Reinforcement Learning Configuration Interaction

Joshua Goings, Hang Hu, Chao Yang, Xiaosong Li

Theory Group Meeting

Monday, May 17, 2021

UNIVERSITY *of* WASHINGTON

eScience Institute

Motivation

Many molecules and materials with complex electronic structures and unique and useful catalytic and electromagnetic properties

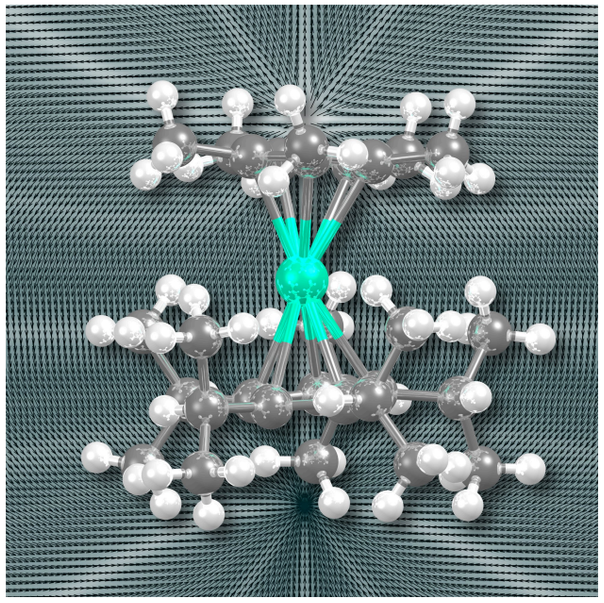

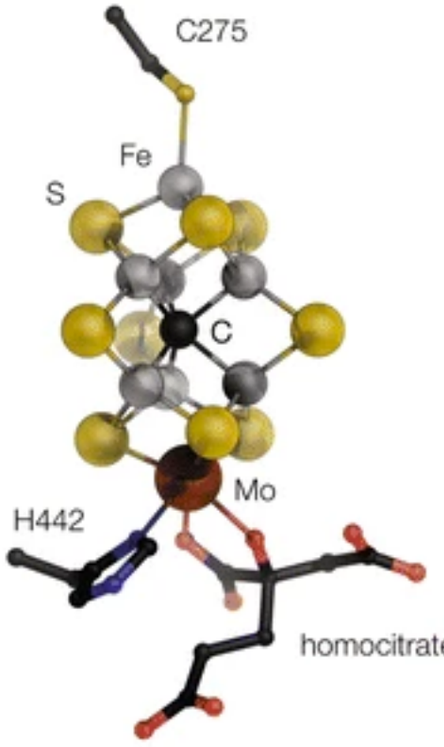

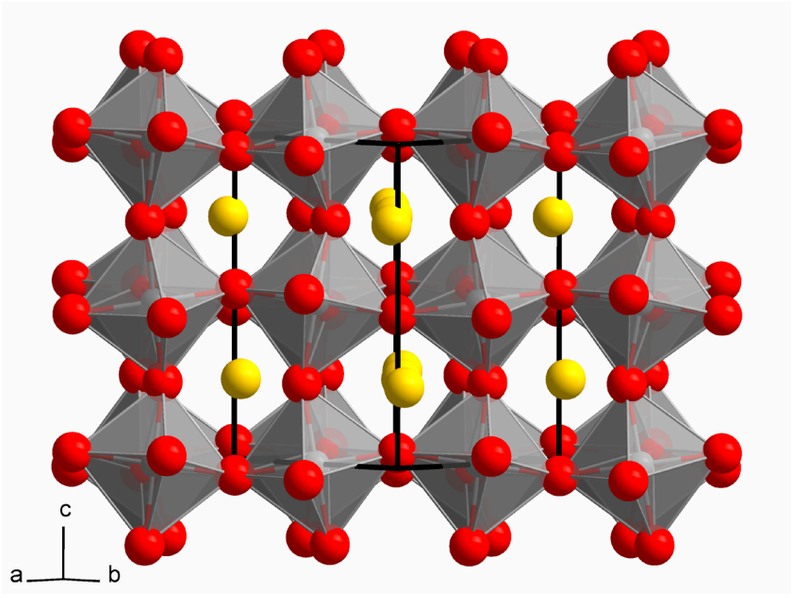

- Complex catalysis with multiple metal centers, e.g. FeMoCo in nitrogenase
- Single molecule magnets and magnetic materials
- Materials for photovoltaics and high-T_c superconductivity

Electronic structure is strongly correlated — conventional methods fail

Einsle, O. Nitrogenase FeMo cofactor: an atomic structure in three simple steps. *J Biol Inorg Chem* **19,** 737–745 (2014).
Richard Layfield, https://phys.org/news/2018-10-scientists-high-temperature-single-molecule-magnet.html
Solid State, CC BY-SA 3.0 <https://creativecommons.org/licenses/by-sa/3.0>, via Wikimedia Commons

Electron correlation

Electron correlation often leads to interesting / useful chemistry.

- **Electron correlation:** measure of how behavior of one electron is influenced by behavior of other electrons
- **Electron correlation energy:** difference in energy between Hartree-Fock and exact ground state energy
- Electron correlation is divided into dynamical and non-dynamical (static)
  - **dynamical correlation**: correlation of movement of electrons
  - **static correlation:** need more than one Slater determinant (i.e., Hartree-Fock is qualitatively incorrect)

Strong correlation: *lots* of Slater determinants required to get qualitatively correct electronic structure

Uncorrelated (independent)	Correlated
$P(AB) = P(A)P(B)$	$P(AB) \neq P(A)P(B)$

Correlated methods: Multiple determinants

- Total wave function is linear combination of determinants $|\Psi\rangle = \sum C_k |\Phi_k\rangle$
- Reference (Hartree-Fock) $|0\rangle = |\phi_1 \phi_2 \ldots \phi_N\rangle$
- Single substitutions / excitations $|S_{i\to a}\rangle = |\phi_1 \ldots \phi_{i-1} \phi_a \phi_{i+1} \ldots \phi_N\rangle$
- Double substitutions / excitations $|D_{ij\to ab}\rangle = |\phi_1 \ldots \phi_{i-1} \phi_a \ldots \phi_{j-1} \phi_b \ldots \phi_N\rangle$

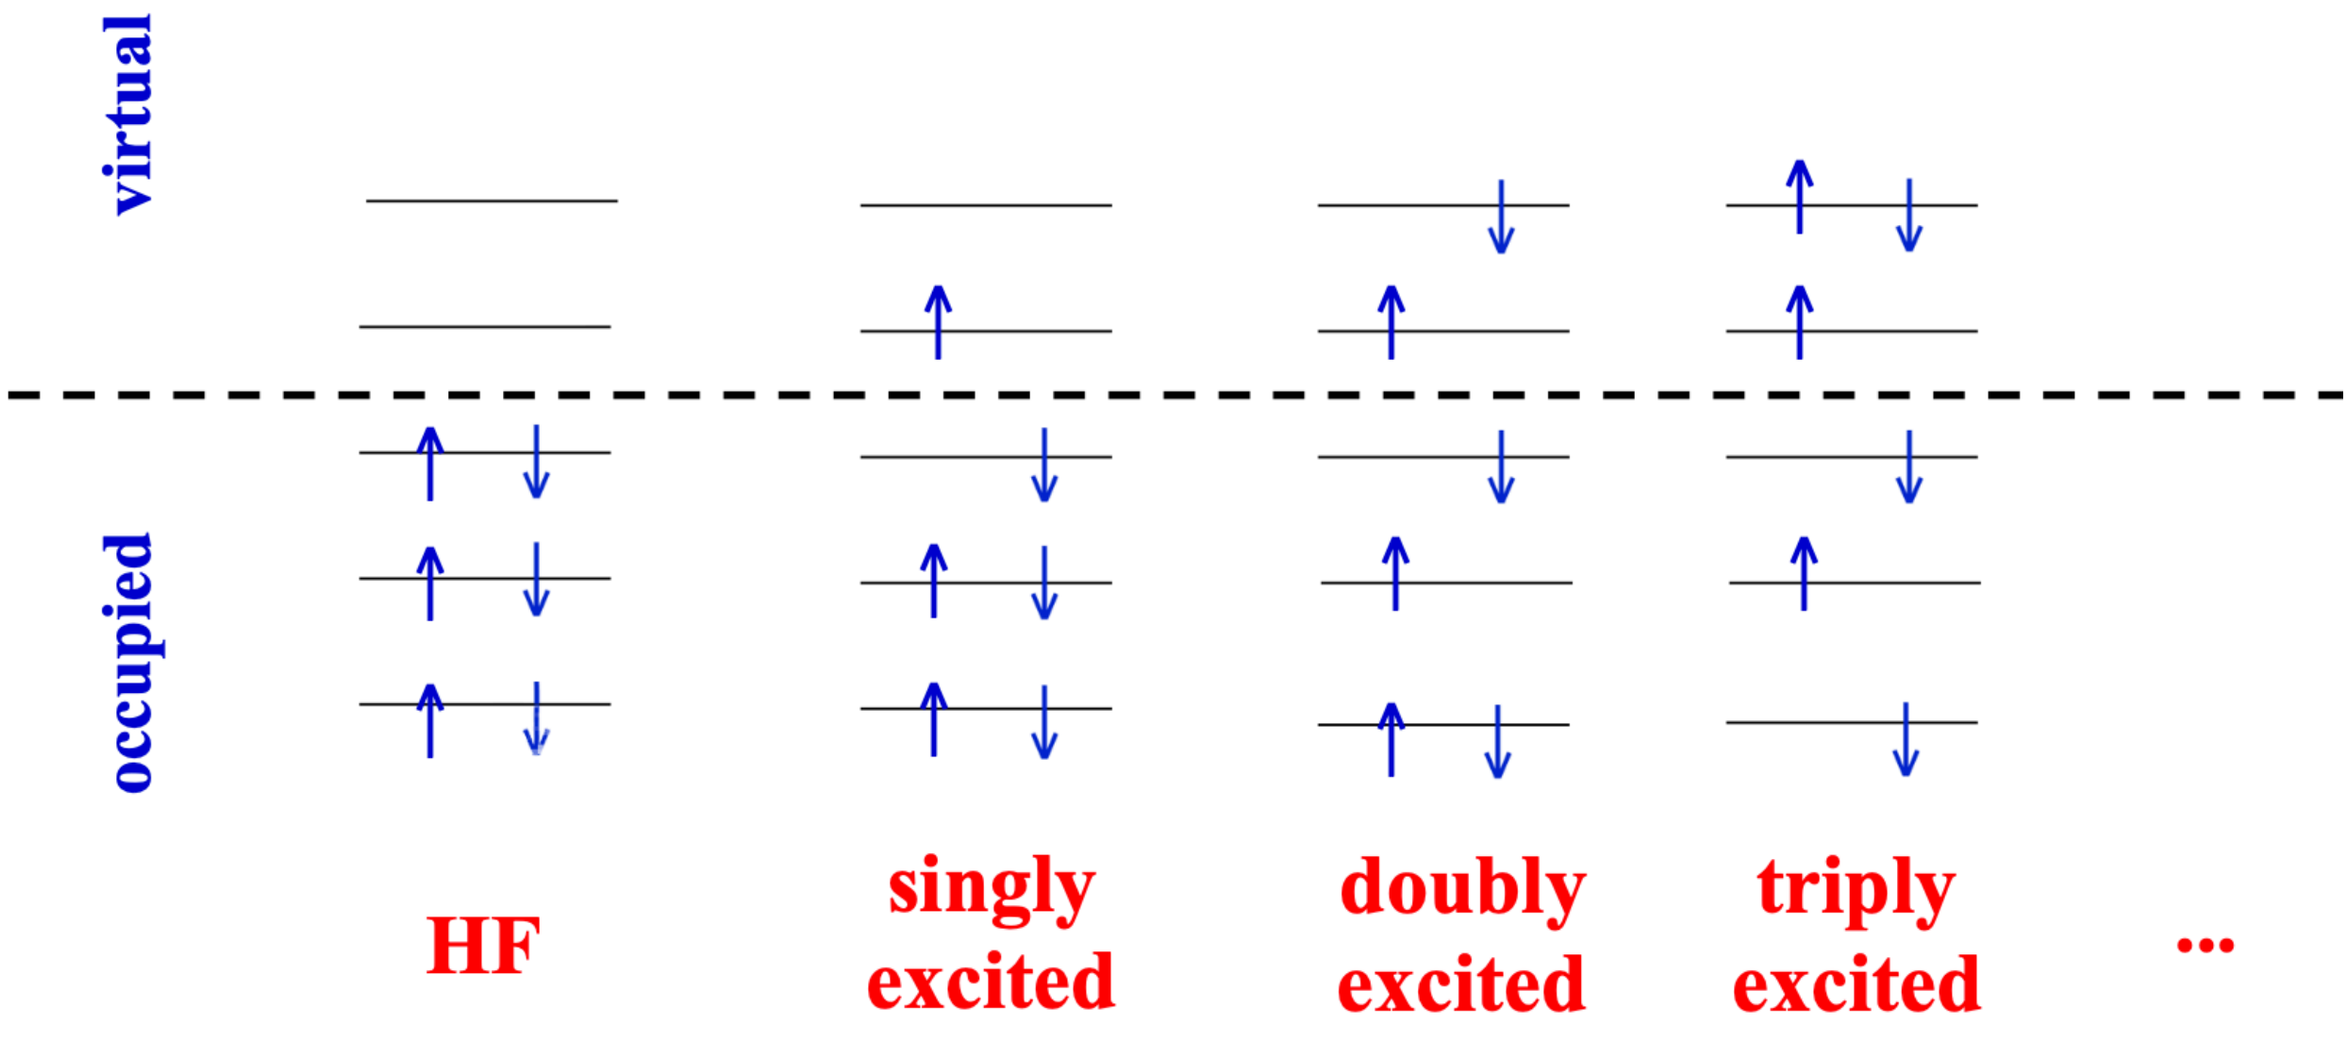

Configuration interaction (CI)

- Total wave function is linear combination of Slater determinants

$$|\Psi_{\mathrm{CI}}\rangle = |0\rangle + \sum_S C_S |S\rangle + \sum_D C_D |D\rangle + \cdots \qquad \hat{H}|\Psi_{\mathrm{CI}}\rangle = E|\Psi_{\mathrm{CI}}\rangle$$

- Determinants can be classified by excitation level
  - $|0\rangle$ is no excitations,
  - $|S\rangle$ is all single excitations,
  - $|D\rangle$ is all doubles,
  - $|T\rangle$ all triples… and so on.
- Determine coefficients C_i variationally
- Full CI: all possible substitutions. $|0\rangle, |S\rangle, |D\rangle, |T\rangle, \ldots$
  - **Exact** for a given atomic orbital basis set
  - **Very expensive**, scales as $N!$ (N orbitals choose k electrons)

Configuration interaction (CI): Matrix form

	$\lvert 0\rangle$	$\lvert S\rangle$	$\lvert D\rangle$	$\lvert T\rangle$
$\langle 0 \rvert$	E_{HF}	0	$\langle 0 \rvert \hat{H} \lvert D\rangle$	0
$\langle S \rvert$	0	$\langle S \rvert \hat{H} \lvert S\rangle$	$\langle S \rvert \hat{H} \lvert D\rangle$	$\langle S \rvert \hat{H} \lvert T\rangle$
$\langle D \rvert$	$\langle D \rvert \hat{H} \lvert 0\rangle$	$\langle D \rvert \hat{H} \lvert S\rangle$	$\langle D \rvert \hat{H} \lvert D\rangle$	$\langle D \rvert \hat{H} \lvert T\rangle$
$\langle T \rvert$	0	$\langle T \rvert \hat{H} \lvert S\rangle$	$\langle T \rvert \hat{H} \lvert D\rangle$	$\langle T \rvert \hat{H} \lvert T\rangle$

Hamiltonian with matrix elements

$$\langle X \rvert \hat{H} \rvert Y\rangle$$

CI is one big eigenproblem, lowest eigenvalue is **ground state energy**

$$\hat{H} \rvert \Psi_{\mathrm{CI}}\rangle = E \rvert \Psi_{\mathrm{CI}}\rangle$$

- Total wave function: $\lvert \Psi_{\mathrm{CI}}\rangle = \lvert 0\rangle + \sum_S C_S \lvert S\rangle + \sum_D C_D \lvert D\rangle + \cdots$

Truncated configuration interaction

	$\lvert 0\rangle$	$\lvert S\rangle$	$\lvert D\rangle$	$\lvert T\rangle$
$\langle 0 \rvert$	E_{HF}	0	$\langle 0 \rvert \hat{H} \rvert D\rangle$	0
$\langle S \rvert$	0	$\langle S \rvert \hat{H} \rvert S\rangle$	$\langle S \rvert \hat{H} \rvert D\rangle$	$\langle S \rvert \hat{H} \rvert T\rangle$
$\langle D \rvert$	$\langle D \rvert \hat{H} \rvert 0\rangle$	$\langle D \rvert \hat{H} \rvert S\rangle$	$\langle D \rvert \hat{H} \rvert D\rangle$	$\langle D \rvert \hat{H} \rvert T\rangle$
$\langle T \rvert$	0	$\langle T \rvert \hat{H} \rvert S\rangle$	$\langle T \rvert \hat{H} \rvert D\rangle$	$\langle T \rvert \hat{H} \rvert T\rangle$

