# Sublinear Classical and Quantum Algorithms for General Matrix Games 

Tongyang Li, ${ }^{* 1,2}$ Chunhao Wang, ${ }^{* 3,4}$ Shouvanik Chakrabarti, ${ }^{1}$ Xiaodi Wu ${ }^{1}$<br>${ }^{1}$ Joint Center for Quantum Information and Computer Science, Department of Computer Science, and Institute for Advanced Computer Studies, University of Maryland<br>${ }^{2}$ Center for Theoretical Physics, Massachusetts Institute of Technology<br>${ }^{3}$ Department of Computer Science and Engineering, Pennsylvania State University<br>${ }^{4}$ Department of Computer Science, University of Texas at Austin<br>tongyang@mit.edu, cwang@psu.edu, \{shouv,xwu\}@cs.umd.edu


#### Abstract

We investigate sublinear classical and quantum algorithms for matrix games, a fundamental problem in optimization and machine learning, with provable guarantees. Given a matrix, sublinear algorithms for the matrix game were previously known only for two special cases: (1) the maximizing vectors live in the L1-norm unit ball, and (2) the minimizing vectors live in either the L1- or the L2-norm unit ball. We give a sublinear classical algorithm that can interpolate smoothly between these two cases: for any fixed $q$ between 1 and 2 , we solve, within some additive error, matrix games where the minimizing vectors are in an Lq-norm unit ball. We also provide a corresponding sublinear quantum algorithm that solves the same task with a quadratic improvement in dimensions of the maximizing and minimizing vectors. Both our classical and quantum algorithms are optimal in the dimension parameters up to poly-logarithmic factors. Finally, we propose sublinear classical and quantum algorithms for the approximate Carathéodory problem and the Lq-margin support vector machines as applications.


## Introduction

Motivations. Minimax games between two parties, i.e., $\min _{x} \max _{y} f(x, y)$, is a basic model in game theory and has ubiquitous connections and applications to economics, optimization and machine learning, theoretical computer science, etc. Among minimax games, one of the most fundamental cases is the bilinear minimax game, also known as the matrix game, with the following form:

$$
\begin{equation*}
\min _{x \in \mathcal{X}} \max _{y \in \mathcal{Y}} y^{\top} A x, \text { where } A \in \mathbb{R}^{n \times d}, \mathcal{X} \subset \mathbb{R}^{d}, \mathcal{Y} \subset \mathbb{R}^{n} \tag{1}
\end{equation*}
$$

Matrix games are fundamental in algorithm design due to their equivalence to linear programs (Dantzig 1998), and also in machine learning because they contain classification (Novikoff 1963; Minsky and Papert 1988) as a special case, and many other important problems.

For many common domains $\mathcal{X}$ and $\mathcal{Y}$, matrix games can be solved efficiently within approximation error $\epsilon$, i.e., to output $x^{\prime} \in \mathcal{X}$ and $y^{\prime} \in \mathcal{Y}$ such that $\left(y^{\prime}\right)^{\top} A x^{\prime}$ is $\epsilon$-close to the

[^0]optimum in (1). For some specific choices of $\mathcal{X}$ and $\mathcal{Y}$, the matrix game can even be solved in sublinear time in the size $n d$ of $A$. When $\mathcal{X}$ and $\mathcal{Y}$ are both $\ell_{1}$-norm unit balls, Grigoriadis and Khachiyan (1995) can solve the matrix game in time $O\left((n+d) \log (n+d) / \epsilon^{2}\right)$. When $\mathcal{X}$ is the $\ell_{2}$-norm unit ball in $\mathbb{R}^{d}$ and $\mathcal{Y}$ is the $\ell_{1}$-norm unit ball in $\mathbb{R}^{n}$, Clarkson, Hazan, and Woodruff (2012) can solve the matrix game in time $O\left((n+d) \log n / \epsilon^{2}\right)$.

As far as we know, the $\ell_{1}-\ell_{1}$ and $\ell_{2}-\ell_{1}$ matrix games are the only two cases where sublinear algorithms are known. However, there is general interest of solving matrix games with general norms. For instance, matrix games are closely related to the Carathéodory problem for finding a sparse linear combination in the convex hull of given data points, where all the $\ell_{p}$-metrics with $p \geq 2$ have been well-studied (Barman 2015; Mirrokni et al. 2017; Combettes and Pokutta 2019). In addition, matrix games are common in machine learning especially support vector machines (SVMs), and general $\ell_{p}$-margin SVMs have also been considered by previous literature, see e.g. the book by Deng, Tian, and Zhang (2012). In all, it is a natural question to investigate sublinear algorithms for general matrix games. In addition, quantum computing has been rapidly advancing and current technology has reached "quantum supremacy" for some specific tasks (Arute et al. 2019); since previous works have given sublinear quantum algorithms for $\ell_{1}-\ell_{1}$ matrix games $(\mathrm{Li}$, Chakrabarti, and Wu 2019; Apeldoorn and Gilyén 2019) and $\ell_{2}-\ell_{1}$ matrix games (Li, Chakrabarti, and Wu 2019 ) with running time $(\sqrt{n}+\sqrt{d})$ poly $(1 / \epsilon)$, it is also natural to explore sublinear quantum algorithms for general matrix games.

Contributions. We conduct a systematic study of $\ell_{q}-\ell_{1}$ matrix games for any $q \in(1,2]$ which corresponds to $\ell_{q^{-}}$ margin SVMs and the $\ell_{p}$-Carathéodory problem for any $p \geq$ 2. We use the following entry-wise input model, the standard assumption in the sublinear algorithms in Grigoriadis and Khachiyan (1995); Clarkson, Hazan, and Woodruff (2012):

Input model: Given any $i \in[n]$ and $j \in[d]$, the $j^{\text {th }}$ entry of $A_{i}$ can be recovered in $O(1)$ time.

Quantumly, we consider an almost same oracle:
Quantum input model: Given any $i \in[n]$ and $j \in[d]$, the $j^{\text {th }}$ entry of $A_{i}$ can be recovered in $O(1)$ time coherently.

The only difference is to allow coherent queries, which
give quantum algorithms the ability to query different locations in superposition, and have been the standard quantization of the classical inputs and commonly adopted in previous works (Li, Chakrabarti, and Wu 2019; Apeldoorn and Gilyén 2019).

Theorem 1 (Main Theorem). Given $q \in(1,2]$. Define $p \geq 2$ such that $\frac{1}{p}+\frac{1}{q}=1$. Consider the $\ell_{q}-\ell_{1}$ matrix game ${ }^{1}$ :

$$
\begin{equation*}
\sigma:=\max _{x \in \mathbb{B}_{q}^{d}} \min _{\mathbf{p} \in \Delta_{n}} \mathbf{p}^{\top} A x \tag{2}
\end{equation*}
$$

where $\mathbb{B}_{q}^{d}$ is the $\ell_{q}$-unit ball in $\mathbb{R}^{d}$ and $\Delta_{n}$ is the $\ell_{1}$-simplex in $\mathbb{R}^{n}$. Then we can find an $\bar{x} \in \mathbb{B}_{q}^{d}$ s.t. ${ }^{2}$

$$
\begin{equation*}
\min _{i \in[n]} A_{i} \bar{x} \geq \sigma-\epsilon \tag{3}
\end{equation*}
$$

with success probability at least $2 / 3$, using

- $O\left(\frac{(n+d)(p+\log n)}{\epsilon^{2}}\right)$ classical queries (Theorem 2 ); or
- $\tilde{O}\left(\frac{p^{2} \sqrt{n}}{\epsilon^{4}}+\frac{p^{3.5} \sqrt{d}}{\epsilon^{7}}\right)$ quantum queries ${ }^{3}$ (Theorem 3).

When $p=\Omega(\log d / \epsilon)$, the above bounds can be improved (by Lemma 1) to respectively

- $O\left(\frac{(n+d)\left(\frac{\log d}{\epsilon}+\log n\right)}{\epsilon^{2}}\right)$ queries to the classical input model; - $\tilde{O}\left(\frac{\sqrt{n}}{\epsilon^{6}}+\frac{\sqrt{d}}{\epsilon^{10.5}}\right)$ queries to the quantum input model.

