

MITOCW | 3. From many-body to single-particle: Quantum modeling of molecules

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

MICHELE: Welcome to lecture three on quantum mechanical methods. As you may have guessed, I'm not Jeff, as much as I look like him, I know. I'm actually a graduate student of his. I'm Michele.

So feel free to interrupt me with any questions you have during this. And this is a really exciting class today because today, we're actually going to start talking about some of the quantum methods that we've been doing in the background up till now. So first of all, I guess you guys know that you don't have class on Thursday, but I guess for those of you doing projects, you can go meet with Professor Buehler then.

So on our outline, we are here. We're going to be talking a little bit about modeling, the beginnings of modeling today, and a little bit about molecules. So last time, we talked about a single electron system.

We talked about hydrogen, and actually, that gave us a lot. We were able to explain how spectral lines came about, just from understanding the basic structure of hydrogen. But today, we're going to work on more-than-one electron systems. And you'll see that this is where the computation really becomes necessary.

So here's the outline. So you can see today, we're going to talk a little bit about Hartree and Hartree-Fock methods, and density functional theory. So I believe you guys are going to be using density functional theory to your next problem set.

So just some review from last time-- here's the Schrodinger equation that I'm sure you're all sick of seeing by now. In general, if our system is time independent, then our Hamiltonian is going to be time independent. And we can separate the variables into the spatial component and the time component. And then by dividing out the time component, we are left with this stationary, time-independent Schrodinger equation. And that's the equation that we've been solving and we will continue to solve.

So as I said, last time we talked about the hydrogen atom. We just solved its Hamiltonian. The Hamiltonian itself is, of course, the kinetic energy and the potential energy.

The kinetic energy is this. The potential energy is actually just that Coulomb attraction between the electron and the proton in that hydrogen atom. And this, we saw, was pretty simple to solve.

I guess we didn't actually do it, but the way this is solved is you switch into spherical coordinates. You solve for a radial component of the wave function, and an angular component of the wave function. So to any of you who have had a lot of chemistry, these probably look very familiar.

So this is the s, these three are p, these five are d. And these are basically just the spherical harmonics that you get from solving the theta and phi parts of the Hamiltonian. And when we combine these, we get the total spatial wave function for the hydrogen atom.

So the spatial wave function is described by these quantum numbers. And the principle quantum number, l , the angular momentum component, and m_l , which is the angular momentum component, projected along the z axis. So you guys have probably seen these a million times-- that leads to these spatial wave functions.

So this is increasing principal quantum number and this is increasing angular momentum. All right, so we've all seen this graph now a couple of times. Can somebody tell me what this graph represents?

AUDIENCE: [INAUDIBLE]

MICHELE: Right. And why is it important that there are these specific lines that it can go between? I mean, what phenomenon are we capturing with this idea?

AUDIENCE: We can [INAUDIBLE]

MICHELE: Right. So this is why people observed spectral lines instead of seeing a continuous spectrum. and so atoms, instead of emitting a continuous spectrum, emit specific lines. And that puzzle was solved by-- we actually already solved it last time with the solution to the hydrogen atom, because we saw that the hydrogen atom had these discrete energy levels.

And electrons could only exist in those specific energy levels. And so when an electron went between those two energy levels, it would emit the energy of that difference in energy levels as a photon. And that photon would be a specific wavelength corresponding to a specific color, and that's exactly where this emission line spectrum comes from.

So let's try to add one more electron. So this is our Hamiltonian now for helium. We still have the kinetic part, the T for electrons 1 and 2. We have the potential part for electrons 1 and 2. That's the interaction of the electrons with the helium nucleus.

But now we have this term that's the Coulomb repulsion. It's the two electrons interacting with each other. And this makes the problem suddenly unable to be solved analytically, just this addition of one more electron. I guess you could think of it as sort of the three-body problem in classical mechanics.

So last time, we also talked about spin. So we know that electrons are either spin up or spin down. Does anybody remember the experiment that was done to show this?

AUDIENCE: Stern-Gerlach.

MICHELE: Yeah, good job. So the Stern-Gerlach experiment, basically it was a stream of electrons that were shot through this magnetic field and then were observed on a plate behind it. So electrons being shot through this magnetic field is kind of like the spin component of the wave function being observed. So now that we're observing this wave function, it can only be in its eigenstates.

So those eigenstates are here and here. So you think if you have an electron that can interact with a magnetic field, its magnetic moment could be pointing sort of any direction. So if it's pointing perpendicular to the magnetic field, why would you expect there to be any kind of interaction? You just expect it to end up in the middle here.

And similarly, it could be anywhere between up or down. But this is quantum mechanics. This magnetic field is making a measurement. And so we all know that once a measurement is made in quantum mechanics, the wave function collapses to one of its observable eigenstates.

