Lecture 4 · Approximating the Hamiltonian I: Hartree–Fock and the SCF

The mean-field idea, the self-consistent field, and what it leaves out

Table of contents

1 Where we are 2
2 The mean-field idea 2
2.1 Hartree's average field 2
2.2 The Slater determinant 3
2.3 The energy of a Slater determinant 3
2.4 The Hartree–Fock equations and the Fock operator 4
2.5 Closed shells: restricted Hartree–Fock 5
3 The Roothaan–Hall equations and the SCF loop 5
3.1 From integro-differential equations to matrices 5
3.2 Why it must be solved iteratively 6
3.3 Convergence, DIIS and the SCF table in your output 8
3.4 Restricted and unrestricted Hartree–Fock 9
4 Orbital energies and Koopmans' theorem 9
5 What Hartree–Fock misses: electron correlation 10
5.1 The correlation energy 10
5.2 Dynamic and static correlation 11
6 Ways forward 12
Key points 13
Study questions 14
Further reading 14
Bibliography 14

Learning goals

After studying this chapter you should be able to:

- explain the **mean-field** idea and why it makes the many-electron problem solvable;
- write the energy of a Slater determinant and say what the one-electron, Coulomb and exchange terms mean;
- write the **Fock operator** and explain why exchange has no classical analogue and why Hartree–Fock is free of self-interaction;
- write the **Roothaan–Hall equations** $\mathbf{FC} = \mathbf{SC}\boldsymbol{\varepsilon}$ and describe every step of the **SCF loop**, including what the DIIS error in a Q-Chem output means;
- tell restricted (RHF) from unrestricted (UHF) calculations and know when each is used;
- use **Koopmans' theorem** to estimate an ionization energy, and know its limits;
- define the **correlation energy**, tell dynamic from static correlation, and explain why it matters for reaction energies even though it is only about 1% of the total energy.

Review from P-Chem 2

This chapter builds on the instructor's P-Chem 2 textbook. Review these sections first if they feel rusty:

- Many-electron atoms, §10.1–10.2: many-electron wave functions, the variational method, the helium atom.
- Postulate 6: the Pauli exclusion principle and Spin.
- Solving the electronic eigenvalue problem, §11.3 and the Hartree–Fock method, §11.3.1: LCAO, the secular determinant, the idea of a self-consistent field.
- Hartree–Fock calculation for water, §13.2: a complete minimal-basis calculation, with the energy printed at every SCF cycle.

In P-Chem 2 you saw the *recipe*. Here we open the box and look at the equations behind it.

1 Where we are

Every calculation in this course makes two approximations. In Lecture 1 you wrote the electronic Hamiltonian (atomic units, nuclei clamped):

$$\hat{H}_{\mathrm{el}} = -\frac{1}{2}\sum_{i=1}^{N}\nabla_i^2 - \sum_{i=1}^{N}\sum_{A}\frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A|} + \sum_{i=1}^{N}\sum_{j>i}^{N}\frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}. \tag{1}$$

The three terms are the kinetic energy of the electrons, their attraction to the nuclei, and the electron–electron repulsion. The last term couples every electron to every other, and that is why no analytic solution exists for $N \geq 2$.

In Lecture 2 you approximated the **wave function's flexibility** with a basis set: orbitals are written as a finite sum of basis functions χ_μ. That is the first approximation. This lecture and the next two are about the second one: how we approximate the **Hamiltonian**, or more precisely, how we treat the electron–electron repulsion. Together, the two choices make a *model chemistry*, written method/basis, such as HF/def2-TZVP or B3LYP/def2-SVP.

Hartree–Fock (HF) is the starting point for everything else. It is a method in its own right (you are running it in Experiment 1), it is the reference for the correlated wave-function methods that produce the benchmark numbers of this lab, and its machinery (orbitals, the Fock matrix, the SCF loop) is reused almost unchanged by density functional theory in Lecture 5.

ℹ Experiment 1 check-in

By now you should have submitted your first jobs for Experiment 1. If anything did not work (the server, IQmol, a job that never finished), check the guides on setting up IQmol, your first job and reading the output, and ask in class. We will discuss the results of Experiment 1 in Lecture 7, after everyone has submitted.

2 The mean-field idea

2.1 Hartree's average field

In 1928 Douglas Hartree proposed a simple way out of the many-electron problem [1]. Instead of following every electron as it dodges every other electron, let each electron move in the **average** electrostatic field created by the nuclei and by the *charge cloud* of all the other electrons. Each electron then obeys its own one-electron Schrödinger equation, with an effective potential that depends on where the other electrons are on average.

This is a **mean-field** approximation. It replaces the instantaneous repulsion $1/|\mathbf{r}_i - \mathbf{r}_j|$ by an interaction with a smeared-out charge density. The catch is that the average field depends on the orbitals of the other electrons, which are themselves the solutions we are looking for. Hartree solved the problem iteratively: guess the orbitals, compute the field, solve for new orbitals, compute a new field, and repeat until the orbitals that come out produce the same field that went in. He called the result a **self-consistent field** (SCF). The name survives in every quantum chemistry program, including the "SCF" lines of your Q-Chem outputs.

Hartree wrote the many-electron wave function as a simple product of orbitals, $\Psi = \varphi_1(\mathbf{r}_1)\varphi_2(\mathbf{r}_2)\cdots\varphi_N(\mathbf{r}_N)$. A product like this is not antisymmetric. It violates the Pauli principle, and it misses an important energy term, as we will see.

