

PROFESSOR: OK. Let me also remind you that we slightly changed the schedule for next week. On one of the older schedules, Tuesday appears as a lab date, but it's actually Thursday. So Tuesday will be a regular lecture here and Thursday will be the lab. And for the lab, we meet in 1115. That's on the handout for the first day.

Also, if you're not registered for class, you don't have automatic access to Stellar, the website. So we can manually add you, but you'll have to let us know. Just send an email to Professor [? Marzari ?] and myself. And even better is if you send an email to both of us. And then we'll manually add that. And we'll put links there to things like articles and things that may have restricted access, copies of the lecture notes, things like that will appear there.

OK. So what I want to go do is go back to their potential models. And what we're going to do today, is talk a little more in detail about the sort of formal or conceptual failures of pair potential models, and then how to get addressed with other empirical models. But before I do that, I wanted to-- let's see if I can get this into place, guess not-- go through a few practical issues, since the first lab is actually on using pair potentials to model things so you get a bit of an idea.

So in a potential model, your energy is written pairwise. And so essentially what you need to do between every pair of atoms-- every couple ij -- essentially to evaluate the distance between them and then put that in some kind of energy function. So this is clearly an N^2 operation where N is the number of atoms, since you have to look for the distance between every atom and every other atom. So if systems get extremely big-- and that's often the reason why you want to use the pair potentials rather than, say, a quantum mechanical method-- N^2 will get very large.

And people do kind of tricks. If you have millions and millions of atoms in your simulation, if you know ahead of time that some of them are very far from the atom on which J you want to calculate-- energies or forces-- then you really never want to look further distance. So people do things like keeping neighbor lists. So literally, for every atom you would know, well, these other items or possibly in its neighbor environment. So then you really only have to look for distances with those atoms.

So if you have a solid, that's great. Because if you have a solid, then the topological relation between atoms will not change much during the simulation. Of course, if you're simulating a gas or a liquid, they will. And then what you have to do is essentially update the neighbor lists regularly.

And so there's overhead associated with that. But the benefit you get is that you essentially have an order N system rather than an N^2 system. I'll show you in a second some large scale simulations.

Very often, you will want to relax a system-- so get it to its lowest energy. And very often, that's done by calculating the forces on atoms and literally stepping down the force. If you want to do dynamics, Professor [? Marzari ?] will later teach you molecular dynamics. You also need the force because the force will tell you how much atoms accelerate or decelerate. You will need to calculate the force, which is a potential derivative. So you're essentially calculating the derivative of the total energy with respect to a given position.

Same deal. That's an N^2 operation. You have to sum over pairs of atoms. We have to sum over 1 all atoms for the force on a given atom. But since they're N atoms, it's an N^2 operation.

In a lot of codes, the minimization is done by very standard schemes-- things like conjugate gradient, Newton-Raphson, sometimes often trivial line minimizations where you literally just calculate the force and assume some kind of elasticity along the force path. And so if you know the force, you assume you know what the curvature is of the energy in that direction. You kind of can make an approximation to its minimum when you do that iteratively. So you can do that with very large simulations.

And I wanted to show you sort of one of the leading edge ones. I got to get out of PowerPoint for that if I can make that work. OK, here we go.

So what I'm going to show you is a order million atom simulation. This is essentially a cube of 1,000 by 1,000 by 1,000 atoms. That would make more than a million. I forgot to write down what it was. Yeah, I wrote down 1,000 by 1,000 by 1,000 atoms. That would be a billion atoms.

And essentially what's being done-- and what you see is a notch on top and on bottom because there's so many atoms you don't see the discrete resolution anymore [? of ?] [? the ?] [? atoms, ?] [? so ?] this thing starts to look like a continuum. And it's being pulled from the side. So you'll see the cracked grow and dislocation spread.

This is a simulation from Farid Abraham when he was at IBM. And of course, if you have big computers, you can do big simulations. So let me show it to you. So what they actually did was they showed just the dislocations by counting under-coordinated or over-coordinated atoms.

So what you're actually seeing is the dislocation lines. I mean, if we plotted all the atoms, you wouldn't see anything. You'd just kind of see green everywhere. And that's sort of zoom up. So this is a molecular dynamic simulation. And you learn later more about that, where you essentially just do Newtonian mechanics on atoms. So this is one of these very large scale things that you can actually do with potentials.

Now, you could ask yourself, what did I learn from this? This is the kind of work that my personal opinion is scientifically not necessarily all that relevant. But it's pioneering work, you know. This is like Lewis and Clark who went to the West. You could always say, what did we learn from them? Maybe not a lot, except that you could get the Pacific. But it's kind of pioneering work.

And when people do these kind of simulations, it's not as much for the science as it is probably for pushing the envelope and seeing what can you do with enormous computational resources. You really have to fine tune your algorithms, your parallelization scheme. So it's kind of like flying to the moon which may have not brought us much. But there was a lot of technology fine tuned for that. So that's one of the reasons I wanted to show it.

OK, let's go back. Another very practical issue that I'm sure a lot of you are familiar with. Even when you have a billion atoms, it's actually still very small. I actually think if you do the calculation, a million atoms is, for a lot of materials, a cube of just about 200 Angstrom on the side so you're still, if these were finite systems, would be very small systems.

So to get rid of boundary effects, you almost always use what's called periodic boundary conditions in simulations in as many directions as you can. And of course, what a periodic boundary condition means is that if you have some unit, if you have some unit, you essentially repeat it next to it. So this end here is the same as that end, and this end here is the same as that end.

So people always ask you how many atoms do you have in your simulation, and the answer is infinite. It's just that you only have a finite number of degrees of freedom because it's these atoms here do exactly the same as these here. But the answer is infinite. So we don't really work with small systems. Or, rarely. I mean, sometimes we work with finite systems.

