

Differential Privacy Amplification in Quantum and Quantum-inspired Algorithms

Armando Angrisani¹, Mina Doosti², Elham Kashefi^{1,2}

¹*LIP6, CNRS, Sorbonne Université, 75005 Paris, France*

²*School of Informatics, University of Edinburgh, EH8 9AB Edinburgh, United Kingdom*

Abstract

Differential privacy provides a theoretical framework for processing a dataset about n users, in a way that the output reveals a minimal information about any single user. Such notion of privacy is usually ensured by noise-adding mechanisms and amplified by several processes, including subsampling, shuffling, iteration, mixing and diffusion. In this work, we provide privacy amplification bounds for quantum and quantum-inspired algorithms. In particular, we show for the first time, that algorithms running on quantum encoding of a classical dataset or the outcomes of quantum-inspired classical sampling, amplify differential privacy. Moreover, we prove that a quantum version of differential privacy is amplified by the composition of quantum channels, provided that they satisfy some mixing conditions.

1 Introduction

Differential Privacy (DP) [1, 2] is a rigorous mathematical framework for preserving the information of each individual in a dataset while enabling to analyse and process the dataset. Intuitively, a differentially private algorithm can learn a statistical property of a dataset consisting of n users, yet it leaks *almost* nothing about each individual user. Such mechanisms are of great interest and importance when dealing with sensitive data like hospital data, banks, social media, etc. Apart from privacy-preserving data analysis, differential privacy has also found several applications in other fields of computer science such as machine learning [3, 4, 5, 6], statistical learning theory [7, 8, 9, 10], mechanism design [11].

Since its introduction, multiple analytical tools for the design of private data analyses have been developed [12, 13, 14, 15]. Most commonly, these mechanisms exploit techniques like adding noise to the final output or randomizing the input. A loose analysis of complex mechanisms built out of these blocks can be conducted with simple tools, such as basic composition rules and robustness to post-processing. However, the inherent trade-off between privacy and utility in practical applications ignited the development of more refined rules leading to tighter privacy bounds. A trend in this direction is to show that several sources of randomness amplify the guarantees of standard DP mechanisms. In particular, DP amplification results have been shown for subsampling, iteration, mixing and shuffling [16, 17, 18, 19].

Given the major influence of quantum computing and quantum information in the past decades over different areas of computer science, an interesting question is whether quantum and quantum-inspired algorithms can enhance differential privacy. This question becomes specifically more relevant with the availability of Noisy Intermediate Scale Quantum devices (NISQ) today [20]. The noisy nature of these devices (also previously exploited by [21]) on one hand, and the potential capabilities of quantum algorithms, on the other hand, make such quantum or hybrid quantum-classical mechanisms, an interesting subject of study from the point of view of differential privacy. Furthermore, the connection between machine learning and differential privacy suggests that answering this question can lead to intriguing insights into the capabilities of quantum machine learning.

Differential privacy has been extended to quantum computation in [22] and [23]. One of the main challenges in translating the definition of DP in the quantum setting is to characterise the notion of *neighbouring quantum states*. Recall that, in the classical setting, two neighbouring datasets differ in at most one single entry. In the two mentioned works, the adopted definitions of neighbouring quantum states are significantly different, and are respectively based on bounded trace distance and single-register measurements. For the purpose of this paper, we follow the definition of [22]. Moreover, quantum private PAC learning has been defined in [24], and a quantum analog of the equivalence between private classification and online prediction has been shown in [10].

Our contributions. In this paper, we initiate a systematic study of differential privacy amplification in quantum and quantum-inspired algorithms. We provide three types of results. **Section 3** provides privacy amplification results when the (classical) data is encoded into a quantum state. Informally, we first show that quantum encoding of classical datasets leads to approximate classical differential privacy. Moreover, we show that a quantum DP operation performed on the quantum encoding of a classical dataset satisfies also classical DP, under some suitable assumptions. These two primary results show a general application of quantum information and quantum encoding for differential privacy and can be employed further to design sophisticated differentially private mechanisms both in the classical and quantum setting. Moreover, we prove that the composition of quantum encoding with noisy mechanisms such as the Laplace and Gaussian mechanisms amplifies differential privacy.

Similarly, **Section 4** investigates the case of *quantum-inspired* algorithms, a family of classical algorithms equipped with an ℓ_2 -norm sampling oracle. We show that differential privacy, both in the exact and approximate setting, is amplified via quantum-inspired subsampling, establishing the concrete amplification bounds.

Finally, our last results concern quantum differential privacy. As for classical DP, quantum DP is preserved under post-processing [22]. In **Section 5**, we show an amplification result for the post-processed quantum channel $S \circ T$, provided that T satisfies a quantum analog of the Dobrushin or the Doeblin condition. This is yet another general result that relies on the contraction property of quantum channels and can be exploited to enhance differential privacy. We expand this result by finding explicit bounds for differential privacy under the mentioned condition for the composition of two well-known quantum channels, namely the *generalized amplitude damping channel* and *depolarizing channel*. Furthermore, we show that the Dobrushin condition together with unitality provides pure quantum differential privacy.

2 Preliminaries

We start by introducing notation and concepts that will be used throughout the paper.

Quantum information. We briefly review the basic concepts in quantum information that we will use throughout the paper. A d -dimensional *pure state* is a unit vector in \mathbb{C}^d , written in ket notation as

$$|\psi\rangle = \sum_{i=1}^d \alpha_i |i\rangle.$$

Here $|1\rangle, \dots, |d\rangle$ is an orthonormal basis for \mathbb{C}^d , and the α_i 's are complex numbers called *amplitudes* satisfying $|\alpha_1|^2 + \dots + |\alpha_d|^2 = 1$. The notation $|\cdot\rangle$ reminds of the fact that the Hilbert space has an inner product $\langle \cdot, \cdot \rangle$, which for Hilbert spaces describing quantum systems is denoted

as $\langle \cdot | \cdot \rangle$. The left side of the inner product $\langle \psi |$ is the conjugate transpose of the quantum state $|\psi\rangle$. Then the unit-norm condition can be expressed succinctly as $\langle \psi | \psi \rangle = 1$.

