



Science and
Technology
Facilities Council

Scientific Computing

Density-functional theory

Adam J. Jackson

Theoretical and Computational Physics Group,
Scientific Computing, STFC Rutherford Appleton Laboratory

Outline

- Atomistic black boxes
- Peeking inside the black box: a bit of theory
- What can DFT do for me?
- Getting started: practical considerations

Atomistic simulations

In atomistic simulations, the fundamental modelling units are *atoms*. As opposed to:

- Finite-element analysis (typically divide by small volumes of solid or liquid)
- Coarse-grained simulations (protein, polymer units)
- Subatomic particle physics


We can loosely group atomistic simulations into:

- Interatomic force fields (parametrised models, atom positions)
- Electronic structure (solve behaviour of electrons)

A few electronic structure approaches exist: while simplified semi-empirical and tight-binding methods have their uses, the “workhorse” DFT is *ab initio*


An *ab initio* black box

“Ab initio” calculations are based on fundamental physics; they don’t need special information about your system. How do we interact with these methods?



The reality of *ab initio* methods


There seem to be quite a lot of parameters to play with...



(Photo: Franz Schuier, CC-by 2.0)

Precision and accuracy

High precision,
high accuracy



High precision,
low accuracy



Low precision,
high accuracy



Low precision,
low accuracy



- Some parameters concern numerical precision and can be systematically improved towards a “correct” result
 - E.g. numerical grid sizes
 - Usually a compromise with computational cost
- Other parameters will affect overall accuracy
 - E.g. physics approximations
 - effects are less predictable
 - May require more knowledge of problem, testing, system-dependence

Inside the black box: electronic structure methods

$$E\Psi = \hat{\mathcal{H}}\Psi$$

- Looking to solve the time-independent Schrödinger equation (above)
- The nuclei define some Coulomb interactions in the Hamiltonian operator
- Solve eigenvalue problem for energy states (eigenvalues) and electronic wavefunctions (eigenvectors)
- For one nucleus and one electron, solutions give hydrogen-like orbitals
- Born–Oppenheimer approximation: nuclei are slow compared to electrons

Inside the black box: electronic structure methods

$$E\Psi = \hat{\mathcal{H}}\Psi$$

- Things get more tricky when dealing with multiple electrons
 - Work with a set of one-electron wavefunctions to keep maths tractable (Hartree approximation)
 - Let electrons respond to a “mean field” of other electrons
- As well as Coulombic interactions, electrons interact through an anti-symmetry property related to Pauli exclusion principle

$$\Psi(\dots, \mathbf{r}_1, \mathbf{r}_2, \dots) = -\Psi(\dots, \mathbf{r}_2, \mathbf{r}_1, \dots)$$

- Putting this together yields the Hartree–Fock method: solve iteratively by minimising energy sum over occupied orbitals.

$$\epsilon_a \psi_a = \left[h + \sum_{b \neq a} \mathcal{I}_b - \sum_{b \neq a} \mathcal{X}_b \right] \psi_a$$

Inside the black box:

Problems with Hartree–Fock

$$\epsilon_a \psi_a = \left[h + \sum_{b \neq a} \mathcal{I}_b - \sum_{b \neq a} \mathcal{X}_b \right] \psi_a$$

- Somewhat expensive $O(n^4)$
- Biggest problem is the neglect of *correlation*: mean-field sum over other electrons assumes they are independent
- Various “post-HF” methods address this by building further calculations on top of a HF starting-point. They are expensive and have their uses, but for most of the applications discussed in this lecture DFT is preferred.
- Alternatively we can work towards a similar Hamiltonian including corrections:

$$\mathcal{H}\Psi = \left(-\frac{1}{2} \nabla^2 + \mathcal{U}_{\text{ext}} + \mathcal{U}_{\text{H}} + \mathcal{U}_{\text{X}} + \mathcal{U}_{\text{C}} \right) \Psi = E\Psi$$

Density functionals

A landmark paper by Hohenberg & Kohn (1964) proves that every external potential (i.e. arrangement of nuclei) yields a unique electron density field $\rho(\mathbf{r})$.

- Therefore each density defines a unique external potential
- Therefore density defines the Hamiltonian and all ground-state properties

$$E, E', \dots = F[\rho(\mathbf{r})]$$

- Operators like F are a function on a function: known as a *functional*
- The good news: $\rho(\mathbf{r})$ is a 3-D field, however many electrons we have
- The bad news: H–K tell us nothing about the form of these functionals

Kohn–Sham DFT

The Kohn–Sham method introduces density functionals while still using single-electron “orbitals” for some purposes.

$$\mathcal{H} = -\frac{1}{2}\nabla^2 + \mathcal{U}_{\text{ext}} + \mathcal{U}_{\text{H}} + \mathcal{U}_{\text{XC}}$$

$$E = -\frac{1}{2} \sum_i \phi_i^*(\mathbf{r}) \nabla^2 \phi_i(\mathbf{r}) + E_{\text{ext}}(\rho(\mathbf{r})) + E_{\text{H}}(\rho(\mathbf{r})) + E_{\text{XC}}(\rho(\mathbf{r}))$$

These orbitals ϕ_i are variously described “auxiliary”, “fictitious” or “non-interacting”. They are summed to create a density and differentiated to obtain a kinetic energy.