So what's the repeat unit you use? Well, if you're studying a periodic material, then it's easy. If you study, say, a perfect crystalline material, here's a perovskite unit, so that's what you repeat. Because the material already has periodicity implied in it. So at that point, you're actually not making any approximation. You're not forcing anything on the material-- not when you do a static calculation that it already doesn't have.

Of course, when you do dynamics, you're enforcing something because you're saying that in dynamics, the translational symmetry is temporarily broken because this atom can vibrate independent of that one. And so, then, when you improve impulse periodicity, you are imposing that these vibrate together. So in some sense, you're imposing a cutoff on the frequencies of vibrations that can occur.

If systems don't have periodicity, you typically tend to enforce it. If I took this perovskite and I made one defect-- say I wanted to study oxygen vacancy so I take out the oxygen vacancy, I lose all translational periodicity now. But what you do is you tend to impose it back by repeating the defect. It's going to be slow.

So let's say I've taken out an oxygen here. I repeat the defect so that, in essence, I just have a bigger unit cell. But it's one with a defect in it. This is now my unit cell. And that's called a supercell approximation. So you make a cell that's big enough so that you hope that these defects-- I don't know how you call it-- don't interact because then, if you have N cells, your calculation just has the energy of N defects. OK?

And the critical assumption is, of course, that they don't interact. Because if they interact, then you're not modeling an isolated defect. You're modeling a defect that's interacting with other ones.

So how do they interact typically? It's kind of important to know. This is one of these things that you're going to do in the lab. You should always check conversions if you can. The problem is that often you already work at your computational limits for size of supercells, so it's easier said than done. In an ideal world, you say, well, I calculate bigger and bigger, and if the answer doesn't change, I take it. But usually, you work so much at your computational edge already that you can't make them bigger and bigger.

A lot of people forget that [? to ?] [? check ?] [? conversions, ?] you also can make it smaller. And if it doesn't change much, you were probably already converged. But it's useful to know how these things usually interact with each other. So that gives you some sense of how far you have to go.

There's often the direct sort of interaction. If you work with a pair potential, then you can actually see the range of the pair potential. That's usually not the problem. Rarely do you have potentials that are relevant, large, and over large distances that you need to make really large supercells. But there can be other energetic effects. If you have electrostatics in your system, that will give you very long range interaction. If that's a charge defect-- which it is in a perovskite-- you will have a very strong one.

Actually, if we truly did a charge defect in that supercell, what would happen? If that supercell became charged-- so let's say I took out an oxygen 2 minus, which means the supercell now has a net positive charge-- what would happen? You'd have an infinite energy because you would essentially have a positive charge interacting with only other net positive charges. And at any distance, if you do an infinite system, that gives you an infinite energy. So that's something that you would have to explicitly take out.

Now you say, well, I'll compensate with a negative charge because that's what happens in reality. If I take out an oxygen, those two electrons of that or O2 minus are going to go somewhere, so it's going to create a net positive charge. So then you have a dipole. So let's say I have sort of plus, minus, and I have a dipole now. So now I have dipoles interacting. Dipoles will not give you infinite energy, but they will give you a fairly long range energy tail.

And if you're smart about it-- so the brute force way to convert supercell calculation is always to push him as hard as you can and hope and sort of see the energy converge of the defect. But if you're smart about it, you can often do much smaller cell calculations and explicitly take out the term that's the defect-defect interaction. Which is, if you know that this is a dipole, you know what the form is of it, and you can take it out.

Other things that often cause havoc is relaxation. If I were to put a really big ion in the center of that perovskite unit cell, I would build up a strain field and the different defects would interact through their strain field. So, rarely do you actually interact through the direct energetics. You often interact through secondary effects.

Another way they often interact which is very indirect. If you put in your defect and you relax the volume of the supercell, then that will interfere with the relaxation of the supercell next to it. So you get a strain field just through the homogeneous. You get an extra energy term just through the homogeneous strain interaction. So these are all things to consider when you do supercell calculations.

We're trying something new where we save the annotations we make on the slides, but it seems to be really slow. So I may not be able to keep that up. OK, so this I went over-- the kind of convergence.

You know, here's a typical example. I don't know if this is typical. This is the vacancy formation energy in aluminum. Now, this is done with quantum mechanics, versus the number of atoms in the supercell.

And you know, this is one of these that you really don't know if you're converged. I mean, it depends on how accurate you want it. If you really only care that it's between 0.7 and 0.75, you're probably OK. But if you did this point and that point, you'd probably think you're converged. But then, if you look at the next one, you're actually going back up a little bit again. So it's sometimes very hard to figure out whether you're actually converged.

OK. So that was a few practical issues to sort of start gearing you up for the lab. What I want to do now is talk explicitly about some of the formal failures of pair potentials and go on to correct them. And the first one is the vacancy formation energy which is, if you look at historically when people started doing simulations late 1950s, 1960s, into the '70s, people thought that whenever there was something wrong with the outcome that they just needed better parameters in the potential.

And it's actually rather recent-- sort of 1980s-- that people figured out that there were really formal problems and that there were certain things you could not get right-- never-- with a potential because the problem of the form rather than of the parameters. And a critical number in that work was actually the vacancy formation energy which nobody could ever get right. And I'll show in a second why.

Here's an FCC crystal. So let's calculate the pair approximation to the vacancy formation energy. What do I do when I make a vacancy? What I do is that I take this central atom and I move it somewhere else in the bulk. So in some sense, I make my system one side bigger. So let's calculate the energy balance of that. So the vacancy formation energy.

OK. What do I destroy? I destroy 12 bonds because this atom is coordinated 12-fold in FCC. So I destroy 12 bonds. I lose 12 times the bond energy. But then, of course, I put it back somewhere in the bulk. And so I gained the cohesive energy of one atom. OK?