Limit Hamiltonian matrix to certain excitation “classes”

Truncated configuration interaction

	$\|0\rangle$	$\|S\rangle$	$\|D\rangle$	$\|T\rangle$
$\langle 0\|$	E_{HF}	0	$\langle 0\|\hat{H}\|D\rangle$	0
$\langle S\|$	0	$\langle S\|\hat{H}\|S\rangle$	$\langle S\|\hat{H}\|D\rangle$	$\langle S\|\hat{H}\|T\rangle$
$\langle D\|$	$\langle D\|\hat{H}\|0\rangle$	$\langle D\|\hat{H}\|S\rangle$	$\langle D\|\hat{H}\|D\rangle$	$\langle D\|\hat{H}\|T\rangle$
$\langle T\|$	0	$\langle T\|\hat{H}\|S\rangle$	$\langle T\|\hat{H}\|D\rangle$	$\langle T\|\hat{H}\|T\rangle$

Truncated CI: Limit Hamiltonian matrix to certain excitation "classes"

e.g. CISD
(CI with Singles and Doubles)
Keep $|0\rangle, |S\rangle, |D\rangle$ submatrix

Minimum eigenvalue is an upper bound on the "true" minimum

Truncated configuration interaction

	$\|0\rangle$	$\|S\rangle$	$\|D\rangle$	$\|T\rangle$
$\langle 0\|$	E_{HF}	0	$\langle 0\|\hat{H}\|D\rangle$	0
$\langle S\|$	0	$\langle S\|\hat{H}\|S\rangle$	$\langle S\|\hat{H}\|D\rangle$	$\langle S\|\hat{H}\|T\rangle$
$\langle D\|$	$\langle D\|\hat{H}\|0\rangle$	$\langle D\|\hat{H}\|S\rangle$	$\langle D\|\hat{H}\|D\rangle$	$\langle D\|\hat{H}\|T\rangle$
$\langle T\|$	0	$\langle T\|\hat{H}\|S\rangle$	$\langle T\|\hat{H}\|D\rangle$	$\langle T\|\hat{H}\|T\rangle$

Truncated CI: Limit Hamiltonian matrix to certain excitation "classes"

e.g. CISD
(CI with Singles and Doubles)
Keep $|0\rangle, |S\rangle, |D\rangle$ submatrix

Minimum eigenvalue is an upper bound on the "true" minimum

Truncated CI is still **very large** and **often inadequate** for interesting molecules

Selected CI methods (sCI)

Can we do better than truncated CI?

	$\|0\rangle$	$\|S\rangle$	$\|D\rangle$	$\|T\rangle$
$\langle 0\|$	E_{HF}	0	dense	0
$\langle S\|$	0	dense	sparse	very sparse
$\langle D\|$	dense	sparse	sparse	extremely sparse
$\langle T\|$	0	very sparse	extremely sparse	extremely sparse

Selected CI methods (sCI)

Can we do better than truncated CI?

- Full CI matrix is very sparse
- ~95% of matrix elements do not contribute significantly to the total energy
- "Configurational deadwood"
- sCI methods truncate at **individual determinant** level

	$\|0\rangle$	$\|S\rangle$	$\|D\rangle$	$\|T\rangle$
$\langle 0\|$	E_{HF}	0	dense	0
$\langle S\|$	0	dense	sparse	very sparse
$\langle D\|$	dense	sparse	sparse	extremely sparse
$\langle T\|$	0	very sparse	extremely sparse	extremely sparse

Selected CI methods (sCI)

Can we do better than truncated CI?

- Full CI matrix is very sparse
- ~95% of matrix elements do not contribute significantly to the total energy
- "Configurational deadwood"
- sCI methods truncate at **individual determinant** level

What are the best determinants to keep? Not obvious *a priori*

	$\|0\rangle$	$\|S\rangle$	$\|D\rangle$	$\|T\rangle$
$\langle 0\|$	E_{HF}	0	dense	0
$\langle S\|$	0	dense	sparse	very sparse
$\langle D\|$	dense	sparse	sparse	extremely sparse
$\langle T\|$	0	very sparse	extremely sparse	extremely sparse

sCI: A trivial example

pretend CI matrix

(note: symmetric)

1	1	0
1	2	1
0	1	3

sCI: A trivial example

pretend CI matrix

(note: symmetric)

1	1	0
1	2	1
0	1	3

$$\lambda_{\min} = 2 - \sqrt{3} \approx 0.268$$

sCI: A trivial example

1	1	0
1	2	1
0	1	3

$$\lambda_{\min} \approx 0.268$$

sCI: A trivial example

Goal: Choose a 2×2 submatrix with $\lambda'_{\min}$ that best approximates $\lambda_{\min}$

1	1	0
1	2	1
0	1	3

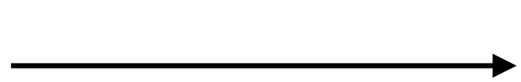

$\lambda_{\min} \approx 0.268$

sCI: A trivial example

Goal: Choose a 2×2 submatrix with $\lambda'_{\min}$ that best approximates $\lambda_{\min}$

1	1	0
1	2	1
0	1	3

$\lambda_{\min} \approx 0.268$

$s = \{1,2\}$

index 1 and 2

2	1
1	3

$\lambda'_{\min} \approx 1.382$

sCI: A trivial example

Goal: Choose a 2 × 2 submatrix with $\lambda'_{\min}$ that best approximates $\lambda_{\min}$

1	1	0
1	2	1
0	1	3

$\lambda_{\min} \approx 0.268$

$s = \{0,2\}$

index 0 and 2

1	0
0	3

$\lambda'_{\min} = 1$

sCI: A trivial example

Goal: Choose a 2×2 submatrix with $\lambda'_{\min}$ that best approximates $\lambda_{\min}$

1	1	0
1	2	1
0	1	3

$\lambda_{\min} \approx 0.268$

$s = \{0,1\}$

index 0 and 1

1	1
1	2

$\lambda'_{\min} \approx 0.382$

sCI: A trivial example

- Approximate minimum eigenvalue upper bounds "true" minimum eigenvalue
- Variational theorem
- $\lambda'_{\min} \geq \lambda_{\min}$
- Matrix must be symmetric

1	1	0
1	2	1
0	1	3

$\lambda_{\min} \approx 0.268$

$$s = \{0,1\} \rightarrow \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}, \quad \lambda_{\min} = \approx 0.382$$

$$s = \{0,2\} \rightarrow \begin{pmatrix} 1 & 0 \\ 0 & 3 \end{pmatrix}, \quad \lambda_{\min} = 1$$

$$s = \{1,2\} \rightarrow \begin{pmatrix} 2 & 1 \\ 1 & 3 \end{pmatrix}, \quad \lambda_{\min} \approx 1.382$$

- Challenge is finding the optimal submatrix
- For $k \times k$ submatrix of $N \times N$ matrix, there are N choose k combinations

$$\binom{N}{k} = \frac{N!}{k!(N-k)!}$$

Optimal selection of determinants in sCI

Given k, what is optimal* subset of determinants?

$$\min_{\|\Psi_{\mathrm{CI}}\|_0 \leq k} \frac{\langle \Psi_{\mathrm{CI}} | \hat{H} | \Psi_{\mathrm{CI}} \rangle}{\langle \Psi_{\mathrm{CI}} | \Psi_{\mathrm{CI}} \rangle}$$

Essentially a combinatorial optimization problem

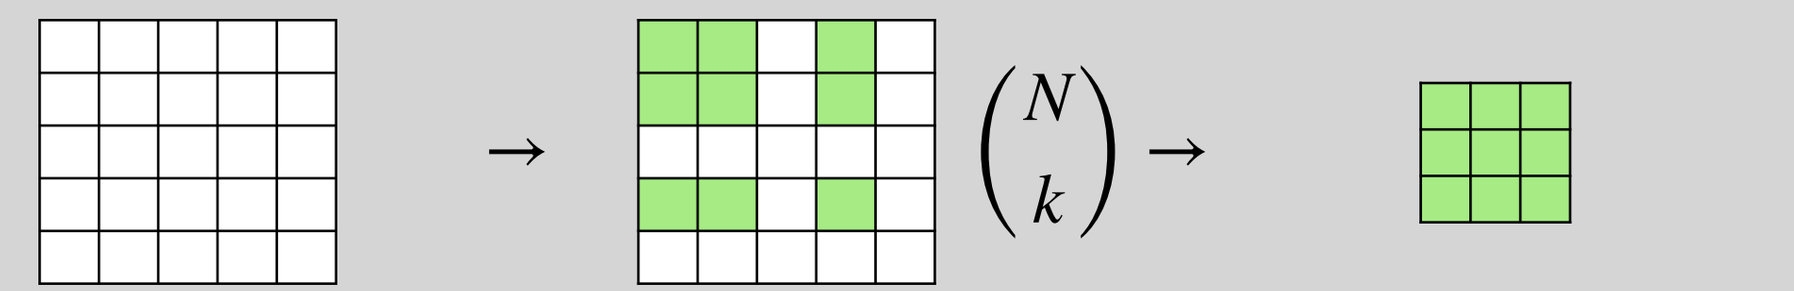

(5×5); λ_{min} $\qquad$ $k = 3$ $\qquad$ (3×3); $\lambda'_{\mathrm{min}} \approx \lambda_{\mathrm{min}}$ $(\lambda'_{\mathrm{min}} \geq \lambda_{\mathrm{min}})$

*Related question: Is k sufficient to yield reasonable approximation?

Conventional selected CI (sCI)

- Combinatorial problems are hard (generally no solution in polynomial time)
- Most selected CI algorithms iteratively build up the determinant space

Conventional selected CI (sCI)

- Combinatorial problems are hard (generally no solution in polynomial time)
- Most selected CI algorithms iteratively build up the determinant space

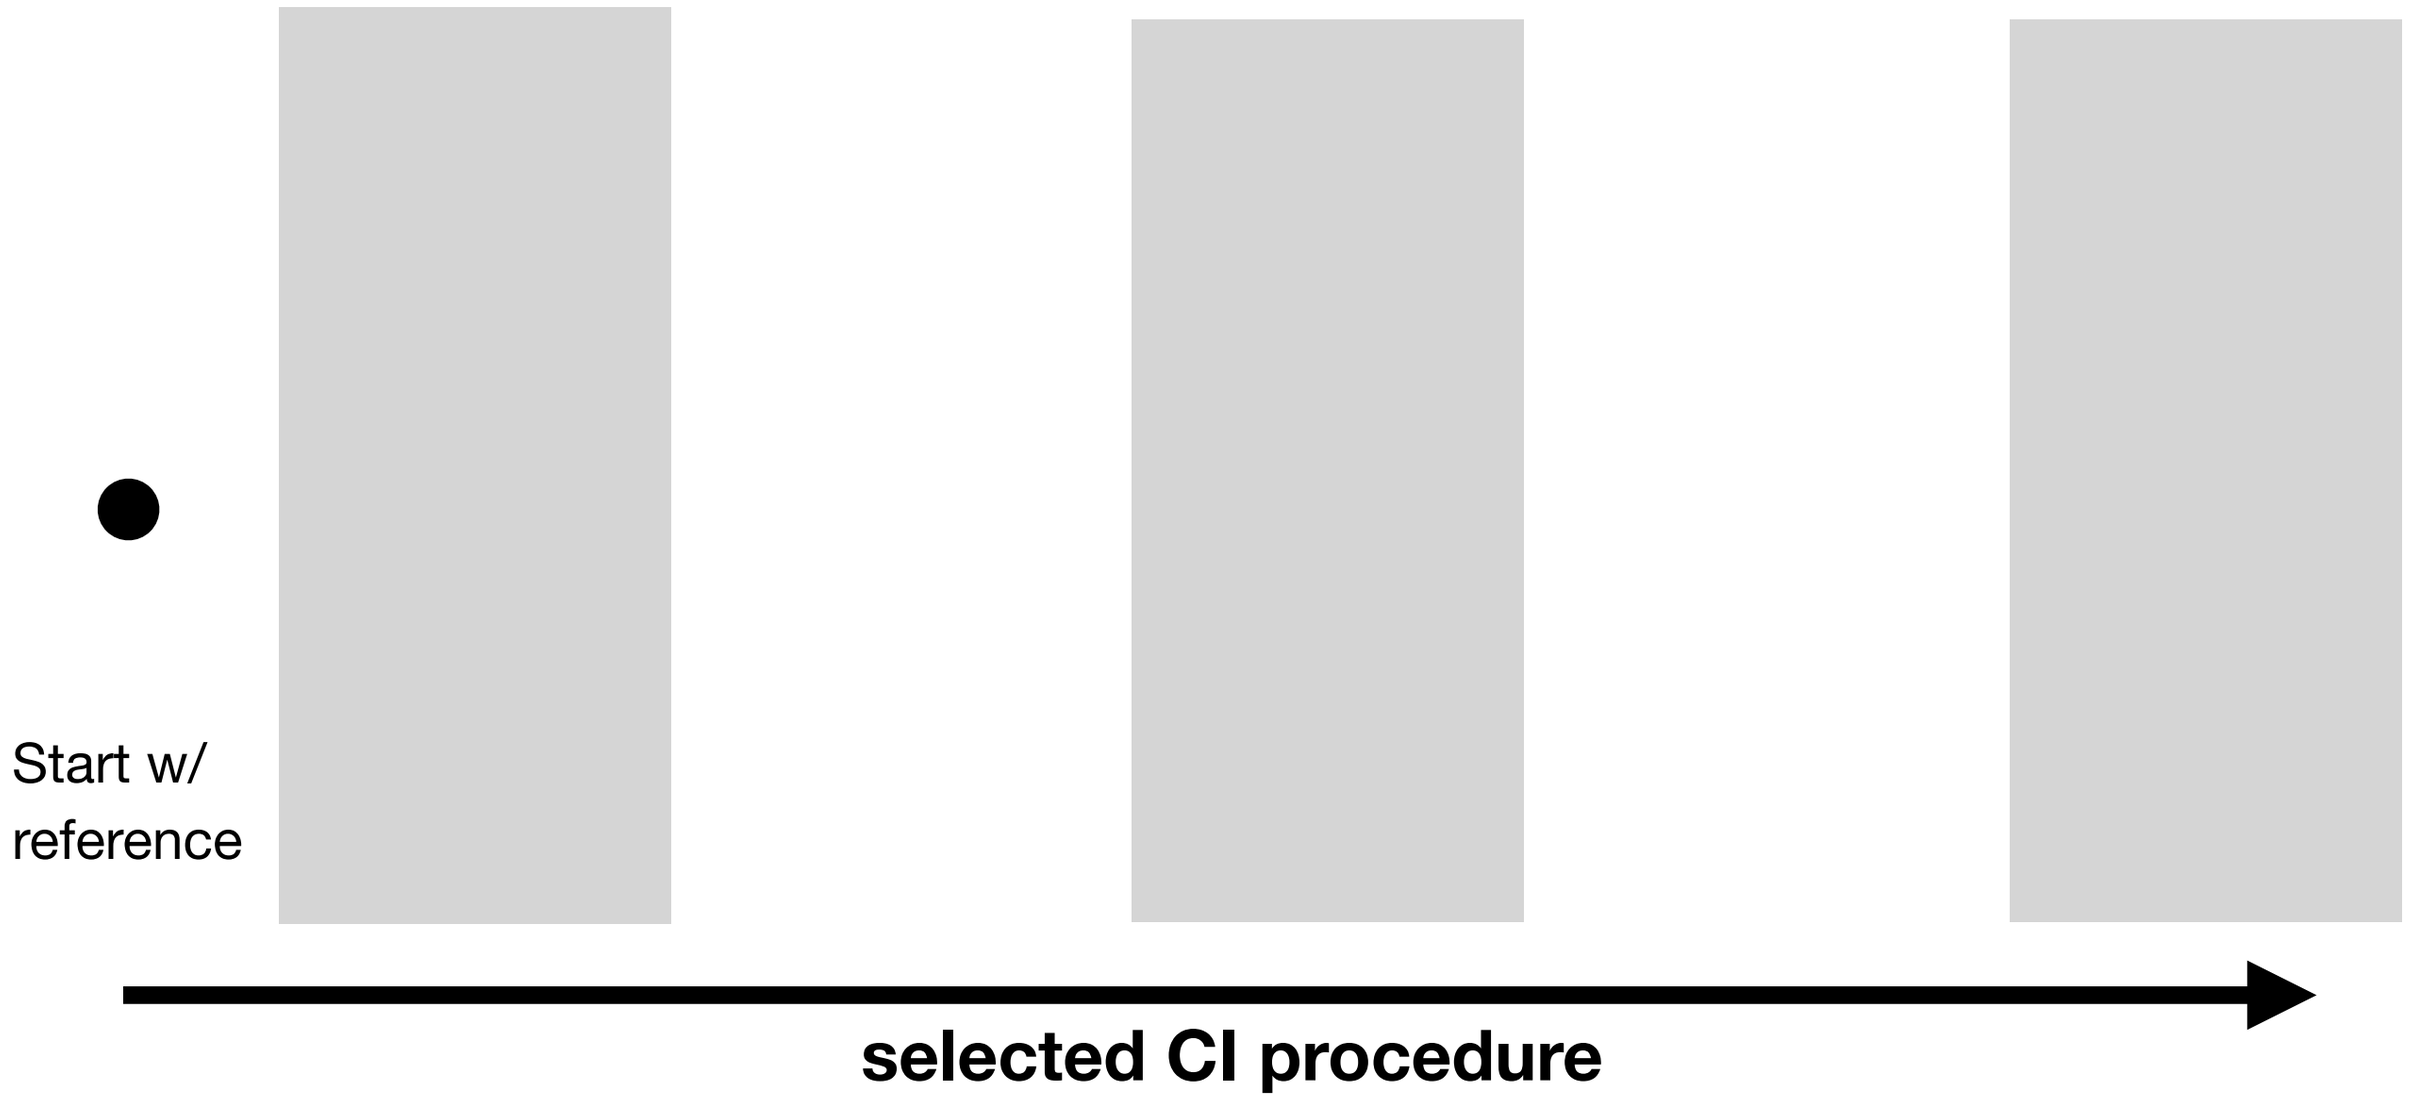

Conventional selected CI (sCI)