In general, we may also have classical probability distributions over pure states. This scenario is captured by the formalism of *mixed* states, which generalises all the states in quantum mechanics including pure states. Mixed states are conveniently described by density matrices. Formally, a d -dimensional mixed state ρ is a $d \times d$ positive semidefinite matrix that satisfies $\text{Tr}(\rho) = 1$. Equivalently, ρ is a convex combination of outer products of pure states with themselves:

$$\rho = \sum_{i=1}^d p_i |\psi_i\rangle \langle \psi_i|,$$

where $p_i \geq 0$ and $\sum_i p_i = 1$. In the special case where $p_i = 1$, we obtain a pure state $\rho = |\psi_i\rangle \langle \psi_i|$. Several norms and distance measures can also be defined for general quantum states. Given a Hermitian matrix A with eigenvalues $\lambda_1, \dots, \lambda_k$, its trace norm is defined as $\|A\|_{\text{tr}} := \frac{1}{2} \sum_{i=1}^d \lambda_i$. The trace norm induces the *trace distance* $\|\rho - \sigma\|_{\text{tr}} = \text{Tr}(|\rho - \sigma|)/2$. For a pair of pure states, the trace distance can be linked to their inner product,

$$\| |\psi\rangle \langle \psi| - |\phi\rangle \langle \phi| \|_{\text{tr}} = \sqrt{1 - |\langle \psi | \phi \rangle|^2}. \quad (1)$$

A *superoperator* S maps a mixed state ρ to the mixed state $S(\rho) = \sum_{i=1}^k B_i \rho B_i^\dagger$, where B_1, \dots, B_k can be any matrices satisfying $\sum_{i=1}^k B_i^\dagger B_i = \mathbb{I}$. This is the most general (norm-preserving) mapping from mixed states to mixed states allowed by quantum mechanics. If we drop the norm-preserving condition and we let $\sum_{i=1}^k B_i^\dagger B_i \leq \mathbb{I}$, then we call S a *quantum operation*. Quantum operations act linearly on mixed states, in the sense that for any $a, b \in \mathbb{C}$, $S(a\rho + b\sigma) = aS(\rho) + bS(\sigma)$. Moreover, quantum operations are *non trace-increasing*. For any hermitian matrix A , $\|S(A)\|_{\text{tr}} \leq \|A\|_{\text{tr}}$. In particular, $\|S(\rho) - S(\sigma)\|_{\text{tr}} \leq \|\rho - \sigma\|_{\text{tr}}$.

The most general class of measurements to perform on mixed states are the POVMs (Positive Operator Valued Measures). In the POVM formalism, a measurement M with possible outcomes $1, 2, \dots, k$ is given by a list of $d \times d$ positive semidefinite matrices E_1, \dots, E_k , which satisfy $\sum_i E_i = \mathbb{I}$. The measurement rule is:

$$\Pr[M(\rho) \text{ returns outcome } i] = \text{Tr}(E_i \rho).$$

Notably, trace distance has also the following physical interpretation

$$\|\rho - \sigma\|_{\text{tr}} = \max_M \Pr[M(\rho) \text{ accepts}] - \Pr[M(\sigma) \text{ accepts}], \quad (2)$$

where the maximum is taken over all possible two-outcome measurements. We also define the infinity norm of a matrix A as the maximum of the absolute row sum value as follows,

$$\|A\|_{\infty} = \max_i \sum_j^n |a_{ij}|, \quad (3)$$

Furthermore, we need to define the operator infinity norm or ℓ^∞ -norm. The vector space ℓ^∞ is a sequence space whose elements are the bounded sequences. The ℓ^∞ space in a Banach space with respect to the following norm,

$$\|x\|_{\infty} = \sup_n |x_n|, \quad (4)$$

The operator norm on the Hilbert space is defined over the space of bounded linear operators as,

$$\|O\|_{\infty} = \sup \|Ox\| : \forall \|x\| \leq 1, \quad (5)$$

We also note that for the operator norms, $\|\cdot\|_1$ is the dual norm of $\|\cdot\|_{\infty}$ [25].

Differential privacy. In the standard model of differential privacy, a trusted curator collects the raw data of the individuals and is responsible for their privacy. On the contrary, in the *local* model the curator is possibly malicious, and hence each individual submits their own privatized data. More formally, consider a statistical dataset, i.e. a vector $x = (x_1, \dots, x_n)$ over a domain X , where each entry $x_i \in X$ represents information contributed by a single individual. datasets x and x' are neighbors if $x_i \neq x'_i$ for exactly one $i \in [n]$. We denote the neighbor relation with $x \sim x'$. A randomized algorithm \mathcal{A} is (ϵ, δ) -differentially private if for any two neighbor datasets x, x' and for every subset F of the possible outcomes of \mathcal{A} we have

$$\Pr[\mathcal{A}(x) \in F] \leq e^\epsilon \Pr[\mathcal{A}(x') \in F] + \delta.$$

We denote as *pure* differential privacy the special case where $\delta = 0$, while in the most general case we have *approximate* differential privacy. One popular method to ensure $(\epsilon, 0)$ -DP is the *Laplace mechanism*. Given a function $f : \mathcal{X}^n \rightarrow \mathbb{R}^k$ we define its ℓ_1 -sensitivity as $\Delta = \max_{x \sim x'} |f(x') - f(x)|$. The Laplace mechanism consists in adding a random perturbation η to $f(x)$, where $\eta \sim \text{Laplace}(\Delta/\epsilon) := \frac{\epsilon}{2\Delta} \exp(-|\eta| \frac{\epsilon}{\Delta})$.

An additional widely-used method is the Gaussian mechanism, that ensures (ϵ, δ) -DP. Given a function f as defined above, the Gaussian mechanism consists in adding a random perturbation η to $f(x)$, where $\eta \sim \mathcal{N}(0, \sigma^2) := \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{\eta^2}{2\sigma^2}\right)$ and $\sigma^2 = 2 \ln(1.25/\delta) \Delta^2 / \epsilon^2$.

We now turn our attention to the *local model*. Following the notation used in [7], we say that a randomized algorithm over datasets is (ϵ, δ) -local differentially private if it's an (ϵ, δ) -differentially private algorithm that takes in input a dataset of size $n = 1$.

The most common mechanism for local differential privacy is *randomized response* (RR). For a dataset $x = (x_1, \dots, x_n) \in \{0, 1\}^n$, each user x_i outputs a random bit z_i , such that $z_i = x_i$ with probability $\frac{e^\epsilon + \delta}{1 + e^\epsilon}$ and $z_i = 1 - x_i$ with probability $\frac{1 - \delta}{1 + e^\epsilon}$. It is easy to see that any algorithm run on $z = (z_1, \dots, z_n)$ is (ϵ, δ) -local differentially private.

