Lecture notes:

Characterization, certification, and validation of quantum systems

Martin Kliesch

Tex document compiled on April 30, 2020

Quantum simulations and quantum computing are among the most exciting applications of quantum mechanics. More generally, in the quantum technology research field one aims to develop new devices using quantum superposition and entanglement. In a popular wording, these anticipated developments will lead to the second quantum revolution.

A main milestone is the use of quantum capabilities to solve a (computational) problem that cannot practically be solved otherwise. Theoretical proposals include integer factoring (Shor's algorithm), speed-ups for optimization and machine learning algorithms, the simulation of complex quantum systems, and certain sampling experiments specifically tailored to that milestone.

But if one cannot obtain the output of a quantum simulation or computation by conventional means how can one make sure that the outcome is correct? The output of integer factorization can efficiently be checked but, for instance, for the estimation of energies in quantum many-body systems, or outcomes of dynamical simulations, the situation is much less clear. Hence, for the development of trusted quantum technologies special characterization and verification techniques are urgently required.

This course gives an introduction to the research field, to the problems of characterization, validation, and verification, and first ways to solve them. More specifically, quantum state tomography, quantum states certification, quantum process tomography, and randomized benchmarking will be covered. In particular, the course provides an overview of the latest developments in this still young and very active research field. The approaches of the course are mainly of conceptual and mathematical nature.

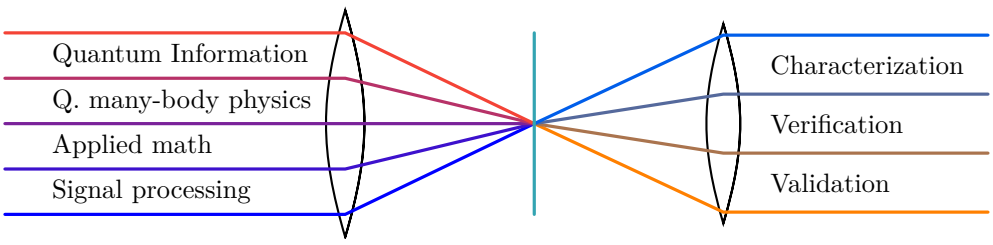

Contents

1. Introduction

Everything not explicitly covered in the lecture is written in gray and sometimes indicated as "auxiliary information".

1.1. Motivation

A central ultimate goal in quantum science is the development of a *scalable universal quantum computer* allowing to run quantum algorithms. However, real quantum systems are difficult to precisely control and subject to unavoidable noise. The noise prevents the direct implementation of such algorithms. Fortunately, if the noise is below a certain threshold then quantum error correction can be used in order to still be able to successfully run quantum algorithms. This desired property is called *fault tolerance*.

But still, for universal quantum computation, one needs to reduce the noise below the fault tolerance threshold and, simultaneously, implement thousands of qubits. These two seemingly conflicting requirements are expected to prevent the development of universal quantum computers in the near future.

So for the time being, many researchers are following more modest goals. As a prove of principle one would like to demonstrate the following[1].

Milestone (quantum supremacy [2])

By using quantum recourses, solve some (computational) problem that cannot practically be solved otherwise.

Here, "practically" means that a corresponding classical computation would have an infeasible long runtime. There are several proposal [3–7] to solve a presumably useless sampling problem in order to demonstrate quantum supremacy.

There is already a caveat included in the definition of quantum supremacy.

The certification problem

If the solved problem cannot be solved otherwise, how can one make sure that the outcome is actually correct? And if one has such a device that can —in

[1] After some criticism *some* researchers have stopped using the term "quantum supremacy". But due to its well-established technical meaning we stick to it.

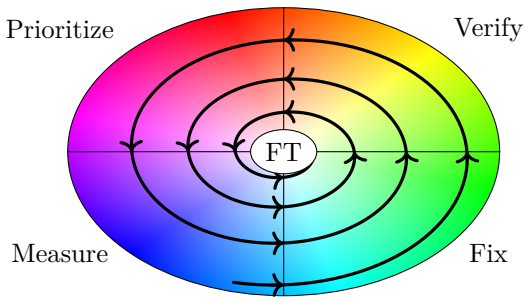

Figure 1.1.: The Smolin Spiral: In order to achieve fault tolerance (FT) "we must iteratively improve devices by estimating sources of smaller and smaller errors, prioritize them, measure them accurately, fix them, and verify they've been fixed" [1].

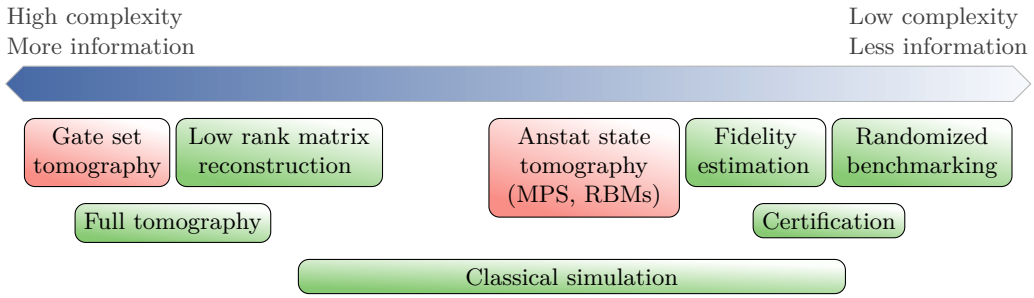

Figure 1.2.: There is a range of characterization and validation techniques. The more complex they are, in terms of measurement and computational effort, the more information they tend to provide. The ones with and without full stable performance guarantees are marked in green and red, respectively.

principle— demonstrate quantum supremacy but one cannot fully check that it functions correctly can one then rightfully claim the achievement of quantum supremacy?

We will introduce tools that can potentially resolve this problem. In particular, for the case of trusted measurements *quantum state certification* can be used.

There are also practically relevant problems where it is expected that full universal quantum computation is not required in order to achieve an advantage over classical computing. In particular, when a quantum circuit is *shallow* enough so that it can be completed before the built-up of noise reaches a critical level then it can be implemented without error correction. There are already many examples suggesting that the era of such Noisy Intermediate-Scale Quantum (NISQ) technology is starting now [8]. Moreover, companies including D-Wave, Google, IBM, and Intel have already heavily invested into building such devices.

In order to be able to improve such devices (see Figure 1.1) and to be able to fairly compare different implementations it is crucial (i) to fully characterize their components and (ii) to find heuristic noise estimation schemes that can be implemented with as little requirements as possible. This course includes methods specifically targeted at these goals: (i) *quantum state* and *process tomography* and (ii) *randomized benchmarking* for estimating the noise in quantum gate implementations. As more fundamental motivation we ask:

- What are the ultimate limits of measuring quantum systems?
- Where do Hamiltonians come from?

Many of us might have wondered about the point on Hamiltonians already in their first quantum mechanics course, where specific Hamiltonians are mostly only postulated. Quantum field theory introduces several *quantization* methods that give quantum Hamiltonians from a classical understanding of physics. However, quantization is typically mathematically not well-defined. Usually uncontrolled approximations are used to derive effective Hamiltonians most physicists are working with. This type of non-rigorous top-down approach might be unsatisfying in some aspects. Therefore, it is very desirable to have a complementary bottom-up approach:

Operational approach

Find a description that explains all possible observations.

In quantum tomography one does exactly that: to learn the full quantum state or process from measured data.

We outlined certification and tomography to give two different examples. In fact, there is a whole range of characterization and verification methods, see Figure 1.2.

The method of choice depends e.g., on the amount and type of information one aims to obtain, the system size, the type of physical correlation that one should expect, the measurement resources, the computational resources, and the type and strength of the noise.

2. Preliminaries

In this chapter we introduce some notation and preliminary results that will be used later. The reader may start skip this chapter first and can come back to it later when needed.

Mathematical quantum information related background information can be found with all details, e.g. in Watrous' lecture notes [9]. All technical terms that are neither explained nor referenced can be found on Wikipedia.

2.1. Notation and math

We will use the following notation and preliminaries:

- $[n] \coloneqq \{1, 2, \ldots, n\}$ for $n \in \mathbb{N}$.
- *Landau symbols* (a.k.a. *Bachmann-Landau notation* or *big-O-notation*) O and Ω are used for asymptotic upper and lower bounds, respectively. These bounds can be relaxed to only hold up to logarithmic factors, which is denoted by $\tilde{\mathrm{O}}$ and $\tilde{\Omega}$, respectively. Θ and $\tilde{\Theta}$ are used to denote that the corresponding upper and lower bound hold simultaneously. For instance, for f defined by $f(x) = x \ln(x)^2$ the following holds: $f(x) \in \mathrm{O}(x^2)$, $f(x) \in \tilde{\mathrm{O}}(x)$, $f(x) \in \Omega(x)$, and hence $f(x) \in \tilde{\Theta}(x)$.
- Throughout this course we assume all vector spaces to be finite dimensional except when explicitly stated otherwise.
- We use bra-ket notation. In particular, we denote the canonical basis of $\mathbb{C}^d$ by $\{|i\rangle\}_{i\in[d]}$ and tensor products of vectors by $|\psi\rangle\,|\phi\rangle \coloneqq |\psi\rangle \otimes |\phi\rangle$.
- We set $|\mathbf{1}_n\rangle \coloneqq (1, 1, \ldots, 1)^\intercal \in \mathbb{R}^n$ to be the vector of ones and often just write $|\mathbf{1}\rangle$ instead of $|\mathbf{1}_n\rangle$.
- We denote the Euclidean unit sphere in $\mathbb{C}^d$ by $\mathbb{S}^{d-1} \coloneqq \{|\psi\rangle \in \mathbb{C}^d : \langle\psi|\psi\rangle = 1\}$ (we index manifolds by their own dimension and not by the dimension of an ambient space).
- We denote the *Pauli matrices* by

$$\sigma_x \coloneqq \sigma_1 \coloneqq \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y \coloneqq \sigma_2 \coloneqq \begin{pmatrix} 0 & -\mathrm{i} \\ \mathrm{i} & 0 \end{pmatrix}, \quad \sigma_z \coloneqq \sigma_3 \coloneqq \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{2.1}$$

  and $\sigma_0 \coloneqq \mathbb{1}_{2\times 2}$.
- Tensor products of the Pauli matrices $\sigma_{s_1} \otimes \cdots \otimes \sigma_{s_n}$ with $s \in \{0, 1, 2, 3\}^n$ are called *Pauli strings*. They are operator basis that is orthogonal in the Hilbert-Schmidt inner product defined below.
- $\mathrm{L}(V, W)$ denotes the vector space of linear (bounded) operators from a vector space V to a vector space W. We set $\mathrm{L}(V) \coloneqq \mathrm{L}(V, V)$.
- $\mathrm{Herm}(\mathcal{H}) \subset \mathrm{L}(\mathcal{H})$ denotes the subspace self-adjoint operators on a Hilbert space $\mathcal{H}$.

- $\mathrm{Pos}(\mathcal{H}) \subset \mathrm{Herm}(\mathcal{H})$ denote the convex cone of positive semidefinite operators. For $X, Y \in \mathrm{Herm}(\mathcal{H})$ we write $X \succeq Y$ if $X - Y \in \mathrm{Pos}(\mathcal{H})$.

- Let $P, Q \in \mathrm{Pos}(\mathbb{C}^d)$. Then
$$\mathrm{Tr}[PQ] \geq 0\,, \tag{2.2}$$
which can be seen by writing the trace as a sum over an eigenbasis of one of the operators.

- A *probability vector* is a vector $p \in [0,1]^d$ that is normalized, i.e., $\sum_i p_i = 1$.

- $\mathcal{S}(\mathcal{H}) \coloneqq \{\rho \in \mathrm{Pos}(\mathcal{H}) : \ \mathrm{Tr}[\rho] = 1\}$ is the convex set of *density operators*. It coincides with the set of convex combinations of the form $\sum_i p_i \, |\psi_i\rangle\langle\psi_i\,|$, where p is a probability vector.

- An operator $X \in \mathrm{L}(\mathbb{C}^d)$ can be diagonalized with unitaries if and only if X is *normal*, i.e., $[X, X^\dagger] = 0$.

- For a normal operator X, which is diagonalized as $X = UDU^\dagger$, and function $f : \mathbb{C} \to \mathbb{C}$ we define
$$f(X) \coloneqq U f(D) U^\dagger\,, \tag{2.3}$$
where $f(D)$ is the diagonal matrix obtained from D by applying f to all its diagonal elements.

- The *Hilbert-Schmidt inner product* on operators $\mathrm{L}(\mathcal{H}_1, \mathcal{H}_2)$ between Hilbert spaces $\mathcal{H}_1$ and $\mathcal{H}_2$ is defined by $\langle X, Y\rangle \coloneqq \mathrm{Tr}[X^\dagger Y]$.

- For a subspace $S \subset \mathrm{L}(\mathcal{H})$ we denote the *orthogonal complement* w.r.t. the Hilbert-Schmidt inner product by $S^\perp$.

- *(Column) vectorization* is a map $|\cdot\rangle : \mathbb{C}^{n_1 \times n_2} \to \mathbb{C}^{n_1 n_2}$ that stacks the columns of a matrix $A \in \mathbb{C}^{n_1 \times n_2}$ on top of each other. For all matrices A, B C with fitting dimensions holds that
$$|ABC\rangle = C^\intercal \otimes A \, |B\rangle\,, \tag{2.4}$$
where $X \otimes Y \cong (X_{i,j} Y)_{i,j}$ (defined by a block matrix) denotes the *Kronecker product*. This formula is of great use numerically, e.g. to avoid explicity basis expansions. There is a similar notion of *row vectorization* with a similar formula.

- The *unitary group* of a Hilbert space $\mathcal{H}$ is denoted by $\mathrm{U}(\mathcal{H}) \subset \mathrm{L}(\mathcal{H})$ and we set $\mathrm{U}(d) \coloneqq \mathrm{U}(\mathbb{C}^d)$.

- The *singular values decomposition* (SVD) of a matrix $X \in \mathbb{C}^{n_1 \times n_2}$ is given by $X = U\Sigma V^\dagger$, where $U \in \mathrm{U}(n_1)$ and $V \in \mathrm{U}(n_2)$ are unitary and the only non-zero entries of Σ are the positive *singular values* $\Sigma_{i,i}$ for $i \in [\mathrm{rank}(X)]$. The unitaries U and V can be obtained by diagonalizing $XX^\dagger$ and $X^\dagger X$, respectively.

  Sometimes one uses a different form of the SVD $X = \tilde{U}\tilde{\Sigma}\tilde{V}^\dagger$ where $\tilde{\Sigma}$ is the positive $(\mathrm{rank}(X) \times \mathrm{rank}(X))$-matrix with the non-zero singular values of X on the diagonal and $\tilde{U}$ and $\tilde{V}$ are obtained from U and V by removing the appropriate columns.

- For $p \in \mathbb{R}_+$ the *ℓ_p-norm* of a vector $x \in \mathbb{C}^d$ is $\|x\|_{\ell_p} \coloneqq (\sum_i |x_i|^p)^{1/p}$.
  For $p = \infty$ it is $\|x\|_{\ell_\infty} \coloneqq \max_i |x_i|$. These norms satisfy *Hölder's inequality*,
$$|\langle x, y\rangle| \leq \|x\|_{\ell_p} \, \|y\|_{\ell_q} \quad \forall x, y \in \mathbb{C}^d \tag{2.5}$$
whenever $1/p + 1/q = 1$ (with $1/\infty \coloneqq 0$). Moreover, the norm inequalities
$$\|x\|_{\ell_q} \leq \|x\|_{\ell_p} \leq s^{1/p - 1/q} \, \|x\|_{\ell_q} \tag{2.6}$$

are satisfied for any $0 < p < q \leq \infty$ and $x \in \mathbb{C}^d$ with s non-zero elements; see e.g. [10, Appendix A.1].

- For $p \in (0, \infty]$ the Schatten p-norm $\|X\|_p$ of an operators X given by the ℓ_p-norm of the vector of its singular values. A similar Hölder's inequality and similar norm inequalities as for the ℓ_p-norms hold for the Schatten p-norms (with s replaced by the rank). In particular, for any operator X,

$$\|X\|_\infty \leq \|X\|_2 \leq \|X\|_1 \leq \sqrt{\operatorname{rank}(X)}\, \|X\|_2 \leq \operatorname{rank}(X)\, \|X\|_\infty \tag{2.7}$$

  and any operator Y from the same space as X,

$$|\operatorname{Tr}[X^\dagger Y]| \leq \|X\|_1 \, \|Y\|_\infty \,. \tag{2.8}$$

- Of particular importance are
  - the Schatten ∞-norm (coincides with the *spectral norm* $\|\cdot\|_{\mathrm{op}}$ a.k.a. *operator norm*),
  - the Schatten 2-norm (coincides with the Frobenius norm $\|\cdot\|_{\mathrm{F}}$, which is induced by the Hilbert-Schmidt inner product and which is invariant under vectorizations maps);
  - and the Schatten 1-norm (coincides with the *trace norm* $\|\cdot\|_1$ a.k.a. *nuclear norm*).

- The *pseudo-inverse* or *Moore–Penrose inverse* is a generalization of the matrix inverse to all matrices. It can be conveniently calculated as follows. Let is write a singular value decomposition of a matrix $X \in \mathbb{C}^{n_1 \times n_2}$ as

$$X = U \Sigma V^\dagger \,. \tag{2.9}$$

  Then the pseudo-inverse Σ^+ of Σ can be obtained component-wise by setting

$$\Sigma^+_{i,j} \coloneqq \begin{cases} 0 & \text{if } \Sigma_{i,j} = 0 \\ \frac{1}{\Sigma_{i,j}} & \text{otherwise} \end{cases} \tag{2.10}$$

  for all i, j. Then the pseudo-inverse of X is

$$X^+ \coloneqq U \Sigma^+ V^\dagger \,. \tag{2.11}$$

- For subsets $S, T \subset V$ of a vector space V we be denote by $S + T \coloneqq \{s + t : \ s \in S, \ t \in T\}$ the *Minkowski sum* of S and T. The sets $-T$ and $S - T$ are define similarly. For $v \in V$ we set $v + S \coloneqq \{v + s : \ s \in S\}$. By $\operatorname{cone}(S) \coloneqq \{\lambda s : \ \lambda \in \mathbb{R}_+, \ s \in S\}$ we denote the *cone generated by* S. The *linear span* $\operatorname{span}_K(S) \subset$ is the linear subspace of all linear combinations of vectors in S and coefficients in K. The subscript K is usually omitted when K is the underlying field of V. The *dimension* of S is $\dim(S) \coloneqq \dim(\operatorname{span}(S))$.

- A subset $C \subset \mathbb{R}^n$ is called *convex* if $(1 - \lambda)x + \lambda y \in C$ for all $x, y \in C$ and $\lambda \in [0, 1]$. Similarly, a subset $V \subset \mathbb{R}^n$ is called *affine set* or *affine subspace* if $(1 - \lambda)x + \lambda y \in V$ for all $x, y \in V$ and $\lambda \in \mathbb{R}$. The set $V_0 \coloneqq V - V \subset \mathbb{R}^n$ is the unique *linear subspace parallel to* V and for any $x_0 \in V$ one has $V = x_0 + V_0$; see Rockafellar's book on convex analysis [11, Section 1] for more details.

- For a series of events $A_1, A_2, \ldots$ the *union bound* (a.k.a. Boole's inequality) hods

$$\mathbb{P}[A_1 \text{ OR } A_2 \text{ OR} \ldots] \leq \sum_i \mathbb{P}[A_i] \,. \tag{2.12}$$

- A *(standard) Gaussian random variable* (RV) is a normally distributed RV with zero mean and unit variance. A *complex (standard) Gaussian RV* is a RV of

which the real and imaginary part are each iid. standard Gaussian RVs. A *complex (standard) Gaussian (random) vector* is a vector in $\mathbb{C}^d$ the components of which are iid. complex standard Gaussian RVs. A *Hermitian (standard) Gaussian (random) matrix* is a matrix in $\mathrm{Herm}(CC^d)$, of which the diagonal components are iid. Guassian RVs and the ohter components iid. complex Gaussian RVs.

- There is a unitarily invariant probability measure on the unitary group $\mathrm{U}(\mathcal{H})$ of a Hilbert space $\mathcal{H}$ called the *Haar measure*. For instance, unitary matrices diagonalizing Hermitian Gaussian matrices are distributed Haar randomly. This fact can be used to numerically sample unitaries from the Haar measure.

2.2. Representation theory

Let us start with the most basic definitions. For a proper introduction we refer to Simon's book [12] and to Goodman and Wallach's book [13] for the representation theory of the standard matrix groups.

Definition 2.1 (Basic definitions from representation theory):

Let G and H be groups.

- $f : G \to H$ is a *(group) homomorphism* if $f(g_1 g_2) = f(g_1) f(g_2)$ for all $g_1, g_2 \in G$. Note that this condition implies that $f(e_G) = e_H$ and $f\big(g^{-1}\big) = f(g)^{-1}$ for all $g \in G$.
- Let $\mathcal{H} \cong \mathbb{C}^\nu$ be a Hilbert space with the unitary group $\mathrm{U}(\mathcal{H}) \subset \mathrm{L}(\mathcal{H})$. A *unitary representation* of G on $\mathcal{H}$ is a homomorphism $R : G \to \mathrm{U}(\mathcal{H})$. Such representations are instances of *linear group representations*, which are defined similarly with $\mathrm{U}(\mathcal{H})$ replaced by the group of invertible operators on $\mathcal{H}$. We will only be concerned with unitary representations and, hence, often omit the word "unitary".
- A subspace $V \subset \mathcal{H}$ is said to be *invariant* if $R(g)V \subseteq V$ for all $g \in G$.
- R is called *irreducible* if the only invariant subspaces are $\{0\}$ and $\mathcal{H}$ itself. *Irreducible representations* are also called irreps.
- Two representations $R : G \to \mathrm{U}(\mathcal{H})$ and $\tilde{R} : G \to \mathrm{U}(\tilde{\mathcal{H}})$ are said to be *unitarily equivalent* if there is a unitary operator $W : \mathcal{H} \to \tilde{\mathcal{H}}$ such that $\tilde{R}(g) = W R(g) W^\dagger$ for all $g \in G$.

If $R_i : G \to \mathcal{H}_i$ for $i = 1, 2$ are two representations of G then $(R_1 \oplus R_2)(g) \coloneqq R_1(g) \oplus R_2(g)$ defines another representation $R_1 \oplus R_2 : G \to \mathcal{H}_1 \oplus \mathcal{H}_2$. This representation has $\mathcal{H}_1$ and $\mathcal{H}_2$ as invariant subspaces. Conversely, if a representation R has a non-trivial invariant subspace V then it can be decomposed as $R = R|_V \oplus R_{V^\perp}$. By iterating this insight, we have the following statement (see e.g. [12, Theorem II.2.3]).

Proposition 2.2 (Decomposition into irreps):

Let $R : G \to \mathrm{L}(\mathcal{H})$ be a unitary representation of a group G on a finite-dimensional Hilbert space $\mathcal{H}$. Then $(R, \mathcal{H})$ can be decomposed into a direct sum of irreps $(R_i, \mathcal{H}_i)$ of G as

$$\mathcal{H} = \bigoplus_i \mathcal{H}_i \quad \text{and} \quad R(g) = \bigoplus_i R_i(g)\,. \tag{2.13}$$

Several irreps $R_{i_1}, \ldots, R_{i_m}$ might be unitarily equivalent to each other. The maximum such m is called the *multiplicity* of that irrep. The space $\mathbb{C}^m$ in the resulting

identification

$$\bigoplus_{j=1}^{m} R_{i_j}(g) \cong R_{i_1}(g) \otimes \mathbb{1}_m \quad \in \mathrm{L}(\mathcal{H}_1 \otimes \mathbb{C}^m)\,. \tag{2.14}$$

is called the *multiplicity space* of R_{i_1}. The decomposition (2.13) is called *multiplicity-free* if all irreps R_i are inequivalent, i.e., not isomorphic.

Theorem 2.3 (Schur's lemma):

Let $R : G \to \mathrm{U}(\mathcal{H})$ be an irrep of G on $\mathcal{H}$. If $A \in \mathrm{L}(\mathcal{H})$ satisfies

$$AR(g) = R(g)A \qquad \forall g \in G \tag{2.15}$$

then $A = c\,\mathbb{1}$ for some $c \in \mathbb{C}$.

Proof. The condition (2.15) implies that $R(h)A^\dagger = A^\dagger R(h)$ for all $h = g^{-1} \in G$. Hence, this condition also holds for $\mathrm{Re}(A) := \frac{1}{2}(A + A^\dagger)$ and $\mathrm{Im}(A) := \frac{1}{2\mathrm{i}}(A - A^\dagger)$ and A is a constant if they both are. Hence, it is sufficient to prove the theorem for $A \in \mathrm{Herm}(\mathcal{H})$.

Let $|\psi\rangle$ be an eigenvector with $A|\psi\rangle = \lambda|\psi\rangle$ and $\mathrm{Eig}_\lambda(A) := \{|\psi\rangle : A|\psi\rangle = \lambda\psi\}$ the full eigenspace. Then $R(g)|\psi\rangle \in \mathrm{Eig}_\lambda(A)$ for all $g \in G$ because $AR(g)|\psi\rangle = R(g)A|\psi\rangle = \lambda R(g)|\psi\rangle$. So, $\mathrm{Eig}_\lambda(A)$ is an invariant subspace. Since $\mathrm{Eig}_\lambda(A) \neq \{0\}$ and R is an irrep, $\mathrm{Eig}_\lambda(A) = \mathcal{H}$ follows.

□

Corollary 2.4 (Irreps of abelian groups):

If G is abelian then every irrep has dimension 1.

Proof. Let R be an irrep of G on $\mathcal{H}$. Theorem 2.3 implies that each $g \in G$ has representation $R(g) = c_g\mathbb{1}$ for some constant c_g. Hence, every subspace of $\mathcal{H}$ is invariant under R. Since R is an irrep this is only possible if $\dim(\mathcal{H}) = 1$. □

There is also a slightly more general version of Schur's lemma:

Theorem 2.5 (Schur's lemma II):

Let $R : G \to \mathrm{U}(\mathcal{H})$ and $\tilde{R} : G \to \mathrm{U}(\tilde{\mathcal{H}})$ be two irreps of G on finite dimensional Hilbert spaces $\mathcal{H}$ and $\tilde{\mathcal{H}}$. If $A \in \mathrm{L}(\mathcal{H}, \tilde{H})$ satisfies

$$AR(g) = \tilde{R}(g)A \qquad \forall g \in G \tag{2.16}$$

then either $A = 0$ or R_1 and R_2 are unitarily equivalent up to a constant factor.

Proof. The condition 2.16 implies that for all $g \in G$

$$R(g)A^\dagger = A^\dagger\tilde{R}(g) \tag{2.17}$$

and, hence,

$$R(g)A^\dagger A = A^\dagger AR(g) \tag{2.18}$$

$$\tilde{R}(g)AA^\dagger = AA^\dagger\tilde{R}(g)\,. \tag{2.19}$$

Schur's lemma (Theorem 2.3) implies that $A^\dagger A = c\,\mathbb{1}$ and $AA^\dagger = \tilde{c}\,\mathbb{1}$ for constants $c, \tilde{c}$. Obviously, $c = \tilde{c}$, as can be seen from the eigenvalues. Either $c = 0$ so that $A = 0$ or $W = A/\sqrt{c}$ is a unitary. In the latter case

$$WR(g) = \tilde{R}(g)W \tag{2.20}$$

for all $g \in G$, i.e., R and $\tilde{R}$ are unitarily equivalent. □

A unitary W relating two representations R and $\tilde{R}$ as in (2.20) is called an *intertwining* unitary of R and $\tilde{R}$.

Particularly interesting in the context of k-designs are two group representations on $\mathcal{H} = (\mathbb{C}^d)^{\otimes k}$. Remember that the *symmetric group* Sym_k is the group of permutations on k elements.

Two group representations on $(\mathbb{C}^d)^{\otimes k}$

The symmetric group Sym_k and the unitary group $\mathrm{U}(d)$ both have a canonical representation on $(\mathbb{C}^d)^{\otimes k}$:

$$\pi_k : \mathrm{Sym}_k \to \mathrm{U}\big((\mathbb{C}^d)^{\otimes k}\big)\,, \tag{2.21}$$
$$\Delta_d : \mathrm{U}(d) \to \mathrm{U}\big((\mathbb{C}^d)^{\otimes k}\big)\,. \tag{2.22}$$

For $\sigma \in \mathrm{Sym}_k$ and $U \in \mathrm{U}(d)$ they are given by

$$\pi_k(\sigma)(|\psi_1\rangle \otimes \cdots \otimes |\psi_k\rangle) \coloneqq |\psi_{\sigma^{-1}(1)}\rangle \otimes \cdots \otimes |\psi_{\sigma^{-1}(k)}\rangle \tag{2.23}$$
$$\Delta_d(U)(|\psi_1\rangle \otimes \cdots \otimes |\psi_k\rangle) \coloneqq (U|\psi_1\rangle) \otimes \cdots \otimes (U|\psi_k\rangle)\,. \tag{2.24}$$

We note that $\pi_k(\sigma)$ and $\Delta_d(U)$ commute for any $\sigma \in \mathrm{Sym}_k$ and $U \in \mathrm{U}(d)$.

The *symmetric* and *anti-symmetric subspace* of $(\mathbb{C}^d)^{\otimes k}$ are defined to be

$$\mathcal{H}_{\mathrm{sym}^k} \coloneqq \left\{|\Psi\rangle \in (\mathbb{C}^d)^{\otimes k} : \ \pi_k(\sigma)|\Psi\rangle = |\Psi\rangle \ \ \forall \sigma \in \mathrm{Sym}_k\right\}, \tag{2.25}$$
$$\mathcal{H}_{\wedge^k} \coloneqq \left\{|\Psi\rangle \in (\mathbb{C}^d)^{\otimes k} : \ \pi_k(\sigma)|\Psi\rangle = \mathrm{sign}(\sigma)|\Psi\rangle \ \ \forall \sigma \in \mathrm{Sym}_k\right\}. \tag{2.26}$$

By P_{sym^k} and $P_{\wedge^k}$ we denote the orthogonal projectors onto these two subspaces, respectively.

Let us briefly consider the case $k = 2$. It is easy to see that any matrix can be decomposed into a symmetric and an anti-symmetric part, which are orthogonal to each other. This implies that

$$(\mathbb{C}^d)^{\otimes 2} = \mathcal{H}_{\mathrm{sym}^2} \oplus \mathcal{H}_{\wedge^2}\,. \tag{2.27}$$

Note that due to Corollary 2.4, both the symmetric and the antisymmetric subspace are isomorphic to some $\mathbb{C}^m$, where m_{sym^2} and $m_{\wedge^2}$ are the multiplicities of the two different one-dimensional irreps of Sym_2.

For large k there is a similar decomposition with more summands called *Schur-Weyl* decomposition. In ordert to state it we write $\lambda \vdash k$ for an integer partition of the form

$$k = \sum_{i=1}^{l(\lambda)} \lambda_i \tag{2.28}$$

of k into integers $\lambda_i \geq 1$.

Theorem 2.6 (Schur-Weyl duality):

The action of $\mathrm{Sym}_k \times \mathrm{U}(d)$ on $(\mathbb{C}^d)^{\otimes k}$ given by the commuting representations (2.23) and (2.24) is multiplicity-free and $(\mathbb{C}^d)^{\otimes k}$ decomposes into irreducible components as

$$(\mathbb{C}^d)^{\otimes k} \cong \bigoplus_{\lambda \vdash k, l(\lambda) \leq d} W_\lambda \otimes S_\lambda\,. \tag{2.29}$$

For any $k \geq 2$, both $\mathcal{H}_{\mathrm{sym}^k}$ and $\mathcal{H}_{\wedge^k}$ occur as component in the direct sum (2.29).

The spaces W_λ are called *Weyl modules* and S_λ *Specht modules*. Schur-Weyl duality implies that the Weyl modules are the multiplicity spaces of the irreps of Sym_k and, similarly, the Specht modules are the multiplicity spaces of the irreps $\mathrm{U}(d)$.

The last ingredient we need in order to prove Lemma 3.12 is the dimension of the symmetric subspace.

Exercise 2.1 (Symmetric subspace):

- Calculate $P_{\mathrm{sym}^k} |\psi\rangle$ for a product state $|\psi\rangle = |\psi_1\rangle \otimes \cdots \otimes |\psi_k\rangle$.
- Show that the dimension of the symmetric subspace $P_{\mathrm{sym}^k}(\mathbb{C}^d)^{\otimes k}$ is

$$\mathrm{Tr}[P_{\mathrm{sym}^k}] = \binom{d+k-1}{k} . \tag{2.30}$$

Hint: Argue first that this is the number of ways to distribute k indistinguishable particles (bosons) into d boxes (modes).

Often it is helpful to write P_{sym^2} in terms of the *flip operator* (a.k.a. *swap operator*) $\mathbb{F} \in \mathrm{L}(\mathcal{H}^{\otimes 2})$, which is defined by

$$\mathbb{F} |\psi\rangle |\phi\rangle \coloneqq |\phi\rangle |\psi\rangle \tag{2.31}$$

and can hence be written as

$$\mathbb{F} = \sum_{i,j=1}^{d} |i,j\rangle\langle j,i| . \tag{2.32}$$

Then

$$P_{\mathrm{sym}^2} = \frac{1}{2}(\mathbb{1} + \mathbb{F}) . \tag{2.33}$$

From (2.32) one can see that $\mathrm{Tr}[\mathbb{F}] = d$, so that indeed $\mathrm{Tr}[P_{\mathrm{sym}^2}] = \frac{1}{2}d(d+1)$.

Proposition 2.7 (Invariant operators for $k = 2$):

Let $E \in \mathrm{L}(\mathbb{C}^d \otimes \mathbb{C}^d)$ be an operator such that

$$(U \otimes U)E = E(U \otimes U) \tag{2.34}$$

for all $U \in \mathrm{U}(d)$. Then

$$E = c_1 P_{\mathrm{sym}^2} + c_2 P_{\wedge^2} \tag{2.35}$$

for some constants $c_1, c_2 \in \mathbb{C}$ depending on E.

Proof. Let us denote the representation of $\mathrm{U}(d)$ by $\Delta : \mathrm{U}(d) \to \mathrm{L}(\mathbb{C}^d \otimes \mathbb{C}^d)$. Schur-Weyl duality (Theorem 2.6) for $k = 2$ tells us that irreps of Δ as

$$(\mathbb{C}^d)^{\otimes 2} \cong W_{\mathrm{sym}^2} \otimes S_{\mathrm{sym}^2} \oplus W_{\wedge^2} \otimes S_{\wedge^2} , \tag{2.36}$$

where S_{sym^2} and $S_{\wedge^2}$ carry the multiplicities of the irreps of Δ. But Sym_2 is abelian, so according to Corollary 2.4 $\dim(S_{\mathrm{sym}^2}) = \dim(S_{\wedge^2}) = 1$, i.e., the irreps W_{sym^2} and $W_{\wedge^2}$ of Δ are multiplicity-free.

Now we can write E as a block matrix corresponding to the decomposition $(\mathbb{C}^d)^{\otimes 2} \cong W_{\mathrm{sym}^2} \oplus W_{\wedge^2}$,

$$E =: \begin{pmatrix} E_{1,1} & E_{1,2} \\ E_{2,1} & E_{2,2} \end{pmatrix} . \tag{2.37}$$

As similar decomposition for Δ is $\Delta = \Delta_{\mathrm{sym}} \oplus \Delta_\wedge$ (the off-diagonal blocks are zero).

The invariance $\Delta(U)E = E\Delta(U)$ implies

$$\Delta_{\mathrm{sym}}(U)E_{1,1} = E_{1,1}\Delta_{\mathrm{sym}}(U) \tag{2.38}$$
$$\Delta_{\wedge}(U)E_{2,2} = E_{2,2}\Delta_{\wedge}(U) \tag{2.39}$$

for all $U \in \mathrm{U}(d)$. Hence, thanks to Schur's lemma (Theorem 2.3), $E_{i,i} = c_i \mathbb{1}$ for $i = 1, 2$ and some constants c_1 and c_2. Similarly, we obtain

$$\Delta_{\mathrm{sym}}(U)E_{1,2} = E_{1,2}\Delta_{\wedge}(U) \tag{2.40}$$
$$\Delta_{\wedge}(U)E_{2,1} = E_{2,1}\Delta_{\mathrm{sym}}(U) \tag{2.41}$$

for all $U \in \mathrm{U}(d)$. According to Schur's lemma (the second version, Theorem 2.5) $E_{1,2}$ and $E_{2,1}$ are each either zero or an intertwining unitary for Δ_{sym} and $\Delta_{\wedge}$ just as W in (2.20). Since Δ_{sym} and $\Delta_{\wedge}$ are irreps they cannot be intertwining unitaries and must hence be zero. Together,

$$E = \begin{pmatrix} c_1 \mathbb{1} & 0 \\ 0 & c_2 \mathbb{1} \end{pmatrix} = c_1 P_{\mathrm{sym}^2} + c_2 P_{\wedge^2} \,. \tag{2.42}$$

□

2.3. Random variables and tail bounds

Tail bounds for random variables are bounds to the probability that a random variable assumes a value that deviates from the expectation value, as visualized by the marked area in Figure 2.1. Indeed, for any random variable X it is unlikely to assume values that are much larger than the expectation value $\mathbb{E}[|X|]$:

Theorem 2.8 (Markov's inequality):

Let X be a non-negative random variable and $a > t$. Then

$$\mathbb{P}[X \geq t] \leq \frac{\mathbb{E}[X]}{t} \,. \tag{2.43}$$

Proof. The indicator function $\mathbf{1}_A$ of a set A is defined by

$$\mathbf{1}_A(x) \coloneqq \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{otherwise.} \end{cases} \tag{2.44}$$

To prove Markov's inequality we choose A to be the set $\{X \geq t\} \coloneqq \{\omega : \ X(\omega) \geq t\}$, observe that

$$t\, \mathbf{1}_{\{X \geq t\}} \leq X \tag{2.45}$$

and take the expectation value of both sides of this inequality. □

The variance can give tighter bounds on the tail than just the expectation values $\mathbb{E}[|X|]$:

Theorem 2.9 (Chebyshev's inequality):

Let X be a mean zero random variable with finite variance $\sigma^2 \coloneqq \mathbb{E}[X^2]$. Then

$$\mathbb{P}[|X| \geq t] \leq \frac{\sigma^2}{t^2} \tag{2.46}$$

for all $t \geq 0$.

