**PHYS-EV0007 — Machine Learning from and for Quantum Science**

Project: Neural quantum state tomography.
From the transverse-field Ising chain to the ANNNI model

*Presentation:* G. Torlai et al., Nat. Phys. **14**, 447 (2018)
*Simulation:* NetKet Ground-State Ising tutorial

Anouar Moustaj
anouar.moustaj@aalto.fi

**Topic nr 2**
**Project nr 1**

## Project description

Neural quantum states (NQS) parametrise the amplitudes of a many-body wave function by a neural network, yielding a family of variational states whose expressive power grows smoothly with the number of hidden units. Beyond the variational ground-state problem they also offer a route to *state tomography*: given projective measurements of an unknown state in several bases, a generative neural network can be trained to reproduce the underlying probability distribution and, from it, physical observables that were not directly measured. This project is structured along two **independent** axes:

- **Presentation axis (scientific paper).** The student reads and presents the paper G. Torlai et al., Nat. Phys. **14**, 447 (2018): *Neural-network quantum state tomography*. The paper shows how a Restricted Boltzmann Machine (RBM) can be trained from measurements in several bases to reconstruct the wave function of interacting spin systems (1D and 2D TFIM, XXZ), and how observables that were never measured directly (including the entanglement entropy) can be recovered from the learned generative model.
- **Simulation axis (hands-on with NetKet).** The student reproduces the NetKet tutorial *Ground-State: Ising model*, which trains a sequence of neural-network ansätze (mean-field, short-range Jastrow, feed-forward, and a translation-symmetric architecture) on the 1D transverse-field Ising model (TFIM), and then *extends* the tutorial by adding a competing next-nearest-neighbour Ising coupling $J_2$ to obtain the *axial next-nearest-neighbour Ising* (ANNNI) chain. The $(h, \kappa)$ phase diagram of the ANNNI model, with $\kappa = J_2/J_1$, contains a ferromagnetic phase, a paramagnetic phase, and a period-4 "antiphase" (and, in a narrow window, a floating / incommensurate phase), giving a genuinely two-dimensional phase diagram to reconstruct with NQS.

## Learning outcomes

- Gain fluency with NQS for spin systems: how a neural network parametrises a quantum wave function and how VMC with stochastic reconfiguration optimises it.
- Understand the idea of *neural-network quantum state tomography*: training a generative neural network on measurements in multiple bases to reconstruct a full quantum state, and the role of the KL / likelihood objective.

- Learn how to use NetKet's spin primitives to define a TFIM Hamiltonian, train a neural-network ansatz, and measure expectation values of physical observables.
- Construct the $(h, \kappa)$ phase diagram of the ANNNI chain using NQS, identifying the ferromagnetic, paramagnetic and period-4 antiphase regions.

## Presentation

- Explain the formalism used in the paper: the complex RBM wave function $\psi_{\boldsymbol{\lambda},\boldsymbol{\mu}}(\sigma) = \sqrt{p_{\boldsymbol{\lambda}}(\sigma)}\, e^{i\phi_{\boldsymbol{\mu}}(\sigma)/2}$ built from two RBMs for amplitude and phase, and the conditional independence of the hidden units that makes sampling tractable.
- Understand the role of measurements in multiple bases $\{U_b\}$ and of the KL-divergence / log-likelihood objective in reconstructing both the magnitudes and the phases of the wave function.
- Explain the benchmarks reported in the paper (1D TFIM across its critical point, 1D XXZ at the isotropic point, 2D TFIM and 2D XXZ; reconstruction of two-point correlators and of the entanglement entropy).
- Explain the advantages and limitations of the method (scalability of the RBM ansatz, access to unmeasured observables; sensitivity to the sign / phase structure, volume-law entanglement, and dependence on the choice of measurement bases).

Useful complementary readings include G. Carleo and M. Troyer, *Solving the quantum many-body problem with artificial neural networks*, Science **355**, 602 (2017), for background on NQS in general, and the arXiv companion paper by the same group (arXiv:1703.05334).

## Simulations

The simulation axis is a hands-on exercise on the 1D transverse-field Ising chain, following the NetKet tutorial *Ground-State: Ising model* and extending it to the ANNNI chain. The ANNNI model is chosen because it is non-trivial (a period-4 ordered phase and, in a narrow window, a floating / incommensurate phase) and yet reduces to the exactly-solvable TFIM in the $\kappa \to 0$ limit, providing a clean benchmark for every NQS data point.

