Extreme quantum mechanics in MATLAB

Ilya Kuprov, University of Southampton

What is quantum mechanics?

- A mathematical description of reality at small scales
- All that is knowable about a system is contained in one function
- Squared amplitude of that function is probability density

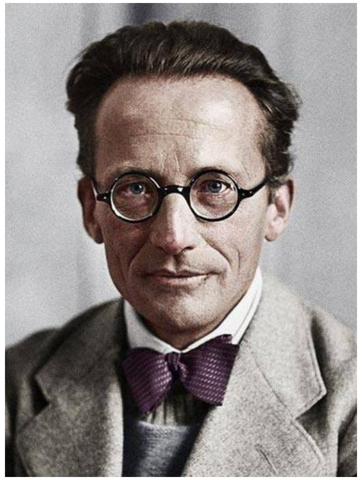

Erwin
Schrödinger

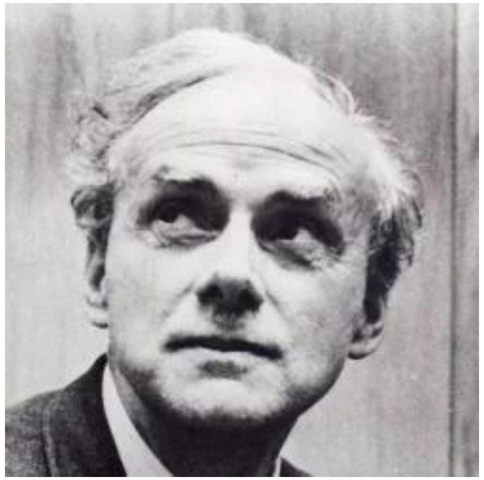

Paul
Dirac

matter potential momentum

$$\frac{\partial}{\partial t}\begin{bmatrix}\Psi^+ \\ \Psi^-\end{bmatrix} = -\frac{i}{\hbar}\begin{bmatrix}\varphi + mc^2 & c\boldsymbol{\sigma}\cdot\boldsymbol{\pi} \\ c\boldsymbol{\sigma}\cdot\boldsymbol{\pi} & \varphi - mc^2\end{bmatrix}\begin{bmatrix}\Psi^+ \\ \Psi^-\end{bmatrix}$$

antimatter **spin** rest mass

What is spin?

- A relativistic symmetry that gives charged particles a magnetic moment
- No classical analogue, mathematically similar to angular momentum
- Responsible for most magnetic properties of matter

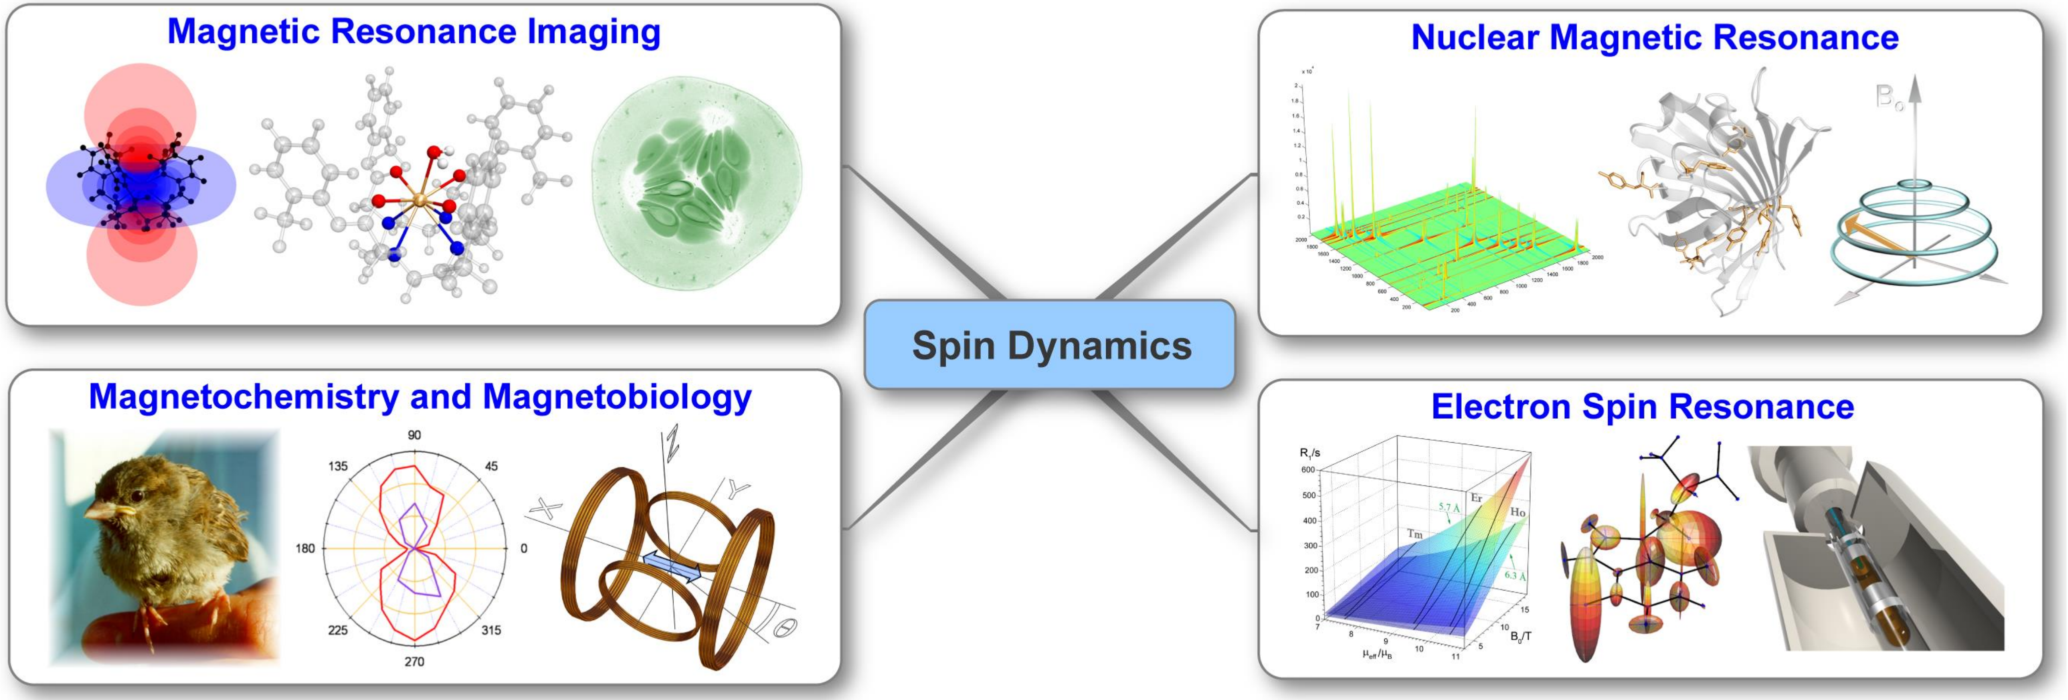

Magnetic resonance industry

"By 2020, nuclear magnetic resonance (NMR) spectrometry is likely to lead the spectroscopy market in terms of annual revenue, which is expected to reach approximately $1.7 billion worldwide."

Transparency Market Research Report, Aug 2014

"Estimates suggest that the global MRI market was worth about £4.3 billion in 2010 and is expected to grow to around £6.2 billion by 2015, equivalent to an annual growth of 7.7% a year."

Oxford Economics Report, Nov 2012

NMR statistics, worldwide

Number of instruments at research organizations:	~16,000
Number of instruments at industrial companies:	~7,000
Market capitalization of the three major vendors:	$78.5B
Number of academic publications in 2014:	~55,000

MRI statistics, worldwide

Number of instruments at research organizations:	~3,000
Number of instruments at medical institutions:	~31,000
Market capitalization of the five major vendors:	$4.65T
Number of academic publications in 2014:	~59,000

Photosynthesis involves generation of spin-correlated radical pairs.

DOI: 10.1038/nature06834

Birds use a spin-selective chemical reaction to sense the direction of the Earth's magnetic field.

DOI: 10.1021/bi060330h

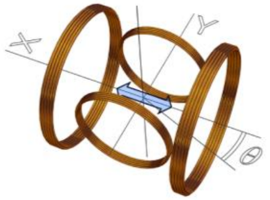

Spin is at the core of every magnetic interaction and every technology that uses magnetism.

