École Normale Supérieure
Department of Mathematics and their Applications (DMA)

Error estimation in maximum likelihood reconstruction for quantum state tomography

Internship report

Thibaut Pérami

supervised by
Igor Dotsenko
and
Pierre Rouchon
in
LKB, Collège de France

September 2, 2019

Abstract

In this report, I explain the content of my internship with Igor Dotsenko and Pierre Rouchon at the Laboratoire Kastler Brossel in the Collège de France. I studied quantum electrodynamics in cavities, and in particular an experiment conducted by another intern, Luis Najera. He developed a way to transfer energy from a colder atom to a hotter cavity by using quantum manipulation with resonant Rabi oscillations. In order to extract main results from this experiment, we use quantum state tomography, in particular maximum log-likelihood estimation. In their original state of the art, when I arrived, the results of such estimations were only proven when observing usual quantum observable. However, Luis was studying quantum thermodynamics and thus needed to use among other the von-Neumann entropy of the reconstructed quantum states. I thus extended the proofs of such estimation to handle that case. Additionally, I also experimented with a method to solve a problem of insufficient measurement on a quantum system.

Notations

- λ, μ, x, w, n: scalar, generic vector and function are normal letters. Usually, I try to use greek letters for scalars, but this is not always the case.
- A, B, H: quantum operator or more generally matrices are capital letters
- $\boldsymbol{u}, \boldsymbol{E}$, 3D vectors or 3D quantum operator are in bold, capital letter for operators.
- $\mathcal{E}$: curved letter are generally for sets.
- $\mathbb{H}, \mathbb{E}$: Apart from usual sets ($\mathbb{R}, \mathbb{C}, \mathbb{N}, \mathbb{Z}$) doubled letters represent super-operator i.e operators from matrices to matrices.
- x^*: complex conjugate of x.
- A^t: transposed matrix of A.
- $A^\dagger$: transpose conjugate of A: $A^\dagger = (A^t)^*$.
- $x \cdot y$: The real scalar product $\sum_i x_i y_i$. On complex vectors it will still be the same formula and thus $x^* \cdot y = \langle x|y\rangle$.
- $\langle x|y\rangle$: complex scalar product, semi-linear to left i.e. physicist convention.
- $|x\rangle$, $\langle x|$: I use Dirac bra-ket notation. In addition, if A is a matrix, $|A\rangle$ will be a column vector representing the matrix.
- $x > y$: for two vector of $\mathbb{R}^n$, $x > y$ means component wise inequality
- $A > B$: $A - B$ is Hermitian positive definite. (semidefinite for $A \geqslant B$)
- $\mathcal{H}$: A Hilbert space
- $\mathcal{L}(\mathcal{H})$: The set of operator on $\mathcal{H}$
- $\mathcal{L}(\mathcal{H}, \mathcal{H}')$: The set of operator from $\mathcal{H}$ to $\mathcal{H}'$.
- $\mathcal{U}(\mathcal{H})$: The set of unitary matrice on $\mathcal{H}$ (orthogonal if $\mathcal{H}$ is a real vector space)
- $\mathcal{O}(\mathcal{H})$: The set of observable i.e of Hermitian matrice on $\mathcal{H}$
- $\mathcal{S}(\mathcal{H})$: The set of real symmetric matrices.
- $\mathcal{D}(\mathcal{H})$: The set of density matrices.
- $\mathbb{R}_+$, $\mathcal{O}_+(\mathcal{H})$: The set of non-negative real, of positive semi-definite hermitian matrices.
- ∇f: gradient on mono-result function, but also jacobian matrix on multi-result function
- $\nabla^2 f$: Hessian matrix
- $\frac{\partial f}{\partial x}, \partial_x f$: Partial derivative. If x is a vector, it is a partial gradient or a partial jacobian matrix.
- $\frac{\partial^2 f}{\partial x^2}$: 2nd partial derivative, but may be a partial Hessian if x is a vector

Contents

Introduction

This report details my internship with Pierre Rouchon and Igor Dotsenko on quantum state tomography by maximum-likelihood reconstruction (MaxLike). It was done mainly at the Laboratory Kastler-Brosel (LKB) in the Collège de France.

The goal of quantum state tomography is to find a way to reconstruct the state of a quantum system from several direct or indirect measurements on it. In some cases these measurements do not modify the state of the system and are called quantum non-destructive (QND) measurement.

The specific physical project on which I worked with Igor Dotsenko is the internship project of Luis Najera on quantum thermodynamics in the case of atom-cavity interaction in quantum optics. More precisely, we send a circular Rydberg atom in a resonant cavity containing several photon field. The state of the cavity is thermal i.e it follows the Bose-Einstein distribution. The atom then interacts with the cavity. If the atom is in a thermal state of its two atomic levels, that are at the same frequency than the cavity, then a thermal exchange will occur in the expected way. However if we pump one of the atom's state in a hidden state, like a Maxwell's demon, we can apparently break the second law of thermodynamics and make a cold atom give heat to a hotter cavity.

In order to study this experiment, we need to measure the state of the cavity. That state cannot be measured directly, so we reconstruct it by quantum state tomography using maximum-likelihood estimation. This construction takes into account all the measurements made and deduces a density matrix. To find it, one needs to solve a convex optimisation problem on the set of density matrices (Hermitian positive definite matrices of trace 1). This was done using a gradient descent method with projection of the gradient on the domain.

The estimated state is usually not a pure state because the reconstruction comes from a state that has been prepared several times. All the prepared states are slightly different. Furthermore, other factors like measurement imperfections and decoherent relaxation decrease the precision of the reconstruction.

With these measurements, we can now get and interpret the results of our Maxwell's demon experiment. However, to be able to interpret measurement, we need to know the error on the reconstructed state. The original proofs in [8] on tomography only provide the standard deviation of usual quantum observables which are linear in the density matrix. However, to perform our analysis of the second law, we need to evaluate the error of the estimated entropy.

In order to do that, I first tried a Monte-Carlo algorithm and used the distribution around the estimated value to sample the examined function. I could then deduce a mean and a standard deviation. Later, I improved the proofs of [8] to show that evaluating the function on the estimator is indeed asymptotically the average of the function around the estimator. I could then also get an analytic value of the standard deviation around any function of the density matrix. Additionally I also worked on fixing convergence problems when the information given by the experiment is incomplete.

In the end we get error bars on our plots and we are able to determine which points were, within their error bars, consistent with the theoretical description of our Maxwell's demon experiment and which were not.

This report alternates maths and physics. The odd chapters are about the physical part of the internship and the even ones about maths. The five first chapters are about my understanding of prior work that I had to do in order to do what I did. The last two describe my personal contributions.

In Chapter 1, I present the experiment, its formal modelisation and why we need the quantum state tomography. In Chapter 2, I explain the basics of maximum likelihood reconstruction and then I show in Chapter 3 how this method is applied to our particular case. In Chapter 4, I show how to solve the convex optimization problem that appeared. Then I present the results obtained before I arrived in Chapter 5. In Chapter 6, I explain what I did for computing error bars of a non-linear function and some other work. The final results are given in Chapter 7.

Chapter 1

Experimental setup and goal

In this chapter I expect the reader to know the basics of quantum mechanics, Hilbert spaces and the Schrödinger equation. I will present the basic quantum electrodynamics concepts needed to understand the experiment, and the experiment itself. Everywhere in my report I will use the usual bra-ket notation and, therefore, a scalar product linear to the right.

I will first schematically describe the experiment setup and then quickly dive into some experimental details. I'll finish by explaining why the state reconstruction is needed.

1.1 Basic setup description

In this section I describe the experiment I worked with at the LKB. This experiment studies the interaction between Rydberg circular atoms and light in superconducting cavities. We produce Rydberg atoms in a known circular state, then we send them through a cavity containing a very small number of photons (usually less than eight). The atom then goes through a detector that ideally performs projective measurement of the state. The atoms can be tuned via Stark effect so the interaction can be resonant (the cavity frequency correspond to a gap between atomic states) or dispersive (the frequency and the gap do not match).

I will present the formalism for each component and then detail the formalism of their interaction. The goal of the experiment is to study quantum thermodynamics. In particular, with the right manipulation of the quantum state of an atom, we can force it to give heat to the cavity (emit a photon) even if the atom is colder than the cavity. I will define the notion of temperature of a single atom later.

Most of the formulas come from the chapter 3 of [3] and the first part of [6].

A diagram of the experiment can be found in fig. 1.1.

1.1.1 Circular Rydberg atoms

The atoms are prepared in a specific circular Rydberg state called "ground" state $|g\rangle$. It corresponds to the energy level $n = 50$. We will also use $|e\rangle$ called "excited" state which is the circular state at $n = 51$ and $|f\rangle$, the "fundamental" state with $n = 49$. Ideally, the atom should never get in any other state during our protocol. The gap between $|e\rangle$ and $|g\rangle$ is approximately 51 GHz, whereas the gap between $|g\rangle$ and $|f\rangle$ is around 54 GHz. These are microwave frequencies.

In most of our calculation we will manipulate our states by pairs, and not as the whole triplet, because the atom will only interact with our setup at, or near, a resonance frequency. We'll use the base $(|e\rangle, |g\rangle)$ in that order, but the same is true for $(|g\rangle, |f\rangle)$. In such case, the Hamiltonian we will use is

$$H = \frac{\hbar\omega}{2}\sigma_Z,$$

where ω is the gap frequency and σ_Z in one of the three Pauli matrices used to describe a two-level system:

$$\sigma_Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \qquad \sigma_X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad \sigma_Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$$

It is important to note that the X, Y, Z axis here have nothing to do with the real world spatial axis. This is just a convention due to the fact that this two-level system behaves similarly to a spin-1/2 system whose axis do have spatial orientation.

From these operators we may also build the raising and lowering operators:

$$\sigma_+ = \frac{\sigma_X + i\sigma_Y}{2} = |e\rangle\langle g| = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \qquad \sigma_- = \frac{\sigma_X - i\sigma_Y}{2} = |g\rangle\langle e| = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$$

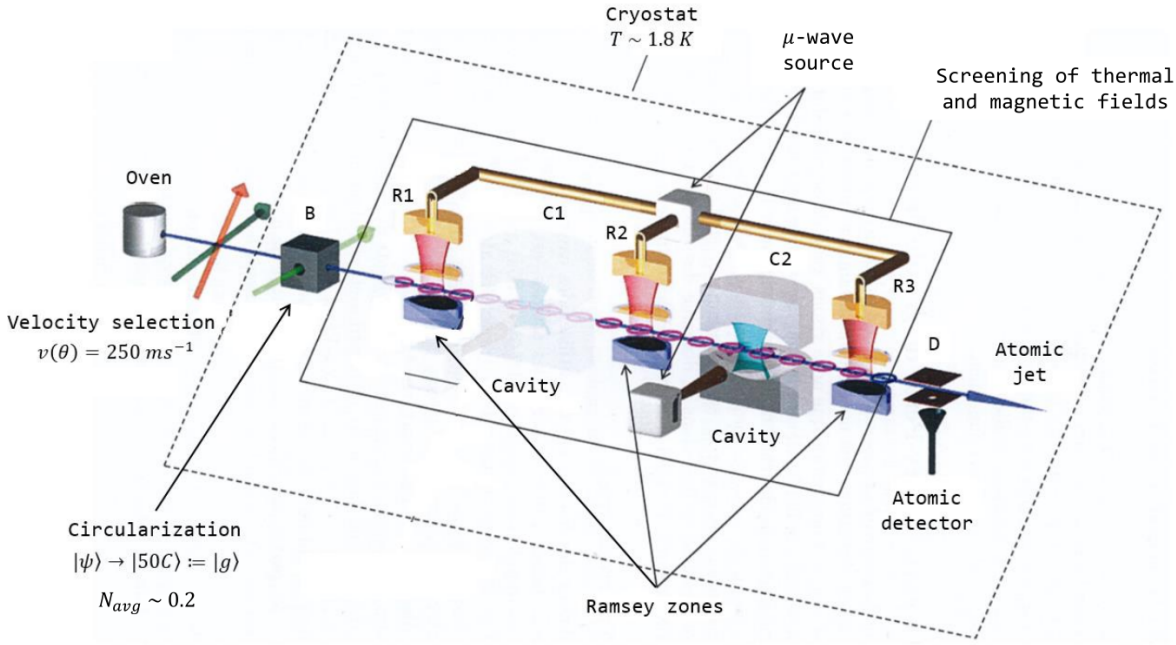

Figure 1.1: A symbolic representation of the experiment (Credit: Luis Najera). The atoms leave from the oven, then they pass through a velocity selection mechanism that ensure that only atoms moving at the target speed are selected. Then, there is a circularization procedure that put all atoms in the state $|g\rangle$ that will be explained in section 1.1.1. On average 0.2 atoms go through in each batch. Most unselected atoms will still go trough the experiment but as they are not Rydberg atoms, they will not interact with anything including the detector. The three Ramsey zones are areas where we manipulate the atomic state through a classical micro-wave field. In the cavity C2 (C1 is unused), a certain state of the electromagnetic field is trapped with sufficiently few photons to observe quantum effects in its interaction with the atom going through. At the exit, atoms are detected by a projective measurement that give the energy level of the atom.

They are equivalent to the annihilation and creation operators of a usual harmonic oscillator (that I will present in section 1.1.3). We can even express the Hamiltonian in a similar way: $H = \hbar\omega(\sigma_+\sigma_- - \frac{1}{2})$.

1.1.2 Ramsey zone: Interaction with classical field

In what we call Ramsey zone, the atom will be submitted to a classical oscillating field at pulsation ω_f, resonant (or near the resonance) with the atomic transition. The field at the position of the atom will be denoted by

$$\boldsymbol{E} = 2\mathcal{E}\begin{pmatrix} u_x\cos(\omega_f t + \varphi + \varphi_x) \\ u_y\cos(\omega_f t + \varphi + \varphi_y) \\ u_z\cos(\omega_f t + \varphi + \varphi_z) \end{pmatrix}$$

where $u_x^2 + u_y^2 + u_z^2 = 1$. We do not merge the global phase φ into the component phase to be able to use it later. This can be represented more cleanly using $\boldsymbol{u}_f = \left(u_x e^{-i\varphi_x}, u_y e^{-i\varphi_y}, u_z e^{-i\varphi_z}\right)$:

$$\boldsymbol{E} = \mathcal{E}(\boldsymbol{u}_f\, e^{-i\omega_f t-\varphi} + \boldsymbol{u}_f^*\, e^{+i\omega_f t+\varphi}). \tag{1.1}$$

With that field we can now express the Hamiltonian H_f of the interaction between the electric dipole of the atom and the electric field. It is $H_f = -\boldsymbol{D}\cdot\boldsymbol{E}$ as proven in section 3.1.2 of [3]. The dipole operator can written in function of the position operator $\boldsymbol{R}$ of the electron relative to the nucleus. The exact expression is $\boldsymbol{D} = q\boldsymbol{R}$. In our circular states having cylindrical symmetry, we have $\langle e|\boldsymbol{R}|e\rangle = 0$ and $\langle g|\boldsymbol{R}|g\rangle = 0$, so we only care about the off-diagonal terms. We write the dipole operator as

$$\boldsymbol{D} = d\begin{pmatrix} \boldsymbol{0} & \boldsymbol{u}_a^* \\ \boldsymbol{u}_a & \boldsymbol{0} \end{pmatrix} = d(\boldsymbol{u}_a\sigma_- + \boldsymbol{u}_a^*\sigma_+), \tag{1.2}$$

such that $\boldsymbol{u}_a$ has unit length and expresses the direction and d is the value of the electric dipole moment. They are defined by $q\langle g|\boldsymbol{R}|e\rangle = d\boldsymbol{u}_a$. The field Hamiltonian has now the expression:

$$H_f = -d\mathcal{E}\begin{pmatrix} \boldsymbol{0} & \boldsymbol{u}_a^*\cdot(\boldsymbol{u}_f\, e^{-i\omega_f t-\varphi} + \boldsymbol{u}_f^*\, e^{i\omega_f t+\varphi}) \\ \boldsymbol{u}_a\cdot(\boldsymbol{u}_f\, e^{-i\omega_f t-\varphi} + \boldsymbol{u}_f^*\, e^{i\omega_f t+\varphi}) & \boldsymbol{0} \end{pmatrix}$$

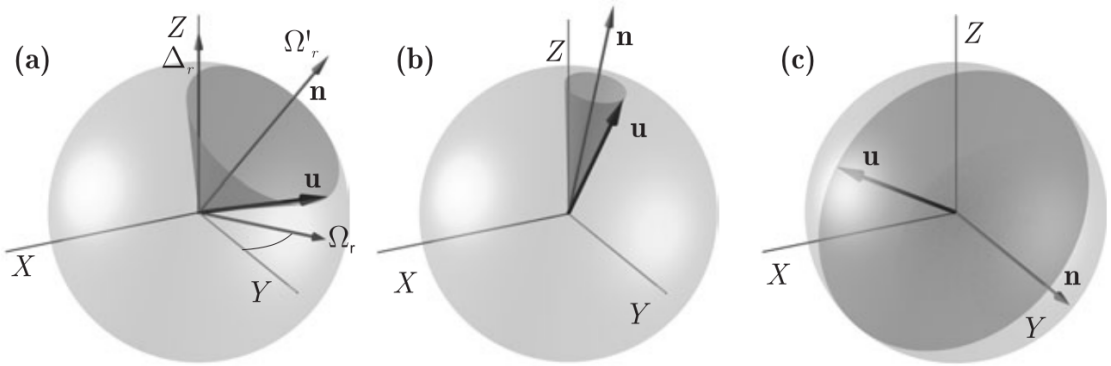

Figure 1.2: Precession of the state $\boldsymbol{u}$ on the Bloch sphere. (a) When Ω_r and Δ_r are of same magnitude. The angle at which Ω_r is placed is φ. (b) When Δ_r is large : dispersive case. (c) When $\Delta_r = 0$: full resonance case ($\varphi = 0$).

The free Hamiltonian of the system being $H_a = \frac{\hbar\omega_a}{2}\sigma_Z$, we'll need to express the evolution of the system by separating both evolutions. We will place ourselves in the interaction picture (explanation in appendix A). Both the free and interaction Hamiltonians make the state turn around, and we want to see the difference between the two of them. Therefore, we remove $H_0 = \frac{\hbar\omega_f}{2}\sigma_Z$. We remove the Hamiltonian with the frequency of the field and not of the atom to simplify the field frequencies elsewhere. The frequency of the atom in this new frame will be $\Delta_r = \omega_a - \omega_f$. Our new Hamiltonian $H_1 = e^{i\frac{H_0}{\hbar}t}(H_f + H_a - H_0)e^{-i\frac{H_0}{\hbar}t}$ is thus

$$H_1 = \frac{\hbar\Delta_r}{2}\sigma_Z - d\mathcal{E}\begin{pmatrix} \mathbf{0} & \boldsymbol{u}_a^* \cdot (\boldsymbol{u}_f\, e^{-i\omega_f t-\varphi} + \boldsymbol{u}_f^*\, e^{i\omega_f t+\varphi})e^{i\omega_f t} \\ \boldsymbol{u}_a \cdot (\boldsymbol{u}_f\, e^{-i\omega_f t-\varphi} + \boldsymbol{u}_f^*\, e^{i\omega_f t+\varphi})e^{-i\omega_f t} & \mathbf{0} \end{pmatrix}$$

The two terms oscillating at frequency $2\omega_f$ are too fast to have lasting effects, so we can neglect them: this is called the secular approximation. The new Hamiltonian is thus

$$H_1 = \frac{\hbar\Delta_r}{2}\sigma_Z - \hbar\frac{\Omega_r}{2}\big(e^{-i\varphi}\sigma_+ + e^{i\varphi}\sigma_-\big), \tag{1.3}$$

where $\Omega_r = \frac{2d}{\hbar}\mathcal{E}\boldsymbol{u}_a^* \cdot \boldsymbol{u}_f = \frac{2d}{\hbar}\mathcal{E}\langle \boldsymbol{u}_a | \boldsymbol{u}_f \rangle$. We can put some of φ in $\boldsymbol{u}_f$ in eq. (1.1) to make Ω_r real positive. Thus, the phase φ is specific to the phase difference between the atom and the field. It can be tuned by changing the phase of the field that we control.

When using the vector operator $\boldsymbol{\sigma} = (\sigma_X, \sigma_Y, \sigma_Z)$, our Hamiltonian can be written as

$$H_1 = \hbar\frac{\Omega'_r}{2}\boldsymbol{\sigma} \cdot \boldsymbol{n}, \tag{1.4}$$

where

$$\Omega'_r = \sqrt{\Delta_r^2 + \Omega_r^2} \qquad \text{and} \qquad \boldsymbol{n} = \frac{1}{\Omega'_r}(-\Omega_r \cos\varphi,\ -\Omega_r \sin\varphi,\ \Delta_r).$$

One can show that H_1 represents a precession of the state of the atom on the Bloch sphere around $\boldsymbol{n}$ at frequency Ω'_r. This precession is named "Rabi oscillation". and is represented on the Bloch sphere in fig. 1.2. The pulsation Ω_r which is the frequency of precession when the the frequencies match is named "Rabi frequency"

By controlling the relative phase between the field and the atom as well as the interaction duration (the vector rotates by an angle $\Omega'_r \Delta t$), we can realize all possible rotations i.e nearly all unitary operation on this two-level system. In practice, both parameters need to be carefully calibrated because they can have huge effects with very small variations.

1.1.3 Resonant cavities

A field mode in a cavity of a specific frequency ω_c can be modeled as a usual one-dimensional harmonic oscillator. Its Hilbert space is $\mathcal{H}_C = \mathbb{C}^{\mathbb{N}}$. The basis of this space is composed of vectors $\{|n\rangle\}$, called *Fock states*, with n non-negative integer. One of the important operators in this space is the number operator $N = \sum_n n|n\rangle\langle n|$. The state $|n\rangle$ represents a state with exactly n photons. There are also the annihilation a and creation $a^\dagger$ operators:

$$a = \sum_n \sqrt{n}|n-1\rangle\langle n|.$$

Therefore, $N = a^\dagger a$ and the energy of the cavity is then

$$H_c = \hbar\omega_c\big(N + \tfrac{1}{2}\big).$$

We need to find how to represent a field oscillating at frequency ω_c. A general formula for a specific mode of the classical field is

$$\boldsymbol{e}(r) = \mathcal{E}\big(\boldsymbol{f}(\boldsymbol{r})\alpha e^{i\omega_c t} + \boldsymbol{f}^*(\boldsymbol{r})\alpha^* e^{-\omega_c t}\big),$$

where $\mathcal{E}$ is a general normalisation factor in units of V/m, that we'll compute later. A polarization function $\boldsymbol{f}$ describes the spacial distribution of the electric field inside the cavity which depends on the cavity shape. We normalize it to $\|\boldsymbol{f}\|_\infty = 1$. Finally, α is the complex amplitude of the field.

We know that on classical harmonic oscillatiors, the eigenvectors of a, named *coherent state* represent directly classical states. I remind that for each complex number $\alpha \in \mathbb{C}$, there is a unique state (up to phase shift) $|\alpha\rangle$ such that $a|\alpha\rangle = \alpha|\alpha\rangle$. Let's prove that in the case of field mode oscillator, the coherent state also represents a classical state.

Property 1.1. *A coherent state $|\phi(0)\rangle = |\alpha\rangle$ evolves as:*

$$|\phi(t)\rangle = |\alpha e^{i\omega_c t}\rangle$$

Proof. In the section *time evolution* of the section 3.1.2 of [3] □

Therefore, we define the electric field operator to be

$$\boldsymbol{E}(\boldsymbol{r}) = \mathcal{E}\big(\boldsymbol{f}(\boldsymbol{r})a + \boldsymbol{f}^*(\boldsymbol{r})a^\dagger\big). \tag{1.5}$$

It will behave exactly as a classical field for coherent state and be a superposition of classical field for other states. We can determine the value of $\mathcal{E}$ from the field energy operator:

$$H = \int \varepsilon_0 |\boldsymbol{E}|^2 \mathrm{d}^3\boldsymbol{r}$$

Property 1.2. *The value we get for $\mathcal{E}$ is*

$$\mathcal{E} = \sqrt{\frac{\hbar\omega_c}{2\varepsilon_0 \mathcal{V}}}$$

where $\mathcal{V} = \int |f(\boldsymbol{r})|^2 \mathrm{d}^3\boldsymbol{r}$ is the effective mode volume

Proof. Just do $\langle n|H|n\rangle = \langle n| \int \varepsilon_0 |\boldsymbol{E}|^2 \mathrm{d}^3\boldsymbol{r}|n\rangle$ and expand everything. □

The value $\mathcal{E}$ is some kind of the average (more like r.m.s) field per photon in the cavity.

