Efficient Hamiltonian reconstruction from equilibrium states

Adam Artymowicz

California Institute of Technology, Pasadena, CA 91125, USA

March 28, 2024

Abstract

We describe a novel algorithm that recovers the Hamiltonian and temperature of a quantum state in thermal equilibrium using a restricted set of measurements. The algorithm works by imposing a semidefinite constraint based on free energy minimization, which we interpret in terms of thermodynamic stability. We benchmark the algorithm on the problem of learning a nearest-neighbour Hamiltonian on a 100-qubit spin chain.

1 Introduction

A defining feature of the Gibbs state $\rho = e^{-h/T} / \operatorname{tr}(e^{-h/T})$ is that it minimizes the free energy:

$$\rho = \operatorname{argmin} F \tag{1}$$

where $F(\sigma) = -TS(\sigma) + tr(\sigma h)$ and S is the von Neumann entropy. Conversely, given ρ , the condition (1) uniquely specifies h, and this fact can in principle be used to recover the Hamiltonian of a Gibbs state [1]. Thanks to the strict convexity of F, (1) is equivalent to the first-order condition

$$dF[\delta\rho] \ge 0 \tag{2}$$

for any $\delta \rho$ in the tangent space at ρ . However, not only is computing the derivative of F difficult, (2) also involves imposing a number of conditions that scales quadratically with the dimension of the physical Hilbert space.

This work is based on a hierarchy of relaxations of the local minimality condition (2). These relaxations recently appeared in [2] and use a matrix version of the Araki-Sewell inequality [3]. For each relaxation we give a classical algorithm that attempts to reconstruct the Hamiltonian H and the temperature T from the expectation values of certain operators. Given a set of variational Hamiltonian terms $h_1, \ldots h_s$, the algorithm either returns a candidate Hamiltonian $h \in \text{span}(h_1, \ldots, h_s)$ and candidate temperature T, or returns a certificate that the state is not a Gibbs state of any Hamiltonian in the span of h_1, \ldots, h_s .

To analyze the algorithm, we introduce a condition which we call restricted thermodynamic stability (RTS) that generalizes the Gibbs variational condition (2). Indeed, it is equivalent to (2) where $\delta\rho$ is restricted to a certain conic region in the tangent space of ρ . Physically, the Gibbs variational condition is related to stability of the state against perturbations. Correspondingly, we interpret the RTS condition as stability with respect to a certain class of open system dynamics.

Although we prove correctness of the algorithm in the absence of measurement noise (Corollary 1), we do not give any bounds on sample complexity (number of copies of ρ needed for accurate reconstruction) or classical computational complexity, leaving this to future work. Instead, we

include numerical evidence that the algorithm achieves good reconstruction accuracy using modest computational resources and in the presence of measurement noise.

Many other proposals for Hamiltonian reconstruction from Gibbs states [4, 5, 6, 7, 8, 9, 10] have recently been advanced. Of these, [7, 8, 9, 10] include numerical tests. In comparison to these approaches, algorithm described in the current work proves to be highly scalable while being robust to noise. Some other advantages are its generality (the choice of variational Hamiltonian terms h_1, \ldots, h_s is unrestricted), ease of implementation, and the fact that it is physically motivated and involves few hyperparameters.

We begin in Section 2 with a description of the algorithm and a guide to interpreting its output. Section 3 concerns a theoretical analysis of the algorithm, including proof of correctness. In Section 4 we describe the results of numerical simulations on a 100-qubit spin chain. We conclude in Section 5 with some future directions for research.

The python implementation of the learning algorithm used in this work, as well as all routines used for testing it, are available for use at

https://github.com/artymowicz/hamiltonian-learning

In the remainder of this section we establish some notation and recall some basic facts about Gibbs states. Let \mathcal{H} be the physical Hilbert space of a system, and assume that dim $\mathcal{H} < \infty$. Write \mathcal{A} for the algebra of all linear operators on \mathcal{H} . As a rule, we will use lowercase letters to denote elements of \mathcal{A} , i.e. operators acting on the physical Hilbert space (this will be to differentiate them from operators on the GNS Hilbert space which we introduce in later sections). We will denote the adjoint of an operator $a \in \mathcal{A}$ by a^* .

A state on \mathcal{A} is a linear map $\omega : \mathcal{A} \to \mathbb{C}$ satisfying $\omega(1) = 1$ and $\omega(a^*a) \ge 0$ for all $a \in \mathcal{A}$. These are in one-to-one correspondence with density matrices ρ via $\omega(a) = \operatorname{tr}(\rho a)$. A state ω is called faithful if $\omega(a^*a) = 0$ implies a = 0, or equivalently, if its density matrix is invertible. Given a selfadjoint operator $h \in \mathcal{A}$, we say h is a symmetry of ω if $\omega([h, a]) = 0$ for all $a \in \mathcal{A}$, or equivalently if $[\rho, h] = 0$. Given a positive number T, the Gibbs state corresponding to h at temperature T is the state given by the density matrix $\rho = e^{-h/T} / \operatorname{tr}(e^{-h/T})$. A Gibbs state is always faithful, and any self-adjoint operator h is a symmetry of its own Gibbs state.