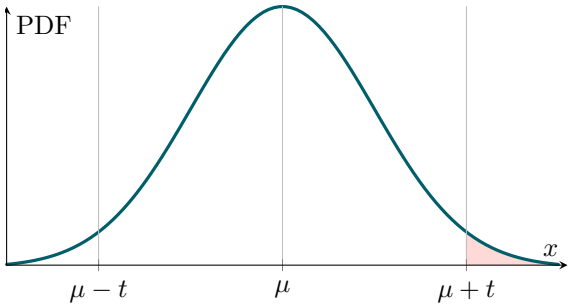

Figure 2.1.: The (upper) *tail* of a random variable X is the probability of X being greater than some threshold t. This probability is given by the corresponding area under the graph of the probability density function (PDF) of X.

Proof. Apply Markov's inequality to the random variable X^2. □

Note that in the case of a random variable Y that does not necessarily have a zero mean Chebyshev's inequality yields a tail bound by applying it to $X \coloneqq Y - \mathbb{E}[Y]$; see also Figure 2.1. The same argument can be made for the tails bounds that follow.

When a random variable is bounded then its empirical mean concentrates much more than a naive application of Markov's inequality suggests. More generally, the following holds (see, e.g., [10, Theorem 7.20]):

Theorem 2.10 (Höffdings inequality):

Let $X_1, \dots, X_n$ be independent random variables with $a_i \leq X_i \leq b_i$ almost surely for all $i \in [n]$ and $S_n \coloneqq \sum_{i=1}^n X_i$. Then

$$\mathbb{P}[S_n - \mathbb{E}[S_n] \geq t] \leq \exp\left(-\frac{2\,t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right), \tag{2.47}$$

$$\mathbb{P}[|S_n - \mathbb{E}[S_n]| \geq t] \leq 2\exp\left(-\frac{2\,t^2}{\sum_{i=1}^n (b_i - a_i)^2}\right) \tag{2.48}$$

for all $t \geq 0$.

Proof. The second statement directly follows from the first one. In order to prove the first one, let $s > 0$, apply Markov's inequality to

$$\mathbb{P}[S_n - \mathbb{E}[S_n] \geq t] = \mathbb{P}[\mathrm{e}^{s(S_n - \mathbb{E}[S_n])} \geq \mathrm{e}^{s\,t}], \tag{2.49}$$

use the independence of the X_i to factorize the exponential, use the bounds on X_i, and optimize over s. □

When one can control the variance of the random variables then the following tail bound can give a better concentration, especially for small values of t.

Theorem 2.11 (Bernstein inequality [10, Corollary 7.31]):

Let $X_1, \dots, X_n$ be independent zero mean random variables with $|X_i| \leq a$ almost surely and $\mathbb{E}[|X_i|^2] \leq \sigma_i^2$ for some $a > 0$, $\sigma_i > 0$, and $i \in [n]$. Set $S_n \coloneqq \sum_{i=1}^n X_i$. Then, for all $t > 0$,

$$\mathbb{P}\big[\big|S_n - \mathbb{E}[S_n]\big| \geq t\big] \leq 2\exp\left(-\frac{t^2/2}{\sigma^2 + at/3}\right) \tag{2.50}$$

with $\sigma^2 \coloneqq \sum_{i=1}^n \sigma_i^2$.

Another related tail bound is *Azuma's inequality*, which allows for a relaxation on the independence assumption (super-martingales with bounded differences).

2.4. Monte Carlo integration

Traditionally, Monte Carlo methods are used in numerical integration, optimization, and sampling from a distribution. Here, we will introduce *importance sampling*, a method to estimate a sum or integral.

Specifically, we aim to compute an integral F that is written as an expectation value

$$F \coloneqq \mathbb{E}_{X\sim p}[f(X)] = \int f(x)p(x)\,\mathrm{d}x \tag{2.51}$$

of some function f. In this setting, we can iid. sample $X^{(1)}, \dots, X^{(m)} \sim p$ and take the empirical average

$$\hat{F} \coloneqq \frac{1}{m}\sum_{i=1}^{m} f(X^{(i)}) \tag{2.52}$$

as estimator of F. It is not difficult to see that $\hat{F}$ is *unbiased*, i.e., that $\mathbb{E}[\hat{F}] = F$. If $\mathrm{Var}[f(X)] < \infty$ then $\hat{F}$ can be proven to be *consistent*, i.e., $\hat{F}$ converges to F for $m \to \infty$ in an appropriate sense. Moreover,

$$\mathrm{Var}[\hat{F}] = \frac{\mathrm{Var}[f(X)]}{m}\,, \tag{2.53}$$

i.e., the empirical variance also gives an estimate of the estimation error. Thanks to the *central limit theorem* $\hat{F}$ converges to a normal random variable with mean F and variance $\frac{\mathrm{Var}[f(X)]}{m}$, which allows for the simple estimation of confidence intervals. However, everything relies on the ability to sample from p. Two popular methods to make such sampling efficient are *importance sampling* and *Markov chain Monte Carlo* sampling.

2.4.1. Importance sampling

The idea is to rewrite the integrand $f\,p$ in (2.51) as

$$fp = \frac{fp}{q}q \tag{2.54}$$

for some probability distribution q. Then we can apply the Monte Carlo sampling idea (2.52) w.r.t. q, i.e., we draw $X^{(1)}, \dots, X^{(m)} \sim q$ to obtain the estimator

$$\hat{F}_q \coloneqq \frac{1}{m}\sum_{i=1}^{m} f(X^{(i)})\frac{p(X^{(i)})}{q(X^{(i)})}\,. \tag{2.55}$$

It holds that $\mathbb{E}_q[\hat{F}_q] = F$ and

$$\mathrm{Var}_q[\hat{F}_q] = \frac{1}{m}\,\mathrm{Var}_q[fp/q] = \frac{1}{m}\int \frac{f^2p^2}{q}\,. \tag{2.56}$$

It can be shown that choosing q as

$$q^* \coloneqq \frac{p\,|f|}{Z} \tag{2.57}$$

with normalization factor Z yields minimum variance. In particular, for $f \geq 0$ we even have $\mathbb{E}_{q^*}[(fp/q^*)^2] = \mathbb{E}_p[f]^2 = \mathbb{E}_{q^*}[fp/q^*]^2$, i.e., $\mathrm{Var}_{q^*}[\hat{F}_{q^*}] = 0$. So, if f does not change its sign then a single sample is sufficient for the estimation! But in order

to obtain Z one needs to solve the integration problem first. However, finding non-optimal but good choices for q can already speed up the integration.

2.5. Basic convex optimization problems

Convex optimization problems are minimization problems where a convex *objective function* is optimized over a convex set, the *feasible region*. They have a rich duality theory [11] that leads to polynomial-time algorithms for many important subclasses of convex optimization problems [14]. Two important subclasses are *linear programs* and *semidefinite programs*. They can be solved efficiently with standard software, e.g., by using CVX(PY) [15, 16].

2.5.1. Linear programs (LPs)

LPs are convex optimizatio problems with linear objective function, linear equality, and linear inequality constraints. These are exactly those problems where a linear function is optimized over a *convex polytope*. A *convex polytope* is defined as a set that is an intersection of finitely many half spaces.

As every convex optimization problem, an LP comes along with a *primal* and *dual* formulation. Let us denote entrywise non-negativity of a vector $x \in \mathbb{R}^n$ by $x \geq 0$.

Definition 2.12 (LPs):

A *linear program* (LP) is given by a tripel (A, c, b) with $c \in \mathbb{R}^n$, $b \in \mathbb{R}^m$, and $A \in \mathbb{R}^{m \times n}$ and comes along with the following pair of optimization problems:

Primal:	*Dual:*
maximize $c^\intercal x$	minimize $b^\intercal y$
subject to $Ax = b\,,$	subject to $A^\intercal y \geq c\,,$
$x \geq 0\,.$	$y \geq 0\,.$

The variables x and y are called *primal* and *dual* variable respectively. A *primal feasible point* is a point x that satisfies the constraints of the primal LP. A *primal optimal point* is a primal feasible point $x^\sharp$ so that $c^\intercal x^\sharp$ is the outcome of the maximization in the primal LP. *Dual feasible points* and *dual optimal points* $y^\sharp$ are defined similarly via the dual LP.

Weak duality states that for an primal feasible point x and dual feasible point y we have $c^\intercal x \leq b^\intercal y$, which directly follows from the definition of the primal and dual problem. The *strong duality theorem* for LPs states that $c^\intercal x^\sharp = b^\intercal y^\sharp$ for all optimal primal and dual points $x^\sharp$ and $y^\sharp$.

An LP for the ℓ_1-norm

The ℓ_1 norm of a vector $z \in \mathbb{R}^n$ can be written as

$$\begin{aligned} \|z\|_{\ell_1} = \min \quad & \mathbf{1}^\intercal y \\ \text{subject to} \quad & y_i \geq z_i \text{ and } y_i \geq -z_i \ \forall i, \\ & y \geq 0\,, \end{aligned} \tag{2.58}$$

where $\mathbf{1}$ is the vector with all components being 1. We note that the ℓ_1-norm of complex vectors is not an LP but a so-called *second-order cone program*, which is a special kind of SDP.

2.5.2. Semidefinite programs (SDPs)

SDPs are generalization of an LPs. Here, vectors are replaced by matrices and the entrywise conic order is replaced by the semidefinite conic ordering on matrices. To

be more explicity, we folllow [17, Chapter 1.2.3] and definine SDPs for either real of complex inner product spaces V and W. A linear operator $\Theta : \mathrm{L}(V) \to \mathrm{L}(W)$ is said to be *Hermiticity-preserving* if $\Theta(X) \in \mathrm{Herm}(W)$ for all $X \in \mathrm{Herm}(V)$.

Definition 2.13 (SDPs):

A *semidefinite program* is specified by a triple (Ξ, C, B), where $C \in \mathrm{Herm}(V)$ and $B \in \mathrm{Herm}(W)$ are self-adjoint operators and $\Xi : \mathrm{L}(V) \to \mathrm{L}(W)$ is a Hermiticity-preserving linear map. With such a triple, we associate the following pair of optimization problems:

Primal:	*Dual:*
maximize $\mathrm{Tr}(CX)$	minimize $\mathrm{Tr}(BY)$
subject to $\Xi(X) = B\,,$	subject to $\Xi^\dagger(Y) \succeq C\,,$
$X \succeq 0\,.$	$Y \in \mathrm{Herm}(W)\,.$

Primal and dual feasible and optimal points are defined similarly as for LPs.

SDPs that characterized as in this definition are said to be in *standard form*. For specific SDPs, equivalent formulations might often be more handy.

Weak duality refers to the fact that the value of the primal SDP cannot be larger than the value of the dual SDP, i.e., that $\mathrm{Tr}(CX) \leq \mathrm{Tr}(DY)$ for any primal feasible point X and dual feasible point Y. This fact follows directly from the definition of the primal and dual problem:

$$\mathrm{Tr}[BY] = \mathrm{Tr}[\Xi(X)Y] = \mathrm{Tr}[\Xi^\dagger(Y)X] \geq \mathrm{Tr}[CX]\,. \tag{2.59}$$

An SDP is said to satisfy *strong duality* if the optimal values coincide, i.e., if for some optimal primal feasible and dual feasible point $X^\sharp$ and $Y^\sharp$ it holds that $\mathrm{Tr}(CX^\sharp) = \mathrm{Tr}(DY^\sharp)$. In fact, from a weak condition, called *Slater's condition*, strong duality follows. Slaters condition is that either one of the two following conditions is fulfilled:

(i) The primal problem is bounded above and there is a strictly feasible point of the dual problem, i.e., there is $Y \in \mathrm{Herm}(W)$ with $\Xi^\dagger(Y) \succ 0$.

(ii) The dual problem is bounded below and there is a strictly feasible point of the primal problem, i.e., there is $X \succ 0$ with $\Xi(X) = B$.

There are efficient solvers for SDPs such as CVX(PY) [15, 16]. The underlying numerical routines (interior point methods) come along with convergence proofs. This means that one can view SDPs as functions in a similar sense as the sine function is a function: both have a rigorous analytic characterization and polynomial time algorithms for their evaluation with convergence guarantees.

An SDP for the spectral norm

For any matrix Z and $y \geq 0$ holds that

$$\|Z\|_{\mathrm{op}} \leq y \quad \Leftrightarrow \quad y^2\mathbb{1} - Z^\dagger Z \geq 0 \quad \Leftrightarrow \quad \begin{pmatrix} y\mathbb{1} & Z \\ Z^\dagger & y\mathbb{1} \end{pmatrix} \succeq 0\,, \tag{2.60}$$

where the last equivalence can be checked, e.g., by using an SVD of Z and reducing it to a 2×2 matrix problem. It follows that

$$\|Z\|_{\mathrm{op}} = \min\left\{ y \in \mathbb{R} : \begin{pmatrix} y\mathbb{1} & Z \\ Z^\dagger & y\mathbb{1} \end{pmatrix} \succeq 0 \right\} \tag{2.61}$$

and dualization yields

$$\begin{aligned} \|Z\|_{\mathrm{op}} = \max \quad & \mathrm{Tr}(Z\,X_2) + \mathrm{Tr}(Z^\dagger X_2^\dagger) \\ \text{subject to} \quad & \mathrm{Tr}(X_1) + \mathrm{Tr}(X_3) = 1\,, \\ & \begin{pmatrix} X_1 & X_2 \\ X_2^\dagger & X_3 \end{pmatrix} \succeq 0\,. \end{aligned} \tag{2.62}$$

An SDP for the trace norm

An SDP that yields the trace norm of a matrix Z is given by the dual formulation

$$\|Z\|_1 = \min_Y \left\{ \mathrm{Tr}(Y_1)/2 + \mathrm{Tr}(Y_2)/2 : \begin{pmatrix} Y_1 & -Z \\ -Z^\dagger & Y_2 \end{pmatrix} \succeq 0 \right\}. \tag{2.63}$$

The corresponding primal formulation is

$$\|Z\|_1 = \max_X \left\{ \mathrm{Tr}(Z\,X)/2 + \mathrm{Tr}(Z^\dagger X^\dagger)/2 : \begin{pmatrix} \mathbb{1} & X \\ X^\dagger & \mathbb{1} \end{pmatrix} \succeq 0 \right\}. \tag{2.64}$$

Since the spectral norm is dual to the spectral norm the equivalence (2.60) implies that the optimal value of the primal problem for the trace norm (2.64) is indeed $\|Z\|_1$.

2.6. Quantum mechanics

Quantum theory is build on basic postulates that formalize the mathematical model of nature. A general version of the *static* postulates (dynamics will be covered in Chapter II), so the ones concerning states and measurements, can be stated as follows.

Postulate (quantum states and measurements, general form):

- Every quantum system is associated with a separable complex Hilbert space $\mathcal{H}$.
- A *measurement* is given by a *positive operator valued measure* (POVM), which in turn is given by a set of positive semidefinite operators (*effects*) $\{E_i\}$ (discrete version only) on $\mathcal{H}$ such that $\sum_i E_i = \mathbb{1}$. (Hence $\|E_i\|_{\mathrm{op}} \leq 1$ for all i.)
- A *quantum state* described by a density operator, i.e., a positive[a] normalized element of the dual space of $(\mathrm{L}(\mathcal{H}), \|\cdot\|_{\mathrm{op}})$.

[a] A functional ρ on $\mathrm{L}(\mathcal{H})$ is called *positive* if $\rho(A^\dagger A) \geq 0$ for all $A \in \mathrm{L}(\mathcal{H})$.

The *outcomes* of a POVM measurement are given by the labels i of the effects E_i.

Exercise 2.2 (Measurement postulate):

1. Show that for finite dimensional Hilbert spaces the dual space of $(\mathrm{L}(\mathcal{H}), \|\cdot\|_{\mathrm{op}})$ is indeed given by $(\mathrm{L}(\mathcal{H}), \|\cdot\|_1)$.
2. Show that $\rho \in \mathrm{Herm}(\mathbb{C}^d)$ is positive semidefinite if and only if $\mathrm{Tr}[\rho A^\dagger A] \geq 0$ for all $A \in \mathrm{L}(\mathbb{C}^d)$. Also note that $\mathrm{Tr}[\rho] = \|\rho\|_1$ for all $\rho \in \mathrm{Pos}(\mathcal{H})$.
3. Explain how the measurement of an observable, i.e., *von Neumann measurements* (a.k.a. *projective measurements*) are related to POVM measurements.
4. That quantum states are dual to observables already tells us that the trace norm is the canonical norm for quantum states. Find an operational interpretation of the trace norm related to measurements.

A density matrix $\rho \in \mathcal{S}(\mathbb{C}^d)$ is called *pure* if there is a state vector $|\psi\rangle \in \mathbb{C}^d$ such that $\rho = |\psi\rangle\langle\psi|$.

Exercise 2.3 (Pure state condition):

Show that a state $\rho \in \mathcal{S}(\mathbb{C}^d)$ is pure if and only if $\mathrm{Tr}[\rho^2] = 1$.

Given two single quantum systems, their joint system should also be a quantum system. This expectation is captured by the following.

Postulate (composite quantum systems):

The Hilbert space of two quantum systems with Hilbert spaces $\mathcal{H}_1$ and $\mathcal{H}_2$, respectively, is the tensor product $\mathcal{H}_1 \otimes \mathcal{H}_2$.

This construction induces an embedding from $\mathrm{L}(\mathcal{H}_1)$ into $\mathrm{L}(\mathcal{H}_1 \otimes \mathcal{H}_2)$ by

$$A \mapsto A \otimes \mathbb{1}\,. \tag{2.65}$$

Dually to that, for any state $\rho \in \mathcal{S}(\mathcal{H}_1 \otimes \mathcal{H}_2)$

$$\mathrm{Tr}[\rho\,(A \otimes \mathbb{1})] = \mathrm{Tr}[\rho_1 A]\,, \tag{2.66}$$

where ρ_1 is the that ρ *reduced* to system 1. The reduced state captures all information of ρ that can be obtained from measuring system 1 alone and can be explicitly obtained by the *partial trace* over the second subsystem

$$\begin{aligned} \mathrm{Tr}_2 : \mathrm{L}(\mathcal{H}_1 \otimes \mathcal{H}_2) &\to \mathrm{L}(\mathcal{H}_1) \qquad \text{(linear)} \\ X \otimes Y &\mapsto \mathrm{Tr}_2[X \otimes Y] \coloneqq X\,\mathrm{Tr}[Y] \end{aligned} \tag{2.67}$$

as $\rho_1 \coloneqq \mathrm{Tr}_2[\rho]$.

Exercise 2.4 (The swap-trick):

Let $\mathbb{F} \in \mathrm{L}(\mathcal{H} \otimes \mathcal{H})$ be the *flip operator* (or *swap operator*), i.e., the linear extension of the map $|\psi\rangle\,|\phi\rangle \mapsto |\phi\rangle\,|\psi\rangle$. Show that

$$\mathrm{Tr}_1[(X \otimes \mathbb{1})\mathbb{F}] = X \tag{2.68}$$

for any $X \in \mathrm{L}(\mathcal{H})$.

Part I

Quantum states

Quantum state tomography is the task to (re)construct a full description of a given quantum state from measurement data. Due to the nature of quantum physics one cannot measure a quantum state from just one single copy but many copies of the same state are required for that task. As the quantum measurements are of probabilistic nature one can also only hope to be able to reconstruct the state up to some finite precision. Of course, the precision depends on the precise reconstruction method, the number of measurements, and the type of measurements at hand. Also the computational efficiency of the actual reconstruction algorithm is crucial. For instance, in one of the largest state tomography implementations [18] an 8-qubit state was reconstructed, which required two weeks post-processing.

This illustrates the need of tomographic schemes that are user friendly and work with an optimal amount of resources at the same time.

3. Quantum state tomography (QST)

Let us start with a common method for QST, which is often used in theory when then details concerning the number of measurements and accuracy do not matter.

A naive QST method

Let us assume that we are given n_ρ copies of an a priori unknown state $\rho \in \mathcal{S}(\mathbb{C}^d)$. Then we can fix a basis of Hilbert-Schmidt orthonormal observables $(M_i)_{i=0,\dots,d^2-1}$ with $M_0 = \mathbb{1}/\sqrt{d}$. Then the other observables are traceless and we can expand ρ in that basis as

$$\rho = \frac{\mathbb{1}}{d} + \sum_{i=1}^{d^2-1} \mathrm{Tr}[\rho M_i] M_i \,. \tag{3.1}$$

The expectation values $\mathrm{Tr}[\rho M_i]$ can, at least in principle, be approximated from experimental data.

This method has several drawbacks:

- $d^2 - 1$ many *measurement settings* are required, which is much more than necessary and infeasible for intermediate-sized quantum systems.
- The errors on the estimates of $\mathrm{Tr}[\rho M_i]$ tend to add up in an unfavorable way.
- Moreover, they will typically result in an estimate of ρ that is not a positive semi-definite operator.

In the following sections, we will address the drawbacks of the naive QST scheme and introduce methods that (partially) resolve them. In order to be able to fairly compare different strategies we will use the following tomography framework.

Sequential QST for iid. state preparations

In *sequential* QST a number n_ρ of copies of a state $\rho \in \mathcal{P}$ are given with $\mathcal{P} \subset \mathcal{S}(\mathbb{C}^d)$ capturing *prior knowledge* on the state. The i-th copy of ρ is measured with a POVM $\mathsf{M}^{(i)}$ with m_i *measurement outcomes*. Based on measurement values and a description of $(\mathsf{M}^{(i)})_{i\in n_\rho}$ an estimate (*reconstruction*) $\hat{\rho}$ of ρ is generated by a *reconstruction algorithm*. The error term is $\hat{\rho} - \rho$ and $\|\hat{\rho} - \rho\|_p$ is the *reconstruction error (in p-norm)*. The combination of measurement data acquisition and reconstruction algorithm is also called *tomography scheme*. Any tomography scheme is probabilistic resulting in some *failure probability*, i.e., probability of the reconstruction being successful, which can be captured by $\|\hat{\rho} - \rho\|_p \leq \epsilon$ for some predetermined $\epsilon > 0$.

Usually one does not use a new POVM for each copy of ρ so that many of the POVMs $\mathsf{M}^{(i)}$ are the same.

The randomness in any tomographic scheme comes at least from the measurements being probabilistic. However, in general, also the measurement setup and potentially even the reconstruction algorithm can contain probabilistic elements, each of which can result in some failure probability of the measurement scheme.

In the end, there will always be a trade-off between *resources*:

- the number of *samples* n_ρ,
- the effort of implementing the *measurement setup* $(\mathsf{M}^{(i)})_{i\in n_\rho}$,
- prior knowledge $\mathcal{P}$, and
- the computational cost of the reconstruction algorithm

and *performance*:

- $\|\hat{\rho} - \rho\|_p$ and
- the failure probability.

There are two paths of how our tomography setting could be generalized. The generalization of sequential measurements are *parallel measurements* where $\rho^{\otimes n_\rho}$ is measured with one large joint POVM. This concept implicitly includes the possibility of processing $\rho^{\otimes n_\rho}$ in a quantum computer with unbounded circuit complexity (see [19, 20] for QST including quantum computations). However, in order to keep things practical and NISQ-area oriented we focus on *sequential* measurements where the $(i+1)$-th copy of ρ is measured after the i-th copy.

Another path of generalization could be to relax iid. assumption, i.e., the assumption that the measured total state is of the form $\rho^{\otimes n_\rho}$. But to keep our considerations simple we will keep the very convenient iid. assumption. When the total state is ρ_{tot} instead of $\rho^{\otimes n_\rho}$ the *iid. mismatch error* $\|\rho_{\text{tot}} - \rho^{\otimes n_\rho}\|_1$ would add to the reconstruction error in the worst case.

3.1. Informational completeness of measurements

In this section we put aside issues concerning sampling errors, noise and accuracy. In such a simplified setting we address the following questions.

- How many expectation values are needed for QST?
- How many measurement settings (in terms of POVM measurements) are required?
- What is the relation between the number of outcomes in POVM measurements and expectation values of observables?
- How many measurements are needed for pure states?

The real dimension of the space spanned by $\mathcal{S}(\mathbb{C}^d)$ is $\dim_{\mathbb{R}}(\mathcal{S}(\mathbb{C}^d)) = d^2 - 1$. Hence, one would expect that $d^2 - 1$ many operators are necessary and sufficient to measure a state $\rho \in \mathcal{S}(\mathbb{C}^d)$. Largely following Heinosaari, Mazzarella, and Wolf [21], we will

now justify this intuition and extend it to the situation where the measured state is known to be pure.

For a set of operators $\mathsf{M} = \{M_i\}_{i\in[m]} \subset \mathrm{Herm}(\mathcal{H})$ we define the *measurement map* $\mathcal{M} : \mathrm{L}(\mathcal{H}) \to \mathbb{R}^m$ by its components

$$\mathcal{M}(\rho)_i \coloneqq \mathrm{Tr}[\rho M_i]\,. \tag{3.2}$$

component-wise for $i \in [m]$. Notice that if M is a POVM with m outcomes then it is already determined by $m-1$ operators; the last operator M_m is given by $M_m = \mathbb{1} - \sum_{i=1}^{m-1} M_i$. In the following we will stick to the notation of denoting by M a set of operators M_i and by $\mathcal{M}$ the associated measurement map with components (3.2). Generally one says that measurements are *informationally complete* (a.k.a. *tomographically complete*) if they can uniquely identify any state in an idealized setting. More precisely:

Definition 3.1 ($\mathcal{P}$-informational completeness):

Let $\mathcal{P} \subseteq \mathcal{S}(\mathcal{H})$ be a subset of states. A collection of POVMs $\{\mathsf{M}^{(i)}\}$ is called *informationally complete w.r.t.* $\mathcal{P}$ if the concatenated measurement map $\bigoplus_i \mathcal{M}^{(i)}$ is injective on $\mathcal{P}$, i.e.,

$$\mathcal{M}^{(i)}(\rho) = \mathcal{M}^{(i)}(\sigma)\ \forall i \ \Rightarrow\ \rho = \sigma \tag{3.3}$$

for all $\rho, \sigma \in \mathcal{P}$. For the case $\mathcal{P} = \mathcal{S}(\mathcal{H})$ we just call $\{\mathsf{M}^{(i)}\}$ *informationally complete*.

In short, informational completeness means that the "exact outcomes" uniquely determine the state.

As we will see in the next proposition, informational completeness of a set of POVMs can be related to informational completeness of just one POVM. The quantity that matters here is

$$\mathfrak{S}(\mathsf{M}^{(1)}, \ldots, \mathsf{M}^{(n_{\mathsf{M}})}) \coloneqq \mathrm{span}_{\mathbb{C}} \bigcup_{i=1}^{n_{\mathsf{M}}} \mathsf{M}^{(i)}\,, \tag{3.4}$$

which is the set of all possible complex linear combinations of all operators occurring in the POVMs $(\mathsf{M}^{(1)}, \ldots, \mathsf{M}^{(n_{\mathsf{M}})})$. A set of operators $\mathfrak{S}$ is called *operator system* if $\mathbb{1} \in \mathfrak{S}$ and if it is closed under taking adjoints, i.e., $\mathfrak{S} = \mathfrak{S}^\dagger$. Clearly, $\mathfrak{S}(\mathsf{M}^{(1)}, \ldots, \mathsf{M}^{(n_{\mathsf{M}})})$ is an operator system (assuming $n_{\mathsf{M}} \geq 1$).

Proposition 3.2 (POVMs and operator systems [21, Proposition 1]):

Let $\mathfrak{S} \subseteq \mathrm{L}(\mathcal{H})$ be an operator system. Then there exists a POVM M such that $\mathfrak{S} = \mathfrak{S}(\mathsf{M})$ and M has $\dim(\mathfrak{S})$ outcomes. Any POVM B satisfying $\mathfrak{S} = \mathfrak{S}(\mathsf{B})$ has at least $\dim(\mathfrak{S})$ outcomes.

Proof. Exercise. □

It seems that operator systems capture the information contained in POVMs they are generated by. Indeed, this is formalized by [21, Proposition 2]:

Proposition 3.3 (Operator systems and informational completeness):

Let $\mathfrak{S} \subseteq \mathrm{L}(\mathcal{H})$ be an operator system and let $\mathcal{P} \subseteq \mathcal{S}(\mathcal{H})$ be a set of states. Then a POVM M satisfying $\mathfrak{S}(\mathsf{M}) = \mathfrak{S}$ is informationally complete w.r.t. $\mathcal{P}$ iff

$$(\mathcal{P} - \mathcal{P}) \cap \mathfrak{S}^\perp = \{0\}\,. \tag{3.5}$$

Proof. For all states $\rho, \sigma \in \mathcal{P}$, we have:

$$\begin{aligned} \mathcal{M}(\rho) = \mathcal{M}(\sigma) &\Leftrightarrow \mathrm{Tr}[\rho A] = \mathrm{Tr}[\sigma A] \quad \forall A \in \mathfrak{S} \\ &\Leftrightarrow \mathrm{Tr}[(\rho - \sigma) A] = 0 \quad \forall A \in \mathfrak{S} \\ &\Leftrightarrow \rho - \sigma \in \mathfrak{S}^{\perp} . \end{aligned}$$

□

One may only considers measurement data that comes in the form of expectation values of observables rather than from full POVM statistics. We adjust the definition of informational completeness also to that case.

Definition 3.4 ($\mathcal{P}$-informational completeness for observables):

Let $\mathcal{P} \subseteq \mathrm{L}(\mathcal{H})$ be a subset of states. A collection $\mathsf{M} = \{M_i\}_{i \in [m]} \subset \mathrm{Herm}(\mathcal{H})$ of self-adjoint operators is called *informationally complete w.r.t.* $\mathcal{P}$ if its measurement map (3.2) is injective on $\mathcal{P}$, i.e.,

$$\mathcal{M}(\rho) = \mathcal{M}(\sigma) \;\Rightarrow\; \rho = \sigma \tag{3.6}$$

for all $\rho, \sigma \in \mathcal{P}$. For the case $\mathcal{P} = \mathcal{S}(\mathcal{H})$ we just call M *informationally complete.*

Now let us compare the informational completeness of POVM measurements with the one of expectation values of observables. First, the normalization constraint of POVM measurements give rise to the identity operator in the generated operator system, which does not yield any information on a quantum state ($\mathrm{Tr}[\rho \mathbb{1}] = 1 \; \forall \rho \, \mathcal{S}(\mathcal{H})$). Second, the operators of a POVM are positive semidefinite, which observables do not need to be. This positivity constraint, however, does not affect the informational completeness:

Proposition 3.5 (POVMs vs. observables [21, Proposition 3]):

Let $\mathcal{P} \subseteq \mathcal{S}(\mathcal{H})$. The following are equivalent.

(i) There exists a POVM M with m outcomes that is informationally complete w.r.t. $\mathcal{P}$.

(ii) There exists a set $\{A_1, A_2, \ldots, A_{m-1}\} \subset \mathrm{Herm}(\mathcal{H})$ of $m-1$ self-adjoint operators that are informationally complete w.r.t. $\mathcal{P}$.

Proof. Exercise. □

Hence, the discussion of informational completeness can be reduced to observables.

New concepts and notation (short outline): Minkowski sum, cones, affine (sub)spaces, convex sets (see Section 2) and $\mathrm{Herm}(\mathbb{C}^d)$, $\mathrm{Pos}(\mathbb{C}^d)$, unit trace matrices and density matrices as examples.

The measurement map can be analyzed in a topological setting. We adjust [21, Proposition 6] from POVMs to observables:

Proposition 3.6 (Inf. compl. and topological embeddings):

Let $\mathcal{P} \subseteq \mathcal{S}(\mathcal{H})$ be a *closed* subset. A set of self-adjoint operators $\{M_1, \ldots, M_m\} \subset \mathrm{Herm}(\mathcal{H})$ is informationally complete w.r.t. $\mathcal{P}$ iff the measurement map is a topological embedding[a] of $\mathcal{P}$ into $\mathbb{R}^m$.

[a] A *topological embedding* is an injective continuous map with continuous inverse on its domain. It is a good exercise to prove that the inverse of a continuous bijection is automatically continuous if the domain is compact.

Proof. The following proof of [21, Proposition 6] also holds here. Also note that continuous embeddings with compact domains automatically have a continuous inverse (see a post on StackExchange), i.e., are topological embeddings. □

Pure states

Let us consider the special but relevant case of $\mathcal{P}$ the set of pure states

$$\mathcal{P}_d := \{\, |\psi\rangle\langle\psi| : \ |\psi\rangle \in \mathbb{C}^d, \ \||\psi\rangle\|_{\ell_2} = 1 \,\}\,, \tag{3.7}$$

i.e. rank-1 density matrices. Before making statements on general dimension d it is instructive to have a look at qubits, i.e., $d = 2$. Here, the following parametrization is particularly useful.

Bloch sphere

A qubit state $\rho \in \mathcal{S}(\mathbb{C}^2)$ can be uniquely written as

$$\rho = \frac{\mathbb{1}}{2} + \sum_{i=1}^{3} r_i \sigma_i \,, \tag{3.8}$$

where $\|r\|_{\ell_2} \leq 1$ is the *Bloch vector* and $(\sigma_i)_{i\in[3]}$ are the Pauli matrices (2.1). A sate ρ is *pure*, i.e., $\mathrm{rank}(\rho) = 1$, iff it is of the form $\rho = |\psi\rangle\langle\psi|$, or equivalently, iff $\|r\|_{\ell_2} = 1$.

Exercise 3.1:

Prove the claims made on the uniqueness and purity.

The Bloch representation tells us that the set of pure one-qubit states is isomorphic to the 2-dimensional unit sphere in $\mathbb{R}^3$. However, there is no topological embedding of the 2D sphere into $\mathbb{R}^2$. Hence, Proposition 3.6 tells us that two expectation values are not enough for the recovery of states in $\mathcal{P}_2$. On the other hand, $\dim_{\mathbb{R}}(\mathcal{S}(\mathbb{C}^2)) = 3$ with the isomorphism $\mathrm{Herm}(\mathbb{C}^d) = \mathbb{R}^{d^2}$ the normalization condition $\mathrm{Tr}[\rho] = 1 \ \forall \rho \in \mathcal{S}$ implies that 3 expectation values are sufficient to recover any state in $\mathcal{P}_2$. Together, this yields the first entry in the list (3.11).

In general, the minimum number of observables that are informationally complete w.r.t. $\mathcal{P}$ can be bounded as follows.

Theorem 3.7 (Informational completeness for pure states [21]):

Let α denote the number of 1's in the binary expansion of $d-1$. Then

(i) there exists a collection of m self-adjoint operators which is informationally complete w.r.t. pure states in $\mathcal{S}(\mathbb{C}^d)$, if

$$m = \begin{cases} 4d - 3 - \alpha & \text{for odd} \quad d\,, \\ 4d - 4 - \alpha & \text{for even} \quad d \geq 4\,. \end{cases} \tag{3.9}$$

(ii) there exists no such collection if

$$m \leq \begin{cases} 4d - 2\alpha - 4 & \forall\, d > 1\,, \\ 4d - 2\alpha - 2 & d \text{ odd, and } \alpha = 3 \mod 4\,, \\ 4d - 2\alpha - 3 & d \text{ odd, and } \alpha = 2 \mod 4\,. \end{cases} \tag{3.10}$$

This theorem implies that for $d = 4, \ldots, 10$, the minimum number $m_{\min}$ of observables that are informationally complete is known to be as follows — the cases $d = 2, 3$

are covered by [21, Theorems 2]:

d	2	3	4	5	6	7	8	9	10
$m_{\min}$	3	7	9 or 10	15	17 or 18	22 or 23	$23-25$	31	33 or 34

(3.11)

Proof idea of Theorem 3.7. For the proof of the upper bound (3.9) an informationally complete set of observables is constructed based on Milgram's work [22]. The subset of pure states $\mathcal{P} \subset \mathcal{S}(\mathbb{C}^{d+1})$ is diffeomorphic[1] to the complex projective space $\mathbb{C}\mathbf{P}^d$. Milgram [22] constructed immersions[2] of the complex projective space $\mathbb{C}\mathbf{P}^d$ into $\mathbb{R}^{4d-\alpha-c}$, where $c = 1$ for even d and $c = 0$ otherwise. Then an identification of the real vector space of observables as $\mathrm{Herm}(\mathbb{C}^d) \cong \mathbb{R}^{d^2}$ can be used to obtain an informationally complete set.

For the proof of the lower bound (3.10) an extension of the argument leading to Proposition 3.3 is derived that takes the non-trivial topology and curvature of the set of pure state $\mathcal{P} \cong \mathbb{C}\mathbf{P}^d$ into account. First it is observed that M is informationally complete w.r.t. $\mathcal{P}$ iff $\mathcal{M}$ is a *topological embedding* of $\mathcal{P}$ into $\mathbb{R}^m$. Next, measurement maps are considered as smooth topological embeddings from $\mathcal{P} \cong \mathbb{C}\mathbf{P}^d$ into $\mathbb{R}^m$ with a derivative that is injective everywhere. In particular, it is shown that if

(a) the measurement map $\mathcal{M} : \mathcal{P} \to \mathbb{R}^m$ is injective and

(b) for all $P \in \mathcal{P}$ the inclusion $T_P(\mathcal{P}) \subseteq \mathrm{cone}(\mathcal{P} - \mathcal{P})$ of the tangent space $T_P(\mathcal{P})$ at P into the cone generated by differences of pure density matrices holds

then the measurement map $\mathcal{M}$ is a smooth embedding of $\mathcal{P}$ into $\mathbb{R}^m$. Condition (b) is proven to always hold. Then a *non*-embedding result by Mayer [23] is used, which states that for the values of m in the lower bound (3.10) no such embedding and hence no injective measurement map can exist. □

In fact, Heinosaari et al. [21] also provide a similar upper bound for generic (random) measurement maps. For all $m > 2(2d-2)$ it essentially states that observables $M_1, \ldots, M_m \in \mathrm{Herm}(\mathbb{C}^d)$ drawn at random from a continuous non-vanishing distribution are informationally complete with probability one. We will make use of similar randomized strategies in Section 3.6 on compressed sensing.

How many different projective measurements, i.e., measurements of observables are required for quantum state tomography? Theorem 3.7 implies that at least 4 observables are necessary for most dimensions. But in order to make the POVM measurements projective, some overhead might be required. Goyeneche et al. [24] provide a construction of 5 observables that allow for quantum state tomography, which also puts a close-by upper bound on that number.