1.1.4 Atom-cavity interaction

In this section we'll see what happens when we put a two-level atomic system in a cavity at or near resonance frequencies. The systems are then coupled during all interaction time, in particular the states $|g, n+1\rangle$ and $|e, n\rangle$, blend together as if absorbing and emitting photons continuously. The model we use is called the James-Cumming model.

1.1.4.1 General case

In the general case, we just reuse the dipolar approximation Hamiltonian for the interaction:

$$H_i = -\boldsymbol{D} \cdot \boldsymbol{E}$$

By using eq. (1.2) and eq. (1.5), and assuming that the atom is at position $\boldsymbol{r}$, we get :

$$H_i = -d(\boldsymbol{u}_a\sigma_- + \boldsymbol{u}_a^*\sigma_+) \cdot \mathcal{E}\big(\boldsymbol{f}(\boldsymbol{r})a - \boldsymbol{f}^*(\boldsymbol{r})a^\dagger\big).$$

For sake of simplification, we assume the atom is more or less in the middle of the cavity and thus that $\boldsymbol{f}(r) = \boldsymbol{u}_c$ of unit length. If we do the same Heisenberg picture as in section 1.1.2, we'll see that terms associated with pure energy gain ($\sigma_+ a^\dagger$) or loss ($\sigma_- a$) oscillates very quickly at the order of $2\omega_c$, whereas the one corresponding to photon absorbtion ($\sigma_+ a$) or emission ($\sigma_- a^\dagger$) oscillates at a more reasonable frequency. We can thus do the secular approximation again to obtain:

$$H_i = \frac{\hbar\Omega_0}{2}(\sigma_+ a + \sigma_- a^\dagger), \tag{1.6}$$

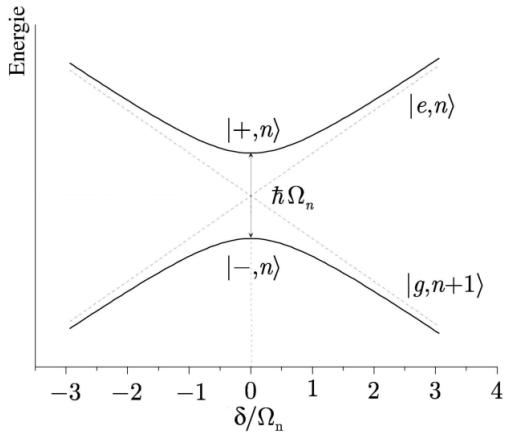

Figure 1.3: The energy of dressed states at level n as a function of the atom-cavity detuning δ (in units of $/\Omega_n$).

where the *vacuum Rabi frequency* reads:

$$\Omega_0 = -2\frac{\mathcal{E}d\langle \boldsymbol{u}_a|\boldsymbol{u}_c\rangle}{\hbar}.$$

Here again, for the sake of simplification, the cavity field and the atom are considered to be in phase, and thus $\langle \boldsymbol{u}_a|\boldsymbol{u}_c\rangle$ is real negative so that $\Omega_0 > 0$. The full Hamiltonian is then

$$H_{ac} = \frac{\hbar\omega_a}{2}\sigma_Z + \hbar\omega_c\big(N + \tfrac{1}{2}\big) + \frac{\hbar\Omega_0}{2}(\sigma_+ a + \sigma_- a^\dagger). \tag{1.7}$$

With this Hamiltonian, the states $|e,n\rangle$ and $|g,n+1\rangle$ are coupled 2 by 2, except for $|g,0\rangle$ which is stationary. On the space $\mathcal{S}_n = \mathrm{Span}(|e,n\rangle, |g,n+1\rangle)$, the Hamiltonian is

$$H_n = \hbar\begin{pmatrix}\omega_c(n+1) + \frac{\delta}{2} & \frac{\Omega_n}{2} \\ \frac{\Omega_n}{2} & \omega_c(n+1) - \frac{\delta}{2}\end{pmatrix}$$

with $\Omega_n = \sqrt{n+1}\,\Omega_0$. Its eigenvectors are given by

$$|+,n\rangle = \cos\frac{\theta_n}{2}\,|e,n\rangle + \sin\frac{\theta_n}{2}|g,n+1\rangle$$
$$|-,n\rangle = \sin\frac{\theta_n}{2}\,|e,n\rangle - \cos\frac{\theta_n}{2}|g,n+1\rangle$$

with $\theta_n \in [0,\pi]$ and $\tan\theta_n = \frac{\Omega_n}{\delta}$. These states are called *dressed states*. We can define again $\Omega'_n = \sqrt{\Omega_n^2 + \delta^2}$ to have the energies

$$E_{\pm,n} = \hbar\omega_c(n+1) \pm \frac{\hbar\Omega'_n}{2}.$$

Those energies are represented in fig. 1.3.

1.1.4.2 Effective interaction time

Let's go back to the hypothesis that $\boldsymbol{f}(\boldsymbol{r})$ is of unit length. According to [6], it is proven in [3] that, if the atom crosses the cavity horizontally, instead of actually caring about the exact form of the field, we can assume that everything behave like the atom traversed a full blown field with $\|\boldsymbol{f}(\boldsymbol{r})\| = 1$ that has an effective width w_{eff}. For an atom at speed v, all happens as if it is submitted to the Hamiltonian H_{ac} during

$$t_{\mathit{eff}} = \frac{w_{\mathit{eff}}}{v}.$$

1.1.4.3 Resonant case

If $\delta = 0$, the Hamiltonian is of the form $H = \hbar\omega_c(n+1) + H_1$ where H_1 is the Rabi Hamiltonian in eq. (1.4) and $\boldsymbol{n} = \boldsymbol{u}_x$. Therefore each pair of states at level n will oscillate around $\boldsymbol{u}_X$ at pulsation Ω_n. In particular if the cavity is empty and the atom arrives in state $|e\rangle$, the atom will absorb and emit a photon at pulsation Ω_0 which is thus the *Vacuum Rabi pulsation.*

1.1.4.4 Dispersive case

In the dispersive case i.e far off the resonance point we have a completely different behaviour. The energy the atom cannot really change because, the energy of the photons do not match the interstate gap. However, if we do a pertubative analysis in $\delta \gg \Omega_n$, we get some more usable results. For simplicity, we assume that $\delta > 0$.

In this analysis the new energies are

$$E_{\pm,n} = \hbar\omega_c(n+1) \pm \frac{\hbar}{2}\left(\delta + \frac{\Omega_n^2}{2\delta}\right),$$

but the in this situation we have:

$$|+,n\rangle \approx |e,n\rangle \qquad \text{and} \qquad |-,n\rangle \approx |g,n+1\rangle$$

If $E_{e,n} = \hbar\omega_c(n+\frac{1}{2}) + \frac{\hbar}{2}\omega_a$ is the normal energy of state $|e,n\rangle$, the new one will be

$$E_{+,n} = E_{e,n} + \frac{\Omega_0^2}{4\delta}(n+1).$$

and similarly for the other case:

$$E_{-,n} = E_{g,n+1} - \frac{\Omega_0^2}{4\delta}(n+1)$$

If the cavity is in state $|n\rangle$, The sum of the difference in pulsation of $|e,n\rangle$ and $|g,n\rangle$ is thus:

$$\Delta\omega_a = \frac{\Omega_0^2}{2\delta}\left(n + \frac{1}{2}\right)$$

Up to a global shift, all happens as if for an effective interaction time t_{eff}, we have a shift of phase $\phi_0 = \frac{\Omega_0^2}{2\delta}t_{eff}$ per photon. It is important to note that the most controllable parameter here is δ via the stark effect on the atom.

1.2 QND Measurement

Quantum Non-Destructive measurement(QND) is a kind of quantum measurement that tries to preserve the state of the observed system. By having the observed system interacting weakly with other systems, we can sometimes gather enough information to reconstruct the state.

The QND measurement in the current experiment is realized by sending a sequence of individual atoms prepared in a superposition state $|g\rangle + |e\rangle$ through the cavity and interacting with it dispersively. One can see in section 1.1.4.4 that an atom in this case will undergo a phase shift between its states proportional to the number of photons stored in the cavity field. If we calibrate our system to have an interesting ϕ_0, the atomic states will be dispersed accordingly on the equator of the Bloch sphere depending on the number of photons in the cavity. We can then do a $\frac{\pi}{2}$ Rabi-pulse at a specific phase to rotate the equator to a polar circle on the Bloch sphere. Then, when we measure the atom, the probabilities of being observed in $|e\rangle$ or $|g\rangle$ depends on the number of photons in the cavity. This setup is called a Ramsey interferometer. By varying the phase of the last Rabi pulse we can gather enough statistical information to recreate the state of the cavity.

However the only concrete output we get is a long list of state detections in $|e\rangle$ and $|g\rangle$, so we need a statistical tool to reconstruct a corresponding density matrix from this data.

Chapter 2

Maximum likelihood reconstruction

I put some reminders about multivariate covariances in appendix B, which could be useful in this chapter.

2.1 Basics of the estimation theory

2.1.1 Estimator

The goal of estimation is to estimate some parameters using observations or measurements. In the simplest case, we measure a random output $X \in \mathcal{X}$ from an unknown input $\theta \in \Theta$. We place ourselves in a model where the law of X conditionally to θ is known. Formally, we have for each θ a probability measure P_θ on $\mathcal{X}$. If P_θ has the right form (in particular, if X is a vector of repetition of the same experiment), we can extract information on θ from X. For that we can use an estimator.

Definition 2.1. An *estimator* of parameter θ from X is just a deterministic function $\hat{\theta}(X)$ that provides the most likely value of θ that could have resulted in X.

Remark. The estimator notation $\hat{\theta}$ has noting to do with the notation of quantum operator in quantum mechanics. It should usually be clear from the context which one it is.

The most frequent case of parameter is $\theta \in \mathbb{R}$. In that case we can easily compare the original θ with its estimator $\hat{\theta}$ with a subtraction:

Definition 2.2. For a given θ, the *bias* of the estimator $\hat{\theta}$ is

$$B(\hat{\theta}) = E_\theta(\hat{\theta}(X) - \theta),$$

where the expectation E_θ is taken on the law P_θ. An estimator is *biased* if the bias is non zero and it is *unbiased* otherwise.

Definition 2.3. For a given θ, the *variance* of the estimator $\hat{\theta}$ is

$$V(\hat{\theta}) = E_\theta\Big(\big(\hat{\theta}(X) - \theta\big)^2\Big).$$

2.1.2 Likelihood

In the case where all probability measures P_θ are absolutely continuous against a "canonincal" measure μ on $\mathcal{X}$, we can define a function $f(x;\theta)$ such that

$$P_\theta(A) = \int_A f(x;\theta)\,\mathrm{d}\mu(x).$$

Definition 2.4. For a given observation $x \in \mathcal{X}$, we define the *likelihood* function $\mathcal{L}_x(\theta) = f(x;\theta)$.

Definition 2.5. For a given observation $x \in \mathcal{X}$, we define the *log-likelihood* function $\ell_x(\theta) = \ln \mathcal{L}_x(\theta)$.

2.1.3 Fischer information and Cramér-Rao bound

We now assume a set of real parameter $\theta \in \mathcal{D} \subset \mathbb{R}^n$. A way to know if we could improve our estimation at a specific point is the gradient of the log-likelihood function. This is called the *score*, $s_x(\theta) = \partial_\theta \ell_x(\theta)$. The lower the score (its norm, actually), the better our estimation.

Property 2.1. *The expectation of the score is 0.*

We can then study the variance of score and its covariances with other things. In particular:

Lemma 2.2. *The covariance of the score and any unbiased estimator is the identity:*

$$\operatorname{cov}(\hat{\theta}, s_x(\theta)) = I$$

In the light of the theorem about bounds of appendix B.2, we can give special interest to the variance of the score and its link with the variance of any estimator.

Definition 2.6. The *Fischer information* of the parameter θ is defined as the variance of the score:

$$I(\theta) = V\big(s(\theta)\big).$$

If $\theta \in \mathbb{R}$, this is just the expectation of the square of the score, but otherwise, $I(\theta)$ is a covariance matrix thus at least positive semi-definite. This information value puts a bound on the minimum variance of any unbiased estimator, which is called the Cramér–Rao bound.

Theorem 2.3 (Cramér–Rao bound)**.** *For any unbiased estimator $\hat{\theta}$ and if $I(\theta) > 0$, we have*

$$V(\hat{\theta}) \geq I(\theta)^{-1}.$$

Remark. In the multivariate case, the $^{-1}$ obviously means the inverse of the matrix

Proof. The bound on multivariate correlation is proven in theorem B.5. The bound is written as

$$\operatorname{cov}(X, Y)V(Y)^{-1}\operatorname{cov}(Y, X) \leqslant V(X).$$

Since $I(\theta) = V(s_x(\theta))$ and from lemma 2.2, the bound holds. □

In order to have a better interpretation of the Fischer information, if the likelihood is at least $\mathscr{C}^2$ in θ, we can rewrite it as a hessian:

Property 2.4. *Under sufficient regularity assumption, the Fischer information can also be written as the opposite of the Hessian of the log-likelihood function:*

$$I(\theta) = -\frac{\partial^2 \ell}{\partial \theta^2}(\theta).$$

2.2 Maximum likelihood estimator

2.2.1 Definition

Now that we have proven the best variance possible, we still have to build an estimator approaching this bound. There are many different kinds of estimator. The most used one is the average when $X = (X_1, \dots, X_n)$, $E_\theta(X_i) = \theta$ and $V_\theta(X_i)$ is small enough. However, in more complex cases this won't work. The goal of the estimator is to produce the more likely θ, ideally the maximum of a density function $p(\theta|X)$.

If all measures can be expressed as density functions, using Bayes law, we obtain

$$p(\theta|x) = p_\theta(x)\frac{p(\theta)}{p(x)} = p_\theta(x)\frac{p(\theta)}{\int p_{\theta'}(x)p(\theta')\,\mathrm{d}\theta'}. \tag{2.1}$$

We aim now to maximize $p(\theta|x)$. However, that formula only makes sense if we have a prior law on θ. If the prior distribution is more or less uniform, it is the same as maximizing $p_\theta(x) = \mathcal{L}_x(\theta)$.

Definition 2.7. The *maximum likelihood estimator* θ_{ML} is defined by:

$$\theta_{ML}(x) = \arg\max \mathcal{L}_x.$$

The maximum likelihood have no excessively nice properties in its own. But when $X = (X_1, \dots, X_n)$, where all the X_i come from the same law $P_\theta(X_0)$ and are independent, and $n \to \infty$, it has nice convergence properties. In particular, it will converge in probability to the right θ and saturate the Cramér-Rao bound.

In practice, as most probability laws are log-concave (in particular the normal law), we use the log-likelihood, ℓ_x instead of $\mathcal{L}_x$. Apart from concavity, this also has the advantage that multiple repetition are additive instead of multiplicative:

$$\mathcal{L}_x(\theta) = \prod_i \mathcal{L}_{x_i}(\theta) \qquad \ell_x(\theta) = \sum_i \ell_{x_i}(\theta).$$

In particular, the probability goes down to zero very quickly. When computing on a 64-bit machine floating point numbers, there is a high risk of $\mathcal{L}$ to be simply evaluated to zero when it is just really small. Note that the natural logarithm of the minimum positive 64-bit floating point number is -740.

Chapter 3

Computation of effect matrices

In this chapter, we stay in Hilbert spaces of finite dimension. Most of the results can be extended to Hilbert spaces of countable dimension by adding some restrictions. Since we won't need them in the rest of the report, I won't mention infinite dimension anymore. The space of operators on $\mathcal{H}$ is denoted by $\mathcal{L}(\mathcal{H})$. The set of observables (Hermitian operators) is denoted by $\mathcal{O}(\mathcal{H})$. The set of unitary operators on $\mathcal{H}$ is denoted by $\mathcal{U}(\mathcal{H})$.

3.1 Density matrices

3.1.1 Justification

In order to make significant observation with an imperfect experimental setup, we need to make averages and statistics on quantum state. One way of doing that would be to manipulate probability measure over the state space which is the unit sphere S of the model Hilbert space $\mathcal{H}$. However, There are a lot of distributions that are indistinguishable by any measurement. In first approximation, we say that two probability distribution λ and μ on S are indistinguishable, if for any observable A on $\mathcal{H}$, we have

$$\int_S \langle\phi|A|\phi\rangle \,\mathrm{d}\lambda(\phi) = \int_S \langle\phi|A|\phi\rangle \,\mathrm{d}\mu(\phi).$$

We can rewrite it in

$$\int_S \langle\phi|A|\phi\rangle \,\mathrm{d}\lambda(\phi) = \int_S \mathrm{Tr}\left(|\phi\rangle\langle\phi|A\right) \mathrm{d}\lambda(\phi) = \mathrm{Tr}\left(\int_S |\phi\rangle\langle\phi| \,\mathrm{d}\lambda(\phi) A\right).$$

Definition 3.1. We define the density matrix of a probability distribution λ on quantum states:

$$\rho = \int_S |\phi\rangle\langle\phi| \,\mathrm{d}\lambda(\phi).$$

Property 3.1. *The average of an observable A in any probability distribution λ is given by* $\mathrm{Tr}(\rho A)$.

Remark. This is sufficient to fully characterise the distribution, because any observable A can be written as $A = \sum_a a P_a$, where P_a are orthogonal projector and are also observables. The average of P_a is $\mathrm{Tr}(\rho P_a)$ and is the probability of getting output a when measuring A. Therefore this probability is the same for all distributions having density matrix ρ.

Since two distribution that share the same density matrix are completely indistinguishable by any measurement, in the rest of this report, we will use density matrix as the only representation of quantum probability distribution. In practice, Quantum physicist mainly manipulate probabilistic quantum state and not that much usual *ket* state, so the term quantum state is used for a statistical mixture described by a density matrix. A state of the form $|\phi\rangle$ is thus called a *pure* state.

3.1.2 Properties

We can now study what are density matrices and find necessary and sufficient conditions for a matrix to be a density matrix.

Property 3.2. *A density matrix is Hermitian positive ($\rho > 0$) of trace 1.*

Proof. $|\phi\rangle\langle\phi| > 0$, so $\int_S |\phi\rangle\langle\phi| \,\mathrm{d}\lambda(\phi) > 0$. $\mathrm{Tr}(|\phi\rangle\langle\phi|) = 1$, so $\mathrm{Tr}(\rho) = \lambda(S) = 1$. □

If we want to prove that this is sufficient we need to look at the spectral decomposition. In finite dimension n, a spectral decomposition of ρ give a collection of eigen values $p_1, \dots, p_n$ which sum up to 1. They are the analogous of probabilities in classical probability theory. The density matrix can then be written as

$$\rho = \sum p_i |\phi_i\rangle\langle\phi_i|.$$

Theorem 3.3. *The condition 3.2 is necessary and sufficient i.e any Hermitian positive matrix of trace 1 is a density matrix.*

Proof. Any matrix satisfying the condition 3.2 can be decomposed in $\rho = \sum p_i |\phi_i\rangle\langle\phi_i|$, and thus the atomic measure $\lambda(\{|\phi_i\rangle\}) = p_i$ give exactly that density matrix. □

The set of density matrix will be denoted by $\mathcal{D}(\mathcal{H})$ (or just $\mathcal{D}$ when there is no ambiguity). We have the inclusions:

$$\mathcal{D}(\mathcal{H}) \subset \mathcal{O}(\mathcal{H}) \subset \mathcal{L}(\mathcal{H}).$$

As a density matrix is an observable, we can thus study its behavior as an observable. In state $|\Psi\rangle$, the average $\langle\rho\rangle$ is the probability of state $|\Psi\rangle$ if we did a projective measurement in an orthonormal base that contains $|\Psi\rangle$. ρ is thus like an observable of classical probability.

3.1.3 Rank and purity

In finite dimension, we can also look at the rank of the density matrix that gives a lot of information.

Property 3.4. *The rank of the density matrix ρ is the dimension of the support of any probability distribution that can give this matrix (with the convention $\dim A = \dim \mathrm{Span}(A)$).*

Definition 3.2. If a density matrix ρ has rank one, it only represents one unique state and can be written $\rho = |\phi\rangle\langle\phi|$. It is thus called a *pure state*. No difference is made between the state $|\phi\rangle$ and the density matrix $|\phi\rangle\langle\phi|$ as both are called pure states.

On the other hand if $\mathrm{rk}\,\rho > 1$, the density matrix represent a *mixed state*: a statistical mixture of states.

Remark. It is very important to differentiate a quantum superposition for a statistical mixture. For example on a qubit, $|+\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$ is a quantum superposition while a statistical mixture would be 50 % of $|0\rangle$ and 50% of $|1\rangle$. When the state in measured in base $|0\rangle, |1\rangle$, the results will be the same but in an other basis they could be different. In particular, in basis $\{|+\rangle, |-\rangle\}$, the first state will be measured as 100 % $|+\rangle$ and 0 % $|-\rangle$, but the second will still be 50 – 50 of $|+\rangle$ and $|-\rangle$.

Property 3.5. *A state ρ is pure iff* $\mathrm{Tr}\,\rho^2 = 1$.

Proof. $\mathrm{Tr}(\rho^2) = \sum p_i^2 < 1$ unless one p_i is one and the others 0. □

The value $\mathrm{Tr}\,\rho^2$ has in fact much more interesting properties, and thus deserves a name.

Definition 3.3. The *purity* of a quantum state represented by ρ is defined as $\gamma = \mathrm{Tr}\,\rho^2$.

Property 3.6. *In dimension d, the purity is always between $\frac{1}{d}$ and 1. If $\gamma = \frac{1}{d}$, then $\rho = \frac{I}{d}$, the represented mixture is uniform and we have absolutely no information on the state. If $\gamma = 1$, the state is pure and we thus have complete information on the quantum state.*

Proof. $\mathrm{Tr}(\rho^2) = \sum p_i^2 \leqslant 1$. But the projection of 0 on the hyperplane of sum 1 is $\frac{I}{d}$ which has therefore the lowest norm of that plane which is $\frac{1}{d}$. □

The main interest of purity is that it remains constant during unitary evolution ($\mathrm{Tr}(U^\dagger \rho U) = \mathrm{Tr}\,\rho^2$), whereas it will change in any other evolution (relaxation, decoherence, measurement etc) as we lose or gain information.

3.1.4 Matrix coefficients

When we express the density matrix in a specific base, we gain a lot of information on the state projected in that basis, in particular on how projective measurements in that basis takes place.

Definition 3.4. When ρ is express in the basis $|1\rangle, \dots, |d\rangle$, the diagonal coefficients, which are real and non-negative, are called *populations*. The coefficient $\langle i|\rho|i\rangle$ equals the probability to obtain state i, when measuring ρ in this basis.

The non-diagonal coefficients are complex and called *coherences*. The coherence between states $|i\rangle$ and $|j\rangle$ is $\langle i|\rho|j\rangle$. It is conjugate to the coherence between $|j\rangle$ and $|i\rangle$.

The point of this distinction is that any projective measurement made in that basis will have the same statistics whatever the value of coherences. Their values will only matter when measuring value in other bases or when the state is evolving. This also means that, reciprocally, in the context of tomography, Any measurements made in that basis can bring no information on the coherence.

As we will reconstruct ρ numerically and will use this matrix representation, we need to study the structure of $\mathcal{D}(\mathcal{H})$ in this form. In particular:

Property 3.7. $\dim \mathcal{D}(\mathcal{H}) = d^2 - 1$ *as a real vector space.*

Proof. The (real) dimension of $\mathcal{O}(\mathcal{H})$ is d^2 because there are d diagonal coefficient and $\frac{d(d-1)}{2}$ complex off-diagonal coefficient so $d + d(d-1)$ coefficients in total. The constraint $\operatorname{Tr}\rho = 1$ removes one degree of freedom, so we only have $d^2 - 1$ dimensions. □

3.1.5 The case of a qubit

In the case of a qubit, with $\dim \mathcal{H} = 2$ (complex dimension), we know that the pure states, up to a phase shift, can be represented on the Bloch sphere. We would like to know how to represent mixed states. The dimension of $\mathcal{D}(\mathcal{H})$ is 3, so we can hope for a simple representation. Any density matrix from $\mathcal{H}$ can be written as

$$\rho = \frac{1}{2}I + S,$$

where $\operatorname{Tr} S = 0$. A basis of the hyperplane of zero-trace hermitian matrice can be given by Pauli matrices:

$$\sigma_Z = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad \sigma_X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad \sigma_Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}.$$

If we write $\boldsymbol{\sigma} = (\sigma_X, \sigma_Y, \sigma_Z)$, we obtain

$$\rho = \frac{1}{2}I + \boldsymbol{r} \cdot \boldsymbol{\sigma},$$

where $\boldsymbol{r} \in \mathbb{R}^3$ is a usual 3D vector. Furthermore, we have:

Property 3.8. $\rho \geqslant 0$ *if and only if* $\|\boldsymbol{r}\| \leqslant 1$.