- Combinatorial problems are hard (generally no solution in polynomial time)
- Most selected CI algorithms iteratively build up the determinant space

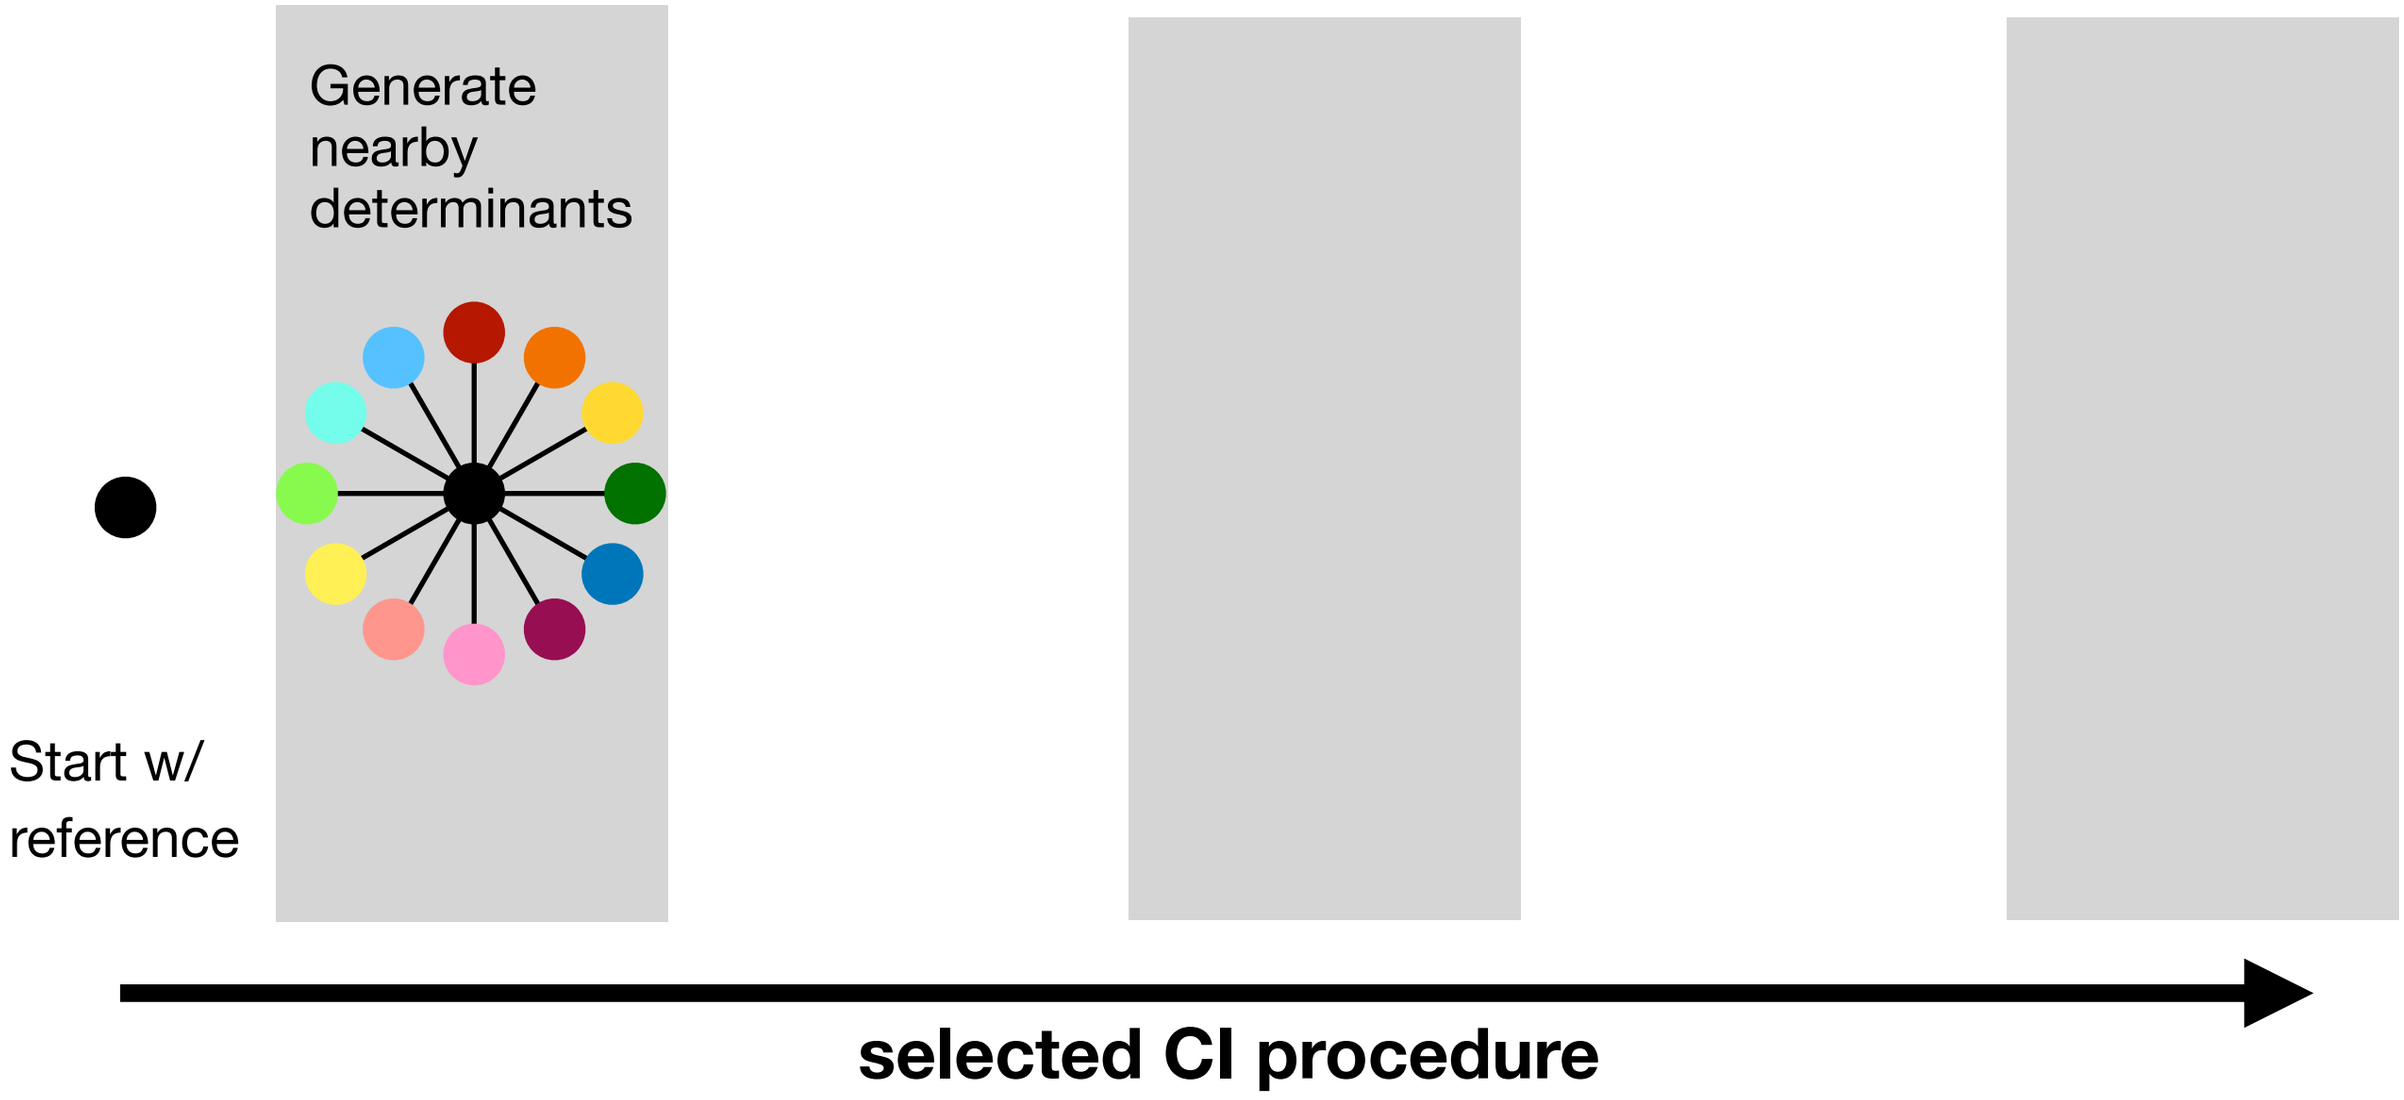

Conventional selected CI (sCI)

- Combinatorial problems are hard (generally no solution in polynomial time)
- Most selected CI algorithms iteratively build up the determinant space

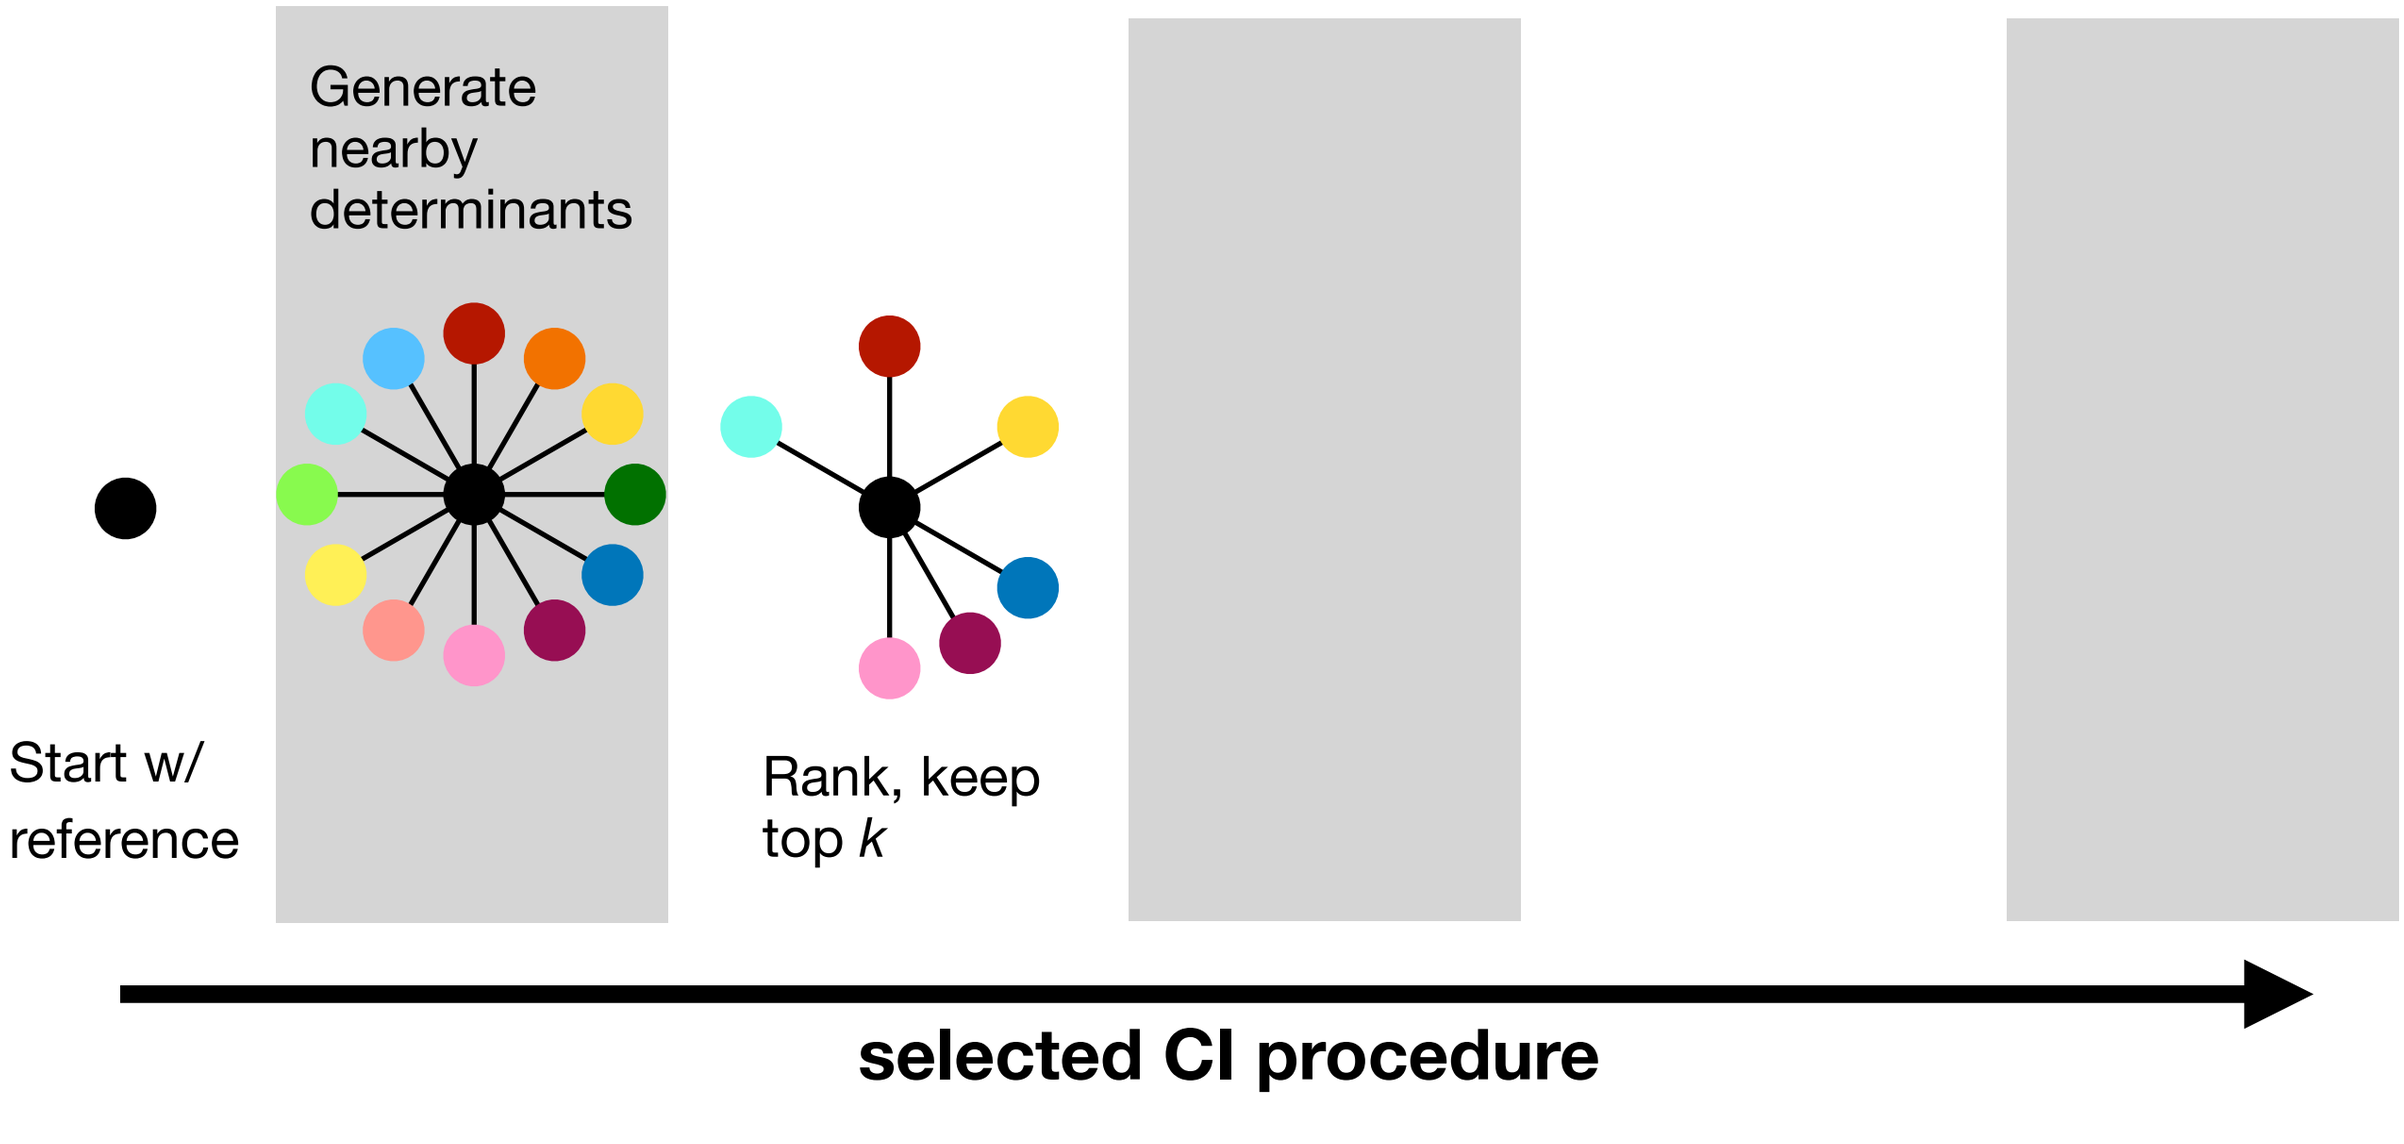

Conventional selected CI (sCI)

