Undecidability and algorithmic randomness for spin chains

André Nies
School of Computer Science
THE UNIVERSITY OF AUCKLAND
NEW ZEALAND

LQCAI conference, IRFS, Arak University (online)
July 1, 2021

Spectral gap problem undecidable (Cubitt et al.)

Bausch, Cubitt, Lucia and Perez-Gracia, 2020, starting from work of Cubitt, Perez-Garcia and Wolf in Nature 2015, showed that the spectral gap problem is undecidable for (1D) spin chains.

For each input w to a universal Turing machine, they construct Hamiltonians $H_n(w)$ via nearest neighbour interactions so that asymptotically there is a spectral gap if and only if the machine does not halt on w.

Randomness for infinite qubit sequences (N. and Scholz, J. Math. Physics 2019)

In work (much less known) with Volkher Scholz, we have considered the notion of Martin-Loef randomness for infinite bit sequences, and extended it to infinite spin chains of qubits.

We showed that there is a universal algorithmic test for randomness, and worked towards a characterisation of this randomness notion via incompressibility of the initial segments, similar to the Levin-Schnorr theorem.

A 2021 PhD thesis and corresponding three publications (JMP, TCS, ENTCS) by Tejas Bhojraj at U Madison have further advanced research on this topic.
We will discuss these two directions of research, and in concluding remarks speculate about their potential connections.

Finite spin chains

Spin chains and spin lattices

Spin chains were introduced to understand magnetism. A classical spin chain consists of N dipoles arranged linearly:

$$\underbrace{\uparrow\downarrow\downarrow\uparrow\ \dots\ \uparrow}_{N}$$

Higher-dimensional arrangements of dipoles have also be studied, in particular square lattices.

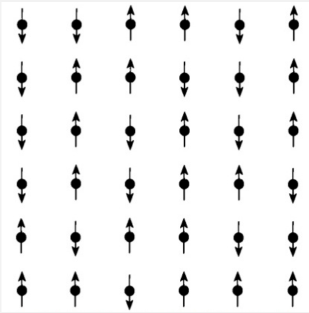

Ising model in 1D: Hamiltonian

The 1D Ising model is due to Lenz (1920), and was “solved” by his student Ising in his thesis (1925).
The positions $i = 1, \dots, N$ in a spin chain are called sites.
The energy of a state of the system is given by a Hamiltonian.

For the 1D Ising model with N sites, the Hamiltonian is

$$H_N = -J \sum_{i=1}^{N-1} \sigma_i \sigma_{i+1} - h \sum_{j=1}^{N} \sigma_j$$

- J is the interaction strength between neighbours,
- h is the strength of the external magnetic field,
- $\sigma_i = 1$ for $\uparrow$ at site i, and $\sigma_i = -1$ for $\downarrow$ at site i.

Quantum setting: Heisenberg (1928) model

- n-chain, each site contains a spin $1/2$ particle (e.g., electron).
- state is unit vector in $(\mathbb{C}^2)^{\otimes n}$
- Spins in x, y, z directions, corresponding to observables given by the Pauli matrices $\sigma^x, \sigma^y, \sigma^z$ (certain 2×2 matrices over $\mathbb{C}$).
- Physicists write $\vec{\sigma} = (\sigma^x, \sigma^y, \sigma^z)$ and
$$\vec{\sigma}_k = \mathbb{I} \otimes \dots \otimes \mathbb{I} \otimes \vec{\sigma} \otimes \mathbb{I} \otimes \dots \otimes \mathbb{I},$$
where $\mathbb{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, and the $\vec{\sigma}$ is in position k.

The Hamiltonian is now a Hermitian operator on $(\mathbb{C}^2)^{\otimes n}$:

$$H = \sum_{i=1}^{n-1} h_{i,i+1}^{(2)} \text{ where } h_{i,i+1}^{(2)} = \frac{J}{4}(\vec{\sigma}_i \cdot \vec{\sigma}_{i+1} - \mathbb{I}^{\otimes n}).$$

$J \in \mathbb{R}$ is a coupling constant, and the local Hamiltonian $h_{i,i+1}^{(2)}$ describes the interaction of neighbouring sites.

Abstract spin chains

For $d \geq 2$ (sometimes suppressed), a qudit is a unit vector in d-dimensional Hilbert space $\mathbb{C}^d$.

- An abstract spin chain is a system of n qudits, arranged linearly. The positions are referred to as sites.
- The state of such a system is given by a vector in the d^n-dimensional Hilbert space $(\mathbb{C}^d)^{\otimes n}$.