2 Description of the algorithm

We begin now with a description of the algorithm. Let ω be a faithful state of \mathcal{A} . The algorithm requires the following input:

- 1. A set of selfadjoint traceless operators $h_1, \ldots, h_s \in \mathcal{A}$,
- 2. A choice of linearly independent operators $b_1 \dots b_r \in \mathcal{A}$ such that $\operatorname{span}(b_1, \dots, b_r) = \operatorname{span}(b_1^*, \dots, b_r^*)$,
- 3. All expectation values of the form $\omega(b_i^*b_j)$ for $1 \le i, j \le r$ and $\omega(b_i^*[h_\alpha, b_j])$ for $1 \le i, j \le r$ and $1 \le \alpha \le s$.

The operators h_1, \ldots, h_r will be the variational Hamiltonian terms. We call the operators $b_1 \ldots b_r$ the *perturbing operators* – physically, these determine the class of perturbations with respect to which we will enforce stability (see section 3). A reasonable choice for spin systems is the set of all geometrically k-local Pauli operators for some k > 0.

Step 1

The first step is to orthonormalize the b_i in an appropriate sense. Let $a_1, \ldots, a_r \in \text{span}(b_1, \ldots, b_r)$ satisfy $\omega(a_i^*a_j) = \delta_{ij}$. Such a set can be found by diagonalizing the quadratic form $\omega(b_i^*b_j)$ (which is necessarily nondegenerate because ω is faithful).

Step 2

For each $\alpha = 1, \ldots, s$ let \boldsymbol{H}_{α} to be the $r \times r$ matrix

$$(\boldsymbol{H}_{\alpha})_{ij} := \omega(a_i^*[h_{\alpha}, a_j]), \tag{3}$$

and define the $s \times s$ matrix

$$W_{\alpha\beta} := \operatorname{tr}((\boldsymbol{H}_{\alpha}^{\dagger} - \boldsymbol{H}_{\alpha})(\boldsymbol{H}_{\beta} - \boldsymbol{H}_{\beta}^{\dagger})).$$
(4)

If W is invertible, then the algorithm terminates. Otherwise, let $\tilde{h}_1, \ldots, \tilde{h}_q$ be a basis for the kernel of W, and let $(\widetilde{H}_{\alpha})_{ij} = \omega(a_i^*[\tilde{h}_{\alpha}, a_j]).$

Step 3

Define the following $r \times r$ matrix:

$$\Delta_{ij} := \omega(a_j a_i^*),\tag{5}$$

and solve the semidefinite optimization problem:

$$\max_{y \in \mathbb{R}^{q}, T \in \mathbb{R}_{>0}} \mu$$

$$(6)$$

$$T\log(\mathbf{\Delta}) + \sum_{\alpha=1}^{q} y_{\alpha} \widetilde{\boldsymbol{H}}_{\alpha} - \mu I \succeq 0, \tag{7}$$

$$\sum_{\alpha=1}^{q} y_{\alpha} \omega(\tilde{h}_{\alpha}) = -1.$$
(8)

Here (7) is the matrix Araki-Sewell inequality (33) with a regularization parameter, and (8) is a normalization that removes the multiplicative gauge of the Hamiltonian¹.

Let y^*, μ^*, T^* be the optimal values of the semidefinite program. Then the algorithm outputs the candidate Hamiltonian $h^* := \sum_{\alpha} y^*_{\alpha} \tilde{h}_{\alpha}$, candidate temperature T^* and the parameter μ^* which is used in the interpretation of the ouput.

Interpretation

Below is a summary of the possible outputs of the algorithm and their interpretations.

- 1. In step 2, if W is invertible then ω is not a stationary state of any operator in the span of h_1, \ldots, h_s .
- 2. In step 3:
 - (a) If the optimization (6) terminates with $\mu^* < 0$ then ω is not a Gibbs state of any Hamiltonian in the span of h_1, \ldots, h_s^2 .
 - (b) Otherwise, $\sum y_{\alpha}^* \tilde{h}_{\alpha}$ and T^* are candidates for the Hamiltonian and temperature of the state ω .

¹A similar, but not equivalent, algorithm could be obtained by removing the normalization condition (8) and setting T = 1.

²However, numerics suggest that if $\mu^* < 0$ but the magnitude of μ^* is small, then ω is close to being a Gibbs state of $\sum y_{\alpha}^* \tilde{h}_{\alpha}$ (see Section 4).

3 Restricted thermodynamic stability

The theoretical justification of the algorithm will be based on two notions. The first is a relaxation of the notion of a symmetry of a state. For a collection of operators $b_1, \ldots, b_r \in \mathcal{A}$, we say a selfadjoint operator $h \in \mathcal{A}$ is a quasi-symmetry of ω with respect to b_1, \ldots, b_r if $\omega([b^*b, h]) = 0$ for all $b \in \text{span}\{b_1, \ldots, b_r\}$. Any symmetry of ω is a quasi-symmetry.