- Combinatorial problems are hard (generally no solution in polynomial time)
- Most selected CI algorithms iteratively build up the determinant space

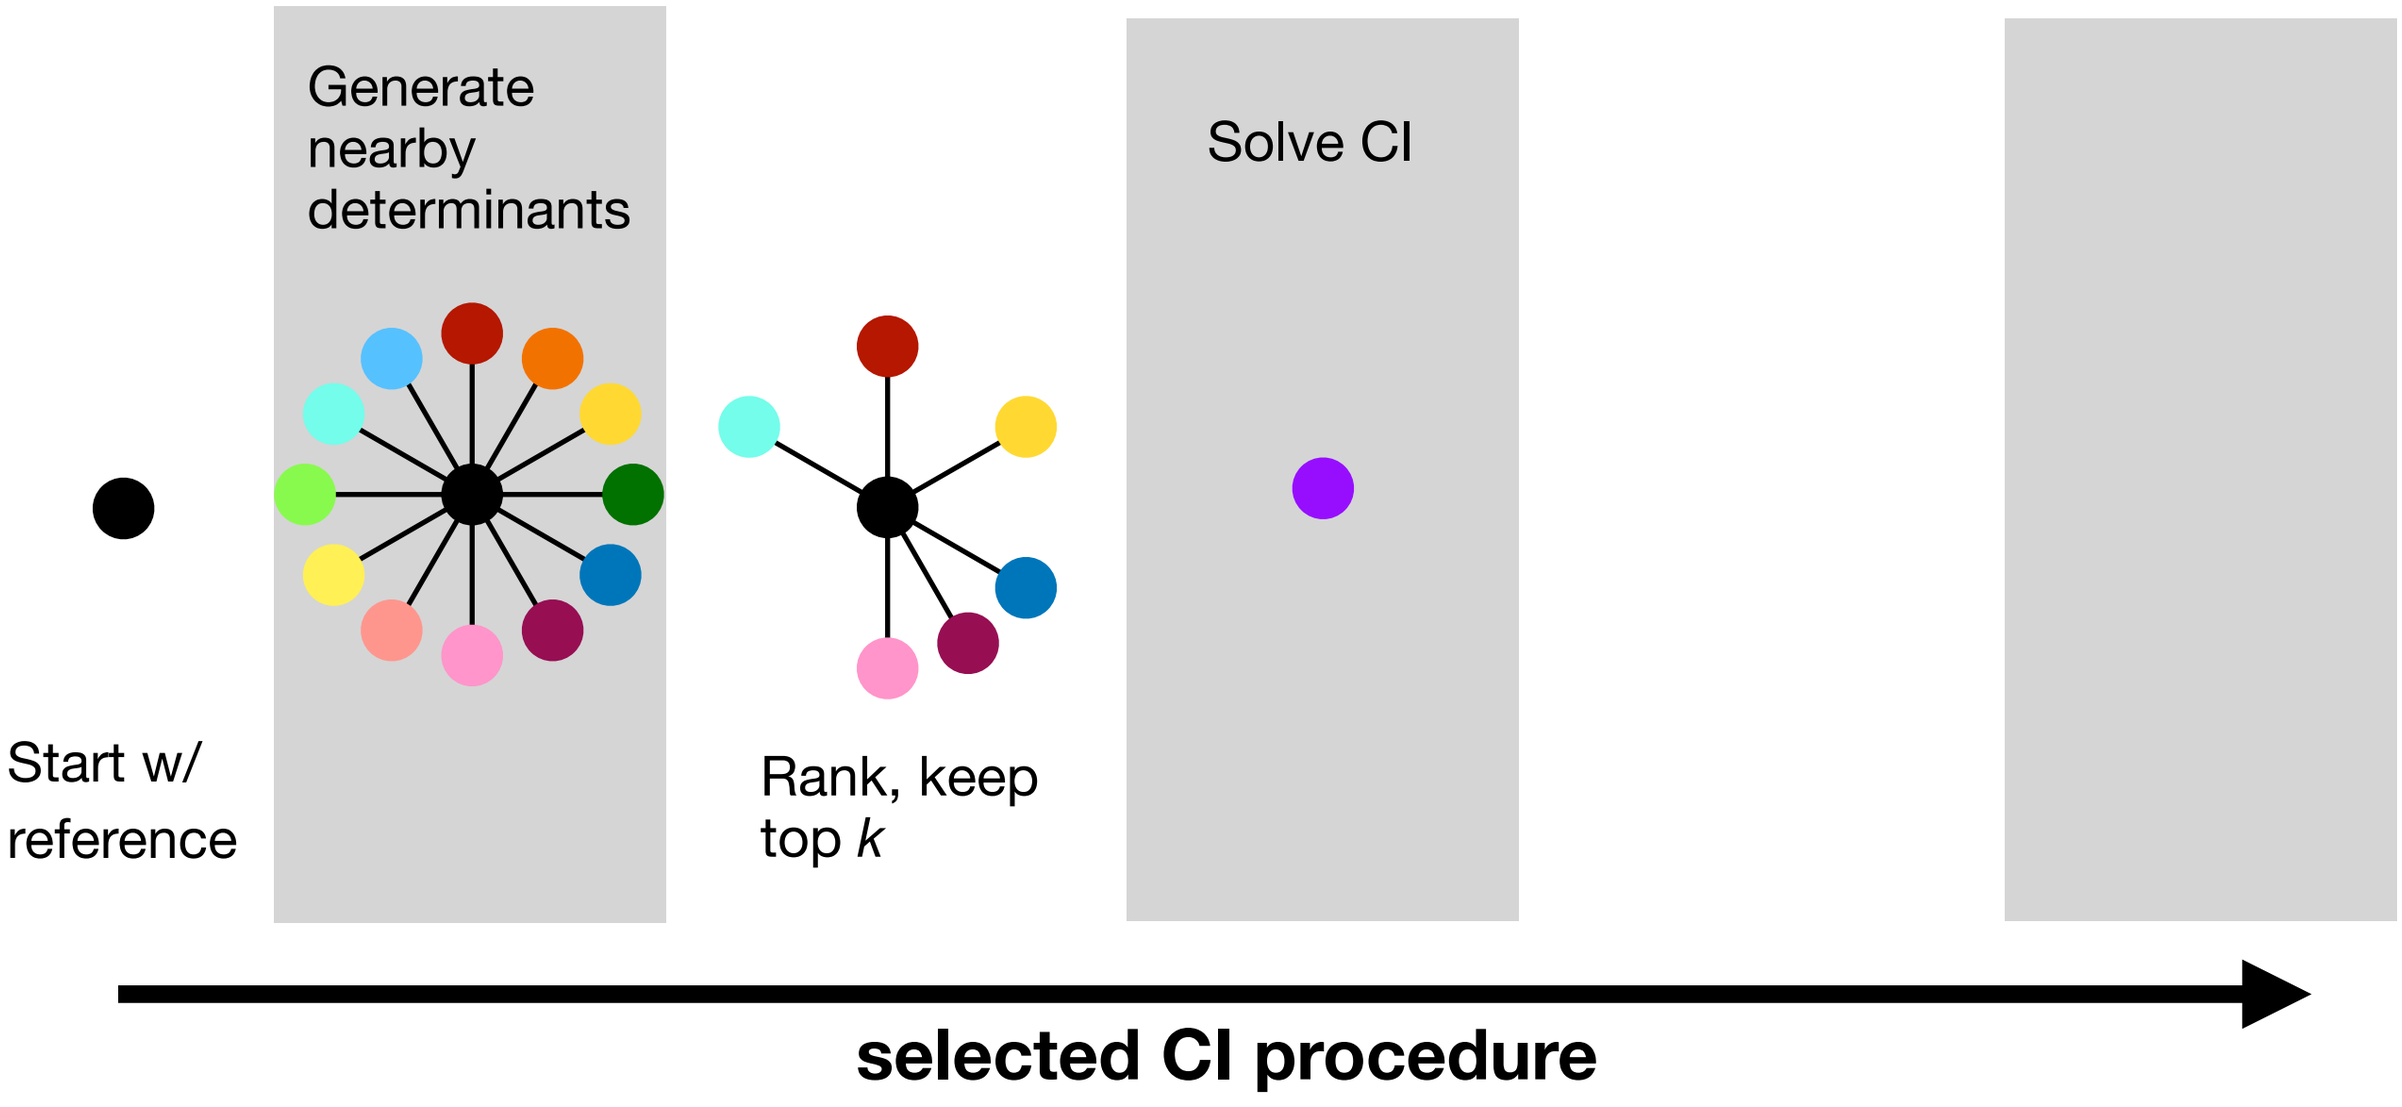

Conventional selected CI (sCI)

- Combinatorial problems are hard (generally no solution in polynomial time)
- Most selected CI algorithms iteratively build up the determinant space

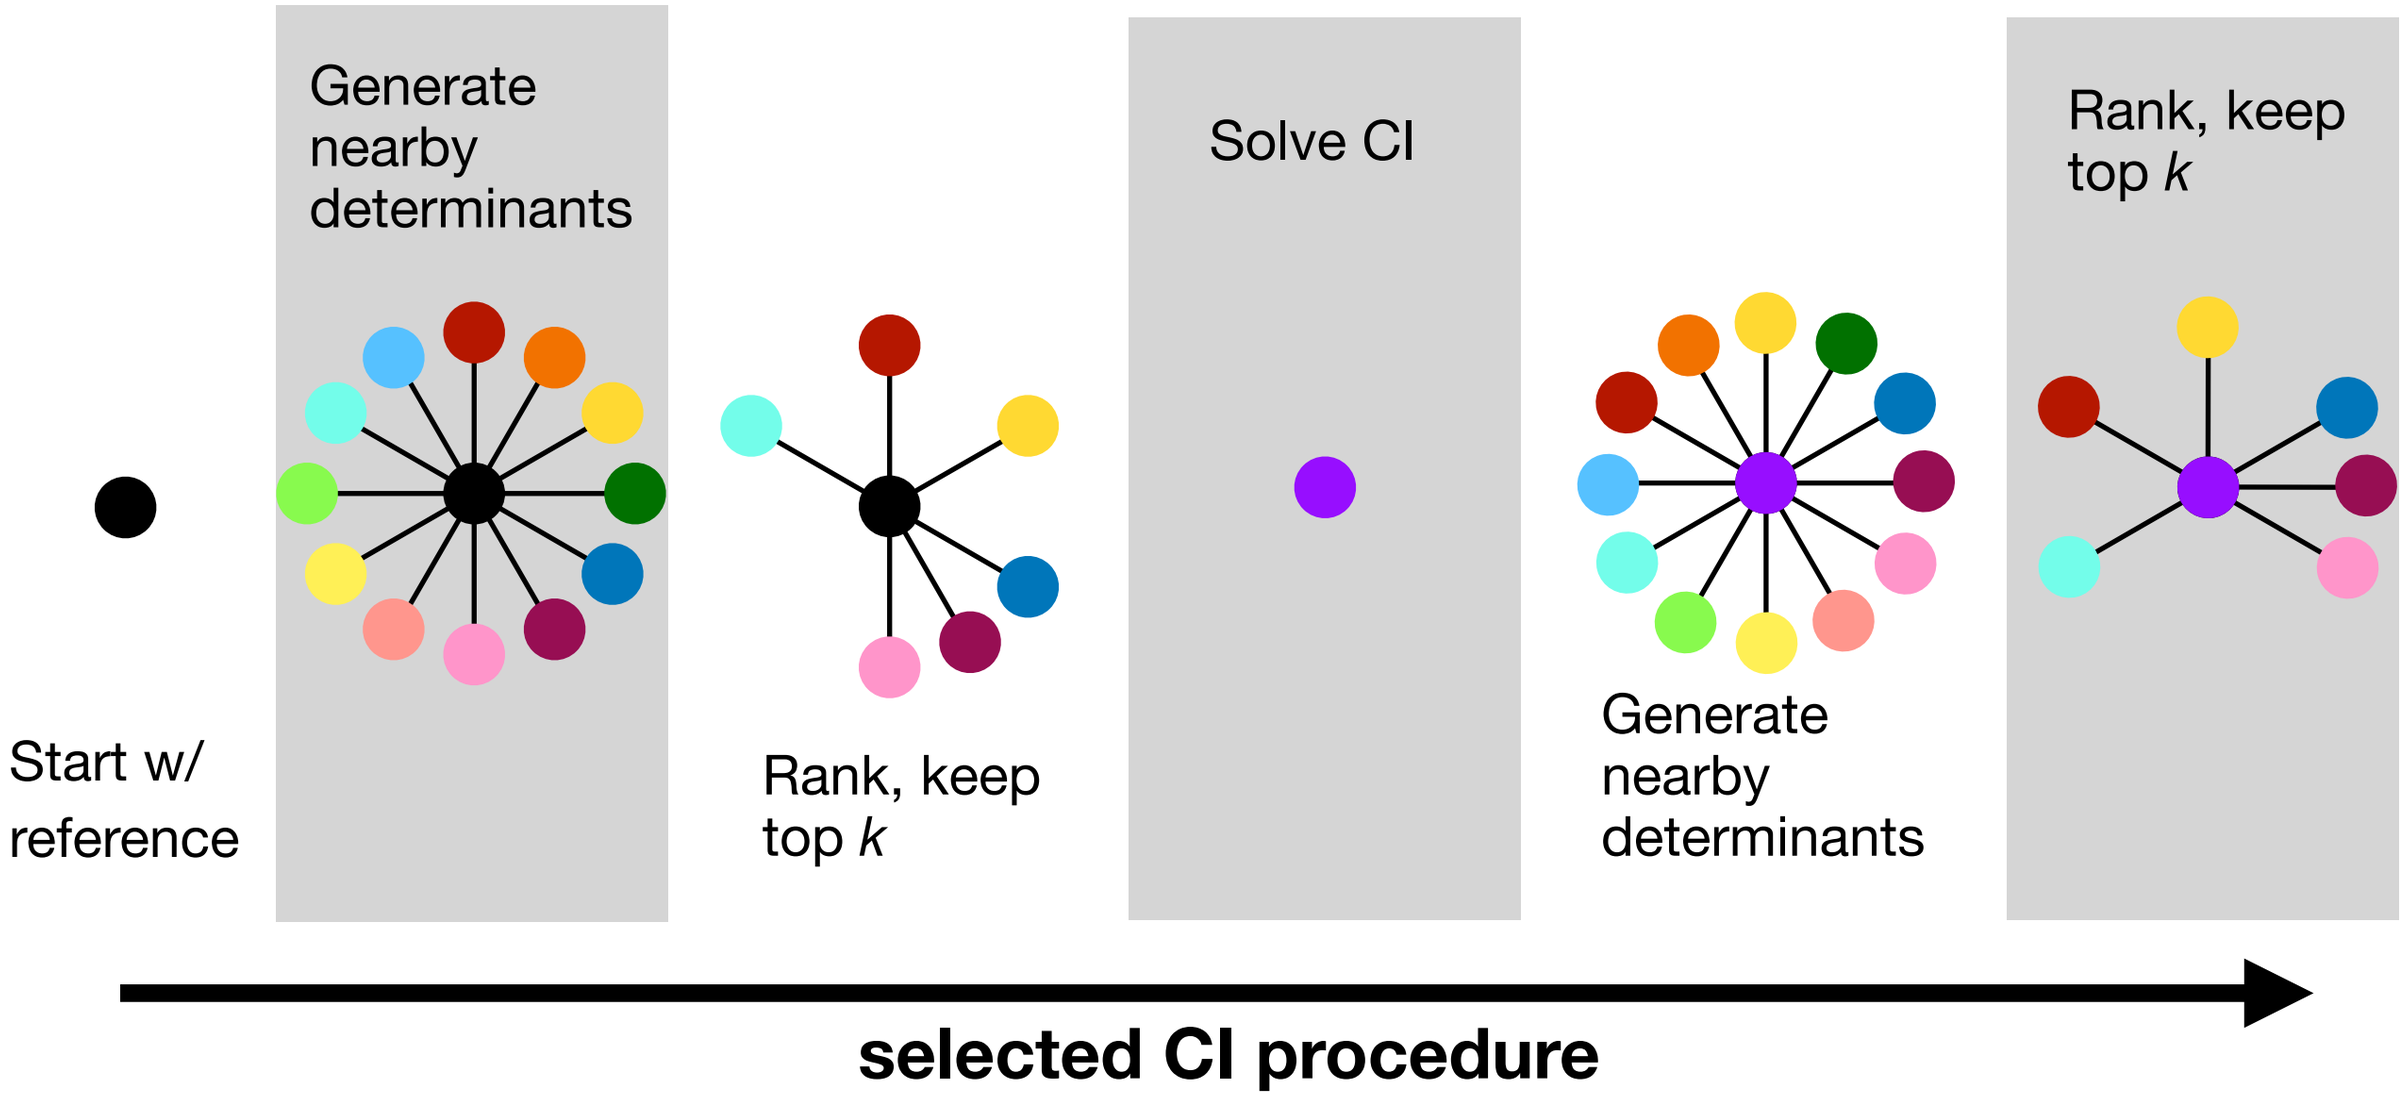

Conventional selected CI (sCI)