And what's the cohesive energy per atom in FCC? It's 6 times the bond energy because there are 12 bonds around an atom, but every bond is shared between two. OK? So there's six per atom. So the vacancy formation is really minus 12 times the bond energy plus 6 times the bond energy. So what it is is minus 6 times the bond energy. So essentially, it's minus the cohesive energy per atom. And that's essentially what you will always find repair potentials.

Now, I did a static approximation. I assumed that all of these bonds were kind of unchanged as I took the atom out. In reality, if I take that central atom out, the atoms around it will relax somewhat and change their bond energy. But you'll see this when you do this in the lab. That will make it from one time the cohesive energy to something like 0.9 or 0.95 times the cohesive energy. Essentially, you find in pair potentials that the vacancy formation energy is about the cohesive energy.

If you look, on the other hand, at experimental results for the cohesive energy-- [INAUDIBLE]-- here's some experimental data. So this is for FCC metals, and this is for Leonard-Jones, and for noble gases. The ratio of the vacancy formation energy to the cohesive energy.

And what you see is that, in the noble gases-- which you think of as well-described by pair potentials-- you're reasonably close. Because remember, noble gases are just inert shells interacting through Van der Waals interaction, and that's how we came up with the Leonard-Jones potential.

But if you look at the metals, the vacancy formation energy is actually only a small fraction of the cohesive energy. And we can never get that low with pair potentials. Never. OK? Unless you maybe really have the most odd form of a pair potential. But the reason is because there are certain physics missing from pair potentials.

So actually, let me go back for a second to the pair potentials. What's the physics here of why the vacancy formation energy is so low in real materials? Remember what I said before, that bonding in metals, the bonding energy goes not linear in the coordination number. It's more something like a rough approximation would be the square root of the coordination number. But essentially, when you go from 1 to 2 bonds, your energy goes down a lot more than when you go from 11 to 12 bonds.

So why is the vacancy formation energy so low? Well, if you think about what you do when you take out an atom, you make a lot of the atoms around it go from 12 coordination to 11. So they lose a lot less than 1/12 of the cohesive energy because the cohesive energy you could think of as the average of all bonds. But around the vacancy, you are just reducing the coordination from 12 to 11.

You can actually make a simple model if you said that the cohesive energy was some constant times the square root of the coordination. So rather than linear in the coordination, which is what we would get in pair of potentials, if you calculate the vacancy formation energy in this model-- let's do it for FCC. In FCC, Z is 12.

The vacancy formation energy, what would it be? Well, if we put the atom somewhere back in, we would gain Z times the square root of 12. And then, as we take it out, what do we lose? For 12 atoms, we go from coordination 12 to 11.

So this is the atom in the bulk. This is actually the change around where we made the hole. And if you actually calculate that-- the vacancy formation energy over the cohesive energy-- you find that that's 1 plus 12 times square root of 11 over square root of 12 minus 1. And it's about 0.49. So in a square root model, you're already much closer to the experiment where the vacancy formation energy is a much smaller fraction of the cohesive energy.

Another critical piece of information was surface relaxation, which people had been seeing with [? LEEDs. ?] If you look at a metal, typically if you create a surface, the difference between the first and the second layer becomes smaller. So in a metal, the first layer tends to relax inward.

And can somebody explain to me why that happens? And typically, you don't see this. In pair potentials, you very often have about the same distance. So the first to second neighbor distance is often bulk-like and in some cases actually relaxes outwards because of competing forces. But in almost all real systems, it tends to relax inward. Why do you think that is? Yes, [INAUDIBLE]?

AUDIENCE: [INAUDIBLE]

PROFESSOR: True. But I'm not sure that that necessarily makes it relax inward. The reason it relaxes inward is because the bonds between the first and the second layer strengthen. That's really what is a sign of them. And why do they strengthen? Because the surface is under-coordinated. Remember, if you go to lower coordination, you're effectively-- although you're losing energy and that's why there's surface energy-- but the energy per bond is actually becoming stronger.

Remember that I showed you the potentials for copper? Remember that for the copper II molecule-- so where you just had one bond-- had the strongest potential. The equation of state gave you where you were 12-fold coordinate gave you much weaker potential. So you see that in surfaces as well. If you cut away the bonds here, the remaining ones strengthen and they pull the layer in.

The other problem that's common in pair potential which I won't say much about is what's called the Cauchy problem. If you write down the relation between the stress and the strain tensor and so this is the matrix of elastic constant that for pair potential C12 and C44 are always the same. And this has something to do with there actually being a zero-energy strain mold.

Remember that I showed you this? It's related to the same issue that in a cubic or in a square model, you can actually kind of slide it without any energy change. And that actually puts a constraint on the elastic constant, which essentially translates to C12 being C44.

And the last one-- if you're a practicing material scientist-- may be one of the important ones. But you can essentially not predict crystal structure in anything that's covalent. And metals are covalent. I know they teach you there's metallic, covalent, then ionic bonding. But metallic is just delocalized covalent.

And the reason is-- I think you can already see-- first of all, potentials tend to go to close packing always because they just they want to have as many atoms around them as possible because that's how you bring the cohesive energy down. And they don't count that cost, that if you bring more around you, your bond strength weakens. And they have no directional dependence, so they miss those two important components to get crystal structure right.

You can actually do some theory about this, which was done already in the '70s that you actually need what's called at least the fourth moment of the density of states to get crystal structure right. And that turns out to be at least a four-body effect. So if you actually thought of potentials as an expansion-- so I'm going to write the energy as a constant plus an expansion in two-body terms, three-body terms, four-body terms, et cetera, which we don't even know would converge-- you would have to go at least up to the fourth order to get crystal structure differences right.

OK, so how do you fix the problem? There are essentially, if you look at the literature of empirical models, two avenues to go. You can either make what's called cluster potentials. And so, if you don't like a two-body potential, you can make a three-body potential.

