \mathbf{r} - (\gamma' - \gamma) \mathbf{P} \mathbf{v}_{n_{\text{outer}}}.$$

Now, note that for all $s \in \mathcal{S}$,

$$\mathbf{v}_{\gamma, \mathbf{r}}^{\pi_{n_{\text{outer}}+1}}(s) = \mathbf{r}(s, \pi_{n_{\text{outer}}+1}(s)) + \gamma \mathbf{p}(s, \pi_{n_{\text{outer}}+1}(s))^\top \mathbf{v}_{\gamma, \mathbf{r}}^{\pi_{n_{\text{outer}}+1}}$$

$$= \mathbf{r}(s, \pi_{n_{\text{outer}}+1}(s)) - (\gamma' - \gamma) \mathbf{p}(s, \pi_{n_{\text{outer}}+1}(s))^\top \mathbf{v}_{\gamma, \mathbf{r}}^{\pi_{n_{\text{outer}}}} + \gamma' \mathbf{p}(s, \pi_{n_{\text{outer}}+1}(s))^\top \mathbf{v}_{\gamma, \mathbf{r}}^{\pi_{n_{\text{outer}}+1}}.$$

Above, in the first line we used the Bellman formulation for values of policies (Definition 29).

Thus, $\mathbf{v}_{\gamma, \mathbf{r}}^{\pi_{n_{\text{outer}}+1}} = \mathbf{v}_{\gamma', \mathbf{r} - (\gamma' - \gamma) \mathbf{P} \mathbf{v}_{\gamma, \mathbf{r}}^{\pi_{n_{\text{outer}}+1}}}$. Since $(\mathbf{I} - \gamma' \mathbf{P}^{\pi_{n_{\text{outer}}+1}})^{-1}$ and $\mathbf{P}^{\pi_{n_{\text{outer}}+1}}$ are positive matrices and $(\gamma - \gamma') / (1 - \gamma') \leq 1$ we have that

$$\begin{aligned} \mathbf{v}_{\gamma', \mathbf{r} - (\gamma' - \gamma) \mathbf{P} \mathbf{v}_{\gamma, \mathbf{r}}^{\pi_{n_{\text{outer}}+1}}}^{\pi_{n_{\text{outer}}+1}} - \mathbf{v}_{\gamma', \mathbf{r}^{\pi_{n_{\text{outer}}+1}}}^{\pi_{n_{\text{outer}}+1}} &= (\mathbf{I} - \gamma' \mathbf{P}^{\pi_{n_{\text{outer}}+1}})^{-1} ((\gamma - \gamma') \mathbf{P}^{\pi_{n_{\text{outer}}}} (\mathbf{v}_{\gamma, \mathbf{r}}^{\pi_{n_{\text{outer}}}} - \mathbf{v}_{n_{\text{outer}}})) \\ &\leq (\mathbf{I} - \gamma' \mathbf{P}^{\pi_{n_{\text{outer}}+1}})^{-1} ((\gamma - \gamma') \mathbf{P}^{\pi_{n_{\text{outer}}}} (\mathbf{v}_{\gamma, \mathbf{r}}^* - \mathbf{v}^{(n_{\text{outer}})})) \\ &\leq \max_{(s, a) \in \mathcal{A}} (\mathbf{v}_{\gamma, \mathbf{r}}^*(s, a) - \mathbf{v}_{n_{\text{outer}}}(s, a)) \\ &\leq (\gamma - \gamma') / (1 - \gamma') \epsilon / 2 \\ &\leq \frac{\epsilon}{2}, \end{aligned}$$

where the second to last step used that $\|\mathbf{I} - \gamma' \mathbf{P}^\pi\|_\infty \leq 1 / (1 - \gamma')$ for any policy π . Consequently, combining with (14), we conclude that $\mathbf{v}_{\gamma, \mathbf{r}}^* - \epsilon \leq \mathbf{v}_{\gamma, \mathbf{r}}^{\pi_{n_{\text{outer}}}} \leq \mathbf{v}_{\gamma, \mathbf{r}}^*$, which completes the proof. \blacksquare

5. Application: Matrix games and minimax problems

In this section, we consider minimax problems, including ℓ_2 - ℓ_1 and ℓ_2 - ℓ_2 matrix-games and finite-sum minimax problems. In Section 5.1 we discuss general preliminaries. In Section 5.2, we discuss ℓ_2 - ℓ_1 matrix games in Section 5.2.1 and ℓ_2 - ℓ_2 matrix games in Section 5.2.2. Section 5.3 discusses applications of our ℓ_2 - ℓ_1 matrix games improvements for two computational geometry problems: maximum inscribed ball and minimum enclosing ball.

5.1. Preliminaries

We first outline preliminaries of the minimax problems we consider. The notation in this subsection is consistent with that of Carmon et al. (2019).

Problem setup. A *setup* is the triplet $(\mathcal{Z} = \mathcal{X} \times \mathcal{Y}, \|\cdot\|, r)$ where we use $\mathcal{Z} := \mathcal{X} \times \mathcal{Y}$, where (1) \mathcal{X} is a compact and convex subset of \mathbb{R}^n and \mathcal{Y} is a compact and convex subset of \mathbb{R}^m ; (2) $\|\cdot\|$ is a norm on \mathcal{Z} , and (3) r is a 1-strongly convex function with respect to \mathcal{Z} and $\|\cdot\|$. We call r a *distance generating function* and denote the associated Bregman divergence as

$$V_{\mathbf{z}}(\mathbf{z}') := r(\mathbf{z}') - r(\mathbf{z}) - \langle \nabla r(\mathbf{z}), \mathbf{z}' - \mathbf{z} \rangle \geq \frac{1}{2} \|\mathbf{z}' - \mathbf{z}\|^2.$$

We also denote $\Theta := \max_{\mathbf{z}'} r(\mathbf{z}') - \min_{\mathbf{z}} r(\mathbf{z})$ and assume it is finite. We use $\|\cdot\|_*$ to denote the dual norm of $\|\cdot\|$. For $\mathbf{z} \in \mathcal{Z}$, we often write $\mathbf{z}^{\mathcal{X}} \in \mathcal{X}$, $\mathbf{z}^{\mathcal{Y}} \in \mathcal{Y}$ to denote the first n and last m coordinates of \mathbf{z} , respectively. We use $d = m + n$ throughout this section.

We assume that \mathcal{Z} is sufficiently simple such that given any $\mathbf{z}' \in \mathbb{R}^d$, one can compute the projection of \mathbf{z}' onto \mathcal{Z} with respect to $\|\cdot\|$, denoted $\text{proj}_{\mathcal{Z}}(\mathbf{z}')$, in $\tilde{O}(d)$ -time. (This is true, for example, for the Euclidean unit ball, or the probability simplex, which are the relevant cases for ℓ_2 - ℓ_2 matrix-games and ℓ_2 - ℓ_1 matrix games.) Finally, we assume that $c\|\cdot\|_\infty \leq \|\cdot\| \leq C\|\cdot\|_\infty$ over \mathcal{Z} , for some $c, C = \text{poly}(d)$. (This is true, for example, for the ℓ_1 and ℓ_2 norms, which are the relevant cases for ℓ_2 - ℓ_2 matrix-games and ℓ_2 - ℓ_1 matrix games.)

Minimax problems. We consider minimax (saddle-point) problems of the form

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} f(x, y),$$

for some setup $(\mathcal{Z} = (\mathcal{X}, \mathcal{Y}), \|\cdot\|, r)$. We use $g(z) := (\nabla_x f(z), -\nabla_y f(z)) \in \mathbb{R}^d$ to denote the *gradient mapping* of f . We use L to denote the Lipschitz constant of g , $D := \max_{z, z' \in \mathcal{Z}} \|z - z'\|$, and $G := \max_{z \in \mathcal{Z}} \|g(z)\|$. We assume that D, L, G are finite and hide polylogarithmic dependencies in these parameters inside of $\tilde{O}(\cdot)$ notation.

Next, we define the minimax problem we consider in this Section.

Definition 47 (Minimax problem) *In the minimax problem, we are given $f : \mathcal{Z} = (\mathcal{X} \times \mathcal{Y}) \rightarrow \mathbb{R}$, $\epsilon > 0$, and $\delta \in (0, 1)$. We must compute \mathbf{x}, \mathbf{y} such that*

$$\max_{\mathbf{y}' \in \mathcal{Y}} f(\mathbf{x}, \mathbf{y}') - \min_{\mathbf{x}' \in \mathcal{X}} f(\mathbf{x}', \mathbf{y}) \leq \epsilon.$$

We will later discuss matrix-games and finite-sum minimax problems as special cases of Definition 47 and discuss the relevant batch and sample query models therein. However, in this section we discuss the outer-solver and sub-problem structure for general minimax problems following the conceptual proximal point framework of Carmon et al. (2019); Nemirovski (2004).

Conceptual proximal point. Carmon et al. (2019); Nemirovski (2004) showed how to solve minimax problems of the form of Definition 47 by iteratively solving a series of α -regularized sub-problems. This method is inspired by Nemirovski’s “conceptual prox point method” (Nemirovski, 2004). Correspondingly, we define the minimax sub-problem as follows.

Definition 48 (Minimax sub-problem) *Let f be as in Definition 47. Let $\mathbf{z}_0 \in \mathcal{Z}, \alpha > 0, \epsilon > 0$. Let \mathbf{z}^α be the solution to the problem $\min_{\mathbf{x}' \in \mathcal{X}} \max_{\mathbf{y}' \in \mathcal{Y}} f(\mathbf{z}) + \alpha V_{\mathbf{z}}(\mathbf{x}', \mathbf{y}')$. In the $(\mathbf{z}_0, \alpha, \epsilon)$ -sub-problem we define $f^{\text{sub}}(\mathbf{z}) := \mathbf{z}_\alpha$ to be the unique point in \mathcal{Z} such that*

$$\langle g(\mathbf{z}_\alpha) + \alpha \nabla V_{\mathbf{z}}(\mathbf{z}_\alpha), \mathbf{z}_\alpha - \mathbf{u} \rangle, \text{ for all } \mathbf{u} \in \mathcal{Z}.$$

We must output a $\mathbf{z}' \in \mathbb{R}^d$ such that $\|\mathbf{z}' - f^{\text{sub}}(\mathbf{z})\|_\infty \leq \epsilon$.

Carmon et al. (2019); Nemirovski (2004) showed how to solve the minimax problem (Definition 47) by solving a sequence of sub-problems of the form of Definition 48 using Conceptual Proximal Point (CPP) as the outer-solver. Here, the outer process is a projection step (to ensure feasibility) followed by an extragradient step.

Definition 49 (Minimax post-process) *Let f be as in Definition 48, and $\alpha > 0$. For any $\mathbf{z}, \mathbf{z}' \in \mathbb{R}^d \times \mathbb{R}^d$, we define $\mathbf{z}'' := \text{proj}_{\mathcal{Z}}(\mathbf{z}')$*

$$\zeta(\mathbf{z}, \mathbf{z}') := \underset{\tilde{\mathbf{z}} \in \mathcal{Z}}{\text{argmin}} (\langle g(\mathbf{z}''), \tilde{\mathbf{z}} \rangle + \alpha V_{\mathbf{z}}(\tilde{\mathbf{z}})).$$

We now specify the outer-solver framework using CPP as follows. We slightly restate a variant of Proposition 4 of Carmon et al. (2019).