DOI: 10.1063/1.2010287

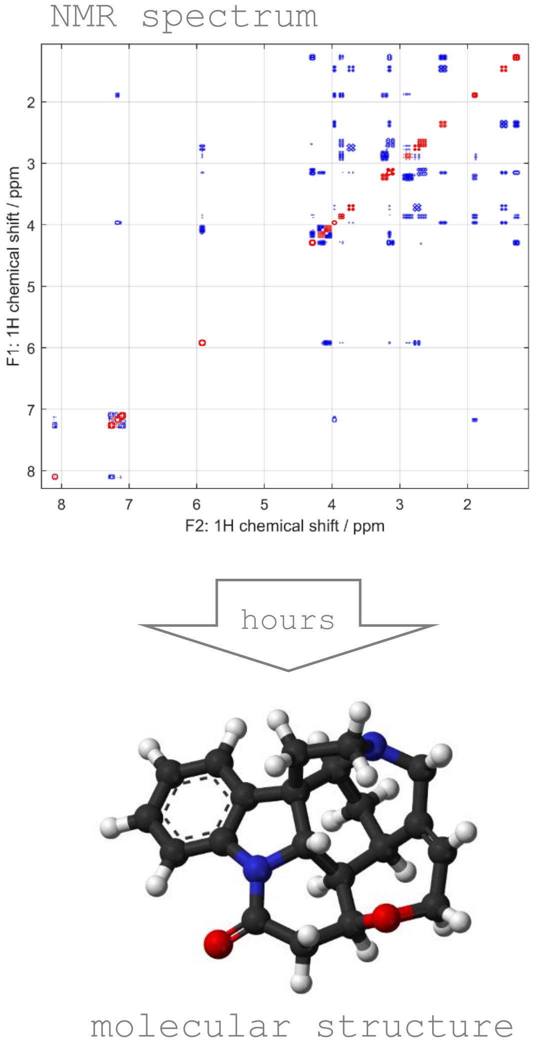

Magnetic resonance simulations

quantum degrees of freedom

$\hat{\rho}$ matrix of probabilities

$\hat{H}$ matrix of energies

sizeable (dim > 10^3) and very sparse complex matrices

classical degrees of freedom

$\{x, y, z\}$ coordinates

$\{\alpha, \beta, \gamma\}$ orientations

$\{a, \varphi\}$ radio waves

$\{\varphi_1, \varphi_2, ...\}$ sample spinning

+ conformations, concentrations, *etc.*

spatial derivative operators are turned into matrices using finite difference approximations

equation of motion

$$\frac{\partial}{\partial t}\hat{\rho}(t) = -i\left[\hat{H}(t), \hat{\rho}(t)\right]$$

Liouville - von Neumann equation

(basically Schrödinger equation for ensembles)

equation of motion

$$\frac{\partial}{\partial t}p(...) = -div\left[flux\left[p(...)\right]\right]$$

Fokker-Planck equation

(probability balance with a great number of special cases)

dim ~ 1000 — quantum degrees of freedom $\otimes$ classical degrees of freedom — dim ~ 100×100×100

dim ~ 1000×100×100×100

Principal problem: sum-of-direct-products-times-a-vector (can just about store the vector).

Polyadic object in MATLAB

A sum of matrix krons acting on a vector:

$$\left(\alpha\left[\mathbf{A}\otimes\mathbf{B}\otimes\ldots\right]+\beta\left[\mathbf{C}\otimes\mathbf{D}\otimes\ldots\right]+\ldots\right)\mathbf{x}$$

matrices scalar vector

A *short* sum of krons of *small* matrices! Times a vector…

dim[**A**] = 1000
dim[**B**] = 1000
dim[**A**(x)**B**] = 10^6
numel(**v**) = 10^6
dim[**V**] = 1000

$$\left[\mathbf{A}\otimes\mathbf{B}\right]\mathbf{v}=\mathrm{vec}\left[\mathbf{B}\mathbf{V}\mathbf{A}^{\mathrm{T}}\right]$$

~10^{12} FLOP ~10^9 FLOP

+ various technicalities: sums, products by scalars, *etc.*

```
% Bring forward n-th dimension
dims=1:numel(x_dims); dims(n)=[];
dims=[n,dims]; x=permute(x,dims);

% Unroll other dimensions
x=reshape(x,[col_dims(n),numel(x)/col_dims(n)]);

% Run multiplication and update dimension map
x=Q{nmats-n+1}*x; x_dims(n)=row_dims(n);

% Roll other dimensions back up
x=reshape(full(x),[row_dims(n),x_dims(dims(2:end))]);

% Put the current dimension back
x=ipermute(x,dims);
```

```
>> A=randn(1000); B=randn(200); C=randn(500);
>> H=polyadic({{A,B,C}})

H =

  100000000×100000000 polyadic array with properties:

     cores: {{1×3 cell}}
    prefix: {}
    suffix: {}

>> cheap_norm(H)

ans =

   5.4739e+07

>> tic; (H+H')*rand(1e8,1); toc
Elapsed time is 9.662600 seconds.
```

Polyadic object in MATLAB

Addition is implemented as buffering of terms:

$$\mathbf{A}\otimes\mathbf{B}+\mathbf{C}\otimes\mathbf{D}\otimes\mathbf{E}+\ldots \quad\Leftrightarrow\quad \{\{\mathbf{A},\mathbf{B}\},\{\mathbf{C},\mathbf{D},\mathbf{E}\},\ldots\}$$

The buffer is replayed every time an action on a vector is needed:

$$\left(\alpha\left[\mathbf{A}\otimes\mathbf{B}\otimes\ldots\right]+\beta\left[\mathbf{C}\otimes\mathbf{D}\otimes\ldots\right]+\ldots\right)\mathbf{x}=\alpha\left[\mathbf{A}\otimes\mathbf{B}\otimes\ldots\right]\mathbf{x}+\beta\left[\mathbf{C}\otimes\mathbf{D}\otimes\ldots\right]\mathbf{x}+\ldots$$

The same applies to pre- and post-multiplication:

$$\mathbf{P}_1\cdot\ldots\cdot\mathbf{P}_N\cdot\left[\mathbf{A}\otimes\mathbf{B}+\mathbf{C}\otimes\mathbf{D}\otimes\mathbf{E}+\ldots\right]\cdot\mathbf{Q}_1\cdot\ldots\cdot\mathbf{Q}_M$$

$$\Updownarrow$$

$$\{\mathbf{P}_1,\ldots,\mathbf{P}_N\}\{\{\mathbf{A},\mathbf{B}\},\{\mathbf{C},\mathbf{D},\mathbf{E}\},\ldots\}\{\mathbf{Q}_1,\ldots,\mathbf{Q}_M\}$$

A Matlab object that pretends to be a matrix.

Result: *massive reduction in memory and CPU time (can skip unit matrices).*

Wall clock time, polyadic rep	Wall clock time, explicit rep
0.37 ± 0.01 ms	0.88 ± 0.12 ms
1.8 ± 0.3 ms	Out of RAM
97 ± 14 ms	Out of RAM
0.21 ± 0.01 ms	0.05 ± 0.01 ms
2.1 ± 0.3 ms	11.4 ± 1.6 ms
105 ± 16 ms	Out of RAM

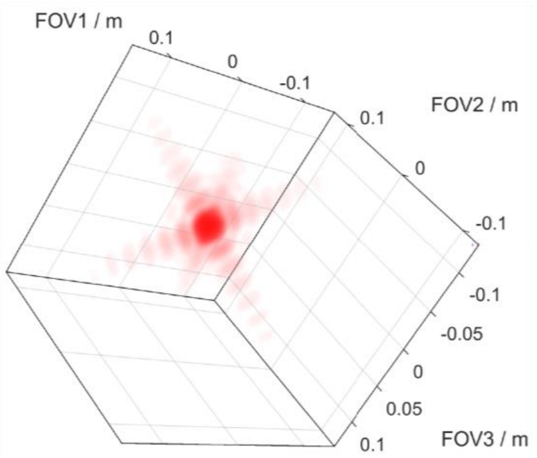

A. Allami, M.G. Concilio, P. Lally, I. Kuprov, *Science Advances*, 2019, 5(7), eeaw8962.

MRI simulation using polyadics

Result: arbitrary spatial dynamics with quantum mechanical description of spin.