$E_{\text{XC}}(\rho(\mathbf{r}))$ is the “exchange-correlation functional”

XC functionals

The XC functional's job is to correct for all the other errors and give correct results!

Based on local density and derivative terms... as we move up the ladder we gain more parameters and make more choices

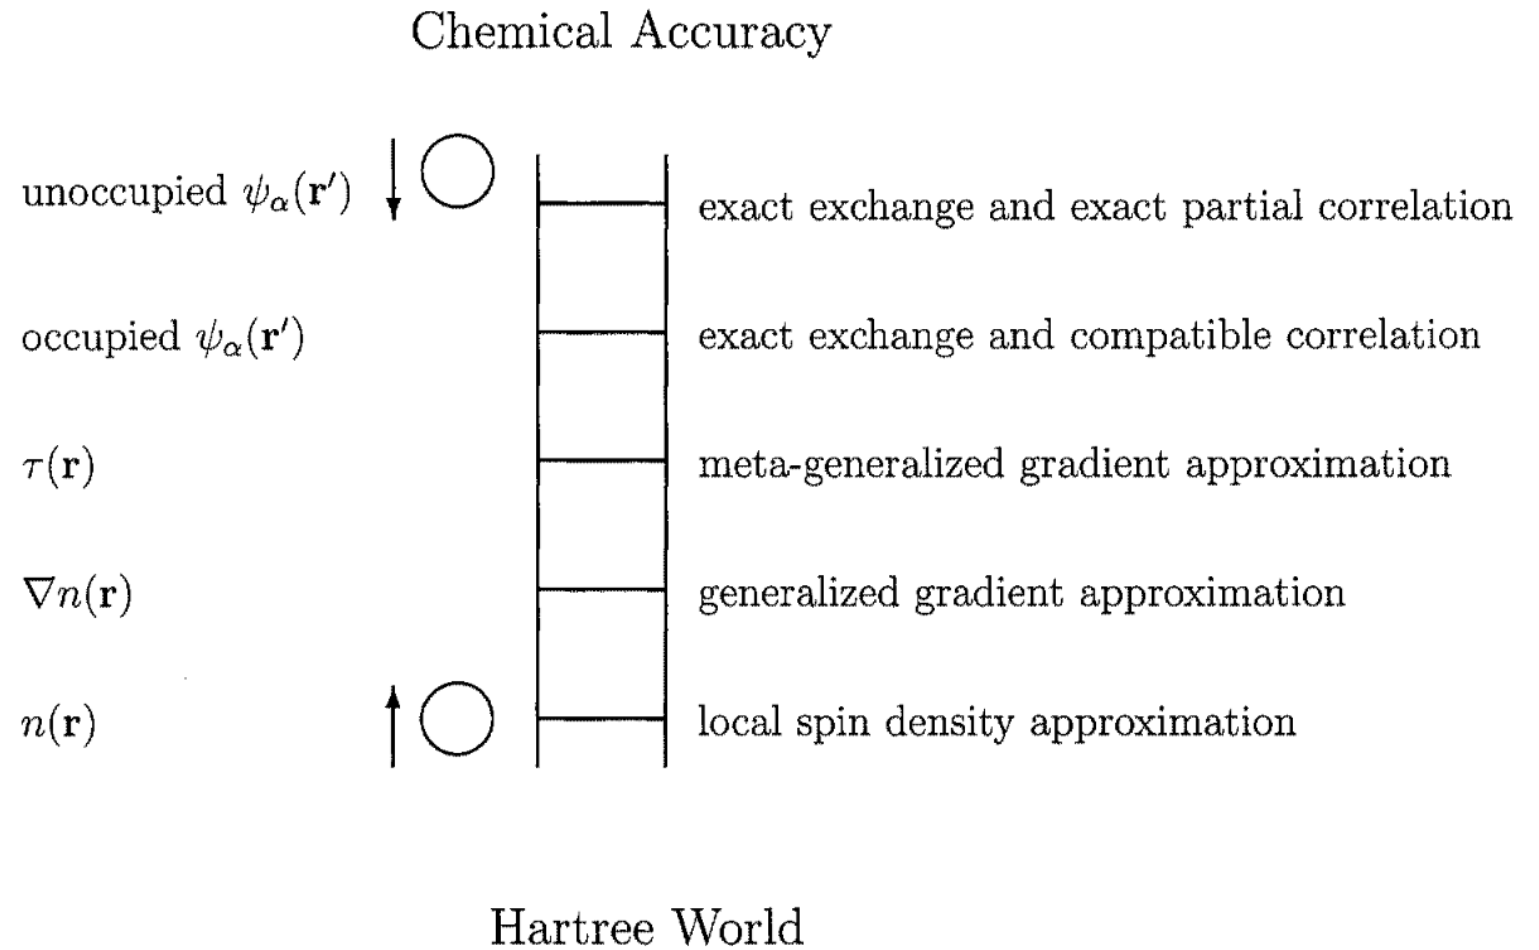

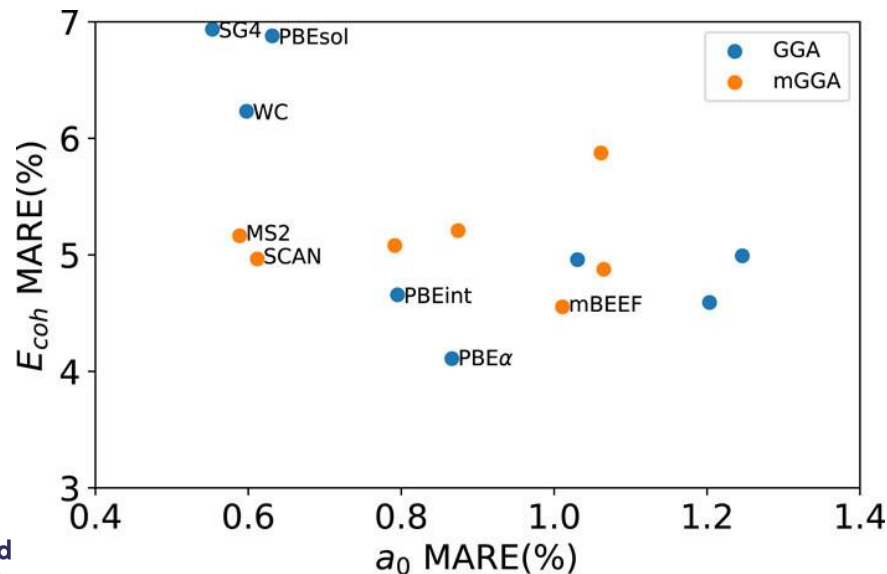


FIGURE 1. Jacob's ladder of density functional approximations. Any resemblance to the Tower of Babel is purely coincidental. Also shown are angels