

Title: Lecture - Machine Learning, PHYS 777

Speakers: Mohamed Hibat Allah

Collection/Series: Machine Learning (Elective), PHYS 777, February 24 - March 28, 2025

Subject: Condensed Matter, Other

Date: March 28, 2025 - 9:00 AM

URL: <https://pirsa.org/25030044>

Lecture 13

Today:

↳ Neural Quantum States.

↳ Metropolis-Hastings sampling.

↳ Recurrent Neural Network Wavefunctions
(RNN)

Last time:

(Variational Monte Carlo (VMC))

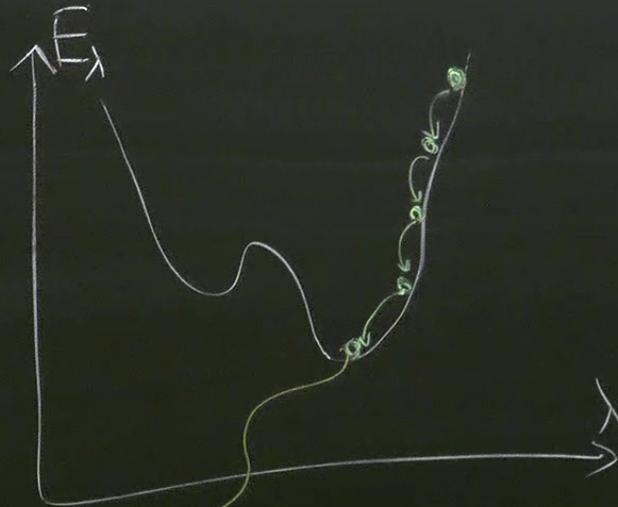
$$E_{\lambda} = \frac{\langle \psi_{\lambda} | \hat{H} | \psi_{\lambda} \rangle}{\langle \psi_{\lambda} | \psi_{\lambda} \rangle} \geq E_G$$

Annotations:
- ψ_{λ} is labeled "Ansatz"
- E_G is labeled "Ground State energy of \hat{H} "
- The entire fraction is labeled "Variational energy"
- $E_{loc}(\vec{\sigma})$ is labeled "Local energies"

$$E_{\lambda} = \sum_{\vec{\sigma}} \frac{|\psi_{\lambda}(\vec{\sigma})|^2}{\langle \psi_{\lambda} | \psi_{\lambda} \rangle} E_{loc}(\vec{\sigma})$$

Importance Sampling

$$E_{\lambda} \approx \frac{1}{M} \sum_{i=1}^M E_{loc}(\vec{\sigma}_i)$$



$$\lambda \leftarrow \lambda - \eta \frac{\partial E}{\partial \lambda}$$

(Gradient descent)

$$|\psi_{\lambda^*}\rangle \approx |\psi_G\rangle$$

$$E_{\lambda^*} \approx E_G$$

(RNN)

Metropolis-Hastings Sampling:

$$\vec{J} \sim \frac{|\psi_x(\vec{J})|^2}{\langle \psi_x | \psi_x \rangle}$$

amplitude $\rightarrow E_x \approx \frac{1}{M} \sum_{i=1}^M E_{x(i)}$ (bc)



Number of spins

Number of samples

$$A(\vec{\sigma}^{(i)} \rightarrow \vec{\sigma}^{(i+1)}) = \min \left(1, \frac{|\Psi_x(\vec{\sigma}^{(i+1)})|^2}{|\Psi_x(\vec{\sigma}^{(i)})|^2} \right)$$

Normalization is not important

Take samples $\int \vec{\sigma}^{(i \times N)} \prod_{j=1}^M$

Last time:

(Variational Monte Carlo (VMC))

$$E_\lambda = \frac{\langle \psi_\lambda | \hat{H} | \psi_\lambda \rangle}{\langle \psi_\lambda | \psi_\lambda \rangle} \geq E_G$$

Annotations:
- Ansatz ψ_λ
- Variational energy E_λ
- Ground State energy of \hat{H} E_G
- Local energies E_{loc}

$$E_\lambda = \sum_{\vec{\sigma}} \frac{|\psi_\lambda(\vec{\sigma})|^2}{\langle \psi_\lambda | \psi_\lambda \rangle} E_{loc}(\vec{\sigma})$$

Importance Sampling

$$E_\lambda \approx \frac{1}{M} \sum_{i=1}^M E_{loc}(\vec{\sigma}_i) = \sum_{\vec{\sigma}_i} H_{\vec{\sigma}_i} \frac{\psi_\lambda(\vec{\sigma}_i)}{\psi_\lambda(\vec{\sigma}_i)}$$

Number of spins

Firefox File Edit View History Bookmarks Tools Window Help

Hydrogen_Atom_Demo x Ground-State: Heisenberg mod: x gs-heisenberg.ipynb - Colab x +

localhost:8888/notebooks/Hydrogen_Atom_Demo.ipynb 150%

jupyter Hydrogen_Atom_Demo Last Checkpoint: 2 years ago

File Edit View Run Kernel Settings Help Trusted

JupyterLab Python 3 (ipykernel)

Based on Giuseppe Carleo Tutorials. Edited by Mohamed Hibat-Allah.

```
[4]: from optimization import *
```

Hamiltonian of the Hydrogen Atom

$$\hat{H} = -\frac{1}{2} \frac{\partial^2}{\partial r^2} + \frac{l(l+1)}{2r^2} - \frac{1}{r}$$

We consider $n = 0$ (principal quantum number), $l = 0$ (angular momentum) to get the ground state. Physical constants are absorbed in r .

- Ground state energy: $E_G = -\frac{1}{2}$
- Ground state: $\psi_G(r) = r \exp(-r)$



Firefox File Edit View History Bookmarks Tools Window Help

Hydrogen_Atom_Demo x Ground-State: Heisenberg mod: x gs-heisenberg.ipynb - Colab x +

localhost:8888/notebooks/Hydrogen_Atom_Demo.ipynb 150%

jupyter Hydrogen_Atom_Demo Last Checkpoint: 2 years ago

File Edit View Run Kernel Settings Help Trusted

JupyterLab Python 3 (ipykernel)

we consider $n = 0$ (principal quantum number), $l = 0$ (angular momentum) to get the ground state. Physical constants are absorbed in r .

- Ground state energy: $E_G = -\frac{1}{2}$
- Ground state: $\psi_G(r) = r \exp(-r)$

1. Define the variational wave function (ansatz)

$r=$ should not exist at nucleus

$\psi(r)$

r^2

$\exp(-\frac{1}{2}r)$

Ansatz should decay exponentially

R Typical radius of atom

r

Firefox File Edit View History Bookmarks Tools Window Help

Hydrogen_Atom_Demo x Ground-State: Heisenberg mod: X gs-heisenberg.ipynb - Colab x +

localhost:8888/notebooks/Hydrogen_Atom_Demo.ipynb 150% ☆

jupyter Hydrogen_Atom_Demo Last Checkpoint: 2 years ago

File Edit View Run Kernel Settings Help Trusted

JupyterLab Python 3 (ipykernel)

```
[5]: psi_ansatz=State([8.1, 6.2]) #Define the variational wavefunction
```

2. Running the optimization with gradient descent

$$\min_{\lambda} E_{\lambda} \rightsquigarrow \lambda \longrightarrow \lambda + \underbrace{(-\epsilon \partial_{\lambda} E_{\lambda})}_{\delta \lambda}$$

Firefox File Edit View History Bookmarks Tools Window Help

Hydrogen_Atom_Demo x Ground-State: Heisenberg mod: X gs-heisenberg.ipynb - Colab x +

localhost:8888/notebooks/Hydrogen_Atom_Demo.ipynb

Jupyter Hydrogen_Atom_Demo Last Checkpoint: 2 years ago

File Edit View Run Kernel Settings Help Trusted

JupyterLab Python 3 (ipykernel)