2.2 The Slater determinant

In 1930 Fock [2] and Slater [3] fixed this by using a determinant of **spin orbitals** instead of a product. A spin orbital $\phi_i(\mathbf{x})$ is a spatial orbital times a spin function, $\phi_i(\mathbf{x}) = \varphi_i(\mathbf{r})\,\sigma(\omega)$ with $\sigma = \alpha$ or β, and $\mathbf{x} = (\mathbf{r}, \omega)$ collects the space and spin coordinates of one electron. The **Slater determinant** is

$$\Psi(\mathbf{x}_1, \ldots, \mathbf{x}_N) = \frac{1}{\sqrt{N!}} \begin{vmatrix} \phi_1(\mathbf{x}_1) & \phi_2(\mathbf{x}_1) & \cdots & \phi_N(\mathbf{x}_1) \\ \phi_1(\mathbf{x}_2) & \phi_2(\mathbf{x}_2) & \cdots & \phi_N(\mathbf{x}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_1(\mathbf{x}_N) & \phi_2(\mathbf{x}_N) & \cdots & \phi_N(\mathbf{x}_N) \end{vmatrix}. \tag{2}$$

Swapping two electrons swaps two rows, which changes the sign of the determinant: the wave function is antisymmetric, as the Pauli principle requires. Putting two electrons in the same spin orbital makes two columns equal, and the determinant vanishes: no two electrons can share all their quantum numbers. You met this determinant in Lecture 1 and in P-Chem 2. We will write it compactly as $\Psi = |\phi_1\phi_2\cdots\phi_N|$.

The Hartree–Fock approximation is a single statement: *the wave function is one Slater determinant, and the best determinant is the one with the lowest energy.* "Best" comes from the variational principle, $E[\Psi_{\text{trial}}] \geq E_0$, which you used for basis sets in Lecture 2. Everything else in this chapter follows from that statement.

2.3 The energy of a Slater determinant

Take the expectation value of the Hamiltonian (Equation 1) with a normalized determinant of orthonormal spin orbitals. The algebra is long but standard (see A. Szabo and N. S. Ostlund [4], chapter 2). The result is short:

$$E_{\text{HF}} = \sum_{i=1}^{N} h_i + \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N}(J_{ij} - K_{ij}) + V_{nn}. \tag{3}$$

Here V_{nn} is the constant nuclear repulsion, and the three kinds of integrals are:

- the **one-electron (core) energy** of spin orbital i, its kinetic energy plus its attraction to all nuclei:

$$h_i = \int \phi_i^*(\mathbf{x}) \left[-\frac{1}{2}\nabla^2 - \sum_A \frac{Z_A}{|\mathbf{r} - \mathbf{R}_A|} \right] \phi_i(\mathbf{x})\, d\mathbf{x};$$

- the **Coulomb integral**, the classical repulsion between the charge clouds $|\phi_i|^2$ and $|\phi_j|^2$:

$$J_{ij} = \iint \frac{|\phi_i(\mathbf{x}_1)|^2\, |\phi_j(\mathbf{x}_2)|^2}{r_{12}}\, d\mathbf{x}_1\, d\mathbf{x}_2;$$

- the **exchange integral**, which has the two orbitals swapped between the electrons on one side:

$$K_{ij} = \iint \frac{\phi_i^*(\mathbf{x}_1)\phi_j(\mathbf{x}_1)\, \phi_j^*(\mathbf{x}_2)\phi_i(\mathbf{x}_2)}{r_{12}}\, d\mathbf{x}_1\, d\mathbf{x}_2.$$

The first two terms are what you would write for classical charge clouds. The third is new. It comes entirely from the antisymmetry of the determinant: a simple Hartree product gives only h and J. Because the spin parts must match on both sides of K_{ij}, the exchange integral is zero unless electrons i and j have the **same spin**. Exchange is always positive ($K_{ij} \geq 0$), so it always lowers the energy.

> **i** Exchange has no classical analogue
>
> Exchange is not a new force. It is a correction to the Coulomb repulsion that appears because two electrons with the same spin cannot be at the same place (the determinant vanishes when $\mathbf{x}_1 = \mathbf{x}_2$). Around each electron there is a region, the **Fermi hole** or exchange hole, where same-spin electrons are unlikely to be found. Less overlap means less repulsion, and $-K_{ij}$ is that energy saving. Nothing in classical electrostatics gives this term. It is a pure consequence of the Pauli principle.

> **!** Remember this: no self-interaction in Hartree–Fock
>
> Look at the double sum in Equation 3. It includes the terms with $i = j$. The term J_{ii} is the repulsion of electron i with *its own* charge cloud, which is unphysical: an electron does not repel itself. But from the definitions, $K_{ii} = J_{ii}$, so the two cancel exactly: $J_{ii} - K_{ii} = 0$. Hartree–Fock is **free of self-interaction**. Hartree's original product method was not, and, as you will see in Lecture 5, approximate density functionals are not either. This single fact explains why mixing some Hartree–Fock exchange into a functional helps (Lecture 6).

2.4 The Hartree–Fock equations and the Fock operator

We want the spin orbitals that minimize Equation 3 while staying orthonormal. This is a constrained minimization, done with Lagrange multipliers. After a unitary rotation that makes the multiplier matrix diagonal (the *canonical* orbitals), the condition for a minimum becomes a set of one-electron eigenvalue equations, the **Hartree–Fock equations**:

$$\hat{f}\,\phi_i = \varepsilon_i\,\phi_i, \qquad i = 1, \ldots, N. \tag{4}$$

The **Fock operator** $\hat{f}$ is an effective one-electron Hamiltonian:

$$\hat{f}(\mathbf{x}_1) = \hat{h}(\mathbf{x}_1) + \sum_{j=1}^{N} \left[\hat{J}_j(\mathbf{x}_1) - \hat{K}_j(\mathbf{x}_1)\right]. \tag{5}$$

- $\hat{h} = -\tfrac{1}{2}\nabla^2 - \sum_A Z_A/|\mathbf{r} - \mathbf{R}_A|$ is the **core Hamiltonian**: kinetic energy plus nuclear attraction of one electron.
- $\hat{J}_j$ is the **Coulomb operator**: it multiplies the orbital by the electrostatic potential of the charge cloud of electron j,

$$\hat{J}_j(\mathbf{x}_1)\,\phi_i(\mathbf{x}_1) = \left[\int \frac{|\phi_j(\mathbf{x}_2)|^2}{r_{12}}\,d\mathbf{x}_2\right]\phi_i(\mathbf{x}_1).$$

- $\hat{K}_j$ is the **exchange operator**: it swaps the orbitals inside the integral,

$$\hat{K}_j(\mathbf{x}_1)\,\phi_i(\mathbf{x}_1) = \left[\int \frac{\phi_j^*(\mathbf{x}_2)\,\phi_i(\mathbf{x}_2)}{r_{12}}\,d\mathbf{x}_2\right]\phi_j(\mathbf{x}_1).$$

$\hat{K}_j$ is a **nonlocal** operator: the result at point $\mathbf{x}_1$ depends on the value of ϕ_i everywhere in space, not only at $\mathbf{x}_1$. Keep this in mind. Density functional approximations replace this nonlocal operator by a *local* potential, and that is the root of several of their errors.