So it may seem strange to you that if you know something about where magnetic moments usually come from, usually we talk about magnetic fields being produced by electrons physically spinning in space. But this is just some property of an electron that creates a magnetic moment, without any kind of movement. So that's still something that we don't really understand today, what exactly that spin means.

But we observe it and it seems necessary, so right now, all we can say is that we know electrons either spin up or they're spin down. And Jeff showed you last time this letter to Pauli about how even Pauli was a little skeptical of this spin quantum number, but it's actually crucial for his exclusion principle. So does anybody remember the Pauli exclusion principle? Yeah.

AUDIENCE: [INAUDIBLE]

MICHELE: Exactly. So all the electrons have to have-- all four quantum numbers have to be different between any pair of electrons. So the four quantum numbers, again, are the three spatial components, and now we've added one more, which is spin.

Pauli himself said, "Already in my original paper I stressed the circumstance that I was unable to give a logical reason for the exclusion principle or to deduce it from more general assumptions. I had always the feeling, and I still have it today, that this is a deficiency." So Pauli came up with this. He used it.

It explained the periodic table, but he had no idea why this had to be true. And doesn't it bother you guys that there's just this arbitrary rule that we're saying no two electrons can have the same four quantum numbers? Well, today, we're actually going to explain this. We're going to explain what Pauli could not.

So let's move on to new stuff, unless there are any questions so far. OK, so today, we're going to be talking about what happens when we have more than one electron in a system. So we could be talking about helium with two electrons, iron, 26 electrons.

We'll move up to molecules, solids, things with 100s, 1,000s of electrons. Dirac said in 1929, "The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble."

So basically, he was feeling pretty cocky here. It's like we understand the theory of everything. We understand the periodic table.

There's nothing left to do in chemistry. We understand now. We can write the equations for how all molecules work, how all solids work. And it's the complication of these mathematical formulas that turned out to be the problem.

So many years later, he still was struggling with these. And he says, "If there's no complete agreement between the results of one's work and the experiment, one should not allow himself to be too discouraged." So even after all this time struggling to solve this equation, he still could not agree with experiment. And even today, we still cannot always agree with experiment, but that's no reason not to try.

So this was our basic Hamiltonian that we had for just the hydrogen system. It was just the kinetic and potential energy. It's as simple as that.

Well, now our equation looks a little bit more like this. So I don't want you to be too intimidated by this equation. It seems to have a lot of terms in it, so let's go through them for a simple case.

So let's say we'd have the simplest molecule you can think of, H_2 . So let's say you have a nucleus here, nucleus here, and an electron that originally was associated with this one, and another electron originally associated with this. So what are these terms in the Hamiltonian? So what do we need to consider?

So this first term here is the kinetic energy of the nuclei themselves. So it's how this nucleus and this nucleus are moving around. So in this very simple case, we'd have two terms in this sum.

This next is the interaction [? C-C ?] of Z_i and a Z_j . So it's the interaction of the nuclei with each other. It's the Coulomb repulsion of nucleus 1 and nucleus 2. So in this case, that's actually just one term.

This next term is the kinetic energy of the electrons. So it's, in this case, two terms. We have electron 1 and electron 2 that are zooming around, and that's their kinetic energy.

The next term is there's only one Z , so it's the interaction of the nucleus and the electrons. So here, we have four of those interactions. We have electron 1 interacting with nucleus 1 and 2, electron 2 interacting with nucleus 1 and 2.

And finally, we have electrons interacting with themselves, so the Coulomb repulsion of electrons. So in this case, again, there's just one term because there are only two electrons. So I think all of these terms are labeled properly.

So our equation now has become rather complicated. Instead of just having one coordinate, which was the relative coordinate between the nucleus and the electron that we had in the hydrogen case, now we have this Hamiltonian that depends on the position of all the nuclei and the position of all the electrons, and our wave function that depends on the position of all the nuclei and all the electrons. So this has become a massively complicated equation.

So at this point, what are we going to do? You might think to yourself that you're stuck, that this equation is too complicated. You'll never get anywhere.

And just as you're despairing, this man walks in, pulls down the back of his shirt. Might take you a little while to stare at these equations, because you'll recognize this as being the equation we just had. And he's giving you the Born-Oppenheimer approximation.

So let's look at this approximation, maybe not on somebody's back. So Born was a pretty foundational guy in quantum mechanics. He did a lot of work in quantum mechanics. And I'm sure you all recognize the name Oppenheimer.

So these two guys, what they decided to do, they looked at this picture, and they thought about the fact that a proton is 2,000 times the mass of an electron. And so because this electron is so heavy, it's just moving so slowly with respect to these electrons that are zipping around. And so basically, what they decided to do is just ignore that motion of the nucleons.