in the spherical approximation, ascending and descending. Users are free to choose the rungs appropriate to their accuracy requirements and computational resources. However, at present their safety can be guaranteed only on the two lowest rungs.

XC functionals

XC functional design can be seen as a parameter optimisation

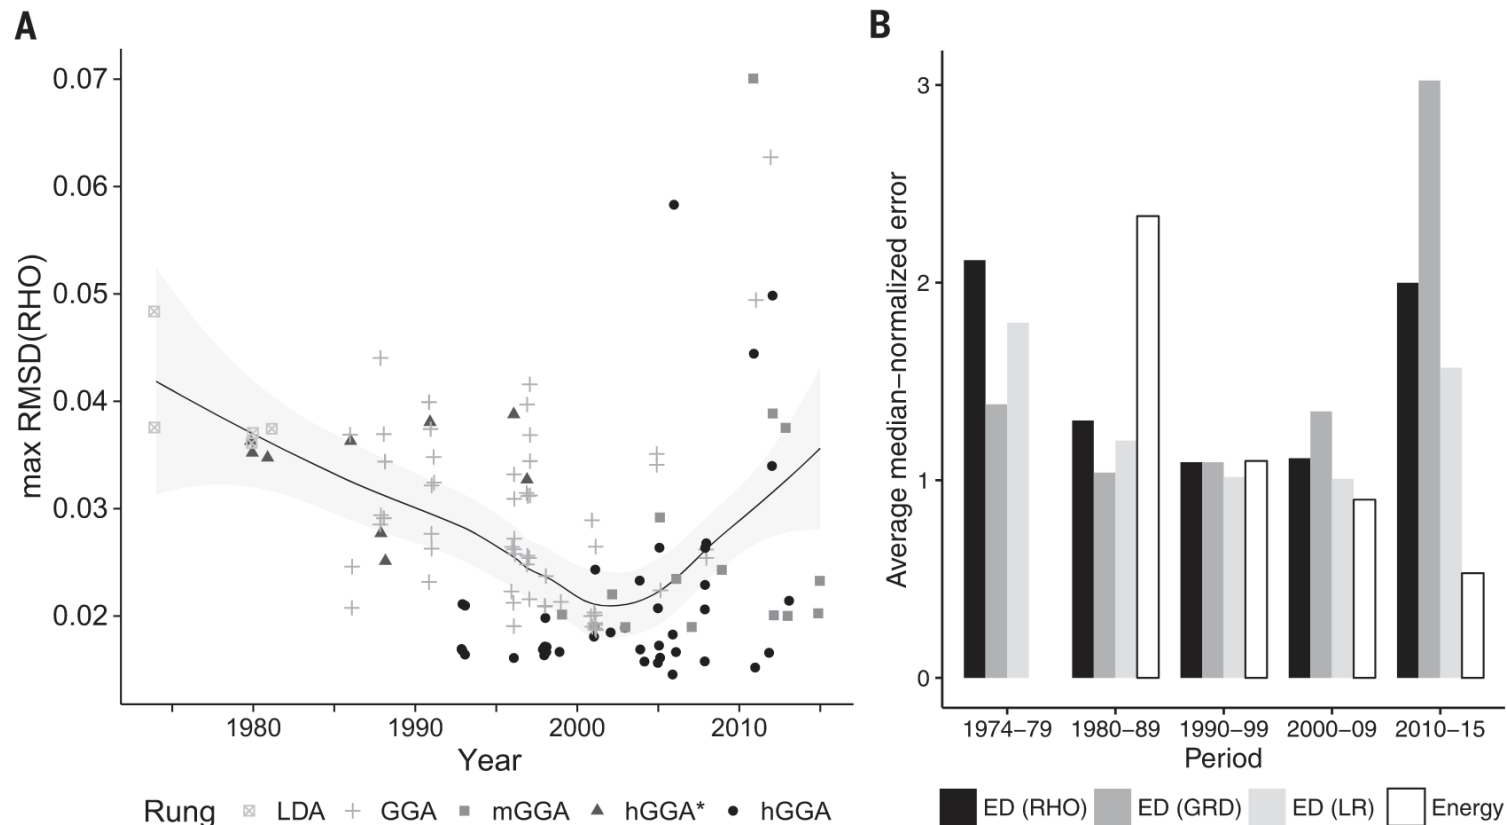
e.g. recent paper “*What is the optimal mGGA exchange functional for solids?*”
P. Kovács, F. Tran, P. Blaha, G. K. H. Madsen (2022) *J. Chem. Phys.* **157**, 094110