Theorem 50 (CPP, Proposition 4 of Carmon et al. (2019), adapted - Minimax outer-solver) *Let f be as in Definition 47 and $\alpha > 0$. Consider Meta-Algorithm 1 with the following instantiation of parameters. For fixed $\alpha > 0$,*

- Initialize \mathbf{u}_0 with $\operatorname{argmin}_{\mathbf{z} \in \mathcal{Z}} r(\mathbf{z})$.
- Define ζ as in Definition 49.
- Set $\mathbf{w}(t) := 1/n_{\text{outer}}$ for each $t \in [n_{\text{outer}}]$.

Then, there is an $\epsilon' = \text{poly}(G, L, D, \Theta, \epsilon, d)$ such that the following holds. Suppose that in each iteration $t \in [n_{\text{outer}}]$ of Algorithm 1, $\mathbf{u}_{t-1/2}$ is a solution to the $(\mathbf{u}_{t-1}, \alpha, \epsilon')$ -sub-problem; then $\mathbf{u}_{n_{\text{outer}}}$ is a solution to the minimax problem.

Proof To prove this, we appeal to Proposition 4 of Carmon et al. (2019). By Carmon et al. (2019)'s Proposition 4, it is sufficient to show that $\mathbf{u}_{t-1/2}$ satisfies

$$\langle g(\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2})), \operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}) - \mathbf{u} \rangle - \alpha V_{\mathbf{u}_{t-1}}(\mathbf{u}) \leq \epsilon \quad \text{for all } \mathbf{u} \in \mathcal{Z}. \quad (15)$$

Consequently, to prove the theorem, it suffices to show that whenever $\mathbf{u}_{t-1/2}$ is a solution to the $(\mathbf{u}_{t-1}, \alpha, \epsilon')$ -sub-problem, (15) holds. To this end, let \mathbf{z}_α be the unique point in \mathcal{Z} such that

$$\langle g(\mathbf{z}_\alpha) + \alpha \nabla V_{\mathbf{u}_{t-1}}(\mathbf{z}_\alpha), \mathbf{z}_\alpha - \mathbf{u} \rangle \leq 0, \quad \text{for all } \mathbf{u} \in \mathcal{Z}.$$

By the three-point-equality for Bregman divergences, we have that

$$\langle g(\mathbf{z}_\alpha), \mathbf{z}_\alpha - \mathbf{u} \rangle - \alpha V_{\mathbf{u}_{t-1}}(\mathbf{u}) \leq -\alpha V_{\mathbf{z}_\alpha}(\mathbf{u}) - \alpha V_{\mathbf{u}_{t-1}}(\mathbf{z}_\alpha), \quad \text{for all } \mathbf{u} \in \mathcal{Z}.$$

Now, since $\|\mathbf{u}_{t-1/2} - \mathbf{z}_\alpha\|_\infty \leq \epsilon$ and $\mathbf{z}_\alpha \in \mathcal{Z}$, we have that

$$\|\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}), \mathbf{z}_\alpha\| \leq \|\mathbf{u}_{t-1/2}, -\mathbf{z}_\alpha\|$$

Consequently, by equivalence of norms,

$$c\|\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}), \mathbf{z}_\alpha\|_\infty \leq \|\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}), \mathbf{z}_\alpha\| \leq \|\mathbf{u}_{t-1/2}, -\mathbf{z}_\alpha\| \leq C\|\mathbf{u}_{t-1/2}, -\mathbf{z}_\alpha\|_\infty.$$

Thus,

$$\|\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}), \mathbf{z}_\alpha\|_\infty \leq \|\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}), \mathbf{z}_\alpha\| \leq \|\mathbf{u}_{t-1/2}, -\mathbf{z}_\alpha\| \leq C/c \cdot \|\mathbf{u}_{t-1/2}, -\mathbf{z}_\alpha\|_\infty.$$

Thus, for any $\mathbf{u} \in \mathcal{Z}$ we can write

$$\begin{aligned} & \langle g(\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2})), \operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}) - \mathbf{u} \rangle \\ &= \langle g(\mathbf{z}_\alpha), \operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}) - \mathbf{u} \rangle + \langle g(\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2})) - g(\mathbf{z}_\alpha), \operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}) - \mathbf{u} \rangle \\ &= \langle g(\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2})), \mathbf{z}_\alpha - \mathbf{u} \rangle + \langle g(\mathbf{z}_\alpha), \operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}) - \mathbf{z}_\alpha \rangle \\ &\quad - \langle g(\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2})) - g(\mathbf{z}_\alpha), \operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}) - \mathbf{u} \rangle \\ &\leq \langle g(\mathbf{z}_\alpha), \mathbf{z}_\alpha - \mathbf{u} \rangle + \|g(\mathbf{z}_\alpha)\|_1 C/c \cdot \epsilon' + \|g(\mathbf{z}_\alpha) - g(\operatorname{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}))\|_* D \end{aligned}$$

where in the last line we used Holder's inequality. Now, using the Lipschitzness of g and equivalence of norms,

$$\|g(\mathbf{z}_\alpha) - g(\text{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}))\|_* \leq L\|\mathbf{z}_\alpha - \mathbf{u}_{t-1/2}\| \leq LC\|\mathbf{z}_\alpha - \mathbf{u}_{t-1/2}\|_\infty \leq LC\epsilon'.$$

And again, using Consequently, taking ϵ' to be a sufficiently small polynomial in L, D, ϵ, d is enough to ensure that (15) holds, i.e., that

$$\langle g(\text{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2})), \text{proj}_{\mathcal{Z}}(\mathbf{u}_{t-1/2}) - \mathbf{u} \rangle - \alpha V_{\mathbf{u}_{t-1}}(\mathbf{u}) \leq GC/c \cot \epsilon' + DLC\epsilon' \leq \epsilon.$$

■

In our applications, we often leverage the following fact about Bregman divergences. (Recall that c is defined in Section 5.1).

Fact 51 *Let $r : \mathcal{Z} \rightarrow \mathbb{R}$ be a 1-strongly convex function with respect to \mathcal{Z} and $\|\cdot\|$. Then,*

$$V_{\mathcal{Z}}(\mathbf{z}') \geq \frac{1}{2} \|\mathbf{z}' - \mathbf{z}\|^2 \geq \frac{c^2}{2} \|\mathbf{z}' - \mathbf{z}\|_\infty^2.$$

The specific sub-problem solvers will vary between ℓ_2 - ℓ_2 matrix-games, ℓ_2 - ℓ_1 matrix-games, and finite-sum minimax problems. However, observe that Theorem 50 already establishes that CPP is robust to bounded ℓ_∞ error in the sub-problem solutions, in the sense of Definition 2. Consequently, it is amenable to using our sample-reuse framework developed in Section 2. We discuss this in the following sections.

5.2. Matrix games

We consider a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and use $\mathbf{a}_i, \mathbf{a}^j$ to denote the i -th row and j -th column of \mathbf{A} , respectively. We use $A_{i,j}$ to denote the (i, j) -th entry of \mathbf{A} . We use $\|\mathbf{A}\|_{2 \rightarrow \infty} := \max_{i \in [m]} \|\mathbf{a}_i\|_2$ and $\|\mathbf{A}\|_F$ to denote the Frobenius norm.

To discuss the application of our pseudo-independence results for improved oracle complexity trade-offs on matrix games, we first restrict to minimax problems on functions f corresponding to composite matrix games.

Definition 52 (MG problem) *In the (composite) matrix-game problem, we are given a setup ($\mathcal{Z} = \mathcal{X} \times \mathcal{Y}, \|\cdot\|, r$); a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$; convex, differentiable functions $\phi : \mathcal{X} \rightarrow \mathbb{R}, \psi : \mathcal{Y} \rightarrow \mathbb{R}$; $\epsilon > 0$; and $\delta \in (0, 1)$. We must solve the minimax problem (as in Definition 47) for the function*

$$f(\mathbf{x}, \mathbf{y}) = \mathbf{y}^\top \mathbf{A} \mathbf{x} + \phi(\mathbf{x}) - \psi(\mathbf{y}).$$

Next, we define the relevant full batch and sample query oracles for the matrix games problems we consider. As in Carmon et al. (2019) we assume that ϕ and ψ are explicit and that \mathbf{A} can be accessed via oracle queries. Concretely, we define one batch oracle and two types of sample oracles for accessing \mathbf{A} .

Definition 53 (Matrix-vector oracle - MG batch oracle) *When queried with $(\mathbf{x}, \mathbf{y}) \in \mathcal{Z}$, a matrix-vector oracle for \mathbf{A} returns $(\mathbf{A} \mathbf{x}, \mathbf{A}^\top \mathbf{y})$.*

Definition 54 (Row/column oracle - MG sample oracle, Type I) When queried with (i, \cdot) for $i \in [m]$, the oracle returns \mathbf{a}_i , the i -th row of \mathbf{A} . When queried with (\cdot, j) for $j \in [n]$, the oracle returns \mathbf{a}^j , the j -th column of \mathbf{A} .

Definition 55 (Entry oracle - MG sample oracle Type II) When queried with $(i, j) \in [m] \times [n]$, the entry oracle returns $A_{i,j}$.

5.2.1. COMPOSITE ℓ_2 - ℓ_1 MATRIX GAMES

In this section, we discuss the composite ℓ_2 - ℓ_1 matrix games setting. Throughout this section, we use the setup $(\mathcal{Z} = (\mathcal{X}, \mathcal{Y}), \|\cdot\|, r)$ where

- $\mathcal{X} = \mathbb{B}^n := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq 1\}$ is the Euclidean ball of radius 1 centered at the origin; and $\mathcal{Y} = \Delta^m$ is the m -dimensional simplex.
- $\|\cdot\| : \mathcal{Z} \rightarrow \mathbb{R}$ is given by $\|\mathbf{z}\| = \sqrt{\|\mathbf{z}^{\mathcal{X}}\|_2^2 + \|\mathbf{z}^{\mathcal{Y}}\|_1^2}$.
- $r : \mathcal{Z} \rightarrow \mathbb{R}$ is given by $r(\mathbf{z}) = \frac{1}{2} \|\mathbf{z}^{\mathcal{X}}\|_2^2 + \sum_{i \in [m]} \mathbf{z}^{\mathcal{Y}}(i) \log(\mathbf{z}^{\mathcal{Y}}(i))$.

In this case, the relevant constants defined in Section 5.1 are trivially given by

- $L \leq \|\mathbf{A}\|_{2 \rightarrow \infty}, G \leq \max_{i,j} |A_{i,j}|$
- $D = 1$
- $c = 1, C = d^2$

The function r is known to be 1-strongly convex with respect to $\mathcal{Z}, \|\cdot\|$ (see, e.g., Section 4.2 of Carmon et al. (2019)). To begin, we define the distributions which we will sample from in order to make sample queries.