The second notion is a relaxation of the notion of a Gibbs state. It is inspired by *local thermodynamic stability*, which was introduced by Araki and Sewell in [3] and generalizes the usual notion of thermal equilibrium. Namely, Araki and Sewell considered states in infinite volume for which any *local* perturbation of the state increases the free energy. They showed that these states are characterized by a correlation inequality which is alternatively known as the Araki-Sewell, Roestropff-Araki-Sewell, or energy-entropy balance inequality. Below we introduce a variant of the local thermodynamic stability condition (Definition 1), and prove that it implies a matrix version of the Araki-Sewell inequality first appeared in [2], although a similar inequality can be found in the 1985 article [11].

Let $b_1, \ldots, b_r \in \mathcal{A}$ be any collection of operators (not necessarily self-adjoint). Given an anti-Hermitean matrix M and a positive semidefinite matrix Λ , the Lindbladian superoperator $L_{M,\Lambda}$: $\mathcal{A} \to \mathcal{A}$ is a linear map defined as

$$L_{\boldsymbol{M},\boldsymbol{\Lambda}}(a) := \sum_{ij} \left\{ -\frac{1}{2} \boldsymbol{M}_{ij}[b_i^* b_j, a] + \boldsymbol{\Lambda}_{ij}(b_i^* a b_j - \frac{1}{2}(b_i^* b_j a + a b_i^* b_j)) \right\}, \quad a \in \mathcal{A}.$$
 (9)

It generates a time-evolution $e^{tL_{M,\Lambda}} : \mathcal{A} \to \mathcal{A}$ which describes open dynamics of the system weakly coupled to its environment[12, 13]. The matrics M_{ij} and Λ_{ij} describe the internal couplings and couplitngs to the environment, respectively. Given a self-adjoint $h \in \mathcal{A}$ and a temperature T > 0the free energy of a state ω is

$$F(\omega) = -TS(\omega) + \omega(h) \tag{10}$$

where $S(\omega)$ is the von Neumann entropy (ie. $-\operatorname{tr}(\rho \log \rho)$ of the corresponding density matrix).

Definition 1. Let ω be a faithful state and h a selfadjoint operator. We say the pair (ω,h) satisfies restricted thermodynamic stability (RTS) at the temperature T > 0 with respect to the operators b_1, \ldots, b_r if for any Lindbladian $L_{\mathbf{M}, \mathbf{\Lambda}}$ of the form (9) we have

$$\left. \frac{d}{dt} \right|_{t=0} F(\omega_t) \ge 0,\tag{11}$$

where $\omega_t = \omega \circ e^{tL_{M,\Lambda}}$.

We now give a characterization of the RTS condition in terms of the Gelfand-Naimark-Segal construction (see Appendix A for an exposition). Let \mathcal{H}_{ω} be the GNS Hilbert space of ω , and π_{ℓ} and π_r the left- and right- representations of \mathcal{A} on \mathcal{H}_{ω} . For any self-adjoint $h \in \mathcal{A}$ we write H for the corresponding GNS Hamiltonian:

$$H := \pi_{\ell}(h) - \pi_{r}(h) \in \mathcal{B}(\mathcal{H}_{\omega}).$$
(12)

Note that even though h is self-adjoint, H need not be. In fact one can check that H is self-adjoint if and only if h is a symmetry of ω .

The first-order change in the expectation value of h under the Lindblad evolution (9) has a straight-

forward expression in terms of the matrix elements of H. With $\omega_t = \omega \circ e^{tL_{M,\Lambda}}$, we have

$$\left. \frac{d}{dt} \right|_{t=0} \omega_t(h) = \omega(L_{\boldsymbol{M},\boldsymbol{\Lambda}}(h)) \tag{13}$$

$$= \frac{1}{2} \sum_{ij} \boldsymbol{M}_{ij} \omega(b_i^*[h, b_j] - [b_i^*, h]b_j) + \boldsymbol{\Lambda}_{ij} \omega(b_i^*[h, b_j] + [b_i^*, h]b_j)$$
(14)

$$= \frac{1}{2} \sum_{ij} \boldsymbol{M}_{ij} \langle b_i | \boldsymbol{H} - \boldsymbol{H}^{\dagger} | b_j \rangle + \boldsymbol{\Lambda}_{ij} \langle b_i | \boldsymbol{H} + \boldsymbol{H}^{\dagger} | b_j \rangle.$$
(15)

Let Δ be the modular operator of ω . Analogously to (15) above, the matrix elements of $\log(\Delta)$ give the first-order changes of the entropy of ω under Lindbladian evolution, as we now show. Using a prime to indicate time-derivative at t = 0, we have

$$-\operatorname{tr}(\rho\log(\rho))' = -\operatorname{tr}(\rho'\log(\rho)) - \operatorname{tr}(\rho\log(\rho)').$$
(16)