[9]: `run_optimization(psi_ansatz, epsilon=0.4)`

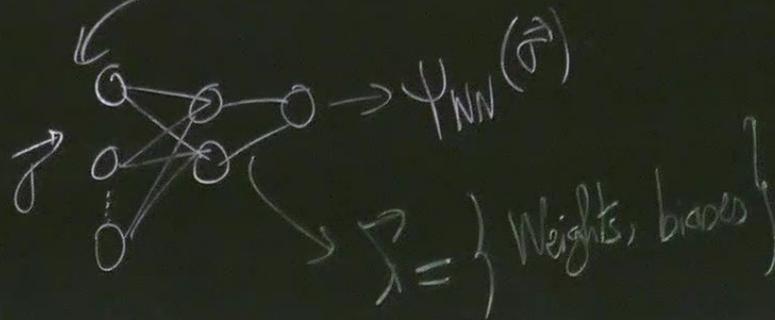
$\lambda_1 = 1.0, \lambda_2 = 1.0$

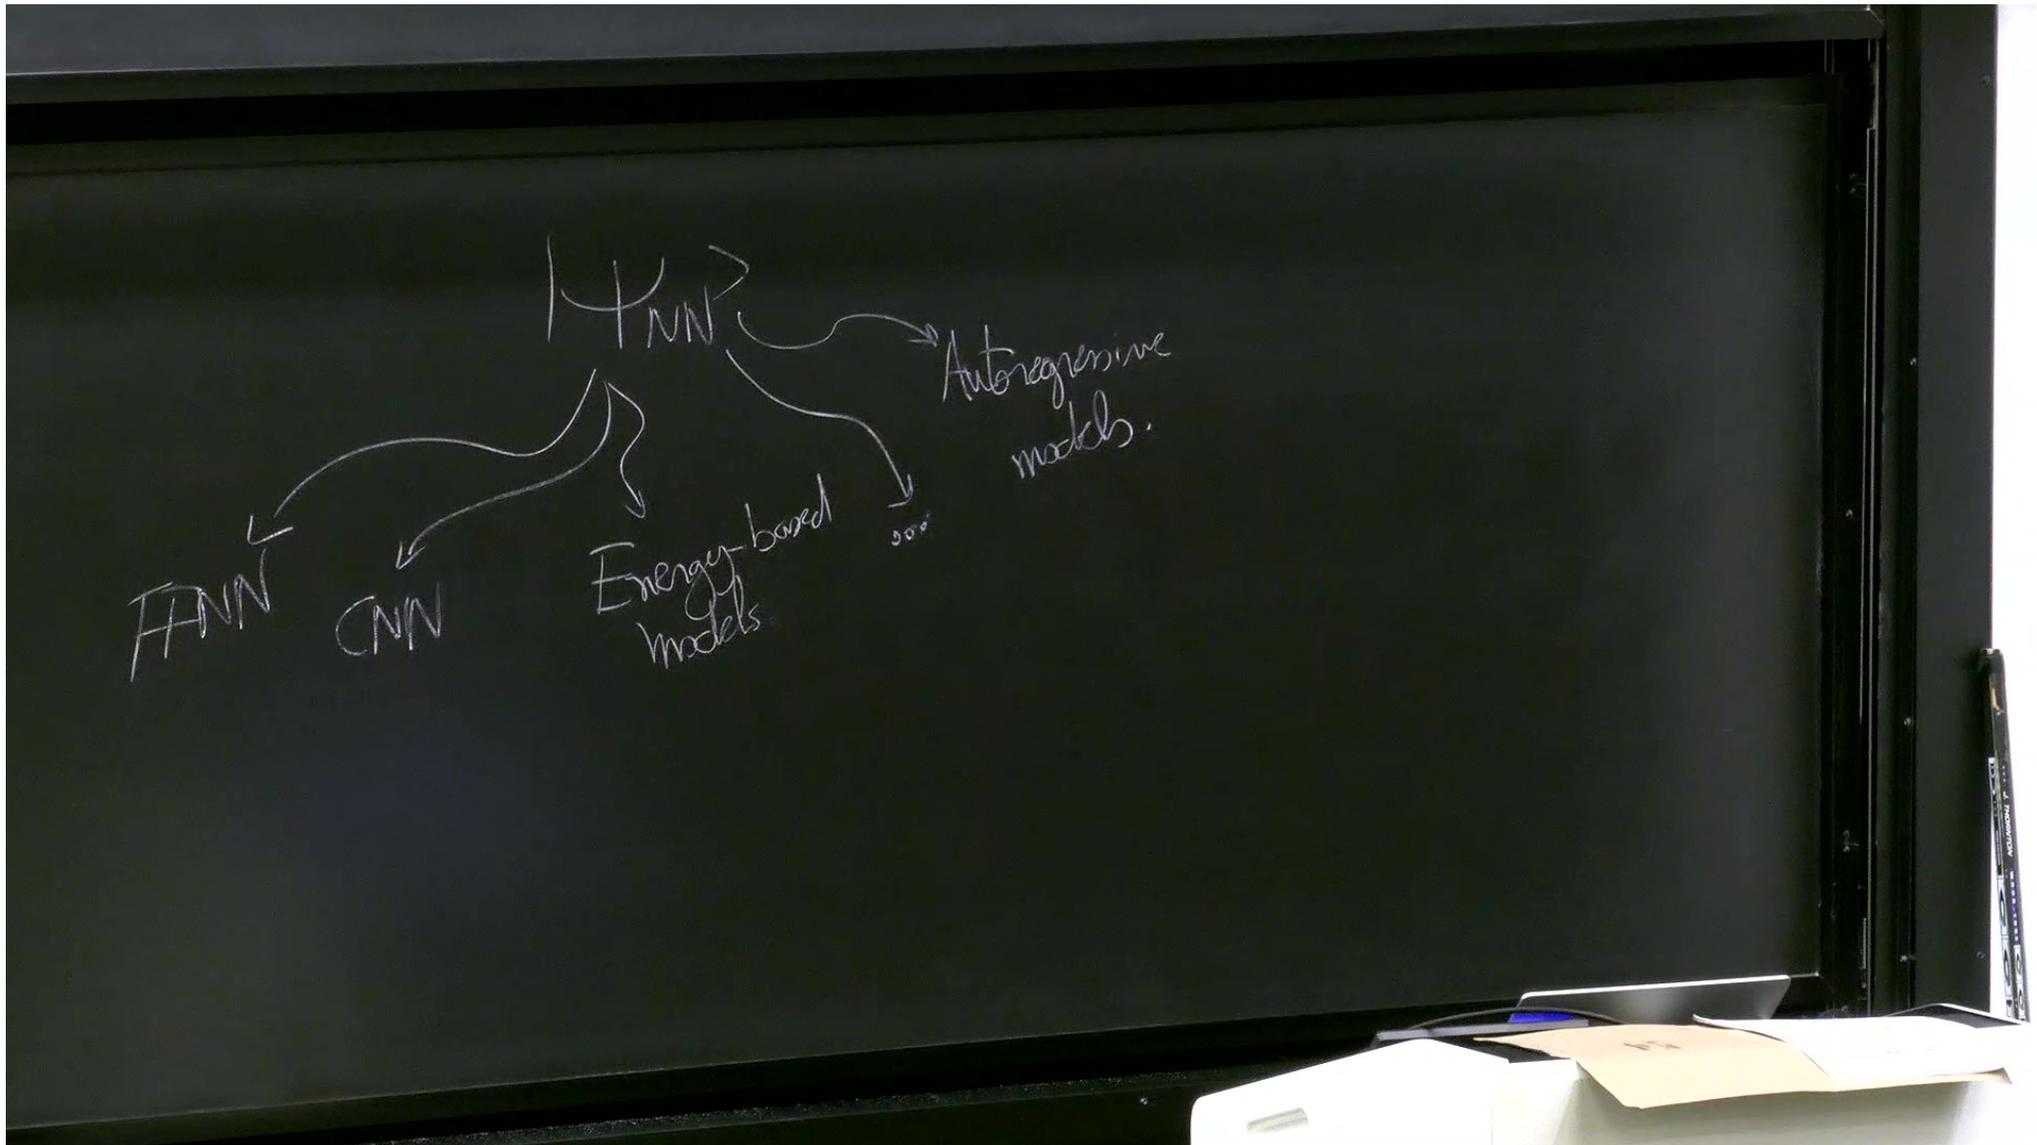
----- Exact energy = -0.5
-•- Variational energy

step 2000
lambda1 = 1.0000000000000004, lambda2 = 1.0000000000000004
energy = -0.5

Neural Quantum States

$$|\Psi_{NN}\rangle = \sum_{\vec{\sigma}} \underbrace{\Psi_{NN}(\vec{\sigma})}_{\text{NN}} |\vec{\sigma}\rangle$$





Variational principle is stat mechanics

$$F_p = \langle \hat{H} \rangle_p - T S_p$$

$$\langle \hat{H} \rangle_p = \sum_{\vec{\sigma}} P(\vec{\sigma}) E(\vec{\sigma})$$

$$S_p = \sum_{\vec{\sigma}} -P(\vec{\sigma}) \log(P(\vec{\sigma})) = \langle -\log(P(\vec{\sigma})) \rangle_p$$

→ Take samples $\{ \vec{\sigma}^{(i)} \}_{i=1}^M$

$\forall P, F_P \geq F_{P_T}$ Boltzmann prob. $\frac{1}{T}$

Use $P = P_\lambda$

$$P_T(\vec{\sigma}) = \frac{\exp(-\beta E(\vec{\sigma}))}{\sum_{\vec{\sigma}} \exp(-\beta E(\vec{\sigma}))}$$

$$F_\lambda = \langle H \rangle_{P_\lambda} + T \langle \log(P_\lambda(\vec{\sigma})) \rangle_{P_\lambda(\vec{\sigma})}$$

Variational free energy (Cost function).

$\approx P_T$ (near convergence)

all terms $(\sigma_1, \sigma_2, \dots, \sigma_N)$

Firefox File Edit View History Bookmarks Tools Window Help

Hydrogen_Atom_Demo x Ground-State: Heisenberg mod: x gs-heisenberg.ipynb - Colab x +

https://colab.research.google.com/drive/1t7388MOTclWiiSJqx7h4klxXJ1PBLEin#scrollTo=e7vgp7hzBMxE&uniqifier=1

gs-heisenberg.ipynb ☆ Share Gemini M

File Edit View Insert Runtime Tools Help

Commands + Code + Text Reconnect T4

Ground-State: Heisenberg model

Author: Giuseppe Carleo and Filippo Vicentini (EPFL-CQSL)

The goal of this tutorial is to review various neural network architectures available in NetKet, in order to learn the ground-state of a paradigmatic spin model: the spin-1/2 Heisenberg antiferromagnetic chain.