3.2. Least squares fitting and linear inversion

In the last section we have discussed the number measurement outcomes or expectation values that are necessary to reconstruct a quantum state. Moreover, we have discussed how this number changes when one knows that the state is pure. However, so far we have only discussed injectivity of the measurement map. In this section we discuss how the measurement map can be practically be inverted (without using prior knowledge on the state).

Let us fix the Hilbert space to be $\mathcal{H} = \mathbb{C}^d$. Moreover, let us assume that we are given a measurement map $\mathcal{M}$ with measurement operators $\mathsf{M} = \{M_1, \ldots, M_m\}$ (see (3.2)) that is tomographically complete. Now let us assume that we have estimates y_i of $\mathrm{Tr}[\rho M_i]$ from measurement data, i.e.,

$$y = \mathcal{M}(\rho) + \epsilon\,, \tag{3.12}$$

[1] A *diffeomorphism* is a smooth bijection with smooth inverse.

[2] An *immersion* is a differentiable function between differentiable manifolds whose derivative is everywhere injective. Due to the inverse function theorem, immersions are local embeddings, i.e., locally injective differentiable functions.

where $\epsilon \in \mathbb{R}^m$ is additive noise (e.g, including the unavoidable statistical estimation error). Now we wish to compute a reasonable first estimate $\hat{\rho}$ of ρ from y and $\mathcal{M}$.

Without loss of generality we can assume that $\mathcal{M} : \mathrm{Herm}(\mathbb{C}^d) \to \mathbb{R}^m$ is injective, since, if it is not, we can modify it by extending M by $\mathbb{1}$ and setting the corresponding estimate $y_j := 1$. Note that $\mathcal{M}$ is *not surjective* in general. In particular, whenever $m > d^2$ the image $\mathcal{M}(\mathrm{Herm}(\mathbb{C}^d)) \subset \mathbb{R}^m$ is a zero-set in $\mathbb{R}^m$. Hence, for a continuously distributed sampling error ϵ (e.g. with components from a normal distribution $\mathcal{N}(0, \sigma^2)$) the probability that there exists some $X \in \mathrm{Herm}(\mathbb{C}^d)$ such that $\mathcal{M}(X) = y$ is also zero. In such a situation it is a common and practical strategy to use the *least squares* estimate:

Least squares estimator

The least squares estimator of $\mathcal{M} : \mathrm{Herm}(\mathbb{C}^d) \to \mathbb{R}^m$ and data vector $y \in \mathbb{R}^m$ is

$$\hat{\rho}_{\mathrm{LS}} := \operatorname*{arg\,min}_{X \in \mathrm{Herm}(\mathbb{C}^d)} \|\mathcal{M}(X) - y\|_{\ell_2}^2 . \tag{3.13}$$

Note that the minimum is unique whenever $\mathcal{M}$ is injective.

Later, in Section 3.7, we will apply a certain post-processing to this estimate to obtain an estimate that is close to an optimal estimate in many situations, also w.r.t. trace norm errors. But first, we calculate useful closed form expressions of the least squares estimate (3.13). The basic one is the following, which is called *linear inversion*.

Proposition 3.8 (Linear inversion):

Let $\mathcal{M} : \mathrm{Herm}(\mathbb{C}^d) \to \mathbb{R}^m$ be an in injective linear map. Then the least squares estimate (3.13) is

$$\hat{\rho}_{\mathrm{LS}} = \left(\mathcal{M}^\dagger \mathcal{M}\right)^{-1} \mathcal{M}^\dagger(y) . \tag{3.14}$$

Proof. Exercise. □

The linear map $\left(\mathcal{M}^\dagger \mathcal{M}\right)^{-1} \mathcal{M}^\dagger$ is the *pseudo-inverse* $\mathcal{M}^+$ or *Moore–Penrose inverse* of $\mathcal{M}$ whenever $\mathcal{M}$ is injective (or equivalently, a matrix representation of $\mathcal{M}$ has linearly independent columns, which is equivalent to $\mathcal{M}^\dagger \mathcal{M}$ being invertible). It can be calculated via an SVD, see Eq. (2.11).

Linear inversion on POVM data

Let $\mathsf{M} := \{M_1, \ldots, M_m\}$ be an informationally complete POVM and that copies of a state ρ are measured. Let us assume that there are no noise sources, so that the probability of observing outcome $i \in [m]$ is $p_i := \mathrm{Tr}[\rho M_i]$. By n_i we denote the number of times outcome i is observed out of $n_\rho = \sum_{i=0}^m n_i$ measurements in total. We estimate the probabilities p_i by the *frequency* $\hat{p}_i := n_i / n_\rho$, which is also called the *empirical estimate of* p_i. Thanks to informational completeness, the estimate $\hat{\rho}_{\mathrm{LS}} := \mathcal{M}^+(\hat{p})$ converges to ρ for $n_\rho \to \infty$, which is also implied by the following.

Exercise 3.2 (Sampling error):

Prove that

$$\mathbb{E}\, \|\hat{\rho}_{\mathrm{LS}} - \rho\|_2 \leq \frac{1}{\lambda_{\min}(\mathcal{M})\sqrt{n_\rho}} , \tag{3.15}$$

where $\lambda_{\min}(\mathcal{M})$ is the minimum singular value of $\mathcal{M}$.

The linear inversion estimator (3.14) can be calculated more explicitly for many relevant measurement settings that have additional structure. As we will see, this can also allow to control $\lambda_{\min}(\mathcal{M})$. Two important mathematical tools to capture such additional structure are *frame theory* and *complex projective k-designs*.

In order for a POVM to be informationally complete its elements need to span the full vector space $\mathrm{Herm}(\mathbb{C}^d)$. However, the POVM elements cannot constitute a basis, see Exercise 3.3. Such generating sets of vectors are generally refereed to as frames. There are several important properties that frames can have and that can, e.g., indeed be used to further simplify the linear inversion estimator (3.14).

Exercise 3.3 (No PSD basis):

Prove that there is no basis of positive-semidefinite operators.

3.3. Frame theory

Frame theory is the theory of expanding vectors into a generating system that does not need to be a basis. The development of this theory was very much motivated by applications from signal processing and image processing in particular. For instance, an important class of example is given by the wavelet transform, which is the mathematical tool JPEG compression is based on. Frame theory is naturally also used in quantum state tomography, see e.g., [25, 26] for references in quantum information theory and [27] for the mathematical duality theory.

Let us consider an inner product space $(V, \langle\,\cdot\,|\,\cdot\,\rangle)$. A set of vectors $\{|v_i\rangle\} \subset V$ is called a *frame* for V if there are constants $0 < A \leq B < \infty$ such that

$$A \,\||x\rangle\|^2 \leq \sum_i |\langle v_i|x\rangle|^2 \leq B \,\||x\rangle\|^2 \quad \forall\, |x\rangle \in V\,. \tag{3.16}$$

The constants A and B are called the *lower* and *upper frame bounds*, respectively. They are of practical importance in numerical applications. The self-adjoint operator

$$S \coloneqq \sum_i |v_i\rangle\langle v_i| \tag{3.17}$$

is called *frame operator*. For any $|x\rangle \in V$ we have

$$\langle x|\, S\, |x\rangle = \sum_i |\langle v_i|x\rangle|^2 \tag{3.18}$$

Hence, the bounds (3.16) imply that the eigenvalues of S are contained in the interval $[A, B]$. In particular, S is a positive operator. In fact, the best frame bounds are

$$A = \min(\mathrm{spec}(S))\,, \quad B = \max(\mathrm{spec}(S))\,, \tag{3.19}$$

where $\mathrm{spec}(S) \subset \mathbb{R}$ is the set of eigenvalues of S.

The vectors $\{|\tilde{v}_i\rangle\} \subset V$ given as

$$|\tilde{v}_i\rangle \coloneqq S^{-1}\,|v_i\rangle \tag{3.20}$$

are called *(canonical) dual frame*. It is easy to see that the following frame expansions hold for any $|x\rangle \in V$,

$$|x\rangle = \sum_i \langle v_i|x\rangle\,|\tilde{v}_i\rangle \quad \text{and} \quad |x\rangle = \sum_i \langle \tilde{v}_i|x\rangle\,|v_i\rangle\,. \tag{3.21}$$

Note that the frame operator of the dual frame is

$$\tilde{S} = \sum_i |\tilde{v}_i\rangle\langle\tilde{v}_i| = S^{-1} \sum_i |v_i\rangle\langle v_i| \, (S^{-1})^\dagger = S^{-1} S S^{-1} = S^{-1} \,. \tag{3.22}$$

In particular, the dual of the dual frame is the frame itself. Moreover, since the eigenvalues of S^{-1} are contained in $[1/B, 1/A]$ it follows that the upper and lower frame constants of the dual frame are $1/A$ and $1/B$, respectively.

The frame $\{|v_i\rangle\} \subset V$ is called a *tight frame* for V if upper and lower frame bounds coincide, i.e., if $A = B$. A tight frame is called a *Parseval frame* if it satisfies Parseval's identity $\sum_i |\langle v_i|x\rangle|^2 = \|x\|^2$ for all x, i.e., if it is a tight frame with unit frame constant. Note that the dual frame of a tight (Parseval) frame is also a tight (Parseval) frame with the inverse frame constant. Let us summarize these insights:

Proposition 3.9 (Dual frames):

Let us consider a frame with frame operator S and frame constants $0 < A \leq B$. Then the frame operator of the dual frame is S^{-1} and has frame constants $0 < 1/B \leq 1/A$. In particular, the dual frame of a tight frame is again a tight frame with inverse frame constant.

The *frame potential* of $\{|v_1\rangle, \ldots, |v_n\rangle\}$ is $\mathrm{Tr}[S^2]$. *Normalized* tight frames are the ones that minimize the frame potential:

Theorem 3.10 (See [28, Section I.D] for a discussion and references):

Let $\{|v_1\rangle, \ldots, |v_n\rangle\} \subset \mathbb{S}^{d-1}$ be a frame for $\mathbb{C}^d$ of *unit norm* vectors. Then

$$\mathrm{Tr}[S^2] \geq \frac{n^2}{d} \,. \tag{3.23}$$

Moreover, the bound is saturated iff $\{|v_1\rangle, \ldots, |v_n\rangle\}$ is a tight frame.

Proof. We remember that S is a positive operator. Let $\{\lambda_i\}_{i\in[d]}$ be the eigenvalues of S. Since S is positive we have

$$\mathrm{Tr}[S] = n = \sum_{j=1}^{d} \lambda_j = \|S\|_1 \quad \text{and} \quad \mathrm{Tr}[S^2] = \sum_{j=1}^{d} \lambda_j^2 = \|S\|_2^2 \tag{3.24}$$

and the inequality $\|S\|_1 \leq \sqrt{d}\,\|S\|_2$ establishes the bound. Moreover, this inequality is only saturated when all singular values of S are the same. □

3.4. Complex spherical/projective k-designs

Finite k-designs on a sphere are sets of evenly distributed points. More precisely, a (finite) *spherical k-design* is a set of normalized vectors $\{|\psi_i\rangle\}_{i\in[n]} \subset \mathbb{S}^{d-1}$ such that the average value of certain k-th order polynomials $g(|\psi_i\rangle)$ (equal order k in $\langle i|\psi\rangle$ and $\langle i|\psi\rangle^*$) over the set $\{|\psi_i\rangle$ coincides with the average of $g(|\psi\rangle)$ over all normalized vectors $\mathbb{S}^{d-1}$. Switching to density matrix representation we make the following definition, which can be shown to be equivalent to the polynomial one above [25, 29].

Definition 3.11 (k-design):

We call a probability measure μ on the complex unit sphere $\mathbb{S}^{d-1} \subset \mathbb{C}^d$ a *(complex spherical) k-design* if

$$\int_{\mathbb{S}^{d-1}} (|\psi\rangle\langle\psi|)^{\otimes k} \mathrm{d}\mu(\psi) = \int_{\mathbb{S}^{d-1}} (|\psi\rangle\langle\psi|)^{\otimes k} \mathrm{d}\psi =: K_k \,, \tag{3.25}$$

where $\mathrm{d}\psi$ denotes the uniform $\mathrm{U}(d)$-invariant probability measure on $\mathbb{S}^{d-1}$. The corresponding distribution of $|\psi\rangle\langle\psi|$ is called a *(complex projective k-design)*. A *finite* set of states is called a *k-design* if the uniform distribution over $\{|\psi_i\rangle\langle\psi_i|\}$ is a k-design.

See also Refs. [29, 30] for related definitions.

Exercise 3.4 (k-designs):

Prove that a k-design is also a $k-1$ design for $k \geq 2$.

In order to more conveniently characterize k-designs we calculate K_k more closely. This average is also called *Haar-average* because it can be obtained from the Haar measure, either on $\mathbb{C}\mathbf{P}^{d-1}$ or on $\mathrm{U}(d)$. We summarize statements made by Renes et al. [25] with the following lemma. For this lemma it is helpful to remember that the symmetric subspace of $(\mathbb{C}^d)^{\otimes k}$ is the subspace spanned by vectors of the form $|\phi\rangle^{\otimes k}$, where $|\phi\rangle \in \mathbb{C}^d$.

Lemma 3.12 (k-th moment):

The operator K_k from Definition 3.11 is

$$K_k = \frac{k!(d-1)!}{(k+d-1)!} P_{\mathrm{sym}^k} \,, \tag{3.26}$$

where P_{sym^k} is the projector onto the symmetric subspace (2.25) of $(\mathbb{C}^d)^{\otimes k}$.

The required material on representation theory (Section 2.2) was also covered in Lecture 5.] Lecture 5 In order to prove Lemma 3.12 we will need a bit of representation theory, see Section 2.2 for the required preliminaries.

Proof of Lemma 3.12. Note that K_k is only supported on the symmetric subspace and, since $K_k = K_k^\dagger$, that also its range is contained in the symmetric subspaces, i.e. $\ker(K_k) = \mathrm{ran}(K_k) \subseteq \mathcal{H}_{\mathrm{sym}^k}$.

Next we observe that for any $U \in \mathrm{U}(d)$ we have $\Delta_d(U)^\dagger K_k \Delta_d(U) = K_k$ and, hence, $K_k \Delta_d(U) = \Delta_d(U) K_k$. Schur's lemma (Theorem 2.3) implies that $K_k|_{\mathcal{H}_{\mathrm{sym}^k}} = c\,\mathbb{1}|_{\mathcal{H}_{\mathrm{sym}^k}}$. Together with $\ker(K_k) = \mathrm{ran}(K_k) \subseteq \mathcal{H}_{\mathrm{sym}^k}$ identity implies that $K_k = c\,P_{\mathrm{sym}^k}$.

In order to obtain c, note that $\mathrm{Tr}[K_k] = 1$ and $\mathrm{Tr}[P_{\mathrm{sym}^k}]$ is given Eq. (2.30). □

Getting back to tight frames, we can make the following statement.

Proposition 3.13:

Complex spherical 1-designs are tight frames.

Proof. Let $\{|\psi_i\rangle\}_{i\in n}$ be a 1-design. We note that the frame potential S of $\{|\psi_i\rangle\}_{i\in n}$ is

$$S = nK_1 = \frac{n}{d}\mathbb{1} \,. \tag{3.27}$$

Moreover,

$$\mathrm{Tr}[S^2] = \frac{n^2}{d^2}\,\mathrm{Tr}[\mathbb{1}] = \frac{n^2}{d}\,. \tag{3.28}$$

Hence, by Theorem 3.10, $\{|\psi_i\rangle\}_{i\in n}$ is a tight frame. □

More generally, one can can lift Theorem 3.10 on frame potentials to k-designs. For this purpose, we define the *frame operator of order k* of $\{|\psi_i\rangle\}_{i\in n}$ to be

$$S_k := \sum_{i=1}^{n} |\psi_i\rangle\langle\psi_i|^{\otimes k}\,. \tag{3.29}$$

Theorem 3.14 (k-designs and tight frames [25]):

A set of states $\{|\psi_i\rangle\}_{i\in[n]}$ with $n \geq \binom{k+d-1}{d-1}$ is a k-design iff $\{|\psi_i\rangle^{\otimes k}\}_{i\in[n]}$ is a tight frame of the symmetric subspace, i.e., iff

$$\mathrm{Tr}[S_k^2] = \frac{n^2 k!(d-1)!}{(k+d-1)!}\,, \tag{3.30}$$

This is the global minimum of $\{|\psi_i\rangle\}_{i\in[n]} \mapsto \mathrm{Tr}[S_k^2]$.

Proof. By the definition of K_t and S_k we have that $\{|\psi_i\rangle\}_{i\in[n]}$ is a k-design iff

$$K_k = \frac{1}{n} S_k\,. \tag{3.31}$$

With $\left|\Psi_i^k\right\rangle := |\psi_i\rangle^{\otimes k}$ we can write S_k as

$$S_k = \sum_{i=1}^{n} \left|\Psi_i^k\right\rangle\!\left\langle\Psi_i^k\right|\,. \tag{3.32}$$

Since $\mathrm{Tr}[K_k] = 1$, Theorem 3.10 together with Lemma 3.12 finish the proof. □

3.4.1. Examples for 2-designs

Stabilizer states (STABs)

STABs are states that are ubiquitous in quantum information and are defined as follows. An n-qubit *Pauli string* is $\sigma_{s_1}\otimes\cdots\otimes\sigma_{s_n}$, where $s \in \{0,1,2,3\}^n$ and $\{\sigma_i\}$ are the Pauli matrices (2.1). Then the *Pauli group* $\mathcal{P}_n \subset \mathrm{U}(2^n)$ is the group generated by all n-qubit Pauli strings and $i\mathbb{1}$. An n-qubit state $|\psi\rangle$ is a *stabilizer state* if there is an abelian subgroup $\mathcal{S} \subset \mathcal{P}_n$, called *stabilizer (subgroup)*, that stabilizes $|\psi\rangle$ and only $|\psi\rangle$, i.e., $|\psi\rangle$ is the unique joint eigenvalue-1 eigenstate of all elements in that subgroup. Such subgroups turn out to be generated by n elements. Note that they cannot contain the element $-\mathbb{1}$.
An example of such a subgroup is the one of all Pauli strings made of $\mathbb{1}$'s and σ_z's.
The set of all stabilizer states is known to be a 2-design [31, 32], actually even a 3-design but not a 4-design [28, 33, 34].

Exercise 3.5 (Stabilizer states):

Let ρ be an n-qubit stabilizer state with stabilizer $\mathcal{S}$. Show that

$$\rho = \frac{1}{2^n} \sum_{S \in \mathcal{S}} S \,. \tag{3.33}$$

Mutually unbiased bases (MUBs)

MUBs are sets of bases with minimal overlaps. More explicitly, two orthonormal bases $\{|\psi_i\rangle\}_{i\in[d]} \subset \mathbb{C}^d$ and $\{|\phi_i\rangle\}_{i\in[d]} \subset \mathbb{C}^d$ are said to be *mutually unbiased* if $|\langle\psi_i|\phi_j\rangle|^2 = \frac{1}{d}$ for all $i, j \in [d]$. For instance, if $U \in \mathrm{U}(d)$ is the discrete Fourier transform the $\{|i\rangle\}_{i\in[d]} \subset \mathbb{C}^d$ and $\{U|i\rangle\}_{i\in[d]} \subset \mathbb{C}^d$ are mutually unbiased. The number of MUBs in $\mathbb{C}^d$ is upper bounded by $d+1$ and in prime power dimensions (e.g., for qubits) there are exactly $d+1$ MUBs [35]. However, it is a well-known open problem to exactly obtain this number for all d. Klappenecker and Roettler [36] showed that maximal sets of MUBs are 2-designs.

SIC POVMs

A *symmetric, informationally complete (SIC) POVM* is given by (see Eq. (3.41)) a set of d^2 normalized vectors $\{|\psi_j\rangle\}_{j\in[d^2]} \subset \mathbb{S}^{d-1} \subset \mathbb{C}^d$ satisfying

$$|\langle\psi_i|\psi_j\rangle|^2 = \frac{1}{d+1} \qquad \forall i \neq j \,. \tag{3.34}$$

"Symmetric" refers to the inner products being all equal. Renes et al. [25] have shown that SIC POVMs are indeed 2-designs and have explicitly constructed them for small dimensions.

3.5. Symmetric measurements and the depolarizing channel

The *depolarizing (quantum) channel* models isotropic noise in quantum processes. However, it also appears in the frame operator of several important measurement settings. In this section and the next ones we follow Guta et al. [37, Appendix, Section VI.A]. We will explicitly provide the pseudo-inverse of the measurement map for several types of measurements.

Definition 3.15 (Quantum depolarizing channel):

The *(quantum) depolarizing channel* $\mathcal{D}_p : \mathrm{L}(\mathbb{C}^d) \to \mathrm{L}(\mathbb{C}^d)$ with parameter $p \in [0, 1]$ is the linear map defined by

$$\mathcal{D}_p(X) \coloneqq pX + (1-p)\,\mathrm{Tr}[X]\frac{\mathbb{1}}{d} \,. \tag{3.35}$$

Proposition 3.16 (The inverse of the depolarizing channel):

For any $p > 0$ the depolarizing channel is invertible (as a map) and the inverse is given by

$$\mathcal{D}_p^{-1}(X) = \frac{1}{p}X - \frac{1-p}{p}\,\mathrm{Tr}[X]\frac{\mathbb{1}}{d} \,. \tag{3.36}$$

Proof. Exercise. □

Proposition 3.17 (Symmetric measurements):

Let $\mathcal{M} : \mathcal{S}(\mathbb{C}^d) \to \mathbb{R}^m$ be given by measurement operators $\mathsf{M} = \{M_1, \dots, M_m\} \subset \mathrm{Herm}(\mathbb{C}^d)$. Suppose that the frame operator, which is given as $S_{\mathsf{M}} = \mathcal{M}^\dagger \mathcal{M}$ is $S_{\mathsf{M}} = c\,\mathcal{D}_p$ for constants $c > 0$ and $p > 0$ (possibly depending on d and m). Then the dual frame is

$$\tilde{M}_i = \frac{1}{cp} M_i - \frac{1-p}{cpd} \operatorname{Tr}[M_i]\, \mathbb{1} \tag{3.37}$$

and the pseudo-inverse of $\mathcal{M}$ given by

$$\mathcal{M}^+(y) = \frac{1}{cp} \left(\mathcal{M}^\dagger(y) - (1-p) \operatorname{Tr}[\mathcal{M}^\dagger(y)] \frac{\mathbb{1}}{d} \right) . \tag{3.38}$$

Proof. We remember the adjoint measurement map (3.40) and the definition (3.17) of the frame operator to note that indeed

$$\mathcal{M}^\dagger \mathcal{M}(X) = \sum_{i=1}^m \operatorname{Tr}[X M_i]\, M_i = S_{\mathsf{M}} . \tag{3.39}$$

Since $\mathcal{D}_p$ is invertible for $p \neq 0$ the measurement operators M are a frame for $\mathrm{Herm}(\mathbb{C}^d)$.

The inverse frame operator is obtained as $S_{\mathsf{M}}^{-1} = \frac{1}{c}\mathcal{D}_p^{-1}$ with the inverse depolarizing channel (3.36). The dual frame is obtained by $\tilde{M}_i = S_{\mathsf{M}}^{-1}(M_i)$ and noting that the Hilbert-Schmidt adjoint of $\mathcal{M}$, with the component-wise definition (3.2), is given by

$$\mathcal{M}^\dagger(y) = \sum_{i=1}^m y_i M_i . \tag{3.40}$$

In order to obtain the pseudo-inverse of $\mathcal{M}$ we observe that $\mathcal{M}^\dagger \mathcal{M} = S_{\mathsf{M}}$ and that this operator is invertible. Hence, the pseudo-inverse can be obtained with the linear inversion formula (3.14). □

The pseudo-inverse (3.38) is given by scaling and shifting $\mathcal{M}^\dagger$, so up to this scaling and shifting it acts like a unitary. This justifies calling such measurements *symmetric*. This is a weaker form of symmetry as the one required for SIC POVMs [25], which is an instance of POVMs based on 2-designs.

3.5.1. 2-design based POVMS

There are two different approaches to connect a POVM $\mathsf{M} = \{M_1, \dots, M_m\}$ to frame theory: (i) to consider M as a frame for $\mathrm{Herm}(\mathbb{C}^d)$ [38] and (ii) to construct M from a frame for the Hilbert spaces $\mathbb{C}^d$. While the idea (i) seems to be more straightforward, (ii) can provides a lot of structure that can be exploited in various ways.

Throughout this section we consider a POVM $\mathsf{M} = \{M_1, \dots, M_m\}$ that is proportional to a complex projective 2-design $\{|\psi_i\rangle\langle\psi_i|\}_{i\in[m]} \subset \mathbb{C}^d$,

$$M_i = \frac{d}{m} |\psi_i\rangle\langle\psi_i| \tag{3.41}$$

for $i \in [m]$.

Proposition 3.18:

M with measurement operators (3.41) is indeed a POVM.

Proof. Obviously, $M_i \succeq 0$. Taking the partial trace of K_2, using Lemma 3.12 and remembering the basis expansion of the flip operator (2.32) yields

$$\begin{aligned} \mathrm{Tr}_2[K_2] &= \frac{2}{d(d+1)} \mathrm{Tr}_2[P_{\mathrm{sym}^2}] = \frac{1}{d(d+1)} \mathrm{Tr}_2[\mathbb{1} + \mathbb{F}] \\ &= \frac{1}{d(d+1)}(d\mathbb{1} + \mathbb{1}) = \frac{1}{d}\mathbb{1}\,. \end{aligned} \tag{3.42}$$

Noting that $\sum_{i=1}^m M_i = d\ \mathrm{Tr}_2[K_2]$ shows that M is indeed a POVM. □

The POVM M is a frame with a frame operator given in terms of the measurement map (3.2) and dual frame as follows.

Lemma 3.19 (Frame from 2-design POVMs):

Any POVM M given by a 2-design as in (3.41) is a frame for $\mathrm{Herm}(\mathbb{C}^d)$ with frame operator

$$S_{\mathsf{M}} = \frac{d}{m}\mathcal{D}_p \tag{3.43}$$

and parameter $p = \frac{1}{d+1}$.

Proof. By the definition of the measurement operators (3.41) the frame operator can be written in terms of the 2-design as

$$S_{\mathsf{M}}(X) = \frac{d^2}{m^2}\sum_{i=1}^m \mathrm{Tr}[X\ |\psi_i\rangle\langle\psi_i|]\ |\psi_i\rangle\langle\psi_i| \tag{3.44}$$

$$= \frac{d^2}{m}\mathrm{Tr}_1\left[(X\otimes\mathbb{1})\frac{1}{m}\sum_{i=1}^m |\psi_i\rangle\langle\psi_i| \otimes |\psi_i\rangle\langle\psi_i|\right] \tag{3.45}$$

$$= \frac{d^2}{m}\mathrm{Tr}_1[(X\otimes\mathbb{1})K_2]\,. \tag{3.46}$$

Using again Lemma 3.12 and the flip operator (2.33),

$$\begin{aligned} S_{\mathsf{M}}(X) &= \frac{d^2}{m}\frac{2}{d(d+1)}\mathrm{Tr}_1[(X\otimes\mathbb{1})P_{\mathrm{sym}^2}] \\ &= \frac{d}{m(d+1)}\left(\mathrm{Tr}[X]\mathbb{1} + \mathrm{Tr}_1[(X\otimes\mathbb{1})\mathbb{F}]\right)\,. \end{aligned} \tag{3.47}$$

Now we can use the swap-trick (2.68) to obtain

$$\begin{aligned} S_{\mathsf{M}}(X) &= \frac{d}{m(d+1)}(X + \mathrm{Tr}[X]\mathbb{1}) \\ &= \frac{d}{m}\mathcal{D}_{1/(d+1)}(X)\,. \end{aligned} \tag{3.48}$$

This identity implies that S_{M} is invertible. Hence, M is indeed a frame. □

This lemma allows for a full frame characterization of the POVMs that are given by spherical 2-designs.

Corollary 3.20 (Frame characterization and linear inversion for POVMs from 2-designs):

M given by (3.41) is a frame for $\mathrm{Herm}(\mathbb{C}^d)$ with frame operator given by

$$S_{\mathsf{M}}(X) = \frac{d}{m(d+1)}(X + \mathrm{Tr}[X]\mathbb{1})\,. \tag{3.49}$$

The inverse frame operator is given by

$$S_{\mathsf{M}}^{-1}(X) = \frac{m}{d}\big((d+1)X - \mathrm{Tr}[X]\,\mathbb{1}\big) \tag{3.50}$$

and the dual frame $\tilde{\mathsf{M}} \coloneqq \{\tilde{M}_1, \ldots, \tilde{M}_m\}$ by

$$\tilde{M}_i = \frac{m(d+1)}{d} M_i - \mathbb{1}\,. \tag{3.51}$$

The pseudo-inverse of the measurement map $\mathcal{M}$ is given by

$$\begin{aligned} \mathcal{M}^+(y) &= (d+1)\sum_{i=1}^m y_i\,|\psi_i\rangle\langle\psi_i| - \sum_{i=1}^m y_i\,\mathbb{1} \\ &= \frac{m(d+1)}{d}\,\mathcal{M}^\dagger(y) - \langle \mathbf{1}, y\rangle\,\mathbb{1}\,, \end{aligned} \tag{3.52}$$

where $\mathbf{1} \in \mathbb{R}^m$ is the vector containing only ones.

Proof. Having characterized the frame operator in Lemma 3.19 this corollary directly follows with Proposition 3.17 and Proposition 3.16. □

3.5.2. Pauli measurements

Let us denote the n-qubit Pauli strings by $W_0, \ldots, W_{d^2-1}$ with $W_0 = \mathbb{1}$, where $d = 2^n$. Note that the spectrum of each non-identity Pauli string is $\{-1, 1\}$, each with degeneracy $d/2$. So, each Pauli string is $W_i = P_i^+ - P_i^-$, where

$$P_i^\pm = \frac{1}{2}(\mathbb{1} \pm W_i) \tag{3.53}$$

is the projector onto the eigenvalue ± 1 eigenspace (note that $P_0^+ = \mathbb{1}$ and $P_0^- = 0$). Each Pauli string (observable) W_i is associated with a two-outcome POVM given by $\mathsf{M}_i \coloneqq \{P_i^+, P_i^-\}$. Now we consider a measurement setting given by the union of all these POVMs plus the trivial measurement given by $W_0 = \mathbb{1}$, i.e., by

$$\mathsf{M} \coloneqq \bigcup_{i=0}^{d^2-1} \{P_i^+, P_i^-\}\,. \tag{3.54}$$

The frame operator is given by

$$\begin{aligned} S_{\mathsf{M}}(X) = \mathcal{M}^\dagger\,\mathcal{M}(X) &= \sum_{i=0}^{d^2-1} \big(\mathrm{Tr}[P_i^+ X]P_i^+ + \mathrm{Tr}[P_i^- X]P_i^-\big) \\ &= \sum_{i=0}^{d^2-1} \frac{1}{2}\,(\mathrm{Tr}[W_i X]W_i + \mathrm{Tr}[X]\,\mathbb{1}) \\ &= \frac{d}{2}\,(X + d\,\mathrm{Tr}[X]\,\mathbb{1})\,, \end{aligned} \tag{3.55}$$

where we have used that $\{W_i/\sqrt{d}\}_{i=0}^{d^2-1}$ is an orthonormal basis of $\mathrm{Herm}(\mathbb{C}^d)$ w.r.t. the Hilbert-Schmidt inner product. In terms of the depolarizing channel (Definition 3.15) this result reads as

$$S_{\mathsf{M}} = \frac{d(d^2+1)}{2}\mathcal{D}_{\frac{1}{d^2+1}}\,. \tag{3.56}$$

Each POVM $\mathsf{M}_i = \{P_i^+, P_i^-\}$ is associated to measurement outcomes $\pm$. For a state $\rho \in \mathcal{S}(\mathbb{C}^d)$ the ideal measurement outcomes are $\mathrm{Tr}[\rho P_i^\pm]$ and a measurement vector $y_i^\pm \in \mathbb{R}^2$ is the corresponding actually measured quantity. Together, this yields a measurement vector $y = (y_i^o)_{i\in[d^2],o\in\{\pm\}} \in \mathbb{R}^m$ with $m = 2d^2$.

Now use use Proposition 3.17 with $c = \frac{1}{2}d(d^2+1)$ and $p = 1/(d^2+1)$ to calculate the pseudo-inverse of the measurement map. Note that $1-p = \frac{d^2}{d^2+1}$ and $\frac{1}{cp} = \frac{2}{d}$ Using (3.53), $y_i^+ + y_i^- = 1$, and Proposition 3.17 with $c = \frac{1}{2}d(d^2+1)$ and $p = 1/(d^2+1)$, the liner inversion estimate (3.14) simplifies as

$$\begin{aligned}
\mathcal{M}^+(y) &= \frac{2}{d}\left(\sum_{i=0}^{d^2-1}\sum_{o=\pm} y_i^o P_i^o - \frac{d^2}{d^2+1}\sum_{i=0}^{d^2-1}\sum_{o=\pm} y_i^o \,\mathrm{Tr}[P_i^o]\,\frac{\mathbb{1}}{d}\right)\\
&= \frac{2}{d}\left(\sum_{i=0}^{d^2-1}\left(\frac{1}{2}\left(y_i^+ - y_i^-\right)W_i + \frac{\mathbb{1}}{2}\right) - \frac{d^2}{d^2+1}\sum_{i=0}^{d^2-1}\frac{d}{2}\frac{\mathbb{1}}{d}\right)\\
&= \frac{1}{d}\sum_{i=0}^{d^2-1}\left(y_i^+ - y_i^-\right)W_i\,.
\end{aligned} \tag{3.57}$$

This formula has a nice interpretation: The difference $y_i^+ - y_i^-$ is the empirical expectation of W_i and the pseudo-inverse is the corresponding empirical estimate of the state ρ (the $1/d$ factor comes from the normalization of W_i).

3.5.2.1. Full Pauli basis measurements

Often the measurement of a Pauli string $W_i = \sigma_{s_1^{(i)}} \otimes \ldots \sigma_{s_n^{(i)}}$ yields one binary measurement outcome per qubit even though the spectrum of W_i is degenerate. The corresponding post-measurement states are the tensor products of the eigenstates of $\sigma_{s_1^{(i)}}, \ldots, \sigma_{s_n^{(i)}}$ denoted by $|b_{\boldsymbol{s}}^{\boldsymbol{o}}\rangle \coloneqq |b_{s_1}^{o_1}\rangle \otimes \cdots \otimes |b_{s_n}^{o_n}\rangle$. The set of all tensor products of all single qubit Pauli basis states

$$\mathsf{M} \coloneqq \left\{|b_{\boldsymbol{s}}^{\boldsymbol{o}}\rangle\right\}_{\boldsymbol{o}\in\{\pm\}^n,\boldsymbol{s}\in\{x,y,z\}^n} \tag{3.58}$$

defines the full set of Pauli basis measurements. For $n \geq 2$ qubits M^{PB} is not a 2-design, as can be checked via Eq. (3.30). However, the pseudo-inverse of the measurement map can still be inverted in a similar fashion as before.

Exercise 3.6 (Pauli basis measurements):

Find the frame operator, the dual frame, and the pseudo-inverse of the measurement of Pauli-basis measurements.

This exercise yields the linear inversion estimator, i.e., the pseudo-inverse of the measurement map as

$$\mathcal{M}^+(y) = \frac{1}{3^n}\sum_{\boldsymbol{s}\in\{x,y,z\}^n}\sum_{\boldsymbol{o}\in\{\pm\}^n} y_{\boldsymbol{s}}^{\boldsymbol{o}}\bigotimes_{i=1}^{n}\left(3\,|b_{\boldsymbol{s}}^{\boldsymbol{o}}\rangle\langle b_{\boldsymbol{s}}^{\boldsymbol{o}}| - \mathbb{1}\right)\,. \tag{3.59}$$

3.6. Compressed sensing

The new material on linear programming (Section 2.5.1) was also covered in Lecture 13.] Lecture 13 Compressed sensing (a.k.a. compressive sensing) is still quite a young research field concerned with the reconstruction of signals from few measurements, where a *signal* is anything one might want to recover from measurements. The field was mainly initialized by works of Candès, Romberg, and Tao [39, 40] and Donoho [41]. It has an extremely wide range of applications. In particular, it is of great use in many technological applications including imaging, acoustics, radar, and mobile network communication.

The problems considered in compressed sensing are typically of the following form. There is a *signal* $x \in V \cong \mathbb{R}^n$ that is known to be compressible (e.g., a sparse vector or low-rank matrix). Then one is given access to linear measurements

$$y_i = \langle a_i, x \rangle + \epsilon_i \quad \text{for } i \in [m] \tag{3.60}$$

that are given by *measurement vectors* $a_i \in V$ and additive *noise* $\epsilon_i \in \mathbb{R}$ that might arise in the measurement process. Now the task is to reconstruct x from y and $(a_i)_{i\in[m]}$. We wish the reconstruction to be efficiently implementable and guaranteed to work for m being as small as possible.

In matrix notation we can rewrite (3.60) as

$$y = \mathcal{A}(x) + \epsilon \tag{3.61}$$

where the i-th component of the *measurement map* $\mathcal{A}$ is $\mathcal{A}(x)_i := \langle a_i, x \rangle$. For a simplified setting with $m \geq n$ and without noise ($\epsilon = 0$) we know from linear algebra that the inverse problem (3.61) has a solution for all signals X iff $\mathcal{A}$ is injective. For $m < n$, however, $\mathcal{A}$ is given by a *short fat matrix*, which cannot be injective.

Compressed sensing exploits a known compressibility of X in order to still practically solve the inverse problem (3.61) for several simple forms of compressibility, the simplest one being sparsity. Indeed, signals that are sparse in a known basis (or frame) have received the most attention [10]. Let us denote by $\|x\|_{\ell_0} := |\{i : x_i \neq 0\}|$ the *sparsity* or ℓ_0-norm (which is not a norm) of a signal $x \in \mathbb{R}^n$. Moreover, let us consider again the noiseless case, i.e. $\epsilon = 0$. If the inverse problem given by (3.61) has a unique solution then the solution must be the minimizer of the optimization problem

$$\text{minimize } \|z\|_{\ell_0} \quad \text{subject to } \mathcal{A}(z) = y\,. \tag{3.62}$$

While this reconstruction is clearly optimal in terms of the required number of measurements m, it cannot be implemented efficiently. Indeed, even for any fixed $\eta \geq 0$, the more general problem

$$\text{minimize } \|z\|_{\ell_0} \quad \text{subject to } \|\mathcal{A}(z) - y\|_{\ell_2} \leq \eta \tag{3.63}$$

is NP-hard, as can be proven by a reduction from *exact cover by 3-sets problem* [42] (see also [10, Theorem 2.17]). In fact, the NP-hardness still holds when the ℓ_0-norm is replaced by an ℓ_q-norm for any $q \in [0, 1)$ [43] (which are *quasi-norms*).