Using the power series of log about the identity operator and the cyclicity of the trace, the second term can be seen to equal $-\operatorname{tr}(\rho') = 0$. The first term, meanwhile, equals

$$-\operatorname{tr}(\rho'\log(\rho)) = -\operatorname{tr}(\rho L_{\boldsymbol{M},\boldsymbol{\Lambda}}(\log(\rho)))$$
(17)

$$= -\omega(L_{\boldsymbol{M},\boldsymbol{\Lambda}}(\log\rho)) \tag{18}$$

$$= \boldsymbol{M}_{ij} \langle b_i | \frac{\log(\Delta) - \log(\Delta)^{\dagger}}{2} | b_j \rangle + \boldsymbol{\Lambda}_{ij} \langle b_i | \frac{\log(\Delta) + \log(\Delta)^{\dagger}}{2} | b_j \rangle,$$
(19)

where in the last line we used (15) and the expression

$$\log(\Delta) = \pi_{\ell}(\log(\rho)) - \pi_{r}(\log(\rho))$$
(20)

from Appendix A. Since ρ and $\log(\rho)$ commute, $\log(\Delta)$ is Hermitean and the first term in (19) vanishes, leaving us with

$$\frac{d}{dt}\bigg|_{t=0} S(\omega_t) = -\sum_{i,j} \mathbf{\Lambda}_{ij} \langle b_i | \log(\Delta) | b_j \rangle.$$
(21)

Equations (15) and (21) allow us to characterize the RTS and quasi-symmetry conditions as follows:

Proposition 1. Let ω be a faithful state, $h \in \mathcal{A}$ a selfadjoint operator, and H given by (12). Let $b_1, \ldots, b_r \in \mathcal{A}$ a collection of operators and write $P : \mathcal{H}_{\omega} \to \operatorname{span}(|b_1\rangle, \ldots, |b_r\rangle)$ for the orthogonal projection.

- i) h is a quasi-symmetry of ω with respect to the operators b_1, \ldots, b_r if and only if PHP^{\dagger} is self-adjoint.
- ii) The pair (ω, h) has RTS at temperature T > 0 with respect to the operators b_1, \ldots, b_r if and only if

$$P(T\log\Delta + H)P^{\dagger} \ge 0. \tag{22}$$

Proof. i) PHP^{\dagger} is self-adjoint if and only if

$$0 = \langle b|H - H^{\dagger}|b\rangle \tag{23}$$

$$=\omega(b^*[H,b]) - \omega([b^*,H]b)$$
(24)

$$= 2\omega([b^*b,h]) \tag{25}$$

for any $b \in \operatorname{span}\{b_1, \ldots, b_r\}$.

ii) Suppose $P(T \log \Delta + H)P^{\dagger} \ge 0$. Since $\log \rho$ commutes with ρ , $\log \Delta$ is hermitian, and so PHP^{\dagger} hermitian too. It follows that for any anti-hermitian M and any positive-semidefinite Λ we have

$$\frac{d}{dt}\Big|_{t=0} F(\omega \circ e^{tL_{M,\Lambda}}) = \sum_{i,j} \Lambda_{ij} \langle b_i | T \log(\Delta) + H | b_j \rangle \ge 0.$$
(26)

Conversely, suppose ω satisfies RTS. Setting M = 0 in (9) gives

$$0 \le \frac{d}{dt} \bigg|_{t=0} F(\omega \circ e^{tL_{\Lambda}}) = \sum_{i,j} \Lambda_{ij} \langle b_i | T \log(\Delta) + H | b_j \rangle.$$
⁽²⁷⁾

Since this holds for any positive-semidefinite $r \times r$ matrix $\mathbf{\Lambda}$, it follows that $P(T \log(\Delta) + H)P^{\dagger} \geq 0$.

We can in short order deduce from Proposition 1 that the RTS condition depends only on the span of $|b_1\rangle, \ldots, |b_r\rangle$, and that if (ω, h) satisfy RTS then h is necessarily a quasi-symmetry of ω . The next Proposition relates the RTS condition to the Gibbs condition.

Proposition 2. Let $b_1, \ldots, b_r \in \mathcal{A}$ be a collection of operators, $h \in \mathcal{A}$ a selfadjoint operator, and T > 0. If ω is the Gibbs state of h at temperature T, then the pair (ω, h) has RTS at temperature T with respect to the operators b_1, \ldots, b_r . The converse holds if $\operatorname{span}(b_1, \ldots, b_r) = \mathcal{A}$.

Proof. If ω is the Gibbs state of h at temperature T then it is easy to check using the explicit expression (20) that $T \log(\Delta) + H = 0$, so thanks to part *ii*) of Proposition 1, ω automatically satisfies RTS for any set of perturbing operators.

For the second statement, suppose that ω and h satisfy RTS with respect to b_1, \ldots, b_r and that $\operatorname{span}(b_1, \ldots, b_r) = \mathcal{A}$. By part ii) of Proposition 1 we have $T \log \Delta + H \ge 0$. We will show that this implies $T \log \Delta + H = 0$. Let J be the modular involution, which is defined as

$$J|a\rangle := |\rho^{1/2}a^*\rho^{-1/2}\rangle.$$
(28)

For any $a \in \mathcal{A}$ we have

$$\langle Ja|H|Ja\rangle = \operatorname{tr}(\rho \,\rho^{-1/2} a \rho^{1/2} [h, \rho^{1/2} a^* \rho^{-1/2}]) \tag{29}$$

$$= \operatorname{tr}(a\rho^{1/2}h\rho^{1/2}a^*) - \operatorname{tr}(\rho^{1/2}a\rho a^*\rho^{-1/2}h)$$
(30)

$$= -\omega(a^*[h,a]) \tag{31}$$

$$= -\langle a|H|a\rangle,\tag{32}$$

where in the third line we used the fact that $[\rho, h] = 0$. The same calculation with $\log(\rho)$ replacing h shows that $\langle Ja | \log(\Delta) | Ja \rangle = -\langle a | \log(\Delta) | a \rangle$. It follows that for every eigenvalue λ of $T \log(\Delta) + H$, $-\lambda$ is also an eigenvalue. Combined with the fact that $T \log \Delta + H \ge 0$, we conclude that $T \log \Delta + H = 0$, as claimed. From this we see that $\log(\rho) - h/T$ is in the center of \mathcal{A} , which means it is a multiple of the identity, and so $\rho = e^{-h/T} / \operatorname{tr}(e^{-h/T})$.