Proof. $\det \rho = 1 - \|r\|^2$ □

We can then look at the purity γ in function of $\boldsymbol{r}$. We have:

Property 3.9. $\gamma = \operatorname{Tr}(\rho^2) = \frac{1}{2}(1 + \|\boldsymbol{r}\|^2)$

Corollary 3.9.1. ρ *is pure if and only if* $\|\boldsymbol{r}\| = 1$

Mixed states are thus represented by the Bloch "ball". The pure states are laying on its surface i.e. on the sphere. One can check that this representation matches the usual Block sphere representation of pure states. The maximally mixed state $\frac{I}{2}$ is in the center of the ball.

3.1.6 Composite systems and partial trace

Next, we model a density matrix of a composite system. For a pure separable state, this is a simple tensor product $|\Psi\Phi\rangle = |\Psi\rangle \otimes |\Phi\rangle = |\Psi\rangle|\Phi\rangle$. Luckily for us we have

$$(|\Psi\rangle|\Phi\rangle) \otimes (\langle\Psi|\langle\Phi|) = |\Psi\rangle\langle\Psi| \otimes |\Phi\rangle\langle\Phi|,$$

where $\otimes$ designate the tensor product of operators and matrices. Therefore the density matrix of the composite system is the tensor product of the density matrices of each system. We still have the same property of entanglement, i.e. the state is entangled if and only if the density matrix cannot be written as a tensor product.

What is interesting now is to study the statistical result of measurement performed on one of the two quantum subsystems $\mathcal{H}_1$ and $\mathcal{H}_2$. Let's take an observable A on $\mathcal{H}_1$ and its extension to $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$ which is $\tilde{A} = A \otimes I$. If we have a density matrix $\rho = \rho_1 \otimes \rho_2$, then the average of A is $\operatorname{Tr}(\rho\tilde{A}) = \operatorname{Tr}(\rho_1 A)\operatorname{Tr}(\rho_2) = \operatorname{Tr}(\rho_1 A)$. On a tensor product density matrix, the evaluation of $\tilde{A}$ only depends on the part of ρ in $\mathcal{H}_1$ and we would like to express that on any density matrix. In order to do that, we rewrite our computation as $\operatorname{Tr}(\rho\tilde{A}) = \operatorname{Tr}((\rho_1 \operatorname{Tr}(\rho_2))A)$ and see that $\rho_1 \operatorname{Tr}(\rho_2)$ is the part of ρ in $\mathcal{H}_1$.

Definition 3.5. Given a tensor product $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2$, The *partial trace over* $\mathcal{H}_2$, is a linear operator $\mathrm{Tr}_{\mathcal{H}_2} : \mathcal{L}(\mathcal{H}) \to \mathcal{L}(\mathcal{H}_1)$ defined on tensor product by

$$\mathrm{Tr}_{\mathcal{H}_2}(\rho_1 \otimes \rho_2) = \rho_1 \,\mathrm{Tr}(\rho_2)$$

and extended by linearity on the full $\mathcal{L}(\mathcal{H})$.

Property 3.10. *For any operator $\rho \in \mathcal{L}(\mathcal{H})$, and any observable $A \in \mathcal{O}(\mathcal{H}_1)$, we have:*

$$\mathrm{Tr}(\rho(A \otimes I)) = \mathrm{Tr}(\mathrm{Tr}_{\mathcal{H}_2}(\rho)A)$$

Proof. It's true if ρ is a tensor product and so it is true on any ρ by linearity. □

An important remark to be made is that if a pure state in entangled, its partial trace may be a mixed state. For example if we have two qubits, and we have a Bell state $|\Psi\rangle = \frac{|00\rangle+|11\rangle}{\sqrt{2}}$, then

$$\mathrm{Tr}_2\left(|\Psi\rangle\langle\Psi|\right) = \frac{1}{2}I.$$

3.2 Quantum operations

Even if a closed quantum system follows a unitary evolution, quantum systems are in general open. They undergo interaction with their environment, whether it is from external measurement or simply thermal relaxation. The goal of this section is to construct a representation of acceptable transformations for open systems similar to those describing unitary evolution of a closed system. This section is strongly inspired from [5].

3.2.1 Projective measurement

The first and simplest example of external interaction in quantum theory is a measurement. When a quantum system interacts with a macroscopic measurement device, it is measured and collapsed in a specific state. The usual framework for measurement is projective measurement. It is made by measuring a observable that fully describes the measurement process.

Let $A \in \mathcal{O}(\mathcal{H})$ be an observable on $\mathcal{H}$. The spectral decomposition of A tells us that $A = \sum_a aP_a$, where a are eigenvalues of A and P_a are orthogonal projectors on the corresponding eigenvector spaces. When starting in state $|\Psi\rangle$, we measure a with probability $\langle\Psi|P_a|\Psi\rangle$ and we get the output state

$$\frac{P_a|\Psi\rangle}{\sqrt{\langle\Psi|P_a|\Psi\rangle}}.$$

In fact, P_a is also a quantum measurement operator which detects if we measure a or not, and thus its average value is the probability of measuring a. We can then express the same thing in the space of density matrices. When starting in ρ, we expect a with probability $\mathrm{Tr}(\rho P_a)$ and get as an output state

$$\frac{P_a\rho P_a}{\mathrm{Tr}(\rho P_a)}.$$

An important question when dealing with probability is whether they always sum to one, and in this case, this is answered by the closure relation

$$\sum_a P_a = I$$

and as $\langle\Psi|I|\Psi\rangle = \mathrm{Tr}(\rho I) = 1$, the probabilities are normalized.

3.2.2 Generalized measurement

In some case quantum detectors can perform other kinds of measurement, not projective. For example, a measurement that completely destroys the object instead of projecting it in a specific vector space. The main example is a photon counting detector. When a cavity is in a Fock state $|n\rangle$, and we insert a photon counting detector, the photons "crash" into the detector and are counted to be n. The remaining state is always $|0\rangle$, as there is no photon left. This measurement is obviously non projective as $|0\rangle$ is orthogonal to any $|n\rangle$ with $n > 0$. We thus need a framework to handle this kind of measurement.

While a projective measurement of operator N, would have left the state in state $|n\rangle$ because we applied $P_n = |n\rangle\langle n|$. The operation the detector apply to the system when measuring n is $M_n = |0\rangle\langle n|$. If we apply it to any state $|\Psi\rangle$, we have to normalize $M_n|\Psi\rangle$. The squared norm is $\langle\Psi|M_n^\dagger M_n|\Psi\rangle$ and by analogy, this looks like our probability. However, not any group of M_n can represent a generalized measurement. We need a normalization condition which is quite simple given the probability expression: $\sum M_n^\dagger M_n = I$.

Definition 3.6. On a quantum Hilbert space $\mathcal{H}$, a *generalized measurement* is a set of outcomes μ of the measurement (there can be more outcomes that the dimension of $\mathcal{H}$), and for each of them a measurement operator M_μ. For the measurement to be complete, we require the normalization condition:

$$\sum_\mu M_\mu^\dagger M_\mu = I.$$

In that case the probability of outcome μ in state $|\Psi\rangle$ is $\langle\Psi|M_\mu^\dagger M_\mu|\Psi\rangle$ and the resulting state is:

$$\frac{M_\mu|\Psi\rangle}{\sqrt{\langle\Psi|M_\mu^\dagger M_\mu|\Psi\rangle}}.$$

Remark. A projective measurement is just a special case of generalized measurement with $M_\mu = P_\mu$ as $P_\mu^\dagger P_\mu = P_\mu$.

Property 3.11. *If we have a mixed state ρ, the action of this generalized measurement is to measure μ with probability* $\mathrm{Tr}(M_\mu \rho M_\mu^\dagger)$. *The outcome would then be*

$$\frac{M_\mu \rho M_\mu^\dagger}{\mathrm{Tr}\left(M_\mu \rho M_\mu^\dagger\right)}.$$

The question that is important now, is what happens when we forget the value of the measurement. We can have each outcome μ with the corresponding probability, so the resulting state is

$$\rho_f = \sum_\mu \mathrm{Tr}(M_\mu \rho M_\mu^\dagger)\frac{M_\mu \rho M_\mu^\dagger}{\mathrm{Tr}\left(M_\mu \rho M_\mu^\dagger\right)} = \sum_\mu M_\mu \rho M_\mu^\dagger.$$

3.2.3 Quantum operations

The goal of this section is to describe a formalism to express the evolution of an open quantum system in terms of density operators. Another order of presention with deeper analysis can be found in [5].

For the sake of generality, the input state $\mathcal{H}_i$ and the output system $\mathcal{H}_o$ will not be the same. In most cases, they will, but in some case the system under study is causally related to an input system but is not the same. What we want is an operator $\mathbb{E} : \mathcal{D}(\mathcal{H}_i) \to \mathcal{D}(\mathcal{H}_o)$. It maps the input mixed state into the output mixed state and represents the evolution.

First, we remark that if ρ is mixed state representing ρ_1 with probability p_1 and ρ_2 with probability p_2, then the output state must be $\mathbb{E}(\rho_1)$ with probability p_1 and $\mathbb{E}(\rho_2)$ with probability p_2. Therefore, in the general case, we have

$$\mathbb{E}(\sum p_i\rho_i) = \sum p_i\mathbb{E}(\rho_i)$$

and $\mathbb{E}$ thus preserves convex combinations. A better way of representing this is to say that $\mathbb{E} \in \mathcal{L}(\mathcal{O}(\mathcal{H}_i), \mathcal{O}(\mathcal{H}_o))$ i.e. $\mathbb{E}$ is linear on Hermitian operators. Any operator that preserves convex combinations on $\mathcal{D}(\mathcal{H})$ can be extended to a linear operator on Hermitian matrices (the proof is easy and not interesting). The opposite is true if the operator preserves the positiveness and the value of the trace.

However, in order to keep positivity in all cases when we have two maps on different systems, we must have $\mathbb{E}_1 \otimes \mathbb{E}_2$ to also preserve positivity. We now have all conditions for a good definition.

Definition 3.7 (Trace-preserving quantum operation)**.** Given an input space $\mathcal{H}_i$ and an output space $\mathcal{H}_o$, a linear map $\mathbb{E} : \mathcal{O}(\mathcal{H}_i) \to \mathcal{O}(\mathcal{H}_o)$ is said to be a *trace preserving quantum operation* (or quantum map) if:

- It preserves trace i.e $\mathrm{Tr} = \mathrm{Tr} \circ \mathbb{E}$.
- It is completely positive i.e, for any other space $\mathcal{H}_0$ with identity quantum operation $\mathbb{I}_0$, The map $\mathbb{I}_0 \otimes \mathbb{E} : \mathcal{O}(\mathcal{H}_0 \otimes \mathcal{H}_i) \to \mathcal{O}(\mathcal{H}_0 \otimes \mathcal{H}_o)$ preserves the positivity of Hermitian operators.

This definition is nice and clean, but not very easy to manipulate, so we'll need a more efficient representation to be able to manipulate them on computers. We will see that in section 3.2.5. Afterwards, in section 3.2.6, we'll see why this formalism can represent all sorts of open quantum evolutions, by giving a more precise existence to the environment.

3.2.4 Non-trace preserving quantum operations

Restraining ourselves to an operator preserving the trace can be a problem for some kinds of evolution. For example, if we look at a generalized measurement $(M_\mu)_\mu$, the operation of measuring and forgetting is

$$\mathbb{M}(\rho) = \sum_\mu M_\mu \rho M_\mu^\dagger.$$

However, when we remember the outcome μ, the operation is $\rho \mapsto \frac{M_\mu \rho M_\mu^\dagger}{\mathrm{Tr}(M_0 \rho M_0^\dagger)}$ which is *not* linear. The associated linear operation is $\mathbb{M}_\mu(\rho) = M_\mu \rho M_\mu^\dagger$ which satisfies all properties of definition 3.7 except preserving the trace. In fact $\mathbb{M}_\mu$ describes a process that is not sure to happen. The probability that $\mathbb{M}_\mu$ happens when we do the measurement on ρ is $p_\mu = \mathrm{Tr}(\mathbb{M}_\mu(\rho))$ and the resulting density matrix is $\mathbb{M}_\mu(\rho)/p_\mu$. Non-trace preserving quantum operation therefore represents operations where the observer has learnt some information on the system. Notice that here, the notion of an observer is a mix of the notion of drawing a value from the probability measure that the density matrix represents and the usual observer associated to quantum collapses. Therefore the generic definition for quantum operators is

Definition 3.8 (quantum operation)**.** Given an input space $\mathcal{H}_i$ and an output space $\mathcal{H}_o$, a linear map $\mathbb{E} : \mathcal{O}(\mathcal{H}_i) \to \mathcal{O}(\mathcal{H}_o)$ is said to be a *quantum operation* (or quantum map) if:

- It decreases the trace i.e $\mathrm{Tr} \circ \mathbb{E} < \mathrm{Tr}$.
- It is completely positive i.e, for any other space $\mathcal{H}_0$ with identity quantum operation $\mathbb{I}_0$, The map $\mathbb{I}_0 \otimes \mathbb{E} : \mathcal{O}(\mathcal{H}_0 \otimes \mathcal{H}_i) \to \mathcal{O}(\mathcal{H}_0 \otimes \mathcal{H}_o)$ preserves the positiveness of Hermitian operators.

Trace preserving operation represent deterministic processes that do not leak any information the observer while non trace preserving operations describe processes that may happen with some probability (the trace of the output) and for which the observer has learnt some classical information on the system. A *complete system of quantum operation* (or *measurement model*) $(\mathbb{M}_\mu)_\mu$ is a set of quantum operation such that $\sum_\mu \mathbb{M}_\mu$ is trace preserving. We can evaluate this system like a generalized measurement, The probability of the outcome μ is $\mathrm{Tr}(\mathbb{M}_\mu(\rho))$ and the output state is

$$\frac{\mathbb{M}_\mu(\rho)}{\mathrm{Tr}\left(\mathbb{M}_\mu(\rho)\right)}$$

We can now look at two complete system in sequence to get the following property.

Property 3.12. *Given $(\mathbb{M}_\mu)$ and $(\mathbb{K})_i$ two measurement models, The probability or measuring μ then i in state ρ is $\mathrm{Tr}(\mathbb{K}_i(\mathbb{M}_\mu(\rho)))$ and the output state is*

$$\frac{\mathbb{K}_i(\mathbb{M}_\mu(\rho)))}{\mathrm{Tr}\left(\mathbb{K}_i(\mathbb{M}_\mu(\rho))\right))}$$

This obviously works for more than two systems and that's why we can compose partial operator without renormalizing each time and only care about renormalizing at the end. The behavior is exactly the same as if we had a single system of operator indexed on (μ, i) where $\mathbb{E}_{\mu,i} = \mathbb{K}_i \circ \mathbb{M}_\mu$.

3.2.5 Kraus operators, Kraus Maps

Until now, we have only manipulated abstract definition for quantum operations, let's see how to represent them numerically. In fact we already know a numerical formula for a certain kind of operation: unread generalized measurements. The form we know is $\mathbb{M}(\rho) = \sum \mathbb{M}_\mu(\rho) = \sum M_\mu \rho M_\mu^\dagger$. In fact we can prove this form is general and is called *Krauss decomposition* or *Krauss map*.

Theorem 3.13 (Krauss)**.** *An application $\mathbb{E} : \mathcal{O}(\mathcal{H}_i) \to \mathcal{O}(\mathcal{H}_o)$ is a quantum operation if and only if there exist a set of operators $E_i \in \mathcal{L}(\mathcal{H}_i, \mathcal{H}_o)$ such that $\sum E_i E_i^\dagger \leq I$ and*

$$\mathbb{E}(\rho) = \sum_i E_i \rho E_i^\dagger.$$

Proof. Can be found in [5]. □

Property 3.14. *The operation is trace preserving if and only if* $\sum_i E_i E_i^\dagger = I$.

Any trace preserving quantum operators can thus be expressed as an unread quantum measurement. It is like if all the decoherence processes were due to unread measurements from the environment.

3.2.6 Interpretation with environment

In fact we can also express quantum operations as a unitary evolution with a environment and then a partial trace. This approach is explained at the beginning of [5].

3.3 Likelihood in terms of quantum maps

Now it is time to use our new tools to give an usable expression for the likelihood function. This transformation comes from [7].

The experiment can be represented as a starting state ρ that we want to determine with tomography. This state evolves according to the experimental setup and thanks to some measurements made by the experimenter. For a given run of the experiment, the result are $y_1, \ldots, y_m$. The state ρ_i will be the state just after the measurement i. The evolution from ρ_{i-1} to ρ_i is given by a quantum operation that depends on the result of the measurement:

$$\rho_i = \frac{\mathbb{K}_{i,y_i}(\rho_{i-1})}{\mathrm{Tr}\left(\mathbb{K}_{i,y_i}(\rho_{i-1})\right)}$$

The system of $\mathbb{K}_{i,y_i}$ is complete for a given i. According to property 3.12, The probability of getting the results $y_1, \ldots, y_n$ is

$$P(y_,, \ldots, y_n) = \mathrm{Tr}(\mathbb{K}_{m,y_m} \circ \ldots \circ \mathbb{K}_{1,y_1}(\rho)).$$

We can use the adjoint maps for the Frobenius product (defined by $\mathrm{Tr}(A\mathbb{K}(B)) = \mathrm{Tr}(\mathbb{K}^*(A)B)$) to get

$$P(y_,, \ldots, y_n) = \mathrm{Tr}(\rho\mathbb{K}^*_{1,y_1} \circ \ldots \circ \mathbb{K}^*_{m,y_m}(I)) = \mathrm{Tr}(\rho E_{y_1,\ldots,y_n}).$$

The matrix E is called the *effect matrix* of the measurement $y_1, \ldots, y_n$.

Property 3.15. *The adjoint of a quantum operation is completely positive so all the E_i are positive hermitian matrices.*

If we have n sequence of measurement, we obtain n effect matrices $E_1, \ldots, E_n$ that represent each measurement. The log-likelihood function is thus

$$\ell(\rho) = \sum_i \log(\mathrm{Tr}(\rho E_i)). \tag{3.1}$$

We just need to be able to compute the $\mathbb{K}_{i,y_i}$. To compute the adjoint, we simply use the representation as Krauss operators, indeed if $\mathbb{K}(\rho) = \sum_\mu M_\mu \rho M_\mu^\dagger$ then

$$\mathbb{K}^*(\rho) = \sum_\mu M_\mu^\dagger \rho M_\mu.$$

3.4 Computation of Krauss operators for QND measurement

The actual computation of the Krauss operators is done in the supplementary information section of [4]. They take into account the detection errors, the decoherence of the cavity and of course all the theory about quantum Rabi oscillations.

Chapter 4

Convex optimization

Now we need to find the maximum ρ_{ML} of ℓ, given that ρ_{ML} is in $\mathcal{D}$, the set of density matrix. The problem we want to solve is

$$\begin{cases} \textbf{maximize} & \ell(\rho) & \textbf{on } \rho \\ \textbf{subject to} & \rho \in \mathcal{D} \end{cases}$$

This problem is "easy" because it is convex and as we'll see, convex problems have fast solutions.

4.1 Reminders on convex optimization

We must now define what is a convex problem, I assume the readers know what is a convex set. We must now define convex cones. This presentation is based on lecture notes from Alexandre d'Aspremont, in particular [2]

4.1.1 Convex cones

Definition 4.1. A set $K \subset E$ where E is a vector space on $\mathbb{R}$ is a *cone* if

$$\forall x \in K. \forall \lambda \in \mathbb{R}_+, \lambda.x \in K$$

The λ can be 0 in this definition i.e we must have $0 \in K$.

A convex cone is simply a cone that is convex. We are particularly interested in proper cones:

Definition 4.2. A cone $K \subset E$ is a *proper cone* if
- It is closed ($K = \overline{K}$)
- Is is solid ($\mathring{K} \neq \emptyset$)
- Is is pointed (It contains no lines)

In particular the set of vector with non negative coordinate (the positive orthant) or the set of positive hermitian matrices ($\mathcal{O}_+(\mathcal{H})$) are proper cones. With this notion of proper cone we can define a generalized inequality:

Definition 4.3. Two vectors a and b are said to be *inferior according to* K which is denoted by $a \leqslant_K b$ if $b - a \in K$. We say that $a <_K b$ if $b - a \in \mathring{K}$

In particular, the usual componentwise inequality on vectors is the the inequality of the positive orthant and the usual inequality on matrices is the inequality of $\mathcal{O}_+(\mathcal{H})$.

The last concept we need is the one of duals cones:

Definition 4.4. The *dual cone* K^* of a cone $K \subset E$ where E is euclidean or hermitian is the set

$$K^* = \{y \in \mathbb{R}^n,\ \forall x \in K,\ \langle y|x\rangle \geqslant 0\}$$

In the hermitian case, the $\geqslant 0$ mean real and non negative.

A cone is self dual if it is its own dual. In the literature, the hermitian case is sometimes called *internal dual cone*

Property 4.1. *The dual of a proper cone is a proper cone*

Property 4.2. *The cone of positive hermitian matrices is self dual.*

Proof. Suppose H is hermitian and $\langle H|A\rangle \geqslant 0$ for all $A \in \mathcal{O}_+(\mathcal{H})$. Then if we diagonalize H, and if there is any negative eigenvalue, the matrice which has as diagonal one in front of that eigenvalue and 0 in front of all the others will make the Frobenius product negative which is impossible, thus $H \geqslant 0$.

For a general matrice M, we have the Cartesian decomposition $M = A + iB$ with A and B hermitian. The fact that the Frobenius product is always real positive is sufficient to conclude that $B = 0$ and $A \geqslant 0$ □

4.1.2 Convex problems

I suppose that the reader knows the usual definition of a convex function. This definition can be generalized, in the multivariate case, to proper cone equality:

Definition 4.5. A function $f : \operatorname{dom} f \subset E \to F$ where E and F are real vector space is said to be K-convex for K a proper cone of F if $\operatorname{dom} f$ is convex and

$$\forall \lambda \in [0,1],\ \forall x, y \in \operatorname{dom} f, \lambda f(x) + (1-\lambda) f(y) \geqslant_K f(\lambda x + (1-\lambda) y).$$

If $K = \mathbb{R}_+$, this is the usual convexity. A convex problem is then a problem of the form

$$\begin{cases} \textbf{minimize} & f_0(x) & \textbf{on } x \in \mathbb{R}^n \\ \textbf{subject to} & f_i(x) \leqslant_{K_i} 0 & \\ & h_j(x) = 0 & \end{cases} \tag{4.1}$$

where f_0 is convex, each $f_i : \mathbb{R}^n \to E_i$ is K_i-convex for $i = 1, \ldots, m$ and each $h_j : \mathbb{R}^n \to \mathbb{R}$ are affine for $j = 1, \ldots, p$. Obviously, the domain isn't really $\mathbb{R}^n$ but $\bigcap_i \operatorname{dom} f_i$. However all convex functions can be extended by setting their values to ∞ outside their domain, so the distinction isn't really relevant (affine functions are always defined everywhere).

Definition 4.6. A point x is said to be *feasible* if it satisfies all the constraints in the "subject to".

Property 4.3. *The set of feasible points is convex.*

The *optimum* value of the problem is the minimal value reached or approached by f_0. It is noted f_{opt} in the rest of this section. In the approached case, it is the infimum of all the values taken by f_0 on the set of feasible points. Any x that reaches f_{opt} is called a *solution point*.

Property 4.4. *The set of solution points is always convex (but may be empty).*

We can also express maximization concave problems as is the case of the log-likelihood estimation:

$$\begin{cases} \textbf{maximize} & f_0(x) & \textbf{on } x \in \mathbb{R}^n \\ \textbf{subject to} & f_i(x) \geqslant_{K_i} & \\ & h_j(x) = 0 & \end{cases}$$

Where f_0 is concave, each f_i is K_i-concave and each h_j is affine.

4.1.3 Strong convexity

Definition 4.7. A function, $f : \mathbb{R}^n \to \mathbb{R}$ is strongly convex on a set A, if it is of class $\mathscr{C}^2$ on A and if there is a constant μ such that for any point of A, we have

$$\nabla^2 f \geqslant \mu . I$$

A problem is strongly convex if f_0 is strongly convex on the feasible set and this set is closed. This happens if all constraint functions are continuous for example.

Property 4.5. *A strongly convex problem has exactly one solution that is reached in a single point*

$$f_{opt} = f(x_{opt})$$

4.1.4 Dual problem

A dual problem for a given problem is like the problem viewed from the other side. If our main (primal) problem is a minimization, the goal of the dual problem is to produce the best possible lower bound. The dual will thus be a maximization problem. In order to do that, we first need to define the Lagrangian of the problem (which has nothing to do with the Lagrangian of classical mechanics).

