

---

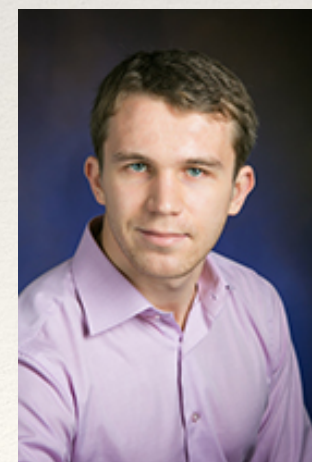
# Variational Wavefunctions in the era of AI

*Virtually in the ether*

---

Bryan Clark ([bkclark@illinois.edu](mailto:bkclark@illinois.edu))

*University of Illinois at Urbana Champaign*



Dmitrii Kochkov



Di Luo



# The Variational Approach

*The variational approach is one of the key approaches to solving the quantum many-body problem.*

Given a Hamiltonian  $H$ , find a good approximation  $\Psi$  to the ground state...

$$R \implies \Psi \implies \#$$

Choose a set of wave-functions you're going to consider...

$$\{\Psi_1, \Psi_2, \Psi_3, \dots, \Psi_N\}$$
$$\{\Psi(\vec{\theta})\}$$

Pick the best one from that list.



# Outline

Choose a set of wave-functions you're going to consider...

$$\begin{array}{l} \{\Psi_1, \Psi_2, \Psi_3, \dots, \Psi_N\} \\ \{\Psi(\vec{\theta})\} \end{array}$$



## Part 1

Wavefunctions

*Neural Network Backflow*

Pick the best one from that list.



## Part 2

Optimization

How is the age of AI changing these two steps?



# Wave-Functions through History



The age of (dressed) mean field.

Slater Determinants (i.e. Hartree Fock)



The age of tensor networks



The age of AI



# The Age of Mean Field



Slater Determinants

BDG States

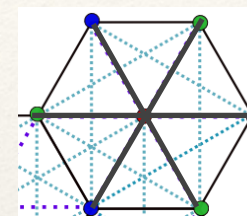
Pfaffians

Jastrow



$$\Psi = \det M \exp[-U(R)]$$

$$\Psi = P[\det M] \text{ (spin liquids)}$$



J1-J2 triangular spin liquid:  
93% overlap with projected SD

$$\Psi(r_1, r_2, \dots, r_n) = \det M;$$

$$M_{ij} = \phi_i(r_j)$$

Single particle orbitals

Fractional Chern Insulators



# The Age of Mean Field



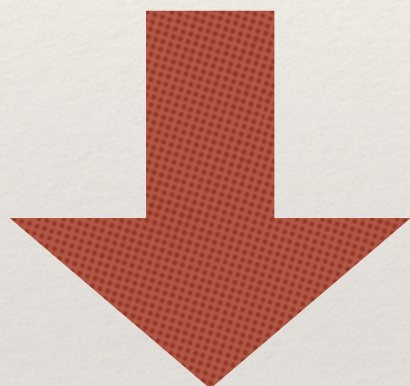
Slater Determinants  
BDG States  
Pfaffians

Jastrow



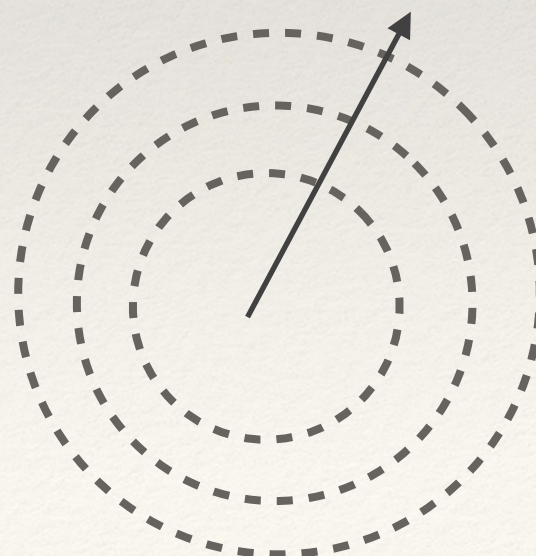
$$\Psi = \det M \exp[-U(R)]$$
$$\Psi = P[\det M] \text{ (spin liquids)}$$

$$\Psi(r_1, r_2, \dots, r_n) = \det M;$$
$$M_{ij} = \phi_i(r_j)$$



Multi-Determinants

More determinants  
Enough determinants give you everything

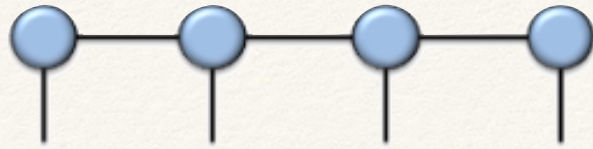




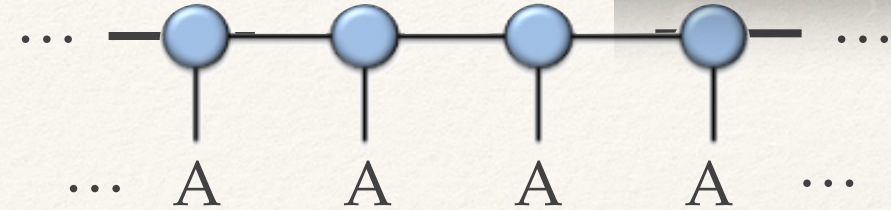
# The Age of Tensor Networks

Low entanglement ansatz

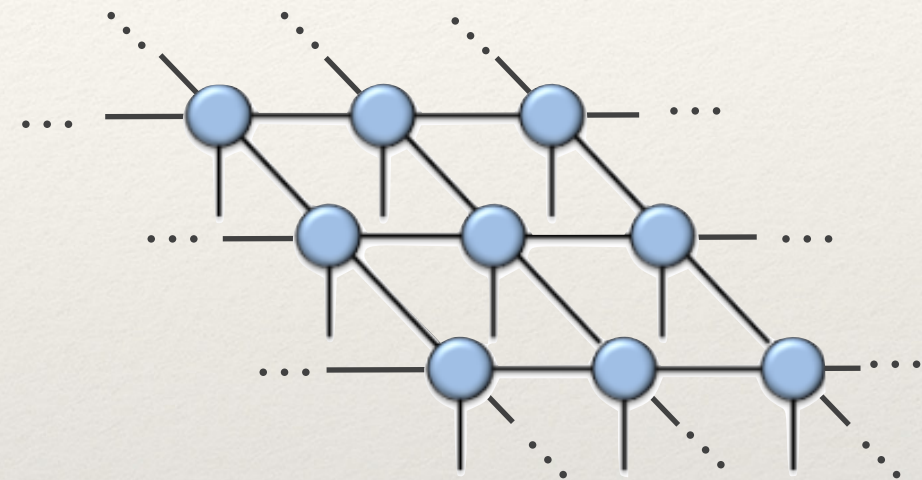
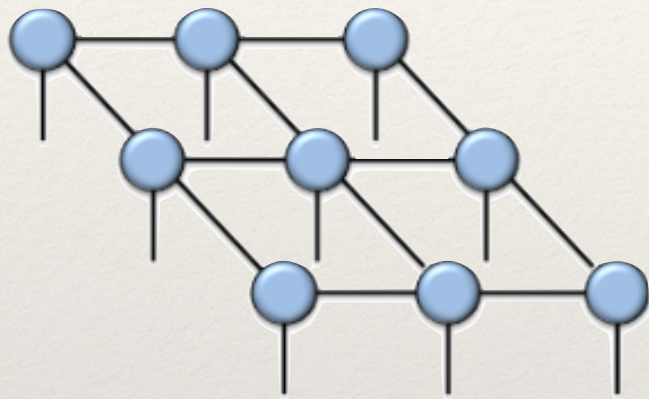
Matrix Product States



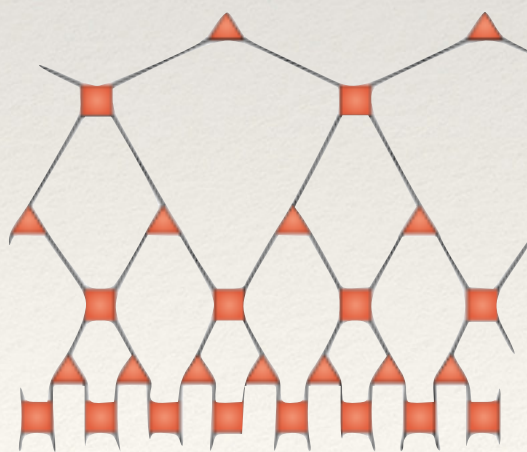
Infinite Ansatz



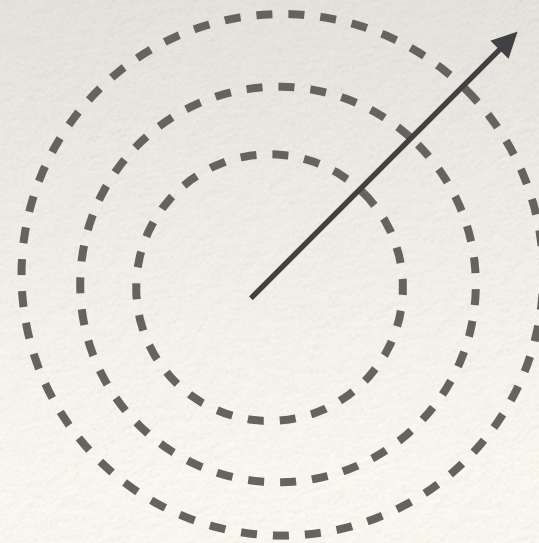
PEPS



MERA



Larger Bond-dimension  
Enough bond-dimension gives you everything

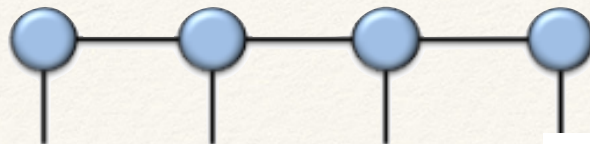




# The Age of Tensor Networks

Low entanglement ansatz

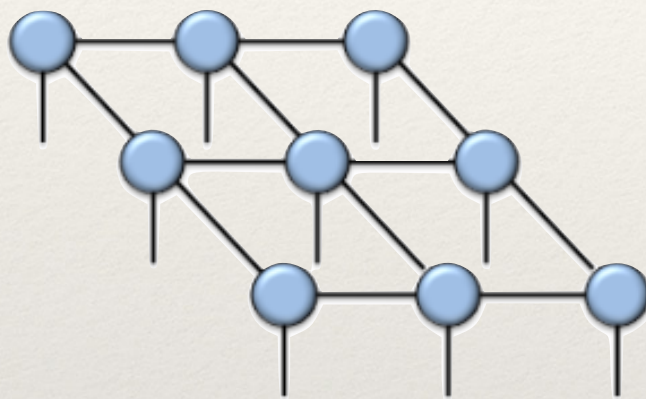
Matrix Product States



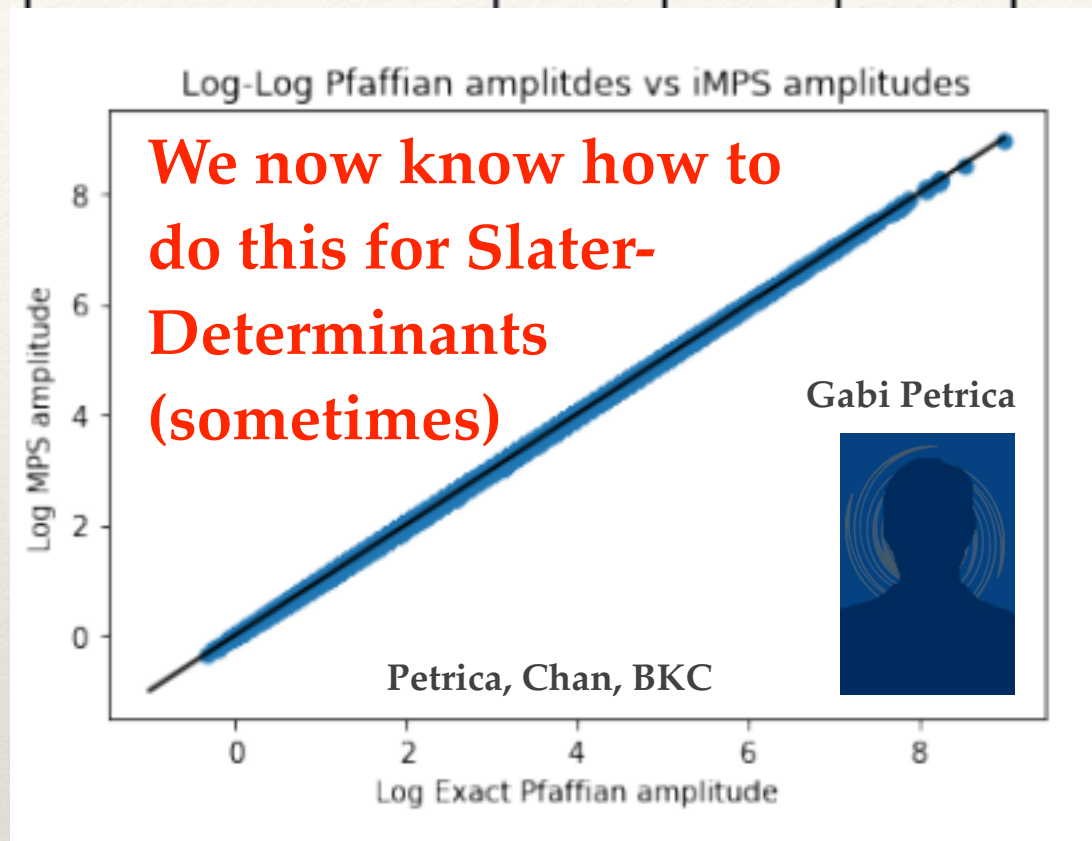
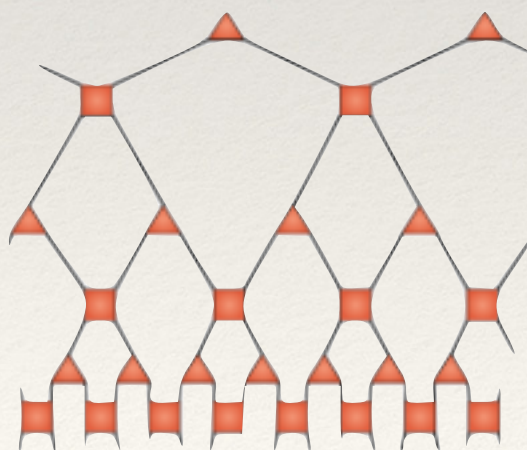
Infinite Ansatz