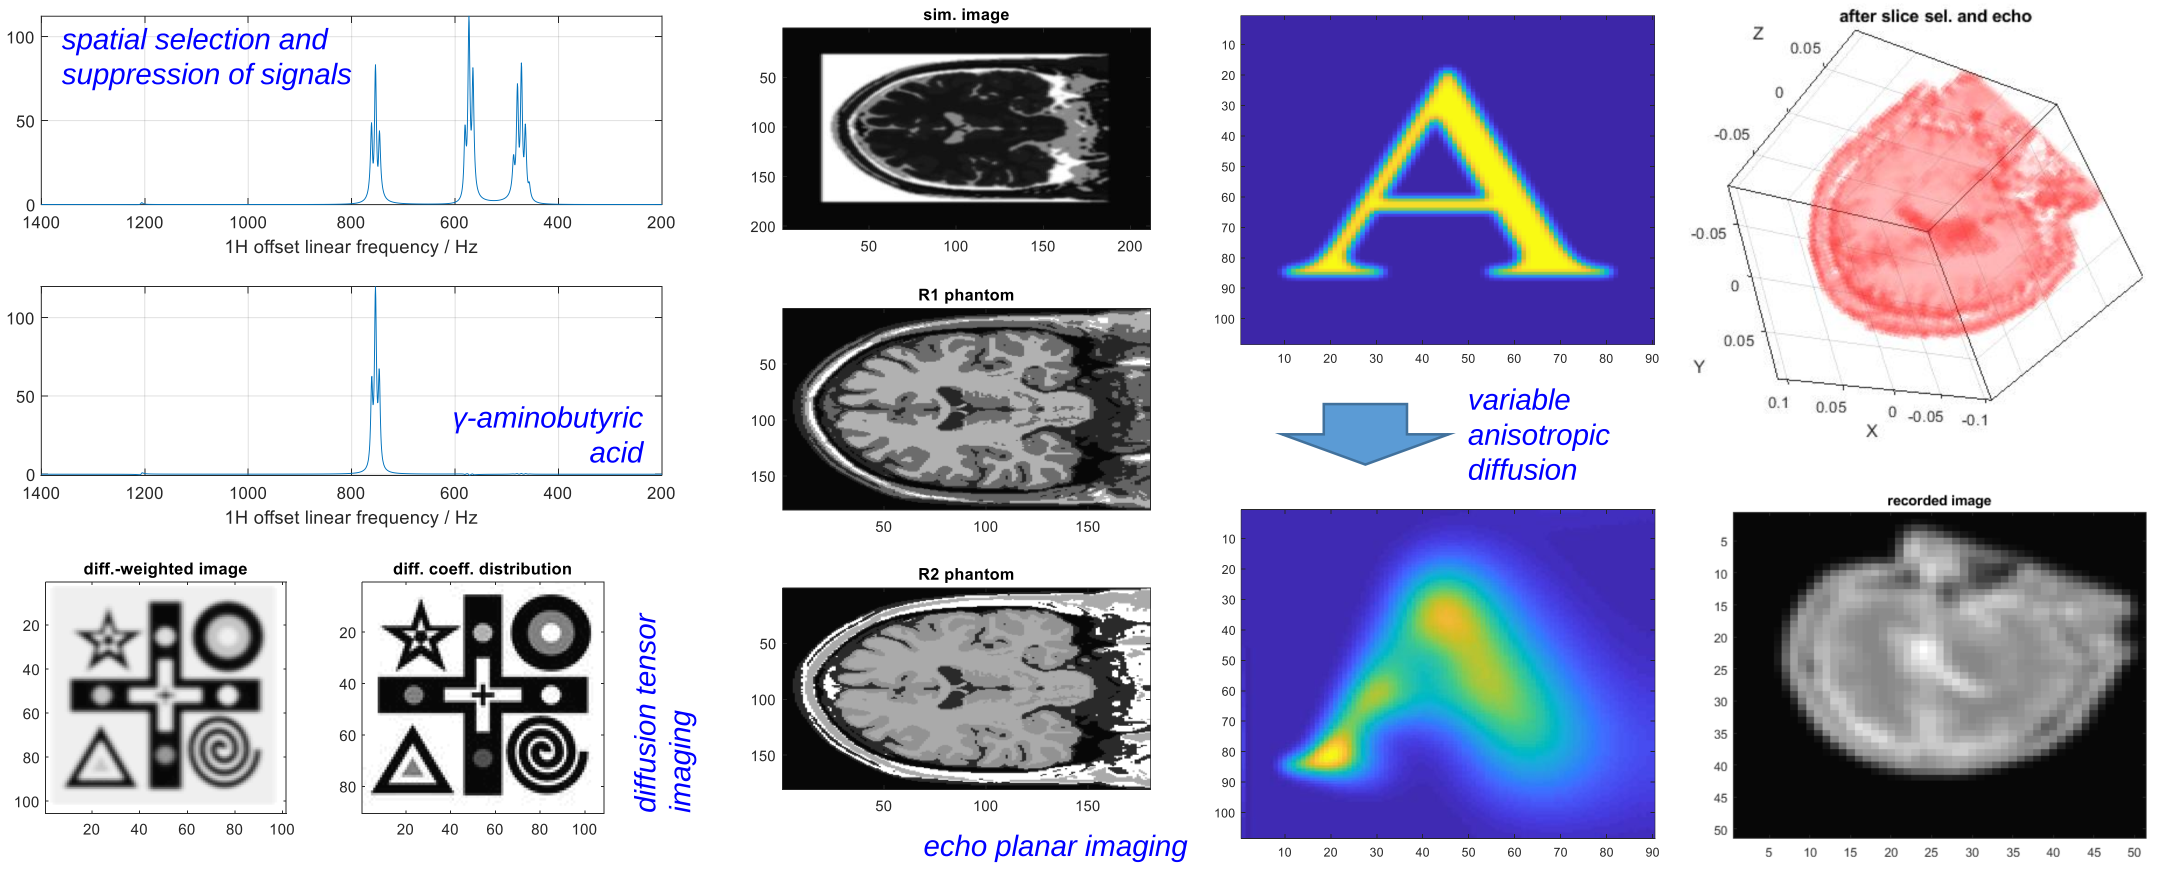

A. Allami, M.G. Concilio, P. Lally, I. Kuprov, *Science Advances*, 2019, 5(7), eeaw8962.

Spinach toolbox

- Magnetic resonance theory library for **large-scale** time-domain simulation work
- All types of magnetic resonance (NMR, EPR, MRI, DNP, PHIP, SQUID, *etc.*)
- Over 600 pages of docs and tutorials, over 100 real-life simulation examples
- Well-annotated open-source code, clear variable names, informative error messages
- Parallel processing, GPU support, tensor structured object support
- Over 50 developers and contributors, 12 years of full-time programming

