Fast randomized iterative numerical linear algebra for quantum chemistry and other applications

Jonathan Weare

Courant Institute, New York University

with

Tim Berkelbach, Sam Greene, Lek-Heng Lim, and Rob Webber

November 13, 2018

The ground state eigenproblem

Compute the smallest eigenvalue E_0 corresponding to an **anti-symmetric** (with respect to particle exchange) eigenfunction, ψ_0, of

$$\mathcal{H} = \sum_{j=1}^{N_{\text{el}}} \left[-\frac{1}{2}\Delta_j + V_{\text{ext}}(x_j) - \sum_{\alpha=1}^{N_{\text{nuc}}} \frac{Z_\alpha}{|x_\alpha - x_j|} + \sum_{i<j}^{N_{\text{el}}} \frac{1}{|x_i - x_j|} \right] = -\frac{1}{2}\Delta + V$$

E_0 determines the energy of any configuration of the molecule and therefore also the structure of the molecule.

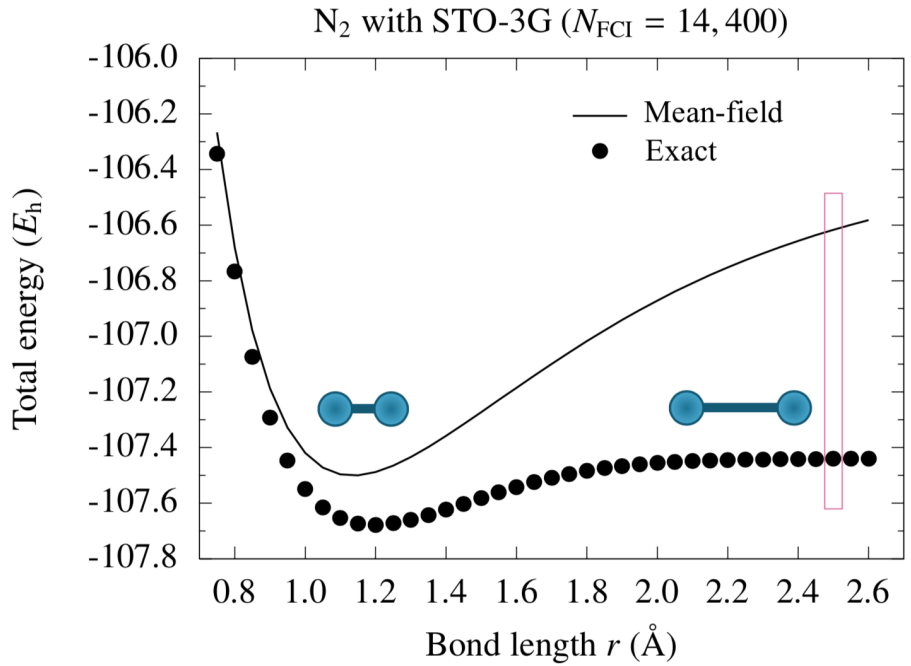

Full configuration interaction (FCI)

First project the eigenproblem $\mathcal{H}\psi_0 = E_0\psi_0$ onto an orthonormal, antisymmetric, basis of **Slater determinants**

$$\Phi_I(x) = \frac{1}{\sqrt{N_{\text{el}}!}} \begin{vmatrix} \phi_{i_1}(x_1) & \phi_{i_2}(x_1) & \cdots & \phi_{i_{N_{\text{el}}}}(x_1) \\ \phi_{i_1}(x_2) & \phi_{i_2}(x_2) & \cdots & \phi_{i_{N_{\text{el}}}}(x_2) \\ \vdots & & \ddots & \vdots \\ \phi_{i_1}(x_{N_{\text{el}}}) & \phi_{i_2}(x_{N_{\text{el}}}) & \cdots & \phi_{i_{N_{\text{el}}}}(x_{N_{\text{el}}}) \end{vmatrix}$$

$\phi_{i_1}, \phi_{i_2}, \ldots, \phi_{i_{N_{\text{el}}}}$ is a selection of N_{el} (number of electrons) distinct **orbitals** from among a collection $\phi_1, \phi_2, \ldots, \phi_M$.

So

$$\psi_0 \approx \sum_I c_I \Phi_I(x)$$

$\mathbf{Hc} = \lambda_0\mathbf{c}$ and $\lambda_0 \approx E_0$ the smallest eigenvalue of $\mathbf{H} \in \mathbb{R}^{N_{\text{FCI}} \times N_{\text{FCI}}}$ with

$$H_{IJ} = \langle \Phi_I, \mathcal{H}\Phi_J \rangle$$

DMC in discrete space (FCIQMC)

Alavi and co-workers propose approximating $\mathbf{c}$ by

$$\Psi^k = \frac{1}{N} \sum_{i=1}^{N} S^k_{(i)} \, \mathbf{e}_{X^k_{(i)}}, \qquad (\mathbf{e}_I)_J = \delta_{IJ}$$

each "walker" $X^k_{(i)}$ is the index of a determinant and $S^k_{(i)} \in \{-1, 1\}$

To generate the next iterate Ψ^{k+1} from Ψ^k, each walker $X^k_{(i)}$

1. creates new walkers on new determinants J for which $H_{X^k_i J} \neq 0$. Sign of new walkers depends on $S^k_{(i)}$ and sign of $H_{X^k_i J}$.
2. is duplicated or removed according to the size of $H_{X^k_{(i)} X^k_{(i)}}$ and current estimate of λ_0.

Pairs of walkers with $X^k_j = X^k_i$ and $S^k_j = -S^k_i$ annihilate one another.

DMC in discrete space (FCIQMC)

(ignoring signs)
similar to **diffusion Monte Carlo**:

1. walkers **move**,
2. are assigned a **weight**,
3. then are **duplicated** or **removed**

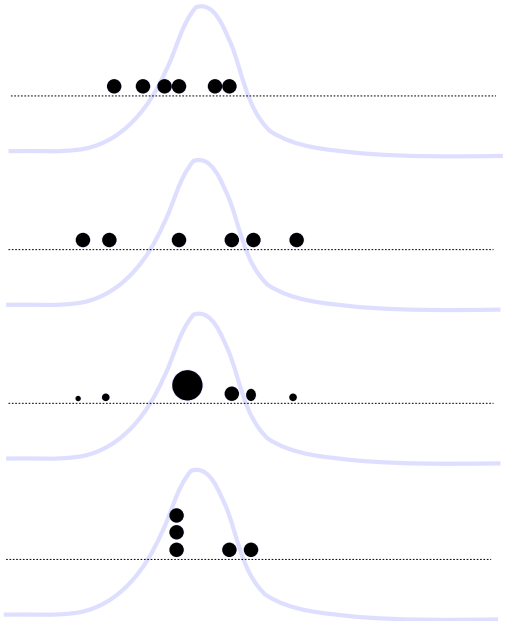

$\mathbf{E}\left[\Psi^k\right]$ approximates $\mathbf{c}^k$ produced by the power method

$$\mathbf{c}^{k+1} = \frac{\mathbf{P}_\varepsilon \mathbf{c}^k}{\|\mathbf{P}_\varepsilon \mathbf{c}^k\|_1} \qquad \text{and} \qquad \mathbf{P}_\varepsilon = (\mathbf{I} - \varepsilon \mathbf{H})$$

The ground state can be estimated from

$$E_0 \approx \lim_{K\to\infty} \frac{1}{K} \sum_{k=0}^{K} \frac{\mathbf{u}^\mathsf{T}\mathbf{H}\Psi^k}{\mathbf{u}^\mathsf{T}\Psi^k}$$

DMC in discrete space (FCIQMC)