Interestingly, differential privacy is related to several desired learnability properties, including robustness, stability and generalization. Concerning robustness to adversarial examples, we recall here the result stated in (Lemma 1, [26]). Let $B_p(r) := \{\alpha \in \mathbb{R}^n : \|\alpha\|_p \leq r\}$ be the p -norm ball of radius r . For a given classification model f and a fixed input $x \in \mathbb{R}^n$, an attacker is able to craft a successful adversarial example of size L for a given p -norm if they find $\alpha \in B_p(L)$ such as $f(x + \alpha) \neq f(x)$. The attacker thus tries to find a small change to x that will significantly change the predicted label. Now, suppose that a randomized function A , with bounded output $A(x) \in [0, b]$, $b \in \mathbb{R}^+$, satisfies (ϵ, δ) -DP. Then the expected value of its output meets the following property:

$$\forall \alpha \in B_p(1) : \mathbb{E}(A(x)) \leq e^\epsilon \mathbb{E}(A(x + \alpha)) + b\delta$$

where the expectation is taken over the randomness in A .

Following the approach proposed in [22], we say that a quantum operation \mathcal{E} is (τ, ϵ, δ) -quantum differentially private (QDP), if for every POVM M , for all subset S of the possible outcomes, and for all inputs ρ, σ such that $\|\rho - \sigma\|_{\text{tr}} \leq \tau$,

$$\Pr[M(\mathcal{E}(\rho)) \text{ output is in } S] \leq e^\epsilon \Pr[M(\mathcal{E}(\sigma)) \text{ output is in } S] + \delta. \quad (6)$$

Theorem 1 (Proposition 1, [22]). *Let \mathcal{E} be a quantum operation that is (τ, ϵ, δ) -QDP. Let \mathcal{F} be an arbitrary quantum operation. Then the composition of \mathcal{E} and \mathcal{F}*

$$\mathcal{F} \circ \mathcal{E} : \rho \mapsto \mathcal{F}(\mathcal{E}(\rho))$$

is (τ, ϵ, δ) -QDP.

3 Amplification by quantum encoding

In quantum computation, a classical dataset $x \in \mathcal{X}$ can be mapped to a quantum state with a *data-encoding feature map*, also referred as quantum encoding, that is a classical-to-quantum transformation

$$\phi(x) = |\phi(x)\rangle \langle\phi(x)| = \rho(x).$$

Given a dataset $x = (x_1, \dots, x_n)$, where each x_i is a binary string, one of the most common information encoding strategy is the *basis encoding*, which is described by a uniform superposition of computational basis states

$$x \mapsto \frac{1}{\sqrt{n}} \sum_{i=1}^n |x_i\rangle.$$

For a complex value dataset $x \in \mathbb{C}^n$, it's convenient to adopt the *amplitude encoding*,

$$x \mapsto \sum_{i=1}^n x_i |i\rangle.$$

We always assume that the input vector x is normalised as $\|x\|^2 = \sum_i |x_i|^2 = 1$. For convenience, we set the following parameter, that will be employed in the following.

$$\Gamma(x) := \max_j |x_j|^2. \quad (7)$$

For a dataset $x \in [0, 2\pi]^n$ we can define the *rotation encoding*,

$$x \mapsto \sum_{q_1 \dots q_n = 0}^1 \prod_{k=1}^n \cos(x_k)^{q_k} \sin(x_k)^{1-q_k} |q_1, \dots, q_n\rangle.$$

As noted in [27], a quantum encoding gives rise to a *quantum kernel*, which is the inner product between two data-encoding feature vectors. For any $x, x' \in \mathcal{X}$, the quantum kernel induced by ϕ is

$$\kappa_\phi(x, x') = \|\rho(x)\rho(x')\|_{\text{tr}} = |\langle\phi(x)|\phi(x')\rangle|^2. \quad (8)$$

Since we are dealing with differential privacy, quantum kernels evaluated on neighbor inputs are of particular interest. To this end, we define the *quantum minimum adjacent kernel*

$$\hat{\kappa}_\phi := \min_{x \sim x'} \kappa_\phi(x, x').$$

The expressions of κ_ϕ and $\hat{\kappa}_\phi$ for the quantum encodings defined above can be found in [Table \(1\)](#). We refer to [27] for more details and more examples of quantum kernels.

We observe that the minimum adjacent kernels allows us to connect the quantum and classical definition of differential privacy.

Lemma 1 (Quantum-to-classical DP). *Let $x \in \mathcal{X}$ and let \mathcal{A} be a quantum algorithm that takes as input only $\rho(x) = |\phi(x)\rangle \langle\phi(x)|$ and perform a $(\sqrt{1 - \hat{\kappa}_\phi}, \varepsilon, \delta)$ -QDP quantum operation \mathcal{E} on $\rho(x)$. Then $\mathcal{A} \circ \rho$ is (ε, δ) -DP.*

Proof. Let x, x' two neighbouring datasets and $|\phi(x)\rangle, |\phi(x')\rangle$ their corresponding encodings. By definition, their trace distance is upper bounded by the *minimum adjacent kernel*,

$$\|\rho(x) - \rho(x')\|_{\text{tr}} \leq \sqrt{1 - |\langle\phi(x)|\phi(x')\rangle|^2} := \sqrt{1 - \kappa_\phi(x, x')} \leq \sqrt{1 - \hat{\kappa}_\phi}.$$

Table 1: Quantum kernels and quantum minimum adjacent kernels for several quantum encodings. Here $\delta_{x,y}$ is the Kronecker function and $\Gamma(x)$ is the parameter defined in Eq. (7).

ENCODING ϕ	$\kappa_\phi(x, x')$	$\hat{\kappa}_\phi$
BASIS ENCODING	$\sum_i \delta_{x_i, x'_i}$	$1 - 1/n$
AMPLITUDE ENCODING	$ x^\dagger x' ^2$	$1 - \Gamma(x)$
ROTATION ENCODING	$\prod_i \cos(x_i - x'_i) ^2$	0

Algorithm 1 Composition of quantum encoding and a noise-adding mechanism

Input: a dataset $x = (x_1, \dots, x_n)$, a POVM M with outcomes in $\{0, 1\}$, a distribution \mathcal{D}
for $i = 1$ **to** m **do**
 Perform the feature map $x \mapsto \rho(x) = |\phi(x)\rangle \langle \phi(x)|$
 Apply M to $\rho(x)$ and store the output in y_i
end for
Compute the mean $\mu = \frac{1}{m} \sum_{i=1}^m y_i$ and output $O = \mu + \eta$, where $\eta \sim \mathcal{D}$.