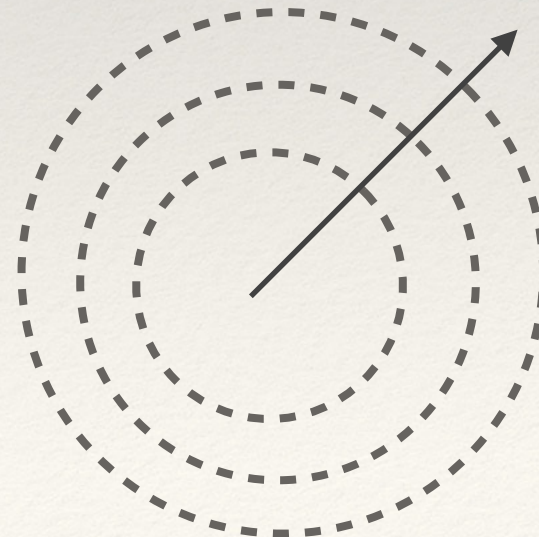
PEPS



MERA

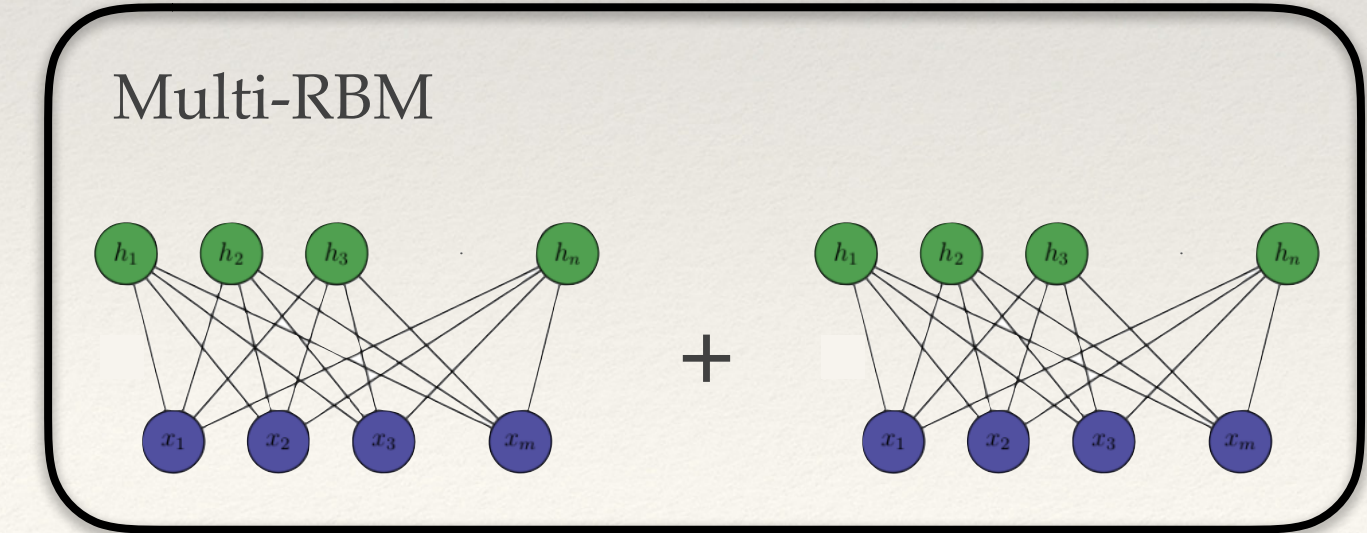
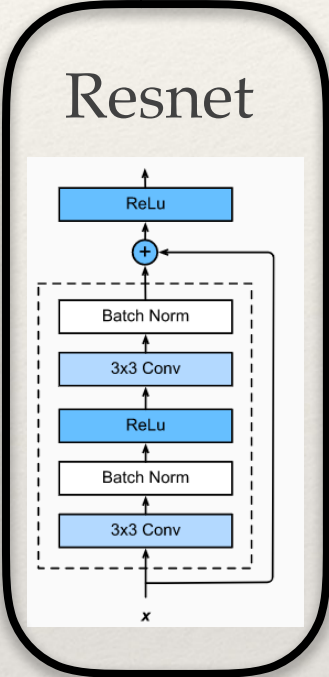
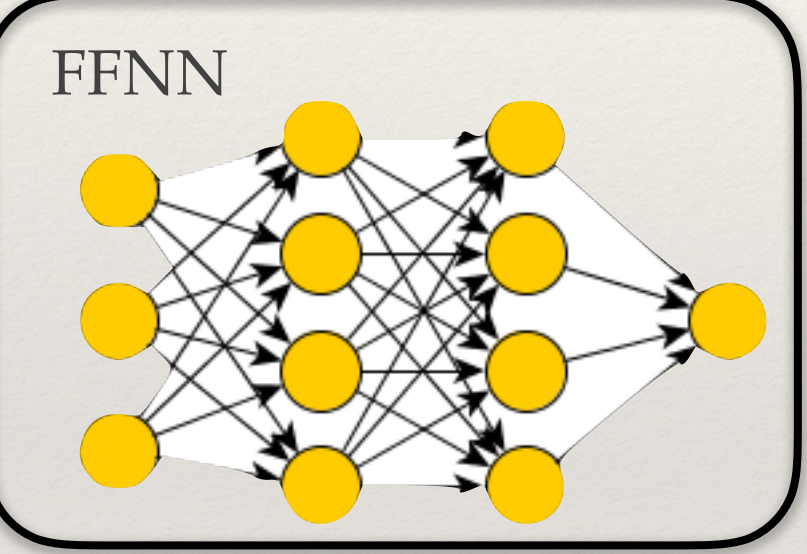
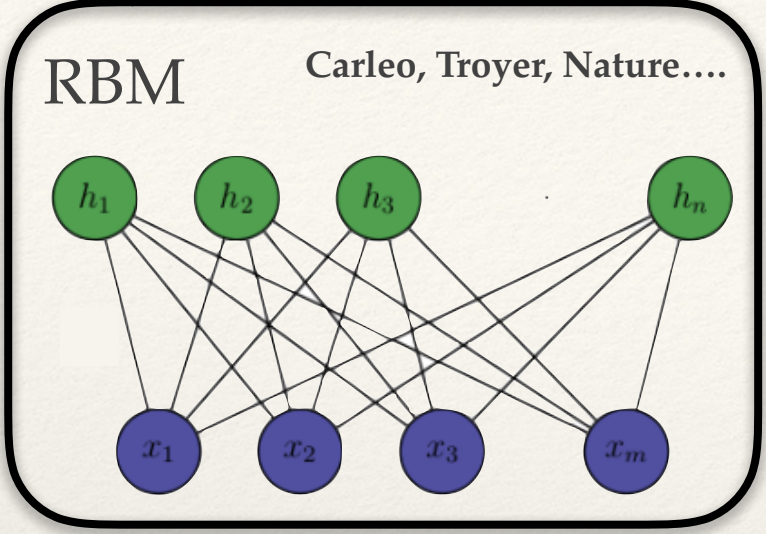
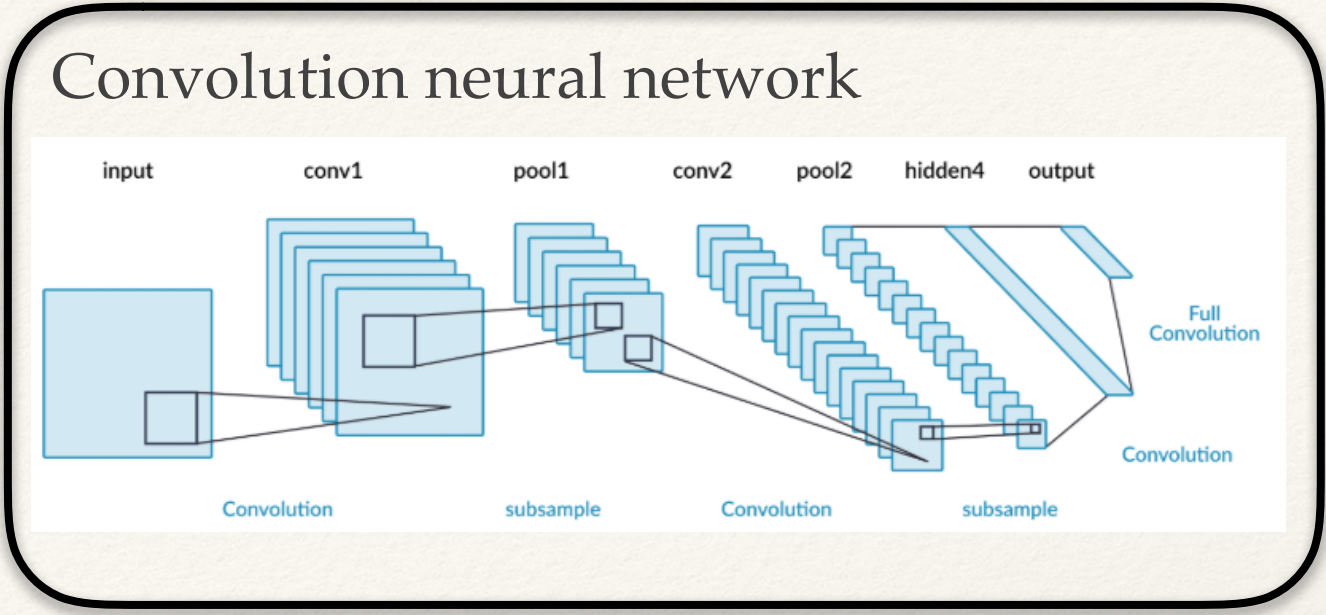


enough bond-dimension gives you everything

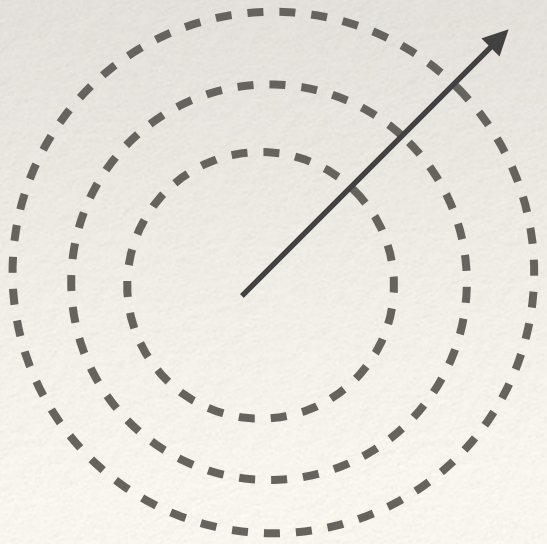




# The Age of AI



*Larger NN,  
Enough 'neurons' gives you everything*

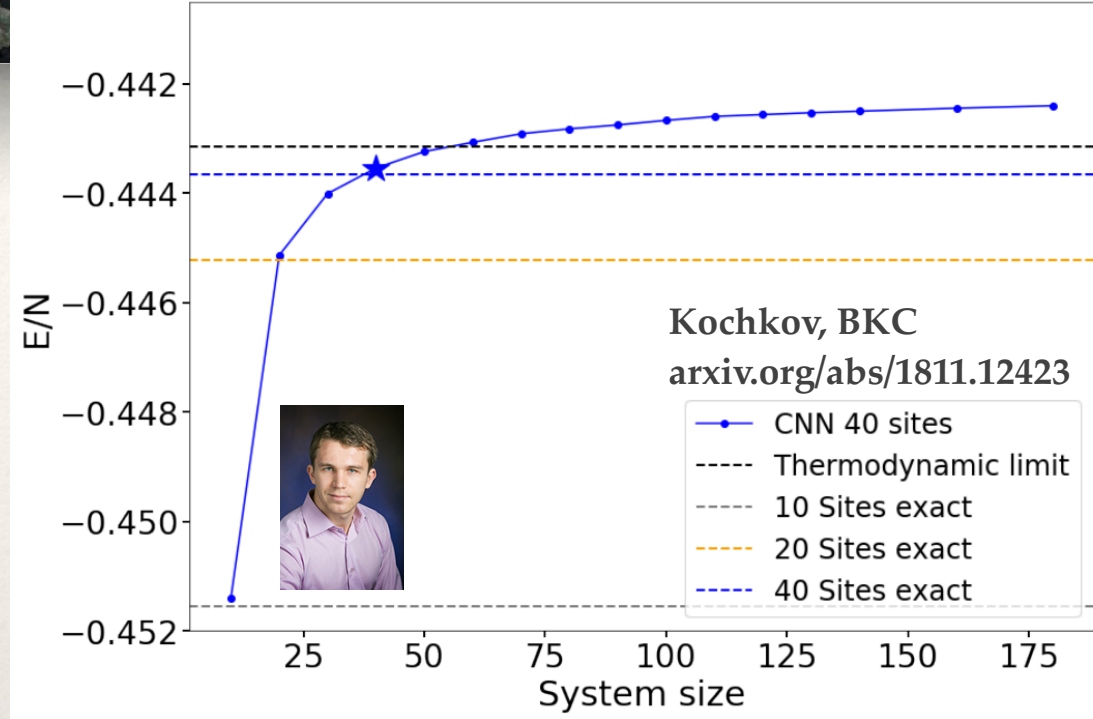
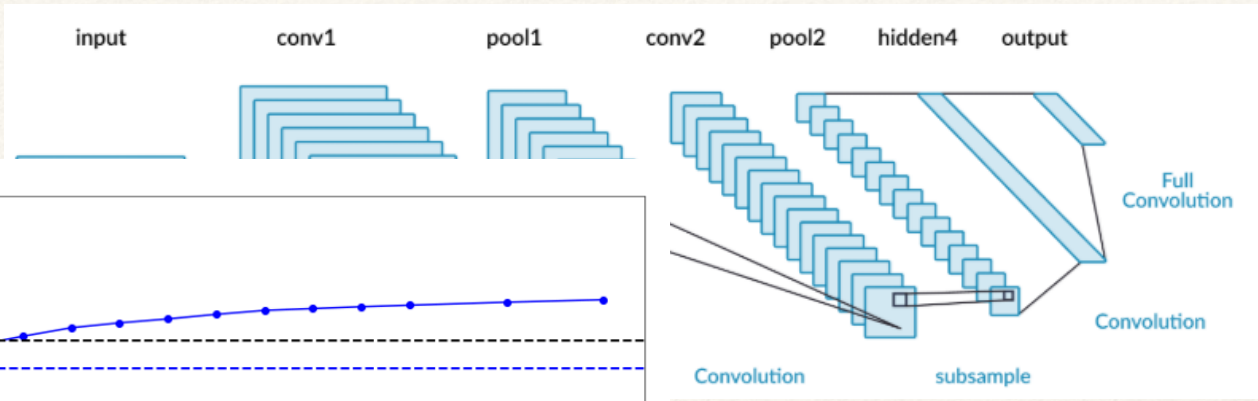