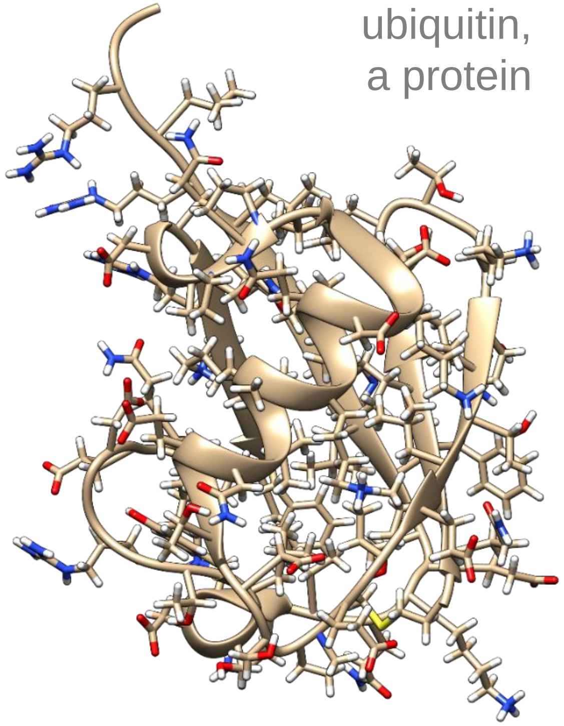

ubiquitin, a protein

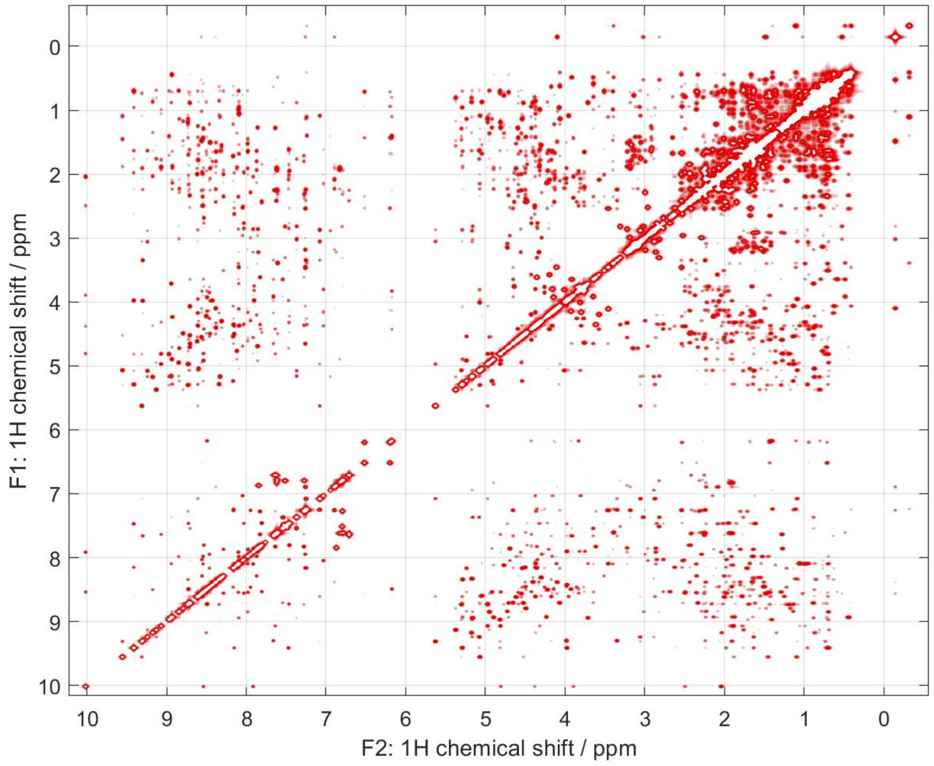

+ kinetics, diffusion, hydrodynamics, spatial encoding, off-resonance soft pulses, etc.

Parallelisation strategies

- Some strategies are trivial (parfor over ensembles, *etc.*)
- SPMD permits more sophisticated techniques

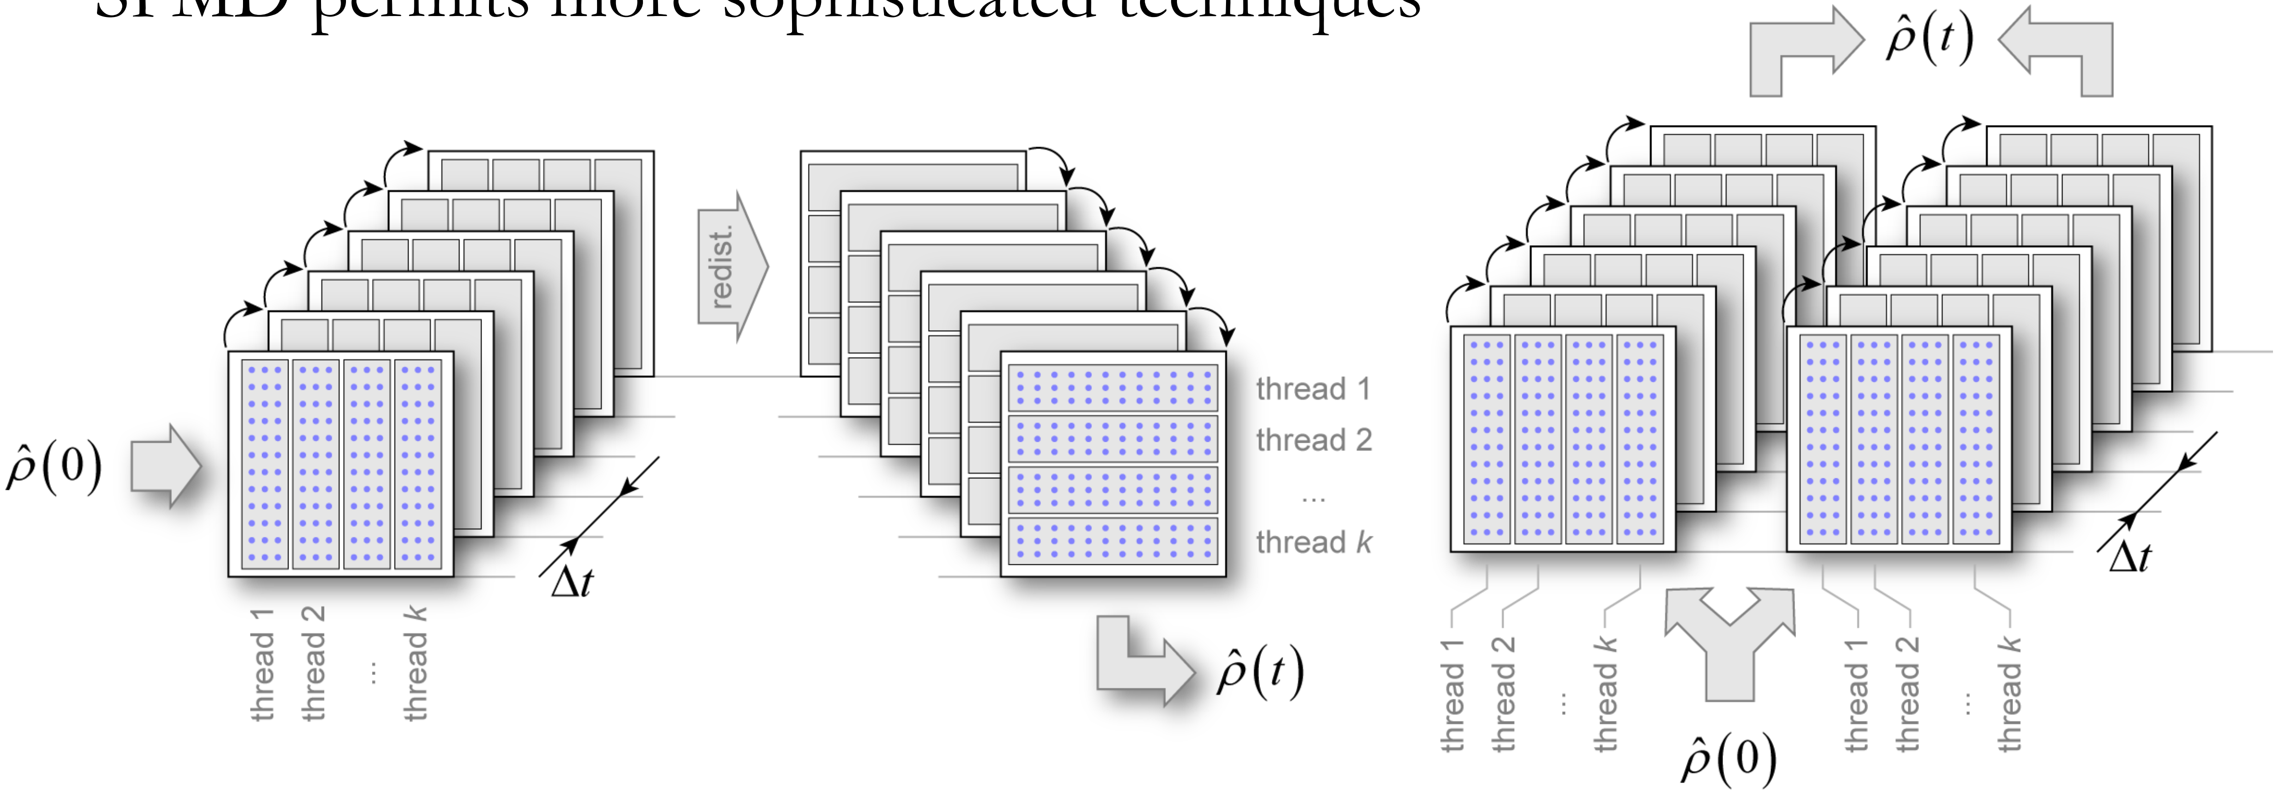

L.J. Edwards, I. Kuprov, *Journal of Chemical Physics*, 2012, 136(4), 044108.

Nested switchable parallelisation

Extended parfor syntax passes parallelisation opportunities down:

```
% Powder averaging loop
parfor (n=1:numel(weights),nworkers)

    % Localise the parameter array
    localpar=parameters;

    % Get the full Hamiltonian at the current orientation
    H=I+orientation(Q,[alphas(n) betas(n) gammas(n)]); H=(H+H')/2;
```

Table 1. Scaling behaviour of the parallel propagation algorithms.