By definition of *quantum differential privacy*, for any measurement M , and for any subset S of the possible outcomes,

$$\Pr[M(\mathcal{E}(\rho(x))) \text{ output is in } S] \leq e^\epsilon \Pr[M(\mathcal{E}(\rho(x'))) \text{ output is in } S] + \delta.$$

Since quantum DP is robust to post-processing ([Theorem 1](#)), the former inequality implies that the algorithm $\mathcal{A} \circ \rho$ is (ϵ, δ) -DP. \square

Moreover, we can show that if $\hat{\kappa}_\phi$ is larger than 0, then any quantum algorithm that accesses solely the quantum encoding of a classical dataset satisfies approximate differential privacy.

Lemma 2 (Approximate DP by quantum encoding). *Let $x \in \mathcal{X}$ and let \mathcal{A} be a quantum algorithm that takes as input only $\rho(x) = |\phi(x)\rangle \langle \phi(x)|$. Then $\mathcal{A} \circ \rho$ is $(0, \sqrt{1 - \hat{\kappa}_\phi})$ -DP.*

Proof. Since differential privacy is preserved under post-processing, we can assume that \mathcal{A} consists of a quantum operation S followed by a POVM measurement M . Let F be a subset of the possible outcomes of M . We define the two-outcome measurement M' such that M' runs M and accepts if the resulting outcome is in F , otherwise it rejects. Plugging [Eq. \(8\)](#) and [Eq. \(1\)](#) into [Eq. \(2\)](#), we get the following bound:

$$\begin{aligned} & \Pr[M(S|\phi(x)) \in F] - \Pr[M(S|\phi(x')) \in F] \\ &= \Pr[M'(S|\phi(x)) \text{ accepts}] - \Pr[M'(S|\phi(x')) \text{ accepts}] \\ &\leq \|S(\rho(x)) - S(\rho(x'))\|_{\text{tr}} \leq \|\rho(x) - \rho(x')\|_{\text{tr}} = \sqrt{1 - |\langle \phi(x) | \phi(x') \rangle|^2} \leq \sqrt{1 - \hat{\kappa}_\phi}. \end{aligned}$$

Thus the algorithm $\mathcal{A} \circ \rho$ is $(0, \sqrt{1 - \hat{\kappa}_\phi})$ -DP. \square

So far we have shown that quantum encodings inherently provide approximate differential privacy, under some suitable assumptions. In the following, we study the interaction of quantum encoding and classical noise-adding mechanisms, as sketched in [Algorithm \(1\)](#). In particular, we show DP amplification results for the Laplace and Gaussian mechanisms.

Theorem 2 (Composition of quantum encoding and Laplace mechanism). *Let x, m, M, ρ be as in [Algorithm \(1\)](#). Let $\mathcal{D} = \text{Laplace}(\frac{1}{\varepsilon}(\sqrt{1 - \hat{\kappa}_\phi} + t))$ for any $t \geq 0$. Then [Algorithm \(1\)](#) is $(\varepsilon, 0)$ -DP with exponentially high probability in t and m .*

Proof. Let x, x' be two neighbouring datasets. Denote as y'_i and μ' the outcomes of M on input $\rho(x')$ and their mean, respectively. First, we make the following observation

$$|\mathbb{E}[M(\rho(x)) - M(\rho(x'))]| = |\Pr[M(\rho(x)) = 1] - \Pr[M(\rho(x')) = 1]| \leq \sqrt{1 - \hat{\kappa}_\phi}.$$

We apply now the Chernoff-Hoeffding's bound

$$\Pr \left[\left| \frac{1}{m} \sum_{i=1}^m y_i - \mathbb{E}[M(\rho(x))] \right| \geq \frac{t}{2} \right] \leq 2e^{-mt^2}.$$

Thus with probability $1 - 4e^{-mt^2}$ the means μ and μ' are within an additive factor $t + \sqrt{1 - \hat{\kappa}_\phi}$. Since $\eta \sim \text{Laplace}((t + \sqrt{1 - \hat{\kappa}_\phi})/\varepsilon)$, we have that for any $z \in \mathbb{R}$,

$$\Pr[\mu + \eta = z] \leq e^\varepsilon \Pr[\mu' + \eta = z]$$

with probability at least $1 - 4e^{-mt^2}$. In other words, a Laplace perturbation of parameter $(t + \sqrt{1 - \hat{\kappa}_\phi})/\varepsilon$ ensures $(\varepsilon, 0)$ -DP with high probability in t and m . \square

Theorem 3 (Composition of quantum encoding and Gaussian mechanism). *Let x, m, M, ρ be as in [Algorithm \(1\)](#). Let $\mathcal{D} = \mathcal{N}(0, \sigma^2)$ where $\sigma^2 = 2 \ln(1.25/\delta)(\sqrt{1 - \hat{\kappa}_\phi} + t)^2/\varepsilon^2$ for any $t \geq 0$. Then [Algorithm \(1\)](#) is (ε, δ) -DP with exponentially high probability in t and m .*

Proof. Let x, x' be two neighbouring datasets. Denote as y'_i and μ' the outcomes of M on input $\rho(x')$ and their mean, respectively. By applying the same arguments of the proof of [Theorem 2](#), we show that, with probability $1 - 4e^{-mt^2}$ the means μ and μ' are within an additive factor $t + \sqrt{1 - \hat{\kappa}_\phi}$.

Since $\eta \sim \mathcal{N}(0, \sigma^2)$ with $\sigma^2 = 2 \ln(1.25/\delta)(\sqrt{1 - \hat{\kappa}_\phi} + t)^2/\varepsilon^2$, we have that for any $z \in \mathbb{R}$,

$$\Pr[\mu + \eta = z] \leq e^\varepsilon \Pr[\mu' + \eta = z] + \delta.$$

with probability at least $1 - 4e^{-mt^2}$. So we proved that a Gaussian perturbation of parameter σ^2 ensures (ε, δ) -DP with high probability in t and m . \square

We can obtain explicit privacy amplification bounds by combining the values of $\hat{\kappa}_\phi$ in [Table \(1\)](#) with the results of this section. Remark that, unlike the basis and the amplitude encoding, the rotation encoding provides no privacy amplification, as $\hat{\kappa}_\phi = 0$ in that case.

4 Amplification by quantum-inspired sampling

In quantum-inspired algorithms [[28](#), [29](#), [30](#), [31](#)], we simulate the measurement of $|x\rangle^{\otimes m}$ in the computational basis and process the outcomes with a classical algorithm. Quantum-inspired subsampling generalizes the uniform subsampling. Indeed, uniform subsampling can be recovered as a special case when $\Gamma(x) = 1/n$.