# The Age of AI

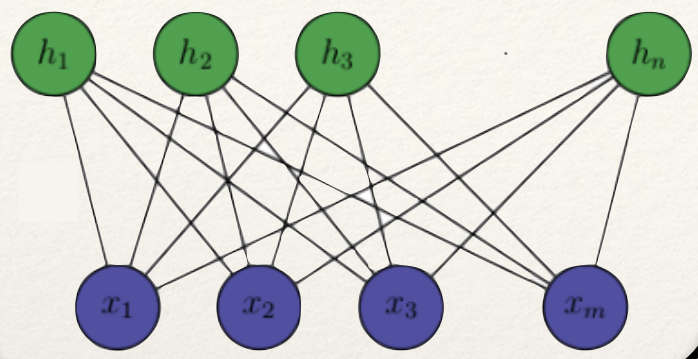


## Convolution neural network

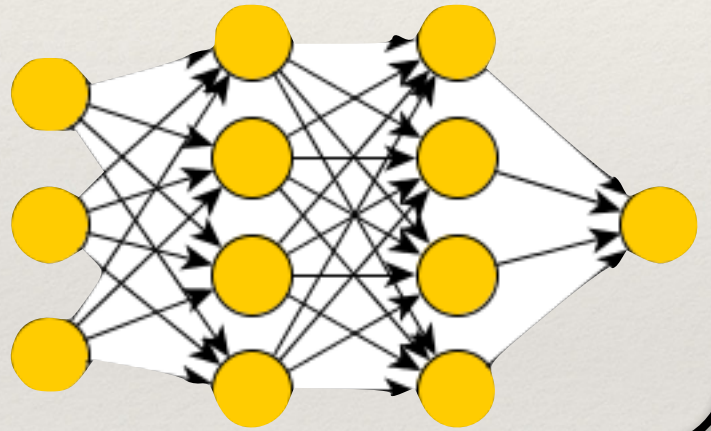


### A New Infinite Ansatz

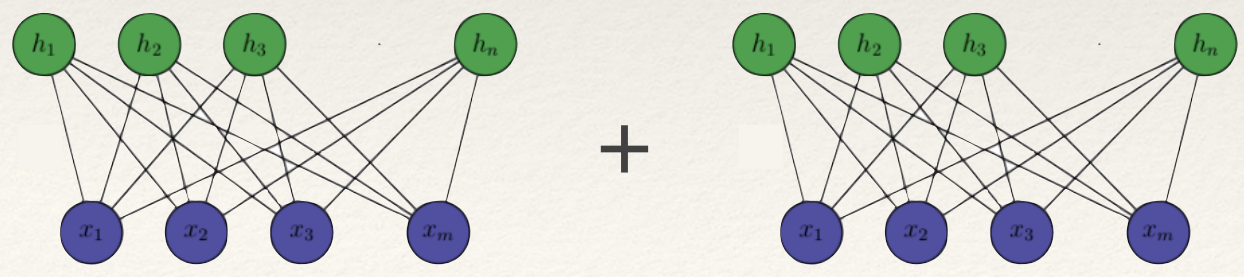
## RBM



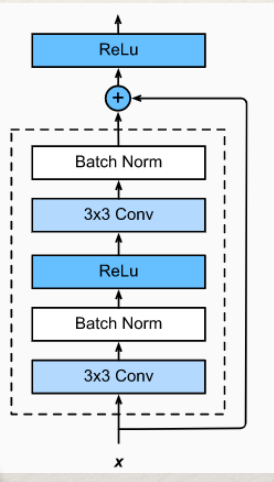
## FFNN



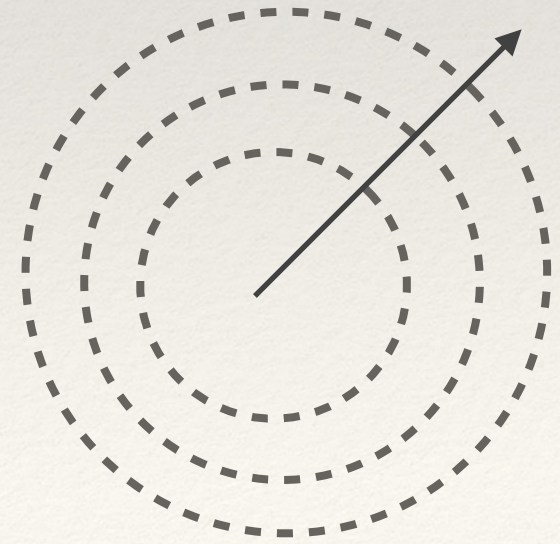
## Multi-RBM



## Resnet

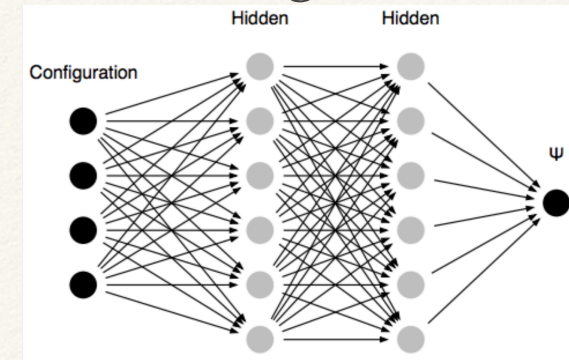


Larger NN,  
 Enough 'neurons' gives you everything

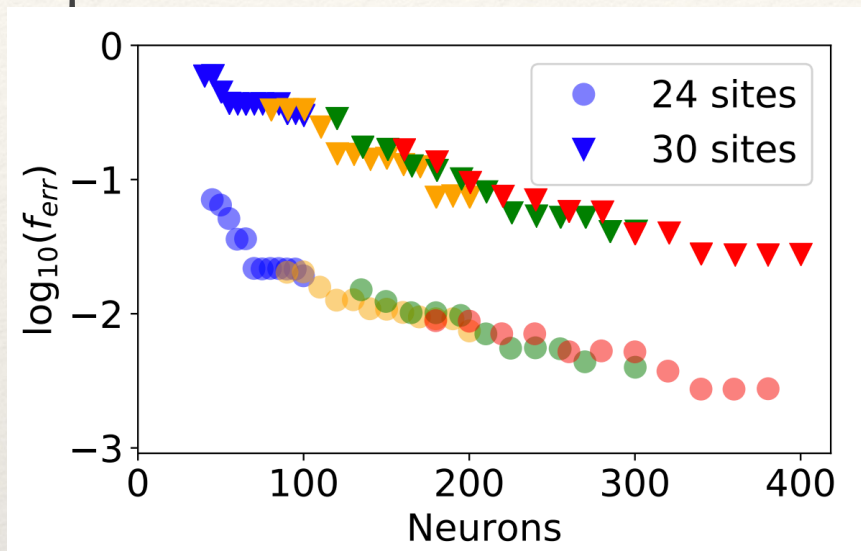




# How good are neural networks..



Square

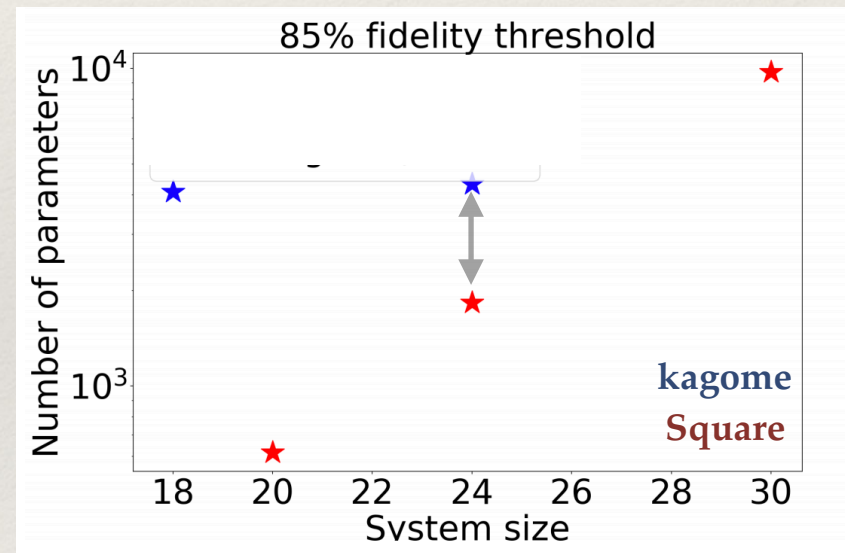
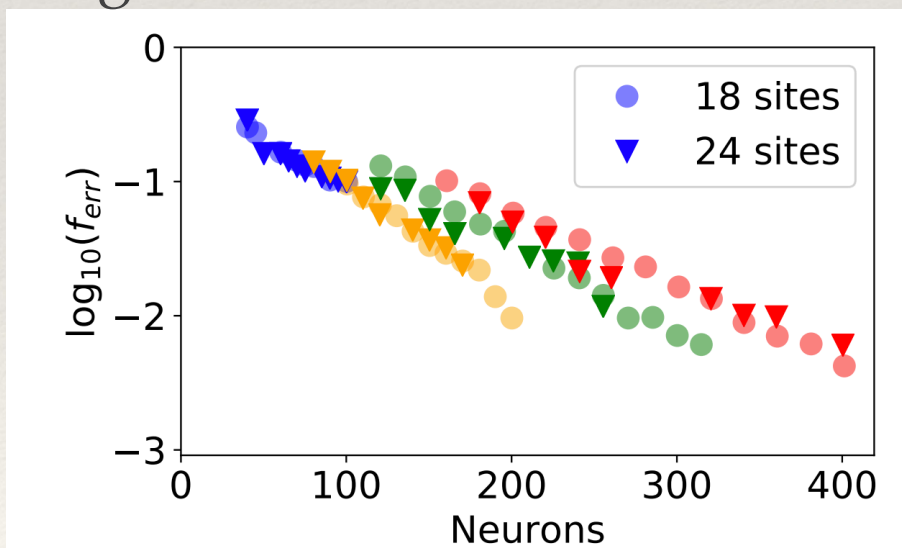


Exponential improvement with neuron number

Depth doesn't help (hurts a little?)

24-30 site systems require a lot of parameters ( $\sim 10^4$ )

Kagome



Square: Much fewer parameters...





# Wave-Functions with Signs

Neural Nets do well with wave-functions with simple sign structures (but so does QMC).  
What about wave-functions with complicated sign structure?

RBM: Can't even get sign structures without complex weights (and this hasn't worked out yet)

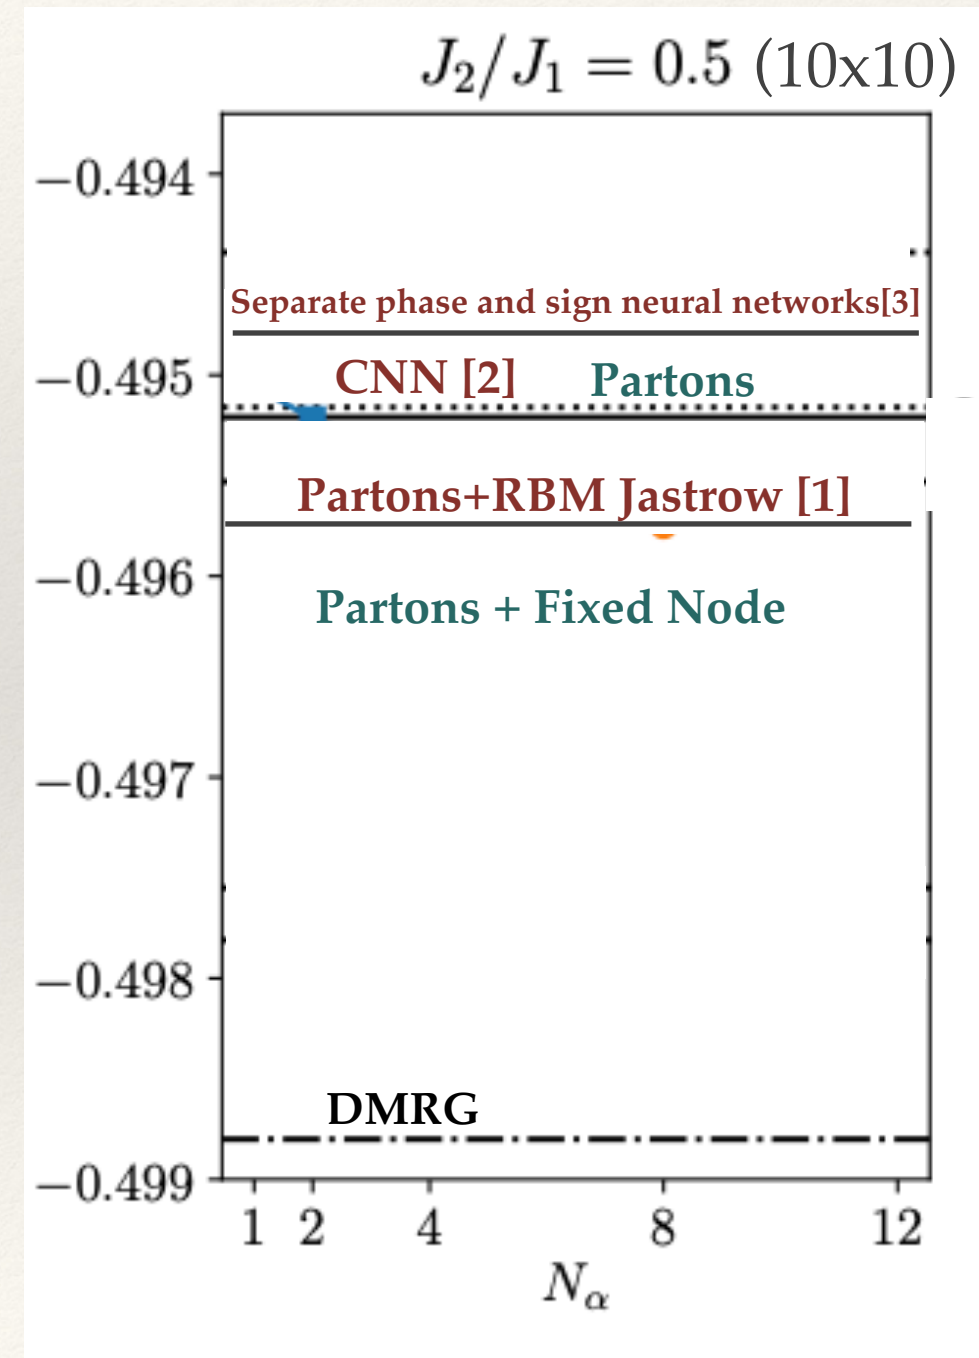
## Square Heisenberg

Neural-networks competitive with partons...