Number of CPU cores	Time steps per wall clock second		
	Algorithm A (observable)	Algorithm B1 (final state)	Algorithm B2 (final state)
1	1.2	3.1	1.9
2	2.5	6.2	3.7
4	4.9	12.5	7.4
8	9.9	25.1	14.8
16	18.9	49.7	29.8
32	29.4	72.7	48.1
64	48.4	112.8	78.6
128	68.0	151.7	110.9

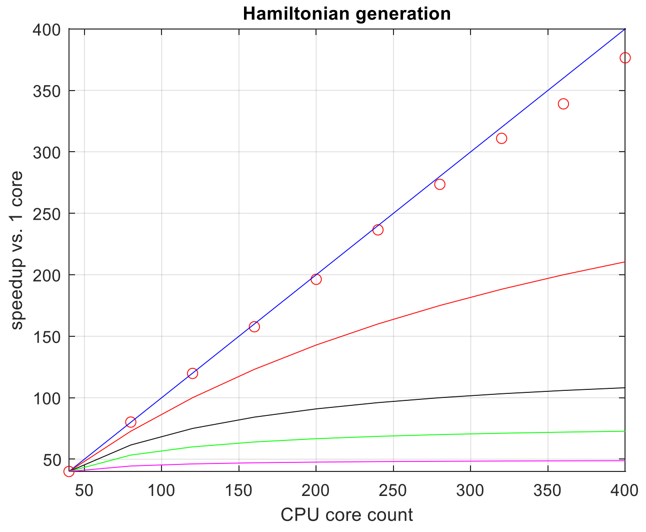

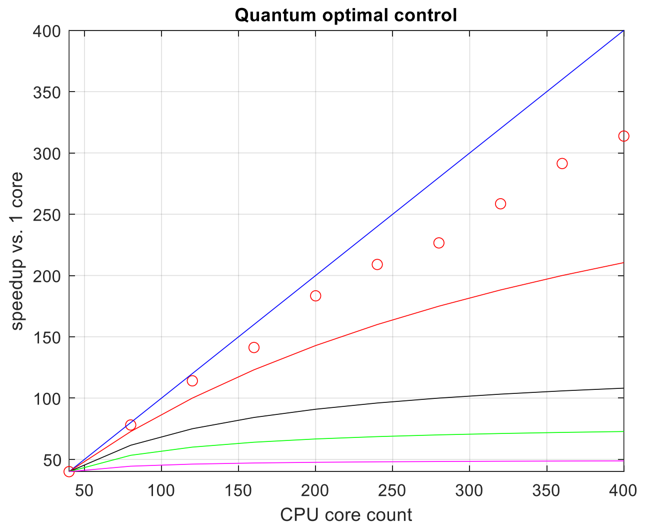

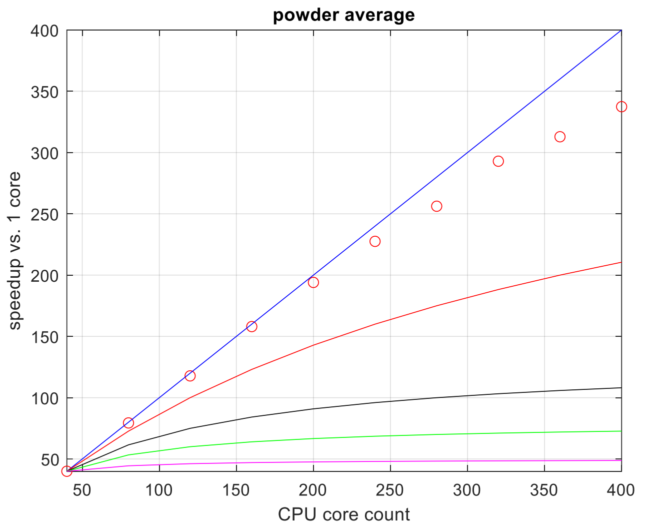

Writing a MATLAB package

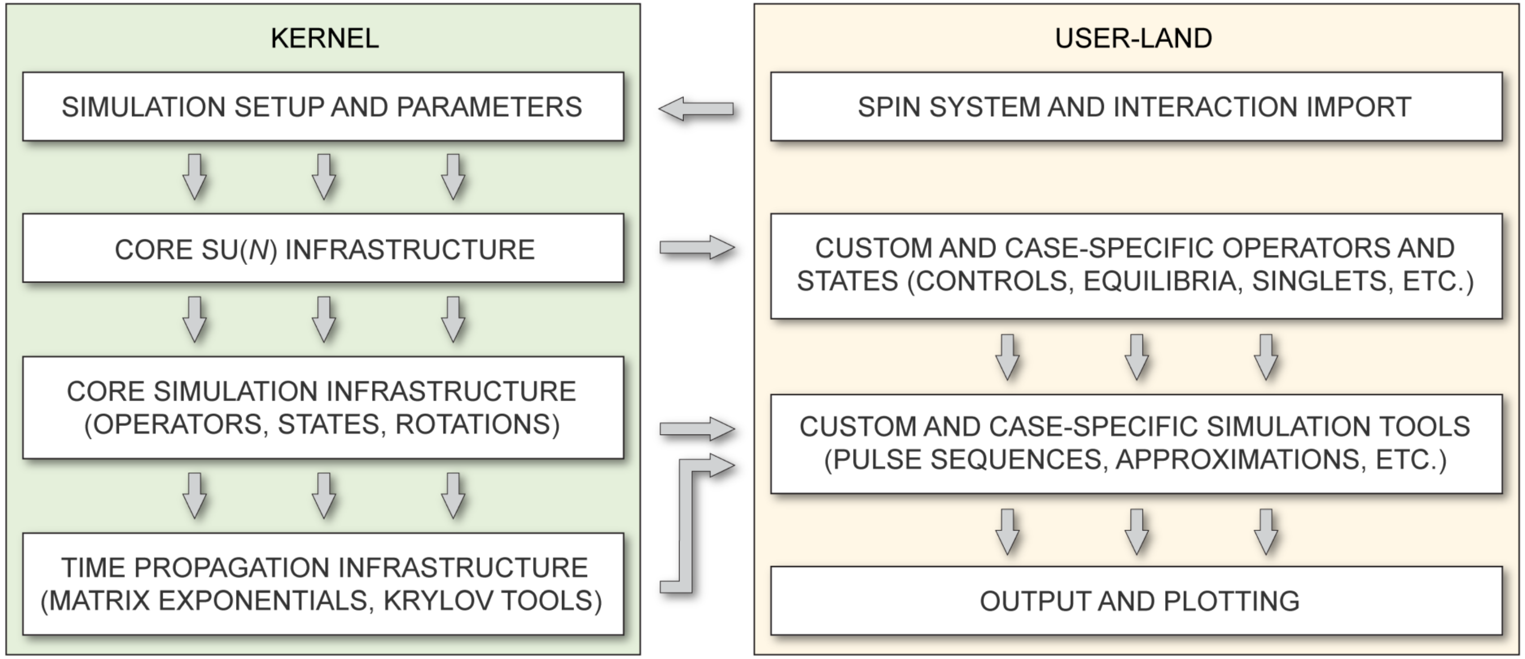

Why MATLAB?

1. *Brain time vs CPU time*
2. *Ease of parallelisation*
3. *Ease of GPU deployment*
4. *Code sustainability*
5. *Audience (chemists!)*
6. *Low specific torment*

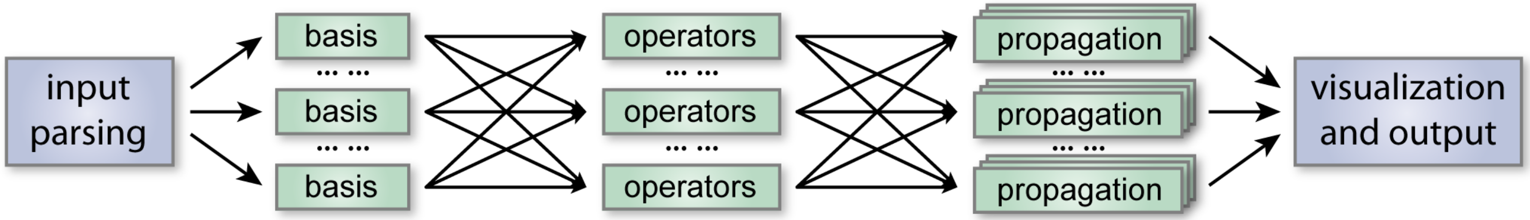

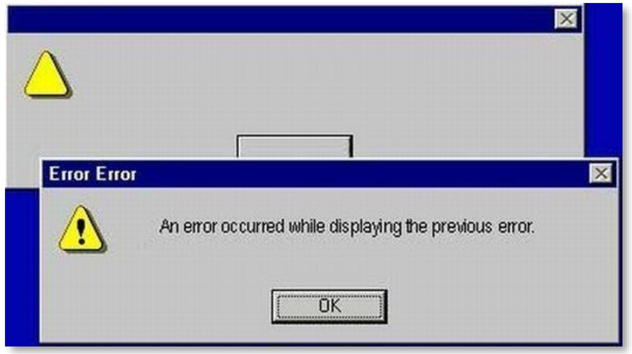

H.J. Hogben, M. Krzystyniak, G.T.P. Charnock, P.J. Hore, I. Kuprov, *Journal of Magnetic Resonance*, 2011, 208(2), 179-194.