- Combinatorial problems are hard (generally no solution in polynomial time)
- Most selected CI algorithms iteratively build up the determinant space

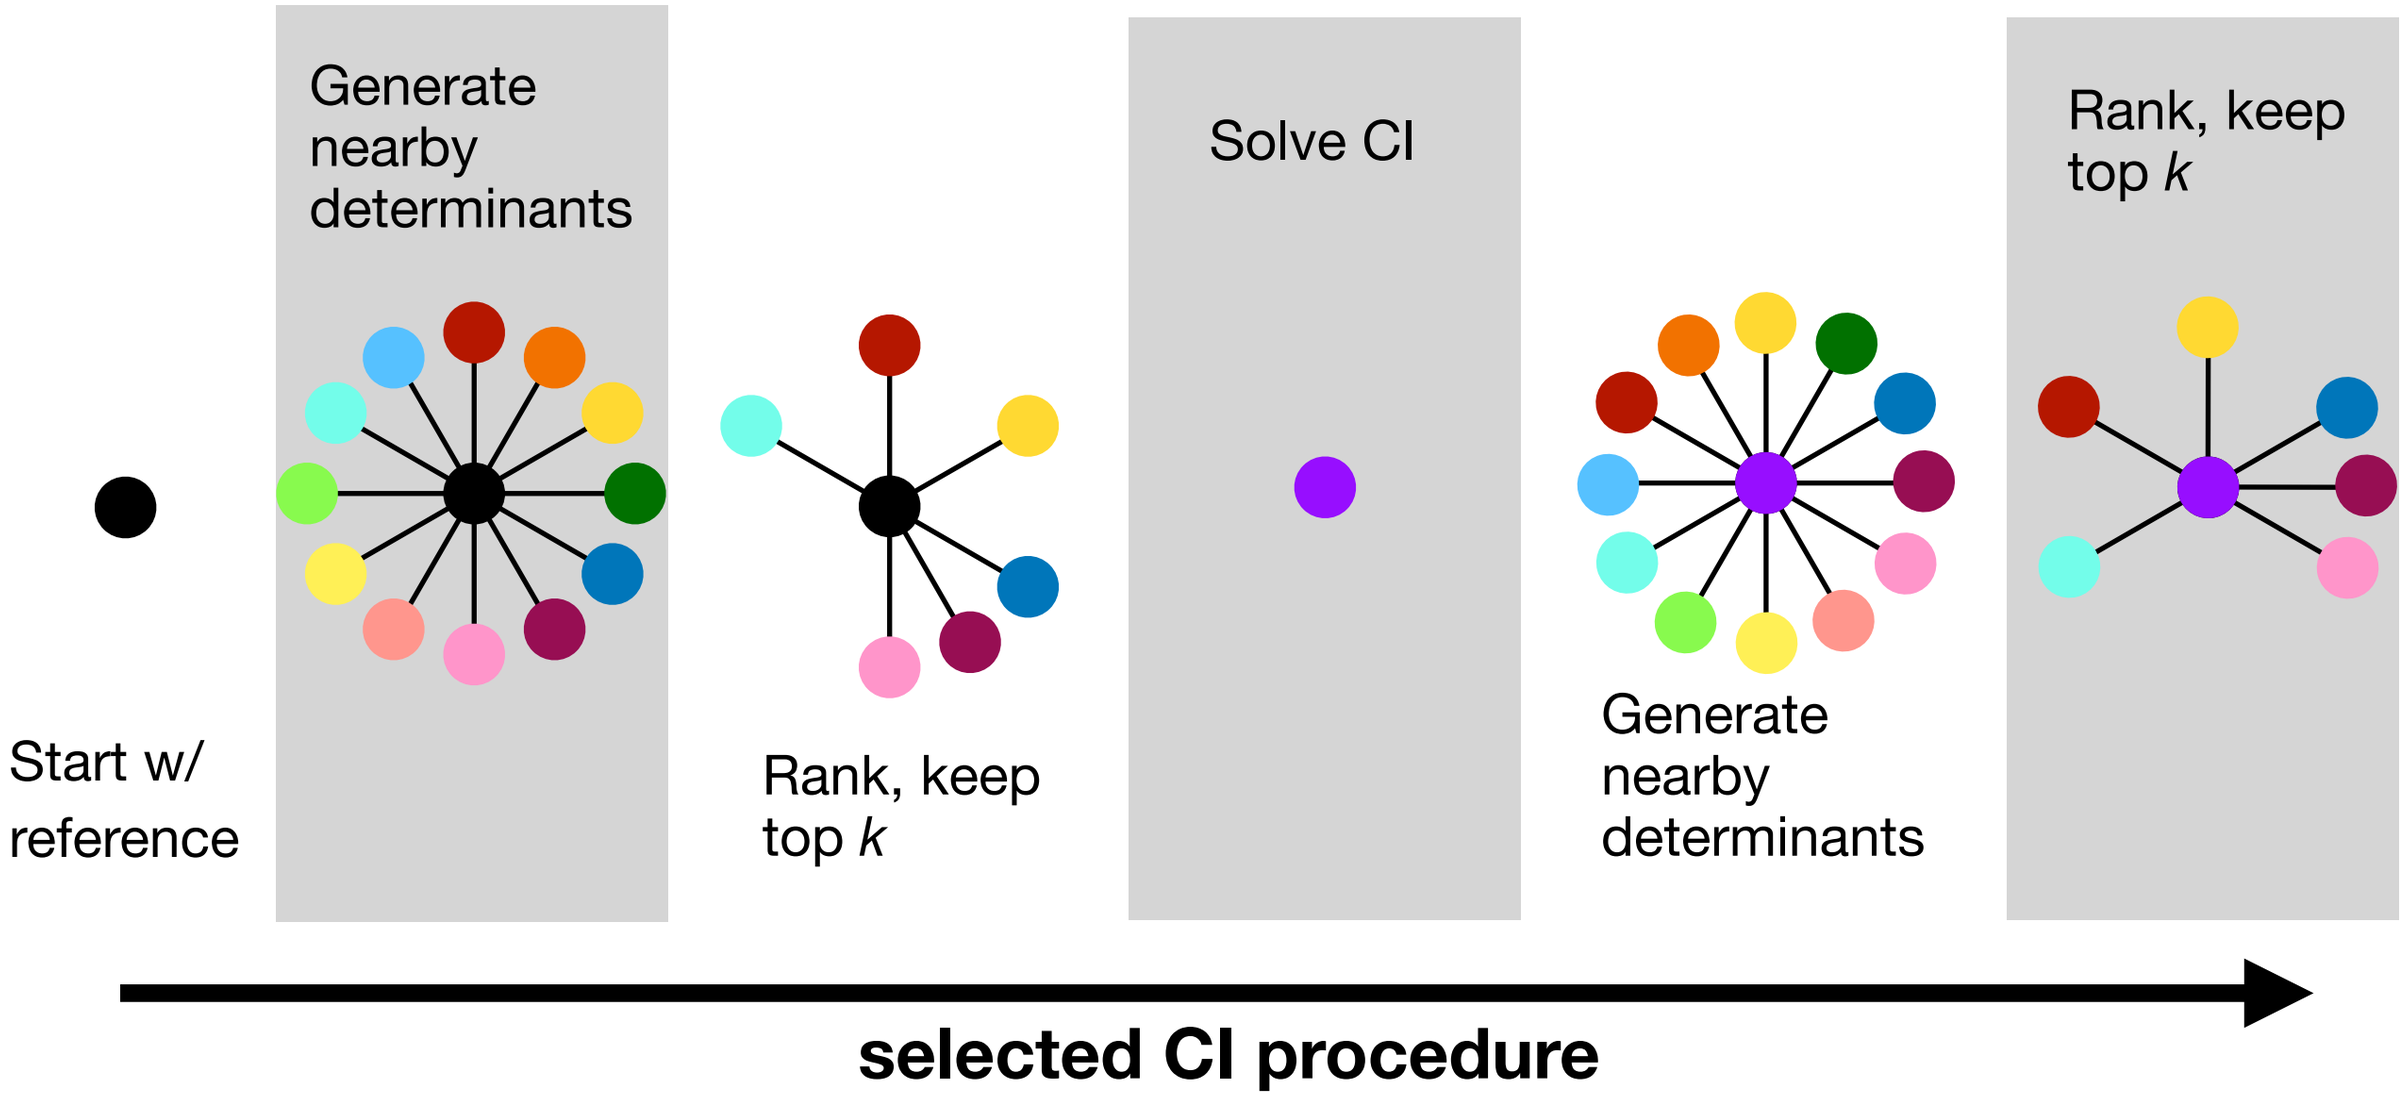

- Ranking procedure generally based on **heuristics** (e.g. perturbative est.)
- Other heuristics to explore relevant determinant / CI matrix structure

Improving on selected CI

- Selected CI methods rely on heuristics to evaluate the quality of determinants to add to the subspace
  - Deterministic (e.g., adaptively selected CI, SVD-CI) variants
  - (Semi-)Stochastic (e.g. Full CI Quantum Monte Carlo) variants
- Because they do not know the underlying structure of the CI Hamiltonian, the solutions found may be far from optimal

Improving on selected CI

- Selected CI methods rely on heuristics to evaluate the quality of determinants to add to the subspace
  - Deterministic (e.g., adaptively selected CI, SVD-CI) variants
  - (Semi-)Stochastic (e.g. Full CI Quantum Monte Carlo) variants
- Because they do not know the underlying structure of the CI Hamiltonian, the solutions found may be far from optimal

Can we use machine learning to identify
important determinants for
selected configuration interaction methods?

Machine learning for sCI

Goals:

#1: Infer the underlying structure to a CI Hamiltonian

#2: Use structure to yield best approximate eigenvalue

Machine learning for sCI

Goals:

#1: Infer the underlying structure to a CI Hamiltonian

#2: Use structure to yield best approximate eigenvalue

Supervised Learning

- Learn from labelled data
- Results depend on training data (efficiency + bias)
- If data sufficient, should work for #1 and #2

Machine learning for sCI

Goals:

#1: Infer the underlying structure to a CI Hamiltonian

#2: Use structure to yield best approximate eigenvalue

Supervised Learning

- Learn from labelled data
- Results depend on training data (efficiency + bias)
- If data sufficient, should work for #1 and #2

Unsupervised Learning

- Learn from unlabelled data
- Patterns and data groupings
- Could work for #1, unclear for #2

Machine learning for sCI

Goals:

#1: Infer the underlying structure to a CI Hamiltonian

#2: Use structure to yield best approximate eigenvalue

Supervised Learning	Unsupervised Learning	Reinforcement Learning
• Learn from labelled data • Results depend on training data (efficiency + bias) • If data sufficient, should work for #1 and #2	• Learn from unlabelled data • Patterns and data groupings • Could work for #1, unclear for #2	• Data generated and incorporated on-the-fly • Goal-oriented (#2) • Exploration informed by learned structure (#1)

Machine learning for sCI

Supervised Learning

Coe, J. P. *JCTC.* **2018**, 5739–5749.

Coe, J. P. *JCTC.* **2019**, 6179–6189.

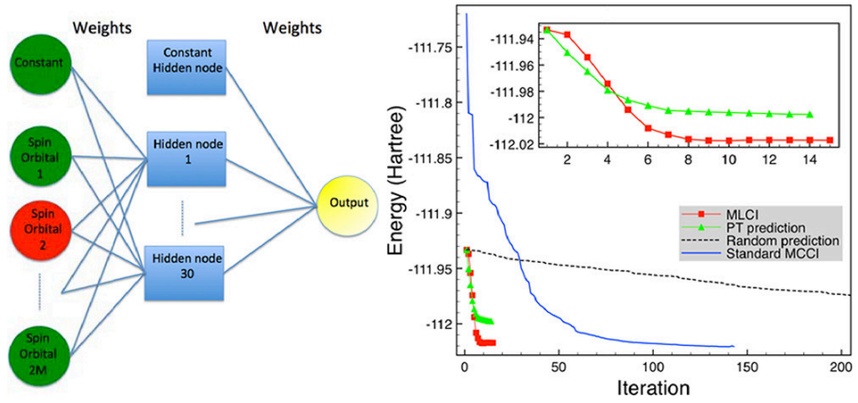

Learns "importance" of determinants from magnitude of coefficients $|C_i|$ during CI subspace diagonalization. Requires user-defined labels $|C_i| > C_{\min}$.

Unsupervised Learning

Taylor, P. R. *JCP.* **2013**, 074113.

Do an SVD on the CI problem, but cost of SVD outweighs benefits;
Large vectors apparently still needed to get target accuracy.

Reinforcement Learning

Goings, J.J., *et al. ChemRxiv preprint*. **2021**,10.26434/chemrxiv.14342234

This work: learn "importance" of determinants from changes in energy during CI subspace diagonalization. No user-defined labels.

Reinforcement learning

- Agent learns what actions to take in order to maximize a reward signal.
- **Trial-and-error search**: Agent is not told what actions to take, but learns by interacting with its environment by trial and error.
- **Delayed reward**: Actions may not just affect immediate reward, but also influence rewards next time, and so on into the future.

Examples:

Recycling robot: Look for more trash, or go back to recharge?

- What is current battery level?
- How quickly/easily has it been able to find the charger in the past?

Playing chess: what is the next move?

- What is the current state of the board?
- What opportunities open up? What vulnerabilities are created?