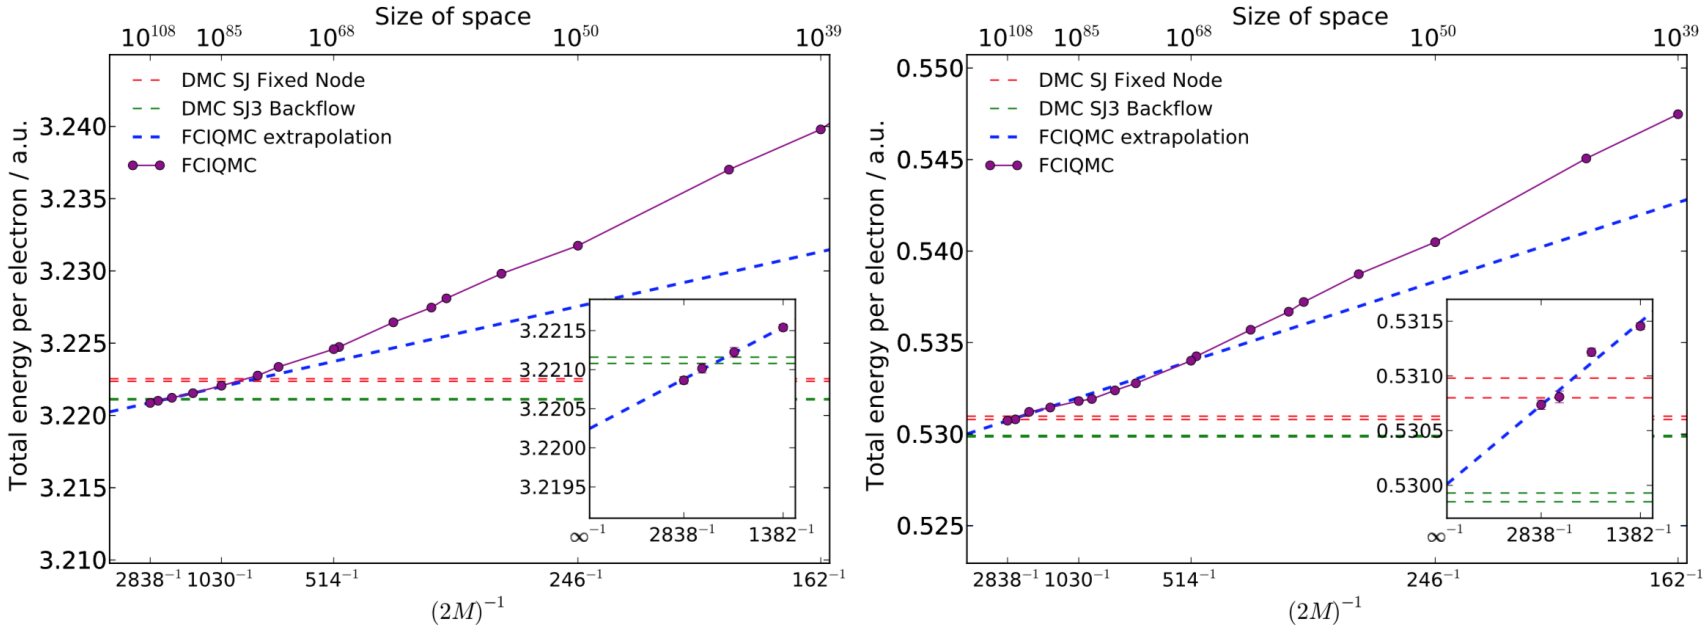

(a) r_s=0.5 (b) r_s=1.0

FIG. 3: i-FCIQMC total energies for a basis of $2M$ spin orbitals. Each basis set corresponds to a kinetic energy cutoff, with $2M = 2838$ corresponding to 208 Ryd at r_s=0.5 a.u. and 52.1 Ryd at r_s=1.0 a.u. Each calculation used 40 million walkers for $r_s = 0.5$ a.u. and 100 million walkers for $r_s = 1.0$ a.u.. The blue dashed line is an extrapolation to $M \to \infty$ based on the expected form $1/M$ using the data set with the largest number of walkers, shown with error bars in the inset. The DMC results, taken from Ríos *et al.*[6], do not suffer from basis set error and are shown as two horizontal lines representing the mean plus and minus one standard deviation. Almost identical backflow results can be found for $r_s = 1.0$ a.u. in a study by Kwon *et al.*[4].

[J.J. Shepherd, G.H. Booth, A. Gruneis, and A. Alavi, Phys. Rev. B 85, 081104–R (2012)]

Fast Randomized Iteration (FRI)

We want variants of classical schemes like power iteration

$$\mathbf{v}^{k+1} = \frac{\mathbf{A}\mathbf{v}^k}{\|\mathbf{v}^k\|_1}$$

to find the dominant eigenpair $(\lambda_*, \mathbf{v}_*)$ of $\mathbf{A}$.

and more general iterations

$$\mathbf{v}^k = \mathcal{M}(\mathbf{v}^k)$$

cost dominated by a matrix vector multiply $\mathbf{A}\mathbf{v}^k$ (e.g. linear system solves and matrix exponentials) for massive $\mathbf{A}$.

Fast Randomized Iteration (FRI)

Key to the efficiency of FCIQMC is the fact that

$$\psi^k = \frac{1}{N} \sum_{i=1}^{N} S^k_{(i)} \, \mathbf{e}_{X^k_{(i)}}$$

is sparse.

If $\mathbf{v} \in \mathbb{R}^n$ is any vector with $\|v\|_0 \leq m$ and $\mathbf{A} \in \mathbb{R}^{n \times n}$ has at most $b \leq n$ non-zero entries per column:

The cost to evaluate $\mathbf{Av}$ is $\mathcal{O}(bm)$.

Fast Randomized Iteration (FRI)

But if we continue to multiply by $\mathbf{A}$ result becomes dense quickly.

We'll replace $\mathbf{v}$ by a perturbed vector to make sure $\mathbf{A}$ always multiplies a sparse vector.

$$\mathbf{A}\,\mathbf{v} \quad \text{is replaced by} \quad \mathbf{A}(\mathbf{v} + \eta(\mathbf{v}))$$

where

$$\mathbf{E}\left[\eta(\mathbf{v})\right] = 0, \qquad \|\mathbf{v} + \eta(\mathbf{v})\|_0 \leq m$$

The cost to generate $\mathbf{v} + \eta$ is typically $\mathcal{O}\left(\|\mathbf{v}\|_0\right)$.

So the cost to evaluate $\mathbf{A}(\mathbf{v} + \eta)$ is

$$\mathcal{O}\left(\|\mathbf{v}\|_0\right) + \mathcal{O}(bm)$$

Fast Randomized Iteration (FRI)

The recursion

$$\Psi^{k+1} = \frac{\mathbf{A}(\Psi^k + \eta^k(\Psi^k))}{\|\mathbf{A}(\Psi^k + \eta^k(\Psi^k))\|_1} \quad \text{instead of} \quad \mathbf{c}^{k+1} = \frac{\mathbf{A}\mathbf{v}^k}{\|\mathbf{A}\mathbf{v}^k\|_1}$$

reduces operations and storage to $\mathcal{O}(bm)$ per iteration.

But we cannot hope that

$$\|\mathbf{v}^k - \Psi^k\|_1 \longrightarrow 0 \quad \text{unless} \quad m \sim n \quad \text{or} \quad \|\mathbf{v}^k\|_0 \ll n$$

(think of $\mathbf{v}_i^k = 1/n$)

What is the right notion of accuracy?

A few words on the accuracy of FRI

Often we only want a few dot products $\mathbf{f}^T\mathbf{c}$ for $\mathbf{f} \in \mathbb{R}^n$.