### 1. Model: TFIM and ANNNI chains

The tutorial targets the 1D transverse-field Ising chain with periodic boundary conditions, written in the same ferromagnetic sign convention as the tutorial,

$$H_{\mathrm{TFIM}} \;=\; -h \sum_{i=1}^{L} \sigma_i^x \;-\; J \sum_{i=1}^{L} \sigma_i^z \sigma_{i+1}^z, \tag{1}$$

with $\sigma_{L+1}^\alpha \equiv \sigma_1^\alpha$, $h \geq 0$ the transverse field and $J > 0$ the nearest-neighbour Ising coupling (the tutorial sets $\Gamma = V = -1$, i.e. $h = J = 1$). The model is analytically solvable (e.g. by Jordan–Wigner transformation to a Fermion problem) and hosts a quantum phase transition at $h/J = 1$. We fix $J = +1$ and treat $h$ as the control parameter; the natural order parameter is the uniform magnetization $\langle \sigma^z \rangle$.

For the *extension* we add a second-neighbor Ising coupling that competes with the nearest-neighbor one to obtain the ANNNI chain,

$$H_{\mathrm{ANNNI}}(h,\kappa) \;=\; -h \sum_{i=1}^{L} \sigma_i^x \;-\; J_1 \sum_{i=1}^{L} \sigma_i^z \sigma_{i+1}^z \;+\; \kappa J_1 \sum_{i=1}^{L} \sigma_i^z \sigma_{i+2}^z, \tag{2}$$

with $\kappa = J_2/J_1$ the frustration ratio (a positive $\kappa$ costs aligned second neighbours and frustrates the ferromagnet). The $(h, \kappa)$ phase diagram contains: a **ferromagnetic phase** for small $h$ and $\kappa < 1/2$; a **paramagnetic phase** for large $h$; a **period-4 $\langle 2 \rangle$ "antiphase"** ($\langle 2 \rangle$ denoting 2-by-2 spins $\uparrow\uparrow\downarrow\downarrow \cdots$) for small $h$ and $\kappa > 1/2$; and, between the antiphase and the paramagnet for $\kappa > 1/2$, a narrow **incommensurate phase** with a continuously varying wave vector. We scan the phase diagram in the window $(h, \kappa) \in [0, 2] \times [0, 1]$, which covers all four regions.

## 2. Reproducing the NetKet GS-Ising tutorial

As a first step, the student reproduces the NetKet tutorial *Ground-State: Ising model* on a 1D chain with periodic boundary conditions at the quantum-critical point $h/J = 1$, retracing the sequence of ansätze discussed there and benchmarking the converged variational energies against exact diagonalisation.

## 3. Extension: phase diagram of the ANNNI chain

Once the tutorial runs correctly at the TFIM critical point ($\kappa = 0$, $h = 1$), the student extends the Hamiltonian by adding the competing second-neighbour $\sigma^z\sigma^z$ term of Eq. (2) (built from the same `sigmaz` primitives) and scans $(h, \kappa)$ across the phase diagram.

**Scan.** Work on a 1D chain with periodic boundary conditions. The recommended chain sizes are $L \in \{16, 20, 24\}$ (reasonable sizes for a regular laptop; you may want to increase `n_samples` and `n_chains` as $L$ grows). Scan along two 1D cuts through the phase diagram:

- Cut A (fixed $\kappa = 0$, pure TFIM): $h \in \{0.0, 0.25, 0.5, 0.75, 1.0, 1.25, 1.5, 2.0\}$ (for example), reproducing the known QPT at $h_c = 1$;
- Cut B (fixed $\kappa = 0.75$, antiphase regime): same set of fields, exposing the antiphase $\rightarrow$ floating $\rightarrow$ paramagnet crossover;
- Cut C (fixed $h = 0.3$, scan in e.g. $\kappa$): $\kappa \in \{0.0, 0.2, 0.4, 0.5, 0.6, 0.8, 1.0\}$, crossing the FM $\rightarrow$ antiphase boundary near $\kappa \simeq 1/2$.

Three extra points can be added to refine the immediate vicinity of the TFIM critical point and of the multicritical point.

Finally, construct the *full 2D phase diagram* on a coarse grid covering $(h, \kappa) \in [0, 2] \times [0, 1]$ using, for instance,

$$q^*(h, \kappa) \;=\; \arg \max_{q \in \{0, \pi/2, \pi\}} S(q)/L.$$

This yields a categorical (FM / antiphase / paramagnet) colour map of the $(h, \kappa)$-plane on top of which the three 1D cuts can be overlaid as consistency checks.

**Compute plan.** Run cuts A, B, C with the *feed-forward* ansatz from the tutorial (a single hidden layer, a few tens of parameters, a few minutes per point on a laptop). Run a second, translation-symmetric ansatz (the `DenseSymm`-based architecture of the tutorial's Section 7) at four representative points $(h, \kappa) \in \{(0.5, 0), (1.0, 0), (0.3, 0.75), (0.3, 0.4)\}$, spanning the four regions of the phase diagram.