Sutton & Barto (2018)

Elements of reinforcement learning

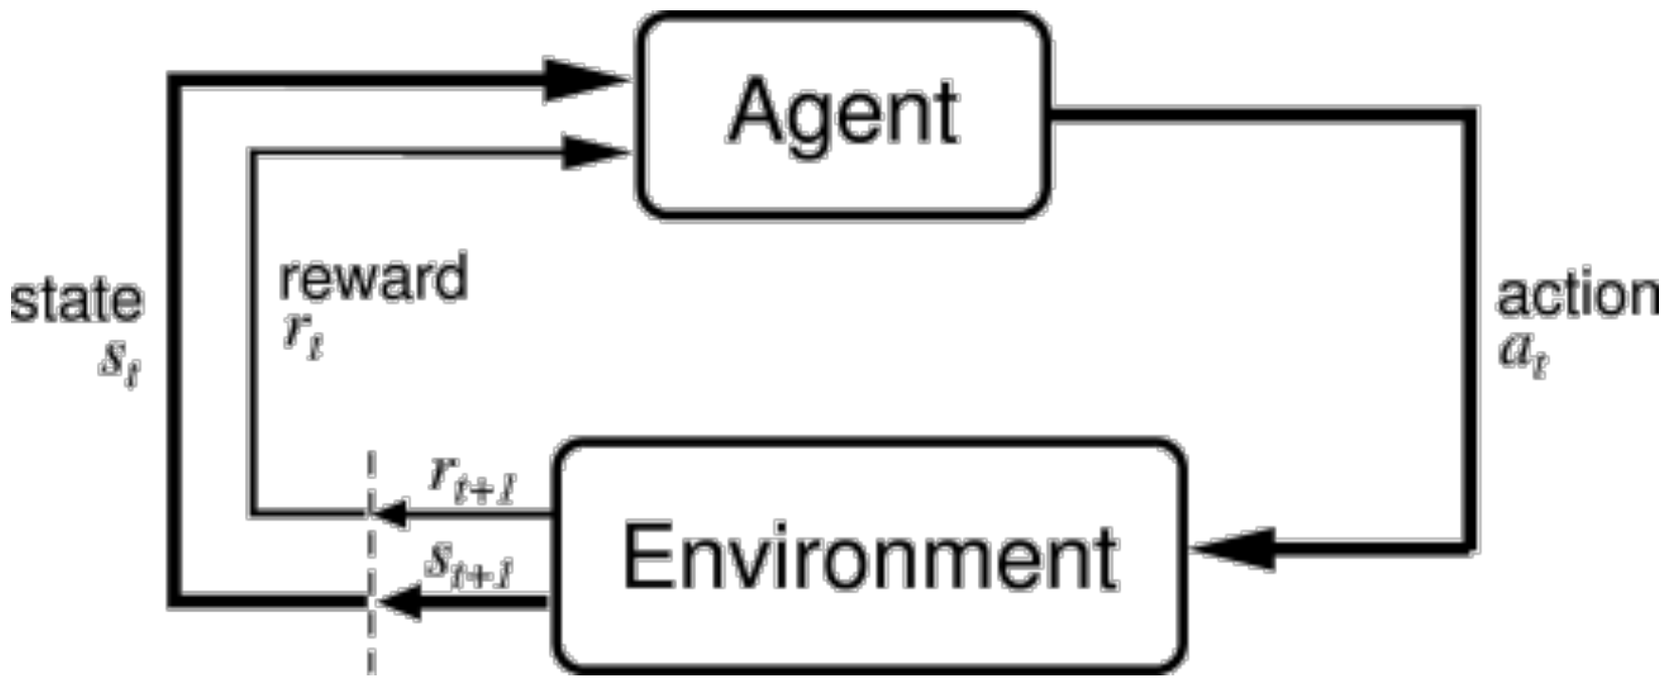

Five elements: **agent, state, reward, actions, environment**

Agent observes information about its current **state** s_t and its **reward** r_t

Agent takes an **action** at a given time a_t

Action on the **environment** yields a **new state** s_{t+1} with a **new reward** r_{t+1}

Sutton & Barto (2018)

Elements of reinforcement learning

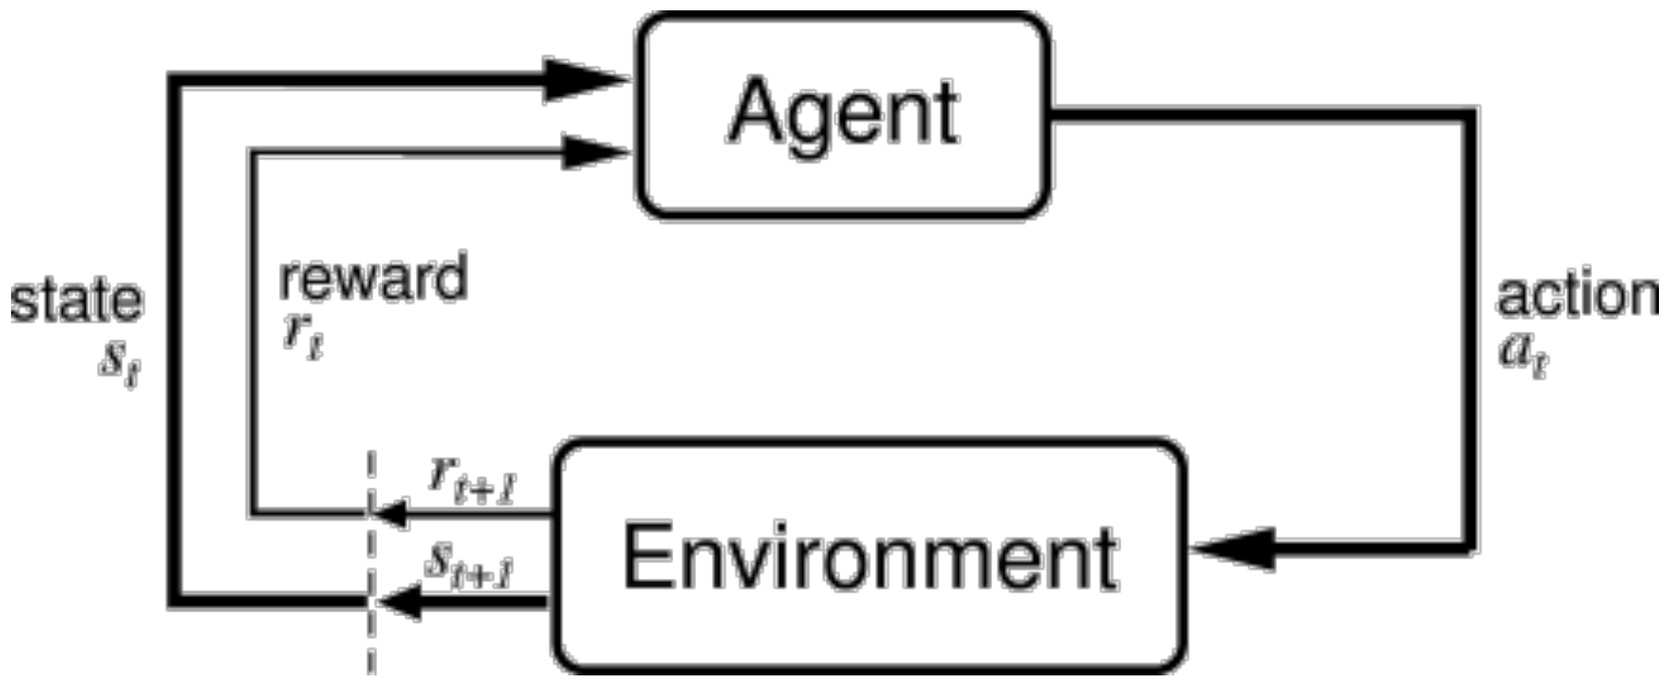

Five elements: **agent, state, reward, actions, environment**

Josh observes information about my **empty stomach** s_t and **hunger** r_t

I **eat a sandwich** at a given time a_t

Action on the **my body** yields a **full stomach** s_{t+1} with **happiness** r_{t+1}

Sutton & Barto (2018)

Mapping RL to sCI

Learn what actions
to take in order to:

$$\min_{\|\Psi_{\mathrm{CI}}\|_0 \le k} \frac{\langle \Psi_{\mathrm{CI}} | \hat{H} | \Psi_{\mathrm{CI}} \rangle}{\langle \Psi_{\mathrm{CI}} | \Psi_{\mathrm{CI}} \rangle}$$

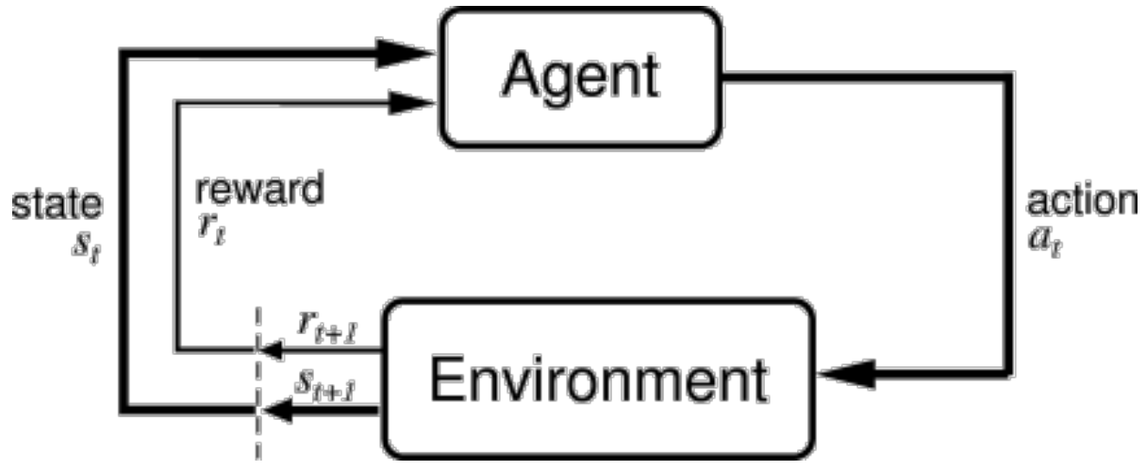

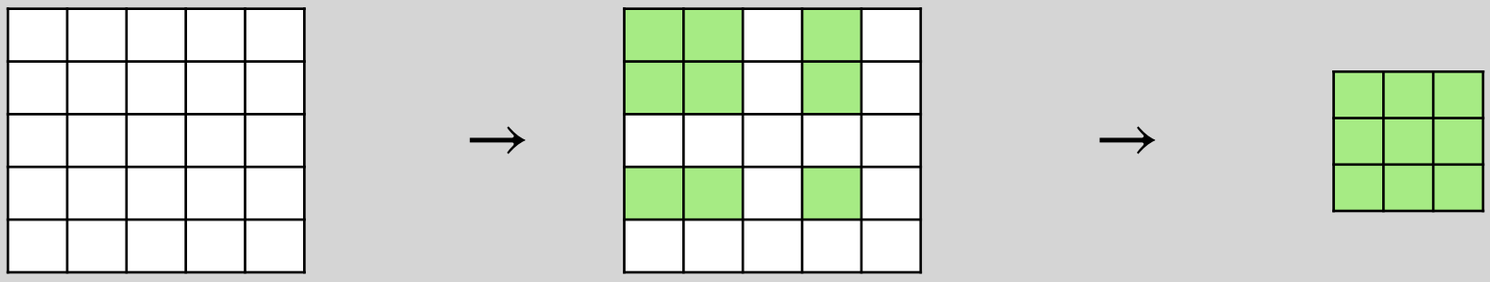

1. **Agent:** the program
2. **State:** current sub-matrix of the CI Hamiltonian
3. **Reward:** (change in) minimum eigenvalue of the CI Hamiltonian
4. **Actions:** swap SDs between “internal” and “external” space
5. **Environment:** CI Hamiltonian

Q-learning

Policy (actions) governed by Q-function $Q(s, a)$ — "quality function"

What is the quality (value) of taking an action a out of a state s?

Optimal policy $\pi(s) = \arg\max_a Q(s, a)$

$Q(s, a)$ learned incrementally for any state/action pair encountered

$$Q(s, a) \leftarrow Q(s, a) + \alpha \left[R + \gamma \max_{a'} Q(s', a') - Q(s, a) \right]$$

Improved estimate of $Q(s, a)$ is difference between reward + discounted sum of future rewards and what you expected

"What you got vs what you expected"

Better than expected? Do that more. Worse than expected? Do that less.

Align expectations with reality.

Q-learning with linear approximation

Q is **big** ($Q \in S \times A$) — represent Q function with linear approximation:

$$Q_w(s,a) = \sum_{i=1}^{m} w_i f_i(s,a)$$

Choose $f_i(s,a)$ so that Q reduces to a difference in learned weights for each determinant.

$$Q_w(s,a) = \sum_{i \in s'} w_i - w_p$$

Example:

action:
swap p with q
(det 3 with det 4)

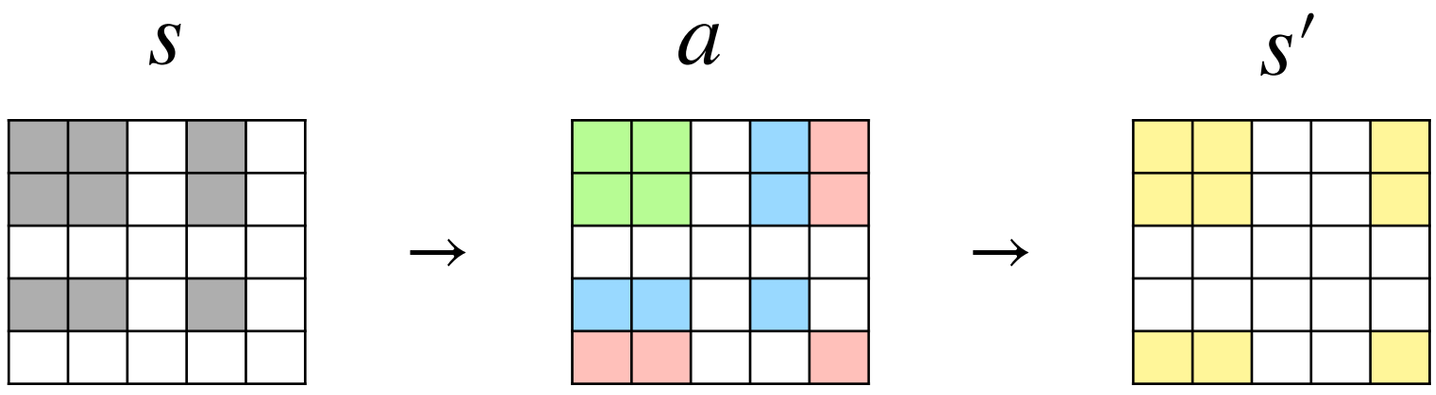

$$Q(s,a) = w_0 + w_1 + w_4 - w_3$$

Smart exploration

- Not all determinant swaps are likely to be worthwhile.
- We don’t want to spend inordinate amounts of time on useless actions.
- Estimate external determinants that may be significant with perturbation theory (similar to existing sCI methods).

$$\left| c_i^{(1)} \right| = \left| \frac{\sum_{j \neq i} H_{ij} c_j^{(0)}}{(E^{(0)} - H_{ii})} \right|$$

- Consider swapping significant external determinants with low-ranked internal determinants—can identify these based on weight in $Q_w(s, a)$.
- For efficiency, don’t consider all of the external space! Just explore over the top several candidates.
- To allow for exploration, want to accept sub-optimal actions.
  - Basically similar to a Metropolis-Hastings acceptance ratio, with greater initial exploration. Later episodes exploit learned Q.