For example if $\mathbf{A}\,\mathbf{v}_* = \lambda_* \mathbf{v}_*$ then

$$\lambda_* = \frac{\mathbf{u}^T \mathbf{A} \mathbf{v}_*}{\mathbf{u}^T \mathbf{v}_*}$$

as long as $\mathbf{u}^T \mathbf{v}_* \neq 0$.

We consider the error

$$\left|\left|\left| \mathbf{v}^k - \psi^k \right|\right|\right| = \sup_{\|\mathbf{f}\|_\infty \leq 1} \sqrt{\mathbf{E}\left[|\mathbf{f}^T \mathbf{v}^k - \mathbf{f}^T \psi^k|^2 \right]}$$

A few words on the accuracy of FRI

A few useful properties of

$$|||\mathbf{X}||| = \sup_{\|\mathbf{f}\|_\infty \leq 1} \sqrt{\mathbf{E}\left[|\mathbf{f}^T\mathbf{X}|^2\right]}$$

for random $\mathbf{X} \in \mathbb{R}^n$:

1. $|||\mathbf{X}||| = \sup_{\|\mathbf{G}\|_{\infty,*}} \sqrt{\mathbf{E}\left[\|\mathbf{G}\mathbf{X}\|_1^2\right]}$

   where for $\mathbf{G} \in \mathbb{R}^n$, $\|\mathbf{G}\|_{\infty,*} = \sum_{i=1}^n \max_{j \leq n} |G_{ij}|$
2. If $\mathbf{X}$ is not random then $|||\mathbf{X}||| = \|\mathbf{X}\|_1$
3. If $\mathbf{X}$ has independent components then $|||\mathbf{X}|||^2 = \mathbf{E}\left[\|\mathbf{X}\|_2^2\right]$

In general

$$\mathbf{E}\left[\|\mathbf{X}\|_2^2\right] \leq |||\mathbf{X}|||^2 \leq \mathbf{E}\left[\|\mathbf{X}\|_1^2\right]$$

A few words on the accuracy of FRI

Consider the general recursions

$$\Psi^{k+1} = \mathcal{M}(\Psi^k + \eta^k(\Psi^k)) \quad \text{and} \quad \mathbf{v}^{k+1} = \mathcal{M}(\mathbf{v}^k)$$

where the η^k are independent conditioned on Ψ^k and

$$\mathbf{E}\left[\eta^k(\mathbf{v})\right] = 0, \qquad \left|\left|\left|\eta^k(\mathbf{v})\right|\right|\right| \leq \frac{\gamma}{\sqrt{m}}\|\mathbf{v}\|_1$$

Can adapt lots of tools for perturbed dynamical systems to produce e.g. global in time error bounds when $\mathcal{M}$ is contractive in the appropriate sense.

Special attention should be paid to the dependence on n.

A few words on the accuracy of FRI

Corollary

If $\mathbf{A}$ *has non-negative entries and is irreducible and aperiodic, then the power iteration*

$$\mathcal{M}(\mathbf{v}) = \frac{\mathbf{A}\,\mathbf{v}}{\|\mathbf{A}\,\mathbf{v}\|_1}$$

is contractive (in the appropriate sense) as long as the entries of $\mathbf{v}$ *are non-negative. So the randomized iteration* $\Psi^k = \mathcal{M}(\Psi^k + \eta^k)$ *satisfies*

$$\left|\left|\left|\Psi^k - \mathbf{v}^k\right|\right|\right| \leq \frac{C}{\sqrt{m}}$$

where C is independent of k and does not (explicitly) depend on n.

Fixed accuracy in constant cost is possible (at least for $b \sim 1$).

Surprising from an NLA perspective. Not from an MCMC (or DMC) perspective.

What about more general matrices?

Perturbations of identity

Recall that $\mathbf{P}_\varepsilon = (\mathbf{I} - \varepsilon\mathbf{H})$.

If $\mathcal{M}(\mathbf{v}) = \mathbf{v} + \varepsilon\mathcal{R}(\mathbf{v})$ we hope our bounds remain stable as $\varepsilon \to 0$ and $k \sim \varepsilon^{-1}$.

This requires a good compression scheme. If $\|\mathbf{v}\|_0 \leq m$, we require

$$\left|\left|\left|\eta^k(\mathbf{v}+\mathbf{w})\right|\right|\right| \leq \frac{\gamma}{\sqrt{m}}\|\mathbf{w}\|_1^{\frac{1}{2}}\|\mathbf{v}+\mathbf{w}\|_1^{\frac{1}{2}}$$

This rules out some possible compression rules. E.g. we cannot use

$$\mathbf{v} + \eta(\mathbf{v}) = \frac{\|\mathbf{v}\|_1}{m}\sum_{i=1}^{n} N_i \frac{v_i}{|v_i|}\mathbf{e}_i$$

where $(N_1, \ldots, N_n) \sim \text{Multinomial}\Big(m, \|\mathbf{v}\|_1^{-1}(|v_1|, \ldots, |v_n|)\Big)$

Vector compression

Some basic principles:

1. No benefit to perturbing an entry that will remain non-zero
2. Important to correlate the N_j (sampling with/without replacement or multinomial/systematic resampling).

Stopping rule selects largest entries to preserve exactly.

Remaining entries $j \in R$ set to

$$(\mathbf{v} + \eta(\mathbf{v}))_j = \begin{cases} \frac{v_j}{p_j} & \text{w.p. } p_j \\ 0 & \text{otherwise} \end{cases}$$

$$p_j = \frac{(m - |R^c|)|v_j|}{\sum_{i \in R} |v_i|} \leq 1 \quad \text{for} \quad j \in R$$