Definition 4.8. Given a problem in form 4.1, We define its *Lagrangian* by the function

$$\begin{aligned} L: \quad & \mathbb{R}^n \times \prod E_i \times \mathbb{R}^p \ \to \ \mathbb{R} \\ & (x, \eta_1, \dots, \eta_m, \lambda_1, \dots, \lambda_p) \ \mapsto \ f_0(x) + \sum_{i=1}^{m} \langle \eta_i | f_i(x) \rangle + \sum_{j=1}^{p} \lambda_j h_j(x) \end{aligned}$$

If we want it to be under the minimum of the primal independently of x when the constraints are satisfied, we need to minimize on x.

Definition 4.9. The *Lagrange dual function* or simply *dual function* is defined by

$$g(\eta, \lambda) = \inf_{x \in \mathbb{R}^n} L(x, \eta, \lambda).$$

Property 4.6. *g is concave.*

Proof. Left to the reader or can be found in a convex optimization book. □

We can now reach our goal to have a lower bound independent of x:

Property 4.7. *If $\forall i \in \{1, \dots, n\}, \eta_i \geqslant_{K_i^*} 0$, and the optimum of the primal is f_{opt}, then*

$$g(\eta, \lambda) \leqslant f_{opt}$$

Proof. Let's take x any value respecting the constraints. We have

$$g(\eta, \lambda) \leq L(x, \eta, \lambda) \leq f_0(x)$$

The first inequality comes from minimization over all possible x. The second inequality comes from the fact that $f_i(x) \leqslant_{K_i} 0$ and $\eta_i \geqslant_{K_i^*} 0$ so, by definition of the dual cone, $\langle \eta_i | f_i(x) \rangle \leqslant 0$. Furthermore all $h_j(x) = 0$. If we minimize over all x, we get the property. □

If we want to find the best lower bound, we thus have the following convex problem to solve which is called the *dual problem*

$$\begin{cases} \textbf{maximize} & g(\eta, \lambda) \qquad \textbf{on } \eta, \lambda \\ \textbf{subject to} & \eta_i \geqslant_{K_i^*} 0 \end{cases} \tag{4.2}$$

4.1.5 Strong duality and KKT conditions

Here we study the link between the dual and the primal problem and try to find conditions for solutions and various other properties. Suppose we have a problem in standard form 4.1 with its optimum $\tilde{f}$ and its canonical dual 4.2 with its optimum $\tilde{g}$. What we have proven in the previous section is that

$$\tilde{g} \leqslant \tilde{f}$$

This is called *weak duality*. In some specific cases, we have $\tilde{g} = \tilde{f}$, this is called *strong duality*. The fact that a problem is strongly dual is very important for it's study and impact various properties. One way of proving that a convex problem is strongly dual is:

Property 4.8 (Slaters' constraint qualification)**.** *If the primal convex problem has a point that is in the interior of the set of feasible points i.e. a point x that satisfies $f_i(x) <_{K_i} 0$ and $h_j(x) = 0$, then the problem is strongly dual.*

Property 4.9 (Complementary slackness)**.** *If strong duality holds and $\tilde{x}$ is optimal for the primal and $\tilde{\eta}, \tilde{\lambda}$ are optimal for the dual, then:*

$$g(\tilde{\eta}, \tilde{\lambda}) = L(\tilde{x}, \tilde{\eta}, \tilde{\lambda}) = f(\tilde{x})$$

and thus $\langle \tilde{\eta}_i | f_i(\tilde{x}) \rangle = 0$ for any i.

Proof.

$$f(\tilde{x}) = g(\tilde{\eta}, \tilde{\lambda}) \leqslant L(\tilde{x}, \tilde{\eta}, \tilde{\lambda}) \leqslant f(\tilde{x})$$

□

Theorem 4.10 (KKT conditions)**.** *The Karush-Kuhn-Tucker conditions, given x, η, λ are:*

- *x is feasible*
- *For all i, $\eta_i \geqslant_{K_i^*} 0$*
- *For all i, $\langle \eta_i | f_i(x) \rangle = 0$*
- *The gradient of the Lagrangian $\frac{\partial L}{\partial x}$ is 0:*

$$\nabla f_0(x) + \sum_{i=1}^{m} \nabla \langle \eta_i | f_i(x) \rangle + \sum_{j=1}^{p} \lambda_j \nabla h_j(x) = 0.$$

For a strongly dual convex problem, they are equivalent to the optimality of x for the primal and the optimality of η, λ for the dual.

Property 4.11. *If we have Slater's condition (4.8), x is optimal if and only if there exists a η, λ satisfying the KKT conditions.*

4.2 Solution characterization

4.2.1 Characterization

Now we can move on to the optimal solutions of the maximum likelihood problem:

$$\begin{cases} \textbf{maximize} & \ell(\rho) \quad \textbf{on } \rho \in \mathcal{D}(\mathcal{H}) \end{cases}$$

which literally means, "maximize the log-likelihood on density matrices". In standard form, this is:

$$\begin{cases} \textbf{maximize} & \sum_i \log(\operatorname{Tr}(\rho E_i)) \quad \textbf{on } \rho \in \mathcal{O}(\mathcal{H}) \\ \textbf{subject to} & \rho \geqslant 0 \\ & \operatorname{Tr} \rho = 1 \end{cases}$$

The optimization is on the vector space $\mathcal{O}(\mathcal{H})$. This problem satisfies Slater's condition (4.8) as the identity matrix is alway in the interior of $\mathcal{D}$. Therefore all results for previous section are available, in particular we can characterize the solution using the KKT conditions (4.10).

Property 4.12. *ρ_{ML} is an optimal solution if and only if there exists $\eta \geqslant 0$ and λ such that*

$$\langle \eta | \rho_{ML} \rangle = 0 \quad \textit{and} \quad \nabla \ell + \eta = \lambda I$$

Proof. Using property 4.11 and $\dfrac{\partial \langle \eta | \rho \rangle}{\partial \rho} = \eta$. □

4.2.2 Uniqueness

The simplest case to ensure the solution exists and is unique, is to check if the problem is strongly convex:

Property 4.13. *If $\operatorname{Span}_i(E_i) = \mathcal{O}(\mathcal{H})$, then the problem is strongly convex.*

Proof. In the interior of $\mathcal{D}$, the gradient of ℓ is:

$$\nabla \ell = \sum_i \frac{E_i}{\operatorname{Tr}(\rho E_i)} \tag{4.3}$$

and so its hessian is:

$$\nabla^2 \ell(X, Y) = -\sum_i \frac{\operatorname{Tr}(X E_i) \operatorname{Tr}(Y E_i)}{\operatorname{Tr}(\rho E_i)^2} \tag{4.4}$$

which is obviously negative ($\nabla^2 \ell(X, X) < 0$ for any X). □

However if $\operatorname{Span}(E_i)$ is not full, we still have some properties: exact solutions will exist and fill an affine sub-space of $\mathcal{D}$ orthogonal to $\operatorname{Span}(E_i)$. Actually it will be exactly the orthogonal of $\operatorname{Span}(E_i)$ projected on $\mathcal{D}$. Which solution to choose among those was one of the problem that was still unsolved when I arrived, I propose some potential solutions in section 6.3.

4.3 Projected gradient method

In this section the goal is to actually compute a solution ρ_{ML} to the max-log-likelihood problem up to numerical errors.

4.3.1 Algorithm

On unconstrained convex problems, the simplest way to find the optimum is to do a "gradient descent". It means that we go downward (toward lower f_0) following the gradient. The algorithm start from $x = x_0$ and is:

- Compute $\nabla f_0(\rho)$.
- Choose an ε
- Move by $h = -\varepsilon \nabla f_0(x)$, i.e. $x \leftarrow x + h$.

The choice of ε is complex and has important consequences. Some examples are:

- Choose always the same value
- Choose a slowly decreasing value independent of the rest.
- Choose the best value that minimize $t \mapsto f_0(x + t\nabla f_0(x))$.
- Choose the best value that minimize the second order approximation

$$t \mapsto f_0(x) + t\|\nabla f_c(x)\|^2 + \frac{t^2}{2}\nabla f_0^t \cdot \nabla^2 f_0 \cdot \nabla f_0$$

In practice the last one was chosen.

Property 4.14. *If f_0 is strongly convex, the gradient descent algorithm with second order approximation converges toward a solution for any unconstrained problem*

However out problem is not unconstrained. The solution must in the feasible set $\mathcal{D}$. One way of dealing with that is to project the new point $x + h$ back to $\mathcal{D}$. Projecting a point on a compact convex set is always doable. We'll see how to do it numerically in section 4.3.2

Property 4.15. *For a constrained strongly convex problem, the projected gradient descent method with second order approximation converges*

In practice we'll do "gradient ascent" as in our problem as we are maximizing a concave function. The full algorithm is thus:

- Compute $\nabla f_0(\rho)$ and $\nabla^2 f_0(\rho)$
- Choose an ε that maximizes (can be done in $O(1)$ time):

$$t \mapsto f_0(x) + t\|\nabla f_c(x)\|^2 + \frac{t^2}{2}\nabla f_0^t \cdot \nabla^2 f_0 \cdot \nabla f_0$$

- Move by $h = \varepsilon \nabla f_0(x)$.
- Project on $\mathcal{D}$. $x \leftarrow P_{\mathcal{D}}(x + h)$.

4.3.2 Projection

This algorithm is an adaptation of work done by Pierre Six in [9], made by Pierre Rouchon in the actual codebase.

For a given point $A \in \mathcal{O}(\mathcal{H})$, its projection on $\mathcal{D}$ is the point of $\mathcal{D}$ that is the nearest from A. We want to:

$$\begin{cases} \textbf{minimize} & \|A - \rho\| \quad \textbf{on } \rho \\ \textbf{subject to} & \rho \in \mathcal{D} \end{cases}$$

Let first diagonalize $A \in \mathcal{O}(\mathcal{H})$ into $A = UDU^\dagger$.
We have $\|A - \rho\| = \|D - U^\dagger \rho U\|$ but we have $U^\dagger \rho U \in \mathcal{D} \iff \rho \in \mathcal{D}$, thus our problem is equivalent to:

$$\begin{cases} \textbf{minimize} & \|D - \rho\| \quad \textbf{on } \rho \\ \textbf{subject to} & \rho \in \mathcal{D} \end{cases}$$

But any outdiagonal coefficient of ρ would only increase the distance, as we can remove all of them and stay in $\mathcal{D}$ (Warning: removing only part of them would not guaranty to stay in $\mathcal{D}$). Our problem is thus equivalent to (given that $D = \text{diag}(d)$):

$$\begin{cases} \textbf{minimize} & \|d - x\| \quad \textbf{on } x \in \mathbb{R}^n \\ \textbf{subject to} & x > 0 \\ & \|x\|_1 = 1 \end{cases}$$

This is therefore just projecting a point on a simplex of dimension $n-1$ in $\mathbb{R}^n$, a much simpler problem but still not trivial. This problem is still convex, and satisfies Slater's condition (4.8) because $x = \left(\frac{1}{n}, \dots, \frac{1}{n}\right)$ is strictly in the interior of the feasible set. If we minimize the squared, norm, we have a $\mathscr{C}^\infty$ problem, and thus we can characterize the solution with the KKT conditions:

- x is feasible
- $n \leqslant 0$. (dual feasibility)
- $n \cdot x = 0$. (complementary slackness)
- $2(x_i - d_i) + n + 2c\mathbf{1} = 0$ ($\nabla \mathcal{L} = 0$).

There are two cases: either $x_i > 0$ and $n_i = 0$ and thus $x_i = d_i - c$, or $x_i = 0$ and $d_i \leqslant c$. The c value cut the set of axis in two parts: those above and those below.

Theorem 4.16. *The following algorithm computes the right projection:*

- *Sort x into $d_{p(1)} \geqslant \cdots \geqslant d_{p(n)}$.*
- *Find i_0, the first index such that $c_{i_0} > d_{p(i_0+1)}$ with*

$$c_k = \frac{\sum_{i=1}^{k} d_{p(i)} \ - 1}{|S|}$$

- *For $i \leqslant i_0$, $x_{p(i)} = d_{p(i)} - c$*
- *For $i > i_0$, $x_{p(i)} = 0$.*

Proof. Property 4.11 ensures the uniqueness of the solution. We just need to find one point that satisfies the KKT conditions. One can check condition by condition, that the solution computed by this algorithm satisfies all KKT conditions. The only non-trivial point is that $d_{p(i_0)} \geqslant c_{i_0}$, and this is true because it is equivalent to $d_{p(i_0)} \geqslant c_{i_0-1}$ which is true because $i_0 - 1$ didn't satisfy the condition. □

Chapter 5

First Results

In this chapter I will present the actual experimental protocol we realized and the first result that were done before I arrived. Most of the work in this section is still unpublished and is the content of the internship of Luis Najera.

5.1 Extracting results and basic errors

Thanks to results of the previous chapter, we can now reconstruct ρ_{ML}, the density matrix that maximizes the log-likelihood of our measurement. In the next section, I'll explain the exact experimental protocol we used. But first let's see how to use the result of the reconstruction. [8] proves that for any observable, as the number of measurements grows, the average value of an observable A is $\mathrm{Tr}(\rho_{ML}A) + O(\frac{1}{N})$. The standard deviation due to the dispersion of measurements around this value tends to be:

$$V_A = \mathrm{Tr}(A_{||}\mathbb{H}^{-1}(A_{||})) + O(\frac{1}{N^2}) \tag{5.1}$$

where:

- $\mathcal{D}_r$ is the submanifold of $\mathcal{D}$ of matrices with rank r
- P_{ML} is the orthogonal projector on the range of ρ.
- $A_{||}$ is the orthogonal projection on the tangent space to $\mathcal{D}_{\mathrm{rk}\,\rho}$ in ρ:

$$A_{||} = A - \frac{\mathrm{Tr}(AP_{ML})}{\mathrm{Tr}(P_{ML})}P_{ML} - (I - P_{ML})A(I - P_{ML})$$

- $\mathbb{H}$ is the hessian of $\ell_{|\mathcal{D}_{\mathrm{rk}\,\rho}}$ in ρ:

$$\mathbb{H}(A) = \sum_i \frac{\mathrm{Tr}(AE_{i||})}{\mathrm{Tr}^2(\rho_{ML}E_i)}E_{i||} + (\lambda I - \nabla\ell(\rho_{ML}))A\rho_{ML}^+ + \rho_{ML}^+A(\lambda I - \nabla\ell(\rho_{ML}))$$

  The ρ^+ is the moore-penrose pseudo inverse (inversion of non-zero eigen values). The λ is the one from property 4.12. The $\mathbb{H}$ is the Fischer information on the submanifold $\mathcal{D}_{\mathrm{rk}\,\rho}$ as proven in property 2.4

Please take care that this is not the standard deviation on quantum measurement due to quantum uncertainty, nor the deviation due to the probabilistic nature of ρ_{ML}. This deviation comes from the dispersion of the effect matrices E_i. If all the E_i are the same, $\rho_{ML} = \frac{E_i}{\mathrm{Tr}(E_i)}$ and $V_A = 0$ for any A. Furthermore it is a known problem of max-likelihood method that V_A does not measure the deviation due to how few experiment there are. If we have only one E_i, the max-likelihood algorithm will tell that $\rho_{ML} = \frac{E_1}{\mathrm{Tr}(E_1)}$ with 0 deviation. I didn't had time to address that in this internship and no one in my team did it before, so we just assume we have enough measurements so that the $O(\frac{1}{N})$ is small enough i.e significantly smaller than V_A.

Remark. The individual coefficients can be observed with some specific matrices. We can thus compute the deviation on each population and coherence individually. Those deviations are *not* independent.

5.2 Maxwell Demon experiment

5.2.1 Overview

The goal of this experiment is to study heat exchange and thermodynamics at a individual quantum level. In particular, with some quantum manipulations, we create a Maxwell Demon context where a colder atom can give heat to a hotter cavity.

The protocol is the following:

- The cavity is cleared by sending a bunch of Rydberg atoms in groud state $|g\rangle$ into it so that they absorb any remaining photons. The temperature of the cavity is then $T = 0\,\mathrm{K}$.
- A thermal field is built inside the cavity at a given initial temperature T_C. This temperature is different from the cryostat temperature, so it isn't stable (but the relaxing time is long enough).
- A single Rydberg atom is sent through the setup.
- In the first Ramsey zone, it is set to temperature T_a.
- In the second Ramsey zone, it may or may not undergo an action from the Maxwell demon system that make it interact with the fundamental level $|f\rangle$.
- The atom interacts with the second cavity at temperature T_C, emiting and absorbing photons.
- The atom is observed to be in either $|e\rangle$, $|g\rangle$ or $|f\rangle$.
- A bunch of non-resonant QND atoms are sent to observe the state of the cavity.
- We reconstruct the state of the cavity by maxlike estimation thanks to the QND measurements.

5.2.2 Thermal state

We first need to define what is a thermal state at temperature T. Ideally it is the statistical state that a quantum system reaches when interacting with a thermostat (a macroscopic system) at temperature T. According the Maxwell-Bolzman distribution, the probability to be in state i of energy E_i is:

$$p_i = \frac{e^{-\beta E_i}}{Z}$$

where $\beta = \frac{1}{k_B T}$, and Z is the partition function $Z = \sum e^{-\beta E_i}$. The thermal state is thus

$$\zeta_\beta = \sum_i p_i |i\rangle\langle i|$$

This can be expressed more compactly with

$$\zeta_\beta = \frac{e^{\beta H}}{\mathrm{Tr}\,(e^{\beta H})} \tag{5.2}$$

where H is the Hamiltonian. This is true for any quantum system of finite dimension.

In the case of a thermal qubit with energy gap E_g, we let $p_\beta(E_g) = \frac{e^{-\beta E_g}}{1+e^{-\beta E_g}}$ be the probability of being in the top state .

5.2.3 Atom state

The Rydberg atom is prepared in state ζ_β^S. This state is a thermal mix of $|e\rangle$ and $|g\rangle$. However, in practice atoms are in pure states. The Ramsey zone is calibrated to do just a rotation with the right angle to bring $|g\rangle$ to:

$$(1 - p_\beta(\hbar\omega_{eg}))|g\rangle + e^{i\phi} p_\beta(\hbar\omega_{eg})|e\rangle$$

The angle and thus the temperature is chosen by the experimenter. However the axis of the rotation is random, and thus the phase ϕ is also random. The average of those states is thus ζ_β^S.

However the goal of the experiment is to make the $|e\rangle$-$|g\rangle$ system interact with $|f\rangle$ in order to decompose those interaction. We'll artificially split our system into two qubits $\mathcal{H}^Q$ and $\mathcal{H}^D$, for the main system e–g and the demon system g-f respectively. The atom is thus represented by the System $\mathcal{H}^Q \otimes \mathcal{H}^D$. The mapping is:

$$|e\rangle = |1^Q 1^D\rangle \qquad |g\rangle = |0^Q 1^D\rangle \qquad |f\rangle = |0^Q 0^D\rangle$$

The system is built in such a way that the first qubit is only about the e-g interaction and the demon qubit is only about the g-f interaction. The state $|1^Q 0^D\rangle$ is purely fictitious and always has probability zero in all our calculations. This decomposition allows to speak about entanglement and other concepts that need multiple systems on a single atom.

The initial atom state is thus:

$$\zeta_\beta^Q \otimes |1^D\rangle\langle 1^D|$$

5.2.4 Maxwell demon: readout

The demon action happens just before the cavity to avoid spontaneous emissions or absorbtion of random thermal photons.

The goal of the demon is simply to send the state $|g\rangle$ to $|f\rangle$ in order to make only the excited state interact with the cavity. That way, the atom can only emit energy to the cavity and cannot receive energy from it. This is called readout because in terms of qubit, the system qubit is "copied" on the demon qubit. In order to make that unitary we also need to send $|f\rangle$ on $|g\rangle$ but no atom should be in $|f\rangle$ in the initial state. The corresponding quantum gate in the 2-qubit system is an inverted C-Not: The second bit is flipped if the first on is 0.

5.2.5 Cavity interaction: Feadback

Inside the cavity the atom undergo an adiabatic passage that forces the absorption or the emission of an atom. As the passage is done only around ω_{eg}, the state $|f\rangle$ does not interact at all with the cavity:

$$\begin{aligned} |g, n+1\rangle &\mapsto |e, n\rangle \\ |g, 0\rangle &\mapsto |g, 0\rangle \\ |e, n\rangle &\mapsto |g, n+1\rangle \\ |f, n\rangle &\mapsto |f, n\rangle \end{aligned}$$

If the demon qubit is set to one, the unitary operation applied to the QC system is:

$$U = |00\rangle\langle 00| + |1\rangle\langle 0| \otimes a + |0\rangle\langle 1| \otimes a^\dagger$$

where a is the annihilation operator of the cavity.

5.3 Quantum information

The thermodynamical analysis of a quantum system requires certain tools in particular Von-Neumann entropy:

Definition 5.1. The Von-Neumann entropy is a function $S : \mathcal{D} \to \mathbb{R}$ defined by:

$$S(\rho) = -\operatorname{Tr}(\rho \ln \rho)$$

Where ln is the principal logarithm on Hermitian positive definite matrices. It can be prolonged by continuity on semi-definite positive Hermitian matrices.

This quantity is the analog of Shanon entropy for classical systems ($H(p) = \sum_i -p_i \log p_i$). From that quantity we can build all the other information theory quantity like conditional entropy or mutual information with S as a building block. The equivalent of the marginal law in classical theory is the partial trace. In particular, the relative entropy is the analogous of the Kullback–Leibler divergence.

For more details I suggest consulting a book like [5].

5.4 Thermodynamical analysis

Luis Bajera, whose internship project motivated the work of my internship, had done some interesting work on the second law of thermodynamics. In particular he related the second law of thermodynamics to the notions of quantum information theory. His main result, that we wanted to test, is a rewrite of the second law of thermodynamics. Usually if two systems have temperature T_1 and T_2, are put into contact and exchange some heat Q from 1 to 2, we'll have:

$$\frac{Q}{T_2} - \frac{Q}{T_1} \geqslant 0.$$

If we speak about the inverse temperature β, we'll write:

$$\Delta\beta Q \geqslant 0.$$

But this inequality does not hold when the Maxwell demon is activated, because the fact that the demon learns information about the system somehow reduces the entropy in some way. The new inequality named *Weak second law* is then

$$\Delta\beta Q - \Delta I_{QC:D} \geqslant 0 \tag{5.3}$$

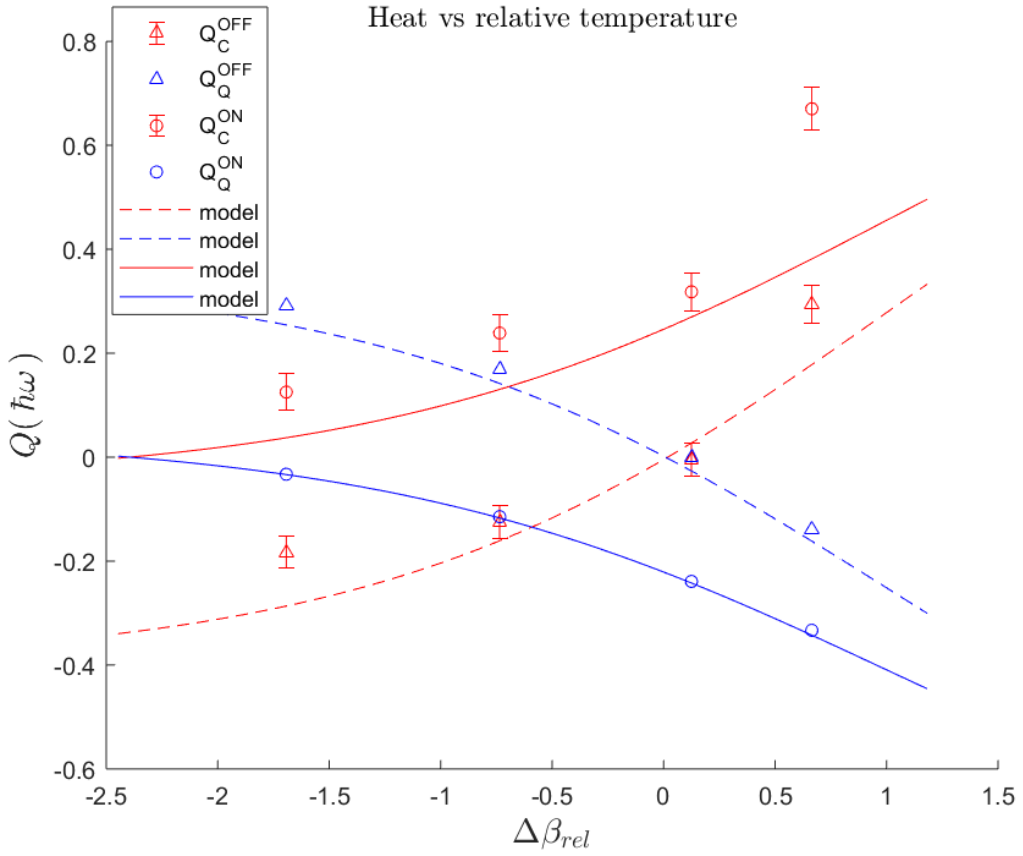

Figure 5.1: Plot of the exchanged heat depending on the Maxwell demon state (ON or OFF)

where $I_{QC:D}$ is the mutual information between, on one side the main atomic qubit Q and the cavity C, and on the other side the demon D. In fact Luis managed to find the exact value of this quantity to get what we call *strong second law*

$$\Delta\beta Q - \Delta I_{QC:D} = D_{QC} \tag{5.4}$$

where D_{QC} is the relative entropy of the system QC compared to its original thermal state $\zeta^Q_{\beta_Q} \otimes \zeta^C_{\beta_C}$.