Definition 56 Let $\mathbf{A} \in \mathbb{R}^{m \times n}$. For each $i \in [m]$, define the distribution $\mathcal{D}_{\text{entry}}(i)$ as follows:

$$\mathbb{P}_{b \sim \mathcal{D}_{\text{entry}}(i)} \{b = j\} = \frac{A_{i,j}^2}{\|\mathbf{a}_i\|_2^2}.$$

The following theorem states the guarantees of the variance-reduced mirror descent (Algorithm 4 of Carmon et al. (2019), VRMD1) solves a minimax-subproblem in the ℓ_2 - ℓ_1 matrix-games setting using $\tilde{O}(1)$ full batch queries, a number of non-oblivious, i.e., adaptive sample queries of Type I (see Definition 54), and a number of *oblivious* sample queries of Type II (see Definition 55).

Theorem 57 (VRMD1, Theorem 2 of Carmon et al. (2019), adapted - ℓ_2 - ℓ_1 MG sub-solver) Let f be as defined in Definition 52 and $\mathbf{z} \in \mathcal{Z}$. There is a randomized algorithm $\mathcal{A}_{\xi, \chi}^{\text{VRMD1-HP}|\alpha, \epsilon, \delta}$ such that the following holds.

- The algorithm takes in the random seeds ξ, χ distributed as follows, for some sufficiently large $T = \tilde{O}(\|\mathbf{A}\|_{2 \rightarrow \infty} / \alpha^2)$. The first random seed $\xi \sim \{\{(i_{1t}, j_{1t}), \dots, (i_{mt}, j_{mt})\}_{t=1}^T\} \subset [m] \times [n]$ where for each $q \in [m], t \in [T]$, $i_{qt} = q$ and $j_{qt} \sim \mathcal{D}_{\text{entry}}(q)$. The second random seed χ is sampled adaptively based on \mathbf{z} and is used to make some additional adaptive sample queries of Type I (row/column oracle queries).

- If $\mathbf{z}' = \mathcal{A}_{\xi, \chi}^{\text{VRMD1-HP}|\alpha, \epsilon, \delta}(\mathbf{u})$, then wp. $1 - \delta$ over the draw of the random seeds ξ and χ , \mathbf{z}' solves the $(\mathbf{z}, \alpha, \epsilon)$ -minimax-subproblem.
- $\mathcal{A}_{\xi, \chi}^{\text{VRMD1-HP}|\alpha, \epsilon, \delta}$ makes only $\tilde{O}(1)$ batch queries, makes sample queries of Type II only the indices encoded in the seed ξ , and makes sample queries of Type I only on the indices encoded in the seed χ .

Proof The proof follows directly from Theorem 2 of [Carmon et al. \(2019\)](#) and Fact 51. ■

By instantiating CPP (Theorem 50) with VRMD1 (Theorem 57 as the sub-problem solver in the ℓ_2 - ℓ_1 setup, we see that we can obtain the following query complexity trade-off of $\tilde{O}(\alpha/\epsilon)$ full batch queries; along with $\tilde{O}(\|A\|_{2 \rightarrow \infty}^2 \alpha^{-1} \epsilon^{-1})$ non-oblivious sample queries of Type I; and $\tilde{O}(m \|A\|_{2 \rightarrow \infty}^2 \alpha^{-2})$ oblivious sample queries of Type II. This, corresponds to instantiating Meta-Algorithm 1 with $S = \text{Standard}$ and:

- $\mathbf{u}_0 = \operatorname{argmin}_{z \in \mathcal{Z}} r(z)$
- $n_{\text{outer}} = \tilde{O}(\alpha/\epsilon)$
- ζ as defined in Definition 49
- $\mathcal{A}_{\xi, \chi}^{\text{VRMD1-HP}}$ as the sub-solver
- \mathcal{D}_ξ and \mathcal{D}_χ as in Theorem 57
- $f^{\text{sub}}(z)$ as defined in Definition 48

Moreover, observe that Theorem 50 ensures that CPP is ℓ_∞ robust with respect to f^{sub} (in the sense of Definition 2), and VRMD1-HP solves sub-problems to high-precision in the ℓ_∞ norm (in the sense of Definition 1). Thus, using our sample-reuse framework, Theorem 6 implies that we can *reuse* the oblivious sample queries of Type II across all n_{outer} iterations of CPP. We obtain the following improved trade-off.

Theorem 58 (ℓ_2 - ℓ_1 MG trade-off improvement) *For $\alpha > 0$, there is an algorithm that wp. $1 - \delta$ solves the MG problem in the ℓ_2 - ℓ_1 setup using $\tilde{O}(\alpha/\epsilon)$ full batch queries; $\tilde{O}(\|A\|_{2 \rightarrow \infty}^2 \alpha^{-1} \epsilon^{-1})$ sample queries of Type I; and $\tilde{O}(m \|A\|_{2 \rightarrow \infty}^2 \alpha^{-2})$ sample queries of Type II.*

5.2.2. COMPOSITE ℓ_2 - ℓ_2 MATRIX GAMES

In this section, we discuss the composite ℓ_2 - ℓ_2 matrix games setting. Throughout this section, we use the setup $(\mathcal{Z} = (\mathcal{X}, \mathcal{Y}), \|\cdot\|, r)$ where

- $\mathcal{X} = \mathbb{B}^n := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq 1\}$ and $\mathcal{Y} = \mathbb{B}^m := \{\mathbf{y} \in \mathbb{R}^m : \|\mathbf{y}\|_2 \leq 1\}$ are the Euclidean balls of radius 1 centered at the origin.
- $\|\cdot\| : \mathcal{Z} \rightarrow \mathbb{R}$ is given by $\|\mathbf{z}\| = \|\mathbf{z}\|_2$.
- $r : \mathcal{Z} \rightarrow \mathbb{R}$ is given by $r(\mathbf{z}) = \frac{1}{2} \|\mathbf{z}\|_2^2$.

In this case, the relevant constants defined in Section 5.1 are trivially given by

- $L, G \leq \|\mathbf{A}\|_2$
- $D = 1$
- $c = 1, C = d$

The function r is known to be 1-strongly convex with respect to \mathcal{Z} , $\|\cdot\|$ (see, e.g., Section 4.3 of Carmon et al. (2019)). To begin, we define the distributions which we will sample from in order to make sample queries.

Definition 59 Let $\mathbf{A} \in \mathbb{R}^{m \times n}$. For each $i \in [m]$, define the distributions \mathcal{D}_{row} and \mathcal{D}_{col} as follows:

$$\mathbb{P}_{a \sim \mathcal{D}_{\text{row}}} \{a = i\} = \frac{\|\mathbf{a}_i\|_2^2}{\|\mathbf{A}\|_F^2}, \quad \text{and} \quad \mathbb{P}_{a \sim \mathcal{D}_{\text{col}}} \{a = j\} = \frac{\|\mathbf{a}^j\|_2^2}{\|\mathbf{A}\|_F^2}.$$

The following theorem states the guarantees of the variance-reduced mirror descent. VRMD-2 solves a subproblem using $\tilde{O}(1)$ full batch queries and *oblivious* sample queries of Type I (see Definition 54).

Theorem 60 (VRMD2, Lemma 5 of Carmon et al. (2019), adapted - ℓ_2 - ℓ_1 MG sub-solver) Let f be as defined in Definition 52 and $\mathbf{z} \in \mathcal{Z}$. There is a randomized algorithm $\mathcal{A}_{\xi, \chi}^{\text{VRMD2-HP}|\alpha, \epsilon, \delta}$ such that the following holds.

- The algorithm takes in the random seeds $\xi = (\xi_1, \xi_2)$ distributed as follows, for some sufficiently large $T = \tilde{O}(\|\mathbf{A}\|_F^2 / \alpha^2)$. $\xi_1 \sim \{i_1, \dots, i_T\} \subset [m]$ and $\xi_2 \sim \{j_1, \dots, j_T\} \subset [n]$, where for each $t \in [T]$, $i_t \sim \mathcal{D}_{\text{row}}$ and $j_t \sim \mathcal{D}_{\text{col}}$. The second random seed $\chi \sim \text{Unif}[0]$ is a zero-bit random seed.
- If $\mathbf{z}' \mathcal{A}_{\xi, \chi}^{\text{VRMD2-HP}|\alpha, \epsilon, \delta}(\mathbf{u})$, then w.p. $1 - \delta$ over the draw of the random seeds ξ and χ , \mathbf{z}' solves the $(\mathbf{z}, \alpha, \epsilon)$ -minimax-subproblem.
- $\mathcal{A}_{\xi, \chi}^{\text{VRMD2-HP}|\alpha, \epsilon, \delta}$ makes only $\tilde{O}(1)$ batch queries, and makes row/column queries (sample queries of Type I) only the indices encoded in the seed ξ . The algorithm makes no entry oracle queries (sample queries of Type II).

Proof The proof follows directly from Lemma 5 of Carmon et al. (2019) and Fact 51. ■

By instantiating CPP (Theorem 50) with VRMD2 (Theorem 60) as the sub-problem solver in the ℓ_2 - ℓ_2 setting, we see that we can obtain the following query complexity trade-off of $\tilde{O}(\alpha/\epsilon)$ full batch queries along with $\tilde{O}(m \|\mathbf{A}\|_F^2 \alpha^{-1} \epsilon^{-1})$ oblivious sample queries of Type II. This, corresponds to instantiating Meta-Algorithm 1 with $S = \text{Standard}$ and:

- $\mathbf{u}_0 = \text{argmin}_{z \in \mathcal{Z}} r(z)$
- $n_{\text{outer}} = \tilde{O}(\alpha/\epsilon)$
- ζ as defined in Definition 49
- $\mathcal{A}_{\xi, \chi}^{\text{VRMD2-HP}}$ as the sub-solver

- \mathcal{D}_ξ and \mathcal{D}_χ as in Theorem 60
- $f^{\text{sub}}(z) := z_\alpha$ as defined in Definition 48

Moreover, observe that Theorem 50 ensures that CPP is ℓ_∞ robust with respect to f^{sub} (in the sense of Definition 2), and VRMD2-HP solves sub-problems to high-precision in the ℓ_∞ norm (in the sense of Definition 1). Thus, using our sample-reuse framework, Theorem 6 implies that we can *reuse* the oblivious sample queries of Type I across all n_{outer} iterations of CPP. We obtain the following improved trade-off.

Theorem 61 (ℓ_2 - ℓ_2 MG trade-off improvement) *For $\alpha > 0$, there is an algorithm that wp. $1 - \delta$ solves the MG problem in the ℓ_2 - ℓ_2 setup using $\tilde{O}(\alpha/\epsilon)$ full batch queries and $\tilde{O}(\|\mathbf{A}\|_F^2 \alpha^{-2})$ sample queries of Type II.*

Optimal Frobenius-norm-dependent query complexities for ℓ_2 - ℓ_2 matrix games. In the case of ℓ_2 - ℓ_2 games, note that the row/column oracle queries (sample queries of Type I) are *strictly* less powerful than the batch queries (matrix-vector oracle queries) in the sense that one can always use a matrix-vector oracle to implement a row/column oracle. Consequently, setting $\alpha = \|\mathbf{A}\|_F^{2/3} \epsilon^{1/3}$ in Theorem 61, we obtain an overall matrix-vector oracle query complexity of $\tilde{O}(\|\mathbf{A}\|_F^{-2/3} \alpha^{-2/3})$.