So the nuclei are moving so much slower than the electrons that basically, you can assume that the electrons will figure out their ground state, figure out where they need to be, by the time any kind of nucleus has moved even a fraction of an Angstrom. So because we're now neglecting the kinetic energy of the nucleons, we can also just calculate the ion-ion interaction classically. So that really simplifies our Hamiltonian.

So now, we're just left with these three terms. So it's the kinetic energy of the electrons only, the interactions of the electrons with the background positive charge-- so where the ions are-- and the interaction of the electrons with themselves. So we're from now on, we're mostly going to call this term just the "external potential," so "external" because we don't care about the motion of these nuclei anymore.

So this leads us to actually starting to talk about how we do these approximations. So traditionally, there have been two pathways that have been followed. There's what a lot of quantum chemists have taken, and we'll talk about that first, but there's also density functional theory, which is what a lot of physicists have traditionally used.

And I'm not going to talk about either of these two in detail. But I'll just mention, the Moller-Plesset perturbation theory is a perturbation theory. so it's based on changing the Hamiltonian, so assuming that the Hamiltonian is pretty much something we can solve, and then just adding a small correction to it.

And the coupled cluster approach is more of a traditional quantum chemistry approach, where instead of modifying the Hamiltonian, we just modify the wave functions. And we play around with the wave functions until they become something that's more close to something that would exist in reality. So let's talk about the basis of the quantum chemistry approach.

So this is Hartree. He was working right after World War I. I think this was actually his dissertation, and he got his PhD for doing this very important work. So what he decided is we have this system that we want to solve.

We have a wave function that's made up of all of these electrons that are zooming around in space. Well, what if that kind of acted like a bunch of one-electron wave functions? What if we could just assume that we could know where one electron was, and another electron was, and we just could just multiply all those together to give us the total wave function? So all of a sudden, this really complicated wave function becomes something that's pretty easy to separate.

So because we can separate them like this, we can separate them into a set of-- so if we have electrons 1 through n, we'll have a set of Schrodinger equations 1 through n, for each electron by itself. So this is great. We've separated our Hamiltonian. We now can solve the system, right?

So there is a problem here. This is the density of electron j. We're doing a sum over j from electron 1 through n, and just skipping the electron that we're considering in this case. So all of us is saying this electron that we'd like to solve depends on the positions of every single other electron in that system.

So even in this most simplified version of the Schrodinger equation, we still have the one electron depending on all of the other electrons. So how do we solve something like this? Do you guys know? I think this is even a little more complicated than something Mathematica can solve.

AUDIENCE: [INAUDIBLE].

MICHELE: Pretty much. Well-- what we do is we try a self-consistent approach. We basically just guess what these wave functions look like. We say, well in this case, we know that it's a sigma bond that's forming between these two, so probably the electron is somewhere in here.

And so if I guess some structure that looks something like this is my wave function, that's maybe a good starting point. And so we put input that into this equation. We take out the first electron, and we're going to now solve the Schrodinger equation for the first electron.

So we plug in what we've guessed as being the wave functions for all the other electrons into this term. And now we solve for this one electron term, for electron number 1. And then we do that again for electron 2, and so on until we've gone through all of the electrons.

And then when we're done, we have a new set of wave functions. So maybe my guess was wrong and it looks something more like this. And so we look at that, and say, wow, those two really don't look anything like each other. I need to try this again.

And so we go through this process over and over again until finally, the solutions to these wave functions prior to doing this and after going through and solving all of these eigenfunctions look pretty much the same. And once we've done that, we know that we've probably gotten something at least relatively close to the ground-state solution, or the solution that we're looking for. So does that make sense to you guys? Are there any questions?

OK, because this is actually pretty close to how DFT works, too. This is the self-consistent method that is pretty common in solving these problems. So this is great. We have a way of doing this.

Given that we can come up with some reasonable guess for our input wave functions, we now can solve the Schrodinger equation, right? Yes and no. This picture is nice and simple and soluble, but it takes out all of the interactions that you get, or all of the effects that you get from interactions beyond just that coulomb repulsion, of electron-electron Coulomb repulsion.

Because remember, this is a quantum system. These electrons don't behave like you expect they would in some kind of classical Newtonian system. These electrons are weird. We'll see later on some of the symmetries that we can look at. Actually, I think that's next.

But these electrons-- you wouldn't expect to just be able to just average over all of the effects. Because this truly was a multi-electron function. We're missing some critical physics here by just separating it out into single electron functions and just multiplying those back together.

So we're missing two important terms. They're going to be called the "exchange" and the "correlation" term. And the fix to at least part of this problem brings us back to spin.