5.5 Experimental results

The goal of the experiment was thus to check Luis' results, in particular the strong second law. The plots I will give here are the one that were made before I arrived to the lab. First we plot in fig. 5.1 the exchange of energy in function of the relative temperature $\Delta\beta$ from the point of view both systems. In theory $Q_C = -Q_Q$. In particular, when the demon is `OFF` the thermal equilibrium is at $\Delta\beta = 0$ whereas when it is activated it is shifted to the left. In that case, we are forcing the heat to move from the atom to the cavity. We can see that there are error bars on the heat because heat is simple quantum observable and thus its error can be evaluated thanks to the results of [8].

We can now plot the mutual information in fig. 5.2. In this one we see that demon acquires a lot of information in the readout phase and then somehow "use" it when arriving in the feedback phase.

We can then plot the weak second law in fig. 5.3, and the strong second law in fig. 5.4. In particular, we need error bars to determine it the point below 0 in fig. 5.3 and all the points off the 0 axis in fig. 5.4 are just random variations or actually mean something.

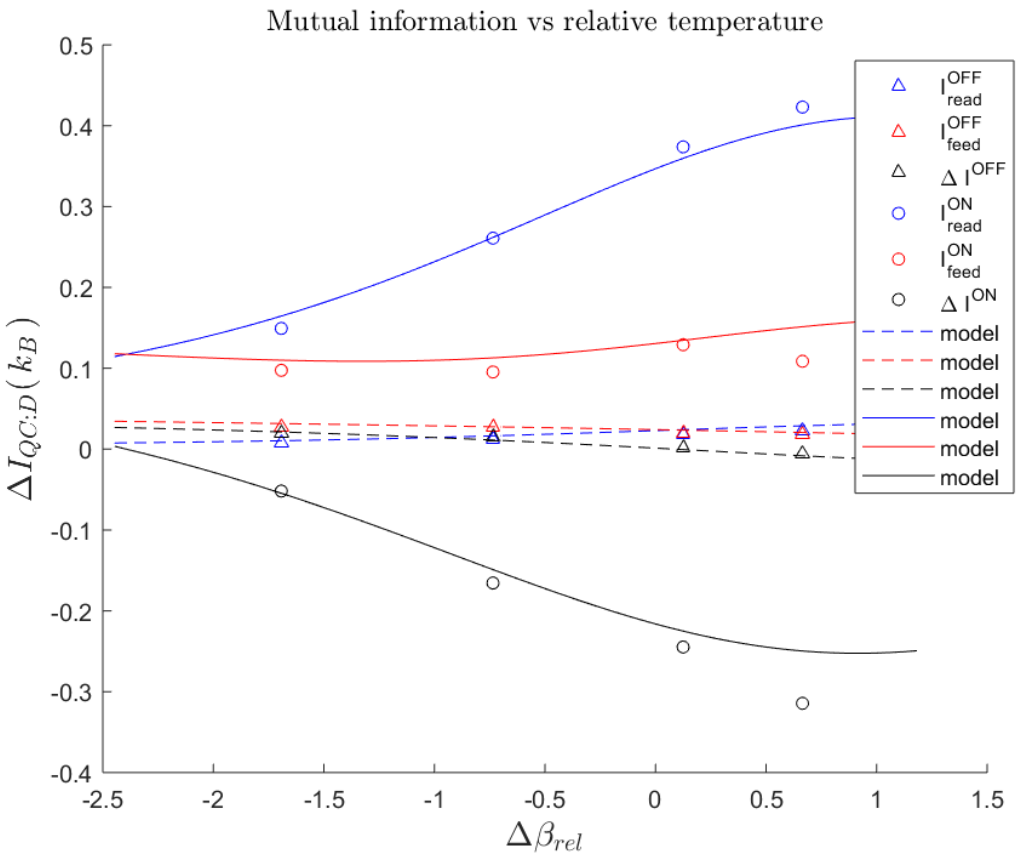

Figure 5.2: Plot of the mutual information in the readout phase vs the feedback phase

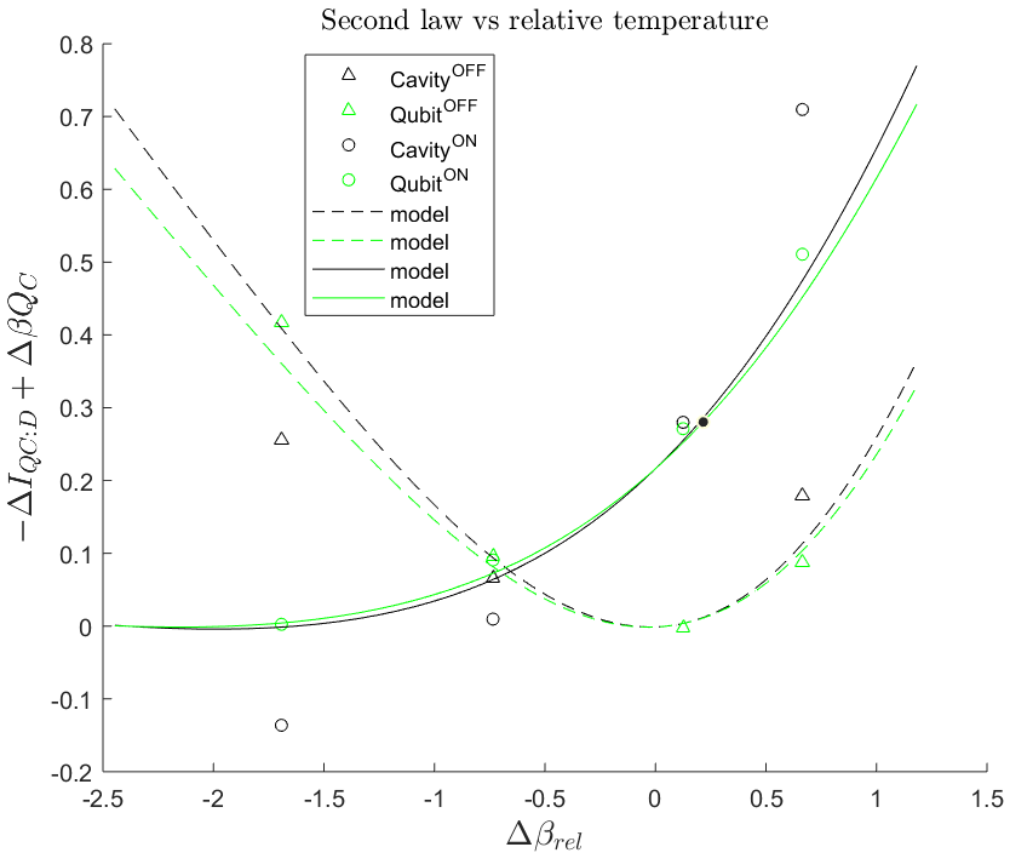

Figure 5.3: Plot of the weak second law in eq. (5.3)

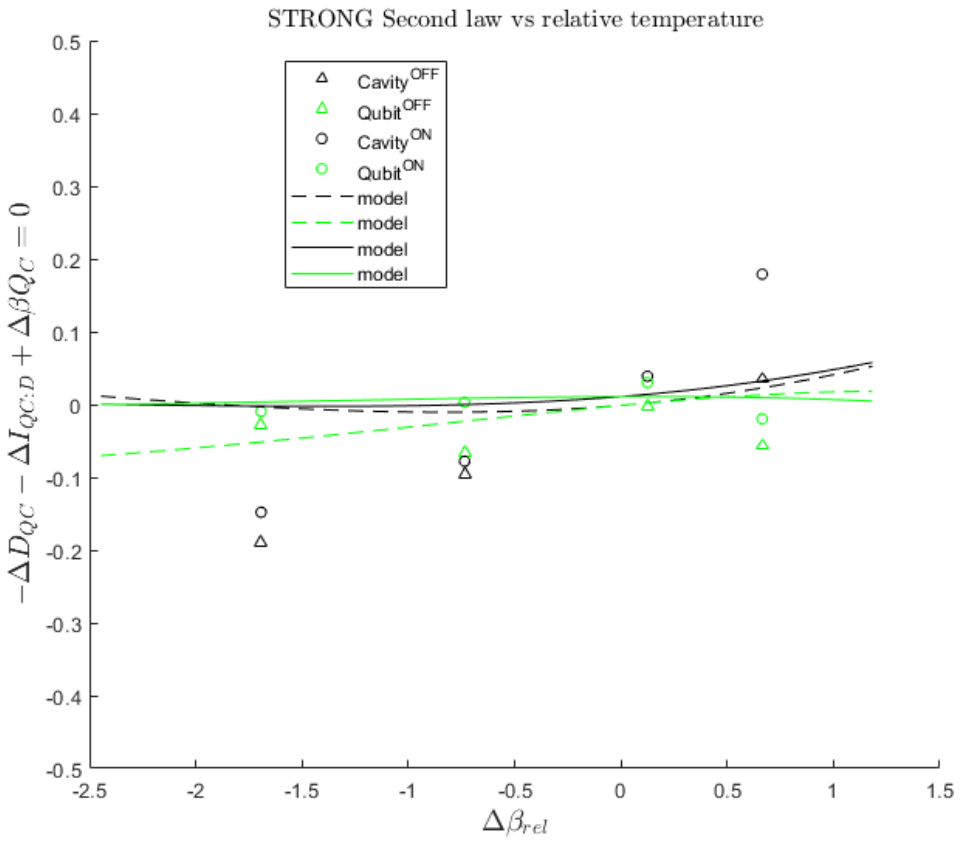

Figure 5.4: Plot of the strong second law in eq. (5.4)

Chapter 6

Error estimation and validity proofs

As seen in the previous chapter, most of our thermodynamics quantities involve directly or indirectly entropy or other statistical quantities of the density matrix. However the proof of validity in [8] only covers the case of observables, i.e. functions such as $f(\rho) = \mathrm{Tr}(\rho A)$ with $A \in \mathcal{O}(\mathcal{H})$. In our case the entropy is not an observable, it is not even linear. It depends not on quantum properties but on the statistical part of the density matrix. And even worse, its derivative goes to infinity on the edge of $\mathcal{D}$.

The two problems I have to solve are then: for a generic function $f : \mathcal{D} \to \mathbb{R}$,

- Is the average value of $f(\rho)$ indeed $f(\rho_{ML})$?
- How to compute the error on the value of such general f on the reconstructed matrix as eq. (5.1) does not work?

My first approach to solve the second problem was a Monte-Carlo simulation explained in section 6.1. It gave quite good results especially since, before this, the team had absolutely no idea of the error of various entropy-related values. It was thus impossible to evaluate how a point a bit off the theoretical curve fitted there.

However this first approach was filled with more or less wrong independence hypothesis (I was forced to assume independence nearly everywhere). So I dived into [8] and proved the first point to be true and found a way to adapt eq. (5.1) to fulfill the second point.

In the end another problem was that, as explained in section 4.2.2, if $\mathrm{Span}(E_i)$ is not the full space, the solution is not unique. We need a way to pick the "best solution". As this is not well defined I propose two possible definitions, both based on maximizing a score function, in section 6.3.

6.1 Quick approach: Monte-Carlo estimation

At the beginning of my internship, I designed a Monte-Carlo algorithm to sample the probability space according to the distribution deduced from the deviation on each coefficient. The goal was to have a quick and dirty method to quickly get error bars that would have been a little but not too much wrongly estimated.

In order to do that, I had to create a probability distribution that matched the deviation on each coefficient. But this distribution cannot be sampled easily because it is the projection of an other probability law on a subvector-space. Let's look at the problem more formally.

In Luis' setup, we have only information on the population from the reconstruction and we force the coherences to be 0. The problem is thus just about probability vectors on the diagonal. So our input is a set of values p_i and deviations Δp_i such that each p_i marginal law is a normal law. In theory the probability distribution on $\mathbb{R}^n$ is:

$$\mathrm{d}p(x_1, \ldots, x_n) = K \exp\left(\sum \frac{(x_i - p_i)^2}{2\Delta p_i^2}\right) \mathrm{d}(x_1, \ldots, x_n)$$

where K is a normalization constant. However, the vector $x = (x_1, \ldots, x_n)$ must be positive and sum to 1. We'll call $\mathcal{P}_n$ the set of probability vectors of dimension n. We will define our new law on $\mathcal{P}_n$ by the same density function but on top of the Lebesgue measure of dimension $n - 1$. This distribution is much harder to simulate because it is both projected and truncated ($\mathcal{P}_n$ is not the whole hyperplane of sum 1).

In order to do the sampling, I use an adapted hit and run method to build a Markov chain on $\mathcal{P}_n$ whose stationary distribution is the wanted distribution. The base idea is very simple: I know how to truncate a Gaussian law to a segment and sample it on that segment. So an iteration of the Markov chain is to go from the previous point x_{n-1} and sample a random direction in the hyperplane of sum 1. This direction and x_{n-1} give us an affine line of that hyperplane.

We can then compute the intersection of that line and $\mathcal{P}_n$. Then we can project the multivariate Gaussian law on the line. As we are working with normal laws, the projected law will be normal and we just have to figure out the center and the standard deviation. If we scale the whole space on each axis by the corresponding Δp_i, this is just an orthogonal projection.

Once we have the law on the line and the segment, we sample the truncated law on the segment and this gives us x_n. The wanted distribution is indeed a stationary distribution of this Markov chain, and the chain fully mixes the set $\mathcal{P}_n$. On discrete probability sets, this is sufficient to say that there is only one stationary distribution and that the chain converges towards it. I hope that this is still true for continuous chains. I decided that this is an inexact and informal way to get the deviation on our functions and thus spending more time on it to have strong convergence results instead of doing the work of the next section was not a good idea.

If I assume that the chain indeed converges toward the right probability distribution, then by sampling this chain long enough, I can get the statistic average and standard deviation of any real function of $\mathcal{P}_n$. This give good enough error bars for our purposes.

6.2 Complete first order error-bars

In this section, I try to formalize and prove formulas for getting the expectation and standard deviation of any function on the set of density matrices $\mathcal{D}$ following the probability given by the maximum-likelihood reconstruction.

All of the theorems in this section are generalizations, adaptations, and improvements from [8]. The goal here is to compute the expected value and the variance of a generic function $h : \mathcal{D} \to \mathbb{R}$ on a variable ρ following a prior distribution P_0 knowing all the measurement we did. If $\operatorname{Span} E_i$ is the full space, we will prove that when the number N of measurement goes to infinity, the expectation is

$$E(h(\rho)) = h(\rho_{ML}) + O(N^{-1})$$

and the variance is

$$V(h(\rho)) = \operatorname{Tr}(\nabla h_{||} \mathbb{H}^{-1} (\nabla h_{||})) + O\Big(\frac{1}{N^2}\Big)$$

where ∇h is the gradient of h on the set of Hermitian matrices of null-trace, and $A_{||}$ and $\mathbb{H}$ are the same as below eq. (5.1). Both asymptotic value are independent of P_0. It is important to note that the term $O(\frac{1}{N^2})$ in the variance formula is still only one order of magnitude below the main term because the variance will decrease in $\frac{1}{N}$ when the number of experiments increase.

One can check that if we set $h(\rho) = \operatorname{Tr}(A\rho)$, we'll find back the result of section 5.1.

Obviously some regularity conditions, that will be given later, are needed on h for these results to work. I tried and achieved to have quite low regularity conditions, so that people could use these results with low-regularity functions. However these conditions are not sufficient to deal with the entropy on the edge of $\mathcal{D}$. In section 6.2.4, I explain how one can achieve similar results with the entropy, but with other terms less bounded e.g. $E(h(\rho)) = h(\rho_{ML}) + O(N^{-\varepsilon})$ for $\varepsilon < 1$.

Let's formalize how to prove those asymptotic. For simplicity we'll assume that P_0 can be expressed as a density function on the Lebesgue measure of the $\operatorname{Tr} = 1$ hyperplane. With that out of the way, we can study the modification of P_0 thanks to the measurements we did by the means of the likelihood function. $\mathcal{L}(\rho) = e^{\ell(\rho)}$ is the density of probability of having the measurements we had, given the density matrix ρ. Therefore, according to Bayes law (see eq. (2.1)), we have

$$E(h(\rho)) = \frac{I_h}{I_1}$$

where

$$I_h = \int h(\rho) e^{\ell(\rho)} P_0(\rho) \mathrm{d}\rho.$$

However to do asymptotic expansions we need to grow a single variable. Thus, we'll write our asymptotic development as if we repeated the experiment and got again the same results. The log-likelihood function will then be $N\ell$. We'll assume that this give the same asymptotic behavior as if we grew the number of measurements inside ℓ. The number of repetition that'll grow to infinity will thus be N, so we will study what happens when $N \to \infty$ on

$$E_h(N) = \frac{I_h(N)}{I_1(N)}$$

where

$$I_h(N) = \int h(\rho) e^{N\ell(\rho)} P_0(\rho) \mathrm{d}\rho.$$

Then, to compute variances, we only need to get the expectation of:

$$h_V = \left(h - \lim_{N\to\infty} E_h(N)\right)^2$$

6.2.1 Full rank

We'll start by studying what happens when $\operatorname{rk}\rho_{ML} = n$ i.e when ρ_{ML} is in the interior of $\mathcal{D}$. To do that, we first need a few technical theorems.

Theorem 6.1. *Let $g : U \to \mathbb{R}$ be a function of class $\mathscr{C}^1$ on U a neighborhood of 0. If $g(0) \neq 0$, we have, as $N \to \infty$:*

$$\int_{z\in U} g(z)e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z = g(0)\left(\frac{2\pi}{N}\right)^{\frac{n}{2}} + O\left(N^{-\frac{n}{2}-1}\right)$$

Proof. We can decompose g in $g(z) = g(0) + h(z)$ with $h(z) = O(\|z\|)$ as g is $\mathscr{C}^1$ in 0, thus:

$$\int_{z\in U} g(z)e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z = g(0)\int_{z\in U} e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z + \int_{z\in U} h(z)e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z.$$

Furthermore,

$$\begin{aligned}
\left|\int_{z\in U} h(z)e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z\right| &\leqslant C\int_{z\in U}\|z\|e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z \\
&\leqslant C\int_{z\in\mathbb{R}^n}\|z\|e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z \\
&\leqslant CA_n\int_{r\in\mathbb{R}_+} r^{n+1}e^{-\frac{N}{2}r^2}\mathrm{d}r \\
&\leqslant CA_n\int_{r\in\mathbb{R}_+} r^{n+1}e^{-\frac{N}{2}r^2}\mathrm{d}r \\
&\leqslant O\left(N^{-\frac{n}{2}-1}\right)
\end{aligned}$$

where C is the constant such that $|h(z)| \leqslant C\|z\|$, A_n is the surface of the hypersphere of $\mathbb{R}^n$ and the last line comes from eq. (C.2). Up to exponentially small terms ($O(e^{N\delta})$ for a given δ), we have (using eq. (C.1)):

$$g(0)\int_{z\in U} e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z \approx g(0)\int_{z\in\mathbb{R}^n} e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z = g(0)\left(\frac{2\pi}{N}\right)^{\frac{n}{2}}$$

□

However when $g(0) = 0$, the previous theorem does not give anymore the first order term but just a bound on the integral which may not be very useful in certain cases. We'll thus look in more details at the case $g(0) = 0$. We'll assume additionally that $\nabla g = 0$, because that is what happens for our cases that require that theorem.

Theorem 6.2. *Let $g : U \to \mathbb{R}$ be a function of class $\mathscr{C}^3$ on U a neighborhood of 0. If $g(0) = 0$ and $\nabla g(0) = 0$, we have, as $N \to \infty$:*

$$\int_{z\in U} g(z)e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z = \frac{\operatorname{Tr}(\nabla^2 g(0))}{2N}\left(\frac{2\pi}{N}\right)^{\frac{n}{2}} + O\left(N^{-\frac{n}{2}-2}\right).$$

Proof. g can be written $g(z) = \frac{1}{2}\langle z|\nabla^2 g(0)|z\rangle + O(\|z\|^3)$ because it is $\mathscr{C}^3$. The integral on the $O(\|z\|^3)$ part gives $O\left(N^{-\frac{n}{2}-2}\right)$ with eq. (C.2) with the same method as in the precedent proof.

We can then use lemma C.1 to get the result. □

Up until now, we used a simple Gaussian in the exponential, but in practice, what we want in there is our log-likelihood function ℓ, let's do a generic theorem then:

Theorem 6.3. *Let U be an open set containing 0. Let $f : U \to \mathbb{R}$ be a $\mathscr{C}^4$ function. Let $g : U \to \mathbb{R}$ be a $\mathscr{C}^1$ function. Assume that f has a unique global non critical maximum in 0 ($\nabla f = 0$ and $\nabla^2 f < 0$). We have at $N \to \infty$:*

$$\int_{z\in U} g(z)e^{Nf(z)}\mathrm{d}z = g(0)\left(\frac{2\pi}{N}\right)^{\frac{n}{2}}\frac{e^{Nf(0)}}{\sqrt{|\det\nabla^2 f(0)|}} + O(e^{Nf(0)}N^{-\frac{n}{2}-1})$$

Additionally, If $g(0) = 0$ and $\nabla g(0) = 0$, and $f \in \mathscr{C}^6$ and $g \in \mathscr{C}^3$, we have:

$$\int_{z\in U} g(z)e^{Nf(z)}\mathrm{d}z = \frac{\mathrm{Tr}\left(-\nabla^2 g(0)\left(\nabla^2 f(0)\right)^{-1}\right)}{2N}\left(\frac{2\pi}{N}\right)^{\frac{n}{2}}\frac{e^{Nf(0)}}{\sqrt{|\det \nabla^2 f(0)|}} + O\left(e^{Nf(0)}N^{-\frac{n}{2}-2}\right)$$

Proof. First we assume $f(0) = 0$ without loss of generality (it suffice to divide everything by $e^{Nf(0)}$). The Morse lemma (C.4) tells us that there is a neighborhood V of 0 and a diffeomorphism $\psi : V \to V'$ with $\psi(0) = 0$ such that

$$f(z) = -\frac{1}{2}\|\psi(z)\|^2 \qquad \text{and} \qquad \nabla\psi(0) = \sqrt{-\nabla^2 f}.$$

Because the hessian of f is negative definite in 0 and the maximum is unique, all values of f outside of V are below a constant $-\delta$. So all the part of the integral outside of V are $O(e^{-N\delta})$ and thus negligible. We can now study the integral only on V. We can then make a change of variable on ψ:

$$\int_{z\in V} g(z)e^{Nf(z)} = \int_{z\in V'} g(\psi^{-1}(z))e^{-\frac{N}{2}\|z\|^2}J(z)\mathrm{d}z$$

with $J(z) = |\det \nabla\psi^{-1}|$. We can then set $g_0(z) = g(\psi^{-1}(z))J(z)$, and look at the two cases:

- As $f \in \mathscr{C}^4$, we have $\psi \in \mathscr{C}^2$ (and thus also ψ^{-1}) and finally $J \in \mathscr{C}^1$, but g is also in $\mathscr{C}^1$, so g_0 is in $\mathscr{C}^1$. We can then apply theorem 6.1 and get the result we want with:

$$g_0(0) = g(\psi^{-1}(0))J(0) = g(0)\frac{1}{\det \nabla\psi(0)} = \frac{g(0)}{\sqrt{|\det \nabla^2 f|}}$$

- As $f \in \mathscr{C}^6$, we have $\psi \in \mathscr{C}^4$ (and thus also ψ^{-1}) and finally $J \in \mathscr{C}^3$, but g is also in $\mathscr{C}^3$, so g_0 is in $\mathscr{C}^3$. We can then apply theorem 6.2 and get the result we want with:

$$\nabla g_0^t = \nabla g^t \nabla\psi^{-1} \times J + g \circ \psi^{-1} \times \nabla J$$

but as $g(0) = 0$, we have $g(\psi^{-1}(0)) = 0$ and furthermore $\nabla g(0) = 0$ so $\nabla(g \circ \psi^{-1})(0) = 0$. Finally, the only remaining term is:

$$\begin{aligned}\nabla^2 g_0(0)) &= \nabla^2(g \circ \psi^{-1}) \times J(0) \\ &= (\nabla\psi^{-1}(0))^t \nabla^2 g(0)(\nabla\psi^{-1}(0)) \times \frac{1}{\sqrt{|\det \nabla^2 f|}}\end{aligned}$$

If we put that to the trace, we get the result we wanted.