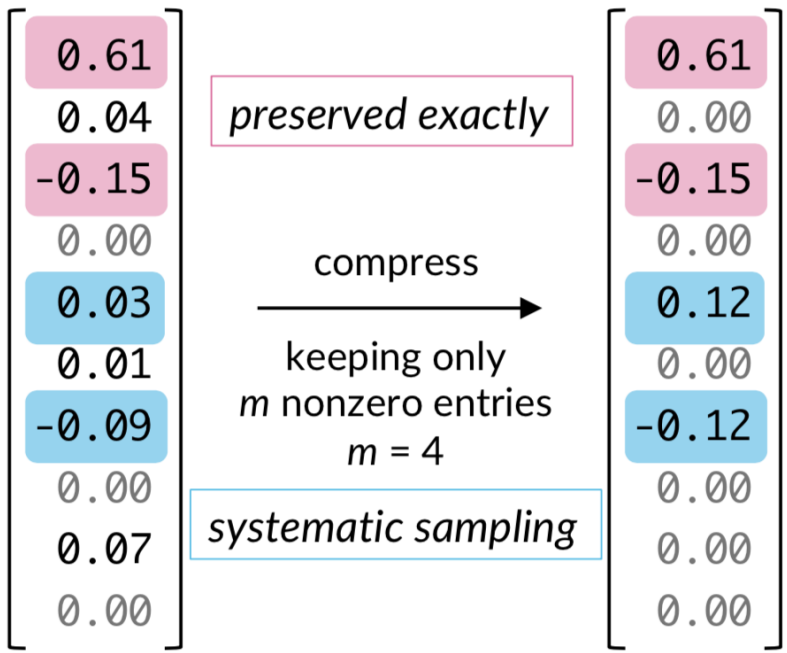

N_2

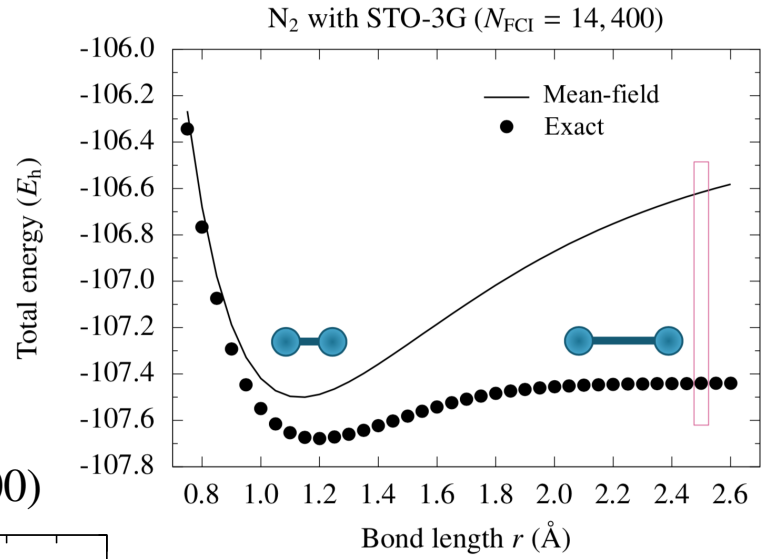
N2 with STO-3G (NFCI = 14, 400)
Total energy (Eh)
Bond length r (Å)
Mean-field
Exact
-106.0
-106.2
-106.4
-106.6
-106.8
-107.0
-107.2
-107.4
-107.6
-107.8
0.8
1.0
1.2
1.4
1.6
1.8
2.0
2.2
2.4
2.6

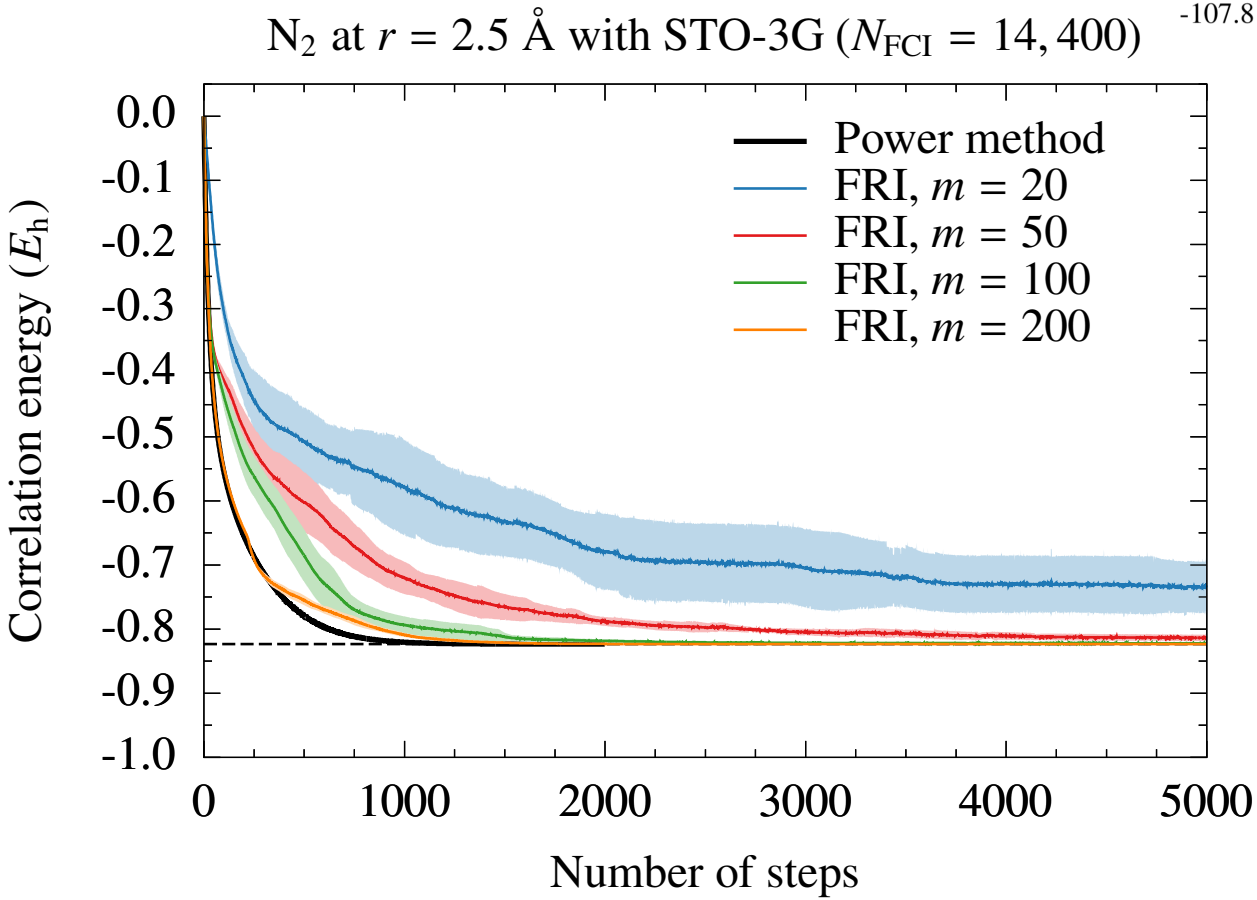
N2 at r = 2.5 Å with STO-3G (NFCI = 14, 400)
Correlation energy (Eh)
Number of steps
Power method
FRI, m = 20
FRI, m = 50
FRI, m = 100
FRI, m = 200
0.0
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
-0.7
-0.8
-0.9
-1.0
0
1000
2000
3000
4000
5000

N_2

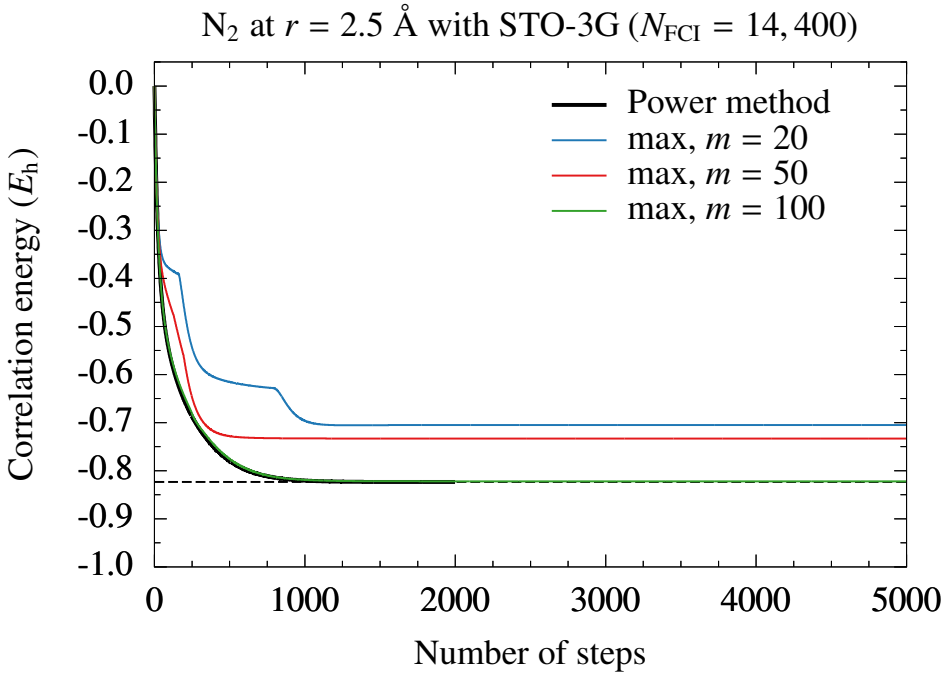
N2 at r = 2.5 Å with STO-3G (NFCI = 14, 400)
Power method
max, m = 20
max, m = 50
max, m = 100
Correlation energy (Eh)
Number of steps
0.0
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
-0.7
-0.8
-0.9
-1.0
0
1000
2000
3000
4000
5000