Now we come to the matrix analog of the Araki-Sewell inequality.

Proposition 3. If (ω, h) satisfies RTS at T > 0 then

$$T\log(\mathbf{\Delta}) + \mathbf{H} \ge 0. \tag{33}$$

The converse holds if span $(|b_1\rangle, \ldots, |b_r\rangle)$ is invariant under the modular flow of ω .

Proof. The matrix $\log(\Delta)$ is related to $\log(\Delta)$ by an operator version of Jensen's inequality [14]:

$$\log(\mathbf{\Delta}) = \log(P \Delta P^{\dagger}) \ge P \log(\Delta) P^{\dagger}, \tag{34}$$

and so we have:

$$T\log(\mathbf{\Delta}) + \mathbf{H} \ge P(T\log\Delta + H)P^{\dagger},\tag{35}$$

which, thanks to Proposition 1, proves the first statement. For the second statement, $\operatorname{span}(|b_1\rangle, \ldots, |b_r\rangle)$ is invariant under the modular flow of ω precisely when P commutes with Δ , in which case we have equality in (35).

We conclude this section with a correctness result for the algorithm.

Corollary 1. Suppose ω is a Gibbs state of a Hamiltonian $h = \sum_{\alpha=1}^{s} y_{\alpha}h_{\alpha}$ at temperature T, and let (h^*, T^*, μ^*) be the output of the semidefinite program (6) with variational Hamiltonian terms h_1, \ldots, h_s .

i) (Feasibility) The triple (h, T, 0) is a feasible point of the program (6). In particular $\mu^* \ge 0$.

ii) (Recoverability) Suppose b_1, \ldots, b_r span the entire algebra \mathcal{A} . Then $(h^*, T^*, \mu^*) = (h, T, 0)$.

Proof. i): by Proposition 2, ω satisfies RTS. The statement then follows from Proposition 3.

ii): if b_1, \ldots, b_r span \mathcal{A} then any solution of the constraint (7) with $\mu \geq 0$ satisfies (33), which by Propositions 2 and 3 implies that ω is the Gibbs state h^* at temperature T^* , and so $(h, T) = (h^*, T^*)$ and $T \log(\mathbf{\Delta}) + \mathbf{H} = 0$. This last equality implies in turn that $\mu = 0$.

Theoretical justification of algorithm

Now we make explicit the connection between the results proved in this section and the algorithm. The operators a_1, \ldots, a_r constructed in step 1 are chosen so that $|a_1\rangle, \ldots, |a_r\rangle$ form an orthonormal basis of span $(|b_1\rangle, \ldots, |b_r\rangle)$, and we have

$$\boldsymbol{\Delta}_{ij} := \omega(a_j a_i^*) = \langle a_i | \Delta | a_j \rangle \tag{36}$$

and

$$\boldsymbol{H}_{ij} := \omega(a_i^*[h, a_j]) = \langle a_i | H | a_j \rangle.$$
(37)

for any selfadjoint $h \in \mathcal{A}$. In other words we have $\Delta = P\Delta P^{\dagger}$ and $H = PHP^{\dagger}$.

In step 2, the kernel of W computes the space of quasi-symmetries of ω that lie in the span of h_1, \ldots, h_α . Indeed, by Proposition 1 *i*), a self-adjoint $h \in \mathcal{A}$ is a quasi-symmetry if and only if H is self-adjoint, and we have have $W = F^{\dagger}F$, where $F : \mathbb{R}^s \to Mat_{r \times r}(\mathbb{C})$ is the map taking $x \mapsto \sum_{\alpha=1}^s x_\alpha (H_\alpha - H_\alpha^{\dagger})$.

Correctness of step 3 is given by Corollary 1.

4 Numerical results

We now turn to the results of numerical simulations of the algorithm.

Corrections to ideal algorithm

In the presence of noise in the expectation values of ω , it is inappropriate to use the exact kernel of the matrix W constructed in step 2. Instead, a threshold $\epsilon_W > 0$ was chosen as a hyperparameter, and we let $\tilde{h}_1, \ldots, \tilde{h}_q$ be the eigenvectors of W with eigenvalue smaller than ϵ_W . Since the matrices \widetilde{H}_{α} can no longer be assumed to be self-adjoint, we symmetrize their definition:

$$(\widetilde{\boldsymbol{H}}_{\alpha})_{ij} := \frac{\omega(a_i^*[\widetilde{h}_{\alpha}, a_j]) + \omega(a_j^*[\widetilde{h}_{\alpha}, a_i])}{2}.$$
(38)

The proof of feasibility in Corollary 1 still holds, since for $\epsilon_W = 0$ this reduces to the original algorithm, while increasing ϵ_W (for a fixed ω) increases μ^* (indeed, increasing ϵ_W only increases the feasible region of the program (6)).