□

6.2.2 Partial rank

When ρ_{ML} is not of full rank, we are on the edge of $\mathcal{D}$. We must thus study what happens to the results of the previous section, when we are on the edge. Luckily, we can just study the case of edges on which only one dimension is missing. Indeed when more dimension are missing, there is a trick to bring it back to only one dimension missing.

Remark. All formulas here work with only x i.e if n, the dimension of z, is 0. It suffice to note that the Lebesgue measure of the point is 1 and all the proofs work.

Theorem 6.4. *Let U be an open set of $\mathbb{R}_+ \times \mathbb{R}^n$ where $(0,0) \in U$. The $\mathbb{R}_+$ coordinate will be x and the $\mathbb{R}^n$ coordinate will be z. Let $g : U \to \mathbb{R}$ be a $\mathscr{C}^1$ function. The derivatives against the edge i.e on x are taken as one-sided derivatives. We have as $N \to \infty$:*

$$\int_{(x,z)\in U} x^m g(x,z)e^{-N(x+\frac{1}{2}\|z\|^2)}\,\mathrm{d}x\,\mathrm{d}z = g(0,0)\,m!(2\pi)^{\frac{n}{2}}N^{-\frac{n}{2}-m-1} + O(N^{-\frac{n}{2}-m-2})$$

Proof. With the now usual argument, we can restrict ourselves to a rectangular set $U' = [0,\eta) \times U_z \subset U$. Here we decompose g in $g(x,z) = h(z) + xg_1(x,z)$ with $h(z) = g(0,z)$ (We can do that because g is $\mathscr{C}^1$). As g_1 is bounded on U, we have:

$$\begin{aligned}\int_{(x,z)\in U'} x^{m+1} g_1(x,z)e^{-N(x+\frac{1}{2}\|z\|^2)}\,\mathrm{d}x\,\mathrm{d}z &\leqslant \|g_1\|_\infty \int_{(x,z)\in U'} x^{m+1}e^{-N(x+\frac{1}{2}\|z\|^2)}\,\mathrm{d}x\,\mathrm{d}z \\ &= O(N^{-\frac{n}{2}-m-2})\end{aligned}$$

This is because we can split the integral and use eq. (C.1) and eq. (C.3).
On the other hand, thanks to eq. (C.1) and theorem 6.1, we have:

$$\begin{aligned}\int_{(x,z)\in U'} x^m h(z) e^{-N(x+\frac{1}{2}\|z\|^2)}\,\mathrm{d}x\,\mathrm{d}z &= \int_0^\eta x^m e^{-Nx}\,\mathrm{d}x \times \int_{z\in U_z} h(z) e^{-\frac{N}{2}\|z\|^2}\,\mathrm{d}z \\ &= m! N^{-m-1} \times \left(h(0)\left(\frac{2\pi}{N}\right)^{\frac{n}{2}} + O\left(N^{-\frac{n}{2}-1}\right)\right)\end{aligned}$$

□

Theorem 6.5. *If we add to the previous theorem the hypothesis that $g(0) = 0$, $\nabla g(0) = 0$ and g is $\mathscr{C}^2$ and in particular $\mathscr{C}^3$ on the variable z (i.e. for any x, $z \mapsto g(x,z) \in \mathscr{C}^3$), we have:*

$$\int_{(x,z)\in U} x^m g(x,z) e^{-N(x+\frac{1}{2}\|z\|^2)}\,\mathrm{d}x\,\mathrm{d}z = \frac{m!\,\mathrm{Tr}\left(\frac{\partial^2 g}{\partial z^2}\right)}{2N^{m+2}}\left(\frac{2\pi}{N}\right)^{\frac{n}{2}} + O\left(N^{-\frac{n}{2}-m-3}\right)$$

Proof. We use the same U' as in last proof. By applying lemma C.2 twice, we get

$$g(x,z) = h(z) + x^2 g_2(x,z)$$

For that same reason as in previous proof, The integral on $x^2 g_2$ is $O(N^{-\frac{n}{2}-m-3})$. We can then, like the previous proof, split the integral and use eq. (C.1) and theorem 6.2 to get the result. □

Like in the previous section we now need to take care about a generic function f in the exponential. We can't use the theorem 6.3 because it assume a null gradient for f but here $\frac{\partial f}{\partial x}$ could be negative.

Theorem 6.6. *Let U be an open set of $\mathbb{R}_+ \times \mathbb{R}^n$ where $(0,0) \in U$. The $\mathbb{R}_+$ coordinate will be x and the $\mathbb{R}^n$ coordinate will be z. Let $f : U \to \mathbb{R}$ be a function of class $\mathscr{C}^4$ and $g : U \to \mathbb{R}$ be a function of class $\mathscr{C}^1$. The derivative against the edge i.e on x are taken as one-sided derivative. Assume that f has a global non-critical maximum in 0 so that $\frac{\partial f}{\partial z}(0) = 0$ and $\frac{\partial^2 f}{\partial z^2}(0) < 0$. Furthermore, we assume that $\frac{\partial f}{\partial x}(0) < 0$. We have as $N \to \infty$:*

$$\int_{(x,z)\in U} x^m g(x,z) e^{Nf(x)}\,\mathrm{d}x\,\mathrm{d}z = \frac{g(0,0)\,m!(2\pi)^{\frac{n}{2}}}{N^{\frac{n}{2}+m+1}} \times \frac{e^{Nf(0,0)}}{\sqrt{\left|\det \frac{\partial^2 f}{\partial z^2}\right|}\left(-\frac{\partial f}{\partial x}\right)^{m+1}} + O(e^{Nf(0,0)} N^{-\frac{n}{2}-m-2})$$

Additionally, if $g(0) = 0$, $\nabla g(0) = 0$, g is $\mathscr{C}^2$ and in particular g is $\mathscr{C}^3$ on variable z and f is $\mathscr{C}^6$, We have as $N \to \infty$:

$$\int_{(x,z)\in U} x^m g(x,z) e^{Nf(x)}\,\mathrm{d}x\,\mathrm{d}z = \frac{m!\,\mathrm{Tr}\left(\frac{\partial^2 g}{\partial z^2}\left(\frac{\partial^2 f}{\partial z^2}\right)^{-1}\right)}{2N^{m+2}} \times \frac{\left(\frac{2\pi}{N}\right)^{\frac{n}{2}} e^{Nf(0,0)}}{\sqrt{\left|\det \frac{\partial^2 f}{\partial z^2}\right|}\left(-\frac{\partial f}{\partial x}\right)^{m+1}} + O\left(e^{Nf(0,0)} N^{-\frac{n}{2}-m-3}\right)$$

Proof. We start by supposing $f(0,0) = 0$ by dividing everything by $e^{Nf(0,0)}$.

We now do a change a variable to bring back f to be $-x - \frac{1}{2}\|z\|^2$ like in the proof of theorem 6.3. To that change of variable we decompose f in $f(x,z) = h(z) + x f_1(x,z)$ with $h(z) = f(0,z)$. The change of variable $\tilde{x} = x f_1(x,z)$ and $\tilde{z} = \psi(z)$ where ψ is the Morse lemma decomposition of h will give the right result. Doing the actual variable change and computing the Jacobian is left to the reader. The regularity analysis is the same as in theorem 6.3 □

6.2.3 Application

Here we can finally apply all those asymptotic theorems to our case. I first recall the notations. For $h : \mathcal{D} \to \mathbb{R}$, we define

$$I_h(N) = \int_{\rho\in\mathcal{D}} h(\rho) e^{N\ell(\rho)} P_0(\rho)\mathrm{d}\rho \tag{6.1}$$

where P_0 is the prior probability distribution on density matrices. So that we have its expectation

$$E_h(N) = \frac{I_h(N)}{I_1(N)}. \tag{6.2}$$

Its variance is the expectation of:

$$h_V = \left(h - \lim_{N\to\infty} E_h(N)\right)^2 \tag{6.3}$$

We make some technical assumptions such as:

- $P_0(\rho_{ML}) > 0$: If our estimator has a null-probability, the prior distribution doesn't make much sense.
- P_0 is of class $\mathscr{C}^3$: This is necessary for the various asymptotic developments.
- $\mathcal{D} = \mathcal{D}(\mathbb{C}^d)$: We use simple d-dimensional density matrices.

Lemma 6.7. *For any h of class $\mathscr{C}^1$ on $\mathcal{D}$, we have:*

$$E_h(N) = h(\rho_{ML}) + O\left(\frac{1}{N}\right)$$

and for any h of class $\mathscr{C}^2$ such that $h(\rho_{ML}) = 0$, $\nabla h(\rho_{ML}) = 0$, and such that, if we name z the variable on the space tangent to $\mathcal{D}_{\mathrm{rk}\,\rho_{ML}}$ in ρ_{ML}, g is of class $\mathscr{C}^3$ in z, we have:

$$E_h(N) = \frac{\mathrm{Tr}\left(-\dfrac{\partial^2 h}{\partial z^2}(\rho_{ML})\left(\dfrac{\partial^2 \ell}{\partial z^2}(\rho_{ML})\right)^{-1}\right)}{2N} + O\left(\frac{1}{N^2}\right)$$

Proof if ρ_{ML} has full rank. In this case, D is a neighborhood of ρ_{ML} in the hyperplane of Hermitian matrices of trace 1. This hyperplane is itself an euclidean real vector space isometric to $\mathbb{R}^n$ for a given n. ℓ is smooth so it is $\mathscr{C}^4$, and it has a global unique (by convexity) non-critical (by strong convexity) maximum in ρ_{ML}. As both h and P_0 are of class $\mathscr{C}^1$, we can just do a translation to apply theorem 6.3. We have then:

$$I_h(N) = h(\rho_{ML})P_0(\rho_{ML})\left(\frac{2\pi}{N}\right)^{\frac{n}{2}} \frac{e^{N\ell(\rho_{ML})}}{\sqrt{|\det \nabla^2 f(\rho_{ML})|}} + O(e^{N\ell(\rho_{ML})}N^{-\frac{n}{2}-1})$$

and

$$I_1(N) = P_0(\rho_{ML})\left(\frac{2\pi}{N}\right)^{\frac{n}{2}} \frac{e^{N\ell(\rho_{ML})}}{\sqrt{|\det \nabla^2 f(\rho_{ML})|}} + O(e^{N\ell(\rho_{ML})}N^{-\frac{n}{2}-1})$$

Everything then simplifies in the result we wanted:

$$E_h(N) = \frac{I_h(N)}{I_1(N)} = h(\rho_{ML}) + O\left(\frac{1}{N}\right)$$

In second case $h(\rho_{ML}) = 0$, z just span the full space and so represent the same variable as ρ. Therefore h is just plainly $\mathscr{C}^3$, we can thus just apply the second part of theorem 6.3, and get:

$$I_h(N) = \frac{\mathrm{Tr}\left(-\nabla^2 h(\rho_{ML})\left(\nabla^2 \ell(\rho_{ML})\right)^{-1}\right)}{2N}\left(\frac{2\pi}{N}\right)^{\frac{n}{2}} \frac{e^{N\ell(\rho_{ML})}}{\sqrt{|\det \nabla^2 f(\rho_{ML})|}} + O\left(e^{Nf(0)}N^{-\frac{n}{2}-2}\right)$$

And finally:

$$E_h(N) = \frac{I_h(N)}{I_1(N)} = \frac{\mathrm{Tr}\left(-\nabla^2 h(\rho_{ML})\left(\nabla^2 \ell(\rho_{ML})\right)^{-1}\right)}{2N} + O\left(\frac{1}{N^2}\right)$$

□

Proof if ρ_{ML} has partial rank and $\nabla\ell \neq \lambda I$. This proof is a bit more subtle, because in theorem 6.6, the edge dimension x is of only one dimension but here, we may miss several dimensions. We thus need to mount a change of variable that reduces several dimension to one. This change of variable came form [8].

Before starting let's just name $h_0(\rho) = h(\rho)P_0(\rho)$.

First let r be the rank of ρ_{ML}. Then let's suppose that ρ_{ML} is diagonal by rotating everything along a unitary U such that $\rho_{ML} = UDU^\dagger$. D will then be $\mathrm{diag}(0, \dots, 0, p_1, \dots, p_r)$. We'll define Δ by:

$$D = \begin{pmatrix} 0 & 0 \\ 0 & \Delta \end{pmatrix}$$

We can then pose the following change of variable:

$$\Psi(\xi, \zeta, \omega) = \exp\begin{pmatrix} 0 & \omega \\ -\omega^\dagger & 0 \end{pmatrix}\begin{pmatrix} \xi & 0 \\ 0 & \Delta + \zeta - \mathrm{Tr}\,\xi\frac{I}{r} \end{pmatrix}\exp\begin{pmatrix} 0 & -\omega \\ \omega^\dagger & 0 \end{pmatrix} \tag{6.4}$$

Where $\xi \in \mathcal{O}(\mathbb{C}^{d-r})$, $\omega \in \mathcal{M}_{(d-r),r}$ and $\zeta \in \mathcal{O}(\mathbb{C}^r)$ but with $\mathrm{Tr}\,\zeta = 0$. First, let's show that it is a diffeomorphism. We have:

$$\nabla\Psi(D)\cdot(\delta\xi, \delta\zeta, \delta\omega) = \begin{pmatrix} \delta\xi & \delta\omega\,\Delta \\ \Delta\,\delta\omega & \delta\zeta - \mathrm{Tr}(\delta\xi)\frac{I}{r} \end{pmatrix} \tag{6.5}$$

This is bijective in 0 (Δ is invertible), so by local inversion theorem $\Psi : U \to V$ is a diffeomorphism on U a neighborhood of D. By the same argument as usual, we can focus our study of the various integral to only $U' = U \cap \mathcal{D}$ and even to $U'' = \mathring{U'}$ because $U' \setminus U''$ is of null measure. We can thus apply the change of variable

$$\int_{\rho \in U''} h_0(\rho) e^{N\ell(\rho)} \mathrm{d}\rho = \int_{(\xi,\zeta,\omega) \in V''} h_0(\Psi(\xi,\zeta,\omega)) e^{N\ell(\Psi(\xi,\zeta,\omega))} J(\xi,\zeta,\omega) \mathrm{d}\xi \, \mathrm{d}\zeta \, \mathrm{d}\omega \tag{6.6}$$

Where $J((\xi,\zeta,\omega))$ is the determinant of the Jacobian of Ψ and $V'' = \Psi(U'')$.

In fact we can characterize what the edge of $\mathcal{D}$ is in the (ξ,ζ,ω)-space:

$$\rho = \Psi(\xi,\zeta,\omega) > 0 \iff \xi > 0$$

So we would like ξ to be our x variable and (ζ,ω) to be our z variable. On one hand the second part is easy: let's name z the variable $z = (\zeta,\omega)$. One the other hand the first part is not that easy because ξ is multidimensional (except if $r = n-1$).

Luckily If we write $\xi = x\sigma$ with $\sigma \in \mathcal{D}(\mathbb{C}^{d-r})$, we'll have $\xi > 0 \iff x > 0$ and we are brought back to a single one sided variable. If we name $\Phi(x,\sigma) = x\sigma$ our function, it is a diffeomorphism on the whole $\mathbb{R}_{>0} \times \mathcal{D}(\mathbb{C}^{d-r})$ to the positive definite Hermitian matrices of size r. Furthermore it's Jacobian can easily be computed and it is x^r.

There just one problem before doing a new change of variable: In order to do a proper change of variable, we'll need to split ξ from z, Thus we'll build a rectangle neighborhood V''' of $(0,0,0)$ that separates ξ on one side and z on the other, such that $V''' = V'''_\xi \times V'''_z$. We'll also use $W''' = \Phi^{-1}(V'''_\xi)$ We'll thus have:

$$\int_{\rho \in U'''} h_0(\rho) e^{N\ell(\rho)} \mathrm{d}\rho = \int_{(x,\sigma) \in W'''} \int_{z \in V'''_z} x^m h_0(\Psi(\Phi(x,\sigma),z)) e^{N\ell(\Psi(\Phi(x,\sigma),z))} J(\Phi(x,\sigma),z) \, \mathrm{d}x \, \mathrm{d}\sigma \, \mathrm{d}z \tag{6.7}$$

Where $m = \dim \mathcal{D}(\mathbb{C}^{d-r}) = (d-r)^2 - 1$. If we split again W''' in a rectangular sub-neighborhood of 0, namely $W''''_x \times W''''_\sigma$, we have:

$$\int_{\rho \in U''''} x^m h_0(\rho) e^{N\ell(\rho)} \mathrm{d}\rho = \int_{\sigma \in W''''_\sigma} \left(\int_{(x,z) \in W''''_x \times V'''_z} h_0(\Psi(\Phi(x,\sigma),z)) e^{N\ell(\Psi(\Phi(x,\sigma),z))} J(\Phi(x,\sigma),z) \, \mathrm{d}x \, \mathrm{d}z \right) \mathrm{d}\sigma \tag{6.8}$$

We now want to apply theorem 6.6 on the integral inside the parenthesis. Let's check all the hypothesis. $U = \overline{W''''_x \times V'''_z}$ is a neighborhood of $(0,0)$ in $\mathbb{R}_+ \times \mathbb{R}^n$. ℓ, Ψ, Φ are smooth so $f(x,z) = \ell(\Psi(\Phi(x,\sigma),z))$ will be $\mathscr{C}^4$. Being a non-critical global maximum is conserved when pre-composing with any differentiable function so $(0,0)$ is global maximum of f on U.

We now only need to prove about f that $\frac{\partial f}{\partial x} < 0$:

$$\frac{\partial f}{\partial x} \times \delta x = \nabla \ell \cdot \begin{pmatrix} \delta x \, \sigma & 0 \\ 0 & -\delta x I \end{pmatrix}$$

According to property 4.12 If $\nabla f \neq \lambda I$, we have $\eta \neq 0$ with $\eta \geqslant 0$ such that $\mathrm{Tr}(\eta D) = 0$ and $\nabla f + \eta = \lambda I$, Thus we have:

$$\eta = \begin{pmatrix} \nu & 0 \\ 0 & 0 \end{pmatrix} \tag{6.9}$$

with $\nu \geqslant 0$ and $\nu \neq 0$. Therefore:

$$\frac{\partial f}{\partial x} = -\mathrm{Tr}(\eta\sigma) < 0.$$

Additionally, h and P_0 are $\mathscr{C}^1$, and J is smooth so $g(x,z) = h_0(\Psi(\Phi(x,\sigma),z)) J(\Phi(x,\sigma),z)$ is of class $\mathscr{C}^1$. We can then apply theorem 6.6, and when $h(\rho_{ML}) \neq 0$ we have:

$$I_h(N) = \frac{h(\rho_{ML}) P_0(\rho_{ML}) J(0,0) \, m! (2\pi)^{\frac{n}{2}}}{N^{\frac{n}{2}+m+1}} \times \int_{\sigma \in W'''_\sigma} \frac{e^{N\ell(\rho_{ML})}}{\sqrt{\left|\det \frac{\partial^2 f}{\partial z^2}\right|} \left(-\frac{\partial f}{\partial x}\right)^{m+1}} \mathrm{d}\sigma + O\left(\frac{e^{Nf(0,0)}}{N^{\frac{n}{2}+m+2}}\right) \tag{6.10}$$

We can then simplify everything and get the result we want.

In the case $h(\rho_{ML}) = 0$. We can easily check that the various regularity conditions are directly mapped on the similar conditions on g. Because $\frac{\partial h}{\partial z} = 0$ by hypothesis, we have:

$$\frac{\partial^2 g}{\partial z^2}(0,0) = J(0,0)P_0(\rho_{ML})\frac{\partial^2 h}{\partial z^2}(\rho_{ML})$$

Therefore when we apply theorem 6.6 and simplify by eq. (6.10), we get the second result we wanted. □

The corner case where the ρ_{ML} is on the edge of $\mathcal{D}$ but the gradient toward the exterior is 0 is not done in this report like in the original paper [8]. The probability that such an event happens is really low (if it is not 0) in particular if we take into account the numerical errors of the implementation. However it can still probably be done by splitting again the x and z variable but keeping x as a quadratic variable in the function in the exponential.

Theorem 6.8. *For any $h : \mathcal{D} \to \mathbb{R}$ of class $\mathscr{C}^1$, we have its expectation:*

$$E_h(N) = h(\rho_{ML}) + O\left(\frac{1}{N}\right)$$

Additionally, if it is of class $\mathscr{C}^3$ in the inside and along the edges and $\mathscr{C}^2$ otherwise with $\frac{\partial^2 h}{\partial x^2} = o(\frac{1}{x})$ when leaving the edge (where x is a scalar variable not tangent to the edge), we'll have the variance:

$$V_h(N) = E_{h_V}(N) = \frac{\operatorname{Tr}(\nabla h_{||}\mathbb{H}^{-1}(\nabla h_{||}))}{N} + O(\frac{1}{N^2})$$

where:

- *$\mathcal{D}_r$ is the submanifold of $\mathcal{D}$ of matrices with rank r*
- *P_{ML} is the orthogonal projector on the range of ρ.*
- *$A_{||}$ is the orthogonal projection on the tangent space to $\mathcal{D}_{\operatorname{rk}\rho}$ in ρ:*

$$A_{||} = A - \frac{\operatorname{Tr}(AP_{ML})}{\operatorname{Tr}(P_{ML})}P_{ML} - (I - P_{ML})A(I - P_{ML})$$

- *$\mathbb{H}$ is the hessian of $\ell_{|\mathcal{D}_{\operatorname{rk}\rho}}$ in ρ:*

$$\mathbb{H}(A) = \sum_i \frac{\operatorname{Tr}(AE_{i||})}{\operatorname{Tr}^2(\rho_{ML}E_i)}E_{i||} + (\lambda I - \nabla\ell(\rho_{ML}))A\rho_{ML}^+ + \rho_{ML}^+A(\lambda I - \nabla\ell(\rho_{ML}))$$

Where ρ^+ is the Moore-Penrose pseudo inverse.

Proof. The expectation proof is just a direct application of lemma 6.7. For the variance, it's a bit more tricky. First, let's check the regularity conditions. We want to apply lemma 6.7 on $h_V = \big(h - h(\rho_{ML})\big)^2$. On the inside h is of class $\mathscr{C}^3$, so we have

$$\nabla h_V = 2(h - h(\rho_{ML}))\nabla h \qquad \text{and} \qquad \nabla^2 h_V = 2(h - h(\rho_{ML}))\nabla^2 h + 2|\nabla h\rangle\langle\nabla h| \tag{6.11}$$

And we find out, that if ρ_{ML} is on the edge, $\nabla^2 h_V$ can be prolonged to the edge because $\nabla^2 h$ will be dominated by $(h - h(\rho_{ML}))$. If we are inside, we have a $\mathscr{C}^3$ neighborhood, so we don't care. Furthermore, we have that $\nabla h_V = 0$. Therefore, we can apply lemma 6.7 and get:

$$V_h(N) = E_{h_V}(N) = \frac{\operatorname{Tr}\left(-\frac{\partial^2 h_V}{\partial z^2}(\rho_{ML})\left(\frac{\partial^2 \ell}{\partial z^2}(\rho_{ML})\right)^{-1}\right)}{2N} + O\left(\frac{1}{N^2}\right)$$

As seen in eq. (6.11), $\frac{\partial^2 h_v}{\partial z^2} = 2|\frac{\partial h}{\partial z}\rangle\langle\frac{\partial h}{\partial z}^t|$. We just have some equalities to show:

- $\frac{\partial h}{\partial z} = \nabla h_{||}$: If we take an hermitian matrix A written as

$$A = \begin{pmatrix} A_0 & A_{0,r} \\ A_{0,r}^\dagger & A_r \end{pmatrix}$$

one can check that we have

$$A_{||} = \begin{pmatrix} 0 & A_{0,r} \\ A_{0,r}^\dagger & A_r - \operatorname{Tr}(A_r)\frac{I}{r} \end{pmatrix}$$

which is exactly the tangent space to $\mathcal{D}_{\operatorname{rk}\rho_{ML}}$ at ρ_{ML}. Furthermore that tangent space is exactly spanned by z.