We will show the intuitive fact that quantum-inspired subsampling amplifies DP. The proof closely follows the one of (Problem 1.b, [[32](#)]) for uniform subsampling, but we include it here for completeness. Given a normalised vector $x = (x_1, \dots, x_n) \in \mathbb{C}^n$, let $|x\rangle := \sum_{i=1}^n x_i |i\rangle$ be the amplitude encoding defined in the previous section.

Theorem 4 (DP amplification by quantum-inspired sampling). *For any $x \in \mathbb{C}^n$, let $s = (s_1, \dots, s_m)$ be the measurement outcomes in the computational basis of $|x\rangle^{\otimes m}$. Denote \mathcal{S} as the sampling mechanism that maps x into s . Let \mathcal{A} be a (ε, δ) -DP algorithm that takes only s as input. Then $\mathcal{A}' = \mathcal{A} \circ \mathcal{S}$ is (ε', δ') -DP, with $\varepsilon' = \log(1 + (e^\varepsilon - 1)\Gamma(x)m)$ and $\delta' = \delta\Gamma(x)m$.*

Proof. We will use $T \subseteq \{1, \dots, n\}$ to denote the identities of the m -subsampled elements s_1, \dots, s_m (i.e. their index, not their actual value). Note that T is a random variable, and that the randomness of $\mathcal{A}' := \mathcal{A} \circ \mathcal{S}$ includes both the randomness of the sample T and the random coins of \mathcal{A} . Let $x \sim x'$ be adjacent datasets and assume that x and x' differ only on some row t . Let s (or s') be a subsample from x (or x') containing the rows in T . Let F be an arbitrary subset of the range of \mathcal{A} . For convenience, define $p = \Gamma(x)m$. Note that, by definition of quantum amplitude encoding and by union bound,

$$\Pr[i \in T] \leq m \Pr[|x\rangle \text{ collapses to state } |i\rangle] = |x_i|^2 \leq m\Gamma(x) := p$$

To show $(\log(1 + p(e^\varepsilon - 1)), p\delta)$ -DP, we have to bound the ratio

$$\frac{\Pr[\mathcal{A}'(x) \in F] - p\delta}{\Pr[\mathcal{A}'(x') \in F]} \leq \frac{p\Pr[\mathcal{A}(s) \in F | i \in T] + (1-p)\Pr[\mathcal{A}(s) \in F | i \notin T] - p\delta}{p\Pr[\mathcal{A}(s') \in F | i \in T] + (1-p)\Pr[\mathcal{A}(s') \in F | i \notin T]}$$

by $p(1 + (e^\varepsilon - 1))$. For simplicity, define the quantities

$$\begin{aligned} C &= \Pr[\mathcal{A}(s) \in F | i \in T] \\ C' &= \Pr[\mathcal{A}(s') \in F | i \in T] \\ D &= \Pr[\mathcal{A}(s) \in F | i \notin T] = \Pr[\mathcal{A}(s') \in F | i \notin T]. \end{aligned}$$

We can rewrite the ratio as

$$\frac{\Pr[\mathcal{A}'(x) \in F] - p\delta}{\Pr[\mathcal{A}'(x') \in F]} = \frac{pC + (1-p)D - p\delta}{pC' + (1-p)D}.$$

Now we use the fact that, by (ε, δ) -DP, $C \leq \min\{C', D\} + \delta$. Plugging all together, we get

$$\begin{aligned} pC + (1-p)D - p\delta &\leq p(e^\varepsilon \min\{C', D\}) + (1-p)D \\ &\leq p(\min\{C', D\} + (e^\varepsilon - 1)\min\{C', D\}) + (1-p)D \\ &\leq p(C' + (e^\varepsilon - 1)(pC' + (1-p)D)) + (1-p)D \\ &\leq (pC' + (1-p)D) + p(e^\varepsilon - 1)(pC' + (1-p)D) \leq (1 + p(e^\varepsilon - 1))(pC' + (1-p)D), \end{aligned}$$

where the third-to-last line follow from $\min\{x, y\} \leq \alpha x + (1 - \alpha)y$ for every $0 \leq \alpha \leq 1$. To conclude the proof, we rewrite the ratio and get the desired bound.

$$\frac{\Pr[\mathcal{A}'(x) \in F] - p\delta}{\Pr[\mathcal{A}'(x') \in F]} \leq 1 + p(e^\varepsilon - 1).$$

□

If we don't require \mathcal{A} to be (ε, δ) -DP, we obtain the following corollary as a special case of [Theorem 4](#).

Corollary 1 (Approximate DP by quantum-inspired sampling). *For any $x \in \mathbb{C}^n$, let $s = (s_1, \dots, s_m)$ be the measurement outcomes in the computational basis of $|x\rangle^{\otimes m}$. Denote \mathcal{S} as the sampling mechanism that maps x into s . Let \mathcal{A} be an algorithm that takes only s as input. Then $\mathcal{A} \circ \mathcal{S}$ is $(0, \Gamma(x)m)$ -DP.*

5 Amplification by quantum evolution

In this section, we look at quantum operations and how they can amplify differential privacy. First, we show a general result regarding the QDP amplification for distance-decreasing quantum operations and then we explore some explicit examples for certain classes of quantum channels. To establish our results, we first need to characterize quantum channels, in terms of the quantum analogs of the classical mixing conditions introduced in [33, 34].

Definition 1. Let $T : \mathcal{H} \rightarrow \mathcal{H}'$ be a quantum operation and $\gamma \in [0, 1]$. We say that T is:

1. γ -Dobrushin if

$$\sup_{\rho \neq \sigma} \frac{\|T(\rho) - T(\sigma)\|_{\text{tr}}}{\|\rho - \sigma\|_{\text{tr}}} \leq \gamma.$$

2. γ -Doebelin if there exists a quantum operation $T' : \mathcal{H} \rightarrow \mathcal{H}'$ such that $T'(X) = \text{Tr}[X]Y$ for some $Y \in \mathcal{H}'$ and $T - \gamma T'$ is positive.

We remark that the $\eta^{\text{Tr}}(T) := \inf\{\gamma : T \text{ is } \gamma\text{-Dobrushin}\}$, where $\eta^{\text{Tr}}(T)$ is the quantum Dobrushin coefficient introduced in [35, 25].

Lemma 3 (adapted from [36], Theorem 8.17). Let T be a γ -Doebelin quantum operation with $\gamma \in [0, 1]$. Then T is a $(1 - \gamma)$ -Dobrushin quantum operation.