The Hamiltonian we will consider for this tutorial is the following

$$H = \sum_{i=1}^L \vec{\sigma}_i \cdot \vec{\sigma}_{i+1}$$

L is the length of the chain, and we will use both open and periodic boundary conditions. $\vec{\sigma} = (\sigma^x, \sigma^y, \sigma^z)$ denotes the vector of Pauli matrices. Please note that there is a factor of 2 between Pauli-matrices and spin-1/2 operators (thus a factor of 4 in H).

We will consider in this tutorial 5 possible ways of determining the ground-state of this model.

0. Installing Netket

If you are executing this notebook on Colab, you will need to install netket. You can do so by running the following cell:

```
[ ] 1 %pip install --quiet netket
```

✓ 1m 18s completed at 10:56 PM



Firefox File Edit View History Bookmarks Tools Window Help

Hydrogen_Atom_Demo x Ground-State: Heisenberg mod: x gs-heisenberg.ipynb - Colab x +

https://colab.research.google.com/drive/1t7388MOTclWiiSjQx7h4klxXJ1PBLEin#scrollTo=fghMO0McBMxP&uniqifier=1 133%

gs-heisenberg.ipynb ☆ Share Gemini M

File Edit View Insert Runtime Tools Help

Commands + Code + Text Reconnect T4

```
[ ] 2 # We impose to have a fixed total magnetization of zero
3 hi = nk.hilbert.Spin(s=0.5, total_sz=0, N=g.n_nodes)
```

The final element of the triptych is of course the Hamiltonian acting in this Hilbert space, which in our case is already defined in NetKet. Note that the NetKet Hamiltonian uses Pauli Matrices (if you prefer to work with spin-1/2 operators, it's pretty trivial to define your own custom Hamiltonian, as covered in another tutorial)

```
{x} 1 # calling the Heisenberg Hamiltonian
2 ha = nk.operator.Heisenberg(hilbert=hi, graph=g)
```

2. Exact Diagonalization (as a testbed)

Just as a matter of future comparison, let's compute the exact ground-state energy (since this is still possible for $L = 22$ using brute-force exact diagonalization). NetKet provides wrappers to the Lanczos algorithm which we now use:

```
[ ] 1 # compute the ground-state energy (here we only need the lowest energy, and do not need the eigenstate)
2 # evals = nk.exact.lanczos_ed(ha, compute_eigenvectors=False)
3 # exact_gs_energy = evals[0]
4 # print('The exact ground-state energy is E0=', exact_gs_energy)
5
6 # Just in case you can't do this calculation, here is the result
7 exact_gs_energy = -39.14752260706246
```

1m 18s completed at 10:56 PM

macOS dock with various application icons.

gs-heisenberg.ipynb

```
[ ] 3 # exact_gs_energy = evals[0]
4 # print('The exact ground-state energy is E0=', exact_gs_energy)
5
6 # Just in case you can't do this calculation, here is the result
7 exact_gs_energy = -39.14752260706246
```

3. The Jastrow ansatz

Let's start with a simple ansatz for the ground-state: the Jastrow Ansatz.

$$\log \psi(\sigma) = \sum_i a_i \sigma_i + \sum_{i,j} \sigma_i J_{i,j} \sigma_j$$

To show how it's done, we write this simple ansatz as a `flax.linen` module. We import this module and call it `nn`, as it is customary in all Flax documentation.

You should define a function in the module called `__call__` and decorated with `@nn.compact`. This function is responsible for defining the flow of your model, that is, to evaluate the module for a batch of inputs.

Parameters in the module are specified by calling `self.param(parameter_name, initializer, shape, dtype)`. The first argument is an arbitrary string, the second should be an initializer from `nn.initializers` and the other two are shape and dtype of that parameter.

As the module should work with batches of inputs (therefore the input will be a 2d matrix with shape `(N_inputs, N_sites)`), but we are lazy and find it easier to define the function for a single input `sigma` a 1D vector of shape `(N_sites)`. Therefore, we write a function called

✓ 1m 18s completed at 10:56 PM

Well that's not too bad for a simple ansatz. But we can do better, can't we?

3. Learning with a Restricted Boltzmann Machine (RBM)

We will now consider another celebrated ansatz, the Restricted Boltzmann Machine (RBM). It simply consists of two layers: a visible one representing the L spin 1/2 degrees of freedom, and an hidden one which contains a different number M of hidden units. There are connections between all visible and hidden nodes. The ansatz is the [following](#)

$$\Psi_{\text{RBM}}(\sigma_1^z, \sigma_2^z, \dots, \sigma_L^z) = \exp(\sum_{i=1}^L a_i \sigma_i^z) \prod_{i=1}^M \cosh(b_i + \sum_j W_{ij} \sigma_j^z)$$

a_i (resp. b_i) are the visible (resp. hidden) bias. Together with the weights W_{ij} , they are variational parameters that we (or rather NetKet) will optimize to minimize the energy. NetKet gives you the control on the important parameters in this ansatz, such as M and the fact that you want to use or not the biases. The full explanation is [here](#).

More conveniently (especially if you want to try another L in this tutorial), let's define the hidden unit density $\alpha = M/L$, and invoke the RBM ansatz in NetKet with as many hidden as visible units.

```
[ ] 1 # RBM ansatz with alpha=1
     2 ma = nk.models.RBM(alpha=1)
```

1m 18s completed at 10:56 PM

Firefox File Edit View History Bookmarks Tools Window Help

Hydrogen_Atom_Demo x Ground-State: Heisenberg mod: x gs-heisenberg.ipynb - Colab x +

https://colab.research.google.com/drive/1t7388MOTclWiiSjQx7h4klxXJ1PBLEin#scrollTo=z2nXmYsiBMxR&uniqifier=1 150%

gs-heisenberg.ipynb ☆ Share Gemini M

File Edit View Insert Runtime Tools Help

Commands + Code + Text Reconnect T4

```
4 # Extract the relevant information
5 iters_RBM = data["Energy"]["iters"]
6 energy_RBM = data["Energy"]["Mean"]
7
8 fig, ax1 = plt.subplots()
9 ax1.plot(iters_Jastrow, energy_Jastrow, color='C8', label='Energy (Jastrow)')
10 ax1.plot(iters_RBM, energy_RBM, color='red', label='Energy (RBM)')
11 ax1.set_ylabel('Energy')
12 ax1.set_xlabel('Iteration')
13 plt.axis([0, iters_RBM[-1], exact_gs_energy-0.03, exact_gs_energy+0.2])
14 plt.axhline(y=exact_gs_energy, xmin=0,
15             xmax=iters_RBM[-1], linewidth=2, color='k', label='Exact')
16 ax1.legend()
17 plt.show()
```

1m 18s completed at 10:56 PM

macOS dock with various application icons.

Firefox File Edit View History Bookmarks Tools Window Help

Hydrogen_Atom_Demo x Ground-State: Heisenberg mod: x gs-heisenberg.ipynb - Colab x +

https://colab.research.google.com/drive/1t7388MOTclWiISJqx7h4klxXJ1PBLEIn#scrollTo=z2nXmYsiBMxR&uniqifier=1 150%

gs-heisenberg.ipynb ☆

File Edit View Insert Runtime Tools Help

Commands + Code + Text Reconnect T4

```
15 xmax=iters_RBM[-1], linewidth=2, color='k', label='Exact')
16 ax1.legend()
17 plt.show()
```

Energy

Energy (Jastrow)
Energy (RBM)
Exact

1m 18s completed at 10:56 PM

macOS dock with various application icons.

Firefox File Edit View History Bookmarks Tools Window Help

Hydrogen_Atom_Demo x Ground-State: Heisenberg mod: x gs-heisenberg.ipynb - Colab x +

https://colab.research.google.com/drive/1t7388MOTclWiiSJqx7h4klxXJ1PBLEin#scrollTo=noBPSdTsBMxT&uniqifier=1 150%

gs-heisenberg.ipynb File Edit View Insert Runtime Tools Help

Commands + Code + Text Reconnect T4

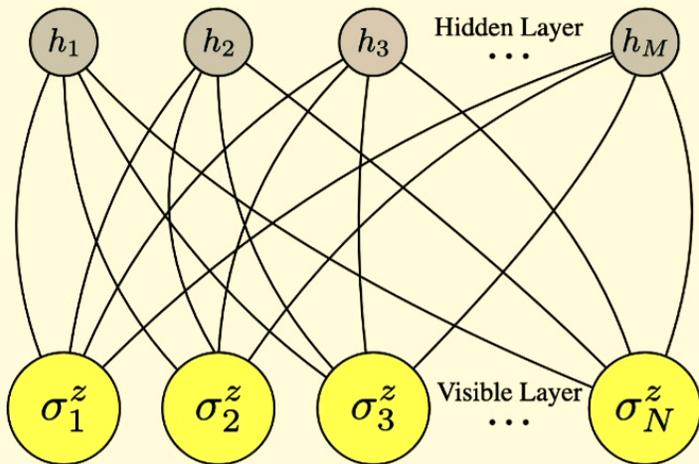
```
18 xmax=iters_RBM[-1], linewidth=2, color='k', label='Exact')
19 ax1.legend()
20 plt.show()
```