The sum $\sum_j (\hat{J}_j - \hat{K}_j)$ is the **Hartree–Fock potential**: the average field that each electron feels from all the others. This is the mean-field idea in equation form. The term $j = i$ is included in the sum, but, as above, its Coulomb and exchange parts cancel when acting on ϕ_i.

Each eigenvalue ε_i is an **orbital energy**. It equals $h_i + \sum_j (J_{ij} - K_{ij})$: the kinetic and nuclear-attraction energy of electron i plus its average repulsion with all other electrons. Note that the total energy is **not** the sum of the orbital energies. Summing ε_i over all occupied orbitals counts every electron pair twice, so

$$E_{\mathrm{HF}} = \sum_i \varepsilon_i - \frac{1}{2} \sum_{i,j} \left(J_{ij} - K_{ij} \right) + V_{nn}.$$

This is exactly the bookkeeping check in §13.2 of P-Chem 2: twice the sum of the occupied orbital energies of water, plus the electron–electron repulsion counted once with the opposite sign, plus the nuclear repulsion, gives back the total energy.

2.5 Closed shells: restricted Hartree–Fock

Most of the molecules in this course have an even number of electrons, all paired. For these **closed-shell** molecules we let each spatial orbital φ_a hold two electrons, one α and one β. This is **restricted Hartree–Fock** (RHF). Doing the spin integration once and for all, the Fock operator for the $N/2$ doubly occupied spatial orbitals becomes

$$\hat{f}(\mathbf{r}) = \hat{h}(\mathbf{r}) + \sum_{a=1}^{N/2} \left[2\hat{J}_a(\mathbf{r}) - \hat{K}_a(\mathbf{r}) \right], \tag{6}$$

where the factor 2 counts the two electrons in each orbital, and the exchange term has no factor 2 because exchange acts only between electrons of the same spin. We come back to open shells in Section 3.4.

3 The Roothaan–Hall equations and the SCF loop

3.1 From integro-differential equations to matrices

Equations (Equation 4) are coupled integro-differential equations. For atoms they can be solved numerically on a radial grid, as Hartree did. For molecules this is impractical. In 1951 Roothaan [5] and, independently, Hall [6] used the LCAO idea from Lecture 2: expand each spatial orbital in a finite basis of K functions,

$$\varphi_i(\mathbf{r}) = \sum_{\mu=1}^{K} c_{\mu i}\, \chi_\mu(\mathbf{r}). \tag{7}$$

Insert this into $\hat{f}\varphi_i = \varepsilon_i \varphi_i$, multiply from the left by χ_ν^* and integrate. The unknown functions become unknown coefficients, and the equations become a matrix problem, the **Roothaan–Hall equations**:

$$\mathbf{F}\mathbf{C} = \mathbf{S}\mathbf{C}\boldsymbol{\varepsilon}. \tag{8}$$

All matrices are $K \times K$:

- $\mathbf{C}$ holds the MO coefficients: column i is orbital φ_i.
- $\boldsymbol{\varepsilon}$ is diagonal, with the orbital energies ε_i.

- $\mathbf{S}$ is the **overlap matrix**, $S_{\mu\nu} = \int \chi_\mu^* \chi_\nu \, d\mathbf{r}$. It is not the identity matrix because basis functions on different atoms overlap (the same point is made in P-Chem 2 §11.3.1). This is why Equation 8 is a *generalized* eigenvalue problem.
- $\mathbf{F}$ is the **Fock matrix**, $F_{\mu\nu} = \int \chi_\mu^* \hat{f} \chi_\nu \, d\mathbf{r}$. For RHF it is

$$F_{\mu\nu} = H_{\mu\nu}^{\text{core}} + \sum_{\lambda=1}^{K} \sum_{\sigma=1}^{K} P_{\lambda\sigma} \left[(\mu\nu \mid \lambda\sigma) - \frac{1}{2} (\mu\lambda \mid \nu\sigma) \right], \tag{9}$$

with the core-Hamiltonian matrix $H_{\mu\nu}^{\text{core}} = \int \chi_\mu^* \hat{h} \chi_\nu \, d\mathbf{r}$, the **two-electron integrals**

$$(\mu\nu|\lambda\sigma) = \iint \frac{\chi_\mu^*(\mathbf{r}_1)\chi_\nu(\mathbf{r}_1)\, \chi_\lambda^*(\mathbf{r}_2)\chi_\sigma(\mathbf{r}_2)}{r_{12}} \, d\mathbf{r}_1 \, d\mathbf{r}_2,$$

and the **density matrix**

$$P_{\lambda\sigma} = 2 \sum_{a=1}^{N/2} c_{\lambda a} \, c_{\sigma a}^*. \tag{10}$$

The density matrix is the bridge between orbitals and electron density: $\rho(\mathbf{r}) = \sum_{\lambda\sigma} P_{\lambda\sigma} \chi_\lambda(\mathbf{r}) \chi_\sigma^*(\mathbf{r})$. In Equation 9 the first bracket term is the Coulomb part and the second, with the swapped indices, is the exchange part. The notation in P-Chem 2 is slightly different ($h = K + V$, G for the two-electron part, V_{ijkl} for the integrals), but the equation is the same.