XC functionals

But we should also beware of overfitting to certain properties: controversial paper
“*Density functional theory is straying from the path toward the exact functional*”
finds energy predictions getting better over time, but density getting worse!

M G. Medvedev I. S. Bushmarinov, J. Sun, J. P. Perdew, K. A. Lyssenko (2017) *Science* **355** (6320) 49




Periodic systems: Bloch's theorem

Infinite crystals are described with periodic boundary conditions.

Long-range wavefunctions can be described using Bloch waves: product of a periodic function (unit-cell-sized!) and plane wave.

With enough \mathbf{k} -points, properties should converge...



Basis sets

How do we actually represent Ψ or ϕ ?

- It's difficult to solve for an unknown *function*
- It's easy to solve for a linear combination of components
- Represent unknown function as a linear combination of known functions:
a basis set

Many different basis sets are used in DFT; technical choice.

Basis sets: pros/cons

Regular grid points



SPARC

Local orbitals



NWCHEM
HIGH-PERFORMANCE COMPUTATIONAL
CHEMISTRY SOFTWARE

Plane waves



QUANTUM ESPRESSO



That's a lot of options!

There are a lot of DFT codes and some questions about the accuracy of some technical choices: basis sets, pseudopotentials etc.

So a study was performed across many different codes, performing equivalent equation-of-state calculations for solids.

Result: they agree pretty well!
(for more recent code versions)
(if you know what you are doing)



The screenshot shows the top of a Science journal article page. The Science logo is in the top left. Navigation links include 'Current Issue', 'First release papers', 'Archive', 'About', and a 'Submit manuscript' button. The article title is 'Reproducibility in density functional theory calculations of solids'. Below the title is a long list of authors. At the bottom of the author list is a 'fewer' button and a link to 'Authors Info & Affiliations'. The footer of the page shows the journal name 'SCIENCE', the date '25 Mar 2016', the volume and issue 'Vol 351, Issue 6280', and the DOI '10.1126/science.aad3000'.

Science

Current Issue First release papers Archive About ▼ Submit manuscript

RESEARCH ARTICLE

Reproducibility in density functional theory calculations of solids

KURT LEJAEGERE, GUSTAV BIHLMAYER, TORBJÖRN BJÖRKMAN, PETER BLAHA, STEFAN BLÜGEL, VOLKER BLUM, DAMIEN CALISTE, IVANO E. CASTELLI, STEWART J. CLARK, ANDREA DAL CORSO, STEFANO DE GIRONCOLI, THIERRY DEUTSCH, JOHN KAY DEWHURST, IGOR DI MARCO, CLAUDIA DRAXL, MARCIN DUŁAK, OLLE ERIKSSON, JOSÉ A. FLORES-LIVAS, KEVIN F. GARRITY, LUIGI GENOVESE, PAOLO GIANNOZZI, MATTEO GIAN TOMASSI, STEFAN GOEDECKER, XAVIER GONZE, OSCAR GRÄNÄS, E. K. U. GROSS, ANDRIS GULANS, FRANÇOIS GYGI, D. R. HAMANN, PHIL J. HASNIP, N. A. W. HOLZWARTH, DIANA IUȘAN, DOMINIK B. JOCHYM, FRANÇOIS JOLLET, DANIEL JONES, GEORG KRESSE, KLAUS KOEPERNIK, EMINE KÜÇÜKBENLİ, YAROSLAV O. KVASHNIN, INKA L. M. LOCHT, SVEN LUBECK, MARTIJN MARSMAN, NICOLA MARZARI, ULRIKE NITZSCHE, LARS NORDSTRÖM, TAISUKE OZAKI, LORENZO PAULATTO, CHRIS J. PICKARD, WARD POELMANS, MATT I. J. PROBERT, KEITH REFSON, MANUEL RICHTER, GIAN-MARCO RIGNANESE, SANTANU SAHA, MATTHIAS SCHEFFLER, MARTIN SCHLIPF, KARLHEINZ SCHWARZ, SANGEETA SHARMA, FRANCESCA TAVAZZA, PATRIK THUNSTRÖM, ALEXANDRE TKATCHENKO, MARC TORRENT, DAVID VANDERBILT, MICHIEL J. VAN SETTEN, VERONIQUE VAN SPEYBROECK, JOHN M. WILLS, JONATHAN R. YATES, GUO-XU ZHANG, AND STEFAAN COTTENIER