One also considers higher dimensional arrangements of qudits, e.g. square lattices.

Cubitt, Perez-Garcia and Wolf (Undecidability of the Spectral Gap, Nature 528, 2015, 5 pages) showed that whether there is a spectral gap is undecidable for the square lattice (2D) case.

The full proof has last been updated on arXiv in 2020 (1502.04573v4), and now stands at 126 pages.

Later on, Bausch, Cubitt, Lucia and Perez-Garcia (Phys. Review X.10, 2020, 20 pages) showed that the existence of a spectral gap is undecidable for the spin chain (1D) case.

Local Hamiltonians

Let $M_n(\mathbb{C})$ denote the algebra of $n \times n$ complex matrices.
As in the case of the Ising and Heisenberg chains, the behaviour of an abstract spin chain is described by local Hamiltonians.
Let $h^{(1)} \in M_d(\mathbb{C})$ and $h^{(2)} \in M_{d^2}(\mathbb{C})$ be Hermitian matrices, where

- $h^{(1)}$ describes the one-site "interactions", and
- $h^{(2)}$ describes the nearest-neighbour interactions.

The global Hamiltonian of a spin chain of n qudits is given by shifting and adding up these interactions as the indices vary:

$$H_n = \sum_{i=1}^{n} h_i^{(1)} + \sum_{i=1}^{n-1} h_{i,i+1}^{(2)}.$$

Spectral gap

The spectral gap of a Hamiltonian H acting on a finite-dimensional Hilbert space is $\Delta(H) = \gamma_1(H) - \gamma_0(H)$, the difference between its least two eigenvalues.
$\langle H_n \rangle_{n\in\mathbb{N}}$ will always denote a sequence such that H_n is a Hamiltonian on the d^n-dimensional Hilbert space.

The asymptotic spectral gap of such a sequence can be defined as

$$\Delta\langle H_n \rangle = \liminf_n \Delta(H_n).$$

(Note that the ground energy $\lambda_0(H_n)$ might increase with n.)
Naively one would say that

the system is gapped if $\Delta\langle H_n \rangle$ is positive, and gapless otherwise.

Gapped and gapless sequences of Hamiltonians

Cubitt et al. (2015) and then Bausch et al. (2020) use definitions making both the gapped and the gapless case more restricted, so that some sequences have neither property.

$\langle H_n \rangle$ is gapped if $\Delta\langle H_n \rangle = \liminf_n \Delta(H_n)$ is positive, moreover, for sufficiently large n, the least eigenvalue $\lambda_0(H_n)$ is non-degenerate, i.e. its eigenspace has dimension 1.

Physically the second condition means that there is a unique ground state of the system (up to phase).

$\langle H_n \rangle$ is gapless if there is some $c > 0$ such that for each $\varepsilon > 0$, for sufficiently large n, each point in the interval $[\lambda_0(H_n), \lambda_0(H_n) + c]$ is ε-close to some eigenvalue of H_n.

(a) $\langle H_n \rangle$ is gapped if $\Delta\langle H_n\rangle = \liminf_n \Delta(H_n)$ is positive and for sufficiently large n, the least eigenvalue $\lambda_0(H_n)$ is non-degenerate, i.e. its eigenspace has dimension 1.
(b) $\langle H_n \rangle$ is gapless if there is some $c > 0$ such that for each $\varepsilon > 0$, for sufficiently large n, each point in the interval $[\lambda_0(H_n), \lambda_0(H_n) + c]$ is ε-close to some eigenvalue of H_n.

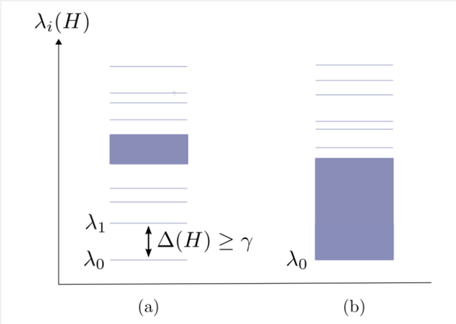

From Cubitt et al., Nature 2015

In the 1-dimensional case, whether there is a spectral gap was shown to be undecidable (Bausch et al., 2020).

Given a Turing machine M, they determine a (large) dimension d. Then, given an input $\eta \in \mathbb{N}$ to M they compute local Hamiltonians $h^{(1)} \in M_d(\mathbb{C})$ and $h^{(2)} \in M_{d^2}(\mathbb{C})$ as above such that

- if $M(\eta)$ halts then the sequence $\langle H_n(\eta) \rangle$ (defined as above by shifting the local interactions) is gapless,
- otherwise the sequence $\langle H_n(\eta) \rangle$ is gapped.