Once the equations are solved, the energy is

$$E_{\text{RHF}} = \frac{1}{2} \sum_{\mu\nu} P_{\nu\mu} \left(H_{\mu\nu}^{\text{core}} + F_{\mu\nu} \right) + V_{nn}. \tag{11}$$

> **i** Where the cost comes from
>
> There are about $K^4/8$ unique two-electron integrals $(\mu\nu|\lambda\sigma)$. With 7 basis functions (water in a minimal basis, as in P-Chem 2 §13.2) this is a few hundred. With the 200 or so functions of a medium-sized molecule in def2-TZVP it is about 2×10^8. Computing and contracting these integrals with $\mathbf{P}$ is the formal K^4 (often written N^4) cost of Hartree–Fock that you met in Lecture 2 and are measuring in Experiment 1. In practice programs skip integrals that are negligibly small (screening) or approximate them (density fitting), so the observed scaling is lower.

3.2 Why it must be solved iteratively

Equation 8 looks like an ordinary eigenvalue problem, but it is not linear: the Fock matrix depends on the density matrix (Equation 9), which depends on the coefficients (Equation 10), which are the solution of the equation. We need $\mathbf{C}$ to build $\mathbf{F}$, and $\mathbf{F}$ to find $\mathbf{C}$. The way out is Hartree's: iterate until the input and output agree. This is the **SCF procedure** (Figure 1).

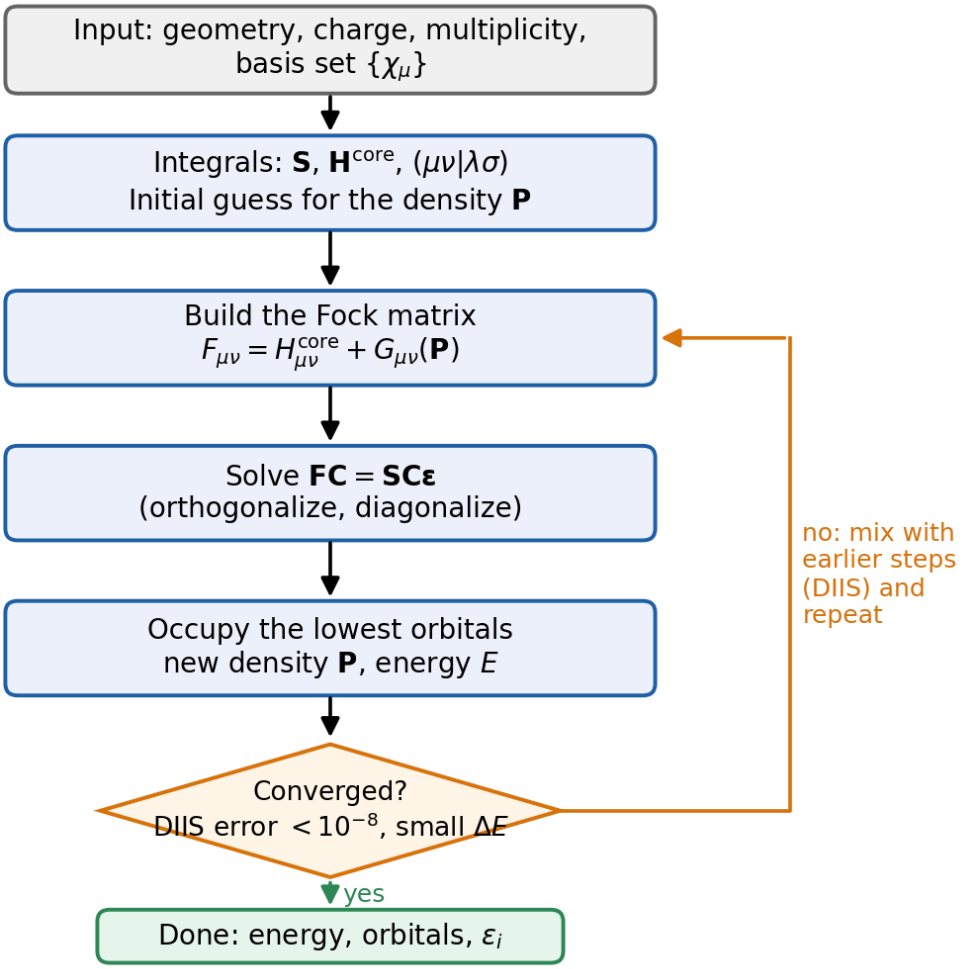


Figure 1: The self-consistent field (SCF) loop used by Hartree–Fock and by Kohn–Sham DFT. Each pass through the loop is one line ("cycle") of the SCF table in a Q-Chem output.

Step by step:

1. **Input.** Geometry, charge, multiplicity and basis set: exactly the `$molecule` and `$rem` sections of a Q-Chem input.
2. **One-time work.** Compute $\mathbf{S}$ and $\mathbf{H}^{\text{core}}$ (cheap), and decide how to handle the two-electron integrals (store them, or recompute them each cycle, which modern programs usually do).
3. **Initial guess.** Make a first density matrix $\mathbf{P}$. A crude choice is to diagonalize $\mathbf{H}^{\text{core}}$ alone (no electron repulsion at all). A much better one, and the default in Q-Chem, is a superposition of precomputed atomic densities (SAD). A good guess saves cycles and makes it more likely that the SCF lands on the right solution.
4. **Build the Fock matrix** from the current $\mathbf{P}$ with Equation 9. This is the expensive step.
5. **Solve $\mathbf{FC} = \mathbf{SC}\varepsilon$.** First transform to an orthonormal basis (for example with $\mathbf{S}^{-1/2}$), so the problem becomes an ordinary eigenvalue problem $\mathbf{F}'\mathbf{C}' = \mathbf{C}'\varepsilon$. Diagonalize, then transform the coefficients back.
6. **Occupy** the $N/2$ orbitals with the lowest ε_i (the *aufbau* principle), build the new $\mathbf{P}$ with Equation 10, and compute the energy with Equation 11.
7. **Check convergence.** If the new density is (almost) the same as the old one, stop. If not, go back to step 4 with the new density.

3.3 Convergence, DIIS and the SCF table in your output

What does "the same" mean in step 7? At self-consistency the Fock matrix and the density matrix commute in the metric of the overlap matrix:

$$\mathbf{FPS} - \mathbf{SPF} = \mathbf{0}. \tag{12}$$

Away from convergence this commutator is not zero, and its size is a natural error measure. Pulay's **DIIS** method (direct inversion in the iterative subspace) [7] uses it twice. First, the largest element of the commutator is the **DIIS error** that Q-Chem prints in the third column of the SCF table. Second, instead of using only the newest Fock matrix, DIIS builds the next one as a linear combination of the Fock matrices from the last several cycles, with coefficients chosen to make the combined error as small as possible. This extrapolation is what makes modern SCF calculations converge in 10–20 cycles instead of hundreds.