Thus, we show that a post-processed quantum channel $\mathcal{S} \circ \mathcal{E}$ amplifies quantum differential privacy, provided that T is γ -Dobrushin.

Theorem 5. Let \mathcal{E} be a γ -Dobrushin quantum operation and let \mathcal{S} be a $(\tau, \varepsilon(\tau), 0)$ -QDP quantum operation, where $\varepsilon : \mathbb{R}^+ \rightarrow [0, 1]$. Then $\mathcal{S} \circ \mathcal{E}$ is $(\tau, \varepsilon(\gamma\tau), 0)$ -QDP.

Proof. Let σ and ρ two states such that $\|\rho - \sigma\|_{\text{tr}} \leq \tau$. By definition of γ -Dobrushin, the trace distance between $\mathcal{E}(\rho)$ and $\mathcal{E}(\sigma)$ can be upper bounded as follows.

$$\|\mathcal{E}(\rho) - \mathcal{E}(\sigma)\|_{\text{tr}} \leq \gamma \|\rho - \sigma\|_{\text{tr}} \leq \gamma\tau.$$

The channel \mathcal{S} is $(\gamma\tau, \varepsilon(\gamma\tau), 0)$ -QDP, thus, for any measurement M ,

$$\Pr[M(\mathcal{S}(\mathcal{E}(\rho))) \in F] \leq \exp(\varepsilon(\gamma\tau)) \Pr[M(\mathcal{S}(\mathcal{E}(\sigma))) \in F].$$

The inequality above shows that $\mathcal{S} \circ \mathcal{E}$ is $(\tau, \varepsilon(\gamma\tau), 0)$ -QDP. □

To better demonstrate the application of this result, we give some explicit examples for different quantum channels. First, we recall the definitions of some relevant quantum operations [37]. The *generalized amplitude damping* channel for a single qubit is defined as

$$\mathcal{E}_{GAD}(\rho) = \sum_{k=0}^3 E_k \rho E_k^\dagger$$

in the 2-dimensional Hilbert space \mathcal{H}_2 , where

$$E_0 = \sqrt{p} \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{bmatrix}, \quad E_1 = \sqrt{p} \begin{bmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{bmatrix}, \\ E_2 = \sqrt{1-p} \begin{bmatrix} \sqrt{1-\gamma} & 0 \\ 0 & 1 \end{bmatrix}, \quad E_3 = \sqrt{1-p} \begin{bmatrix} 0 & 0 \\ \sqrt{\gamma} & 0 \end{bmatrix}.$$

and p and γ are two parameters. The *phase damping* channel for a single qubit is defined by the operator-sum representation

$$\mathcal{E}_{PD}(\rho) = E_0\rho E_0^\dagger + E_1\rho E_1^\dagger$$

in the 2-dimensional Hilbert space \mathcal{H}_2 , where

$$E_0 = \begin{bmatrix} \sqrt{1} & 0 \\ 0 & \sqrt{\lambda} \end{bmatrix}, \quad E_1 = \begin{bmatrix} 0 & 0 \\ 0 & \sqrt{\lambda} \end{bmatrix}.$$

Thus we can compose the channels above to define the *phase-amplitude damping* channel

$$\mathcal{E}_{PAD}(\rho) = \mathcal{E}_{GAD}(\mathcal{E}_{PD}(\rho)).$$

As shown in [22], the phase-amplitude damping channel is $(d, \varepsilon, 0)$ -QDP, with

$$\varepsilon := \ln \left(1 + \frac{2d\sqrt{1-\gamma}\sqrt{1-\lambda}}{1-\sqrt{1-\gamma}\sqrt{1-\lambda}} \right) \quad (9)$$

The *depolarizing* channel corresponds to the quantum operation

$$\mathcal{E}_{Dep} = p \frac{\mathbb{I}}{D} + (1-p)\rho$$

where D is the dimension of the state Hilbert space and p is the probability parameter. As shown in [22], the depolarizing channel is $(d, \varepsilon, 0)$ -QDP, with

$$\varepsilon := \ln \left(1 + \frac{1-p}{p} dD \right).$$

The following result characterizes a family of Dobrushin quantum operation.

Lemma 4 ([25], Theorem 6.1). *Let Φ_T be a quantum operation such that*

1. $\Phi_T(\mathbb{I}) = \mathbb{I}$ (Φ_T is unital),
2. $\Phi_T : \mathbb{I} + w \cdot \sigma \rightarrow \mathbb{I} + (Tw) \cdot \sigma$ where T is a real matrix with $\|T\|_\infty \leq 1$, where $\|\cdot\|_\infty$ is the operator norm.

Then Φ_T is $\|T\|_\infty$ -Dobrushin.

Interestingly, the Dobrushin condition and unitality ensure quantum *pure* differential privacy, as we show in the following theorem.

Theorem 6. *Let Φ be a quantum operation in the 2-dimensional Hilbert space \mathcal{H}_2 , such that*

1. $\Phi(\mathbb{I}) = \mathbb{I}$ (Φ is unital),
2. Φ is γ -Dobrushin.

Then Φ is $(d, \log(1 + 2d\gamma), 0)$ -QDP.

Table 2: The channels below satisfies $(d, \varepsilon, 0)$ -QDP with the ε values shown in the table. We denoted as Φ_γ an arbitrary γ -Dobrushin unital channel.

CHANNEL	ε	REFERENCE
\mathcal{E}_{Dep}	$\ln\left(1 + \frac{1-p}{p}dD\right)$	[22]
\mathcal{E}_{PAD}	$\ln\left(1 + \frac{2d\sqrt{1-\gamma}\sqrt{1-\lambda}}{1-\sqrt{1-\gamma}\sqrt{1-\lambda}}\right)$	[22]
Φ_γ	$\ln(1 + 2d\gamma)$	THEOREM 6
$\mathcal{E}_{PAD} \circ \mathcal{E}_{Dep}$	$(1-p)\ln\left(1 + \frac{2d\sqrt{1-\gamma}\sqrt{1-\lambda}}{1-\sqrt{1-\gamma}\sqrt{1-\lambda}}\right)$	THEOREM 7

Proof. Consider two arbitrary qubit states ρ_1, ρ_2 and set $\|\rho_1 - \rho_2\|_{\text{tr}} := d$. The γ -Dobrushin condition can be restated as follows:

$$\|\Phi(\rho_1) - \Phi(\rho_2)\|_{\text{tr}} \leq \gamma\|\rho_1 - \rho_2\|_{\text{tr}} = \gamma d.$$