For sufficiently small values of σ_{noise} , the spectrum of W was found to have a low-lying part with a spectral gap to the rest of the eigenvalues. ϵ_W was set according to the formula

$$\epsilon_W = 400 \max(\sigma_{noise}^2 \sqrt{m}, \ 10^{-11}),$$
(39)

where *m* denotes the number of terms $a_i^*[h_\alpha, a_j]$ such that $[h_\alpha, a_j] \neq 0$. This formula was found empirically to produce an ϵ_W lying in the spectral gap of *W*, and is not expected to be universal across different values of *n* and choices of perturbing operators. We note that in practice, while choosing ϵ_W to be too low caused the output to be inaccurate, choosing ϵ_W to lie above the gap did not significantly affect the accuracy of the result.

Learning the XXZ Hamiltonian

The MPS purification technique [15] was used to prepare thermal states of the following anisotropic Heisenberg ferromagnet:

$$h = -\sum_{i=1}^{n-1} (\sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y + \frac{1}{2} \sigma_i^y \sigma_{i+1}^y).$$
(40)

with n = 100. Both the set of perturbing operators b_1, \ldots, b_r and the set of variational Hamiltonian terms h_1, \ldots, h_s were chosen to be the 1192 geometrically 2-local Pauli operators. Measurement error was simulated by adding Gaussian noise with variance σ_{noise} to the expectation value of each Pauli operator. The learning algorithm itself was implemented in Python, using the MOSEK solver [16] for the semidefinite optimization³.

Hamiltonian recovery error was quantified using the overlap as in [19]: let $y \in \mathbb{R}^s$ be the vector of recovered Hamiltonian coefficients and $z \in \mathbb{R}^s$ the vector of true Hamiltonian coefficients. The Hamiltonian recovery error is then defined as the relative angle of the two, which for small angles approximately equals the reciprocal of the signal-to-noise ratio:

$$\theta = \arccos\left(\frac{|\langle y|z\rangle|}{\|y\|\|z\|}\right) \approx \frac{\|y-z\|}{\|z\|}.$$
(41)

Note that this metric is not sensitive to the overall scaling of the Hamiltonian. This degree of freedom of the Hamiltonian is fixed by the normalization (8) anyway and instead appears in the temperature T. Interestingly, the algorithm reconstructed the "projective" degrees of freedom of the Hamiltonian terms much more accurately than it did its overall scale (or equivalently, the temperature).

 $^{^{3}}$ The convex modeling language CVXPy [17] and the open source solver SCS [18] were used in prototyping but not in the final code.



Figure 1: Numerical results for the 100-qubit anisotropic Heisenberg model (40) at several temperatures. Left: Recovery error θ as a function of noise amplitude σ_{noise} , averaged over 10 runs. Dotted line is (mean) + (standard deviation). Right: Ratio of recovered temperature to actual temperature, averaged over 10 runs. Shaded region is (mean) \pm (standard deviation).

The Hamiltonian recovery error θ and the recovered temperature T are plotted against σ_{noise} in Figure 1. A temperature-dependent noise threshold is found between $\sigma_{noise} \approx 10^{-5}$ and $\sigma_{noise} \approx 10^{-3}$ above which the matrix Δ ceases to be positive definite. The algorithm could possibly be emended to work for higher noise values by projecting onto the orthocomplement of the nonpositive eigenspace of Δ , but we leave this to future work.

As one shrinks the noise amplitude, the recovery error first decreases (for high temperatures, this decrease is linear to a good approximation). This persists up until, at some temperature-dependent critical value of the noise amplitude, the recovery error plateaus. We interpret this two-stage behaviour as follows. In the limit of zero measurement error, perfect recovery is not guaranteed because the condition (33) is weaker than the Gibbs condition. Instead, (33) defines a convex set of candidate Hamiltonians, and the algorithm picks one of these by maximizing the regularization parameter μ . The recovery error is then on the order of the diameter of this convex set. Thus for low enough levels of measurement noise the recovery error is roughly constant.

We posit that the only way to lower the levels of these plateaux is to enlarge the set of perturbing operators, which tightens the constraint (33). This is relevant if one wants to prove asymptotic bounds on the number of copies of the state and the computational resources needed to specify the Hamiltonian up to an arbitrarily low error.

Note however that for the particular Hamiltonian under consideration, the plateaux start at noise amplitudes σ_{noise} of around 10^{-9} to 10^{-8} . Assuming that expectation values are estimated from independent copies of the state, this would require on the order of 10^{16} to 10^{18} samples, far beyond what is experimentally feasible anyway. So for practical applications it may be more important to understand the high-noise regime rather than the locations of the plateaux.

5 Discussion and outlook

Let us conclude by describing two directions for future research. While Corollary 1 establishes the correctness of the algorithm, it suffers from two important limitations which must be overcome if one is to prove sample complexity and computational complexity bounds. First, neither the feasibility nor the recoverability statements of Corollary 1 take into account measurement noise in

the expectation values of the state, which is unavoidable whenever these are estimated using finitely many copies of the state.