So as I said, in quantum mechanics, symmetry is really, really important. Symmetry tells us a lot about the way nature works. So even in classical mechanics, you guys have probably encountered the symmetry of real space, the symmetry of time, things like that.

And that's what gives you the conservation laws in mechanics. So just like that's true in mechanics, it's also true in quantum mechanics. So we're going to be looking at one particular symmetry, and talking about how that might give us some insight into this problem.

So that symmetry is called "exchange symmetry." So that's based on the fact that electrons are indistinguishable. So what does it mean to be really indistinguishable?

So if I have, maybe, two pieces of chalk that look-- I guess they're not exactly the same, but they look pretty similar. And let's say I want to follow-- I'm going to call this one Fred and this one George. And I'm going to follow their movement throughout space.

So if I give them an initial position and some initial velocity, I know throughout all time which one is Fred and which one is George. Because I knew their exact initial position. I knew their exact initial velocity. I can calculate their trajectories. I know which one is which.

But that's not true for electrons. So does anybody know the uncertainty principle? Can somebody tell me what the uncertainty principle is? Yeah. Did you volunteer?

AUDIENCE: No, I didn't. I [INAUDIBLE] algebraic [INAUDIBLE].

MICHELE: Right.

AUDIENCE: [INAUDIBLE].

MICHELE: Right, so the exact mathematical formulation is something like the uncertainty in the momentum times the uncertainty in the position has to be always greater than or equal to basically some constant. So if I have my two starting electrons, and I have some idea of where it starts, and with some idea of what velocity it starts with, and some other electron-- with similarly not an exact position but some kind of position and some kind of knowledge of the velocity-- it's clear that over time, I'm going to have no idea which one was which. They'll switch and I won't be able to tell that they've switched because I can't follow their trajectories the same way I can with classical objects.

And furthermore, electrons don't have any identifying marks. They don't have like a beauty mark somewhere that can tell us which one was the electron that we really fell in love with first. And so electrons act completely indistinguishable. And so this is a symmetry we're going to exploit.

So let's say we have a system full of electrons, and I pull a curtain in front of you so you can't see the system anymore. I take one electron, and I take another electron, and I swap them exactly. I swap their positions, I swap their momentum, I swap everything about them. And I open that curtain again, you won't be able to tell at all that anything's happened, because these electrons act exactly like each other.

So it doesn't matter if I swapped one with another. So this might sound a little bit dumb. I mean, you might be sitting there thinking, OK, this is obvious. Why is she going on so long about something so obvious?

And the reason is when we formalize this in terms of mathematics, suddenly things become really interesting. So we're going to define an exchange operator, which is this χ_{12} . So χ_{12} acting on a system that has an electron 1 and electron 2, all it's going to do is swap electron 1 with electron 2.

Just so this formalism makes sense, this 1 refers to what particular wave function it is. So this ψ_1 refers to that set of four quantum numbers, or set of whatever quantum numbers describes that system. And similarly, ψ_2 describes a different set of quantum numbers, and r_1 is the position of electron 1, r_2 is the position of electron 2.

So when we use this exchange operator on our system, we get something that should be indistinguishable. And you can see that if we act twice-- so if we have electron 1 and 2, and we swap them, and then we act again with the exchange operator and we swap 1 and 2, we'll switch back. So we know mathematically that this exchange operator, acted twice on a system, leaves you exactly with the original system that you had.

So we have to assume here that this wave function is going to be an eigenstate of the exchange operator. And there is a quantum mechanical reason for that. But what this tells us, assuming that we've now acted with this exchange operator twice, we know that the value of the eigenvalue of the exchange operator squared is equal to 1, which gives us that the eigenvalue itself has to be plus or minus 1.

So just to be clear, what I've said is that when I act on this wave function with this exchange operator-- so I have my total wave function, whatever it is, which is a function of multiple electrons, when I act with this exchange operator on it, I get something where I've switched r_1 and r_2 , and all the other electrons are left the same. And what I said in the previous slide is that these two states-- this state and this state-- have to act exactly the same.

But now I've also told you that this is equal to some eigenfunction times the original wave function. And this eigenfunction, I've said, can be negative 1. So can anybody reconcile this for me? How can this eigenfunction be negative and yet have this original eigenstate not be changed, or have no measurable difference between the electrons being in the original positions or 1 and 2 being switched?

So the key to the question is actually the term "measurable difference." So remember, Jeff has said that in quantum mechanics, we have this wave function. We have no idea what the wave function actually means. It's the square, the absolute value square of the function, that actually is meaningful.

So you can see, even if this is a negative 1, and we have a negative 1 in front of the wave function, when we square that, the negative is going to disappear. So we still have this symmetry because the thing that's measurable is the absolute value squared of the wave function, and not the wave function itself. But this is actually a pretty key piece of quantum mechanics.