Given an arbitrary POVM $M = \{M_m\}$, we want to bound the following quantity

$$\frac{\text{Tr}\{\Phi(\rho_1)M_m\}}{\text{Tr}\{\Phi(\rho_2)M_m\}} - 1$$

First, we upper bound the numerator:

$$\text{Tr}\{\Phi(\rho_1)M_m\} - \text{Tr}\{\Phi(\rho_2)M_m\} = \text{Tr}\{(\Phi(\rho_1) - \Phi(\rho_2))M_m\} \leq d\gamma\text{Tr}\{M_m\}.$$

Since ρ_2 is a qubit state, can write $\rho_2 = \frac{1}{2}(\mathbb{I} + r \cdot \sigma)$ for a Bloch vector r .

$$\text{Tr}\{\Phi(\rho_2)M_m\} = \text{Tr}\left\{\frac{1}{2}\Phi(\mathbb{I} + r \cdot \sigma)M_m\right\} \geq \frac{1}{2}\text{Tr}\{\Phi(\mathbb{I})M_m\} = \frac{1}{2}\text{Tr}\{M_m\},$$

where we used unitality in the last inequality. Putting all together, we get

$$\frac{\text{Tr}\{\Phi(\rho_1)M_m\}}{\text{Tr}\{\Phi(\rho_2)M_m\}} - 1 \leq \frac{d\gamma\text{Tr}\{M_m\}}{\frac{1}{2}\text{Tr}\{M_m\}} = 2d\gamma.$$

Thus the channel Φ is $\log(1 + 2d\gamma)$ -QDP.

$$\frac{\text{Tr}\{\Phi(\rho_1)M_m\}}{\text{Tr}\{\Phi(\rho_2)M_m\}} \leq e^\varepsilon,$$

where $\varepsilon := \ln(1 + 2d\gamma)$. □

Finally, we show how [Theorem 5](#) can be used to derive privacy amplification bounds for the composition of several channels.

Theorem 7. *The composition of the depolarizing channel and the phase-amplitude damping channel $\mathcal{E}_{PAD} \circ \mathcal{E}_{Dep}$ is $(d, \varepsilon, 0)$ -QDP, where*

$$\varepsilon = (1-p)\ln\left(1 + \frac{2d\sqrt{1-\gamma}\sqrt{1-\lambda}}{1-\sqrt{1-\gamma}\sqrt{1-\lambda}}\right).$$

Proof. We use the fact that the depolarizing channel \mathcal{E}_{Dep} satisfies the Dobrushin condition. This can be shown either by direct computation or by observing that \mathcal{E}_{Dep} satisfies the hypothesis of [Lemma 4](#) with $\|T\|_\infty = 1-p$. Then we can combine [Eq. \(9\)](#) with [Theorem 5](#) to derive the desired bound for the composed channel $\mathcal{E}_{PAD} \circ \mathcal{E}_{Dep}$. □

6 Discussion and future work

We have undertaken a systematic study of differential privacy amplification in quantum and quantum-inspired algorithms. Our work is the first to reason about quantum encodings through the lens of differential privacy, laying the foundation for further analysis. Prior to this work, the choice of the encoding was motivated mainly by expressiveness, efficiency and robustness to experimental noise [38, 39]. Due to the intimate relation between DP, algorithmic stability and robustness to adversarial examples, our results suggest new criteria for the choice of quantum encodings. Previous work explored the relation between experimental noise, quantum differential privacy and robustness. The tighter bounds presented in our paper can be used to improve the result of [21], taking into account the composition of several mechanisms.

In the future, it would be interesting to provide similar amplification results for the notion of quantum differential privacy employed in [23]. As previously mentioned, in the quantum setting, different definitions of neighbouring quantum states lead to different notions for quantum differential privacy. Thus, understanding the relationship between these notions is of both theoretical and practical interest. To this end, one could also adopt the variety of quantum distances available in the quantum information literature. One of the best candidates for this purpose is the *quantum Wasserstein distance* introduced in [40], which generalises the notion of neighbouring quantum states of [23]. Furthermore, the relation between the contraction coefficient of quantum channels and DP which we have explored in this paper can also be studied alternatively with this distance due to the contractivity results proved in [40].

Another interesting future direction would be to expand our composition results and the previous work of [22] on the differential privacy of specific quantum channels, to more general classes of quantum operations, in terms of their general characteristics such as contraction coefficients or channel capacity. On this note, a good candidate would be to study the effect of the class of LOCC (Local Operations and Classical Communication) operations on both classical and quantum differential privacy. This class is of particular interest due to its relation to entanglement, which is another non-classical and unique property of the quantum world to be studied in the context of differential privacy.

Finally, as a follow up of our theoretical results, we aim to investigate the experimental implementation, their feasibility, and their application using the available NISQ devices. Experimental noise is among the major limitations of current architectures. Yet, demonstrating that such noise provides beneficial properties, such as privacy and robustness, could shape the pathway for new applications.

Acknowledgements. We thank Vincent Cohen-Addad, Alex B. Grilo, Nai-Hui Chia and Brian Coyle for useful discussions.

References

- [1] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Conference on Theory of Cryptography*, TCC'06, page 265–284, Berlin, Heidelberg, 2006. Springer-Verlag. ISBN 3540327312. doi: 10.1007/11681878_14. URL https://doi.org/10.1007/11681878_14.
- [2] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. 9(3–4):211–407, August 2014. ISSN 1551-305X. doi: 10.1561/04000000042. URL <https://doi.org/10.1561/04000000042>.

- [3] Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(29):1069–1109, 2011. URL <http://jmlr.org/papers/v12/chaudhuri11a.html>.
- [4] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, Oct 2016. doi: 10.1145/2976749.2978318. URL <http://dx.doi.org/10.1145/2976749.2978318>.
- [5] Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data, 2017.
- [6] Raef Bassily, Om Thakkar, and Abhradeep Thakurta. Model-agnostic private learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, page 7102–7112, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [7] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM J. Comput.*, 40(3):793–826, June 2011. ISSN 0097-5397. doi: 10.1137/090756090. URL <https://doi.org/10.1137/090756090>.
- [8] Yu-Xiang Wang, Jing Lei, and Stephen E. Fienberg. Learning with differential privacy: Stability, learnability and the sufficiency and necessity of erm principle. *Journal of Machine Learning Research*, 17(183):1–40, 2016. URL <http://jmlr.org/papers/v17/15-313.html>.
- [9] M. Bun, R. Livni, and S. Moran. An equivalence between private classification and online prediction. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 389–402, Los Alamitos, CA, USA, nov 2020. IEEE Computer Society. doi: 10.1109/FOCS46700.2020.00044. URL <https://doi.ieeecomputersociety.org/10.1109/FOCS46700.2020.00044>.
- [10] Srinivasan Arunachalam, Yihui Quek, and John Smolin. Private learning implies quantum stability. In *Advances in Neural Information Processing Systems 34 pre-proceedings (NeurIPS 2021)*, NIPS’21, 2021.
- [11] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS’07)*, pages 94–103, 2007. doi: 10.1109/FOCS.2007.66.
- [12] Cynthia Dwork, Guy N. Rothblum, and Salil Vadhan. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 51–60, 2010. doi: 10.1109/FOCS.2010.12.
- [13] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, 2017. doi: 10.1109/TIT.2017.2685505.
- [14] Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 61–70, 2010. doi: 10.1109/FOCS.2010.85.