They rely on the methods Cubitt et al. (2015) who showed the spectral gap problem is undecidable in the 2D cas, using square lattices of qudits. The definitions are similar except but there are two types of nearest-neighbor interactions, corresponding to rows and columns.

Remarks

1. Interestingly, in the 2D case the relationship between machines and Hamiltonians is the other way round: if $M(\eta)$ halts then the sequence is gapped, else gapless.
2. The entries of the Hamiltonians are easy "complex" numbers:
   - Let F be the subring of $\mathbb{C}$ generated by
   $$\mathbb{Q} \cup \{\sqrt{2}\} \cup \{\exp(2\pi i \theta) \colon \theta \in \mathbb{Q}\}.$$
   - The entries of the local Hamiltonians, and hence of the $H_n(\eta)$, are all in F.
   - So the undecidability of the spectral gap is not an artefact of the well-known fact that equality of two computable reals is undecidable.

Elements of the proofs in 2D and 1D

2D:

- quantum Turing machines (Bernstein and Vazirani)
- history state Hamiltonian due to Feynman, then Kitaev
- Gottesman and Irani (FOCS 2013): The ground state encodes the whole computation of a QTM up to stage T.
- The QTM is not related to M, rather it is related to the phase estimation algorithm (e.g. Nielsen/Chuang)
- Quasi periodic Wang tiling due to Robinson (Inventiones, 1971).

1D:
the Wang tiling (which needed the second spatial dimension in the lattice) is replaced by a "marker Hamiltonian".

Infinite spin chains

Background

- Martin-Löf randomness (1966) is a key concept to formalize our intuition of infinite bit sequences which “look” random.
- We extend this notion to the setting of infinite qubit sequences.
- First we need to clarify what we mean by such a sequence.
- Can’t use mathematical sequences of qubits because different entries of a sequence can be entangled.

Finite sequences of quantum bits

- The pure state of a system of n qubits is a unit vector in the tensor power $(\mathbb{C}^2)^{\otimes n} := \underbrace{\mathbb{C}^2 \otimes \ldots \otimes \mathbb{C}^2}_{n}$.
- The standard basis of $(\mathbb{C}^2)^{\otimes n}$ is given by n-bit strings: it consists of vectors
$$|a_1 \ldots a_n\rangle := |a_1\rangle \otimes \ldots \otimes |a_n\rangle.$$
- The state $|\psi\rangle$ of a system of n qubits is a linear superposition of them. Example: EPR state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$.
- $|\psi\rangle\langle\psi|$ is the operator projecting $(\mathbb{C}^2)^{\otimes n}$ onto $|\psi\rangle$.
- A mixed state is a statistical combination of orthogonal pure states, i.e. $S = \sum_{i=1}^{2^n} \alpha_i |\psi_i\rangle\langle\psi_i|$ where $0 \leq \alpha_i \leq 1$ and $\sum_i \alpha_i = 1$. One calls S a density matrix.

The C^*-algebra M_∞

- M_d denotes the C^*-algebra of $d \times d$ matrices A over $\mathbb{C}$.
- There is a natural embedding $M_{2^n} \to M_{2^{n+1}}$ via
$$A \to A \otimes I_2 = \begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix}.$$
- The limit M_∞ is the norm completion of $\bigcup_n M_{2^n}$.
- Partial trace operation $T_n \colon M_{2^{n+1}} \to M_{2^n}$.
- If $A \in M_{2^{n+1}}$ is a density matrix then $T_n(A)$ is A with the last qubit erased.

States on M_∞ are sequences of density operators

$\mathcal{S}(M_\infty)$ denotes the sequences $(\rho_n)_{n\in\mathbb{N}}$ of density matrices in M_{2^n} that are coherent in that $T_n(\rho_{n+1}) = \rho_n$ for each n.

- This is the set of states (positive linear functionals of norm 1) on the computable C^* algebra M_∞
- A classical bit sequence Z becomes $(\rho_n)_{n\in\mathbb{N}}$ where the bit matrix $B = \rho_n \in M_{2^n}$ satisfies $b_{\sigma,\tau} = 1 \iff \sigma = \tau = Z \upharpoonright n$.
- If all the ρ_n are diagonal matrices, we describe a measure on Cantor space. The classical bit sequences are Dirac measures.
- A dynamics is given by the "shift" operator T which deletes the first qubit.