fewer Authors Info & Affiliations

SCIENCE • 25 Mar 2016 • Vol 351, Issue 6280 • DOI: 10.1126/science.aad3000

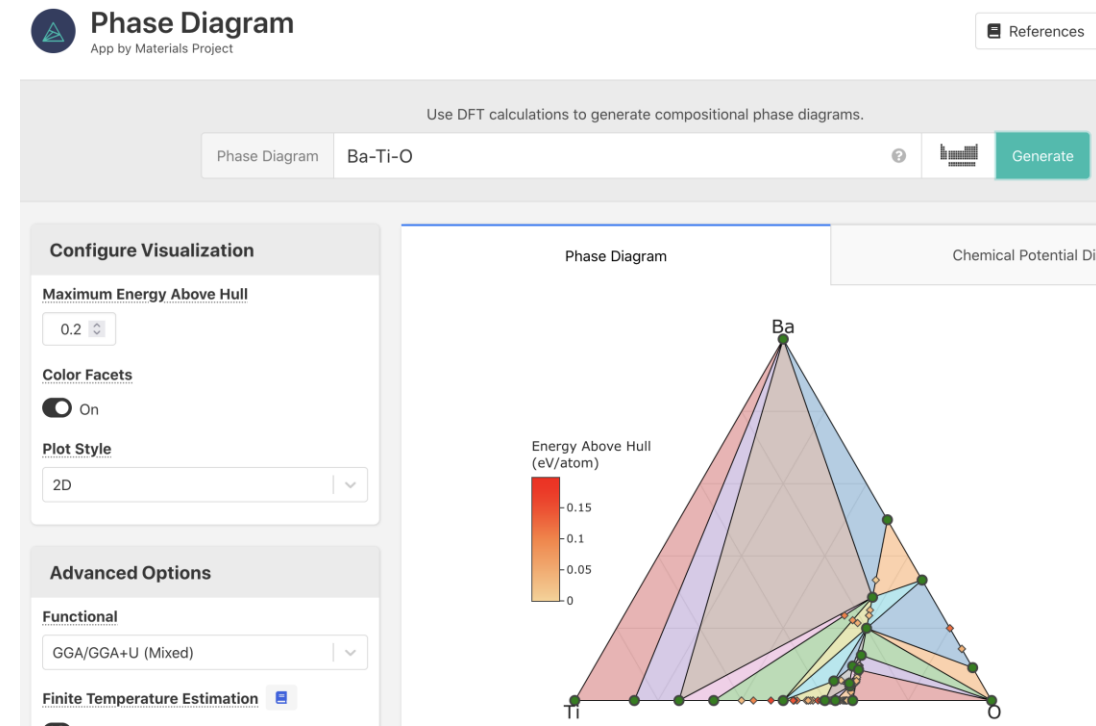
What can DFT do for me? Energies

The most direct output of a DFT calculation is a “total energy”.

- These are not directly useful: need to compare with other energies
- Be careful comparing energies from calculations with different settings
- **Do not compare total energies from different XC functionals**
- DFT energies are “athermal”; neglect vibrations, even ZPE