Both results are optimal in $n$ and $d$ up to poly-log factors as we show $\Omega(n+d)$ and $\Omega(\sqrt{n}+\sqrt{d})$ classical and quantum lower bounds respectively when $\epsilon=\Theta(1)$ (Theorem 4).

Conceptually, our classical and quantum algorithms for general matrix games enjoy quite a few nice properties. On the one hand, they can be directly applied to

- Convex geometry: We give the first sublinear classical and quantum algorithms for the approximate Carathéodory problem (Corollary 1), improving the previous lineartime algorithms of Mirrokni et al. (2017); Combettes and Pokutta (2019);
- Supervised learning: We provide the first sublinear algorithms for general $\ell_{q}$-margin support vector machines (SVMs) (Corollary 2).
On the other hand, our quantum algorithm is friendly for near-term applications. It uses the standard quantum input model and needs not to use any sophisticated quantum data structures. It is classical-quantum hybrid where the quantum part is isolated by pieces of state preparations connected by classical processing. Its output is completely classical.

Technique-wise, we are deeply inspired by Clarkson, Hazan, and Woodruff (2012), which serves as the starting point of our algorithm design. At a high level, Clarkson et al.'s algorithm follows a primal-dual framework where the primal part applies ( $\ell_{2}$-norm) online gradient descent (OGD)

[^1]by Zinkevich (2003), and the dual part applies multiplicative weight updates (MWU) by $\ell_{2}$-sampling. The choice of the $\ell_{2}$-norm metric greatly facilitates the design and analysis of the algorithms for both parts. However, it is conceivable that more sophisticated design and analysis will be required to handle general $\ell_{q}-\ell_{1}$ matrix games.

Classically, our main technical contribution is to expand the primal-dual approach of Clarkson, Hazan, and Woodruff (2012) to work for more general metrics for the $\ell_{q}-\ell_{1}$ matrix game. Specifically, in the primal we replace OGD by a generalized $p$-norm OGD due to Shalev-Shwartz (2012), and in the dual we replace the $\ell_{2}$-sampling by $\ell_{q}$-sampling. We conduct a careful algorithm design and analysis to ensure that this strategy only incurs an $O\left(p / \epsilon^{2}\right)$ overhead in the number of iterations, and the error of the $\ell_{q}-\ell_{1}$ matrix game is still bounded by $\epsilon$ as in (3). In a nutshell, our algorithm can be viewed as an interpolation between the $\ell_{2}-\ell_{1}$ matrix game (Clarkson, Hazan, and Woodruff 2012) and the $\ell_{1}-\ell_{1}$ matrix game (Grigoriadis and Khachiyan 1995): when $q$ is close to 2 the algorithm is more similar to Clarkson, Hazan, and Woodruff (2012), whereas when $q$ is close to $1, p$ is large and the $p$-norm GD becomes closer to the normalized exponentiated gradient (Shalev-Shwartz 2012), which is exactly the update rule in Grigoriadis and Khachiyan (1995).

Quantumly, our main contribution is the systematic improvement of the previous quantum algorithm for $\ell_{2}-\ell_{1}$ matrix games by Li , Chakrabarti, and Wu (2019). They achieved a quantum speedup of $\tilde{O}(\sqrt{n}+\sqrt{d})$ for solving $\ell_{2}-\ell_{1}$ matrix games by levering quantum amplitude amplification and observing that $\ell_{2}$-sampling can be readily accomplished by quantum state preparation as quantum states refer to $\ell_{2}$ unit vectors. For general $\ell_{q}-\ell_{1}$ matrix game ( $q \in(1,2]$ ), we likewise upgrade both primal and dual parts as in our classical algorithm: specifically, in the primal, we apply the $p$-norm OGD in $\tilde{O}(\sqrt{d})$ time, whereas in the dual, we apply the multiplicative weight update via an $\ell_{q}$-sampling in $\tilde{O}(\sqrt{n})$ time. To that end, we contribute to the following technical improvements, which may be of independent interest:

- In our algorithm, we cannot directly leverage quantum state preparation in the $\ell_{q}$ metric because it corresponds to $\ell_{2}$-normalized vectors. Instead, we propose Algorithm 2 for quantum $\ell_{q}$-sampling with $O(\sqrt{n})$ oracle calls which works with states whose amplitudes follow $\ell_{q}$-norm proportion. Measuring such states is equivalent to performing $\ell_{q}$-sampling.
- When $p=q=2$, we improved the $\epsilon$-dependence from the $1 / \epsilon^{8}$ in the prior art by Li, Chakrabarti, and Wu (2019) to $1 / \epsilon^{7}$. This is achieved by deriving a better upper bound on the entries of the vectors in the $p$-norm OGD (i.e., $y_{t, j}$ as in Eq. (38)); see the supplementary material (Eqs. (38)(40)) for details.
- In our lower bounds, although the hard cases are motivated by Li, Chakrabarti, and Wu (2019), the matrix game values are much more complicated in the $\ell_{q}-\ell_{1}$ case. In the supplementary material, we figure out two functions $f_{1}$ and $f_{2}$ that not only separate the game values of two specifically-constructed $\ell_{q}-\ell_{1}$ matrix games but also have
monotone and nonnegative properties, which are crucial factors in our proof.
These improvements together result in Theorem 1.

Related work. Matrix games were probably first studied as zero-sum games by Neumann (1928). The seminal work (Nemirovski and Yudin 1983) proposed the mirror descent method and gave an algorithm for solving matrix games in time $\tilde{O}\left(n d / \epsilon^{2}\right)$. This was later improved to $\tilde{O}(n d / \epsilon)$ by the prox-method due to Nemirovski (2004) and the dual extrapolation method due to Nesterov (2007). To further improve the cost, there have been two main focuses:

- Sampling-based methods: They focus on achieving sublinear cost in $n d$, the size of the matrix $A$. Grigoriadis and Khachiyan (1995); Clarkson, Hazan, and Woodruff (2012) mentioned above are seminal examples; these sublinear algorithms can also be used to solve semidefinite programs (Garber and Hazan 2011), SVMs (Hazan, Koren, and Srebro 2011), etc.
- Variance-reduced methods: They focus on the cost in $1 / \epsilon$, in particular its decoupling with nd. Palaniappan and Bach (2016) showed how to apply the standard SVRG (Johnson and Zhang 2013) technique for solving $\ell_{2}-\ell_{2}$ matrix games; this idea can also be extended to smooth functions using general Bregman divergences (Shi, Zhang, and Yu 2017). Variance-reduced methods for solving matrix games culminate in Carmon et al. (2019), where they show how to solve $\ell_{1}-\ell_{1}$ and $\ell_{2}-\ell_{1}$ matrix games in time $\tilde{O}(\mathrm{nnz}(A)+\sqrt{\mathrm{nnz}(A) \cdot(n+d)} / \epsilon)$, where $\mathrm{nnz}(A)$ is the number of nonzero elements in $A$.

There have been relatively few quantum results for solving matrix games. Kapoor, Wiebe, and Svore (2016) solved the $\ell_{2}-\ell_{1}$ matrix game with cost $\tilde{O}\left(\sqrt{n} d / \epsilon^{2}\right)$ using an unusual input model where the representation of a data point in $\mathbb{R}^{d}$ is the concatenation of $d$ floating point numbers. More recently, Apeldoorn and Gilyén (2019) was able to solve the $\ell_{1}$ $\ell_{1}$ matrix game with cost $\tilde{O}\left(\sqrt{n} / \epsilon^{3}+\sqrt{d} / \epsilon^{3}\right)$ using the standard input model above, and Li , Chakrabarti, and Wu (2019) solved the $\ell_{2}-\ell_{1}$ matrix game with cost $\tilde{O}\left(\sqrt{n} / \epsilon^{4}+\sqrt{d} / \epsilon^{8}\right)$ also using the standard input model.

## Preliminaries and Notations

To facilitate the reading of this paper, we introduce necessary definitions and notations here.