Note that lower bounds of Liu et al. (2023) indicate that $\tilde{\Omega}(\|\mathbf{A}\|_F^{-2/3} \alpha^{-2/3})$ -matrix-vector oracle queries is information-theoretically necessary. To our knowledge, Theorem 61 is the first to get this information-theoretically near-optimal rate for general ℓ_2 - ℓ_2 matrix-games.

5.3. Applications of ℓ_2 - ℓ_1 matrix games

In this section, we discuss the implications of our results for two problems in computational geometry. Allen-Zhu et al. (2014); Carmon et al. (2019) showed how to reduce the minimum enclosing ball problem to ℓ_2 - ℓ_1 matrix games. Much of the notation and presentation in this section is inspired by Allen-Zhu et al. (2014) and Carmon et al. (2019).

Maximum inscribed ball. In the maximum inscribed ball problem, we are given a polyhedron specified by a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and vector $\mathbf{b} \in \mathbb{R}^m$ so that $P = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{A}\mathbf{x} + \mathbf{b} \geq \mathbf{0}\}$. We make the following assumptions, as in Carmon et al. (2019); Allen-Zhu et al. (2014):

- We assume that the polytope P is bounded and hence $m \geq n$.
- We assume that $\|\mathbf{a}_i\|_2 = 1$ for all $i \in [m]$ so that $\|\mathbf{A}_{2 \rightarrow \infty}\| = 1$.
- The origin is inside the polytope. This is without loss of generality, as we may always shift the polytope to satisfy this requirement.

In the maximum inscribed ball problem, we must (approximately) compute $\mathbf{x}^* \in P$ such that

$$\mathbf{x}^* \in \operatorname{argmax}_{\mathbf{x} \in P} \min_{i \in [n]} \frac{\langle \mathbf{a}_i, \mathbf{x} \rangle + b_i}{\|\mathbf{a}_i\|_2}. \quad (16)$$

We define

$$r^* := \max_{\mathbf{x} \in P} \min_{i \in [n]} \frac{\langle \mathbf{a}_i, \mathbf{x} \rangle + b_i}{\|\mathbf{a}_i\|_2}, \text{ and}$$

$$R := \min\{r > 0 : P \subset \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq r\}\}.$$

We use $\rho := R/r^*$ to denote the aspect ratio of the problem. [Allen-Zhu et al. \(2014\)](#) showed that solving (16) is equivalent to solving the following minimax problem:

$$\max_{\mathbf{x} \in \mathbb{R}^n} \min_{\mathbf{y} \in \Delta^m} \mathbf{y}^\top \mathbf{A} \mathbf{x} + \mathbf{y}^\top \mathbf{b}.$$

and formulated the maximum inscribed ball approximation problem as follows.

Definition 62 (Maximum inscribed ball (Max-IB)) *In the maximum inscribed ball problem, we must compute $\hat{\mathbf{x}} \in \mathbb{R}^n$ such that wp. $1 - \delta$,*

$$\min_{\mathbf{y} \in \Delta^m} \mathbf{y}^\top \mathbf{A} \hat{\mathbf{x}} + \mathbf{y}^\top \mathbf{b} \geq (1 - \epsilon) \max_{\mathbf{x} \in \mathbb{R}^n} \min_{\mathbf{y} \in \Delta^m} \mathbf{y}^\top \mathbf{A} \mathbf{x} + \mathbf{y}^\top \mathbf{b}.$$

[Carmon et al. \(2019\)](#) further showed that the Max-IB problem can be solved using a two stage approach. In the first stage, we solve ℓ_2 - ℓ_1 matrix games of the following form for a sequence of μ 's:

$$\max_{\mathbf{x} \in \mathbb{B}^n} \min_{\mathbf{y} \in \Delta^m} \mathbf{y}^\top \mathbf{A} \mathbf{x} + \mathbf{y}^\top \mathbf{b} + \mu \sum_{i \in [m]} \mathbf{y}(i) \log(\mathbf{y}(i)) - \frac{\mu}{2} \|\mathbf{x}\|_2^2 \quad (17)$$

to obtain a constant-multiplicative approximation \hat{r} to r^* (Lemma 10 of [Carmon et al. \(2019\)](#)). In the second stage, we solve an ℓ_2 - ℓ_1 matrix games of the following form to $O(\epsilon \hat{r})$ -accuracy (Theorem 3 of [Carmon et al. \(2019\)](#)):

$$\max_{\mathbf{x} \in \mathbb{B}^n} \min_{\mathbf{y} \in \Delta^m} \mathbf{y}^\top \tilde{\mathbf{A}} \mathbf{x} + \mathbf{y}^\top \mathbf{b} \quad (18)$$

where $\tilde{\mathbf{A}} = 2R \cdot \mathbf{A}$. As our MG Problem definition in the ℓ_2 - ℓ_1 setup (Definition 52) captures both (17) and (18), our methods can be used to obtain improved full batch versus sample query complexity trade-offs for solving the Max-IB problem.

Minimum enclosing ball. In the *minimum enclosing ball problem*, we are given a data matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ such that $\mathbf{a}_1 = \mathbf{0}$, and $\max_{i \in [m]} \|\mathbf{a}_i\| = 1$ so that $\|\mathbf{A}\|_{2 \rightarrow \infty} = 1$. We must (approximately) find R^* such that there exists a point \mathbf{x} with $\|\mathbf{x} - \mathbf{a}_i\|_2 \leq R^*$ for all $i \in [m]$. That is,

$$R^* := \min_{\mathbf{x} \in \mathbb{R}^n} \max_{\mathbf{y} \in \Delta^m} \mathbf{y}^\top \mathbf{A} \mathbf{x} + \mathbf{y}^\top \mathbf{b} + \frac{1}{2} \|\mathbf{x}\|_2^2. \quad (19)$$

Definition 63 (Minimum enclosing ball (Min-EB)) *In the minimum enclosing ball problem, we must solve the minimax problem (Definition 47) for $f(\mathbf{x}, \mathbf{y}) = \mathbf{y}^\top \mathbf{A} \mathbf{x} + \mathbf{y}^\top \mathbf{b} + \frac{1}{2} \|\mathbf{x}\|_2^2$.*

[Carmon et al. \(2019\)](#) showed that the Min-EB problem can be solved to accuracy $\epsilon/8$ by solving the following ℓ_2 - ℓ_1 matrix game to accuracy $\epsilon/16$ (Lemma 11 of [Carmon et al. \(2019\)](#)):

$$\min_{\mathbf{x} \in \mathbb{B}^n} \max_{\mathbf{y} \in \Delta^m} \mathbf{y}^\top \mathbf{A} \mathbf{x} + \mathbf{y}^\top \mathbf{b} - \frac{\epsilon}{32 \log(m)} \sum_{i \in [m]} \mathbf{y}(i) \log(\mathbf{y}(i)) + \frac{1}{2} \|\mathbf{x}\|_2^2. \quad (20)$$

As our MG Problem definition in the ℓ_2 - ℓ_1 setup (Definition 52) captures both (20) our methods can be used to obtain improved full batch versus sample query complexity trade-offs for solving the Min-EB problem.

6. Application: Finite-sum minimization with non-uniform smoothness

In this section, we discuss applications of pseudo-independence for improved full batch versus sample query trade-offs for finite sum minimization (FSM) where the component functions are of non-uniform smoothness. This is a generalization of Section 3.

Definition 64 (Generalized FSM (GFSM) problem) *In the GFSM problem, we are given $c > 1$ and $\mathbf{x}_0 \in \mathbb{R}^d$ and must output $\hat{\mathbf{x}} \in \mathbb{R}^d$ such that $F(\hat{\mathbf{x}}) - \min_{\mathbf{x}} F(\mathbf{x}) \leq 1/c \cdot (F(\mathbf{x}_0) - \min_{\mathbf{z}} F(\mathbf{z}))$ where $F : \mathbb{R}^d \rightarrow \mathbb{R}$ with $F : \mathbb{R}^d \rightarrow \mathbb{R}$ with $F(x) := \frac{1}{n} \sum_{i \in [n]} f_i(x)$, F is μ strongly-convex, and each $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is of known smoothness L_i .*

State-of-the-art query complexities for non-uniform smoothness finite-sum minimization can be achieved by using a primal-dual extra-gradient method (Jin et al., 2022). By using this primal-dual extra-gradient method as our sub-problem solver for solving regularized problems and using APP as the outer-solver (Frostig et al., 2015) as in the uniform-smoothness case (Section 3), we can obtain trade-offs between batch and sample queries that depend on the distribution of the smoothness parameters L_i (rather than bounds that depend only on the worst-case smoothness L as in Section 3).

The gradient (batch) oracle and component (sample) oracle are the same as for FSM (Definitions 16 and 17). The FSM sub-problems and f^{sub} are defined exactly as in the uniform smoothness case (Definition 18.) The only change relative to Section 3 is the choice of sub-solver. For the GFSM problem, we use the primal-dual finite-sum minimization algorithm of Jin et al. (2022).

Theorem 65 (PDFSM, Theorem 2 of Jin et al. (2022), restated - GFSM sub-problem solver) *Let F be as in Definition 64. There is a randomized algorithm $\mathcal{A}_{\xi, \chi}^{\text{PDFSM}|\lambda, c, \delta}(\mathbf{x})$ such that the following holds.*

- *The algorithm takes in the random seeds ξ, χ distributed as follows. The first random seed $\xi \sim \{i_1, \dots, i_T\}$ where each $\mathbb{P}[i_t = j] = \sqrt{L_j} / (\sum_{k \in [n]} \sqrt{L_k})$ for some*

$$T = \tilde{O} \left(\sum_{i \in [n]} \sqrt{\frac{L_i}{n\mu}} \right).$$

The second random seed $\chi \sim \text{Unif}[0]$ is a 0-bit random seed.

- *If $\mathbf{x}' = \mathcal{A}_{\xi, \chi}^{\text{PDFSM}|\lambda, c, \delta}(\mathbf{u})$, then wp. $1 - \delta$ over the draw of the random seeds ξ and χ , \mathbf{x}' is a solution to the (\mathbf{u}, λ, c) -sub-problem.*
- *$\mathcal{A}_{\xi, \chi}^{\text{PDFSM}|\lambda, c, \delta}$ makes only $\tilde{O}(1)$ batch queries and makes sample queries only on the indices contained in ξ .*

As in the nonuniform case, by combining Theorem 65 with Theorem 20, we obtain a standard trade-off of $\tilde{O}(\sqrt{\lambda/\mu})$ batch (gradient oracle) queries and $\tilde{O} \left(\sqrt{\lambda/\mu} \cdot \sum_{i \in [n]} \sqrt{\frac{L_i}{n\mu}} \right)$ sample (component oracle) queries for any $\lambda \geq \mu$.

However, as in Section 3, using our sample-reuse framework, we observe that we can *reuse* randomness across all $n_{\text{outer}} = \sqrt{\lambda/\mu}$ sub-problem solves. More concretely, using identical convexity argument as in the proof of Lemma 23, we obtain the following analog of Theorem 65.