We note that these hardness results extend to matrices, where the sparsity is replaced by the matrix rank, i.e., where the recovery problem is of the form

$$\text{minimize } \operatorname{rank}(Z) \quad \text{subject to } \|\mathcal{M}(Z) - y\|_{\ell_2} \leq \eta \tag{3.64}$$

for some fixed $\eta \geq 0$, where $\mathcal{M}(Z)_i = \operatorname{Tr}[M_i^\dagger Z]$ for *measurement matrices* M_i.

At first sight, these results seem to be bad news for the compressed sensing idea. However, these results just mean that the problems (3.63) are hard to solve for *general* matrices $\mathcal{A}$ and vectors y, i.e., that there cannot be an efficient algorithm that solves *all* instances of the problem. In practical applications, however, one is typically given very specific measurement setups that determine $\mathcal{A}$ and y and has also some control allowing to tune these quantities.

Compressed sensing makes use of this observation. In typical settings, the measurements $\mathcal{A}$ are drawn at random, where hard instances do practically not occur. The function that measures the compressibility of the signal is replaced by a tractable convex function. In particular, for the mentioned examples, the sparsity is replaced by the ℓ_1-norm and the matrix rank by the trace norm, i.e., the problems (3.63) and (3.64) are replaced by

$$\text{minimize } \|z\|_{\ell_1} \quad \text{subject to } \quad \|\mathcal{A}(z) - y\|_{\ell_2} \leq \eta\,, \tag{3.65}$$

$$\text{minimize } \|Z\|_1 \quad \text{subject to } \quad \|\mathcal{M}(Z) - y\|_{\ell_2} \leq \eta\,, \tag{3.66}$$

respectively. Such an idea of replacing a non-convex function by a convex one in an optimization problem is called *convex relaxation* and can be done systematically for a given set of the "most simple" signals [44] by using convex analysis [11]. Then, for the easiest-to-analyze case of fully Gaussian measurement vectors, it turns out that the number of measurements required for recovery (with high probability over the random measurements) is

$$m^* \approx 2s\ln(n/s) + 2s \qquad \text{and} \tag{3.67}$$

$$m^* \approx 3r(n_1 + n_2 - r) \tag{3.68}$$

for s-sparse signals in $\mathbb{R}^n$ and real $n_1 \times n_2$ matrices of rank r, respectively [44–46]. Note that these numbers are much smaller than the dimensions of the ambient space, $m^* \ll n$ and $m^* \ll n_1 n_2$, respectively. In particular, the number of measurements for case of $d \times d$ rank-1 matrices is $6d - 1$. This number is roughly a factor of 1.5 larger than needed to guarantee injectivity of the measurement map $\mathcal{M}$, which is roughly $4d$ according to Theorem 3.7. However, this comparison is not completely fair as (3.68) holds for *real* and non-Hermitian matrices but the complex Hermitian behaves very similarly.

It should be emphasized that these results come along with rigorous recovery guarantees that can be extended to many different types of measurements (not just Gaussian ones). However, for each new type of measurements the proofs need to be adjusted.

The new material on linear programming (Section 2.5.2) was also covered in Lecture 14.] Lecture 14 Let us from now on only focus on low-rank matrix reconstruction, as this is the relevant problem for quantum state tomography. We emphasize that the convex optimization problem (3.65) (and also (3.66)) can be solved efficiently by rephrasing it as SDP (see Section 2.5). Once we can write a constraint of the form $\|w\|_{\ell_2} \leq \eta$ as a PSD constraint we can include $\|\mathcal{M}(Z) - y\|_{\ell_2} \leq \eta$ as additional constraint into the dual SDP formulation of the trace norm (2.63) and obtain an SDP formulation of the constraint trace norm minimization (3.66). But $\|w\|_{\ell_2} \leq \eta$ can be written as

$$\sum_i \lambda_i \leq \eta^2 \quad \text{and} \quad w_i^2 \leq \lambda_i \tag{3.69}$$

and the latter constraints can be rewritten as PSD constraint using (2.60) as

$$|w_i|^2 \leq \lambda_i \quad \Leftrightarrow \quad \begin{pmatrix} \lambda_i & w_i \\ w_i^* & \lambda_i \end{pmatrix} \succeq 0\,. \tag{3.70}$$

There are several variations of the reconstruction (3.66). For instance, there is the *matrix Dantzig selector* where the measured matrix is estimated by the minimizer of

$$\text{minimize } \|Z\|_1 \quad \text{subject to } \left\|\mathcal{M}^\dagger(\mathcal{M}(Z) - y)\right\|_{\text{op}} \leq \lambda\,, \tag{3.71}$$

where λ depends on the noise strength as it needs to satisfy $\lambda \geq \left\|\mathcal{M}^\dagger(\epsilon)\right\|_{\text{op}}$. *matrix Lasso* (least absolute shrinkage and selection operator) estimate obtained as the optimal point of

$$\text{minimize } \mu\|Z\|_1 + \frac{1}{2}\|\mathcal{M}(Z) - y\|_{\ell_2}^2\,, \tag{3.72}$$

where $\mu \geq 2\|\mathcal{M}^\dagger(\epsilon)\|_{\mathrm{op}}$. These optimization problems can all be written as SDPs and yield very similar solutions.

There are also computationally more efficient algorithms to solve these convex optimization problems. For instance, one can use *singular value thresholding*, which is an iterative algorithm where one iteratively (i) updates a matrix by adding $\eta\,\mathcal{M}^\dagger(y)$ for some step size η and (ii) shrinks the singular values to keep it approximately low rank. In fact, such algorithms are an instance of the *alternating direction method of multipliers* (ADMM), see e.g. [47], where one solves a class of convex optimization problems including (3.66) iteratively. These algorithms are relatively fast and also allow for rigorous performance guarantees. Another direction is to us non-convex optimization methods to directly approximately solve the rank minimization problem (3.64). These methods are even faster and often need even fewer measurements but rigorous guarantees very rare.

Finally, let us summarize what a recovery guarantee for low rank matrix reconstruction exactly is. Typical cases are covered by the following definition.

What is a recovery guarantee?

Let $X \in \mathrm{Herm}(\mathbb{C}^d)$ be a *signal* of matrix rank r and $M_i \in \mathrm{Herm}(\mathbb{C}^d)$ for $i \in [m]$ be *measurement matrices* drawn iid. from some measure μ on $\mathrm{Herm}(\mathbb{C}^d)$. Moreover, let

$$y_i = \mathrm{Tr}[M_i X] + \epsilon_i\,, \qquad i \in [m]\,, \tag{3.73}$$

define a *measurement vector*, possibly corrupted by *additive noise* ϵ bounded as $b(\epsilon) \leq \eta$ (typically $b = \|\cdot\|_{\ell_2}$).
Then a *recovery guarantee* provides a threshold m_0 depending on r and d so that for all $m \geq m_0$

$$\|\hat{X} - X\|_p \leq h(\eta) \tag{3.74}$$

(typically $p = 1$) with probability at least $1 - \delta_m$.
Typically, $\delta_m \leq C\mathrm{e}^{-c\,m}$ for some constants $C, c > 0$. The function h captures *stability* of the reconstruction against additive noise. Similarly, one can quantify the *robustness* of the reconstruction against small violations against the low-rank assumption. Typically, the robustness is quantified in terms of $\|X - X_r\|_1$, where X_r is the best rank-r approximation to X in one (and then all) Schatten p-norms.
A recovery guarantee is called *uniform* if w.h.p. a drawn measurement map recovers all matrices satisfying the rank constraint.

It is crucial that a "good" distribution μ of measurements is chosen. For instance, if μ is a projective 4-design one can prove recovery guarantees with an essentially optimal scaling [48] but if μ is a only a projective 2-design –such as a maximum set of MUBs– then the measurement if not injective with high probability [49, Theorem 18]. This shows that the measurement matrices need to be able to cover the matrix space densely enough in order for such randomized strategies to work.

3.6.1. Application to quantum state tomography

In quantum state tomography one often aims at the reconstruction of an (approximately) pure state. Hence, compressed sensing can be naturally applied here. In fact, exactly this application has lead to a great break through in low-rank matrix reconstruction [50, 51] by using matrix concentration inequalities that arose in quantum information theory [52].

In the quantum state tomography setting there is a simple way to improve the reconstruction via the trace norm minimization (3.66). We can simply add the constraint that Z is a density matrix, which is indeed a PSD constraint. But then $\|Z\|_1 = 1$, so the minimization is superfluous and we are left with a feasibility problem, i.e., with the task to find a density matrix Z such that $\|\mathcal{M}(Z) - y\|_{\ell_2} \leq \eta$. But then we can

μ, Reference	m_0	$h(\eta)$	Remarks/properties	Proof technique
Random Paulis, [50]	$\mathrm{O}(r\,d\,\ln(d)^2)$	$\mathrm{O}(\eta\sqrt{r\,d})$		Golfing scheme [51]
Random Paulis, [55]	$\mathrm{O}(r\,d\,\ln(d)^6)$	$\mathrm{O}(r\,\eta)$, where $\eta \geq \|\mathcal{M}^\dagger(\epsilon)\|_{\mathrm{F}}$	Uniform, robust	RIP [56]
$\mathrm{rank}(M_i)=1$, $\|M_i\|_{\mathrm{op}} \lesssim d$, (approx.) proj. 4-designs, [48]	$\mathrm{O}(r\,d\,\ln(d))$	$\mathrm{O}(\eta/\sqrt{m})$, $p=2$	Uniform	Bowling scheme [46]
$\mathrm{rank}(M_i)=1$, $\|M_i\|_{\mathrm{op}} \lesssim d$, (approx.) proj. 4-designs, [54]	$\mathrm{O}(r\,d\,\ln(d))$	$\mathrm{O}\left(\frac{\|\epsilon\|_{\ell_2}\, r^{1/p-1/2}}{\sqrt{m}}\right)$	PSD-fit, uniform, robust	NSP [54]
Random Clifford orbits [57]	$\mathrm{O}(r^3\,d\,\ln(d))$	$\mathrm{O}\left(\frac{r^2 d\|\epsilon\|_{\ell_q}}{m^{1/q}}\right)$	PSD-fit, uniform, robust	NSP [54], Clifford irreps [58]

Table 3.1.: List of recovery guarantees for compressed sensing based quantum state tomography. "Random Paulis" refers to measurements of Pauli string observables. A random Clifford orbit is a set of states obtained by applying all Clifford group operations to a certain fixed state.
There are several proof techniques for low-rank matrix reconstruction: The *Golfing scheme* relies on dual certificates, the *restricted isometry property* quantifies the "distortion" of the relevant signals under the measurement map, the *Bowling scheme* is a combination of geometric proof techniques [44] and *Mendelson's small ball method* [59], and the *(stable and robust rank) null space property* (NSP) is a property of the measurement map that allows to quantify the robustness of a reconstruction against violations of the low-rank assumption.

minimize that quantity instead. This yields the so-called PSD-fit

$$\text{minimize} \quad \|\mathcal{M}(Z) - y\|_{\ell_2}^2 \;\text{ subject to }\; Z \in \mathcal{S}(\mathbb{C}^d)\,, \tag{3.75}$$

which was suggested by Baldwin et al. [53] and made rigorous by Kabanava et al. [54]. Besides having a slightly improved performance the main advantage of this reconstruction method is that it does not require an estimate on the noise strength.

For an overview of recovery guarantees for convex compressed sensing methods in state tomography see Figure 3.1.

3.7. Projected least squares estimation

Extended norm-minimization estimation [60], or now called *projected least squares (PLS) estimation* [37] complements linear inversion in order improve the reconstruction and to analytically give confidence regions (in trace norm) for the reconstructed state. Similarly as in compressed sensing, the PLS estimator can take advantage of low-rankness of the measured state [37].

First confidence guarantees in terms of a conditioning of the measurement map have been provided by Sugiyama et al. [60]. These guarantess have recently been improved and calculated explicitly for the following types of measurements by Guta et al. [37]:

- 2-design based measurements (Section 3.5.1, Eq. (3.41)),
- Pauli observable measurements (Section 3.5.2, Eq. (3.54)),
- Pauli basis measurements (Section 3.5.2.1, Eq. (3.58)).

PLS estimation works in two steps: (i) a linear inversion (3.14) of the measurement data is performed and (ii) the linear inversion estimate is projected onto the set of density matrices. Remember that the least squares estimator can bey calculated as $\hat{\rho}_{\mathrm{LS}} = \mathcal{M}^+(y)$ (Proposition 3.8).

We again denote the measurement matrices by $\mathsf{M} = \{M_1, \ldots, M_m\} \subset \mathrm{Herm}(\mathbb{C}^d)$ and the corresponding measurement map by $\mathcal{M} : \mathrm{Herm}(\mathbb{C}^d) \to \mathbb{R}^m$. Then the PLS estimator of a quantums state $\rho \in \mathcal{S}(\mathbb{C}^d)$ from measurement data $y = \mathcal{M}(\rho) + \epsilon$ is explicitly given as [60]

$$\hat{\rho}_{\mathrm{PLS}} \coloneqq \operatorname*{arg\,min}_{\sigma \in \mathcal{S}(\mathbb{C}^d)} \|\sigma - \hat{\rho}_{\mathrm{LS}}\|_2\,, \tag{3.76}$$

where $\hat{\rho}_{\mathrm{LS}} = \mathcal{M}^+(y)$ and $\mathcal{M}^+$ denotes the pseudo-inverse of $\mathcal{M}$.

Theorem 3.21 (Error bounds for PLS [37]):

Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be a state, fix a total number of measurements of ρ to be n_ρ and let $\varepsilon \in [0,1]$. Consider the following three measurement setups with associated function g,

(i) $g(d) \coloneqq 2d$ for 2-design based POVMs (3.41),

(ii) $g(d) \coloneqq d^2$ for Pauli observable measurements (3.54), and

(iii) $g(d) \coloneqq d^{1.6}$ for Pauli basis measurements (3.58).

Then the PLS estimator (3.76) satisfies

$$\mathbb{P}[\|\hat{\rho}_{\mathrm{PLS}} - \rho\|_1 \geq \varepsilon] \leq d \exp\left(-\frac{n_\rho \varepsilon^2}{43\, g(d) r^2}\right), \tag{3.77}$$

where $r = \min\{\mathrm{rank}(\rho), \mathrm{rank}(\hat{\rho}_{\mathrm{PLS}})\}$.

We will outline the proof only for the case of 2-design based POVMs in Section 3.7.1.

The theorem tells implies that with probability at least $1-\delta$ the PLS estimation yields an estimate of ρ within trace norm error ε whenever the number of measurements is

$$n_\rho \geq 43\, g(d) \frac{\mathrm{rank}(\rho)^2}{\varepsilon^2} \ln(d/\delta)\,. \tag{3.78}$$

The PLS estimation (3.76) can be written as an SDP, which can be seen with the machinery of Section 2.5.2. However, the PLS estimator allows for an analytic solution (up to one parameter), which allows to compute it much faster than the SDP runtime.

Proposition 3.22 (State space projection [61], version of [37]):

Let $P_{\mathcal{S}} : \mathrm{Herm}(\mathbb{C}^d) \to \mathrm{Herm}(\mathbb{C}^d)$ be the Euclidean projection onto the set of density matrices, i.e.,

$$P_{\mathcal{S}}(X) \coloneqq \operatorname*{arg\,min}_{\rho \in \mathcal{S}(\mathbb{C}^d)} \|X - \rho\|_{\mathrm{F}} \tag{3.79}$$

for $X \in \mathrm{Herm}(\mathbb{C}^d)$. Then for $\mathcal{X} \in \mathrm{Herm}(\mathbb{C}^d)$ with $\mathrm{Tr}[X] = 1$ the projection $\rho = P_{\mathcal{S}}(X)$ can be calculated as follows from and eigenvalue decomposition $X = U\,\mathrm{diag}(\lambda)U^\dagger$, with $\lambda \in \mathbb{R}^d$ and $U \in \mathrm{U}(d)$. Set

$$\rho \coloneqq U\,\mathrm{diag}([\lambda - x_0 \mathbf{1}]^+)U^\dagger \tag{3.80}$$

where $[y]_i^+ \coloneqq \max\{y_i, 0\}$, $\mathbf{1} \in \mathbb{R}^d$ is the vector with only 1-entries, and x_0 is chosen such that $\mathrm{Tr}[\rho] = 1$. Moreover, x_0 is the root of the function f defined by[a]

$$f(x) \coloneqq \sum_{i=1}^{d} |\lambda_i - x| - d \cdot x - \mathrm{Tr}[X]\,. \tag{3.81}$$

[a] We obtained a slightly different function than in [37, ArXiv version v1].

Proof. Exercise. □

3.7.1. Proof of Theorem 3.21 for 2-design based POVMs

In order to fully explain the general approach of PLS estimation we review the proof of Theorem 3.21 for the 2-design based measurements in more detail. Let us remember

that the linear inversion on the measurement data is given by $\mathcal{M}^+$ with the closed form expression (3.52).

We note that M is a single POVM and we take

$$y_i = \frac{n_i}{n_\rho} \tag{3.82}$$

where n_i is the number of times (out of n_ρ single measurements) the outcome i corresponding to POVM element

$$M_i = \frac{d}{m} |\psi_i\rangle\langle\psi_i| \tag{3.83}$$

(cp. (3.41)) was measured.

First, we quantify the error that arises in the linear inversion.

Lemma 3.23 ([37, Appendix, Section VII.A]):

For measurements (3.82) set $\hat{\rho}_{\mathrm{LS}} \coloneqq \mathcal{M}^+(y)$ with the pseudo-inverse $\mathcal{M}^+$ from (3.52). Then

$$\mathbb{P}[\|\hat{\rho}_{\mathrm{LS}} - \rho\|_{\mathrm{op}} \geq \tau] \leq d \exp\left(-\frac{3\tau^2 n}{16d}\right) \tag{3.84}$$

for all $\tau \in [0, 2]$.

For the proof we use a matrix version of the Bernstein inequality from Theorem 2.11.

Theorem 3.24 (Matrix Bernstein inequality [62, Theorem 1.4]):

Let $X_1, \ldots, X_\ell \in \mathrm{Herm}(\mathbb{C}^d)$ be independent Hermitian random matrices with

$$\mathbb{E}[X_i] = 0 \quad \text{and} \quad \|X_i\|_{\mathrm{op}} \leq a \quad \text{almost surely} \tag{3.85}$$

for all $i \in [\ell]$ and some $a > 0$. Set

$$\sigma^2 \coloneqq \left\| \sum_{i=1}^{\ell} \mathbb{E}[X_i^2] \right\|_{\mathrm{op}} . \tag{3.86}$$

Then, for all $\tau > 0$,

$$\mathbb{P}\left[\left\| \sum_{i=1}^{\ell} X_i \right\|_{\mathrm{op}} \geq \tau \right] \leq d \exp\left(-\frac{\tau^2/2}{\sigma^2 + a\,\tau/3} \right) . \tag{3.87}$$

Proof of Lemma 3.23. From $\mathrm{Tr}[\rho] = 1$ and M being a single POVM follows that $\sum_{i=1}^m y_i = 1$. Hence, the pseudo-inverse of the measurement operator (3.52) becomes

$$\begin{aligned} \mathcal{M}^+(y) &= (d+1) \sum_{i=1}^m y_i \, |\psi_i\rangle\langle\psi_i| - \sum_{i=1}^m y_i \, \mathbb{1} \\ &= \sum_{i=1}^m \frac{n_i}{n_\rho} \left((d+1) \, |\psi_i\rangle\langle\psi_i| - \mathbb{1} \right) \\ &= \frac{1}{n_\rho} \sum_{k=1}^{n_\rho} Y_k \, , \end{aligned} \tag{3.88}$$

where $\{Y_k\}_{k \in [n_\rho]}$ are iid. copies of a random matrix Y corresponding to the measurement outcome of the POVM M: Y is $(d+1) \, |\psi_i\rangle\langle\psi_i| - \mathbb{1}$ with probability $\mathbb{P}[i] = \mathrm{Tr}[M_i \rho] = \frac{d}{m} \langle\psi_i| \, \rho \, |\psi_i\rangle$. By construction, it holds that $\mathbb{E}[Y] = \mathbb{E}[\mathcal{M}^+(y)]$.

Since M is informationally complete we have

$$\mathbb{E}[\mathcal{M}^+(y)] = \mathcal{M}^+(\mathbb{E}[y]) = \mathcal{M}^+\,\mathcal{M}(\rho) = \rho\,, \tag{3.89}$$

i.e., the linear inversion estimator $\hat{\rho}_{\mathrm{LS}} = \mathcal{M}^+(y)$ is unbiased, which implies that $\mathbb{E}[Y] = \rho$. Together, we have an error term

$$\hat{\rho}_{\mathrm{LS}} - \rho = \sum_{k=1}^{n_\rho} \frac{1}{n_\rho} \left(Y_k - \mathbb{E}[Y_k]\right)\,. \tag{3.90}$$

Now we wish to apply the matrix Bernstein inequality from Theorem 3.24 with $X_k = \frac{1}{n_\rho}\left(Y_k - \mathbb{E}[Y_k]\right)$. So we calculate

$$\begin{aligned}
\frac{1}{n_\rho}\left\|Y - \mathbb{E}[Y]\right\|_{\mathrm{op}} &\leq \frac{1}{n_\rho} \max_{i\in[m]} \left\|(d+1)\,|\psi_i\rangle\langle\psi_i| - \mathbb{1} - \rho\right\|_{\mathrm{op}} \\
&\leq \frac{1}{n_\rho} \max_{i\in[m]} \Big\|\underbrace{d\,|\psi_i\rangle\langle\psi_i|}_{d\mathbb{1}\succeq\ldots\succeq 0} - \underbrace{\left((\mathbb{1} - |\psi_i\rangle\langle\psi_i|) + \rho\right)}_{2\mathbb{1}\succeq\ldots\succeq 0}\Big\|_{\mathrm{op}} \\
&\leq \frac{d}{n_\rho} =: a\,,
\end{aligned} \tag{3.91}$$

since we have implicitly assumed $d \geq 2$. Moreover,

$$\mathbb{E}[Y^2] - \mathbb{E}[Y]^2 = \sum_{i=1}^{m} \frac{d}{m} \langle\psi_i|\,\rho\,|\psi_i\rangle \left((d+1)\,|\psi_i\rangle\langle\psi_i| - \mathbb{1}\right)^2 - \rho^2\,. \tag{3.92}$$

We note that $(d+1)^2 - 2(d+1) = d^2 - 1$, remember the frame operator (3.44), and und use Lemma 3.19 to obtain

$$\begin{aligned}
\mathbb{E}[(Y - \mathbb{E}[Y])^2] = \mathbb{E}[Y^2] - \mathbb{E}[Y]^2 &= \sum_{i=1}^{m} \mathbb{P}[i]\left((d+1)\,|\psi_i\rangle\langle\psi_i| - \mathbb{1}\right)^2 - \rho^2 \\
&= (d^2-1)\sum_{i=1}^{m} \frac{d}{m}\,\langle\psi_i|\,\rho\,|\psi_i\rangle\,|\psi_i\rangle\langle\psi_i| + \mathbb{1} - \rho^2 \\
&= (d^2-1)\frac{m}{d} S_{\mathsf{M}}(\rho) + \mathbb{1} - \rho^2 \\
&= (d^2-1)\left(\frac{1}{d+1}\rho + \frac{d}{d+1}\frac{\mathbb{1}}{d}\right) + \mathbb{1} - \rho^2 \\
&= (d-1)\rho + d\,\mathbb{1} - \rho^2\,.
\end{aligned} \tag{3.93}$$

This leads to

$$\begin{aligned}
\sigma^2 &:= \left\|\mathbb{E}\Big[\sum_{k=1}^{n_\rho}\left(\frac{1}{n_\rho}(Y_k - \mathbb{E}[Y_k])\right)^2\Big]\right\|_{\mathrm{op}} \\
&= \left\|\frac{1}{n_\rho}\sum_{k=1}^{n_\rho}\left(\mathbb{E}[Y_k^2] - \mathbb{E}[Y_k]^2\right)\right\|_{\mathrm{op}} \\
&= \frac{1}{n_\rho}\left\|(d-1)\rho + d\,\mathbb{1} - \rho^2\right\|_{\mathrm{op}} \\
&\leq \frac{2d-1}{n_\rho}\,.
\end{aligned} \tag{3.94}$$

We note that for $\tau \in [0,2]$ we have

$$\frac{\tau^2/2}{\sigma^2 + a\tau/3} \geq \frac{n_\rho\tau^2/2}{2d - 1 + 2d/3} \geq \frac{3\tau^2 n_\rho}{16d} \tag{3.95}$$

and an application of Theorem 3.24 with $X_k = \frac{1}{n_\rho}(Y_k - \mathbb{E}[Y_k])$ yields (3.84). □

Lemma 3.23 tells us that a number of copies of

$$n_\rho \geq \frac{16}{3} \frac{d \ln(d/\delta)}{\tau^2} \tag{3.96}$$

is sufficient to reconstruct any quantum state with error bounded by τ in spectral norm with probability at least $1 - \delta$. However, the distinguishability of quantum states is given by the trace norm, see Proposition 4.8.

The bound (2.7) implies that a the reconstruction error is bounded in trace norm as

$$\|\hat{\rho}_{\mathrm{LS}} - \rho\|_1 \leq \frac{\tau}{\mathrm{rank}(\hat{\rho}_{\mathrm{LS}} - \rho)} =: \varepsilon\,. \tag{3.97}$$

The estimate $\hat{\rho}_{\mathrm{LS}}$ carries a reconstruction error due to the statistical estimation error. It can be expected that the reconstruction error is quite isotropically distributed in $\mathrm{Herm}(\mathbb{C}^d)$. Hence, one would expect that $\mathrm{rank}(\hat{\rho}_{\mathrm{LS}} - \rho) = d$ with high probability, since low rank matrices are a zero set in $\mathrm{Herm}(\mathbb{C}^d)$. This only leads to a sample complexity of $n_\rho \in \tilde{\mathrm{O}}(d^3/\varepsilon^2)$.

Now let us assume that ρ is of low rank r. If the reconstruction $\hat{\rho}_{\mathrm{LS}}$ is close to ρ then $\hat{\rho}_{\mathrm{LS}}$ is approximately of low rank. For any operator $X \in \mathrm{L}(\mathbb{C}^d)$ we denote by X_r the best rank-r approximation of X in trace norm (and then all Schatten p-norms for $p \in [1, \infty)$). The error of this approximation is

$$\sigma_r(X) \coloneqq \min_{\mathrm{rank}(Z) \leq r} \|X - Z\|_1 \tag{3.98}$$

and the minimizer is approximation X_r. Note that $\sigma_r(X)$ and X_r can be calculated using a singular value decomposition.

The low rank of ρ cab be exploited as follows.

Lemma 3.25 (Approximate rank [37, Appendix, Section VIII]):

Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ and $\hat{\rho} \in \mathrm{Herm}(\mathbb{C}^d)$ with $\mathrm{Tr}[\hat{\rho}] = 1$ such that $\|\hat{\rho} - \rho\|_{\mathrm{op}} \leq \tau$ for some $\tau \geq 0$. Then the projection of ρ to the density matrices

$$\hat{\rho}_{\mathrm{P}} \coloneqq \operatorname*{arg\,min}_{\sigma \in \mathcal{S}(\mathbb{C}^d)} \|\sigma - \hat{\rho}\|_2 \tag{3.99}$$

satisfies

$$\|\hat{\rho}_{\mathrm{P}} - \hat{\rho}\|_1 \leq 4r\tau + 2\min\{\sigma_r(\rho), \sigma_r(\hat{\rho})\} \tag{3.100}$$

for all $r \in \mathbb{Z}_+$.

Proof. Proof sketch The proof works in two steps: (i) the threshold value x_0 in the density matrix projection (3.80) satisfies $x_0 \in [0, \tau]$. (ii) for quantum states $\rho_1, \rho_2 \in \mathcal{S}(\mathbb{C}^d)$

$$\|\rho_1 - \rho_2\|_1 \leq 2r \|\rho_1 - \rho_2\|_{\mathrm{op}} + 2\min\{\sigma_r(\rho_1), \sigma_r(\rho_2)\} \tag{3.101}$$

holds for all $r \in \mathbb{Z}_+$. These two statements are relatively straightforward to prove, see [37, Appendix, Section VIII]. □

This leads to the following refined version of the PLS guarantee (Theorem 3.21) for 2-design based POVMs.

Theorem 3.26 (PLS for 2-design based POVMs [37]):

Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be a state, fix a total number of 2-design based POVM measurements (3.41) of ρ to be n_ρ.

Then, for an $r \in \mathbb{Z}_+$ and $\varepsilon \in [0,1]$, the PLS estimator (3.76) satisfies

$$\mathbb{P}\Big[\|\hat{\rho}_{\mathrm{PLS}} - \rho\|_1 \geq \varepsilon + 2\,\sigma_r(\rho)\Big] \leq d \exp\left(-\frac{3\,n_\rho \varepsilon^2}{256\,d\,r^2}\right). \tag{3.102}$$

Proof. We apply Lemma 3.23 with $\tau = \frac{\varepsilon}{4r}$ to obtain

$$\|\hat{\rho}_{\mathrm{LS}} - \rho\|_{\mathrm{op}} \leq \frac{\varepsilon}{4r} \tag{3.103}$$

with probability at least $1 - d\exp\left(-\frac{3\varepsilon n_\rho}{256 d r^2}\right)$. For that "success case" Lemma 3.25 yields

$$\|\hat{\rho}_{\mathrm{PLS}} - \rho\|_1 \leq \varepsilon + 2\min\{\sigma_r(\hat{\rho}_{\mathrm{PLS}}), \sigma_r(\rho)\}\,. \tag{3.104}$$

□

We note that Theorem 3.26 exploits an approximate rank of the measured state. This means that for an increasing number of measurements n_ρ the PLS estimate of the measured state ρ estimates and increasing number of eigenvectors of ρ correctly.

3.8. Lower bounds

In the previous section we have discussed one concrete method for quantum state tomography. This results provide upper bounds on the sample complexity of quantum state tomography. In order to investigate the optimality of the tomographic methods also lower bounds are required.

Naturally, the derivation of lower bounds seems to be much less straightforward. The standard idea to obtain these bounds [63, Supplemental Material/Appendix] is the following. First, one constructs a so-called ϵ-packing net for the subset of states $\mathsf{P} \subset \mathcal{S}(\mathbb{C}^d)$ one is interest in (e.g., states with at most rank r). (Such packing nets are constructed using Levy's lemma.) An *ϵ-packing net* for P is a set of points $\{\rho_i\}_{i\in[s]} \subset \mathsf{P}$ such that (i) $\|\rho_i - \rho_j\|_1 > \epsilon$ for all $i \neq j$ and (ii) s is maximal. Generaically, s scales exponentially in the manifold dimension of P.

One can encode $\log_2(s)$ many bits with an integer $i \in [s]$. This can be used to send a message of $\log_2(s)$ many bits by sending the receiving party the corresponding quantum state ρ_i. Now if one would send instead the measurement outputs of quantum state tomography of ρ_i with trace norm error ϵ to the receiving party then one still needs at least $\log_2(s)$ many bits, i.e., the measurement output must contain at least this many bits. One can make this consideration rigorous by using information theory [64] (such as Holevo's bound, Fano's inequality, and the data processing inequality) to put a lower bound on the number of required measurements.

One way to capture that a number ℓ of single iid. measurements is insufficient for quantum tomography is via the *minimax risk*

$$R^*(\epsilon, \ell) := \inf_{\left(\hat{\rho}_\ell, \{\mathsf{M}^{(i)}\}_{i\in[\ell]}\right)} \sup_{\rho\in\mathsf{P}} \mathbb{P}[\|\hat{\rho}_n(y) - \rho\|_1 > \epsilon]\,, \tag{3.105}$$

where the infimum is taken over ℓ admissible measurements $\mathsf{M}^{(i)}$ (the POVM might be different in every trial) and all estimators $\hat{\rho}_\ell$ taking the measurement outputs $y \in \mathbb{R}^n$, where each y_i is an iid. sample taken from the measurement outcome probabilities $(\mathrm{Tr}[M_j^{(i)}\rho])_j$ [55]. This minimax risk is the failure probability of the best tomographic strategy for the worst possible state.

Theorem 3.27 (Pauli observable measurements [55, Theorem 6]):

Let the set of feasible measurements be the observable measurements with POVM elements (3.53) in $d = 2^n$ dimensions and let the set of possible states P be the subset of states of $\mathcal{S}(\mathbb{C}^d)$ with rank at most r. Moreover, fix $\epsilon \in (0, 1 - r/d)$ and $\delta \in [0, 1)$. Then $R^*(\epsilon, \ell) < \delta$ implies

$$\ell \in \Omega\left(\frac{r^2 d^2}{\ln(d)}\right), \tag{3.106}$$

where the implicit constants depend on δ and ϵ.

The theorem tells us that the scaling in r and d of the number of mesurements in Theorem 3.21(ii) is optimal up to $\ln(d)$-factors.

With similar ideas one can derive lower bounds for the case of a single large POVM acting on ℓ copies of the state, i.e., on $\rho^{\otimes \ell}$. Note that setting includes parallel and interactive measurements. One can prove a lower bound on the number of measurements ℓ required for tomography scaling as [19]

$$\ell \in \Omega\left(\frac{dr}{\epsilon^2} \frac{(1-\epsilon)^2}{\ln[d/(r\epsilon)]}\right), \tag{3.107}$$

where the implicit constant depends on a confidence parameter δ. This bound implies that also the scaling of the number of mesurements in Theorem 3.21(i) is optimal up to $\ln(d)$-factors. In particular, the global measurements on $\rho^{\otimes \ell}$ *do not* yield an improved scaling compared to sequential measurements.

3.9. Maximum likelihood estimation

In the maximum likelihood estimation (MLE) approach to quantum state tomography, the final state estimate $\hat{\rho}$ is the state which is most likely to reproduce the measurement statistics [65, 66].

Let N be the total number of measurements of a POVM $\{M_i\}_{i\in[m]}$ and let n_i be the number of occurrences for outcome i. We further define the measurement frequencies $f_i = n_i/N$ and the probabilities $p_i(\rho) = \mathrm{Tr}[\rho M_i]$. The state estimate $\hat{\rho}$ is then the state that maximizes the log-likelihood function given by

$$\mathcal{L}(\rho) \coloneqq \sum_{i=1}^{m} f_i \ln p_i(\rho), \tag{3.108}$$

subject to the constraints $\rho \succeq 0$ and $\mathrm{Tr}[\rho] = 1$.

We incorporate the constraint $\rho \succeq 0$ by parameterizing the state as $\rho = A^\dagger A$. The constraint $\mathrm{Tr}[\rho] = 1$ is effectively incorporated by introducing the Lagrange multiplier λ to regularize $\mathcal{L}$ towards a small trace $\mathrm{Tr}[\rho]$. The new objective function then reads as

$$F(A) \coloneqq \sum_{i=1}^{m} f_i \ln(\mathrm{Tr}[A^\dagger A M_i]) - \lambda \,\mathrm{Tr}[A^\dagger A]\,. \tag{3.109}$$

Exercise 3.7 (MLE):

1. A necessary condition for A being an extremum is that $F(A + \delta A) = F(A) + \mathrm{O}\big(\|\delta A\|_{\mathrm{op}}^2\big)$. Show that this condition implies

$$\sum_{i=1}^{m} \frac{f_i}{p_i(\rho)} M_i A^\dagger = \lambda A^\dagger. \tag{3.110}$$

Hint: *Use* $\ln(1+x) = x + \mathrm{O}(x^2)$ *for small* x.

2. We define the operators $R_\rho := \sum_{i=1}^m \frac{f_i}{p_i(\rho)} M_i$ for $i \in [m]$. Use Eq. (3.110) to show that $\lambda = 1$ and hence $R_\rho \rho = \rho$, as well as $R_\rho \rho R_\rho = \rho$.

The last result $R_\rho \rho R_\rho = \rho$ can be used to iteratively find the maximum via the update rule

$$\rho_{k+1} = \frac{1}{N} R_{\rho_k} \rho_k R_{\rho_k} \,. \tag{3.111}$$

3. Implement MLE numerically using the above iteration rule for the target state $|\psi\rangle$ and the Pauli basis measurements from Section 3.5.2.1. Plot the reconstruction error $\|\hat{\rho} - |\psi\rangle\langle\psi|\|_1$ and the log likelihood function for $m \in [100]$.

4. Repeat MLE for measurements with added Gaussian noise and plot the reconstruction error over the noise strength $\eta \in [0, 0.1]$ for $m = 100$.

5. Repeat this exercise for linear inversion (Section 3.2), compressed sensing based reconstructions (Section 3.6), and projected least squares (Section 3.7). Compare the results.

This exercise shows that the numerical implementation of MLE can have a relatively slow convergence. However, there are much faster implementations [67] based on more elaborate gradient ascent methods.

For an experimental comparison of MLE, least squares estimation (a different version than considered here), and linear inversion see the work by Schwemmer et al. [68]. They also show the following negative result.

Proposition 3.28 (Biases in QST [68]):

A reconstruction scheme for quantum state tomography that always outputs a density operators is biased.

Hence, quantum state tomography methods using the density matrix structure allow for trace norm error guarantees with better scalings in the dimension (see Section 3.7) but come at the cost of being biased. Projected least squares estimation [37] provides a way to obtain an unbiased estimate of a quantum state so that its nearest density operator has a close to optimal trace norm error bound.

3.10. Confidence regions (additional information)

Confidence polytomes [69], Bayesian region estimates [70] , particle filtering from `QInfer` [71]

3.11. Other methods (additional information)

Adaptive quantum tomography [72]

4. Quantum state certification

The new material on frame theory (Section 3.3, in particular Proposition 3.9) and Monte Carlo integration (Section 2.4) were also covered in Lecture 8. Also Chebyshev's and Höffding's inequality from Section 2.3 will be used in Lecture 9.] Lecture

8 State certification is the task of making sure that a quantum state prepared in an experiment σ is a sufficiently good approximation of a target state ρ. More precisely, we make the following definitions.