Notation: given a state ρ in $\mathcal{S}(M_\infty)$, write ρ_n for $\rho \upharpoonright M_{2^n}$.

Example of a state on M_∞

The EPR state $\beta = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \in M_{2^2}$ is a pure state which turns into the mixed state $\frac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|)$ after erasing the second qubit (taking the partial trace).

Example: letting

$$\rho_{2n} = \beta^{\otimes n}$$

and

$$\rho_{2n+1} = \rho_{2n} \otimes \tfrac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|),$$

we obtain a state such that the initial segments of even length are pure and the initial segments of odd length are mixed.

Defining quantum Martin-Löf randomness

Martin-Löf randomness

A central algorithmic randomness notion for infinite bit sequences Z is the one of Martin-Löf. The simplest of several equivalent definitions is based on "interval Solovay tests": a test is a computable list of dyadic intervals, of which Z eventually escapes.

Definition. Z is Martin-Löf random $\iff$

for every computable sequence $(\sigma_i)_{i\in\mathbb{N}}$ of binary strings with $\sum_i 2^{-|\sigma_i|} < \infty$, there are only finitely many i such that σ_i is an initial segment of Z.

ML-randoms are noncomputable, and satisfy law of large numbers.

Levin-Schnorr theorem: Z is ML-random iff for some fixed b, NO initial segment $Z \upharpoonright n$ can be compressed to less than $n - b$ bits.

Martin-Löf's test notion (bit sequences = reals)

- A Martin-Löf test is an effective sequence $(U_m)_{m\in\mathbb{N}}$ of open sets in $[0,1]$ such that the Lebesgue measure of U_m is at most 2^{-m}.
- Intuitively, U_m is an attempt to approximate a bit sequence (or real) Z with accuracy 2^{-m}.
- Z passes the test if Z is not in all U_m.
- Z is called Martin-Löf random if it passes all ML-tests.

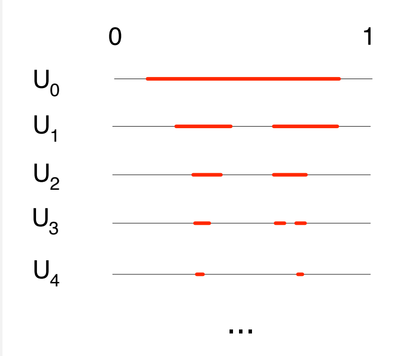

Projections and measurements of states

- A projection in M_{2^n} is a Hermitian matrix p such that $p^2 = p$.
- A special projection in M_{2^n} is a projection with matrix entries in $\mathbb{C}_{\mathtt{alg}}$, the field of algebraic complex numbers.

We can view $\rho(p_n)$ as a measurement of state $\rho \in \mathcal{S}(M_\infty)$ with the observable p_n.

$\rho(p_n)$ is the probability that ρ is "in" p_n.

In the classical case this is simply 1 (in) or 0 (out).

Σ^0_1 probabilistic sets on $S(M_\infty)$

Recall that a special projection in M_{2^n} is a projection with matrix entries in $\mathbb{C}_{\mathtt{alg}}$.
Consider projections $p \in M_{2^n}, q \in M_{2^k}$, $n \le k$. By $p \le q$ we mean that range of p is contained in range of q. (We view p as an element of M_{2^k} via the embedding $M_{2^n} \to M_{2^k}$.)
The trace of an $m \times m$ matrix is $\mathsf{tr}(A) = \sum_i A_{ii}$. Also $\tau(A) = \mathsf{tr}(A)/m$.

- A test component (or quantum Σ^0_1 set) G is given by a computable ascending sequence of special projections (p_n) where $p_n \in M_{2^n}$.
- Let $\tau(G) := \sup_n \tau(p_n) = \sup_n 2^{-n}\mathsf{tr}(p_n)$.
- For $\rho \in \mathcal{S}(M_\infty)$ let $G(\rho) = \sup_n \rho(p_n)$.

Quantum ML randomness (N. and Scholz, 2019)

Recall that we can view $\rho(p_n)$ as a measurement of state $\rho \in \mathcal{S}(M_\infty)$ with the observable p_n. The value $\rho(p_n) = \mathsf{tr}\ (\rho \restriction n)p_n$ is the probability that ρ is "in" p_n. In the classical case this is simply 1 (in) or 0 (out).