So we have two possible values for this wave function. So this wave function can be positive 1 or it can be negative 1. So any quantum mechanical system is going to have a wave function that either has the eigenvalue for this equation positive 1 or negative 1.

When it's positive 1, we call them "bosons." So this leads to a lot of interesting physics, like the Bose-Einstein condensate, which I don't know if you guys have heard about. There's a lot of cool experiments that have been done on helium 4, where at a low enough temperature, all of the helium atoms suddenly find themselves in exactly the same ground state. And there's some really bizarre consequences of that.

However, what we're interested in are the particles with an eigenvalue of minus 1. Those are called "fermions." Electrons are fermions. So electrons, therefore, need to have an eigenvalue of negative 1.

And this is actually the key to the Pauli exclusion principle. So this is why the Pauli exclusion principle is true. So can somebody explain to me that? It's kind of a subtle point. Does anybody see that?

So let's look at a wave function that doesn't follow the Pauli exclusion principle and see what happens. So let's say I have two electrons. I'm going to put them both in the 1s with spin up. So what happens when I operate with this exchange operator on this?

So these two are exactly the same. So I needed a negative 1 here, but I have a positive 1. So because this exchange gives me the wrong sign, I know anything that has some system like this, where I have the same quantum numbers on two different particles, has to be a boson and not a fermion.

So let me just give you an example of a wave function that does satisfy the proper eigenstate for this. So if we have something like $\frac{1}{\sqrt{2}}$ 1s up for particle 1, 1s down for particle 2, minus 1s down for particle 1 times 1s up for particle 2. So can everybody see? If I switch 1 and 2, I'm going to get the negative of this. Because these two are going to switch and become this term, these two are going to switch and become this term.

Is that clear? I'm seeing a lot of blank faces. I hope that means that people are way past this and have mastered this long ago. Are there any questions on this?

To me, this is really exciting. This is great to be able to finally explain something that Paul himself couldn't explain. We actually can understand quantum mechanically why electrons behave the way they do, why we have the Pauli exclusion principle, and therefore, why we have the periodic table. I mean, this is a pretty big consequence.

[DOOR CLOSING]

I guess it was so big, he couldn't take it. All right, let's move on to what the implication of this is in the Hartree method. So a guy named Fock came along, and he has either a very fortunate or very unfortunate name, depending on how you feel about him.

And he basically just took an anti-symmetrized wave function. So anti-symmetrized just means this characteristic, that when you act upon it with this operator, you get a negative 1. So he just put in an anti-symmetrized wave function into Hartree's original equation, which was just, if you remember, just this. It's the kinetic part and then from the ions, the potential part from the other electrons.

So this was the original equation. And when he plugged in an anti-symmetrized wave function, suddenly, this term appeared. So this is the term we call the "exchange term."

If you look very closely at it, you'll see that what it does is quantum mechanically, two electrons with the same spin will never be in the same position. And it basically subtracts out that possibility, of two electrons with the same spin being in the same place. So this actually reduces the energy of the system because suddenly, all electrons with the same spin have to be just a little bit farther apart from each other. So that minimizes the Coulomb interaction just a little bit.

So actually, with this correction, the results seem to come and match up with experiments a lot closer. And this is actually the foundation for molecular orbital theory. So just by adding this spin consideration into the Hartree equation, we suddenly have an equation that's actually pretty functional. So Jeff thinks this is quite an emotional moment.

However, as I said, there were two energy terms that we were missing in that original Hartree equation. So we found the exchange term, but we're still missing the correlation term. So there are some ways to deal with this, but I think the way we're going to deal with this is just by moving on to DFT.

So now, we're going to start talking about DFT, which I think is what you guys are mostly using, so you might find this more interesting. So the Schrodinger equation is just really hard to solve with all of these electrons. As we've said, it takes a lot of time. It takes a lot of computational expense.

So let's divide space into the worst grid you could imagine. I guess the worst would be one point, but the worst grid you can imagine is 2 by 2 by 2. And let's think about the number of points that we need to keep track of if we're calculating something with little n electrons.

So if you only have one electron, it's only eight, and it's not too bad to keep track of. If we go up all the way to 1,000 electrons, suddenly we have this huge number of points that we have to keep track of. And that's just inconceivable, even with today's computers.

Why try to do something this big? That would take forever. So it'd be really nice if we could think about some other aspect of the system that we could calculate, instead of keeping track of every single one of those electrons separately.

And a term like that is density. So if we consider density instead, suddenly, it's pretty much the same scaling no matter how many electrons you have. So the electron density definitely seems to be a lot more manageable than dealing with each electron separately.