Definition 4.1 (Quantum state certification):

Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be a quantum state (target). A *(quantum state) validation test* or $\mathcal{T}_{n_\sigma}$ of ρ consists of a quantum measurement on $\boldsymbol{\sigma} \in \mathcal{S}((\mathbb{C}^d)^{\otimes n_\sigma})$ followed by classical post-processing outputting either `"accept"` or `"reject"` and satisfying the *completeness* condition

$$\boldsymbol{\sigma} = \rho^{\otimes n_\sigma} \;\Rightarrow\; \mathbb{P}[\texttt{"accept"}] \geq \frac{2}{3}\,. \tag{4.1}$$

Let $\epsilon > 0$ (tolerated error) and dist be a distance measure on $\mathcal{S}(\mathbb{C}^d)$. We call validation test $\mathcal{T}_{n_\sigma}$ an *ϵ-certification test of ρ w.r.t.* dist *from n_σ independent samples* if the following *soundness* condition is satisfied: for any sequence of prepared states $\sigma_1, \ldots, \sigma_{n_\sigma} \in \mathcal{S}(\mathbb{C}^d)$

$$\mathrm{dist}(\rho, \sigma_i) > \epsilon \;\forall i \in [n_\sigma] \;\Rightarrow\; \mathbb{P}[\texttt{"reject"}] \geq \frac{2}{3}\,. \tag{4.2}$$

An *ϵ-certification test of ρ w.r.t.* dist *from n_σ iid. samples* is defined similarly with the σ_i being all the same state.
The *sample complexity* of a family of any such a test $\{\mathcal{T}_{n_\sigma}\}$ is (the scaling of) n_σ with d and ϵ.

In *hypothesis testing* one has a *null hypothesis* H_0 (usually the one one hopes to disprove) and an *alternative hypothesis* H_1 and one needs to figure out which is true based on statistical data. In this setting, there are two types of error,

$$\mathbb{P}[\text{accept } H_1 | H_0] \quad \text{(type I error)} \tag{4.3}$$
$$\mathbb{P}[\text{accept } H_0 | H_1] \quad \text{(type II error)}\,. \tag{4.4}$$

In state certification we choose null hypothesis H_0 to be $\|\sigma - \rho\|_1 > \epsilon$ and $\sigma = \rho$ to be the alternative hypothesis H_1. Then $\mathbb{P}[\texttt{"reject"} | \sigma = \rho]$ is the type II error and $\mathbb{P}[\texttt{"accept"} | \, \|\sigma - \rho\|_1 > \epsilon]$ the type I error. So, (4.2) is the requirement that the type I error is bounded by 1/3 and (4.1) corresponds to the type II error being also bounded by 1/3.

A validation test just checks for some errors, which we do not want to specify in the general definition above. However, any meaningful validation test needs to accept the target state with some confidence. As an example for validation, it is common to check for particle loss in experiments with photons.

Often, especially in the computer science community, *certification* is also called *verification*. However, in particular from an epistemological point of view, a physical model or hypothesis can never be fully verified. Therefore, we will mainly stick to the term *certification* but might still use *verification* interchangeably.

Quantum state certification just outputs the minimal information of whether a certain target state has been prepared. It is not difficult to see that the confidence value of 2/3 can be amplified:

Exercise 4.1 (Confidence amplification):

Let $\mathcal{T}_{n_\rho}$ be an ϵ-certification test of a quantum state ρ from n_ρ iid. samples with maximum failure probability $\delta = \frac{1}{3}$. We repeat the certification test N times and obtain a new certification test by performing a majority vote on the outcomes. Show that the new test satisfies the *completeness* and *soundness*

conditions

$$\sigma = \rho \Rightarrow \mathbb{P}[\texttt{"accept"}] \geq 1 - \delta\,, \tag{4.5}$$
$$\mathrm{dist}(\rho, \sigma) > \epsilon \Rightarrow \mathbb{P}[\texttt{"reject"}] \geq 1 - \delta\,, \tag{4.6}$$

for all $\sigma \in \mathcal{S}(\mathbb{C}^d)$, where $\delta = \mathrm{e}^{-c\,N}$ and $c > 0$ is an absolute constant. The parameter $1 - \delta$ is also called the *confidence* of the test.

A certification test is only required to accept the target state. However, in practice, such test will accept states from some region around the target state with large probability. Such a property of a certification test is called *robustness (against deviations from the target states)*. One way of how such a robustness can be guaranteed is by estimating the distance of the targeted state ρ and the prepared state σ, as we will see in Section 4.1 on *fidelity estimation*, which bounds on the distance. In this way, one obtains more information (a distance) than just certification (just `"accept"` or `"reject"`).

Clearly, one can also certify through full quantum state tomography. However, the number of single sequential measurements in general required for tomography of a state $\sigma \in \mathcal{S}(\mathbb{C}^d)$ scales as $\Omega(d\,\mathrm{rank}(\rho))$ and as $\Omega(d^2\,\mathrm{rank}(\rho)^2)$ in the case tow-outcome Pauli string measurements [55]. So, for the relevant case of pure n-qubit states this number scales as least as 2^n. This measurement effort becomes unfeasible already for relatively moderate n.

We will see that fidelity estimation can work with dramatically fewer measurements than full tomography, when the target state has additional structure. In many situations, certification can work with even fewer measurements than fidelity estimation due to an improved ϵ-dependence in the sample complexity.

4.1. Direct fidelity estimation

The *fidelity* is a measure of *closeness* for two quantum states. In order to define it remember that the square root $\sqrt{\rho}$ of a positive semidefinite operator $\rho \in \mathrm{Pos}(\mathbb{C}^d)$ is defined by $(\sqrt{\rho})^2 = \rho$ and $\sqrt{\rho} \succeq 0$ and can be obtained through an eigenvalue decomposition. The fidelity of two quantum state $\rho, \sigma \in \mathcal{S}(\mathbb{C}^d)$ is defined as[1]

$$\mathrm{F}(\rho, \sigma) \coloneqq \left\|\sqrt{\rho}\sqrt{\sigma}\right\|_1^2 . \tag{4.7}$$

Note that

$$\left\|\sqrt{\rho}\sqrt{\sigma}\right\|_1 = \mathrm{Tr}\big[\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}}\big] \tag{4.8}$$

The closeness measured by the fidelity is equivalent to the trace norm distance (see Exercise 2.2.4) as captured by the Fuchs-van-de-Graaf inequalities [73, Theorem 1],

$$1 - \sqrt{\mathrm{F}(\rho, \sigma)} \leq \frac{1}{2}\left\|\rho - \sigma\right\|_1 \leq \sqrt{1 - \mathrm{F}(\rho, \sigma)}\,. \tag{4.9}$$

Moreover, the fidelity is symmetric, i.e., $\mathrm{F}(\rho, \sigma) = \mathrm{F}(\sigma, \rho)$ for all ρ, σ.

When at least one of the states ρ and σ is pure, say $\rho = |\psi\rangle\langle\psi|$ then

$$\mathrm{F}(\rho, \sigma) = \langle\psi|\,\sigma\,|\psi\rangle = \left\|\rho\sigma\right\|_1 , \tag{4.10}$$

which can easily be proven using (4.8). We will mostly encounter that case, i.e., where one of the states is pure.

Indeed, in *direct fidelity estimation* (DFE) [63, 74] one has a target state $\rho \in \mathcal{S}(\mathbb{C}^d)$ and assumes to be given iid. state preparations of some state $\sigma \in \mathcal{S}(\mathbb{C}^d)$. The goal is to estimate the fidelity $\mathrm{Tr}[\sigma\rho]$ for the case where the ρ is a pure state, i.e. of the

[1] Some authors define the fidelity just as $\left\|\sqrt{\rho}\sqrt{\sigma}\right\|_1$ (without the square).

form $\rho = |\psi\rangle\langle\psi|$. This estimation is then solved using Monte Carlo methods; see Section 2.4 for the relevant tools.

For *general* direct fidelity estimation we fix a finite tight frame $\{M_\lambda\}_{\lambda\in\Lambda} \subset \mathrm{Herm}(\mathbb{C}^d)$ with frame constant A; see Section 3.3 for an introduction to frame theory. By $\{\tilde{M}_\lambda\}_{\lambda\in\Lambda} \subset \mathrm{Herm}(\mathbb{C}^d)$ we denote its dual frame, which has $1/A$ as frame constant (Proposition 3.9). We define the maximum norm of the frame as $C \coloneqq \max_{\lambda\in\Lambda} \|M_\lambda\|_{\mathrm{op}}$ and observe that, due to Hölder's inequality (2.8),

$$|W_\sigma(\lambda_i)| \leq C\,. \tag{4.11}$$

Traditionally [63, 74], the frame (and also the dual frame) are taken to be the normalized n-qubit Pauli basis $\{2^{-n/2}\,\sigma_{s_1}\otimes\cdots\otimes\sigma_{s_n}\}_{s\in\{0,1,2,3\}^n}$, which are an orthonormal basis. But it has been proven to be useful to consider more general frames for quantum information tasks [75] and we will follow this trend. In general, Λ can be a continuous set but we assume it to be finite here.

Given any operator $\sigma \in \mathrm{Herm}(\mathbb{C}^d)$ we define its W-function (sometimes called *discrete Wigner function* or *quasi-probability distribution*) and $\tilde{W}$-function $W_\sigma, \tilde{W}_\sigma : \Lambda \to \mathbb{R}$ by

$$W_\sigma(\lambda) \coloneqq \mathrm{Tr}[M_\lambda\sigma]\,, \qquad \tilde{W}_\sigma(\lambda) \coloneqq \mathrm{Tr}[\tilde{M}_\lambda\sigma]\,. \tag{4.12}$$

This allows us to write (see the frame expansion (3.21))

$$\begin{aligned}\mathrm{Tr}[\rho\sigma] &= \mathrm{Tr}\left[\rho\sum_{\lambda\in\Lambda} W_\sigma(\lambda)\tilde{M}_\lambda\right] = \sum_{\lambda\in\Lambda}\mathrm{Tr}\big[\rho\tilde{M}_\lambda\big]W_\sigma(\lambda)\\ &= \sum_{\lambda\in\Lambda}\tilde{W}_\rho(\lambda)W_\sigma(\lambda)\end{aligned} \tag{4.13}$$

for $\rho,\sigma \in \mathrm{Herm}(\mathbb{C}^d)$.

Now, we will use *importance sampling* from *Monte Carlo integration* (see Section 2.4) to estimate the sum (4.13) for the case where the target state $\rho \in \mathcal{S}(\mathbb{C}^d)$ is a pure state and the prepared states $\sigma \in \mathcal{S}(\mathbb{C}^d)$ are arbitrary. For this purpose we rewrite the overlap (4.13) as

$$\mathrm{Tr}[\rho\sigma] = \sum_{\lambda\in\Lambda}\frac{W_\sigma(\lambda)}{A\,\tilde{W}_\rho(\lambda)}\,A\tilde{W}_\rho(\lambda)^2 \tag{4.14}$$

and define

$$q_\lambda \coloneqq A\tilde{W}_\rho(\lambda)^2 \tag{4.15}$$

as importance sampling distribution on the sampling space Λ, where $1/A$ is the frame constant of $\{\tilde{M}_\lambda\}$, which we now argue to be the right normalization constant. The tight frame condition for ρ can be written as

$$\sum_{\lambda\in\Lambda}\tilde{W}_\rho(\lambda)^2 = \sum_{\lambda\in\Lambda}|\langle\tilde{M}_\lambda,\rho\rangle|^2 = \frac{\langle\rho,\rho\rangle}{A}\,. \tag{4.16}$$

For ρ being a pure state, i.e., $\langle\rho,\rho\rangle = \mathrm{Tr}[\rho^2] = 1$ (see Exercise 2.3), we indeed obtain

$$\sum_{\lambda\in\Lambda} q_\lambda = 1\,. \tag{4.17}$$

Next, we consider the estimator based on samples $\lambda \sim q$ given by

$$X_\lambda \coloneqq \frac{W_\sigma(\lambda)}{A\,\tilde{W}_\rho(\lambda)}\,. \tag{4.18}$$

We will exploit that X_λ with $\lambda \sim q$ is an *unbiased estimator* of the fidelity:

$$\mathbb{E}_{\lambda\sim q}[X_\lambda] = \sum_{\lambda\in\Lambda} \frac{W_\sigma(\lambda)}{A\,\tilde{W}_\rho(\lambda)} q_\lambda = \sum_{\lambda\in\Lambda} \tilde{W}_\rho(\lambda) W_\sigma(\lambda) = \mathrm{Tr}[\rho\sigma]\,, \tag{4.19}$$

where the last identity is again Eq. (4.13). Next, we take the empirical estimate of X from ℓ samples,

$$Y \coloneqq \frac{1}{\ell}\sum_{i=1}^{\ell} X_{\lambda_i} \tag{4.20}$$

where $X^{(i)}$ are iid. drawn as (4.18). This is also and unbiased estimator of $\mathrm{Tr}[\rho\sigma]$, the precision of which we can control by increasing ℓ. In order to bound the confidence that we have an estimation error $|Y - \mathrm{Tr}[\rho\sigma]| \leq \epsilon$ for some desired $\epsilon > 0$ we need find a *maximum failure probability* δ so that the tail bound

$$\mathbb{P}[|Y - \mathrm{Tr}[\rho\sigma]| \geq \epsilon] \leq \delta \tag{4.21}$$

is satisfied for some $\epsilon, \delta > 0$ controlled by ℓ; see Figure 2.1 for the idea of tail bounds. Then we have an ϵ-good estimation of $\mathrm{Tr}[\rho\sigma]$ with *confidence* $1-\delta$.

However, we also need to take into account the error estimating X_λ from single measurements. This for this purpose we will use an estimator $\hat{Y}$ of the estimator Y, which uses finitely many measurements, and derive a tail bound of the form

$$\mathbb{P}\Big[\big|\hat{Y} - Y\big| \geq \epsilon\Big] \leq \delta \tag{4.22}$$

for $\epsilon, \delta > 0$. More precisely, we consider the following protocol:

Protocol 4.2 (Extension of DFE [63] to frames):

Let $\{M_\lambda\}_{\lambda\in\Lambda}$ be a finite tight frame for $\mathrm{Herm}(\mathbb{C}^d)$ satisfying $\|M_\lambda\|_{\mathrm{op}} \leq C$ for all $\lambda \in \Lambda$ and some constant C. Denote by A the frame constant of $\{M_\lambda\}$ and by $\{\tilde{M}_\lambda\}_{\lambda\in\Lambda}$ the canonical dual frame.
Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be a target state with respect to which we wish to estimate the fidelity from measurements of the observables $\{M_\lambda\}$ and let $\epsilon > 0$ and $\delta > 0$ be the parameters for the desired estimation accuracy and maximum failure probability.
The protocol consists of the following steps applied to n_σ iid. copies of a prepared state $\sigma \in \mathcal{S}(\mathbb{C}^d)$:

(i) Take iid. samples $\lambda_1, \ldots, \lambda_\ell \sim q$ from the importance sampling distribution (4.15), where $\ell \coloneqq \lceil \frac{1}{\epsilon^2\delta} \rceil$ (or as (4.44) for well-conditioned states).

(ii) Measure each M_{λ_i} a number of m_i times for $i \in [\ell]$ with m_i be chosen as

$$m_i \coloneqq \left\lceil \frac{2\,C^2}{\ell\, A^2 \tilde{W}_\rho(\lambda_i)^2 \epsilon^2} \ln(2/\delta) \right\rceil \tag{4.23}$$

(or as $m_i = 1$ for well-conditioned states).

(iii) Take the empirical estimates $\hat{W}_\sigma(\lambda_i)$ of $W_\sigma(\lambda_i) = \mathrm{Tr}[M_{\lambda_i}\sigma]$ from these measurements for $i \in [\ell]$.

(iv) Estimate $X_{\lambda_i} = \frac{W_\sigma(\lambda_i)}{A\tilde{W}_\rho(\lambda_i)}$ by $\hat{X}_{\lambda_i} \coloneqq \frac{\hat{W}_\sigma(\lambda_i)}{A\tilde{W}_\rho(\lambda_i)}$.

(v) Estimate $Y = \frac{1}{\ell}\sum_{i=1}^{\ell} X_{\lambda_i}$ by $\hat{Y} \coloneqq \frac{1}{\ell}\sum_{i=1}^{\ell} \hat{X}_{\lambda_i}$.

(vi) Output $\hat{Y}$ as estimate of $\mathrm{Tr}[\rho\sigma]$.

Theorem 4.3 (Guarantee for DEF, frame version of [63]):

The fidelity estimate $\hat{Y}$ from Protocol 4.2 satisfies

$$\mathbb{P}[|\hat{Y} - \mathrm{F}(\rho,\sigma)| \leq 2\epsilon] \geq 1 - 2\delta\,. \tag{4.24}$$

The expected required number of copies of σ is

$$\mathbb{E}[n_\sigma] = \mathbb{E}\sum_{i=1}^{\ell} m_i \leq 1 + \frac{1}{\epsilon^2\delta} + \frac{2\,C^2|\Lambda|}{A\,\epsilon^2}\ln(2/\delta)\,. \tag{4.25}$$

We note that a rescaling $M_\lambda \mapsto \nu M_\lambda$ of the frame changes the constants as $A \mapsto \nu^2 A$ and $C \mapsto \nu C$. So, the sample complexity is indeed invariant under such a rescaling. The constant C^2/A can be seen as an *incoherence* parameter [56] of the frame. Indeed, if the frame is an orthonormal basis then A/C^2 measures a maximum effective rank of its elements, see the norm inequalities (2.7).

Proof of Theorem 4.3. We start to bound the estimation error arising by taking the empirical average in Step (v). We note that X_λ defined in (4.18) is an unbounded random variable in general, as $\tilde{W}_\rho(\lambda)$ can be arbitrarily small. Hence, we will use Chebyshev's inequality (2.46) to make the tail bound (4.21) explicit. Using the definitions (4.15) and (4.18) of q and X and that X is the unbiased estimator (4.19), the variance of X becomes

$$\begin{aligned}
\mathrm{Var}_{\lambda\sim q}[X_\lambda] &= \mathbb{E}_{\lambda\sim q}[X_\lambda^2] - \mathrm{Tr}[\rho\sigma]^2 \\
&= \sum_{\lambda\in\Lambda} \frac{W_\sigma(\lambda)^2}{A^2\tilde{W}_\rho(\lambda)^2}\, A\tilde{W}_\rho(\lambda)^2 - \mathrm{Tr}[\rho\sigma]^2 \\
&= \frac{1}{A}\sum_{\lambda\in\Lambda} W_\sigma(\lambda)^2 - \mathrm{Tr}[\rho\sigma]^2 \\
&= \mathrm{Tr}[\sigma^2] - \mathrm{Tr}[\rho\sigma]^2\,,
\end{aligned} \tag{4.26}$$

where we have used the frame condition ($A = B$ in (3.16)) in the last step. Hence,

$$\mathrm{Var}_{\lambda\sim q}[X_\lambda] \leq \mathrm{Tr}[\sigma^2] \leq 1\,. \tag{4.27}$$

Using the basic insight of Monte Carlo estimation (2.53), we obtain

$$\mathrm{Var}_q[Y] = \mathbb{E}_q[(Y - \mathrm{Tr}[\rho\sigma])^2] \leq 1/\ell\,. \tag{4.28}$$

As Y is an unbiased estimator of $\mathrm{Tr}[\rho\sigma]$, i.e., $\mathbb{E}_q[Y - \mathrm{Tr}[\rho\sigma]] = 0$, we can directly apply Chebyshev's inequality (2.46) to arrive at

$$\mathbb{P}\big[\big|Y - \mathrm{Tr}[\rho\sigma]\big| \geq \epsilon\big] \leq \frac{1}{\epsilon^2\ell} \tag{4.29}$$

for any $\epsilon > 0$. Hence, for any $\delta > 0$ and

$$\ell \geq \frac{1}{\epsilon^2\delta} \tag{4.30}$$

this *failure probability* is bounded by δ, i.e., the tail bound (4.21) is satisfied.

Now we bound the statistical error that arises from the estimation of X_{λ_i} from measurement measurement setup $i \in [\ell]$ in Step (iii). For this purpose we write for each λ the eigendecomposition of M_λ as

$$M_\lambda = \sum_\alpha a_{\lambda,\alpha} P_{\lambda,\alpha} \tag{4.31}$$

with $\{P_{\lambda,\alpha}\}$ being the projector onto the eigenspaces and $\{a_{\lambda,\alpha}\} \subset \left[-\|M_\lambda\|_{\text{op}}, \|M_\lambda\|_{\text{op}}\right]$ the eigenvalues of M_λ. We note that the expected measurement outcome is

$$\mathbb{E}[a_{\lambda,\alpha}] = \text{Tr}[M_\lambda \sigma] = W_\sigma(\lambda)\,. \tag{4.32}$$

Denoting by a_{λ_j,α_j} the measurement outcome for measurement $j \in [m_i]$ and consider the following corresponding empirical estimate of X_{λ_i} (see (4.18))

$$\hat{X}_{\lambda_i} \coloneqq \frac{1}{m_i A \tilde{W}_\rho(\lambda_i)} \sum_{j=1}^{m_i} a_{\lambda_i,\alpha_j}\,. \tag{4.33}$$

Then the estimation error of the empirical estimator $\hat{Y}$ of Y from (4.20) becomes

$$\begin{aligned} \hat{Y} - Y &= \frac{1}{\ell} \sum_{i=1}^{\ell} (\hat{X}_{\lambda_i} - X_{\lambda_i}) \\ &= \frac{1}{\ell} \sum_{i=1}^{\ell} \sum_{j=1}^{m_i} \frac{1}{m_i A \tilde{W}_\rho(\lambda_i)} \left(a_{\lambda_i,\alpha_j} - W_\sigma(\lambda_i)\right)\,. \end{aligned} \tag{4.34}$$

Using the bound (4.11) and Höffding's inequality (2.47) with $t = \epsilon\ell$ and

$$b_{i,\lambda_j} = -a_{i,\lambda_j} = \frac{C}{m_i A \tilde{W}_\rho(\lambda_i)} \tag{4.35}$$

we find that (w.l.o.g. we assume that there are no i with $\tilde{W}_\rho(\lambda_i) = 0$)

$$\mathbb{P}[|\hat{Y} - Y| \geq \epsilon] \leq 2 \exp\left(\frac{-\epsilon^2}{\frac{1}{\ell}\sum_{i=1}^{\ell} \frac{2\,C^2}{\ell\, m_i A^2 \tilde{W}_\rho(\lambda_i)^2}} \right). \tag{4.36}$$

We wish that the tail bound (4.22) holds. Therefore, we impose the RHS of (4.36) to be bounded by δ, which is equivalent to

$$\ln(2/\delta) \leq \frac{\epsilon^2}{\frac{1}{\ell}\sum_{i=1}^{\ell} \frac{2\,C^2}{\ell\, m_i A^2 \tilde{W}_\rho(\lambda_i)^2}}\,. \tag{4.37}$$

The choice of m_i as in (4.23) guarantees that this bound it always satisfied, i.e., that the desired tail bound (4.22) holds. Then combination of the tails bounds (4.21) and (4.22) with the union bound (2.12) proves the confidence statement (4.24).

In order to calculate the final sample complexity (4.25) note that m_i is a random variable itself, since $\tilde{W}_\rho(\lambda_i)$ was randomly chosen. By the definition of the sampling (4.15), we have

$$\begin{aligned} \mathbb{E}[m_i] &= \sum_{\lambda_i \in \Lambda} m_i q_{\lambda_i} \\ &\leq 1 + \frac{2\,C^2 |\Lambda|}{\ell\, A\, \epsilon^2} \ln(2/\delta)\,, \end{aligned} \tag{4.38}$$

where the $+1$ comes from the ceiling in (4.23). Using the bound (4.30) on ℓ, the expected total number of measurements, the expected sample complexity, is

$$\mathbb{E} \sum_{i=1}^{\ell} m_i \leq 1 + \frac{1}{\epsilon^2 \delta} + \frac{2\,C^2 |\Lambda|}{A\,\epsilon^2} \ln(2/\delta)\,. \tag{4.39}$$

□

Example 4.4 (Paul measurements):

Let the frame for n qubits be given as

$$\{M_s\} \coloneqq \{\sigma_{s_1} \otimes \cdots \otimes \sigma_{s_n}\}_{s \in \{0,1,2,3\}^n} \tag{4.40}$$

and denote the corresponding frame operator by S. Then

$$d = 2^n\,, \quad |\Lambda| = d^2\,, \quad C = 1\,, \quad \tilde{M}_s = \frac{1}{d} M_s\,, \quad A = d\,, \tag{4.41}$$

where the last two identities can be obtained from the requirements $S(\tilde{M}_s) = M_s$. and $\sum_s |\operatorname{Tr}[M_s X]|^2 = d\,\|X\|_2^2$ for any operator X.
The sample complexity (4.25) hence becomes

$$\mathbb{E}[n_\sigma] \leq 1 + \frac{1}{\epsilon^2 \delta} + \frac{2\,d}{\epsilon^2} \ln(2/\delta)\,, \tag{4.42}$$

which is consistent with its original version [63]. □

Note that the sample complexity scales linearly in the Hilbert space dimension for the case of Pauli measurements. In contrast, the number of Pauli measurements required for state tomography scales as $\tilde{\Omega}(d^2 \operatorname{rank}(\sigma)^2)$ [55].

The main contribution to the number of measurements in the derivation of the sample complexity above can be trace back to the application of Chebyshev's inequality in (4.29). This step can, however, be improved for the following class of states.

Definition 4.5 (Well-conditioned states):

Let $\{\tilde{M}_\lambda\}_{\lambda \in \Lambda} \subset \operatorname{Herm}(\mathbb{C}^d)$ be a frame. Then we call an operator $\rho \in \operatorname{Herm}(\mathbb{C}^d)$ *well-conditioned with parameter* $\tilde{\alpha} > 0$ if for each $\lambda \in \Lambda$ either $\operatorname{Tr}[\tilde{M}_\lambda \rho] \geq \tilde{\alpha}$ or $\operatorname{Tr}[\tilde{M}_\lambda \rho] = 0$.

For example, if the frame $\{\tilde{M}_\lambda\}$ is the dual frame of the Pauli stings, $\{\tilde{M}_s\} = \{2^{-n}\,\sigma_{s_1} \otimes \cdots \otimes \sigma_{s_n}\}$, then each stabilizer state ρ (3.33) with stabilizer $\mathcal{S}$ on n-qubits is well conditioned with parameter $\tilde{\alpha} = 2^{-n} \equiv 1/d$, as

$$\operatorname{Tr}[\tilde{M}_s \rho] = 2^{-n} \sum_{S \in \mathcal{S}} \operatorname{Tr}[\tilde{M}_s S] = 2^{-n}\,\delta_{2^n \tilde{M}_s \in \mathcal{S}}\ \operatorname{Tr}[\tilde{M}_s M_s] = 2^{-n}\,\delta_{M_s \in \mathcal{S}}\,. \tag{4.43}$$

Theorem 4.6 (DFE for well-conditioned states):

Let $\rho \in \operatorname{Herm}(\mathbb{C}^d)$ with $\|\rho\|_2 = 1$ be a target "state" that is well-conditioned with parameter $\tilde{\alpha} > 0$ w.r.t. the tight frame in Protocol 4.2 for fidelity estimation. Moreover, we consider the protocol modified by setting $m_i = 1$ for all $i \in [\ell]$ in Step (ii) and

$$\ell \coloneqq \left\lceil \frac{2\,C^2}{A^2 \tilde{\alpha}^2 \epsilon^2} \ln(2/\delta) \right\rceil \tag{4.44}$$

in Step (i). Then fidelity estimate $\hat{Y}$ from $n_\sigma = \ell$ iid. measurements satisfies

$$\mathbb{P}[|\hat{Y} - \mathrm{F}(\rho, \sigma)| \leq \epsilon] \geq 1 - \delta\,. \tag{4.45}$$

Proof. With probability one we have $\tilde{W}_\rho(\lambda_i) \geq \tilde{\alpha}$ for all $i \in [\ell]$. Moreover, $|\hat{W}_\sigma(\lambda_i)| \leq C$. The estimator from Step (iv) of Protocol 4.2 is bounded as

$$|X_{\lambda_i}| \leq \frac{C}{A\tilde{\alpha}} \tag{4.46}$$

with probability 1.

Hence, the estimator $\hat{Y}$ is also bounded as $|\hat{Y}| \leq \frac{C}{A\tilde{\alpha}}$ almost surely. Höffding's inequality (2.47) with $t = \epsilon/\ell$ yields

$$\mathbb{P}\Big[\big|\hat{Y} - \mathrm{Tr}[\rho\sigma]\big| \geq \epsilon\Big] \leq 2\,\exp\left(-\frac{\ell\, A^2\tilde{\alpha}^2\epsilon^2}{2\,C^2}\right). \tag{4.47}$$

Imposing

$$2\,\exp\left(-\frac{\ell\, A^2\tilde{\alpha}^2\epsilon^2}{2\,C^2}\right) \leq \delta \tag{4.48}$$

and solving for ℓ yields (4.44). □

This theorem tells us direct fidelity estimation has a smaller sampling complexity for well-conditioned states. For instance, well-conditioning in the Pauli basis leads to a constant sample complexity:

Example 4.7 (Pauli measurements and well-conditioned states):

Consider the Pauli observable measurements from Example 4.4 and a pure state ρ that is well-conditioned with parameter $\tilde{\alpha} = \alpha/d$, where $\alpha > 0$ is some constant. This well-conditioning is equivalent to

$$\mathrm{Tr}[\sigma_{s_1} \otimes \cdots \otimes \sigma_{s_n}\rho] \begin{cases} \geq \alpha \\ = 0 \end{cases} \quad \text{or} \qquad \forall s \in \{0,1,2,3\}^n\,. \tag{4.49}$$

Then the sample complexity (4.44) becomes

$$n_\sigma \leq 1 + \frac{2\ln(2/\delta)}{\alpha^2\epsilon^2}\,. \tag{4.50}$$

An example for well-conditioned states are stabilizer states, which are easy to see to be well-conditioned with $\alpha = 1$ using (3.33). □

Exercise 4.2 (Certification w.r.t. the trace distance via DFE):

Fix parameters $\tilde{\epsilon}, \epsilon, \delta > 0$ with $\tilde{\epsilon} \leq \frac{1}{2}\epsilon^2$. Let $\hat{Y}$ be the direct fidelity estimator of the fidelity $\mathrm{F}(\rho,\sigma)$ so that $|\hat{Y} - \mathrm{F}(\rho,\sigma)| \leq \tilde{\epsilon}$ with confidence $1-\delta$. We consider the protocol that *accepts* if $\hat{Y} \geq 1 - \tilde{\epsilon}$ and *rejects* otherwise. As distance we choose the trace distance $\mathrm{dist}_{\mathrm{Tr}}$ defined by $\mathrm{dist}_{\mathrm{Tr}}(\rho,\sigma) := \frac{1}{2}\,\|\rho - \sigma\|_1$.

- Show that this protocol is an ϵ-certification test w.r.t. the trace distance in the sense of Exercise 4.1, i.e., that the completeness and soundness conditions are satisfied with confidence $1-\delta$.
- What is the resulting sampling complexity of DFE in the well-conditioned setting from Example 4.7?
- Let $\epsilon' < \epsilon$. Turn this protocol into a *robust* (ϵ,ϵ')-certification test, i.e., into an ϵ-certification test that is guaranteed to accept all states within an ϵ'-trace norm ball around ρ with confidence $1-\delta$.

4.2. Direct quantum state certification

We start with considering the case of a single measurement. For this purpose, we remember the connection of the trace distance to the distinguishability of states by POVMs.

Proposition 4.8 (Operational interpretation of the trace distance):

Let $\rho, \sigma \in \mathcal{S}(\mathbb{C}^d)$ be states. Then

$$\frac{1}{2}\left\|\rho-\sigma\right\|_1 = \sup_{0\leq P\leq \mathbb{1}} \mathrm{Tr}[P(\rho-\sigma)] \tag{4.51}$$

and the supremum is attained for the projector $P^+ := \mathbf{1}_{\mathbb{R}_+}(\sigma-\rho)$ onto the positive part of $\rho-\sigma$.

Proof. First we show that the supremum is attained for P^+. The self-adjoint operator difference

$$\rho-\sigma = X^+ - X^- \tag{4.52}$$

has a *positive part* $X^+ \in \mathrm{Pos}(\mathbb{C}^d)$ and a *negative part* $X^- \in \mathrm{Pos}(\mathbb{C}^d)$. We note that $\|X^\pm\|_{\mathrm{op}} \leq 1$ and, since $\mathrm{Tr}[X^+ - X^-] = \mathrm{Tr}[\rho-\sigma] = 0$, we have $\mathrm{Tr}[X^+] = \mathrm{Tr}[X^-]$. Moreover, $\|\rho-\sigma\|_1 = \mathrm{Tr}[X^+] + \mathrm{Tr}[X^-]$. The last two statements together yield that the trace distance between the two states is

$$\frac{1}{2}\left\|\rho-\sigma\right\|_1 = \mathrm{Tr}[X^+] = \mathrm{Tr}[P^+(\rho-\sigma)]\,, \tag{4.53}$$

where P^+ is the orthogonal projector onto the support of X^+ and can be obtained as $P^+ = \mathbf{1}_{\mathbb{R}_+}(\rho-\sigma)$ using spectral calculus (2.3), where $\mathbf{1}_{\mathbb{R}_+}$ denotes the indicator function of $\mathbb{R}_+$, see (2.44).

In order to show that the supremum cannot become larger than the trace distance, we consider some operator P with $0 \leq P \leq \mathbb{1}$. Then, indeed,

$$\begin{aligned}\mathrm{Tr}[P(\rho-\sigma)] &= \mathrm{Tr}[PX^+] - \mathrm{Tr}[PX^-] \leq \mathrm{Tr}[PX^+] \\ &\leq \left\|X^+\right\|_1 = \frac{1}{2}\left\|\rho-\sigma\right\|_1\,,\end{aligned} \tag{4.54}$$

where we have used Hölder's inequality (2.8) and (4.53) in the last two steps. □

This proposition means that the trace distance of two states is given by the maximum distinguishability by binary POVM measurements $\{P, \mathbb{1}-P\}$. This distinguishability can be amplified by measuring iid. copies of a quantum state σ with $\{P, \mathbb{1}-P\}$. Next, we turn this basic insight into an ϵ-certification test of pure state

$$\rho = |\psi\rangle\langle\psi|\,, \tag{4.55}$$

where $|\psi\rangle \in \mathbb{C}^d$ is a state vector. However, for practical reasons we consider the *infidelity* $1-\mathrm{F}(\rho,\sigma)$ of states ρ and σ as distance measure instead of the trace distance. We remember, that the Fuchs-van-de-Graaf inequalities (4.9) relate these two distance measures.

We consider the POVM given by $P = |\psi\rangle\langle\psi|$. Then, for any σ we have

$$\mathrm{Tr}[P\sigma] = \mathrm{F}(\rho,\sigma)\,, \tag{4.56}$$

i.e., the probability to obtain the POVM measurement outcome corresponding to ρ is the fidelity of the two states. Next, we will consider several measurements with the same POVM in order to boost the probability to detect deviations for the form $\mathrm{F}(\rho,\sigma) < 1-\epsilon$ with some targeted confidence $1-\delta$.

In order to be able to capture a class of large measurement settings we consider POVM measurements $\{\Omega, \mathbb{1}-\Omega\}$ with $\mathrm{Tr}[\Omega\sigma] = 1$. Moreover, we consider the case of several independently prepared states and sequential measurements. We mostly follow a work by Pallister et al. [76].

Protocol 4.9 (Naive direct quantum state certification):

Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be a pure target state and $\Omega \in \mathrm{Pos}(\mathbb{C}^d)$ with $\|\Omega\|_{\mathrm{op}} \leq 1$. Denote by $\{\Omega, \mathbb{1} - \Omega\}$ the binary POVM given by Ω, call the outcome corresponding to Ω `"pass"` and the one of $\mathbb{1} - \Omega$ `"fail"`.
For state preparations $\sigma_1, \ldots, \sigma_{n_\sigma} \in \mathcal{S}(\mathbb{C}^d)$ the protocol consists of the following steps.

1: **for** $i = 1, \ldots, n_\sigma$ **do**
2: measure σ_i with $\{\Omega, \mathbb{1} - \Omega\}$
3: **if** the outcome is `"fail"` **then**:
4: output `"reject"`
5: end protocol (break)
6: **end if**
7: **end for**
8: output `"accept"`

Proposition 4.10 (Performance guarantee I):

Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be a pure target state, $\epsilon, \delta > 0$, and consider the distance measure on quantum states given by the infidelity defined by $1 - \mathrm{F}(\rho, \sigma)$. The test from Protocol 4.9 with $\Omega = \rho$ is an ϵ-certification test w.r.t. the infidelity from n_σ independent samples for

$$n_\sigma \geq \frac{\ln(1/\delta)}{\epsilon} \tag{4.57}$$

with confidence at least $1-\delta$ according to Definition 4.1. Moreover, the protocol accepts the target state ρ with probability one.

Proof. The probability of measurement outcome `"pass"` in step $i \in [n_\sigma]$ is

$$\mathbb{P}[\texttt{"pass"}|\sigma_i] = \mathrm{Tr}[\Omega\sigma_i] = \mathrm{Tr}[\rho\sigma_i] = \mathrm{F}(\rho, \sigma_i)\,. \tag{4.58}$$

Hence, the final probability that the protocol accepts is

$$\mathbb{P}[\texttt{"accept"}] = \prod_{i=1}^{n_\sigma} \mathrm{F}(\rho, \sigma_i)\,. \tag{4.59}$$

Clearly, if $\sigma_i = \rho$ for all $i \in [n_\sigma]$ then the protocol accepts almost surely. Now let us consider the case that the fidelity is small, i.e.,

$$\mathrm{F}(\rho, \sigma_i) = \langle\psi|\,\sigma_i\,|\psi\rangle \leq 1 - \epsilon \qquad \forall i \in [n_\sigma]. \tag{4.60}$$

Then the probability that the protocol wrongfully accepts is

$$\mathbb{P}[\texttt{"accept"}] \leq (1 - \epsilon)^{n_\sigma}\,. \tag{4.61}$$

Now we wish this probability (type II error) be bounded by $\delta > 0$, i.e.,

$$(1 - \epsilon)^{n_\sigma} \leq \delta\,. \tag{4.62}$$

This maximum type II error is achieved for

$$n_\sigma \geq \frac{\ln\left(\frac{1}{\delta}\right)}{\ln\left(\frac{1}{1-\epsilon}\right)}\,. \tag{4.63}$$

We note that for $\epsilon \in [0, a] \subset [0, 1)$ the following bounds hold

$$\epsilon \leq \ln\left(\frac{1}{1-\epsilon}\right) \leq \ln\left(\frac{1}{1-a}\right)\frac{\epsilon}{a}\,, \tag{4.64}$$

which can be seen by using the fact that $\epsilon \mapsto \ln\left(\frac{1}{1-\epsilon}\right)$ is smooth, has value 0 at 0, its first derivative is lower bounded by 1, and its second derivative is positive. So, the minimum n_σ is

$$n_\sigma \approx \epsilon \ln(1/\delta) \tag{4.65}$$

for small $\epsilon > 0$. Moreover, for any $n_\sigma \geq \frac{\ln(1/\delta)}{\epsilon}$ the required bound (4.62) is satisfied. □

Perhaps surprisingly, the sample complexity of this protocol does not depend on the physical system size at all. It has a zero type I error and one can control the type II error via the parameter δ.