Preliminaries for quantum computing. Quantum mechanics can be formulated in terms of linear algebra. For the space $\mathbb{C}^{d}$, we denote $\left\{\vec{e}_{0}, \ldots, \vec{e}_{d=1}\right\}$ as its computational basis, where $\vec{e}_{i}=(0, \ldots, 1, \ldots, 0)^{\top}$ where 1 only appears in the $(i+1)^{\mathrm{th}}$ coordinate. These basic vectors can be written by the Dirac notation: $\vec{e}_{i}:=|i\rangle$ (called a "ket"), and $\vec{e}_{i}^{\top}:=\langle i|$ (called a "bra"). A $d$-dimensional quantum state is a unit vector in $\mathbb{C}^{d}$ : i.e., $|v\rangle=\left(v_{0}, \ldots, v_{d-1}\right)^{\top}$ such that $\sum_{i=0}^{d-1}\left|v_{i}\right|^{2}=1$.

Tensor product of quantum states is their Kronecker product: if $|u\rangle \in \mathbb{C}^{d_{1}}$ and $|v\rangle \in \mathbb{C}^{d_{2}}$, then

$$
\begin{equation*}
|u\rangle \otimes|v\rangle:=\left(u_{0} v_{0}, u_{0} v_{1}, \ldots, u_{d_{1}-1} v_{d_{2}-1}\right)^{\top} \tag{4}
\end{equation*}
$$

which is a vector in $\mathbb{C}^{d_{1}} \otimes \mathbb{C}^{d_{2}}$.
Quantum access to an input matrix, also known as a quantum oracle, is reversible and allows access to coordinates of the matrix in superposition, this is the essence of quantum speedups. In particular, to access entries of a matrix $A \in \mathbb{R}^{n \times d}$, we exploit a quantum oracle $O_{A}$, which is a unitary transformation on $\mathbb{C}^{n} \otimes \mathbb{C}^{d} \otimes \mathbb{C}^{d_{\text {acc }}}$ ( $d_{\text {acc }}$ being the dimension of a floating-point register) such that

$$
\begin{equation*}
O_{A}(|i\rangle \otimes|j\rangle \otimes|z\rangle)=|i\rangle \otimes|j\rangle \otimes\left|z \oplus A_{i j}\right\rangle \tag{5}
\end{equation*}
$$

for any $i \in[n], j \in[d]$, and $z \in \mathbb{C}^{d_{\text {acc }}}$. Intuitively, $O_{A}$ reads the entry $A_{i j}$ and stores it in the third register as a floating-point number. However, to promise that $O_{A}$ a unitary transformation, $O_{A}$ applies the XOR operation $(\oplus)$ on the third register. This is a natural generalization of classical reversible computation, when each entry of $A$ can be recovered in $O(1)$ time. Subsequently, a common assumption is that a single query to $O_{A}$ takes $O(1)$ cost.

Interpolation for large $p$. If $p$ is large, we prove the following lemma showing that we can restrict without loss of generality to cases where $p$ such that $\frac{1}{p}+\frac{1}{q}=1$ is $O(\log d / \epsilon)$, since in this case the $\ell_{q}-\ell_{1}$ matrix game is $\epsilon$-close to the $\ell_{1}-\ell_{1}$ matrix game in the following sense:
Lemma 1. An $\ell_{q}-\ell_{1}$ matrix game where $p$ such that $\frac{1}{p}+\frac{1}{q}=1$ is greater than $\log d / \epsilon$ can be solved using an algorithm for solving $\ell_{1}-\ell_{1}$ games. This introduces an error $O(\epsilon)$ in the objective value.

Proof. Assume without loss of generality that $\epsilon \leq 1 / 2$. Let $p \geq \log d / \epsilon \geq \log d /(-\log (1-\epsilon))$. It can be easily verified that $\mathbb{B}_{1}^{d} \subset \mathbb{B}_{q}^{d} \subset \mathbb{B}_{1}^{d}+\left(1-d^{-1 / p}\right) \mathbb{B}_{q}^{d}$. Thus $\mathbb{B}_{q}^{d} \subset \mathbb{B}_{1}^{d}+\epsilon \mathbb{B}_{q}^{d}$.

Consider applying an algorithm to solve an $\ell_{1}-\ell_{1}$ matrix game instead of the $\ell_{q}-\ell_{1}$ matrix game as required in (2). Let the optimal solution to (2) be $x^{*} \in \mathbb{B}_{q}^{d}, p^{*} \in \Delta_{n}$. By the previous analysis, there is a point $x \in \mathbb{B}_{1}^{d}$ such that $\left\|x-x^{*}\right\|_{q} \leq \epsilon$. Thus the solution $x, p^{*}$ has an error at most $O(\epsilon)$ from the true objective, and the algorithm for solving $\ell_{1}-\ell_{1}$ games finds a solution at least as good as this.

Notations. Throughout the paper, we denote $p, q>1$ to be two real numbers such that $\frac{1}{p}+\frac{1}{q}=1 ; p \in[2,+\infty)$ and $q \in$ $(1,2]$. For any $s>1$, we use $\mathbb{B}_{s}^{d}$ to denote the $d$-dimensional unit ball in $\ell_{s}$-norm, i.e., $\mathbb{B}_{s}^{d}:=\left\{x: \sum_{i \in[d]}\left|x_{i}\right|^{s} \leq 1\right\}$; we use $\Delta_{n}$ to denote the $n$-dimensional unit simplex $\left\{p \in \mathbb{R}^{n}\right.$ : $\left.p_{i} \geq 0, \sum_{i} p_{i}=1\right\}$, and use $\mathbf{1}_{n}$ to denote the $n$-dimensional all-one vector. We denote $A \in \mathbb{R}^{n \times d}$ to be the matrix whose $i^{\text {th }}$ row is $A_{i}^{\top}$ for all $i \in[n]$. We define sgn: $\mathbb{R} \rightarrow\{-1,0,1\}$ such that $\operatorname{sgn}(x)=-1$ if $x<0, \operatorname{sgn}(x)=1$ if $x>0$, and $\operatorname{sgn}(0)=0$.

## A Sublinear Classical Algorithm for General Matrix Games

For any $q \in(1,2]$, we consider the $\ell_{q}-\ell_{1}$ matrix game:

$$
\begin{equation*}
\sigma:=\max _{x \in \mathbb{B}_{q}^{d}} \min _{\mathbf{p} \in \Delta_{n}} \mathbf{p}^{\top} A x \tag{6}
\end{equation*}
$$

The goal is to find a $\bar{x}$ that approximates the equilibrium of the matrix game within additive error $\epsilon$ :

$$
\begin{equation*}
\min _{i \in[n]} A_{i} \bar{x} \geq \sigma-\epsilon \tag{7}
\end{equation*}
$$