Lemma 66 (GFSM sub-problem solver high-precision) *Let F be as in Definition 64. There is a randomized algorithm $\mathcal{A}_{\xi, \chi}^{\text{PDFSM-HP}|\lambda, c, \delta}(\mathbf{x})$ such that the following holds.*

- *The algorithm takes in the random seeds ξ, χ distributed as follows. The first random seed $\xi \sim \{i_1, \dots, i_T\}$ where each $\mathbb{P}[i_t = j] = \sqrt{L_j} / (\sum_{k \in [n]} \sqrt{L_k})$ for some*

$$T = \tilde{O} \left(\sum_{i \in [n]} \sqrt{\frac{L_i}{n\mu}} \right).$$

The second random seed $\chi \sim \text{Unif}[0]$ is a 0-bit random seed.

- *If $\mathbf{x}' = \mathcal{A}_{\xi, \chi}^{\text{SVRG-HP}|\lambda, c, \delta}(\mathbf{u})$, then wp. $1 - \delta$ over the draw of the random seeds ξ and χ ,*

$$\left\| \mathbf{x}' - \bar{f}_{\lambda'}^{\text{sub}}(\mathbf{y}_{\mathbf{u}}) \right\|_{\infty}^2 \leq \frac{1}{c} \cdot \left(F(\mathbf{y}_{\mathbf{u}}) - \min_{\tilde{\mathbf{x}} \in \mathbb{R}^d} F(\tilde{\mathbf{x}}) + \frac{\lambda}{2} \|\tilde{\mathbf{x}} - \mathbf{y}_{\mathbf{u}}\|_2^2 \right).$$

- $\mathcal{A}_{\xi, \chi}^{\text{PDFSM-HP}|\lambda, c, \delta}$ *makes only $\tilde{O}(1)$ batch queries and makes sample queries only on the indices contained in ξ .*

By combining Lemma 66 with Lemma 22 and applying Theorem 6, we obtain a trade-off of $\tilde{O}(\sqrt{\lambda/\mu})$ batch (gradient oracle) queries and $\tilde{O}(\sqrt{\lambda/\mu} \cdot \sum_{i \in [n]} \sqrt{\frac{L_i}{n\mu}})$ sample (component oracle) queries for any $\lambda \geq \mu$. The pseudo-code is analogous to Algorithm 3 except that the inner-loop sub-problem solver is replaced with PDFSM-HP in place of SVRG-HP.

Theorem 67 (Non-uniform smoothness FSM trade-off improvement) *For $\lambda \geq \mu$, there is an algorithm that solves FSM with non-uniform smoothness (Definition 64) using $\tilde{O}(\sqrt{\lambda/\mu})$ full batch queries and only $\tilde{O}(\sum_{i \in [n]} \sqrt{L_i/(n\lambda)})$ sample queries.*

7. Application: Top eigenvector computation

Here, we discuss the setting of top eigenvector (TopEV) computation. This is an interesting specialized setting in which our improvement for finite-sum optimization (Definition 15) can be applied even if the components f_i of the finite sum might not be convex.

Throughout this section, we use $\mathbf{A} \in \mathbb{R}^{n \times d}$ to denote a matrix, and we use $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathbb{R}^d$ to denote its rows. We use $\|\mathbf{A}\|_F := \sqrt{\sum_{i,j} \mathbf{A}_{i,j}^2}$ and $\|\mathbf{A}\|_2$ to denote the spectral norm of \mathbf{A} . We use $\text{sr}(\mathbf{A}) := \sum_i \frac{\lambda_i}{\lambda_1} = \|\mathbf{A}\|_F^2 / \|\mathbf{A}\|_2^2$, where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$ are the eigenvalues of Σ . Note that we always have $\text{sr}(\mathbf{A}) \leq \text{rank}(\mathbf{A})$. We define the relative eigen-gap $\text{gap}(\mathbf{A}) := \frac{\lambda_1 - \lambda_2}{\lambda_1}$.

Definition 68 (TopEV problem) *In the TopEV problem, we are given a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ and $\epsilon > 0$ and must compute a unit vector $\mathbf{x} \in \mathbb{R}^d$ such that $\mathbf{x}^\top \Sigma \mathbf{x} \geq (1 - \epsilon)\lambda_1$ where $\Sigma := \mathbf{A}^\top \mathbf{A} \in \mathbb{R}^{d \times d}$ and λ_1 is the largest eigenvalue of Σ .*

In typical settings (Garber et al., 2016), the goal is to solve this problem with probability inverse polynomial in d (i.e., with high probability in d .) We define the batch and sample queries for solving the problem as follows.

Definition 69 (Matrix-vector oracle - TopEV batch oracle) When queried with $\mathbf{x} \in \mathbb{R}^d$, the matrix-vector oracle returns $\mathbf{A}\mathbf{x}$.

Definition 70 (Row oracle - TopEV sample oracle) When queried, with $i \in [n]$, a row oracle for \mathbf{A} returns $\mathbf{a}_i \in \mathbb{R}^d$.

Reducing top eigenvector computation to solving linear systems in $\lambda'\mathbf{I} - \mathbf{A}^\top\mathbf{A}$ Garber et al. (2016) showed how to reduce top eigenvector computation to performing an approximate version of the classical inverse power method in the shifted matrix $\lambda'\mathbf{I} - \mathbf{A}^\top\mathbf{A}$, where λ' is an appropriately chosen parameter. This is called the *approximate shift-and-invert* power method and can be implemented given access to an oracle that approximately solves linear systems in $\lambda'\mathbf{I} - \mathbf{A}^\top\mathbf{A}$.

Furthermore, Garber et al. (2016) also showed that selecting $(1 + \text{gap}(\mathbf{A})/150)\lambda_1 \leq \lambda \leq (1 + \text{gap}(\mathbf{A})/100)\lambda_1$ is sufficient to ensure that $\tilde{O}(1)$ iterations of approximate shift-and-invert power method is sufficient to solve the top eigenvector problem. Section 6 of Garber et al. (2016) shows that given an algorithm for approximately solving linear systems in $\lambda'\mathbf{I} - \mathbf{A}^\top\mathbf{A}$ for $\lambda' > \lambda_1 + \text{gap}(\mathbf{A})/120$, there is a method for computing a valid shift parameter λ (such that $(1 + \text{gap}(\mathbf{A})/150)\lambda_1 \leq \lambda \leq (1 + \text{gap}(\mathbf{A})/100)\lambda_1$) with runtime and query complexity overhead that is only lower order relative to the approximate shift-and-invert power method steps.

Consequently, in the remainder of this section, we simply focus on the problem of solving linear systems of the form

$$(\lambda'\mathbf{I} - \mathbf{A}^\top\mathbf{A})\mathbf{x} = \mathbf{b},$$

where $\lambda' > \lambda_1 + \Theta(\text{gap}(\mathbf{A}))$ and \mathbf{b} is known, since this is the fundamental subroutine used in the algorithms of Garber et al. (2016). Concretely, we summarize their reductions as follows.

Theorem 71 (Theorems 5, 8, 15, and 30 of Garber et al. (2016)) Given an algorithm $\text{Solve}(\mathbf{x}_0, \mathbf{A}, \lambda', \mathbf{b})$ that computes \mathbf{x} such that $\text{wp. } 1 - \text{poly}(1/d, \text{gap}(\mathbf{A}))$,

$$\left\| \mathbf{x} - (\lambda'\mathbf{I} - \mathbf{A}^\top\mathbf{A})^{-1}\mathbf{b} \right\|_2^2 \leq \text{poly}(1/d, \text{gap}(\mathbf{A})) \left\| \mathbf{x}_0 - (\lambda'\mathbf{I} - \mathbf{A}^\top\mathbf{A})^{-1}\mathbf{b} \right\|_2^2,$$

for any $\lambda > \lambda_1 + \text{gap}(\mathbf{A})/120$, there is an algorithm that solves the TopEV problem with only $\tilde{O}(1)$ calls to Solve .

Thus, in the remainder of this section, we consider the problem of solving linear systems of the form

$$(\lambda'\mathbf{I} - \mathbf{A}^\top\mathbf{A})\mathbf{x} = \mathbf{b}, \tag{21}$$

for some $\mathbf{b} \in \mathbb{R}^n$ and $\lambda' > \lambda_1 + \text{gap}(\mathbf{A})/120$. This can equivalently be viewed as the following minimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{2} \mathbf{x}^\top (\lambda'\mathbf{I} - \mathbf{A}^\top\mathbf{A})\mathbf{x} - \mathbf{b}^\top \mathbf{x} = \min_{\mathbf{x}} \frac{1}{n} \sum_{i \in [n]} \frac{1}{2} \cdot \mathbf{x}^\top (w_i\mathbf{I} - n\mathbf{a}_i\mathbf{a}_i^\top)\mathbf{x} - \mathbf{b}^\top \mathbf{x}, \tag{22}$$

whenever $\sum_i w_i/n = \lambda'$. This is reminiscent of the FSM problem (Definition 15) with $f_i(\mathbf{x}) = \frac{1}{2} \cdot \mathbf{x}^\top (w_i\mathbf{I} - n\mathbf{a}_i\mathbf{a}_i^\top)\mathbf{x} - \mathbf{b}^\top \mathbf{x}$. However, we have the caveat that the matrices $(w_i\mathbf{I} - \mathbf{a}_i\mathbf{a}_i^\top)$ need

not be positive semi-definite; and consequently, the f_i 's are not necessarily convex *even though* F is $\lambda' - \lambda_1 = \Theta(\text{gap}(\mathbf{A}))$ -strongly convex (and $\|\mathbf{A}\|_F^2$ -smooth).

Interestingly, the outer-solver procedure APP (Theorem 20) still applies in this setting; however, the guarantees of SVRG (Theorem 21) unfortunately do not apply *directly* when the f_i are potentially non-convex. Nonetheless, Garber et al. (2016) leveraged problem structure to show that SVRG can be applied for (22). Consequently, we define the **TopEV sub-problem** similarly as in Definition 21 for $F(\mathbf{x}) := f_i(\mathbf{x})$, where

$$F(\mathbf{x}) := f_i(\mathbf{x}) \text{ where } f_i(\mathbf{x}) := \frac{1}{2} \cdot \mathbf{x}^\top \left(w_i \mathbf{I} - n \mathbf{a}^{(i)} \mathbf{a}^{(i)\top} \right) \mathbf{x} - \mathbf{b}^\top \mathbf{x}; w_i := n \cdot \frac{\lambda' \|\mathbf{a}_i\|_2^2}{\|\mathbf{A}\|_F^2}. \quad (23)$$

Note that from this perspective, the batch oracle call for TopEV (Definition 16) exactly corresponds to a gradient oracle call for F , and a sample oracle call for F exactly corresponds to a component oracle call for F (Definition 17.)

Definition 72 (TopEV sub-problem) *Let F be as in (23) and $\rho \geq \lambda' - \lambda_1$. In the (\mathbf{u}, ρ, c) -sub-problem for TopEV, we must solve the (\mathbf{u}, ρ, c) -FSM-sub-problem (Definition 18.)*

Note the similarity between the TopEV sub-problem and the FSM sub-problem from Definition 18. The only difference is that the f_i need not be *convex*. Nonetheless, Garber et al. (2016) provide the following guarantee, which is analogous to Theorem 21 from the FSM setting in Section 3.