Three “toy” test cases

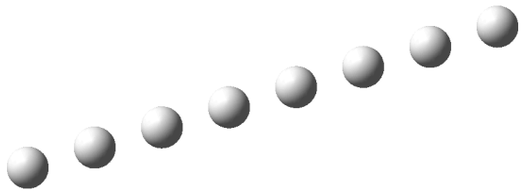

H_8

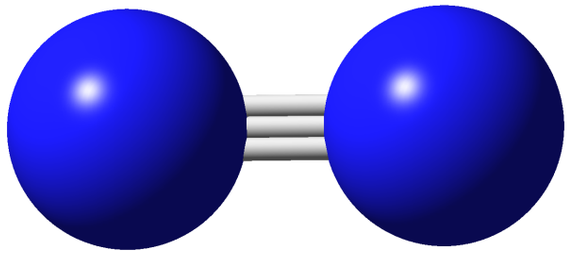

N_2

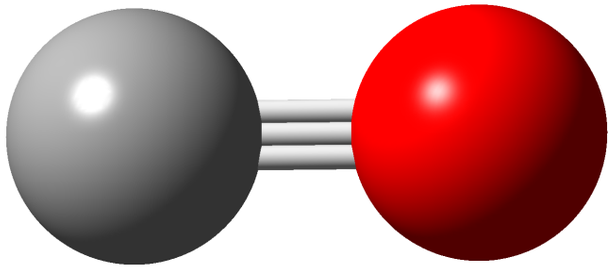

CO

- Compute potential energy surface over bond stretching
- Compare with “exact” FCI: STO-6G basis
  - FCI space: 4900 SDs for H_8, and 14400 SDs for N_2 and CO
- Energy surface allows us to evaluate under different conditions
  - e.g. stretched H_8 chain or N_2 dissociation is strongly correlated
  - CO or N_2 at equilibrium geometry not so much

Molecular nitrogen N_2

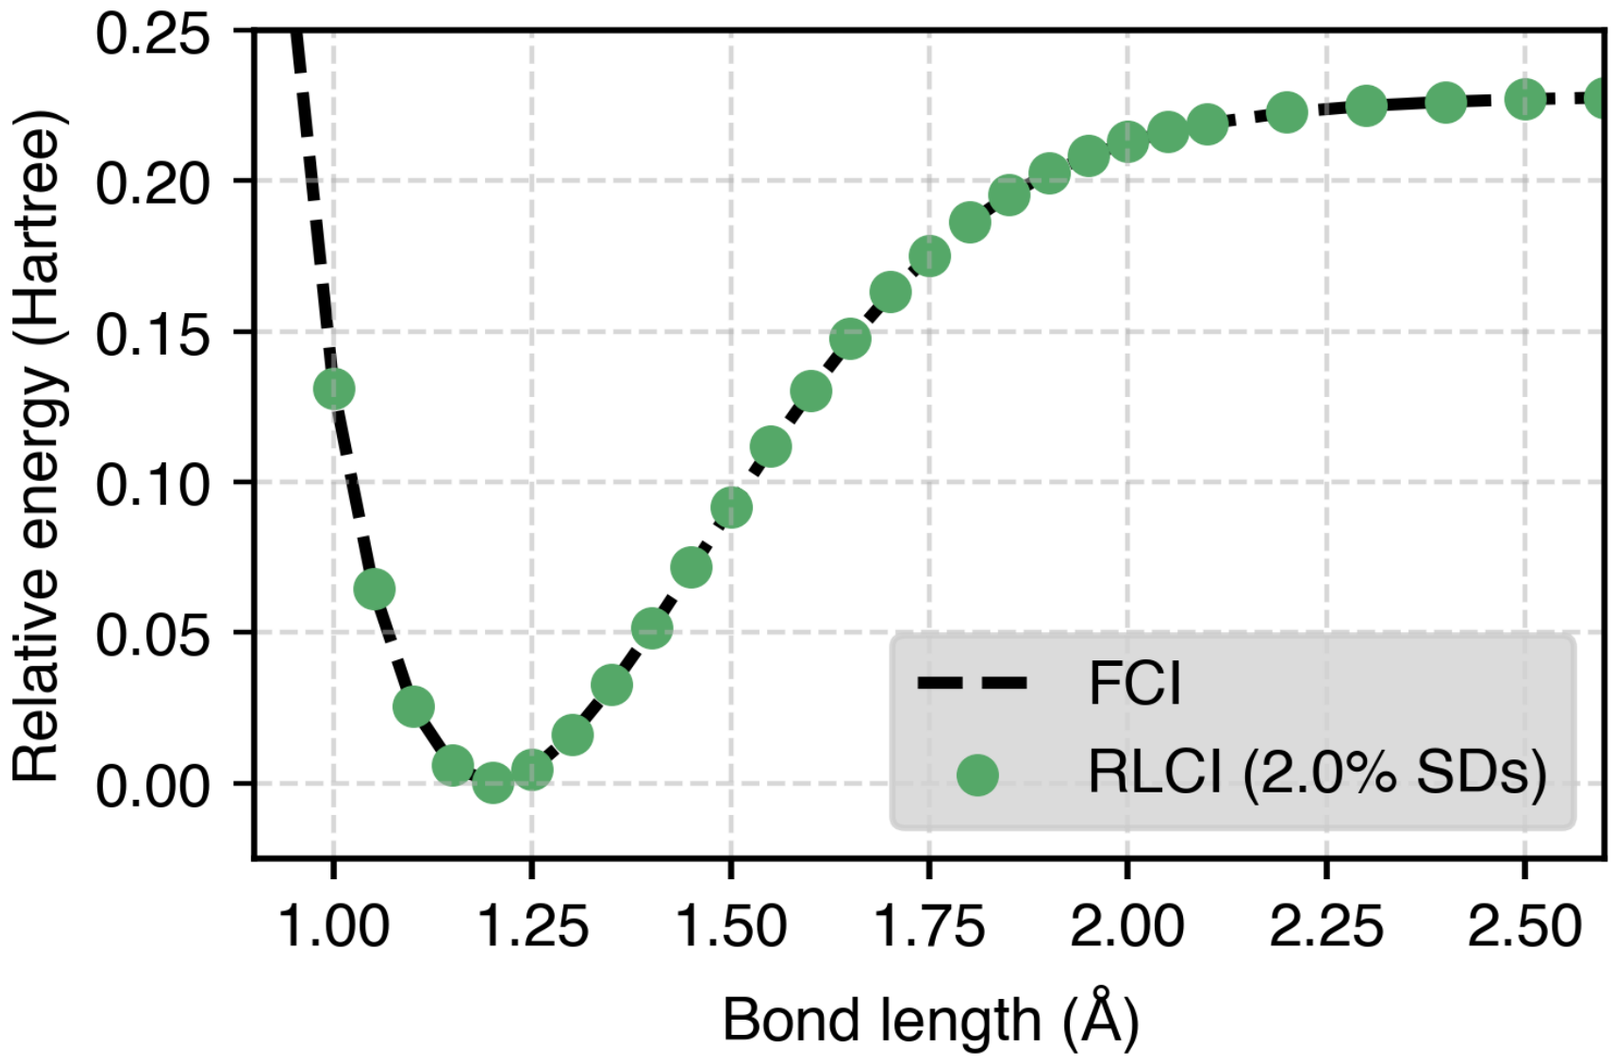

With just **2%** (288/14400) of the Slater determinants (SDs), the dissociation of nitrogen is indistinguishable from Full Configuration Interaction (FCI)

Molecular nitrogen N_2

Compare RLCI to existing sCI methods: **heat-bath CI (HCI)**

- Different free parameters, so not 1-to-1 correspondence
- But, tune a few to yield comparable accuracy (top figure)

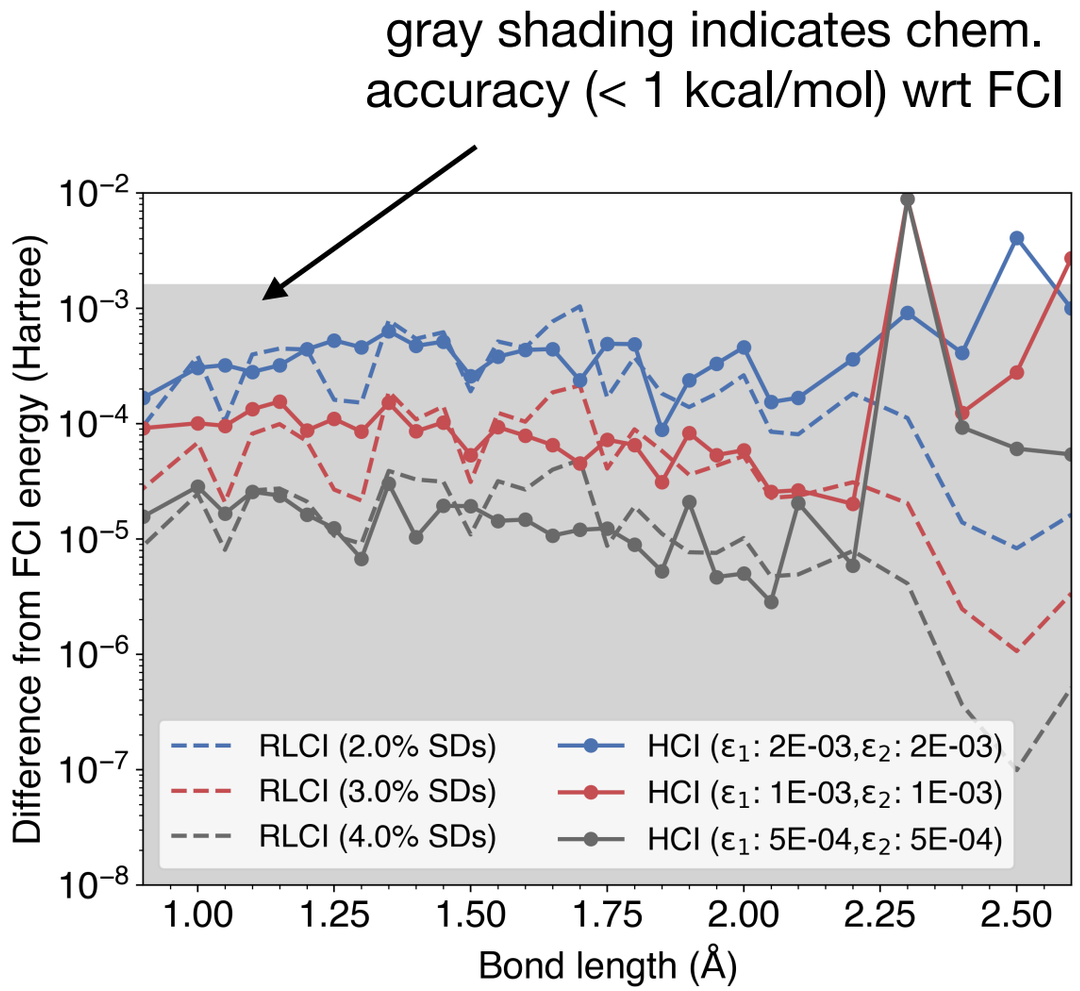

Molecular nitrogen N_2

Compare RLCI to existing sCI methods: **heat-bath CI (HCI)**

- Different free parameters, so not 1-to-1 correspondence
- But, tune a few to yield comparable accuracy (top figure)
- Efficiency ~ accuracy / num. SDs
- RLCI requires < 2—3x the number of determinants compared to HCI (bottom figure)
- **This is a considerable compression!**
- Reducing dimension of matrix by a factor of three can yield 27x faster diagonalization (9x if doing Davidson)