Second, the recoverability statement only holds when the set of perturbing operators is grown to a complete set of operators. The utility of this algorithm depends on approximate recoverability when the set of perturbing operators is far smaller than a complete set. Section 4 gives numerical evidence that this is indeed the case, but a proof is still lacking.

The second point above leads one to consider states that satisfy RTS for some Hamiltonian h, but are not Gibbs states of h. The algorithm described here can never tell these apart from true Gibbs states. Aside from their implications for the analysis of the algorithm, it is an interesting physical question to rule out or characterize such "false Gibbs states", which are thermodynamically stable against some perturbations but not others.

Acknowledgements

The author would like to thank Anton Kapustin, Hsin-Yuan (Robert) Huang, and Elliott Gesteau for helpful discussions and comments on the draft.

References

- Brian Swingle and Isaac H. Kim. "Reconstructing Quantum States from Local Data". In: *Physical Review Letters* 113.26 (Dec. 2014). ISSN: 1079-7114. DOI: 10.1103/physrevlett. 113.260501. URL: http://dx.doi.org/10.1103/PhysRevLett.113.260501.
- Hamza Fawzi, Omar Fawzi, and Samuel O. Scalet. Certified algorithms for equilibrium states of local quantum Hamiltonians. DOI: 10.48550/arXiv.2311.18706. arXiv: 2311.18706
 [quant-ph].
- Huzihiro Araki and Geoffrey L. Sewell. "KMS conditions and local thermodynamical stability of quantum lattice systems". In: *Communications in Mathematical Physics* 52.2 (June 1977), pp. 103-109. ISSN: 1432-0916. DOI: 10.1007/bf01625778. URL: http://dx.doi.org/10. 1007/BF01625778.
- [4] Anurag Anshu et al. "Sample-efficient learning of interacting quantum systems". In: Nature Physics 17.8 (May 2021), pp. 931–935. ISSN: 1745-2481. DOI: 10.1038/s41567-021-01232-0. URL: http://dx.doi.org/10.1038/s41567-021-01232-0.
- [5] Jeongwan Haah, Robin Kothari, and Ewin Tang. "Optimal learning of quantum Hamiltonians from high-temperature Gibbs states". In: 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS). IEEE, Oct. 2022. DOI: 10.1109/focs54457.2022.00020. URL: http://dx.doi.org/10.1109/F0CS54457.2022.00020.
- [6] Ainesh Bakshi et al. Learning quantum Hamiltonians at any temperature in polynomial time. 2023. DOI: 10.48550/ARXIV.2310.02243. URL: https://arxiv.org/abs/2310.02243.
- [7] Andi Gu, Lukasz Cincio, and Patrick J. Coles. "Practical Hamiltonian learning with unitary dynamics and Gibbs states". In: *Nature Communications* 15.1 (Jan. 2024). ISSN: 2041-1723. DOI: 10.1038/s41467-023-44008-1. URL: http://dx.doi.org/10.1038/s41467-023-44008-1.
- [8] Tim J. Evans, Robin Harper, and Steven T. Flammia. Scalable Bayesian Hamiltonian learning. 2019. DOI: 10.48550/ARXIV.1912.07636. URL: https://arxiv.org/abs/1912.07636.
- Yotam Y. Lifshitz et al. Practical Quantum State Tomography for Gibbs states. 2021. DOI: 10.48550/ARXIV.2112.10418. URL: https://arxiv.org/abs/2112.10418.