Throughout the paper, we assume $A_{1}, \ldots, A_{n} \in \mathbb{B}_{p}^{d}$, i.e., all the $n$ data points are normalized to have $\ell_{p}$-norm at most 1 .

```
Algorithm 1: A sublinear algorithm for \(\ell_{q}-\ell_{1}\) games.
    Input: \(\epsilon>0 ; p \in[2,+\infty), q \in(1,2]\) such that
            \(\frac{1}{p}+\frac{1}{q}=1 ; A \in \mathbb{R}^{n \times d}\) with \(A_{i} \in \mathbb{B}_{p}^{d} \forall i \in[n]\).
    Output: \(\bar{x}\) that satisfies (7).
    Let \(T=\left\lceil\frac{895 \log n+4 p}{\epsilon^{2}}\right\rceil, y_{1}=\mathbf{0}_{d}, \eta=\sqrt{\frac{11 \log n}{12 T}}\),
    \(w_{1}=\mathbf{1}_{n} ;\)
    for \(t=1\) to \(T\) do
        \(p_{t} \leftarrow \frac{w_{t}}{\left\|w_{t}\right\|_{1}}, x_{t} \leftarrow \frac{y_{t}}{\max \left\{1,\left\|y_{t}\right\|_{q}\right\}} ;\)
        Choose \(i_{t} \in[n]\) by \(i_{t} \leftarrow i\) with probability \(p_{t}(i)\);
        Define \(y_{t+1}\) where for any \(j \in[d]\),
            \(y_{t+1, j} \leftarrow y_{t}+\sqrt{\frac{q-1}{2 T}} \frac{\operatorname{sgn}\left(A_{i_{t}, j}\right)\left|A_{i_{t}, j}\right|^{p-1}}{\left\|A_{i_{t}}\right\|_{p}^{p-2}} ;\)
        Choose \(j_{t} \in[d]\) by \(j_{t} \leftarrow j\) with probability \(\frac{x_{t}(j)^{q}}{\left\|x_{t}\right\|_{q}^{q}}\);
        for \(i=1\) to \(n\) do
            \(\tilde{v}_{t}(i) \leftarrow A_{i}\left(j_{t}\right)\left\|x_{t}\right\|_{q}^{q} / x_{t}\left(j_{t}\right)^{q-1} ;\)
            \(v_{t}(i) \leftarrow \operatorname{clip}\left(\tilde{v}_{t}(i), \frac{1}{\eta}\right)\) where
            \(\operatorname{clip}(v, M):=\min \{M, \max \{-M, v\}\}\)
            \(\forall v, M \in \mathbb{R}\);
            \(w_{t+1}(i) \leftarrow w_{t}(i)\left(1-\eta v_{t}(i)+\eta^{2} v_{t}(i)^{2}\right) ;\)
11 Return \(\bar{x}=\frac{1}{T} \sum_{t=1}^{T} x_{t}\).
```

Theorem 2. The output of Algorithm 1 satisfies (7) with probability at least $2 / 3$, and its total running time is $O\left(\frac{(n+d)(p+\log n)}{\epsilon^{2}}\right)$ where $p \geq 2$ such that $\frac{1}{p}+\frac{1}{q}=1$.

Our sublinear algorithm follows the primal-dual approach of Algorithm 1 of Clarkson, Hazan, and Woodruff (2012), which solves $\ell_{1}-\ell_{2}$ matrix games. Here for $\ell_{q}-\ell_{1}$ matrix games, the solution vector $x$ now lies in $\mathbb{B}_{q}^{d}$. Hence, the most natural adaptations are to use $\ell_{q}$-sampling instead of $\ell_{2}$-sampling in the primal updates, and to use a $p$-norm OGD by Shalev-Shwartz (2012) which generalizes the online gradient descent by Zinkevich (2003) in $\ell_{2}$-norm. In the following, we use various technical tools to show these natural adaptations actually work.
Proposition 1 (Shalev-Shwartz 2012, Corollary 2.18). Consider a set of vectors $u_{1}, \ldots, u_{T} \in \mathbb{R}^{d}$ such that $\left\|u_{i}\right\|_{p} \leq$

1. Set $\iota=\sqrt{\frac{q-1}{2 T}}$. Let $x_{0} \leftarrow \mathbf{0}_{d}, \tilde{x}_{t+1, i} \leftarrow x_{t, i}+$
$\iota \frac{\operatorname{sgn}\left(u_{t, i}\right)\left|u_{t, i}\right|^{p-1}}{\left\|u_{t}\right\|_{p}^{p-2}}$ for all $i \in[d]$, and $x_{t+1} \leftarrow \frac{\tilde{x}_{t+1}}{\max \left\{1,\left\|\tilde{x}_{t+1}\right\|_{q}\right\}}$. Then

$$
\begin{equation*}
\max _{x \in \mathbb{B}_{q}^{d}} \sum_{t=1}^{T} u_{t}^{\top} x-\sum_{t=1}^{T} u_{t}^{\top} x_{t} \leq \sqrt{\frac{2 T}{q-1}} . \tag{8}
\end{equation*}
$$

The analysis of Algorithm 1 uses the following lemma, adapted from the variance multiplicative weight lemma and martingale tail bounds in Clarkson, Hazan, and Woodruff (2012) ${ }^{4}$ :

Lemma 2 (Section 2 of Clarkson, Hazan, and Woodruff 2012). In Algorithm 1, the parameters $p_{t}$ in Line 3 and $v_{t}$ in Line 9 satisfy

$$
\begin{equation*}
\sum_{t \in[T]} p_{t}^{\top} v_{t} \leq \min _{i \in[n]} \sum_{t \in[T]} v_{t}(i)+\eta \sum_{t \in[T]} p_{t}^{\top} v_{t}^{2}+\frac{\log n}{\eta} \tag{9}
\end{equation*}
$$

where $v_{t}^{2}$ is defined as $\left(v_{t}^{2}\right)_{i}:=\left(v_{t}\right)_{i}^{2}$ for all $i \in[n]$, as long as the update rule of $w_{t}$ is as in Line 10 and $\operatorname{Var}\left[v_{t}(i)^{2}\right] \leq 1$ for all $t \in[T]$ and $i \in[n]$. Furthermore, with probability at least $1-O(1 / n)$,

$$
\begin{gather*}
\max _{i \in[n]} \sum_{t \in[T]}\left[v_{t}(i)-A_{i} x_{t}\right] \leq 4 \eta T  \tag{10}\\
\left|\sum_{t \in[T]} A_{i_{t}} x_{t}-\sum_{t \in[T]} p_{t}^{\top} v_{t}\right| \leq 10 \eta T \tag{11}
\end{gather*}
$$

with probability at least $5 / 7, \sum_{t \in[T]} p_{t}^{\top} v_{t}^{2} \leq 7 T$.
We also need to prove the following inequality on different moments of random variables.
Lemma 3. Suppose that $X$ is a random variable on $\mathbb{R}$, and $p \geq 2$. If $\mathbb{E}\left[|X|^{p}\right] \leq 1$, then $\mathbb{E}\left[X^{2}\right] \leq 1$.
Proof. Denote the probability density of $X$ as $\mu$. Then $\int_{-\infty}^{+\infty}|x|^{p} \mathrm{~d} \mu_{x}=\mathbb{E}\left[|X|^{p}\right] \leq 1$. By Hölder's inequality, we have

$$
\begin{align*}
1 & \geq\left(\int_{-\infty}^{+\infty}|x|^{p} \mathrm{~d} \mu_{x}\right)^{2 / p}\left(\int_{-\infty}^{+\infty} 1 \mathrm{~d} \mu_{x}\right)^{1-2 / p} \\
& \geq \int_{-\infty}^{+\infty}|x|^{2} \cdot 1^{1-2 / p} \mathrm{~d} \mu_{x}=\int_{-\infty}^{+\infty} x^{2} \mathrm{~d} \mu_{x} \tag{12}
\end{align*}
$$

hence the result follows.
Now we are ready to prove our main theorem.
Proof of Theorem 2. First, $\tilde{v}_{t}(i)$ is an unbiased estimator of $A_{i} x_{t}$ as

$$
\begin{equation*}
\mathbb{E}\left[\tilde{v}_{t}(i)\right]=\sum_{j_{t}=1}^{d} \frac{x_{t}\left(j_{t}\right)^{q}}{\left\|x_{t}\right\|_{q}^{q}} \cdot \frac{A_{i}\left(j_{t}\right)\left\|x_{t}\right\|_{q}^{q}}{x_{t}\left(j_{t}\right)^{q-1}}=A_{i} x_{t} \tag{13}
\end{equation*}
$$

[^2]Furthermore,

$$
\begin{align*}
\mathbb{E}\left[\left|\tilde{v}_{t}(i)\right|^{p}\right] & =\sum_{j_{t}=1}^{d} \frac{x_{t}\left(j_{t}\right)^{q}}{\left\|x_{t}\right\|_{q}^{q}} \cdot \frac{\left|A_{i}\left(j_{t}\right)\right|^{p}\left\|x_{t}\right\|_{q}^{p q}}{x_{t}\left(j_{t}\right)^{p(q-1)}} \\
& =\left\|A_{i}\right\|_{p}^{p}\left\|x_{t}\right\|_{q}^{p} \leq 1, \tag{14}
\end{align*}
$$

where the second equality follows from the identities $q=$ $p(q-1)$ and $p=q(p-1)$, and the last inequality follows from the assumption that $A_{i} \in \mathbb{B}_{p}^{d} \forall i \in[n]$. By Lemma 3, $\mathbb{E}\left[\tilde{v}_{t}(i)^{2}\right] \leq 1$. Because the clip function in Line 9 only makes variance smaller, this means that the conditions of Lemma 2 are satisfied and we hence have (9), rewritten below:

$$
\begin{equation*}
\sum_{t \in[T]} p_{t}^{\top} v_{t} \leq \min _{i \in[n]} \sum_{t \in[T]} v_{t}(i)+\eta \sum_{t \in[T]} p_{t}^{\top} v_{t}^{2}+\frac{\log n}{\eta} \tag{15}
\end{equation*}
$$