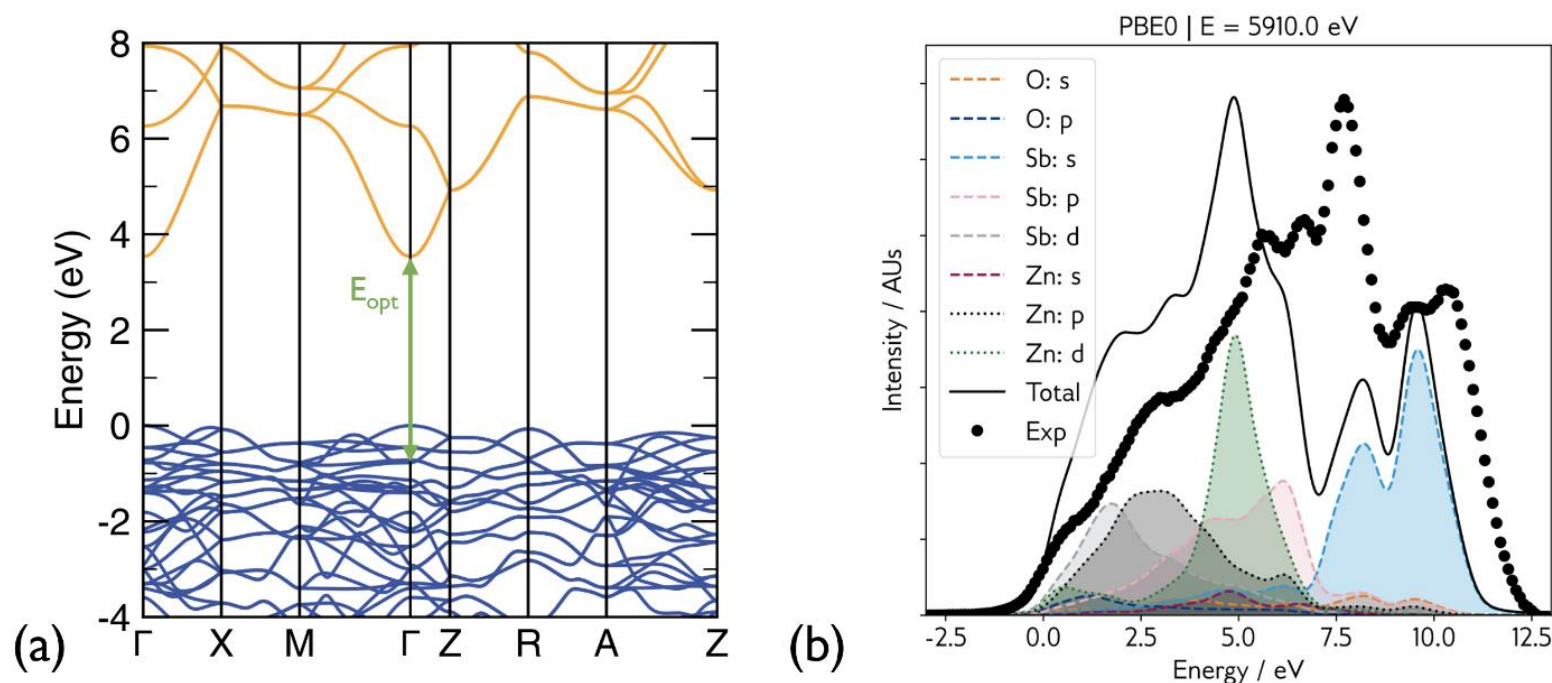
With a large set of energies, we can predict phase stability, choose between magnetic / ionic ordering, estimate heat of reaction...

Try <https://materialsproject.org> for a database of existing solid-state data



What can DFT do for me? Orbital energies

If we examine the shapes and energies of Kohn–Sham orbitals, we get pretty reasonable-looking information about electronic structure.




What can DFT do for me? Forces

The Hellmann–Feynman theorem allows forces to be calculated quite inexpensively.

Typically structures are optimised by adjusting positions and lattice parameters to minimise the energy; this is *much* more efficient when forces are available.

This information also helps explore the “potential energy landscape” searching for transition states (i.e. reaction barriers → kinetics)



*Ethanol molecule optimization
with LBFGS algorithm*


What can DFT do for me? Vibrations

In the *harmonic approximation*, vibrations are driven by linear spring-like interactions between each pair of molecules. This can be collected as a matrix of *force constants*.

$$E = T + E_0 + \cancel{\Phi_i u_i} + \frac{1}{2!} \Phi_{ij} u_i u_j + \frac{1}{3!} \Phi_{ijk} u_i u_j u_k + \dots$$
$$F = -\frac{\partial E}{\partial u_i} = -\cancel{\Phi_i} - \Phi_{ij} u_j - \frac{1}{2!} \cancel{\Phi_{ijk} u_j u_k} - \dots$$

Zero for
"relaxed"
structure



Anharmonic terms



Force constant matrix Φ_{ij} can be filled in a row at a time by making small displacements along each degree of freedom

What can DFT do for me? Phonons

For a periodic system the same principle applies, but we may need to use a large “supercell” to avoid self-interaction and get correct information about low-symmetry q -points



Interactive visualisations:


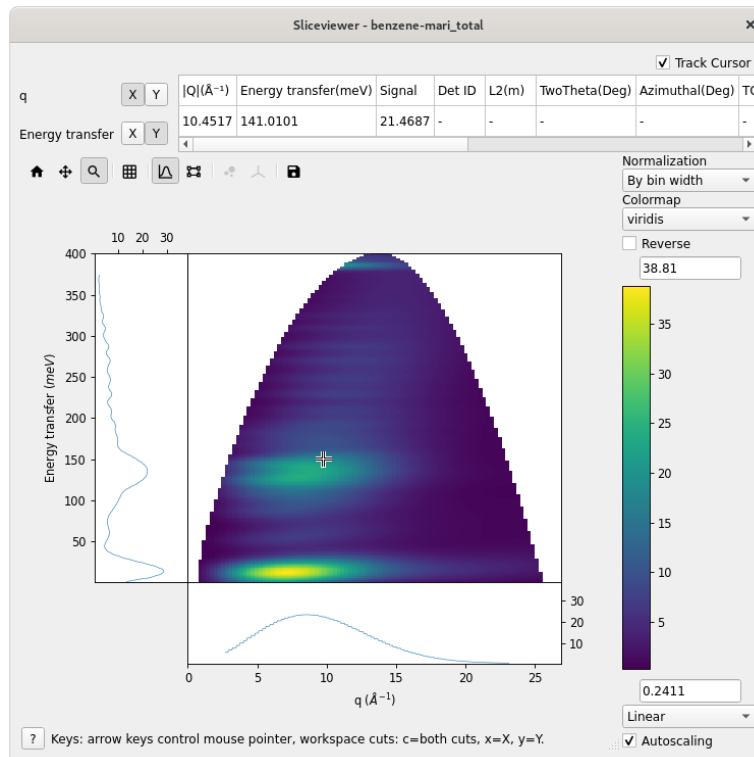
<http://henriquemiranda.github.io/phononwebsite/phonon.html>

What can I do with these phonons?

Codes developed at ISIS for simulation of inelastic neutron scattering spectra

AbINS: incoherent approximation, DOS-like analytic powder-averaging, instrumental resolution functions, multi-phonon scattering (Good for molecular crystals)

Euphonic: coherent scattering, single-crystal and numerical powder-averaging, simple broadening (Good for solids, inorganic)

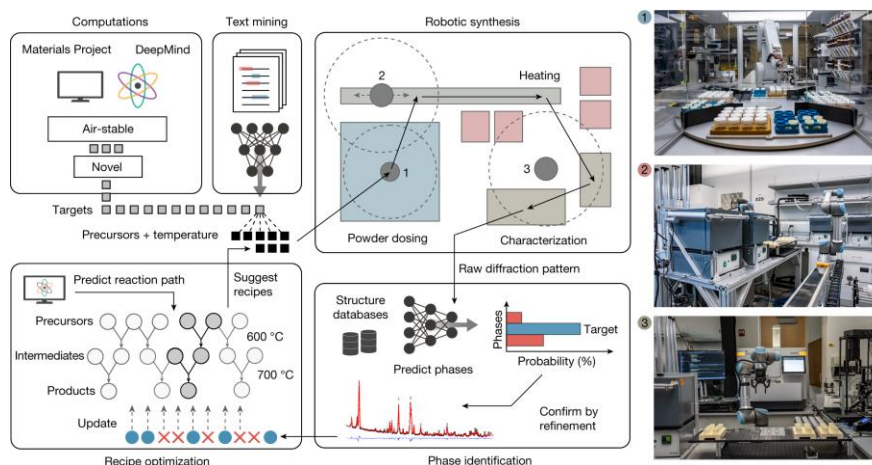


Ok, sounds great. What do I need?

- A suitable problem
- Software
- Computers
- Guidance

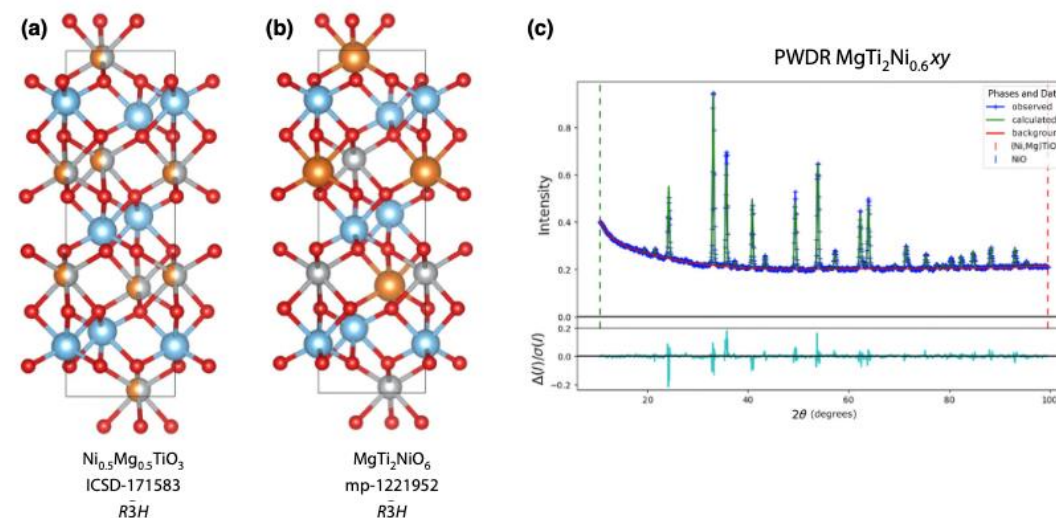
A suitable problem

- **Atomic positions** need to come from somewhere
 - Hypothesised structure
 - Diffraction data (x-ray or neutron)
 - Hydrogen may be missing
 - Average sites may differ from ground state
 - Fractional occupancy is a problem