Energy

Energy (Jastrow)
Energy (RBM)
Energy (Symmetric RBM)
Exact

1m 18s completed at 10:56 PM

Neural Quantum States



Restricted Boltzmann Machine (RBM)

$$\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2 \dots \mathcal{S}_N)$$

$$\Psi_M(\mathcal{S}; \mathcal{W}) = \sum_{\{h_i\}} e^{\sum_j a_j \sigma_j^z + \sum_i b_i h_i + \sum_{ij} W_{ij} h_i \sigma_j^z}$$

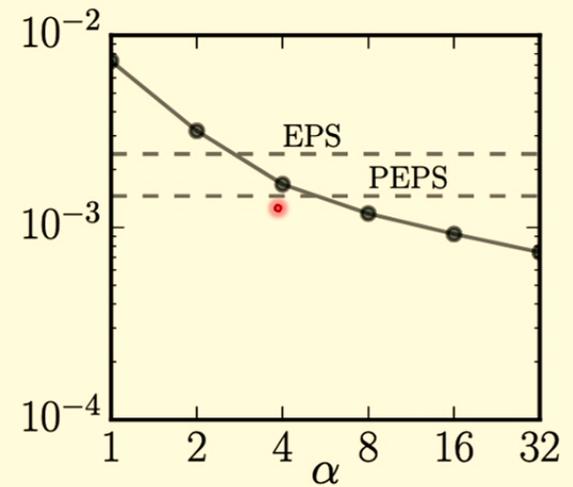
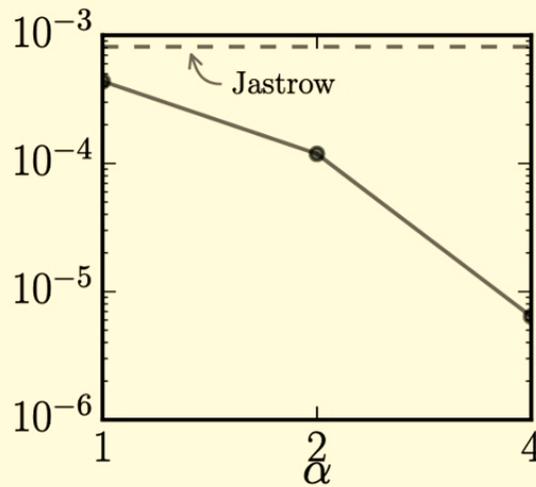
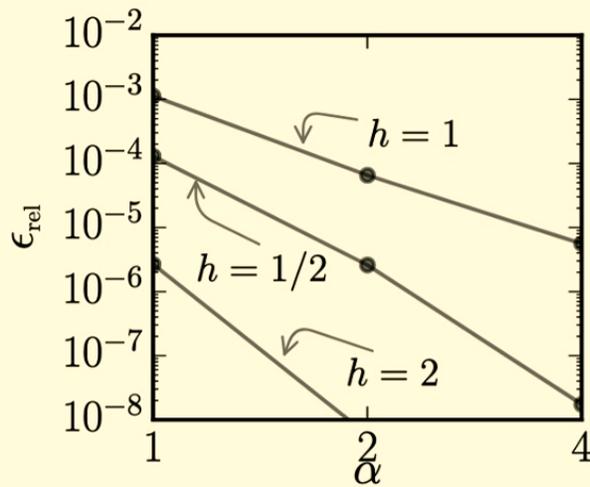
$$\Psi(\mathcal{S}; \mathcal{W}) = e^{\sum_i a_i \sigma_i^z} \times \prod_{i=1}^M F_i(\mathcal{S})$$

$$F_i(\mathcal{S}) = 2 \cosh \left[b_i + \sum_j W_{ij} \sigma_j^z \right]$$

Carleo, Troyer, Science, 2017.

Ground States

$$\alpha = M/N$$



$$\mathcal{H}_{\text{TFI}} = -h \sum_i \sigma_i^x - \sum_{\langle i,j \rangle} \sigma_i^z \sigma_j^z$$

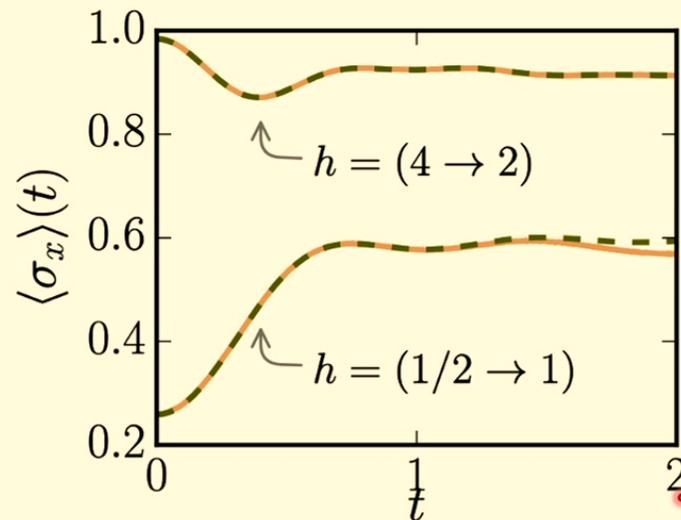
$$\mathcal{H}_{\text{AFH}} = \sum_{\langle i,j \rangle} \sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y + \sigma_i^z \sigma_j^z$$

Carleo, Troyer, Science, 2017.

Neural Quantum States

Time evolution

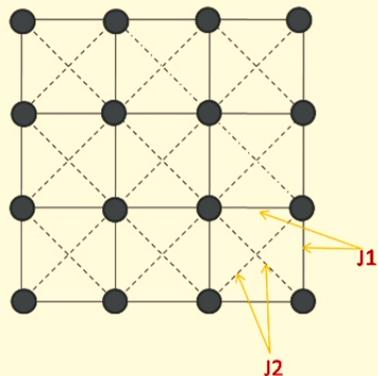
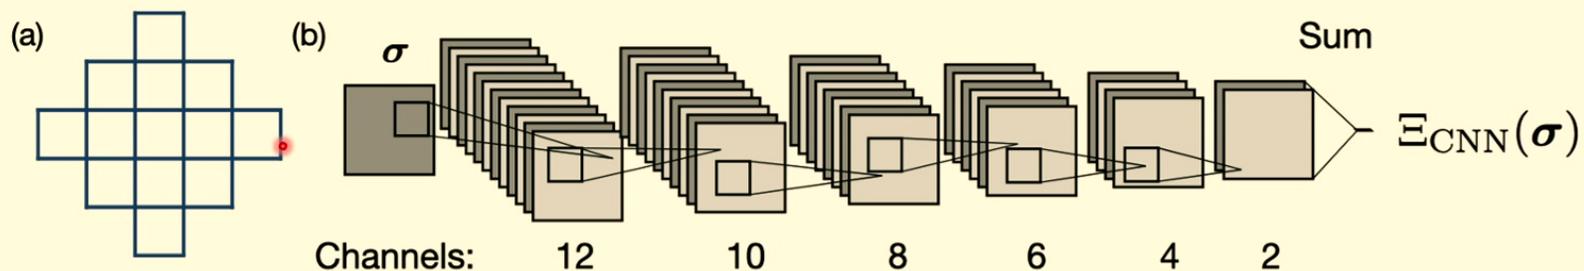
$$R(t; \dot{\mathcal{W}}(t)) = \text{dist}(\partial_t \Psi(\mathcal{W}(t)), -i\mathcal{H}\Psi)$$



$$\mathcal{H}_{\text{TFI}} = -h \sum_i \sigma_i^x - \sum_{\langle i,j \rangle} \sigma_i^z \sigma_j^z$$

Carleo, Troyer, Science, 2017.

Convolutional Neural Network Quantum States

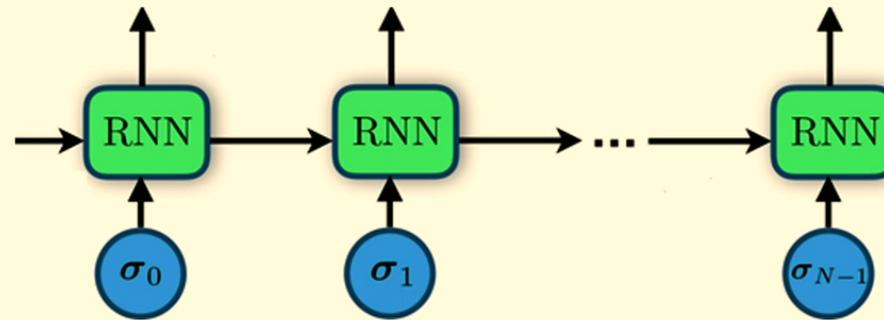


$$\hat{H} = J_1 \sum_{\langle ij \rangle} \hat{S}_i \cdot \hat{S}_j + J_2 \sum_{\langle\langle ij \rangle\rangle} \hat{S}_i \cdot \hat{S}_j$$

Choo, Neupert, Carleo, PRB, 2019.