Furthermore, Lemma 2 implies that with probability 5/7$O(1 / n)$ we have

$$
\begin{equation*}
\sum_{t \in[T]} A_{i_{t}} x_{t} \leq \min _{i \in[n]} \sum_{t \in[T]} v_{t}(i)+17 \eta T+\frac{\log n}{\eta} \tag{16}
\end{equation*}
$$

Moreover, (10) gives $\sum_{t \in[T]}\left[v_{t}(i)-A_{i} x_{t}\right] \leq 4 \eta T$, and hence $\min _{i \in[n]} \sum_{t \in[T]} v_{t}(i) \leq 4 \eta T+\min _{i \in[n]} \sum_{t \in[T]} A_{i} x_{t}$. Plugging this into (16), we have

$$
\begin{align*}
\sum_{t \in[T]} A_{i_{t}} x_{t} & \leq \sum_{t \in[T]} p_{t}^{\top} v_{t}+10 \eta T \\
& \leq \min _{i \in[n]} \sum_{t \in[T]} A_{i} x_{t}+21 \eta T+\frac{\log n}{\eta} \tag{17}
\end{align*}
$$

with probability $(5 / 7-O(1 / n)) \cdot(1-O(1 / n)) \geq 2 / 3$.
On the other hand, by taking $u_{t}=A_{i_{t}}$ in Proposition 1,

$$
\begin{equation*}
T \sigma \leq \max _{x \in \mathbb{B}_{q}^{d}} \sum_{t=1}^{T} A_{i_{t}} x \leq \sum_{t=1}^{T} A_{i_{t}} x_{t}+\sqrt{2 T p} \tag{18}
\end{equation*}
$$

since $\frac{1}{q-1}=\frac{p}{q} \leq p$. Combining (17) and (18), we have

$$
\begin{equation*}
\min _{i \in[n]} \sum_{t \in[T]} A_{i} x_{t} \geq T \sigma-\sqrt{2 T p}-21 \eta T-\frac{\log n}{\eta} \tag{19}
\end{equation*}
$$

Consequently, the return $\bar{x}=\frac{1}{T} \sum_{t=1}^{T} x_{t}$ of Algorithm 1 in Line 11 satisfies

$$
\begin{equation*}
\min _{i \in[n]} A_{i} \bar{x} \geq \sigma-\sqrt{\frac{2 p}{T}}-21 \eta-\frac{\log n}{\eta T} \tag{20}
\end{equation*}
$$

To prove (7), it remains to show that $\sqrt{\frac{2 p}{T}}+21 \eta+\frac{\log n}{\eta T} \leq \epsilon$, which is equivalent to $\sqrt{2 p}+21 \sqrt{\frac{11 \log n}{12}}+\sqrt{\frac{12 \log n}{11}} \leq$ $\sqrt{T} \epsilon$ by the definition of $\eta$. This is true because the AMGM inequality implies that that LHS is at most $2(\sqrt{2 p})^{2}+$ $2\left(21 \sqrt{\frac{11 \log n}{12}}+\sqrt{\frac{12 \log n}{11}}\right)^{2} \leq 4 p+895 \log n \leq T \epsilon^{2}$.

Lemma 1 combined with Theorem 2 yields the classical result in Theorem 1.

## A Sublinear Quantum Algorithm for General Matrix Games

In this section, we give a quantum algorithm for solving the general $\ell_{q}-\ell_{1}$ matrix games. It closely follows our classical algorithm because they both use a primal-dual approach, where the primal part is composed of $p$-norm online gradient descent and the dual part is composed of multiplicative weight updates. However, we adopt quantum techniques to achieve speedup on both.

The intuition behind the quantum algorithm and the quantum speedup is that we measure quantum states to obtain random samples. These quantum states can be efficiently prepared (with cost $\tilde{O}(\sqrt{n})$ and $\tilde{O}(\sqrt{d})$ ). Mathematically, A quantum state can be represented by an $\ell_{2}$-normalized complex vector $\psi$ in the sense that measuring this quantum states yields outcome $i$ with probability $\left|\psi_{i}\right|^{2}$ (thus for every probability distribution there is a quantum state corresponding to it). Let us denote the quantum state for sampling from $w$ by $|w\rangle$ and the quantum state for sampling from $x$ by $|x\rangle$ (different from the notation in Algorithm 3). If we can maintain $|w\rangle$ and $|x\rangle$ in each iteration, then there is no need for classical updates, and preparing $|w\rangle$ and $|x\rangle$ becomes the bottleneck of the quantum algorithm.

The source of our quantum speedup comes from an important subroutine, Algorithm 2, which is designed to prepare states for $\ell_{q}$-sampling. It uses standard Grover-based techniques to prepare states but we carefully keep track of the normalizing factor to facilitate $\ell_{q}$-sampling. We showed (in Proposition 2 in the supplementary material) that preparing $|w\rangle$ costs $\tilde{O}(\sqrt{n})$ and preparing $|x\rangle \operatorname{costs} \tilde{O}(\sqrt{d})$. In the following, we give the high-level ideas of Algorithm 2.

1. We first create a quantum state corresponding to the uniform distribution, which is easy using Hadamard gates.
2. For each entry, we create a state with the desired amplitude associated with 0 , and an undesired amplitude associated to 1 (the unitarity of quantum operations necessitates the existence of this undesired term).
3. Finally we use a technique called amplitude amplification to amplify the portion of the state corresponding to 0 for each entry, to get a state with only the desired amplitudes.
The details of our quantum algorithm for solving the general $\ell_{q}-\ell_{1}$ matrix games, Algorithm 3, is rather technical. To simplify the presentation, we postpone its pseudocode (Algorithm 3) to the supplementary material and highlight how it is different from Algorithm 1 in the following.

- For the primal part, we prepare a quantum state $\left|y_{t}\right\rangle$ for the $q$-norm OGD and measure it (in Line 7) to obtain a sample $j_{t} \in[d]$. The subtlety here is that we need to perform the $\ell_{q}$-sampling to the vector $y_{t}$; this is different from the $\ell_{2}$-sampling in Li , Chakrabarti, and Wu (2019) which uses the fact that pure quantum states are $\ell_{2}$-normalized. To this end, we design Algorithm 2 for $\ell_{q}$-quantum state sampling, which may be of independent interest; this algorithm is built upon a clever use of quantum amplitude amplification, the technique behind the Grover search (Grover 1996). Note that sampling according to $y_{t}$ is equivalent to sampling according to $x_{t}$

```
Algorithm 2: Prepare an \(\ell_{q}\)-pure state given an oracle
to its coefficients.
1 Find \(a_{\|q\|}:=\max _{i \in[n]}\left|a_{i}\right|^{q / 2}\) in \(O(\sqrt{n})\) time by the minimum-finding algorithm (Dürr and Høyer 1996);
2 Prepare the uniform superposition \(\frac{1}{\sqrt{n}} \sum_{r \in[n]}|i\rangle\);
3 Perform the following unitary transformations:
\[
\begin{aligned}
& \frac{1}{\sqrt{n}} \sum_{i \in[n]}|i\rangle \stackrel{O_{a}}{\longmapsto} \frac{1}{\sqrt{n}} \sum_{i \in[n]}|i\rangle\left|a_{i}\right\rangle \\
& \mapsto \frac{1}{\sqrt{n}} \sum_{i \in[n]}|i\rangle\left|a_{i}\right\rangle\left(\frac{a_{i}^{q / 2}}{a_{\|q\|}}|0\rangle+\sqrt{1-\frac{\left|a_{i}\right|^{q}}{a_{\|q\|}^{2}}}|1\rangle\right) \\
& \stackrel{O_{a}^{-1}}{\longmapsto} \frac{1}{\sqrt{n}} \sum_{i \in[n]}|i\rangle|0\rangle\left(\frac{a_{i}^{q / 2}}{a_{\|q\|}}|0\rangle+\sqrt{1-\frac{\left|a_{i}\right|^{q}}{a_{\|q\|}^{2}}}|1\rangle\right)
\end{aligned}
\]
```