Worse than fixed-node+partons

Much worse with DMRG

*Partons+RBM-Jastrow also done in Hubbard by Imada*



### [1] Neural Gutzwiller-projected variational wave functions

Francesco Ferrari,<sup>1,\*</sup> Federico Becca,<sup>2</sup> and Juan Carrasquilla<sup>3,4</sup>

### [2] Study of the Two-Dimensional Frustrated J1-J2 Model with Neural Network Quantum States

Kenny Choo,<sup>1</sup> Titus Neupert,<sup>1</sup> and Giuseppe Carleo<sup>2</sup>

### [3] Neural network wave functions and the sign problem

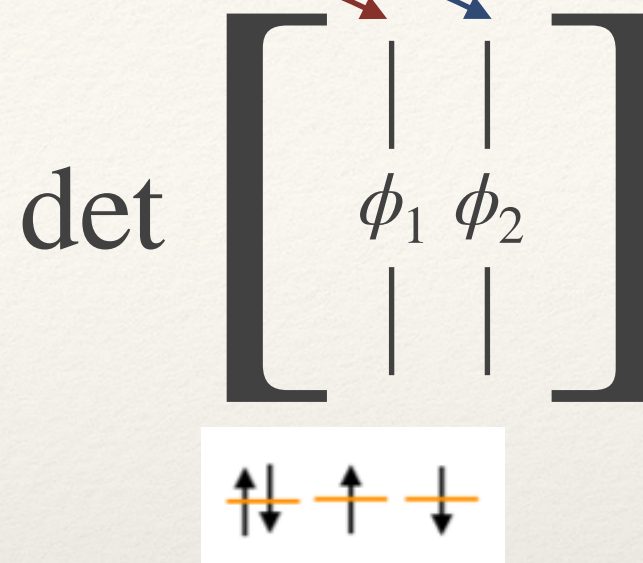
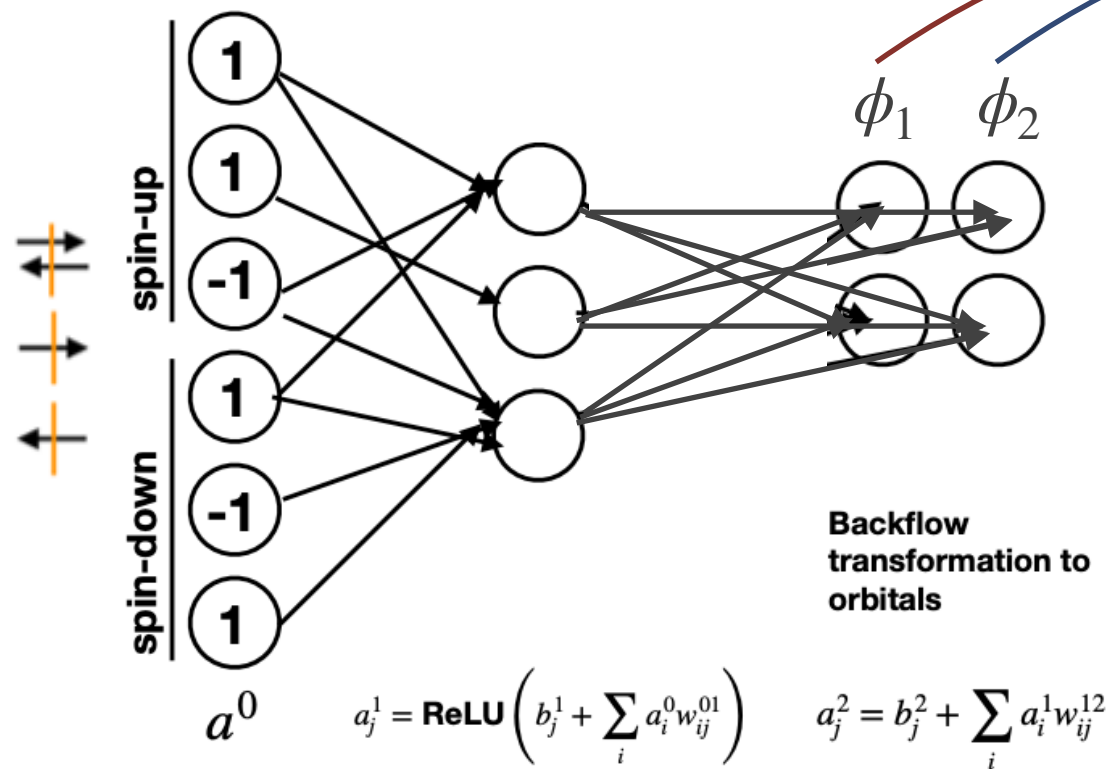
Attila Szabó and Claudio Castelnovo

### [4] Solving frustrated quantum many-particle models with convolutional neural networks

Xiao Liang,<sup>1,2</sup> Wen-Yuan Liu,<sup>1,2,3</sup> Pei-Ze Lin,<sup>1,2</sup> Guang-Can Guo,<sup>1,2</sup> Yong-Sheng Zhang,<sup>1,2,\*</sup> and Lixin He<sup>1,2,†</sup>