And the advantage of that, once you have a three-body potential, you can see angular dependence. Because if I have three atoms and they don't interact just pairwise, but there's a [? two ?] three-body, I can include this bonding angle. A bonding angle is truly a three-body concept. So you can start building an angular dependence.

The other direction to go is go to pair functionals. So you still want to write the [? energies ?] pairwise-- some pairwise summation. You just don't want the cohesive energy per atom to go linear with what you sum pairwise or what you sum as its environment pairwise. So we're going to deal with both in class.

Today, I'm going to mainly do pair functionals. Pair functionals tends to be most applicable to metals. Where I would say the square root depends, that effect is strongest, whereas cluster potentials or many body potentials in general tend to be more favored in the organic literature-- both organic chemistry and polymer chemistry. And so, going to that.

So there's a whole class of methods which are called effective medium theories. The version of it that I'm going to discuss is what's called the embedded atom method, which is probably the most popular one. This was developed by [? Mike ?] [? Baskis ?] [? and ?] [? Murray ?] [? Dahl ?] I think going back to the '80s. Mid '80s, or something like that. And it's probably now the most popular of the sort of empirical energy schemes. If you look back at those papers, those papers have now over 1,000 citations.

And the idea is extremely simple. It's, if you understand the source of the problem-- that our energy should not be linear with coordination-- that's what you're going to fix. So in these models, what you do is you write the energy per atom as a function of the number of bonds around it. It's just that F should not be linear. And so that's why they call it energy functionals rather than a pair potential.

The question is, how do you measure number of bonds? And if you sort think about it, if you think yourself back in the '80s, how do you want to measure number of bonds? You don't really want to do it discretely. Because you could say, well, maybe I'm just going to count atoms within a two-Angstrom radius, or some radius. The problem is, as soon as you measure it discretely, you're going to have to deal with discontinuities.

As atoms move in and out of that radius, suddenly you've got 14 things around you instead of 15. So you would always have super large forces because of those discontinuities around your [? cutoff ?] interface. So what you need is a smooth count of coordination. And the one they came up with in the embedded atom method is literally, you measure the electron density being projected on you from your neighbors.

It's ironic. I'll show you later the formal justification for the embedded atom method which has a lot more to do with quantum mechanics. But that was actually a justification after the fact. The original idea came up with sort of, you really want to have a continuous measure of bumping around you.

And so the idea is that you're going to do something like, I have a central atom. I have other atoms around me. How do I measure how much these contribute to my neighborhood? Well, these have some electron density functions around them. Their atomic electron density I'm sort of drawing. And they all project onto that central atom. You count up how much electron density there and that's somehow a measurement of your neighborhood.

And you see, that's going to roughly have the right behavior. It's continuous. As atoms move farther away, they still contribute, but they contribute less. So it is a continuous measure. And the more density you project, either the closer by the atoms are or the more of them there are.

OK. So this is the idea of the electron density. You're going to actually calculate the electron density on a site as a measurement in some sense of the coordination as a sum of electron densities coming from its neighboring atoms, J. OK? What electron densities you use is kind of less relevant.

In the original embedded atom and in most embedded atom codes still, they use essentially the atomic densities. So people have calculated the atomic electron densities with [INAUDIBLE] methods. And they're actually tabulated in these Clementi and Roetti tables. And you know, I'm just sort of showing you that they're basically parameterized by a bunch of Gaussians, but you can get those. So you can get those functions. We skipped through that. OK. So remember that if you have-- yes, sir.

AUDIENCE: [INAUDIBLE]?

PROFESSOR: Correct. Correct. Well, it's a distance dependent function.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Exactly. Yeah. So it's sort of what people in quantum mechanics think of as sort of non self-consistent methods. You literally treat the electron density of the crystal as the overlap of the atomic densities of the atoms. The reason that saves you here is that you're actually not going to do quantum mechanics on that density, and you're just using that density as some measure of coordination.

So [INAUDIBLE] I skipped some slides. Somehow, I lost the slide in here. Sorry about that. OK. You know what? I put it in the wrong place. Here we go.

So the way you write the energy in the embedded atom method is you write it as the embedding energy, which is you sum over every atom. And you're looking for its cohesive energy, which is some function F of the neighborhood measurements, which is the electron density at the site of i . And that's your summing from these functions F which are the projected densities from all the other atoms J .

So this is the nonlinear part of the energy. Typically, that is supplemented with a pair potential. So the standard embedded atom method is a pair potential which is often just used for having a repulsive part and a nonlinear embedding energy.

So that means we have to solve three problems, really-- the one we've solved. You need to know what the electron densities are. Take the Clementi and Roetti tables. You need to know what the embedding function is, and you need to know what the pair potential is.

The pair potential, to be honest, you can use almost anything. It's mainly used for the hardcore repulsion part. You don't want the atoms to fly into each other. Typically, they use some screened electrostatic form which is essentially the product of charge distribution functions on atoms.

But you know, nothing would stop you if you really wanted to from using a Leonard-Jones. I'll show you in a second that there is ambiguity about the division of energy between the embedding part and the pair potential is not unique. So that means, every time you change your pair potential, you'll have to change your embedding function. I'll show that in a second.

What do you use for the embedding function? There are two ways that it can be done. People have written down analytical forms with some theoretical justification that I'm not going to go into. There's actually an excellent review article on the embedded atom methods which we post just for internal use on the website. It's a 1993 *Materials Science and Engineering* report. It's a really long paper, but it's really an excellent paper.

So there are analytical forms. And here's a few of them given. And people have come up with other ones. More and more now, what people simply do is tabulate the embedding function. And really what they do is they tabulate it so that you exactly reproduce the equation of state.