“A-Lab” project, *Nature* paper claimed to discover 41 novel compounds

[DOI:10.1038/s41586-023-06734-w](https://doi.org/10.1038/s41586-023-06734-w)



Critics observed that XRD cannot distinguish the majority of these from known disordered phases

[DOI:10.1103/PRXEnergy.3.011002](https://doi.org/10.1103/PRXEnergy.3.011002)

A suitable problem

- **Atomic positions** need to come from somewhere
 - Hypothesised structure
 - Diffraction data (x-ray or neutron)
 - Hydrogen may be missing
 - Average sites may differ from ground state
 - Fractional occupancy is a problem
- “Global optimisation” algorithm
 - A more expensive problem...
- **Magnetic (dis)order** can break symmetry and expand the unit cell
- **Known properties** can guide selection of parameters such as XC functional
 - e.g. lattice parameters, optical bandgap, bulk modulus, ...

Software



There are many atomistic “codes”

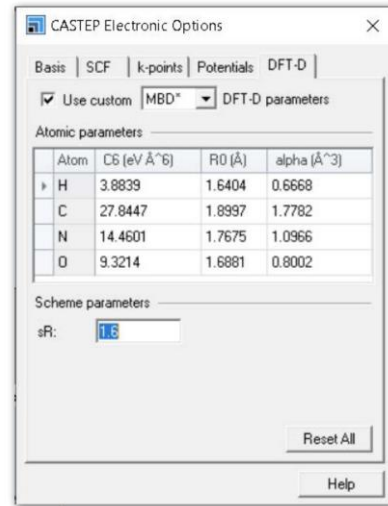
- Some based on different technical choices (*e.g. theory, basis set, parallelism*)
- Some emerging from different communities (*target systems, code interfaces*)
- A lot of “not invented here syndrome” and licensing issues...
(*open vs commercial vs evolved-from-a-1980s-PhD-project*)

Most of them are capable of decent precision, so in practice choose based on

- Required features (*periodicity, XC functionals, spin-orbit coupling, dispersion*)
- Computational performance (*scaling for small or large systems, GPU support*)
- Usability (*documentation quality, included tools, compatibility with external*)
- Access (*licensing, friendly developers, helpful post-docs*)

Software

A few codes come with graphical user interfaces that run on Windows (Materials Studio, GaussView) but this costs \$\$\$ and may limit available functionality.



VS

```
create mode 100644 ase/test/testdata/aims/DFP
create mode 100644 ase/test/testdata/vasp/vas
(gpaw) ajj-mbp-1:~/src/ase (master)$

task                spectral
spectral_task       dos
pdos_calculate_weights true
reuse               si.check
xc_functional        pbe
max_scf_cycles       100
basis_precision     fine
elec_energy_tol      1e-7
elec_eigenvalue_tol  1e-6
spectral_perc_extra_bands 50

UU- ~/Desktop/tmp/sumo-tutorial/si-dos.para
```

Most codes are developed for use in a Linux(-like) command-line environment. To write input files, automate groups of calculations and compile codes from source code, it is recommended to:

- Get a bit of familiarity with the Bash shell
- Learn how to use a text editor effectively
- (Optional) use Python for more advanced manipulation and plotting

Computers

Modern DFT codes will run on a laptop, a 750,000-core supercomputer, or most things in-between.

- Personal computers can generally handle tens of atoms with GGA functionals.
- A high-end workstation will go slightly bigger... or run a lot of those calculations more quickly and reliably.
- Academic clusters become necessary when going to larger unit cells, heavy elements, hybrid DFT, spin-orbit coupling...



UK national supercomputer: ARCHER2

2x cores only means 2x faster if your problem breaks down very well! Check scaling efficiency and test parallelism options.

Guidance


This was just an overview!

- Find out what codes people in your field are using for similar systems
- See if you can get a copy of their parameter files and understand what options were chosen and what implications they have
- Talk to experts to see if what you are trying to do is reasonable: DFT is not always a practical option
 - For larger simulation regions, the practical options might be *trained* against a set of smaller DFT calculations
 - Some systems are known to be tricky: highly-correlated materials, liquids and glasses, unstable phases, excited states...

Recap: how to read a DFT methods section

- “We used DFT” is not enough information to reproduce a calculation!
- Neither is “we used CASTEP”
- Most important information are the “accuracy parameters” defining approximations made: XC functional, spin treatment, dispersion...
- Also useful are “precision” details: basis set, pseudopotentials, **k**-point sampling, convergence thresholds
- Best practice: include parameter files in the SI!


High precision,
high accuracy



High precision,
low accuracy



Low precision,
high accuracy



Low precision,
low accuracy





Science and
Technology
Facilities Council

Scientific Computing

Thank you

scd.stfc.ac.uk

 [@SciComp_STFC](https://twitter.com/SciComp_STFC)