- [15] Moritz Hardt, Katrina Ligett, and Frank McSherry. A simple and practical algorithm for differentially private data release. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'12, page 2339–2347, Red Hook, NY, USA, 2012. Curran Associates Inc.
- [16] Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, page 6280–6290, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [17] Vitaly Feldman, Ilya Mironov, Kunal Talwar, and Abhradeep Thakurta. Privacy amplification by iteration. *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, Oct 2018. doi: 10.1109/focs.2018.00056. URL <http://dx.doi.org/10.1109/FOCS.2018.00056>.
- [18] Borja Balle, Gilles Barthe, Marco Gaboardi, and Joseph Geumlek. *Privacy Amplification by Mixing and Diffusion Mechanisms*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [19] Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. *Lecture Notes in Computer Science*, page 375–403, 2019. ISSN 1611-3349. doi: 10.1007/978-3-030-17653-2_13. URL http://dx.doi.org/10.1007/978-3-030-17653-2_13.
- [20] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018. doi: 10.22331/q-2018-08-06-79. URL <https://quantum-journal.org/papers/q-2018-08-06-79/>. Publisher: Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften.
- [21] Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, Dacheng Tao, and Nana Liu. Quantum noise protects quantum classifiers against adversaries. *Physical Review Research*, 3(2), May 2021. ISSN 2643-1564. doi: 10.1103/physrevresearch.3.023153. URL <http://dx.doi.org/10.1103/PhysRevResearch.3.023153>.
- [22] Li Zhou and Mingsheng Ying. Differential privacy in quantum computation. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 249–262, 2017. doi: 10.1109/CSF.2017.23.
- [23] Scott Aaronson and Guy N. Rothblum. Gentle measurement of quantum states and differential privacy. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 322–333, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367059. doi: 10.1145/3313276.3316378. URL <https://doi.org/10.1145/3313276.3316378>.
- [24] Srinivasan Arunachalam, Alex B. Grilo, and Henry Yuen. Quantum statistical query learning, 2020. URL <https://arxiv.org/abs/2002.08240>.
- [25] Fumio Hiai and Mary Beth Ruskai. Contraction coefficients for noisy quantum channels. *Journal of Mathematical Physics*, 57(1):015211, Jan 2016. ISSN 1089-7658. doi: 10.1063/1.4936215. URL <http://dx.doi.org/10.1063/1.4936215>.
- [26] Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy, 2019.

- [27] Maria Schuld. Supervised quantum machine learning models are kernel methods, 2021. URL <https://arxiv.org/abs/2101.11020>.
- [28] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 217–228, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450367059. doi: 10.1145/3313276.3316310. URL <https://doi.org/10.1145/3313276.3316310>.
- [29] Ewin Tang. Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions. *Physical Review Letters*, 127(6), Aug 2021. ISSN 1079-7114. doi: 10.1103/physrevlett.127.060503. URL <http://dx.doi.org/10.1103/PhysRevLett.127.060503>.
- [30] András Gilyén, Seth Lloyd, and Ewin Tang. Quantum-inspired low-rank stochastic regression with logarithmic dependence on the dimension, 2018.
- [31] Nai-Hui Chia, Han-Hsuan Lin, and Chunhao Wang. Quantum-inspired sub-linear classical algorithms for solving low-rank linear systems, 2018. URL <https://arxiv.org/abs/1811.04852>.
- [32] Jonathan Ullman. Cs7880: Rigorous approaches to data privacy, 2017. URL <https://www.ccs.neu.edu/home/jullman/cs7880s17/HW1sol.pdf>.
- [33] W. Doeblin. Sur les propriétés asymptotiques de mouvements régis par certains types de chaînes simples (suite et fin). *Bulletin mathématique de la Société Roumaine des Sciences*, 39(2):3–61, 1937. ISSN 12203858. URL <http://www.jstor.org/stable/43769812>.
- [34] R. L. Dobrushin. Central limit theorem for nonstationary markov chains. ii. *Theory of Probability & Its Applications*, 1(4):329–383, 1956. doi: 10.1137/1101029. URL <https://doi.org/10.1137/1101029>.
- [35] Stéphane Gaubert and Zheng Qu. Dobrushin ergodicity coefficient for Markov operators on cones. *Integral Equations and Operator Theory*, 1(81):127–150, January 2015. doi: 10.1007/s00020-014-2193-2. URL <https://hal.inria.fr/hal-01099179>. Also arXiv:1307.4649.
- [36] Michael M. Wolf. Quantum channels and operations. guided tour, July 2012. URL <https://www-m5.ma.tum.de/foswiki/pub/M5/Allgemeines/MichaelWolf/QChannelLecture.pdf>.
- [37] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010. doi: 10.1017/CBO9780511976667.
- [38] Ryan LaRose and Brian Coyle. Robust data encodings for quantum classifiers. *Phys. Rev. A*, 102:032420, Sep 2020. doi: 10.1103/PhysRevA.102.032420. URL <https://link.aps.org/doi/10.1103/PhysRevA.102.032420>.
- [39] Maria Schuld, Ryan Sweke, and Johannes Jakob Meyer. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Physical Review A*, 103(3), Mar 2021. ISSN 2469-9934. doi: 10.1103/physreva.103.032430. URL <http://dx.doi.org/10.1103/PhysRevA.103.032430>.

- [40] Giacomo De Palma, Milad Marvian, Dario Trevisan, and Seth Lloyd. The quantum wasserstein distance of order 1. *IEEE Transactions on Information Theory*, 67(10):6627–6643, Oct 2021. ISSN 1557-9654. doi: 10.1109/tit.2021.3076442. URL <http://dx.doi.org/10.1109/TIT.2021.3076442>.