What is the equation of state? The equation of state is energy versus volume or lattice parameter. So say at every level parameter, you could try to get that energy with a better method-- say, quantum mechanics-- and then fit your embedding function so that you exactly reproduce that. And so, often, the embedding function is tabulated. But commonly it has this form. It's decaying, it's sort of decreasing convex, decreasing in energy and convex, and it needs both those properties.

Well, why does it need to be decreasing in energy? So this is the embedding density and this is the value of the embedding function. Of course, as you put more density on a site, that comes from more neighbors. So your cohesive energy should be going down. That's why the energy of the embedding function goes down. But it needs to be convex because convexity gives you sort of the correct behavior.

Convexity really tells you that adding a bond going from here to here-- let's say a bond is adding 0.01. So if we added a bond here, we'd get this change in embedding energy. If we added it here, we'd only get that change. So the function has to be convex in some sense. Decrease has to be per unit, and electron density has to be less for higher densities than for lower. And that's the definition of convex. The slope of it has to be increasing of this function.

And you see, that's going to by construction now get some of our problems right because this shape of the function, the convexity is what tells you that the low coordination bonds are stronger than the high coordination bonds. So you can already guess from this that we're going to get a lot of things right. I'll show you some examples in a second.

What's the physical concept that sort of justifies the embedded atom method? There are things you can do formally in ab initio theory showing that in fairly simple tight binding models. So these are things that have the structure of a proper quantum mechanical Hamiltonian, but they are parameterized forms of a Hamiltonian. But still, what they all tend to show is that the cohesive energy in highly delocalized systems to first order goes like the square root of the coordination. So you can get some justification for this nonlinear behavior.

A somewhat intuitive concept of where the bonding energy comes from is that actually, if you put an atom in a metallic solid, where does it get a cohesive energy from? For the most part, it gets its cohesive energy from the delocalization of its electrons. And that's largely a kinetic energy effect. Remember that the kinetic energy in quantum mechanics is the curvature of the wave function. It's $d^2\psi/dx^2$ if you do it in one dimension.

So highly curved wave functions have high kinetic energy. So if you can actually delocalize, your wave function becomes, in some sense, much smoother with much less curvature, and you have lower kinetic energy. So if you believe that a lot of bonding cohesive energy and delocalized source comes from that, then the measure of cohesive energy is essentially how well you can delocalize it.

So what do you measure with the embedding density? If you measure the electron density coming from atoms, the reason this works is that that's probably a very good measure over the number of states you can delocalize over. Things with higher electron density tend to have a higher number of states-- higher densities of states. And so you can delocalize over more states.

And that's the somewhat intuitive explanation for why the embedded atom works. I mean, people have given somewhat other explanations. They could say, well, the electron density that I see from a neighboring atom measures the number of electrons I can bond with. But that's not totally true because if those electrons come from a filled shell, there's nothing I can do with them. So somehow, I tend to believe more it's really a measure of the delocalization you see. The embedded atom method is-- yes, sir?

AUDIENCE: Did you say [INAUDIBLE]?

PROFESSOR: Correct.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah. Yeah. So in an atom, you actually have a much higher kinetic energy because your state is much more localized. The embedded atom method is probably the most popular one of all these effective medium theories. But you'll see other ones with other names that are in spirit very much the same-- things like the glue model.

Finnis-Sinclair potentials, which slightly predate the embedded atom method, which are also non-linear potentials. So they're really particle potential pair functionals. And the equivalent crystal models Smith and Banerjee. These are all exactly in the same spirit. The brilliance of the embedded atom method was essentially that it was something for nothing. The computational cost is essentially the same as for doing pair potentials because you sum things only pairwise.

If you think of what you need to do an embedded atom calculation, you have your pair potential. That's a pairwise sum. But your embedding energy is also a pairwise sum. Because what you have to do per atom, you just have to sum the electron density coming from a bunch of other atoms around you. So that's a pairwise sum. And then you just have to stick it in F , which is an extra evaluation. But embedded atom fundamentally scales only like N squared.

So that was the brilliant part of it. Or in the end, that's what made it successful. You could essentially do much better things with almost the same effort as pair potentials. And that's why it sort of took off as a method. And pretty much anybody in metals now would use embedded atom if they want to get any physics right. If you want to do demonstration projects or things like want to do billions of atoms, there can be reasons to go through much simpler forces still to evaluate. But in metals, there's really no reason not to use EM over simple pair potentials.

What do people fit to? People fit to anything, like I told you before. The original EM potentials were fit to a standard set of data, which was lattice parameter, sublimation energy, elastic constants, and vacancy formation energies.

Now you'll regularly see papers in the literature claiming, I have a better potential than you for aluminum, and I have a better potential than you for nickel. And some of that is true. So people are not in the game of kind of super optimizing these potentials for what it's worth. But the original database was a set of standard potentials-- largely unknowable metal alloys. People have spent a lot of time finding better potentials on technologically important alloys like nickel aluminum. You'll find at least five potentials who all claim to be better than somebody else's potential on the nickel aluminum system.

So what I want to do now is just show you some results. And we're not going to analyze these in great deal. But I want to show you so you get some idea of what people do with this stuff.

Here's linear thermal expansion for a bunch of late transition metals calculated in EM and calculated in experiment. And these agree pretty well. That's maybe not a big surprise in close packed solids. The thermal expansion is largely a measure of the asymmetry of your equation of state. I mean, in complex materials, that's not true. But in simple, close packed metals. Remember, your equation of state is slightly asymmetric. It curves less towards the high lattice parameters and the low lattice parameter.

So if you draw energy versus volume, it's asymmetric around the minimum. And that's why you have thermal expansion. Think about it. It's slightly easier to fluctuate to the higher volume than to the lower volume. Yes, sir?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Oh, good point. There is, of course, no temperature dependence in the potential itself. So to do something like thermal expansion, you either have to do a dynamic simulation like molecular dynamics-- and we'll come to that-- or you have to explicitly calculate the entropic factors that give you thermal expansion, which is essentially the dependence of the phonon frequencies on volume.