- [10] Eyal Bairey, Itai Arad, and Netanel H. Lindner. "Learning a Local Hamiltonian from Local Measurements". In: *Physical Review Letters* 122.2 (Jan. 2019). ISSN: 1079-7114. DOI: 10.1103/ physrevlett.122.020504. URL: http://dx.doi.org/10.1103/PhysRevLett.122.020504.
- [11] Eberhard E. Muller. "Note on relative entropy and thermodynamical limit". In: Helvetica Physica Acta 58 (Oct. 1985), pp. 622-632. URL: https://www.e-periodica.ch/cntmng?pid= hpa-001:1985:58::1152.
- G. Lindblad. "On the generators of quantum dynamical semigroups". In: Communications in Mathematical Physics 48.2 (June 1976), pp. 119–130. ISSN: 1432-0916. DOI: 10.1007/ bf01608499. URL: http://dx.doi.org/10.1007/BF01608499.
- Robert Alicki and Karl Lendi. Quantum Dynamical Semigroups and Applications. Springer Berlin Heidelberg, 2007. ISBN: 9783540708605. DOI: 10.1007/3-540-70861-8. URL: http: //dx.doi.org/10.1007/3-540-70861-8.
- [14] Frank Hansen and Gert K. Pedersen. "Jensen's Operator Inequality". In: Bulletin of the London Mathematical Society 35.04 (June 2003), pp. 553–564. ISSN: 1469-2120. DOI: 10.1112/s0024609303002200. URL: http://dx.doi.org/10.1112/S0024609303002200.
- [15] Adrian E. Feiguin and Steven R. White. "Finite-temperature density matrix renormalization using an enlarged Hilbert space". In: *Physical Review B* 72.22 (Dec. 2005). ISSN: 1550-235X. DOI: 10.1103/physrevb.72.220401. URL: http://dx.doi.org/10.1103/PhysRevB.72.220401.
- [16] Erling D. Andersen and Knud D. Andersen. "The Mosek Interior Point Optimizer for Linear Programming: An Implementation of the Homogeneous Algorithm". In: *High Performance Optimization*. Springer US, 2000, pp. 197–232. ISBN: 9781475732160. DOI: 10.1007/978-1-4757-3216-0_8. URL: http://dx.doi.org/10.1007/978-1-4757-3216-0_8.
- [17] Steven Diamond and Stephen Boyd. "CVXPY: a python-embedded modeling language for convex optimization". In: J. Mach. Learn. Res. 17.1 (Jan. 2016), pp. 2909–2913. ISSN: 1532-4435.
- Brendan O'Donoghue et al. "Conic Optimization via Operator Splitting and Homogeneous Self-Dual Embedding". In: Journal of Optimization Theory and Applications 169.3 (Feb. 2016), pp. 1042–1068. ISSN: 1573-2878. DOI: 10.1007/s10957-016-0892-3. URL: http://dx.doi. org/10.1007/s10957-016-0892-3.
- [19] Xiao-Liang Qi and Daniel Ranard. "Determining a local Hamiltonian from a single eigenstate". In: *Quantum* 3 (July 2019), p. 159. ISSN: 2521-327X. DOI: 10.22331/q-2019-07-08-159. URL: http://dx.doi.org/10.22331/q-2019-07-08-159.
- [20] Ola Bratteli and Derek W. Robinson. Operator Algebras and Quantum Statistical Mechanics I. Springer Berlin Heidelberg, 1987. ISBN: 9783662025208. DOI: 10.1007/978-3-662-02520-8. URL: http://dx.doi.org/10.1007/978-3-662-02520-8.
- [21] Masamichi Takesaki. "Fundamentals of Banach Algebras and C*-Algebras". In: Theory of Operator Algebras I. Springer New York, 1979, pp. 1–57. ISBN: 9781461261889. DOI: 10.1007/ 978-1-4612-6188-9_1. URL: http://dx.doi.org/10.1007/978-1-4612-6188-9_1.

A The Gelfand-Naimark-Segal construction

In this appendix we give a brief introduction to Gelfand-Naimark-Segal (GNS) construction which is central to the theoretical analysis in section 3. The GNS construction is a fundamental tool in the theory of operator algebras and is well described in standard references [20, 21]. However, these references work in the general setting of infinite-dimensional C*-algebras, where the discussion is plagued with analytic subtleties. We avoid these subtleties because we work in finite dimensions, so it is worthwhile to introduce the concepts we need in a streamlined and self-contained manner here.

Let ω be a faithful state of $\mathcal{A} = \mathcal{B}(\mathcal{H})$, and $\rho \in \mathcal{A}$ its density matrix. The bilinear form $(a, b) \mapsto \omega(a^*b)$ endows \mathcal{A} with the structure of a Hilbert space, which we call the GNS space and denote by the symbol \mathcal{H}_{ω} (although \mathcal{A} and \mathcal{H}_{ω} are isomorphic as vector spaces, it is important to maintain their distinction as mathematical objects). For an operator $a \in \mathcal{A}$ we denote the corresponding vector in the GNS space by $|a\rangle \in \mathcal{H}_{\omega}$. The GNS space \mathcal{H}_{ω} carries two distinguished representations of the algebra \mathcal{A} . These are the *left* representation:

$$\pi_{\ell}(a) : |b\rangle \mapsto |ab\rangle, \tag{42}$$

and the right representation:

$$\pi_r(a): |b\rangle \mapsto |ba^*\rangle. \tag{43}$$

Notice that these representations commute in the sense that $\pi_r(a)\pi_\ell(b)|c\rangle = \pi_\ell(b)\pi_r(a)|c\rangle$ for all $a, b, c \in \mathcal{A}$. The modular operator Δ is the complex-linear operator on \mathcal{H}_ω corresponding to the sequilinear form $(a, b) \mapsto \omega(ba^*)$. That is, for any $a, b \in \mathcal{A}$ we have

$$\langle a|\Delta|b\rangle := \omega(ba^*) \tag{44}$$

$$= \operatorname{tr}(\rho b a^*) \tag{45}$$

$$= \operatorname{tr}(\rho a^*(\rho b \rho^{-1})) \tag{46}$$

$$=\omega(a^*(\rho b \rho^{-1})) \tag{47}$$

$$= \langle a | \pi_{\ell}(\rho) \pi_{r}(\rho^{-1}) | b \rangle, \tag{48}$$

and so Δ can be expressed as

$$\Delta = \pi_{\ell}(\rho)\pi_r(\rho^{-1}). \tag{49}$$

The modular flow of ω is the one-parameter group of unitaries on \mathcal{H}_{ω} given by $t \mapsto \Delta^{it}$. It is generated (in the sense of the Stone's theorem) by the selfadjoint operator $\log(\Delta)$, which can be expressed as

$$\log(\Delta) = \pi_{\ell}(\log(\rho)) - \pi_r(\log(\rho)).$$
(50)