Theorem 73 (SVRG, Theorem 2.2 of Frostig et al. (2015), restated - TopEV sub-problem solver)

Let F be as defined in (23) and $\rho \geq (\lambda' - \lambda_1)$. There is a randomized algorithm $\mathcal{A}_{\xi, \chi}^{\text{SVRG-TopEV}|\rho, c, \delta}$ such that the following holds.

- *The algorithm takes in the random seeds ξ, χ distributed as follows. The first random seed $\xi \sim \{i_1, \dots, i_T\}$ where each $\mathbb{P}(i_j = k) = \|\mathbf{a}_k\|_2^2 / \|\mathbf{A}\|_F^2$ for some $T = \tilde{O}(L/\lambda)$. The second random seed $\chi \sim \text{Unif}[0]$ is a 0-bit random seed.*
- *If $\mathbf{x}' = \mathcal{A}_{\xi, \chi}^{\text{SVRG-TopEV}|\rho, c, \delta}(\mathbf{u})$, then wp. $1 - \delta$ over the draw of the random seeds ξ and χ , \mathbf{x}' is a solution to the (\mathbf{u}, ρ, c) -sub-problem.*
- *$\mathcal{A}_{\xi, \chi}^{\text{SVRG-TopEV}|\rho, c, \delta}$ makes only $\tilde{O}(1)$ batch queries and makes sample queries only on the indices encoded in ξ .*

By Theorem 20, we can instantiate APP using SVRG (Theorem 73) with the to solve problems of the form (21) as required by Theorem 71. The pseudocode is the same as Algorithm 2, replacing SVRG-HP with SVRG-TopEV. Setting δ to be inversely polynomial in d , this solves the TopEV problem with high probability in d and obtains a trade-off of $\tilde{O}(\sqrt{\rho/(\lambda' - \lambda_1)})$ full batch (matrix-vector oracle) queries and

$$\tilde{O} \left(\frac{\rho^2 + 12\lambda_1 \|\mathbf{A}\|_F^2}{(\rho - \lambda_1 + \lambda)^2} \right) \cdot \sqrt{\frac{\rho}{\lambda' - \lambda_1}}$$

sample (row oracle) queries for any $\rho \geq \lambda' - \lambda_1$.

Now, since F in this case remains smooth and strongly convex (even though the f_i 's may not be convex), using the *exact* same arguments as in Section 3 (Lemmas 22 and Lemma 23) we can use our sample reuse framework from Section 2 (Theorem 6) to obtain the following improved trade-off. (Again, the pseudocode is the same as Algorithm 3, replacing SVRG-HP with SVRG-TopEV.)

Theorem 74 (TopEV trade-off improvement) *For $\lambda \geq \mu$, there is an algorithm that solves FSM with high probability in d using only $\tilde{O}(\sqrt{\rho/(\lambda' - \lambda_1)})$ full batch queries and only $\tilde{O}\left(\frac{\rho^2 + 12\lambda_1 \|\mathbf{A}\|_F^2}{(\rho - \lambda_1 + \lambda)^2}\right)$ sample queries.*

Finally, we remark that to obtain the trade-offs reported in Table 2, we simply make the change of variables $\rho = \alpha\lambda_1$. With this change of variables,

$$\sqrt{\frac{\rho}{\lambda' - \lambda_1}} = \sqrt{\frac{\alpha}{\text{gap}(\mathbf{A})}}, \quad \text{and} \quad \frac{\rho^2 + 12\lambda_1 \|\mathbf{A}\|_F^2}{(\rho - \lambda_1 + \lambda)^2} = \tilde{O}\left(\frac{\alpha^2\lambda_1^2 + 12\lambda_1 \|\mathbf{A}\|_F^2}{\alpha^2\lambda_1^2}\right) = \tilde{O}\left(\frac{\text{sr}(\mathbf{A})}{\alpha^2}\right).$$

8. Conclusion

We introduced a sample reuse framework to reuse randomness across multiple subproblem solutions in variance-reduced optimization methods without sacrificing theoretical correctness guarantees. Our results enabled improved query complexity trade-offs for a broad range of optimization problems. One limitation of our current sample reuse framework and analysis of pseudo-independence is that it only allows us to reuse samples across *high-accuracy* subroutines. Although this already enables improvements for several problems, as seen in Table 2, in future work it may be interesting to study whether tighter analysis allows sample reuse even for sub-routines which do not solve to high-accuracy. This could have applications, for instance, to non-convex optimization problems. As mentioned in Section 1.1, another limitation is that our results don't directly yield a worst-case asymptotic-runtime improvement that we are aware of; however it sheds new light on the information needed to solve FSM and could yield faster algorithms depending on caching and memory layout. Building on this approach to obtain efficiency gains, e.g., in distributed computing settings, is an interesting direction for future research.

Acknowledgments

We thank anonymous reviewers for their feedback. Yujia Jin and Ishani Karmarkar were funded in part by NSF CAREER Award CCF-1844855, NSF Grant CCF-1955039, and a PayPal research award. Aaron Sidford was funded in part by a Microsoft Research Faculty Fellowship, NSF CAREER Award CCF-1844855, NSF Grant CCF1955039, and a PayPal research award. Yujia Jin's contributions to the project occurred while she was a graduate student at Stanford.

References

- Alekh Agarwal and Leon Bottou. A lower bound for the optimization of finite sums. In *32nd International Conference on Machine Learning (ICML)*, 2015.
- Zeyuan Allen-Zhu, Zhenyu Liao, and Yang Yuan. Optimization algorithms for faster computational geometry. In *41st International Colloquium on Automata, Languages and Programming (ICALP)*, 2014.

- Hilal Asi and John C Duchi. Near instance-optimality in differential privacy. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, 2020.
- Amos Beimel, Haim Kaplan, Yishay Mansour, Kobbi Nissim, Thatchaphol Saranurak, and Uri Stemmer. Dynamic algorithms against an adaptive adversary: generic constructions and lower bounds. In *54th Annual ACM Symposium on Theory of Computing (STOC)*, 2022.
- Dimitri P Bertsekas. Approximate policy iteration: A survey and some new methods. In *Journal of Control Theory and Applications*, 2011.
- Vladimir Braverman, Robert Krauthgamer, Aditya Krishnan, and Shay Sapir. Lower bounds for pseudo-deterministic counting in a stream. *arXiv preprint arXiv:2303.16287*, 2023.
- Yair Carmon, Yujia Jin, Aaron Sidford, and Kevin Tian. Variance reduction for matrix games. *Advances in Neural Information Processing Systems 32 (NeurIPS)*, 2019.
- Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *63rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2022.
- Yeshwanth Cherapanamjeri and Jelani Nelson. On adaptive distance estimation. In *Advances in Neural Information Processing Systems*, 2020.
- Edith Cohen, Jelani Nelson, Tamás Sarlós, Mihir Singhal, and Uri Stemmer. One attack to rule them all: Tight quadratic bounds for adaptive queries on cardinality sketches. In *arXiv preprint arXiv:2411.06370*, 2024.
- Michael B Cohen, Yin Tat Lee, Cameron Musco, Christopher Musco, Richard Peng, and Aaron Sidford. Uniform sampling for matrix approximation. In *6th Conference on Innovations in Theoretical Computer Science (ITCS)*, 2015.
- Michael B. Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. In *Journal of the ACM*, 2020a.
- Michael B Cohen, Cameron Musco, and Jakub Pachocki. Online row sampling. In *Theory of Computing*, 2020b.
- Peter Dixon, Aduri Pavan, Jason Vander Woude, and NV Vinodchandran. Pseudodeterminism: promises and lowerbounds. In *54th Annual ACM Symposium on Theory of Computing (STOC)*, 2022.
- Matteo Fischetti, Iacopo Mandatelli, and Domenico Salvagnin. Faster sgd training by minibatch persistency. In *arXiv preprint arXiv:1806.07353*, 2018.
- Roy Frostig, Rong Ge, Sham Kakade, and Aaron Sidford. Un-regularizing: approximate proximal point and faster stochastic algorithms for empirical risk minimization. In *32nd International Conference on Machine Learning (ICML)*, 2015.
- Dan Garber, Elad Hazan, Chi Jin, Cameron Musco, Praneeth Netrapalli, Aaron Sidford, et al. Faster eigenvector computation via shift-and-invert preconditioning. In *33rd International Conference on Machine Learning (ICML)*, 2016.

- Erann Gat and Shafi Goldwasser. Probabilistic search algorithms with unique answers and their cryptographic applications. In *Electronic Colloquium on Computational Complexity: ECCC*, 2011.
- Badih Ghazi, Ravi Kumar, Pasin Manurangsi, and Jelani Nelson. Differentially private all-pairs shortest path distances: Improved algorithms and lower bounds. In *33rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2022.
- Shafi Goldwasser and Ofer Grossman. Bipartite perfect matching in pseudo-deterministic nc. In *44th International Colloquium on Automata, Languages and Programming (ICALP)*, 2017.
- Shafi Goldwasser, Ofer Grossman, and Dhiraj Holden. Pseudo-deterministic proofs. In *arXiv preprint arXiv:1706.04641*, 2017.
- Ivo Grondman, Lucian Busoniu, Gabriel AD Lopes, and Robert Babuska. A survey of actor-critic reinforcement learning: Standard and natural policy gradients. In *IEEE Transactions on Systems, Man, and Cybernetics, part C (applications and reviews)*, 2012.
- Ofer Grossman, Meghal Gupta, and Mark Sellke. Tight space lower bound for pseudo-deterministic approximate counting. In *64th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2023.
- Russell Impagliazzo, Rex Lei, Toniann Pitassi, and Jessica Sorrell. Reproducibility in learning. In *Proceedings of the 54th annual ACM SIGACT symposium on theory of computing*, 2022.
- Shunhua Jiang, Zhao Song, Omri Weinstein, and Hengjie Zhang. A faster algorithm for solving general lps. In *53rd Annual ACM Symposium on Theory of Computing (STOC)*, 2021.
- Yujia Jin and Aaron Sidford. Towards tight bounds on the sample complexity of average-reward mdps. In *38th International Conference on Machine Learning (ICML)*, 2021.
- Yujia Jin, Aaron Sidford, and Kevin Tian. Sharper rates for separable minimax and finite sum optimization via primal-dual extragradient methods. In *35th Annual Conference on Computational Learning Theory (COLT)*, 2022.
- Yujia Jin, Ishani Karmarkar, Aaron Sidford, and Jiayi Wang. Truncated variance reduced value iteration. In *arXiv preprint arXiv:2405.12952*, 2024.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems 26 (NeurIPS)*, 2013.
- Ishani Karmarkar, Liam O’Carroll, and Aaron Sidford. Solving zero-sum games with fewer matrix-vector products. In *Foundations of Computer Science*, 2025.
- Michael Kearns and Satinder Singh. Finite-sample convergence rates for q-learning and indirect algorithms. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, 1998.
- Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in $o(\text{vrnk})$ iterations and faster algorithms for maximum flow. In *55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014.