- $\langle X|\left(\frac{\partial^2\ell}{\partial z^2}\right)^{-1}|X\rangle = \operatorname{Tr}(X \times \mathbb{H}^{-1}(X))$: Computing the link between the original hessian at ρ_{ML} truncated to the tangent space and $\frac{\partial^2\ell}{\partial z^2}$ is a bit more complicated because the z variable has a curve. We use again explicitly the change of variable of the previous proof. If we write $\ell(z) = \ell(\Psi(z))$, we have

$$\frac{\partial^2\ell}{\partial z^2} = \langle\nabla\ell|\nabla^2\Psi + \nabla\Psi\nabla^2\ell\nabla\Psi$$

In this formula, $\nabla^2\Psi$ is a not a matrix but a 3 dimensional tensor with two input sides and one output side. The $\langle\nabla\ell|$ applies on the output side.
Since $\nabla\Psi|X\rangle = |X_{||}\rangle$ the second term is simply

$$\nabla\Psi^t\nabla^2\ell\nabla\Psi = \sum_i \frac{|E_{i||}\rangle\langle E_{i||}|}{\operatorname{Tr}^2(\rho_{ML}E_i)}$$

On the other hand, with some calculus, one can prove that

$$\nabla^2\Psi(\delta z, \delta z) = \begin{pmatrix} 2\delta\omega\,\Delta\,\delta\omega & ? \\ ? & ? \end{pmatrix}.$$

We use again the notations of eq. (6.9) and property 4.12. We know that $\langle\lambda I|\nabla^2\Psi = 0$ because $\mathcal{Y}$ in the set of matrices of trace one. So the only remaining part is:

$$\langle\nabla\ell|\nabla^2\Psi(\delta z, \delta z) = -2\operatorname{Tr}(\nu\delta\omega\Delta\delta\omega).$$

If we write $\delta\Psi = \nabla\Psi\delta z$, we have:

$$\begin{aligned}\operatorname{Tr}(\delta\Psi(\lambda I - \nabla\ell(\rho_{ML}))\delta\Psi\rho_{ML}^+ + \delta\Psi\rho_{ML}^+\delta\Psi(\lambda I - \nabla\ell(\rho_{ML}))) &= -\operatorname{Tr}(\delta\Psi\eta\delta\Psi\rho_{ML}^+ + \delta\Psi\rho_{ML}^+\delta\Psi\eta) \\ &= -2\operatorname{Tr}(\nu\delta\omega\Delta\delta\omega)\end{aligned}$$

In the end,

$$\langle X|\frac{\partial^2\ell}{\partial z^2}|X\rangle = \langle X|\mathbb{H}(X) = \operatorname{Tr}(X\mathbb{H}(X))$$

so they are just two version of the same operator.

□

If we remove the N factor and replace $N\ell$ by just ℓ, the N denominator will enter $\mathbb{H}$ and we will have the results that were announced at the beginning of the section 6.2

6.2.4 Regularity of Entropy

If look at the previous regularity constraints, the entropy does not satisfy them: it is not $\mathscr{C}^1$ on the edge, and its variance is not $\mathscr{C}^2$ on the edge. However, if I define intermediate regularity classes as:

Definition 6.1. We say that a function $f : \mathbb{R} \to \mathbb{R}$ is of class $\mathscr{C}^\beta$ in 0 if for $k = \lfloor\beta\rfloor$, the function is $\mathscr{C}^k$ on a neighborhood of 0 and we can write:

$$f^{(k)} = f^{(k)}(0) + o(x^{\beta-k})$$

Remark. f is of class $\mathscr{C}^\beta$ if and only if $f' \in \mathscr{C}^{\beta-1}$.

Remark. This definition extends trivially to multivariate functions, So I'll use the notation for those functions

Property 6.9. *If $f \in \mathscr{C}^\beta$ and $\lfloor k\rfloor = \beta$, we have*

$$f(x) = f(0) + xf'(0) + \cdots + \frac{x^{(k)}}{k!}f^{(k)}(0) + o(x^\beta)$$

Proof. If we call $g(x) = f(x) - \left(f(0) + xf'(0) + \cdots + \frac{x^{(k)}}{k!}f^{(k)}(0)\right)$, then $g^{(i)}(0) = 0$ for any $i \leqslant k$, and $g^{(k)}(x) = o(x^{\beta-k})$. By successive integration's we get $g^{(i)}(x) = o(x^{\beta-i})$ and thus $g(x) = o(x^\beta)$. □

Theorem 6.10. *For any $h : \mathcal{D} \to \mathbb{R}$ of class $\mathscr{C}^\varepsilon$ with $0 < \varepsilon < 1$, we have its expectation:*

$$E_h(N) = h(\rho_{ML}) + O\left(\frac{1}{N^\varepsilon}\right)$$

Proof. It suffice to go over all the proofs and replace $\mathscr{C}^1$ by $\mathscr{C}^\varepsilon$ everywhere, and all the proofs can be adapted: when we decompose g in $g(z) = g(0) + h(z)$ with $O(\|z\|)$ in theorem 6.1, We just replace it by $h(z) = O(\|z\|^\varepsilon)$. In all the proof the change of class propagates well (The product of a $\mathscr{C}^1$ function and $\mathscr{C}^\varepsilon$ function is $\mathscr{C}^\varepsilon$. The same is true for composition $g \circ f$ if g is $\mathscr{C}^\varepsilon$ and f is in $\mathscr{C}^1$).

In theorem 6.4, we also decompose $g(x, z)$ in $h(0, z) + xg_1(x, z)$. In the $\mathscr{C}^\varepsilon$, we can do $g(x, z) = h(0, z) + x^\varepsilon g_1(x, z)$ and the rest of the proof will give what we want. □

Theorem 6.11. *Let's take $h : \mathcal{D} \to \mathbb{R}$ that is $\mathscr{C}^3$ on the interior of $\mathcal{D}$ and on the tangent spaces to the edges. Furthermore we ask that h is $\mathscr{C}^\varepsilon$ on the edge for $0 < \varepsilon < 1$ and that it satisfies $\frac{\partial h}{\partial x} = o(\frac{1}{x^\delta})$ with $0 < \delta < \varepsilon$, for any scalar variable x not tangent to the edge, the variance will be:*

$$V_h(N) = E_{h_V}(N) = \frac{\mathrm{Tr}(\nabla h_{||}\mathbb{H}^{-1}(\nabla h_{||}))}{N} + O(\frac{1}{N^{1+\varepsilon-\delta}})$$

Proof. Let's study the regularity of $h_V = (h - h(\rho_{ML}))^2$. On the interior of $\mathcal{D}$ we have:

$$\nabla h_V = 2(h - h(\rho_{ML}))\nabla h$$

First on any vector variable z tangent to the edge of $\mathcal{D}$, this formula will be true on the edge with partial gradient on z. On any variable x not tangent to the edge, we'll have $\frac{\partial h_v}{\partial x} = 2o(x^\varepsilon)o(\frac{1}{x}^\delta) = o(x^{\varepsilon-\delta})$. From $\varepsilon > \delta$, we deduce that $\frac{\partial h_V}{\partial x}(\rho_{ML})$ exists and is 0, so h_V can be prolonged in a $\mathscr{C}^1$ function on the edge. In fact is prolonged in a $\mathscr{C}^{1+\varepsilon-\delta}$ function.

Now, if we look at lemma 6.7, it only manipulates derivative on z, so its proof works. If we go further back, all theorem require both $\mathscr{C}^3$ in the interior and when tangent to the edge, and $\mathscr{C}^2$ otherwise. That $\mathscr{C}^2$ can be replaced everywhere by $\mathscr{C}^{1+\varepsilon-\delta}$ with the same kind of proof modification as in the previous theorem. □

Property 6.12. *The con-Newman entropy is of class $\mathscr{C}^\varepsilon$ for any $\varepsilon < 1$ on the edge of $\mathcal{D}$, and smooth on the interior.*

Proof. $x^\varepsilon \ln(x) \xrightarrow[x\to 0]{} 0$, for any $0 < \varepsilon$ □

Corollary 6.12.1. *The expectation of the entropy S is*

$$E_S(N) = S(\rho_{ML}) + O(N^{-\varepsilon})$$

Remark. It is likely that as $\varepsilon \to 1$, the hidden constant of the O will explode.

Property 6.13. *The entropy satisfy the hypothesis of theorem 6.11 for any $\varepsilon < 1$ and any $\delta < \varepsilon$. Therefore for any $1 < \beta < 2$, we have*

$$V_S(N) = \frac{\mathrm{Tr}(\nabla S_{||}\mathbb{H}^{-1}(\nabla S_{||}))}{N} + O\left(\frac{1}{N^\beta}\right)$$

I haven't tried them all, but I think all the other entropy-related values like mutual-information will also have similar regularity characteristics and thus will have the same result on the asymptotic expectation and variance.

6.3 Fixing the non full span problem

6.3.1 Centering function

As explained in section 4.2.2, when the E_i do not span the whole space of positive hermitian matrices, the solution to the max-likelihood problem is not unique. We need to decide a way to pick one. The simplest way to do that is to choose a centering function $c : \mathcal{D} \to \mathbb{R}$ such that, the higher c is, the more "centered", the density matrix is. This function must satisfy some properties to make sense

Definition 6.2. A function $c : \mathcal{D}(\mathcal{H}) \to \mathbb{R}$ is a *centering function* if:

1. It is concave, so that being between other matrices is always better.
2. $\arg\max_{\rho\in\mathcal{D}} c = \dfrac{I}{\dim \mathcal{H}}$.

3. c must be unitary invariant: $c(U\rho U^\dagger) = c(\rho)$ for $U \in \mathcal{U}(\mathcal{H})$
4. In dimension 2, the density matrix is inside the Bloch sphere which is linearly isomorphic to $\mathcal{D}$. Any vector space cutting this space gives either a disc or a segment that has an obvious center. c must give that center.

All these constraints are qualitative constraints that ensure that c makes sense as a centering function. If a function does not satisfy those conditions, it is not a good centering function. But satisfying them does not guaranty that the function make sense from the physical point of view. The logarithm of the determinant and the von-Neuman entropy are good candidates. I plot heatmaps of them on random planes to see how good they are in fig. 6.1. However, I don't think there is a perfect centering function waiting to be found, but I obviously can't prove that.

I didn't use any of its results directly in the report, but [1] was very useful for analysing the structure of the density matrix space in my quest for the perfect centering function

6.3.1.1 Log-Determinant

The log-determinant is good choice of strictly concave function. It seems to really keep the matrices geometrically centered. And on some special case, for example if it happens to be quadratic on the vector space studied, It will give the exact center (It the determinant is quadratic on a slice of $\mathcal{D}$, that slice will be an ellipsoid which has an exact center). The reasons to use the log determinant instead of the determinant are twofold: It is strictly concave on $\mathcal{D}$, and it will more likely fit on a machine floating point number on the edges of $\mathcal{D}$.

I haven't time to prove formally that the log-determinant satisfy all centering function properties, but none of those should be hard to do.

6.3.1.2 Von Neumann Entropy

The Von-Neumann entropy is also a good choice of centering function from a probabilistic point of view. Maximizing the entropy means maximizing the incertitude, which makes sense when we are moving on direction on which we have zero information.

The entropy also satisfies the four conditions and thus is a centering function.

6.3.1.3 In practice

That part of my work wasn't directly useful for Luis project, because his effect matrices only span the populations of Fock base and thus the centering simply consists in putting 0 in each correlation, which I think achieve perfect centering whatever the centering function is.

6.3.2 Two objective optimisation

In order to use the centering function, what we want to do is that, after optimizing ℓ and landing on $M = \arg\max_{\rho\in\mathcal{D}} \ell(\rho)$, we can optimize the centering function on M and get the most centered matrix that maximize ℓ.

One way of doing it is to do exactly what I just said: Find a ρ_{int} in M with a first optimization (projected gradient ascent will find one). Then we could optimize c on M. However I have no idea how to project a vector on M: The projection method described in section 4.3.2 only works on the whole $\mathcal{D}$ because $\mathcal{D}$ is unitary invariant which is generally not the case of M.

An other way of doing that is by using a kind of barrier method with c. By optimizing $\ell + \varepsilon c$ and reducing progressively ε, we'll the right $\rho_{ML} = \arg\max_{\rho\in M} c$.

Property 6.14. *If I name ρ_ε the solution of:*

$$\left\{\textbf{maximize} \quad \ell + \varepsilon c \quad \textbf{on } \rho \in \mathcal{D}\right.$$

Then, if c is strictly concave, and uniformly continuous we'll have:

$$\rho_\varepsilon \xrightarrow[\varepsilon\to 0]{} \rho_{ML} = \arg\max_{\rho\in M} c$$

Proof. Let's define $\ell_{ML} = \ell(\rho_{ML})$. ℓ is strictly concave on all the dimension on which there is an E_i i.e all the dimension orthogonal to M. That means that if $\ell(\rho_\varepsilon) \to \ell_{ML}$, then $d(\rho_\varepsilon, M) \to 0$.

Then lets prove that $\ell(\rho_\varepsilon) \to \ell_{ML}$. We now that $\ell_{ML} > \ell(\rho_\varepsilon)$, but we also know that $\ell(\rho_\varepsilon) + \varepsilon c(\rho_\varepsilon) \geqslant \ell_{ML} + \varepsilon c(\rho_{ML})$. Therefore:

$$\ell_{ML} - \ell(\rho_\varepsilon) < \varepsilon(g(\rho_\varepsilon) - g(\rho_{ML}))$$

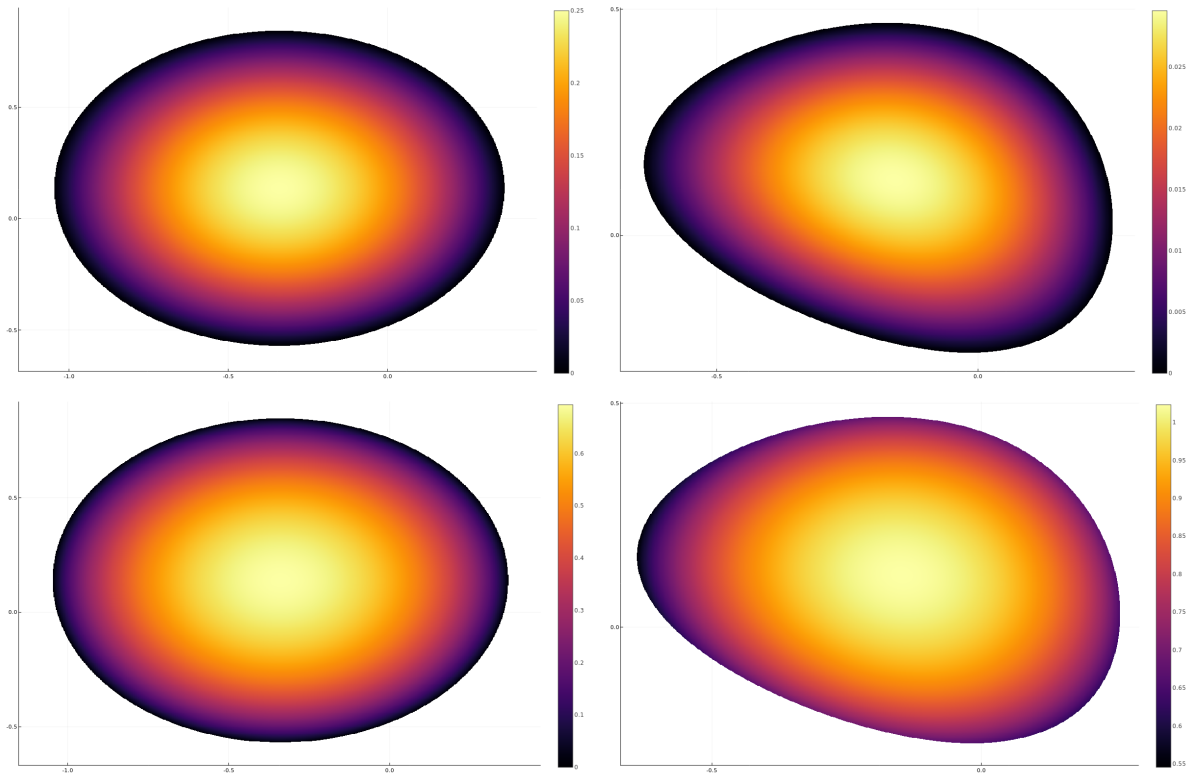

(a) dimension 2

(b) dimension 3

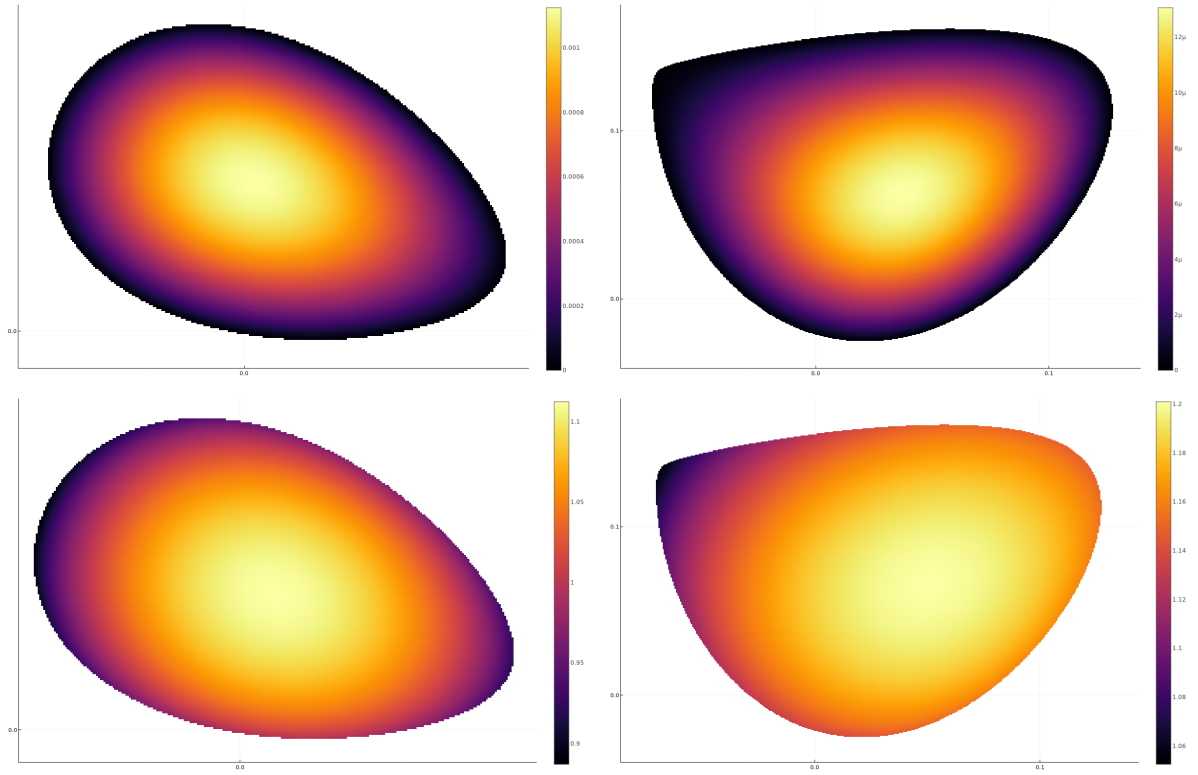

(c) dimension 4

(d) dimension 5

Figure 6.1: In order to compare $\log \circ \det$ and S, I plotted them on random affine planes slicing $\mathcal{D}$. I pick a random density matrix A and two random orthonormal null trace matrices H_1 and H_2. I then plot $\log \circ \det$ and S in the plane $A + xH_1 + yH_2$. In each case $\log \circ \det$ is on top and the entropy is below. The colorless area is outside of $\mathcal{D}$

But c is strictly concave so it has an upper bound. So when $\varepsilon \to 0$, we have $\ell_{ML} - \ell(\rho_\varepsilon) \to 0$ and thus $d(\rho_\varepsilon, M) \to 0$.

Let's name $\rho_{p\varepsilon}$ the projection of ρ_ε on M. As $\rho_{p\varepsilon} \in M$, we have $g(\rho_{p\varepsilon}) < c(\rho_{ML})$.

By uniform continuity, $d(g(\rho_{p\varepsilon}), c(\rho_\varepsilon)) \to 0$ and $g(\rho_{ML})$ is between them so $c(\rho_{p\varepsilon}) \to c(\rho_{ML})$. But as c is strictly concave, this must mean that $\rho_{p\varepsilon} \to \rho_{ML}$.

On the other hand we had $d(\rho_\varepsilon, M) = d(\rho_\varepsilon, \rho_{p\varepsilon}) \to 0$, so in the end we have $\rho_\varepsilon \to \rho_{ML}$. □

This proof works perfectly with the entropy but not with the $\log \circ \det$ as it isn't uniformly continuous. Furthermore, the projected gradient method won't work directly with the entropy because when projecting on the edge, the gradient of the entropy will be infinite. I think both of this problems have solutions, for example when we are on the edge, only output the tangent gradient for the entropy. However, the internship is finished, so I won't have the time to check it properly.

Chapter 7

Final Results

7.1 Approximation and implementation of first order error propagation

Now I have the operator $\mathbb{H}$, and I want to compute the standard deviation on various functions that are mixture of classical observable evaluations and entropy-related functions. For my function $h : \mathcal{D} \to \mathbb{R}$, I know from previous chapter that its standard deviation is:

$$\sqrt{\mathrm{Tr}(\nabla h_{||} \mathbb{H}^{-1} \nabla h_{||})} \tag{7.1}$$

The problem is that the functions I want to evaluate are complex and span multiple files of source code. I would like a method to propagate the error estimation through all the layers of the code, following the way functions are computed. In order to do that, let's look at a more generic way of propagating errors.

Suppose we have a random variable $X \in \mathbb{R}^n$ and its covariance matrix V_X. Suppose we take a function $f : \mathbb{R}^n \to \mathbb{R}^m$ such that $P(X \in \mathrm{dom}\, f) = 1$. We would like to know the covariance matrix of $f(X)$. If X has a probability distribution centered on its expectation X_0, such as a normal law, we can make a first order approximation and assume that f on the domain where X varies looks like $f(X_0) + \nabla f \cdot (X - X_0)$. In such a situation one can prove that:

$$V_{f(X)} = \nabla f^t V_X \nabla f = \langle \nabla f | V_x | \nabla f \rangle. \tag{7.2}$$

This really looks like the eq. (7.1). In fact if we linearize our Hermitian matrices ∇h and thus think of $\mathbb{H}$ as an element of $\mathcal{L}(\mathcal{O}(\mathcal{H}))$, we can write:

$$V_h = \langle \nabla h_{||} | \mathbb{H}^{-1} | \nabla h_{||} \rangle$$

Furthermore the projecting operator $_{||}$ is linear and thus can be put in a matrix form such that

$$V_h = \langle \nabla h | P^t \mathbb{H}^{-1} P | \nabla h \rangle.$$

In fact we have $P^t \mathbb{H}^{-1} P = \mathbb{H}$ but I don't need it so I won't prove it. What is important is that $V_\rho = P^t \mathbb{H}^{-1} P$ looks like the covariance matrix of ρ around ρ_{ML} and thus I proved in the last chapter that the first order error propagation gives the right asymptotic variance of any function.

Therefore, in the code, I compute V_ρ from the results of the reconstruction and then when a function is applied on a variable, I compute its covariance matrix with eq. (7.2). In the end when a function has a single dimension output, the variance matrix is a simple variance real that I can put through a square root to get the standard deviation on that value.

7.2 Plots

Here are the plots of section 5.5 upgraded with the new error bars. Unfortunately I didn't have enough time to finish implementing the error propagation of last section, so those values come from the Monte-Carlo algorithm. We added several new points of data in between. The horizontal error bars are due to Luis and are originated from imperfections in the control and measurement of the temperatures of both the cavity and the atom.