Here is what this looks like in a Q-Chem output (see the guide on reading the output):

```
---------------------------------------
 Cycle       Energy         DIIS error
---------------------------------------
   1    -920.4153225231      4.31e-02
   2    -920.6017766913      7.45e-03
   ...
  11    -920.6431907518      1.12e-07
  12    -920.6431907520      4.05e-09  Convergence criterion met
---------------------------------------
```

Each line is one pass through Figure 1. The energy goes down and settles. The DIIS error shrinks by orders of magnitude. The job stops when the DIIS error falls below 10^{-n}, where n is the value of `SCF_CONVERGENCE` in the `$rem` section. All experiments in this course use `SCF_CONVERGENCE 8`, a tight threshold: the energy is then converged far beyond the digits you need for a reaction energy, because the energy error is roughly the *square* of the density error. That is a consequence of the variational principle: at the minimum, first-order errors in the orbitals give only second-order errors in the energy.

⚠ When the SCF goes wrong

The SCF is a numerical procedure, and it can fail in three ways. Learn to recognize them now. You will meet all three later in the course.

- **Oscillation.** The energy jumps up and down from cycle to cycle and never settles. This happens when two or more orbitals are close in energy near the HOMO–LUMO gap and swap occupation back and forth. Typical for transition-metal compounds and stretched bonds.
- **Slow convergence.** The DIIS error decreases, but very slowly, and the job hits the maximum number of cycles (Q-Chem prints a warning that the SCF failed to converge). The energy of a job that did not converge is meaningless. Never use it.
- **Convergence to the wrong solution.** The SCF converges perfectly, but to a stationary point that is not the lowest-energy solution: an excited configuration, or a saddle point that is a minimum only within the constraints you imposed (for example, RHF for a molecule that wants unpaired electrons). The output looks normal. Only a **stability analysis** can tell. This is the subject of Lecture 12 and Experiment 8.

Programs have many tools against the first two: better initial guesses, damping, level shifting, and second-order algorithms such as geometric direct minimization. You will not need them for the organic molecules of Experiments 1–7.

3.4 Restricted and unrestricted Hartree–Fock

Molecules with unpaired electrons (radicals, the O and S atoms, triplet states, many transition-metal complexes) cannot use Equation 6, because their α and β electrons are not paired up in the same spatial orbitals. There are two common ways to handle them:

- **Unrestricted Hartree–Fock (UHF)** [8] gives the α and β electrons their own sets of spatial orbitals, $\{\varphi_i^\alpha\}$ and $\left\{\varphi_i^\beta\right\}$. There are two Fock matrices and two density matrices, coupled through the Coulomb term (the α electrons repel the β electrons and vice versa). In a basis these are the Pople–Nesbet equations, $\mathbf{F}^\alpha\mathbf{C}^\alpha = \mathbf{S}\mathbf{C}^\alpha\boldsymbol{\varepsilon}^\alpha$ and the same for β. UHF has more freedom than RHF, so its energy is lower or equal. The price is that the UHF determinant is not, in general, an eigenfunction of $\hat{S}^2$: it is contaminated by higher spin states. Q-Chem prints $\langle\hat{S}^2\rangle$ for unrestricted jobs so you can check (a pure doublet has 0.75, a pure triplet 2.0).
- **Restricted open-shell Hartree–Fock (ROHF)** keeps doubly occupied orbitals for the paired electrons and singly occupied orbitals for the unpaired ones. It gives a pure spin state, at a slightly higher energy than UHF.

In this course, open-shell species are run **unrestricted** (Q-Chem does this automatically when the multiplicity is larger than 1). The same choice exists in DFT, where we speak of restricted and unrestricted Kohn–Sham. You will use unrestricted calculations for atoms and radicals in Experiment 5 and later.

Multiplicity reminder

The multiplicity is $2S + 1$, where S is the total spin: 1 for a closed shell (singlet), 2 for one unpaired electron (doublet), 3 for two parallel unpaired electrons (triplet). It is the second number on line 2 of every `.xyz` file of this lab ("charge multiplicity"). A triplet O atom run as a singlet gives a perfectly converged, perfectly wrong energy.

4 Orbital energies and Koopmans' theorem

Diagonalizing the Fock matrix gives K orbitals, but only the lowest $N/2$ (for RHF) are **occupied**. The remaining $K - N/2$ are **virtual** (unoccupied) orbitals. The highest occupied one is the **HOMO** and the lowest unoccupied

one the **LUMO**. The number of virtual orbitals grows with the basis set, while the number of occupied orbitals is fixed by the number of electrons.

Q-Chem prints the orbital energies after the SCF table, under a heading "Orbital Energies (a.u.)", split into "Occupied" and "Virtual" lists (for unrestricted jobs, separately for α and β). Occupied orbital energies are negative: these electrons are bound. Their pattern follows chemical intuition: very low core levels (about -20.5 Eh for the O 1s orbital in water), then valence levels between about -1.5 and -0.3 Eh. Compare the minimal-basis water numbers in P-Chem 2 §13.2.

Koopmans' theorem [9] gives the occupied orbital energies a physical meaning. If we remove an electron from orbital i and keep all the other orbitals frozen, the energy of the resulting cation, minus the energy of the neutral molecule, is exactly $-\varepsilon_i$. So the first ionization energy is approximately

$$\mathrm{IE} \approx -\varepsilon_{\mathrm{HOMO}}. \tag{13}$$

Two errors are hidden in this approximation. **Orbital relaxation**: in the real cation the remaining electrons rearrange to the new field, which lowers the cation energy, so Koopmans' value is too large. **Correlation**: the neutral molecule has one more electron pair to correlate than the cation, so including correlation increases the true ionization energy, and Koopmans' value is too small. The two errors have opposite signs and partly cancel, so Koopmans IEs of closed-shell molecules are often within about 1 eV of experiment. The cancellation is fortunate, not guaranteed. The analogous estimate of the electron affinity from the LUMO, $\mathrm{EA} \approx -\varepsilon_{\mathrm{LUMO}}$, is much worse, because the two errors now add up.