# Neural Network Backflow

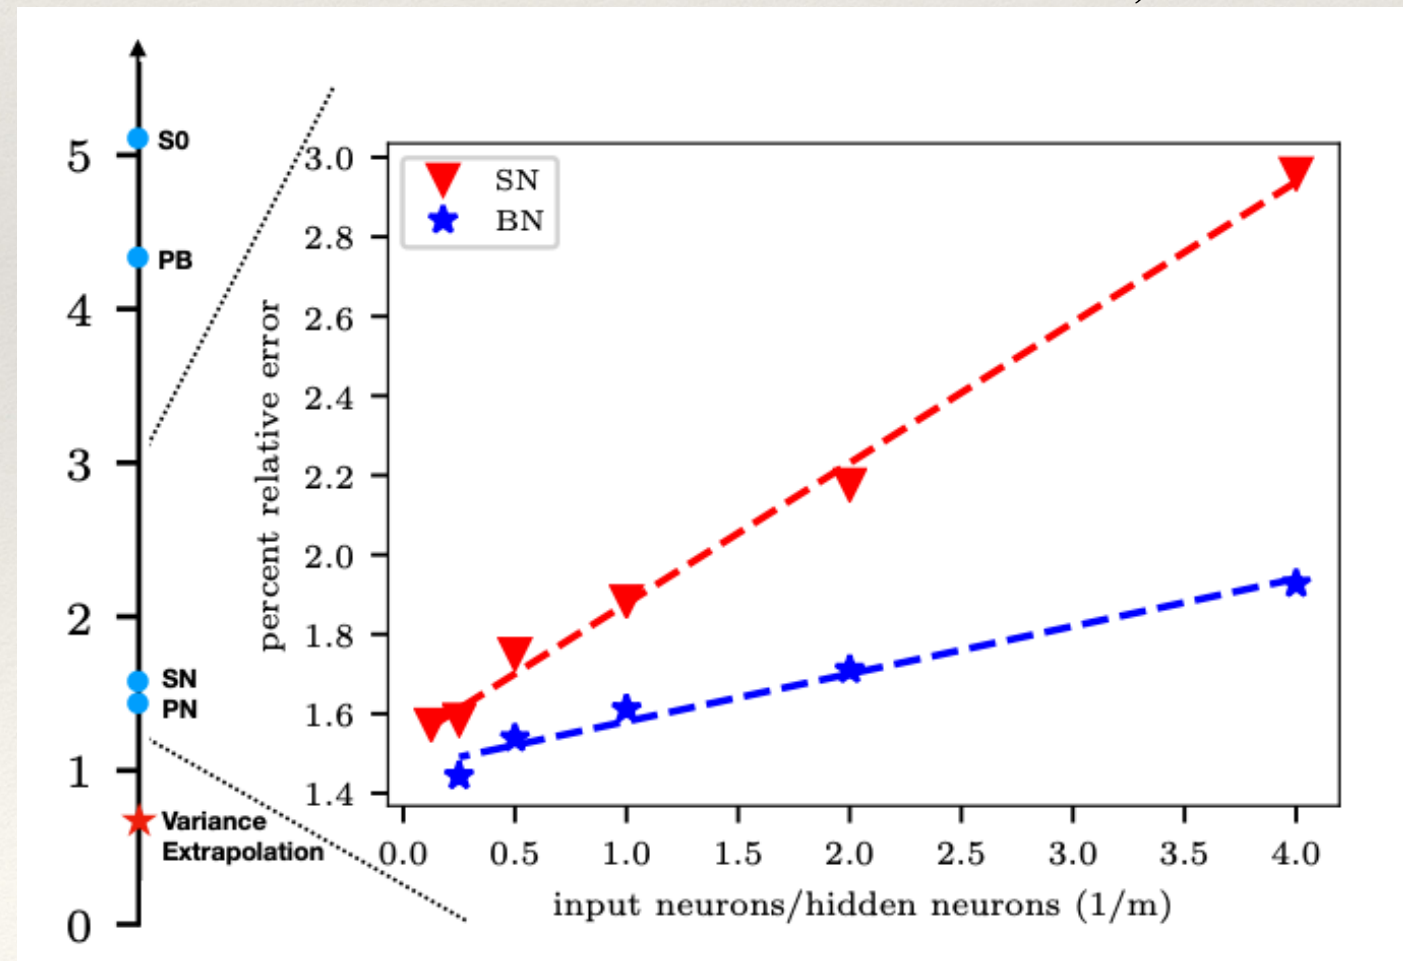


*4x4 Hubbard model at  $U/t=8$ ,  $n=0.875$*

$$\psi_{SD}(\mathbf{r}) = \det[M^{SD,\uparrow}] \det[M^{SD,\downarrow}];$$

$$M_{ik}^{SD,\sigma} = \phi_{k\sigma}(r_{i\sigma})$$

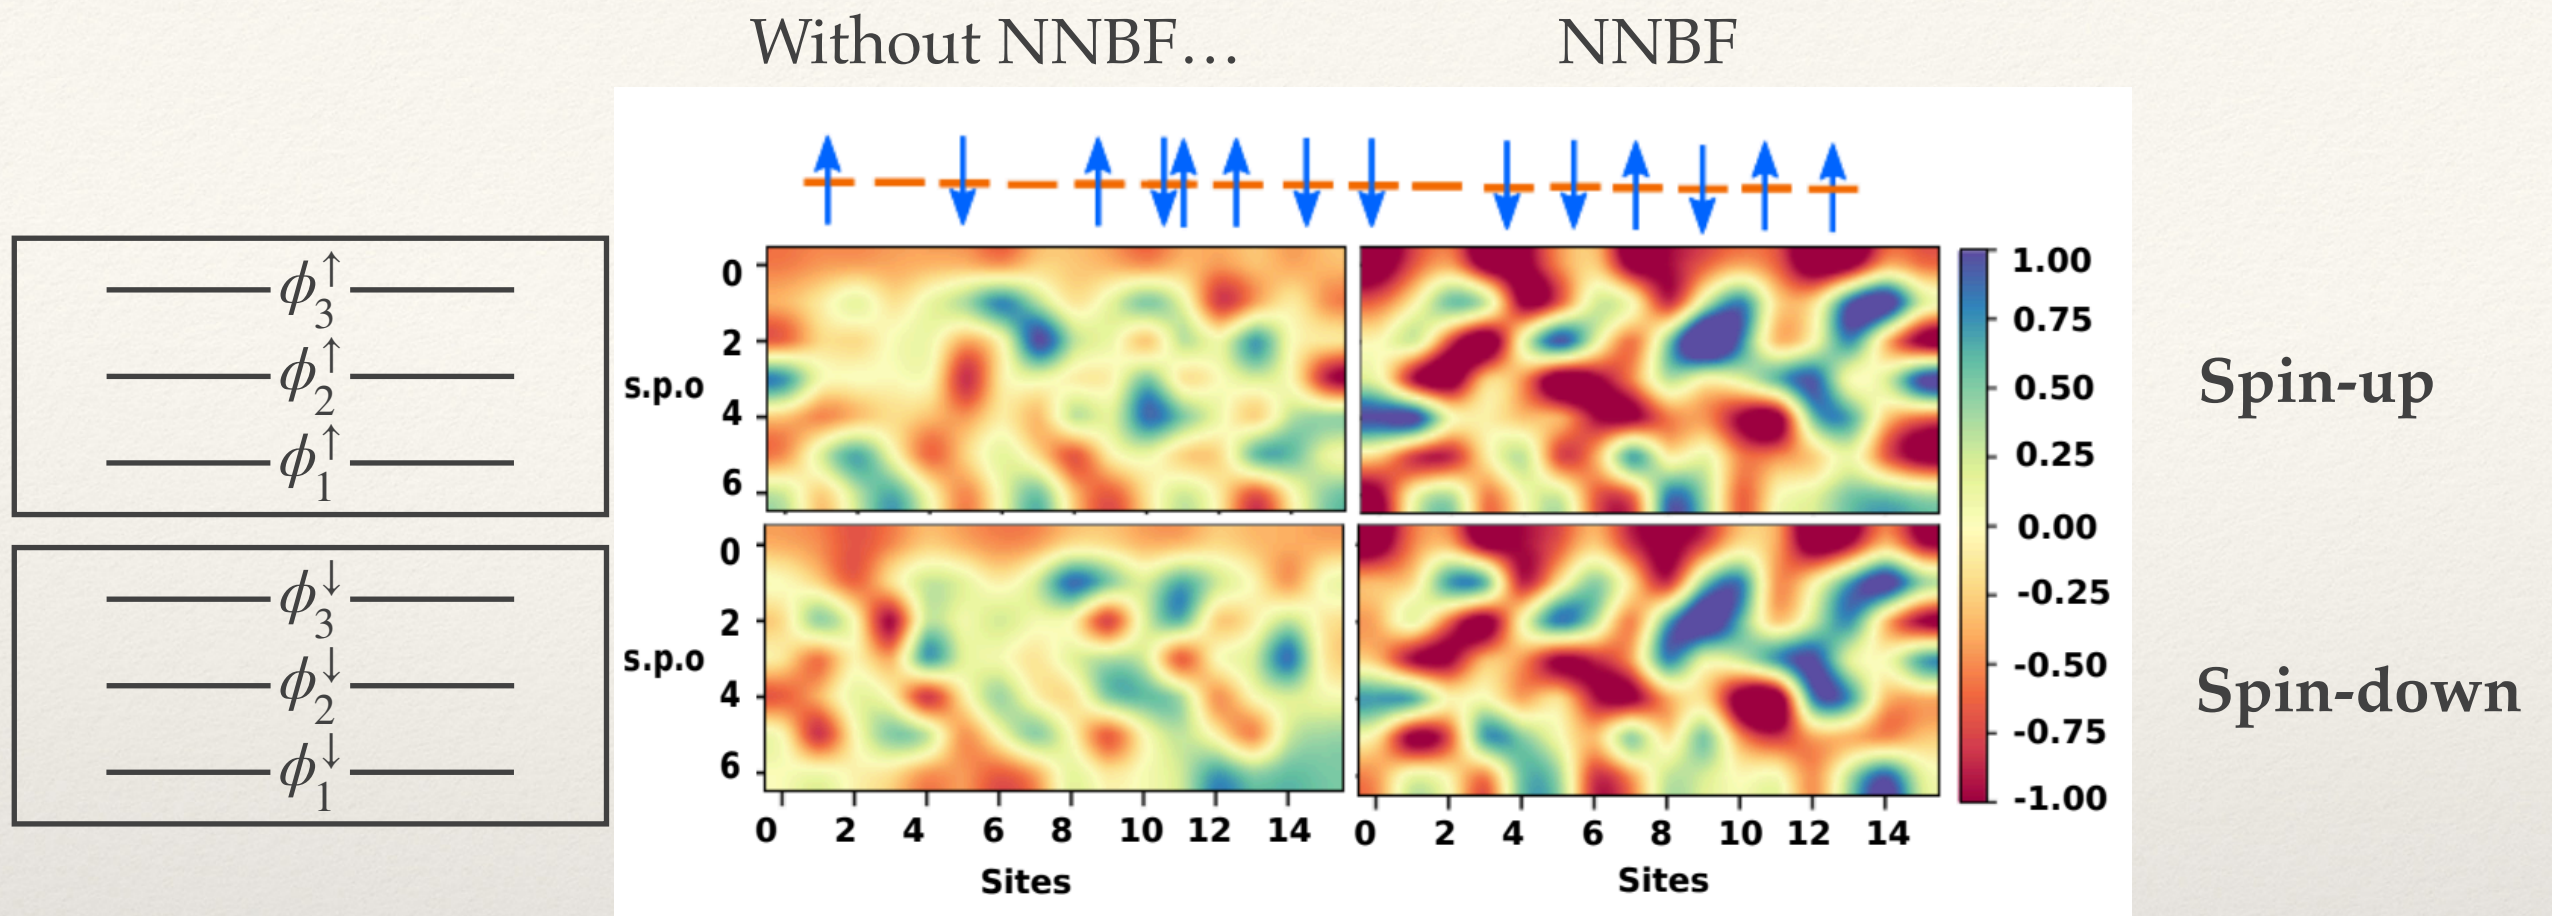
$$\phi_{k\sigma}^b(r_{i,\sigma}; \mathbf{r}) = \phi_{k\sigma}(r_{i,\sigma}) + a_{ki,\sigma}^{NN}(\mathbf{r})$$



Cost:  $O(N^4)$  per sweep



We not only get better energies; we restore the symmetry.



and change the signs...

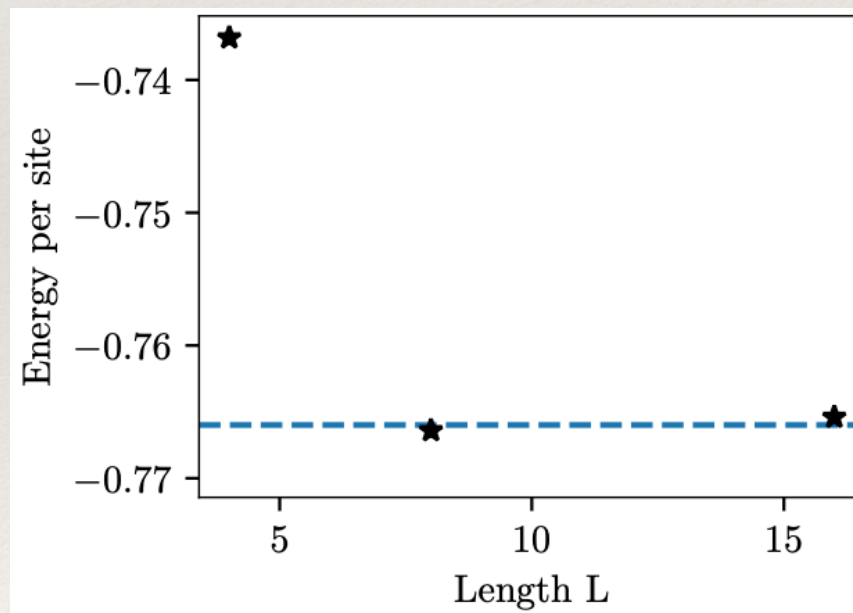
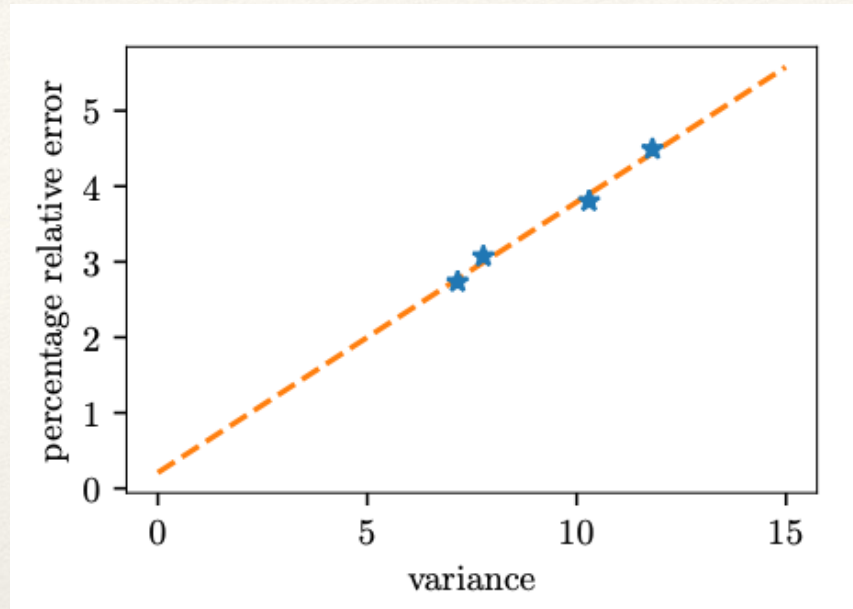
$$\frac{\int |\Psi_{S0}(x)|^2 \text{sgn}(\Psi_{SN}(x)) \text{sgn}(\Psi_{S0}(x)) dx}{\int |\Psi_{S0}(x)|^2 dx} = 0.815$$

9% difference between signs...



Bigger systems....

Relative energy error	Slater-Jastrow	NNB	NNB variance extrapolation
$16 \times 4$ Hubbard, $n=0.875$	$(5.9 \pm 2 \times 10^{-3})\%$	$(2.734 \pm 8 \times 10^{-3})\%$	0.209%
$12 \times 8$ Hubbard, $n=0.875$	$(6.3 \pm 3 \times 10^{-3})\%$	$(3.94 \pm 10^{-2})\%$	0.655%
$4 \times 4 \times 3$ Kagome	$(1.8 \pm 10^{-5})\%$	$(1.093 \pm 4 \times 10^{-3})\%$	0.286%



iDMRG answer

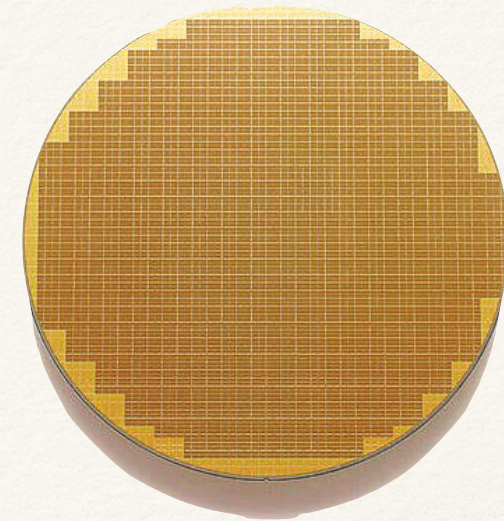
Solutions of the Two-Dimensional Hubbard Model: Benchmarks and Results from a Wide Range of Numerical Algorithms

J. P. F. LeBlanc, Andrey E. Antipov, Federico Becca, Ireneusz W. Bulik, Garnet Kin-Lic Chan, Chia-Min Chung, Youjin Deng, Michel Ferrero, Thomas M. Henderson, Carlos A. Jiménez-Hoyos, E. Kozik, Xuan-Wen Liu, Andrew J. Millis, N. V. Prokof'ev, Mingpu Qin, Gustavo E. Scuseria, Hao Shi, B. V. Svistunov, Luca F. Tocchio, I. S. Tupitsyn, Steven R. White, Shiwei Zhang, Bo-Xiao Zheng, Zhenyue Zhu, and Emanuel Gull (Simons Collaboration on the Many-Electron Problem)  
 Phys. Rev. X **5**, 041041 – Published 14 December 2015



# Beyond the silver age...

(and how close are we to reaching it?)



Silicon Age



We'd like to consider the wave-functions accessible 'quickly' by computers.

Q: What's a good way to represent this space?

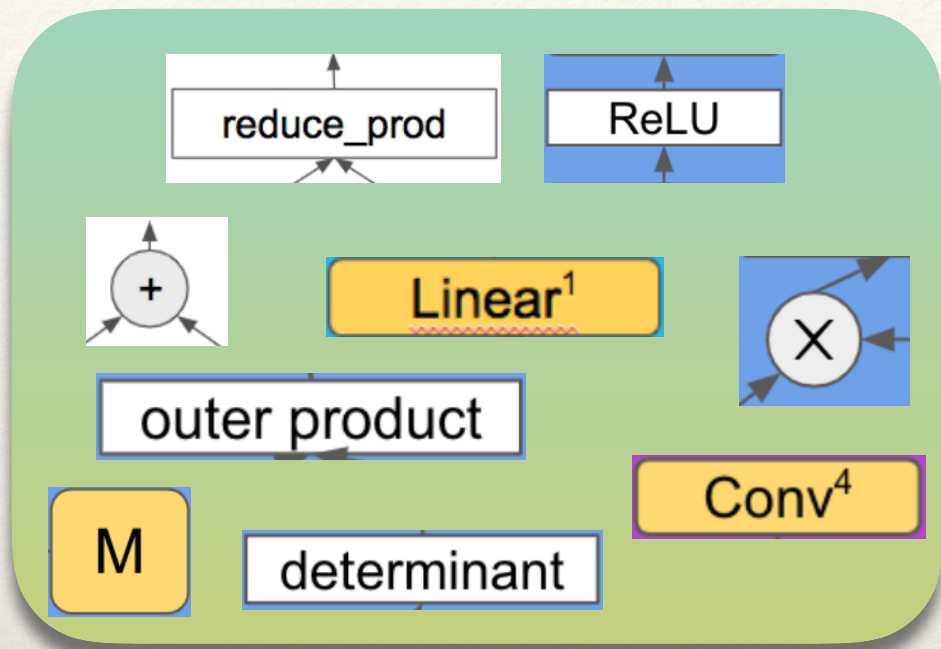
*Our answer (stolen from machine-learning):* computational graph states...