Phosphine ligand chemistry

A di-tert-butylphosphine… 22 spins, meaning matrix dimension 2^{22} (very sparse).

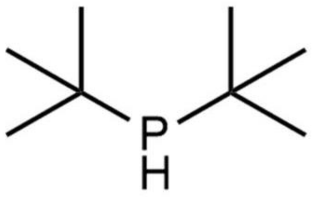

Confused chemists – the pattern of NMR peaks is emphatically not textbook.

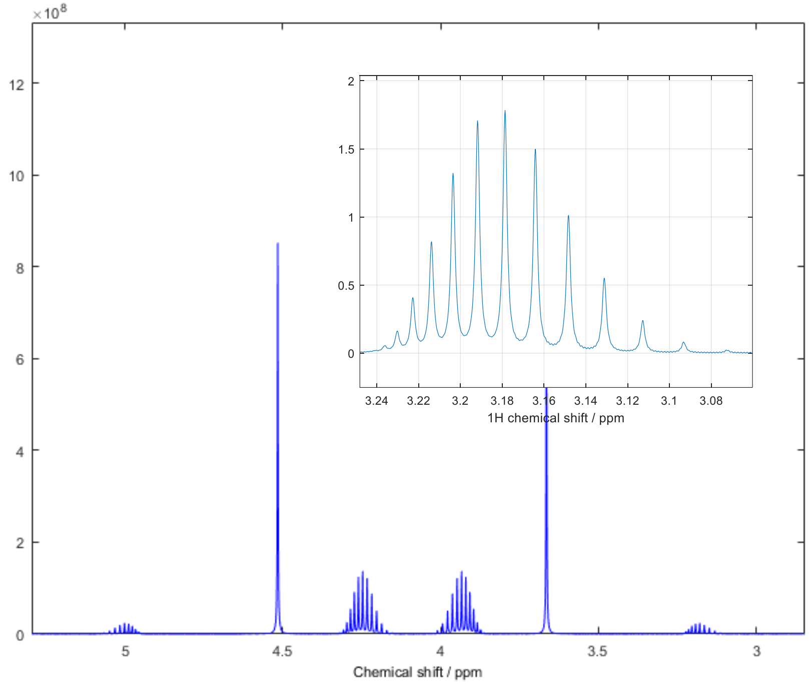

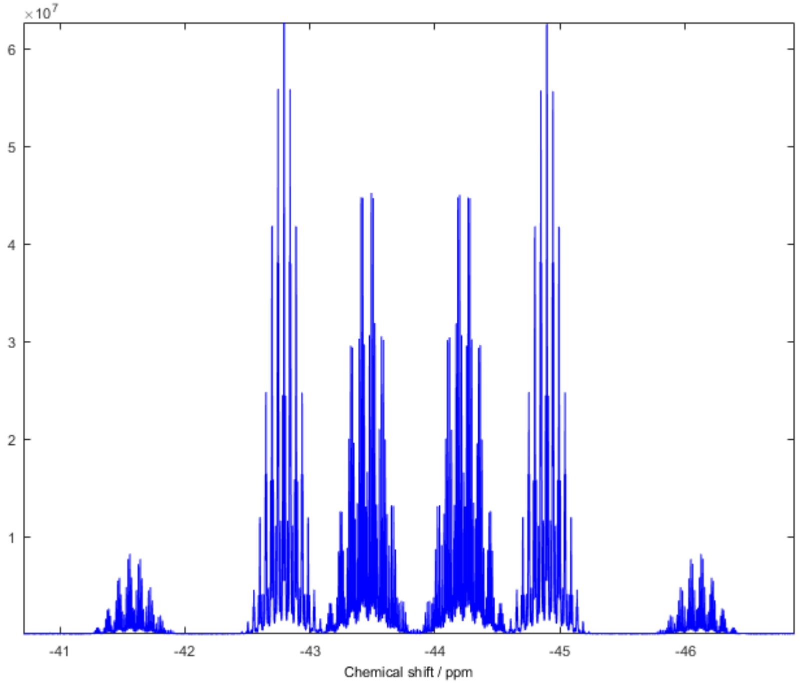

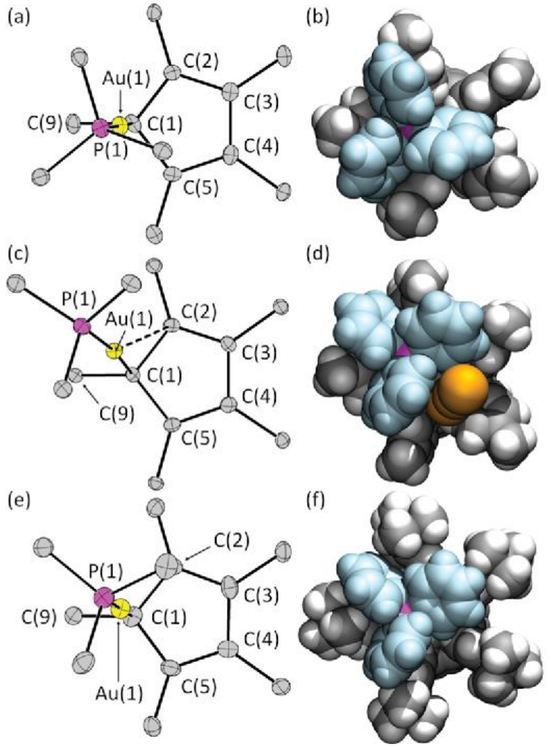

SPMD across 512 cores: 15 minutes, perfect match to the experiment.

Industrial gold extraction, *etc.*

Complex sparse GPU arithmetic

haemoglobin (Fe)

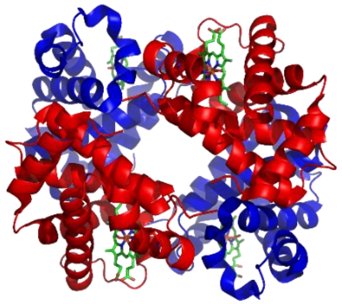

superoxide dismutase (Cu)

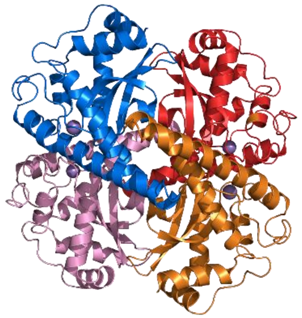

nitrile hydratase (Co)

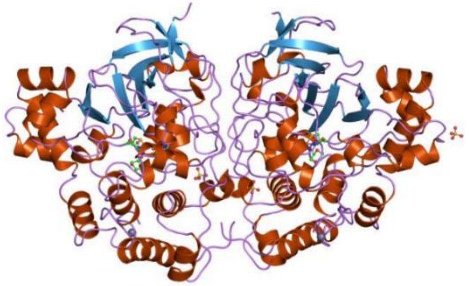

Metal locations in metalloproteins using Tikhonov regularisation:

(least squares gradient) = [2×FFT, 2×DGEMM]

Gradient explicit, but Hessian implicit (hello, Optimisation Toolbox):

(Hessian-times-vector) = [4×FFT, 4×DGEMM]

Quadruple 3D FFT of a 512×512×512 dataset - over 2 GB of data!