5 Discard the second register above; rewrite the state as

$$
\begin{equation*}
\frac{\|a\|_{q}^{q / 2}}{\sqrt{n} a_{\|q\|}}\left(\frac{1}{\|a\|_{q}^{q / 2}} \sum_{i \in[n]} a_{i}^{q / 2}|i\rangle\right)|0\rangle+\left|a^{\perp}\right\rangle|1\rangle, \tag{21}
\end{equation*}
$$

where $\left|a^{\perp}\right\rangle:=\frac{1}{\sqrt{n}} \sum_{i \in[n]} \sqrt{1-\frac{\left|a_{i}\right|^{q}}{a_{\|q\|}^{2}}}|i\rangle ;$
6 Apply amplitude amplification (Brassard et al. 2002) for the state in (21) conditioned on the second register being 0 . Return the output.
in Algorithm 1, because $x_{t}(j)^{q} /\left\|x_{t}\right\|_{q}^{q}=y_{t}(j)^{q} /\left\|y_{t}\right\|_{q}^{q}$. Moreover, it suffices to replace $\left\|x_{t}\right\|_{q}^{q} / x_{t}\left(j_{t}\right)^{q-1}$ with $\left\|y_{t}\right\|_{q}^{q} /\left(y_{t}\left(j_{t}\right)^{q-1} \max \left\{1,\left\|y_{t}\right\|_{q}\right\}\right)$ in Line 8 of Algorithm 3. Similar to preparing $\left|p_{t}\right\rangle$, we use $\tilde{O}(\sqrt{d})$ queries to $O_{A}$ to prepare $y_{t}$, while classically we need to compute all the entries of $y_{t}$, which takes $O(d)$ queries.

- For the dual part, we prepare the multiplicative weight vector as a quantum state $\left|p_{t}\right\rangle$ and measure it (in Line 3 ) to obtain a sample $i_{t} \in[n]$. This adaption enables us to achieve the $\tilde{O}(\sqrt{n})$ dependence by using quantum amplitude amplification in the quantum state preparation: in Line 8, we implement the oracle $O_{t}$ and in Line 9 we use $\tilde{O}(\sqrt{n})$ queries to $O_{t}$ to prepare the state $\left|p_{t+1}\right\rangle$ for the next iteration. In contrast, classically we need to compute all the entries of $w_{t+1}$ to obtain the probability distribution $p_{t+1}$ for the next iteration, which takes $O(n)$ queries.

In general, Algorithm 3 can be viewed as a template for achieving quantum speedups for online mirror descent methods: In this work, we focus on the general matrix games where the primal and dual are in the special relationship of $\ell_{p}$ and $\ell_{q}$ norms, but in principle it may be applicable to study other dualities in online learning.

We summarize the main quantum result as the following theorem, which states the correctness and time complexity of Algorithm 3. The relevant technical proofs are deferred to
the supplementary material.
Theorem 3. Algorithm 3 returns a succinct classical representation ${ }^{5}$ of a vector $\bar{w} \in \mathbb{R}^{d}$ such that

$$
\begin{equation*}
A_{i} \bar{x} \geq \max _{x \in \mathbb{B}_{q}^{d}} \min _{i^{\prime} \in[n]} A_{i^{\prime}} x-\epsilon \quad \forall i \in[n], \tag{22}
\end{equation*}
$$

with probability at least $2 / 3$, and its total running time is $\tilde{O}\left(\frac{p^{2} \sqrt{n}}{\epsilon^{4}}+\frac{p^{3.5} \sqrt{d}}{\epsilon^{7}}\right)$. We can also assume $p=O(\log d / \epsilon)$ (Lemma 1) and result in running time $\tilde{O}\left(\frac{\sqrt{n}}{\epsilon^{6}}+\frac{\sqrt{d}}{\epsilon^{10.5}}\right)$.

Moreover, Algorithm 3 enjoys the following features:

- Simple quantum input: Algorithm 3 uses the standard quantum input model and needs not to use any sophisticated quantum data structures, such as quantum random access memory (QRAM) in some other quantum machine learning applications, to achieve speedups.
- Hybrid classical-quantum feature: Algorithm 3 is also highly classical-quantum hybrid: the quantum part is isolated by pieces of state preparations connected by classical processing. In addition, it only has $O\left(\frac{\log n+p}{\epsilon^{2}}\right)$ iterations, which implies that the corresponding quantum circuit is shallow and can potentially be implemented even on near-term quantum machines (Preskill 2018).
- Classical output: The output of Theorem 3 is completely classical. Compared to quantum algorithms whose output is a quantum state and may incur overheads (Aaronson 2015), Algorithm 3 guarantees minimal overheads and can be directly used for classical applications.


## Applications

We give two applications that generically follow from our classical and quantum $\ell_{q}-\ell_{1}$ matrix game solvers.

## Approximate Carathéodory Problem

The exact Carathéodory problem is a fundamental result in linear algebra and convex geometry: every point $u \in \mathbb{R}^{d}$ in the convex hull of a vertex set $S \subset \mathbb{R}^{d}$ can be expressed as a convex combination of $d+1$ vertices in $S$. Recently, a breakthrough result by Barman (2015) shows that if $S \subset \mathbb{B}_{p}^{d}$, i.e., $S$ is in the $\ell_{p}$-norm unit ball, then there exists a point $u^{\prime}$ s.t. $\left\|u-u^{\prime}\right\|_{p} \leq \epsilon$ and $u^{\prime}$ is a convex combination of $O\left(p / \epsilon^{2}\right)$ vertices in $S$. The follow-up work by Mirrokni et al. (2017) proved a matching lower bound $\Omega\left(p / \epsilon^{2}\right)$, and Combettes and Pokutta (2019) can give better bounds under stronger assumptions on $S$ or $u$.

Currently, the best-known time complexity of solving the approximate Carathéodory problem is $O\left(n d p / \epsilon^{2}\right)$ by Theorem 3.5 of Mirrokni et al. (2017). We give classical and quantum sublinear algorithms:
Corollary 1. Suppose that $S \subset \mathbb{B}_{p}^{d},|S|=n$, and $u$ is in the convex full of $S$. Then we can find a convex combination $\sum_{i=1}^{k} x_{i} v_{i}$ such that $v_{i} \in S$ for all $i \in[k]$,

[^3]$k=O\left((p+\log n) / \epsilon^{2}\right)$, and $\left\|\sum_{i=1}^{k} x_{i} v_{i}\right\|_{p} \leq \epsilon$, using $a$ classical algorithm with running time $O\left(\frac{(n+d)(p+\log n)}{\epsilon^{2}}\right)$ or a quantum algorithm with running time $\tilde{O}\left(\frac{p^{2} \sqrt{n}}{\epsilon^{4}}+\frac{p^{3.5} \sqrt{d}}{\epsilon^{7}}\right)$. We can also assume $p=O(\log d / \epsilon)$ (Lemma 1) and result in running time $O\left(\frac{(n+d)(\log d / \epsilon+\log n)}{\epsilon^{2}}\right)$ and $\tilde{O}\left(\frac{\sqrt{n}}{\epsilon^{6}}+\frac{\sqrt{d}}{\epsilon^{10.5}}\right)$, respectively.