So you can indirectly calculate the thermal expansion. But we sort of haven't got into-- we will go into micro dynamics. But of course, in the potential, there's no temperature dependence.

Here's activation barriers for self-diffusion in metals. These are almost too good to believe it. You actually don't get them this good with quantum mechanics. So it's almost surprising to me that they're this good, if you compare EM to external. So this is the activation barrier for-- I forget. But I thought it was interstitials self-diffusion, but I'm not for sure whether it was interstitial or vacancy self-diffusion.

You know, here's a bunch of surface energies. These are not quite as good as the experiment. If you, for example, look at copper, here's the experimental number in ergs per centimeter squared.

And here is the calculation on different facets. Because of course, often experimentally you measure some average of different facets. Unless you do careful single crystal work, you have some average. On the calculations you have to calculate the energy on a specific surface. And that's actually one of the things you're going to do in the lab.

So these are a little lower. And that seems to be a rather consistent problem that is, for example, also in platinum. These are quite a bit lower than the experiment. So that may tell you that either there are more complicated quantum mechanical effects going on that you don't capture with these simple energy models. But sometimes, it also means the experiment is wrong. I think in this case, platinum is a fairly easy metal to deal with clean. But surface energies of materials is sort of notoriously difficult to get experimental.

Here's a phonon dispersion curve for copper. Again, don't worry about the details. I just wanted to show you the kind of things people do. Calculations are the lines and the points are the measurements. So, reasonably good at some of these. High frequencies there seems to be some problems with.

Melting point. Again, this would have been done with molecular dynamics and then free energy integration. And again, for most materials, these come out remarkably well.

So, look at this. Except again for the very late transition metals like palladium and platinum, there just seems to be some amount of problems. I can tell you in a second about the limitations of these methods. And they may be related to the unusual electronic structure of palladium and platinum. You know, here's a pair correlation function in a liquid. [? People, ?] [? we're ?] [? good? ?]

A grain boundary. This is something you simply could not do with a pair potential in metals because when you calculate a grain boundary, you really are dealing with different coordinated environments. And the painful thing about a grain boundary is that you don't know ahead of time what kind of coordinations you're going to see. And you're going to see multiple coordinations.

See, but people with potentials always get, if I know I'm going to deal with a surface, I can make a pair potential that's pretty good for those coordinations at the surface. If I deal with the bulk, I make a pair potential that's pretty good with bulk coordinations. It's when you work with grain boundaries that you kind of end up with all kinds of coordinations because the atoms in the grain boundary, some are high coordinated, some are low coordinated, depending on how the grains contact each other.

So it's a great method, but it still has its limitations. One is that the bonding is spherical. You literally sum all the electron density coming from around you, and you really just sum it up. You don't keep track of its orientation dependence. So having one atom there next to me and one atom on that side is the same as having both straight in front of me because I just sum up the electron density coming from them So it's purely spherical.

That can be fixed. About in the '90s, [? Mike ?] [? Baskis ?] developed what's called a MEAM, and that just stands for modified embedded atom method, which basically has embedding functions that keep track of the angular dependence of the electron density around you, essentially by mapping the problem on to spherical harmonics.

MEAM is not nearly as popular as the standard EM because it adds a level of complexity to the method. And it's not totally clear that you're making things work because the physics is better or you just have more fitting parameters. [? Mike ?] [? Baskis ?] originally developed the MEAM to work on silicon where angular dependence clearly seems to play an important role. Yes, sir?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Well, your force--

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah. Your force is directional dependent because you're taking the derivative of energy with respect to a vector, which is some distance. But--

AUDIENCE: [INAUDIBLE]

PROFESSOR: But again, the force will be directional dependent. Let me show you an example. Let's say I have a central atom here and I got four atoms around it. The force will be directional dependent because-- what is the force? The force is how the energy of an atom changes as I move it into a particular direction.

So the force along this direction is the energy change as I move the atom along this direction because it is the derivative of the total energy with respect to some unit vector in that direction. See, if I change in this direction, I'm going to significantly increase the overlap with that atom. But if I change it, say, in this direction, that's not necessarily the same derivative than in this direction. So the force will be directional dependent just because of the crystallography, but the energy is not.

So what I mean with that is that-- let's call this some unit length L , and they're all at L . If I do this, and here's my central atom, and these are all L , that gives me exactly the same embedding energy as in the first problem. Of course, the total energy will be different because these guys interact very closely. But the bonding energy of the central atom is exactly the same. And that's what I mean with no angular dependence. It doesn't in the end see what orbitals that electron density comes from.

And when you think about angular dependence, in the end, what it is, it is polarization of orbitals. Think of sp^3 bonding. What is that really telling you? If I do an sp^3 hybrid along this direction, I'm really starting to polarize the orbitals along-- what is it-- 109 degrees in other directions. So I really need to try to get my bonds in those directions. So that's essentially what angular dependence comes from. Does that help?

AUDIENCE: [? I ?] [? think ?] [? so. ?]

PROFESSOR: OK. You think so. Well, let me know if it doesn't. OK.

The other problem is-- which is more a technical problem-- is that the potential is not unique. And this is something you just have to be careful with if you start messing around with potential files yourself. If you're going to use it as a black box, you're probably OK. Let me get you some notes on that.

The way you can see that the potential is not unique. If I have an embedding function F of ρ , if I define another embedding function-- we call it G of ρ -- they're always the electron density. And I say, that's F of ρ plus some linear term in ρ . You could make the total energy function exactly the same with either of the two embedding functions just by changing the pair potential. And let me show you that.