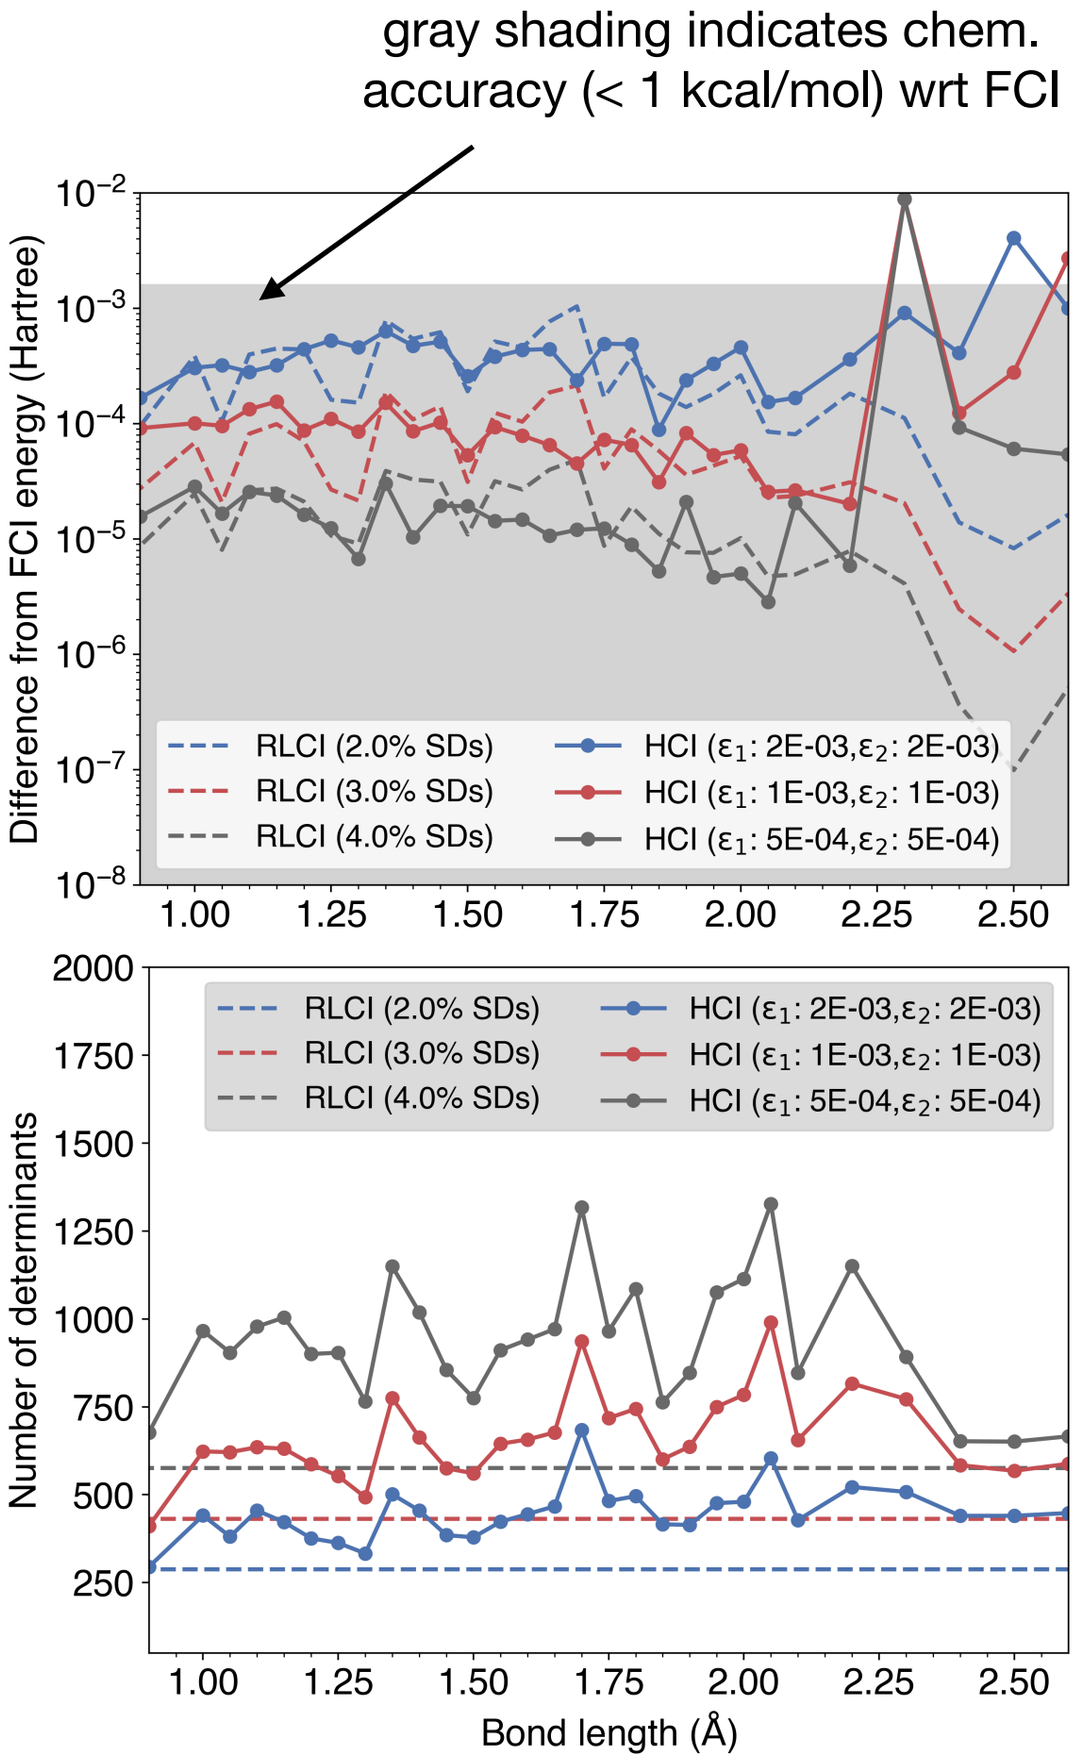

Hydrogen chains and carbon monoxide

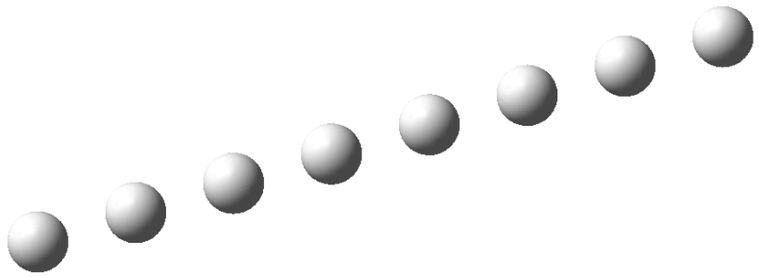

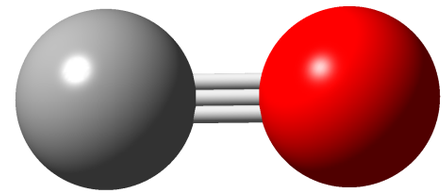

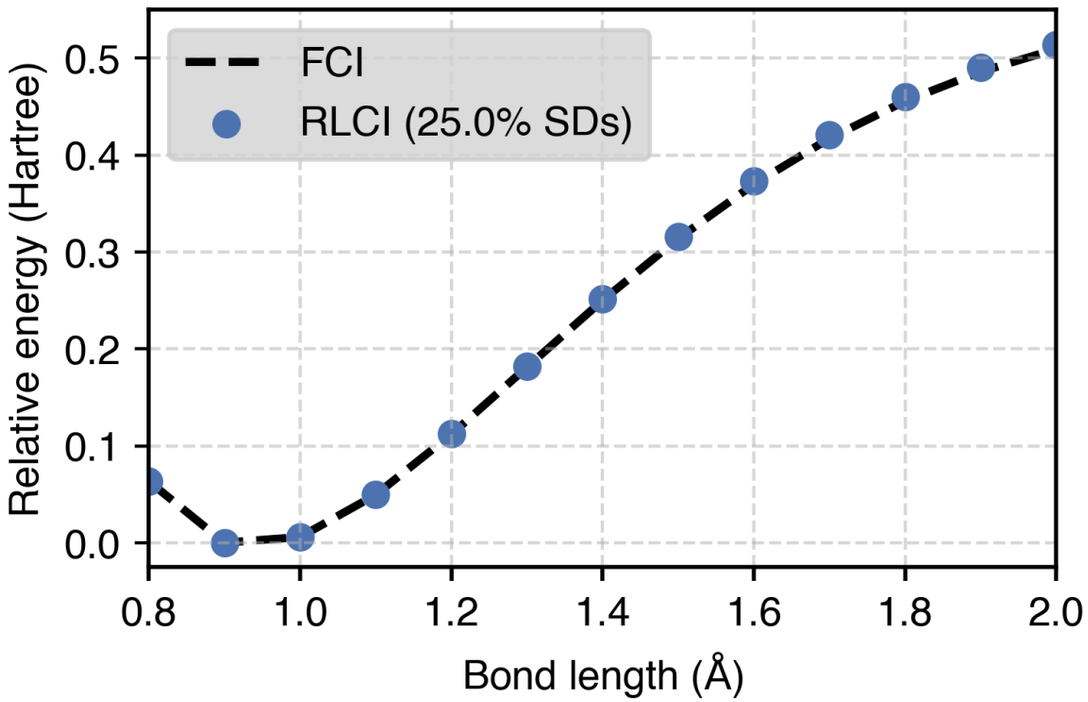

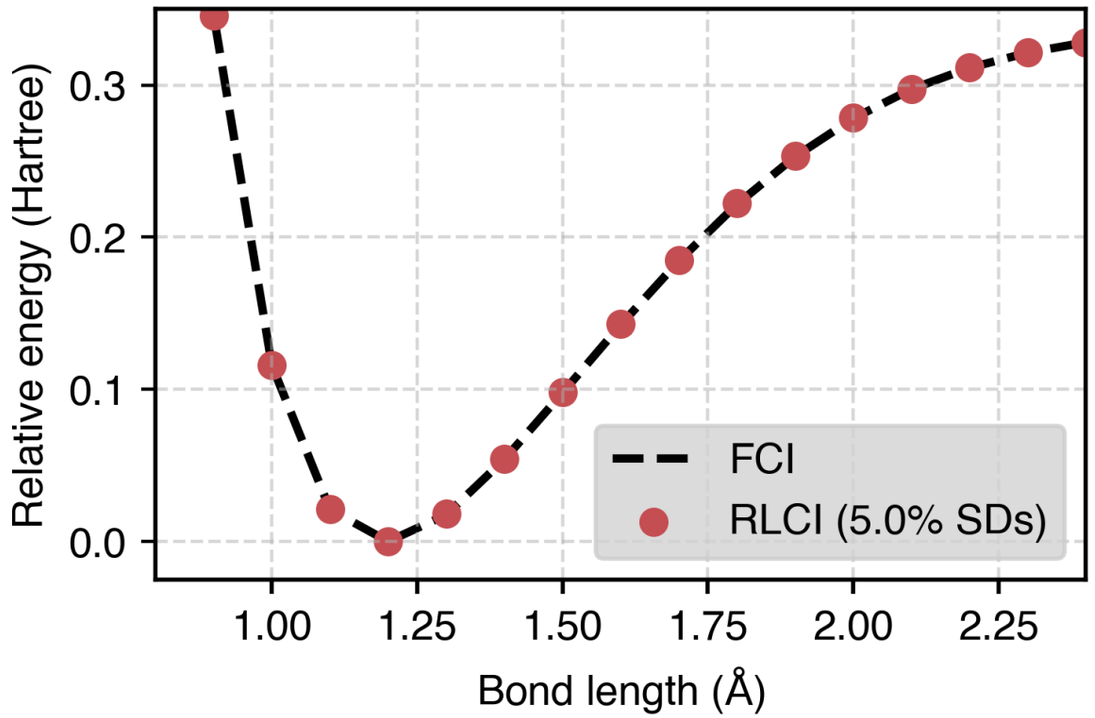

- Similar story with hydrogen chains and carbon monoxide tests.
- Fraction of full space yields FCI accuracy.
- Compared to heat-bath CI (other sCI method),
  generally around 2—3× fewer determinants for similar accuracy.

Convergence analysis

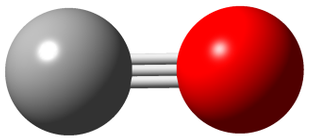

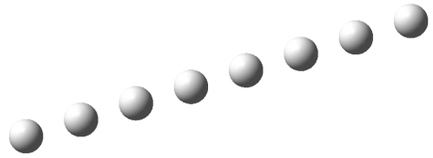

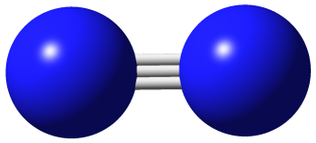

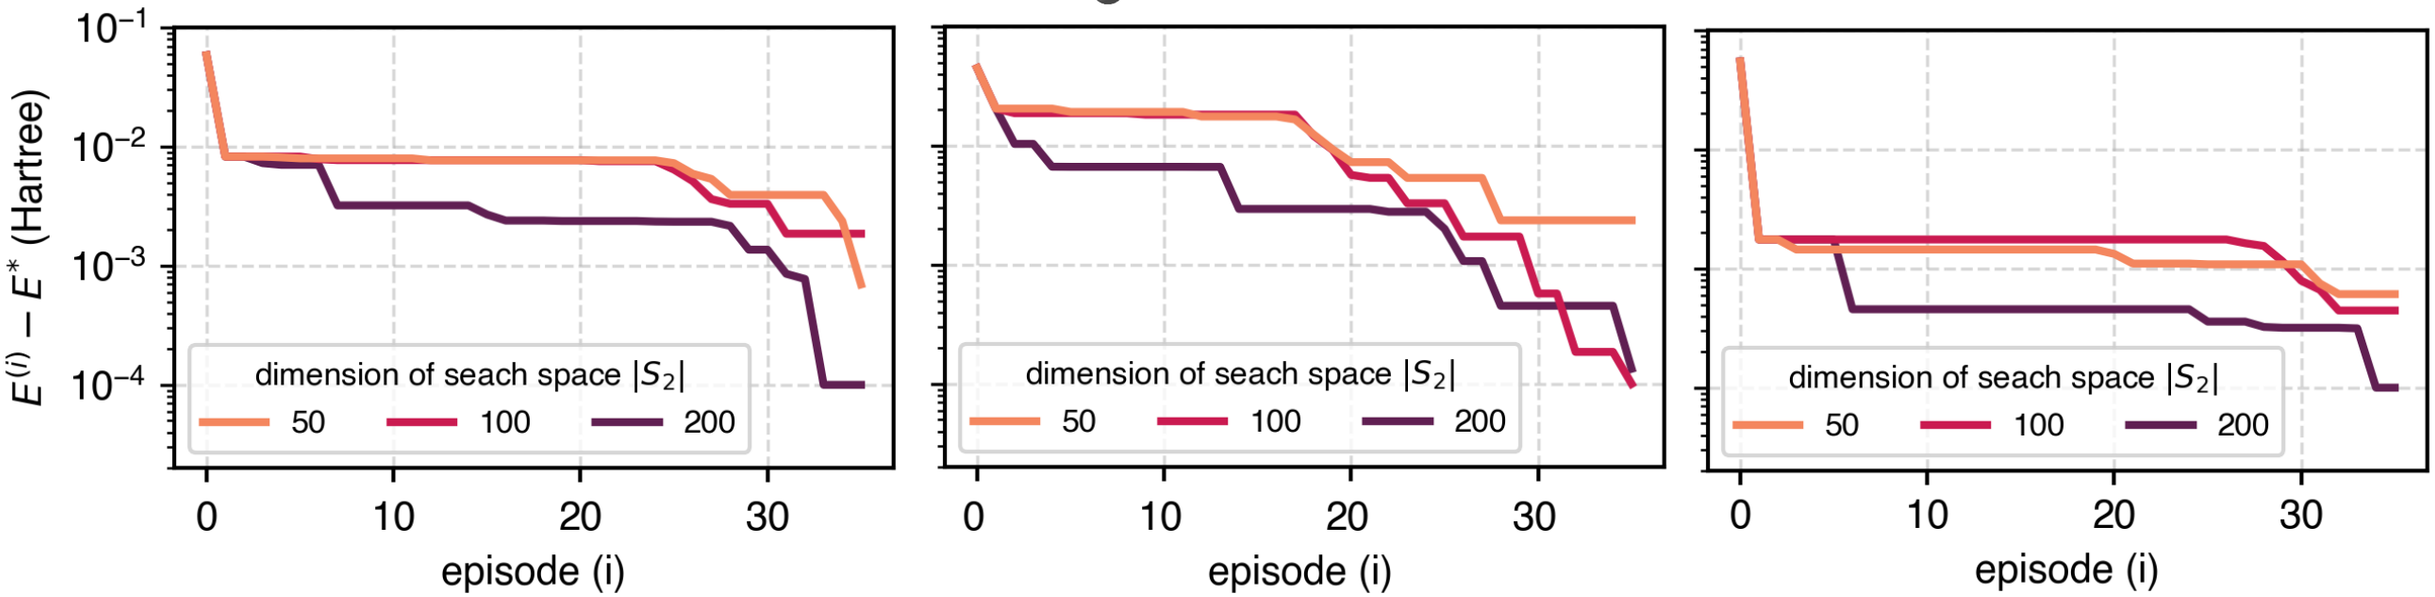

Convergence profile:

1. Rapid initial convergence from initial guess
2. Exploration for next few episodes (little improvement)
3. Exploit previous experience for last bit of energetic benefit

- Case-by-case, but generally converged by ~30 episodes.
- Need to balance search space, but generally larger search space (more choices) can give better overall optimum
- All tests here 1.5 $\mathring{\text{A}}$ with STO-6G basis and $k = 200$

Efficiency considerations

- Implemented methods in ChronusQuantum, allows us to take advantage of efficient infrastructure.
- **Lots of (small) matrix diagonalizations**:
  - Cache solutions to re-use as initial guess.
  - Actions change 1 vector element at a time, so Davidson converges *fast.*
- **Limit exploration space**:
  - Don’t try everything, just try most promising actions.
- **Memory**:
  - Weight vectors for $Q_w(s, a)$ are big, but also sparse and do not need to store all possible elements.

Hang Hu

Larger systems

- Hydrogen rings,1.5 Å apart
- Compare with FCI results https://github.com/evangelistalab/hydrogen-models-data
- RLCI recovers over 91% of FCI correlation energy with only 0.05% of the SDs in H_{14}
- This error is ~ 0.01 Hartree
- For larger rings, need more determinants (0.008% just isn't cutting it)
- Or, instead of adding more determinants in subspace, perturbation theory could be cheap route to rest of correlation (several sCI methods do this)

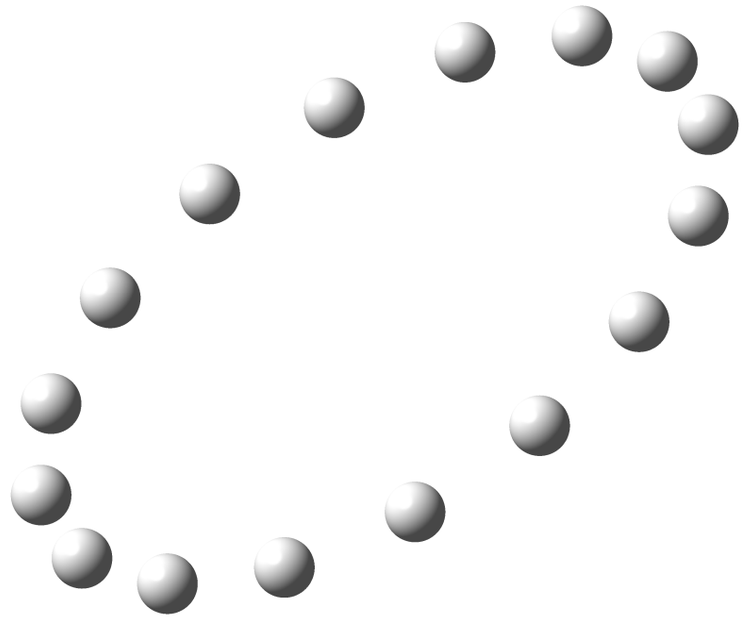

n	k	N_{det}	% N_{det}	% corr.
10	2000	184756	1.083%	96.2%
10	4000	184756	2.165%	98.3%
10	6000	184756	3.248%	99.2%
12	4000	2704156	0.148%	74.3%
12	6000	2704156	0.222%	77.2%
12	10000	2704156	0.370%	95.6%
14	6000	40116600	0.015%	86.2%
14	10000	40116600	0.025%	88.6%
14	20000	40116600	0.050%	91.3%
16	10000	601080390	0.002%	47.4%
16	20000	601080390	0.003%	49.6%
16	50000	601080390	0.008%	52.4%

What's next?

- Linear approximation for $Q(s, a)$ is biggest roadblock
- Representation: what is an efficient yet general way to encode Slater determinants?
- Neural networks? Autoencoders?
- Transfer learning — can the trained agent be deployed on other similar (yet different) problems?
- Can general heuristics be learned?

Much more to explore!

Image: Shutterstock.com

Thank you!

- Hang Hu (UW)
- Prof. Xiaosong Li (UW)
- Dr. Chao Yang (LBNL)

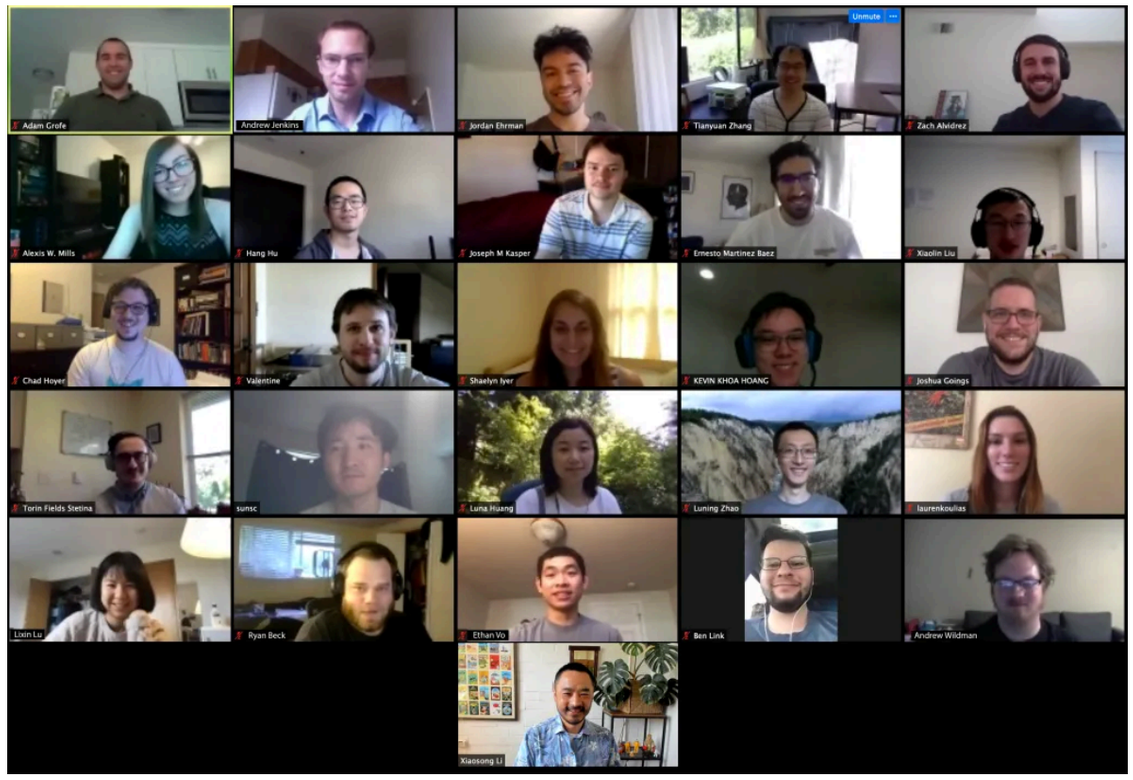

U.S. DEPARTMENT OF ENERGY | Office of Science

NSF

CQ

Goings, J.J., Hu, H., Yang, C., Li, X., Reinforcement Learning Configuration Interaction. *ChemRxiv preprint*. **2021**,10.26434/chemrxiv.14342234