Proof. We denote the matrix $V:=\left(v_{1} ; v_{2} ; \cdots ; v_{n}\right)$ where $v_{i}$ is the $i^{\text {th }}$ element in $S$. Note that the approximate Carathéodory problem can be formed as $\min _{\mathbf{p} \in \Delta_{n}}\left\|V^{\top} \mathbf{p}-u\right\|_{p}$. In addition, by Hölder's inequality $\|y\|_{p}=\max _{x:\|x\|_{q} \leq 1} y^{\top} x$; therefore, we obtain the following minimax matrix game:

$$
\begin{equation*}
\min _{\mathbf{p} \in \Delta_{n}} \max _{x \in \mathbb{B}_{q}^{d}}\left(\mathbf{p}^{\top} V-u^{\top}\right) x \tag{23}
\end{equation*}
$$

We denote $U=(u ; u ; \cdots ; u) \in \mathbb{R}^{n \times d}$, i.e., all the $n$ rows of $U$ are $u$. Then we have $\left(\mathbf{p}^{\top} V-u^{\top}\right) x=2 \mathbf{p}^{\top} \frac{V-U}{2} x$. Furthermore, since $u, v_{i} \in \mathbb{B}_{p}^{d}$ for all $i \in[n]$, each row of $\frac{V-U}{2}$ is also in $u, v_{i} \in \mathbb{B}_{p}^{d}$. Finally, by the Sion's Theorem (Sion 1958) we can switch the order of the min and max in (23). In all, to solve the approximate Carathéodory problem with precision $\epsilon$, it suffices to solve the maximin game

$$
\begin{equation*}
\max _{x \in \mathbb{B}_{q}^{d}} \min _{\mathbf{p} \in \Delta_{n}} \mathbf{p}^{\top} \frac{V-U}{2} x \tag{24}
\end{equation*}
$$

with precision $\frac{\epsilon}{2}$. This is exactly (6), thus the result follows from Theorem 2 and Theorem 3.

Compared to Mirrokni et al. (2017), we pay a $\log n$ overhead in the cardinality of the convex combination, but in time complexity the dominating term $n d$ is significantly improved to $n+d$. We also give the first sublinear quantum algorithm. Note that as Mirrokni et al. (2017) pointed out, the approximate Carathéodory problem has wide applications in machine learning and optimization, including support vector machines (SVMs), rounding in polytopes, submodular function minimization, etc. We elaborate the details of SVMs below, and leave out the details of other applications as the reductions are direct.

## The $\ell_{q}$-Margin Support Vector Machine (SVM)

When we solve the $\ell_{q}-\ell_{1}$ matrix game in Algorithm 1, we apply $\ell_{q}$-sampling where $j_{t}=j$ with probability $x(j)^{q} /\|x\|_{q}^{q}$ for any $j \in[d]$. The key reason of the success of Algorithm 1 is because the expectation of the random variable $A_{i}\left(j_{t}\right)\left\|x_{t}\right\|_{q}^{q} / x_{t}\left(j_{t}\right)^{q-1}$ in Line 8 is $A_{i} x$, which is unbiased.

If we consider some alternate random variables, we can potentially solve a maximin game in $\ell_{q}-\ell_{1}$ norm with respect to some nonlinear functions of the matrix. A specific problem of significant interest is the $\ell_{q}$-margin support vector machine (SVM), where we are given $n$ data points $X_{1}, \ldots, X_{n}$ in $\mathbb{R}^{d}$ and a label vector $y \in\{1,-1\}^{n}$. The goal is to find a separating hyperplane $w \in \mathbb{R}^{d}$ of these data points with the largest margin under the $\ell_{q}$-norm loss, i.e.,

$$
\begin{equation*}
\sigma_{\mathrm{SVM}}:=\max _{w \in \mathbb{R}^{d}} \min _{i \in[n]} 2 y_{i} \cdot X_{i}^{\top} w-\|w\|_{q}^{q} . \tag{25}
\end{equation*}
$$

Without loss of generality, we assume $y_{i}=1$ for all $i \in[n]$, otherwise we take $X_{i} \leftarrow(-1)^{y_{i}} \cdot X_{i}$. In this case, the random variable $2 X_{i}(j)\|w\|_{q}^{q} / w(j)^{q-1}-\|w\|_{q}^{q}$ is unbiased under $\ell_{q^{-}}$ sampling on $j$ :

$$
\begin{equation*}
\mathbb{E}\left[\frac{2 X_{i}(j)\|w\|_{q}^{q}}{w(j)^{q-1}}-\|w\|_{q}^{q}\right]=2 X_{i}^{\top} w-\|w\|_{q}^{q} \tag{26}
\end{equation*}
$$

Note that $\sigma_{\mathrm{SVM}} \geq 0$ since $2 X_{i}^{\top} w-\|w\|_{q}^{q}=0$ for all $i \in[n]$ when $w=0$. For the case $\sigma_{\text {SVM }}>0$ and taking $0<\epsilon<$ $\sigma_{\text {SVM }}$, similar to Theorem 2 and Theorem 3 we have:
Corollary 2. To return a vector $\bar{w} \in \mathbb{R}^{d}$ such that with probability at least $2 / 3$,

$$
\begin{equation*}
\min _{i \in[n]} 2 X_{i} \bar{w}-\|\bar{w}\|_{q}^{q} \geq \sigma_{\mathrm{SVM}}-\epsilon>0 \tag{27}
\end{equation*}
$$

there is a classical algorithm that achieves this with $O\left(\frac{(n+d)(p+\log n)}{\epsilon^{2}}\right)$ time and a quantum algorithm that achieves this with $\tilde{O}\left(\frac{p^{2} \sqrt{n}}{\epsilon^{4}}+\frac{p^{3.5} \sqrt{d}}{\epsilon^{7}}\right)$ time. We can also assume $p=O(\log d / \epsilon)($ Lemma 1) and result in running time $O\left(\frac{(n+d)(\log d / \epsilon+\log n)}{\epsilon^{2}}\right)$ and $\tilde{O}\left(\frac{\sqrt{n}}{\epsilon^{6}}+\frac{\sqrt{d}}{\epsilon^{10.5}}\right)$, respectively.

Notice that classical sublinear algorithms for $\ell_{2}$-SVMs have been given (Clarkson, Hazan, and Woodruff 2012; Hazan, Koren, and Srebro 2011), and there is also a sublinear quantum algorithm for $\ell_{2}$-SVMs in Li , Chakrabarti, and Wu (2019). We essentially generalize their results to the $l_{q}$-norm cases based on our new general matrix game solvers in Theorem 2 and Theorem 3.

## Classical and Quantum Lower Bounds

For both our classical and quantum algorithms for general matrix games, we can prove matching classical and quantum lower bounds in $n$ and $d$ for constant $\epsilon$ :
Theorem 4. Assume $0<\epsilon<0.04$. Then to return an $\bar{x} \in$ $\mathbb{B}_{q}^{d}$ satisfying

$$
\begin{equation*}
A_{j} \bar{x} \geq \max _{x \in \mathbb{B}_{q}^{d}} \min _{i \in[n]} A_{i} x-\epsilon \quad \forall j \in[n], \tag{28}
\end{equation*}
$$

with probability at least $2 / 3$, we need $\Omega(n+d)$ classical queries or $\Omega(\sqrt{n}+\sqrt{d})$ quantum queries.

Due to the space limitation, we postpone the proof details of Theorem 4 to the supplementary material.

## Conclusions

We give sublinear algorithms for solving general $\ell_{q}-\ell_{1}$ matrix games for any $q \in(1,2]$. Our classical and algorithms run in time $O\left(\frac{(n+d)(p+\log n)}{\epsilon^{2}}\right)$ and $\tilde{O}\left(\frac{p^{2} \sqrt{n}}{\epsilon^{4}}+\frac{p^{3.5} \sqrt{d}}{\epsilon^{7}}\right)$, respectively; both bounds are tight up to poly-logarithmic factors in $n$ and $d$. Our results can be applied to solve the approximate Carathéodory problem and the $\ell_{q}$-margin SVMs.

Our paper raises a couple of natural open questions for future work. For instance:

- Can we give sublinear algorithms for $\ell_{p}-\ell_{1}$ matrix games where $p>2$ ? Technically, this will probably require a $q^{\text {th }}$ moment multiplicative weight lemma to replace Lemma 2.
- Can we give quantum algorithms that achieve speedup of variance-reduced methods for solving matrix games, such as the state-of-the-art result in Carmon et al. (2019)?