> **i** Worked example: water
>
> An HF calculation on water with a large basis gives $\varepsilon_{\mathrm{HOMO}} \approx -0.507$ Eh. Then $\mathrm{IE} \approx 0.507 \times 27.211 \approx 13.8$ eV. The experimental vertical ionization energy of water is 12.6 eV. Koopmans' estimate is about 1.2 eV too high, as expected when relaxation wins over correlation. Try it yourself: run water with HF and a triple-ζ basis and look for the last "Occupied" number in the output. Your value will depend slightly on the geometry and basis.

> What the orbital pictures in IQmol show
>
> IQmol can plot any orbital from a Q-Chem job as an isosurface (a surface of constant $|\varphi_i|$, with two colors for the two signs). These pictures are useful for chemical intuition: the HOMO of ethene is the C=C π orbital, the LUMO is the π*. Keep two things in mind. First, the canonical orbitals of Equation 4 are delocalized over the whole molecule and do not look like Lewis bonds and lone pairs (water's two "lone pairs" are two quite different orbitals). Any unitary rotation of the occupied orbitals gives the *same* determinant and the same energy, so localized orbitals are equally valid. Second, the virtual orbitals depend strongly on the basis set, and adding diffuse functions changes them a lot.

5 What Hartree–Fock misses: electron correlation

5.1 The correlation energy

The Hartree–Fock wave function is the best single determinant, but it is still only one determinant. In the mean-field picture, each electron sees the *average* position of the others. Real electrons react to each other's *instantaneous* positions: they avoid each other more than the average field suggests. Their motions are **corre-**

lated. Exchange already keeps same-spin electrons apart (the Fermi hole), but nothing in HF keeps opposite-spin electrons apart.

Löwdin defined the **correlation energy** as the difference between the exact energy and the Hartree–Fock energy [10]:

$$E_{\mathrm{corr}} = E_{\mathrm{exact}} - E_{\mathrm{HF}}. \tag{14}$$

Because of the variational principle, E_{corr} is always negative. Two details matter. "Exact" means the exact nonrelativistic Born–Oppenheimer energy. And both energies should be in the same basis set, or, ideally, both at the complete-basis-set (CBS) limit. Otherwise the basis-set error gets mixed into the correlation energy.

How big is it? For the helium atom, $E_{\mathrm{HF}} = -2.8617$ Eh at the basis-set limit, while the exact energy is -2.9037 Eh. The correlation energy is -0.042 Eh, or -26 kcal/mol: about 1.4% of the total energy. For the water molecule it is roughly -0.37 Eh (about -230 kcal/mol), about 0.5% of the total. As a rule of thumb, the correlation energy is around 1% of the total energy, or very roughly 25 kcal/mol per electron pair.

One percent sounds small. It is not. **Chemistry happens in energy differences.** A typical reaction energy is 1–100 kcal/mol, and the conformational energies of Experiment 3 are 1–3 kcal/mol. The total energy of a small organic molecule is about 100,000–500,000 kcal/mol. If the correlation energy changed by just 1% between reactants and products, it would shift a reaction energy by a few kcal/mol. When bonds are broken or formed, or when the number of electron pairs changes, the change in correlation energy is much larger than that.

> **i** Hartree–Fock has no dispersion at all
>
> London dispersion, the attraction between two neutral, nonpolar molecules (two argon atoms, two methane molecules, the stacked layers of graphite), comes from correlated, instantaneous fluctuations of the electron clouds. A mean-field method cannot describe it: in HF each molecule sees only the *static* average charge of the other. At long range, HF gives no dispersion attraction at all. Weakly bound complexes held together mainly by dispersion are unbound, or far too weakly bound, at the HF level, no matter how large the basis set. Keep this in mind when you look at your Experiment 1 results. Dispersion returns in Lecture 8.

5.2 Dynamic and static correlation

It is useful to split correlation into two kinds, although there is no sharp boundary between them.

Dynamic correlation is the short-range dodging of electrons: each electron creates a small "Coulomb hole" around itself where other electrons, of either spin, are less likely to be. It is present in every molecule with more than one electron. Describing it requires mixing in many determinants, each with a small weight. It converges slowly with the basis set, because the wave function has a cusp where two electrons meet.

Static (nondynamic) correlation appears when one determinant is not even qualitatively right: when two or more electron configurations have nearly the same energy and must be mixed with comparable weights. Typical examples are stretched or breaking bonds, diradicals, ozone, and many transition-metal compounds.

The textbook example is the dissociation of H_2 (Figure 2). Near the equilibrium distance, one doubly occupied bonding orbital σ_g is a good description, and RHF is qualitatively right. The minimum is a little too shallow, and the difference from the exact curve is the dynamic correlation energy. But as the bond is stretched, the RHF determinant $|\sigma_g\alpha\,\sigma_g\beta|$ still forces both electrons into the same spatial orbital. Expanded in atomic orbitals, it is 50% covalent (H· + ·H) and 50% ionic ($H^+ + H^-$), at any distance. At infinite separation it should be 100% covalent. The ionic half costs a lot of energy, so the RHF curve rises far above the energy of two H atoms. The exact wave

function fixes this by mixing in a second determinant, $|\sigma_u\alpha\,\sigma_u\beta|$, with a weight that grows until both have equal weight at dissociation. That second determinant is the static correlation.