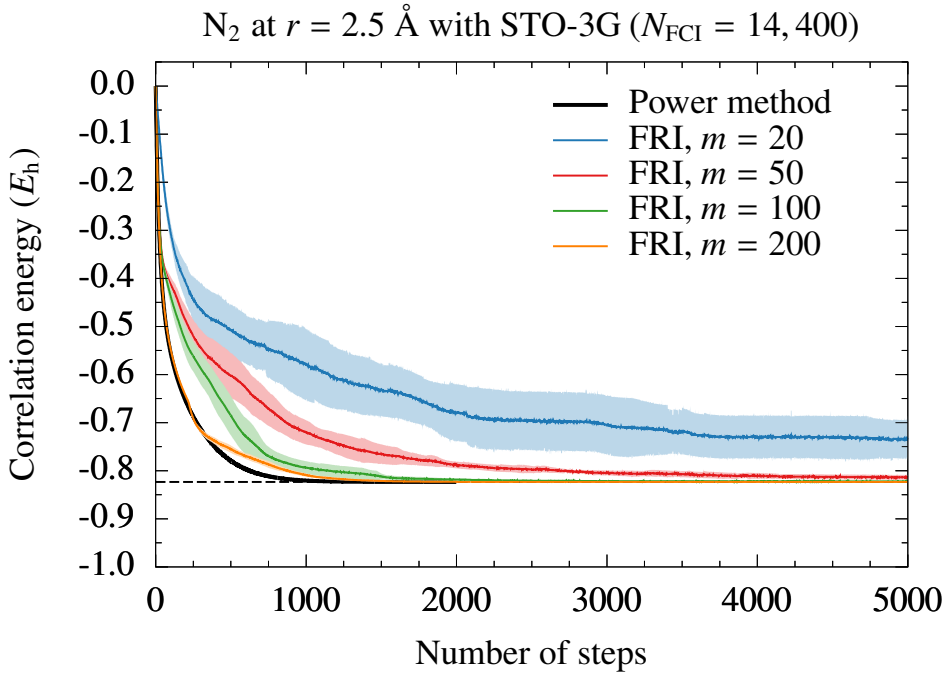
N2 at r = 2.5 Å with STO-3G (NFCI = 14, 400)
Power method
FRI, m = 20
FRI, m = 50
FRI, m = 100
FRI, m = 200
Correlation energy (Eh)
Number of steps
0.0
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
-0.7
-0.8
-0.9
-1.0
0
1000
2000
3000
4000
5000


0.61
0.04
-0.15
0.00
0.03
0.01
-0.09
0.00
0.07
0.00
compress
keeping only
m nonzero entries
m = 4
0.61
0.00
-0.15
0.00
0.00
0.00
-0.09
0.00
0.07
0.00

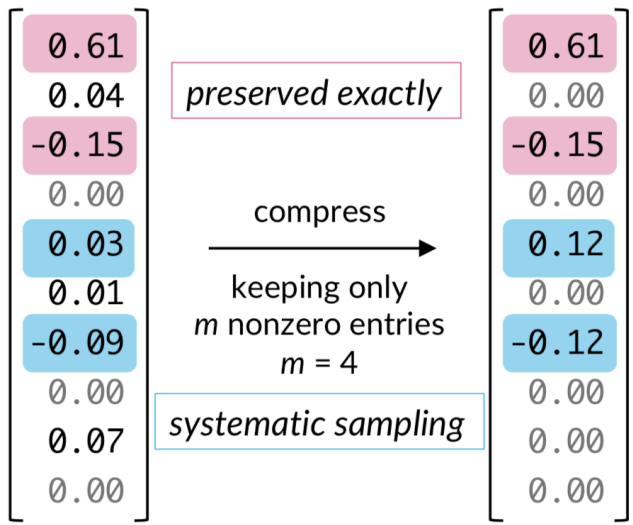
0.61
0.04
-0.15
0.00
0.03
0.01
-0.09
0.00
0.07
0.00
preserved exactly
compress
keeping only
m nonzero entries
m = 4
systematic sampling
0.61
0.00
-0.15
0.00
0.12
0.00
-0.12
0.00
0.00
0.00

Matrix compression

Can't afford to list all entries in a column of $\mathbf{A}$? This is the case for $\mathbf{P}_\varepsilon$.

We can use compression again if we have "cheap" matrix $\mathbf{Q}$ with

$$|A_{ij}| > 0 \Longrightarrow |Q_{ij}| > 0$$

$$(\mathbf{A} + \eta(\mathbf{A}))_{ij} = \begin{cases} \frac{A_{ij}}{Q_{ij}} (\mathbf{Q} + \eta(\mathbf{Q}))_{ij} & \text{if } |Q_{ij}| > 0 \\ 0 & \text{otherwise} \end{cases}$$

Only evaluate A_{ij} when $(\mathbf{Q} + \eta(\mathbf{Q}))_{ij} \neq 0$.

Only compress columns of $\mathbf{A}$ corresponding to non-zero entries of $\mathbf{v} + \eta$.

E.g. $\mathbf{Q}$ might be uniform on the indices of non-zero entries of $\mathbf{A}$.

Matrix compression

Compressing the matrix columns will typically require that column entries of **Q** can be factored in a tree structure.