## Acknowledgements

TL thanks Adrian Vladu for many helpful discussions, as well as Yair Carmon for the discussions about his paper (Carmon et al. 2019). TL was supported by an IBM PhD Fellowship, an QISE-NET Triplet Award (NSF grant DMR-1747426), the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Algorithms Teams program, ARO contract W911NF-17-1-0433, and NSF grant PHY-1818914. CW was supported by Scott Aaronson's Vannevar Bush Faculty Fellowship. SC and XW were partially supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Algorithms Team program, and were also partially supported by the U.S. National Science Foundation grant CCF-1755800, CCF-1816695, and CCF-1942837 (CAREER).

## Ethics Statement

This work is purely theoretical. Researchers working on learning theory and quantum computing may benefit from our results. In the long term, once fault-tolerant quantum computers have been built, our results may find practical applications in matrix game scenarios arising in the real world. As far as we are aware, our work does not have immediate negative ethical impact.

## References

Aaronson, S. 2015. Read the fine print. Nature Physics 11(4): 291.

Apeldoorn, J. v.; and Gilyén, A. 2019. Quantum algorithms for zero-sum games. arXiv:1904.03180
Arute et al., F. 2019. Quantum supremacy using a programmable superconducting processor. Nature 574(7779): 505-510. arXiv:1910.11333
Barman, S. 2015. Approximating Nash equilibria and dense bipartite subgraphs via an approximate version of Carathéodory theorem. In Proceedings of the 47th Annual ACM Symposium on Theory of Computing, 361-369. arXiv:1406.2296
Bennett, C. H.; Bernstein, E.; Brassard, G.; and Vazirani, U. 1997. Strengths and weaknesses of quantum computing. SIAM Journal on Computing 26(5): 1510-1523. arXiv:quant-ph/9701001
Brassard, G.; Høyer, P.; Mosca, M.; and Tapp, A. 2002. Quantum amplitude amplification and estimation. Contemporary Mathematics 305: 53-74. arXiv: quant-ph/0005055
Carmon, Y.; Jin, Y.; Sidford, A.; and Tian, K. 2019. Variance reduction for matrix games. In Advances in Neural Information Processing Systems, 11377-11388. arXiv:1907.02056
Clarkson, K. L.; Hazan, E.; and Woodruff, D. P. 2012. Sublinear optimization for machine learning. Journal of the ACM (JACM) 59(5): 23. arXiv:1010. 4408
Combettes, C. W.; and Pokutta, S. 2019. Revisiting the Approximate Carathéodory Problem via the Frank-Wolfe Algorithm. arXiv:1911.04415

Dantzig, G. B. 1998. Linear programming and extensions, volume 48. Princeton University Press.
Deng, N.; Tian, Y.; and Zhang, C. 2012. Support vector machines: optimization based theory, algorithms, and extensions. CRC press.
Dürr, C.; and Høyer, P. 1996. A quantum algorithm for finding the minimum. arXiv:quant-ph/9607014
Garber, D.; and Hazan, E. 2011. Approximating semidefinite programs in sublinear time. In Advances in Neural Information Processing Systems, 1080-1088.
Grigoriadis, M. D.; and Khachiyan, L. G. 1995. A sublineartime randomized approximation algorithm for matrix games. Operations Research Letters 18(2): 53-58.
Grover, L. K. 1996. A fast quantum mechanical algorithm for database search. In Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing, 212-219. ACM. arXiv:quant-ph/9605043
Hazan, E.; Koren, T.; and Srebro, N. 2011. Beating SGD: Learning SVMs in sublinear time. In Advances in Neural Information Processing Systems, 1233-1241.
Johnson, R.; and Zhang, T. 2013. Accelerating stochastic gradient descent using predictive variance reduction. In Advances in Neural Information Processing Systems, 315-323.
Kapoor, A.; Wiebe, N.; and Svore, K. 2016. Quantum perceptron models. In Proceedings of the 30th Conference on Neural Information Processing Systems, 3999-4007. arXiv:1602.04799
Li, T.; Chakrabarti, S.; and Wu, X. 2019. Sublinear quantum algorithms for training linear and kernel-based classifiers. In Proceedings of the 36th International Conference on Machine Learning, 3815-3824. arXiv:1904.02276
Minsky, M.; and Papert, S. A. 1988. Perceptrons: An introduction to computational geometry. MIT Press.
Mirrokni, V.; Leme, R. P.; Vladu, A.; and Wong, S. C.-w. 2017. Tight bounds for approximate Carathéodory and beyond. In Proceedings of the 34th International Conference on Machine Learning, 2440-2448. arXiv:1512.08602
Nemirovski, A. 2004. Prox-method with rate of convergence $O(1 / t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. SIAM Journal on Optimization 15(1): 229-251.
Nemirovski, A. S.; and Yudin, D. B. 1983. Problem complexity and method efficiency in optimization .
Nesterov, Y. 2007. Dual extrapolation and its applications to solving variational inequalities and related problems. Mathematical Programming 109(2-3): 319-344.
Neumann, J. v. 1928. Zur theorie der gesellschaftsspiele. Mathematische Annalen 100(1): 295-320.
Novikoff, A. B. 1963. On convergence proofs for perceptrons. In Proceedings of the Symposium on the Mathematical Theory of Automata, volume 12, 615-622.
Palaniappan, B.; and Bach, F. 2016. Stochastic variance reduction methods for saddle-point problems. In Advances
in Neural Information Processing Systems, 1416-1424. arXiv:1605.06398
Preskill, J. 2018. Quantum Computing in the NISQ era and beyond. Quantum 2: 79. ISSN 2521-327X. arXiv:1801.00862
Shalev-Shwartz, S. 2012. Online learning and online convex optimization. Foundations and Trends® in Machine Learning 4(2): 107-194.
Shi, Z.; Zhang, X.; and Yu, Y. 2017. Bregman divergence for stochastic variance reduction: saddle-point and adversarial prediction. In Advances in Neural Information Processing Systems, 6031-6041.
Sion, M. 1958. On general minimax theorems. Pacific Journal of Mathematics 8(1): 171-176.
Zinkevich, M. 2003. Online convex programming and generalized infinitesimal gradient ascent. In Proceedings of the 20th International Conference on Machine Learning, 928936.


[^0]:    *Equal contribution. Full version of the paper (including the supplementary material) is at https://arxiv.org/abs/2012.06519
    Copyright (C) 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1}$ Throughout the paper, we use the bold font $\mathbf{p}$ to denote a vector and the math font $p$ to denote a real number.
    ${ }^{2} \bar{x} \in \mathbb{B}_{q}^{d}$ is the standard objective quantity under the $\ell_{q}$-norm. Also note that once we have the $\bar{x}$ in (3), any $\mathbf{p} \in \Delta_{n}$ satisfies $\mathbf{p}^{\top} A \bar{x} \geq \sigma-\epsilon$.
    ${ }^{3}$ Here $\tilde{O}$ omits poly-logarithmic factors.

[^2]:    ${ }^{4}$ The proof follows from the proofs of Lemmas 2.3, 2.4, 2.5, and 2.6 in Section 2 and Appendix B of Clarkson, Hazan, and Woodruff (2012), with only small modifications to fit our new parameter choices. For instance, the original statement requires that $\eta \geq \sqrt{\frac{\log n}{T}}$, but the proofs actually work for $\eta \geq \sqrt{\frac{11 \log n}{12 T}}$.

[^3]:    ${ }^{5}$ The algorithm stores $T=\tilde{O}\left(p / \epsilon^{2}\right)$ real numbers classically: $i_{1}, \ldots, i_{T}$ obtained from Line 3 and $\widetilde{\left\|y_{1}\right\|_{q}}, \ldots, \widetilde{\left\|y_{T}\right\|_{q}}$ obtained from Line 6 . After that, each coordinate of $\bar{x}$ can be computed in time $\tilde{O}\left(p / \epsilon^{2}\right)$.