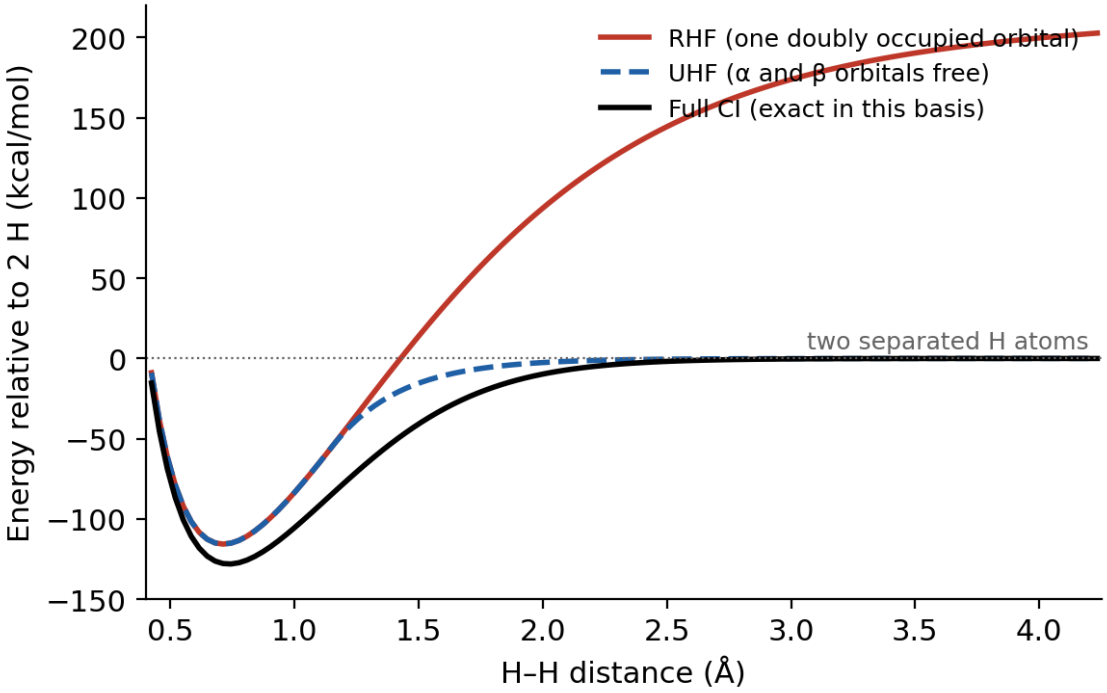


Figure 2: Dissociation of H_2 computed with a minimal basis set (STO-3G), where the full configuration interaction (full CI) curve is the exact answer within this basis. RHF is good near equilibrium but goes to an energy far above two H atoms. UHF breaks the spin symmetry beyond about 1.2 Å and dissociates correctly, at the cost of a spin-contaminated wave function. The gap between RHF and full CI at the minimum is the dynamic correlation energy; the growing gap at long distance is static correlation. Computed with `images/src/l04_l06_figures.py`.

UHF also "fixes" the dissociation: past a certain distance (the Coulson–Fischer point) the α electron localizes on one atom and the β electron on the other, and the energy goes correctly to two H atoms. But the UHF wave function is then a 50/50 mixture of a singlet and a triplet, and the curve has a kink. UHF solutions that are lower than RHF are an example of an **instability** of the RHF solution: the topic of Lecture 12.

> ⚠ Common mistake: "a converged SCF means the method is right"
>
> A converged SCF tells you that the equations were solved, not that the approximation behind them is good. RHF for H_2 at 3 Å converges without trouble to a number that is 200 kcal/mol too high. Convergence is about numerics. Accuracy is about the method.

6 Ways forward

There are two broad families of methods that go beyond Hartree–Fock.

Correlated wave-function methods keep the HF determinant as a starting point (the *reference*) and add excited determinants to it.

- **Møller–Plesset perturbation theory** treats the difference between the real Hamiltonian and the sum of Fock operators as a perturbation [11]. Its second-order version, **MP2**, recovers a large part of the dynamic correlation and describes dispersion, at a cost that scales as N^5.

- **Coupled-cluster theory** includes excitations to all orders in an exponential form. **CCSD(T)** (singles, doubles and a perturbative estimate of triples) [12] is called the "gold standard" of quantum chemistry: for molecules without strong static correlation, CCSD(T) near the CBS limit is typically within about 1 kcal/mol of experiment for reaction energies. Its cost scales as N^7: doubling the size of the molecule makes the calculation about 128 times slower. The reference values of most reactions in this lab come from CCSD(T)/CBS or similar high-level calculations collected in benchmark databases [13], as you will see in Lecture 7.
- **Multireference methods** (CASSCF, MRCI, CASPT2) start from several determinants and are needed when static correlation is strong. They are outside the scope of this course.

Density functional theory takes a different path. It keeps a one-determinant, mean-field machinery, with the same SCF loop and the same N^3–N^4 cost as HF, but it changes the *Hamiltonian* of that machinery so that it includes, approximately, both exchange and correlation. How that is possible, and why the hard part is a single term called the exchange–correlation energy, is the subject of Lecture 5.

Table 1: Hartree–Fock and the methods that go beyond it. "Formal cost" is how the time grows with system size N for a fixed basis-set quality.

Method	Correlation	Formal cost	Dispersion	Typical use
HF	none	N^4	none	reference, starting point
DFT (GGA, hybrid)	approximate, through E_{xc}	N^3–N^4	only if added (-D4, VV10) or fitted	everyday chemistry
MP2	dynamic, second order	N^5	yes (often overestimated)	medium molecules
CCSD(T)	dynamic, nearly complete	N^7	yes	benchmarks, small molecules

Key points

- Hartree–Fock approximates the wave function by **one Slater determinant** and finds the best one with the variational principle. Each electron moves in the **average field** of the others.
- The HF energy contains one-electron terms, classical **Coulomb** terms J_{ij}, and **exchange** terms K_{ij} between same-spin electrons. Exchange comes from antisymmetry and has no classical analogue.
- Because $J_{ii} = K_{ii}$, HF is **free of self-interaction**. Approximate density functionals are not.
- The **Fock operator** is $\hat{f} = \hat{h} + \sum_j (\hat{J}_j - \hat{K}_j)$. In a basis, the HF equations become the **Roothaan–Hall equations** $\mathbf{FC} = \mathbf{SC}\boldsymbol{\varepsilon}$.
- $\mathbf{F}$ depends on its own solution through the density matrix, so the equations are solved **iteratively** (the SCF loop). DIIS accelerates the loop, and the **DIIS error** in the output measures how far from self-consistency you are. `SCF_CONVERGENCE 8` means a DIIS error below 10^{-8}.
- The SCF can oscillate, converge slowly, or converge to the **wrong solution**. A converged SCF is not proof of a correct answer.
- RHF pairs electrons in the same spatial orbitals; **UHF** gives α and β electrons their own orbitals and is used for open shells.
- **Koopmans' theorem**: IE $\approx -\varepsilon_{\mathrm{HOMO}}$, thanks to a partial cancellation of relaxation and correlation errors.
- The **correlation energy** $E_{\mathrm{exact}} - E_{\mathrm{HF}}$ is about 1% of the total energy but as large as or larger than most reaction energies. It has **dynamic** and **static** parts. HF has **no dispersion**.
- Correlated wave-function methods (MP2, CCSD(T)) are accurate but scale as N^5 to N^7. DFT is the cheaper alternative.