Neural Quantum States

Recurrent Neural Networks (RNNs)



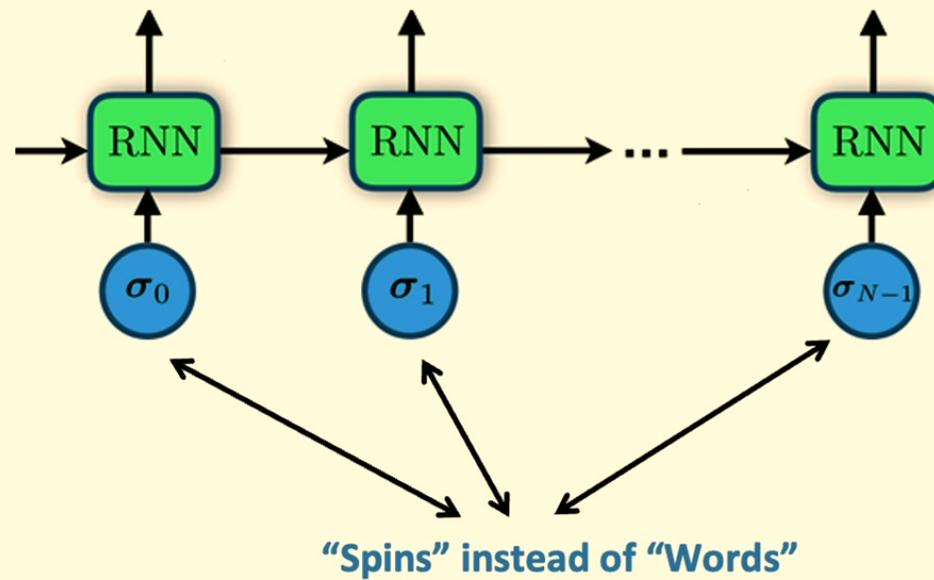
Very powerful at generating sequential data:

- Speech recognition, machine translation,...

An RNN is a Universal Turing Machine: can be seen as a general-purpose computer that can perform any classical computation.

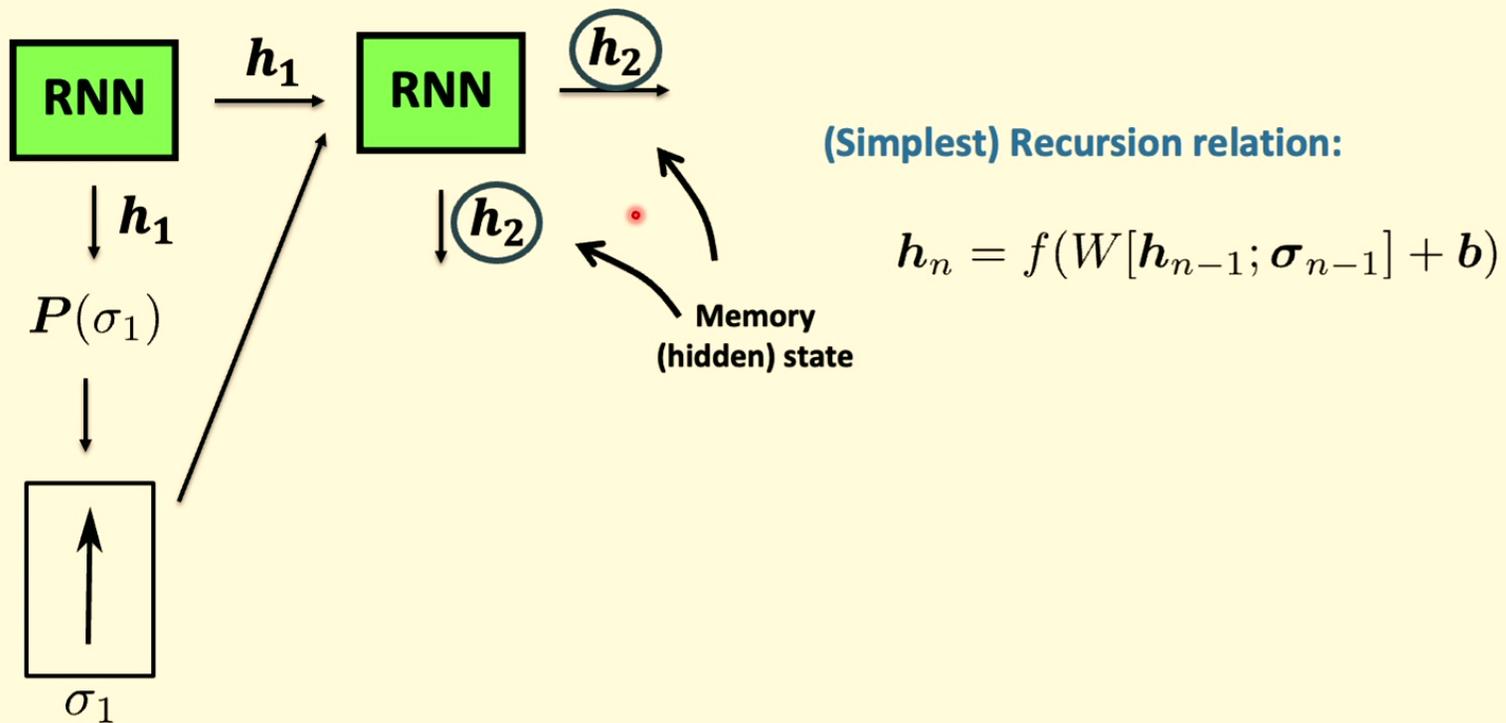
Siegelmann and Sontag, ACM, 1982
Chung and Siegelmann, NeurIPS 2021

RNNs can be also used in many-body Physics

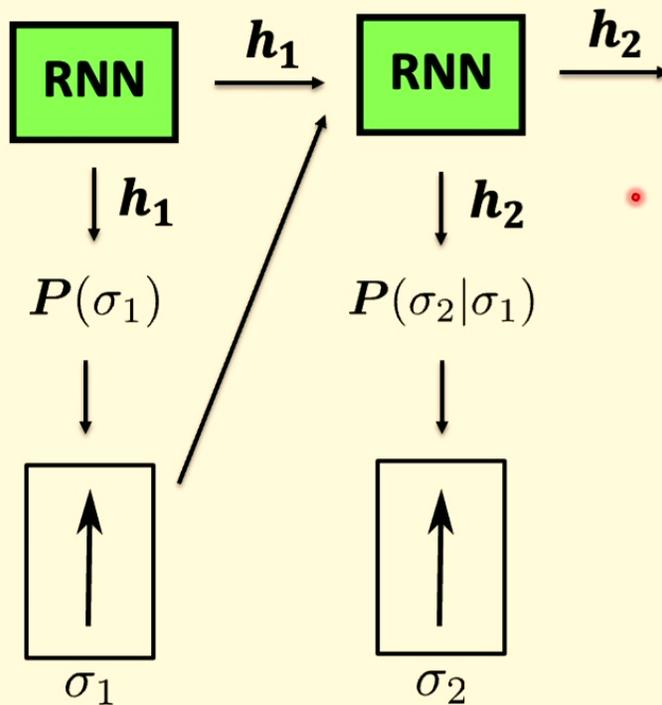


Neural Quantum States

Autoregressive sampling



Autoregressive sampling



(Simplest) Recursion relation:

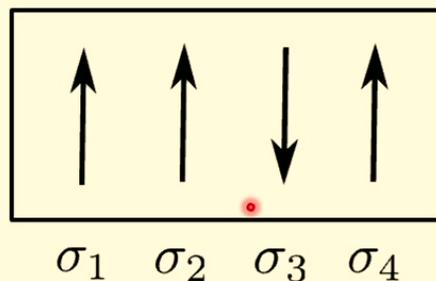
$$h_n = f(W[h_{n-1}; \sigma_{n-1}] + b)$$

Conditional probability:

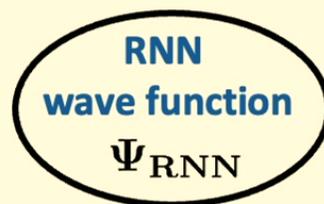
$$P(\sigma_i | \sigma_{<i}) = \text{Softmax}(Uh_n + c) \cdot \sigma_n$$

W, U, b and c are the parameters of the RNN.

RNN wave functions

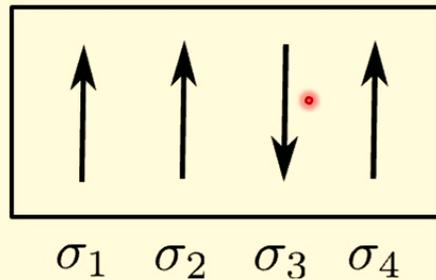


$$|\Psi_{\text{RNN}}(\sigma_1, \sigma_2, \sigma_3, \sigma_4)|^2 = P(\sigma_1)P(\sigma_2|\sigma_1)P(\sigma_3|\sigma_2, \sigma_1)P(\sigma_4|\sigma_3, \sigma_2, \sigma_1)$$



Neural Quantum States

Autoregressive sampling

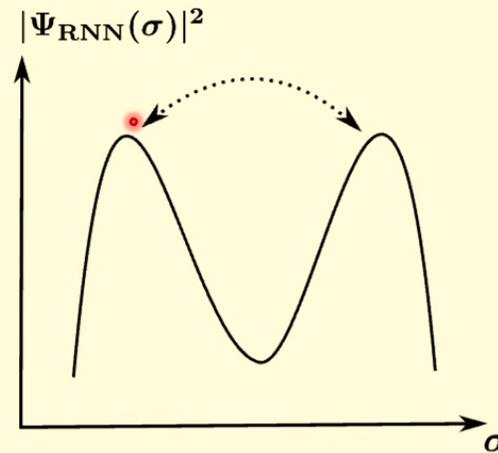


$$|\Psi_{\text{RNN}}(\sigma_1, \sigma_2, \sigma_3, \sigma_4)|^2 = P(\sigma_1)P(\sigma_2|\sigma_1)P(\sigma_3|\sigma_2, \sigma_1)P(\sigma_4|\sigma_3, \sigma_2, \sigma_1)$$