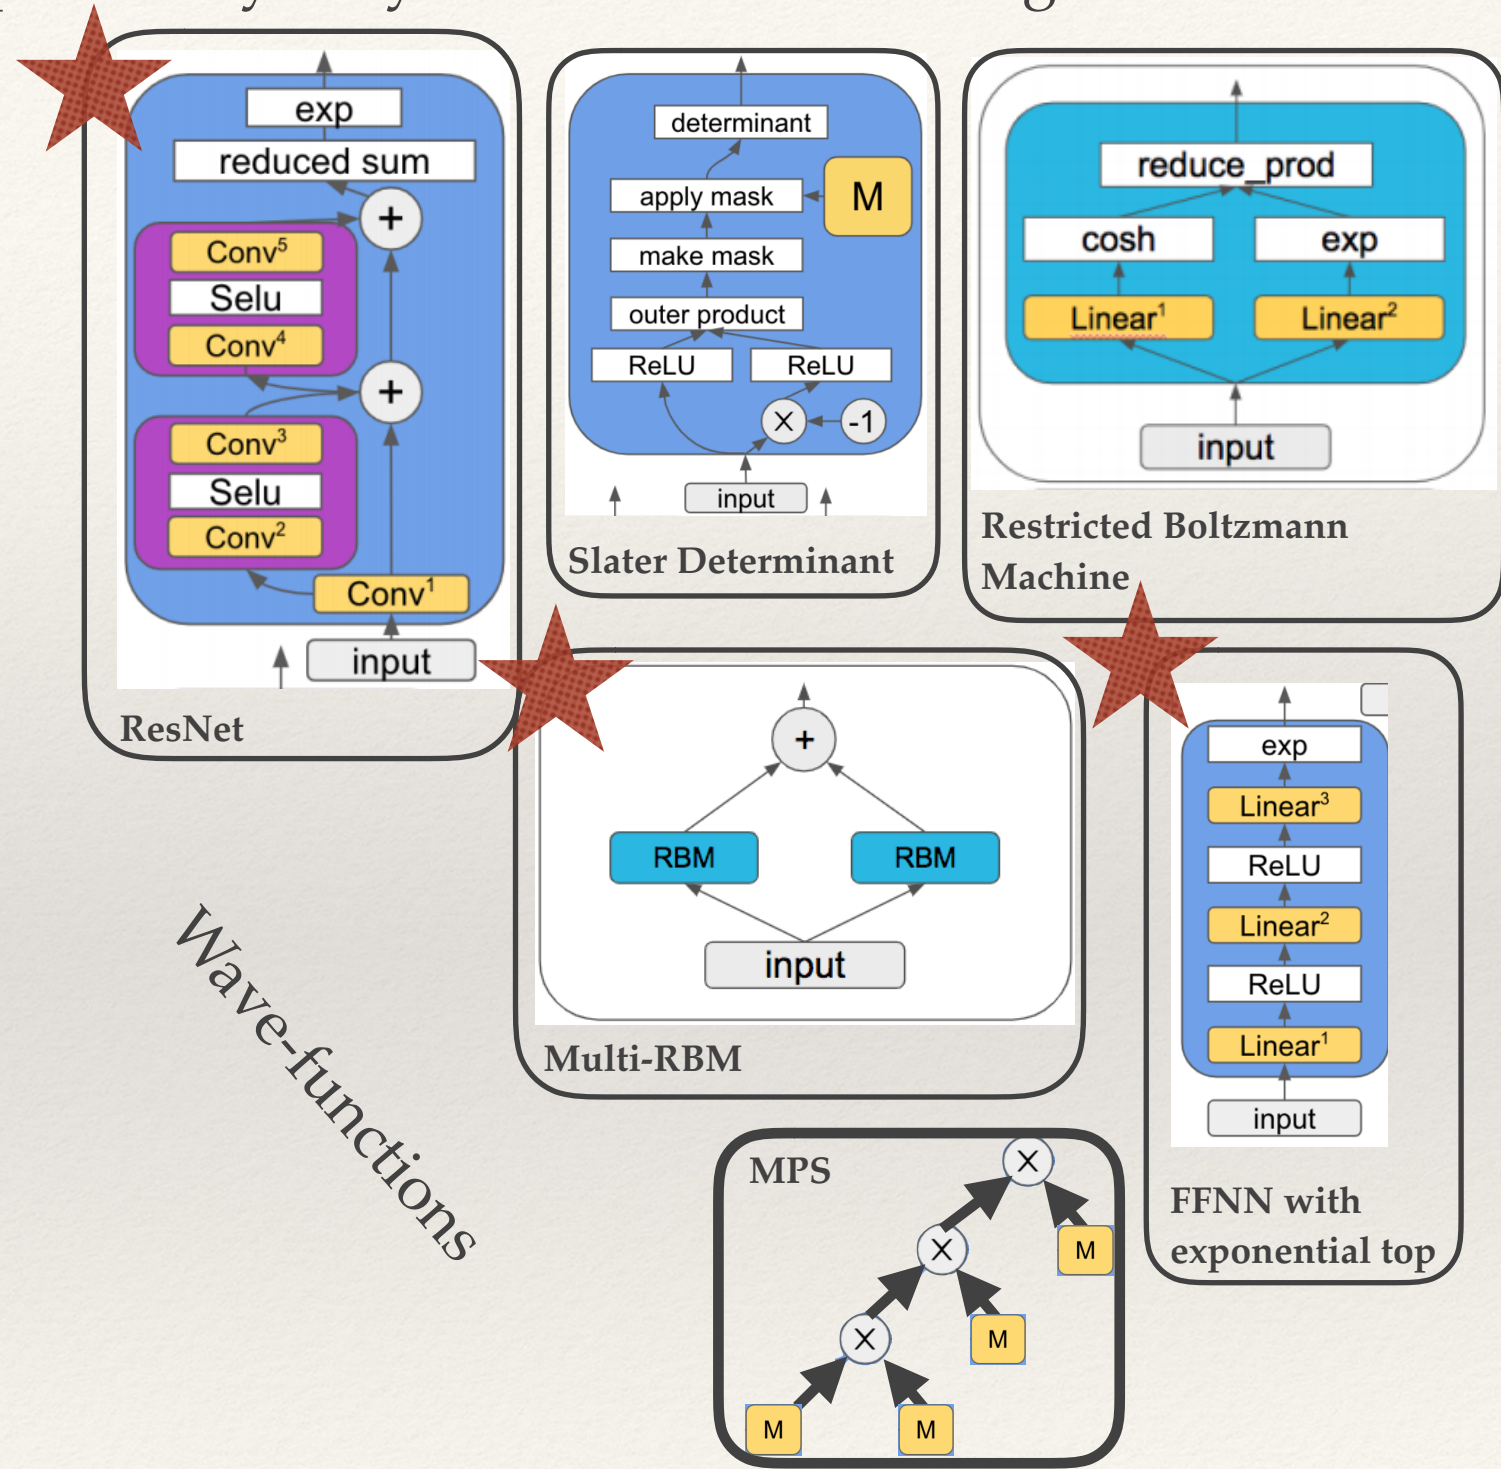
# Computational Graph States

Taking a page, from the machine learning play-book, we can represent all “fast wave-functions” as a computational graph. Everybody is on the same footing now.



Box of operations...

WF	Hidden Layers	Hidden Units	Params 1d - 2d
FCNN	2	80	9800 - 11744
RBM	0	80	3320 - 5264
FC-RBM	2	80	16280 - 18224
Conv1D	5	16 filters (size 5)	5280 -
Conv2D	5	16 filters (5 by 5)	- 26080
ResNet	1 + 2(2)	16 filters (size 5)	5280 - 25760
Multi-RBM	2	16 filters (5 by 5)	6640 - 10528
Multi-FCNN	2	80	19600 - 23488
P-BDG			1600 - 4096
MPS			4800 -



Wave-functions

Good representation...still a question of optimization



```

876 class ProjectedBDG(Wavefunction):
877     """P-BDG module."""
878
879     def __init__(
880         self,
881         num_sites: int,
882         name: str = 'projected_bdg'):
883         """Constructs a projected BDG module.
884
885         Args:
886             num_sites: Number of sites.
887             name: Name of the module.
888         """
889         super(ProjectedBDG, self).__init__(name=name)
890         self._num_sites = num_sites
891         with self._enter_variable_scope():
892             self._pairing_matrix = tf.get_variable(
893                 'pairing_matrix', shape=[1, num_sites, num_sites], dtype=tf.float32)
894
895     def _build(self, inputs: tf.Tensor) -> tf.Tensor:
896         """Connects the P-BDG module into the graph with input `inputs`.
897
898         Args:
899             inputs: Tensor with input values of shape=[batch] and values +/- 1.
900
901         Returns:
902             Wave-function amplitudes of shape=[batch].
903         """
904         batch_size = inputs.shape[0]
905         n_sites = self._num_sites
906         mask = tf.einsum('ij,ik->ijk', tf.nn.relu(inputs), tf.nn.relu(-inputs))
907         bool_mask = tf.greater(mask, tf.zeros([batch_size, n_sites, n_sites]))
908         tiled_pairing = tf.tile(self._pairing_matrix, [batch_size, 1, 1])
909         det_size = [batch_size, n_sites // 2, n_sites // 2]
910         pre_det = tf.reshape(tf.boolean_mask(tiled_pairing, bool_mask), det_size)
911
912         sign, ldet = tf.linalg.slogdet(pre_det)
913         det_value = tf.exp(self.add_exp_normalization(ldet))
914         return sign * det_value
915
916     @classmethod
917     def from_hparams(
918         cls,
919         hparams: tf.contrib.training.HParams,
920         name: str = ''
921     ) -> 'Wavefunction':
922         """Constructs an instance of a class from hparams."""
923         pbdg_params = {
924             'num_sites': hparams.num_sites,
925         }
926         if name:
927             pbdg_params['name'] = name
928         return cls(**pbdg_params)
929

```

```

557     super(Conv2DNetwork, self).__init__(name=name)
558     self._num_layers = num_layers
559     self._num_filters = num_filters
560     self._kernel_size = kernel_size
561     self._nonlinearity = nonlinearity
562     self._output_activation = output_activation
563     self._size_x = size_x
564     self._size_y = size_y
565
566     reduction = functools.partial(tf.reduce_sum, axis=[1, 2, 3])
567     self._components = []
568     with self._enter_variable_scope():
569         for layer in range(num_layers):
570             self._components.append(layers.Conv2dPeriodic(num_filters, kernel_size))
571             if layer + 1 != num_layers:
572                 self._components.append(nonlinearity)
573             if output_activation == tf.exp:
574                 self._components += [reduction, self.add_exp_normalization, tf.exp]
575             else:
576                 self._components += [reduction, output_activation]
577
578     def _build(
579         self,
580         inputs: tf.Tensor,
581     ) -> tf.Tensor:
582         """Builds computational graph evaluating the wavefunction on inputs.
583
584         Args:
585             inputs: Input tensor, must have shape (batch, num_sites, ...).
586
587         Returns:
588             Tensor holding values of the wavefunction on `inputs`.
589
590         Raises:
591             ValueError: Input tensor has wrong shape.
592         """
593         inputs_new_shape = [-1, self._size_x, self._size_y, 1]
594         inputs = tf.reshape(inputs, shape=inputs_new_shape)
595         return snt.Sequential(self._components)(inputs)
596
597     @classmethod
598     def from_hparams(
599         cls,
600         hparams: tf.contrib.training.HParams,
601         name: str = ''
602     ) -> 'Wavefunction':
603         """Constructs an instance of a class from hparams."""
604         conv_2d_params = {
605             'num_layers': hparams.num_conv_layers,
606             'num_filters': hparams.num_conv_filters,
607             'kernel_size': hparams.kernel_size,
608             'size_x': hparams.size_x,
609             'size_y': hparams.size_y,
610             'output_activation': layers.NONLINEARITIES[hparams.output_activation],
611             'nonlinearity': layers.NONLINEARITIES[hparams.nonlinearity],
612         }
613         if name:
614             conv_2d_params['name'] = name
615         return cls(**conv_2d_params)

```



## Part 2

# Optimizing wave-functions in the age of AI





# Wave-Function optimization through history

## Ancient History

### Stochastic Gradient Descent

Given an objective function, walk downhill.

$$E(\vec{\theta}) \equiv \langle \Psi(\vec{\theta}) | H | \Psi(\vec{\theta}) \rangle$$
$$\theta_i \leftarrow \theta_i - \delta \frac{\partial E}{\partial \theta_i}$$

$O(p)$

Slow convergence

(actually, pre-ancient history optimized the variance)

### Imaginary Time Evolution

Stochastic Reconfiguration

*i*TEBD for VMC

$$|\Psi_{new}\rangle = \mathcal{P}(1 - \tau H) |\Psi_{old}\rangle$$

Projection into manifold...

Need:  $S_{ij}^{-1}$  where  $S_{ij} \equiv \left\langle \frac{\partial \Psi}{\partial \theta_i} \middle| \frac{\partial \Psi}{\partial \theta_j} \right\rangle$

$O(p^3)$

Numerically unstable

Undersampling problem

### Linear Method

DMRG for VMC

$$\widetilde{H} |\Psi\rangle = E S |\Psi\rangle$$

$$\widetilde{H}_{ij} \equiv \left\langle \frac{\partial \Psi}{\partial \theta_i} \middle| H \middle| \frac{\partial \Psi}{\partial \theta_j} \right\rangle$$




$O(p^3)$

Numerically unstable

Complicated



# Supervised Learning...

<u>IMAGE</u>	<u>LABEL</u>
	<i>CAT</i>
	<i>DOG</i>
	<i>CAT</i>


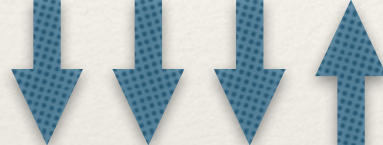

Train your supervised algorithm on these labels.

If you give it a new image, it should be able to tell if it's a cat or dog.

Cost:  $O(p)$



# Supervised Learning...

<u>Configuration</u>	<u>Label</u>
	0.005
	-0.0035
	0.105

Train your supervised algorithm on these labels.