- Yin Tat Lee and Aaron Sidford. Efficient inverse maintenance and faster algorithms for linear programming. In *56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2015.
- Hongzhou Lin, Julien Mairal, and Zaid Harchaoui. A universal catalyst for first-order optimization. In *Advances in Neural Information Processing Systems 28 (NeurIPS)*, 2015.
- Michael L Littman, Thomas L Dean, and Leslie Pack Kaelbling. On the complexity of solving markov decision problems. In *11th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, 1995.
- Yuanshi Liu, Hanzhen Zhao, Yang Xu, Pengyun Yue, and Cong Fang. Accelerated gradient algorithms with adaptive subspace search for instance-faster optimization. In *arXiv preprint arXiv:2312.03218*, 2023.
- Arkadi Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. In *SIAM Journal on Optimization*, 2004.
- Sebastien Roch. Modern discrete probability: An essential toolkit. In *Cambridge Series in Statistical and Probabilistic Mathematics*, 2024.
- Yousef Saad. Numerical methods for large eigenvalue problems: revised edition. In *SIAM*, 2011.
- Scott Sallinen, Nadathur Satish, Mikhail Smelyanskiy, Samantika S Sury, and Christopher Ré. High performance parallel stochastic gradient descent in shared memory. In *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2016.
- Aaron Sidford, Mengdi Wang, Xian Wu, Lin Yang, and Yinyu Ye. Near-optimal time and sample complexities for solving markov decision processes with a generative model. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*, 2018a.
- Aaron Sidford, Mengdi Wang, Xian Wu, and Yinyu Ye. Variance reduced value iteration and faster algorithms for solving markov decision processes. In *29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2018b.
- Aaron Sidford, Mengdi Wang, Xian Wu, and Yinyu Ye. Variance reduced value iteration and faster algorithms for solving markov decision processes. In *Naval Research Logistics (NRL)*, 2023.
- Paul Tseng. Solving h-horizon, stationary markov decision problems in time proportional to $\log(h)$. In *Operations Research Letters*, 1990.
- Jan van den Brand, Yin Tat Lee, Yang P. Liu, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. Minimum cost flows, mdps, and l1-regression in nearly linear time for dense instances. In *53rd Annual ACM Symposium on Theory of Computing (STOC)*, 2021.
- Jan van den Brand, Yu Gao, Arun Jambulapati, Yin Tat Lee, Yang P. Liu, Richard Peng, and Aaron Sidford. Faster maxflow via improved dynamic spectral vertex sparsifiers. In *54th Annual ACM Symposium on Theory of Computing (STOC)*, 2022.

Blake E Woodworth and Nati Srebro. Tight complexity bounds for optimizing composite objectives. In *Advances in Neural Information Processing Systems 29 (NeurIPS)*, 2016.

Blake E Woodworth, Kumar Kshitij Patel, and Nati Srebro. Minibatch vs local sgd for heterogeneous distributed learning. In *Advances in Neural Information Processing Systems*, 2020.

Appendix A. Inducing pseudoindependence numerically stably

The pseudo-independent algorithm constructed in the proof of Theorem 12 is simple; however, it may not be directly implementable in finite precision, as it requires infinite precision to directly implement the addition of uniform random noise. The proof of Theorem 75 below provides an alternative construction that can be implemented in finite precision.

Theorem 75 (Finite precision analog of Theorem 12) *Let $\epsilon, \delta \in (0, 1)$, $\eta > 0$, $\eta' := \min\left(\frac{\eta}{2}, \frac{\eta\epsilon}{p}\right)$, and let $\mathcal{A}_{\xi, \chi}$ be a randomized algorithm that is an (η', δ) -approximation of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}^p$. There is a numerically stable algorithm $\mathcal{A}'_{\xi, \chi'}$ such that $\mathcal{A}'_{\xi, \chi'}$ is an (ϵ, δ) -pseudo-independent of ξ and an (η, δ) -approximation of f with the same runtime and query complexities as $\mathcal{A}_{\xi, \chi}$ up to an additive $O(p)$ in runtime.*

Proof Consider $\mathcal{S} \subset \mathbb{R}^p$ to be a β -covering of \mathbb{R}^p in the ℓ_∞ norm. Define $\text{round} : \mathbb{R}^p \rightarrow \mathcal{S}$ to be the operator that maps $\mathbf{x} \in \mathbb{R}^p$ to some $\mathbf{x}' \in \mathbb{R}^p$ such that $0 \leq \mathbf{x}(i) - \mathbf{x}'(i) \leq \beta$ for all $i \in [d]$. Define $\text{smooth}_{\nu, t}$ to be the operator which uses a random seed $\nu \sim \mathcal{D}_\nu$ to map $\mathbf{x} \in \mathcal{S}$ to a uniformly random $\mathbf{x}' \in \{y \in \mathcal{S} : -t \leq (\mathbf{x}(i) - \mathbf{x}'(i)) \leq t\}$. Here, t is a parameter that will be specified later in the proof.

Let $\mathcal{A}'_{\xi, \chi'}$ be the randomized algorithm which takes input $\mathbf{x} \in \mathbb{R}^d$, random seed $\xi \sim \mathcal{D}_\xi$, and χ where $\chi = (\chi', \nu) \sim \mathcal{D}_\chi^{\mathbf{x}}$ is the concatenation of an independently drawn seed $\chi' \sim \mathcal{D}_{\chi'}$ and seed $\nu \sim \mathcal{D}_\nu$. For any realization s, c, n of ξ, χ', ν , let $\mathcal{A}'_{\xi=s, \chi'=(c, e)}(\mathbf{x}) = \text{smooth}_{\nu=e, t}(\text{round}(\mathcal{A}_{\xi=s, \chi=c}(\mathbf{x})))$.

First, we'll construct a smoothing for $\mathcal{A}'_{\xi, \chi'}$. Let $\bar{\mathcal{A}}_{\chi'}$ be the randomized algorithm which takes input $\mathbf{x} \in \mathbb{R}^d$ and a random seed $\chi' \sim \mathcal{D}_{\chi'}$. For any realization c, n of χ', ν , we define the mapping $\bar{\mathcal{A}}_{\chi'=(c, e)}(\mathbf{x}) = \text{smooth}_{\nu=e, t}(\text{round}(f(\mathbf{x})))$. Now, using that $\mathcal{A}_{\xi, \chi}$ is an (η', δ) -approximation of f , we can conclude that

$$\mathbb{P}_{\xi \sim \mathcal{D}_\xi, \chi \sim \mathcal{D}_\chi^{\mathbf{x}}} (\|\mathcal{A}_{\xi, \chi}(\mathbf{x}) - f(\mathbf{x})\|_\infty \leq \eta') \geq 1 - \delta, \quad (24)$$

and hence

$$\mathbb{P}_{\xi \sim \mathcal{D}_\xi, \chi' \sim \mathcal{D}_{\chi'}} (\|\mathcal{A}'_{\xi, \chi'}(\mathbf{x}) - f(\mathbf{x})\|_\infty \leq (\eta' + t)) \geq 1 - \delta. \quad (25)$$

Thus, $\mathcal{A}'_{\xi, \chi'}$ is an $(\eta' + t, \delta)$ -approximation of f . In particular, setting $t = \tau = \max\left(\frac{\eta'}{2}, \frac{\eta'p}{2\epsilon}\right)$ and recalling $\eta' = \min\left(\frac{\eta}{2}, \frac{\eta\epsilon}{p}\right)$, we have

$$\eta' + t \leq \eta' \max\left(\frac{3}{2}, \frac{p}{\epsilon}\right) \leq \eta$$

and hence $\mathcal{A}'_{\xi, \chi'}$ is an (η, δ) -approximation of f . To bound the total variation distance, note that by (24), we have that with probability $1 - \delta$, $d_{TV} \left(p_{A_{\xi=s, \chi}(x)}, p_{\bar{A}_\chi(x)} \right)$ is the total variation distance between q, q' where q is the distribution of a $\text{smooth}_{\nu, t}(0)$ and q' is the distribution of $\text{smooth}_{\nu, t}(\alpha)$ where α is ℓ_∞ bounded by η' . Using that $t \geq \eta'/2$, we have that

$$\begin{aligned} d_{TV}(q, q') &= 1 - \prod_{i \in [p]} \left(\frac{\max \left(0, \left\lfloor \frac{2t - |\alpha_i|}{\beta} \right\rfloor \cdot \beta \right)}{\lceil 2t/\beta \rceil \cdot \beta} \right) \leq 1 - \prod_{i \in [p]} \frac{\lfloor \frac{2t - \eta'}{\beta} \rfloor}{\lceil 2t/\beta \rceil} \leq 1 - \left(\frac{2t - \eta' - \beta}{2t + \beta} \right)^p \\ &= 1 - \left(1 - \frac{\eta' - 2\beta}{2t + \beta} \right)^p \leq \left(\frac{\eta' + 2\beta}{2t} \right)^p \end{aligned}$$

where the final step used Bernoulli's inequality and that $\beta \geq 0$. Thus, setting $t = \tau = \max \left(\frac{\eta'}{2}, \frac{(\eta' + 2\beta)p}{2\epsilon} \right)$, we have that with probability $1 - \delta$,

$$\mathbb{P}_{s \sim \mathcal{D}_\chi} \left\{ d_{TV} \left(p_{A'_{\xi=s, \chi'}(x)}, p_{\bar{A}_\chi(x)} \right) \leq \epsilon \right\} \geq 1 - \delta.$$

For the second guarantee, note that $\mathcal{A}'_{\xi, \chi'}$ has the same runtime and query complexities as $\mathcal{A}_{\xi, \chi}$ up to an additive $O(p)$ increase in the runtime due to the cost of performing the p -dimensional random perturbation induced by ν . \blacksquare

Appendix B. Pseudoindependence and repeated compositions

In this section, our goal is to prove the following theorem (see Definition 13 for related notation.)

Lemma 76 *Let $\mathcal{A}'_{\xi, \chi'}$ be randomized algorithm which takes an input $\mathbf{u} \in \mathbb{R}^d$ and two random seeds $\xi \sim \mathcal{D}_\xi, \chi' \sim \mathcal{D}_{\chi'}$ and is (ϵ, δ) -pseudo-independent of ξ . Then,*

$$d_{TV} \left(p_{\Phi_{\mathcal{A}'}^T(\mathbf{u}; s, \mathcal{D}_{\chi'})}, p_{\Phi_{\mathcal{A}'}^T(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'})} \right) \leq 2T(\delta + \epsilon).$$

Before proving Lemma 14, we state the following fact about transformations of random variables.