**Observables.** At each $(h, \kappa)$, measure the following:

- **Ground-state energy per site** $E_0/L = \langle H \rangle / L$, compared to exact Lanczos diagonalisation (the latter for reasonable system sizes) or exact solution (derive it or search online).

- **Energy variance** $\sigma_E^2 = \langle H^2 \rangle - \langle H \rangle^2$, as a quality-control diagnostic (vanishing for the true eigenstate).
- **Transverse magnetisation** $m_x = L^{-1} \sum_i \langle \sigma_i^x \rangle$, which approaches 1 deep in the paramagnetic phase.
- **Longitudinal structure factor**

$$S(q) \;=\; \frac{1}{L} \sum_{i,j} e^{iq(i-j)} \big[ \big\langle \sigma_i^z \sigma_j^z \big\rangle - \langle \sigma_i^z \rangle \big\langle \sigma_j^z \big\rangle \big],$$

  evaluated at $q = 0$ (FM / uniform order), $q = \pi$ (Néel order) and $q = \pi/2$ (period-4 antiphase order). The dominant peak identifies the phase.
- **Real-space correlator** $C^{zz}(r) = \langle \sigma_0^z \sigma_r^z \rangle - \langle \sigma_0^z \rangle \langle \sigma_r^z \rangle$ on distances $r = 1, \ldots, L/2$, and its long-distance envelope.

**Plots.** The student should produce:

1. $E_0/L$ vs $h$ along cuts A and B (with the exact Jordan–Wigner reference on cut A), and $E_0/L$ vs $\kappa$ along cut C.
2. $m_x$ and the three structure factors $S(0)/L$, $S(\pi)/L$, $S(\pi/2)/L$ vs the scan parameter, on the same panel for each cut, so that the dominant ordering wave vector in each region is visible.
3. $\log |C^{zz}(r)|$ vs $r$ at four representative points (one deep in each of the four regions: FM, paramagnetic, antiphase, and the intermediate floating / incommensurate window).
4. $\sigma_E^2$ along each cut, one curve per ansatz, to identify the regimes in which each NQS architecture remains reliable.
5. A 2D categorical phase-diagram map of $q^*(h, \kappa)$ on the coarse grid, with the three 1D cuts overlaid, and a continuous backup heatmap of $S(\pi/2)/L$ that highlights the period-4 antiphase region.

**Discuss the following**

- The location of the TFIM quantum-critical point on cut A and the softening / shift of the transition as $\kappa$ is turned on.
- The hallmark physical observations in each phase (dominant $S(q)$ peak, decay of $C^{zz}(r)$, behaviour of $m_x$).
- The relative performance of the two ansätze across the phase diagram: where does the plain feed-forward network already suffice? Where does the translation-symmetric architecture become essential? Tie these empirical observations back to the expressivity of neural-network wave functions discussed in the presentation paper.

## Deliverables

- A concise report on the article in the form of a presentation.
- A repo with organized, well-documented code, and a notebook with working examples for the simulations.

## Working practices and tools

*You are strongly encouraged to recycle existing material.* You are expected to **read, reuse, and adapt** existing reference implementations pipelines rather than re-derive or re-implement every algorithmic detail from scratch; the pedagogical goal of this project is to *understand and apply* the method, not to reproduce boilerplate. Any code you borrow must be clearly attributed in your repository (e.g. in a `README.md` or as in-source comments) and integrated cleanly with your own contributions. *Use of LLMs (ChatGPT, Claude, Gemini, ...) is permitted and encouraged for onboarding.* In particular, LLMs are very effective at:

- explaining unfamiliar programming syntax and idioms;
- summarising sections of the paper or cross-referencing related literature;
- debugging installation issues, bookkeeping, and programming recommendations;
- generating boilerplate (plotting scripts, parameter sweeps, unit tests, diagnostic helpers).

You remain, however, fully responsible for the correctness, clarity, and scientific content of what you submit: LLM output must be checked, understood, and, where appropriate, cited. Treat an LLM as a capable but occasionally wrong collaborator, not as an oracle. **Programming language:** Python

**Packages required:** `NetKet` (for NQS and VMC), `JAX` and `Flax` for automatic differentiation and neural-network layers (dependencies of NetKet). `NumPy`, `SciPy`, `matplotlib`

## References


- G. Torlai et al., Nat. Phys. **14**, 447 (2018).
- G. Carleo and M. Troyer, Science **355**, 602 (2017).
- F. Vicentini et al., SciPost Phys. Codebases **7** (2022).
- NetKet tutorial, Ground-State: Ising model.
- *ANNNI model*: W. Selke, Phys. Rep. **170**, 213 (1988).
- S. Sachdev, *Quantum Phase Transitions*, Cambridge University Press (2nd ed., 2011).