If I write the energy as the sum over atoms, the embedding function on that atom, plus some air potential part. Let me substitute this in. So I get that this is $\sum_i F_i \rho_i + \frac{1}{2} \sum_{ij} \psi_{ij}, R_{ij}$, plus-- let me-- $\sum_i k \sum_j [?$ not equal ?] i of f_{ij} where ρ_i was $\sum_j f_{ij}$. OK.

Essentially, if you want to look through the math, don't worry too much about it. If you do a linear transformation on the embedding function, you see this term is linear in things coming from the atoms around it. It's linear in these electron density functions.

So anything that's linear in functions that I sum from the other atoms is just like a pair potential because that's what a pair potential is. I sum contributions, OK? I add up contributions from the things around me. So I'm linear in that contribution. So any linear part of the embedding function just looks like a potential.

This is a pair potential. This is the sum over [? $2n$. ?] This is i and j , and I add something up. So this I could just add the ψ . So if I add ψ to this and use the embedding function F , I have exactly the same as using the embedding function G and using just the potential ψ . So the linear part you can move arbitrarily between the embedding function and the pair potential.

So that tells you something else that's kind of important. When is a pair potential alone good in metals? It's when the embedding part doesn't change. If my atoms stay roughly at the same density, then all the changes come from the pair potential part. OK?

So what would you do-- [INAUDIBLE]-- if you're having an embedding function as a function of ρ . If you're working, say, at this density? Let's sort of say the density of the bulk, and you're studying bulk problems. You could take the slope of this, take the linear term, make a linear approximation to that and add that to the pair potential, and now you wouldn't have to worry about the embedding function.

So at constant density, pair potentials are actually pretty good. It's when the density around you-- the electron density around you changes which is, in essence, a measure of the coordination change that you run into trouble. And so, people have gone into that game of making density dependent potentials. And so you would have a different potential when you're near a surface than when you're near the bulk.

So my summary on effective medium theories which I showed by embedded atom methods is, in essence, that there's no reason not to use them. If you've decided you're going to use an empirical model for metals, there's really no reason at all to use pair potentials. Embedded atom basically does better at very little cost. Saying that it does better doesn't say that it does everything.

And I remember talking to people like [? Mike ?] [? Baskis ?] who invented this. They're embarrassed for kind of things people now do with the embedded atom method. Like with any successful method, people push it way too hard and try to calculate all kinds of subtle energy effects that it just never was made for. And you see people write papers on, if I change my constant by 1,000 here, I can get this phase correct. But in the end, it gives you these gross coordination effects right, but it does not give you very subtle hybridization effects right.

Why, for example, were platinum and palladium not so good? Because platinum and palladium have almost filled or filled D bands. And as a function of the environment around platinum and palladium, there is very subtle electron transfer between the D band and the S band. OK?

And when you're just measuring electron density, that all looks the same, but those states behave very differently. So there will be always subtle electronic change that you simply won't get. And if your system is really influenced strongly by those, you won't get them right.

For example, people try to get complex crystal structure differences in metals with EM. You will just never get there, period. Let's say you make an alloy-- palladium ruthenium, whatever-- the crystal structure energy difference in that [? material, ?] like in most metals over the order of 5 to 10 millielectron [? volt. ?] You cannot get things right on that scale with these kind of simple models, so you shouldn't try too hard.

But on the other hand, if you're going to study, say, I don't know, fracture in a material in a simple element like one of these big movies I showed you, that's to a large extent a topological event. You will get a lot of essential pieces of that right. Because what is that? That's a coordination change event, largely. You will get a lot of that right by doing EM over potentials.

This is a great website I only added that this year to the lecture I found. There's a group in Japan that called themselves the [? EAMers. ?] They even have a blog now. And so there's a lot of cool stuff you can download. They have potential files. You can get EM code from the Sandia site. It's a little hard to get to, because Sandia being a national lab, they're totally paranoid about downloading anything. But there are sources where you can get the EM code. It is technically freely available. It's just not necessarily trivial to get to it.

So we talked about pair functionals, which was this direction. So the other direction is obviously to go to cluster potentials. Well, you know, I don't like my expansion in just two-body terms. I'll start adding three-body terms. It's the obvious extension. And if you don't like that, you can go to four-body terms. That's definitely done. And for reasons I'll go into in the next lecture on Tuesday, it's extremely popular in the world of organics, where it's actually really, really powerful. And we'll go into why that is the case.

But a typical material in condensed matter where obviously this was done was silicon. Silicon is a low coordinated solid-- only fourfold coordinated in the diamond cubic structure. It's extremely hard to get that with a pair potential because again, your systems always try to form FCC or HCP, essentially.

So what you can do is build a three-body potential that favors these 109-degree tetrahedral angles. And as soon as you do that, if you actually make that force strong enough, you can do almost nothing else than form diamond cubic. Because diamond cubic is basically the way to pack atoms and put them under all these 109-- what is it-- .4 or something degree angles. So if you make a strong enough three-body force, you will by definition end up in the diamond cubic structure.

If you work with three coordinates, you can also transform that and instead work with two distances and an angle. So if I have three atoms, the coordinates of them, I can also always say that to have the relation between them, I just need to have these two distances and this angle that fully defines the three-body problem. And that's usually how it's done.

People have pair potentials that take care of the distance dependence of that. And then they have an angular potential which takes care of the angle between these. And often, the prefactor of the angular potential will have some distance dependence in it. Because of course, maybe you want to enforce this 109-degree angle when your two other atoms are close. But when they're very far, you don't really care anymore.

What's the possible choices for angular potentials? The simplest one is simply a quadratic-- a harmonic function in the angle. If you want to impose an angle θ , you make a potential that's literally a quadratic around θ . So, extremely simple to evaluate. So that function's actually often used. You can also use-- the disadvantage of this one is that it's not periodic. Angle measurements should be periodic. So you can kind of get into trouble with this if you go too far from equilibrium.