However, it is generically not practical to implement the POVM. So, we follow Pallister et al. [76] and allow for more complicated strategies. Say, we have access to a set of POVM elements

$$\mathsf{M} \subset \{M \in \mathrm{Pos}(\mathbb{C}^d) : \ \|M\|_{\mathrm{op}} \leq 1\}\,, \tag{4.66}$$

where $\rho \in \mathcal{S}(\mathbb{C}^d)$ is the target state. As one can only make finitely many measurements, we assume that $|\mathsf{M}| < \infty$ in order to avoid technicalities. Then we pick POVM elements $P_j \in \mathsf{M}$ with some probability and consider the corresponding binary POVMs $\mathsf{M}_j \coloneqq \{P_j, \mathbb{1} - P_j\}$, where all P_j have output `"pass"` and $\mathbb{1} - P_j$ have output `"fail"`. Now we modify Protocol 4.9 by including this probabilistic measurement strategy.

Protocol 4.11 (Direct quantum state certification):

Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be a pure target state and (μ_j, P_j) a measurement strategy, i.e., μ a probability vector and $0 \leq P_j \leq \mathbb{1}$. For each POVM $\{P_j, \mathbb{1} - P_j\}$, call the outcome corresponding to P_j `"pass"` and the one of $\mathbb{1} - P_j$ `"fail"`.
For state preparations $\sigma_1, \ldots, \sigma_{n_\sigma} \in \mathcal{S}(\mathbb{C}^d)$ the protocol consists of the following steps.

1: **for** $i = 1, \ldots, n_\sigma$ **do**
2: Draw $j \sim \mu$ and measure σ_i with $\{P_j, \mathbb{1} - P_j\}$
3: **if** the outcome is `"fail"` **then**:
4: output `"reject"`
5: end protocol (break)
6: **end if**
7: **end for**
8: output `"accept"`

The measurement description $(\mu_j, P_j)_j$ is also called a *(measurement) strategy*. The resulting probability of measuring `"pass"` is

$$\mathbb{P}[\texttt{"pass"}] = \sum_j \mu_j \operatorname{Tr}[P_j \sigma] = \operatorname{Tr}[\Omega\sigma] \tag{4.67}$$

with

$$\Omega \coloneqq \sum_j \mu_j P_j \tag{4.68}$$

being the *effective measurement operator*.

Now we make the constraint that

$$\operatorname{Tr}[\Omega\rho] = 1\,, \tag{4.69}$$

i.e., that the there is no false reject of the target state ρ with probability one. In particular, this means that $\mathrm{Tr}[P_j\rho] = 1$ for each measurement setup j. This constraint still allows for optimal measurement strategies:

Proposition 4.12 ([76, Proposition 8]):

Let $\rho = |\psi\rangle\langle\psi|$ be a target state. Let $0 \leq \Omega' \leq \mathbb{1}$ be an effective measurement operator with $\mathrm{Tr}[\Omega'\rho] < 1$ so that Protocol 4.11 is an ϵ-certification test w.r.t. infidelity from ${n_\sigma}'$ iid. samples. Then there exists an effective measurement operator $0 \leq \Omega \leq \mathbb{1}$ with $\mathrm{Tr}[\Omega\rho] = 1$ so that Protocol 4.11 is an ϵ-certification test w.r.t. infidelity from n_σ iid. samples so that $n_\sigma \leq {n_\sigma}'$ holds for sufficiently small ϵ.

The proof of this statement is a consequence of the Chernoff-Stein lemma, which quantifies the asymptotic distinguishability of two distributions in terms of their relative entropy.

With the constraint (4.69) the only remaining hypothesis testing error is a *false acceptance*, which is the event where a state σ with $\mathrm{F}(\rho,\sigma) < 1-\epsilon$ is accepted. This event has the worst-case probability over all states σ given as

$$\mathbb{P}[\texttt{"pass"}] = \max_{\sigma\colon \mathrm{Tr}[\rho\sigma]\leq 1-\epsilon} \mathrm{Tr}[\Omega\sigma]\,. \tag{4.70}$$

The maximum is given as follows.

Lemma 4.13 ([76], [77, Suppl. material, Section I]):

Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be a pure state, $0 \leq \Omega \leq \mathbb{1}$, $\mathrm{Tr}[\rho\Omega] = 1$, and $\epsilon > 0$. Then

$$\max_{\sigma\colon \mathrm{Tr}[\rho\sigma]\leq 1-\epsilon} \mathrm{Tr}[\Omega\sigma] = 1 - \nu(\Omega)\epsilon\,, \tag{4.71}$$

where $\nu(\Omega) := 1 - \lambda_2(\Omega)$ is the spectral gap between the maximum eigenvalue 1 (corresponding to ρ) and the second largest eigenvalue $\lambda_2(\Omega)$ (among all d eigenvalues).

Proof. We note that $\mathrm{Tr}[\rho\Omega] = 1$ means that a state vector $|\psi\rangle$ with $\rho = |\psi\rangle\langle\psi|$ is an eigenvalue-1 eigenvector of Ω. Moreover, let us write Ω in spectral decomposition,

$$\Omega = \sum_{j=1}^{d} \lambda_j P_j \tag{4.72}$$

with $1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$ and $P_1 = \rho$. For the case $\lambda_2 = 1$ the choice $\sigma = P_2$ yields a maximum of 1 in the maximization. Let us now consider the case $\lambda_2 < 1$. Then for

$$\sigma = (1-\epsilon)\rho + \epsilon P_2 \tag{4.73}$$

we have

$$\begin{aligned} \mathrm{Tr}[\Omega\sigma] &= 1 - \epsilon\,\mathrm{Tr}[\Omega\rho] + \epsilon\,\mathrm{Tr}[\Omega P_2] \\ &= 1 - \epsilon + \epsilon\lambda_2 = 1 - (1-\lambda_2)\epsilon\,, \end{aligned} \tag{4.74}$$

i.e., the claimed maximum in (4.71) is attained for some feasible σ.

To show that the claimed maximum is actually the maximum we consider some state $\sigma \in \mathcal{S}(\mathbb{C}^d)$ with $\mathrm{Tr}[\rho\sigma] \leq 1-\epsilon$. We write σ as convex combination $\sigma = (1-\epsilon')\rho + \epsilon'\rho^\perp$

and observe that $\epsilon' \geq \epsilon$. Then

$$\begin{aligned}
\mathrm{Tr}[\Omega\sigma] &= \mathrm{Tr}[\rho\sigma] + \sum_{j=2}^{d} \lambda_j \,\mathrm{Tr}[P_j\sigma] \\
&\leq \mathrm{Tr}[\rho\sigma] + \lambda_2 \sum_{j=2}^{d} \mathrm{Tr}[P_j\sigma] \\
&= 1 - \epsilon' + \lambda_2\epsilon' \,\mathrm{Tr}\left[\sum_{j=2}^{d} P_j \rho^{\perp}\right] \\
&= 1 - \epsilon' + \lambda_2\epsilon' \,\mathrm{Tr}[\rho^{\perp}] \\
&= 1 - \epsilon' + \lambda_2\epsilon' = 1 - (1-\lambda_2)\epsilon' \\
&\leq 1 - (1-\lambda_2)\epsilon \,.
\end{aligned} \tag{4.75}$$

□

Given a measurement strategy with effective measurement operator Ω this lemma provides a closed form formula for the false acceptance probability (4.70). This allows us to state the following guarantee for Protocol 4.11.

Proposition 4.14 (Performance guarantee II [76]):

Let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be a pure target state and $\epsilon, \delta > 0$. We consider an effective measurement operator $0 \leq \Omega \leq 1$ with $\mathrm{Tr}[\Omega\rho] = 1$, which has a second largest eigenvalue $\lambda_2(\Omega) < 1$ (among the d eigenvalues).
Then the certification test from Protocol 4.11 is an ϵ-certification test w.r.t. the infidelity from n_σ independent samples for

$$n_\sigma \geq \frac{\ln(1/\delta)}{\epsilon\,(1-\lambda_2(\Omega))} \tag{4.76}$$

with confidence at least $1-\delta$. Moreover, the protocol accepts the target state ρ with probability one.

Proof. The proof is mostly analogous to the one of Proposition 4.10.

Thanks to Lemma 4.13, the probability of wrongfully accepting a state $\sigma \in \mathcal{S}(\mathbb{C}^d)$ with $\mathrm{F}(\rho, \sigma_i) \leq 1-\epsilon$ is bounded as

$$\mathbb{P}[\texttt{"pass"}|\sigma_i] \leq 1 - (1-\lambda_2(\Omega))\epsilon \,. \tag{4.77}$$

Hence, the probability that Protocol 4.9 accepts is bounded as

$$\mathbb{P}[\texttt{"accept"}] \leq (1-(1-\lambda_2(\Omega))\epsilon)^{n_\sigma} \,. \tag{4.78}$$

Imposing $(1-(1-\lambda_2(\Omega))\epsilon)^{n_\sigma} \leq \delta$ and solving for n_σ yields

$$n_\sigma \geq \frac{\ln(1/\delta)}{\ln\left(\frac{1}{1-(1-\lambda_2(\Omega))\epsilon}\right)} \tag{4.79}$$

and the bound (4.64) finishes the proof. □

This proposition tells us that as long as Ω has a constant gap between its largest and second larges eigenvalue the sample complexity of the certification protocol has the same scaling as the one where Ω is the target state itself. Now it depends on the physical situation of what feasible measurement strategies Ω are. Given a set M of feasible measurements we can single out an optimal strategy as follows.

Definition 4.15 (Minimax optimization):

Let ρ be a pure state and $\epsilon > 0$. Moreover, let us assume the we have access to a compact set of binary measurements given by the operators $\mathsf{M} \subset \{P : \ 0 \leq P \leq \mathbb{1}\,,\ \mathrm{Tr}[P\rho] = 1\}$.
Then the best strategy Ω for the worst case state preparation σ is

$$\min_{\Omega \in \mathrm{conv}(\mathsf{M})} \max_{\sigma:\, \mathrm{Tr}[\rho\sigma] \leq 1-\epsilon} \mathrm{Tr}[\Omega\sigma]\,. \tag{4.80}$$

This quantity is called *minimax value* and a strategy Ω where the minimum is attained is called *minimax optimal*.

Such minimax optimizations are common in game theory and risk analysis.

If there are no restrictions on the measurements of a pure target state ρ, i.e., $\mathsf{M} = \{P : \ 0 \leq P \leq \mathbb{1}\,,\ \mathrm{Tr}[P\rho] = 1\}$, then $\Omega = \rho$ is minimax optimal.

For number of settings with physically motivated measurement restrictions the minimax strategy, or at least one that is close to it, has been obtained. Popular instances include the following settings:

- Stabilizer states and two qubit states with single qubit measurements [76]
- Ground states of locally interacting Hamiltonians [78]
- Bipartite states [79, 80], qubit case in an LOCC setting [81]
- Hypergraph states [78] with improvements in efficiency by Zhu and Hayashi [77]
- Stabilizer states [76, 78]

Here, we only outline the example of stabilizer states in more detail. Remember the definition of stabilizer states from the box on STABs in Section 3.4.1.

Theorem 4.16 (Minimax optimal Pauli measurements for STABs [76]):

Let $|\psi\rangle$ we an n-qubit stabilizer state with stabilizer group $\mathcal{S} \subset \mathcal{P}_n$ with elements $\mathcal{S} = \{\mathbb{1} = S_0, S_1, \ldots, S_{2^n-1}\}$. For $i \in [2^n - 1]$ denote by $P_i := \frac{1}{2}(\mathbb{1} + S_i)$ the projector onto the positive eigenspace of S_i.
Then the minimax optimal measurement strategy for having Pauli observables $\mathcal{P}_n$ as accessible measurements (see Definition 4.15) is given by measuring S_i with probability $\frac{1}{2^n-1}$. The resulting effective measurement operator $\Omega = \frac{1}{2^n-1}\sum_{i=1}^{2^n-1} P_i$ satisfies $\Omega\,|\psi\rangle = |\psi\rangle$ and has the second largest eigenvalue

$$\lambda_2(\Omega) = \frac{2^{n-1}-1}{2^n-1}\,. \tag{4.81}$$

Proof. By Lemma 4.13, the minimax optimum is

$$\begin{aligned} \min_{\Omega \in \mathcal{X}} \max_{\sigma:\, \mathrm{Tr}[\rho\sigma] \leq 1-\epsilon} \mathrm{Tr}[\Omega\sigma] &= \min_{\Omega \in \mathcal{X}} (1 - \nu(\Omega)\epsilon) \\ &= 1 - \epsilon \max_{\Omega \in \mathcal{X}} \nu(\Omega)\,, \end{aligned} \tag{4.82}$$

where

$$\begin{aligned} \mathcal{X} &:= \{\Omega \in \mathrm{conv}(\mathcal{P}_n) : \ \Omega\,|\psi\rangle = |\psi\rangle\} \\ &= \mathrm{conv}(\mathcal{S})\,. \end{aligned} \tag{4.83}$$

We argue that the minimization over $\mathrm{conv}(\mathcal{S})$ can be replaced by a minimization over $\mathrm{conv}(\mathcal{S}')$ with $\mathcal{S}' := \mathcal{S} \setminus \{\mathbb{1}\}$. To see this, observe that if $\Omega = (1-\alpha)\Omega' + \alpha\mathbb{1}$ for

$\alpha \in [0,1]$ then $\nu(\Omega) \leq \nu(\Omega')$. Then minimax optimal measurement strategies are of the form

$$\Omega = \sum_{i=1}^{2^n-1} \mu_j P_i \,. \tag{4.84}$$

for a probability vector μ. We note that

$$\mathrm{Tr}[\Omega] = 2^{n-1} \tag{4.85}$$

since $\mathrm{Tr}[P_i] = 2^{n-1}$.

Next, since $|\psi\rangle$ is an eigenvalue-1 eigenvector of Ω, we have

$$\Omega = 1 \oplus \tilde{\Omega} \tag{4.86}$$

and, hence

$$\lambda_2(\Omega) = \left\|\tilde{\Omega}\right\|_{\mathrm{op}} . \tag{4.87}$$

Moreover, $\mathrm{Tr}[\tilde{\Omega}] = 2^{n-1} - 1$. The operator $\tilde{\Omega}$ with the minimal norm $\left\|\tilde{\Omega}\right\|_{\mathrm{op}}$ under this constraint is of the form $\tilde{\Omega} = a\mathbb{1}$ for $a > 0$. Taking the trace of that equality, solving for a and denoting the orthogonal projector of $|\psi\rangle\langle\psi|$ by $|\psi\rangle\langle\psi|^{\perp} \coloneqq \mathbb{1} - |\psi\rangle\langle\psi|$ yields

$$\Omega = |\psi\rangle\langle\psi| + \frac{2^{n-1}-1}{2^n-1} \; |\psi\rangle\langle\psi|^{\perp} \tag{4.88}$$

with

$$\lambda_2(\Omega) = \frac{2^{n-1}-1}{2^n-1} \,. \tag{4.89}$$

In order to finish the proof we show that $\Omega \in \mathrm{conv}(\mathcal{S})$, i.e., that this choice of Ω is indeed compatible with (4.84).

We write the stabilizer state $|\psi\rangle\langle\psi|$ as combination of the stabilizers (see (3.33)) and use that $S_j = 2P_j - \mathbb{1}$,

$$\begin{aligned} |\psi\rangle\langle\psi| &= \frac{1}{2^n}\left(\mathbb{1} + \sum_{j=1}^{2^n-1} S_j\right) \\ &= \frac{1}{2^n}\left(\mathbb{1} + 2\sum_{j=1}^{2^n-1} P_j - (2^n-1)\mathbb{1}\right) \\ &= \left(\frac{1}{2^{n-1}} - 1\right)\mathbb{1} + \frac{1}{2^{n-1}} \sum_{j=1}^{2^n-1} P_j \,. \end{aligned} \tag{4.90}$$

With $\mathbb{1} = |\psi\rangle\langle\psi| + |\psi\rangle\langle\psi|^{\perp}$ this implies

$$\sum_{j=1}^{2^n-1} P_j = (2^n-1)\; |\psi\rangle\langle\psi| + (2^{n-1}-1)\; |\psi\rangle\langle\psi|^{\perp} \tag{4.91}$$

and, hence

$$\frac{1}{2^n-1} \sum_{j=1}^{2^n-1} P_j = |\psi\rangle\langle\psi| + \frac{2^{n-1}-1}{2^n-1} \; |\psi\rangle\langle\psi|^{\perp} \,, \tag{4.92}$$

which is the Ω from (4.88) and also the measurement strategy from the theorem statement. □

Corollary 4.17 (Sampling complexity [76]):

Let us call the outcome corresponding to P_i `"pass"` and the one corresponding to $1 - P_i$ `"fail"`. Then Protocol 4.11 is an ϵ-certification test of ρ w.r.t. infidelity from n_σ independent samples for

$$n_\sigma \geq 2\,\frac{\ln(1/\delta)}{\epsilon} \tag{4.93}$$

with confidence $1 - \delta$. Moreover, ρ is accepted with probability 1.

Proof. According to Proposition 4.14 a number of measurements

$$n_\sigma \geq \frac{\ln(1/\delta)}{\epsilon\,\nu(\Omega)} \tag{4.94}$$

is sufficient, where

$$\begin{aligned} \nu(\Omega) &= 1 - \lambda_2(\Omega) \\ &= 1 - \frac{2^{n-1} - 1}{2^n - 1} \\ &= \frac{2^{n-1}}{2^n - 1}\,. \end{aligned} \tag{4.95}$$

This results in

$$n_\sigma \geq \frac{2^n - 1}{2^{n-1}}\,\frac{\ln(1/\delta)}{\epsilon}\,. \tag{4.96}$$

□

So, restricting from all measurements to Pauli measurements results in at most a constant overhead of 2, cp. Proposition 4.10. We note that only very few of the $2^n - 1$ non-trivial stabilizers of ρ are measured. More precisely, the measurements are the ones of randomly subsampled stabilizer observables.

4.3. Other works (additional information)

First version for the certification of ground states of locally interacting Hamiltonians by Cramer et al. [82]; extension by Hangleiter et al. [83], discussing also ground state enabling universal quantum computation. In this line of research, fidelity witnesses [83–85] can be used to measure and estimate on a fidelity lower bounds.

In fact, the work [82] solves the certification problem by an instance of *ansatz state tomography*. Here, the measured state is assumed to be of matrix product form an this form is efficiently reconstructed from measurement data. Similar ideas work for permutationally invariant states [86–88].

Kalev et al. [89] have extended arguments from direct fidelity estimation [63] on stabilizer states using Bernstein's inequality to give an quadratically improved ϵ scaling for small ϵ.

Global von Neumann measurements on multiple iid. copies of the prepared quantum state have also been considered [90] (even with mixed target states), which leads to a sample complexity scaling as $n_\sigma \in \mathrm{O}(d/\epsilon)$ a version of ϵ-certification of quantum states in $\mathcal{S}(\mathbb{C}^d)$.

There is a very helpful survey on *quantum property testing* [91], where several methdos and notions of certification are reviewed.

Part II

Quantum dynamics

Quantum stats can be fully reconstructed tomographically with an essentially optimal number of measurements [37, 55], see Chapter 3. In particular, the reconstruction error can be bounded in the operationally relevant norm, the trace norm. Similar results hold for quantum state certification, where one has implemented a targeted state and is tasked to certify the correctness of the implementation up to some trace norm error, see Chapter 4.

In Part II of these note, we aim to find similar results for quantum processes. In principle, one can use the Choi-Jamiołkowski isomorphism (see Section 5.3) to map a quantum process to a quantum state and apply the characterization and verification methods from Part I. However, this approach has drawbacks: (i) it might result in measurements that are practically infeasible and typically require maximally entangled states as available resource and (ii) the error is controlled in the "wrong" norm. In a similar way as the trace norm is an operationally motivated norm for quantum states (Proposition 4.8) the so-called *diamond norm* is an operationally motivated norm for quantum processes, see Section 5.5. There is even a third potential problem: To (partially) characterize a quantum process one needs to prepare quantum states, to evolve them under the process, and to measure the final states. In this task so-called state preparation and measurement (SPAM) errors can be a serious obstacle for reliable characterization. Therefore, the development of quantum characterization and verification methods for quantum processes requires a significant amount of extra work.

As we will see in Chapter 6, one is able to estimate a weaker error measure than the diamond norm efficiently using *randomized benchmarking*. This is typically done in a way that is robust against SPAM errors. In Chapter 7 we will discuss state-of-the-art methods for quantum process tomography. In particular, we will see that one can reconstruct the most relevant part of a quantum process (the unital part) using similar measurements as in randomized benchmarking. However, the error measure is again a weaker one than the diamond norm. In Chapter 8 we will discuss *gate set tomography*. Here, one reconstruct a full gate set from measurement data and is able to estimate gate errors in the diamond norm. A disadvantage of this method is that it comes at the expense of a large overhead in the measurement effort and the amount of classical post-processing. Further improving those methods and finding lower bounds on the required measurement and computational effort is still subject to ongoing research, which makes this field particularly exciting.

5. Preliminaries II

In this chapter we introduce some more preliminaries required to discuss the characterization and validation of quantum processes. Let us start with a proper introduction to quantum processes.

5.1. Quantum processes

A *quantum processes* is given by linear map taking density operators to density operators and satisfying certain properties. Therefore, we start with introducing some notation related to linear maps between operator spaces.

Let $\mathcal{H}, \mathcal{K}$ be Hilbert spaces.

- The vector space of linear maps from $\mathrm{L}(\mathcal{H})$ to $\mathrm{L}(\mathcal{K})$ is denoted by $\mathbb{L}(\mathcal{H}, \mathcal{K}) \coloneqq \mathrm{L}(\mathrm{L}(\mathcal{H}), \mathrm{L}(\mathcal{K}))$. We set $\mathbb{L}(\mathcal{H}) \coloneqq \mathbb{L}(\mathcal{H}, \mathcal{H})$ and denote the identity by $\mathrm{id}_{\mathcal{H}} \coloneqq \mathbb{1}_{\mathrm{L}(\mathcal{H})} \in \mathbb{L}(\mathcal{H})$. Often we just write id when it is clear from the context what $\mathcal{H}$ is.

- A map $\Phi \in \mathbb{L}(\mathcal{H}, \mathcal{K})$ is called *Hermiticity-preserving* if

$$\Phi(\mathrm{Herm}(\mathcal{H})) \subset \mathrm{Herm}(\mathcal{K}), \tag{5.1}$$

  positive if

$$\Phi(\mathrm{Pos}(\mathcal{H})) \subset \mathrm{Pos}(\mathcal{K}), \tag{5.2}$$

  and *trace-preserving* if

$$\mathrm{Tr}[\Phi(X)] = \mathrm{Tr}[X] \tag{5.3}$$

  for all $X \in \mathrm{L}(\mathcal{H})$. Note that positive maps are also Hermiticity-preserving.

  The map Φ is called *completely positive* (CP) if $\Phi \otimes \mathbb{1}_{\mathrm{L}(\mathcal{H}')}$ is positive for all Hilbert spaces $\mathcal{H}'$ with identity map $\mathbb{1}_{\mathrm{L}(\mathcal{H}')} \in \mathbb{L}(\mathcal{H}')$. The set of CP maps is denoted by $\mathrm{CP}(\mathcal{H}, \mathcal{K}) \subset \mathbb{L}(\mathcal{H}, \mathcal{K})$ and forms a convex cone. We set $\mathrm{CP}(\mathcal{H}) \coloneqq \mathrm{CP}(\mathcal{H}, \mathcal{H})$.

- A completely positive and trace preserving (CPT) map is also called a *quantum channel* or just *channel*. The subset of CPT maps is denoted by $\mathrm{CPT}(\mathcal{H}, \mathcal{K}) \subset \mathrm{CP}(\mathcal{H}, \mathcal{K})$ and forms a convex set. Again, we set $\mathrm{CPT}(\mathcal{H}) \coloneqq \mathrm{CPT}(\mathcal{H}, \mathcal{H})$.

- A map $\Phi \in \mathbb{L}(\mathcal{H}, \mathcal{K})$ is called *unital* if $\Phi(\mathbb{1}_{\mathcal{H}}) = \mathbb{1}_{\mathcal{K}}$. Note that Φ is trace-preserving iff $\Phi^\dagger$ is unital.

So, essentially, quantum channels are maps that take density matrices to density matrices even when applied to a part of a larger system. Usual unitary dynamics is of this form:

Example 5.1 (Unitary channels):

We use calligraphic letters to denote the adjoint representation $\mathcal{U} \in \mathbb{L}(\mathcal{H})$ of a unitary $U \in \mathrm{U}(\mathcal{H})$ given by

$$\mathcal{U}(X) \coloneqq UXU^\dagger . \tag{5.4}$$

These maps are quantum channels and are called *unitary (quantum) channels*.

Unitary channels are invertible and the inverses are again unitary channels.

5.2. Tensor networks

When dealing with higher order tensors, such as linear maps on operators, it is very useful to use tensor network diagrams. A *tensor network* is a set of tensors together with a contraction corresponding to pairs of indices where pairs of contracted indices need to have the same dimension. Tensor networks have diagrammatic representations, which allow to visually track index contractions rather than spelling out these contractions explicitly. Instances of tensor networks are, e.g., the workhorse of powerful simulation techniques for strongly correlated quantum systems [92]. In this course, however, we will only use comparable small tensor networks, i.e., with just a few tensors. But even in this case, tensor networks will help us to dramatically simplify a

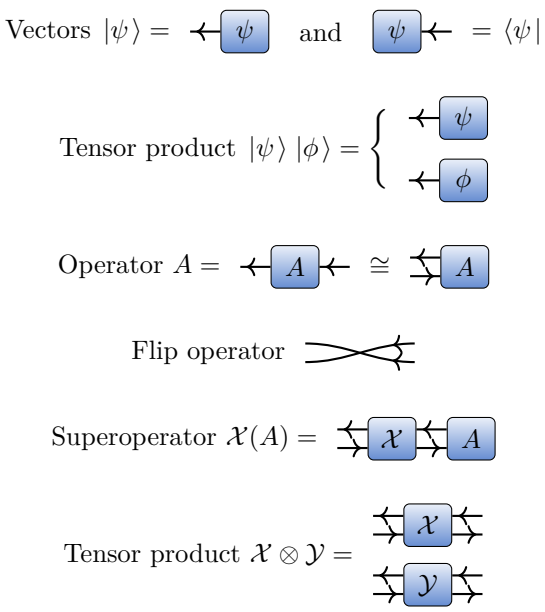

Figure 5.1.: Basic examples for tensor network diagrams: A vector $|\psi\rangle$ in some vector space V, vectorization of an operator A, a map $\mathcal{X} \in \mathbb{L}(V)$ applied to that vectorization, the non-vectorized version of the flip operator $\mathbb{F}$, and a tensor product of two maps on operators $\mathcal{X}$ and $\mathcal{Y}$.

number of calculations. For an introduction to tensor networks from a category theory point of view see the work by Biamonte et al. [93] and a follow-up work on open quantum systems [94].

In the diagrammatic representation tensors are denoted by boxes with one line for each index; Figure 5.1 for examples. A contraction of two indices is represented by connecting their representing lines. One can indicate whether an index corresponds to a vector space or a dual vector space (functionals) by arrows (outgoing or incoming) or the direction of the index (e.g. left/right). Tensor products of smaller tensors are represented simply by drawing the smaller tensors into the same diagram.

Exercise 5.1 (The swap-trick revisited):

Solve Exercise 2.4 using tensor network diagrams.

5.3. The Choi-Jamiołkowski isomorphism

The Choi-Jamiołkowski isomorphism [95, 96] provides a duality between CP maps and bipartite positive semidefinite operators and allows to identify channels with certain states. It has many applications in quantum information theory and related fields. In particular, it allows to practically check whether a given map is a quantum channel.

Throughout the whole section, let $\mathcal{H}$ and $\mathcal{K}$ be Hilbert spaces. For any vector space V, we have the natural isomorphism

$$\mathrm{L}(V) = V \otimes V^* \,, \tag{5.5}$$

where $V^* \coloneqq \mathrm{L}(V, \mathbb{C})$ is the dual space of V.

The *Choi-Jamiołkowski isomorphism*

$$\mathfrak{C} : \mathbb{L}(\mathcal{H}, \mathcal{K}) \to \mathrm{L}(\mathcal{K} \otimes \mathcal{H}) \tag{5.6}$$

is an isomorphism of vector spaces. Let $(|i\rangle)_{i\in[\dim(\mathcal{H})]}$ be a basis of $\mathcal{H}$ and let us denote complex conjugation in w.r.t. that basis by cc. Then the *Choi-Jamiołkowski*

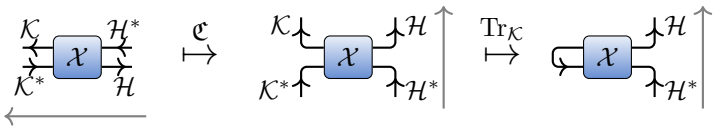

Figure 5.2.: The Choi-Jamiołkowski isomorphism and partial trace in terms of tensor network diagrams (explained in Figure 5.1).
Left: Order-4 tensor $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$ as a map from $\mathrm{L}(\mathcal{H}) \cong \mathcal{H} \otimes \mathcal{H}^*$ to $\mathrm{L}(\mathcal{K}) \cong \mathcal{K} \otimes \mathcal{K}^*$.
Middle: Its Choi-matrix $\mathfrak{C}(\mathcal{X})$ as an operator on $\mathcal{K} \otimes \mathcal{H}$.
Right: Partial trace $\mathrm{Tr}_1[\mathfrak{C}(\mathcal{X})]$ of the Choi matrix $\mathfrak{C}(\mathcal{X})$. This operator corresponds to the functional $\rho \mapsto \mathrm{Tr}[\mathcal{X}(\rho)]$.

isomorphism is given by the following identification:

$$\begin{aligned} \mathbb{L}(\mathcal{H}, \mathcal{K}) &= \mathrm{L}(\mathcal{K}) \otimes \mathrm{L}(\mathcal{H})^* = \mathcal{K} \otimes \mathcal{K}^* \otimes \mathcal{H}^* \otimes \mathcal{H} \\ &\cong \mathcal{K} \otimes \mathcal{H}^* \otimes \mathcal{K}^* \otimes \mathcal{H} = \mathrm{L}(\mathcal{K} \otimes \mathcal{H}^*) \\ &\overset{\mathrm{cc}}{\cong} \mathrm{L}(\mathcal{K} \otimes \mathcal{H})\,, \end{aligned} \tag{5.7}$$

where the natural isomorphism (5.5) is denoted by "=", the isomorphism of changing the order of the vector spaces by "≅", and the last one refers to the conjugate linear Hilbert space isomorphism $\mathcal{H} \cong \mathcal{H}^*$ composed with complex conjugation isomorphism; see Figure 5.2 for a tensor network representation of $\mathfrak{C}$. The Choi-Jamiołkowski isomorphism can be written explicitly. In terms of the unnormalized maximally entangled state

$$|\mathbb{1}\rangle = \sum_{i=1}^{\dim(\mathcal{H})} |i, i\rangle \ \in \mathcal{H} \otimes \mathcal{H} \tag{5.8}$$

the *Choi* matrix of $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$ is

$$\mathfrak{C}(\mathcal{X}) = \mathcal{X} \otimes \mathrm{id}(|\mathbb{1}\rangle\langle\mathbb{1}|)\,. \tag{5.9}$$

We note that $U \otimes U^{\intercal} |\mathbb{1}\rangle = |\mathbb{1}\rangle$ for all $U \in \mathrm{U}(\mathcal{H})$. Hence, the Choi-Jamiołkowski isomorphism is invariant under orthogonal basis changes $U \in \mathrm{O}(\mathcal{H})$, as the dual basis of $\mathcal{H}^*$ changes as $\langle i| \mapsto \langle i| U^{\dagger}$.

Exercise 5.2 (Choi-Jamiołkowski isomorphism):

Show that the characterizations of Choi-Jamiołkowski isomorphism from (5.9), (5.7), and Figure 5.2 coincide. Moreover, show that

$$\mathrm{Tr}[B\mathcal{X}(A)] = \mathrm{Tr}[(B \otimes A^{\intercal})\, \mathfrak{C}(\mathcal{X})] \tag{5.10}$$

for all $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$, $A \in \mathrm{L}(\mathcal{H})$ and $B \in \mathrm{L}(\mathcal{K})$.

Now we can connect the Choi-Jamiołkowski isomorphism to the properties of quantum channels.

Theorem 5.2 (CPT conditions):

For any map $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$ the following equivalences hold:

(i) $\mathcal{X}$ is trace-preserving iff $\mathrm{Tr}_{\mathcal{K}}[\mathfrak{C}(\mathcal{X})] = \mathbb{1}$.

(ii) $\mathcal{X}$ is Hermiticity-preserving iff $\mathfrak{C}(\mathcal{X})$ is Hermitian.

(iii) $\mathcal{X}$ is completely positive iff $\mathcal{X} \otimes \mathrm{id}_{\mathcal{H}}(|\mathbb{1}\rangle\langle\mathbb{1}|)$ is positive semidefinite.

(iv) $\mathcal{X}$ is completely positive iff $\mathfrak{C}(\mathcal{X})$ is positive semidefinite.

(v) $\mathcal{X}$ is a CP map iff there are operators $K_1, \ldots, K_r \in \mathrm{L}(\mathcal{H}, \mathcal{K})$, where $r = \mathrm{rank}(\mathfrak{C}(\mathcal{X}))$ so that

$$\mathcal{X}(A) = \sum_{i=1}^{r} K_i A K_i^\dagger \tag{5.11}$$

for all $A \in \mathrm{L}(\mathcal{H})$. Moreover, show that $\mathcal{X}$ is a CPT map iff (5.11) holds with $\sum_{i=1}^{r} K_i^\dagger K_i = \mathbb{1}$.

(vi) $\mathcal{X}$ is a CPT map iff it has a *Stinespring dilation*, i.e., there is a unitary $U \in \mathrm{U}((\mathcal{K} \otimes \mathcal{H})^{\otimes 2})$ so that $\mathcal{X}(\rho) = \mathrm{Tr}_{2,3}(U\rho \otimes |00\rangle\langle 00| U^\dagger)$, where $\mathrm{Tr}_{2,3}$ is the partial trace over the space the $|00\rangle$ state is from.

Proof. As an exercise or see, e.g., [17, Chapter 2.2]. □

The last statement means that CPT maps are exactly the reductions of unitary channels on a larger system.

There are two different normalization conventions for the Choi-Jamiołkowski isomorphism. For $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$ we set

$$\mathfrak{J}(\mathcal{X}) \coloneqq \frac{1}{\dim(\mathcal{H})} \mathfrak{C}(\mathcal{X}) \,. \tag{5.12}$$

The theorem tells us that $\mathcal{X}$ is a quantum channel iff $\mathfrak{J}(\mathcal{X})$ is a density matrix with the reduction to $\mathcal{H}$ (obtained by tracing over $\mathcal{K}$) being a maximally mixed state. The so-called *Choi state* of a channel $\mathcal{X}$ is

$$\mathfrak{J}(\mathcal{X}) = \mathcal{X} \otimes \mathrm{id}_{\mathcal{H}}(\phi^+) \ \in \mathcal{S}(\mathcal{K} \otimes \mathcal{H}) \,, \tag{5.13}$$

where

$$\phi^+ \coloneqq \frac{1}{\dim(\mathcal{H})} |\mathbb{1}\rangle\langle\mathbb{1}| \quad \in \mathcal{S}(\mathcal{H} \otimes \mathcal{H}) \tag{5.14}$$

is a *maximally entangled state*, i.e., has the strongest bipartite quantum correlations possible in a precise sense. In particular, the Choi state can be prepared by applying the channel to this state.

Also note that not every bipartite state corresponds to a channel. Indeed, the Choi-Jamiołkowski isomorphism is an isomorphism of convex cones, $\mathfrak{C} : \mathrm{CP}(\mathcal{H}, \mathcal{K}) \to \mathrm{Pos}(\mathcal{K} \otimes \mathcal{H})$ but $\mathrm{CPT}(\mathcal{H}, \mathcal{K})$ is mapped to a proper subset of $\mathcal{S}(\mathcal{K} \otimes \mathcal{H})$. The reason is that the trace-preservation constraint of channels corresponds to $\dim(\mathcal{H})^2$ many equalities whereas the trace constraint of states is just one equality.

Exercise 5.3 (Depolarizing channel):

Remember the depolarizing channel $\mathcal{D}_p \in \mathbb{L}(\mathbb{C}^d)$ from Definition 3.15. Show that $\mathcal{D}_p \in \mathrm{CPT}(\mathbb{C}^d)$ iff

$$-\frac{1}{d+1} \leq p \leq 1 \,. \tag{5.15}$$

For which of those values of p is $\mathcal{D}_p$ also invertible and when is the inverse also a channel?