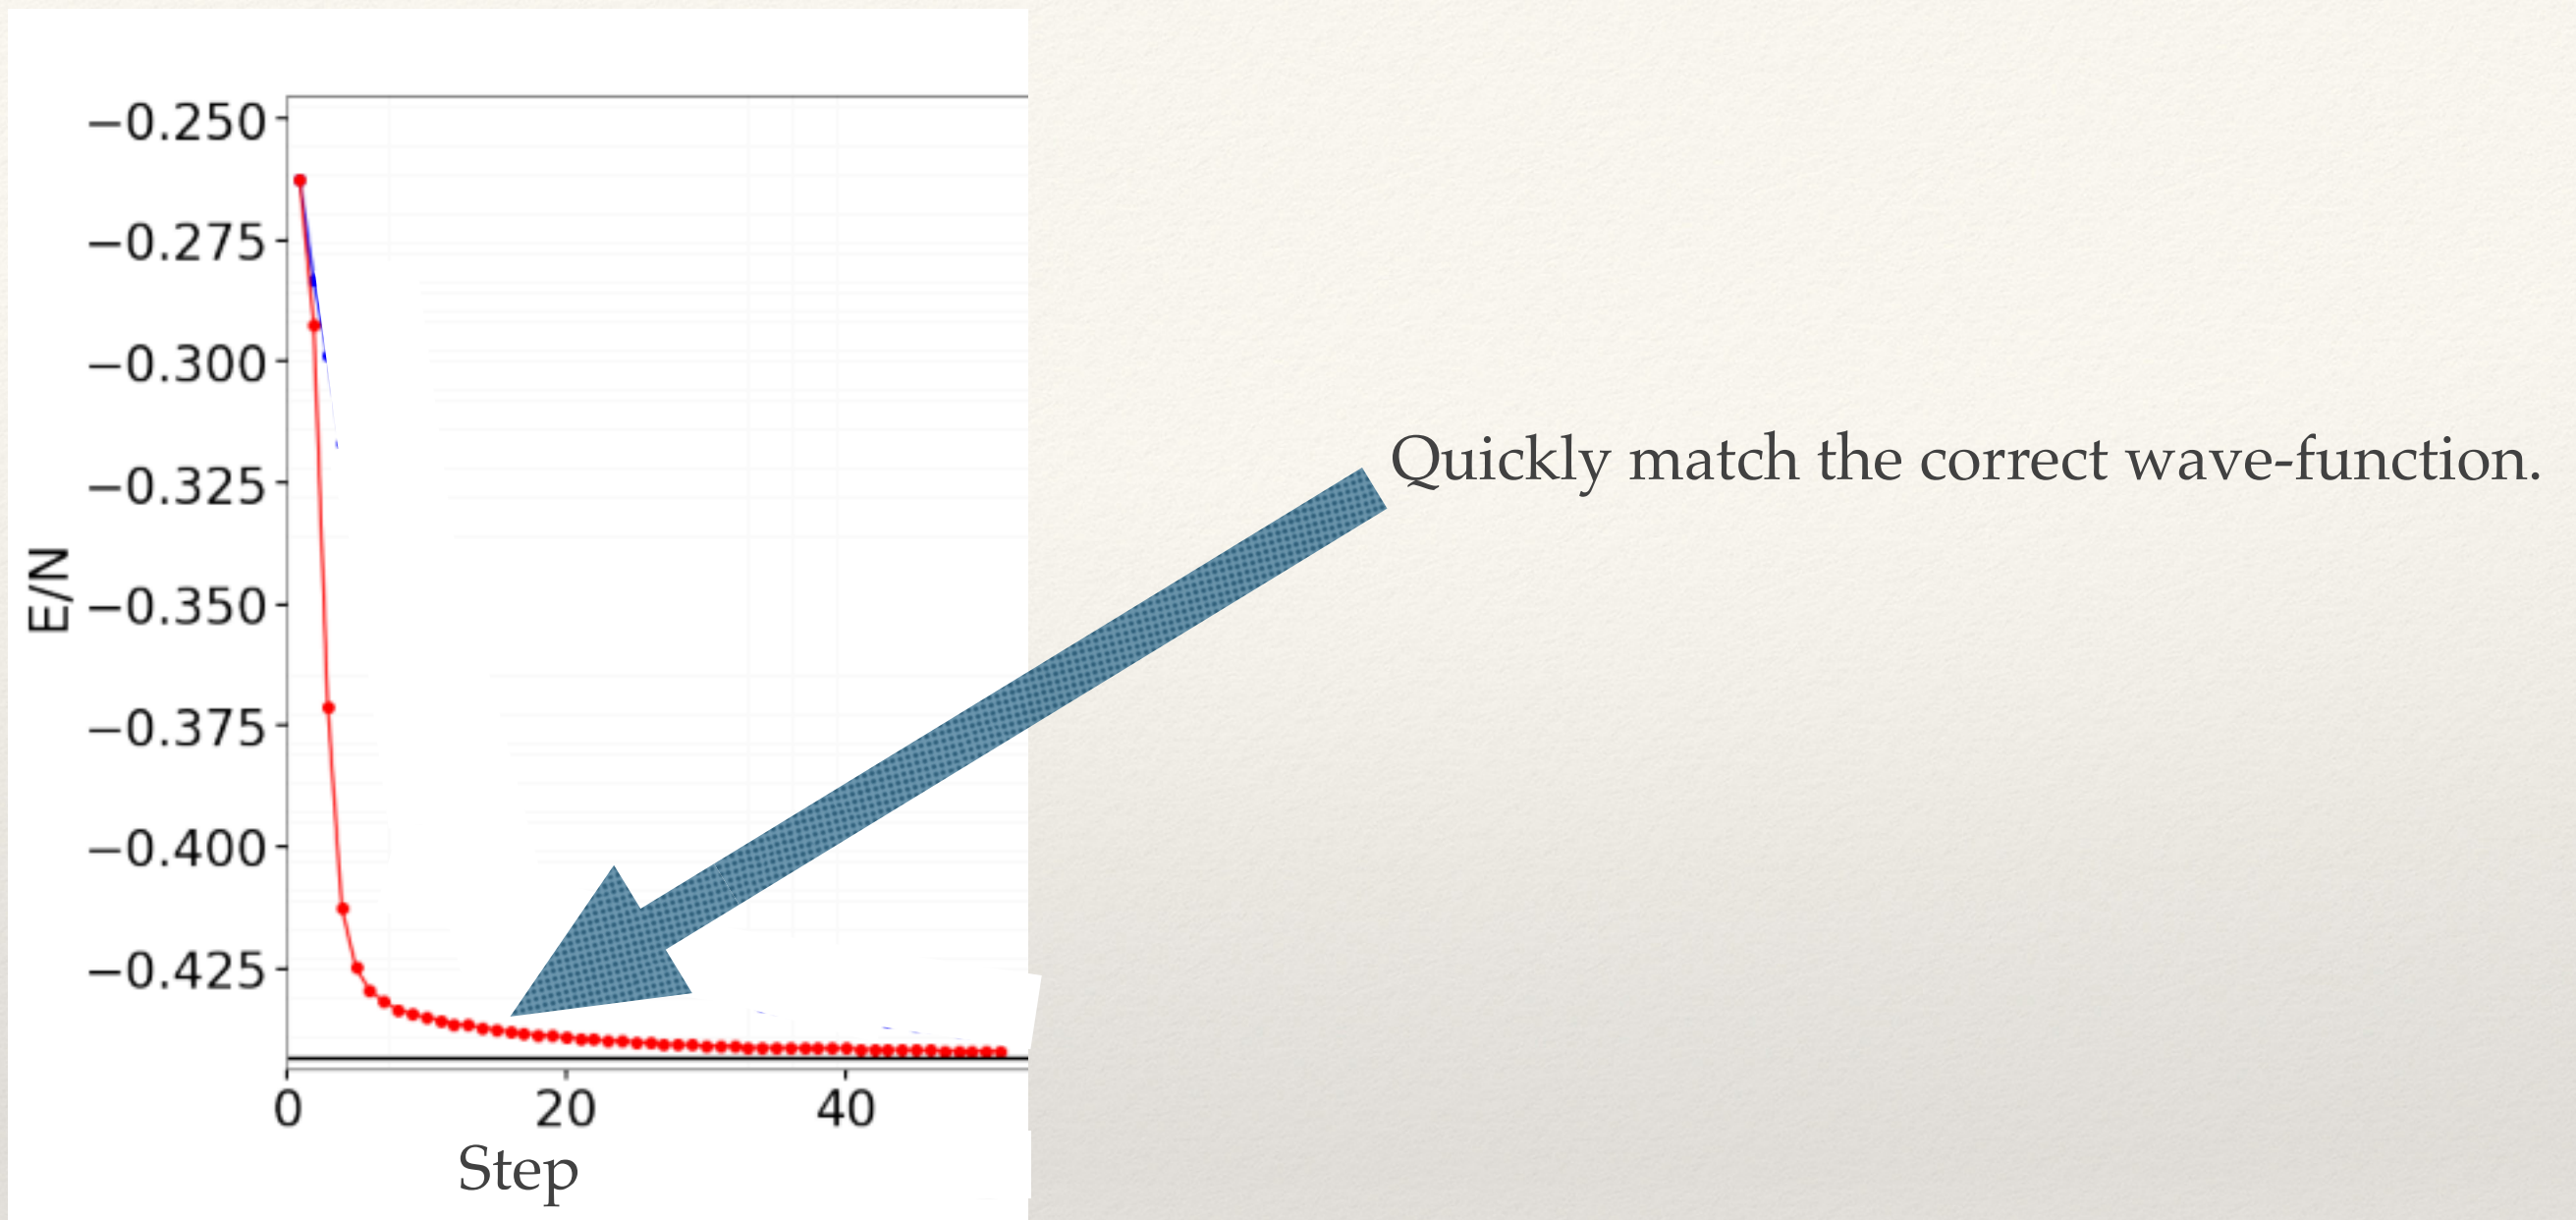
If you give it a new image, it should be able to tell if it's a cat or dog.

Cost:  $O(p)$

We call this approach, supervised wave-function optimization (**SWO**)



All one needs now is a labelled wave-function to match....



Unfortunately, getting the exact wave-function is difficult....

Instead...our goal will be to get a better wave-function.



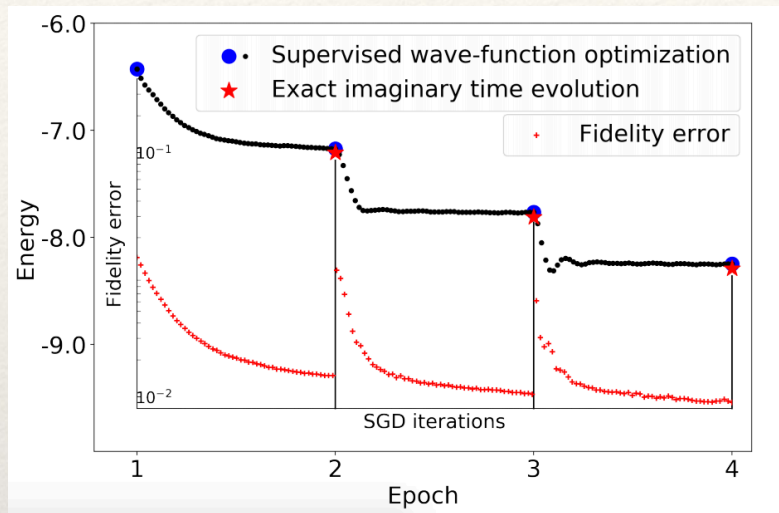
Better Wave-functions:

*Imaginary Time Evolution*  $(1 - \tau H) |\Psi_{NN}\rangle$

IT-SWO

Current best wave-function

This gives us an  $O(p)$  approach which avoids lots of other problems.  
 (no ill-defined inverse; no over-parameterization; no under sampling)

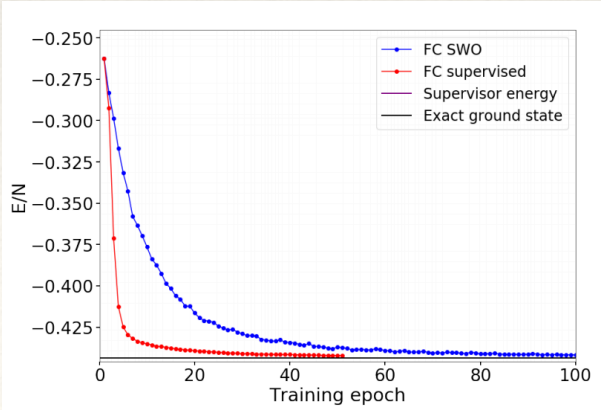


*Lanczos Steps*  $|\Psi_{NN}\rangle + \alpha H |\Psi_{NN}\rangle + \beta H^2 |\Psi_{NN}\rangle$

Lanczos-SWO

*Previously Optimized (but stuck) states...* MPS, simpler NN, etc.