So another thing you can use is a cosine function which has the proper periodicity. If you take-- I don't know-- 109 degrees or whatever angle you want, or you take its complement-- 360 minus that-- you should get the same answer. And so, cosine functions will do that right. And this one, $\cos(\theta - \theta_0)$, this part is minimal when θ is 109 with tetrahedral angle. [CLEARS THROAT] Excuse me. Losing my voice.

So it works working with three-body potentials. It's, of course, more work. It's not $O(N^3)$. Because, for every atom, you now have to look at distance and angles with pairs of atoms. OK so it becomes an $O(N^2)$ operation, so it's more work. And calculating the forces as the derivatives of the energy is also quite a bit more work.

The most famous three-body potential is probably the Stillinger Webber potential which, as you can see, has this cosine θ form for the-- this is the angular part. This is the distance dependence of the angular part. So how much should you penalize angular deviations [INAUDIBLE]. And then this is the two-body part. Just sort of a more $O(N^2)$ potential.

Originally, Stillinger Webber, when they actually wrote down the function, they didn't even fit it very carefully. I mean, it was a great form, but they just kind of did a rough fit. And so, the original Stillinger Webber parameters and the ones people now use sometimes more recently are not the same anymore. The potential has the same form. But some people have done recently better fits. So when somebody says that Stillinger Webber potential, there are variations of it for the constants.

So what do you get right with a [? sale ?] something like a Stillinger Webber potential? Well, you can kind of guess why people did this. Because silicon is so important, they wanted to be able to model silicon right. So it's important to understand which pieces you get right.

One of the critical issues was surface reconstruction. If you take silicon 111-- sorry, 100. If you truncate it, it actually undergoes a 2 times 1 reconstruction. What essentially happens if you take this top layer atoms, they're only bonded to two atoms below them, and so they dimerize on the surface. So basically, these move together in sort of double rows. I need a better color.

And then these will move together. And so they double the periodicity in one direction. And so you have these dimers that form. And because of that, the surface atom is bonded to three other atoms-- two below it, one all the way in the surface. And so they're better bonded. And that's the 2 by 1 reconstruction. The 2 by 1 reconstruction you get right with the Stillinger Webber potential. You'd actually get it right with almost anything. It's essentially a topological reconstruction.

You know, atoms like to have more atoms around them to bond, period. And so that's what drives a 2 by 1 reconstruction. At the top layer, atoms now are at least threefold coordinated and they're actually almost fourfold coordinated instead of twofold coordinated at just an unreconstructed terminated surface.

But here's the reconstruction on the 111 surface. The reconstruction on the 111 surface is 7 by 7. So it forms a much larger unit cell on the surface. And that is not reproduced by the Stillinger Webber potential. So clearly, that's a more subtle effect.

Silicon, being such an important material, has a multitude of potentials. Everybody sort of has made one and claims theirs is better than yours. And I just show this graph not for the details, but to show you that this isn't even a tenth of the potentials made for silicon, but some of the famous ones. SW-- Stillinger Webber-- [? bismuth, ?] [? Harman, ?] and the Tersoff potentials, which are quite famous for silicon, are not even on here.

But if you see, this is the angular part of these potentials. It's very different. They all give the same answer, unlike static properties like the elastic properties of silicon, the cohesive energy, the lattice parameter. But it's sort of divided differently between pair parts and angular parts. This one clearly wouldn't give you the same stiffness of bond bending around the typical dihedral angle. And that shows up.

There's what I thought was a really nice paper. Oh, and I'm sorry, I-- oh, no. The reference [? isn't ?] here. This is the [? Norman ?] [? and ?] [? paper ?] from *Physical Review B*. They did what I thought was brilliant. They did exactly the same simulation with three different potentials.

Left is the Stillinger Webber potential, and then middle column is Tersoff [? 2. ?] This is the Tersoff potential, and then Tersoff [? 3. ?] So this is exactly the same Monte Carlo simulation with the three potentials. And this is the [? 001 ?] or [? 100 ?] surface of silicon. So remember, that's the one I showed you that dimerizes.

And so the top line is at low temperature. Tersoff [? 3 ?] doesn't dimerize. Tersoff 2 does sort of dimerize. And so does Stillinger Webber, but they look slightly different. But these are in reasonable agreement. If you go higher in temperature, then Tersoff [? 3 ?] dimerizes. So does the Stillinger Webber. But Tersoff [? 2 ?] starts kind of looking quite different already.

So these are potentials that, if you looked at how they reproduce static properties, they all do extremely well. But it's extremely hard to fit to dynamical properties because you don't have that as quantitative information very often. Like, first of all, we don't know what the right answer is here. But even if we did, let's say, somebody does brilliant STM and we know the right answer, it's very hard to feed that back into a potential.

And so the reason I'm telling you this is not to dissuade you from using [INAUDIBLE] when you need them. It's just to make sure you're cautious. There's a big reason that people do ab initio calculations these days. It's not because they necessarily like it, but it's because the results tend to be more reliable, or at least predictable. The errors are more predictable, let's put it that way, which is a useful thing to have. Having error is one thing. Having unpredictable error is a lot worse.

But still, there will always be problems where they're just too complex, or you want to do them fast where you want to use pair potentials. Just don't always believe everything you see. Try to fit as close to where you're going to explore the physics that's essentially [? meaning. ?]

OK, here's a bunch of references. I'll actually put these on the Stellar website. We'll finish the empirical potential lecture. We'll go into organics and a little bit of polymers on Tuesday. And then Professor [? Marzari ?] will start giving you introduction to basic quantum mechanics. And then we'll start sort of ab initio energy methods, and Professor [? Marzari ?] will do that. So I'll be still lecturing on Tuesday. So have a good weekend, and I'll see you on Tuesday.