And so at some point, someone's, like, well, wouldn't it be nice if we could just calculate with density instead of with wave functions? Because think about it-- these wave functions don't mean anything anyway until you've squared them to get the density. That's the only thing that we can measure anyway.

Why don't we just look at the density? So what if I could come up with a function of energy that was dependent on the density of electrons? And Walter Kohn actually won the Nobel Prize in Chemistry in 1998 with somebody else for coming up with this idea.

So as I said, one reason that we want to go to DFT-- it scales so much better than the quantum chemistry methods. And we're going to have to spend so much less time, especially as the systems get larger. So let's look at an example.

So if we have two atoms of silicon, DFT and the two sample quantum chemistry codes all take about the same amount of time. If we try to calculate 100 atoms of silicon, suddenly DFT is 5 hours, which might sound long. But the Moller-Plesset is a year, and the coupled cluster approach is 2,000 years.

I don't think anybody's tested that, but I wouldn't recommend it. But suddenly, you can see why if you want to do anything more than just two atoms, maybe DFT would be a good way to go about it. So this is a graph just showing, again, that DFT scales better than all the other approaches.

So again, the basic tenet of DFT is we've taken this really complicated, multi-electron wave function, which we already had trouble enough calculating in the quantum chemistry method. I mean, we already found problems that if we tried to simplify this by making this a function of each individual electron, we suddenly lost all the electron correlation and exchange, which is pretty crucial to getting these calculations right. And we know that the square of this wave function is the electron density anyway.

So let's just think about the density. So does anybody know what a functional is, any math people?

AUDIENCE: [INAUDIBLE].

MICHELE: Great. Great. So generally a function, you input a number and output a number. A functional, you input a function and output a number. So this is called "density functional theory," because we are writing a functional of the function density.

Density is a function. At every point in space, there is a particular value of the density. And so we would like to be able to solve this energy as a function of the density or I guess a functional of the density.

So of course, we're going to see all the familiar terms showing up-- the kinetic energy, the ion potential, the ion-electron interaction, electron-electron interaction. We're going to see all the same terms showing up. And we'll just have to deal with them slightly differently.

So the way we actually go about solving this is we write some equation of all of the different terms, as far as we know what they look like. We write out the density in terms of its wave functions, and then do a functional derivative of one of these wave functions. And so the functional derivative works-- it's basically just a variational principle.

Are you guys familiar with the variational principle in quantum mechanics? The lowest energy possible for a system is the ground state. And so any time that you can find any system that gives you a lower energy, you're getting closer to the ground state.

So I have my basic Schrodinger equation. If I can change something about this wave function such that when I solve this equation I get a lower energy, I probably have something that's closer to the true ground state than I used to have. Because the smallest possible value for this energy that I can get out of this equation is going to be the ground state.

So that's actually used a lot in quantum mechanics. So maybe I don't know exactly what this function looks like, but I can kind of guess what this function looks like. And I can minimize this energy with respect to what exactly this wave function looks like.

And that's exactly what we're doing here. We're minimizing this energy the best that we can, to try to get something that's as close to the ground state as we can. So that leads to the Kohn-Sham equations.

Again, this is pretty much the same Hamiltonian you've been looking at all class, all last week, too. We have the kinetic term. We have the potential term, and that acts on a wave function, and gives us some eigenvalue times the wave function.

So in this case, our potential is broken up into a couple of different terms. So we've already made the Born-Oppenheimer approximation. So we're ignoring the ion-ion potential. We can add that in whenever. That doesn't change.

We have the ion interaction with the electrons. We have the electrons interacting with themselves. And now we have this new term that's called the "exchange correlation potential." So this is the energy right here that we were missing before in the quantum chemistry approach.

But there's a problem-- we actually don't know what this exchange correlation function looks like. So people use a bunch of different approximations. The two most commonly used are the local density approximation and the general gradient approximation.

And they're pretty self-explanatory. The local density approximation, you just approximate the local density as being a flat in a local area. And the general gradient approximation, you assume a first derivative.

So if you have some potential that looks like this, the local density approximation might approximate it as looking something like that, whereas the general gradient approximation would get the slopes for each. So generally, the general gradient approximation ends up being a little bit more accurate, but they both have their strengths and weaknesses. And depending on what problem you want to solve, you really need to pay attention to what you should expect from either one.

So how does this process actually work? So this is a lot like the solution to the Hartree equation. We start with our guess at the position of the ions, which may or may not change.

From there, we also need to start with a couple initial parameters that we didn't talk about for the Hartree case. So in this case, we're going to pick a cut-off for the plane wave basis. So that's something I'm going to talk about in a minute. I haven't talked about that yet.

Then we guess what our electrons look like in the system. From that guess of the electrons, we calculate the terms in the Hamiltonian that we're missing. We solve the Hamiltonian. We just diagonalize it. And that diagonalization gives us a new density.