5.4. Inner products of superoperators and the χ process matrix

Let $\mathcal{H} \cong \mathbb{C}^d$ and $\mathcal{K} \cong \mathbb{C}^{d'}$ be Hilbert spaces and $E_0, E_1, \ldots, E_{dd'-1} \subset \mathrm{L}(\mathcal{H}, \mathcal{K})$ be a Hilbert-Schmidt orthonormal basis for the linear operators from $\mathcal{H}$ to $\mathcal{K}$. We note

that vector space of linear maps $\mathbb{L}(\mathcal{H})$ is also equipped with a canonical inner product (the Hilbert-Schmidt inner product for superoperators) given by

$$\langle \mathcal{X}, \mathcal{Y} \rangle = \mathrm{Tr}[\mathcal{X}^\dagger, \mathcal{Y}] \tag{5.16}$$

for any $\mathcal{X}, \mathcal{Y} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$, where the trace can be calculated as

$$\mathrm{Tr}[\mathcal{X}] = \sum_{i=1}^{dd'-1} \langle E_i, \mathcal{X}(E_i) \rangle = \sum_{i=1}^{dd'-1} \mathrm{Tr}[E_i^\dagger \mathcal{X}(E_i)] \,. \tag{5.17}$$

We also note that for any $\mathcal{X}, \mathcal{Y} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$

$$\langle \mathcal{X}, \mathcal{Y} \rangle = \langle \mathfrak{C}(\mathcal{X}), \mathfrak{C}(\mathcal{Y}) \rangle \,, \tag{5.18}$$

where $\mathfrak{C}$ denotes the Choi-Jamiołkowski isomorphism (5.9). Moreover, a map $\mathcal{X} \in \mathbb{L}(\mathcal{H})$ can be expanded into the induced product basis $\{\mathcal{E}_{i,j}\} \subset \mathbb{L}(\mathcal{H})$, which is given by

$$\mathcal{E}_{i,j}(A) = E_i A E_j \,. \tag{5.19}$$

The expansion is

$$\mathcal{X} = \sum_{i,j=0}^{dd'-1} x_{i,j} \mathcal{E}_{i,j} \tag{5.20}$$

with

$$x_{i,j} = \langle E_i, \mathcal{X}(E_j) \rangle = \mathrm{Tr}[E_i^\dagger \mathcal{X}(E_j)] \,. \tag{5.21}$$

For $\mathcal{H} = \mathcal{K}$ one typically uses basises with $E_0 \propto \mathbb{1}$. Moreover, it is common to use a different normalization convention. For qubits, i.e. $d = 2^n$, it is common to use Pauli strings $\mathbb{1} = P_0, P_1, \ldots, P_{d^2-1}$ with $\|P_i\|_{\mathrm{op}} = 1$ for all $i = 0, \ldots, d^2 - 1$. Then $\mathcal{X} \in \mathbb{L}(\mathbb{C}^d)$ can be written as

$$\mathcal{X}(\rho) = \sum_{i,j=0}^{d^2-1} \chi_{i,j}^{\mathcal{X}} P_i \rho P_j \tag{5.22}$$

in terms of an argument $\rho \in \mathrm{L}(\mathbb{C}^d)$. This representation is called the χ *process matrix representation* or just χ *matrix* of $\mathcal{X}$. Taking the normalization into account, the χ matrix elements are given as

$$\chi_{i,j}^{\mathcal{X}} = \frac{1}{d^2} \mathrm{Tr}[P_i \, \mathfrak{C}(P_j)] \,. \tag{5.23}$$

Exercise 5.4 (CPT conditions and χ-matrix):

Complement the list in Theorem 5.2 by expressing the CPT conditions in terms of the χ process matrix (5.23).

5.5. The diamond norm

The diamond norm is a norm on maps that quantifies distances of quantum channels in an operationally meaningful way.

We start with defining the *$(1 \to 1)$-norm* on $\mathbb{L}(\mathcal{H}, \mathcal{K})$ to be the operator norm induced by the trace norm, i.e., by

$$\|\mathcal{X}\|_{1\to 1} := \sup_{\|A\|_1 \leq 1} \|\mathcal{X}(A)\|_1 \,. \tag{5.24}$$

Note that if $\mathcal{X}$ is Hermiticity-preserving then the supremum is attained for a Hermitian

operator since in that case $\mathcal{X}(A) = \frac{1}{2}(\mathcal{X}(A) + \mathcal{X}(A)^\dagger) = \mathcal{X}\left(\frac{1}{2}(A + A^\dagger)\right)$. Moreover, due to convexity, we have for any $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$

$$\|\mathcal{X}\|_{1\to 1} = \sup\left\{\|\mathcal{X} \otimes \mathrm{id}(|\psi\rangle\langle\phi|)\|_1 : \; \||\psi\rangle\|_{\ell_2} = \||\psi\rangle\|_{\ell_2} = 1\right\}, \tag{5.25}$$

where one can take $|\psi\rangle = |\phi\rangle$ if $\mathcal{X}$ is Hermiticity-preserving. This means that the supremum is attained for rank-1 operators. As density operators are normalized in trace norm this implies that channels are normalized in $(1 \to 1)$-norm, i.e.,

$$\|\mathcal{X}\|_{1\to 1} = 1 \qquad \forall \mathcal{X} \in \mathrm{CPT}(\mathcal{H}, \mathcal{K}). \tag{5.26}$$

In order to distinguish quantum channels one can use ancillary systems. This motivates the definition of the diamond norm as a so-called *CB-completion* of the $(1 \to 1)$-norm, which is justified by Theorem 5.3 below. To begin with, we define *diamond norm* by

$$\|\mathcal{X}\|_\diamond := \|\mathcal{X} \otimes \mathrm{id}_{\mathcal{H}}\|_{1\to 1}. \tag{5.27}$$

Note that the diamond norm inherites the above mentioned properties from the $(1 \to 1)$-norm.

Theorem 5.3 (Complete boundedness and (sub)multiplicativity):

For any $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$

$$\|\mathcal{X}\|_\diamond := \sup_{\mathcal{H}'} \|\mathcal{X} \otimes \mathrm{id}_{\mathcal{H}'}\|_{1\to 1}, \tag{5.28}$$

where the supremum is taken over all finite dimensional Hilbert spaces $\mathcal{H}'$. Moreover,

$$\|\mathcal{X} \otimes \mathcal{Y}\|_\diamond = \|\mathcal{X}\|_\diamond \|\mathcal{Y}\|_\diamond \tag{5.29}$$
$$\|\mathcal{X}\mathcal{Z}\|_\diamond \leq \|\mathcal{X}\|_\diamond \|\mathcal{Z}\|_\diamond \tag{5.30}$$

for all $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$, $\mathcal{Y} \in \mathbb{L}(\mathcal{H}', \mathcal{K}')$ and $\mathcal{Z} \in \mathbb{L}(\mathcal{H}', \mathcal{H})$.

Proof. For the proof we refer e.g. to [17, Chapter 3.3] or recommend to prove it as an exercise. □

The theorem tells us that the diamond norm is the maximum distinguishability of quantum channels in the following sense. Let $\Phi = \mathcal{X} - \mathcal{Y}$ with $\mathcal{X}, \mathcal{Y} \in \mathrm{CPT}(\mathcal{H}, \mathcal{K})$ be the difference of two quantum channels. One can prepare copies of a state $\rho \in \mathcal{S}(\mathcal{H} \otimes \mathcal{H}')$ and apply either $\mathcal{X}$ or $\mathcal{Y}$ to the parts on $\mathcal{H}$ to obtain states on $\mathcal{K} \otimes \mathcal{H}'$. Then Proposition 4.8 tells us that $\frac{1}{2}\|\Phi \otimes \mathrm{id}_{\mathcal{H}'}(\rho)\|_1$ is the distinguishability of the output states. Taking the supremum over all (pure) states ρ yields the distinguishability of $\mathcal{X}$ and $\mathcal{Y}$, which is given by the diamond norm distance $\frac{1}{2}\|\mathcal{X} - \mathcal{Y}\|_\diamond$. In particular, the theorem tells us that optimal distinguishability can be obtained by choosing $\mathcal{H}' = \mathcal{H}$ in a similar sense as it can be detected when a map is not CP just using $\mathcal{H}' = \mathcal{H}$.

Another way to distinguish quantum states is to prepare their Choi states and distinguish them, as characterized by Proposition 4.8 via the trace norm. The following statements provides a relation of the two notions of distinguishability of quantum channels.

Proposition 5.4 (Diamond norm and trace norm):

Let $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$ and $d := \dim(\mathcal{H})$. Then

$$\|\mathfrak{J}(\mathcal{X})\|_1 \leq \|\mathcal{X}\|_\diamond \leq \dim(\mathcal{H}) \|\mathfrak{J}(\mathcal{X})\|_1, \tag{5.31}$$

where $\mathfrak{J}$ denotes the Choi-Jamiołkowski isomorphism (5.13).

Remark: The upper bound can be improved. For a Hermitian-preserving map $\mathcal{X} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$ the improved bound implies [97, Corollary 2]

$$\|\mathcal{X}\|_\diamond \leq \dim(\mathcal{H}) \, \|\mathrm{Tr}_2[|\mathfrak{J}(\mathcal{X})|]\|_\infty \, . \tag{5.32}$$

Proof of Proposition 5.4. We prove the proposition in terms of $\mathfrak{C}(\mathcal{X}) = \dim(\mathcal{H})\,\mathfrak{J}(\mathcal{X})$. It holds that

$$\|\mathcal{X}\|_\diamond = \sup\big\{\|(\mathbb{1}\otimes A)\,\mathfrak{C}(\mathcal{X})(\mathbb{1}\otimes B)\|_1 : \ \|A\|_\mathrm{F} = \|B\|_\mathrm{F} = 1\big\}\,, \tag{5.33}$$

as can be seen from (5.25) and using tensor network diagrams. Choosing $A = B = \mathbb{1}/\sqrt{\dim(\mathcal{H})}$ (corresponding to the maximally entangled state (5.14)) establishes the lower bound. The upper bound follows using Hölder's inequality,

$$\begin{aligned}\|(\mathbb{1}\otimes A)\,\mathfrak{C}(\mathcal{X})(\mathbb{1}\otimes B)\|_1 &\leq \|\mathbb{1}\otimes A\|_\mathrm{op}\,\|\mathfrak{C}(\mathcal{X})\|_1\,\|\mathbb{1}\otimes B\|_\mathrm{op} \\ &= \|\mathbb{1}\|_\mathrm{op}\,\|A\|_\mathrm{op}\,\|\mathfrak{C}(\mathcal{X})\|_1\,\|\mathbb{1}\|_\mathrm{op}\,\|B\|_\mathrm{op} \\ &\leq \|A\|_\mathrm{F}\,\|B\|_\mathrm{F}\,\|\mathfrak{C}(\mathcal{X})\|_1\,.\end{aligned} \tag{5.34}$$

□

Exercise 5.5 (The diamond norm/trace norm inequalities are tight):

Show that the bounds in Proposition 5.4 are tight, i.e., that there are $\mathcal{X}, \mathcal{Y} \in \mathbb{L}(\mathcal{H}, \mathcal{K})$ so that $\|\mathfrak{J}(\mathcal{X})\|_1 = \|\mathcal{X}\|_\diamond$ and $\|\mathcal{X}\|_\diamond = \dim(\mathcal{H})\,\|\mathfrak{J}(\mathcal{X})\|_1$.

These results tell us that the distinguishing quantum channels via their Choi states is in general not optimal.

It is non-obvious how the diamond norm can actually be computed in practically. Watrous has shown that the diamond norm can be computet efficiently via an SDP [98]. However, for the relevant case where the map is a difference of two unitary channels the computation is much simpler:

Proposition 5.5 (Diamond norm distance of unitary channels):

For any $U, V \in \mathrm{U}(d)$ the diamond norm distance of the corresponding unitary channels is

$$\frac{1}{2}\,\|\mathcal{U} - \mathcal{V}\|_\diamond = \sqrt{1 - \mathrm{dist}\big(0, \mathrm{conv}\{\lambda_i\}_{i\in[d]}\big)^2}\,, \tag{5.35}$$

where λ_i are the eigenvalues of $U^\dagger V$ and $\mathrm{dist}(\,\cdot\,,\,\cdot\,)$ denotes the Euclidean distance and $\mathrm{conv}(\,\cdot\,)$ the convex hull, both in the complex plane.

This proposition reflects that the diamond norm distance is a worst-case qunantity, where here the worst-case optimization is done over the spectrum of the unitary "difference" $U^\dagger V$.

Proof of Proposition 5.5. Starting with (5.25) and, e.g. by using tensor network diagrams or using the Choi-Jamiołkowski isomorphism, (5.10) and the vectorization rules for matrix products (2.4), we can write the diamond norm of the channel difference as

$$\begin{aligned}\|\mathcal{U} - \mathcal{V}\|_\diamond &= \max\{\|(\mathbb{1}\otimes A)(|U\rangle\langle U| - |V\rangle\langle V|)(\mathbb{1}\otimes A)\|_1 : \ \|A\|_2 = 1\} \\ &= \max\{\||AU\rangle\langle AU| - |AV\rangle\langle AV|\|_1 : \ \|A\|_2 = 1\} \\ &= \max\big\{\big\||A\rangle\langle A| - |AU^\dagger V\rangle\langle AU^\dagger V|\big\|_1 : \ \|A\|_2 = 1\big\}\end{aligned} \tag{5.36}$$

According to Watrous' lecture notes [9, Example 2.3], normalized vectors $|\psi\rangle, |\phi\rangle \in \mathbb{S}^{d-1} \subset \mathbb{C}^d$ satisfy

$$\| |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_p = 2^{1/p} \sqrt{1 - |\langle\psi|\phi\rangle|^2}\,. \tag{5.37}$$

This yields

$$\begin{aligned}
\frac{1}{2}\|\mathcal{U} - \mathcal{V}\|_\diamond &= \max\left\{ \sqrt{1 - |\langle A|AU^\dagger V\rangle|^2} : \ \|A\|_2 = 1 \right\} \\
&= \max\left\{ \sqrt{1 - |\mathrm{Tr}[A^2 U^\dagger V]|^2} : \ \|A\|_2 = 1 \right\} \\
&= \max\left\{ \sqrt{1 - |\mathrm{Tr}[\rho\, U^\dagger V]|^2} : \ \rho \in \mathcal{S}(\mathbb{C}^d) \right\} \\
&= \sqrt{1 - \min_{\rho \in \mathcal{S}(\mathbb{C}^d)} |\mathrm{Tr}[\rho\, U^\dagger V]|^2} \\
&= \sqrt{1 - \min\left\{ \left| \sum_i p_i \lambda_i \right|^2 : \ p \in [0,1]^d, \sum_i p_i = 1 \right\}} \\
&= \sqrt{1 - \mathrm{dist}(0, \mathrm{conv}\{\lambda_i\})}\,.
\end{aligned} \tag{5.38}$$

□

The required further material on representation theory (Section 2.2) was also covered in Lecture 16 including an a characterization of irreps of abelian groups (Corollary 2.4), an important extension of Schur's lemma (Theorem 2.5), some more details on Schur-Weyl duality (Theorem 2.6), and a characterization of certain invariant operators (Proposition 2.7).] Lecture 16

5.6. Unitary k-designs

Unitary k-designs are distributions of unitaries that match the first k moments of the Haar measure on the unitary group. In this sense, they are similarly defined as the projective k-designs from Section 3.4 and similar results also hold here.

Definition 5.6 (Unitary k-desing):

The k-th moment operator $\mathcal{M}_\mu^{(k)} \in \mathbb{L}(\mathbb{C}^d)^{\otimes k}$ of a distribution μ on $\mathrm{U}(d)$ is defined by

$$\mathcal{M}_\mu^{(k)}(X) \coloneqq \int_{\mathrm{U}(d)} U^{\otimes k} X U^{\otimes k\,\dagger}\, \mathrm{d}\mu(U) \tag{5.39}$$

and we write $\mathcal{M}^{(k)}$ for the k-th moment operator of the Haar measure. Then a distribution μ on the unitary group is an *unitary k-design* if $\mathcal{M}_\mu^{(k)} = \mathcal{M}^{(k)}$. A subset $\{U_1, \ldots, U_{n_\mathsf{G}}\} \subset \mathrm{U}(d)$ is called an *unitary k-design* if its uniform distribution is one.

Example 5.7 (Clifford group):

The n-qubit *Clifford group* $\mathrm{Cl}_n \subset \mathrm{U}(2^n)$ is the stabilizer of the Pauli group $\mathcal{P}_n$ (see Section 3.4.1),

$$\mathrm{Cl}_n \coloneqq \{U \in \mathrm{U}(2^n) : \ U\mathcal{P}_n U^\dagger \subset \mathcal{P}_n\}\,. \tag{5.40}$$

The Clifford group is a unitary 3-design but not a unitary 4-design [33, 34, 58].

The k-th moment operator of the Haar measure can be calculated using representation theory. The following identity can be seen as a generalization of Proposition 2.7

since $\mathcal{M}^{(k)}(A)$ commute with the representations (2.24) of $U(d)$ and (2.23) of Sym_k (see e.g. [99, integration formula section]),

$$\mathcal{M}^{(k)}(A) = \frac{1}{k!} \sum_{\sigma \in \mathrm{Sym}_k} \mathrm{Tr}[A\sigma]\sigma^{-1} \sum_{\lambda \vdash k} \frac{d_\lambda}{D_\lambda} P_\lambda \,, \tag{5.41}$$

where Sym_k is the symmetric group on k elements, $d_\lambda = \dim(S_\lambda)$ and $D_\lambda = \dim(W_\lambda)$ are the dimensions of the Weyl modules W_λ and the Specht modules S_λ (in the Schur-Weyl decomposition 2.29), respectively, σ acts on $(\mathbb{C}^d)^{\otimes k}$ by permuting the tensor factors (we have omitted the π_k of the representation (2.23)), and $\lambda \vdash k$ denotes integer partitions of k.

The *(k-th) frame potential* of μ on $\mathrm{U}(d)$ is defined to be

$$\mathcal{F}_\mu^{(k)} := \mathbb{E}_{U,V\sim\mu} \left|\mathrm{Tr}[U^\dagger V]\right|^{2k} \tag{5.42}$$

and we again drop the subscript μ if μ is the Haar measure. For $k \leq d$ one can show that

$$\mathcal{F}^{(k)} = k! \,. \tag{5.43}$$

Then it holds that [100, Theorem 5.4] $\mathcal{F}_\mu^{(k)} \geq \mathcal{F}^{(k)}$ for any finite distribution μ on $\mathrm{U}(d)$ and [101, Eq. (47)]

$$\left\|\mathcal{M}_\mu^{(k)} - \mathcal{M}^{(k)}\right\|_{\mathrm{F}}^2 = \mathcal{F}_\mu^{(k)} - \mathcal{F}^{(k)} \,. \tag{5.44}$$

As important example, for $k = 2$ the only irreps are given by $\lambda = (2)$ and $\lambda = (1,1)$. It holds that $d_{(2)} = 1 = d_{(1,1)}$. Thanks to $P_{(2)} = \frac{1}{2}(\mathbb{1} + \mathbb{F})$ and $P_{(1,1)} = \frac{1}{2}(\mathbb{1} - \mathbb{F})$ it turns out that $D_{(2)} = \frac{1}{d(d+1)}$ and $D_{(1,1)} = \frac{1}{d(d-1)}$, where $\mathbb{F}$ denotes again the flip operator.

A straightforward simplification of (5.41) yields

$$\mathcal{M}^{(2)}(A) = \frac{\mathrm{Tr}[A]}{(d-1)(d+1)}\mathbb{1} - \frac{\mathrm{Tr}[A]}{(d-1)d(d+1)}\mathbb{F} + \frac{\mathrm{Tr}[A\mathbb{F}]}{(d-1)(d+1)}\mathbb{F} - \frac{\mathrm{Tr}[A\mathbb{F}]}{(d-1)d(d+1)}\mathbb{1} \tag{5.45}$$

for $d \geq 2$. Note that this formula for the second moment operator is consistent with Proposition 2.7. Indeed, $\mathcal{M}^{(2)}(A)$ satisfies the invariance condition of this statement and can be written as linear combination of $P_{\mathrm{sym}^2} = P_{(2)}$ and $P_{\wedge^2} = P_{(1,1)}$.

6. Randomized benchmarking

Randomized benchmarking (RB) can be used to practically measure the average error rate of targeted quantum channels. It does not quantify the operationally best motivated error measure –the diamond norm distance– but it can be practically measured in a comparatively cheap way that is robust against *state preparation and measurement* (SPAM) errors. The original version of RB [32, 102–105] quantifies the average error of Clifford gates. With *interleaved randomized benchmarking* [106, 107] one can measure in a similar way the average gate fidelity of a single Clifford gate. There are several extensions [108–116] of these basic setups.

6.1. The average gate fidelity

The average error quantified in RB is given via the average fidelity of the output of a channel. We define the *average gate fidelity* (AGF) between maps $\mathcal{X}, \mathcal{Y} \in \mathbb{L}(\mathbb{C}^d)$ to

be

$$\mathrm{F}_{\mathrm{avg}}(\mathcal{X},\mathcal{Y}) \coloneqq \int \mathrm{d}\psi \, \langle \mathcal{X}(|\psi\rangle\langle\psi|), \mathcal{Y}(|\psi\rangle\langle\psi|)\rangle \tag{6.1}$$

where the integral is taken according to the uniform Haar-invariant probability measure on state vectors. So, the average gate fidelity $\mathrm{F}_{\mathrm{avg}}(\mathcal{X},\mathcal{Y})$ is a measure of closeness of $\mathcal{X}$ and $\mathcal{Y}$.

Let us list some properties of the average gate fidelity.

- For any $\mathcal{X},\mathcal{Y}$

$$\mathrm{F}_{\mathrm{avg}}(\mathcal{X},\mathcal{Y}) = \mathrm{F}_{\mathrm{avg}}(\mathcal{Y}^\dagger\mathcal{X},\mathrm{id})\,. \tag{6.2}$$

  This motivates the definition

$$\mathrm{F}_{\mathrm{avg}}(\mathcal{X}) \coloneqq \mathrm{F}_{\mathrm{avg}}(\mathcal{X},\mathrm{id})\,. \tag{6.3}$$

- The average gate fidelity is *skew symmetric*, $\mathrm{F}_{\mathrm{avg}}(\mathcal{X},\mathcal{Y}) = \mathrm{F}_{\mathrm{avg}}(\mathcal{Y},\mathcal{X})^*$.
- When $\mathcal{X}$ and $\mathcal{Y}$ are Hermiticity-preserving then their average gate fidelity is real, $\mathrm{F}_{\mathrm{avg}}(\mathcal{X},\mathcal{Y}) \in \mathbb{R}$, and hence *symmetric*,

$$\mathrm{F}_{\mathrm{avg}}(\mathcal{X},\mathcal{Y}) = \mathrm{F}_{\mathrm{avg}}(\mathcal{Y},\mathcal{X}) \tag{6.4}$$

- The distance measure corresponding to the AGF

$$\mathrm{r}(\mathcal{X},\mathcal{Y}) \coloneqq 1 - \mathrm{F}_{\mathrm{avg}}(\mathcal{X},\mathcal{Y}) \tag{6.5}$$

  is called the *average error rate* and inherites those two properties. We set $\mathrm{r}(\mathcal{X}) \coloneqq 1 - \mathrm{F}_{\mathrm{avg}}(\mathcal{X})$ for $\mathcal{X} \in \mathbb{L}(\mathcal{H})$.

The average gate fidelity is related to the diamond norm as follows.

Proposition 6.1 (AGF and diamond norm [109, Proposition 9]):

For any $\mathcal{X} \in \mathrm{CPT}(\mathbb{C}^d)$

$$\frac{d+1}{d}\left(1 - \mathrm{F}_{\mathrm{avg}}(\mathcal{X})\right) \leq \frac{1}{2}\left\|\mathrm{id} - \mathcal{X}\right\|_\diamond \leq \sqrt{d(d+1)}\sqrt{1 - \mathrm{F}_{\mathrm{avg}}(\mathcal{X})}\,. \tag{6.6}$$

Proof. The proof follows from Proposition 5.4 and [117, 118]

$$(d+1)\,\mathrm{F}_{\mathrm{avg}}(\mathcal{X}) = d\,\mathrm{F}(\phi^+, \mathfrak{J}(\mathcal{X})) \tag{6.7}$$

with ϕ^+ being the maximally entangled state (5.14). □

A refinedment of these bounds taking the so-called *unitarity* into account has been derived by Kueng et al. [119].

With the next theorem we will derive identities that are crucial for RB.

Theorem 6.2 (Twirling of channels [102, 118]):

Let $\mathcal{X} \in \mathbb{L}(\mathbb{C}^d)$ be trace-preserving and μ be a unitary 2-design. Then

$$\mathbb{E}_{U\sim\mu}[U\mathcal{X}(U^\dagger\rho U)U^\dagger] = \mathcal{D}_p(\rho)\,, \tag{6.8}$$

where

$$p = \frac{d\,\mathrm{F}_{\mathrm{avg}}(\mathcal{X}) - 1}{d-1} \tag{6.9}$$

$$= \frac{\mathrm{Tr}[\mathcal{X}] - 1}{d^2 - 1}\,. \tag{6.10}$$

We note that the $(0,0)$ component of the chi process matrix (5.23) of $\mathcal{X}$ is

$$\chi_{0,0} = \frac{1}{d^2}\operatorname{Tr}[\mathcal{X}] \tag{6.11}$$

and that $\operatorname{Tr}[\mathcal{X}]$ real if $\mathcal{X}$ is Hermiticity-preserving. Often one only considers qubits and then (6.10) is sometimes stated in terms of $\chi_{0,0}$.

Proof. The map $\mathrm{tw} : \mathrm{L}(\mathbb{C}^d) \to \mathrm{L}(\mathbb{C}^d)$ given by

$$\mathrm{tw}(\mathcal{X}) \coloneqq \mathbb{E}_{U\sim\mu}[U\mathcal{X}(U^\dagger(\,\cdot\,)U)U^\dagger] \tag{6.12}$$

is isomorphic to the second moment operator (5.39)

$$\mathcal{M}^{(2)}_\mu(A) \coloneqq \mathbb{E}_{U\sim\mathrm{U}(d)}[(U\otimes U)A(U\otimes U)^\dagger]\,, \tag{6.13}$$

since in both cases an order-4 tensor is multiplied by $U\otimes U\otimes U^\dagger\otimes U^\dagger$ with properly matched indices. Hence, it is sufficient to proved the statement for the case where μ is the Haar measure.

It is not difficult to see that $(U\otimes U)\mathcal{M}^{(2)}(A) = \mathcal{M}^{(2)}(A)(U\otimes U)$ for all $A \in \mathrm{L}(\mathbb{C}^d\otimes\mathbb{C}^d)$ and $U\in\mathrm{U}(d)$. Hence, Proposition 2.7 implies that the range of $\mathcal{M}^{(2)}$ is two-dimensional and, hence,

$$\dim(\mathrm{ran}(\mathrm{tw})) = 2\,. \tag{6.14}$$

We note that $\mathrm{tw}(\mathrm{id}) = \mathrm{id}$ and $\mathrm{tw}(\frac{\mathbb{1}}{d}\operatorname{Tr}[\,\cdot\,]) = \frac{\mathbb{1}}{d}\operatorname{Tr}[\,\cdot\,]$. Hence, id and $\frac{\mathbb{1}}{d}\operatorname{Tr}[\,\cdot\,]$ are in the range of tw. Also note that $\mathrm{tw}(\mathcal{X})$ is trace-preserving and, hence, $\mathrm{tw}(\mathcal{X})$ must be an affince combination of these two maps. This proves that (6.8) holds for some $p\in\mathbb{C}$.

In order to derive (6.9) we observe that

$$\mathrm{F}_{\mathrm{avg}}(\mathcal{X}) = \mathrm{F}_{\mathrm{avg}}(\mathrm{tw}(\mathcal{X}))\,. \tag{6.15}$$

Hence, we only need to calculate the AGF of $\mathcal{D}_p$,

$$\begin{aligned}\mathrm{F}_{\mathrm{avg}}(\mathcal{X}) &= \mathrm{F}_{\mathrm{avg}}(\mathcal{D}_p)\\ &= p + (1-p)\frac{1}{d}\,,\end{aligned} \tag{6.16}$$

which implies (6.9).

A similar argument yields (6.10): We take the trace of (6.8) to obtain

$$\operatorname{Tr}[\mathcal{X}] = \operatorname{Tr}[\mathrm{tw}(\mathcal{X})] = \operatorname{Tr}[\mathcal{D}_p] = pd^2 + (1-p)\,, \tag{6.17}$$

which is equivalent to (6.10).

□

We note that (6.10) is equivalent to

$$\mathrm{F}_{\mathrm{avg}}(\mathcal{X}) = \frac{\operatorname{Tr}[\mathcal{X}]-1}{d(d+1)} + \frac{1}{d}\,, \tag{6.18}$$

and

$$\operatorname{Tr}[\mathcal{X}] = d(d+1)\,\mathrm{F}_{\mathrm{avg}}(\mathcal{X}) - d\,, \tag{6.19}$$

where $\mathcal{X}\in\mathbb{L}(\mathbb{C}^d)$ was assumed to be trace-preserving in Theorem 6.2.

This implies that the average gate fidelity can be connected to the cononical inner product on $\mathbb{L}(\mathbb{C}^d)$ as [117, 118] (see also [119])

$$\langle\mathcal{Y},\mathcal{X}\rangle = \operatorname{Tr}[\mathcal{Y}^\dagger\mathcal{X}] = d(d+1)\,\mathrm{F}_{\mathrm{avg}}(\mathcal{X},\mathcal{Y}) - \langle\mathcal{X}(\mathbb{1}),\mathcal{Y}(\mathbb{1})\rangle \tag{6.20}$$

for any $\mathcal{X},\mathcal{Y}\in\mathbb{L}(\mathbb{C}^d)$. Indeed, if $\mathcal{Y}^\dagger\mathcal{X}$ is trace-preserving then (6.20) simplifies to (6.19). Note that this identity also connects the average gate fidelity to the Frobenius

norm meaning that the Frobenius norm is an average case error measure as well.

Also not that for a unitary channel $\mathcal{U} \in \mathrm{CPT}(\mathbb{C}^d)$ with $U \in \mathrm{U}(d)$

$$\mathrm{F}_{\mathrm{avg}}(\mathcal{U}) = \frac{\mathrm{Tr}[\mathcal{U}] - 1}{d(d+1)} + \frac{1}{d} = \frac{|\,\mathrm{Tr}[U]\,|^2 - 1}{d(d+1)} + \frac{1}{d}\,. \tag{6.21}$$

This equality reflects that the average gate fidelity measures how close U is to $\mathbb{1}$ on average, where here the average is taken over it's spectrum.

6.2. The standard RB protocol

Standard RB [105, 120] aims to provide an estimate of the AGF averaged over a gate set $\mathsf{G} = \{G_1, \ldots, G_{n_\mathsf{G}}\} \subset \mathrm{U}(d)$ that is a subgroup and a unitary 2-design. Usually, G is taken to be the Clifford group.

The quantity we would like to estimate ind standard RB is

$$\bar{F} \coloneqq \frac{1}{n_\mathsf{G}} \sum_{G \in \mathsf{G}} \mathrm{F}_{\mathrm{avg}}(\tilde{\mathcal{G}}, \mathcal{G})\,, \tag{6.22}$$

where $\tilde{\mathcal{G}}$ denotes the implementation of the gate $\mathcal{G}$. In the analysis of the following protocol we will see that RB provides indeed a consistent estimate of this quantity for the case of *gate independent noise*. For an extension to gate dependent noise see the initial work by Magesan et al. [105] for a first perturbative analysis and for a more recent and more rigorous analysis the work by Wallman [121].

Protocol 6.3 (Standard RB):

Let $\mathsf{G} = \{G_1, \ldots, G_{n_\mathsf{G}}\} \subset \mathrm{U}(d)$ be a subgroup. Moreover, let $\rho \in \mathcal{S}(\mathbb{C}^d)$ be the initial state, and $\mathsf{M} = \{M, \mathbb{1} - M\} \subset \mathrm{Pos}(\mathbb{C}^d)$ be the measurement (usually, $\rho = |0\rangle\langle 0| = M$). For any sequence $s \in [n_\mathsf{G}]^m$ set $\mathcal{G}^{(s)} \coloneqq \mathcal{G}_{s_m} \ldots \mathcal{G}_{s_1}$ to be the *gate sequence* and

$$F_{m,s} \coloneqq \mathrm{Tr}[M \mathcal{G}^{(s)\,-1} \mathcal{G}(s)(\rho)] \tag{6.23}$$

to be the *sequence fidelity*.
Then the standard RB protocol consist of the following steps.

- Draw sequence $s \in [n_\mathsf{G}]^m$ uniformly at random, which we denote by $s \sim [n_\mathsf{G}]^m$.
- Implement the gate sequence $\mathcal{S}^{(s)} \coloneqq \mathcal{G}_{s_{m+1}} \mathcal{G}_{s_m} \ldots \mathcal{G}_{s_1}(\rho)$, where the last gate $\mathcal{G}_{s_{m+1}} \coloneqq \mathcal{G}^{(s)\,-1}$ is the inverse of the gate $G^{(s)} \in \mathsf{G}$.
- Obtain an estimate $\hat{F}_{m,s}$ of $F_{m,s}$ my measuring $\mathcal{S}^{(s)}(\rho)$ with the measurement M a number of n_s times.
- For each $m = 1, \ldots, m_{\max}$ repeat this estimation for sequences $s^{(1)}, \ldots, s^{(n_m)} \in [n_\mathsf{G}]^m$ and set $\hat{F}_m$ to be the corresponding empirical estimate of the *average sequence fidelity*

$$\bar{F}_m \coloneqq \mathbb{E}_{s \sim [n_\mathsf{G}]^m} \left[F_{m,s} \right]. \tag{6.24}$$

- Fit parameters $A, B, p \in \mathbb{R}$ to the model

$$\bar{F}_m = A p^m + B \tag{6.25}$$

to obtain estimates $\hat{p}$, $\hat{A}$, $\hat{B}$.

- Obtain an estimate $\hat{\bar{F}}$ of the AGF (6.22) via

$$p = \frac{d\bar{F} - 1}{d - 1}\,, \tag{6.26}$$

(remember the relation (6.9) of p and $\mathrm{F}_{\mathrm{avg}}$).

Analysis for gate independent noise. We denote the noisy implementations of the initial state ρ, the gates $\mathcal{G}_i$, and the measurement M by $\tilde{\rho}$, $\tilde{\mathcal{G}}_i$, and $\tilde{M}$, respectively.

We restrict our analysis to gate-independent noise, i.e.,

$$\tilde{\mathcal{G}}_i = \Lambda \mathcal{G}_i \tag{6.27}$$

for some channel $\Lambda \in \mathrm{CPT}(\mathbb{C}^d)$. Setting $\mathcal{C}_{s_j} \coloneqq \mathcal{G}_{s_j} \dots \mathcal{G}_{s_1}$ for $j \in [m]$, this assumption allows us to rewrite the implemented gate sequence as

$$\begin{aligned} \tilde{\mathcal{S}}^{(s)} &= \Lambda\, \mathcal{G}_{s_{m+1}} \Lambda\, \mathcal{G}_{s_m} \Lambda\, \mathcal{G}_{s_{m-1}} \dots \Lambda \mathcal{G}_{s_1} \\ &= \Lambda \big(\mathcal{C}_{s_m}^\dagger \Lambda \mathcal{C}_{s_m}\big) \big(\mathcal{C}_{s_{m-1}}^\dagger \Lambda \mathcal{C}_{s_{m-1}}\big) \dots \big(\mathcal{C}_{s_1}^\dagger \Lambda \mathcal{C}_{s_1}\big)\,, \end{aligned} \tag{6.28}$$

where we have used that $\mathcal{G}_{s_{m+1}} = \mathcal{C}_m^\dagger$. Since the gates $\mathcal{G}_{s_i}$ are drawn iid. from a unitary 2-design, Theorem 6.2 implies that

$$\begin{aligned} \mathbb{E}_{s \sim [n_{\mathsf{G}}]^m} \big[\tilde{\mathcal{S}}^{(s)}\big] &= \Lambda \mathbb{E}_{C \in \mathsf{G}} [\mathcal{C}^\dagger \Lambda \mathcal{C}] \\ &= \Lambda \mathcal{D}_p^m = \Lambda \mathcal{D}_{p^m} \end{aligned} \tag{6.29}$$

with p given by (6.9), which matches the estimation (6.26). Moreover, the average sequence fidelity is

$$\begin{aligned} \bar{F}_m &= \mathrm{Tr}[\tilde{E} \Lambda \mathcal{D}_{p^m}(\tilde{\rho})] \\ &= p^m \,\mathrm{Tr}[\tilde{E} \Lambda(\tilde{\rho})] + (1 - p^m)\,\mathrm{Tr}[\tilde{E} \Lambda(\mathbb{1}/d)] \\ &= p^m \,\mathrm{Tr}[\tilde{E} \Lambda(\tilde{\rho} - \mathbb{1}/d)] + \mathrm{Tr}[\tilde{E} \Lambda(\mathbb{1}/d)]\,. \end{aligned} \tag{6.30}$$

This expression matches the fitting model (6.25) with

$$A \coloneqq \mathrm{Tr}[\tilde{E} \Lambda(\tilde{\rho} - \mathbb{1}/d)] \qquad \text{and} \qquad B \coloneqq \mathrm{Tr}[\tilde{E} \Lambda(\mathbb{1}/d)] \tag{6.31}$$

being the so-called *SPAM constants*. □

Note that the resulting estimate of (6.22) is robust against *state preparation and measurement* (SPAM) errors, which are absorbed in the SPAM constants A and B.

In order to make RB *scalable* it is important to have an efficiently tractable group structure so that the inverse of the gate sequence can be computed. For the important example of the Clifford group the Gottesman-Knill theorem allows to compute the inverse of $\mathcal{G}^{(s)}$ in polynomial time in the number of qubits. Using probabilistic tails bounds (see Section 2.3), one can prove that the estimation of the involved quantities can also be done efficiently, even when using just two different sequence lengths m [122].

6.3. Interleaved randomized benchmarking

It is also an important task to estimate the quality of the implementation of one specific gate $\mathcal{G}_t$. *Interleaved RB* [106] solves this task by providing an estimate of the average gate fidelity between the targeted gate $\mathcal{G}_t$ and its implementation $\tilde{\mathcal{G}}_t$. This is achieved in a way that is robust against SPAM errors as well as against errors in the RB gate sequences used to extract the average gate fidelity.

Let us consider again gate independent noise $\Lambda \in \mathrm{CPT}(\mathcal{H})$ that again acts on every gate $\mathcal{G} \in \mathsf{G}$, i.e., the implemented gates are

$$\tilde{\mathcal{G}} = \Lambda\mathcal{G} \tag{6.32}$$

and a noisy target gate $G_t \in \mathsf{G}$ with possibly different noise $\Lambda_t \in \mathrm{CPT}(\mathcal{H})$

$$\tilde{\mathcal{G}}_t = \Lambda_t\mathcal{G}_t\,. \tag{6.33}$$

The idea of interleaved RB is to insert $\tilde{\mathcal{G}}_t$ after every gate in the gate sequence $\mathcal{G}^{(s)}$ in standard RB, so to interleave the sequence with multiple applications of $\tilde{\mathcal{G}}_t$. However, the gate independent noise Λ also needs to be estimated.