- A quantum Martin-Löf test is an effective sequence $\langle G_r \rangle_{r\in\mathbb{N}}$ of test components such that $\tau(G_r) \le 2^{-r}$ for each r.
- ρ passes the test if $\inf_r G_r(\rho) = 0$.
- ρ is quantum ML random if it passes each quantum ML test.

Think of the $\langle G_r \rangle_{r\in\mathbb{N}}$ as forming a sequence of measurements, with asymptotic value $\inf_r G_r(\rho)$ at ρ. If all the pieces are classical we re-obtain the usual definition of Martin-Löf tests.

Physical intuition

Given: One-way infinite spin chain

- We have access to the first few of them
- Can we make any predictions about expectation values of observables defined on bigger parts of the state?

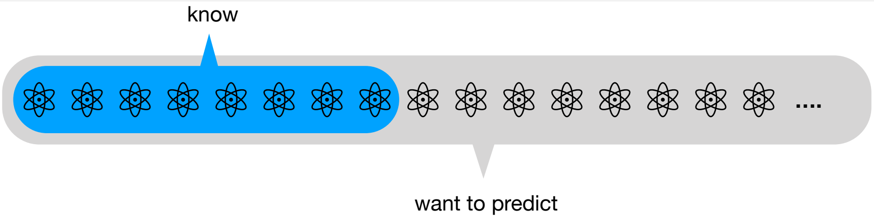

If state is quantum Martin-Löf-random, this is not possible.

Tracial state is random according to this definition

A probability measure on Cantor space can be seen as a state of the form $\rho = (\rho_n)_{n \in \mathbb{N}}$ where all the ρ_n are diagonal matrices.

- The tracial state τ corresponds to the usual Lebesgue measure. This is the case where each diagonal entry of τ_n is 2^{-n}.
- This state is random according to our definition, because $\tau(G_m) \to 0$ for each qML test (G_m).

This is compatible with our intuition because each τ_n is a fully mixed state, so has no structure.
Replacing the term "random" by "unstructured" might be more appropriate for infinite sequences of qubits; but then it doesn't agree with the accepted term for bit sequences.

Facts about quantum Martin-Löf randomness

One of Martin-Löf's results was the construction of a universal test for bit sequences.
Prop. There is a universal quantum ML-test.

If one wants to test classical bit sequences, the additional power of quantum ML-tests doesn't help.

Thm. Suppose $Z \in \{0,1\}^{\mathbb{N}}$. Then Z is ML-random $\iff$ Z viewed as an element of $\mathcal{S}(M_{2^\infty})$ is qML-random.

Special quantum Solovay tests

Interval Solovay tests were the ones used in our initial definition of ML-randomness for bit sequences. The quantum analog:

Definition. A special quantum Solovay test is a computable sequence $\langle p_r \rangle_{r \in \mathbb{N}}$ of special projections such that $\sum_r \tau(p_r) < \infty$.
We say that ρ is weakly quantum Solovay-random if it passes each special quantum Solovay test $\langle p_r \rangle_{r \in \mathbb{N}}$ in the sense that $\lim_r \rho(p_r) = 0$.

It's unknown whether there is a state that passes each special Solovay test but is not qML random. In fact this is unknown even when the state is a measure. For Dirac measures, i.e. bit sequences, both test notions have equal strength. They also have equal strength for states that are computable (Barmpalias and N., see Bhojraj thesis Thm 3.32)

Unitary machines and descriptive complexity

Definition (N. and Scholz, 2019)

- A unitary machine L is given by a computable sequence $\langle L_n \rangle$ of unitary (algebraic) matrices such that $L_n \in M_{2^n}$.
- For an input z which is a density matrix in M_{2^n}, its output is $L(z; n) := L_n z L_n^\dagger$.
- If z is a pure state $|\psi\rangle$, with the usual identifications the output is $L_n|\psi\rangle$.

$D(x, y)$ denotes trace distance of two density matrices x, y.
The L-quantum Kolmogorov complexity $QC_L^\epsilon(x \mid n)$ of a (possibly mixed) state x on n qubits is the least natural number k such that there exists a (mixed) state $y \in M_{2^k}$ with

$$D(x, L(y \otimes |0^{n-k}\rangle\langle 0^{n-k}|; n)) < \epsilon.$$

Weak quantum version of Levin-Schnorr theorem

Let ρ be a state on M_∞.