And then again, we ask, is this density close enough to our input density or is it pretty far away? And if it's pretty far away, we go back up, use that new calculated density, again to calculate the terms in the Hamiltonian that we were missing, diagonalize the Hamiltonian again, get a new density. And we just keep doing this again and again until we've found something that we say is close enough to the original density that we've probably found something that's a good approximation.

So you'll notice that this is actually different from the Hartree. So this looks very similar to the solution of the Hartree equation, but you'll notice that instead of solving one electron wave function at a time, every time we have to do the self-consistent cycle. All we have to do here is just diagonalize this Hamiltonian. So this is why the process speeds up so much.

There are a bunch of software that people use for DFT calculations. Some of them are free. Some of them are not. I think you guys are using Siesta, which is not on this table.

But all of these different software have their strengths and weaknesses. And depending on what you want to calculate, you should probably think about which model is the best for you. But we're going to be looking at a PWscf input file, just to give you a general idea of what these input files look like. and I guess pretty soon you'll find out anyway, since you'll have to write some of them.

So I promised you I was going to talk about basis functions. That was that one parameter in the self-consistent cycle that you have to pick. So a basis function is basically just whatever function you like, what function you want to deal with mathematically, that you can write your wave functions in terms of.

So all we do is we take our wave functions and we expand them. So we get some linear combination of whatever basis we like. So usually, people pick Gaussians, or plane waves, or things that are really easy to calculate, because you can write any wave function in terms of a linear combination of these, as long as you have a complete basis.

So we pick these spaces such that it's orthonormalized. What we can do from that is when you multiply one of these by another one, you'll get a delta function. So basically, if I multiply through by-- so I have to integrate over all space because these functions exist over all space.

So I integrate over all space. I pick one of these particular basic functions that I multiply by. And that picks out exactly the j basis function here.

So this whole term, when I integrate out, because it's orthogonal, I only get the j term here. Because it's normalized, I get 1. So I'm just left with this E and c_j on this side.

On this side, when I do that same integration over all space and multiplication by this j basis function, I get something more complicated because I have this Hamiltonian in between. So this Hamiltonian, this H_{ij} function, is actually not quite as simple as it looks. It's an integration over all space of the Hamiltonian acting on the basis set i , and then the inner product of that with basis set j , the complex conjugate.

But this is something we can calculate. This actually ends up being the Hamiltonian that we diagonalize in the end. So you can see that the choice of these basis functions, to make it something that's easy to calculate this term, is going to be really crucial for getting the shortest possible computation time.

One that's commonly used is plane waves. This is a plane wave. The exact form of the ones we use, $e^{i \mathbf{G} \cdot \mathbf{r}}$. You guys know the e^{ix} is equal to $\cos x + i \sin x$.

So you can see why these are called "plane waves," because in real space, they're basically just some wave. So this \mathbf{r} is the real space vector. What we have to do is pick what \mathbf{j} 's we're going to use in our simulation.

So really, if we're going to expand any function in terms of this basis, we need the spaces to be complete, which means basically what it says. It's "complete," meaning you can expand anything in terms of it. But we don't really want a complete basis set, because for this to be complete, this \mathbf{G} would have to go to infinity.

We'd have to go for every number for \mathbf{G} between 0 and infinity, and we don't have the computational time to do that, because the size of our Hamiltonian that we diagonalize depends on the number of these functions that we put in. So what we hope to do is pick a maximum \mathbf{G} that's sufficient to capture all of the physics that happens, but not so big that our computational time takes forever.

So as I increase G , I get shorter and shorter wavelengths. And at some point, this wavelength is going to be so short that it's going to represent a really high kinetic energy system. And it's going to be just un-physical. I mean, if I have an atom here, an atom here, an atom here, probably the variation in the electron density is not this great right around that nucleus.

And so at some point, we can just cut off G , and say we have enough. We can probably capture all the physics we need just going up to the G that we have. So you can see actually how this plane wave basis works really well for periodic crystals.

I just drew nuclei here. So when you have an infinite crystal, you have a nucleus here, a nucleus here, a nucleus here. And because they're all the same, you don't expect that the wave function or the density should be different at this nucleus than it is at this nucleus.

So this plane wave is actually a great representation of that because the density is the same here as it is here, as it'll be here. And so these plane waves are actually ideal for calculating periodic crystals. So how are we going to attack molecules?

What we're going to do is put these molecules in a box. We're going to put them in a box with some amount of vacuum on either side, and then repeat that box in all directions. In x , in y , in z , we're going to infinitely repeat this box. And so now instead of having an atom here, an atom here, an atom here, we're going to have a molecule here, and a molecule here, and a molecule here. So you can see how maybe with this, these plane waves might represent that system a little bit better.