Fact 77 *Let A and B be random variables in \mathcal{Q} and let ζ be a deterministic function $\zeta : \mathcal{Q} \rightarrow \mathcal{Q}$. Then, $d_{TV} (p_{\zeta(A)}, p_{\zeta(B)}) \leq d_{TV} (p_A, p_B)$.*

Proof For any deterministic function f , let $f^{-1}(\cdot)$ denote the preimage of f . That is, for any $T \subset \mathcal{Q}$, let $f^{-1}(T) := \{\omega' \in \Omega : f(\omega') \in T\}$. Then, by the definition of the total variation distance, we have

$$\begin{aligned} d_{TV} (p_{\zeta(A)}, p_{\zeta(B)}) &= \sup_{S \subset \mathcal{Q}} |\mathbb{P} \{\zeta(A) \in S\} - \mathbb{P} \{\zeta(B) \in S\}| \\ &= \sup_{S \subset \mathcal{Q}} |\mathbb{P} \{A \in \zeta^{-1}(S)\} - \mathbb{P} \{B \in \zeta^{-1}(S)\}| \\ &\leq \sup_{T \subset \mathcal{Q}} |\mathbb{P} \{A \in T\} - \mathbb{P} \{B \in T\}| = d_{TV} (p_A, p_B). \end{aligned}$$

■

Now, to prove Lemma 14, we first bound the TV distance between $p_{\Phi_{\mathcal{A}'}}^T(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'})$ and $p_{H_{\bar{\mathcal{A}}}}^T(\mathbf{u}; \mathcal{D}_{\chi'})$ (recall Definition 13.)

Lemma 78 *Let $\mathcal{A}'_{\xi, \chi'}$ be a randomized algorithm which takes an input $\mathbf{u} \in \mathbb{R}^d$ and two random seeds $\xi \sim \mathcal{D}_\xi, \chi' \sim \mathcal{D}_{\chi'}$. Suppose $\mathcal{A}'_{\xi, \chi'}$ is (ϵ, δ) -pseudo-independent and that $\bar{\mathcal{A}}_{\chi'}$ is an (ϵ, δ) -smoothing of \mathcal{A}' with respect to ξ . Let $s \sim \mathcal{D}_\xi$. Then,*

$$d_{TV} \left(p_{\Phi_{\mathcal{A}'}}^T(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}), p_{H_{\bar{\mathcal{A}}}}^T(\mathbf{u}; \mathcal{D}_{\chi'}) \right) \leq T(\delta + \epsilon).$$

Proof Induct on T . If $T = 1$, the statement essentially follows from the definition of pseudo-independence, Fact 7, and a union bound, as we elaborate in the next few sentences. Concretely, we have that with probability $1 - \delta$ over the draw of $s \sim \mathcal{D}_\xi$,

$$d_{TV} \left(p_{\mathcal{A}'_{\xi=s, \chi'}}(\mathbf{u}), p_{\bar{\mathcal{A}}_{\chi'}}(\mathbf{u}) \right) \leq \epsilon.$$

Since ζ is a deterministic function, using Fact 77, we have

$$d_{TV} \left(p_{\zeta(\mathbf{u}, \mathcal{A}'_{\xi=s, \chi'})}(\mathbf{u}), p_{\zeta(\mathbf{u}, \bar{\mathcal{A}}_{\chi'})}(\mathbf{u}) \right) \leq d_{TV} \left(p_{\mathcal{A}'_{\xi=s, \chi'}}(\mathbf{u}), p_{\bar{\mathcal{A}}_{\chi'}}(\mathbf{u}) \right) \leq \epsilon.$$

Consequently, by Fact 7 and union bound,

$$d_{TV} \left(p_{\Phi_{\mathcal{A}'}}^1(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}), p_{H_{\bar{\mathcal{A}}}}^1(\mathbf{u}; \mathcal{D}_{\chi'}) \right) \leq (\delta + \epsilon).$$

This completes the base case.

For the inductive step, suppose the theorem holds up to $T - 1$. That is, suppose that

$$d_{TV} \left(p_{\Phi_{\mathcal{A}'}}^{T-1}(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}), p_{H_{\bar{\mathcal{A}}}}^{T-1}(\mathbf{u}; \mathcal{D}_{\chi'}) \right) \leq (T - 1)(\delta + \epsilon).$$

Then, by Fact 7, there exists an event E_1 and random variables C, D, F such that conditioned on E_1 , $\Phi_{\mathcal{A}'}^{T-1}(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}) \stackrel{D}{=} C \stackrel{D}{=} H_{\bar{\mathcal{A}}}^{T-1}(\mathbf{u}; \mathcal{D}_{\chi'})$ and $\mathbb{P}\{E_1\} \geq 1 - (T - 1)(\delta + \epsilon)$. Consequently, for any event E , we have

$$\begin{aligned} \left| p_{\Phi_{\mathcal{A}'}}^T(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'})(E) - p_{H_{\bar{\mathcal{A}}}}^T(\mathbf{u}; \mathcal{D}_{\chi'})(E) \right| &\leq \left| p_{\Phi_{\mathcal{A}'}}^T(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'})|_{E_1}(E) - p_{H_{\bar{\mathcal{A}}}}^T(\mathbf{u}; \mathcal{D}_{\chi'})|_{E_1}(E) \right| + \mathbb{P}\{\neg E_1\} \\ &= \left| p_{\Phi_{\mathcal{A}'}}^T(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'})|_{E_1}(E) - p_{H_{\bar{\mathcal{A}}}}^T(\mathbf{u}; \mathcal{D}_{\chi'})|_{E_1}(E) \right| + (T - 1)(\delta + \epsilon) \end{aligned}$$

To bound the first term, we observe that by the definition of Φ and H , we have

$$\begin{aligned} &\left| p_{\Phi_{\mathcal{A}'}}^T(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'})|_{E_1}(E) - p_{H_{\bar{\mathcal{A}}}}^T(\mathbf{u}; \mathcal{D}_{\chi'})|_{E_1}(E) \right| \\ &= \left| p_{\zeta(\Phi_{\mathcal{A}'}^{T-1}(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}), \mathcal{A}'_{\xi, \chi'}(\Phi_{\mathcal{A}'}^{T-1}(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'})))|_{E_1}(E) - p_{\zeta(H_{\bar{\mathcal{A}}}^{T-1}(\mathbf{u}; \mathcal{D}_{\chi'}), \bar{\mathcal{A}}_{\chi'}(H_{\bar{\mathcal{A}}}^{T-1}(\mathbf{u}; \mathcal{D}_{\chi'})))|_{E_1}(E)} \right| \\ &\leq \sup_{\mathbf{c} \in \mathbb{R}^d} \left| p_{\zeta(\mathbf{c}, \mathcal{A}'_{\xi, \chi'}(\mathbf{c}))} - p_{\zeta(\mathbf{c}, \bar{\mathcal{A}}_{\chi'}(\mathbf{c}))} \right|, \end{aligned}$$

where in the last line, we used the fact that conditional on E_1 , $\Phi^{T-1}(x; \mathcal{D}_\xi, \mathcal{D}_{\chi'}) \stackrel{\mathcal{D}}{=} C \stackrel{\mathcal{D}}{=} H^{T-1}(x; \mathcal{D}_{\chi'})$. By Fact 77 and the definition of pseudo-independence (using an identical argument to the base case) we have

$$\sup_{\mathbf{c} \in \mathbb{R}^d} \left| p_{\zeta}(\mathbf{c}; \mathcal{A}'_{\xi, \chi'}(\mathbf{c})) - p_{\zeta}(\mathbf{c}; \bar{\mathcal{A}}_{\chi'}(\mathbf{c})) \right| \leq (\delta + \epsilon).$$

Thus, by substitution, we conclude that for any event E ,

$$\left| p_{\Phi_{\mathcal{A}'}}(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'})|_{E_1}(E) - p_{H_{\bar{\mathcal{A}}}}(\mathbf{u}; \mathcal{D}_{\chi'})|_{E_1}(E) \right| \leq (\delta + \epsilon),$$

and consequently,

$$\left| p_{\Phi_{\mathcal{A}'}}(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'})(E) - p_{H_{\bar{\mathcal{A}}}}(\mathbf{u}; \mathcal{D}_{\chi'})(E) \right| \leq T(\delta + \epsilon).$$

■

The following lemma obtains the analogous bound as Lemma 78 when the same random seed ξ is reused across iterations.

Lemma 79 *Let $\mathcal{A}'_{\xi, \chi'}$ be a randomized algorithm which takes an input $x \in \mathbb{R}^d$ and two random seeds $\xi \sim \mathcal{D}_\xi, \chi' \sim \mathcal{D}_{\chi'}^{\mathbf{u}}$. Let $\bar{\mathcal{A}}_{\chi'}$ be a (ϵ, δ) -smoothing of \mathcal{A}' with respect to ξ . Let $s \sim \mathcal{D}_\xi$. Then,*

$$d_{TV} \left(p_{\Phi_{\mathcal{A}'}}(\mathbf{u}; s, \mathcal{D}_{\chi'}), p_{H_{\bar{\mathcal{A}}}}(\mathbf{u}; \mathcal{D}_{\chi'}) \right) \leq T(\delta + \epsilon).$$

Proof By Fact 7 and Fact 77, with probability $1 - T\delta$ over the draw of s , there exist random variables C^t and events E_t for $t = \{1, \dots, T\}$ so that $\mathbb{P}\{E_t\} \geq 1 - \epsilon$, and conditioned on E_1, \dots, E_t ,

$$\zeta(\mathbf{u}, \mathcal{A}'_{\xi=s, \chi'}(\mathbf{u})) \stackrel{\mathcal{D}}{=} C^1 \stackrel{\mathcal{D}}{=} \zeta(\mathbf{u}, \bar{\mathcal{A}}_{\chi'}(\mathbf{u})),$$

and

$$\begin{aligned} \Phi_{\mathcal{A}'}^t(\mathbf{u}; s, \mathcal{D}_{\chi'}) &\stackrel{\mathcal{D}}{=} \zeta(C^{t-1}, \mathcal{A}'_{\xi=s, \chi'}(C^{t-1})) \stackrel{\mathcal{D}}{=} C^t, \\ H_{\bar{\mathcal{A}}}^t(\mathbf{u}; \mathcal{D}_{\chi'}) &\stackrel{\mathcal{D}}{=} \zeta(C^{t-1}, \bar{\mathcal{A}}_{\chi'}(C^{t-1})) \stackrel{\mathcal{D}}{=} C^t. \end{aligned}$$

for every $t \in [T]$. So, let $E' := \bigwedge_{t \in [T]} E_t$, and note that $\mathbb{P}\{E'\} \geq 1 - \epsilon T$. Consequently,

$$d_{TV} \left(p_{\Phi_{\mathcal{A}'}}(x; s, \mathcal{D}_{\chi'}), p_{H_{\bar{\mathcal{A}}}}(x; \mathcal{D}_{\chi'}) \right) \leq T\delta + (1 - \mathbb{P}\{E'\}) \leq T(\delta + \epsilon).$$

■

Lemma 80 *Let $\mathcal{A}'_{\xi, \chi'}$ be randomized algorithm which takes an input $\mathbf{u} \in \mathbb{R}^d$ and two random seeds $\xi \sim \mathcal{D}_\xi, \chi' \sim \mathcal{D}_{\chi'}^{\mathbf{u}}$ and is (ϵ, δ) -pseudo-independent of ξ . Then,*

$$d_{TV} \left(p_{\Phi_{\mathcal{A}'}}(\mathbf{u}; s, \mathcal{D}_{\chi'}), p_{\Phi_{\mathcal{A}'}}(\mathbf{u}; \mathcal{D}_\xi, \mathcal{D}_{\chi'}) \right) \leq 2T(\delta + \epsilon).$$

Proof The result follows directly by applying triangle inequality, Lemma 79, and Lemma 78. ■