1. Let L be a unitary machine. Let $1 > \epsilon > 0$ and suppose ρ passes each qML-test at order $1 - \epsilon$. Then for each computable function f satisfying $\sum_n 2^{-f(n)} < \infty$, for almost every n
$$QC_L^\epsilon(\rho_n \mid n) \geq n - f(n).$$
2. For each special quantum Solovay test $\langle p_r \rangle_{r \in \mathbb{N}}$, there exists a total computable function $f\colon \mathbb{N} \to \mathbb{N}$ with $\sum_n 2^{-f(n)} \leq 4$ and a unitary machine L such that the following holds.
If ρ fails $\langle p_r \rangle$ at order $1 - \epsilon$ where $1 > \epsilon > 0$, then there are infinitely many n such that
$$QC_L^{\sqrt{\epsilon}}(\rho_n \mid n) < n - f(n).$$

An initial segment condition for randomness, based on entropy

- The von Neumann entropy of a density matrix $S \in M_n$ is $H(S) = -\mathrm{tr}(S \log_2 S)$.
- This is the usual entropy of the distribution that S induces on its eigenvectors.
- Its maximum value is n, when the distribution is uniform.

Theorem (Bhojraj, Thesis, 2021)

Let ρ be a state on M_∞. Suppose there is a constant $b \in \mathbb{N}$ such that $H(\rho_n) \geq n - b$ for infinitely many n.
Then ρ is quantum ML-random.

For instance, this applies to the tracial state τ; but Bhojraj constructs further examples.

Effective SMB theorem

- A 1950s theorem due to Shannon, McMillan and Breiman says that the entropy of an ergodic measure μ on $\{0,1\}^{\mathbb{N}}$ can be obtain as the empirical entropy along almost every trajectory Z.
- If μ is computable, it suffices that Z be ML-random relative to μ to determine the entropy of μ (Hochman, 2009).
- Bjelakovic et al. (2004) have a version of the Shannon-McMillan theorem for ergodic quantum lattice systems. but they use an ad hoc, finitary definition of "for almost", given that there is no notion of null set of states.
- we present a conjecture that attempts to remedy this, using that is a quantum analog of effective null sets of states.

Effective quantum SMB theorem?

- A state μ on M_∞ is called ergodic if it is an extreme point on the convex set of shift invariant states.
- $h(\mu) = \lim_n \frac{1}{n} H(\mu_n)$ is the von Neumann entropy of μ.
- Define qML-randomness relative μ as before, but with the condition $\mu(G_r) \le 2^{-r}$ when defining tests.

Conjecture (with Marco Tomamichel, Logic Blog 2017, 2020)

Suppose that for some $D > 0$, for each n, the diagonal entries of μ_n are bounded below by 2^{-nD}. Let ρ be a state that is quantum ML-random with respect to μ. Then

$$h(\mu) = -\lim \frac{1}{n} \mathsf{tr}(\rho_n \log_2 \mu_n).$$

This is known when μ is a diagonal state (measure). See Logic Blog 2020 Prop 9.3 (arxiv.org/abs/2101.09508).

Conclusions

- We have seen that notions from computability theory and algorithmic randomness interact meaningfully with the study of finite and infinite abstract spin chains.
- We've discussed work of Cubitt and others showing that the existence of a spectral gap is undecidable.
- We have defined infinite qubit sequences, and extended Martin-Löf randomness for classical bit sequences to the quantum setting.
- To find connections between the two approaches to spin chains, one would have to first formulate a version of the Cubitt et al. results for infinite qudit chains.
- Hamiltonians have been studied in this case, but they are usually not bounded, and only defined on a dense Hilbert subspace. See papers and books by Nachtergaele, Naaijkens, Sims, also Bjelakovic

J. Bausch, T. Cubitt, AngeA. Lucia, D. Perez-Garcia, and M. Wolf.
Size-driven quantum phase transitions.
arXiv preprint arXiv:1512.05687, 2015.

I. Bjelaković, T. Krüger, R. Siegmund-Schultze, and A. Szkoła.
The Shannon-McMillan theorem for ergodic quantum lattice systems.
Inventiones mathematicae, 155(1):203–222, 2004.

T. Bjojraj.
Algorithmic Randomness and Kolmogorov Complexity for Qubits.
Thesis, 2021
Arxiv 2106.14280.

A. Nies.
Computability and Randomness, volume 51 of *Oxford Logic Guides*.
Oxford University Press, Oxford, 2009.
444 pages. Paperback version 2011.

A. Nies and V. Scholz.
Martin-Löf random quantum states.
Journal of Mathematical Physics, 60(9):092201, 2019.

available at doi.org/10.1063/1.5094660.