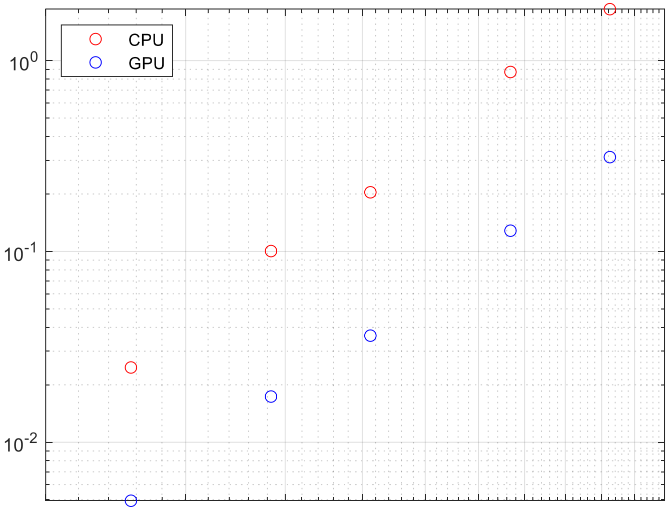

Complex sparse GPU arithmetic

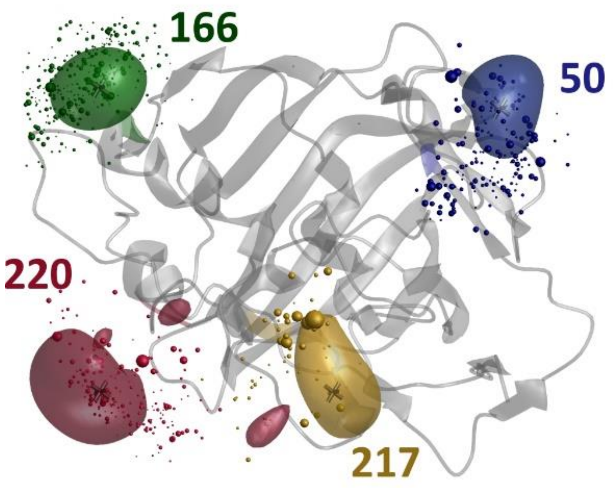

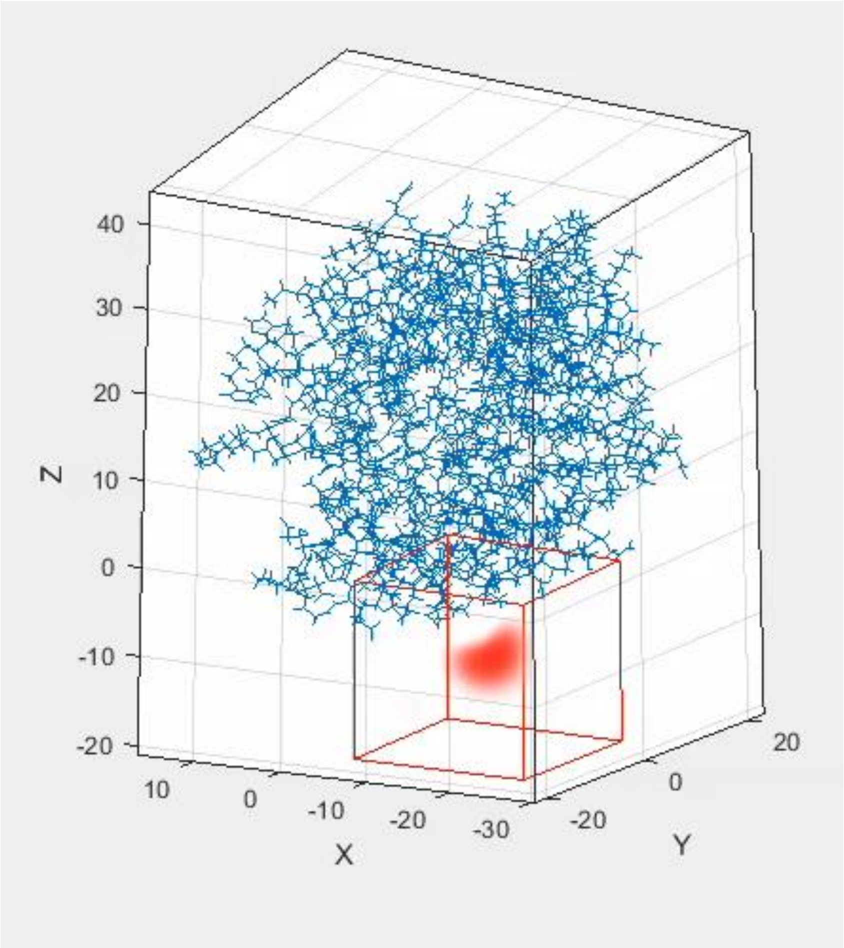

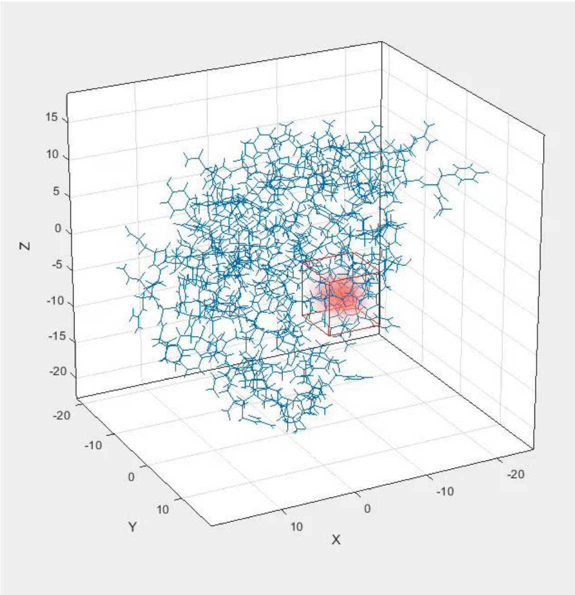

A module for *Spinach* library:

> run time: minutes

> good match to DEER data

> sensible match to MD data

> new use for old PCS data

> some structural insight

Relevant operations:

> *matrix exponentiation*

> *time propagation*

> *very large Fourier transforms*

> *element-wise operations*

E.A. Suturina, D. Häussinger, K. Zimmermann, L. Garbuio, M. Yulikov, G. Jeschke, I. Kuprov, *Chemical Science*, 2017, 8(4), 2751-2757.

Molecular distance measurement

Signal modulation by dipolar coupling:

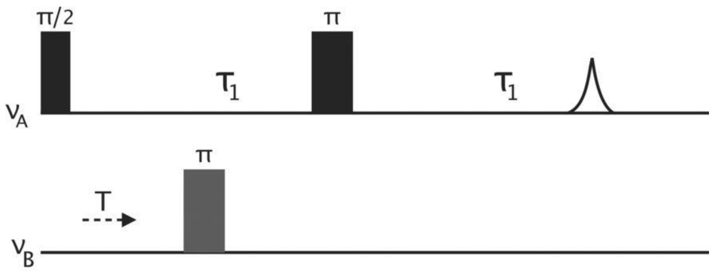

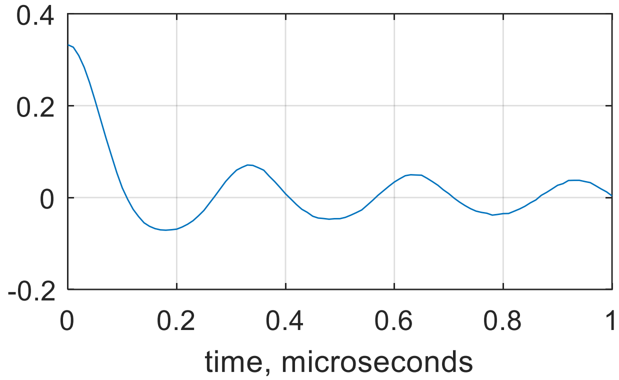

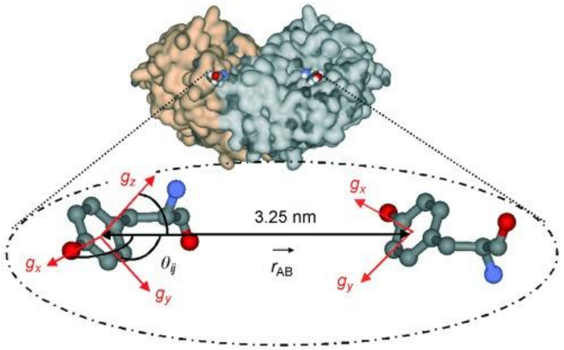

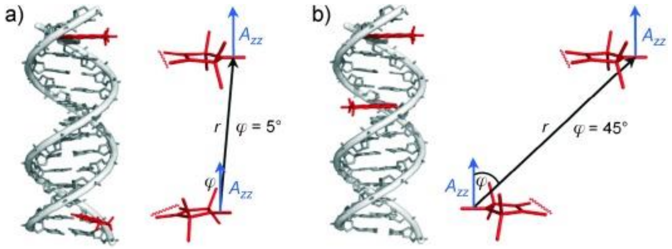

Straightforward theory for two point spin ½ particles:

$$\gamma(r,t) \sim \frac{1}{\sqrt{Dt}} \cdot \mathrm{trig}(Dt) \cdot \mathrm{fresnel}\left(\sqrt{Dt}\right) \qquad D \sim \frac{1}{r^3}$$