But there are some things you have to take into consideration with this. Since there's a molecule here and some distance away there's another molecule, they might interact with each other. And if you want just the properties of that molecule by itself in vacuum, you don't want them interacting with other molecules.

So you need to make sure that you've included enough vacuum between these molecules, between the molecule and the edge of its unit cell, that the interaction will be negligible or will be on the order of the error that you have from everything else in your calculation. So once we've done this, we can then start our self-consistent loop. So what happens?

We put in where the ions are. We calculate our potential. We've guessed at some original electron density.

We've calculated all the terms in the Hamiltonian. We've diagonalized that Hamiltonian. That diagonalization gives us our energy.

And then the program said, well, that new density wasn't really close to the old density. So let's do that again. And it does that again, and calculates energy.

And you can see how the energy always goes down. So this is a negative number. The energy always goes down.

So that's the variational principle-- as you approach the correct ground state, you get lower in energy. And you can see that the difference between these successively gets smaller and smaller. And that's what you really want to have happen, because that means that your energy is converging, and you'll end up with a pretty good approximation of the ground state energy.

You've also ended up with the charge density, because that's what you've been calculating this whole time to calculate the energy. So just with those two properties, and being kind of clever, you can actually calculate a huge range of important molecular and solid properties. So you calculate at least you can guess at a structure.

You can calculate bulk modulus, shear modulus, elastic constants. You can calculate vibrational properties and sound velocity. Binding energy is actually just probably going to be this energy that you've calculated.

You can kind of guess at reaction pathways. Those are a little bit harder to do with this approach. You can calculate forces, pressure, stress. You can calculate all of these properties just with this simple calculation that we can now do.

But how do we know that these calculations are meaningful? How do we know that we've actually said anything useful? Dirac was complaining that he found that his calculations never really matched experiment, but don't lose heart.

But when should you lose heart? Or when should you maybe check to make sure that what you're doing is as accurate as you can get? We need to talk about convergence.

So there are a couple of different properties that you need to worry about. We talked about is the basis big enough. So that's, did I pick a high enough G that I got enough wiggleness of my wave functions, my input basis wave functions, that I've captured all the relevant physics?

And so this is the trade-off. It takes more time if you have a larger basis set, but you also lose some accuracy by decreasing the size of your basis set. And at some point, when you've taken out too many basic functions, you suddenly get something that's not going to be meaningful at all because you just don't have the right structure of the potential.

We talked about is the box big enough. Did you include enough vacuum so that you don't get interactions between molecules in neighboring unit cells? Because if you want the properties of the molecule by itself in vacuum, you don't want the interaction of those two molecules to play a part.

But here we get, actually, a trade-off with the basis set. So the bigger you make your box, the bigger you need your basis set to be. So you can see if this is the size of my box here, and this is my cut-off-- G has a wavelength of this-- if I suddenly make my box much bigger, this wavelength is going to expand and it's going to be a much bigger wavelength.

And so maybe now this is actually an important wavelength to my system. And maybe even a smaller wavelength is an important wavelength to my system. Does that make sense? OK, so making your box bigger means that you need a bigger basis set, which means more computational time, but on the other hand, you don't want to include the effects of intermolecular interaction.

Finally, you can ask if you exited the self-consistent loop at the right time. So I keep saying, well, we compare the old density to the new density and if they're close enough, then we assume it's right. But what's that definition of "close enough?" And how do you know that that definition of "close enough" is close enough?

So that's another property that you're going to have to converge along. So let's just take a look at an example input file, and see how some of these properties that you need to converge along, and the other properties you might care about, show up in this system. So we might care about putting in the right atoms, talk about how the atoms are related to each other spatially, and then these three parameters that we talked about that we need to converge along.

So this is just a sample of water. And so here, this is the size of the cell. This is the number of atoms, number of types of atoms. This is E_{cut} , this is the energy cutoff, which refers to the size of your G , and then your atoms and the atomic positions.

And these are just the pseudo potentials. So we haven't talked at all about pseudo potentials, but just so you know, we talked about the fact that when we think about atoms we only care about the valence electrons. So all those core electrons still exist, but we don't want to calculate them.

So we use pseudo potentials to sort of capture those inner electrons a little bit, and tell us how those inner electrons are going to affect the valence electrons and everything else. Because when we have an atom that has some core of electrons-- so let's say this has 27 protons in it. An electron out here is not going to feel that 27 proton charge. It's going to feel that 27 protons shielded by some number of core electrons.

And so to get the Coulomb interaction right there, we need to have some kind of modified version of what exactly our atoms' potential looks like. So that's what the pseudo potentials are about. So that's it.

We've covered the basis of quantum chemistry and density functional theory today. Are there any questions? All right, thank you.