The samples are **independents** and can be generated in **parallel** $\{\sigma^{(j)}\}_{i=1}^{N_s}$

Importance sampling: $E_{\text{RNN}} \equiv \Psi_{\text{RNN}}^\dagger H \Psi_{\text{RNN}} \approx \frac{1}{N_s} \sum_{i=1}^{N_s} E_{\text{loc}}(\sigma^{(j)})$

Autoregressive sampling

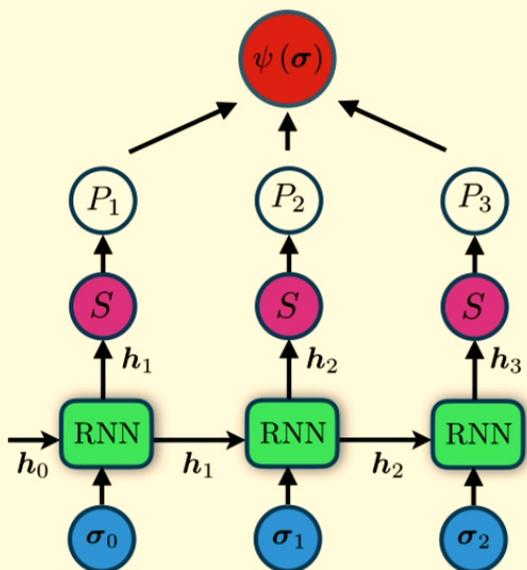


The **RNN** can sample from **multiple modes**

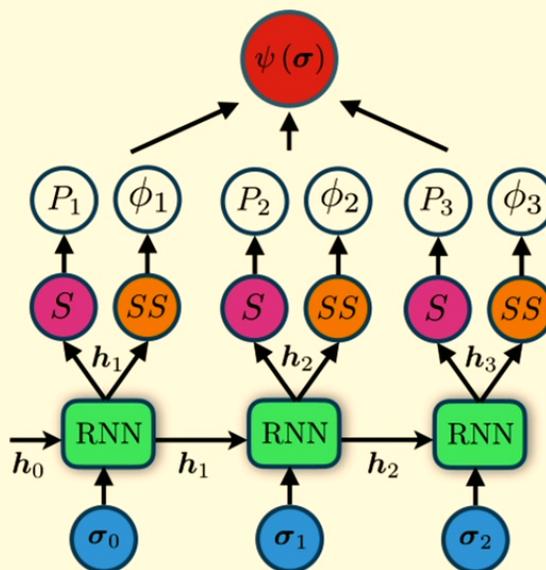
$$|\Psi_{\text{RNN}}(\sigma_1, \sigma_2, \sigma_3, \sigma_4)|^2 = P(\sigma_1)P(\sigma_2|\sigma_1)P(\sigma_3|\sigma_2, \sigma_1)P(\sigma_4|\sigma_3, \sigma_2, \sigma_1)$$

RNN wave functions

Positive RNN wave function



Complex RNN wave function



$$P(\boldsymbol{\sigma}) \equiv \prod_{n=1}^N P_n$$

$$\phi(\boldsymbol{\sigma}) \equiv \sum_{n=1}^N \phi_n$$

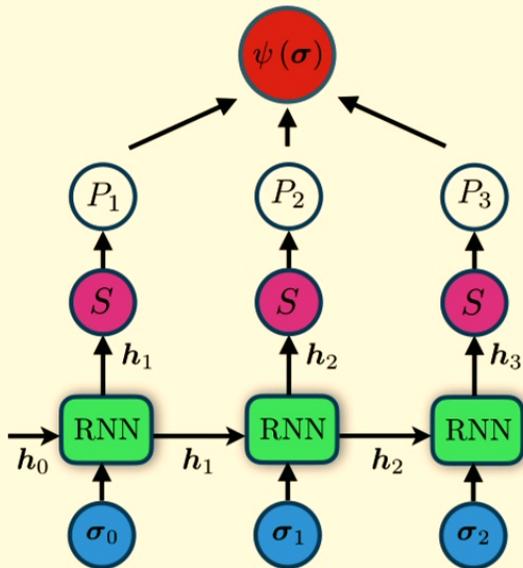
$$\Psi_{\text{RNN}}(\boldsymbol{\sigma}) = \sqrt{P(\boldsymbol{\sigma})} \exp(i\phi(\boldsymbol{\sigma}))$$

M.H., M. Ganahl, L. Hayward, R. Melko, J. Carrasquilla, PRRResearch, 2020.

Neural Quantum States

RNN wave functions

Positive RNN wave function



$$P(\sigma) \equiv \prod_{n=1}^N P_n$$

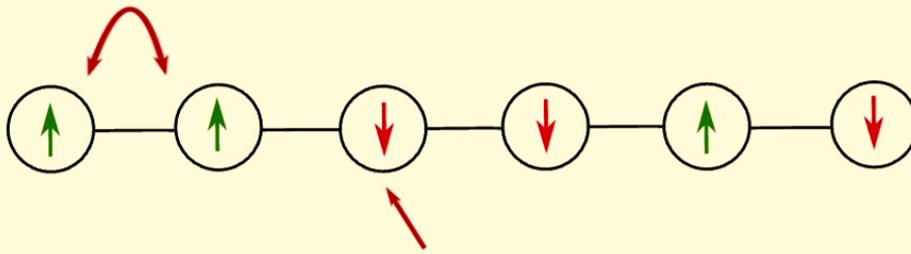
$$\Psi_{\text{RNN}}(\sigma) = \sqrt{P(\sigma)}$$

M.H., M. Ganahl, L. Hayward, R. Melko, J. Carrasquilla, PRResearch, 2020.

Neural Quantum States

Transverse-field Ising model in 1D

Ferromagnetic coupling (order)



Spin flip coupling (disorder)

Hamiltonian:

$$\hat{H} = - \sum_{i=1}^{N-1} \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z - h \sum_{i=1}^N \hat{\sigma}_i^x$$

Order Disorder

Pauli matrices:

$$\hat{\sigma}^z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad \hat{\sigma}^x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

Transverse-field Ising model in 1D

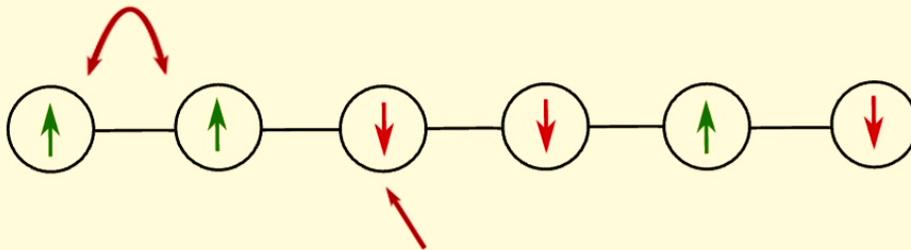
Hamiltonian:

$$\hat{H} = - \sum_{i=1}^{N-1} \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z - h \sum_{i=1}^N \hat{\sigma}_i^x$$

Order

Disorder

Ferromagnetic coupling (order)

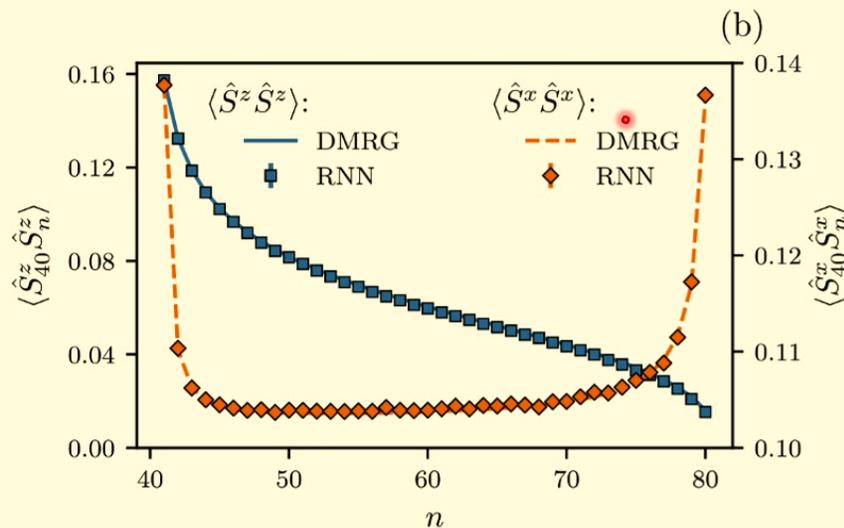


Spin flip coupling (disorder)