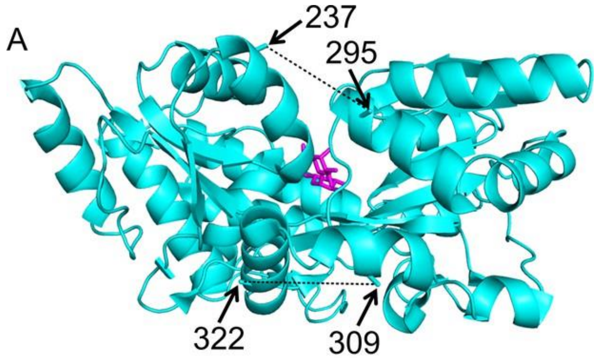

Can extract distance *distributions*:

$$DEER(t) = \int_0^\infty p(r)\gamma(r,t)\,dr$$

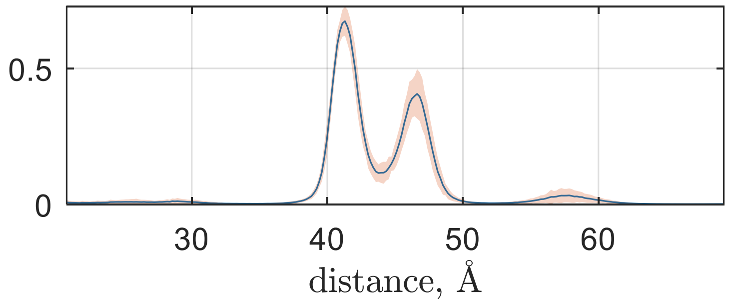

Inverse QM with neural nets

Neural networks are surprisingly good at getting distance distributions:

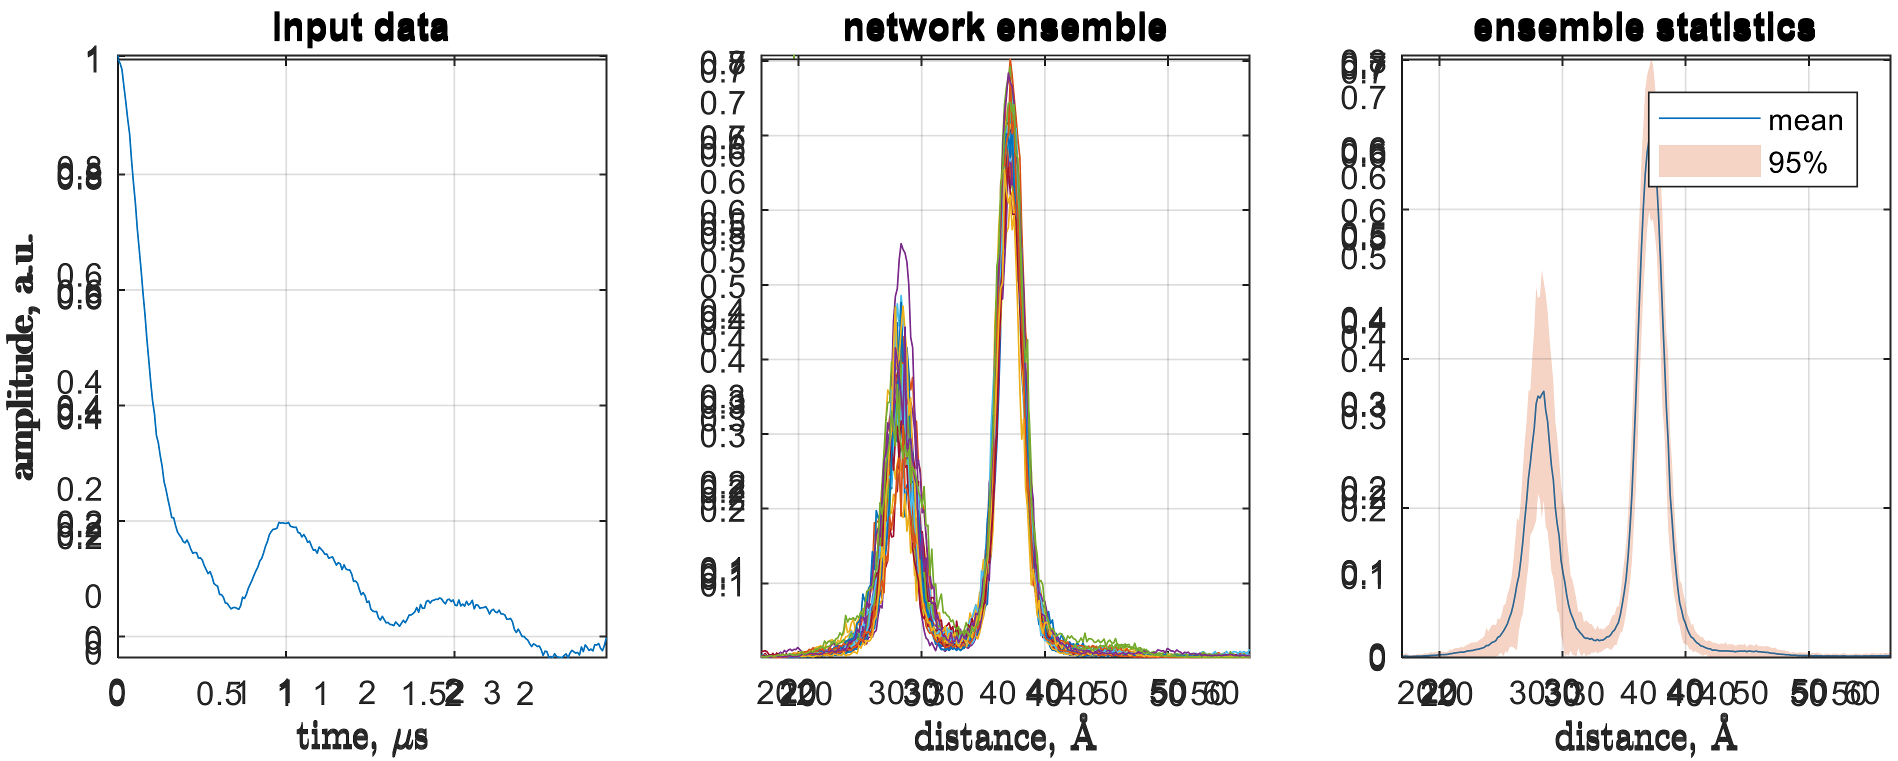

10 hours of unattended training vs. 10 years of programming!

Philosophical matters…

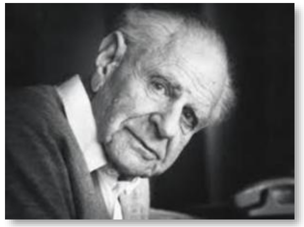

A black box neural network is not compatible with Descartes/Popper framework of science!

issues: *interpretability* and *trust*

We managed to find out how a two-layer DEERNet works…

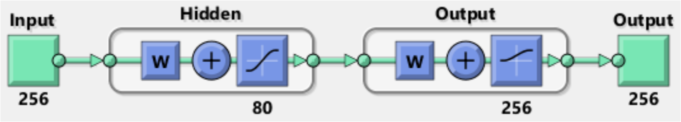

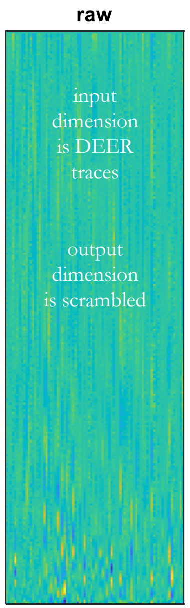

The net spontaneously evolved:

1. A bandpass filter *vs* the noise
2. A notch filter *vs* the baseline
3. A frequency axis rectifier from cubic to linear: $\omega \sim 1/r^3$ in DEER, but the plot is linear.

All packed into a single linear (!) transform in one (!!) layer.

Jake Amey

Steven Worswick

Summary

- Simulations in spin physics are stupidly hard!
- They can be done with some tensor algebra tricks…
- …but you need a language that understands tensor structures!
- …which chemists know how to use
- …that is compatible with version control
- …and of which there is only one dialect.

Acknowledgements

Jake Amey

Steven Worswick

Gunnar Jeschke

Liza Suturina

NVIDIA

The Leverhulme Trust

ETH zürich

MathWorks

Jos Martin

Edric Ellis

Joss Knight

Heiko Weichelt

Alison Eele

Ben Tordoff