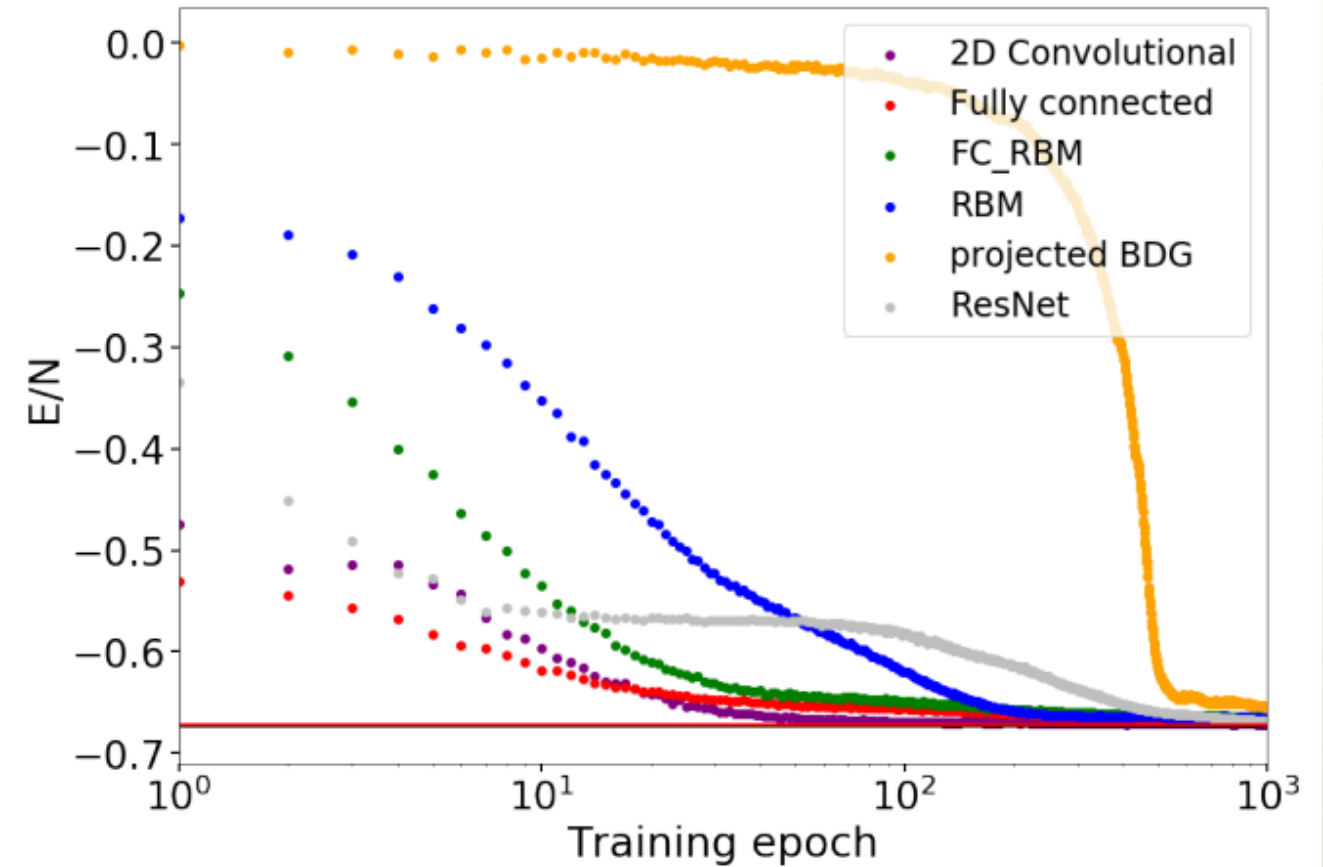
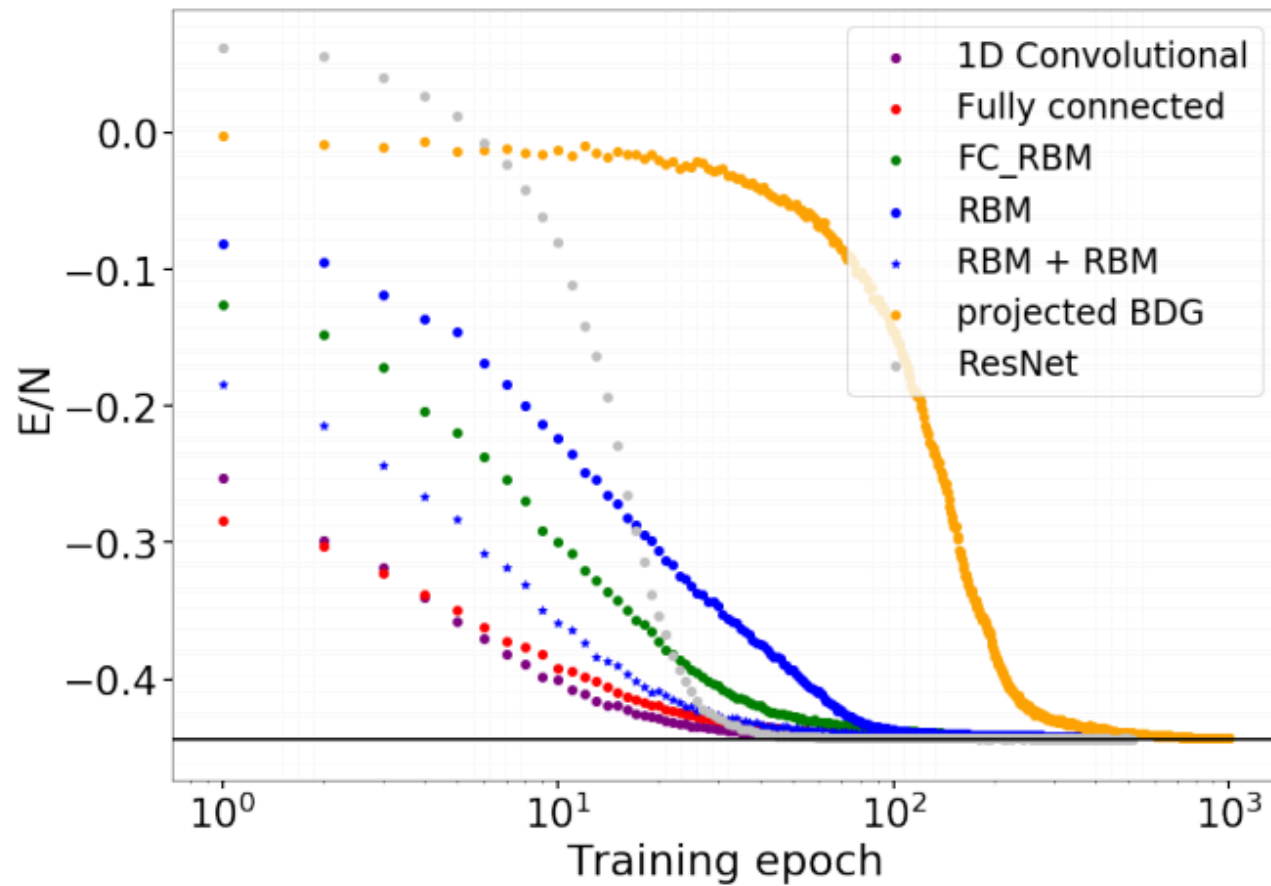
Better tunneling, ....



Matching-SWO



# Putting it all together...



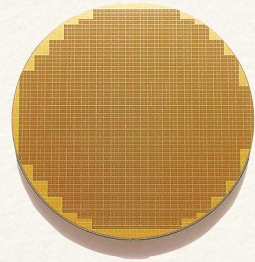
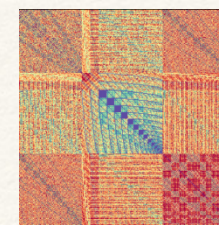
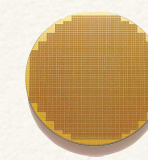
WF	Hidden Layers	Hidden Units	Params 1d - 2d
FCNN	2	80	9800 - 11744
RBM	0	80	3320 - 5264
FC-RBM	2	80	16280 - 18224
Conv1D	5	16 filters (size 5)	5280 -
Conv2D	5	16 filters (5 by 5)	- 26080
ResNet	1 + 2(2)	16 filters (size 5)	5280 - 25760
Multi-RBM	2	16 filters (5 by 5)	6640 - 10528
Multi-FCNN	2	80	19600 - 23488
P-BDG			1600 - 4096
MPS			4800 -

Can play with this yourself:

[github.com/ClarkResearchGroup/cgs-vmc/](https://github.com/ClarkResearchGroup/cgs-vmc/)



# The next steps (for the variational approach)....



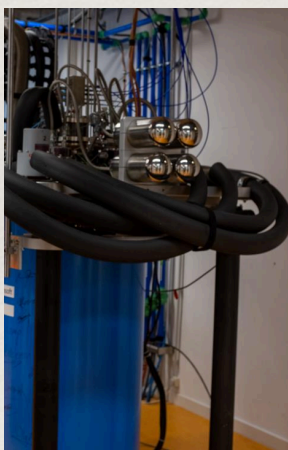
Fulfill the promise of the silicon age....

Computational graphs are a great structure, but how do you pick the correct graph?

*Current Solution:* Graph optimization by graduate student.

*Future Solution:* Graph optimization by computer.

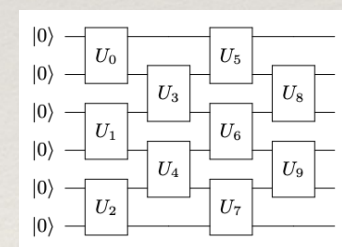
Matching SWO is a key piece of this puzzle. Once you have a new graph you want to quickly get its parameters to match the old graph. You can't afford to do this through Hamiltonian optimization.



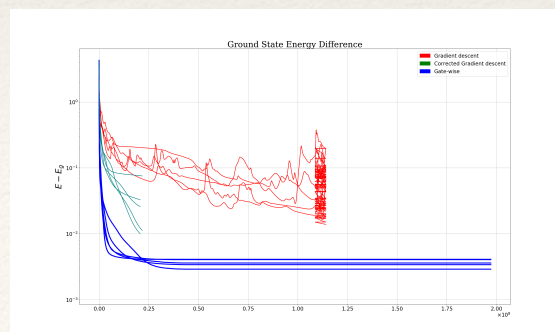
The quantum age...

Variational Quantum Eigensolvers

A standard quantum circuit is the moral equivalent of computational graph states.



But we are currently in ancient history as far as optimization is concerned...



*DMRG (or linear method) for quantum circuits....*



Cherkov, BKC, ....  
Cherkov, BKC, ....







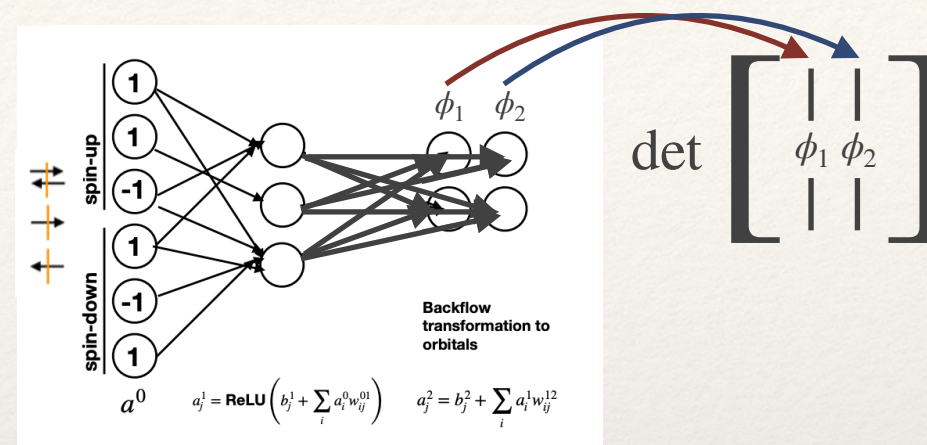
# Summary (the main pieces)...

There's a lot of hype in variational wave-functions + AI. Despite what you might have heard, they don't solve all problems.

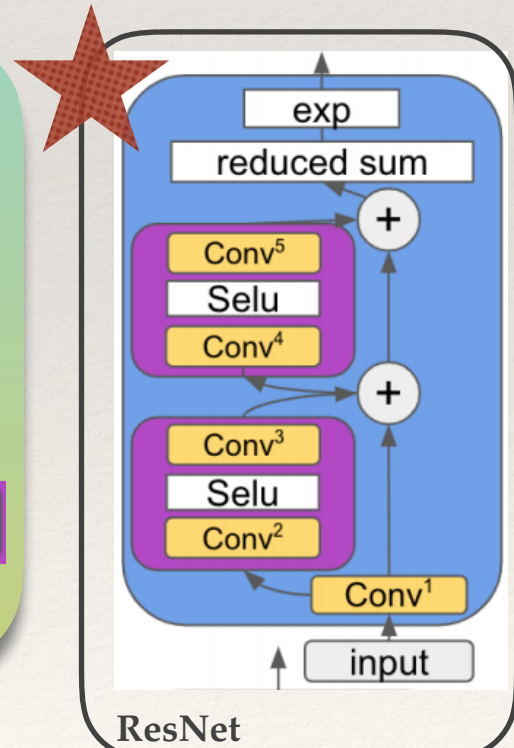
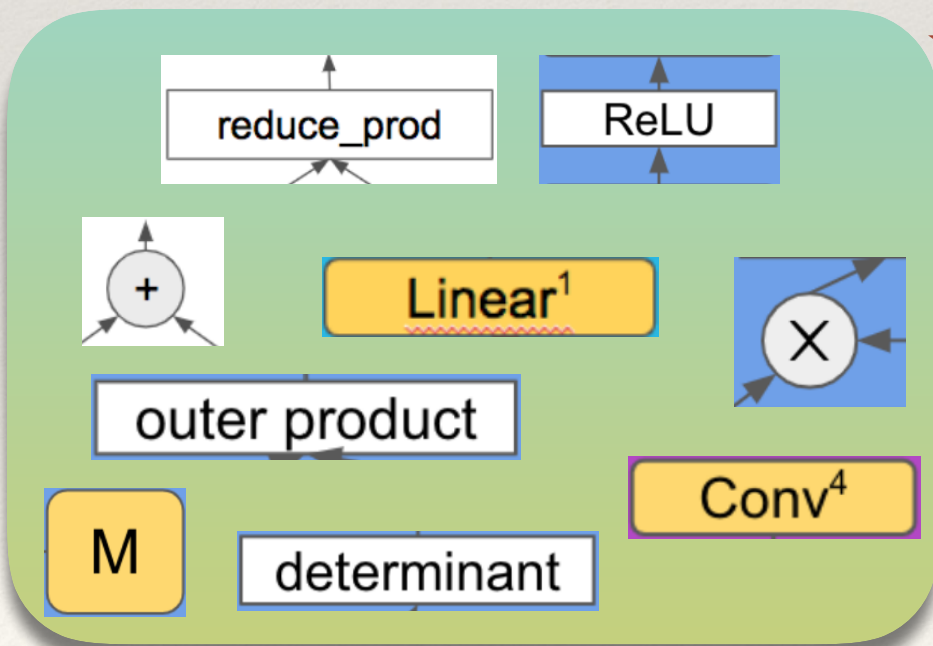
That said, smart choices are making real progress, optimization is improving, etc.

## Neural Net Backflow

*Fermions with Neural Net @ DMRG accuracy*

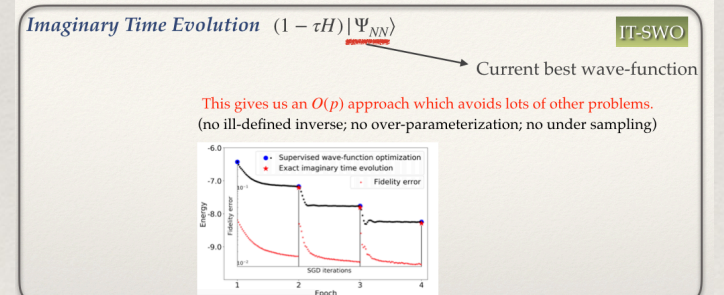


## Computational Graphs



Box of operations...

## SWO



Lanczos Steps  $|\Psi_{NN}\rangle + \alpha H |\Psi_{NN}\rangle + \beta H^2 |\Psi_{NN}\rangle$  Lanczos-SWO

Previously Optimized (but stuck) states... MPS, simpler NN, etc.