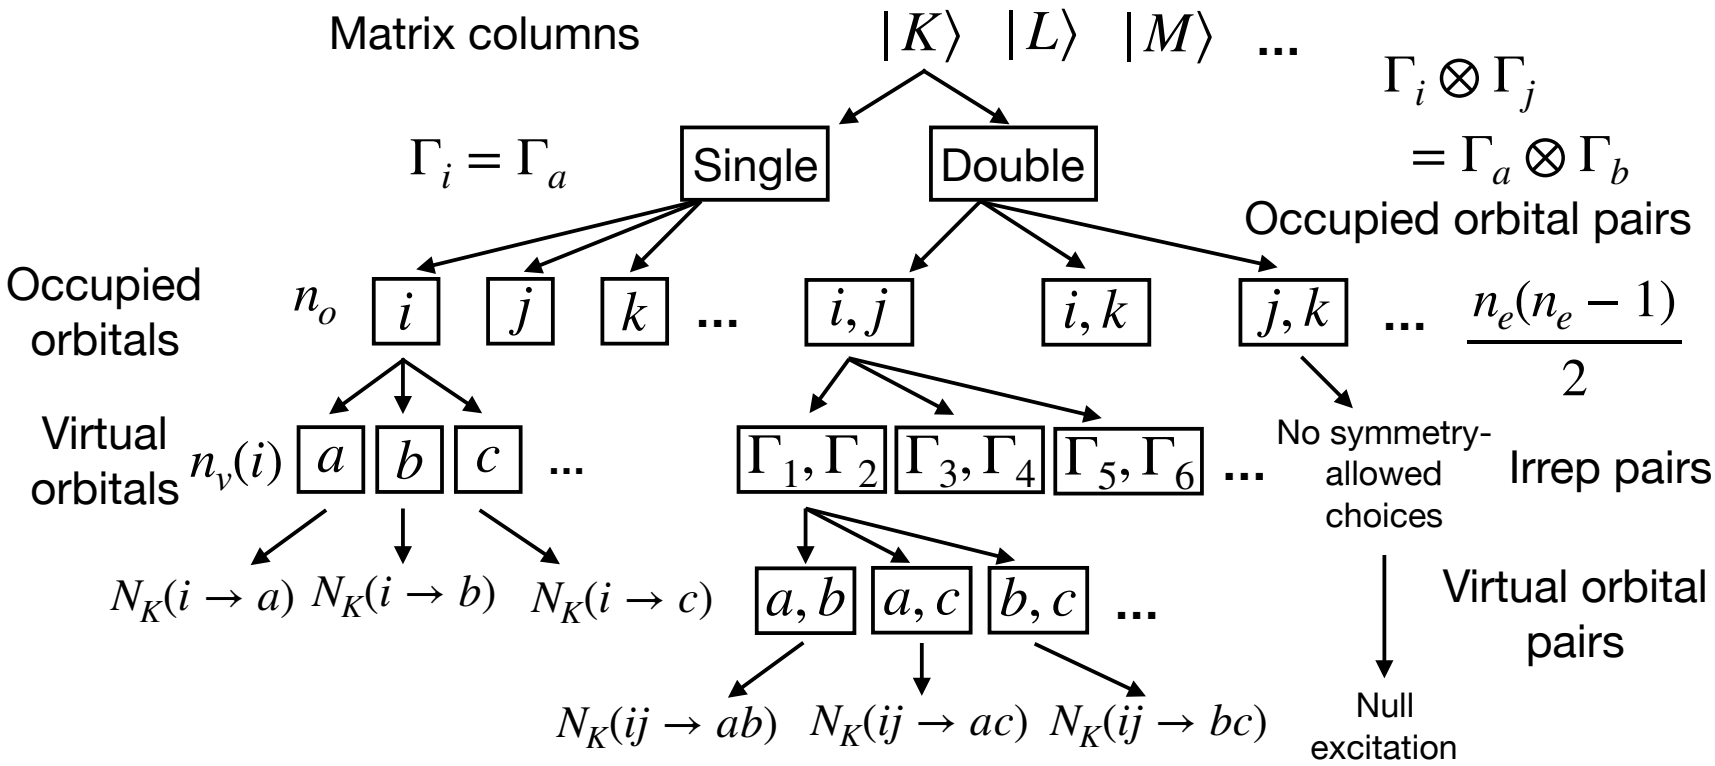

It can get complicated.

Ne

aug-cc-pVDZ basis (6.9 million determinants)

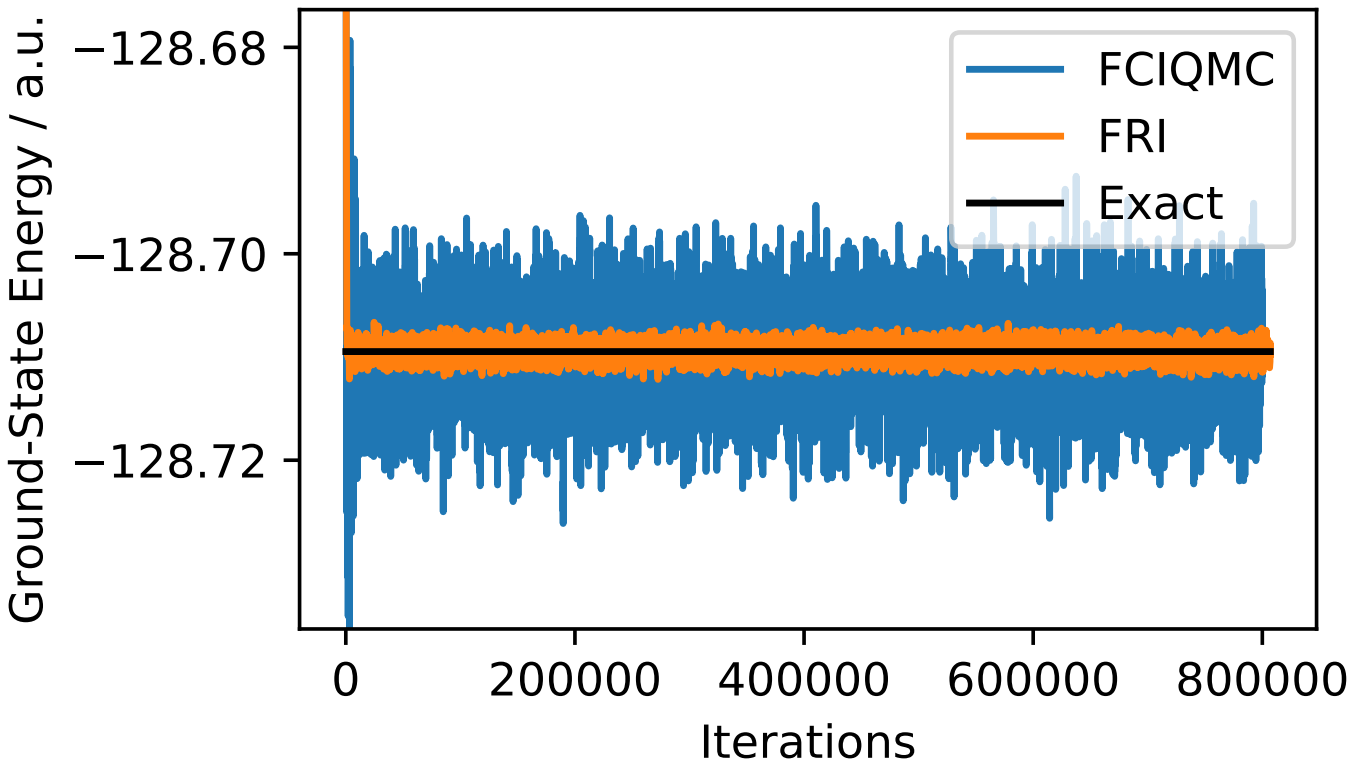

$m = 243K$

Ne

aug-cc-pVDZ basis (6.9 million determinants)