Study questions

1. Write the Hartree product and the Slater determinant for a two-electron system. Show that only the determinant changes sign when the two electrons are exchanged, and that it vanishes if both electrons occupy the same spin orbital.
2. In Equation 3, the double sum runs over all i and j, including $i = j$. Why is this harmless in Hartree–Fock? What would go wrong if the exchange term were approximated by something that does not exactly cancel J_{ii}?
3. Explain in words why $K_{ij} = 0$ when electrons i and j have opposite spins. What does this imply about how HF treats opposite-spin electrons?
4. The Roothaan–Hall equations look like an eigenvalue problem, but they must be solved iteratively. Explain why, in two or three sentences, using the words "Fock matrix" and "density matrix".
5. In the SCF table of a Q-Chem output, what is printed in the "DIIS error" column, and what does the job need to see there to stop with `SCF_CONVERGENCE 8`? Why can the energy be converged to more digits than the density?
6. A colleague's job printed a final energy, but no "Convergence criterion met" line. Can the energy be used in a reaction energy? Why or why not?
7. Use Koopmans' theorem to estimate the ionization energy (in eV) of a molecule with $\varepsilon_{\mathrm{HOMO}} = -0.40$ Eh. Name the two approximations hidden in this estimate and the direction of each error.
8. Using Figure 2, explain why RHF fails to dissociate H_2 into two H atoms. What does the exact wave function contain that the RHF wave function lacks? Is this dynamic or static correlation?
9. The correlation energy of a molecule is "only" about 1% of its total energy. Explain why it is still essential for computing reaction energies, using an order-of-magnitude argument.
10. Argon dimer is bound by about 0.3 kcal/mol. What would an HF calculation with a very large basis set predict for it, and why?

Further reading

- P-Chem 2 textbook: Introduction to Molecules, §11.3 and Hartree–Fock calculation for water, §13.2.
- A. Szabo and N. S. Ostlund [4], chapters 2–3: the classic, step-by-step derivation of everything in this chapter, including a complete minimal-basis H_2 calculation like the one in Figure 2.
- T. Helgaker, P. Jørgensen, and J. Olsen [14]: the advanced reference on wave-function methods, for those who want the full theory.
- P. Pulay [7]: the original DIIS paper, short and readable.
- W. Koch and M. C. Holthausen [15], chapter 1: a compact review of Hartree–Fock and correlation written as a lead-in to DFT.

Bibliography

[1] D. R. Hartree, *Math. Proc. Camb. Philos. Soc.*, vol. 24, pp. 89–110, 1928.

[2] V. Fock, *Z. Phys.*, vol. 61, pp. 126–148, 1930.

[3] J. C. Slater, *Phys. Rev.*, vol. 35, pp. 210–211, 1930.

[4] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. Mineola, NY: Dover, 1996.

[5] C. C. J. Roothaan, "New developments in molecular orbital theory," *Rev. Mod. Phys.*, vol. 23, pp. 69–89, 1951, doi: 10.1103/RevModPhys.23.69.

[6] G. G. Hall, "The molecular orbital theory of chemical valency. VIII. A method of calculating ionization potentials," *Proc. R. Soc. London, Ser. A*, vol. 205, pp. 541–552, 1951, doi: 10.1098/rspa.1951.0048.

[7] P. Pulay, "Convergence acceleration of iterative sequences. The case of SCF iteration," *Chem. Phys. Lett.*, vol. 73, pp. 393–398, 1980, doi: 10.1016/0009-2614(80)80396-4.

[8] J. A. Pople and R. K. Nesbet, "Self-consistent orbitals for radicals," *J. Chem. Phys.*, vol. 22, pp. 571–572, 1954, doi: 10.1063/1.1740120.

[9] T. Koopmans, "Über die Zuordnung von Wellenfunktionen und Eigenwerten zu den einzelnen Elektronen eines Atoms," *Physica*, vol. 1, pp. 104–113, 1934, doi: 10.1016/S0031-8914(34)90011-2.

[10] P.-O. Löwdin, "Quantum theory of many-particle systems. III. Extension of the Hartree–Fock scheme to include degenerate systems and correlation effects," *Phys. Rev.*, vol. 97, pp. 1509–1520, 1955, doi: 10.1103/PhysRev.97.1509.

[11] C. Møller and M. S. Plesset, "Note on an approximation treatment for many-electron systems," *Phys. Rev.*, vol. 46, pp. 618–622, 1934, doi: 10.1103/PhysRev.46.618.

[12] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon, "A fifth-order perturbation comparison of electron correlation theories," *Chem. Phys. Lett.*, vol. 157, pp. 479–483, 1989, doi: 10.1016/S0009-2614(89)87395-6.

[13] L. Goerigk, A. Hansen, C. Bauer, S. Ehrlich, A. Najibi, and S. Grimme, *Phys. Chem. Chem. Phys.*, vol. 19, pp. 32184–32215, 2017.

[14] T. Helgaker, P. Jørgensen, and J. Olsen, *Molecular Electronic-Structure Theory*. Chichester: Wiley, 2000.

[15] W. Koch and M. C. Holthausen, *A Chemist's Guide to Density Functional Theory*, 2nd ed. Weinheim: Wiley-VCH, 2001.