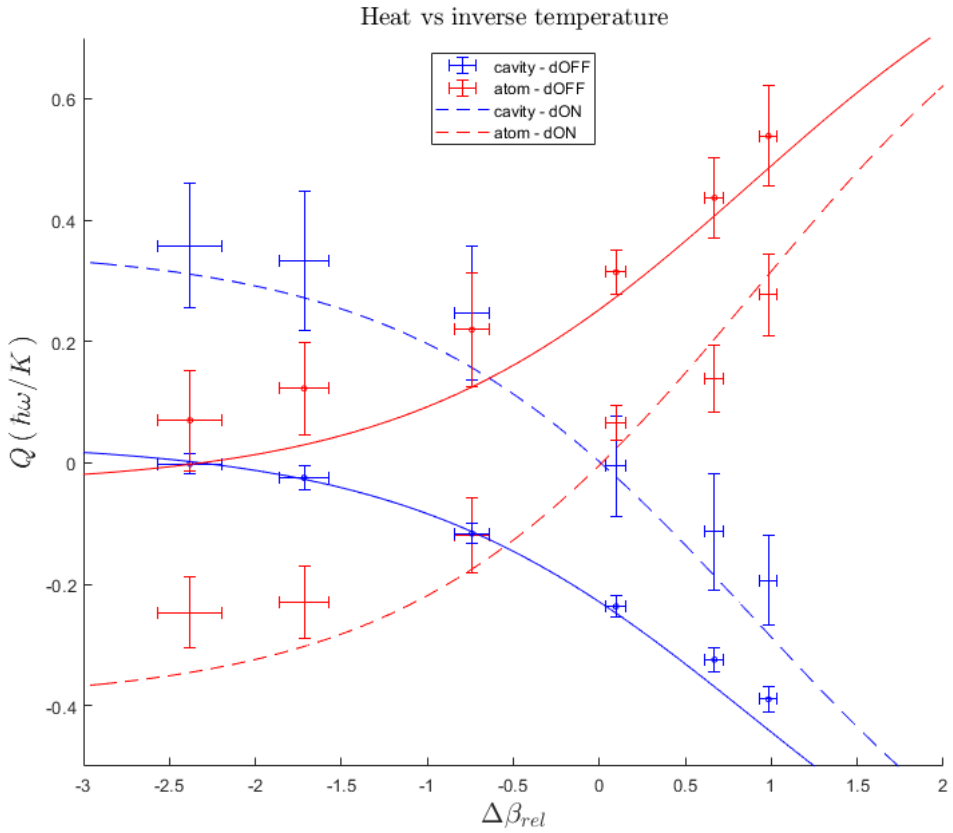

Figure 7.1: Plot of the exchanged heat depending on the Maxwell demon state (ON or OFF)

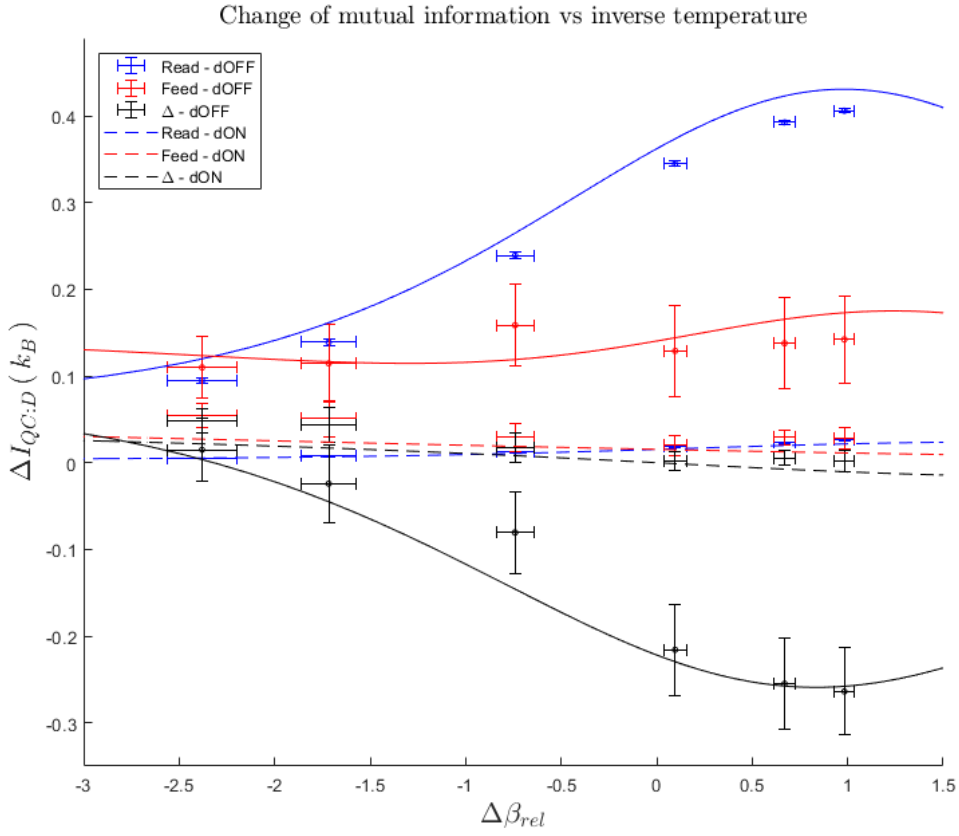

Figure 7.2: Plot of the mutual information in the readout phase vs the feedback phase

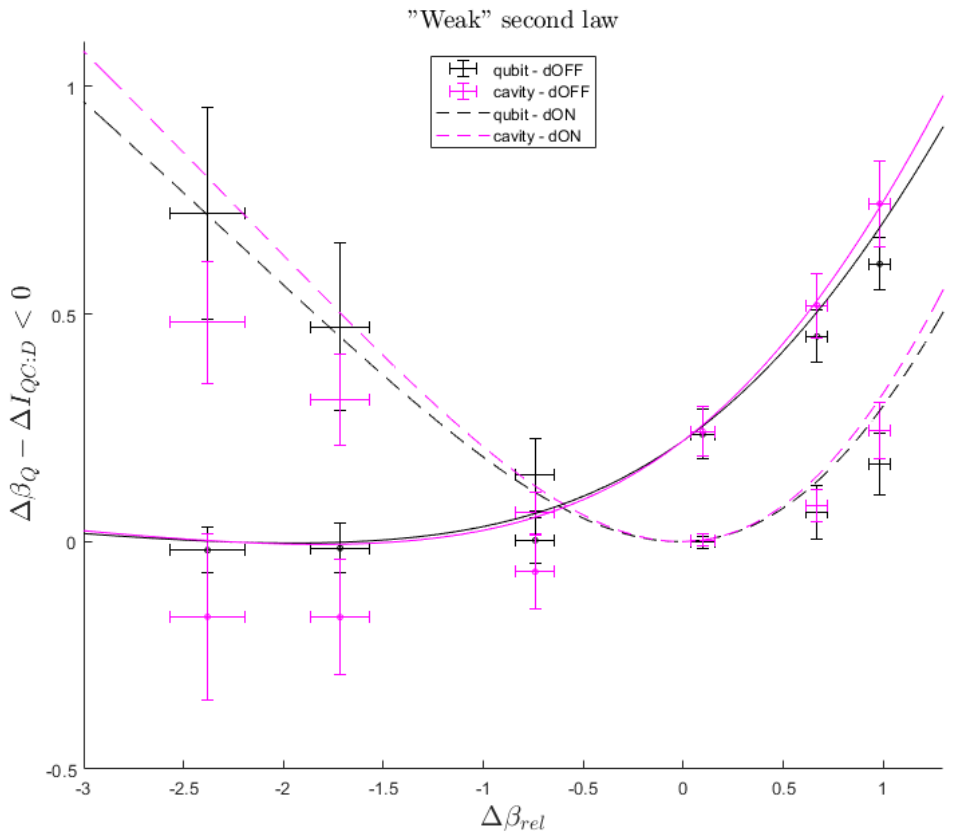

Figure 7.3: Plot of the weak second law in eq. (5.3)

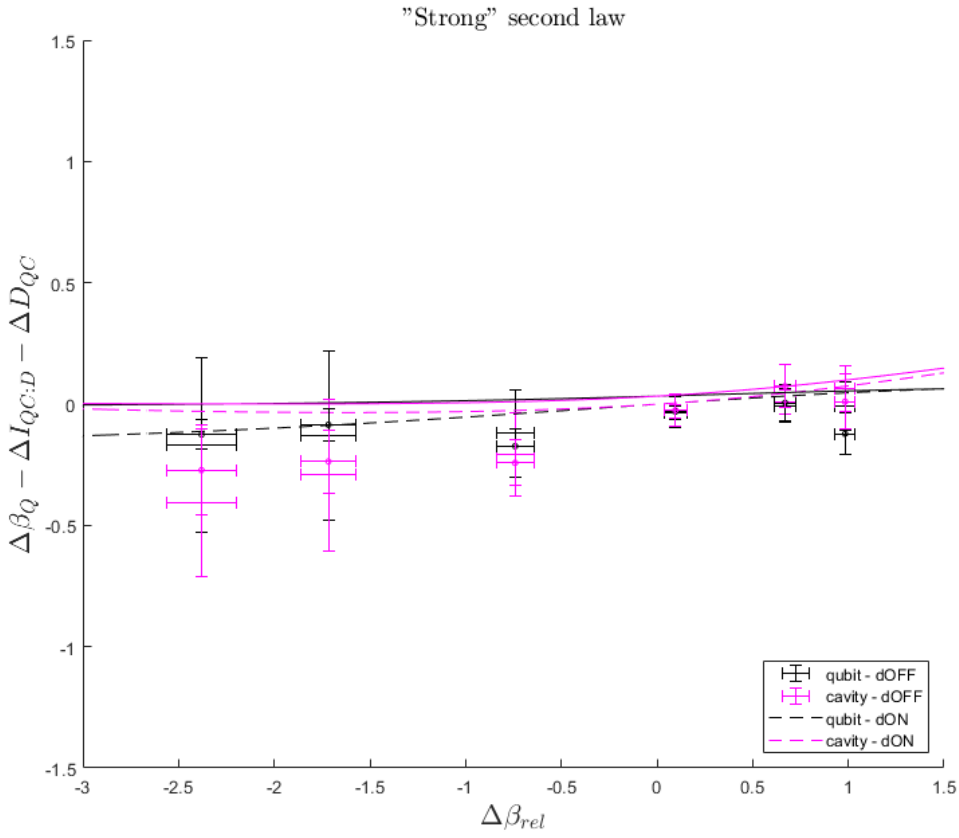

Figure 7.4: Plot of the strong second law in eq. (5.4)

Conclusion

This internship was really interesting; I had to understand a lot of various mathematical and physical theory to understand the huge amount of work that my predecessor did. Once I did that, I could have the pleasure to make a small but nevertheless important contribution to the field of quantum state tomography.

It was really interesting to see the inner working of such a physics lab, and to see the amazing experimental setup they have at the LKB. In particular to see them running. Before this internship I couldn't imagine that physicists were able to control atoms and photons one by one to make experiments.

On a more theoretic point of view, this internship greatly improved my comprehension of quantum mechanics both on the calculus side and the interpretation side including my way of interpreting classical mechanics and thermodynamics. It has also improved my knowledge on asymptotic integrals, but that is a bit less prone to philosophy.

In the end, this was a great experience where I learnt a lot of things and met a lot of interesting people, some of which I want to thank :

Thanks

- Thanks to Igor Dotsenko for welcoming me in his lab, showing me their amazing experimental setup and giving me interesting things to do around it.

- Thanks to Pierre Rouchon for giving me interesting corner-case problems to solve around log-likelihood maximization and helping me along the way when I was stuck.

- Thanks to Luis Najera for having a great internship project thus allowing me to contribute to something very interesting. Thanks also to him for answering all my question about the codebase and helping me to implement some parts of my work. Also thank to him for answering various physical questions when I had trouble understanding quantum electro-dynamics or just quantum information theory.

- Thanks to Valentin Metillon for helping me to understand various thing about physics and statistics.

Bibliography

[1] Rajendra Bhatia. *Positive Definite Matrices*. Princeton Series in Applied Mathematics. Princeton University Press, Princeton, NJ, USA, 2007. http://press.princeton.edu/titles/8445.html.

[2] Alexandre d'Aspremont. Lecture notes on convex optimization, first session. https://www.di.ens.fr/~aspremon/PDF/ENS/ConvexityENS.pdf.

[3] Serge Haroche and Jean Michel Raimond. *Exploring the Quantum: Atoms, Cavities, and Photons*. Oxford Univ. Press, Oxford, 2006. doi:10.1093/acprof:oso/9780198509141.001.0001.

[4] V. Métillon, S. Gerlich, M. Brune, J.M. Raimond, P. Rouchon, and I. Dotsenko. Benchmarking maximum-likelihood state estimation with an entangled two-cavity state. https://arxiv.org/abs/1904.04681.

[5] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*, chapter 8.2 Quantum operations, pages 356–373. Cambridge University, 2001.

[6] Clément Sayrin. *Préparation et stabilisation d'un champ non classique en cavité par rétroaction quantique*. PhD thesis, 2011.

[7] P. Six, Ph. Campagne-Ibarcq, I. Dotsenko, A. Sarlette, B. Huard, and P. Rouchon. Quantum state tomography with noninstantaneous measurements, imperfections, and decoherence. *Phys. Rev. A*, 93:012109, Jan 2016. doi:10.1103/PhysRevA.93.012109.

[8] P. Six and P. Rouchon. *Feedback Stabilization of Controlled Dynamical Systems*, chapter Asymptotic expansions of Laplace integrals for quantum state tomography, pages 307–327. Lecture notes in Control and Information Sciences 473. Springer, 2017. https://arxiv.org/abs/1607.00948.

[9] Pierre Six. *Estimation d'état et de paramètres pour les systèmes quantiques ouverts*. PhD thesis, 2016. http://www.theses.fr/2016PSLEM019/document.

Appendix A

Quantum mechanics pictures

Pictures are different way to perceive and express quantum mechanics. They vary in how the system state and the various operators change over time. All these representations yield the same mechanics and are equivalent up to certain change of Hilbert space basis dependent on time. Going from one to the other is just up to a unitary $U(t)$.

A.1 Schrödinger picture

The Schrödinger picture is the usual representation of quantum mechanics. In this context, only the state of the system varies and the way to observe it do not change. When the system evolves under Hamiltonian $H(t)$, the various operators relevant to the system are constant (unless they vary because of a external source which is not the system) and the state follow the Schrödinger equation:

Equation A.1.

$$i\hbar\frac{\partial|\Psi\rangle}{\partial t}(t) = H(t)|\Psi(t)\rangle$$

Property A.2. *If H is constant over time, the solution of equation A.1 is:*

$$|\Psi(t)\rangle = e^{-i\frac{H}{\hbar}t}|\Psi(0)\rangle$$

On very important theorem that makes a link with the Heisenberg picture is the following:

Theorem A.3 (Ehrenfest)**.**

$$\frac{\mathrm{d}}{\mathrm{d}t}\langle A\rangle = \frac{i}{\hbar}\langle[H,A]\rangle + \left\langle\frac{\partial A}{\partial t}\right\rangle$$

A.2 Heisenberg picture

In the Heisenberg picture however, the state in $\mathcal{H}$ only corresponds to the initial state and represent the whole trajectory. However to measure an observable at time t, you use its expression $A(t)$ at this point in time. The equation that rules this evolution is

Equation A.4.

$$\frac{\mathrm{d}A}{\mathrm{d}t} = \frac{i}{\hbar}[H,A] + \left(\frac{\partial A}{\partial t}\right)_H$$

Here the $[\cdot,\cdot]$ is the commutator and the partial derivative is the variation of A due to external elements out of the system (i.e. the variation of A in Schrödinger picture). This equation is actually the not-averaged version of Ehrenfest theorem (A.3)

Property A.5. *The Schrödinger picture and Heisenberg picture describe the same mechanics. In particular $\langle\Psi|A|\Psi\rangle$ is the same for any operator in both pictures.*

Proof. The value of the average in Schrödinger picture is given by Ehrenfest theorem. Let's compute the same average in Heisenberg picture. On the right side, the averaging give instantly the right hand of Ehrenfest theorem. We only have left to prove:

$$\left\langle\frac{\mathrm{d}A}{\mathrm{d}t}\right\rangle = \frac{\mathrm{d}}{\mathrm{d}t}\langle A\rangle$$

But this is still true by linearity. □

In fact, The Schrödinger equation has a general solution which is the unitary evolution operator $U(t)$. From there, the evolution of the state in Schrödinger's picture in $|\Psi(t)\rangle = U(t)|\Psi(0)\rangle$, whereas the evolution of a operator in Heisenberg picture is $A(t) = U^\dagger(t)A_0(t)U(t)$, where $A_0(t)$ expresses the variation from external elements.

If the Hamiltonian is time-independent, $U(t) = e^{-i\frac{H}{\hbar}t}$

A.3 Interaction picture

Sometimes, in order to compare two phenomenon, it is useful to compare two Hamiltonian in certain way. Usually, this is to compare the free evolution of a system to the interaction with another system, this is why this is called the interaction picture. Will split our Hamiltonian in two parts:

$$H = H_0 + H_1$$

In most case we'll try to keep H_0 time independent and simple, whereas H_1 will contains all the complexity of the system. That way, $U_0(t)$, will still be $e^{-i\frac{H_0}{\hbar}t}$. We'll then make operators evolve with H_0, when the state evolves with H_1. The equations are:

$$i\hbar\frac{\partial|\Psi\rangle}{\partial t}(t) = H_1(t)|\Psi(t)\rangle$$
$$A(t) = U_0^\dagger(t)A_0(t)U_0(t)$$

This equivalent to manipulating state in a moving frame that moves according to U_0. The Schrödinger frame would then be the fixed frame and the Heisenberg frame would be the one that move exactly with the state.

Appendix B

Covariance

B.1 Covariance and variance

Definition B.1. Given two random vector $X \in \mathbb{R}^n$ and $Y \in \mathbb{R}^m$, we let the covariance matrix be defined by:

$$\operatorname{cov}(X,Y) = \big(\operatorname{cov}(X_i, Y_j)\big)_{i\leqslant n, j\leqslant m}$$

The fundamental goal of the covariance matrix is to have $\langle a|\operatorname{cov}(x,y)|b\rangle = \operatorname{cov}(a\cdot x, b\cdot y)$

Property B.1. *If $X \in \mathbb{R}^n$ and $Y \in \mathbb{R}^m$ are two random vectors, we have* $\operatorname{cov}(X,Y) = \operatorname{cov}(Y,X)^t$.

Property B.2. *If A is a deterministic matrix, we have* $\operatorname{cov}(AX,Y) = A\operatorname{cov}(X,Y)$

Definition B.2. Given a random vector $X \in \mathbb{R}^n$, we define its variance by

$$V(X) = \operatorname{cov}(X,X)$$

Property B.3. *For $X \in \mathbb{R}^n$ a random vector, we have $V(X) > 0$*

B.2 Correlation

Definition B.3. If X and Y are two real valued variables that are not fixed (non-zero variance), we have:

$$\rho = \left|\frac{\operatorname{cov}(X,Y)}{\sigma(X),\sigma(Y)}\right|$$

where $\sigma(X) = \sqrt{V(X)}$

Property B.4. *We always have: $|\rho| \leq 1$*

We can extend the correlation to the multivariate case by setting

$$\rho = V(X)^{-\frac{1}{2}}\operatorname{cov}(X,Y)V(Y)^{-\frac{1}{2}}$$

The norm of ρ now becomes the symmetric part of the polar decomposition: $|\rho| = (\rho\rho^t)^{\frac{1}{2}}$. We can now get the theorem that give the bound on the correlation:

Theorem B.5. *If $X \in \mathbb{R}^n$ and $Y \in \mathbb{R}^m$ are random vectors, and $V(Y)$ is invertible, we have:*

$$\operatorname{cov}(X,Y)V(Y)^{-1}\operatorname{cov}(Y,X) \leqslant V(X)$$

Corollary B.5.1. *If $V(X)$ is also invertible, we have $|\rho| < I$*

Proof. We take $a \in R^n$ such that $\langle a|V(X)|a\rangle > 0$ and any $b \in \mathbb{R}^m$. By using property B.4, we get:

$$\frac{\langle a|\operatorname{cov}(X,Y)|b\rangle}{\sqrt{\langle a|V(X)|a\rangle}\sqrt{\langle b|V(Y)|b\rangle}} \leqslant 1$$

To saturate the inequality, we can maximize the numerator on b, while keeping the denominator constant. If we use the Lagrange multiplier method, we want to find a critical point of:

$$\mathcal{L}(b,\lambda) = \langle a|\operatorname{cov}(X,Y)|b\rangle + \lambda\langle b|V(Y)|b\rangle$$

We thus need to have:

$$\langle a| \operatorname{cov}(X,Y) + 2\lambda\langle b|V(Y) = 0$$

and thus:

$$\langle b| = -\frac{1}{2\lambda}\langle a| \operatorname{cov}(X,Y)V(Y)^{-1}$$

By substituting in the first equation we get:

$$\frac{\langle a| \operatorname{cov}(X,Y)V(Y)^{-1} \operatorname{cov}(Y,X)|a\rangle}{\sqrt{\langle a|V(X)|a\rangle}\sqrt{\langle a| \operatorname{cov}(X,Y)V(Y)^{-1} \operatorname{cov}(Y,X)|a\rangle}} \leqslant 1$$

We thus get for all a such that $\langle a|V(X)|a\rangle > 0$:

$$\langle a| \operatorname{cov}(X,Y)V(Y)^{-1} \operatorname{cov}(Y,X)|a\rangle < \langle a|V(X)|a\rangle$$

As for other a, both terms are 0, this inequality holds for all a, thus the theorem holds. □

Appendix C

Some useful calculus

$$\int_{z\in\mathbb{R}^n} e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z = \left(\frac{2\pi}{N}\right)^{\frac{n}{2}} \tag{C.1}$$

$$\int_{r\in\mathbb{R}_+} r^{n+1}e^{-\frac{N}{2}r^2}\mathrm{d}r = 2^{\frac{n}{2}}N^{-\frac{n}{2}-1}\Gamma\left(\frac{n}{2}+1\right) \tag{C.2}$$

On integers, we have a simple value for the gamma-like integral:

$$\int_{r\in\mathbb{R}_+} r^m e^{-Nr}\mathrm{d}r = m!N^{-m-1} \tag{C.3}$$

Lemma C.1. *Let S be an real symmetric matrix of dimension n, we have:*

$$\int_{z\in\mathbb{R}^n} \langle z|S|z\rangle e^{-\frac{N}{2}\|z\|^2}\mathrm{d}z = \frac{\mathrm{Tr}(S)}{N}\left(\frac{2\pi}{N}\right)^{\frac{n}{2}}$$

Proof. If we take the spectral decomposition of S and split the coordinate along S eigenvectors, the result comes after some calculus. □

C.1 Morse lemma

Lemma C.2. *Let $f : U \subset \mathbb{R}^n \to \mathbb{R}^m$, be $\mathscr{C}^n$ with U a neighborhood of 0. Suppose that $f(0) = 0$. Then we have $g : U \to \mathcal{L}(\mathbb{R}^n, \mathbb{R}^m)$ of class $\mathscr{C}^{n-1}$ such that:*

$$f(x) = g(x)x$$

and with $g(0) = \nabla f(0)$

Proof. $g(x) = \int_0^1 \nabla f(tx)\mathrm{d}t$ □

Lemma C.3. *Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$, be $\mathscr{C}^n$ with U a neighborhood of 0. Suppose that $f(0) = 0$ and $\nabla f(0) = 0$. Then we have $h : U \to \mathcal{S}(\mathbb{R}^n)$ of class $\mathscr{C}^{n-2}$ such that $h(0) = \nabla^2 f(0)$ and:*

$$f(x) = \langle x|h(x)|x\rangle$$

Proof. We apply lemma C.2 on f to get g, that we see as a gradient $g : U \to \mathbb{R}^n$. Then we apply it again on g to get h_0, then we take the symmetric part:

$$h = \frac{h_0 + h_0^t}{2}$$

We also have: $h_0(0) = \nabla g(0) = \nabla^2 f(0) = h(0)$ □

Lemma C.4 (Morse lemma)**.** *Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$, be $\mathscr{C}^n$ with U a neighborhood of 0. Suppose that $f(0) = 0$, $\nabla f(0) = 0$ and $\nabla^2 f > 0$. Then we have a neighborhood V of 0 and $\psi : V \to \mathbb{R}^n$ of class $\mathscr{C}^{n-2}$, a such that on V:*

$$f(x) = \|\psi(x)\|^2$$

If $n \geqslant 3$, we can have ψ a diffeomorphism with $\nabla\psi(0) = \sqrt{\nabla^2 f(0)}$

Proof. We take h as in lemma C.3. As h is continuous, we can take V such that $h > 0$ on V ($h(0) = \nabla^2 f > 0$).

We can then define $\psi(x) = \sqrt{h(x)}x$. With the square root being the $\mathscr{C}^\infty$ square root on the positive definite matrices. If $n \geqslant 3$, then h is $\mathscr{C}^1$, thus we have: $\nabla\psi(0) = \sqrt{h(0)} = \sqrt{\nabla^2 f(0)}$. As $\nabla^2 f(0) > 0$, we have $\nabla\psi(0)$ invertible, so by local inversion theorem, we can reduce the size of V to make ψ a diffeomorphism on V.

□