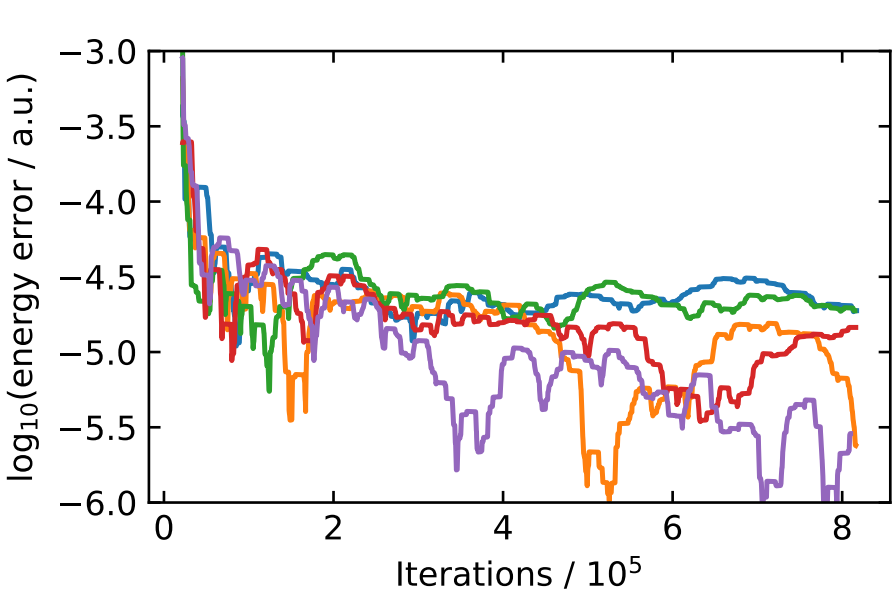

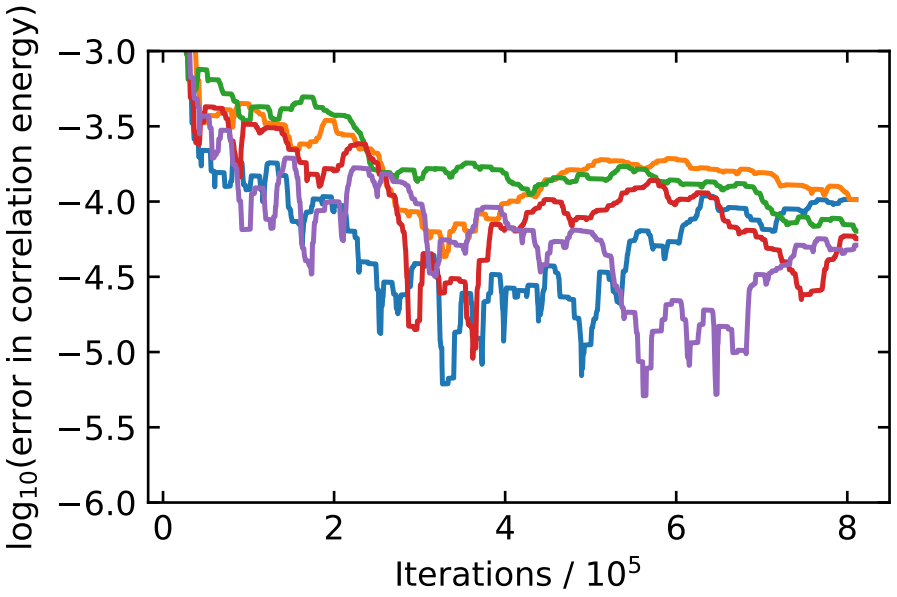

For QMC applications FRI **may** offer a significant performance improvement over FCIQMC at similar cost.

This is still a very small problem and we're testing against an early version of FCIQMC. Parallel scaling is a key question.

Is FRI just an efficient search

Does FRI quickly find a sparse vector that is a good approximation of the true ground state?

For Neon, the energy of the iterates produced by FRI is well above the final FRI estimate.

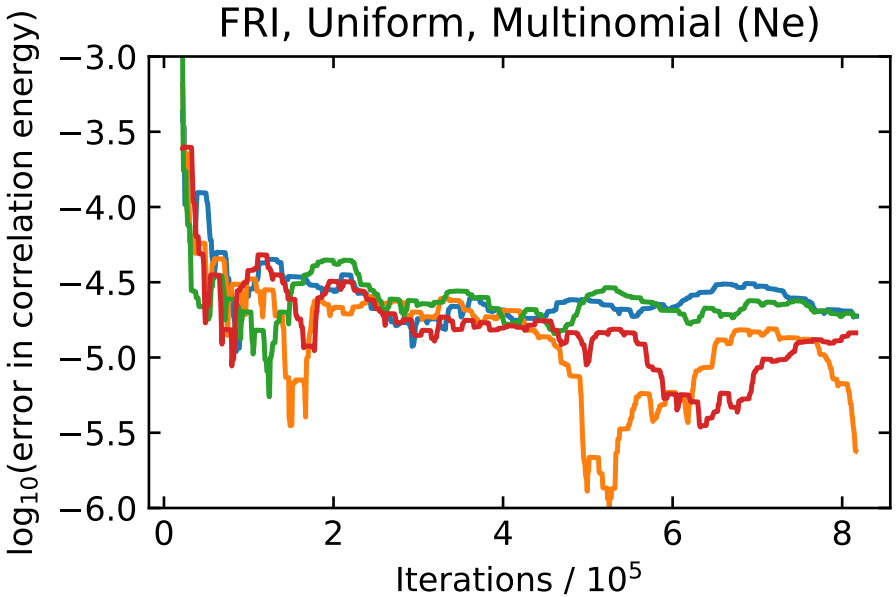

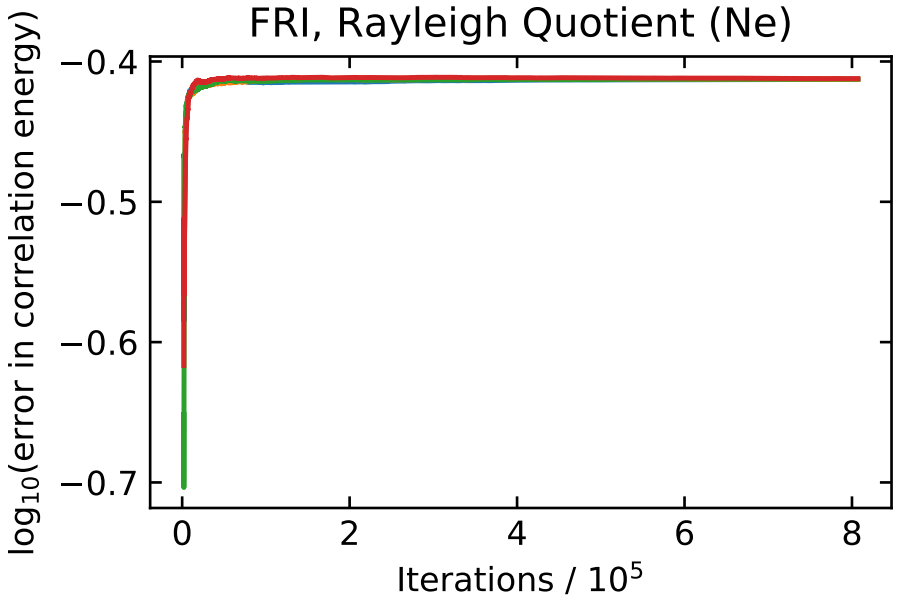

Recall: at least for a non-negative matrix, constant cost is possible even when the desired eigenvector is not at all sparse.

Why not just truncate?

Consider the overdamped operator

$$Lf = -\nabla V \nabla^{\mathsf{T}} f + \Delta f, \qquad (x_j, y_j) \in [-1,1) \times [-1,1) \text{ for } j \leq \ell$$

$$V(x,y) = \frac{1}{2} \sum_{j \leq \ell} \cos(2\pi x_j) \cos(2\pi y_j) + 2 \sum_{\substack{j \leq 4 \\ k > j}} \cos\left(\pi(x_j - x_k)\right) \cos\left(\pi(y_j - y_k)\right)$$

ℓ attractive particles each experience a corrugated external potential

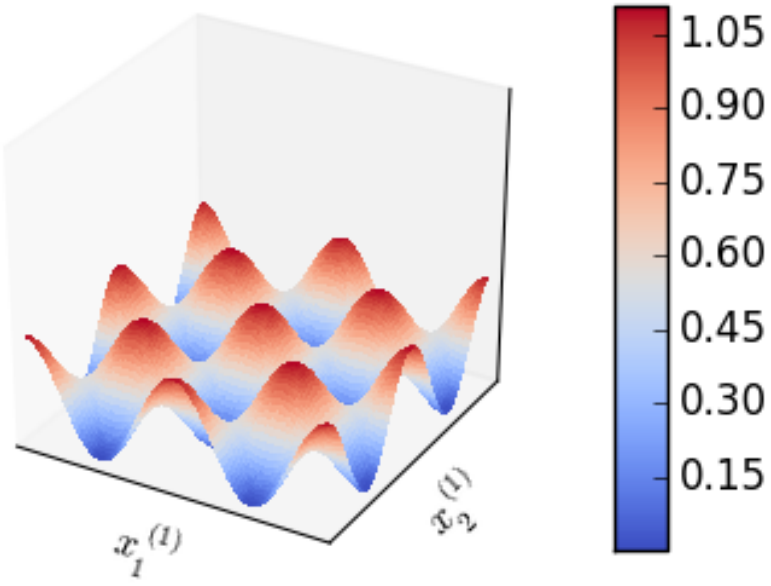

Why not just truncate?

For 5 particle (10D) we discretize with a Fourier basis of size $101^{10}/2 \approx 10^{20}/2$ and find the spectral gap (second largest eigenvalue of L).