In more detail, in interleaved RB one estimates $\mathrm{F}_{\mathrm{avg}}(\tilde{\mathcal{G}}_t, \mathcal{G}_t)$ by applying standard RB twice, (i) to obtain an estimate on $\mathrm{F}_{\mathrm{avg}}(\Lambda)$ and (ii) to obtain an estimate on $\mathrm{F}_{\mathrm{avg}}(\tilde{\mathcal{G}}_t\Lambda, \mathcal{G}_t) = \mathrm{F}_{\mathrm{avg}}(\mathcal{G}_t^\dagger\tilde{\mathcal{G}}_t\Lambda)$; see Protocol 6.4 for an RB method to achieve (i).

The idea is that one can extract $\mathrm{F}_{\mathrm{avg}}(\tilde{\mathcal{G}}_t, \mathcal{G}_t)$ from $\mathrm{F}_{\mathrm{avg}}(\tilde{\mathcal{G}}_t\Lambda, \mathcal{G}_t)$ once the noise strength given by $\mathrm{F}_{\mathrm{avg}}(\Lambda)$ is known. In order extract and estimation of $\mathrm{F}_{\mathrm{avg}}(\tilde{\mathcal{G}}_t, \mathcal{G}_t)$ from these two quantities an approximation of the form

$$\mathrm{F}_{\mathrm{avg}}(\mathcal{X}\mathcal{Y}) \approx \mathrm{F}_{\mathrm{avg}}(\mathcal{X})\,\mathrm{F}_{\mathrm{avg}}(\mathcal{Y}) \tag{6.34}$$

is used. Then one obtains the desired average gate fidelity by taking estimates corresponding to

$$\begin{aligned}\mathrm{F}_{\mathrm{avg}}(\tilde{\mathcal{G}}_t, \mathcal{G}_t) &= \mathrm{F}_{\mathrm{avg}}(\tilde{\mathcal{G}}_t\mathcal{G}_t^\dagger) = \mathrm{F}_{\mathrm{avg}}(\mathcal{G}_t^\dagger\tilde{\mathcal{G}}_t)\\ &\approx \frac{\mathrm{F}_{\mathrm{avg}}(\mathcal{G}_t^\dagger\tilde{\mathcal{G}}_t\Lambda)}{\mathrm{F}_{\mathrm{avg}}(\Lambda)}\,.\end{aligned} \tag{6.35}$$

Interleaved RB has been improved and simplified by Kimmel et al. [110, Section 6A]. They have found a bound on the apprximation error that is tighther than the prvious bound [106]. In terms of $\chi_{0,0}$ from (6.11) it reads as

$$|\chi_{0,0}^{\mathcal{X}\mathcal{Y}} - \chi_{0,0}^{\mathcal{X}}\chi_{0,0}^{\mathcal{Y}}| \leq 2\sqrt{(1-\chi_{0,0}^{\mathcal{X}})\chi_{0,0}^{\mathcal{X}}(1-\chi_{0,0}^{\mathcal{Y}})\chi_{0,0}^{\mathcal{Y}}} + (1-\chi_{0,0}^{\mathcal{X}})(1-\chi_{0,0}^{\mathcal{Y}}) \tag{6.36}$$

for any $\mathcal{X}, \mathcal{Y} \in \mathrm{CPT}((\mathbb{C}^2)^{\otimes n})$. This bound yields bounds on the error in (6.34) via (6.10).

Protocol 6.4 (Modified RB):

For a target gate $G_t \in \mathsf{G}$ this protocol is obtained from Protocol 6.3 by replacing the gate sequence $\mathcal{G}^{(s)}$ by

$$\mathcal{G}_{G_t}^{(s)} := \mathcal{G}_t\mathcal{G}_{s_m}\mathcal{G}_t\mathcal{G}_{s_{m-1}}\ldots\mathcal{G}_t\mathcal{G}_{s_1}\,. \tag{6.37}$$

Everyting is now done w.r.t. this modified gate sequence. For instance, the last gate is $\mathcal{G}_{s_{m+1}} := \mathcal{G}_{G_t}^{(s)\,-1}$.

Analysis. We assume the noise model given by (6.32) and (6.33) and that G is a unitary 2-design.

It is not difficult to see that the implemented gate sequence can be written as

$$\tilde{\mathcal{S}}^{(s)} = \Lambda\big(\mathcal{C}_{s_m}^\dagger\,\Phi\,\mathcal{C}_{s_m}\big)\big(\mathcal{C}_{s_{m-1}}^\dagger\,\Phi\,\mathcal{C}_{s_{m-1}}\big)\ldots\big(\mathcal{C}_{s_1}^\dagger\,\Phi\,\mathcal{C}_{s_1}\big) \tag{6.38}$$

with $\Phi = \mathcal{G}_t^\dagger\tilde{\mathcal{G}}_t\Lambda$ and $C_{s_i} \sim \mathsf{G}$ iid. Hence, applying the same arguments as in the analysis of the standard RB protocol yields

$$\mathbb{E}_{s\sim[n_{\mathsf{G}}]^m}\big[\tilde{\mathcal{S}}^{(s)}\big] = \Lambda\mathcal{D}_{p^m} \tag{6.39}$$

with the estimated average gate fidelity $\hat{\bar{F}}_{\mathsf{G}}$ being an estimate of

$$\mathrm{F}_{\mathrm{avg}}(\Phi) = \mathrm{F}_{\mathrm{avg}}(\mathcal{G}_t^\dagger \tilde{\mathcal{G}}_t \Lambda) \tag{6.40}$$

as desired. □

7. Process tomography

A quantum state can be reconstructed from measurement data using quantum state tomography (Section 3). In a similar way, quantum process tomography can be used to fully reconstruct quantum channels from measurement data. However, more complicated measurement setups are required for process tomography.

In the most simple version, one can use linear inversion [123] to obtain the channel's χ process matrix (5.23). But this can be challenging to implement and leads to a non-optimal sampling complexity (number of invocations of the channel). If the χ matrix is sparse, then compressed sensing 1.0 can be used [124] to dramatically reduce the measurement effort, cp. the reconstruction program (3.65). However, in most situations the χ matrix cannot be expected to be sparse.

Another way to process tomography is to reduce it to state tomography by preparing the channel's Choi state [125]. But this method requires maximally entangled states, see (5.13), which are often difficult prepare with high fidelity.

Flammia et al. [55] presents a process tomography protocol that is based on low-rank matrix reconstruction with random Pauli measurements. These low rank recovery guarantees can be applied to the channel's Choi state, since the rank of this matrix representation equals the Kraus rank of the original channel. On first sight, such an approach requires the use of an ancilla in order to implement the Choi state physically in a concrete application as in the work by Altepeter et al. [125]. However, Ref. [55] also provides a more direct implementation of their protocol that does not require any ancillas. Valid for multi-qubit processes this trick exploits the tensor-product structure of (multi-qubit) Pauli operators. The demerit of this approach is that the number of individual channel measurements required to evaluate a single Pauli expectation value scales with the dimension of the underlying Hilbert space.

The channel version of the PSD-fit (3.75), which is later called the CPT-fit [126], has been suggested by Baldwin et al. [53] and numerically compared to full tomography and compressed sensing by Rodionov et al. [127]. First recover guarantees for the CPT-fit and other CS recovery gurantees for quantum process tomography wree proven in Ref. [126]. Here, the channel's input states are sampled from an (approximate) projective 4-design and the output states are measured with random observables (approximate) unitary 4-design eigenbases.

Minimizing the diamond norm as a reguralizer in compressed sensing has been investigated in Ref. [126, 128]. It is argued that for certain signals including low Kraus rank quantum channels [128, 129] this recovery perform at least as well in terms of measurement effort compared to the conventional trace-norm minimization.

Another approach to quantum process tomography is the use of RB methods pioneered by Kimmel et al. [110] and explained in more detail in the following sections. The advantage of this approach is some —not yet rigorously quantified— roboustness against SPAM errors, which is inherited from randomized benchmarking.

7.1. Randomized benchmarking tomography

What kind of information on a channel $\mathcal{X}$ can be extracted from average gate fidelities? One can use interleaved RB to learn about non-Clifford gates. On can see that the

deviations of a channel $\mathcal{X}$ from being unital is not "seen" by average gate fidelities. But everything else can be reconstructed. More precisely, one can learn the *unital part* of X, which is given by

$$\mathcal{X}_u(\rho) \coloneqq \mathcal{X}\left(\rho - \mathrm{Tr}[\rho]\frac{\mathbb{1}}{d}\right), \tag{7.1}$$

from average gate fidelities [110]. The following result extends and simplifies results by Scott [100] and Kimmel et al. [110].

Theorem 7.1 ([99, Theorem 38]):

Let $\mathsf{G} \subset \mathrm{U}(d)$ be a finite unitary 2-design. The orthogonal projection $P_V : \mathbb{L}(\mathbb{C}^d) \to \mathbb{L}(\mathbb{C}^d)$ onto the linear hull of unital and trace- and Hermiticity preserving maps $V \subset \mathbb{L}(\mathbb{C}^d)$ is given by

$$P_V(\mathcal{X}) = \frac{1}{|\mathsf{G}|}\sum_{U\in\mathsf{G}} c_{\mathcal{U}}(\mathcal{X})\mathcal{U} \tag{7.2}$$

with coefficients

$$c_{\mathcal{U}}(\mathcal{X}) = \alpha\, \mathrm{F}_{\mathrm{avg}}(\mathcal{U},\mathcal{X}) - \beta\, \mathrm{Tr}[\mathcal{X}(\mathbb{1})]\,, \tag{7.3}$$

where $\alpha = d(d+1)(d^2-1)$ and $\beta = \frac{1}{d}\left(\frac{\alpha}{d}-1\right)$.

The average gate fidelities $\mathrm{F}_{\mathrm{avg}}(\mathcal{U},\mathcal{X})$ can be estimated using interleaved RB [110]. Hence, RB methods can be used to tomographically reconstruct the unital part of any quantum channel.

In the case where G is the Clifford group and $\mathcal{X}$ is (close to) a unitary channel one can use compressed sensing for the reconstruction [99]. It is sufficient to subsample the U's from the Clifford group to recover the unitary (and hence unital) $\mathcal{X}$ from a number of average gate fidelities scaling as $\tilde{\mathrm{O}}(d^2)$. This caling is clearly optimal. However, it is still non-obvious what the sample complex in terms of the number of invocations of the channel $\mathcal{X}$ is when RB is used [99].

8. Gate set tomography (additional information)

In gate set tomography (GST) [130, 131] one considers an initial state $\rho \in \mathcal{S}(\mathbb{C}^d)$, a finite gatset $\mathsf{G} = \{G_1,\ldots,G_{n_\mathsf{G}}\} \subset \mathrm{U}(d)$, and a finite POVM $\mathsf{M} = \{M_1,\ldots,M_m\} \subset \mathrm{Pos}(\mathbb{C}^d)$. Then one reconstructs a description of the triple $(\mathsf{M},\mathsf{G},\rho)$ –up to some gauge invariance– from measuring $\mathcal{G}^{(s)}(\rho)$ with M for different gate sequences

$$\mathcal{G}^{(s)} \coloneqq \mathcal{G}_{s_\ell}\ldots\mathcal{G}_{s_1}(\rho) \tag{8.1}$$

with $s \in [n_\mathsf{G}]^\ell$.

The gauge freedom is given by linear invertible transformations on $(\mathsf{M},\mathsf{G},\rho)$ that preserve the output distributions

$$p(j|s) \coloneqq \mathrm{Tr}\big[M_j\mathcal{G}^{(s)}(\rho)\big]\,. \tag{8.2}$$

The actual reconstruction of a representation of $(\mathsf{M},\mathsf{G},\rho)$ from measurement data is based on several estimation steps applied to the measurement data [132]. Since no prior knowledge on $(\mathsf{M},\mathsf{G},\rho)$ is assumed the measurement effort much larger as, e.g. in

randomized benchmarking tomography (Section 7.1). Moreover, the involved computations are challenging and do not (yet) come along with any theoretical guarantees.

On the upside, however, one can estimate diamond norm errors (Section 5.5) of the implementation of G by optimizing over the gauge freedom. So far, this seems to be the only quantum characterization and verification that has been used to estimate error in diamond norm.

Bibliography

[1] S. Flammia, *Characterization of quantum devices*, QIP tutorial 2017, Seattle (2017).

[2] J. Preskill, *Quantum supremacy now?* (2012).

[3] D. Shepherd and M. J. Bremner, *Temporally unstructured quantum computation*, Proc. Roy. Soc. A **465**, 1413 (2009), arXiv:0809.0847.

[4] M. J. Bremner, A. Montanaro, and D. J. Shepherd, *Average-case complexity versus approximate simulation of commuting quantum computations*, Phys. Rev. Lett. **117**, 080501 (2016), arXiv:1504.07999 [quant-ph].

[5] S. Aaronson and A. Arkhipov, *The computational complexity of linear optics*, in *STOC'11: Proc. 43rd Ann. ACM Symp. Theor. Comput.* (ACM, 2011) pp. 333–342, arXiv:1011.3245 [quant-ph].

[6] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, *Characterizing quantum supremacy in near-term devices*, Nature Physics **14**, 595 (2018), arXiv:1608.00263 [quant-ph].

[7] J. Bermejo-Vega, D. Hangleiter, M. Schwarz, R. Raussendorf, and J. Eisert, *Architectures for quantum simulation showing a quantum speedup*, Phys. Rev. X **8**, 021010 (2018), arXiv:1703.00466 [quant-ph].

[8] J. Preskill, *Quantum computing in the NISQ era and beyond*, Quantum **2** (2018), 10.22331/q-2018-08-06-79, arXiv:1801.00862 [quant-ph].

[9] J. Watrous, *John watrous's lecture notes*, `https://cs.uwaterloo.ca/~watrous/LectureNotes.html`, [accessed 2019-March-31].

[10] S. Foucart and H. Rauhut, *A mathematical introduction to compressive sensing* (Springer, 2013).

[11] R. T. Rockafellar, *Convex analysis*, 2nd ed. (Princeton university press, 1970).

[12] B. Simon, *Representations of finite and compact groups*, 10 (Am. Math. Soc., 1996).

[13] R. Goodman and N. R. Wallach, *Representations and invariants of the classical groups*, Vol. 68 (Cambridge University Press, 2000).

[14] M. Grant and S. Boyd, in *Recent advances in learning and control*, Lecture Notes in Control and Information Sciences, edited by V. Blondel, S. Boyd, and H. Kimura (Springer-Verlag Limited, 2008) pp. 95–110, `http://stanford.edu/~boyd/graph_dcp.html`.

[15] M. Grant and S. Boyd, *CVX: Matlab software for disciplined convex programming, version 2.1*, `http://cvxr.com/cvx` (2014).

[16] *Python package cvxpy*, https://www.cvxpy.org/install/index.html, accessed: 2019-06-10.

[17] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, 2018).

[18] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-Al-Kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt, *Scalable multiparticle entanglement of trapped ions*, Nature **438**, 643 (2005), arXiv:quant-ph/0603217 [quant-ph].

[19] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, *Sample-optimal tomography of quantum states*, IEEE Trans. Inf. Theory **63**, 5628 (2017), arXiv:1508.01797 [quant-ph].

[20] R. O'Donnell and J. Wright, *Efficient quantum tomography*, arXiv:1508.01907 [quant-ph].

[21] T. Heinosaari, L. Mazzarella, and M. M. Wolf, *Quantum tomography under prior information*, Commun. Math. Phys. **318**, 355 (2013), arXiv:1109.5478 [quant-ph].

[22] R. J. Milgram, *Immersing projective spaces*, Ann. Math. **85**, 473 (1967).

[23] K. H. Mayer, *Elliptische differentialoperatoren und ganzzahligkeitssätze für charakteristische zahlen*, Topology **4**, 295 (1965).

[24] D. Goyeneche, G. Cañas, S. Etcheverry, E. S. Gómez, G. B. Xavier, G. Lima, and A. Delgado, *Five measurement bases determine pure quantum states on any dimension*, Phys. Rev. Lett. **115**, 090401 (2015), arXiv:1411.2789 [quant-ph].

[25] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves, *Symmetric informationally complete quantum measurements*, J. Math. Phys. **45**, 2171 (2004), quant-ph/0310075.

[26] A. J. Scott, *Tight informationally complete quantum measurements*, J. Phys. A Math. Gen. **39**, 13507 (2006), quant-ph/0604049.

[27] Z. Fan, A. Heinecke, and Z. Shen, *Duality for frames*, J. Fourier Anal. Appl. **22**, 71 (2016).

[28] R. Kueng and D. Gross, *Qubit stabilizer states are complex projective 3-designs*, arXiv:1510.02767 [quant-ph].

[29] A. Ambainis and J. Emerson, *Quantum t-designs: t-wise independence in the quantum world*, in *Computational Complexity, 2007. CCC '07. Twenty-Second Annual IEEE Conference on* (2007) pp. 129–140, quant-ph/0701126.

[30] A. Roy and A. J. Scott, *Weighted complex projective 2-designs from bases: Optimal state determination by orthogonal measurements*, J. Math. Phys. **48**, 072110 (2007), arXiv:quant-ph/0703025 [quant-ph].

[31] D. Gross, K. M. R. Audenaert, and J. Eisert, *Evenly distributed unitaries: on the structure of unitary designs*, J. Math. Phys. **48**, 052104 (2007), quant-ph/0611002.

[32] C. Dankert, R. Cleve, J. Emerson, and E. Livine, *Exact and approximate unitary 2-designs and their application to fidelity estimation*, Phys. Rev. A **80**, 012304 (2009), arXiv:quant-ph/0606161 [quant-ph].

[33] H. Zhu, *Multiqubit clifford groups are unitary 3-designs*, Phys. Rev. A **96**, 062336 (2017), arXiv:1510.02619 [quant-ph].

[34] Z. Webb, *The clifford group forms a unitary 3-design*, Quantum Info. Comput. **16**, 1379 (2016), arXiv:1510.02769 [quant-ph].

[35] J. Benhelm, G. Kirchmair, U. Rapol, T. Körber, C. F. Roos, and R. Blatt, *Generation of hyperentangled photon pairs*, Phys. Rev. A **75**, 032506 (2007).

[36] A. Klappenecker and M. Roetteler, *Mutually unbiased bases are complex projective 2-designs*, in *Proc. IEEE International Symposium on Information Theory, ISIT, 2005* (IEEE, 2005) pp. 1740–1744, arXiv:quant-ph/0502031 [quant-ph].

[37] M. Guta, J. Kahn, R. Kueng, and J. A. Tropp, *Fast state tomography with optimal error bounds*, arXiv:1809.11162 [quant-ph].

[38] G. M. D'Ariano and P. Perinotti, *Optimal data processing for quantum measurements*, Phys. Rev. Lett. **98**, 020403 (2007), arXiv:quant-ph/0610058 [quant-ph].

[39] E. J. Candes and T. Tao, *Near-optimal signal recovery from random projections: Universal encoding strategies?* IEEE T Inform Theory **52**, 5406 (2006), arXiv:math/0410542 [math.CA].

[40] E. J. Candes, J. Romberg, and T. Tao, *Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information*, IEEE Trans. Inform. Theor. **52**, 489 (2006), arXiv:math/0409186 [math.NA].

[41] D. L. Donoho, *Compressed sensing*, IEEE Trans. Inf. Th. **52**, 1289 (2006).

[42] B. K. Natarajan, *Sparse approximate solutions to linear systems*, SIAM J. Comp. **24**, 227 (1995).

[43] D. Ge, X. Jiang, and Y. Ye, *A note on the complexity of lpminimization*, Mathematical Programming **129**, 285 (2011).

[44] V. Chandrasekaran, B. Recht, P. Parrilo, and A. Willsky, *The convex geometry of linear inverse problems*, Found. Comput. Math. **12**, 805 (2012), arXiv:1012.0621 [math.OC].

[45] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp, *Living on the edge: Phase transitions in convex programs with random data*, Information and Inference: A Journal of the IMA **3**, 224 (2014), arXiv:1303.6672 [cs.IT].

[46] J. A. Tropp, *Convex recovery of a structured signal from independent random linear measurements*, in *Sampling Theory, a Renaissance*, edited by E. G. Pfander (Springer, 2015) pp. 67–101, arXiv:1405.1102.

[47] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Found. Trends Mach. Learn. **3**, 1 (2011).

[48] R. Kueng, H. Rauhut, and U. Terstiege, *Low rank matrix recovery from rank one measurements*, Appl. Comp. Harm. Anal. (2015), arXiv:1410.6913 [cs.IT].

[49] R. Kueng, D. Gross, and F. Krahmer, *Spherical designs as a tool for derandomization: The case of phaselift*, in *2015 International Conference on Sampling Theory and Applications (SampTA)* (IEEE, 2015) pp. 192–196.

[50] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, *Quantum state tomography via compressed sensing*, Phys. Rev. Lett. **105**, 150401 (2010), arXiv:0909.3304 [quant-ph].

[51] D. Gross, *Recovering low-rank matrices from few coefficients in any basis*, IEEE Trans. Inf. Th. **57**, 1548 (2011), arXiv:0910.1879 [cs.IT].

[52] R. Ahlswede and A. Winter, *Strong converse for identification via quantum channels*, IEEE Trans. Inform. Theory **48**, 569 (2002), arXiv:quant-ph/0012127 [quant-ph].

[53] C. H. Baldwin, A. Kalev, and I. H. Deutsch, *Quantum process tomography of unitary and near-unitary maps*, Phys. Rev. A **90**, 012110 (2014), arXiv:1404.2877.

[54] M. Kabanava, R. Kueng, H. Rauhut, and U. Terstiege, *Stable low-rank matrix recovery via null space properties*, Information and Inference: A Journal of the IMA **5**, 405 (2016), arXiv:1507.07184 [cs.IT].

[55] S. T. Flammia, D. Gross, Y.-K. Liu, and J. Eisert, *Quantum tomography via compressed sensing: error bounds, sample complexity and efficient estimators*, New J. Phys. **14**, 095022 (2012), arXiv:1205.2300 [quant-ph].

[56] Y.-K. Liu, *Universal low-rank matrix recovery from Pauli measurements*, Adv. Neural Inf. Process. Syst. **24**, 1638 (2011), arXiv:1103.2816.

[57] R. Kueng and D. Zhu, H. an, *Low rank matrix recovery from Clifford orbits*, arXiv:1610.08070 [cs.IT].

[58] H. Zhu, R. Kueng, M. Grassl, and D. Gross, *The Clifford group fails gracefully to be a unitary 4-design*, arXiv:1609.08172 [quant-ph].

[59] S. Mendelson, *Learning without concentration*, J. ACM **62**, 21:1 (2015), arXiv:1401.0304 [cs.LG].

[60] T. Sugiyama, P. S. Turner, and M. Murao, *Precision-guaranteed quantum tomography*, Phys. Rev. Lett. **111**, 160406 (2013), arXiv:1306.4191 [quant-ph].

[61] J. A. Smolin, J. M. Gambetta, and G. Smith, *Efficient method for computing the maximum-likelihood quantum state from measurements with additive gaussian noise*, Phys. Rev. Lett. **108**, 070502 (2012), arXiv:1106.5458 [quant-ph].

[62] J. A. Tropp, *User-friendly tail bounds for sums of random matrices*, Found. Comput. Math. **12**, 389 (2012), arXiv:1004.4389 [math.PR].

[63] S. T. Flammia and Y.-K. Liu, *Direct fidelity estimation from few Pauli measurements*, Phys. Rev. Lett. **106**, 230501 (2011), arXiv:1104.4695 [quant-ph].

[64] T. M. Cover and J. A. Thomas, *Elements of information theory* (John Wiley and Sons, New York, 2012).

[65] Z. Hradil, J. Řeháček, J. Fiurášek, and M. Ježek, *3 maximum-likelihood methodsin quantum mechanics*, in *Quantum State Estimation*, Lecture Notes in Physics No. 649, edited by M. Paris and J. Řeháček (Springer Berlin Heidelberg, 2004) pp. 59–112.

[66] J. Řeháček, Z. Hradil, E. Knill, and A. I. Lvovsky, *Diluted maximum-likelihood algorithm for quantum tomography*, Phys. Rev. A **75**, 042108 (2007), arXiv:quant-ph/0611244 [quant-ph].

[67] J. Shang, Z. Zhang, and H. K. Ng, *Superfast maximum-likelihood reconstruction for quantum tomography*, Phys. Rev. A **95**, 062336 (2017), arXiv:1609.07881 [quant-ph].

[68] C. Schwemmer, L. Knips, D. Richart, H. Weinfurter, T. Moroder, M. Kleinmann, and O. Gühne, *Systematic errors in current quantum state tomography tools*, Phys. Rev. Lett. **114**, 080403 (2015), arXiv:1310.8465 [quant-ph].

[69] J. Wang, V. B. Scholz, and R. Renner, *Confidence polytopes in quantum state tomography*, Phys. Rev. Lett. **122**, 190401 (2019), arXiv:1808.09988 [quant-ph].

[70] P. Faist and R. Renner, *Practical and reliable error bars in quantum tomography*, Phys. Rev. Lett. **117**, 010404 (2016), arXiv:1509.06763 [quant-ph].

[71] C. Granade, C. Ferrie, I. Hincks, S. Casagrande, T. Alexander, J. Gross, M. Kononenko, and Y. Sanders, *QInfer: Statistical inference software for quantum applications*, Quantum **1**, 5 (2017), arXiv:1610.00336 [quant-ph].

[72] C. Granade, C. Ferrie, and S. T. Flammia, *Practical adaptive quantum tomography*, New J. Phys. **19**, 113017 (2017), arXiv:1605.05039 [quant-ph].

[73] C. A. Fuchs and J. van de Graaf, *Cryptographic distinguishability measures for quantum mechanical states*, IEEE Trans. Inf. Th. **45**, 1216 (1999), arXiv:quant-ph/9712042 [quant-ph].

[74] M. P. da Silva, O. Landon-Cardinal, and D. Poulin, *Practical characterization of quantum devices without tomography*, Phys. Rev. Lett. **107**, 210404 (2011), arXiv:1104.3835 [quant-ph].

[75] H. Pashayan, J. J. Wallman, and S. D. Bartlett, *Estimating Outcome Probabilities of Quantum Circuits Using Quasiprobabilities*, Phys. Rev. Lett. **115**, 070501 (2015), arXiv:1503.07525 [quant-ph].

[76] S. Pallister, N. Linden, and A. Montanaro, *Optimal Verification of Entangled States with Local Measurements*, Phys. Rev. Lett. **120**, 170502 (2018), arXiv:1709.03353 [quant-ph].

[77] H. Zhu and M. Hayashi, *Efficient verification of pure quantum states with applications to hypergraph states*, arXiv:1806.05565 [quant-ph].

[78] Y. Takeuchi and T. Morimae, *Verification of many-qubit states*, Phys. Rev. X **8**, 021060 (2018), arXiv:1709.07575 [quant-ph].

[79] Z. Li, Y.-G. Han, and H. Zhu, *Efficient verification of bipartite pure states*, arXiv:1901.09783 [quant-ph].

[80] X.-D. Yu, J. Shang, and O. Gühne, *Optimal verification of general bipartite pure states*, arXiv:1901.09856 [quant-ph].

[81] K. Wang and M. Hayashi, *Optimal verification of two-qubit pure states*, arXiv:1901.09467 [quant-ph].

[82] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, S. D. Bartlett, O. Landon-Cardinal, D. Poulin, and Y.-K. Liu, *Efficient quantum state tomography*, Nat. Commun. **1**, 149 (2010).

[83] D. Hangleiter, M. Kliesch, M. Schwarz, and J. Eisert, *Direct certification of a class of quantum simulations*, Quantum Sci. Technol. **2**, 015004 (2017), arXiv:1602.00703 [quant-ph].

[84] L. Aolita, C. Gogolin, M. Kliesch, and J. Eisert, *Reliable quantum certification of photonic state preparations*, Nat. Commun. **6**, 8498 (2015), arXiv:1407.4817 [quant-ph].

[85] M. Gluza, M. Kliesch, J. Eisert, and L. Aolita, *Fidelity witnesses for fermionic quantum simulations*, Phys. Rev. Lett. **120**, 190501 (2018), arXiv:1703.03152 [quant-ph].

[86] G. Tóth, W. Wieczorek, D. Gross, R. Krischek, C. Schwemmer, and H. Weinfurter, *Permutationally invariant quantum tomography*, Phys. Rev. Lett. **105**, 250403 (2010), arXiv:1005.3313 [quant-ph].

[87] T. Moroder, P. Hyllus, G. Tóth, C. Schwemmer, A. Niggebaum, S. Gaile, O. Gühne, and H. Weinfurter, *Permutationally invariant state reconstruction*, New J. Phys. **14**, 105001 (2012), arXiv:1205.4941 [quant-ph].

[88] C. Schwemmer, G. Tóth, A. e. Niggebaum, T. Moroder, D. Gross, O. Gühne, and H. Weinfurter, *Experimental comparison of efficient tomography schemes for a six-qubit state*, Phys. Rev. Lett. **113**, 040503 (2014), arXiv:1401.7526 [quant-ph].

[89] A. Kalev, A. Kyrillidis, and N. M. Linke, *Validating and certifying stabilizer states*, Phys. Rev. A **99**, 042337 (2019), arXiv:1808.10786 [quant-ph].

[90] C. Bădescu, R. O'Donnell, and J. Wright, *Quantum state certification*, arXiv:1708.06002 [quant-ph].

[91] A. Montanaro and R. de Wolf, *A survey of quantum property testing*, Theory of Computing Graduate Surveys, **7**, 1 (2016), arXiv:1310.2035 [quant-ph].

[92] U. Schollwöck, *The density-matrix renormalization group in the age of matrix product states*, Ann. Phys. **326**, 96 (2011), arXiv:1008.3477 [cond-mat.str-el].

[93] J. D. Biamonte, S. R. Clark, and D. Jaksch, *Categorical tensor network states*, AIP Advances **1**, 042172 (2011), arXiv:1012.0531 [quant-ph].

[94] C. J. Wood, J. D. Biamonte, and D. G. Cory, *Tensor networks and graphical calculus for open quantum systems*, Quant. Inf. Comp. **15**, 0579 (2015), arXiv:1111.6950 [quant-ph].

[95] A. Jamiolkowski, *Linear transformations which preserve trace and positive semidefiniteness of operators*, Rep. Math. Phys. **3**, 275 (1972).

[96] M.-D. Choi, *Completely positive linear maps on complex matrices*, Lin. Alg. App. **10**, 285 (1975).

[97] I. Nechita, Z. Puchala, L. Pawela, and K. Zyczkowski, *Almost all quantum channels are equidistant*, J. Math. Phys. **59**, 052201 (2018), arXiv:1612.00401 [quant-ph].

[98] J. Watrous, *Simpler semidefinite programs for completely bounded norms*, Chicago J. Theo. Comp. Sci. **2013**, 1 (2013), arXiv:1207.5726.

[99] I. Roth, R. Kueng, S. Kimmel, Y. K. Liu, D. Gross, J. Eisert, and M. Kliesch, *Recovering quantum gates from few average gate fidelities*, Phys. Rev. Lett. **121**, 170502 (2018), arXiv:1803.00572 [quant-ph].

[100] A. J. Scott, *Optimizing quantum process tomography with unitary 2-designs*, J. Phys. A **41**, 055308 (2008), arXiv:0711.1017 [quant-ph].

[101] N. Hunter-Jones, *Unitary designs from statistical mechanics in random quantum circuits*, arXiv:1905.12053 [quant-ph].

[102] J. Emerson, R. Alicki, and K. Życzkowski, *Scalable noise estimation with random unitary operators*, J. Opt. B **7**, S347 (2005), arXiv:quant-ph/0503243.

[103] B. Lévi, C. C. López, J. Emerson, and D. G. Cory, *Efficient error characterization in quantum information processing*, Phys. Rev. A **75**, 022314 (2007), arXiv:quant-ph/0608246 [quant-ph].

[104] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, *Randomized benchmarking of quantum gates*, Phys. Rev. A **77**, 012307 (2008), arXiv:0707.0963 [quant-ph].

[105] E. Magesan, J. M. Gambetta, and J. Emerson, *Scalable and robust randomized benchmarking of quantum processes*, Phys. Rev. Lett. **106**, 180504 (2011), arXiv:1009.3639 [quant-ph].

[106] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, M. B. Ketchen, and M. Steffen, *Efficient measurement of quantum gate error by interleaved randomized benchmarking*, Phys. Rev. Lett. **109**, 080505 (2012), arXiv:1203.4550 [quant-ph].

[107] J. P. Gaebler, A. M. Meier, T. R. Tan, R. Bowler, Y. Lin, D. Hanneke, J. D. Jost, J. P. Home, E. Knill, and D. Leibfried, *Randomized benchmarking of multiqubit gates*, Phys. Rev. Lett. **108**, 260503 (2012), arXiv:1203.3733 [quant-ph].

[108] J. Emerson, M. Silva, O. Moussa, C. Ryan, M. Laforest, J. Baugh, D. G. Cory, and R. Laflamme, *Symmetrized characterization of noisy quantum processes*, Science **317**, 1893 (2007).

[109] J. J. Wallman and S. T. Flammia, *Randomized benchmarking with confidence*, New J. Phys. **16**, 103032 (2014), arXiv:1404.6025 [quant-ph].

[110] S. Kimmel, M. P. da Silva, C. A. Ryan, B. R. Johnson, and T. Ohki, *Robust extraction of tomographic information via randomized benchmarking*, Phys. Rev. X **4**, 011050 (2014), arXiv:1306.2348 [quant-ph].

[111] J. Wallman, C. Granade, R. Harper, and S. T. Flammia, *Estimating the coherence of noise*, New J. Phys. **17**, 113020 (2015), arXiv:1503.07865 [quant-ph].

[112] J. J. Wallman, M. Barnhill, and J. Emerson, *Robust characterization of loss rates*, Phys. Rev. Lett. **115**, 060501 (2015), arXiv:1412.4126.

[113] J. J. Wallman, M. Barnhill, and J. Emerson, *Robust characterization of leakage errors*, New J. Phys. **18**, 043021 (2016), arXiv:1412.4126 [quant-ph].

[114] A. W. Cross, E. Magesan, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, *Scalable randomised benchmarking of non-clifford gates*, npj Quant. Inf. **2**, 16012 (2016), arXiv:1510.02720 [quant-ph].

[115] A. Carignan-Dugas, J. J. Wallman, and J. Emerson, *Characterizing universal gate sets via dihedral benchmarking*, Phys. Rev. A **92**, 060302 (2015), arXiv:1508.06312 [quant-ph].

[116] E. Onorati, A. H. Werner, and J. Eisert, *Randomized benchmarking for individual quantum gates*, arXiv:1811.11775 [quant-ph].

[117] M. Horodecki, P. Horodecki, and R. Horodecki, *General teleportation channel, singlet fraction, and quasidistillation*, Phys. Rev. A **60**, 1888 (1999).

[118] M. A. Nielsen, *A simple formula for the average gate fidelity of a quantum dynamical operation*, Phys. Lett. A **303**, 249 (2002), quant-ph/0205035.

[119] R. Kueng, D. M. Long, A. C. Doherty, and S. T. Flammia, *Comparing experiments to the fault-tolerance threshold*, Phys. Rev. Lett. **117**, 170502 (2016), arXiv:1510.05653 [quant-ph].

[120] E. Magesan, J. M. Gambetta, and J. Emerson, *Characterizing quantum gates via randomized benchmarking*, Phys. Rev. A **85**, 042311 (2012), arXiv:1109.6887.

[121] J. J. Wallman, *Randomized benchmarking with gate-dependent noise*, Quantum **2**, 47 (2018), arXiv:1703.09835 [quant-ph].

[122] R. Harper, I. Hincks, C. Ferrie, S. T. Flammia, and J. J. Wallman, *Statistical analysis of randomized benchmarking*, Phys. Rev. A **99**, 052350 (2019), arXiv:1901.00535 [quant-ph].

[123] I. L. Chuang and M. A. Nielsen, *Prescription for experimental determination of the dynamics of a quantum black box*, J. Mod. Opt. **44**, 2455 (1997), quant-ph/9610001.

[124] A. Shabani, R. L. Kosut, M. Mohseni, H. Rabitz, M. A. Broome, M. P. Almeida, A. Fedrizzi, and A. G. White, *Efficient measurement of quantum dynamics via compressive sensing*, Phys. Rev. Lett. **106**, 100401 (2011), arXiv:0910.5498 [quant-ph].

[125] J. B. Altepeter, D. Branning, E. Jeffrey, T. C. Wei, P. G. Kwiat, R. T. Thew, J. L. O'Brien, M. A. Nielsen, and A. G. White, *Ancilla-assisted quantum process tomography*, Phys. Rev. Lett. **90**, 193601 (2003), quant-ph/0303038.

[126] M. Kliesch, R. Kueng, J. Eisert, and D. Gross, *Guaranteed recovery of quantum processes from few measurements*, Updated reference, key: KliKueEis19.

[127] A. V. Rodionov, A. Veitia, R. Barends, J. Kelly, D. Sank, J. Wenner, J. M. Martinis, R. L. Kosut, and A. N. Korotkov, *Compressed sensing quantum process tomography for superconducting quantum gates*, Phys. Rev. B **90**, 144504 (2014), arXiv:1407.0761 [quant-ph].

[128] M. Kliesch, R. Kueng, J. Eisert, and D. Gross, *Improving compressed sensing with the diamond norm*, IEEE Trans. Inf. Th. **62**, 7445 (2016), arXiv:1511.01513 [cs.IT].

[129] U. Michel, M. Kliesch, R. Kueng, and D. Gross, *Note on the saturation of the norm inequalities between diamond and nuclear norm*, IEEE Trans. Inf. Th. **64**, 7443 (2016), arXiv:1612.07931 [cs.IT].

[130] S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen, *Self-consistent quantum process tomography*, Phys. Rev. A **87**, 062119 (2013), arXiv:1211.0322 [quant-ph].

[131] R. Blume-Kohout, J. King Gamble, E. Nielsen, J. Mizrahi, J. D. Sterk, and P. Maunz, *Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit*, arXiv:1310.4492 [quant-ph].

[132] R. Blume-Kohout, J. K. Gamble, E. Nielsen, K. Rudinger, J. Mizrahi, K. Fortier, and P. Maunz, *Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography*, Nat. Comm. **8**, 14485 (2017), arXiv:1605.07674 [quant-ph].