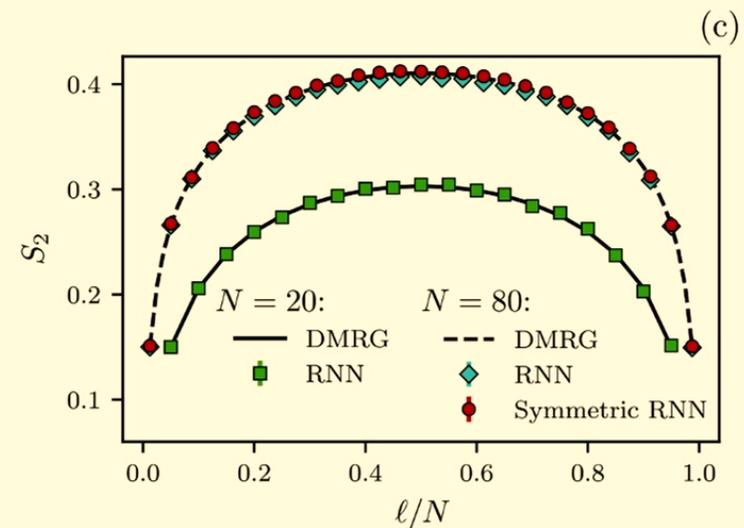
$$\mathbf{H} = \begin{pmatrix} * & \cdots & * \\ \vdots & \ddots & \vdots \\ * & \cdots & * \end{pmatrix} \begin{matrix} \leftarrow 2^N \rightarrow \\ \updownarrow 2^N \end{matrix}$$

Neural Quantum States

RNNs can estimate **observables** and **entanglement**



Correlations functions

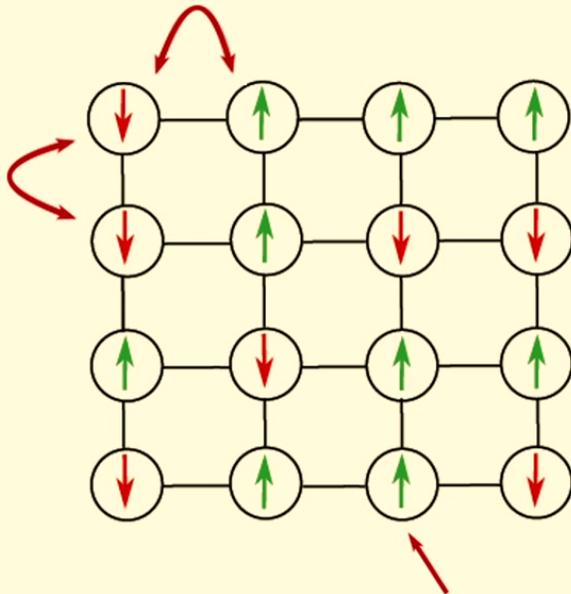


Entanglement entropies

M.H., M. Ganahl, L. Hayward, R. Melko, J. Carrasquilla, PRRresearch, 2020.

Transverse-field Ising model in 2D

Ferromagnetic coupling (order)

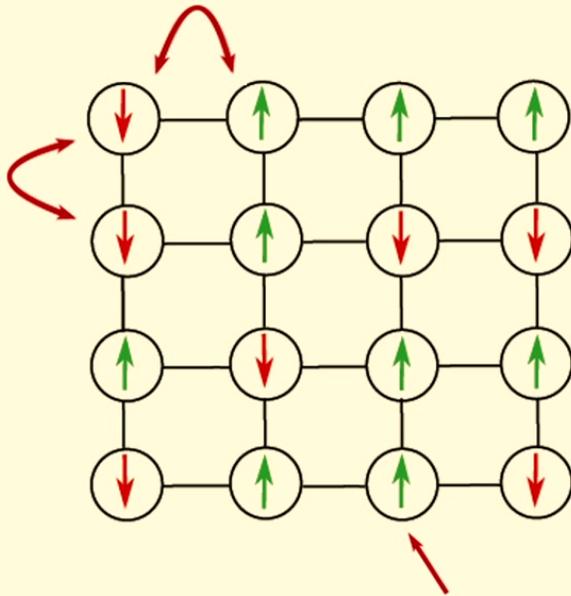


Spin flip coupling (disorder)

- In 2D, the exact ground state energy is **not known** theoretically.

Transverse-field Ising model in 2D

Ferromagnetic coupling (order)



Spin flip coupling (disorder)

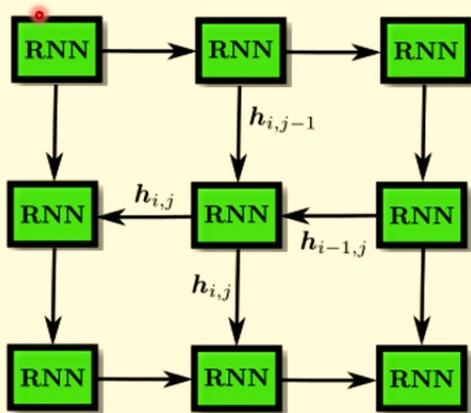
- Hamiltonian:

$$\hat{H} = - \sum_{\langle i,j \rangle} \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z - h \sum_{i=1}^N \hat{\sigma}_i^x$$

↑
↑
 Order Disorder

$$\mathbf{H} = \begin{pmatrix} * & \dots & * \\ \vdots & \ddots & \vdots \\ * & \dots & * \end{pmatrix} \begin{matrix} \leftarrow 2^N \\ \rightarrow 2^N \\ \uparrow 2^N \\ \downarrow 2^N \end{matrix}$$

2D recurrent neural networks



~~1D recursion relation:~~

~~$$\mathbf{h}_n = f(W[\mathbf{h}_{n-1}; \sigma_{n-1}] + \mathbf{b})$$~~

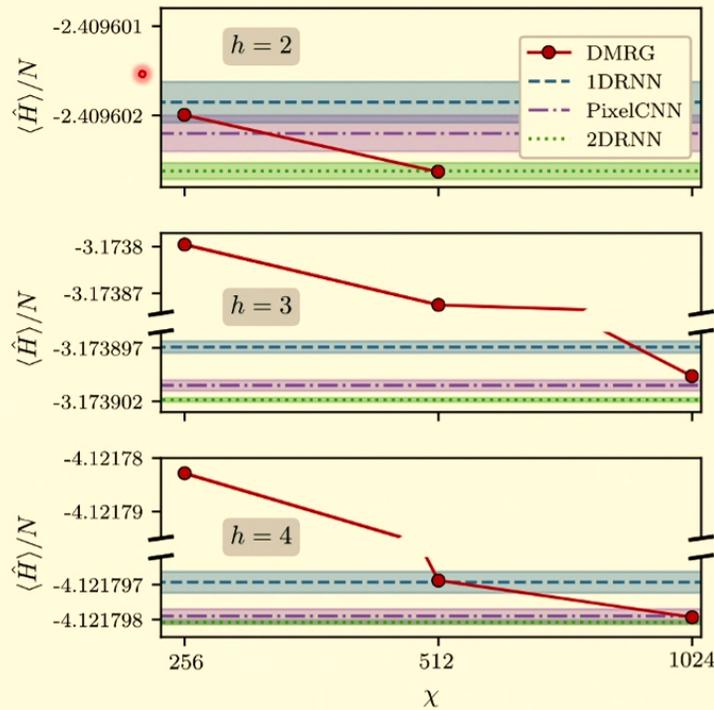
2D recursion relation:

$$h_{i,j} = f\left(W^{(h)}[h_{i-1,j}; \sigma_{i-1,j}] + W^{(v)}[h_{i,j-1}; \sigma_{i,j-1}] + \mathbf{b}\right)$$

M.H., M. Ganahl, L. Hayward, R. Melko, J. Carrasquilla, RNN Wave functions, PRResearch, 2020.

A. Grave et al., Multi-dimensional recurrent neural networks, 2007.

Comparison with the state-of-the art in 2D



- Hamiltonian (N = 12x12):

$$\hat{H} = - \sum_{\langle i,j \rangle} \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z - h \sum_{i=1}^N \hat{\sigma}_i^x$$

DMRG: Density Matrix Renormalization Group.

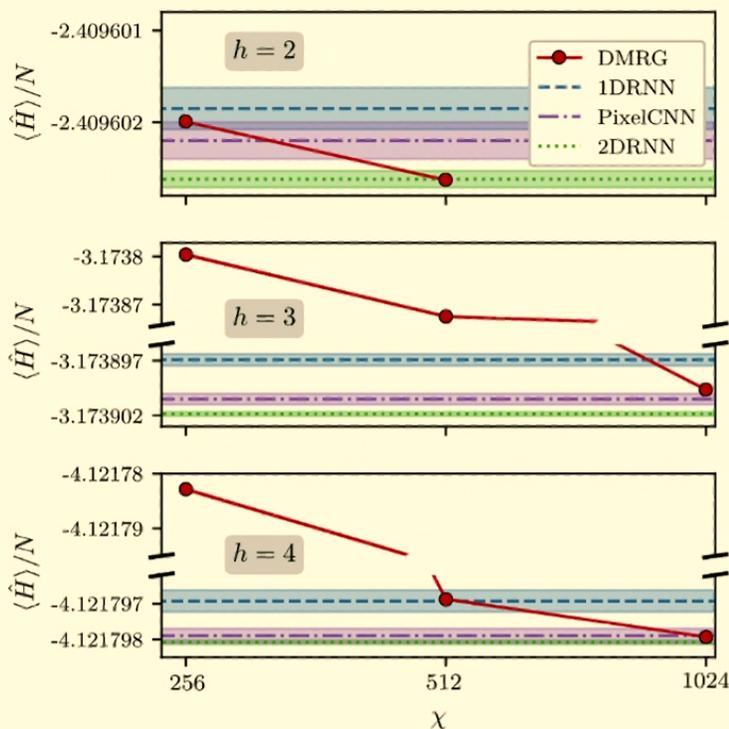
2DRNN and **PixelCNN:** Neural Network architectures.