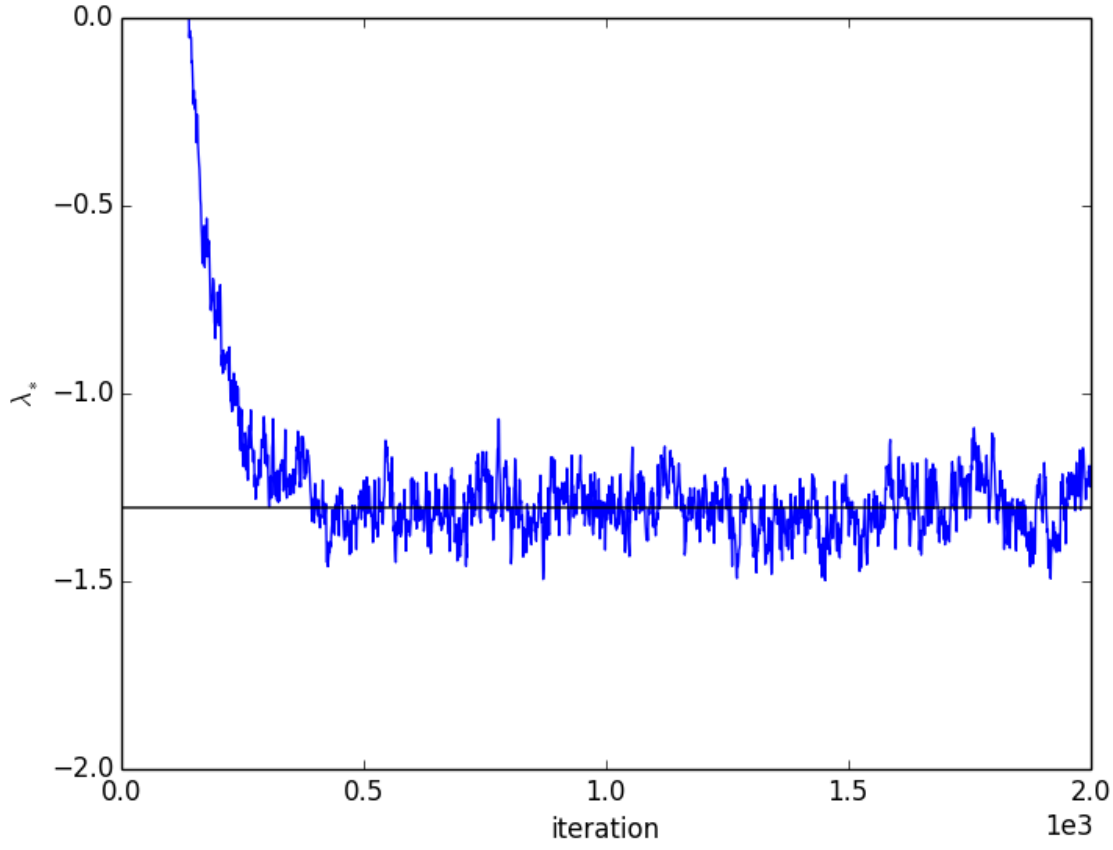

$m = 10^6$, no matrix compression

Why not just truncate?

For 4 particles (8D) with a basis of size $101^8/2 \approx 10^{16}/2$ we compare against truncation:

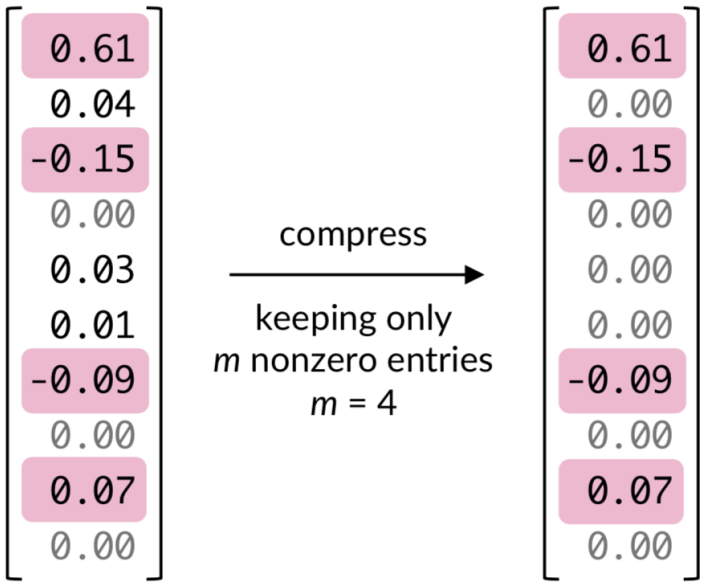

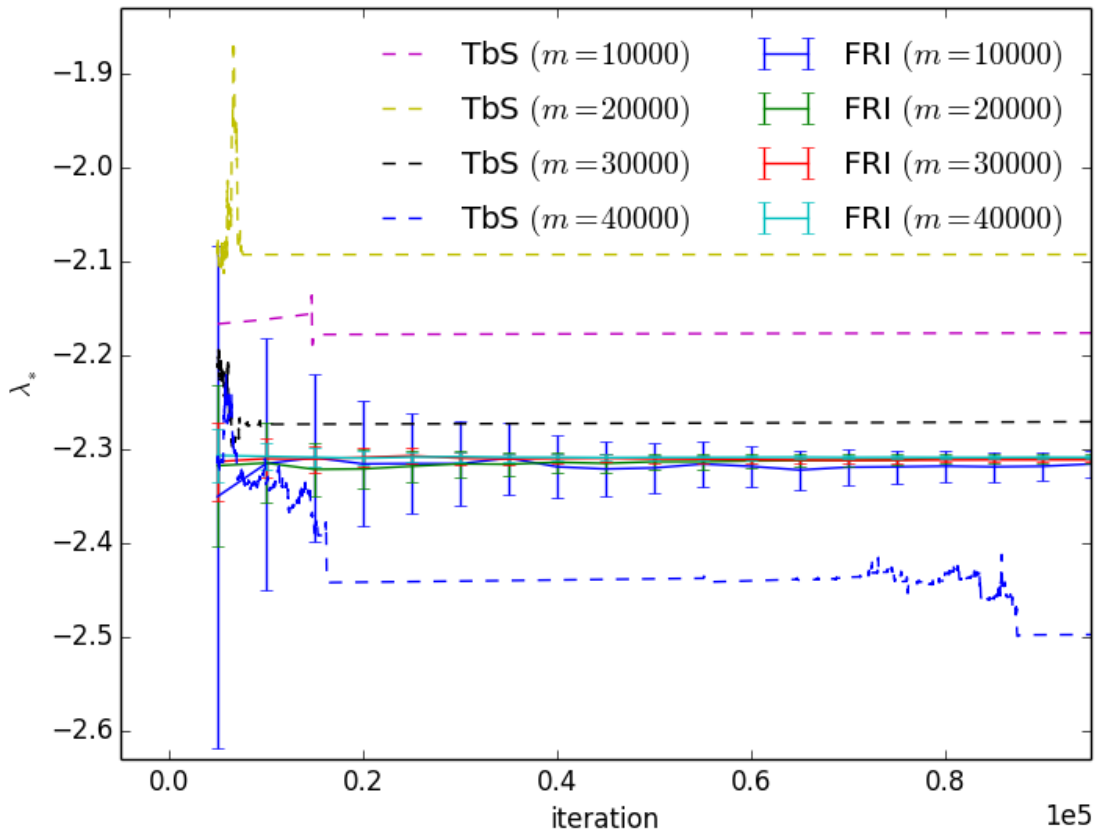

no matrix compression