- The lower the energy the better!

2DRNN: [M.H., M. Ganahl, L. Hayward, R. Melko, J. Carrasquilla, PRResearch, 2020.](#)

Pixel CNN: [Sharir et al. 2020, PRL, 2020.](#)

Comparison with the state-of-the-art in 2D



- Hamiltonian (N = 12x12):

$$\hat{H} = - \sum_{\langle i,j \rangle} \hat{\sigma}_i^z \hat{\sigma}_{i+1}^z - h \sum_{i=1}^N \hat{\sigma}_i^x$$

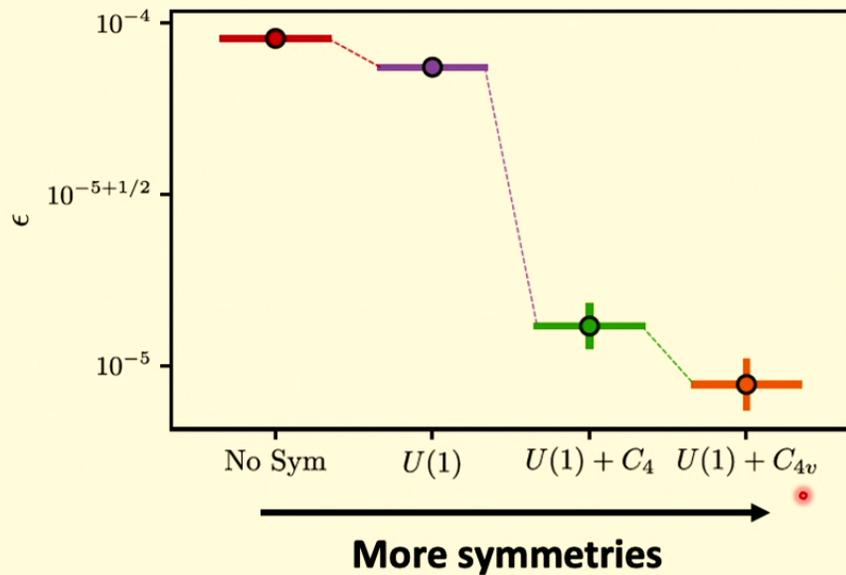
- RNNs are more compact!

0.03% of the variational parameters of **DMRG**.

2.6% of the variational parameters of **PixelCNN**.

2DRNN: [M.H., M. Ganahl, L. Hayward, R. Melko, J. Carrasquilla, PRResearch, 2020](#)
 Pixel CNN: [Sharir et al. 2020, PRL, 2020.](#)

Embedding physical symmetries into 2D RNNs

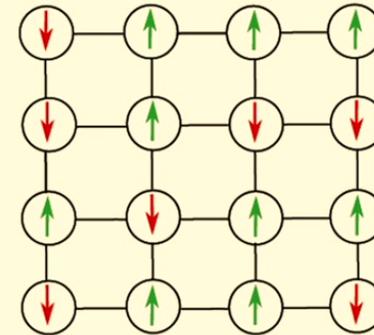


- Relative error :

$$\epsilon = |E_{\text{RNN}} - E_{\text{Exact}}| / |E_{\text{Exact}}|$$

2D Heisenberg model on square lattice

$$\hat{H} = \frac{1}{4} \sum_{\langle i,j \rangle} (\hat{\sigma}_i^x \hat{\sigma}_j^x + \hat{\sigma}_i^y \hat{\sigma}_j^y + \hat{\sigma}_i^z \hat{\sigma}_j^z)$$



M.H, R. Melko, J. Carrasquilla, ML for Physical Sciences workshop, NeurIPS 2021.

Neural Quantum States

What is next?

<https://doi.org/10.1038/s42005-024-01584-y>

Variational Monte Carlo with large patched transformers

 Check for updates

Kyle Sprague & Stefanie Czisczek  

Large language models, like transformers, have recently demonstrated immense powers in text and image generation. This success is driven by the ability to capture long-range correlations between elements in a sequence. The same feature makes the transformer a powerful wavefunction ansatz that addresses the challenge of describing correlations in simulations of qubit systems. Here we consider two-dimensional Rydberg atom arrays to demonstrate that transformers reach higher accuracies than conventional recurrent neural networks for variational ground state searches. We further introduce large, patched transformer models, which consider a sequence of large atom patches, and show that this architecture significantly accelerates the simulations. The proposed architectures reconstruct ground states with accuracies beyond state-of-the-art quantum Monte Carlo methods, allowing for the study of large Rydberg systems in different phases of matter and at phase transitions. Our high-accuracy ground state representations at reasonable computational costs promise new insights into general large-scale quantum many-body systems.

Attention Is All You Need

Ashish Vaswani*
Google Brain
avaswani@google.com

Noam Shazeer*
Google Brain
noam@google.com

Niki Parmar*
Google Research
nikip@google.com

Jakob Uszkoreit*
Google Research
usz@google.com

Llion Jones*
Google Research
llion@google.com

Aidan N. Gomez* †
University of Toronto
aidan@cs.toronto.edu

Lukasz Kaiser*
Google Brain
lukaszkaizer@google.com

Illia Polosukhin* †
illia.polosukhin@gmail.com

What is next?

RydbergGPT

David Fitzek,^{1,2} Yi Hong Teoh,³ Hin Pok Fung,³ Gebremedhin A. Dagnew,³ Ejaaz Merali,³ M. Schuyler Moss,³ Benjamin MacLellan,³ and Roger G. Melko^{3,4}

¹Department of Microtechnology and Nanoscience, Chalmers University of Technology, 412 96 Gothenburg, Sweden

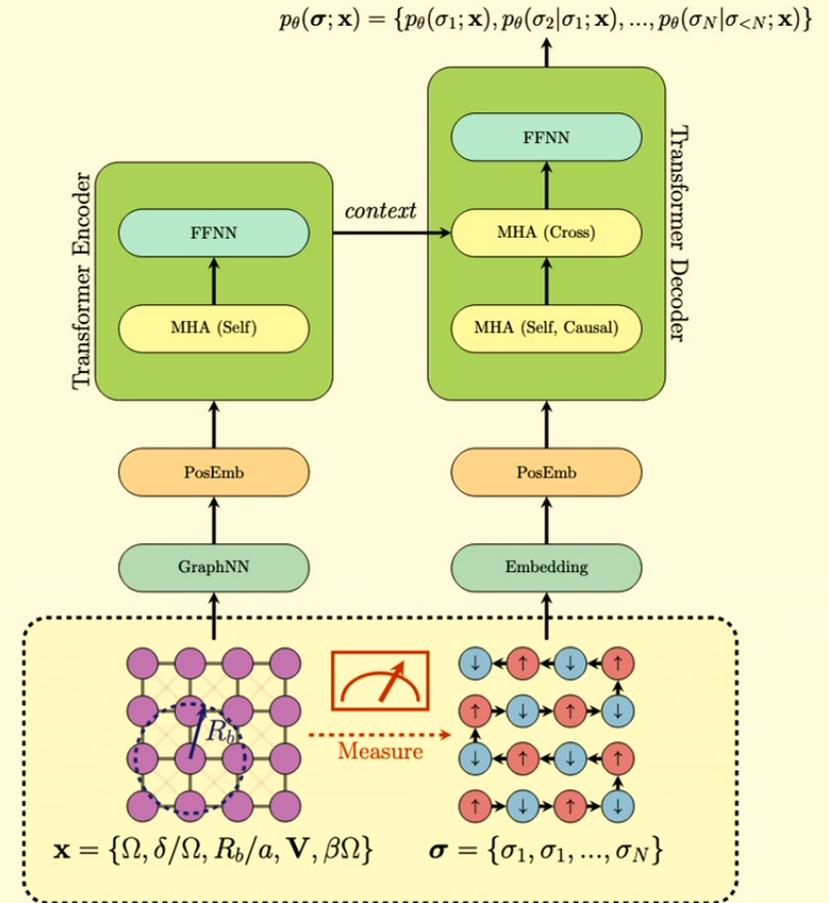
²Volvo Group Trucks Technology, 405 08 Gothenburg, Sweden

³Department of Physics and Astronomy, University of Waterloo, 200 University Ave. West, Waterloo, Ontario N2L 3G1, Canada

⁴Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L 2Y5, Canada

(Dated: June 3, 2024)

We introduce a generative pretrained transformer (GPT) designed to learn the measurement outcomes of a neutral atom array quantum computer. Based on a vanilla transformer, our encoder-decoder architecture takes as input the interacting Hamiltonian, and outputs an autoregressive sequence of qubit measurement probabilities. Its performance is studied in the vicinity of a quantum phase transition in Rydberg atoms in a square lattice array. We explore the ability of the architecture to generalize, by producing groundstate measurements for Hamiltonian parameters not seen in the training set. We focus on examples of physical observables obtained from inference on three different models, trained in fixed compute time on a single NVIDIA A100 GPU. These can act as benchmarks for the scaling of larger RydbergGPT models in the future. Finally, we provide RydbergGPT open source, to aid in the development of foundation models based off of a wide variety of quantum computer interactions and data sets in the future.



Thank you!

1. **Natural Language Processing** meets **many-body physics**.
2. **RNN wave functions** are **competitive** with **the state-of-the-art**.
3. The future of **< Machine Learning | many-body physics >** is **promising!**